ANNUAL REVIEWS Further
Click here to view this article's online features:
- Download figures as PPT slides
- Navigate linked references
- Download citations
- Explore related articles
- Search keywords

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

Proton–Lead Collisions at the CERN LHC

Carlos A. Salgado[1] and Johannes P. Wessels[2]

[1]Instituto Galego de Física de Altas Enerxías (IGFAE), Universidade de Santiago de Compostela, E-15782 Santiago de Compostela, Spain

[2]Institut für Kernphysik, Westfälische Wilhelms-Universität Münster, D-48149 Münster, Germany

Annu. Rev. Nucl. Part. Sci. 2016. 66:449–73

First published online as a Review in Advance on August 8, 2016

The *Annual Review of Nuclear and Particle Science* is online at nucl.annualreviews.org

This article's doi:
10.1146/annurev-nucl-102014-022110

Copyright © 2016 by Annual Reviews.
All rights reserved

Keywords

QCD, LHC, heavy-ion collisions, proton–nucleus collisions, quark–gluon plasma

Abstract

The single proton–lead collision run performed at the LHC has had profound consequences for the entire field of hot and dense QCD. The prior, 20-year-long experimental exploration of high-energy nucleus–nucleus collisions indicated that a hot QCD medium is created in a very short time, a fast equilibration. The striking similarities found in these nucleus–nucleus data and the a priori simpler proton–lead collisions, where no hot medium was expected, are refining our understanding of the whole process of thermalization at the most fundamental level. We review the most relevant sets of data taken during the proton–lead collisions of Run 1 at the LHC and the implications they have for our understanding of the underlying physical mechanisms and characterization of hot and dense QCD systems.

Contents

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

1. INTRODUCTION AND MAIN PHYSICS GOALS

New regimes and domains of quantum chromodynamics (QCD), the theory of the strong interaction, are being explored and characterized in collisions involving nuclear targets at high energies (see, e.g., Reference 1). The energy frontier is currently set by the Large Hadron Collider (LHC) at CERN in Geneva, where lead–lead (Pb + Pb) collisions at $\sqrt{s_{NN}} = 5$ TeV started in 2015, continuing the previous program at $\sqrt{s_{NN}} = 2.76$ TeV during Run 1. Brookhaven National Laboratory in the United States has a dedicated facility, the Relativistic Heavy Ion Collider (RHIC), that has delivered gold–gold (Au + Au) collisions at energies of up to $\sqrt{s_{NN}} = 0.2$ TeV since the year 2000. During these 16 years of operation, several different colliding systems have been studied at various energies.

These experiments are designed to deposit a large energy density distributed over a macroscopic (in terms of QCD scales) spatial region around the center of mass of the collision, complementing the traditional high-energy physics goal of producing well-localized, compact, large energy densities in which to study the properties of contact interactions or the production of new particles. Such experiments date back more than 20 years (see, e.g., Reference 2 for a brief historical perspective). They aim to study the collective or emerging properties of the QCD Lagrangian that are out of the reach of other experimental conditions. The temperature of the created system is expected to grow with both the energy of the collision and the size of the nuclei (i.e., their atomic number). In this sense, the high-energy collisions of two heavy nuclei would generate a temperature high enough to create a deconfined QCD medium. Smaller colliding systems were expected to follow

the simpler physics of dilute, noninteracting systems, following the well-known phenomenology of proton–proton (pp) collisions, with some possible cold nuclear matter modifications in the case of proton–nucleus (pA) collisions; here, we use the term cold matter to distinguish it from truly thermalized hot matter, which is the main object of study.

The basic methodology used to study hot and dense media was established during the past 20 years. The short lifetime of the QCD matter produced—some tens of fm/c at most—forbids a direct detection of its properties, which can be measured only through indirect signals. Typical examples of these signals include the suppression of quarkonia states (bound states of a heavy quark and its antiquark) due to the screening of the potential at high temperature; the suppression of particles produced at high transverse momentum (p_T), known as jet quenching; the angular structure of the particle spectrum, compatible with hydrodynamical simulations; or the relative particle-type abundances (hadrochemistry) of the produced particles. In order to identify a genuine hot matter effect, we need a benchmark with simpler collisions in which this matter is expected not to be formed. For hard processes such as the first two in the above list, a convenient way to proceed is to define the ratio of the observable cross section in nucleus–nucleus (AA) collisions to that in pp collisions normalized by the number of elementary pN collisions (where N is a proton or a neutron in the nucleus) taking place in the nuclear case. The departure from unity of this ratio thus indicates the presence of nuclear effects. Benchmarking with pA collisions is also essential for a correct interpretation of the results because cold nuclear matter effects could mimic some of the signals described above.

The situation changed profoundly following the first proton–lead (p + Pb) run at the LHC. This run, which took place at the end of the first LHC running period (Run 1) in January–February 2013, was devoted mostly to benchmarking the cold nuclear matter effects in order to properly extract the hot medium QCD parameters (3). Some indications of new interesting collective effects in small systems had been found previously; these include the presence of long-range two-particle angular correlations (known as the ridge, discussed in detail in the following sections) in those rare pp collisions in which a very large number of particles were present (i.e., high-multiplicity events). These first indications appeared during Run 1 in 2010 (4), and new results at $\sqrt{s} = 13$ TeV from Run 2 are now available (5, 6). Note that nuclear-like effects in high-multiplicity pp collisions at the LHC were proposed before the data appeared (7–12). Thus, it was not surprising that similar correlations were found in p + Pb collisions. What was surprising, however, was the strength of the effect: It was much stronger than in pp collisions and similar to that in nucleus–nucleus (AA) collisions. Interpretations favoring a small, proton-size-like, thermal system—a mini quark–gluon plasma—soon appeared (13). This possibility is one of the main topics of discussion in the high-energy physics community because it addresses one of the most fundamental open questions, namely how the QCD thermal medium is created in laboratory. Indeed, the precise way thermalization happens in AA collisions is not yet known from first principles, despite extraordinary progress in recent years (see, e.g., Reference 14). The presence of thermalization in the small systems discussed here would strongly constrain the different possible mechanisms involved, or could address the question of whether local thermalization is reached at all. These studies address the core question of how collectivity develops at the most fundamental level in QCD.

Benchmarking is still an essential part of the pA program, and it works well for the hard sector. Observables involving large scales, such as $W^{\pm}/Z$ boson production and jet or inclusive hadron production at high p_T, may provide new constraints on the partonic structure of the nucleus and the corresponding global fits to nuclear parton distribution functions (nPDFs). So far, no departure from the linear Dokshitzer–Gribov–Lipatov–Altarelli–Parisi (DGLAP) (15–17) evolution has been observed, and a description of the data in terms of universal PDFs provides excellent

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

results. Nevertheless, the physics of saturation of partonic densities, encoding the nonlinearities in evolution equations, is at the center of phenomenological analyses, especially as a tool to compute the initial stage of the collision in the soft part of the spectrum, where DGLAP evolution cannot be applied. The information contained in *pA* collisions is essential to constrain both the nonperturbative input in these analyses and the relevance of the nonlinear terms. A formalism known as the Color Glass Condensate (CGC; see, e.g., Reference 18 for a review) provides a general framework that has seen substantial progress in the past few years.

This review is organized as follows. Section 2 reports on the running conditions of the LHC as a p + Pb collider during Run 1. Section 3 discusses the results from one of the main motivations to study these systems prior to experimental data taking: the study of partonic densities in the nucleus at high energy. Section 4 reports on studies of the particle spectra at low p_T, a region of interest for hydrodynamics. Section 5 discusses the main issue emerging from the first p + Pb run at the LHC: the possibility that a hot medium is already produced in these small collision systems, contrary to previous expectations, and the implications for the solution of the problem of thermalization; a surprising effect known as the ridge is discussed at length in this section. Section 6 discusses some of the findings useful for benchmarking with p + Pb data, especially those concerning quarkonia suppression and jet quenching. Finally, Section 7 summarizes the main observations of and lessons learned from the p + Pb run.

2. THE LHC AS A PROTON–LEAD COLLIDER: RUNNING CONDITIONS IN 2013

p + Pb collisions, a mode of operation at the LHC that was not foreseen in the original proposal for the accelerator, started in September 2012 with a short pilot run in a single fill of approximately 9 h (19). This run was designed to test the operational procedures of two constant-rigidity, different-energy beams and to allow the experiments to set up triggers well in advance of the main physics run planned for January–February 2013. However, this short run provided two of the main experimental results in small systems: (*a*) the large long-range correlations in pseudorapidity ($\Delta\eta$) on the near side ($\Delta\phi \simeq 0$) of two-particle $\Delta\eta - \Delta\phi$ correlation functions, known as the ridge (20–22), and (*b*) the compatibility of the average nuclear modification factor, $R_{p+\mathrm{Pb}}$, with unity (23). The small instantaneous luminosity in this pilot run, $\mathcal{L} \sim 10^{-26}$ cm^{-2} s^{-1}, led to an integrated luminosity of several μb^{-1}. The peak instantaneous luminosity increased by a factor of $\sim 10^3$ during the main physics run of 2013, leading to an integrated luminosity of $\sim$30 nb^{-1} for ATLAS, ALICE, and CMS, whereas LHCb, participating for the first time in an LHC heavy-ion run, recorded fewer collisions by a factor of $\sim$20.

The center-of-mass energy of the collision can be derived from the equal-rigidity acceleration of the LHC. For a proton beam energy of 4 TeV (the energy of the 2012 *pp* run), the corresponding energy per nucleon of the lead nuclei ($Z = 82, A = 208$) is $Z/A \times 4\,\mathrm{TeV} \simeq 1.58$ TeV, leading to a center-of-mass energy of $\sqrt{s_{NN}} = 5.02$ TeV. The equal rigidity of the magnets also implies that there is a rapidity shift between the center of mass of the p + Pb collision and the laboratory frame of $\Delta y \approx (1/2)\log(Z/A) \simeq -0.465$—that is, a shift toward the proton-going side.

In order to account for asymmetries in the acceptance of the experiments, and to optimally explore the phase space for the proton- and lead-going sides, the beams were reversed about halfway through the run; in other words, the LHC was operated in both p + Pb and Pb + p modes. The LHC also operated for some time at a much lower luminosity ($\mathcal{L} \approx 0.5 \times 10^{28}$ cm^{-2} s^{-1}) dedicated to minimum-bias measurements, the careful exploration of possible trigger biases, and van der Meer scans of the total visible cross section (24). Reference 19 offers details of the machine operation

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

during the first *pA* run. Operation of the LHC at higher energy and with more protons in the proton beam during Run 2 would increase the luminosity by approximately one order of magnitude.

3. INITIAL STATE: THE HIGH-ENERGY STRUCTURE OF THE NUCLEUS IN TERMS OF QUARKS AND GLUONS

The goal of nuclear collisions is to study dense and hot QCD systems, which translates into complex theoretical and experimental conditions dominated by the large number of particles involved. A useful way to separate the different effects is to classify them according to the timeline of the collision process in the center-of-mass frame, in which the initial state would refer to the structure of the colliding objects in terms of quarks and gluons (see, e.g., Reference 25 for a phenomenological review). Whether the dynamics of these quarks and gluons inside the nuclei is dominated by linear or nonlinear physics is one of the most relevant questions to be addressed. In the linear regime, only splittings of one parton into two (a gluon splitting into two, for example) are allowed, whereas the nonlinear regime allows both splittings and mergings. In the first case, the usual DGLAP evolution applies, whereas the second case uses an implementation generically known as the Color Glass Condensate (CGC), which describes the phenomenon of saturation of partonic densities. Interestingly, the CGC approach has the flexibility to include most of the dynamics that are relevant in dense systems. We review the status of both in this section. This initial state, by contrast, controls the dynamics of the system right after the collisions in the eventual process toward thermalization. The study of these initial stages, which would include thermalization, is where *pA* collisions provide the most essential information.

3.1. Status of Nuclear Parton Distribution Functions After Run 1

The use of PDFs to compute hard cross sections in perturbative QCD is a mature and reliable technique. These nonperturbative objects contain all relevant information about the structure of the hadrons in terms of the quarks and gluons that scatter in elementary collisions to produce other particles. Thus, the study of PDFs is key to the phenomenological interpretation of data obtained at hadron colliders. Extracting them from experimental data—although PDFs are nonperturbative objects, their perturbative evolution is given by the DGLAP equations—is one of the most standardized methods in particle physics. It follows a global fitting procedure to a large number of data wherein the initial, nonperturbative input distributions $f_i(x, Q_0^2)$ are obtained by χ^2 minimization. Hessian or Bayesian analyses are used, and the way of parameterizing the initial conditions differs from group to group. As a result, the available PDF sets allow one to propagate the experimental uncertainties in their determination to observables not included in the fit. Interest in the study of nuclear PDFs is twofold: On one hand, it enables benchmarking of some of the cold nuclear matter effects that need to be subtracted to identify genuine hot matter signals, and on the other hand, it serves as a test of the linear dynamics that is expected to break when the gluon densities reach the saturation level.

The PDFs of protons bound in nuclei are known to be modified with respect to the free case, with suppressions or enhancements at different values of the fraction of momentum x (see, e.g., Reference 26 for a review). A relevant question is whether a universal set of nPDFs is possible within the usual collinear factorization, which symbolically reads

$$\sigma^{AB\to C} = f_i^A(x_1, Q^2) \otimes f_j^B(x_2, Q^2) \otimes \hat{\sigma}^{ij\to k} \otimes D^{k\to C}(z, Q^2) \qquad 1.$$

In this equation, $\sigma^{ij\to k}$ is the short-distance perturbative scattering of two partons i and j to produce a partonic state k that eventually decays into a final state C; x_1 and x_2 are the corresponding fractions

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

of momentum. Although no formal proof of factorization exists for nuclear targets, it is expected to work for large-enough virtualities, and its range of validity is also expected to be smaller than for the proton case due to enhancement factors $\sim A^{1/3}$ of the higher-twist terms. In the present situation, no departure from universality has been found for the nPDFs, whereas a departure from universality of the hadronization process has been clearly observed, as in the jet quenching phenomenon in AA collisions (see, e.g., Reference 27 for a recent review). All available global fits for nPDFs (28–31) include data from deep-inelastic scattering (DIS) on nuclear targets and pA Drell–Yan production. Both sets of data are quite old and have a very limited kinematical reach in x and Q^2. Moreover, the DIS data were measured with a variety of targets, only some of which are heavy enough to constrain nuclei such as lead or gold (i.e., those used in LHC and RHIC experiments). Some global fits include other sets of data, such as inclusive pion production produced at high p_T in deuteron–gold (d + Au) data at RHIC or neutrino DIS (neutrino DIS data are usually taken with either lead or iron targets). No global fit including LHC data is available at present, but this situation may soon change.

Reweighting is a powerful way to quantify the agreement of the PDFs with a new set of data, or the relevance of this particular set of data to provide further constraints to the PDFs (32–34). This technique explores how well the new data can be accommodated within the parameter space allowed by the error analyses performed in the global fits. Most reweighting studies are performed using a Bayesian procedure (33–35), although a purely Hessian method is also possible (36). For simplicity, we provide the most important components of the former method.

The method starts with a series of Monte Carlo replicas of the PDFs, namely an ensemble of $N_{\rm rep}$ different sets of PDFs, $[f_k], k = 1, \ldots, N_{\rm rep}$, which is a sample of the PDF probability density, $\mathcal{P}$. How this ensemble is computed depends on the actual method used for the PDF fit; it is natural within the NNPDF Collaboration's approach (33, 34) and is generated a posteriori with the more frequent Hessian analyses. Any quantity depending on the PDFs, be it a physical observable or not, can be evaluated by the average

$$\langle \mathcal{O} \rangle = \frac{1}{N_{\rm rep}} \sum_{k=1}^{N_{\rm rep}} \mathcal{O}[f_k]. \qquad 2.$$

Bayesian interference allows one to update the old probability distribution $\mathcal{P}_{\rm old}[f]$ to a new one, $\mathcal{P}_{\rm new}[f]$, by computing a new weight w_k for each of the PDF replicas f_k. The computation of these weights is not especially complicated, so we refer the reader to the original publications. Thus, new quantities depending on the PDFs can be computed as

$$\langle \mathcal{O} \rangle_{\rm new} = \frac{1}{N_{\rm rep}} \sum_{k=1}^{N_{\rm rep}} w_k \mathcal{O}[f_k]. \qquad 3.$$

With a large-enough number of replicas, this procedure would be completely equivalent to performing a new PDF fit, for the ideal case of perfectly compatible data, symmetric Hessian errors, and so on. However, in practice, the accuracy of the representation of the real distribution $\mathcal{P}$ provided by the replicas may be not good enough, and the final result could suffer from these deficiencies. In order to determine the accuracy of the results, the method provides quantitative tests of the degree of compatibility of the new data with the existing global fits and their extra constraining power. These analyses are crucial for the nuclear case, in which the lack of experimental data induces parameterization-bias effects due to assumptions made at the level of the initial conditions of the evolution (e.g., the behavior at small x). **Figure 1** shows an example of these analyses, including p + Pb data.

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

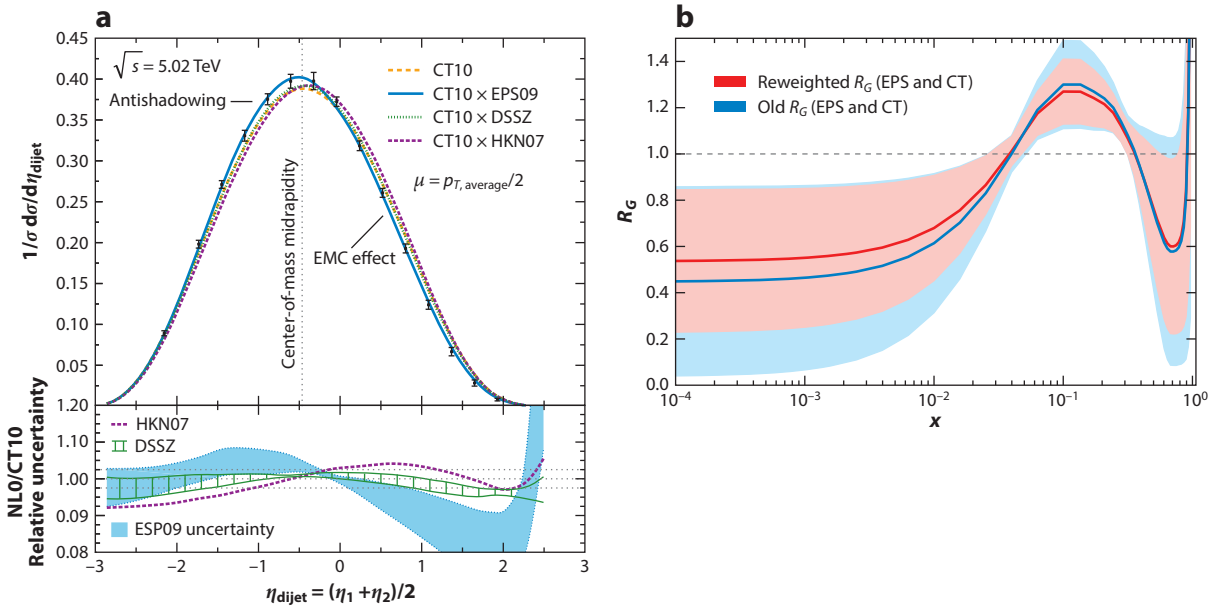

Figure 1

(*a*) CMS dijet data (37) presented as differences between the data and the theory calculations. The dashed lines correspond to the nPDF uncertainty. (*b*) The corresponding reweighted gluon distribution (actually, the ratio of the nuclear to proton gluon PDF) including LHC p + Pb data. In panel *a*, the HKN07 line is from Reference 31, the DSSZ line is from Reference 29, and the EPS09 line is from Reference 28. Abbreviations: EMC, European Muon Collaboration; NLO, next-to-leading order; nPDF, nuclear parton distribution function. Panel *a* modified from Reference 38. Panel *b* modified from Reference 39.

The only comprehensive study of the relevance of nPDF analyses in the description of the LHC p + Pb data to date can be found in Reference 39. This paper analyzes most of the available observables that could potentially provide further constraints. The findings can be summarized as follows.

1. Collinear perturbative QCD factorization works very well with the two available sets of nPDFs considered [EPS09 (28) and DSSZ (29)] for all data included in the analysis.
2. The systematic errors in some of the data make the analysis inconclusive with respect to the relevance and constraints of these data for nPDF fits.
3. However, some of the data sets, in particular the dijet data from CMS (**Figure 1**), do provide constraints, favoring the EPS09 set.
4. The *pp* data taken during Run 2 at the reference energy of $\sqrt{s_{NN}} = 5.02$ TeV could significantly improve the situation and render these p + Pb data much more useful due to better control over the systematic uncertainties.

3.2. Dense Gluonic Systems in the Initial State: Saturation of Partonic Densities

Collinear perturbative QCD is expected to work when the colliding objects are dilute, meaning that linear dynamics dominates and strong multiparton correlations are irrelevant in the partonic wave functions of the high-energy colliding hadrons. At large partonic densities, the linear approximation ceases to be valid, and a phenomenon of saturation eventually appears. This saturation

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

can be understood as arising from gluon fusion corrections to the usual splittings in the evolution equations, which slow the growth of the parton distributions toward small x and ensure that unitarity is not violated (40). These effects are expected to appear at a small fraction of momentum x, when the partonic (mainly gluonic) distributions become very large. Simple geometrical considerations imply that an easy way to enhance these densities is through collisions of heavy nuclei, which provide an enhancement factor $\sim A^{1/3}$ with respect to the proton. In the (current) absence of lepton–nucleus colliders, which would be the ideal experimental setting for the study of this physics, pA collisions may be an excellent way to study dilute-dense collisions, assuming the proton can be described by the usual linear equations.

The theoretical progress in this field over the past 20 years has been remarkable. Most of this research has been performed within a generic framework known as the CGC, which originated in seminal research by McLerran & Venugopalan (41, 42). These authors computed the nuclear parton densities at small x, solving the color Yang–Mills equations with appropriate sources. The small-x evolution of this setup was then computed (43–48) and is known as the Balitsky–Jalilian-Marian–Iancu–McLerran–Weigert–Leonidov–Kovner (B-JIMWLK) equation, whose mean-field approximation, the Balitsky–Kovchegov (BK) equation (43, 49), can be written in the simple form

$$\frac{\partial \phi(x, k_t)}{\partial \log x} = \mathcal{K} \otimes \phi(x, k_t) - \phi(x, k_t)^2 \qquad 4.$$

Here, $\phi(x, k_t)$ is the unintegrated gluon distribution at small x and $\mathcal{K}$ is the Balitsky–Fadin–Kuraev–Lipatov (BFKL) kernel. The linear part of Equation 4 is the BKFL equation, which diffuses in the IR region, leading to an exponential growth of soft gluons and an eventual violation of the unitarity bound. This growth is controlled here by the last, nonlinear term, ensuring that unitarity is fulfilled. These and similar equations have been extensively used in phenomenological applications of single-particle or multiparticle production in proton (deuteron)–nucleus collisions at RHIC and the LHC (see, e.g., Reference 25). In general terms, these formalisms predict (*a*) a reduction of the particle production due to the suppression that the nonlinear terms impose on the partonic densities and (*b*) an enhancement of the multiparticle correlations due to color correlations in the transverse plane among different partons. Consequences of the latter effect are discussed further below, in the discussion of the ridge in Section 5.

A simple way to estimate when these effects should start to be sizable involves defining a saturation scale, whose simplest estimate gives

$$Q_{\rm sat}^2 \simeq Q_0^2 \left(\frac{x_0}{x}\right)^{\lambda} A^{1/3}, \qquad 5.$$

with the phenomenological value $\lambda \sim 0.3$ (50). This scale encodes the energy and geometry dependences (through x and A) of the problem. Saturation of partonic densities is expected to become dominant when $Q^2 \leq Q_{\rm sat}^2$. Numerical results give $Q_{\rm sat,Pb}^2 \simeq 1...2$ GeV^2 for central rapidities at the LHC and larger values at forward rapidities. Thus, the LHC provides good kinematical conditions for these effects to become visible.

The nonlinear BK equation (Equation 4) is known at next-to-leading order (NLO) (51–53) and is used to extract the unintegrated gluon distribution [or the scattering amplitude $N(\mathbf{r}, y)$ in the more common coordinate-space representation of the problem]. This equation has been used in phenomenological studies analogously to the global fit procedure for DGLAP analyses of PDFs (see, e.g., Reference 54, which uses a simplified version of the NLO equations). Note, however, that although the knowledge of the evolution equation itself (e.g., Equation 4) is quite precise and rigorous, the same cannot be said for the observables. For example, for single inclusive production of particles at moderately high p_T, traditionally one of the most important phenomenological

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

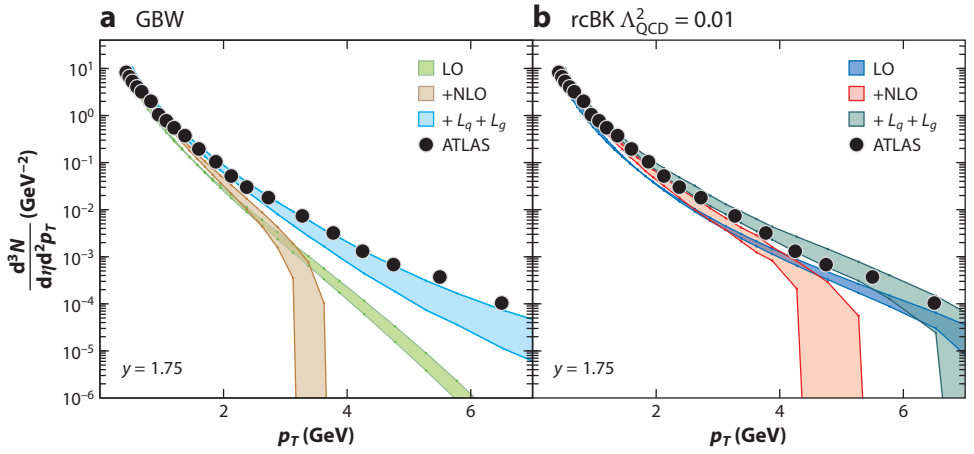

Figure 2

Comparison of NLO calculations in the Color Glass Condensate framework using different unintegrated gluon distributions (GBW or rcBK; see Reference 57 for definitions) with experimental data from p + Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV with rapidity $y = 1.75$. Abbreviations: LO, leading order; NLO, next-to-leading order. Modified from Reference 57.

observables, the corresponding theoretical treatment usually relies on a factorized formula, which has a much less satisfactory range of applicability to actual experimental conditions.

One problem involves how to treat the two different colliding objects, one in the dilute regime and the other in the dense regime, with a rigorous resummation of the different logarithms and kinematic limits involved. In the past few years, researchers have developed NLO calculations in a hybrid approach, in which the proton is basically treated in DGLAP approximation (55, 56). As an example, a recent paper (57) includes these refinements and compares the calculations with experimental data from RHIC d + Au and LHC p + Pb collisions (**Figure 2**). Recent developments of the solution of the BK equation with full NLO terms and improvements on its stability and kinematics can be found elsewhere (58–63).

Saturation approaches can also reproduce the rapidity dependence of the multiplicity in p + Pb data. **Figure 3** compares predictions from the ALICE Collaboration (64) with corresponding data from several models (67). Similar calculations have also been applied to Pb + Pb collisions, with good results.

A paradigmatic example of a "smoking gun" indicating the presence of saturation dynamics is the suppression of inclusive particles at moderately high p_T, $p_T \sim Q_{sat}$: Color screening, which reduces the gluon densities in the partonic wave function, was expected to strongly suppress the particle yields with respect to the pp case (for a review, see Reference 3 and references therein). However, more careful analyses led to a more complicated situation in which strong suppression is present only at quite forward rapidities, $y \sim 4$, whereas moderate suppression and even moderate enhancement appear at central rapidities in p + Pb collisions at the LHC (25). In fact, the actual situation is still unclear due to the limitations on the theoretical control of NLO corrections and the treatment of the transverse dynamics and geometry. It seems clear, however, that the results for suppression are similar to those predicted by nPDFs in the collinear factorization, making the interpretation of the data inconclusive. **Figure 4** compares the nuclear modification factor

$$R_{p+\mathrm{Pb}} = \frac{\mathrm{d}\sigma^{p+\mathrm{Pb}}/\mathrm{d}p_T}{A\mathrm{d}\sigma^{p+\mathrm{Pb}}/\mathrm{d}p_T} \qquad 6.$$

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

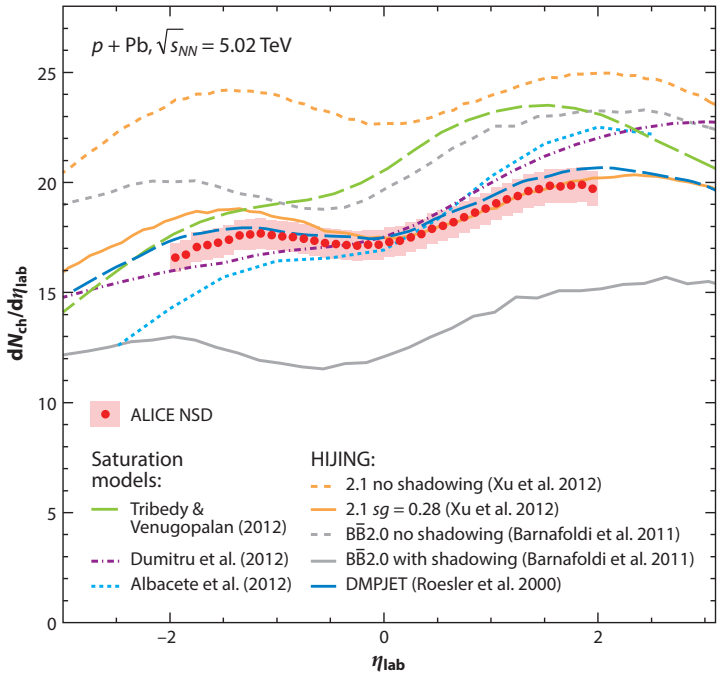

Figure 3

The p + Pb charged-particle multiplicity density, measured by the ALICE Collaboration, compared with several models, including Color Glass Condensate predictions. Abbreviation: NSD, non–single diffractive. Modified from Reference 67.

for inclusive charged-particle production, measured by ALICE (23), with several theoretical expectations. This ratio is defined to be unity in the absence of nuclear effects, when the p + Pb cross section is an independent superposition of incoherent pp collisions. Forward rapidities, where the suppression predicted by the CGC framework can be clearly identified, can be accessed by the LHCb experiment. In this region, the clear differences between the two approaches should become visible (3) and the data discriminating.

4. ARE THE PARTICLE SPECTRA IN PROTON–LEAD COLLISIONS THERMAL?

The degree of thermalization of the system created in AA collisions is frequently studied by particle species spectra at low p_T. In these studies, the thermal and (hadro)chemical equilibrium translates into well-defined predictions for these spectra, with a single temperature and baryochemical potential as fitting parameters. Contrary to initial expectations, the early p + Pb data revealed that hadron production can be successfully described in terms of thermal models. Although there is a nontrivial relationship between the collision geometry (69, 70) and the observed charged-particle multiplicity, the hadron ratios approach the values observed in Pb + Pb collisions for the case of the p + Pb collisions with the highest multiplicity (71). These ratios are rather well reproduced by grand-canonical distributions with vanishing baryochemical potential and a chemical freeze-out temperature of $T_{\rm ch} \simeq 156$ MeV, with no need to invoke empirical saturation parameters for particular quark species (71).

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

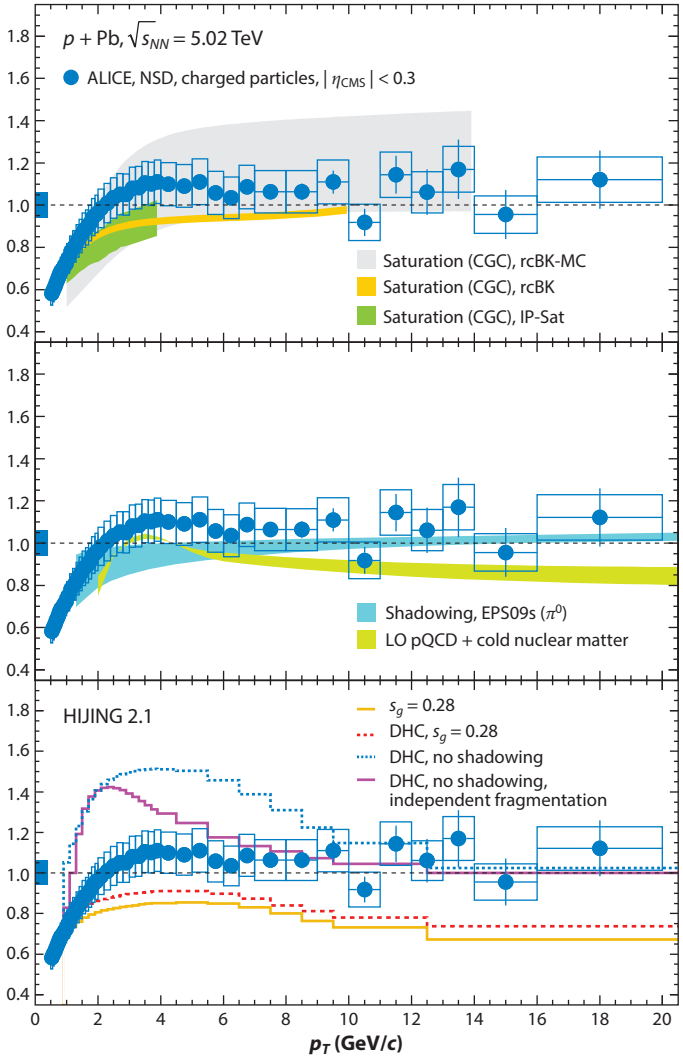

Figure 4

p_T dependence of the nuclear modification factor $R_{p+\mathrm{Pb}}$ of charged particles measured in p + Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV ($|\eta_{\mathrm{CMS}}| < 0.3$), compared with model calculations. In the top panel, the rcBK-MC curve is from Reference 64 and the rcBK and Ip-Sat curves are from Reference 68. In the bottom panel, the HIJING 2.1 data are from Reference 66. Abbreviation: NSD, non–single diffractive. Modified from Reference 23.

p_T spectra are also compatible with an interpretation favoring a strong collective expansion (72), also known as radial flow. The average transverse momentum, $\langle p_T \rangle$, increases with the multiplicity of the event and scales with the mass of the hadron at the same multiplicity (**Figure 5**). In addition, the individual hadron spectra have been fitted using a blast-wave formalism developed in Reference 73, and the data are compatible with strong radial flow. Whereas the extracted temperatures are similar at the same multiplicity density, $\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}\eta$, the common radial expansion velocity is slightly larger in p + Pb versus Pb + Pb collisions at the same multiplicity density. **Figure 6** shows a detailed comparison of the spectra to the fits and to model calculations.

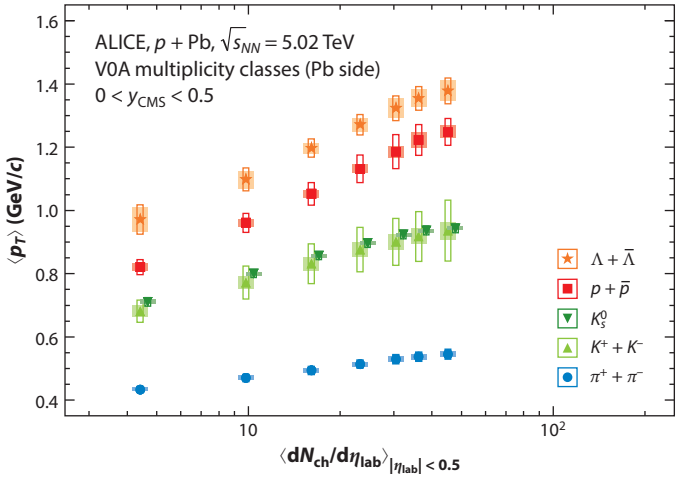

Figure 5

Average p_T as a function of $dN_{ch}/d\eta$ for different particle species measured in the rapidity interval $0 < y_{CMS} < 0.5$ (values for K^0_s are shifted for visibility). Empty boxes show the total systematic uncertainty; filled boxes indicate the contribution uncorrelated across multiplicity bins (72).

Figure 6 also shows the p_T distributions of identified particles in high-multiplicity events (the top 5–10% measured on the lead-going side) and compares them with predictions from three models: the QCD-inspired DPMJET (74) generator, the Kraków viscous hydrodynamic model with fluctuating initial conditions (75), and EPOS LHC 1.99 v3400 (76). Although DPMJET can reproduce the overall pseudorapidity distribution of charged particles, it does not reproduce the p_T spectra. The inclusion of final-state effects may be needed in order to reach better agreement with the data (72). In particular, at low p_T, where hydrodynamics drives the shape of the spectra, the data are well described by the Kraków model, whereas the agreement is weaker at high p_T, where hard scatterings contribute. EPOS quantitatively describes the pion and proton data, but the agreement for strange particles, especially kaons, is worse. When final-state interactions are turned off in simulations of both pp and p + Pb reactions (72), tension with the data is much greater (76).

5. THE RIDGE IN PROTON–NUCLEUS COLLISIONS: A GATEWAY TO THE STUDY OF THERMALIZATION?

The measurement of particle correlations plays a central role in the identification of collective behavior in high-energy collisions. In recent years, a special type of correlation, measured in both the azimuthal and polar angles with respect to the beam axis, has played a prominent role. Instead of the polar angle, a variable known as pseudorapidity (η) is normally considered: $\eta = -\ln[\tan(\theta/2)]$. **Figure 7*a*** shows the typical structure of the correlation function for two particles, in terms of the differences in their azimuthal angle and pseudorapidity variables, when only short-range correlations, typical for back-to-back particle production in perturbative QCD, are present (the broad distribution in the away side is a consequence of the difference in the center-of-mass frame of the partonic collision with respect to the laboratory value). **Figure 7*b*** presents an excellent visualization of the phenomenon known as the ridge. The near-side peak is now accompanied by a long ridge, extending several units in rapidity. Such a long-range correlation can be created only if some collective behavior is present.

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

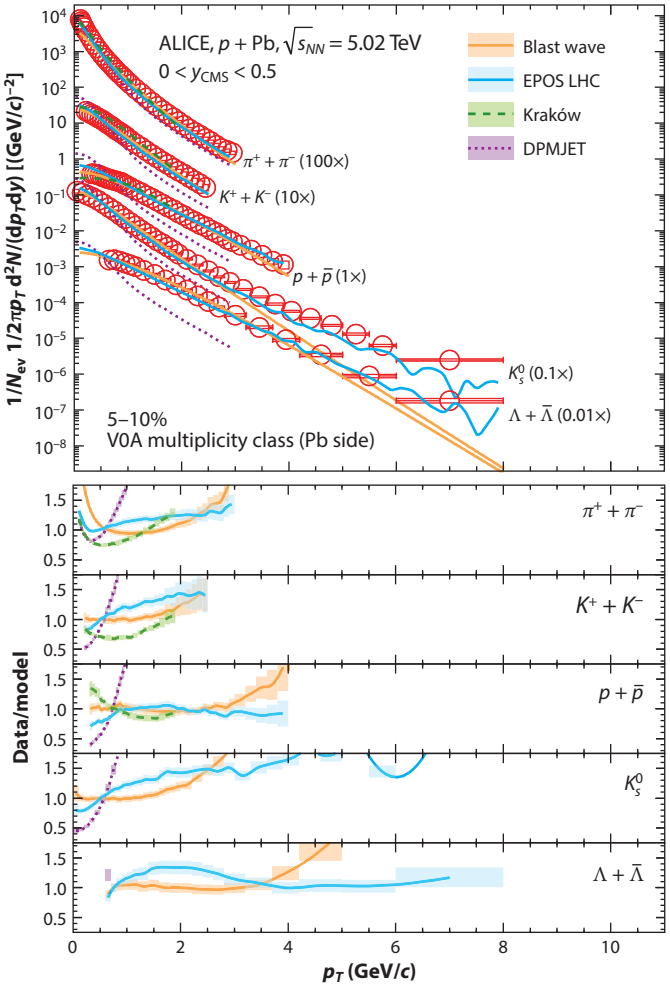

Figure 6

Pion, kaon, proton, and Λ p_T distributions in high-multiplicity events measured in the rapidity interval $0 < y_{CMS} < 0.5$, compared with predictions from several models (72): the QCD-inspired DPMJET (73) generator, the Kraków viscous hydrodynamic model with fluctuating initial conditions (75), and EPOS LHC 1.99 v3400 (76).

The ridge was first observed at RHIC in Au + Au collisions at $\sqrt{s_{NN}} = 200$ GeV (79–81). During the first few months of the LHC *pp* run in 2010, the small instantaneous luminosities enabled an experimental study of events with large multiplicities in which, surprisingly, a ridge was also identified (4); the ridge is absent in *pp* collisions except in these very rare high-multiplicity events. However, the corresponding strength of the signal is significantly smaller than the one in *AA* collisions. The great effort expended during the first Pb + Pb run to study the ridge led to the identification of interesting features, particularly the presence of sizable odd harmonics in the Fourier decomposition of the signal in the azimuthal angle (see below). The presence of these odd harmonics is especially relevant because, according to the symmetry of the collision, they should cancel unless event-by-event fluctuations break this global geometrical symmetry.

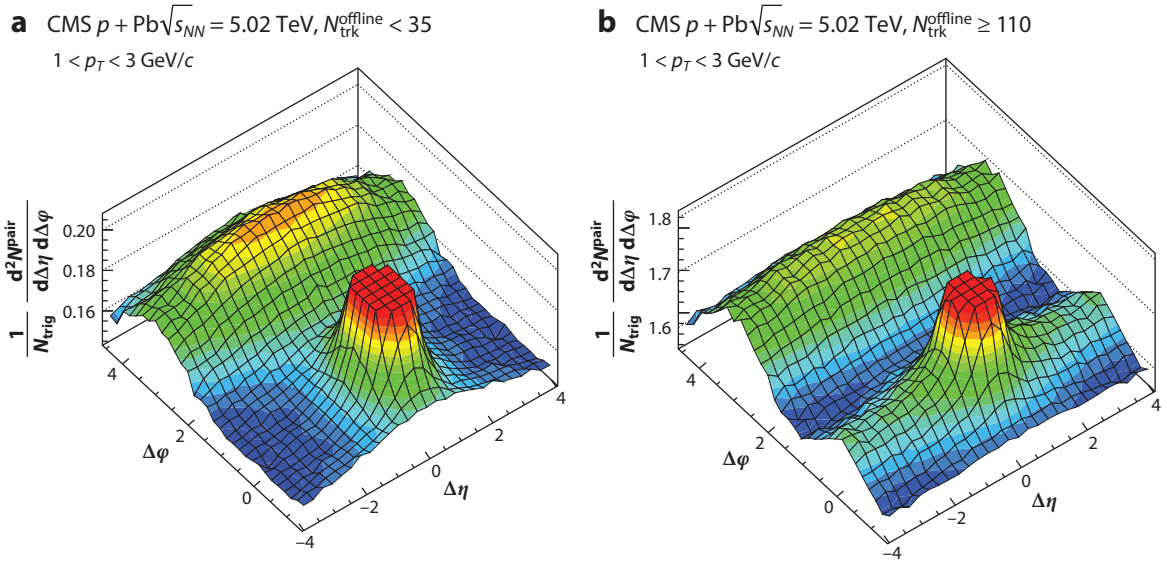

Figure 7

($\Delta\eta$, $\Delta\phi$) two-particle correlation function for two different event classes, with a (*a*) small or (*b*) large number of tracks, where a clear ridge structure appears for the second. The data were taken by the CMS Collaboration in p + Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV. Modified from Reference 20.

One of the most relevant observations from the p + Pb run at the LHC has been the strength of the ridge signal (20–22). These data have generated a lot of discussion on the origin and the underlying physics of the collective behavior needed to generate the ridge. The first ridge results from p + Pb collisions were published with the very few data collected during the pilot run in December 2012. The much larger statistics of the full p + Pb run enabled measurements of different properties of this structure, such as the hadronic composition (71), the presence of a double ridge (21, 22), and the azimuthal angle harmonic decomposition of n-particle correlations (82–84). Recently, the LHCb Collaboration (85), which participated in the nuclear program of the LHC for the first time for the p + Pb run, released its data on the ridge at forward and backward rapidities. The wealth of these data reveals a very rich structure and opens new lines of research into what are known as small systems due to the possibility of producing some collective behavior at transverse sizes much smaller than the nuclear radius.

What makes the ridge extremely interesting is that, according to causality, the long-range correlations in rapidity can be built up only in the very early stages of the collision, at times comparable to or shorter than the ones estimated for the thermalization of the system. Studies of these structures in small systems are expected to shed light on the collective dynamics at work. In the following subsections, we discuss the experimental observables and the competing theoretical explanations.

5.1. The Experimental Characterization of the Ridge

The ridge is found in correlations between a so-called trigger particle with p_T values of the order of several GeV and another associated particle with similar, or slightly smaller, p_T values. The

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

associated yield per trigger particle is defined as

$$\frac{1}{N_{\text{trig}}}\frac{\mathrm{d}^2 N_{\text{assoc}}}{\mathrm{d}\Delta\eta \mathrm{d}\Delta\phi} = \frac{S(\Delta\eta, \Delta\phi)}{B(\Delta\eta, \Delta\phi)}, \quad 7.$$

where N_{trig} is the total number of trigger particles in the event class and the $p_{T,\text{trig}}$ interval; the signal distribution $S(\Delta\eta, \Delta\phi)$ is the associated yield per trigger for particle pairs from the same event,

$$S(\Delta\eta, \Delta\phi) = \frac{1}{N_{\text{trig}}}\frac{\mathrm{d}^2 N_{\text{same}}}{\mathrm{d}\Delta\eta \mathrm{d}\Delta\phi}; \quad 8.$$

and the distribution $B(\Delta\eta, \Delta\phi)$ corrects for random combinatorial background and pair-acceptance effects. **Figure 7** presents the correlation function (Equation 7) measured by the CMS Collaboration (4) for p + Pb collisions. Experimental results indicate that both short- and long-range correlations are present in the signal. Short-range correlations are typical of $2 \rightarrow 2$ jetlike processes, a physics that is well understood according to perturbative QCD. To isolate the long-range signal, the ALICE Collaboration (21) studied the ridge structure by subtracting the jetlike short-range correlations. The result is an interesting double-ridge structure, with similar ridge signals in both the near- and away-side hemispheres (**Figure 8**). The ATLAS Collaboration (22) obtained similar results.

The correlation function (Equation 7) is usually characterized by a study of the coefficients of the Fourier expansion in the azimuthal angle. This process follows the well-known example of flow measurements in AA collisions through these coefficients, the best known of which is the elliptic flow, v_2. Indeed, Fourier coefficients, or other derived quantities as cumulants, have been extensively employed in AA collisions both at RHIC and at the LHC as a powerful tool to identify and characterize the hydrodynamical behavior of the systems created in these collisions. Although different actual definitions of the Fourier decomposition are available (and mathematically related), for simplicity we use the following:

$$\frac{\mathrm{d}N_{\text{pair}}}{\mathrm{d}\Delta\phi} \propto 1 + 2\sum_n v_{n,n}(p_{T,\text{trig}}, p_{T,\text{assoc}})\cos(n\Delta\phi). \quad 9.$$

Under some conditions, $v_{n,n}(p_{T,1}, p_{T,2}) = v_n(p_{T,1})v_n(p_{T,2})$, and a single value of v_n is often extracted from this decomposition (see, e.g., Reference 86 and the original experimental papers for precise definitions). The n-odd harmonics would vanish in the case of an azimuthally symmetric collision; in early analyses, this was usually considered to be the case. The collision of two nuclei, or of a proton or deuteron with a nucleus, does indeed present this symmetry, on average. Event-by-event fluctuations, however, break this azimuthal symmetry, and the odd components do not vanish. The study of these harmonics has been proposed as a powerful way to pin down the initial conditions to be used in hydrodynamical computations (see, e.g., Reference 87 for a recent review). The study of these initial stages has been a primary topic of research during the last few years. This increasing interest was triggered by the first LHC Pb + Pb data, in which odd harmonics were first measured, as well as by the later p + Pb data.

Experimental observations of the ridge have revealed sizable v_n coefficients; n values of two to five have been measured. The ordering and shape of the coefficients in p_T resemble those found in large systems; for example, there is significant signal even at high p_T. **Figure 9** depicts experimental measurements of the v_2 coefficient in p + Pb collisions at the LHC. Importantly, a large v_3 was found, which is a direct consequence of the event-by-event geometry fluctuations in the initial state. The precise origin of these fluctuations is not completely understood, although several models can reproduce the data. The systematics in the particle species dependence of the v_2 measured from the ridge is also very similar to that in large AA systems: v_2 is larger for baryons

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

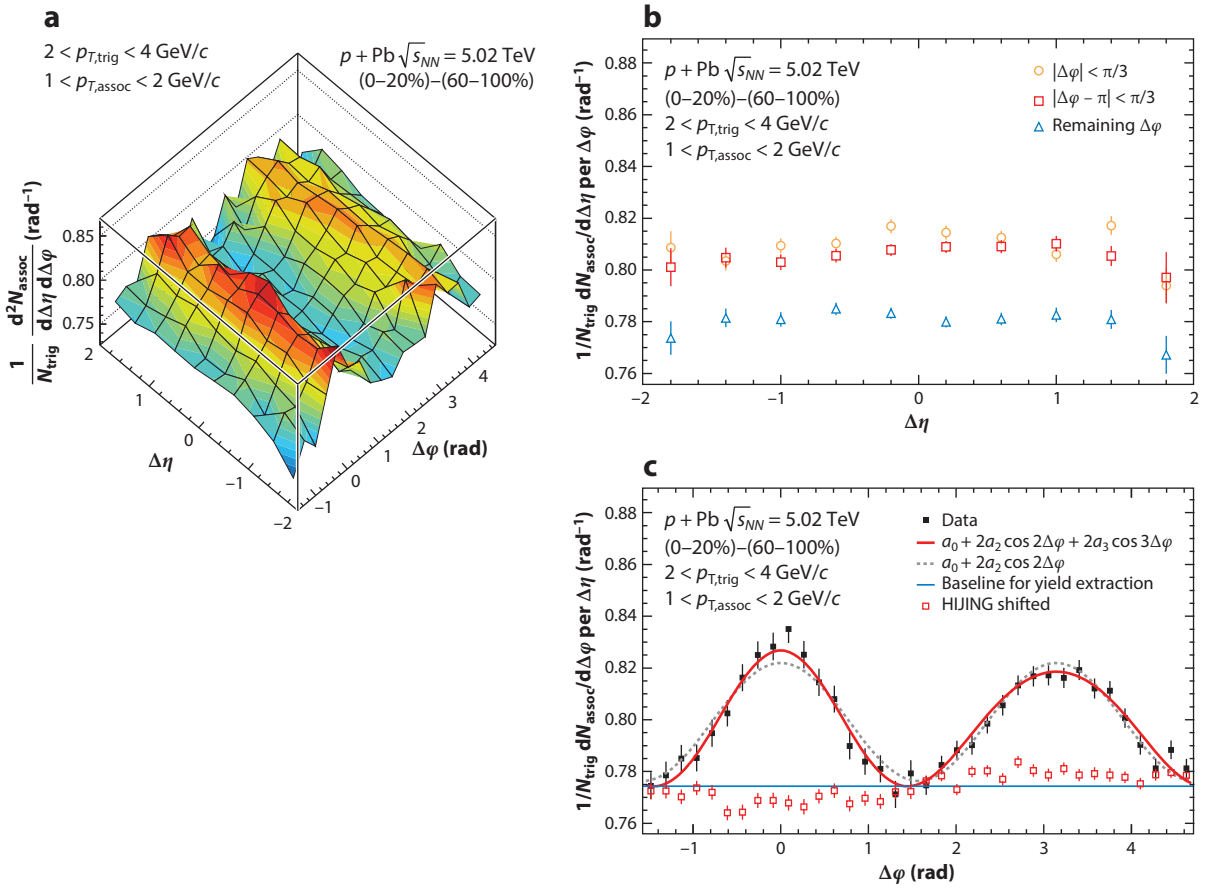

Figure 8

(*a*) Associated yield per trigger particle in $\Delta\phi$ and $\Delta\eta$ for pairs of charged particles with $2 < p_{T,\mathrm{trig}} < 4$ GeV/c and $1 < p_{T,\mathrm{assoc}} < 2$ GeV/c in p + Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV for the 0–20% multiplicity class, after subtraction of the associated yield obtained in the 60–100% event class. (*b*) The associated yield per trigger after subtraction (as shown in panel *a*) projected onto $\Delta\eta$ averaged over $|\Delta\phi| < \pi/3$ (*orange circles*), $|\Delta\phi - \pi| < \pi/3$ (*red squares*), and the remaining area ($\Delta\phi < -\pi/3$, $\pi/3 < \Delta\phi < 2\pi/3$, and $\Delta\phi > 4\pi/3$; *blue triangles*). (*c*) Data are as in panel *b* but projected onto $\Delta\phi$ averaged over $0.8 < |\Delta\eta| < 1.8$ on the near side and $|\Delta\eta| < 1.8$ on the away side. Superimposed are fits containing a $\cos(2\Delta\phi)$ shape alone (*black dashed line*) and a combination of $\cos(2\Delta\phi)$ and $\cos(3\Delta\phi)$ shapes (*red solid line*). The blue horizontal line shows the baseline obtained from the latter fit, which is used for the yield calculation. Also shown for comparison is the subtracted associated yield when the same procedure is applied on HIJING shifted to the same baseline. Only statistical uncertainties are shown. Systematic uncertainties are mostly correlated and affect the baseline. Uncorrelated uncertainties are less than 1%. Modified from Reference 21.

than for mesons at small p_T, with a crossing point at $p_T \sim 2$ GeV. ALICE has measured the relative behavior of protons and pions, as well as and Λ mesons and kaons.

Stronger tests of the underlying dynamics are expected to be accessed through measurements of multiparticle correlations; a generalization of Equation 9 can be used to compute $v_n\{N\}$, where n is the order of the harmonics extracted from the N-particle cumulant. Measurements performed by CMS on $v_2\{N\}$ (**Figure 9**) (82) show that all of the harmonics are of the same order: $v_2\{4\} \simeq v_2\{6\} \simeq v_2\{8\} \simeq v_2\{\mathrm{LYZ}\}$, where LYZ stands for Lee–Yang zeros, a method that was devised to eliminate nonflow components (88) and that can be understood as a correlation among an

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

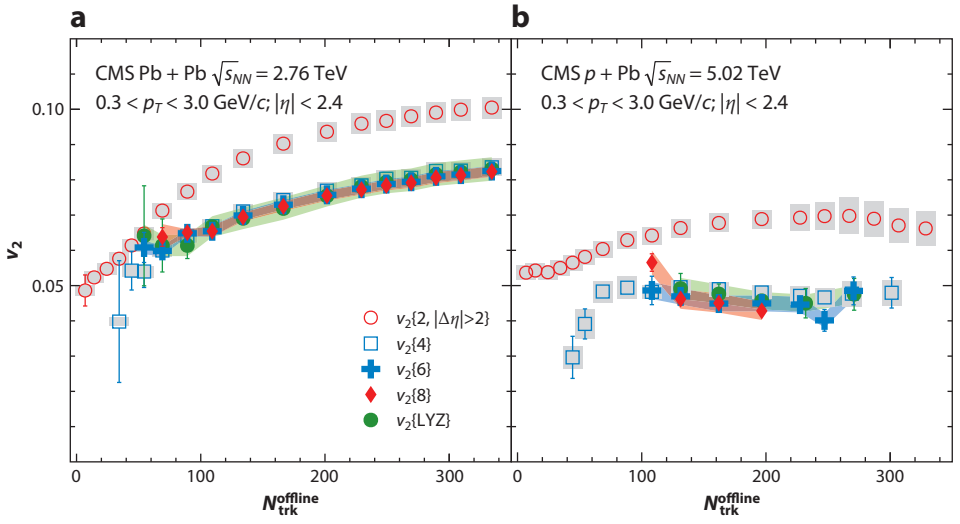

Figure 9

The v_2 values for different n-particle correlations as a function of the number of tracks, measured by the CMS Collaboration. (*a*) Results for Pb + Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. (*b*) Results for p + Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV. Abbreviation: LYZ, Lee–Yang zero. Modified from Reference 82.

asymptotically large number of particles. This measurement is a strong indication of the collective nature of the correlations. The similarity among $v_n\{N\}$ for a given n and different values of N is very natural in hydrodynamical approaches, in which these coefficients can easily be related to one another due to the independent hadronization during the freeze-out (as explained in the next section). Computing these multiparticle correlations in the CGC, the competing explanation, is not straightforward, because doing so involves the $2N$-correlation functions of Wilson lines, which are an unsolved problem.

The ridge, and its characterization, is one of the most important observations from the LHC studies of hot and dense QCD, especially that found in p + Pb collisions. As yet, there is no consensus regarding the actual interpretation in the different systems. In the remainder of this section, we discuss the two main explanations of the experimental findings, either in terms of hydrodynamics or in terms of color correlations present in the initial wave function.

5.2. Hypothesis 1: The Ridge Originates from a Locally Thermalized Hydrodynamical Medium

Hydrodynamical approaches provide a very successful description of soft and semisoft particle production data in nuclear collisions at both RHIC and LHC energies. The main assumption is that the system reaches local thermal equilibrium at some stage of the collision, so that the subsequent evolution can be described by the hydrodynamical equations. The basic equation to be solved in hydrodynamics is the conservation of the energy-momentum tensor

$$\partial_\mu T^{\mu\nu} = 0, \qquad 10.$$

which in the ideal case simply reads

$$T^{\mu\nu}_{\text{ideal}} = (\varepsilon + p)u^\mu u^\nu - pg^{\mu\nu}, \qquad 11.$$

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

for a timelike flow four-vector u^μ; ε is the energy density, p is the pressure, and $g^{\mu\nu}$ is the metric tensor. This system of four coupled partial differential equations needs an extra condition in order to be closed; this condition is provided by the equation of state $p = p(\varepsilon)$. Viscosity corrections to Equation 11 are known, and state-of-the-art approaches employ viscous hydrodynamics to reproduce experimental data (see, e.g., Reference 87). The most common way to include these corrections is to expand the energy-momentum tensor around $T^{\mu\nu}_{\text{ideal}}$ and compute the corrections $\Pi^{\mu\nu}$. For example, different gradient combinations of the flow quantities appear to be multiplied by coefficients as the shear or bulk viscosities (see, e.g., Reference 89 for a recent pedagogical introduction).

The hydrodynamical equations are usually solved numerically up to some freeze-out temperature. Note that no correlations are taken into account in the freeze-out process, so that the n-particle distribution is simply the product of n single-particle distributions. This independence of the particle production leads to constraints on the final multiparticle correlations that arise only from the global geometry of the system at freeze-out and translate into relations for the different harmonics or cumulants, such as the one mentioned above: $v_2\{4\} \simeq v_2\{6\} \simeq v_2\{8\} \simeq v_2\{\text{LYZ}\}$. To solve the partial differential equations, the boundary conditions are taken at an initial time τ_0, which is assumed to be the time at which a local thermal equilibrium has been reached. Most of the numerical analyses require a very small time, $\tau_0 < 1$ fm/c, to describe the experimental data. The initial conditions are usually taken either from a simple Glauber model or from a more sophisticated calculation, such as one of the different implementations of the CGC. These conditions at the time of the collision are, however, very far from thermal equilibrium, as there are strong anisotropies between the longitudinal momentum and the p_T, and in the absence of a good description of the thermalization process, some matching procedure is employed.

The underlying dynamics is hidden in the hydrodynamical equations, which do not address the origin or mechanism of thermalization or the magnitude of the transport coefficients. This dynamics is more a tool to characterize these properties by fitting the experimental data, as well as a good check of the degree of thermalization. Microscopic scales are not relevant when the typical mean free path of the particles in the medium is much smaller than the medium size, the main assumption of hydrodynamics. For these reasons, hydrodynamics was not expected to be useful in small systems such as those created in pp or p + Pb collisions. Nevertheless, hydrodynamics has been employed to describe these collisions and to predict some of the data with very reasonable results; it has, in fact, provided perhaps the best description of the data in the soft sector to date (75, 90–94). The ridge has been the main focus of these studies.

In order to obtain a ridge structure from hydrodynamics, one must have a suitable set of initial conditions for the partial differential (Equation 10). In modern calculations, different event-by-event initial conditions are evolved through the hydrodynamical equations, naturally leading to nontrivial angular structures as the ridge. Long-range rapidity correlations arise from the freeze-out of objects extending in the longitudinal direction. Again, the origin of these objects is not provided by hydrodynamics, in which several different models are employed for these initial conditions. Also, note that the results depend strongly on the assumed granularity of the initial conditions; fine-grained initial conditions would generally present larger gradients and, hence, larger effects than coarse-grained ones.

5.3. Hypothesis 2: The Ridge Originates from the QCD Correlations Present in the Gluon-Saturated Nucleus Before Collision

Although QCD, being a quantum field theory with pointlike particles and interactions, predicts that long-range correlations are suppressed in elementary collisions, researchers have long known

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

that collective behavior, including unitarity corrections, may lead to long-range rapidity correlations (see, e.g., Reference 95). Long-range rapidity correlations were also proposed to be present in the CGC approach (96, 97) and at the origin of the long-range angular correlations, the ridge (98–100). In all of these proposals, the effect originates in the partonic correlations, which already existed in the wave functions of the colliding objects. In the CGC/glasma picture, the stages immediately after the collision are dominated by strong longitudinal color fields that extend in rapidity by a length $\Delta y \sim 1/\alpha_s$—flux tubes similar to the old Lund string models of multiparticle production (101). This elongated structure was proposed to be at the origin of the ridge after subsequent hydrodynamical evolution takes place. So, the angular structure, strictly speaking, was not discussed in the framework of the CGC in these early papers, in which the role of the CGC was merely to provide longitudinally elongated initial conditions for the hydrodynamical evolution to form the ridge. The discovery of the ridge in high-multiplicity *pp* collisions at the LHC, with a signal much weaker than the one observed at RHIC in Au + Au collisions, indicated that the CGC dynamics could be responsible for at least part of the effect; soon thereafter, it was found that the CGC can indeed generate the characteristic angular structure of the ridge (102) without requiring final-state interactions. The first proposal in this direction was an application of the so-called glasma graphs (98, 102), which correlate particles produced by different flux tubes that are separated by less than the typical color-correlation length, $1/Q_{\rm sat}$, in the transverse plane of the saturated object. Another proposal, or rather a different picture of the same physics, suggested that local anisotropies of the target fields (103) are the source of the ridge structure. Regardless of the actual physical picture, the presence of long-range rapidity correlations with the angular structure of the ridge is very general, and arises naturally, as the saturation of partonic densities cause partons to be color correlated when they are closer than a transverse distance of $1/Q_{\rm sat}$, the color-correlation length, so that two partons from the projectile that happen to be within one of these domains receives a similar p_T kick, $\mathbf{q}$. The result is a correlation between rapidity and azimuthal angle.

In summary, the ridge in *AA* collisions is generally thought to be dominated by final-state hydrodynamical evolution; it seems plausible that the much weaker ridge in high-multiplicity *pp* collisions does not need final-state effects, and that it is purely a reflection of the initial correlations in the partonic wave function of the colliding protons. The main question is whether or not the p + Pb data need the presence of a small thermalized medium. Both the CGC and hydrodynamics can reproduce the strength of the correlation as measured by the experiments, although not all experimental data are equally well reproduced. The question is still open and remains one of the most interesting problems in the phenomenology of hot and dense QCD.

6. PROTON–LEAD COLLISIONS AS A BENCHMARK OF COLD NUCLEAR MATTER EFFECTS

As discussed in Section 1, the characterization of the QCD medium created in *AA* collisions needs control experiments to subtract the cold nuclear matter effects from the genuine hot QCD effects. The main motivation to do experiments involving *pA* collisions was to provide this needed benchmark. In this sense, the *pA* collision program was, to some extent, subsidiary to the *AA* program. Even within this role, the historical relevance of *pA* collisions in benchmarking the probes used to characterize the medium produced in *AA* collisions cannot be overemphasized. Possibly the most important example is the suppression of charmonia, first found in Pb + Pb collisions at the CERN SPS, but later also found, with smaller strength, in simpler *pA* systems. Subtracting these cold nuclear matter effects is essential in order to correctly interpret the data, especially when the effects expected from the hot QCD medium can also be produced, to some

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

extent, with the cold nuclear medium (as in the case of quarkonia suppression, mentioned above). Note that this is not always the case; for example, effects on jets or high-p_T particles do not follow this rule. In this section, we provide a brief overview of the experimental situation, with some discussion of the theoretical implications of the findings for two classes of observables: quarkonia and jets.

6.1. Suppression of Quarkonia States in Proton–Lead Collisions

Quarkonia are expected to be strongly suppressed in the presence of a hot, deconfined medium (104). Moreover, the melting of the different quarkonia states would be stronger for the excited than for the ground states, providing an excellent thermometer with which to measure the medium temperature (105). However, the actual extraction of the medium properties with these probes suffers from the lack of a well-controlled theoretical framework to describe the suppression in pA collisions observed from fixed-target to collider energies. In this situation, the use of phenomenological models has been crucial, and the new data from the LHC p + Pb run (as well as from the d + Au run at RHIC) are further constraining the relevant dynamics (see Reference 106 for a comprehensive review of experimental data in quarkonia and heavy flavor, as well as theoretical interpretations).

The total J/Ψ (107–110) or Υ (111–113) experimental data from the p + Pb collisions at the LHC show a suppression that is similar in size to the one measured at the lower RHIC energies, but with large error bars—a fact that does not contradict the present models at the level of precision of the experimental data. Perhaps the model that best describes the data from fixed-target to LHC energies is the coherent energy loss model (**Figure 10**) proposed by Arleo & Peigné (116). The main component of this model is a coherent energy loss proportional to the energy of the probe, $\Delta E \sim E$, which is a direct consequence of the opening of phase space due to the breaking of coherence between the initial- and final-state QCD radiation.

There are interesting indications of stronger suppression for the excited states than for the ground states, both in the charm sector (J/Ψ versus Ψ') (117, 118) and in the beauty sector (Υ, $\Upsilon' \ldots$) (111–113). Fixed-target data, starting at center-of-mass energies of $\sqrt{s_{NN}} \sim 20$ GeV, do not present sizable differences between the excited and the ground states—the experimental errors are large, however—and the reference pp data are still not available at $\sqrt{s_{NN}} = 5.02$ TeV. The signal seems to be solid at RHIC but is especially strong in the LHC p + Pb collisions. A natural explanation for this difference is the presence of final-state effects. In fact, the only model that can reproduce the experimental data includes the effect of suppression by comoving particles (119), which would have a larger breakup cross section for the larger excited states than for the more compact ground states. These would indeed be very interesting measurements for future, higher-luminosity runs; a proper pp reference should improve the interpretation.

As mentioned above, from the point of view of a reference for the Pb + Pb collisions it would be important to have good control over these effects. For example, the sequential suppression of the excited states has been proposed as a good probe of the properties of the hot QCD medium.

6.2. Effects on Jet Structure and Particle Spectra at High Transverse Momentum

Jet quenching is one of the most important probes of the hot QCD medium created in AA collisions (see, e.g., Reference 27 for a recent review). The basic idea is that the jets produced by the parton shower generated from a highly energetic quark or gluon produced in elementary

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

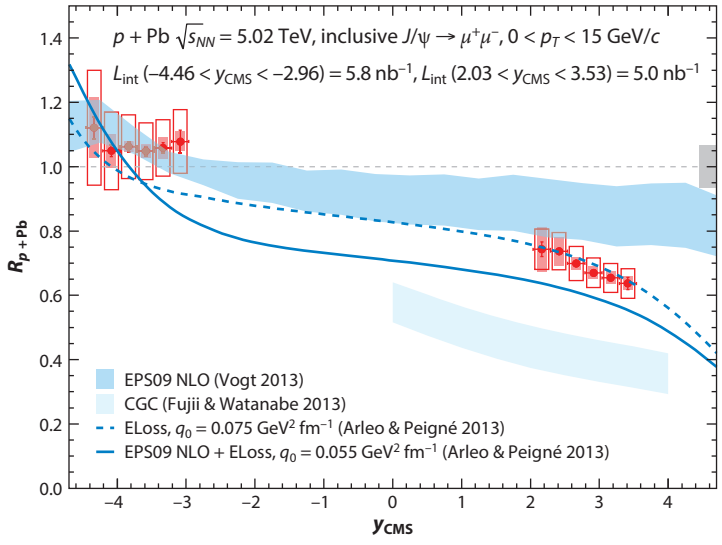

Figure 10

Nuclear modification factor, R_{p+Pb}, for the J/Ψ measured at forward and backward rapidities by the ALICE Collaboration, compared with several different calculations, modifications due to nPDFs, suppression due to the CGC, different combinations of nPDFs, and energy loss. Abbreviations: CGC, Color Glass Condensate; NLO, next-to-leading order; nPDF, nuclear parton distribution function. Modified from Reference 107.

scattering interact with the surrounding QCD matter, losing energy and eventually modifying their internal structure. Measuring these modifications allows one to characterize the properties of the medium produced in the collision. Different observables measured during the p + Pb run include high-p_T inclusive particle (23, 120) and inclusive jet (121–123) nuclear modification factors, (no) modification of the fragmentation functions (124, 125), (no) modification of the jet azimuthal decorrelation (37, 126), and dijet rapidity distributions (37). All these observables can be understood through modification of the nPDFs with respect to those of free protons, as explained in Section 3.1. The agreement between the data and the nPDFs provides a strong cross-check for the factorization hypothesis (Equation 1) at the scales probed. Note, however, that the enhancement found in the nuclear modification function of charged hadrons, R_{p+Pb}, at $p_T > 30$ GeV (120), if confirmed, is impossible to absorb in a universal nPDF. This enhancement, not observed by ALICE, should be clarified with proper *pp* reference data.

As mentioned in Section 3.1, above, some of these observables could eventually be used to further constrain the nPDFs in global analyses. At present, these new constraints are not strong (39), although this situation may improve with the new reference *pp* data or with the higher statistics expected in the p + Pb run of LHC Run 2.

7. CONCLUSIONS

The LHC p + Pb run is profoundly changing our understanding of hot dense matter. Although some of the most striking observations are still not completely understood and competing interpretations are unresolved, novel tools are becoming available. The *pA* programs at the LHC, RHIC, and other, lower-energy facilities provide access to some of the most fundamental questions in hot dense QCD, especially the question of thermalization. It has long been known that

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

high-energy collisions of heavy atomic nuclei produce a hot medium, with temperatures hundreds of thousands times higher than those at the center of the Sun—indeed, temperatures beyond the QCD transition temperature to the quark–gluon plasma. The hydrodynamical description indicates that the details of the underlying microscopic dynamics do not play a fundamental role in the bulk evolution of the system. However, we still do not understand exactly how this medium is formed. Thus, the study of small systems such as *pA* collisions is arguably the best way to study this particular question.

Some of the most important observations from the p + Pb program at LHC Run 1 can be summarized as follows:

1. Universal nPDFs in the usual collinear factorization approach provide a good description of all the available data at large Q^2, where hadronization is not expected to be modified by the surrounding cold nuclear matter (as in the case of quarkonia, for example). More constraints will require better control over the systematic uncertainties, which could be achieved with the new reference ($\sqrt{s} = 5.02$ TeV) *pp* data, recorded in December 2015, and future data at higher energies and luminosities.
2. The soft or semisoft part of the spectrum shares properties with *AA* collisions that are usually associated with the presence of a hot medium. The most important of them refer to the azimuthal angle asymmetries measured with the Fourier decomposition of the correlation functions.
3. Especially important is the presence of a very strong ridge signal. The long range of the correlations in rapidity indicates that this signal should be formed very early during the collisions, close to the time estimated for thermalization. Moreover, the ridge is sensitive to the event-by-event fluctuations of the initial conditions and has, for the first time, provided access to them. Two competing explanations exist: Either the correlations stem uniquely from those present in the wave function of the colliding objects, as in the CGC interpretation, or they are formed during the hydrodynamical evolution of the system. The truth may well be a combination of the two, and measurements could point to a system approaching thermal equilibrium, which is not reached. This would be, in fact, the most interesting situation, as it would indicate the path to the theory of thermalization of hot and dense QCD matter.
4. Finally, interesting features have been observed in the quarkonia sector; for example, the excited states are more suppressed than the ground states. Again, this is an effect expected for a hot medium that also appears in the simpler, smaller systems. The systematics of the quarkonia suppression could give rise to a theory that accommodates all these observations and, eventually, extrapolates them to the *AA* case.

The increase in energy and luminosity during Run 2 will enable further progress in answering these fundamental questions, in studying the energy dependence of the effects, and in accessing new tools to probe the dynamics underlying the building of collectivity in QCD.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

We thank Néstor Armesto for a critical reading of the manuscript. C.A.S. acknowledges support from the European Research Council through grant HotLHC ERC-2011-StG-279579, Ministerio de Ciencia e Innovación of Spain through grant FPA2014-58293-C2-1-P, and Xunta de

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

Galicia (Consellería de Educación). C.A.S. is a member of Strategic Unit AGRUP2015/11. J.P.W. acknowledges support from the Bundesministerium für Bildung und Forschung under contract 05P15PMCA within BMBF-FSP 202.

LITERATURE CITED

1. Armesto N, Scomparin E. arXiv:1511.02151 [nucl-ex] (2015)
2. Müller B. arXiv:1501.06077 [nucl-th] (2015)
3. Salgado CA, et al. *J. Phys. G* 39:015010 (2012)
4. Khachatryan V, et al. (CMS Collab.) *J. High Energy Phys.* 1009:091 (2010)
5. Aad G, et al. (ATLAS Collab.) arXiv:1509.04776 [hep-ex] (2015)
6. Khachatryan V, et al. (CMS Collab.) arXiv:1510.03068 [nucl-ex] (2015)
7. Cunqueiro L, Dias de Deus J, Pajares C. *Eur. Phys. J. C* 65:423 (2010)
8. Bautista I, Cunqueiro L, Dias de Deus J, Pajares C. *J. Phys. G* 37:015103 (2010)
9. d'Enterria D, et al. *Eur. Phys. J. C* 66:173 (2010)
10. Prasad SK, Roy V, Chattopadhyay S, Chaudhuri AK. *Phys. Rev. C* 82:024909 (2010)
11. Bozek P. *Acta Phys. Polon. B* 41:837 (2010)
12. Casalderrey-Solana J, Wiedemann UA. *Phys. Rev. Lett.* 104:102301 (2010)
13. Bozek P, Broniowski W. *Phys. Lett. B* 718:1557 (2013)
14. Kurkela A. arXiv:1601.03283 [hep-ph] (2016)
15. Altarelli G, Parisi G. *Nucl. Phys. B* 126:298 (1977)
16. Gribov VN, Lipatov LN. *Sov. J. Nucl. Phys.* 15:438 (1972)
17. Dokshitzer YL. *Sov. Phys. JETP* 46:641 (1977)
18. Gelis F, Iancu E, Jalilian-Marian J, Venugopalan R. *Annu. Rev. Nucl. Part. Sci.* 60:463 (2010)
19. Jowett JM, et al. In *Proceedings of the 4th International Particle Accelerator Conference* (*IPAC13*), ed. Z Dai, C Petit-Jean-Genaz, VRW Schaa, C Zhang, 1:49. Geneva: JACoW.org. **http://accelconf.web.cern.ch/AccelConf/IPAC2013/papers/moodb201.pdf** (2013)
20. Chatrchyan S, et al. (CMS Collab.) *Phys. Lett. B* 718:795 (2013)
21. Abelev BB, et al. (ALICE Collab.) *Phys. Lett. B* 719:29 (2013)
22. Aad G, et al. (ATLAS Collab.) *Phys. Rev. Lett.* 110:182302 (2013)
23. Abelev BB, et al. (ALICE Collab.) *Phys. Rev. Lett.* 110:082302 (2013)
24. Abelev BB, et al. (ALICE Collab.) *J. Instrum.* 9:P11003 (2014)
25. Albacete JL, Dumitru A, Marquet C. *Int. J. Mod. Phys. A* 28:1340010 (2013)
26. Armesto N. *J. Phys. G* 32:R367 (2006)
27. Mehtar-Tani Y, Milhano JG, Tywoniuk K. *Int. J. Mod. Phys. A* 28:1340013 (2013)
28. Eskola KJ, Paukkunen H, Salgado CA. *J. High Energy Phys.* 0904:065 (2009)
29. de Florian D, Sassot R, Zurita P, Stratmann M. *Phys. Rev. D* 85:074028 (2012)
30. Kovarik K, et al. arXiv:1509.00792 [hep-ph] (2015)
31. Hirai M, Kumano S, Nagai T-H. *Phys. Rev. C* 76:065207 (2007)
32. Giele WT, Keller S. *Phys. Rev. D* 58:094023 (1998)
33. Ball RD, et al. (NNPDF Collab.) *Nucl. Phys. B* 849:112 (2011); Ball RD, et al. (NNPDF Collab.) *Nucl. Phys. B* 854:926 (2012); Ball RD, et al. (NNPDF Collab.) *Nucl. Phys. B* 855:927 (2012)
34. Ball RD, et al. (NNPDF Collab.) *Nucl. Phys. B* 855:608 (2012)
35. Watt G, Thorne RS. *J. High Energy Phys.* 1208:052 (2012)
36. Paukkunen H, Zurita P. *J. High Energy Phys.* 1412:100 (2014)
37. Chatrchyan S, et al. (CMS Collab.) *Eur. Phys. J. C* 74:2951 (2014)
38. Paukkunen H, Eskola KJ, Salgado C. *Nucl. Phys. A* 931:331 (2014)
39. Armesto N, et al. arXiv:1512.01528 [hep-ph] (2015)
40. Gribov LV, Levin EM, Ryskin MG. *Phys. Rep.* 100:1 (1983)
41. McLerran LD, Venugopalan R. *Phys. Rev. D* 49:2233 (1994)
42. McLerran LD, Venugopalan R. *Phys. Rev. D* 49:3352 (1994)
43. Balitsky I. *Nucl. Phys. B* 463:99 (1996)

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

44. Jalilian-Marian J, Kovner A, Leonidov A, Weigert H. *Phys. Rev. D* 59:014014 (1998)
45. Jalilian-Marian J, Kovner A, Leonidov A, Weigert H. *Nucl. Phys. B* 504:415 (1997)
46. Kovner A, Milhano JG, Weigert H. *Phys. Rev. D* 62:114005 (2000)
47. Iancu E, Leonidov A, McLerran LD. *Nucl. Phys. A* 692:583 (2001)
48. Ferreiro E, Iancu E, Leonidov A, McLerran L. *Nucl. Phys. A* 703:489 (2002)
49. Kovchegov YV. *Phys. Rev. D* 60:034008 (1999)
50. Golec-Biernat KJ, Wusthoff M. *Phys. Rev. D* 59:014017 (1998)
51. Balitsky I, Chirilli GA. *Phys. Rev. D* 77:014019 (2008)
52. Balitsky I, Chirilli GA. *Phys. Rev. D* 88:111501 (2013)
53. Kovner A, Lublinsky M, Mulian Y. *Phys. Rev. D* 89:061704 (2014)
54. Albacete JL, et al. *Eur. Phys. J. C* 71:1705 (2011)
55. Chirilli GA, Xiao BW, Yuan F. *Phys. Rev. D* 86:054005 (2012)
56. Altinoluk T, et al. *Phys. Rev. D* 91:094016 (2015)
57. Watanabe K, Xiao BW, Yuan F, Zaslavsky D. *Phys. Rev. D* 92:034026 (2015)
58. Lappi T, Mäntysaari H. *Phys. Rev. D* 91:074016 (2015)
59. Lappi T, Mäntysaari H. arXiv:1601.06598 [hep-ph] (2016)
60. Motyka L, Stasto AM. *Phys. Rev. D* 79:085016 (2009)
61. Beuf G. *Phys. Rev. D* 89:074039 (2014)
62. Iancu E, et al. *Phys. Lett. B* 750:643 (2015)
63. Iancu E, et al. *Phys. Lett. B* 744:293 (2015)
64. Albacete JL, Dumitru A, Fujii H, Nara Y. *Nucl. Phys. A* 897:1 (2013)
65. Barnafoldi GG, et al. *Phys. Rev. C* 95:024903 (2011)
66. Xu R, Deng WT, Wang XN. *Phys. Rev. C* 86:051901 (2012)
67. Abelev B, et al. (ALICE Collab.) *Phys. Rev. Lett.* 110:032301 (2013)
68. Tribedy P, Venugopalan R. *Phys. Lett. B* 701:125 (2011); Tribedy P, Venugopalan R. Erratum. *Phys. Lett. B* 718:1154 (2013)
69. Toia A. (ALICE Collab.) *Nucl. Phys. A* 931:315 (2014)
70. Adare A, et al. (PHENIX Collab.) *Phys. Rev. C* 90:034902 (2014)
71. Floris M. *Nucl. Phys. A* 931:103 (2014)
72. Abelev BB, et al. (ALICE Collab.) *Phys. Lett. B* 728:25 (2014)
73. Schnedermann E, Sollfrank J, Heinz UW. *Phys. Rev. C* 48:2462 (1993)
74. Roesler S, Engel R, Ranft J. In *Advanced Monte Carlo for Radiation Physics, Particle Transport Simulation and Applications: Proceedings of the Monte Carlo 2000 Conference*, ed. A Kling, et al., p. 1033. Berlin: Springer (2000)
75. Bozek P. *Phys. Rev. C* 85:014911 (2012)
76. Pierog T, et al. *Phys. Rev. C* 92:034906 (2015)
77. Dumitru A, Kharzeev DE, Levin EM, Nara Y. *Phys. Rev. C* 85:044920 (2012)
78. Tribedy P, Venugopalan R. *Phys. Lett. B* 710:125 (2012)
79. Adams J, et al. (STAR Collab.) *Phys. Rev. Lett.* 95:152301 (2005)
80. Abelev BI, et al. (STAR Collab.) *Phys. Rev. C* 80:064912 (2009)
81. Alver B, et al. (PHOBOS Collab.) *Phys. Rev. Lett.* 104:062301 (2010)
82. Khachatryan V, et al. (CMS Collab.) *Phys. Rev. Lett.* 115:012301 (2015)
83. Aad G, et al. (ATLAS Collab.) *Phys. Rev. C* 90:044906 (2014)
84. Abelev BB, et al. (ALICE Collab.) *Phys. Rev. C* 90:054901 (2014)
85. Aaij R, et al. (LHCb Collab.) arXiv:1512.00439 [nucl-ex] (2015)
86. Luzum M. *Phys. Lett. B* 696:499 (2011)
87. Gale C, Jeon S, Schenke B. *Int. J. Mod. Phys. A* 28:1340011 (2013)
88. Bhalerao RS, Borghini N, Ollitrault JY. *Nucl. Phys. A* 727:373 (2003)
89. Strickland M. *Acta Phys. Polon. B* 45:2355 (2014)
90. Bzdak A, Schenke B, Tribedy P, Venugopalan R. *Phys. Rev. C* 87:064906 (2013)
91. Qin GY, Müller B. *Phys. Rev. C* 89:044902 (2014)
92. Werner K, et al. *Phys. Rev. Lett.* 112:232301 (2014)
93. Bozek P, Broniowski W, Torrieri G. *Phys. Rev. Lett.* 111:172303 (2013)

94. Kozlov I, et al. arXiv:1405.3976 [nucl-th] (2014)
95. Capella A, Krzywicki A. *Phys. Rev. D* 18:4120 (1978)
96. Kovchegov YV, Levin E, McLerran LD. *Phys. Rev. C* 63:024903 (2001)
97. Armesto N, McLerran LD, Pajares C. *Nucl. Phys. A* 781:201 (2007)
98. Dumitru A, Gelis F, McLerran L, Venugopalan R. *Nucl. Phys. A* 810:91 (2008)
99. Gavin S, McLerran L, Moschelli G. *Phys. Rev. C* 79:051902 (2009)
100. Gelis F, Lappi T, Venugopalan R. *Phys. Rev. D* 79:094017 (2009)
101. Andersson B, Gustafson G, Ingelman G, Sjöstrand T. *Phys. Rep.* 97:31 (1983)
102. Dumitru A, et al. *Phys. Lett. B* 697:21 (2011)
103. Kovner A, Lublinsky M. *Phys. Rev. D* 83:034017 (2011)
104. Matsui T, Satz H. *Phys. Lett. B* 178:416 (1986)
105. Digal S, Petreczky P, Satz H. *Phys. Rev. D* 64:094015 (2001)
106. Andronic A, et al. arXiv:1506.03981 [nucl-ex] (2015)
107. Abelev BB, et al. (ALICE Collab.) *J. High Energy Phys.* 1402:073 (2014)
108. Adam J, et al. (ALICE Collab.) *J. High Energy Phys.* 1506:055 (2015)
109. Aad G, et al. (ATLAS Collab.) *Phys. Rev. C* 92:034904 (2015)
110. Aaij R, et al. (LHCb Collab.) *J. High Energy Phys.* 1402:072 (2014)
111. Chatrchyan S, et al. (CMS Collab.) *J. High Energy Phys.* 1404:103 (2014)
112. Abelev BB, et al. (ALICE Collab.) *Phys. Lett. B* 740:105 (2015)
113. Aaij R, et al. (LHCb Collab.) *J. High Energy Phys.* 1407:094 (2014)
114. Vogt R. *Int. J. Mod. Phys. E* 22:1330007 (2013)
115. Fujii H, Watanabe K. *Nucl. Phys. A* 915:1 (2013)
116. Arleo F, Peigné S. *J. High Energy Phys.* 1303:122 (2013)
117. Aaij R, et al. (LHCb Collab.) arXiv:1601.07878 [nucl-ex] (2016)
118. Abelev BB, et al. (ALICE Collab.) *J. High Energy Phys.* 1412:073 (2014)
119. Ferreiro EG. *Phys. Lett. B* 749:98 (2015)
120. Khachatryan V, et al. (CMS Collab.) *Eur. Phys. J. C* 75:237 (2015)
121. Aad G, et al. (ATLAS Collab.) *Phys. Lett. B* 748:392 (2015)
122. Adam J, et al. (ALICE Collab.) *Phys. Lett. B* 749:68 (2015)
123. Khachatryan V, et al. (CMS Collab.) arXiv:1601.02001 [nucl-ex] (2016)
124. CMS Collab. *CMS-PAS-HIN-15-004* (2015)
125. ATLAS Collab. *ATLAS-CONF-2015-022* (2015)
126. Adam J, et al. (ALICE Collab.) *Phys. Lett. B* 746:385 (2015)

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

Annual Review of Nuclear and Particle Science

Volume 66, 2016

Contents

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.

Errata

An online log of corrections to *Annual Review of Nuclear and Particle Science* articles may be found at http://www.annualreviews.org/errata/nucl

Annu. Rev. Nucl. Part. Sci. 2016.66:449-473. Downloaded from www.annualreviews.org
Access provided by CERN Library on 02/14/20. For personal use only.