ANNUAL REVIEWS

Annual Review of Nuclear and Particle Science

Heavy Ion Collisions: The Big Picture and the Big Questions

Wit Busza,[1] Krishna Rajagopal,[1,2] and Wilke van der Schee[2,3]

[1]Laboratory for Nuclear Science and Department of Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA

[2]Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA

[3]Institute for Theoretical Physics and Center for Extreme Matter and Emergent Phenomena, Utrecht University, 3584 CE Utrecht, Netherlands

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

Annu. Rev. Nucl. Part. Sci. 2018. 68:339–76

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

https://doi.org/10.1146/annurev-nucl-101917-020852

Copyright © 2018 by Annual Reviews.
All rights reserved

ANNUAL REVIEWS CONNECT
www.annualreviews.org
- Download figures
- Navigate cited references
- Keyword search
- Explore related articles
- Share via email or social media

Keywords

quark–gluon plasma, heavy ion collisions, relativistic hydrodynamics, jets, multiparticle production

Abstract

Heavy ion collisions quickly form a droplet of quark–gluon plasma (QGP) with a remarkably small viscosity. We give an accessible introduction to how to study this smallest and hottest droplet of liquid made on Earth and why it is so interesting. The physics of heavy ion collisions ranges from highly energetic quarks and gluons described by perturbative QCD to a bath of strongly interacting gluons at lower energy scales. These gluons quickly thermalize and form QGP, while the energetic partons traverse this plasma and end in a shower of particles called jets. Analyzing the final particles in various ways allows us to study the properties of QGP and the complex dynamics of multiscale processes in QCD that govern its formation and evolution, providing what is perhaps the simplest form of complex quantum matter that we know of. Much remains to be understood, and throughout the review big open questions are encountered.

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

Contents

1. INTRODUCTION

In the past 50 years, as beams of ultrarelativistic protons and nuclei have become available, the collisions of protons with nuclei and nuclei with nuclei, at higher and higher relative velocities (or total collision energy), have been studied in greater and greater detail. This review provides some answers to the questions "Why do such studies?" and "What have we learned from them so far?" and "What are the big questions that they may illuminate in the future?"

We start with a qualitative description, in the center-of-mass frame (the lab frame at a collider), of the sequence of events that occur when two ultrarelativistic nuclei collide head on. This picture follows from the observed phenomenology (summarized in Section 3), relativity, and our understanding of the workings of QCD.[1]

Each incident nucleus is a Lorentz-contracted disc. For large nuclei such as Pb or Au, the diameter of the disc is about 14 fm (short for femtometer, or Fermi) and its thickness is about $14/\gamma$ where, at the highest beam energies attainable at the Relativistic Heavy Ion Collider (RHIC) and the Large Hadron Collider (LHC), the relativistic γ factors are approximately 100 and 2,500, respectively, corresponding to beam rapidities of $y = 5.3$ and 8.5. Each disc includes many colored quarks and antiquarks, with three more quarks than antiquarks per nucleon in the incident nuclei and with $q\bar{q}$ pairs coming from quantum fluctuations in the initial-state wave functions that are "almost real," as a consequence of time dilation. These quarks and antiquarks are, in turn, sources of strong, almost completely transverse, color fields and corresponding field quanta, the gluons, which also carry color.

Momentum rapidity y: $\cosh(y) \equiv \gamma$, with $\gamma = 1/\sqrt{1 - v_z^2}$, with v_z the velocity along the beam direction in units of the speed of light

Space-time rapidity y_s: $\tanh(y_s) \equiv z/t$, where z and t are space-time coordinates centered at the collision. y_s cannot be measured experimentally; however, in a boost-invariant flow in which $v_z = z/t$, $y_s = y$

[1] We recommend the books in References 1–3 for the reader interested in a more comprehensive introduction to heavy ion collisions. Note also that because we are required to limit the number of citations, throughout this review where possible we cite articles in which additional citations can be found.

The area density of quarks, antiquarks, and gluons (partons for short, in the language of Feynman) increases with the velocity of the nuclei. It is not uniform across the area of the disc, and fluctuates from nucleus to nucleus. The spatial variation of the partons primarily reflects the instantaneous distribution of the nucleons inside the nuclei and of the partons inside the nucleons. Overall, the incident nuclei are highly complex systems of partons with a longitudinal momentum distribution (referred to as a structure function) that is close to being a superposition of that in the individual nucleons but with small modifications coming from the proximity, and motion, of nucleons in nuclei.

When the two discs, each a tiny fraction of a femtometer thick, overlap or collide, most of the incident partons lose some energy but are not deflected by any large angle. Most of these interactions are soft, meaning that they involve little transverse momentum (p_T) transfer. These strong interactions can be described in terms of interacting fields or slabs of energy. In the language of fields and particles, as the two discs of strongly interacting transverse color fields and associated color charges collide, some color charge exchange occurs between the discs, and longitudinal color fields are produced, which fill the space between the two receding discs, reducing the energy in the discs themselves, and then gradually decay into $q\bar{q}$ pairs and gluons. A small fraction of the incident partons suffer hard perturbative interactions as the discs overlap initially, which, as we discuss further below, leads to a relatively improbable but very important production of particles with high p_T.

In high-energy heavy ion collisions, the maximum energy density occurs just as the two highly Lorentz-contracted nuclei collide. Clearly this system is very far from equilibrium, and its very high energy density is really just a consequence of Lorentz contraction. It is much more interesting to ask what we can say in a generic way about the average energy density, say, 1 fm/c after the collision, by which time the two discs are 2 fm apart. The expanding high-energy-density system produced around the midpoint between the two discs, where the collision occurred, has an energy density at that time that is still far in excess of 500 MeV fm^{-3}, the energy density inside a typical hadron. A rough estimate can be obtained from the available data for head-on LHC collisions with $\sqrt{s_{NN}} = 2.76$ TeV (corresponding to $\gamma = 1{,}400$ and $y = 8.0$) by noting that the total transverse energy in particles with pseudorapidity between -0.5 and 0.5 (so longitudinal velocity $-0.46 < v < 0.46$) is measured to be 1.65 ± 0.1 TeV (4), meaning that the average energy density at 1 fm/c after the collision is greater than 1.65 TeV/$[\pi(7\text{ fm})^2(0.92\text{ fm})] = 12$ GeV fm^{-3}, about 20 times the energy density of a hadron. The entropy produced in these collisions is also enormous; to get a sense of its size, note that before the collision the entropy of the two incident nuclei is essentially zero, whereas the final state after the collision can contain as many as 30,000 particles, and hence has a very large entropy. We return to this topic in Section 3, below, and in particular we see that most of this entropy is produced quickly, in the initial moments after the collision.

To get a further sense of the magnitude of the average energy density 1 fm/c after the collision, note that, as discussed in Section 3, lattice calculations of QCD thermodynamics show that matter in thermal equilibrium at a temperature of 300 MeV has an energy density $\approx 12\,T^4 = 12.7$ GeV/fm^3. Thus, the quarks and gluons produced in the collision cannot be described as a collection of distinct individual hadrons. Nevertheless, as discussed in subsequent sections, the quarks and gluons in this high-energy-density matter are far from independent. They are so strongly coupled to one another that they form a collective medium that expands and flows as a relativistic hydrodynamic fluid with a remarkably low ratio of viscosity to entropy density, $\eta/s \approx 1/4\pi$ (5, 6), in units with $\hbar = k_B = 1$, within a time that can be shorter than or of order 1 fm/c in the rest frame of the fluid. This form of matter has been named quark–gluon plasma, or QGP for short. Even if the transverse velocity of the fluid is small initially, say, 1 fm/c after the collision, the pressure-driven hydrodynamic expansion rapidly builds up transverse velocities of

$\sqrt{s_{NN}}$: the total collision energy per nucleon–nucleon pair in the center-of-mass frame

Pseudorapidity η: $\equiv -\log[\tan(\theta/2)]$, where θ is the polar angle in momentum space relative to the beam direction. η is a standard proxy for rapidity y because $\eta = y$ for massless particles, and should not be confused with shear viscosity, which is also written as η

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

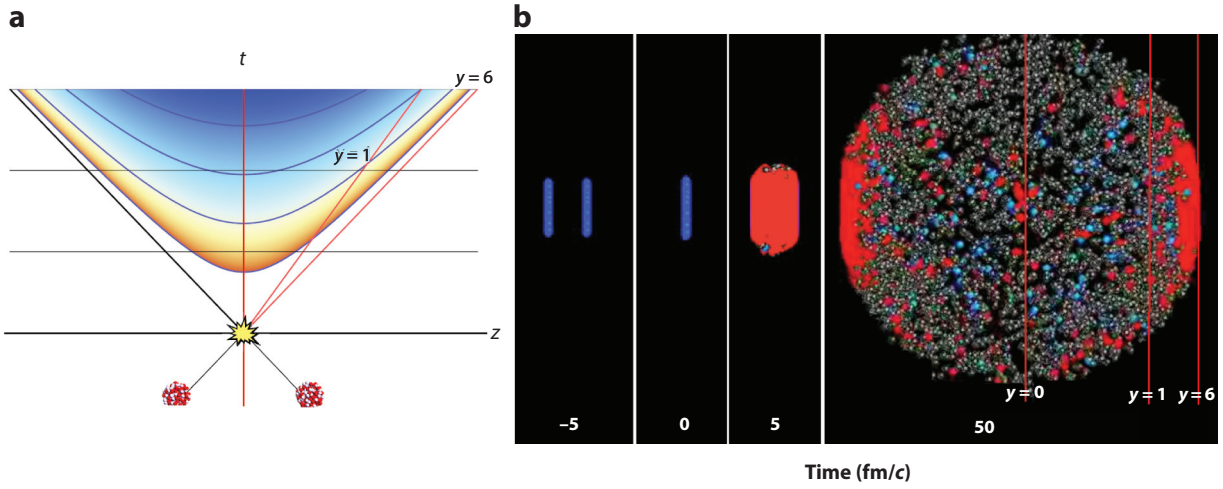

Figure 1

(*a*) Space-time picture of a heavy ion collision, where the color indicates the temperature of the plasma formed. Dynamics takes place as a function of proper time (*blue curves*), which is why plasma forms later at higher rapidities. (*b*) Snapshots of a central 2.76 TeV PbPb collision at different times (different horizontal slices of the space-time picture on the left) with hadrons (*blue and gray spheres*) as well as quark–gluon plasma (*red*) (see **http://web.mit.edu/mithig/movies/LHCanmation.mov**). In both panels, at a given time the hottest regions can be found at high rapidity close to the outgoing remnants of the nuclei, and the red lines indicate the approximate longitudinal location of particles with rapidity $y = 0$, $y = 1$, and $y = 6$. Panel *a* adapted from Reference 7. Panel *b* adapted from Reference 8.

order half the speed of light. As the discs recede from one another and the QGP produced between them is expanding and cooling, new QGP is continually forming in the wake of each receding disc; see **Figure 1**. This happens because the quarks and gluons produced at high rapidity are moving at almost the speed of light in one of the beam directions, meaning that when enough time has passed in their frame for them to form QGP a long time has passed in the lab frame, around 330 fm/c for rapidity $y = 6.5$. Throughout this QGP production process, each disc gradually loses energy as partons with higher and higher rapidity separate from it and form QGP. In contrast, the occasional high-p_T particles seen in some collisions are produced by large-angle scattering at very early times, when the incident nuclei collide.

The process ends once QGP has formed at the rapidities where most of the baryon number from the incident nuclei ends up, which is expected to be about two units of rapidity less than that of the incident nuclei, based upon measurements made in low-energy proton–nucleus (pA) collisions (9). So, the discs lose about 85% of their energy while varying amounts of QGP form at varying rapidities over a range that extends between $y = -6.5$ and $y = 6.5$ in collisions at the LHC. A good way to visualize the QGP production process described above is to consider the production of each volume element of QGP in its own local rest frame, where the two colliding nuclei have an asymmetric rapidity and energy, and then boost this volume of QGP back to the lab frame.

After production, each elemental volume of QGP expands in all directions. Looked at overall, the droplet of fluid flows hydrodynamically, as its initial high pressure drives fluid motion, expansion, and consequent cooling. This picture holds until the energy density at a given location in the fluid drops below that within an individual hadron, at which point the fluid falls apart into a mist of hadrons that scatter off one another a few times and then stream away freely. This mechanism of particle production, via an intermediate epoch during which a hydrodynamic fluid forms and expands, is quite different from the current understanding of particle production in elementary collisions in which only a few new particles are created.

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

Meanwhile, remnants of the original nuclei [excited nuclear matter compressed by a factor of about 5–10 (10)] progress in the forward and backward directions. This high-baryon-density system then expands and hydrodynamizes, forming hot quark-rich QGP after a time of order 1 fm/c in its own rest frame, corresponding to a time of approximately 330 fm/c in the lab frame. After further expansion, it falls apart into hadrons. Unfortunately, none of the LHC detectors are adequately instrumented around $y = 6.5$, which would be an almost impossible task; as a result, the debris formed from this hot, high baryon density, QGP, has not yet been studied.

So far we have considered head-on ("central") collisions. What about noncentral collisions? In the overlap region, the process is the same as described above, except that the droplet of QGP is formed with an initial approximately lenticular shape in the transverse plane. In reality, because nuclei are made of individual nucleons, the energy density of the QGP that forms is lumpy in the transverse plane, making it neither perfectly circular in head-on collisions nor perfectly lenticular in noncentral collisions. Deviations from circular symmetry in the initial shape of the QGP, whether due to off-center collisions or the lumpiness and fluctuations of the incident nuclei, result in anisotropies in the pressure of the hydrodynamic fluid, which in turn drive anisotropies in the expansion velocity and hence in the azimuthal momentum distribution of the finally produced particles.

In an off-center (noncentral) collision, the parts of the incident nuclei that do not collide are referred to as spectators. At very early times, they create a magnetic field in the collision zone whose possible effects are the subject of much discussion that we do not have space to review (11). Later, they fragment into excited nuclei and hadrons, moving with almost the full rapidity of the incident nuclei. The extreme limit of an off-center collision is one in which the nuclei themselves miss one another but the Lorentz-contracted discs of electromagnetic fields around them do interact. These ultraperipheral collisions give rise to copious $\gamma\gamma$ and γA interactions (12), which we do not pursue except to note here that they dominate the total nucleus–nucleus interaction cross section.

Finally, we briefly discuss the hard collisions between two partons in the incident nuclei. Such collisions, especially where particles with very large transverse momenta (say, greater than a few tens of GeV/c) are produced, are rare but very important. They lead to the production, essentially from a point at the earliest times of the overall collision process, of high-energy parton pairs and electroweak bosons. The high-energy partons evolve, decay, radiate, and finally produce a cone-shaped spray or "jet" of hadrons and/or high-energy photons, leptons, or heavy $Q\bar{Q}$ pairs, all while traversing a region where QGP is in the process of being produced and evolving. They thus contain a wealth of information about the produced medium (in essence, they "X-ray" the medium) and on how partons lose energy or disturb the medium as they interact with it. We return to this topic in Section 7.

As is evident from the above description, collisions of ultrarelativistic nuclei are complex, consisting of several distinct stages, each probing different aspects of QCD. What makes them interesting is that the regimes of QCD that they give us a means to explore are places where, because of the strength of the QCD interactions, we would not even have a zeroth-order understanding of the properties of the matter that QCD describes, let alone the dynamical phenomena, without having seen what happens in these collisions. Heavy ion collisions are a laboratory that is rich with unique ways to probe fundamental aspects of QCD empirically, with some control over varying conditions. Our description of a relativistic heavy ion collision brings us directly to several questions: Do we have, at least qualitatively, an understanding of all the stages of a heavy ion collision? Have any fundamentally new phenomena been seen? Any unexplained phenomena? What new insights have we obtained, or can we obtain, about the workings of QCD from analysis of heavy ion collision data and corresponding theoretical calculations? How, and how well, do we

Central and off-central: head-on collisions between ions are referred to as central collisions, whereas collisions between ions that only partially overlap are called noncentral or peripheral

Particle nomenclature: collisions are often described with the notation p, A, γ, π, and K for protons, heavy ions, photons, pions, and kaons, respectively

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

understand the initial stages of the collision process, up to the creation of QGP? Which aspects involve weakly coupled dynamics, and which strongly coupled? What are the properties of QGP? From the study of jets and high momentum particles, what have we learned about the properties of strongly interacting matter, and about the dynamics of fast particles as they traverse strongly coupled matter? What new insights have we obtained about the formation of hadrons?

Beginning in Section 3, we attempt to answer some of these questions. But first, in the next section we expand our perspective.

2. WHY DO WE STUDY ULTRARELATIVISTIC HEAVY ION COLLISIONS?

The overarching answer with which we ended Section 1, albeit formulated as a suite of questions, is that studying ultrarelativistic heavy ion collisions may give us a more complete understanding of how particles are produced in high-energy collisions in QCD. This is a fundamental question that, in fact, long predates QCD: Heisenberg (13) and Heitler and colleagues (14) wrestled with it in the 1930s and 1940s, Fermi (15) and Landau (2, 16) did so in the 1940s and 1950s, and Feynman (17) tried his hand in the 1960s. We can now gain new purchase on these old questions by studying high-energy collisions in a new regime in which experimenters have new knobs to dial, including the size of each of the colliding nuclei, (proxies for) the impact parameter, the final-state multiplicity, and more.

In this section, we formulate variants of the "Why are we doing this?" question that take us beyond the subject of ultrarelativistic heavy ion collisions per se. Are there insights that we hope to gain from studying these collisions that go beyond understanding the dynamics of these collisions, or even of ultrarelativistic collisions in QCD more generally? The affirmative answers to this question, which we divide into three groups below, motivate much of the experimental and theoretical effort that we describe in this review.

2.1. QCD in Cosmology

Heavy ion collisions recreate droplets of the matter that filled the Universe a microsecond or so after the Big Bang. And it has been understood since the mid 1970s (23, 24) that when the Universe was only a few microseconds old it was filled with matter at temperatures above $\Lambda_{\rm QCD}$ (the fundamental energy scale in QCD, of order a few hundred MeV, which is best thought of as the inverse of the size of a hadron in QCD) and was too hot for protons, neutrons, or any hadrons to have formed. This direct and tangible connection to the earliest moments of the Universe, together with the insight that the primordial matter found at these temperatures had to be some new form of matter not made of hadrons, provides two powerful motivations for studying ultrarelativistic heavy ion collisions. Historically, these were the motivations underlying much of the initial impetus for the field.

In the 1980s, a fair amount of work was done on possible observable consequences in cosmology of a first-order phase transition between hot primordial matter and ordinary hadronic matter. These all relied upon presuming a strong first-order phase transition that occurred via the nucleation of widely separated bubbles of the low-temperature hadronic phase. As the walls of these putative bubbles plowed through the microseconds-old Universe over distances as long as centimeters or meters, they would have left behind matter that was inhomogeneous over these long length scales (25, 26). If this had happened, it would have modified the synthesis of light nuclei that occurred when the Universe was minutes old. Starting in the late 1990s, and culminating in classic work in the 2000s (27, 28), it became clear from first-principles lattice QCD calculations of the pressure and energy density of hot QCD matter containing equal densities of quarks

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

and antiquarks that the transition from primordial hot QCD matter to hadronic matter in the first few microseconds after the Big Bang proceeded via a continuous crossover, not a first-order phase transition. This picture is consistent with the modern understanding of Big Bang nucleosynthesis, which is in accord with cosmological observations without any of the disruption that a strong first-order phase transition would have introduced (29). A continuous crossover does not introduce any fluctuations on length scales much longer than the femtometer-scale natural length scales of QCD, meaning that it left no imprint in the microseconds-old Universe that survived so as to be visible in some way today. That is, we now understand that we cannot use cosmological observations to "see" the primordial hot QCD matter that filled the microseconds-old Universe, or the crossover transition at which ordinary protons and neutrons first formed.

A central goal of ultrarelativistic heavy ion collisions, then, is to use these experiments to recreate droplets of Big Bang matter in the laboratory—where we can learn about its material properties as well as about its phase diagram in ways that we will never be able to do via observations made with telescopes or satellites. What can we learn from such studies? What have we learned so far? One of the most important discoveries made via studying ultrarelativistic heavy ion collisions is that matter that is a few trillions of degrees hot is a liquid. The early ideas that motivated the field turned out to be half right: Primordial matter at these temperatures is not made of hadrons, as anticipated; however, at the temperatures that have been achieved in heavy ion collisions to date, it is not a weakly coupled plasma of quarks and gluons, as originally expected. Instead, when the hadrons that make up ordinary nuclei are heated to these extraordinary temperatures, the resulting matter is better thought of as a soup of quarks and gluons, in which there are no hadrons to be found but in which every quark and gluon is always strongly coupled to its neighbors, with no quasiparticles that can travel long distances between discrete scatterings. We describe how this insight was obtained in Section 4.

The material property that quantifies the liquidness of a liquid made up of ultrarelativistic constituents is the ratio of its shear viscosity η to its entropy density s. The ratio η/s is dimensionless in units in which $\hbar$ and k_B have been set to one. This ratio plays a central role in the equations of hydrodynamics, where it governs the amount of entropy produced within the fluid as a sound wave propagates through it, or more generally as it flows in any nontrivial way. It is the natural dimensionless measure of the effects of shear viscosity in a relativistic fluid, and we refer to it as the "specific viscosity." In Section 4 we sketch how the combination of data from ultrarelativistic collisions and hydrodynamic calculations is being used to constrain η/s, and even its temperature dependence. We also see that η/s for the liquid of quarks and gluons produced in heavy ion collisions is close to the value $1/4\pi$. Although, because of its extraordinarily high temperature, this liquid has extraordinarily large values of both η and s relative to those of any quotidian fluid, its specific viscosity η/s is smaller than that of any other known fluid. Interestingly, $1/4\pi$ is the value of the ratio η/s in the plasma of infinitely strongly coupled gauge theories that are cousins of QCD that have a dual gravitational description in terms of a black hole horizon in 4+1-dimensional anti–de Sitter space (20, 30), a horizon whose undulations encode the hydrodynamic motion of the plasma (31). This connection between the properties of the primordial matter recreated in heavy ion collisions, via a duality first discovered in string theory, and properties of black hole horizons certainly provides strong motivation for pushing the determination of η/s, and recently also the bulk viscosity, of QGP to greater accuracy.

The zeroth-order input that a hydrodynamic calculation needs to get from the microscopic theory of whatever hydrodynamic fluid it is seeking to describe is the equation of state, relating the pressure of the fluid to its energy density. In QCD, the equation of state, and any other thermodynamic property of a static volume of QGP in thermal equilibrium, can be reliably calculated via implementing the standard path integral formulation of thermodynamics on a discretized lattice

Shear viscosity: the larger the shear viscosity η the more easily momentum can be exchanged between distant fluid cells and, consequently, the faster a gradient in fluid velocity (or a sound wave) dissipates into heat

Specific viscosity: the ratio of shear viscosity to entropy density, η/s; the natural dimensionless measure of the effects of shear viscosity in a relativistic fluid, where lower values correspond to more liquid fluids

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

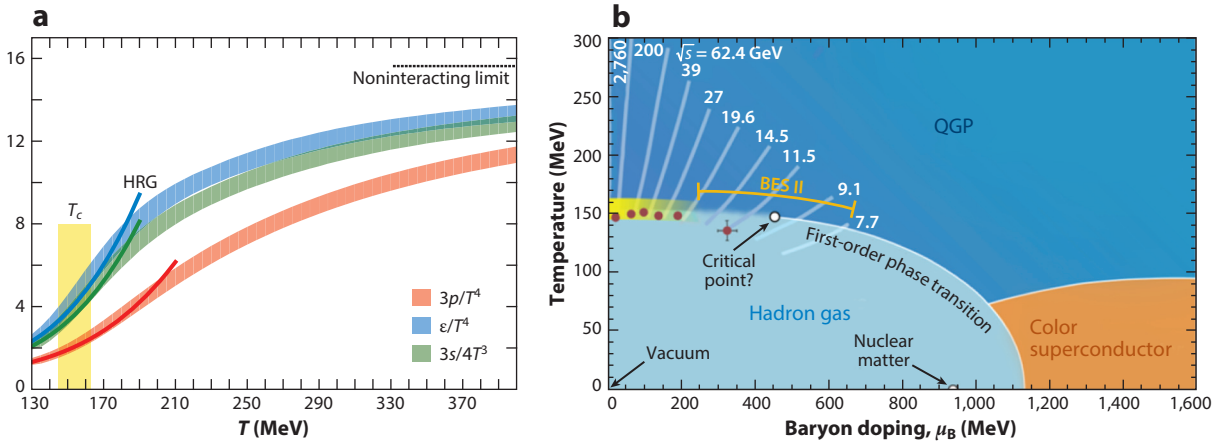

Figure 2

(*a*) Lattice QCD calculations (*colored bands*) of the pressure p, energy density ε, and entropy density s of hot QCD matter in thermal equilibrium at temperature T (18, 19) show a continuous crossover around $T \sim 150$ MeV, from a hadron resonance gas (HRG; *colored lines*) at lower temperatures to quark–gluon plasma (QGP) at higher temperatures. Because QCD is asymptotically free, thermodynamic quantities will reach the Stefan–Boltzmann limit (weakly coupled quarks and gluons; the noninteracting limit is indicated) at extremely high temperature. At the range shown, however, they are around 20% below their Stefan–Boltzmann values, which is consistent with simple estimates for strongly coupled plasma based on holography (20). The rise in ε/T^4 and s/T^3 shown in the figure is a direct manifestation of the crossover from a hadron gas to QGP, which has more thermodynamic degrees of freedom because color is deconfined. Using experimental data to constrain ε/T^4 remains an outstanding challenge: Comparing hydrodynamic calculations to various experimental measurements gives us information about ε versus time, but although some information about T can be obtained by analyzing measurements of photons, electrons, and muons from heavy ion collisions (21), at present T cannot be determined with sufficient accuracy to constrain ε/T^4 well enough to see the rise in the number of degrees of freedom in QGP. (*b*) This sketch illustrates our current understanding of the expected features of the phase diagram of QCD as a function of temperature and baryon doping, the excess of quarks over antiquarks, parameterized by the chemical potential for baryon number μ_B. The lattice calculations in panel *a* were done with $\mu_B = 0$, corresponding to the vertical axis of the phase diagram. The regions of the phase diagram traversed by the expanding cooling droplets of QGP formed in heavy ion collisions with varying energy $\sqrt{s_{NN}}$ are sketched. The transition from QGP to hadrons is a crossover near the vertical axis; the thermodynamics of this crossover is well understood from lattice QCD calculations that are quantitative and controlled in the yellow region. At higher doping, the transition may become first order at a critical point. A central goal of the coming second phase of the RHIC Beam Energy Scan (BES II) is to determine whether such a critical point exists in the region of the phase diagram that can be explored using heavy ion collisions. At higher baryon density and lower temperature, cold dense quark matter is expected to be a color superconductor. This form of matter may be found at the centers of neutron stars. Panel *a* adapted from Reference 19. Panel *b* adapted from Reference 22.

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

of points in space and Euclidean time, and doing so for a series of lattices with smaller and smaller spacing between the points, thus taking the "continuum limit" (32). These lattice calculations have taught us many lessions, including that the transition from a gas of hadrons to the hot, liquid QGP with zero baryon number is a crossover (28), as in **Figure 2**, with no further transitions anticipated as QGP gradually goes from liquid-like to gas-like at higher and higher temperatures (33). Because lattice calculations are built upon the Euclidean formulation of equilibrium thermodynamics, it is much more challenging to use them to gain information about transport coefficients including the shear and bulk viscosities, which describe the time-dependent processes via which infinitesimal perturbations away from equilibrium relax, producing entropy. Pioneering attempts in these directions have been made (34). Lattice calculations of more dramatically time-dependent phenomena, including the quenching of jets in the liquid plasma or the initial formation of the plasma from a far-from-equilibrium collision, are beyond the horizon.

The big picture that has emerged over the past 15 years, namely that the hot matter produced in ultrarelativistic heavy ion collisions rapidly forms a strongly coupled hydrodynamic liquid with a strikingly small value of η/s, has posed new open questions that motivate much experimental and theoretical investigation today. For example, how, and how quickly, does the hydrodynamic liquid form from the nonhydrodynamic initial conditions at the moment of the collision? Or, how does a hydrodynamic liquid emerge at its natural length scales (of order $1/T$ and longer, with T the temperature) in an asymptotically free gauge theory in which all matter, when resolved at short length scales, must be made of weakly coupled quarks and gluons? Or, what is the smallest droplet of this stuff that can sensibly be described using the language of hydrodynamics? We return to the first and second of these big questions below, in Section 6 and Section 7, respectively. As to the third big question, it is currently the subject of intense investigation, both experimental and theoretical, having been put squarely on the agenda for the field by measurements made in pPb and pp collisions at the LHC and dAu collisions at RHIC, which indicate that even proton-sized droplets of hot QCD matter can exhibit liquid-like behavior. In response to this discovery, theorists have shown that, in the cousins of QCD with a dual gravitational description, the dynamics of a droplet of strongly coupled plasma with temperature T that is $\sim 1/T$ in size or larger can be described hydrodynamically (35–37). This suggests that hydrodynamic behavior should not persist in pA collisions at lower collision energies and hence lower multiplicity. Noting that this question is at the center of a different article in this volume (38), we keep our discussion brief.

Length scale: according to the uncertainty principle, the characteristic microscopic length scale is proportional to the inverse of the characteristic momentum, which in a relativistic plasma means it is inversely proportional to the temperature

2.2. Phase Diagram of QCD

Among the most important reasons for studying ultrarelativistic heavy ion collisions is the expectation that doing so will teach us about the phase diagram of hot QCD matter, in thermal equilibrium, as a function of both temperature and baryon doping (**Figure 2**). By baryon doping (or net baryon number density) we mean the excess of quarks over antiquarks in the hot matter. The standard parameter used to characterize the degree of baryon doping is the baryon number chemical potential μ_B. To this point, we have set μ_B to zero, which describes matter with equal densities of quarks and antiquarks. This is a very good approximation for the matter produced at midrapidity in the highest-energy heavy ion collisions at RHIC, an even better approximation at the LHC, and an exceedingly good approximation in the early Universe. In all these cases, ordinary hadronic matter forms via a continuous crossover as the liquid QGP expands and cools. However, matter with $\mu_B = 0$ and varying temperature is only one edge of a phase diagram. A substantial component of our understanding of the nature of any complex material in condensed matter physics is mapping its full phase diagram, and the same is true in QCD.

One way to study QGP doped with a significant excess of quarks over antiquarks would be to study the debris produced at very high rapidity in the highest-energy heavy ion collisions, the rapidities where QGP forms from the compressed remnants of the incident nuclei. Neither RHIC nor the LHC features detectors that can do so, at present. Instead, we can scan a region of the phase diagram of QCD by looking at heavy ion collisions with lower and lower collision energies in which the initial baryon number found in the incident nuclei makes a larger and larger contribution to the matter formed in the collisions: Decreasing the collision energy increases μ_B, scanning the phase diagram. Lower-energy AA studies are under way at the Super Proton Synchrotron (SPS) (39) and in the RHIC Beam Energy Scan (BES) (40), where tantalizing early results are at hand from the first phase of the BES program with relatively low statistics per collision energy (22). There is a second, high-statistics phase of the BES program planned for 2019–2020. Extensions of this program to even lower collision energies (and hence even higher μ_B, albeit at lower temperature) are planned at the Facility for Antiproton and Ion Research (FAIR) in Darmstadt, Germany (41), at

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

the Nuclotron-Based Ion Collider Facility (NICA) in Dubna, Russia (42), and at the Japan Proton Accelerator Research Complex (J-PARC) in Tōkai, Japan. One of the central questions that these experiments aim to answer is whether the continuous crossover between liquid QGP and hadronic matter turns into a first-order phase transition above some nonzero, critical value of μ_B, meaning in heavy ion collisions below some collision energy. There are many models for QCD in which the phase diagram features a critical point like this (43). [In QCD with two massless quarks—the "chiral limit"—the crossover at $\mu_B = 0$ becomes a sharp second-order phase transition at which chiral symmetry is restored, and a point at $\mu_B > 0$ where the transition becomes first order is a tricritical point (43, 44).] Furthermore, a critical point has also been seen in some pioneering efforts to explore physics at nonzero μ_B using lattice techniques, although because lattice calculations at nonzero μ_B suffer from a "sign problem," these calculations typically require small μ_B/T, and to date it has not been possible to take the continuum limit (45). There are also tantalizing indications of increased non-Gaussian fluctuations (46–48) in exactly the observable that has been predicted to be most sensitive to critical fluctuations in RHIC collisions near the low end of the beam energy scan, but these indications are inconclusive given the presently available statistics. Do we know whether there is a critical point in the phase diagram at nonzero baryon doping? No. Are there strong motivations for the experimental program that aims to answer this question within the next few years? Yes. We have been relatively brief here, anticipating that the forthcoming data and analyses coming soon will warrant a focused review of their own before too long.

2.3. Emergence of Complex Quantum Matter

In the history of the Universe, liquid QGP was the earliest complex matter to form. At much earlier times, when the temperature was a few orders of magnitude hotter than those of interest to us here, the matter that filled the Universe was a weakly coupled plasma of quarks and gluons. We

HIGH BARYON DENSITY AT LOW TEMPERATURES IN THE COSMOS

Pushing to very high baryon doping while staying at low temperature (i.e., squeezing nuclei without heating them) takes us into another interesting region of the QCD phase diagram. Matter that is sufficiently dense cannot be made of well-separated nucleons, even at low temperatures: The nucleons are crushed into one another. Because quarks attract one another, cold dense matter in which quarks fill momentum space up to some high Fermi momentum is a color superconductor in which a condensate of correlated Cooper pairs of quarks creates a superfluid and yields the QCD analog of a Meissner effect. Extensive theoretical analyses of the phase diagram and consequent properties of color superconducting quark matter (49) have been performed; they are well understood at asymptotic densities, but at densities of order 10 times that of nuclei they turn out to be sensitive to the ratio of the s quark mass to the Fermi momentum as well as to the strength of the Cooper pairing, making them hard to pin down quantitatively. Experimental data are sorely needed. Unfortunately, the only place in the Universe where cold dense quark matter may be found is in the centers of neutron stars. Remarkably, the first collision between two neutron stars has just been observed by the LIGO and VIRGO Collaborations, via the gravitational waves it produced (50)! Although the gravitational waves from this discovery event seen by LIGO reveal only the inspiraling incident neutron stars, with coming improvements to LIGO's sensitivity future events will give us a view of the collision itself, making it possible to learn about the compactness and density profile of the incident neutron stars and, conceivably, whether or not they feature dense quark matter cores. If they do, present constraints on heat transport in neutron stars coming from X-ray observations of how they cool will turn into constraints on the transport properties of cold dense quark matter.

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

know this because QCD is asymptotically free, meaning that quarks and gluons interact with one another only weakly when they scatter off one another with large enough momentum transfer. Moreover, there is also a sense in which it is the simplest form of complex matter that we know of, namely the complex matter that is "closest," or most directly connected to, the fundamental laws that govern all matter in the Universe, in this case the fundamental theory of QCD. Again because QCD is asymptotically free, we know that if we could hold a droplet of the liquid QGP with temperature T in place and study its microscopic structure with a spatial resolution that is much finer than $1/T$, for example, via scattering high-energy electrons off it in this thought experiment, what we would see is weakly coupled quarks and gluons. This is the genesis of the strongest motivation for developing experimental techniques for probing the structure of the liquid QGP on varying length scales. We know that at the shortest length scales we must see weakly coupled quarks and gluons. We also know that at length scales of order $1/T$ and longer we see a liquid in which neighboring "unit cells" are tightly coupled to one another, meaning that the liquid flows hydrodynamically with a small η/s. If we can probe both these length scales and scales in between, for example, by studying how jets (which are intrinsically multiscale probes), or heavy quarks with varying mass, or tightly bound quarkonium mesons with varying sizes "see" the plasma and how the plasma responds to their passage through it, we will have a chance to probe, and maybe even understand, how the simplest form of complex matter that we know of emerges from weakly coupled, asymptotically free constituents at short length scales. The question of how the almost infinite variety of complex forms of matter that we see in the world around us emerge from laws of nature that are so simple that they can easily fit on a T-shirt is one of the great quests of modern physics. If we can answer it for the case of liquid QGP, which we have a chance to do by virtue of this simple form of complex matter being so close to its laws-of-nature underpinnings, we may have a chance of shedding light on the larger more general question.

Asymptotic freedom: the QCD coupling α_s weakens for interactions between quarks that are close together (or scatter at high energy), and is strong and nonperturbative at length (or energy) scales of order (the inverse of) the size of a hadron. In gas-like QGP at asymptotically high temperatures, the interaction energy is small compared with the kinetic energy

3. PHENOMENOLOGY OF HEAVY ION COLLISIONS

In the study of heavy ion collisions, experimenters have only two quantities under their direct control: which two nuclei they collide and at what energies. The energies are known to high precision. However, knowing the colliding nuclei is not the same as knowing the colliding systems. Neither the impact parameter b (the transverse distance between the center of masses of the two nuclei) nor the location and motion of the nucleons in the nuclei, let alone those of the quarks and gluons in the nucleons, are measurable quantities. They have to be inferred, as well as possible or as needed, event by event, from the observed outcome of the collision. This then makes it possible, after the fact, to select an ensemble of collisions with a relatively narrow distribution of impact parameters.

Based on nuclear and particle physics studies we know that, from the point of view of relativistic heavy ion collision studies, the nuclei can instantaneously be reasonably well approximated by a collection of nucleons, distributed on average according to a well-determined three-dimensional distribution. We also know the average quark and gluon content of the nucleons in the nuclei in terms of parton distribution functions, or PDFs, and find that the PDFs in nuclei differ only mildly from those describing free nucleons (51). Furthermore, we can use the measured energy-dependent total inelastic pp cross sections $\sigma_{pp}(\sqrt{s})$ (52) to model the nucleons in the nucleus as hard spheres with a radius that depends on energy.

It turns out to be useful to do a "gedanken experiment" where we imagine the colliding nuclei to be composed of A (transparent) spheres of radius $\sqrt{\sigma_{pp}/4\pi}$, where $A_{\mathrm{L,R}}$ is the number of nucleons inside the left- and right-moving nuclei. We then call those nucleons that do not encounter any nucleon from the other nucleus spectators (dashed circles in **Figure 3**). These nucleons continue

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

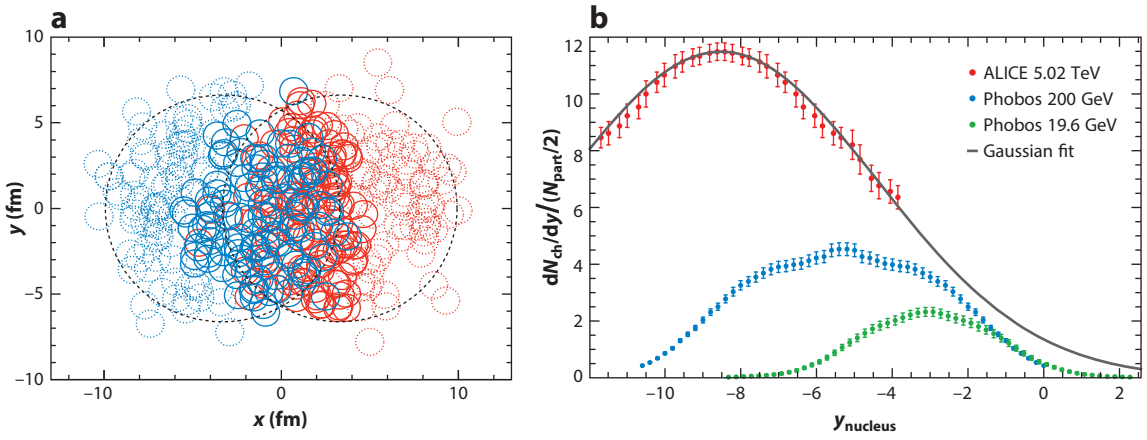

Figure 3

(*a*) An example of a PbPb collision at the LHC with impact parameter $b \approx 7$ fm. The number of participants (*solid*) is counted by nucleons that collide with any nucleon, whereas the number of binary collisions counts all overlapping blue/red nucleon pairs. Spectators (*dashed*) do not collide. (*b*) Rapidity distributions of charged hadrons, in the rest frame of one of the nuclei, for AuAu collisions at 19.6 and 200 GeV (converted from pseudorapidity from Reference 54 using a simplified Jacobian) and for PbPb collisions at 5.02 TeV (55). Panel *a* adapted from Reference 53.

traveling down the beam pipe, and the number of spectators $N_{\rm spec}$ can hence in principle be measured directly, although in practice this is usually hard. In this gedanken experiment, all other "wounded" or participating nucleons collide with at least one other nucleon and make up the number $N_{\rm part}$ (solid circles in **Figure 3**; by definition, $N_{\rm spec} + N_{\rm part} = A_{\rm L} + A_{\rm R}$). It is unfortunate that $N_{\rm spec}$ is not measurable in practice, since if it were then $N_{\rm part}$ could be determined directly. Lastly, if we imagine the spheres as transparent we can also count the total number of encounters between left- and right-moving nucleons, which we call the number of binary collisions $N_{\rm coll}$. For example, if one "nucleus" consists of seven nucleons lined up in a row and it collides head on with a "nucleus" consisting of four nucleons in a row, $N_{\rm part} = 11$, $N_{\rm coll} = 28$, $N_{\rm spec} = 0$, and the impact parameter $b = 0$. In a real central heavy ion collision, a nucleon at the center of one nucleus will, on average, hit about 12 nucleons from the other, but fewer if it is located at the edge of the collision. So $N_{\rm coll}$ will be much larger than $N_{\rm part}$, and even more so for the more central collisions.

In a *pA* collision, the probability of the proton hitting another nucleon is given by the ratio σ_{pp}/σ_{pA} of inelastic scattering cross sections. This makes it possible to determine that on average $N_{\rm coll} = N_{\rm part} - 1 = A\sigma_{pp}/\sigma_{pA}$, which can be measured directly. Experimental data on *pA* collisions with widely varying A and collision energies [going back to the 1970s (56)] show that the number of particles produced in such collisions is proportional to $N_{\rm part}$ to a good approximation. Although in *AA* collisions $N_{\rm part}$ and $N_{\rm coll}$ cannot be determined directly from measured cross sections, there is a well-defined theoretical procedure [called a Glauber model calculation (57)] for determining these abstract measures, at least on average within "centrality classes." This procedure generates many configurations with different b, as illustrated in **Figure 3**, and thereby generates Monte Carlo distributions of $N_{\rm part}$ (as well as $N_{\rm coll}$). It is then assumed that there is a monotonic relation between the number (or energy) of the produced particles and $N_{\rm part}$. For example, we assume that events in which the number (or energy) of particles falls into the highest 5% class correspond to the 5% most central collisions, with $N_{\rm part}$ or $N_{\rm coll}$ (from the Glauber model calculation) in the highest 5% category. The bases for this prescription are, first, the experience from *pA* collisions that we mentioned above and, second, the observation that the shape of the measured probability distribution

Participant: a nucleon that collides with at least one other nucleon

Spectator: a nucleon that does not collide and hence keeps moving along the beam direction

Binary collisions: total number of nucleon pairs that collide, assuming transparency of the collision

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

for the number (or energy) of particles in an ensemble of AA events is similar to the probability distribution for N_{part} obtained from a Glauber model calculation. Most importantly, the participant scaling observed for collisions of nuclei with widely varying A (discussed in Section 3.3, below) provides a strong indication that these abstract measures in some way reflect a physical reality.

We now describe, as a function of energy, N_{part} and N_{coll}, the most general features observed when two heavy ions collide at relativistic velocities. In order to give the "big picture," in this discussion we ignore small differences and subtle effects. For useful summaries of and references to RHIC and LHC data, we refer the reader to References 58–61 and 62–64, respectively (see also recent proceedings of Quark Matter conferences and References 22 and 65 for an overview of theoretical and experimental work).

3.1. Hard Collisions

High-p_T γ and Z^0 boson production have been studied in pp and AA collisions (66–68). The measured pp cross sections are well understood and are in excellent agreement with predictions based on the known PDFs and perturbative QCD (pQCD) theory. The measured AA production rates, in turn, are in excellent agreement with the product of N_{coll} and the yield in a single pp collision, taking into account the measured modifications of the PDFs of nucleons inside nuclei and uncertainties in the determination of N_{coll}. Since the hard γs and Z^0 bosons are not affected by the postcollision AA environment (both are colorless), these results show that we have a good understanding of the initial hard (high-p_T) parton–parton interactions in AA collisions and of the determination of N_{coll}, meaning that these results provide independent confirmation from experimental data of the Glauber model calculations described above. This, in turn, implies that AA–pp comparisons of outgoing strongly interacting (colored) hard probes (jets, high-p_T hadrons, heavy quarks, etc.) can be used to give us valuable information about the nature of the medium produced in heavy ion collisions and on how the colored hard probes themselves are modified as they traverse the medium produced in AA collisions. Experimental measurements of jets themselves, as well as of high-p_T hadrons and heavy quarks coming from jets (69), all show that jets lose considerable energy as they propagate through QGP, with the "lost" energy ending up as many soft particles [$\lesssim$3 GeV/c (70)] moving at large angles relative to the original jet direction (71), suggesting that the jet leaves a wake behind in the liquid QGP. This suite of results and phenomena, collectively referred to as jet quenching, is very important since they give us direct evidence of very strong interactions occurring after the collision, strong interactions between the jet and the liquid, as well as strong interactions within the perturbed liquid; we return to them in Section 7.

3.2. Baryon Stopping Power

It has long been known that in lower-energy pp collisions the longitudinal momentum of forward-going protons in the final state has a flat distribution, evenly distributed between zero and the incident energy (72, 73). This implies that, on average, a proton loses half its energy, which is about one unit of rapidity. In pA and AA collisions, the energy lost by the incident nucleons is higher on average, and more narrowly distributed. On average, in heavy ion collisions, each participant loses about two units of rapidity (9), which is to say that 85% of its energy goes into the creation and kinetic energy of a very large number of particles, up to 30,000 in central PbPb collisions at the LHC. The net proton rapidity distribution in AA collisions has a double-hump structure (74), each hump consisting of hot baryonic matter moving at a speed of about two units of rapidity below that of the incident beam, and having a net baryon density of about 5–10 times that of normal nuclear matter (10). [At the LHC the beam rapidity is $y = 8.5$. As seen in the frame in which $y = 6.5$ is at rest, an incident disc that is Lorentz contracted by a factor of about cosh(2) is

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

hit by a disc that is Lorentz contracted by a factor of about cosh(15) and brought approximately to rest, compressed by roughly 2 cosh(2) ≈ 7.5.]

A further consequence is that the maximum value of the net baryon density at midrapidity is produced when heavy ions collide with $\sqrt{s_{NN}} \approx 7$ GeV. Above this collision energy, the midrapidity net baryon density, and therefore also the baryon chemical potential in the QGP produced at midrapidity, decreases with energy. By top RHIC collision energies, and even more so for LHC energies, both are essentially zero (75).

3.3. Energy and Centrality Dependence of Multiparticle Production

For practical reasons we have the most information about charged particles, which correspond to about two-thirds of all the produced particles. However, there is no reason to doubt that the picture obtained from the charged particles is anything other than the whole picture!

From the lowest energies measured (76), through RHIC (54) to LHC (55) energies, for all *AA*, *pA*, πA, and KA collisions, the total number of charged particles produced is approximately proportional to the number of participants. This phenomenon is known as "participant scaling" and is not well understood. Even more surprisingly, provided that one takes into account the fraction of energy that is taken away by the forward-going baryons and not available for particle production (mentioned above; also see Reference 72), the total number of charged particles produced per participant in *AA* collisions is the same as that in *pp* and e^+e^- collisions (77). However this arises, it suggests that, on average, in *AA* collisions most of the entropy production, which is proportional to the number of produced particles[2] and hence the number of participants, occurs early in the collision and that there is little, if any, late-stage entropy production. Below, we see that there are powerful arguments in support of this conclusion: After the early stage of the collision when entropy production is copious, a hydrodynamic fluid forms, and because this fluid has low specific viscosity, little entropy is produced subsequently, as the liquid flows.

For a given number of participants, the total number of produced particles N increases with energy [as $N \propto s_{NN}^{0.15} \log(s_{NN})$] (79). Except for close to the receding discs, the longitudinal rapidity distributions look crudely like wide Gaussians with a width that increases as $\log(s_{NN})$ (i.e., as the beam rapidity or longitudinal phase space; see **Figure 3**) (55), and the width of the distribution increases weakly from central to peripheral collisions (80). The produced particle density $\mathrm{d}N/\mathrm{d}y$ hence has no wide boost-invariant plateau, and for central collisions the maximum at midrapidity increases with energy as ($\mathrm{d}N/\mathrm{d}y \propto s_{NN}^{0.15}$) (81). At first glance, the simplicity of these empirical facts seems to be at odds with the complex sequence of stages that precede particle production. In the rest frame of the produced particles, the finally observed particle density $\mathrm{d}N/\mathrm{d}y$ is the result of a local history that includes the initial impact of the nuclei, followed by the creation, expansion, and flow of a hot medium, and its eventual hadronization into particles. At one level, the simplicity of the empirical facts can be explained by noting that the number of particles in the final state is proportional to its entropy, and concluding that at any rapidity most of the entropy is produced very early in the collision, making $\mathrm{d}N/\mathrm{d}y$ insensitive to all that occurs later. Nevertheless, these facts are not fully understood. For example, the energy dependence and centrality dependence are surprisingly independent of one another from the lowest to the highest energies studied (54, 82, 83).

[2] At the chemical freeze-out temperature, the multiplicity of each of the hadron species present in QCD is given to a good approximation by a thermal distribution (discussed further in this section and in Section 4). This makes a direct connection between the entropy at this moment and the number of charged particles possible. The contribution from any single species in a thermal distribution to N_{ch} is proportional to S, with a proportionality constant that decreases with increasing mass. Adding up all the known species of hadrons yields $N_{\mathrm{ch}} \approx S/7.25$ at freeze-out (78).

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

This means that the naïvely expected increase with energy resulting from the increase of hard (proportional to N_{coll}) relative to soft (proportional to N_{part}) processes does not play a leading role in determining the number of produced particles.

Another interesting observation is the so-called extended longitudinal scaling. If the rapidity of one nucleus is kept constant and that of the other is gradually increased, at first $\mathrm{d}N/\mathrm{d}y$ increases but then reaches a limiting value (see **Figure 3**). If we think of the second nucleus as a wall of gluons, boosting these gluons more and more seems to have no effect on particle production in the collision around the rapidity of the first nucleus. This phenomenon has been observed for all systems studied (84) and is direct evidence that a kind of saturation occurs in the fast nucleus (85).

Finally, we point out that all these facts do not support Landau's (2, 16) and Fermi's (15) early models, in which they postulated that the two colliding systems completely stop each other and then (in the case of Landau, after a period of hydrodynamic expansion from rest) break up into particles according to thermodynamic laws. And they are also inconsistent with Feynman's (17) intuition that $\mathrm{d}N/\mathrm{d}y$ at midrapidity would not increase with increasing collision energy. Feynman expected the rapidity distribution of the produced particles to broaden with increasing collision energy; this does happen, but because of the rapid rise of the gluon PDF that Feynman did not anticipate, the total particle production increases fast enough that $\mathrm{d}N/\mathrm{d}y$ at midrapidity nevertheless increases.

3.4. Particle Correlations

Strong correlations are observed between particles produced with momenta in different directions. They are much stronger than expected from the superposition of independent *pp* collisions, and are evidence that the products of the initial collision act collectively.

Correlations between particles that are widely separated in rapidity are observed (87, 88). By causality, they must have their origin in early times and thus give information about correlations present at the earliest stages of the collision of the two nuclei. Azimuthal correlations, as in **Figure 4*b*** in particular, have a very pronounced and rich structure and have been extensively studied as a function of the centrality of the collision, produced particle type, rapidity, p_T, and expected event-by-event geometrical fluctuations of the nuclei (5, 89). As discussed in detail in Section 4, they can be remarkably well explained by relativistic hydrodynamics, if one assumes that in high-energy heavy ion collisions, before the final production of free-streaming particles, some kind of a relativistic liquid is formed, which expands and flows radially at about half the speed of light and in which pressure-driven anisotropies in the flow velocity form and persist because the liquid has an incredibly low viscosity-to-entropy ratio, in fact—lower than that of any other known liquid. It is for this reason that we know that QGP is a strongly coupled liquid.

A good way to see that the medium produced in a heavy ion collision indeed behaves as a low-viscosity hydrodynamic liquid is to note the following. Like all nuclei, those that collide in heavy ion collisions are lumpy, meaning that the energy density of the matter produced in the earliest moments of the collision must also be lumpy. If that matter were a tenuous gas-like plasma, made of lots of particles that fly around while interacting only rarely with one another, the initial lumpiness would quickly disappear as the particles fly around in random directions, and at the end of the day all one would see is an isotropic explosion of particles, with just as many particles going in any one direction as in any other. If, instead, the matter that is produced is a liquid whose motion is governed by hydrodynamics, the initial lumpiness will mean initial pressure gradients, and these pressure gradients will drive anisotropic flow in the liquid. If the viscosity of the liquid is high, these anisotropic flows will damp out. Instead, what is seen in heavy ion collisions is substantial anisotropies in the azimuthal distribution of particles in the final state (as in **Figure 4**), which

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

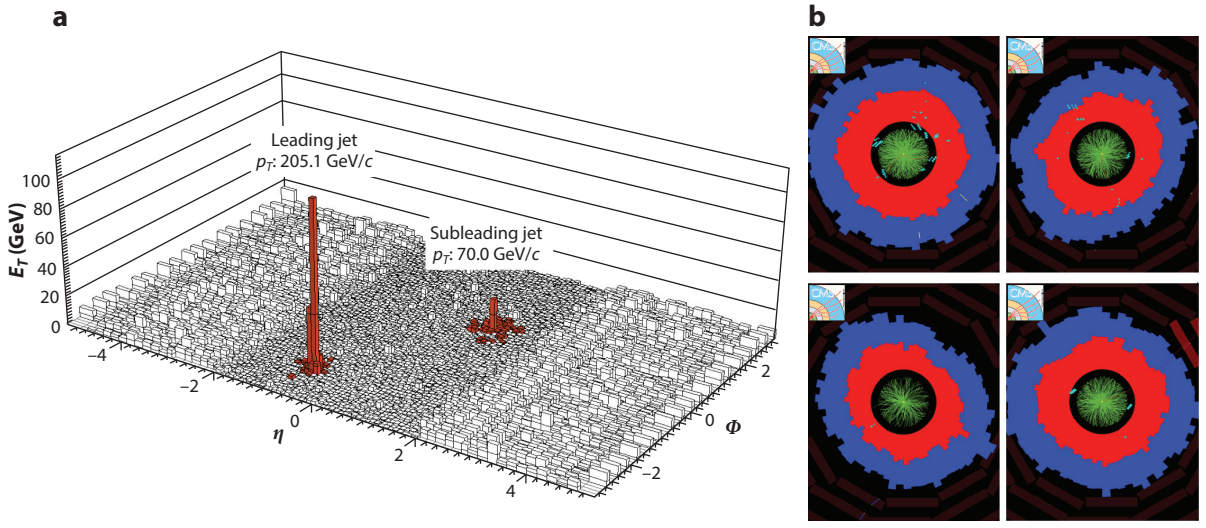

Figure 4

(*a*) This event display (86) shows energy deposited in the CMS calorimeter in a heavy ion collision as a function of azimuthal angle ϕ and pseudorapidity η, a proxy for rapidity that is more easily measured. Two jets of very different energies are apparent, suggesting that one jet lost more energy as it traversed the droplet of quark–gluon plasma (QGP). (*b*) CMS event displays showing azimuthal distribution of charged tracks (*green*) and energy in the electromagnetic (*red*) and hadronic (*blue*) calorimeters from four heavy ion collision events as seen by the CMS detector. The azimuthal anisotropies are apparent, with the upper-right and lower-left events showing marked ellipticity and the bottom-right event showing a substantial anisotropy in a higher harmonic. It is remarkable that the strongly coupled character (*left*) and the liquid nature (*right*) of the QGP formed in these collisions can be seen so clearly in individual events.

reflect azimuthal anisotropies in the geometry of the overlap region of the colliding nuclei. This means that the matter produced in the collisions must be a fluid with low specific viscosity.

3.5. Medium Properties

Azimuthal correlations provide information about the relativistic hydrodynamic nature of the medium, about its transport coefficients, and about the fluctuations in the initial state from which the medium forms (5, 6, 90). Jet quenching studies (69, 91, 92) give us a wealth of information on both how the medium responds when a high-energy quark or gluon jet produced in an initial hard scattering traverses it, and how a fast quark or gluon jet is modified by the medium as it passes through it. As mentioned in the preceding section, jet quenching shows that QGP is extremely strongly interacting and provides us with insight into the workings of QCD; we return to this important topic in Section 7. There are no known measurements that give us direct and unambiguous information about the nature of the low-viscosity fluid produced, whether it is in equilibrium, how it forms and equilibrates, its equation of state and phase diagram, the best degrees of freedom for its description, how many thermodynamic degrees of freedom it has compared with a hadron gas, and whether it is a liquid of deconfined quarks and gluons. However, there are indirect measurements that give us insight into these questions and, together with theoretical studies (particularly lattice gauge calculations), a consistent picture of the nature of QGP is emerging.

For example, near midrapidity, in RHIC or LHC central *AA* collisions, the ratios of the hadrons containing the lighter up, down, and even the strange valence quarks (93) are well represented by a system in chemical equilibrium at a temperature of approximately 155 MeV (94, 95). (Note that the number densities of charm and bottom quarks do not reach chemical equilibrium because the

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

temperature is not high enough, meaning that their multiplicities retain memory of their initial production. Top quarks are not relevant here because of their short lifetimes.) In contrast, the p_T spectra are consistent with a system in equilibrium with a lower temperature of about 95 MeV and substantial radial flow (75, 96). Consistent with QGP being a strongly coupled liquid that behaves hydrodynamically as it expands and cools, there are no indications of any abnormal production of very low momentum pions (with wavelength approximately the same size as a QGP droplet) (97), for example, from the formation of a region of disoriented chiral condensate (98).

These facts, combined with the observed azimuthal correlations, participant scaling, and jet quenching, are consistent with the following interpretation. Very early in the collision of the two Lorentz-contracted nuclei, a thin cylindrical volume of QGP liquid is formed, with an entropy that is determined early, before the fluid hydrodynamizes. At first this liquid has a nonuniform energy density and temperature distribution determined by the lumpiness of the colliding nuclei. It expands and cools in accordance with relativistic hydrodynamics, and because its specific viscosity is so small it does so almost isentropically. When the temperature of the system locally falls below about 155 MeV, the QGP goes through a crossover phase transition and hadronizes. It is not known whether the hadrons are produced in chemical equilibrium or chemically equilibrate quickly, after the phase transition. All this is the so-called chemical freeze-out. The produced hadronic system then continues to interact, expand, and cool until the temperature falls to about 95 MeV, when thermal freeze-out occurs. After thermal freeze-out, the hadrons stream outwards freely, eventually reaching the detectors. At the thermal freeze-out time, in addition to thermal motion the hadrons have radial and anisotropic velocities inherited from the flow of the expanding liquid that came before.

Measurements of quarkonia (J/ψ and Υ mesons made from moderately heavy 1.3-GeV c quarks and heavy 4.2-GeV b quarks, respectively) production in heavy ion collisions compared with that in pp collisions (64) provide further information about the properties of the QGP medium, in two different ways. Consider first the case in which the production of a heavy $Q\bar{Q}$ pair in the hard collision at the very beginning of the collision process is rare, for example, as for $b\bar{b}$ pairs in LHC collisions, ideally meaning that in each heavy ion collision there is zero or one $b\bar{b}$ pair. The $b\bar{b}$ pair finds itself immersed in the QGP medium, which, via Debye screening, weakens the attractive force between the pair. The smallest, most tightly bound, $\Upsilon(1S)$ state has a size comparable to or even smaller than the Debye length of QGP, meaning that the b and $\bar{b}$ may be close enough together to remain bound even when immersed in QGP. The $\Upsilon(3S)$, by contrast, is comparable in size to ordinary hadrons, meaning that a b and a $\bar{b}$ with this separation do not attract each other when screened by QGP, and drift apart. The $\Upsilon(2S)$ is an intermediate case. **Figure 5** is a beautiful example of data that show that Υ states with different sizes and binding strengths do indeed have different probabilities of surviving in QGP, supporting this picture.

J/ψ production in LHC heavy ion collisions is interestingly different (64). These collisions are sufficiently energetic that, on average, about 30 $c\bar{c}$ pairs are produced in each heavy ion collision (100). In N_{coll} independent pp collisions, in which the same number of $c\bar{c}$ pairs is produced, any J/ψs that form originate from the c and $\bar{c}$ produced in a single hard scattering. In a heavy ion collision, those primordial J/ψs are expected to fall apart in QGP, as above. However, it now becomes possible for a J/ψ to form via a new process in which a c and a $\bar{c}$ from different initial hard scatterings thermalize in, and diffuse through, the QGP formed in the collision and then happen to find each other at the time of hadronization. $c\bar{c}$ production is so copious at LHC energies that there are more J/ψs produced in heavy ion collisions via this recombination process than are produced in the standard fashion in N_{coll} independent pp collisions. This confirms that the c and $\bar{c}$ quarks produced in heavy ion collisions wander independently of each other, and is thus a direct confirmation that quarks in QGP are not confined within hadrons.

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

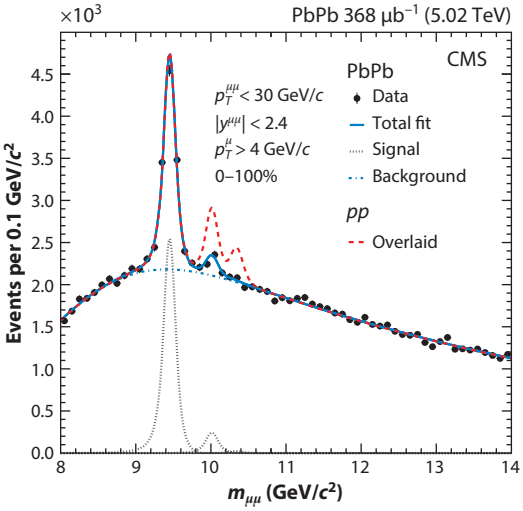

Figure 5

The dimuon invariant mass distribution shows the different ϒ states, where the red dashed line shows the *pp* result added to the PbPb background and normalized to the ϒ(1*S*) state (99). Clearly the ϒ(2*S*) and ϒ(3*S*) states in PbPb collisions are much less pronounced, which is interpreted as the melting of these larger and less strongly bound $b\bar{b}$ states when they find themselves immersed in quark–gluon plasma.

3.6. Comparing *AA* Collisions with *pp* and *pA* Collisions

Unlike in *AA* collisions, the jet quenching phenomenon is not seen in *pA* collisions: At midrapidity the number of jets seen is just what one would expect from N_{coll} *pp* collisions (101). [At large forward and backward rapidities there are deviations from this pattern, deviations that are understood as coming from differences between nuclear and nucleon initial states (101).] This absence of jet quenching came as no surprise, since *pA* collisions produce an energetic final state that is small in transverse extent, and because in ultrarelativistic collisions the incident nucleus is highly contracted in the longitudinal direction, meaning that the nascent jets quickly emerge from the energetic final state and cannot encounter the spectators from the incident nucleus. What did come as a surprise is how many other phenomena are similar in *AA* and *pA* collisions, and even in *pp* collisions, in particular when the comparison is done between collisions in different systems with the same final-state particle density $\mathrm{d}N/\mathrm{d}y$. Examples include the rapidity distribution (102), particle spectra (103), particle ratios including those involving strangeness (104), and most significantly, the azimuthal anisotropies (105, 106) encoded in multiparticle correlations that were once thought to be unique to *AA* collisions. Although these similarities are not yet well understood and are currently topics of intense debate, it is tempting to interpret them as indicating that proton-sized droplets of QGP can be formed in those *pp* and *pA* collisions that produce final states with sufficiently large $\mathrm{d}N/\mathrm{d}y$. This has prompted a recent theoretical focus on the question of how small the smallest droplet of QGP that can be described hydrodynamically can be, and the realization that in the case of a strongly coupled liquid-like QGP the answer seems to be around $1/T$ (35–37). This makes it plausible after the fact that some sufficiently energetic *pp* or *pA* collisions can create droplets of QGP with temperatures well in excess of the inverse of the proton size. A full discussion of the "heavy ion" features observed in small collision systems can be found in another article in this volume (38).

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

4. A HYDRODYNAMIC FLUID

A crucial feature of our description of a heavy ion collision and the interpretation of the observed facts is that, shortly after the initial impact of the heavy ions and before the hadronization process, the system (QGP) is in the form of a near-perfect (extremely low specific viscosity) liquid. We now address in more detail and rigor, and to the extent current understanding allows, questions such as: How and to what extent do we understand the state of this system? Is it in equilibrium? Is it hydrodynamized and locally isotropic? What do we know about its transport properties?

As explained in Section 2, above, we know that at its peak the energy density of the system far exceeds that of hadrons, let alone nuclei. There is no way that the system could be a tightly packed collection of hadrons. Instead, it has to be described in terms of the quarks and gluons themselves. The interplay between two crucial features of QCD determines the nature of this state of matter. First, because of asymptotic freedom and the high energies probed at RHIC and the LHC, it could be that the interactions between the quarks and gluons are so weak that an equilibrium thermal state of matter would never be reached. Second, at energy scales within an order of magnitude of the confinement/deconfinement energy scale, QCD is strongly coupled. The implication of this was not fully realized before experiments at RHIC began (110, 111), as the most common expectation was the formation of an equilibrated gas of quarks and gluons with a temperature somewhat above the confinement/deconfinement scale. We now realize that in this temperature range QCD describes a relativistic fluid consisting of quarks and gluons that are so strongly coupled to their neighbors that the resulting liquid cannot even be described in terms of quasiparticles. The weak coupling picture must be correct at early times in collisions with exceedingly high energy; even in these collisions, the strong coupling picture would become applicable later, after a hydrodynamic fluid has formed. The question of how long during the initial moments of a RHIC or LHC collision a weakly coupled picture can be applied remains open.

The crucial distinction between these two scenarios can be found by measuring the anisotropy of particles produced in heavy ion collisions. Qualitatively this is easy to understand, as we have seen in Section 3: In the case of a weakly interacting gas of particles, scatterings are rare, the directions of the momenta of the gas particles are close to random, the initial spatial anisotropy in the collision zone is largely washed out by random motion, and the azimuthal distribution of particles in the final state ends up close to isotropic. In this case, the measured two-particle correlations are close to trivial, coming almost exclusively from effects like momentum conservation in late-time decays of hadrons. Alternatively, if the quarks and gluons form a strongly coupled liquid soon enough, while the distribution of energy density produced in the collision remains anisotropic, this noncircular and lumpy drop of fluid will expand in a hydrodynamic fashion, yielding faster expansion in the direction of larger gradients: Hydrodynamics converts spatial anisotropies into momentum anisotropy. For perfectly circular collisions this would not lead to an interesting distinction, but in the hydrodynamic picture we would expect anisotropy to arise because the incident nuclei are made of nucleons, and hence lumpy, as well as an increasing anisotropy in the particle spectrum as we probe less central, less circular collisions.[3]

[3]It is also worth noting that when researchers have modeled the bulk dynamics of the matter produced in heavy ion collisions via a system of colliding particles, fitting such models to empirical observations inevitably requires unphysically large scattering cross sections [e.g., parton–parton inelastic scattering cross sections 15 times larger than in pQCD (112) or values of α_s as large as 0.6 (113) or unphysically short mean-free paths]. For example, in both the BAMPS (113) and AMPT (114) approaches, the particles in the model have mean-free paths that are much shorter than their de Broglie wavelengths. Although these approaches differ from hydrodynamics in detail, at a qualitative level what is happening in these models is that interactions in a particulate model are being dialed up to a sufficient degree that the model describes a fluid with low specific viscosity. [This has been shown explicitly for BAMPS (113).]

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

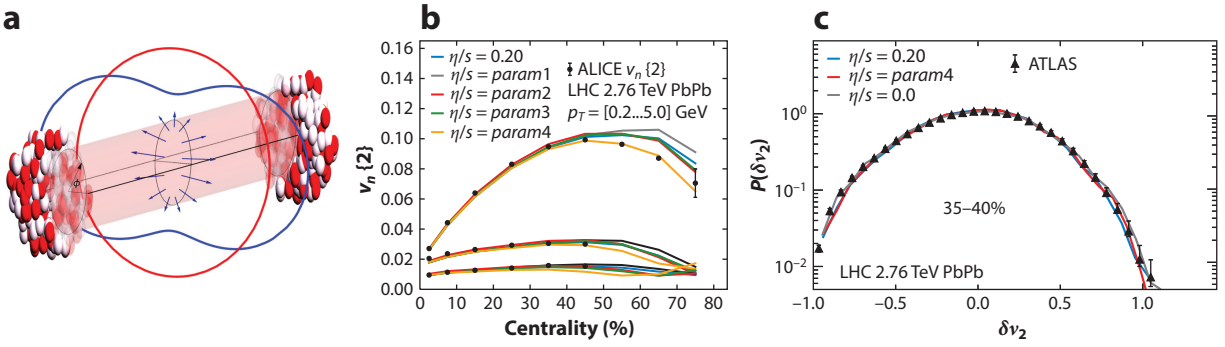

Figure 6

(*a*) A peripheral heavy ion collision produces an approximately elliptical collision region (*red shading*). A gas of weakly interacting particles would give a more or less isotropic distribution of final particles (*red*), whereas a fluid would give rise to an anisotropic distribution (*blue*), due to the difference in pressure gradients in the transverse directions. (*b*) In Reference 108, a hydrodynamic model with several temperature-dependent parameterizations of η/s is compared with ALICE measurements of the anisotropy, as obtained by the integrated Fourier coefficients v_n ($n = 2$–4 from top to bottom), for charged particles with transverse momenta p_T between 0.2 and 5.0 GeV in $\sqrt{s_{NN}} = 2.76$-TeV collisions as a function of the centrality class (0% being head-on collisions) (107). For more off-central collisions there is an increasing and large v_2, giving a hint of the importance of hydrodynamic evolution. (*c*) This panel shows event-by-event distributions of the v_2 distribution for off-central collisions from Reference 107 compared with ATLAS measurements (109). The results shown in panels *b* and *c* are two among many recent comparisons between increasingly precise measurements of the anisotropy and increasingly sophisticated hydrodynamic calculations.

To quantify the measurement of the azimuthal momentum anisotropy, we perform a Fourier transformation on the angular distribution of (charged) hadrons in the final state of the collision (115), which results in the anisotropic flow coefficients $\bar{v}_n$, defined from

$$\frac{\mathrm{d}\bar{N}}{\mathrm{d}\varphi} = \frac{\bar{N}}{2\pi}\left(1 + 2\sum_{n=1}^{\infty} \bar{v}_n \cos[n(\varphi - \bar{\Psi}_n)]\right), \qquad 1.$$

where φ is the angle in the transverse plane, $\bar{\Psi}_n$ are the event plane angles (the first angle where the nth harmonic component has its maximum multiplicity), and $\bar{N}$ is the average number of particles of interest per event. All these observables can in principle be measured as a function of rapidity, centrality, and p_T and, around midrapidity (in collider experiments), also differentially for different particle species. The second to fourth harmonics are shown as a function of centrality in **Figure 6*b***, as extracted from the two-particle correlator with particles separated by a large gap in rapidity.[4] We return to the hydrodynamic curves below.

As anticipated, the system before hadronization indeed requires a full hydrodynamic simulation in order to generate the sizable anisotropies found. Hydrodynamics is a gradient expansion, assuming that a fluid is everywhere close to thermal equilibrium but allowing for small gradients in both temperature and velocity field. In ideal (zeroth-order) hydrodynamics, these gradients are ignored, which by assumption yields an isotropic plasma in the plasma's local rest frame. For

[4]There are several ways to measure the v_n found in Equation 1, most notably via measuring correlations among four, six, eight, or more particles or via analyzing particles separated in rapidity. Both techniques are designed to exclude jet-like correlations between nearby particles that come from the same jet shower or nearly back-to-back correlations from pairs of jets. We do not review the (by now quite sophisticated) methods for extracting the v_n (116). We also do not review the dependence of the v_n on p_T or on hadron species (60), even though their dependence on particle momentum and mass provides important evidence in support of their origin from a single hydrodynamic fluid with a common flow velocity, or their distribution around their average value in each centrality class, which also supports a consistent picture (e.g., 63, 117, 118).

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

viscous (first-order) hydrodynamics, the gradients lead to an anisotropic stress tensor $T_{\mu\nu}$, according to

$$T_{\mu\nu} = \varepsilon u_\mu u_\nu + p[\varepsilon]\Delta_{\mu\nu} - \eta[\varepsilon]\sigma_{\mu\nu} - \zeta[\varepsilon]\Delta_{\mu\nu}\nabla_\mu u^\mu + \mathcal{O}(\partial^2), \text{ where} \quad 2.$$

$$\sigma_{\mu\nu} = \Delta_{\mu\alpha}\Delta_{\nu\beta}(\nabla^\alpha u^\beta + \nabla^\beta u^\alpha) - \frac{2}{3}\Delta_{\mu\nu}\Delta_{\alpha\beta}\nabla^\alpha u^\beta, \quad 3.$$

$$\Delta_{\mu\nu} = g_{\mu\nu} + u_\mu u_\nu, \quad 4.$$

where ε is the energy density and u_μ is the velocity field, both depending on the full space-time coordinates. In the local fluid rest frame, where $u^\mu_{\text{LRF}} = (1, 0, 0, 0)$, the projector is given by $\Delta^{\mu\nu}_{\text{LRF}} = \text{diag}(0, 1, 1, 1)$, and in any frame, $\Delta_{\mu\nu}u^\mu = \Delta_{\mu\nu}u^\nu = 0$. The first two terms in Equation 2 are simply ideal hydrodynamics, whereby the stress–energy tensor is given by an isotropic fluid with energy density ε that is boosted with a velocity u_μ. This fluid has a pressure that is given by the equation of state $p[\varepsilon]$, which is an input into hydrodynamics that depends on the microscopic properties of the theory under consideration. For heavy ion collisions, this is the QCD equation of state, which is usually obtained from lattice calculations like those of **Figure 2** (119; see, however, Reference 120). Lattice calculations are also used to relate the energy density to the temperature.

Beyond ideal hydrodynamics, one needs to include corrections proportional to gradients and consistent with the symmetries present. For scale-invariant viscous relativistic hydrodynamics, it turns out that the only transport coefficient possible at first order in gradients is the shear viscosity $\eta[\varepsilon]$, which accompanies the $\sigma_{\mu\nu}$ of Equation 3, containing first derivatives of the fluid velocity. Close to the deconfinement/confinement transition, QCD is definitely not scale invariant, and there it is also necessary to include the term proportional to the bulk viscosity $\zeta[\varepsilon]$. Just like $p[\varepsilon]$, the viscosities depend on the microscopic properties of the theory, but these transport properties are notoriously difficult to determine from a lattice calculation because they describe the (time-dependent) process by which small deviations from equilibrium relax, whereas what is calculated directly on the lattice is (time-independent) derivatives of the equilibrium partition function. We return to the determination of transport properties below.

Hydrodynamic evolution follows from the conservation of the stress–energy tensor after specifying the equation of state, the transport coefficients, and the energy and velocity profiles at an initial time.[5] In the hydrodynamic evolution equations, $\nabla_\mu T^{\mu\nu} = 0$, the shear viscosity arises in the combination $\eta/(\varepsilon + p) = \eta/(Ts)$, which is proportional to the length scale over which momentum can be transported in the fluid (6). At weak coupling, when the hydrodynamic fluid is made up of quasiparticles with a well-defined mean-free path λ_{mfp}, it can be shown that $\eta/(\varepsilon + p) \propto \lambda_{\text{mfp}}$, meaning that $\eta/s \propto T\lambda_{\text{mfp}}$ (6, 20, 123). In a strongly coupled fluid, η/s is well defined and small, but quasiparticles with mean-free paths cannot be defined since attempting to do so would result in a λ_{mfp} comparable to or smaller than the de Broglie wavelength $1/T$. Whether the fluid is weakly or strongly coupled, η arises in the hydrodynamic equations in this combination, and it is

[5]In practice, solving the equations of viscous hydrodynamics is a bit more involved, since when they are discretized they contain modes with wavelengths of order the discretization scale that propagate faster than light. These modes are unphysical and are outside the regime of the hydrodynamic gradient expansion, but because they are acausal they make the numerical scheme unstable. This makes it necessary in practice to solve a version of second-order hydrodynamics and verify that the choice of second-order terms does not much affect the final results, as must be the case if the gradient expansion is under control (e.g., 5, 6, 121). We also note that we only review the application of hydrodynamics to collisions at LHC and top RHIC energies and, at these energies, for production of QGP far from the fragmentation regions. Extending such calculations outside these regions, as is relevant for the exploration of the QCD phase diagram via the RHIC BES mentioned in Section 2, requires extending Equation 2 to incorporate the time evolution of the conserved baryon number current. It is well known how to do this (122), but less is known about the QCD equation of state and transport coefficients at nonzero baryon chemical potential. There is an additional complication in that when the Lorentz contraction of the incident nuclei is only moderate the dynamics is intrinsically three dimensional.

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

the specific viscosity η/s that controls how rapidly sound waves, shear stress, or gradients of any sort introduced in the initial conditions are dissipated into heat, meaning that it is this quantity that is ultimately constrained by comparing hydrodynamic calculations with data. To proceed further, we must model the initial energy and velocity profile at some proper time τ_0. Fortunately, we find that the insights we present do not depend strongly on simplifying assumptions that we have to make to solve our equations. A simple model used is to take two discs of heavily Lorentz-contracted nuclei to collide at some impact parameter b and from this construct an initial energy profile that follows the overlap of the two discs (the Glauber model) with an overall amplitude as a free parameter. The velocity profile is often taken to be zero in the transverse plane, and in the longitudinal direction the evolution is assumed to be boost invariant around the collision point at $t = z = 0$. Since this assumption implies that the longitudinal velocity is given by $v_z = z/t$, this gives a simple and convenient model for an expanding plasma where all physics just depends on proper time $\tau = \sqrt{t^2 - z^2}$ and the transverse coordinates. In state-of-the-art hydrodynamic calculations that do not assume boost invariance, $v_z = z/t$ remains a good approximation, but the initial distribution of energy density does depend on rapidity.

Having specified the initial conditions and the hydrodynamic equations, the latter by choosing η/s and taking $p[\varepsilon]$ and $T[\varepsilon]$ from lattice calculations, it is possible to start a simulation of the hydrodynamic evolution of this putative hydrodynamic QGP. This simulation evolves the hydrodynamic variables describing an expanding and cooling droplet of matter forward in time up to a freeze-out hypersurface in space-time where the fluid temperature has dropped to a specified value of order the temperature T_c, where the crossover from QGP to hadrons occurs. At the freeze-out hypersurface, the fluid is converted into a thermal distribution of hadrons, conserving energy and momentum (124). Subsequent evolution is described via a gas of hadrons, which interact with one another as further expansion and cooling occur until all scattering ceases at a lower "kinetic freeze-out" temperature. The resultant ratios between the numbers of different hadron species, single-particle spectra for various hadron species, and anisotropy coefficients v_n can all be directly compared with experimental data.

In this model, in order to generate as much transverse flow (both isotropic, or radial, flow and anisotropic flow as described by the v_n) as observed in data, it is necessary to take τ_0 to be smaller than 1.0 fm/c—in some calculations, as small as 0.2 fm/c. [In more advanced models that include the growth of the transverse velocity before τ_0, this constraint can be somewhat weaker (125).] The amplitude of the initial energy density profile for central collisions ($b = 0$) is fitted to yield the observed total particle multiplicity per unit rapidity. The multiplicity as a function of impact parameter b is then a prediction of the model, which can be compared with the experimental results and used to obtain a constraint on η/s.

The precise magnitude of the anisotropies v_n then depends quite sensitively on the viscosity of the plasma. Already from the relatively straightforward simulation with smooth initial conditions described above, it can be estimated that $\eta/s \sim 0.08$–0.20 during the hydrodynamic phase in heavy ion collisions at RHIC energies (126; see Reference 127 for a full three-dimensional hydrodynamic simulation). This is one of the greatest discoveries of the heavy ion programs at RHIC and the LHC: The experimental data are well described by the hydrodynamic evolution of a droplet of QGP with a specific viscosity lower than that of any other fluid known in nature. For this reason, QGP is sometimes referred to as the most perfect liquid.

As an example, and to give a sense of how well the system is currently understood, **Figure 6*b*** presents a more precise comparison between experimental data and one particular hydrodynamic calculation. A crucial ingredient in this computation is the initial condition for the (lumpy, fluctuating) transverse profile, which is taken from a Monte Carlo Glauber model, with fluctuating positions of individual protons and neutrons, convolved with fluctuations of the energy density

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

within a single nucleon that is based on a saturation model in which color fields are large in magnitude but weakly coupled. Fluctuations in the initial state are necessary to obtain agreement with current, precise data, including, in particular, the v_n anisotropies in head-on collisions and the odd harmonics such as v_3 in **Figure 6*b***. Without such fluctuations, the collisions would be perfectly symmetric under parity in the $\vec{b}$ direction, which would imply that all harmonics v_n with n odd would vanish. In actual collisions v_3 is in fact larger than the higher harmonics, showing that fluctuations break this parity symmetry (128). The authors of Reference 107 evolved hydrodynamics with five different assumptions for η/s as a function of temperature (**Figure 6*b***). Current analyses of data are beginning to yield some constraints on this temperature dependence. It is also possible to introduce a bulk viscosity at temperatures near the QCD phase transition (129), where it is expected to be important. The bulk viscosity in Reference 129 is needed to get an accurate fit of the p_T spectrum. The authors of Reference 129 found that introducing bulk viscosity improves the fit to single-particle p_T spectra without spoiling the quality of the fit for the elliptic flow v_2, but the optimal value of η/s for the matter produced in LHC heavy ion collisions changes from 0.16 to 0.095.

In addition to providing a more accurate description of the systematic dependences of the measured flow coefficients upon averaging, an event-by-event analysis of a large ensemble of events with fluctuating initial conditions also makes it possible to compare hydrodynamic calculations of the distributions of the v_n coefficients to experimental distributions. It turns out that the distribution $\delta v_n \equiv (v_n - \langle v_n \rangle)/\langle v_n \rangle$ is largely independent of the hydrodynamic transport coefficients but is instead sensitive to the initial shape of the energy density (**Figure 6*c***) (130), including its lumpiness. Hence, these distributions are an excellent way to constrain the hydrodynamic initial conditions, after which other observables can be used with greater confidence to constrain transport coefficients such as η/s. The correlations between different event plane angles Ψ_n also turn out to be a useful event-by-event observable. These correlations not only are sensitive to the average η/s during the hydrodynamic evolution but also can begin to constrain different hypotheses for the temperature dependence of η/s (107).

The question of how conclusions about the shear viscosity depend on the model used is an essential one, especially considering the uncertainty in the initial profiles as well as in the bulk viscosity and the temperature dependence of the shear viscosity. It has recently become possible to study the model dependence more systematically by doing a Bayesian analysis over a space of model parameters that include most models available, with recent estimates obtained via fitting to many different kinds of data from both RHIC and the LHC giving $\eta/s \approx (0.07^{+0.05}_{-0.04}) + c(T - T_c)$ (131) for temperatures $T > T_c = 154$ MeV, with T_c corresponding to the crossover between QGP and hadrons, but where the constant c is at present only constrained to be between 0/GeV and 1.58/GeV.

The small value for the dimensionless shear viscosity ratio η/s is especially interesting. At weak coupling, this ratio is proportional to the ratio of the quasiparticle mean-free path to the mean spacing between quasiparticles. A larger value of the ratio means that momentum can more easily be transported over significant distances, which is what is required in order to dissipate shear stress into heat. And weaker coupling means larger values of this ratio. [For a gas of gluons, in fact $\eta \sim 1/g^4 \log(1/g)$ in the weakly coupled $g \to 0$ limit (132). Here, g is the QCD coupling, namely the QCD analog of e in electromagnetism.] At strong coupling, by contrast, each volume element of the QGP fluid is so strongly coupled to its neighbors that very little (net) momentum can be transferred to nearby fluid elements, meaning that velocity gradients remain, shear stress does not dissipate, and the specific viscosity is small. The measured value of η/s for QGP turns out to be so small, however, that the fluid cannot be described in terms of quasiparticles with mean-free paths since to do so would require mean-free paths that are smaller than $1/T$. Strikingly, for

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

infinitely strongly coupled quantum theories with a large number of degrees of freedom that are described by a holographic dual gravitation theory, it can be computed that $\eta/s = 1/4\pi \approx 0.08$ (30), which is conspicuously close to the (average) viscosity found in hydrodynamic calculations used to model the dynamics of droplets of QGP produced in heavy ion collisions. Although QCD itself is not known to have a holographic dual, this motivates using gauge theories that do have dual gravitational descriptions to model dynamics in heavy ion collisions, as we elaborate below.

5. THERMALIZATION, HYDRODYNAMIZATION, AND ISOTROPIZATION

The success of the hydrodynamic paradigm begs the questions of when, why, and how the colliding debris begins to be accurately described by hydrodynamics, which is to say, hydrodynamizes. We would like to understand the underlying physics behind these questions, preferably from a QCD point of view. A related question is whether, if hydrodynamics is applicable for most of the evolution, this implies that the QGP formed in a heavy ion collision also thermalizes fast. The answer need not be yes, since a thermal state, by definition, is locally isotropic and free of gradients, whereas hydrodynamics can work well when contributions that are first order in gradients are significant while those coming from higher-order gradients are small. [Note that recent research (6, 133, 134) has succeeded in resumming all orders in gradients in a simplified geometry.] We hence have to ask how the pressure anisotropy evolves during the hydrodynamic evolution.

The droplet of QGP formed in a heavy ion collision is expanding rapidly, and even though it has the smallest dimensionless specific viscosity η/s found in nature, the gradients present due to the fast expansion (initially mostly in the longitudinal direction) imply that the viscous corrections, which are first order in gradients, are in fact sizable. Indeed, when $v_z = z/t$, as for a boost-invariant velocity profile, it is clear that at early times the gradients of the velocity field are large. When extracting the pressure anisotropy, as in **Figure 7**, it can be seen that the gradient corrections are important for proper times between 0.3 and 6 fm/c and the plasma only becomes approximately isotropic after a proper time of roughly $\tau = 6$ fm/c (125).

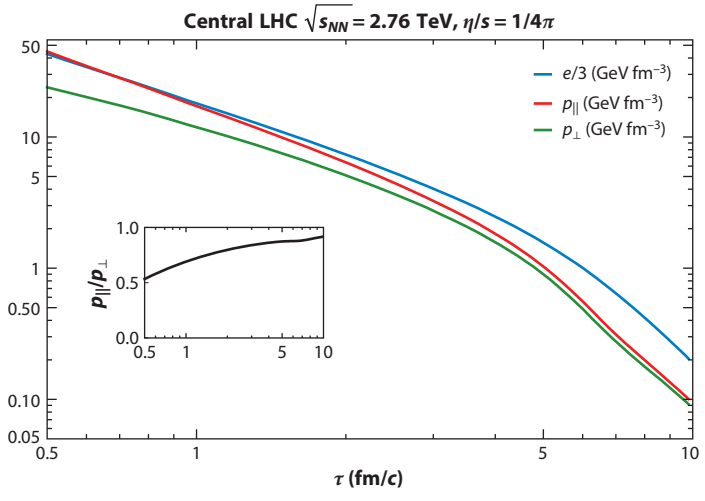

Figure 7

Typical energy density e, longitudinal pressure $p_{||}$, and transverse pressure $p_{\perp}$ as a function of proper time τ for a central collision at the LHC, at the center of the transverse plane. The inset shows the pressure anisotropy. Adapted from Reference 125.

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

So, even if hydrodynamics is indeed a good description around times as early as 0.5 fm/c, because of the significant initial gradients and the smallness of the specific viscosity, the fluid does not fully isotropize and hence thermalize before a much later time of around 6 fm/c. We say that the fluid hydrodynamizes rapidly, within a proper time of around 0.2–0.6 fm/c or 0.4–1.0 fm/c at LHC or RHIC energies, respectively (appropriately, these are typical starting times used in the hydrodynamic simulations described in Section 4), with this hydrodynamization followed by an extended period of hydrodynamic evolution with significant gradients in the fluid, before isotropization and complete thermalization at a substantially later time.

The estimates of the hydrodynamization time that we have quoted are often based upon assuming that when QGP hydrodynamizes in a heavy ion collision it does so without any initial transverse fluid velocity. This extra assumption is, however, unnecessary, and in fact, as we see in Section 6, below, all theoretical frameworks would predict the generation of some transverse flow already during the far-from-equilibrium, prehydrodynamization stage of the collision, which can hence resemble hydrodynamic evolution. The question of when the QGP formed in a heavy ion collision hydrodynamizes is hence intricately linked to how hydrodynamics becomes applicable, and in particular how much the far-from-equilibrium prehydrodynamic dynamics resembles hydrodynamics itself. We return to this in the next section.

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

6. INITIAL STAGE

The hydrodynamic model described above works well, perhaps surprisingly well, explaining many features of the particle spectra and the anisotropic flow coefficients. This poses three urgent questions: First, how does the debris left after a heavy ion collision evolve into an almost-perfect hydrodynamic fluid so fast? Second, how should this initial nonhydrodynamic stage in the dynamical evolution be described, and in what initial conditions for the hydrodynamic stage does this result? We sketch the present understanding of both questions. The third question is: How is entropy produced? This question provides a further reason for interest in the initial stage because almost all of the entropy produced during a heavy ion collision is produced before hydrodynamization: Because the specific viscosity of the hydrodynamic liquid is so small, very little additional entropy is produced in the later, longer hydrodynamic expansion. This means that the multiplicity of particles produced in the final state of a heavy ion collision is controlled by the dynamics occurring during its initial stage.

From a purely QCD point of view these questions are unfortunately hard to answer because nonperturbative real-time dynamics cannot be studied on the lattice. During the initial stage, and in particular during its earliest moments, many of the important scattering processes involve high-p_T transfer and hence can be described using pQCD. Soft, strongly coupled interactions are also important, in particular later in the initial stage as the matter hydrodynamizes, namely as it becomes a strongly coupled fluid. It is therefore reasonable that various authors have developed entirely weak coupling descriptions of the initial stage while at the same time other authors have modeled far-from-equilibrium dynamics nonperturbatively using holography. We consider the two approaches in turn.

In the context of pQCD, the starting point for the description of the initial stages involves the phenomenon of "saturation" in the gluon wave function of the incident nuclei (139). When colliding ions at higher and higher energies, the gluons that collide and end up near midrapidity after the collision are gluons from the PDF of the incident nuclei with smaller and smaller momentum fraction $x = p_z/P$, defined with respect to the momentum of the nucleon P. In a perturbative analysis, at large $Y = \log(1/x)$, the gluon PDF increases rapidly with increasing Y. At midrapidity in collisions at higher and higher energy, meaning smaller and smaller values of x, there will be

more and more gluons. However, this persists until, above some gluon density, gluon merging becomes as important as gluon splitting as x is decreased further. The occupation number of gluon modes in momentum space with this value of x and below is of order $1/\alpha_s$, and this component of the wave function of the nucleus is referred to as saturated. The typical p_T of these saturated gluons is referred to as the saturation scale Q_s, and the number density of these gluons per unit area in the transverse plane is then given by Q_s^2/α_s. The premise of the quantitative version of this analysis is that $\sqrt{\alpha_s(Q_s)}$ (where α_s is the running QCD coupling constant, which becomes small at high momentum transfer) is small when evaluated at the scale Q_s. It is at the saturation scale where we find the low-x gluons that dominate the interaction in an ultrarelativistic heavy ion collision (see Reference 140 for an accessible introduction).

't Hooft coupling: the smallness of α_s controls perturbative corrections in QCD at weak coupling, while the largeness of the 't Hooft coupling $\lambda \equiv 4\pi\alpha_s N_c$ controls finite coupling corrections in holography. $\alpha_s = 0.3$ corresponds to $\lambda \approx 11$, meaning that this coupling is neither close to nor far from both zero (perturbative methods) and infinity (holography)

The above perturbative analysis implies that just after a heavy ion collision one ends up at midrapidity with gluon modes with transverse momenta up to $\sim Q_s$ that are overoccupied. Making this analysis more quantitative leads to the conclusion that Q_s is of order 1 or 3 GeV for collisions at RHIC or the LHC—not so high as to make the assumptions of the perturbative treatment incontrovertible. Next, these gluons with transverse momenta of order Q_s radiate softer gluons and scatter with the growing bath of softer gluons until hydrodynamization is achieved (141, 142). These processes are somewhat involved, and can be described via weakly coupled classical field theory or an effective kinetic theory of weakly coupled partons in overlapping regions of parameter space. (The first can be used when there are modes with occupation numbers that are $\gg 1$; the second works for occupation numbers that are smaller than $1/\alpha_s$.) Plasma instabilities can play a role in the classical approach, although to leading order the classical evolution is self-similar due to the rapid longitudinal expansion (143). This expansion also drives the occupation numbers down, though, and at later times during the prehydrodynamic stage the effective kinetic theory must be used (see **Figure 8**). The earliest analyses of these processes yielded the conclusion that in the limit of very weak coupling the parametric dependence of the hydrodynamization time is $\tau_{\rm hydro} \gtrsim \alpha_s^{-13/5} Q_s^{-1}$ (141, 142). As numerical analyses of both the classical and the kinetic evolution have advanced, currently the most quantitative estimate is that, if one assumes $\alpha_s = 0.3$, the kinetic theory description of the energy density, transverse pressure, and longitudinal pressure hydrodynamizes after a time that is about, or even a little less than, 1 fm/c (144).

We know the specific viscosity is small and the coupling strong in the hydrodynamic liquid. This motivates exploring strongly coupled analyses of hydrodynamization as an alternative path to insights. The option that has been pursued most successfully has been to analyze the complete far-from-equilibrium initial stage assuming that the dynamics is strongly coupled throughout using holography (see the sidebar), which provides a dual gravitational description for certain gauge theories around infinitely strong coupling (20). This duality is truly remarkable, as it maps intractable real-time far-from-equilibrium nonperturbative quantum field theory problems onto equivalent, but tractable, computations within classical general relativity in anti–de Sitter space, a (4+1)-dimensional space-time with a negative cosmological constant. Due to the strong interactions, the hydrodynamization time can be much shorter than at weak coupling. An early hint of this was the discovery that small perturbations around an equilibrium thermal state (equivalent to exciting quasinormal modes of the dual black hole horizon) relax exponentially with characteristic time $\tau \sim 1/\pi T$ (145). Computations of the relaxation of many far-from-equilibrium disturbances to boost-invariant expanding flows (146, 147) have shown that hydrodynamization occurs within a time $\tau_{\rm hydro} \sim 0.7/T_{\rm hydro}$, where $T_{\rm hydro}$ is the temperature at which hydrodynamization occurs, and furthermore a remarkably broad applicability of the quasinormal mode analysis (148).

More advanced calculations permit the complete and rigorous simulation of the collision of sheets or discs of energy density in the infinitely strongly coupled super-Yang–Mills theory, which is a cousin of QCD with a dual holographic description (138, 150), from the moment of collision

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

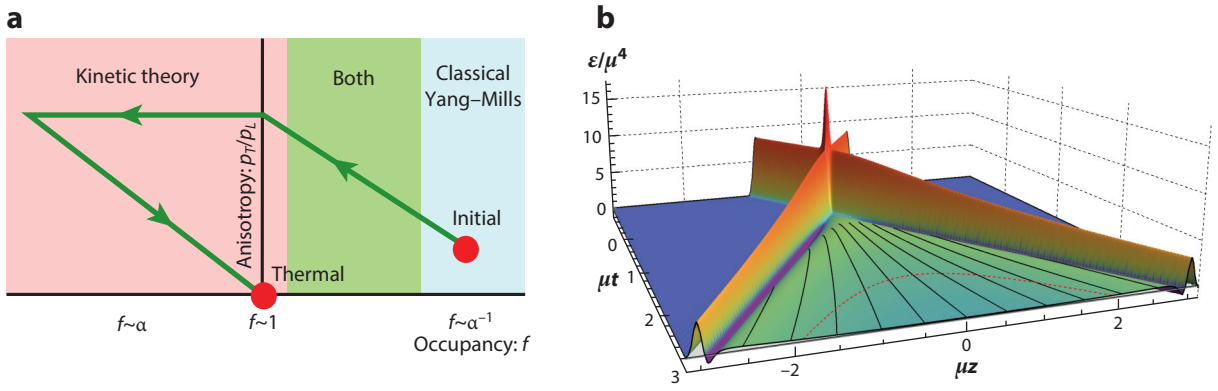

Figure 8

(*a*) Cartoon of how the pressure anisotropy (*vertical axis*) and f, the typical occupation number of modes with momentum $\sim Q_s$ in the gluon wave function, evolve during the initial stages of a heavy ion collision if one assumes that this can be described entirely at weak coupling. The weakly coupled dynamics can be described in terms of classical Yang–Mills fields if $f \gg 1$, and in terms of kinetic theory if $f \ll 1/\alpha_s$, meaning that as long as α_s is small enough there is a regime in which both descriptions are valid. The evolution begins with classical gluon fields at high occupancy $f \sim 1/\alpha_s$ and moderate pressure anisotropy, and proceeds to kinetic theory at low occupancy and large pressure anisotropy, after which the matter thermalizes, meaning that the pressures become isotropic and the occupancy of modes in the kinetic theory reaches $f \sim 1$. (*b*) Energy density as a function of time and longitudinal coordinate z in units of μ for a collision of two highly contracted parallel sheets of energy in strongly coupled super Yang–Mills theory colliding head on along the z direction at $t = z = 0$. Here, $N_c^2\mu^3/2\pi^2$ is the energy per transverse area of the incident sheets of energy, with N_c the number of colors in the theory. The energy density contains far-from-equilibrium regions, even including regions where energy density is negative and a rest frame cannot even be defined (135). After the red dashed line, the evolution of the plasma (*green*) is hydrodynamic within 5% accuracy (136). Panel *a* adapted from Reference 137. Panel *b* adapted from Reference 138.

through hydrodynamization and subsequent hydrodynamic expansion and cooling, including the development of radial and elliptic flow (35, 151). Such calculations enable direct and quantitative analyses of the hydrodynamization process after a collision, analyses that yield an affirmation of the hydrodynamic picture sketched above. In this context, the most important conclusion is that a system that begins with an ultrarelativistic collision can become hydrodynamic quickly, with collisions starting from a wide range of initial conditions yielding values for $\tau_{\rm hydro} T_{\rm hydro}$ between

HOLOGRAPHY

Holography started with a seminal paper by Maldacena (149) that provides an exact equivalence between certain string theories and certain (supersymmetric) gauge theories. In one direction, this exact equivalence has led to a much better understanding of quantum gravity by using gauge theory dynamics. In order to use the equivalence in the other direction, it is also possible to take the limit where string theory becomes a theory of ordinary classical gravity in a curved space-time with a negative cosmological constant and one extra dimension. In that case, the gauge theory has many colors and is infinitely strongly coupled. The equivalence can then provide reliable insights into complex dynamical questions in a strongly coupled gauge theory. Position in the extra dimension encodes the length scale of excitations in the gauge theory. For example, the position of a horizon in the gravitational space-time corresponds to $1/T$ in the gauge theory, where T is the temperature of the strongly coupled plasma with $\eta/s = 1/4\pi$ (20, 30). Because all aspects of a one-higher-dimensional gravitational theory are encoded in features of the gauge theory, the mapping is referred to as a holographic duality.

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

1/4 and 1, as well as a hydrodynamic fluid that is initially strongly anisotropic, with significant gradients. For the hydrodynamic calculation in **Figure 7**, solving the equation $\tau\, T(\tau) = 1$ leads to $\tau_{\rm hydro} \approx 0.35$ fm/c, whereby at that time $T_{\rm hydro} \approx 560$ MeV. Hydrodynamization may occur at an even earlier time and hotter temperature if $\tau_{\rm hydro} T_{\rm hydro} \lesssim 1$.

These calculations yield other qualitative insights about the prehydrodynamic stage in a collision. For example, they show that the far-from-equilibrium dynamics of the collision yields a hydrodynamic fluid whose longitudinal velocity profile is to a very good approximation boost invariant but whose energy (and entropy) density profile is far from boost invariant, taking on a shape that is approximately Gaussian in rapidity with a width of 0.98 (136). This is qualitatively in line with what is seen empirically, but is too narrow. Calculations have also been performed that follow the collision of strongly coupled sheets of energy density that carry "baryon number" (a conserved quantum number introduced by hand in the holographic gauge theory), showing that after the collision the "baryon number" distribution is also centered on midrapidity (152), rather than losing only a few units of rapidity as in QCD. This, and the narrowness of the energy/entropy distribution, is almost certainly a consequence of the fact that the gauge theory used in these calculations is not asymptotically free. The fact that in QCD the coupling is weak at the earliest moments of the collision is indeed important. This provides strong motivation for recent developments in the holographic framework, including collisions in theories that are not scale invariant (153) and that feature weaker-than-infinitely-strong coupling (154, 155), which give a shear viscosity that is larger than canonical, a nonzero bulk viscosity, and somewhat longer hydrodynamization times. It will be quite interesting to see how the distributions of energy, entropy, and "baryon number" change in these collisions. As a final example, these calculations permit the assessment of how much radial transverse flow already develops before hydrodynamization. While hydrodynamic gradients, of course, generate this flow later, early far-from-equilibrium evolution can do so too, and in fact, at strong coupling it is found that more prehydrodynamic flow is generated than would arise if this earliest epoch were instead hydrodynamic (35, 156). Similar results have also been obtained in an equivalent study at weak coupling (157).

A question that holographic calculations have not yet addressed (because to date they have not included any representation of the fact that in QCD the Lorentz-contracted incident nuclei are made of nucleons) is how the lumpiness of the energy density is distributed over the transverse plane at the start of the hydrodynamic stage in a heavy ion collision. In order to make comparisons to the increasingly precise measurements of azimuthal anisotropies and correlations described in Section 3, the fluctuations in the energy density across the transverse plane must be included. All phenomenological modeling includes the lumpiness coming from the initial positions of the participating (or "wounded") nucleons inside the colliding nuclei, via the Monte Carlo Glauber model described in Section 3. The simplest models just assume that each wounded nucleon contributes a Gaussian blob of energy density, but the precision of present data is sufficient that fluctuations that are somewhat smaller than a nucleon must be included in order to optimize model predictions. Refined models translate the density of wounded nucleons into a locally varying saturation scale, and then use this scale as a guide to placing fluctuating color sources in the transverse plane. These sources in turn drive the numerical evolution of classical Yang–Mills fields, whose stress–energy tensor is then used to initialize hydrodynamics (158, 159).

Much research remains to be done, including implementing an intermediate kinetic theory description, introducing lumpiness into the holographic calculations to provide a strong coupling benchmark, and in the long run testing the predictions of saturation calculations for the gluon distribution across the transverse plane in the incident nuclei against measurements at a future electron–ion collider (160). However, the best available calculations that begin with an initially lumpy energy density and follow its hydrodynamic evolution give an excellent simultaneous fit for

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

RHIC and the LHC to the probability distributions of the v_n's even for very off-central collisions (107).

7. JETS IN QUARK–GLUON PLASMA

In occasional heavy ion collisions, partons from the incident nuclei scatter off one another at very large momentum transfer, creating two or more quarks or gauge bosons with very high p_T (many tens of GeV at RHIC; as high as 100 or even 1,000 GeV at the LHC). When such a hard scattering occurs in a *pp* collision, each hard parton that is produced showers into a spray of softer partons within some irregular cone in momentum space, called a jet. Jet production and showering in vacuum are well described by pQCD (161). When a jet is produced in a heavy ion collision, the partons in the shower must plow through the droplet of QGP produced in the same collision. As this happens, the jet partons (*a*) lose energy and forward momentum, (*b*) pick up momentum transverse to their original direction, and (*c*) deposit energy and momentum into the droplet of QGP, creating a wake. The first of these phenomena is well established experimentally (162, 163), and there are strong indications of the third (70, 71). The second, which is referred to as momentum broadening since it can broaden the shape of a jet in momentum space, is apparent in all theoretical approaches but has not yet been observed experimentally (162, 164).

Nuclear modification factor: ratio of the number of some countable objects (e.g., jets defined via a specified reconstruction procedure with a given p_T, hadrons of a specified type with a given p_T) found in nuclear collisions divided by the (theoretical) value that would be expected from an analogous number of *pp* collisions without the presence of a medium

In the longer term, and in particular once we have high-statistics jet data at RHIC from the future sPHENIX detector (165) and from higher-luminosity running at the LHC in the early 2020s, the motivation for precision analyses of how jets are modified via their passage through QGP is that this may teach us about the inner workings of QGP. This is the closest we can ever come to probing QGP by doing a scattering experiment, and as discussed in Section 2, this is the best possible path toward addressing one of the big open questions in the field: How does a strongly coupled liquid emerge from an asymptotically free gauge theory? When the short-distance structure of QGP is resolved, it must consist of weakly coupled quarks and gluons. And yet, at length scales of order $1/T$ and longer, they become so strongly correlated as to form a liquid. Just as Rutherford found nuclei within atoms and Friedman, Kendall, and Taylor found quarks within protons by doing scattering experiments, in the longer term experimentalists hope to see the short-distance particulate structure of QGP by observing rare events in which a jet parton resolves, and scatters off, a parton in a droplet of QGP.

There are many physics questions (involving larger or more common effects) that are very interesting in their own right that must be understood quantitatively before realizing the vision of using jets as microscopes trained upon a droplet of QGP. This program is well under way, and could easily be the subject of an entire review of its own (see, e.g., 69, 91, 92). The most basic observation is that jets lose a substantial amount of energy, often 10 GeV or more, as they traverse a droplet of QGP. Noting that losing this amount of energy over only a few femtometers of distance corresponds to an enormous dE/dx, this provides a direct, and completely independent, confirmation that the matter produced in a heavy ion collision is strongly coupled. This energy loss can be observed in many ways, such as simply by counting the number of jets with a given (high) p_T: It is suppressed in heavy ion collisions relative to what would be observed in $N_{\rm coll}$ *pp* collisions, which is to say relative to the expected number of jets if there were no interaction with the medium, as explained in Section 3. This suppression is quantified by the nuclear modification factor (166)

$$R_{AA}(p_T) = \frac{dN^{AA}/dp_T}{\langle N_{\rm coll}\rangle dN^{pp}/dp_T}, \qquad 5.$$

where dN^{xx}/dp_T is the number of jets (or, in other contexts, particles of a specified type) produced in *AA* or *pp* collisions. Indeed, **Figure 9** shows a large suppression of these jets, especially for

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

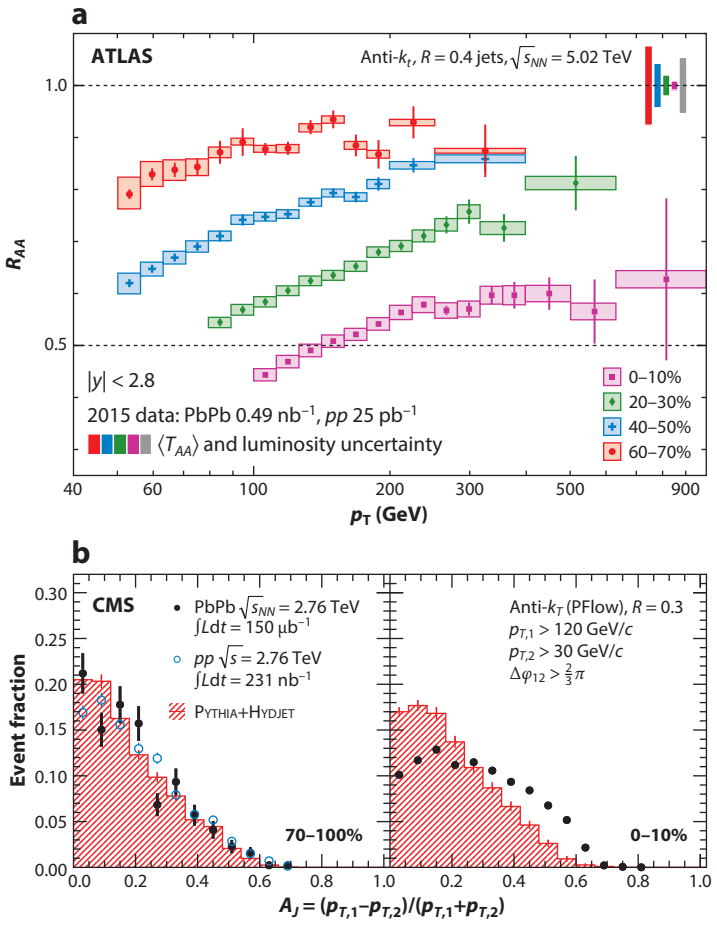

Figure 9

(*a*) The nuclear modification factor R_{AA} for jets for four different centralities as a function of jet transverse momentum p_T (163). (*b*) The dijet asymmetry A_J for pp collisions and for peripheral (*left*) and central (*right*) heavy ion collisions. The Monte Carlo–generated distribution (PYTHIA+HYDJET) shows the expected asymmetry if no nuclear effects were present (162).

central collisions in which the droplet of QGP that the jets need to traverse is the largest. A crucial check of this procedure is the fact that high-p_T colorless probes, such as γs or Z bosons, are indeed found to have $R_{AA} = 1$, as expected since they do not interact with QGP.

Throughout the study of jets it is important to realize that high-p_T jets are produced with a probability that drops very rapidly with increasing p_T. The production probability for jets produced at midrapidity with p_T values that are not within an order of magnitude of the beam energy scales roughly as $p_T{}^{-6}$ (167). The steepness of the energy spectrum implies that a small fractional jet energy loss corresponds to a large suppression in R_{AA} for jets. [As a contrafactual example, if we imagine that all jets lose 10% of their energy—that is, jets of 100 GeV started as 110-GeV jets—then, since 110-GeV jets are approximately $(100/110)^6 \approx 56\%$ rarer than 100-GeV jets, it follows that we would observe a nuclear modification factor of $R_{AA} \approx 0.56$.] In reality, different jets with the same initial energy lose very different amounts of energy, as discussed below, meaning that this argument must be made at the ensemble level. However, the conclusion is the same:

Because of the steepness of the jet energy spectrum the suppression in R_{AA} for jets is a sensitive measure of jet energy loss. Note that this argument does not apply in the same way to R_{AA} for high-p_T hadrons, as discussed below.

Figure 9*b* illustrates another way of seeing that jets lose energy, and also provides direct evidence that in a given event some jets lose more energy than others. This arises for two reasons. First, the characteristics of jets with a given energy vary quite considerably, and there are now a variety of theoretical arguments (at both weak and strong coupling) that indicate that a jet that fills a cone with a wide opening angle (and at weak coupling contains many partons) loses much more energy than a narrower jet with the same energy carried by fewer harder partons (164, 168–170). Because of the steepness of the jet spectrum described above, the ensemble of jets that comes out of the droplet of QGP will be dominated by those jets that lost relatively little energy, meaning that the jets that survive in a heavy ion collision with a given energy are likely to be those that started out the narrowest and are on average narrower than typical jets with the same energy in *pp* collisions. There is some evidence for this effect in measured jet shapes (171). Note that measuring R_{AA} for high-p_T hadrons is quite different: In both *pp* and *AA* collisions, a high-p_T hadron is statistically likely to come from a specific, unusual type of jet that contains one very hard parton and is very narrow; selecting (i.e., triggering on) hadrons therefore constitutes selecting an unusual sample of jets that lose less energy, and this selection effect becomes stronger at higher p_T. This is one reason that R_{AA} for hadrons rises at the highest p_T even though R_{AA} for jets remains comparably suppressed.

The second reason why some jets lose more energy than others is that when two or more jets are produced in a collision they each traverse different lengths of QGP. There is evidence for this effect in measurements of a v_2-like anisotropy for particles with high p_T that originate in jets (172): These jets typically lose less energy when moving along the short axis, as measured by the event-plane angle of the v_2 at low p_T (described in Section 4). All of this is to say that parton energy loss is a dominant effect contributing to the modification of many jet observables in heavy ion collisions as compared with *pp* collisions. $\mathrm{d}E/\mathrm{d}x$, the rate of parton energy loss in plasma, is parameterized in different ways for partons that are assumed to be traversing a weakly coupled plasma versus for those that are assumed to be traversing a strongly coupled plasma that behaves as it would in a holographic gauge theory (169, 173–175). In either case, present data are being used to constrain the magnitude of $\mathrm{d}E/\mathrm{d}x$, and in the near future, as the precision of the data improves further, it should become possible to differentiate between different choices for the T, x, and E dependence of $\mathrm{d}E/\mathrm{d}x$.

The energy and momentum "lost" by a jet in a heavy ion collision are, of course, not lost. We now know from experiment that they end up shared among many soft hadrons in the final state of the collision that are spread out over a wide range of angles, up to 60° or even 120°, around the jet direction (70, 71). This is certainly qualitatively consistent with a picture in which the jet excites a wake in the droplet of QGP, namely a region of moving and perhaps heated plasma behind the jet that carries the momentum in the jet direction "lost" by the jet. Like the unperturbed plasma, this wake becomes many soft hadrons after the droplet of QGP falls apart into hadrons. Because they carry net momentum in the jet direction, some of the hadrons from the wake must end up within what experimentalists observe as the jet (164, 176). This means that a quantitative understanding of the wake is a prerequisite for a quantitative understanding of the soft component of jets reconstructed in heavy ion collisions. Quantitative studies of the hydrodynamics of these wakes are now being done (177), and theorists should soon be able to do large-scale Monte Carlo calculations that track jet production in a hard scattering, jet showering, jet quenching, and the hydrodynamics of the specific wake produced by each specific jet (e.g., 178). Although full-scale calculations remain to be done, there are preliminary indications in some calculations (164, 179)

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

that the wakes made by jets shooting through the plasma do not have time to fully hydrodynamize, as they yield more 2–4-GeV hadrons and fewer 0–2-GeV hadrons than they would have if they had completely hydrodynamized (164). This is exciting as it raises the prospect of using jets, specifically their wakes, to obtain experimental access to the physics of hydrodynamization. In this way, analysis of jets in heavy ion collisions may yield insights into how QGP forms as a function of time in addition to, in the longer term, revealing how QGP emerges as we coarsen the resolution scale of the microscope with which we probe it. Achieving this longer-term goal will require having a quantitative understanding built upon precise data of energy loss (which results in an ensemble of narrower jets), jet wakes (which make jets as observed wider), and the accumulation of p_T by the jet partons via their soft interactions with the liquid QGP. Only then will it be possible to look for the rare (power law rare, not exponentially rare) hard scattering events in which a parton within a jet (or, even more rarely, a jet itself) gets kicked by a detectable angle as it resolves, and scatters off, a parton within the liquid.

It is at present an unfortunate aspect of studies of hard probes that experimentally only the final particles can be measured: It is in general not possible to directly compare a probe before and after it passes through QGP. This situation has recently improved through the selection of events with an energetic photon or Z boson and one or more energetic jets (67), which have the advantage that the photon or Z boson is unperturbed by the plasma and hence gives some probabilistic information about the energy of the jet or jets produced in the same event. Nevertheless, measuring the photon yields little information about the width of the jet, which plays an important and perhaps dominant role in determining how much energy it loses. In this context, it is exciting that experimentalists have recently begun to measure a host of different jet substructure observables, beyond the traditional shape, width, fragmentation function, and mass, that are constructed in various ways via grooming jets and obtaining operational measures of their substructure (180–182). Although this has not yet been realized, it may be possible to identify an observable that is (relatively) unmodified by the passage of a jet through QGP and that in pp collisions is in one-to-one correspondence with the width of the jet. If this potential is realized, by measuring other observables that are sensitive to energy loss as a function of this observable it will be possible to study the quenching of jets for which we have some information about what their widths would have been in the absence of quenching.

SUMMARY POINTS

1. We study heavy ion collisions to gain insight into perhaps the simplest form of complex matter, described by the fundamental laws of QCD. This superhot liquid filled the microseconds-old Universe, making it the first complex matter to form as well as the source of all protons and neutrons. Heavy ion collisions are Little Bangs, recreating droplets of Big Bang matter.
2. Within a time of order of 1 fm/c, the matter and entropy produced in a heavy ion collision form a droplet of strongly coupled QGP, evolving according to relativistic hydrodynamics with very small specific viscosity.
3. QGP is neither a collection of hadrons nor a nearly free gas of quarks and gluons. The colored quarks are free to diffuse and are not confined, but at the same time they are always very strongly coupled with their neighbors in the liquid.
4. Hydrodynamics converts spatial anisotropies into momentum anisotropy, yielding a direct experimental probe of both the spatial geometry, which is the source of the

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

anisotropies, and the viscosity, which seeks to dissipate them. The QGP is very lumpy when it forms. As it expands and cools hydrodynamically, the lumps smooth out while the resulting momentum anisotropies persist, because the specific viscosity of QGP is small.

5. The strongly coupled nature of QGP is also seen and probed at a broad range of length scales by jet quenching: the rapid loss of energy by highly energetic partons traversing QGP.
6. There is a wealth of experimental data, from longitudinal rapidity and p_T distributions to quark flavor, and from two-particle to multiparticle correlations, that are surprisingly similar across a variety of colliding systems spanning three orders of magnitude in both volume and energy.

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

FUTURE ISSUES

We frame the future issues facing the field, given the big picture that we have summarized above, as a series of big questions.

1. How does QGP form and hydrodynamize within 1 fm/c? What are the qualitative differences, if any, between the description of hydrodynamization in a heavy ion collision obtained by assuming a weakly coupled initial stage versus a strongly coupled holographic calculation? Note that perturbative calculations typically treat $\alpha_s = 0.3$ as small, while holographic calculations treat the corresponding 't Hooft coupling $\lambda \approx 11$ as large. What can we learn about the timescales and dynamics of hydrodynamization, and hence QGP formation, by analyzing the wakes that jets leave behind as they traverse a droplet of QGP?
2. What are the limits of the applicability of hydrodynamics? Can it be applied even to systems of size 1 fm or less? What is the smallest droplet of QGP that behaves hydrodynamically, and how does the answer to this question change at very high temperatures, where $\eta/s > 1$ and QGP is no longer a strongly coupled liquid?
3. How does a strongly coupled liquid emerge when QGP is analyzed with a spatial resolution of order $1/T$ or coarser, given that because QCD is asymptotically free what we see at much finer resolution is weakly coupled quarks and gluons? How can we use jets to see the inner workings of QGP and answer this question? If we can understand how QGP emerges from an asymptotically free gauge theory, can we use this understanding to learn general lessons about how complex forms of matter emerge from simple underlying laws?
4. How can we relate measurements of the gluon distribution in nuclei made at a future electron–ion collider to the distribution of the energy density across the transverse plane immediately after a heavy ion collision—quantitatively?
5. Can we obtain an experimental determination, even indirectly, of the temperature of the matter produced in a heavy ion collision at a time at which we can also determine its energy density? If so, we could obtain an experimental determination of the number of thermodynamic degrees of freedom, the quantity whose increase reflects the liberation of color above the crossover in the QCD phase diagram.

6. How do the hydrodynamics of QGP and the thermodynamics of its transition to hadronic matter as it cools change as QGP is doped with an excess of quarks over antiquarks? Is there a critical point in the region of the QCD phase diagram that heavy ion collisions can explore, or do all collisions that make QGP explore only a crossover in the phase diagram?
7. Can we explain the distribution of energy and entropy (particle multiplicity) as a function of rapidity in heavy ion collisions over a wide range of collision energies from first-principles computations? And, can we explain why hadronization produces hadrons in chemical equilibrium? More generally, why are many bulk phenomena so similar for AA, pA, pp, πA, and in some cases even e^+e^- collisions, over an enormous range of collision energies?
8. Is there color-superconducting quark matter at the centers of some or all neutron stars? This important question about the phase diagram of QCD cannot be addressed by heavy ion collisions; we hope that observations of binary neutron stars colliding and merging will help.

DISCLOSURE STATEMENT

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

ACKNOWLEDGMENTS

We are pleased to acknowledge helpful comments from Gian Michele Innocenti, Guilherme Milhano, Greg Ridgway, Raju Venugopalan, Jing Wang, Ryan Weller, and Bill Zajc. The writing of this review was supported by the US Department of Energy under contract number DE-SC0011090. W.S. is supported by VENI grant 680-47-458 from the Netherlands Organisation for Scientific Research (NWO).

LITERATURE CITED

1. Yagi K, Hatsuda T, Miake Y. *Camb. Monogr. Part. Phys. Nucl. Phys. Cosmol.* 23:1 (2005)
2. Florkowski W. *Phenomenology of Ultra-Relativistic Heavy-Ion Collisions*. Singapore: World Sci. (2010)
3. Wang XN, ed. *Quark–Gluon Plasma 5*. Singapore: World Sci. (2016)
4. Toia A. *J. Phys. G* 38:124007 (2011)
5. Heinz U, Snellings R. *Annu. Rev. Nucl. Part. Sci.* 63:123 (2013)
6. Romatschke P, Romatschke U. arXiv:1712.05815 [nucl-th] (2017)
7. van der Schee W. PhD thesis, Utrecht Univ., Utrecht, Neth. (2014)
8. Lee YJ, Yoon AS, Busza W. *MIT heavy ion event display: Pb+Pb 2.76 TeV*. Video, MIT, Cambridge, MA. **http://web.mit.edu/mithig/movies/LHCanmation.mov** (2018)
9. Busza W, Ledoux R. *Annu. Rev. Nucl. Part. Sci.* 38:119 (1988)
10. Busza W, Goldhaber AS. *Phys. Lett. B* 139:235 (1984)
11. Kharzeev DE, Liao J, Voloshin SA, Wang G. *Prog. Part. Nucl. Phys.* 88:1 (2016)
12. Baltz AJ, et al. *Phys. Rep.* 458:1 (2008)
13. Heisenberg W. *Nature* 164:65 (1949)
14. Hamilton J, Heitler W, Peng HW. *Phys. Rev.* 64:78 (1943)
15. Fermi E. *Prog. Theor. Phys.* 5:570 (1950)
16. Landau L. *Izv. Akad. Nauk Ser. Fiz.* 17:51 (1953)

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

17. Feynman RP. *Phys. Rev. Lett.* 23:1415 (1969)
18. Borsanyi S, et al. *Phys. Lett. B* 730:99 (2014)
19. Bazavov A, et al. (HotQCD Collab.) *Phys. Rev. D* 90:094503 (2014)
20. Casalderrey-Solana J, et al. arXiv:1101.0618 [hep-th] (2014)
21. Paquet JF, et al. *Phys. Rev. C* 93:044906 (2016)
22. Nucl. Sci. Advis. Comm. *The 2015 long-range plan for nuclear science*. Report, US Dep. Energy/Natl. Sci. Found., Washington, DC. **https://science.energy.gov/~/media/np/nsac/pdf/2015LRP/2015_LRPNS_091815.pdf** (2015)
23. Collins JC, Perry MJ. *Phys. Rev. Lett.* 34:1353 (1975)
24. Linde AD. *Rep. Prog. Phys.* 42:389 (1979)
25. Witten E. *Phys. Rev. D* 30:272 (1984)
26. Applegate JH, Hogan CJ. *Phys. Rev. D* 31:3037 (1985)
27. Karsch F. *Nucl. Phys. A* 698:199 (2002)
28. Aoki Y, et al. *Nature* 443:675 (2006)
29. Thomas D, et al. *Astrophys. J.* 430:291 (1994)
30. Policastro G, Son D, Starinets A. *Phys. Rev. Lett.* 87:081601 (2001)
31. Hubeny VE, Minwalla S, Rangamani M. arXiv:1107.5780 [hep-th] (2011)
32. Lin HW, Meyer HB. *Lattice QCD for Nuclear Physics*. Berlin: Springer (2015)
33. Borsanyi S, et al. *J. High Energy Phys.* 07:056 (2012)
34. Meyer HB. *Eur. Phys. J. A* 47:86 (2011)
35. van der Schee W. *Phys. Rev. D* 87:061901 (2013)
36. Chesler PM. *Phys. Rev. Lett.* 115:241602 (2015)
37. Chesler PM. *J. High Energy Phys.* 03:146 (2016)
38. Nagle JL, Zajc WA. *Annu. Rev. Nucl. Part. Sci.* 68:211 (2018)
39. Aduszkiewicz A, et al. (NA61/SHINE Collab.) *Eur. Phys. J. C* 77:671 (2017)
40. Luo X. *Nucl. Phys. A* 956:75 (2016)
41. Heuser JM. (CBM Collab.) *Eur. Phys. J. Web Conf.* 13:03001 (2011)
42. Toneev V. *Proc. Sci.* CPOD07:057 (2007)
43. Stephanov MA. *Prog. Theor. Phys. Suppl.* 153:139 (2004)
44. Rajagopal K, Wilczek F. In *At the Frontier of Particle Physics: Handbook of QCD*, ed. M Shifman, B Ioffe, 3:2061. Singapore: World Sci. (2000)
45. de Forcrand P. *Proc. Sci.* LAT2009:010 (2009)
46. Stephanov MA. *Phys. Rev. Lett.* 102:032301 (2009)
47. Athanasiou C, Rajagopal K, Stephanov M. *Phys. Rev. D* 82:074008 (2010)
48. Luo X. (STAR Collab.) *Proc. Sci.* CPOD2014:019 (2015)
49. Alford MG, Schmitt A, Rajagopal K, Schafer T. *Rev. Mod. Phys.* 80:1455 (2008)
50. Abbott B, et al. (Virgo Collab., LIGO Collab.) *Phys. Rev. Lett.* 119:161101 (2017)
51. Eskola KJ, Paakkinen P, Paukkunen H, Salgado CA. *Eur. Phys. J. C* 77:163 (2017)
52. Patrignani C, et al. (Part. Data Group) *Chin. Phys. C* 40:100001 (2016)
53. Alver B, Baker M, Loizides C, Steinberg P. arXiv:0805.4411 [nucl-ex] (2008)
54. Alver B, et al. (PHOBOS Collab.) *Phys. Rev. C* 83:024913 (2011)
55. Adam J, et al. (ALICE Collab.) *Phys. Lett. B* 772:567 (2017)
56. Busza W. *Acta Phys. Polon. B* 8:333 (1977)
57. Miller ML, Reygers K, Sanders SJ, Steinberg P. *Annu. Rev. Nucl. Part. Sci.* 57:205 (2007)
58. Arsene I, et al. (BRAHMS Collab.) *Nucl. Phys. A* 757:1 (2005)
59. Back B, et al. (PHOBOS Collab.) *Nucl. Phys. A* 757:28 (2005)
60. Adams J, et al. (STAR Collab.) *Nucl. Phys. A* 757:102 (2005)
61. Adcox K, et al. (PHENIX Collab.) *Nucl. Phys. A* 757:184 (2005)
62. Armesto N, Scomparin E. *Eur. Phys. J. Plus* 131:52 (2016)
63. Foka P, Janik MA. *Rev. Phys.* 1:154 (2016)
64. Foka P, Janik MA. *Rev. Phys.* 1:172 (2016)

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

65. Azaiez F, et al. *NuPECC long range plan 2017: perspectives in nuclear physics*. Report, NuPECC/Eur. Sci. Found., Strasbourg, Fr. **http://www.esf.org/fileadmin/user_upload/esf/Nupecc-LRP2017.pdf** (2017)
66. Aad G, et al. (ATLAS Collab.) *Phys. Rev. Lett*. 110:022301 (2013)
67. Chatrchyan S, et al. (CMS Collab.) *J. High Energy Phys*. 03:022 (2015)
68. Aad G, et al. (ATLAS Collab.) *Phys. Rev. C* 93:034914 (2016)
69. Connors M, Nattrass C, Reed R, Salur S. arXiv:1705.01974 [nucl-ex] (2017)
70. Collab CMS. arXiv:1803.00042 [nucl-ex] (2018)
71. Khachatryan V, et al. (CMS Collab.) *J. High Energy Phys*. 01:006 (2016)
72. Brenner AE, et al. *Phys. Rev. D* 26:1497 (1982)
73. Barton DS, et al. *Phys. Rev. D* 27:2580 (1983)
74. Arsene IC, et al. (BRAHMS Collab.) *Phys. Lett. B* 677:267 (2009)
75. Adamczyk L, et al. (STAR Collab.) *Phys. Rev. C* 96:044904 (2017)
76. Elias JE, et al. (E178 Collab.) *Phys. Rev. D* 22:13 (1980)
77. Back B, et al. (PHOBOS Collab.) *Phys. Rev. C* 74:021902 (2006)
78. Müller B, Rajagopal K. *Eur. Phys. J. C* 43:15 (2005)
79. Abbas E, et al. (ALICE Collab.) *Phys. Lett*. B 726:610 (2013)
80. Back B, et al. (PHOBOS Collab.) *Phys. Rev. Lett*. 91:052303 (2003)
81. Aamodt K, et al. (ALICE Collab.) *Phys. Rev. Lett*. 105:252301 (2010)
82. Aad G, et al. (ATLAS Collab.) *Phys. Lett*. B 710:363 (2012)
83. Adam J, et al. (ALICE Collab.) *Phys. Rev. Lett*. 116:222302 (2016)
84. Busza W. *Nucl. Phys. A* 854:57 (2011)
85. Gelis F, Stasto AM, Venugopalan R. *Eur. Phys. J. C* 48:489 (2006)
86. Chatrchyan S, et al. (CMS Collab.) *Phys. Rev. C* 84:024906 (2011)
87. Chatrchyan S, et al. (CMS Collab.) *Eur. Phys. J. C* 72:2012 (2012)
88. Aamodt K, et al. (ALICE Collab.) *Phys. Lett. B* 708:249 (2012)
89. Li W. *Nucl. Phys. A* 967:59 (2017)
90. Braun-Munzinger P, Koch V, Schäfer T, Stachel J. *Phys. Rep*. 621:76 (2016)
91. Mehtar-Tani Y, Milhano JG, Tywoniuk K. *Int. J. Mod. Phys. A* 28:1340013 (2013)
92. Qin GY, Wang XN. *Int. J. Mod. Phys. E* 24:1530014 (2015)
93. Koch P, Müller B, Rafelski J. *Int. J. Mod. Phys. A* 32:1730024 (2017)
94. Becattini F, et al. *Phys. Rev. Lett*. 111:082302 (2013)
95. Andronic A, Braun-Munzinger P, Redlich K, Stachel J. arXiv:1710.09425 [nucl-th] (2017)
96. Abelev B, et al. (ALICE Collab.) *Phys. Rev. C* 88:044910 (2013)
97. Back B, et al. (PHOBOS Collab.) *Phys. Rev. C* 70:051901 (2004)
98. Rajagopal K. arXiv:hep-ph/9504310 [hep-ph] (1995)
99. Sirunyan AM, et al. (CMS Collab.) arXiv:1706.05984 [hep-ex] (2017)
100. Zhang BW, Ko CM, Liu W. *Phys. Rev. C* 77:024901 (2008)
101. Khachatryan V, et al. (CMS Collab.) *Eur. Phys. J. C* 76:372 (2016)
102. Busza W. *Acta Phys. Polon. B* 35:2873 (2004)
103. Adam J, et al. (ALICE Collab.) *Phys. Lett. B* 760:720 (2016)
104. Adam J, et al. (ALICE Collab.) *Nat. Phys*. 13:535 (2017)
105. Bozek P, Broniowski W. *Phys. Rev. C* 88:014903 (2013)
106. Aaboud M, et al. (ATLAS Collab.) *Phys. Rev. C* 96:024908 (2017)
107. Niemi H, Eskola KJ, Paatelainen R. *Phys. Rev. C* 93:024907 (2016)
108. Aamodt K, et al. (ALICE Collab.) *Phys. Rev. Lett*. 107:032301 (2011)
109. Aad G, et al. (ATLAS Collab.) *J. High Energy Phys*. 1311:183 (2013)
110. Rischke D, Levin G. *Nucl. Phys. A* 750:1 (2005)
111. BRAHMS Collab., PHOBOS Collab., STAR Collab., PHENIX Collab. *Nucl. Phys. A* 757:1 (2005)
112. Molnar D, Gyulassy M. *Nucl. Phys. A* 697:495 (2002); Erratum. *Nucl. Phys. A* 703:893 (2002)
113. Xu Z, Greiner C. *Phys. Rev. Lett.* 100:172301 (2008)
114. Lin ZW, et al. *Phys. Rev. C* 72:064901 (2005)
115. Ackermann KH, et al. (STAR Collab.) *Phys. Rev. Lett*. 86:402 (2001)

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

116. Ollitrault JY, Poskanzer AM, Voloshin SA. *Phys. Rev. C* 80:014904 (2009)
117. Adams J, et al. (STAR Collab.) *Phys. Rev. Lett.* 92:052302 (2004)
118. Chatrchyan S, et al. (CMS Collab.) *Phys. Rev. C* 89:044906 (2014)
119. Huovinen P, Petreczky P. *Nucl. Phys. A* 837:26 (2010)
120. Pratt S, Sangaline E, Sorensen P, Wang H. *Phys. Rev. Lett.* 114:202301 (2015)
121. Baier R, Romatschke P. *Eur. Phys. J. C* 51:677 (2007)
122. Shen C, et al. In *Proceedings of the 26th International Conference on Ultrarelativistic Nucleus–Nucleus Collisions* (*Quark Matter 2017*), ed. U Heinz, O Evdokimov, P Jacobs, p. 796. Amsterdam: Elsevier (2017)
123. York MA, Moore GD. *Phys. Rev. D* 79:054011 (2009)
124. Cooper F, Frye G. *Phys. Rev. D* 10:186 (1974)
125. van der Schee W, Romatschke P, Pratt S. *Phys. Rev. Lett.* 111:222302 (2013)
126. Romatschke P, Romatschke U. *Phys. Rev. Lett.* 99:172301 (2007)
127. Schenke B, Jeon S, Gale C. *Phys. Rev. Lett.* 106:042301 (2011)
128. Alver B, Roland G. *Phys. Rev. C* 81:054905 (2010); Erratum. *Phys. Rev. C* 82:039903 (2010)
129. Ryu S, et al. *Phys. Rev. Lett.* 115:132301 (2015)
130. Niemi H, Denicol GS, Holopainen H, Huovinen P. *Phys. Rev. C* 87:054901 (2013)
131. Bernhard JE, et al. *Phys. Rev. C* 94:024907 (2016)
132. Arnold PB, Moore GD, Yaffe LG. *J. High Energy Phys.* 0011:001 (2000)
133. Heller MP, Spalinski M. *Phys. Rev. Lett.* 115:072501 (2015)
134. Denicol GS, Noronha J. arXiv:1711.01657 [nucl-th] (2017)
135. Arnold P, Romatschke P, van der Schee W. *J. High Energy Phys.* 1410:110 (2014)
136. Chesler PM, Kilbertus N, van der Schee W. *J. High Energy Phys.* 11:135 (2015)
137. Kurkela A. *Nucl. Phys. A* 956:136 (2016)
138. Casalderrey-Solana J, Heller MP, Mateos D, van der Schee W. *Phys. Rev. Lett.* 111:181601 (2013)
139. McLerran LD, Venugopalan R. *Phys. Rev. D* 49:2233 (1994)
140. Iancu E. In *Proceedings of the 2011 European School of High-Energy Physics* (*ESHEP 2011*), ed. C Grojean, M Mulders, p. 197. Geneva: CERN (2014)
141. Baier R, Mueller AH, Schiff D, Son DT. *Phys. Lett. B* 502:51 (2001)
142. Baier R, Mueller AH, Schiff D, Son DT. *Phys. Lett. B* 539:46 (2002)
143. Berges J, Boguslavski K, Schlichting S, Venugopalan R. *Phys. Rev. D* 89:074011 (2014)
144. Kurkela A, Zhu Y. *Phys. Rev. Lett.* 115:182301 (2015)
145. Horowitz GT, Hubeny VE. *Phys. Rev. D* 62:024027 (2000)
146. Chesler PM, Yaffe LG. *Phys. Rev. D* 82:026006 (2010)
147. Heller MP, Janik RA, Witaszczyk P. *Phys. Rev. Lett.* 108:201602 (2012)
148. Heller MP, Mateos D, van der Schee W, Trancanelli D. *Phys. Rev. Lett.* 108:191601 (2012)
149. Maldacena JM. *Int. J. Theor. Phys.* 38:1113 (1999)
150. Chesler PM, Yaffe LG. *Phys. Rev. Lett.* 106:021601 (2011)
151. Chesler PM, Yaffe LG. *J. High Energy Phys.* 10:070 (2015)
152. Casalderrey-Solana J, Mateos D, van der Schee W, Triana M. *J. High Energy Phys.* 09:108 (2016)
153. Attems M, et al. arXiv:1604.06439 [hep-th] (2016)
154. Grozdanov S, van der Schee W. *Phys. Rev. Lett.* 119:011601 (2017)
155. Waeber S, Schaefer A, Vuorinen A, Yaffe LG. *J. High Energy Phys.* 11:087 (2015)
156. Habich M, Nagle J, Romatschke P. arXiv:1409.0040 [nucl-th] (2014)
157. Keegan L, Kurkela A, Mazeliauskas A, Teaney D. *J. High Energy Phys.* 08:171 (2016)
158. Schenke B, Tribedy P, Venugopalan R. *Phys. Rev. Lett.* 108:252301 (2012)
159. Gale C, et al. *Phys. Rev. Lett.* 110:012302 (2013)
160. Accardi A, et al. *Eur. Phys. J. A* 52:268 (2016)
161. Sjostrand T, Mrenna S, Skands PZ. *Comput. Phys. Commun.* 178:852 (2008)
162. Chatrchyan S, et al. (CMS Collab.) *Phys. Lett. B* 712:176 (2012)
163. Aaboud M, et al. (ATLAS Collab.) arXiv:1805.05635 [nucl-ex] (2018)
164. Casalderrey-Solana J, et al. *J. High Energy Phys.* 03:135 (2017)
165. Adare A, et al. arXiv:1501.06197 [nucl-ex] (2015)
166. Adcox K, et al. (PHENIX Collab.) *Phys. Rev. Lett.* 88:022301 (2002)

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

167. Spousta M, Cole B. *Eur. Phys. J. C* 76:50 (2016)
168. Milhano JG, Zapp KC. *Eur. Phys. J. C* 76:288 (2016)
169. Chesler PM, Rajagopal K. *J. High Energy Phys.* 05:098 (2016)
170. Rajagopal K, Sadofyev AV, van der Schee W. *Phys. Rev. Lett.* 116:211603 (2016)
171. Khachatryan V, et al. (CMS Collab.) *J. High Energy Phys*. 11:055 (2016)
172. Sirunyan AM, et al. (CMS Collab.) *Phys. Lett. B* 776:195 (2018)
173. Chesler PM, Rajagopal K. *Phys. Rev. D* 90:025033 (2014)
174. Casalderrey-Solana J, et al. *J. High Energy Phys*. 10:019 (2014); Erratum. *J. High Energy Phys*. 09:175 (2015)
175. Casalderrey-Solana J, et al. *J. High Energy Phys.* 03:053 (2016)
176. Milhano JG, Wiedemann UA, Zapp KC. arXiv:1707.04142 [hep-ph] (2017)
177. Chen W, et al. *Phys. Lett. B* 777:86 (2018)
178. Cao S, et al. (JETSCAPE Collab.) *Phys. Rev. C* 96:024909 (2017)
179. Hulcher Z, Pablos D, Rajagopal K. arXiv:1707.05245 [hep-ph] (2017)
180. Sirunyan AM, et al. (CMS Collab.) arXiv:1708.09429 [nucl-ex] (2017)
181. Acharya S, et al. (ALICE Collab.) *Phys. Lett. B* 776:249 (2018)
182. Kauder K. (STAR Collab.) *Nucl. Phys. A* 967:516 (2017)

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.

Annual Review of Nuclear and Particle Science

Volume 68, 2018

Contents

Annu. Rev. Nucl. Part. Sci. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. 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. 2018.68:339-376. Downloaded from www.annualreviews.org
Access provided by Utrecht University on 02/05/19. For personal use only.