1

Opening remarks

The discovery in the late 1990s of the AdS/CFT correspondence, as well as its subsequent generalizations now referred to as the gauge/string duality, have provided a novel approach for studying the strong coupling limit of a large class of non-Abelian quantum field theories. In recent years, there has been a surge of interest in exploiting this approach to study properties of the plasma phase of such theories at nonzero temperature, including the transport properties of the plasma and the propagation and relaxation of plasma perturbations. Besides the generic theoretical motivation of such studies, many of the recent developments have been inspired by the phenomenology of ultra-relativistic heavy ion collisions. Inspiration has acted in the other direction too, as properties of non-Abelian plasmas that were determined via the gauge/string duality have helped to identify new avenues in heavy ion phenomenology. There are many reasons for this at-first-glance surprising interplay among string theory, finite-temperature field theory, and heavy ion phenomenology, as we shall see throughout this book. Here, we anticipate only that the analysis of data from the Relativistic Heavy Ion Collider (RHIC) had emphasized the importance, indeed the necessity, of developing strong coupling techniques for heavy ion phenomenology. Now, this case is further strengthened by data from the CERN Large Hadron Collider (LHC). For instance, in the calculation of an experimentally accessible transport property, the dimensionless ratio of the shear viscosity to the entropy density, weak and strong coupling results turn out to differ not only quantitatively but parametrically, and data favor the strong coupling result. Strong coupling presents no difficulty for lattice-regularized calculations of QCD thermodynamics, but the generalization of these methods beyond static observables to characterizing transport properties has well-known limitations. Moreover, these methods are quite unsuited to the study of the many and varied time-dependent problems that heavy ion collisions are making experimentally accessible. It is in this context that the very different suite of opportunities provided by gauge/string calculations of strongly coupled plasmas have started to provide a complementary source of insights for heavy ion phenomenology. Although these new methods come with limitations of their own, the results are obtained from first-principles calculations in non-Abelian field theories at nonzero temperature.

This book aims to provide an introductory exposition of the results obtained from the interplay between gauge/string duality, lattice QCD and heavy ion phenomenology within the past decade. It is written in a form accessible to graduate students seeking to enter this field of research from any direction. At the same time, it includes a comprehensive coverage that will be beneficial to established researchers from either community. It is a book about a newly emerging research field at the intersection of two domains that were until recently completely separate. As such, it does not attempt to cover the many aspects of these research fields that are of interest in their own terms, focusing on those aspects of each field that are relevant for understanding their new interplay. The introductions to heavy ion phenomenology (Chapter 2) and lattice QCD (Chapter 3) are thus written for beginners in these fields (whether graduate students or experienced string theorists) and they focus mainly on topics that have recently made contact with techniques from gauge/string duality. Analogously, Chapters 4 and 5 provide a targeted introduction to the principles behind the gauge/string duality with a focus on those aspects relevant for calculations at nonzero temperature. These chapters are for beginners too, again whether these beginners are graduate students or experts in heavy ion physics or lattice QCD. With the groundwork on both sides in place, we then proceed with a comprehensive exposition of gauge/string calculations of bulk thermodynamic and hydrodynamic properties (Chapter 6); of far-from-equilibrium dynamics and its late-time evolution to hydrodynamics (Chapter 7); of the propagation of probes (heavy or energetic quarks, and quark–antiquark pairs) through a strongly coupled non-Abelian plasma and the excitations of the plasma that result (Chapter 8); and a detailed analysis of mesonic bound states and spectral functions in a deconfined plasma (Chapter 9).

This book aims at covering the main developments of the interplay between hot QCD, heavy ion phenomenology and the gauge/string duality in a way that enables the reader to follow also other parts of the literature on applications of gauge/string duality to heavy ion phenomenology and hot QCD. We aimed at a comprehensive exposition but we had to make choices on what to cover. One important decision was to focus on insights that have been obtained from calculations that are directly rooted in quantum field theories analyzed via the gauge/string duality. Consequently, we have omitted the so-called AdS/QCD approach that aims to optimize ansätze for gravity duals that do not correspond to known field theories, in order to best incorporate known features of QCD in the gravitational description [343, 304, 518, 55, 201, 415, 416, 417, 508, 115, 689, 181, 182]. We are confident, however, that a reader of this book will be well-positioned to understand the motivations for research in this vast subject, as well as the main tools and techniques used in this research. Similarly, a reader of this book will have learned the tools needed to follow the broad and rapidly expanding range of string theory approaches in which dual gravitational descriptions are being developed not just for the strongly coupled plasma and its properties, that we focus on in this book, but also for the dynamics of how the plasma forms, equilibrates, expands, and cools after a collision. Our discussion of this vast topic in Chapter 7 highlights only a fraction of the many developments. Last but not least, we have omitted any discussion of the physics of saturation in QCD and its application to understanding the initial conditions for heavy ion collisions [374, 431, 432, 630, 128, 34, 300]. Here, our main reason was that a self-contained introduction to this subject in QCD would have required significant space, while its connection to the gauge/string duality rests at present on relatively few tentative, albeit very interesting, works. In short, neither this book, nor the list of omissions presented in this paragraph, should be regarded as being complete.

The existence of a textbook is a hallmark for the development of a research topic into a research field. There are several good textbooks in the field of finite temperature and lattice QCD, the field of heavy ion phenomenology and the field of gauge/string duality. However, aspects of the intersection between these fields have been covered so far only in scientific reviews, for example focussing on the techniques for calculating finite-temperature correlation functions of local operators from the gauge/string duality [749], or on the phenomenological aspects of perfect fluidity and its manifestation in different systems, including the quark–gluon plasma produced in heavy ion collisions and strongly coupled fluids made of trapped fermionic atoms that are more than twenty orders of magnitude colder [730]. There are also a number of shorter topical reviews that provide basic discussions of the duality and its most prominent applications in the context of heavy ion phenomenology [672, 601, 340, 398, 410]. We hope that beyond serving as an overview of what has been achieved already in this newly emerging field, our book will serve as a springboard for great achievements yet to come, in particular from its readers.

2

A heavy ion phenomenology primer

What macroscopic properties of matter emerge from the fundamental constituents and interactions of a non-Abelian gauge theory? The study of ultra-relativistic heavy ion collisions addresses this question for the theory of the strong interaction, Quantum Chromodynamics, in the regime of extreme energy density. To do this, heavy ion phenomenologists employ tools developed to identify and quantify collective phenomena in collisions that have many thousands of particles in their final states. Generically speaking, these tools quantify deviations with respect to benchmark measurements (for example in proton–proton and proton–nucleus collisions) in which collective effects are absent. In this chapter, we provide details for three cases of current interest: (i) the characterization of azimuthally anisotropic flow, which teaches us how soon after the collision matter moving collectively is formed and which allows us to constrain the value of the shear viscosity of this matter; (ii) the characterization of jet quenching, which teaches us how this matter affects and is affected by a high-velocity colored particle plowing through it; and (iii) the characterization of the suppression of quarkonium production, which has the potential to teach us about the temperature of the matter and of the degree to which it screens the interaction between colored particles.

2.1 General characteristics of heavy ion collisions

In a heavy ion collision experiment, large nuclei, such as gold (at RHIC) or lead (at the CERN SPS and LHC), are collided at an ultra-relativistic center of mass energy $\sqrt{s}$. The reason for using large nuclei is to create as large a volume as possible of matter at a high energy density, to have the best chance of discerning phenomena or properties that characterize macroscopic amounts of strongly interacting matter. In contrast, in energetic elementary collisions (say electron–positron collisions but to a good approximation also in proton–proton collisions) one may find many hadrons in the final state but these are understood to result from a few initial partons that each fragment rather than from a macroscopic volume of interacting matter. Many years ago Phil Anderson coined the phrase "more is different" to emphasize that macroscopic volumes of (in his case condensed) matter manifest qualitatively new phenomena, distinct from those that can be discerned in interactions among few elementary constituents and requiring distinct theoretical methods and insights for their elucidation [53]. Heavy ion physicists do not have the luxury of studying systems containing a mole of quarks, but by using the heaviest ions that nature provides they go as far in this direction as is possible.

The purpose of building accelerators that achieve heavy ion collisions at higher and higher $\sqrt{s}$ is simply to create matter at higher and higher energy density. A simple argument to see why this may be so arises upon noticing that in the center-of-mass frame we have the collision of two Lorentz-contracted nuclei, pancake-shaped, and increasing the collision energy makes these pancakes thinner. Thus, at $t = 0$ when these pancakes are coincident the entire energy of the two incident nuclei is found within a smaller volume for higher $\sqrt{s}$. This argument is overly simple, however, because not all of the energy of the collision is transformed into the creation of matter; much of it is carried by the debris of the two colliding nuclei that spray almost along the beam directions.

The question of how the initial state wave function of the colliding nuclei determines precisely how much matter, containing how much entropy, is produced soon after the collision, and consequently determines the number of particles in the final state, is a subject of intense theoretical interest. We shall not describe this branch of heavy ion phenomenology in any detail, but it is worth having a quantitative sense of just how many particles are produced in a typical heavy ion collision. In Fig. 2.1 we show the multiplicity of charged particles per unit pseudorapidity for RHIC collisions at four different values of $\sqrt{s}$. Recall that the pseudorapidity η is related to the polar angle θ measured with respect to the beam direction by $\eta = -\log\tan(\theta/2)$. Note also that, by convention, the incident ions in these collisions have a velocity such that individual nucleons colliding with that velocity would collide with a center of mass energy of $\sqrt{s}$. Since each gold nucleus has 197 nucleons and each Pb nucleus has 208 nucleons, the total center of mass energy in a heavy ion collision at the top RHIC energy is about 40 TeV and it rises to about 600 TeV at the current LHC energy. By integrating under the curve in Fig. 2.1, one finds that a heavy ion collision at top RHIC energy yields 5060 ± 250 charged particles [94, 95]. The multiplicity measurement is made by counting tracks, meaning that neutral particles (like π^0s and the photons they decay into) are not counted. So, the total number of hadrons is greater than the total number of charged particles. If all the hadrons in the final state were pions, and if the small isospin breaking introduced by the different number of protons and neutrons in a gold nucleus can be neglected, there would be equal numbers of π^+, π^- and π^0 meaning that the total

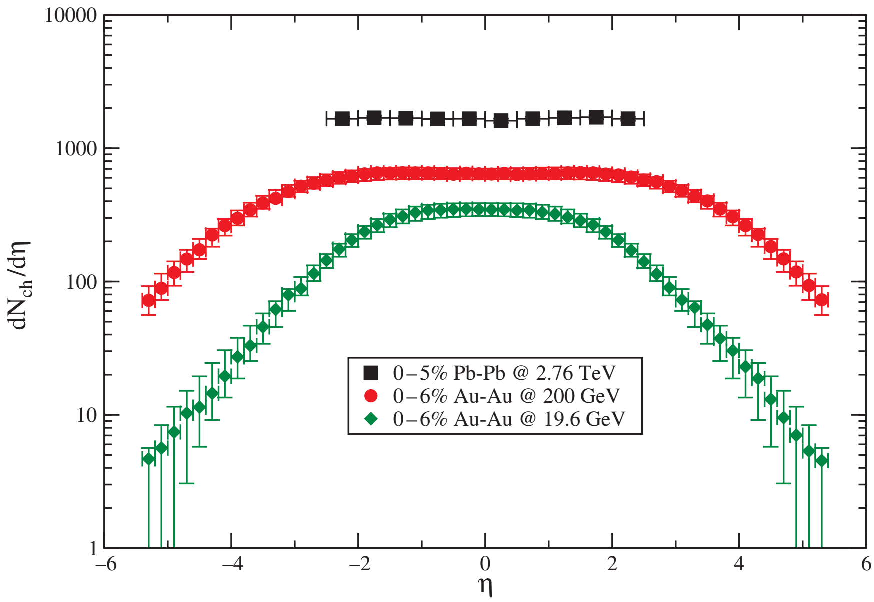

Figure 2.1 Charged particle multiplicity distributions for central nucleus–nucleus collisions (i.e. the 5% or 6% of collisions that have the smallest impact parameter) over more than two orders of magnitude in $\sqrt{s_{\mathrm{NN}}}$. Data taken from Refs. [263] and [94].

multiplicity would be 3/2 times the charged multiplicity. In reality, this factor turns out to be about 1.6 [96], meaning that heavy ion collisions at the top RHIC energy each produce about 8000 hadrons in the final state. At the LHC, the corresponding pseudorapidity distribution is known so far only in a range around mid-rapidity (see Fig. 2.1), with $dN_{\mathrm{ch}}/d\eta = 1584 \pm 4(\mathrm{stat}) \pm 76(\mathrm{sys})$ at $\eta = 0$ in the 5% or 6% of collisions with $\sqrt{s} = 2.76$ TeV that have the smallest impact parameter [4]. We see from Fig. 2.1 that this multiplicity grows with increasing collision energy by a factor of close to 2.5 from the top RHIC energy to LHC at $\sqrt{s} = 2.76$ GeV. The multiplicity per unit pseudorapidity is largest in a range of angles centered around $\eta = 0$, meaning $\theta = \pi/2$. Moreover, the distribution extends with increasing center of mass energy to larger values of pseudorapidity, so that the total event multiplicity at LHC is estimated to be a factor ~ 5 larger than at RHIC, lying in the ballpark of $\sim 25\,000$ charged particles in central collisions. The illustrations in Fig. 2.2 provide an impression of what collisions with these multiplicities look like.

The large multiplicities in heavy ion collisions indicate large energy densities, since each of these particles carries a typical (mean) transverse momentum of several hundred MeV. There is a simple geometric method due to Bjorken [165], that can be used to estimate the energy density at a fiducial early time, conventionally

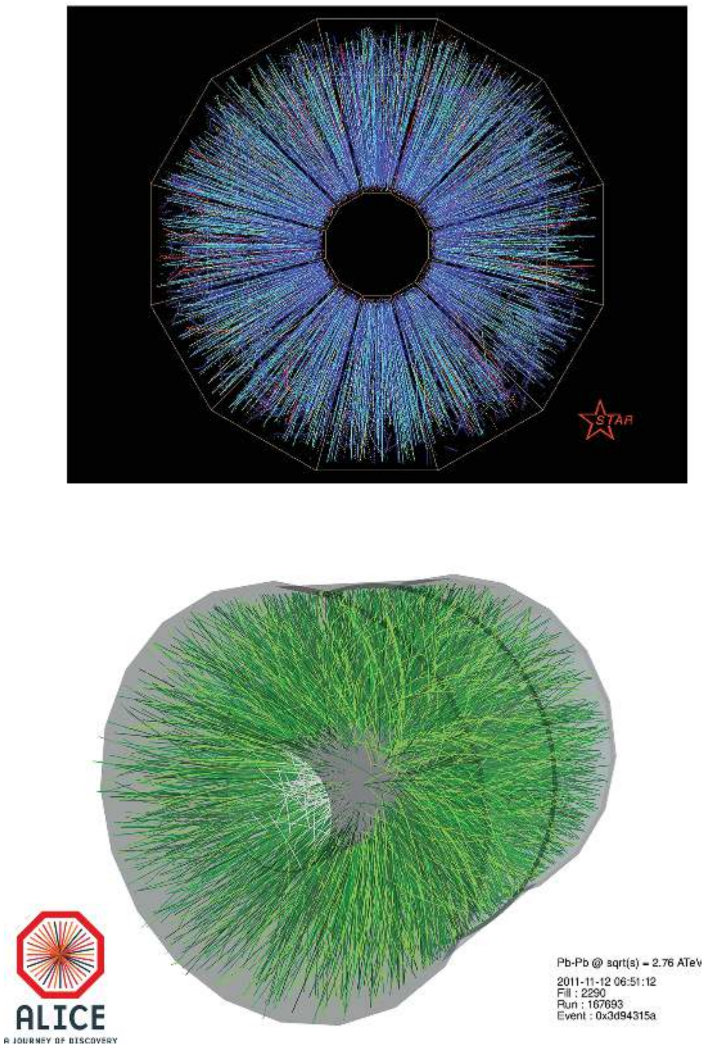

Figure 2.2 Event displays illustrating heavy ion collisions as seen by the STAR detector at RHIC (upper panel) and the ALICE detector at the LHC (lower panel). Nuclei (gold above; lead below) collided at the center of each image, and the resulting tracks made by those charged particles produced in the collision that pass through the STAR and ALICE time-projection chambers and the ALICE inner tracker are shown, projected onto the page in the upper image and in perspective in the lower image. Figures courtesy of Brookhaven National Laboratory (above) and the ALICE Collaboration and CERN (below).

chosen to be $\tau_0 = 1$ fm. The smallest reasonable choice of τ_0 would be the thickness of the Lorentz-contracted pancake-shaped nuclei, for instance $\sim (14\ \text{fm})/107$ at RHIC since gold nuclei have a radius of about 7 fm and the Lorentz factor is set by energy of the incident nucleons and their mass in the center-of-mass frame,

$\gamma \sim m_N/E$. But, at these early times ~ 0.1 fm the matter whose energy density one would be estimating would still be far from equilibrium. We shall see below that data on azimuthally anisotropic flow indicate that by ~ 1 fm after the collision, matter is flowing collectively like a fluid in local equilibrium. The geometric estimate of the energy density is agnostic about whether the matter in question is initial state partons that have not yet interacted and are far from equilibrium or matter in local equilibrium behaving collectively; because we are interested in the latter, we choose $\tau_0 = 1$ fm. Bjorken's geometric estimate can be written as

$$\varepsilon_{Bj} = \left.\frac{dE_T}{d\eta}\right|_{\eta=0} \frac{1}{\tau_0 \pi R^2}, \tag{2.1}$$

where $dE_T/d\eta$ is the transverse energy $\sqrt{m^2 + p_T^2}$ of all the particles per unit rapidity and $R \approx 7$ fm is the radius of the nuclei. The logic is simply that at time τ_0 the energy within a volume $2\tau_0$ in longitudinal extent between the two receding pancakes and πR^2 in transverse area must be at least $2dE_T/d\eta$, the total transverse energy between $\eta = -1$ and $\eta = +1$. At RHIC with $dE_T/d\eta \approx 800\,\text{GeV}$ [95], we obtain $\varepsilon_{Bj} \approx 5\,\text{GeV/fm}^3$. In choosing the volume in the denominator in the estimate (2.1) we neglected transverse expansion because $\tau_0 \ll R$. But, there is clearly an arbitrariness in the range of η used; if we had included particles produced at higher pseudorapidity (closer to the beam directions) we would have obtained a larger estimate of the energy density. Note also that there is another sense in which (2.1) is conservative. If there is an epoch after the time τ_0 during which the matter expands as a hydrodynamic fluid, and we shall later see evidence that this is so, then during this epoch its energy density drops more rapidly than $1/\tau$ because as it expands (particularly longitudinally) it is doing work. This means that by using $1/\tau$ to run the clock backwards from the measured final state transverse energy to that at τ_0 we have significantly underestimated the energy density at τ_0. It is striking that even though we have deliberately been conservative in making this underestimate, we have found an energy density that is about five times larger than the QCD critical energy density $\varepsilon_c \approx 1\,\text{GeV/fm}^3$, where the crossover from hadronic matter to quark–gluon plasma occurs, according to lattice calculations of QCD thermodynamics [129].

As shown in Fig. 2.3, the spectrum in a nucleus–nucleus collision extends to very high momentum, much larger than the mean. However, the multiplicity of high-momentum particles drops very fast with momentum, as a large power of p_T. We may separate the spectrum into two sectors. In the soft sector, spectra drop exponentially with $\sqrt{m^2 + p_T^2}$ as in thermal equilibrium. In the hard sector, spectra drop like power laws in p_T as is the case for hard particles produced by high momentum-transfer parton–parton collisions at $\tau = 0$. The bulk of the particles

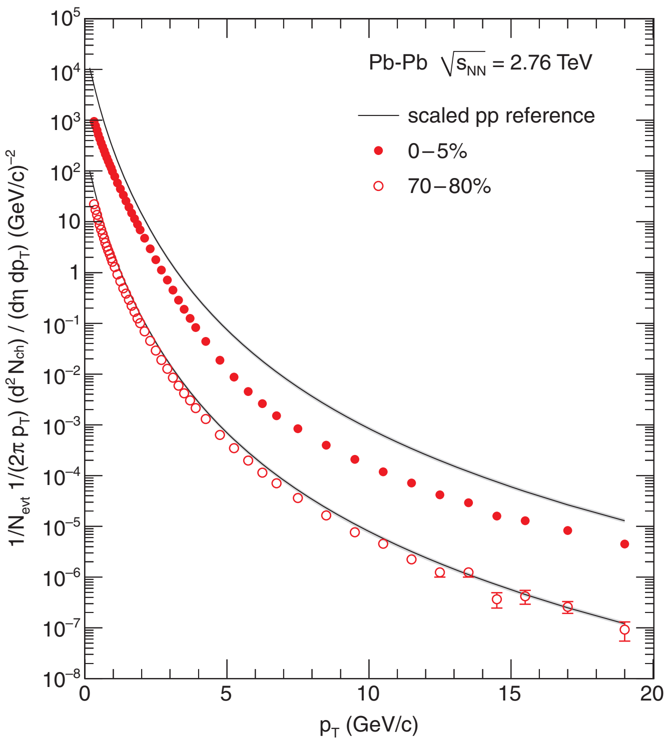

Figure 2.3 Charged particle spectrum as function of p_T in Pb+Pb collisions at LHC energy for nearly head-on (the 5% of collisions with the lowest impact parameter) and grazing collisions, compared to the corresponding spectrum in p+p collisions with an appropriately scaled normalization. Figure taken from Ref. [7].

have momenta in the soft sector; hard particles are rare in comparison. The separation between the hard and the soft sectors, which is by no means sharp, lies in the range of a few (say 3–6) GeV.

There are several lines of evidence that indicate that the soft particles in a heavy ion collision, which are the bulk of all the hadrons in the final state, have rescattered many times and come into local thermal equilibrium. The most direct approach comes via the analysis of the exponentially falling spectra of identified hadrons. Fitting a slope to these exponential spectra and then extracting an "effective temperature" for each species of hadron yields different "effective temperatures" for each species. This species dependence arises because the matter produced in a heavy ion collision expands radially in the directions transverse to the beam axis; perhaps explodes radially is a better phrase. This means that we should expect the p_T spectra to be a thermal distribution boosted by some radial velocity. If all hadrons are boosted by the same *velocity*, the heavier the hadron the more its

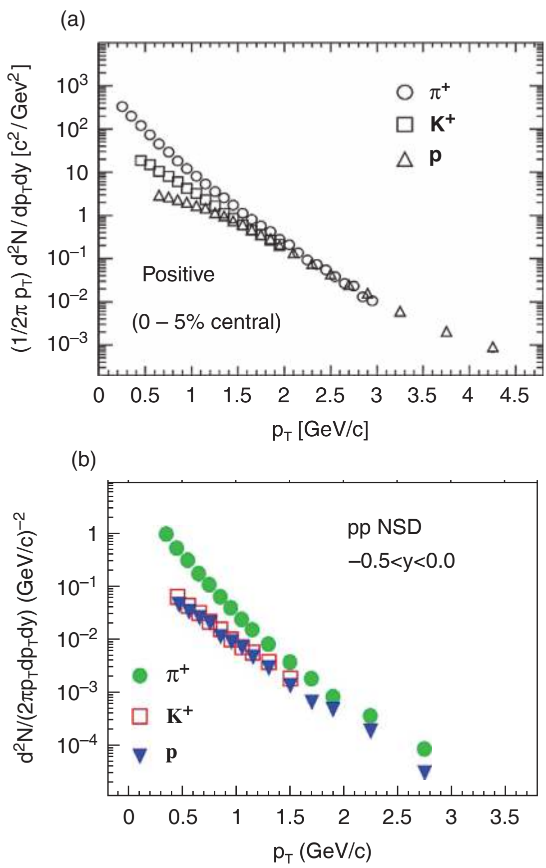

Figure 2.4 (a) Spectra for identified pions, kaons and protons as a function of p_T in head-on gold–gold collisions at top RHIC energy [25]. (b) Spectra for identified pions, kaons and protons as a function of p_T in (non-single-diffractive) proton–proton collisions at the same energy $\sqrt{s} = 200$ GeV [17].

momentum is increased by the radial boost. Indeed, what is found in data is that the effective temperature increases with the mass of the hadron species. This can be seen at a qualitative level in Fig. 2.4a: in the soft regime, the proton, kaon and pion spectra are ordered by mass, with the protons falling off most slowly with p_T, indicating that they have the highest effective temperature. Quantitatively, one uses the data for hadron species with varying masses to first extract the mass-dependence of the effective temperature, and thus the radial expansion velocity, and then to extrapolate the effective "temperatures" to the mass → zero limit, and in this way obtain a measurement of the actual temperature of the final state hadrons. This "kinetic freezeout temperature" is the temperature at the (very late) time at which the gas of hadrons becomes so dilute that elastic collisions between the hadrons cease, and the momentum distributions therefore stop changing as the system expands.