

Gauge/String Duality, Hot QCD and Heavy Ion Collisions

Jorge Casalderrey-Solana,¹ Hong Liu,² David Mateos,^{3,4}
Krishna Rajagopal,² and Urs Achim Wiedemann¹

¹ *Department of Physics, CERN, Theory Unit, CH-1211 Geneva*

² *Center for Theoretical Physics, MIT, Cambridge, MA 02139, USA*

³ *Institució Catalana de Recerca i Estudis Avançats (ICREA), Passeig Lluís Companys 23,
E-08010, Barcelona, Spain*

⁴ *Departament de Física Fonamental (FFN) & Institut de Ciències del Cosmos (ICC),
Universitat de Barcelona (UB), Martí i Franquès 1, E-08028 Barcelona, Spain*

jorge.casalderrey@cern.ch, hong_liu@mit.edu, dmateos@icrea.cat,
krishna@mit.edu, urs.wiedemann@cern.ch

Abstract

Over the last decade, both experimental and theoretical advances have brought the need for strong coupling techniques in the analysis of deconfined QCD matter and heavy ion collisions to the forefront. As a consequence, a fruitful interplay has developed between analyses of strongly-coupled non-abelian plasmas via the gauge/string duality (also referred to as the AdS/CFT correspondence) and the phenomenology of heavy ion collisions. We review some of the main insights gained from this interplay to date. To establish a common language, we start with an introduction to heavy ion phenomenology and finite-temperature QCD, and a corresponding introduction to important concepts and techniques in the gauge/string duality. These introductory sections are written for nonspecialists, with the goal of bringing readers ranging from beginning graduate students to experienced practitioners of either QCD or gauge/string duality to the point that they understand enough about both fields that they can then appreciate their interplay in all appropriate contexts. We then review the current state-of-the art in the application of the duality to the description of the dynamics of strongly coupled plasmas, with emphases that include: its thermodynamic, hydrodynamic and transport properties; the way it both modifies the dynamics of, and is perturbed by, high-energy or heavy quarks passing through it; and the physics of quarkonium mesons within it. We seek throughout to stress the lessons that can be extracted from these computations for heavy ion physics as well as to discuss future directions and open problems for the field.

Contents

1	Preface	5
2	A heavy ion phenomenology primer	8
2.1	General characteristics of heavy ion collisions	8
2.2	Elliptic flow	15
2.2.1	Introduction and motivation	15
2.2.2	The elliptic flow observable v_2 at RHIC	17
2.2.3	Calculating elliptic flow using (ideal) hydrodynamics	20
2.2.4	Comparing elliptic flow in heavy ion collisions and hydrodynamic calculations	27
2.3	Jet quenching	31
2.3.1	Single inclusive high-pt spectra and “jet” measurements	34
2.3.2	Analyzing jet quenching	37
2.4	Quarkonia in hot matter	44
3	Results from lattice QCD	50
3.1	The QCD Equation of State from the lattice	51
3.1.1	Flavor susceptibilities	54
3.2	Transport coefficients from the lattice	56
3.3	Quarkonium Spectrum from the Lattice	61
4	Introducing the gauge/string duality	70
4.1	Motivating the duality	70
4.1.1	An intuitive picture: Geometrizing the renormalization group flow	71
4.1.2	The large- N_c expansion of a non-Abelian gauge theory vs. the string theory expansion	72
4.1.3	Why AdS?	74
4.2	All you need to know about string theory	75
4.2.1	Strings	75
4.2.2	D-branes and gauge theories	78
4.2.3	D-branes as spacetime geometry	83

4.3	The AdS/CFT conjecture	85
5	General aspects of the duality	88
5.1	Gauge/gravity duality	88
5.1.1	UV/IR connection	89
5.1.2	Strong coupling from gravity	89
5.1.3	Symmetries	92
5.1.4	Matching the spectrum: the field/operator correspondence	92
5.1.5	Normalizable vs. non-normalizable modes and mass-dimension relation	93
5.2	Generalizations	96
5.2.1	Nonzero temperature and nonzero chemical potential	96
5.2.2	A confining theory	98
5.2.3	Other generalizations	99
5.3	Correlation functions of local operators	100
5.3.1	Euclidean correlators	100
5.3.2	Real-time thermal correlators	102
5.4	Wilson loops	105
5.4.1	Rectangular loop: vacuum	108
5.4.2	Rectangular loop: nonzero temperature	110
5.4.3	Rectangular loop: a confining theory	111
5.5	Introducing fundamental matter	113
5.5.1	The decoupling limit with fundamental matter	114
5.5.2	Models with fundamental matter	116
6	Bulk properties of strongly coupled plasma	119
6.1	Thermodynamic properties	122
6.1.1	Entropy, energy and free energy	122
6.1.2	Holographic susceptibilities	124
6.2	Transport properties	125
6.2.1	A general formula for transport coefficients	126
6.2.2	Universality of the Shear Viscosity	128
6.2.3	Bulk viscosity	132
6.2.4	Relaxation times and other 2nd order transport coefficients	133
6.3	Quasiparticles and spectral functions	136
6.3.1	Quasiparticles in perturbation theory	137
6.3.2	Absence of quasiparticles at strong coupling	142
7	Probing strongly coupled plasma	144
7.1	Parton energy loss via a drag on heavy quarks	145
7.1.1	Regime of validity of the drag calculation	147

7.2	Momentum broadening of a heavy quark	150
7.2.1	κ_T and κ_L in the $p \rightarrow 0$ limit	151
7.2.2	Direct calculation of the noise term	152
7.2.3	κ_T and κ_L for a moving heavy quark	156
7.2.4	Implications for heavy quarks in heavy ion collisions	160
7.3	Disturbance of the plasma induced by an energetic heavy quark	162
7.3.1	Hydrodynamic preliminaries	165
7.3.2	AdS computation	168
7.3.3	Implications for heavy ion collisions	172
7.3.4	Disturbance excited by a moving quarkonium meson	173
7.4	Stopping light quarks	174
7.5	Calculating the jet quenching parameter	176
7.6	Quenching a beam of strongly coupled synchrotron radiation	184
7.7	Velocity-scaling of the screening length and quarkonium suppression in a hot wind	188
8	Quarkonium mesons	196
8.1	Adding quarks to $\mathcal{N} = 4$ SYM	196
8.2	Zero temperature	198
8.2.1	D7-brane embeddings	198
8.2.2	Meson spectrum	202
8.3	Nonzero temperature	204
8.3.1	D7-brane embeddings	204
8.3.2	Thermodynamics of D7-branes	209
8.3.3	Quarkonium thermodynamics	214
8.4	Quarkonium mesons in motion and in decay	219
8.4.1	Spectrum and dispersion relations	220
8.4.2	Decay widths	225
8.4.3	Connection with the quark-gluon plasma	227
8.5	Black hole embeddings	230
8.5.1	Absence of quasiparticles	230
8.5.2	Meson spectrum from a spectral function	231
8.6	Two universal predictions	235
8.6.1	A meson peak in the thermal photon spectrum	235
8.6.2	A new mechanism of quark energy loss: Cherenkov emission of mesons	237
9	Concluding remarks and outlook	241
A	Green-Kubo formula for transport coefficients	245

B	Hawking temperature of a general black brane metric	247
C	Holographic renormalization, one-point functions and an example of a Euclidean two-point function	248

Chapter 1

Preface

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 non-zero 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 review. 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. 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-principle calculations in non-abelian field theories at non-zero temperature.

The present review aims at providing an overview of the results obtained from this interplay between gauge/string duality, lattice QCD and heavy ion phenomenology within the last decade in a form such that readers from either community can appreciate and understand the emerging synthesis. For the benefit of string theory practitioners, we begin in Sections 2

and 3 with targeted overviews of the data from and theory of heavy ion collisions, as well as of state-of-the art lattice calculations for QCD at non-zero temperature. We do not provide a complete overview of all aspects of these subjects that are of current interest in their own terms, aiming instead for a self-contained, but targeted, overview of those aspects that are keys to understanding the impacts of gauge/string calculations. In turn, Sections 4 and 5 are mainly written for the benefit of heavy ion phenomenologists and QCD practitioners. They provide a targeted introduction to the principles behind the gauge/string duality with a focus on those aspects relevant for calculations at non-zero temperature. With the groundwork on both sides in place, we then provide an in-depth review of the gauge/string calculations of bulk thermodynamic and hydrodynamic properties (Section 6); 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 (Section 7); and a detailed analysis of mesonic bound states and spectral functions in a deconfined plasma (Section 8).

The interplay between hot QCD, heavy ion phenomenology and the gauge/string duality has been a very active field of research in recent years, and there are already a number of other reviews that cover various aspects of these developments. In particular, there are reviews focussing on the techniques for calculating finite-temperature correlation functions of local operators from the gauge/string duality [1]; 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 [2] (we shall not review the connection to ultracold atoms here); and also shorter topical reviews [3–7] that provide basic discussions of the duality and its most prominent applications in the context of heavy ion phenomenology. The present review aims at covering a broader range of applications while at the same time providing readers from either the string theory or the QCD communities with the opportunity to start from square one on ‘the other side’ and build an understanding of the needed context.

Since any attempt at covering everything risks uncovering nothing, even given its length this review does remain limited in scope. Important examples of significant advances in this rapidly-developing field that we have not touched upon include: (i) the physics of saturation in QCD and its application to understanding the initial conditions for heavy ion collisions [8–14]; (ii) various field-theoretical approaches to understanding equilibration in heavy ion collisions [15–19]; and (iii) work from the string side on developing dual gravitational descriptions not just of the strongly-coupled plasma and its properties, as we do review, but of the dynamics of how the plasma forms, equilibrates, expands, and cools after a collision. This last body of literature includes the early work of Refs. [20–23], continues in modern form with the discovery of a dual gravitational description of a fluid that is cooling via boost-invariant longitudinal expansion [24] extends to the very recent discovery of the dual description of the collision of two finite-thickness sheets of energy density and the ensuing plasma formation [25], and includes much in between [26–43]. Also, we focus throughout 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 [44–56]. So, while we have high hopes that the material reviewed here will serve in coming years as a pathway into the field for new practitioners with experience in QCD or string theory,

we have not aimed for completeness. Indeed, we can already see now that this very active field continues to broaden and that adding new chapters to a review like this will soon be warranted.

Chapter 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 Section, we provide details for three cases of current interest: (i) the characterization of elliptic 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 via single inclusive hadron spectra and angular correlations, 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 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 [57]. Heavy ion




Figure 2.1: Left: Charged particle multiplicity distributions for collisions at RHIC with four different energies as a function of pseudorapidity [58]. Very recently, a charged particle multiplicity $dN_{\text{ch}}/d\eta = 1584 \pm 4(\text{stat}) \pm 76(\text{sys})$ has been measured at $\eta = 0$ in heavy ion collisions with $\sqrt{s} = 2.76$ TeV by the ALICE detector at the LHC [59]. Right: Charged particle spectrum as function of p_T in gold-gold collisions at top RHIC energy at different values of the impact parameter, with the top (second-from-bottom) curve corresponding to nearly head-on (grazing) collisions [60]. The bottom curve is data from proton-proton collisions at the same energy.

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. At present, however, calculations using gravity duals have not shed much light on the attendant issues and so we shall not describe this branch of heavy ion phenomenology in any detail. It is, however, worth having a quantitative sense of just how many particles are produced in a typical heavy ion collision. In the left panel of 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, $\sqrt{s} = 200$ GeV (which is the top energy achieved at RHIC) means that 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 200 GeV. Since each gold nucleus has 197 nucleons, the total center of mass energy in a heavy ion collision at the top RHIC energy is about 40 TeV. By integrating under the curve in the left panel of Fig. 2.1, one finds that such a collision yields 5060 ± 250 charged particles [61, 62]. The multiplicity measurement is made by counting tracks, meaning that neutral particles (like π^0 's 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 multiplicity would be 3/2 times the charged multiplicity. In reality, this factor turns out to be about 1.6 [58], meaning that heavy ion collisions at the top RHIC energy each produce about 8000 hadrons in the final state. We see from the left panel of Fig. 2.1 that this multiplicity grows with increasing collision energy, and we see that the multiplicity per unit pseudorapidity is largest in a range of angles centered around $\eta = 0$, meaning $\theta = \pi/2$. Heavy ion collisions at the LHC will have center of mass energies up to $\sqrt{s} = 5.5$ TeV per nucleon. Estimates for the charged particle multiplicity per unit pseudorapidity at this energy that were made before LHC heavy ion collisions began typically range from $dN_{\text{ch}}/d\eta \approx 1000$ to $dN_{\text{ch}}/d\eta \approx 3000$ [63]. This range can now be considerably tightened, now that we know that $dN_{\text{ch}}/d\eta = 1584 \pm 4(\text{stat}) \pm 76(\text{sys})$ at $\eta = 0$ in most central collisions (5 %) with $\sqrt{s} = 2.76$ TeV [59].

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 [64], that can be used to estimate the energy density at a fiducial early time, conventionally 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, namely $\sim (14 \text{ fm})/100$ since gold nuclei have a radius of about 7 fm. 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 elliptic flow data 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

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

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 $dE_T/d\eta \approx 800$ GeV [62], we obtain $\varepsilon_{Bj} \approx 5$ GeV/fm³. In choosing the volume in the denominator in the estimate (2.1) we neglected transverse expansion because $\tau_0 \ll R$.

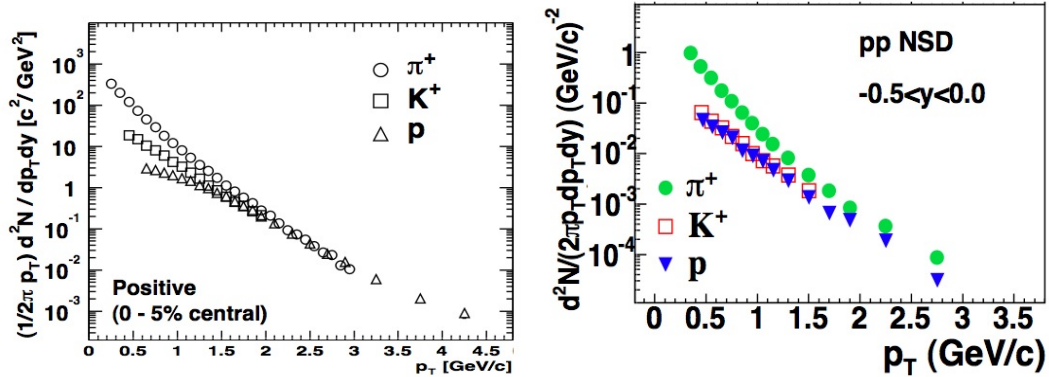


Figure 2.2: Left: Spectra for identified pions, kaons and protons as a function of p_T in head-on gold-gold collisions at top RHIC energy [66]. Right: Spectra for identified pions, kaons and protons as a function of p_T in (non single diffractive) proton-proton collisions at the same energy $\sqrt{s} = 200$ GeV [67].

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 $\varepsilon_c \approx 1 \text{ GeV}/\text{fm}^3$, where the crossover from hadronic matter to quark-gluon plasma occurs, according to lattice calculations of QCD thermodynamics [65].

As shown in the right panel of Fig. 2.1, 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 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 2-5) 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 interacted 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 *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 the left panel of Fig. 2.2: 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 species 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 \rightarrow 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. In heavy ion collisions at the top RHIC energy, models of the kinetic freezeout account for the data with freeze-out temperatures of ≈ 90 MeV and radial expansion velocities of $0.6c$ for collisions with the smallest impact parameters [68]. With increasing impact parameter, the radial velocity decreases and the freeze-out temperature increases. This is consistent with the picture that a smaller system builds up less transverse flow and that during its expansion it cannot cool down as much as a bigger system, since it falls apart earlier.

The analysis just described is unique to heavy ion collisions: in elementary electron-positron or proton-(anti)proton collisions, spectra at low transverse momentum may also be fit by exponentials, but the “temperatures” extracted in this way do not have a systematic dependence on the hadron mass. In fact, they are close to the same for all hadron species, as can be seen qualitatively in the right panel Fig. 2.2. Simply seeing exponential spectra and fitting a “temperature” therefore does not in itself provide evidence for multiple interactions and equilibration. Making that case in the context of heavy ion collisions relies crucially on the existence of a collective radial expansion with a common velocity for all hadron species.

Demonstrating that the final state of a heavy ion collision at the time of kinetic freezeout is a gas of hadrons in local thermal equilibrium is by itself not particularly interesting. It does, however, embolden us to ask whether the material produced in these collisions reaches local thermal equilibrium at an earlier time, and thus at a higher temperature. The best evidence for an affirmative answer to this question comes from the analysis of “elliptic flow” in collisions with nonzero impact parameter. We shall discuss this at length in the next subsection.

We close this subsection with a simpler analysis that lays further groundwork by allowing us to see back to a somewhat earlier epoch than that of kinetic freezeout. If we think of a heavy ion collision as a “little bang”, replaying the history of the big bang in a small volume and with a vastly accelerated expansion rate, then kinetic freezeout is the analogue of the (late) cosmological time at which photons and electrons no longer scatter off each other. We now turn to the analogue of the (earlier) cosmological epoch of nucleosynthesis, namely the time at which the composition of the final state hadron gas stops changing. Experimentalists can measure the abundance of more than a dozen hadron species, and it turns out that all




Figure 2.3: So-called thermal fit to different particle species. The relative abundance of different hadron species is well-described by a two-parameter grand canonical ensemble in terms of temperature T and baryon chemical potential μ_b [69].




Figure 2.4: Left: Chemical potential extracted from thermal fits at different center of mass energies [69]. Right: The number of protons minus number of antiprotons per unit rapidity for central heavy ion collisions [70]. This net proton number decreases with increasing center of mass energy from $\sqrt{s} = 5$ GeV (at the AGS collider at BNL), via $\sqrt{s} = 17$ GeV (at the SPS collider at CERN) to $\sqrt{s} = 200$ GeV (at RHIC). (For each collision energy, y_p indicates the rapidity of a hypothetical proton that has the same velocity after the collision as it did before.)

the ratios among these abundances can be fit upon assuming thermal distributions with some temperature T and some baryon number chemical potential μ , as shown in Fig. 2.3 This is a two parameter fit to about a dozen ratios. The temperature extracted in this way is called the chemical freezeout temperature, since one interpretation is that it is the temperature at which the hadronic matter becomes dilute enough that inelastic hadron-hadron collisions cease to modify the abundance ratios. The chemical freezeout temperature in heavy ion collisions at top RHIC energies is about 155-180 MeV [69, 71]. This is interesting for several reasons. First, it is not far below the QCD phase transition temperature, which means that the appropriateness of a hadron gas description of this epoch may be questioned. Second, within error bars it is the same temperature that is extracted by doing a thermal model fit to hadron production in electron-positron collisions, in which final state rescattering, elastic or inelastic, can surely be neglected. So, by itself the success of the thermal fits to abundance ratios in heavy ion collisions could be interpreted as telling us about the statistical nature of hadronization, as must be the case in electron-positron collisions. However, given that we know that in heavy ion collisions (and not in electron-positron collisions) kinetic equilibrium is maintained down to a *lower* kinetic freezeout temperature, and given that as we shall see in the next subsection approximate local thermal equilibrium is achieved at a *higher* temperature, it does seem most natural to interpret the chemical freezeout temperature in heavy ion collisions as reflecting the temperature of the matter produced at the time when species-changing processes cease.

We have not yet talked about the baryon number chemical potential extracted from the thermal fit to abundance ratios. As illustrated in the left panel of Fig. 2.4, this μ decreases with increasing collision energy \sqrt{s} . This energy dependence has two origins. The dominant effect is simply that at higher and higher collision energies more and more entropy is produced while the total net baryon number in the collision is always 197+197. At top RHIC energies, these baryons are diluted among the 8000 or so hadrons in the final state, making the baryon chemical potential much smaller than it is in lower energy collisions where the final state multiplicity is much lower. The second effect is that, in the highest energy collisions, most of the net baryon number from the two incident nuclei stays at large pseudorapidity (meaning small angles near the incident beam directions). These two effects can be seen directly in the data shown in the left panel of Fig. 2.1 and the right panel of Fig. 2.4: as the collision energy increases, the total number of hadrons in the final state grows while the net baryon number at mid-rapidity drops.¹ This experimental fact that baryon number is not “fully stopped” teaches us about the dynamics of the earliest moments of a hadron-hadron collision. (In this respect, heavy ion collisions are not qualitatively different than proton-proton collisions.) In a high energy proton proton collision, particle production at mid-rapidity is dominated by

¹The data in the right panel of Fig. 2.4 is plotted relative to rapidity

$$y \equiv \frac{1}{2} \ln \left(\frac{E + p_L}{E - p_L} \right), \quad (2.2)$$

where E and p_L are the energy and longitudinal momentum of a proton in the final state. Recall that rapidity and pseudorapidity η (used in the plot in the left panel of Fig. 2.1) become the same in the limit in which E and p_L are much greater than the proton mass. When one plots data for all charged hadrons, as in the left panel of Fig. 2.1, only pseudorapidity can be defined since the rapidity of a hadron with a given polar angle θ depends on the hadron mass. When one plots data for identified protons, pseudorapidity can be converted into rapidity.

the partons in the initial state that carry a small fraction of the momentum of an individual nucleon — small Bjorken x . And, the small- x parton distributions functions that describe the initial state of the incident nucleons or nuclei are dominated by gluons and to a lesser extent by quark-antiquark pairs; the net baryon number is at larger- x . As we have already said, we shall not focus on the many interesting questions related to the early-time dynamics in heavy ion collisions, since gravity dual calculations have not had much to say about them. Instead, in the next subsection we turn to the evidence that local thermal equilibrium is established quickly, therefore at a high temperature. This means that heavy ion collisions can teach us about properties of the high temperature phase of QCD, namely the quark-gluon plasma. And, we shall see later, so can gravity dual calculations. We shall henceforth always work at $\mu = 0$. This is a good approximation as long as $\mu/3$, the quark chemical potential, is much less than the temperature T . The results in the left panel Fig. 2.4 show that this is a very good approximation at top RHIC energies, and will be an even better approximation for heavy ion collisions at the LHC. Using lower energy heavy ion collisions to scan the QCD phase diagram by varying μ is a very interesting ongoing research program, but we shall not address it in this review.

2.2 Elliptic flow

2.2.1 Introduction and motivation

The phrase “elliptic flow” refers to a suite of experimental observables in heavy ion physics that utilize the experimentalists’ ability to select events in which the impact parameter of the collision lies within some specified range and use these events to study how the matter produced in the collision flows collectively. The basic idea is simple. Suppose we select events in which the impact parameter is comparable to the nuclear radius. Now, imagine taking a beam’s eye view of one of these collisions. The two Lorentz-contracted nuclei (think circular “pancakes”) collide only in an “almond-shaped” region, see Fig. 2.5 The fragments of the nuclei outside the almond that did not collide (“spectator nucleons”) fly down the beam pipes. All the few thousand particles at mid-rapidity in the final state must have come from the few hundred nucleon-nucleon collisions that occurred within the almond. If these few thousand hadrons came instead from a few hundred independent nucleon-nucleon collisions, just by the central limit theorem the few thousand final state hadrons would be distributed uniformly in azimuthal angle ϕ (angle around the beam direction). This null hypothesis, which we shall make quantitative below, is ruled out by the data as we shall see. If, on the other hand, the collisions within the almond yield particles that interact, reach local equilibrium, and thus produce some kind of fluid, our expectations for the “shape” of the azimuthal distribution of the final state hadrons is quite different. The hypothesis that is logically the opposite extreme to pretending that the thousands of partons produced in the hundreds of nucleon-nucleon collisions do not see each other is to pretend that what is produced is a fluid that flows according to the laws of ideal, zero viscosity, hydrodynamics, since this extreme is achieved in the limit of zero mean free path. In hydrodynamics, the almond is thought of as a drop of fluid, with zero pressure at its edges and a high pressure at its center. This droplet of course explodes. And, since the pressure gradients are greater across the short extent of the almond than they are across its long direction, the explosion is




Figure 2.5: Sketch of the collision of two nuclei, shown in the transverse plane perpendicular to the beam. The collision region is limited to the interaction almond in the center of the transverse plane. Spectator nucleons located in the white regions of the nuclei do not participate in the collision. Figure taken from Ref. [72].

azimuthally asymmetric. The first big news from the RHIC experimental program was the discovery that, at RHIC energies, these azimuthal asymmetries can be large: the explosions can blast with summed transverse momenta of the hadrons that are twice as large in the short direction of the almond as they are in the long direction. And, as we shall see, it turns out that ideal hydrodynamics does a surprisingly good job of describing these asymmetric explosions of the matter produced in heavy ion collisions with nonzero impact parameter. This has implications which are sufficiently interesting that they motivate our describing this story in considerable detail over the course of this entire Section. We close this introduction with a sketch of these implications.

First, the agreement between data and ideal hydrodynamics teaches us that the viscosity η of the fluid produced in heavy ion collisions must be low. η enters in the dimensionless ratio η/s , with s the entropy density, and it is η/s that is constrained to be small. A fluid that is close to the ideal hydrodynamic limit, with small η/s , requires strong coupling between the fluid constituents. Small η/s means that momentum is not easily transported over distances that are long compared to $\sim s^{-1/3}$, which means that there can be no well-defined quasiparticles with long mean free paths in a low viscosity fluid since if they existed, they would transport momentum and damp out shear flows. No particles with long mean free paths means strongly coupled constituents.

Second, we learn that the strong coupling between partons that results in approximate local equilibration and fluid flow close to that described by ideal hydrodynamics must set in very soon after the initial collision. If partons moved with significant mean free paths for many fm of time after the collision, delaying equilibration for many fm, the almond would circularize to a significant degree during this initial period of time and the azimuthal momentum asymmetry generated by any later period of hydrodynamic behavior would be less than observed. When this argument is made quantitative, the conclusion is that RHIC

collisions produce strongly coupled fluid in approximate local thermal equilibrium within close to or even somewhat less than 1 fm after the collision [73].²

We can begin to see that the circle of ideas that emerge from the analysis of elliptic flow data are what make heavy ion collisions of interest to the broader community of theoretical physicists for whom we are writing this review. These analyses justify the conclusion that only 1 fm after the collision the matter produced can be described using the language of thermodynamics and hydrodynamics. And, we have already seen that at this early time the energy density is well above the hadron-QGP crossover in QCD thermodynamics which is well-characterized in lattice calculations. This justifies the claim that heavy ion collisions produce quark-gluon plasma. Furthermore, the same analyses teach us that this quark-gluon plasma is a strongly coupled, low viscosity, fluid with no quasiparticles having any significant mean free path. Lattice calculations have recently begun to cast some light on these transport properties of quark-gluon plasma, but these lattice calculations that go beyond Euclidean thermodynamics are still in their pioneering epoch. Perturbative calculations of quark-gluon plasma properties are built upon the existence of quasiparticles. The analyses of elliptic flow data thus cast doubt upon their utility. And, we are motivated to study the strongly coupled plasmas with similar properties that can be analyzed via gauge theory / gravity duality, since these calculational methods allow many questions that go beyond thermodynamics to be probed rigorously at strong coupling.

In Section 2.2.2 we first describe how to select classes of events all of whose impact parameters lie in some narrow range. Next, we define the experimental observable used to measure azimuthal asymmetry, and describe how to falsify the null hypothesis that the asymmetry is due only to statistical fluctuations. We then take a first look at the RHIC data and convince ourselves that the asymmetries that are seen reflect collective flow, not statistical fluctuations. In Section 2.2.3 we describe how to do a hydrodynamical calculation of the azimuthal asymmetry, and in Section 2.2.4 we compare hydrodynamic calculations to RHIC data and describe the conclusions that we have sketched above in more quantitative terms. The reader interested only in the bottom line should skip directly to Section 2.2.4.

2.2.2 The elliptic flow observable v_2 at RHIC

We want to study the dependence of collective flow on the size and anisotropy of the almond-shaped region of the transverse plane, as seen in the qualitative beam's eye view sketch in Figure 2.5. To this end, it is obviously necessary to bin heavy ion collisions as a function of

²Reaching approximate local thermal equilibrium and hence hydrodynamic behavior within less than 1 fm after a heavy ion collision has been thought of as “rapid equilibration”, since it is rapid compared to weak coupling estimates [15]. This observation has launched a large effort (that we shall not review) towards explaining equilibration as originating from weakly-coupled processes that arise in the presence of the strong color fields that are present in the initial instants of a heavy ion collision. A very recent calculation indicates, however, that the observed equilibration time may not be so rapid after all. In a strongly coupled field theory with a dual gravitational description, when two sheets of energy density with a finite thickness collide at the speed of light a hydrodynamic description of the plasma that results becomes reliable only ~ 3 sheet-thicknesses after the collision [25]. And, a Lorentz-contracted incident gold nucleus at RHIC has a maximum thickness of only 0.14 fm. So, if the equilibration processes in heavy ion collisions could be thought of as strongly coupled throughout, perhaps local thermal equilibrium and hydrodynamic behavior would set in even more rapidly than is indicated by the data.

this impact parameter. This is possible in heavy ion collisions, since the number of hadrons produced in a heavy ion collision is anticorrelated with the impact parameter of the collision. For head-on collisions (conventionally referred to as “central collisions”) the multiplicity is high; the multiplicity is much lower in collisions with impact parameters comparable to the radii of the incident ions (often referred to as “semi-peripheral collisions”); the multiplicity is lower still in grazing (“peripheral”) collisions. Experimentalists therefore bin their events by multiplicity, using that as a proxy for impact parameter. The terminology used refers to the “0-5% centrality bin” and the “5-10%” centrality bin and . . . , meaning the 5% of events with the highest multiplicities, the next 5% of events with the next highest multiplicity, The correlation between event multiplicity and impact parameter is described well by the Glauber theory of multiple scattering, which we shall not review here. Suffice to say that even though the absolute value of the event multiplicities is the subject of much ongoing research, the question of what distribution of impact parameters corresponds to the 0-5% centrality bin (namely the most head-on collisions) is well established. Although experimentalists cannot literally pick a class of events with a single value of the impact parameter, by binning their data in multiplicity they can select a class of events with a reasonably narrow distribution of impact parameters centered around any desired value. This is possible only because nuclei are big enough: in proton-proton collisions, which in principle have impact parameters since protons are not pointlike, there is no operational way to separate variations in impact parameter from event-by-event fluctuations in the multiplicity at a given impact parameter.

Suppose that we have selected a class of semi-peripheral collisions. Since these collisions have a nonzero impact parameter, the impact parameter vector together with the beam direction define a plane, conventionally called the reaction plane. Directions within the transverse plane of Figure 2.5 specified by the azimuthal angle ϕ now need not be equivalent. We can ask to what extent the multiplicity and momentum of hadrons flying across the short direction of the collision almond (in the reaction plane) differs from that of the hadrons flying along the long direction of the collision almond (perpendicular to the reaction plane).

Let us characterize this dependence on the reaction plane for the case of the single inclusive particle spectrum $\frac{dN}{d^3\mathbf{p}}$ of a particular species of hadron. The three-momentum \mathbf{p} of a particle of mass m is parametrized conveniently in terms of its transverse momentum p_T , its azimuthal angle ϕ , and its rapidity y which specifies its longitudinal momentum. Specifically,

$$\mathbf{p} = \left(p_T \cos \phi, p_T \sin \phi, \sqrt{p_T^2 + m^2} \sinh y \right). \quad (2.3)$$

The energy of the particle is $E = \sqrt{p_T^2 + m^2} \cosh y$. The single particle spectrum can then be written as

$$\frac{dN}{d^2\mathbf{p}_t dy} = \frac{1}{2\pi p_T} \frac{dN}{dp_T dy} [1 + 2v_1 \cos(\phi - \Phi_R) + 2v_2 \cos 2(\phi - \Phi_R) + \dots], \quad (2.4)$$

where Φ_R denotes explicitly the azimuthal orientation of the reaction plane, which we do not know *a priori*. Thus, the azimuthal dependence of particle production is characterized by the harmonic coefficients

$$v_n \equiv \langle \exp [i n (\phi - \Phi_R)] \rangle = \frac{\int \frac{dN}{d^3\mathbf{p}} e^{i n (\phi - \Phi_R)} d^3p}{\int \frac{dN}{d^3\mathbf{p}} d^3p}. \quad (2.5)$$

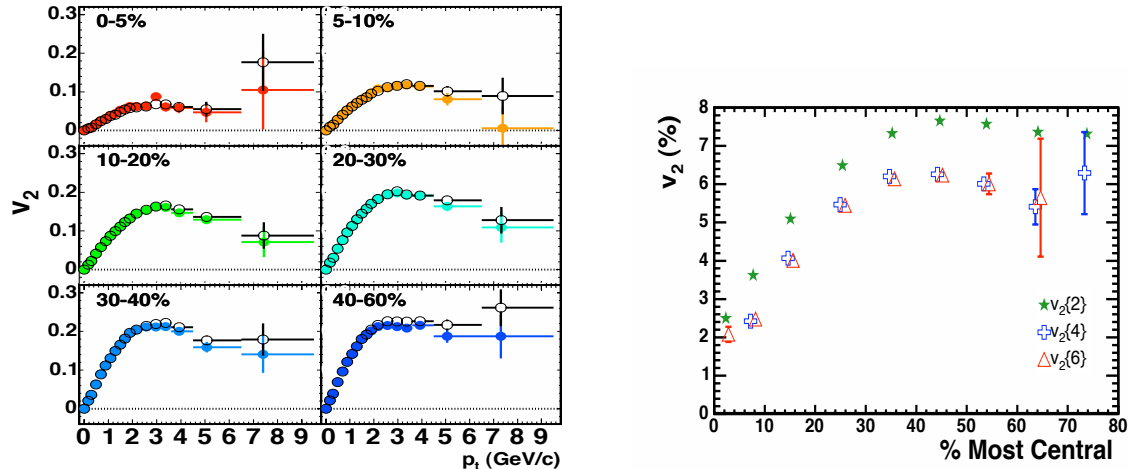


Figure 2.6: Left: Transverse momentum dependence of the elliptic flow $v_2(p_T)$ for different centrality bins. We see v_2 increasing as one goes from nearly head-on collisions to semi-peripheral collisions. Right: the p_T -integrated elliptic flow v_2 as a function of centrality bins, reconstructed with 2nd, 4th and 6th order cumulants. The most head-on collisions are on the left. Figures taken from Ref. [74].

The coefficients v_n are referred to generically as n -th order flow. In particular, v_1 is referred to as “directed flow” and v_2 as “elliptic flow”. In general, the v_n can depend on the transverse momentum p_T , the rapidity y , the impact parameter of the collision, and they can differ for different particle species. In the collision of identical nuclei at mid-rapidity, the collision region is symmetric under $\phi \rightarrow \phi + \pi$ and all odd harmonics vanish. In this case, the elliptic flow v_2 is the first non-vanishing coefficient. We shall focus only on mid-rapidity, $y = 0$. The most common observables used are $v_2(p_T)$ for collisions with varying centrality (i.e. impact parameter) and v_2 integrated over all p_T , which is just a function of centrality. A nonzero v_2 (or $v_2(p_T)$) means that there are more particles (with a given transverse momentum) going left and right than up and down.

In Fig. 2.6, we show data for the transverse momentum dependence of the elliptic flow $v_2(p_T)$, and the transverse momentum integrated quantity, measured for different centrality classes in Au+Au collisions at RHIC. The azimuthal asymmetry v_2 of the final state single inclusive hadron spectrum is maximal in semi-peripheral collisions. v_2 is less for more central collisions, since the initial geometric asymmetry is less — the almond-shaped collision region becomes closer to circular as the impact parameter is reduced. In the idealized case of zero impact parameter, v_2 due to collective effects should vanish. The reason that v_2 is not smaller even in the sample of the 5% most head-on collisions is that this sample includes events with a distribution of impact parameters in the range $0 < b < 3.5$ fm. In principle, v_2 should also vanish in the limit of grazing collisions. Some decrease is visible in the most peripheral collisions shown in the right panel of Fig. 2.6, but the error bars become large long before one reaches collisions which are literally grazing.

It is worth pausing to appreciate the size of the elliptic flow signal shown in Fig. 2.6. According to (2.4), the ratio of $\frac{dN}{d^3\mathbf{p}}$ in whatever azimuthal direction it is largest to $\frac{dN}{d^3\mathbf{p}}$ ninety degrees in azimuth away is $(1 + 2v_2)/(1 - 2v_2)$, which is a factor of 2 for $v_2 = 1/6$. Thus,

a v_2 of the order of magnitude shown in semi-peripheral collisions at RHIC for $p_T \sim 2$ GeV corresponds to a factor of more than 2 azimuthal asymmetry. It is natural to expect a nonzero quadrupolar asymmetry v_2 at nonzero impact parameter due to collective effects occurring in the almond-shaped collision zone, given that the almond shape itself has a quadrupole. Of course, v_4, v_6, v_8, \dots will generically be nonzero as well.

There are two basic experimental problems in determining the coefficients (2.5). The orientation Φ_R of the reaction plane is not known *a priori*, and *a priori* it is unclear whether an azimuthal asymmetry in an event arises as a consequence of collective motion, or whether it is the result of a statistical fluctuation. In practice, both problems are connected. There are two methods in common use that each address both problems. One option is to use data at high rapidity to determine Φ_R , and then to measure the v_n by applying (2.5) to mid-rapidity data. This eliminates the contribution of all statistical fluctuations uncorrelated with the reaction plane unless these fluctuations introduce correlations between particles in the mid-rapidity and high rapidity regions of the detector. The second option uses only the mid-rapidity data. This means that in each event the azimuthal direction of the maximal particle yield depends both on the orientation of the reaction plane and on the statistical fluctuations in that event. If it just so happens that there were 3 jets going up and 3 jets going down, that would certainly skew the determination of the reaction plane Φ_R from the data. However, there is a systematic procedure, known as the cumulant method [75], which makes an unambiguous separation of flow effects from non-flow effects possible. The basic idea is that if two particles are correlated in the azimuthal angle because of a non-flow effect such as resonance decay or back-to-back jets or some other fluctuation, then they are not correlated with all the other particles in the event. In contrast, a flow effect leaves a signature in which all particles know about the reaction plane, and thus are correlated. By a suitable analysis of 2-, 4-, and 6- particle correlations, it is in practice possible to disentangle these two effects [75]. The right-hand side of Fig. 2.6 shows such an analysis. Subtracting non-flow effects in a more accurate 4-th order cumulant analysis, the elliptic flow signal is reduced a bit. Further increasing the accuracy to 6th order, the signal does not change, indicating that all non-flow effects have been corrected. In short, this demonstrates that the large elliptic flow signal cannot arise from statistical fluctuations in an incoherent superposition of nucleon-nucleon collisions or from non-flow effects. It must characterize a *collective* phenomenon, reflecting interactions among the debris produced by all the nucleon-nucleon collisions occurring within the almond-shaped collision region. For more details, we refer the reader to the original literature [75].

2.2.3 Calculating elliptic flow using (ideal) hydrodynamics

We have now seen that the azimuthal asymmetry of the little explosions created by heavy ion collisions with nonzero impact parameter, parametrized by the second harmonic v_2 , characterizes a strong *collective* phenomenon. At time zero, the almond-shaped collision is azimuthally asymmetric *in space*; subsequent collective dynamics must convert that spatial asymmetry into an asymmetry *in momentum space*. After all, v_2 measures the azimuthal asymmetry of the momentum distribution of the final state hadrons. This conversion of initial spatial anisotropy into final momentum anisotropy is characteristic of any explosion — one shapes the explosive in order to design a charge that blasts with greater force in some directions

than others. Hydrodynamics provides the natural language for describing such processes: the initial spatial anisotropy corresponds to anisotropic pressure gradients. Assuming that the pressure is maximal at the center of the almond and zero at its edge, the gradient is greater across the almond (in the reaction plane) than along it (perpendicular to the reaction plane). The elliptic flow v_2 measures the extent to which these pressure gradients lead to an anisotropic explosion with greater momentum flow in the reaction plane. The size of v_2 characterizes the efficiency of translating initial pressure gradients into collective flow. By doing hydrodynamic calculations and comparing the calculated v_2 to that in the data (and also comparing the final state radial flow velocity to that determined from the single-particle spectra as described in the previous section) one can constrain the input quantities that go into a hydrodynamic description.

The starting point in any hydrodynamic analysis is to consider the limit of ideal, zero viscosity, hydrodynamics. In this limit, the hydrodynamic description is specified entirely by an equation of state, which relates the pressure and the energy density, and by the initial spatial distribution of energy density and fluid velocity. In particular, in ideal hydrodynamics one is setting all dissipative coefficients (shear viscosity, bulk viscosity, and their many higher order cousins) to zero. If the equation of state is held fixed and viscosity is turned on, v_2 must decrease: turning on viscosity introduces dissipation that has the effect of turning some of the initial anisotropy in pressure gradients into entropy production, rather than into directed collective flow. So, upon making some assumption for the equation of state and for the initial energy density distribution, setting the viscosities to zero yields an upper bound on the v_2 in the final state. Ideal, inviscid, hydrodynamics has therefore long been used as a calculational benchmark in heavy ion physics. As we shall see below, in heavy ion collisions at RHIC energies ideal hydrodynamics does a good job of describing $v_2(p_T)$ for pions, kaons and protons for transverse momenta p_T below about 1 – 2 GeV. This motivates an ongoing research program in which one begins by comparing data to the limiting case of ideal, inviscid, hydrodynamics and then turns to a characterization of dissipative effects, asking how large a viscosity will spoil the agreement with data. In this subsection, we sketch how practitioners determine the equation of state and initial energy density profile, and we recall the basic principles behind the hydrodynamic calculations on which all these studies rest. In the next subsection, we summarize the current constraints on the shear viscosity that are obtained by comparing to RHIC v_2 data.

The equation of state relates the pressure P to the energy density ε . P is a thermodynamic quantity, and therefore can be calculated using the methods of lattice quantum field theory. Lattice calculations (or fits to them) of $P(\varepsilon)$ in the quark-gluon plasma and in the crossover regime between QGP and hadron gas are often used as inputs to hydrodynamic calculations. At lower energy densities, practitioners either use a hadron resonance gas model equation of state, or match the hydrodynamic calculation onto a hadron cascade model. One of the advantages of focussing on the v_2 observable is that it is insensitive to the late time epoch of the collision, when all the details of these choices matter. This insensitivity is easy to understand. v_2 describes the conversion of a spatial anisotropy into anisotropic collective flow. As this conversion begins, the initial almond-shaped collision explodes with greater momentum across the short direction of the almond, and therefore circularizes. Once it has circularized, no further v_2 can develop. Thus, v_2 is generated early in the collision, say over the first ~ 5 fm of time. By the late times when a hadron gas description is needed, v_2 has

already been generated. In contrast, the final state radial flow velocity reflects a time integral over the pressure built up during all epochs of the collision: it is therefore sensitive to how the hadron gas phase is described, and so is a less useful observable.

The discussion above reminds us of a second sense in which the ideal hydrodynamic calculation of v_2 is a benchmark: ideal hydrodynamics requires local equilibrium. It therefore cannot be valid from time $t = 0$. By using an ideal hydrodynamic description beginning at $t = 0$ we must again be overestimating v_2 , and so we can ask how long an initial phase during which partons stream freely without, starting to circularize the almond-shaped region without generating any momentum anisotropy, can be tolerated without spoiling the agreement between calculations and data.

After choosing an equation of state, an initialization time, and viscosities (zero in the benchmark calculation), the only thing that remains to be specified is the distribution of energy density as a function of position in the almond-shaped collision region. (The transverse velocities are assumed to be zero initially.) In the simplest approach, called Glauber theory, this energy density is proportional to the product of the thickness of the two nuclei at a given point in the transverse plane. It is thus zero at the edge of the almond, where the thickness of one nucleus goes to zero, and maximum at the center of the almond. The proportionality constant is determined by fitting to data other than v_2 , see e.g. [76]. The assumptions behind this Glauber approach to estimating how much energy density is created at a given location as a function of the nuclear thickness at that location are assumptions about physics of the collision at $t = 0$. There is in particular one alternative model parametrization, for which the energy density rises towards the center of the almond more rapidly than the product of the nuclear thicknesses. This parametrization is referred to as the CGC initial condition, since it was first motivated by ideas of parton saturation (called “color glass condensate”) [77]. The Glauber and CGC models for the initial energy density distribution are often used as benchmarks in the hope that they bracket Nature’s choice.

We turn now to a review of the formulation of the hydrodynamic equations of motion. Hydrodynamics is an effective theory which describes the small frequency and long wavelength limit of an underlying interacting dynamical theory [78]. It is a classical field theory, where the fields can be understood as the expectation values of certain quantum operators in the underlying theory. In the hydrodynamic limit, since the length scales under consideration are longer than any correlation length in the underlying theory, by virtue of the central limit theorem all n -point correlators of the underlying theory can be factorized into two point correlators (Gaussian approximation). The fluctuations on these average values are small, and a description in terms of expectation values is meaningful. If the underlying theory admits a (quasi)particle description, this statement is equivalent to saying that the hydrodynamic description involves averages over many of these fundamental degrees of freedom.

The hydrodynamic degrees of freedom include the expectation values of conserved currents such as the stress tensor $T^{\mu\nu}$ or the currents of conserved charges J_B . As a consequence of the conservation laws

$$\partial_\mu T^{\mu\nu} = 0, \tag{2.6}$$

$$\partial_\mu J_B^\mu = 0, \tag{2.7}$$

long wavelength excitations of these fields can only relax on long timescales, since their

relaxation must involve moving stress-energy or charges over distances of order the wavelength of the excitation. As a consequence, these conservation laws lead to excitations whose lifetime diverges with their wavelength. Such excitations are called hydrodynamic modes.

If the only long-lived modes are those from conserved currents, then hydrodynamics describes a normal fluid. However, there can be other degrees of freedom that lead to long lived modes in the long wave length limit. For example, in a phase of matter in which some global symmetry is spontaneously broken, the phase(s) of the symmetry breaking expectation value(s) is (are) also hydrodynamic modes [78]. The classic example of this is a superfluid, in which a global $U(1)$ symmetry is spontaneously broken. Chiral symmetry is spontaneously broken in QCD, but there are two reasons why we can neglect the potential hydrodynamic modes associated with the chiral order parameter [79]. First, explicit chiral symmetry breaking gives these modes a mass (the pion mass) and we are interested in the hydrodynamic description of physics on length scales longer than the inverse pion mass. Second, we are interested in temperatures above the QCD crossover, at which the chiral order parameter is disordered, the symmetry is restored, and this question does not arise. So, we need only consider normal fluid hydrodynamics. Furthermore, as we have discussed in Section 2, at RHIC energies the matter produced in heavy ion collisions has only a very small baryon number density, and it is a good approximation to neglect J_B^μ . The only hydrodynamic degrees of freedom are therefore those described by $T^{\mu\nu}$.

At the length scales at which the hydrodynamic approximation is valid, each point of space can be regarded as a macroscopic *fluid cell*, characterized by its energy density ε , pressure P , and a velocity u^μ . The velocity field can be defined by the energy flow together with the constraint $u^2 = -1$. In the so-called Landau frame, the four equations

$$u_\mu T^{\mu\nu} = -\varepsilon u^\nu. \quad (2.8)$$

determine ε and u from the stress tensor.

Hydrodynamics can be viewed as a gradient expansion of the stress-tensor (and any other hydrodynamic fields). In general, the stress tensor can be separated into a term with no gradients (*ideal*) and a term which contains all the gradients:

$$T^{\mu\nu} = T_{\text{ideal}}^{\mu\nu} + \Pi^{\mu\nu}. \quad (2.9)$$

In the rest frame of each fluid cell ($u^i = 0$), the ideal piece is diagonal and isotropic $T_{\text{ideal}}^{\mu\nu} = \text{diag}(\varepsilon, P, P, P)$. Thus, in any frame,

$$T_{\text{ideal}}^{\mu\nu} = (\varepsilon + P)u^\mu u^\nu + P g^{\mu\nu}, \quad (2.10)$$

where $g^{\mu\nu}$ is the metric of the space.

If there were a nonzero density of some conserved charge n , the velocity field could either be defined in the Landau frame as above or may instead be defined in the so-called Eckart frame, with $J^\mu = nu^\mu$. In the Landau frame, the definition (2.8) of u^μ implies $\Pi^{\mu\nu}u_\nu = 0$ (transversality). Hence, there is no heat flow but there can be currents of the conserved charge. In the Eckart frame, the velocity field is comoving with the conserved charge, but there can be heat flow.

A. Ideal fluid

Ideal hydrodynamics is the limit in which all gradient terms in $T^{\mu\nu}$ are neglected. It can

be used to describe motions of the fluid that occur on macroscopic length scales and time scales associated with how the fluid is “stirred” that are long compared to any internal scales associated with the fluid itself. Corrections to ideal hydrodynamics — namely the gradient terms in $\Pi^{\mu\nu}$ that we shall discuss shortly — introduce internal length and time scales. For example, time scales for relaxation of perturbations away from local thermal equilibrium, length scales associated with mean free paths, etc. Ideal hydrodynamics works on longer length scales than these. At long time scales, the fluid cells are in equilibrium, P is the pressure in the rest frame of each fluid cell, and ε and P are related by the equation of state. This equation of state can be determined by studying a homogeneous system at rest with no gradients, for example via a lattice calculation.

The range of applicability of ideal fluid hydrodynamics can be characterized in terms of the isotropization scale τ_{iso} and the thermalization scale τ_{eq} . The isotropization scale measures the characteristic time over which an initially anisotropic stress tensor acquires a diagonal form in the local rest frame. The thermalization scale sets the time over which entropy production ceases. These scales do not need to be the same, but $\tau_{\text{eq}} \geq \tau_{\text{iso}}$ since thermalization implies isotropy. P and ε are related by the equilibrium equation of state only at times $\tau > \tau_{\text{eq}}$. If there is a large separation of scales $\tau_{\text{iso}} \ll \tau_{\text{eq}}$, then between these timescales the relationship between P and ε can be different and is generically time-dependent. Ideal hydrodynamics can nevertheless be used, as long as $P(\varepsilon)$ is replaced by some suitable non-equilibrium equation of state that will depend on exactly how the system is out of thermal equilibrium. It is worth noting that in the case of a conformal theory, since the trace of the stress tensor vanishes, $P = \frac{1}{3}\varepsilon$ whether or not the system is in equilibrium. In this special case, as long as there are no charge densities the equation of state is unmodified even if the system is isotropic but out of equilibrium, and ideal hydrodynamics with the standard equation of state can be used as long as $\tau \gg \tau_{\text{iso}}$.

B. 1st order dissipative fluid dynamics

Going beyond the infinite wavelength limit requires the introduction of viscosities. To first order in gradients, the requirement that $\Pi^{\mu\nu}$ be transverse means that it must take the form

$$\Pi^{\mu\nu} = -\eta(\varepsilon)\sigma^{\mu\nu} - \zeta(\varepsilon)\Delta^{\mu\nu}\nabla\cdot u, \quad (2.11)$$

where η and ζ are the shear and bulk viscosities, $\nabla^\mu = \Delta^{\mu\nu}d_\nu$, with d_ν the covariant derivative and

$$\Delta^{\mu\nu} = g^{\mu\nu} + u^\mu u^\nu, \quad (2.12)$$

$$\sigma^{\mu\nu} = \Delta^{\mu\alpha}\Delta^{\nu\beta}(\nabla_\alpha u_\beta + \nabla_\beta u_\alpha) - \frac{2}{3}\Delta_{\alpha\beta}\nabla\cdot u. \quad (2.13)$$

The operator $\Delta^{\mu\nu}$ is the projector into the space components of the fluid rest frame. Note that in this frame the only spatial gradients that appear in eq. (2.11) are gradients of the velocity fields. The reason that the derivative of ε does not appear is that it can be eliminated in the first order equations by using the zeroth order equation of motion

$$D\varepsilon = -(\varepsilon + P)\nabla_\mu u^\mu, \quad (2.14)$$

where $D = u^\mu d_\mu$ is the time derivative in the fluid rest frame. (Similarly, time derivatives of the energy density can be eliminated in the second order equations that we shall give below using the first order equations of motion.)

It is worth pausing to explain why we have introduced a covariant derivative, even though we will only ever be interested in heavy ion collisions — and thus hydrodynamics — occurring in flat spacetime. But, it is often convenient to use curvilinear coordinates with a nontrivial metric. For example, the longitudinal dynamics is more conveniently described using proper time $\tau = \sqrt{t^2 - z^2}$ and spacetime rapidity $\xi = \text{arctanh}(z/t)$ as coordinates rather than t and z . In these “Milne coordinates”, the metric is given by $g_{\mu\nu} = \text{diag}(g_{\tau\tau}, g_{xx}, g_{yy}, g_{\xi\xi}) = (-1, 1, 1, \tau^2)$. These coordinates are useful because boost invariance simply translates into the requirement that ε , u^μ and $\Pi^{\mu\nu}$ be independent of ξ , depending on τ only. In particular, if the initial conditions are boost invariant, then the fluid dynamic evolution will preserve this boost invariance, and the numerical calculation reduces in Milne coordinates to a 2+1-dimensional problem. In the high energy limit, one expects that hadronic collisions distribute energy density approximately uniformly over a wide range rapidity. Accordingly, boost-invariant initial conditions are often taken to be a good approximation [64]. But, fully 3+1-dimensional calculations that do not assume boost invariance can also be found in the literature [80–83].

It is often convenient to phrase the hydrodynamic equations in terms of the entropy density s . In the absence of conserved charges, i.e. with baryon chemical potential $\mu_B = 0$, the entropy density is $s = (\varepsilon + P)/T$. Using this and another fundamental thermodynamic relation, $DE = TDS - PDV$ (where $E/V \equiv \varepsilon$) the zeroth order equation of motion (2.14) becomes exactly the equation of entropy flow for an ideal isentropic fluid

$$Ds = -s \nabla_\mu u^\mu \quad (2.15)$$

Repeating this analysis at first order, including the viscous terms, one easily derives from $u_\mu \nabla_\nu T^{\mu\nu} = 0$ that

$$\frac{Ds}{s} = -\nabla_\mu u^\mu - \frac{1}{sT} \Pi^{\mu\nu} \nabla_\nu u_\mu. \quad (2.16)$$

A similar analysis of the other three hydrodynamic equations then shows that they take the form

$$Du_\alpha = -\frac{1}{Ts} \Delta_{\alpha\nu} \left(\nabla^\nu P + \nabla_\mu \Pi^{\mu\nu} \right). \quad (2.17)$$

It then follows from the structure of the shear tensor $\Pi^{\mu\nu}$ (2.11) that shear viscosity and bulk viscosity always appear in the hydrodynamic equations of motion in the dimensionless combinations η/s and ζ/s . The net entropy increase is proportional to these dimensionless quantities. Gradients of the velocity field are measured in units of $1/T$.

In a conformal theory, $\zeta = 0$ since $\Pi^{\mu\nu}$ must be traceless. There are a number of indications from lattice calculations that as the temperature is increased above $(1.5 - 2)T_c$, with T_c the crossover temperature, the quark-gluon plasma becomes more and more conformal. The equation of state approaches $P = \frac{1}{3}\varepsilon$ [84, 85]. The bulk viscosity drops rapidly [86]. So, we shall set $\zeta = 0$ throughout the following, in so doing neglecting temperatures close to T_c . One of the things that makes heavy ion collisions at the LHC interesting is that in these collisions the plasma that is created is expected to be better approximated as conformal than is the case at RHIC, where the temperature at $\tau = 1$ fm is thought to be between $1.5T_c$ and $2T_c$.

Just like the equation of state $P(\varepsilon)$, the shear viscosity $\eta(\varepsilon)$ is an input to the hydrodynamic description that must be obtained either from experiment or from the underlying microscopic

theory. We shall discuss in Section 3.2 how transport coefficients like η are obtained from correlation functions of the underlying microscopic theory via Kubo formulae.

C. Second order dissipative Hydrodynamics

Even though hydrodynamics is a controlled expansion in gradients, the first order expression for the tensor $\Pi^{\mu\nu}$, eq. (2.11), is unsuitable for numerical computations. The problem is that the set of equations (2.6) with the approximations (2.11) leads to acausal propagation. Even though this problem only arises for modes outside of the region of validity of hydrodynamics (namely high momentum modes with short wavelengths of the order of the microscopic length scale defined by η), the numerical evaluation of the first order equations of motion is sensitive to the acausality in these hard modes. This problem is solved by going to one higher order in the gradient expansion. This is known as second order hydrodynamics.

There is a phenomenological approach to second order hydrodynamics due to Müller, Israel and Stewart aimed at explicitly removing the acausal propagation [87–89]. In this approach, the tensor $\Pi^{\mu\nu}$ is treated as a new hydrodynamic variable and a new dynamical equation is introduced. In its simplest form this equation is

$$\tau_{\Pi} D\Pi^{\mu\nu} = -\Pi^{\mu\nu} - \eta\sigma^{\mu\nu}, \quad (2.18)$$

where τ_{Π} is a new (second order) coefficient. Note that as $\tau_{\Pi} \rightarrow 0$, eq. (2.18) coincides with eq. (2.11) with the bulk viscosity ζ set to zero. Eq. (2.18) is such that $\Pi^{\mu\nu}$ relaxes to its first order form in a (proper) time τ_{Π} . There are several variants of this equation in the literature, all of which follow the same philosophy. They all introduce the relaxation time as the characteristic time in which the tensor $\Pi^{\mu\nu}$ relaxes to its first order value. The variations arise from different ways of fixing some pathologies of eq. (2.18), since as written eq. (2.18) does not lead to a transverse stress tensor (although this is a higher order effect) and is not conformally invariant. Since in this approach the relaxation time is introduced *ad hoc*, it may not be possible to give a prescription for extracting it from the underlying microscopic theory.

The systematic extraction of second order coefficients demands a similar analysis of the second order gradients as was done at first order. The strategy is, once again, to write all possible terms with two derivatives which are transverse and consistent with the symmetries of the theory. As before, only spatial gradients (in the fluid rest frame) are considered, since time gradients can be related to the former via the zeroth order equations of motion.

In a conformal theory, second order hydrodynamics simplifies. First, only terms such that $\Pi_{\mu}^{\mu} = 0$ are allowed. Furthermore, the theory must be invariant under Weyl transformations

$$g_{\mu\nu} \rightarrow e^{-2\omega(x)} g_{\mu\nu} \quad (2.19)$$

which implies

$$T \rightarrow e^{\omega(x)} T, \quad u^{\mu} \rightarrow e^{\omega(x)} u^{\mu}, \quad T^{\mu\nu} \rightarrow e^{(d+2)\omega(x)} T^{\mu\nu}, \quad (2.20)$$

where T is the temperature and d the number of spacetime dimensions. It turns out that there are only five operators that respect these constraints [90]. The second order contributions to

the tensor $\Pi^{\mu\nu}$ are linear combinations of these operators, and can be cast in the form [90,91]

$$\begin{aligned}\Pi^{\mu\nu} &= -\eta\sigma^{\mu\nu} - \tau_{\Pi} \left[\langle D\Pi^{\mu\nu} \rangle + \frac{d}{d-1}\Pi^{\mu\nu}(\nabla\cdot u) \right] \\ &\quad + \kappa \left[R^{\langle\mu\nu\rangle} - (d-2)u_{\alpha}R^{\alpha\langle\mu\nu\rangle\beta}u_{\beta} \right] \\ &\quad + \frac{\lambda_1}{\eta^2}\Pi^{\langle\mu}_{\lambda}\Pi^{\nu\rangle\lambda} - \frac{\lambda_2}{\eta}\Pi^{\langle\mu}_{\lambda}\Omega^{\nu\rangle\lambda} + \lambda_3\Omega^{\langle\mu}_{\lambda}\Omega^{\nu\rangle\lambda}.\end{aligned}\tag{2.21}$$

Here, $R^{\mu\nu}$ is the Ricci tensor, the indices in brackets are the symmetrized traceless projectors onto the space components in the fluid rest frame, namely

$$\langle A^{\mu\nu} \rangle \equiv \frac{1}{2}\Delta^{\mu\alpha}\Delta^{\nu\beta}(A_{\alpha\beta} + A_{\beta\alpha}) - \frac{1}{d-1}\Delta^{\mu\nu}\Delta^{\alpha\beta}A_{\alpha\beta} \equiv A^{\langle\mu\nu\rangle},\tag{2.22}$$

and the vorticity tensor is defined as

$$\Omega^{\mu\nu} \equiv \frac{1}{2}\Delta^{\mu\alpha}\Delta^{\nu\beta}(\nabla_{\alpha}u_{\beta} - \nabla_{\beta}u_{\alpha}).\tag{2.23}$$

In deriving (2.21), we have replaced $\eta\sigma^{\mu\nu}$ by $\Pi^{\mu\nu}$ on the right-hand side in places where doing so makes no change at second order. We see from (2.21) that five new coefficients τ_{Π} , κ , λ_1 , λ_2 , and λ_3 arise at second order in the hydrodynamic description of a conformal fluid, in addition to η and the equation of state which arise at first and zeroth order respectively. The coefficient κ is not relevant for hydrodynamics in flat spacetime. The λ_i coefficients involve nonlinear combinations of fields in the rest frame and, thus, are invisible in linearized hydrodynamics. Thus, these three coefficients cannot be extracted from linear response. Of these three, only λ_1 is relevant in the absence of vorticity, as in the numerical simulations that we will describe in the next Section. These simulations have also shown that, for physically motivated choices of λ_1 , the results are insensitive to its precise value, leaving τ_{Π} as the only phenomenologically relevant second order parameter in the hydrodynamic description of a conformal fluid. In a generic, nonconformal fluid, there are nine additional transport coefficients [92].

For more in-depth discussions of 2nd order viscous hydrodynamics and its applications to heavy ion collisions, see e.g. Refs. [93–114].

2.2.4 Comparing elliptic flow in heavy ion collisions and hydrodynamic calculations

For the case of ideal hydrodynamics, the hydrodynamic equations of motion are fully specified once the equation of state $P = P(\varepsilon)$ is given. A second-order dissipative hydrodynamic calculation also requires knowledge of the transport coefficients $\eta(\varepsilon)$ and $\zeta(\varepsilon)$ (although in practice the latter is typically set to zero) and the relaxation time τ_{Π} and the second-order coefficient λ_1 entering Eq. (2.21). (κ would enter in curved space time, and λ_2 and λ_3 would enter in the presence of vorticity.) All these parameters are well-defined in terms of correlation functions in the underlying quantum field theory. In this sense, the hydrodynamic evolution equations are model-independent.

The output of any hydrodynamic calculation depends on more than the evolution equations. One must make model assumptions about the initial energy density distribution. As we have discussed in Section 2.2.2, there are two benchmark models for the energy distribution across the almond-shaped collision region, which serve to give us a sense of the degree to which results are sensitive to our lack of knowledge of the details of this initial profile. Often, the initial transverse velocity fields are set to zero and boost invariance is assumed for the longitudinal velocity field and the evolution. For dissipative hydrodynamic simulations, the off-diagonal elements of the energy-momentum tensor are additional hydrodynamic fields which must be initialized. The initialization time τ_0 , at which these initial conditions are fixed, is an additional model parameter. It can be viewed as characterizing the isotropization time, at which hydrodynamics starts to apply but collective flow has not yet developed. In addition to initial state sensitivity, results depend on assumptions made about how the system stops behaving hydrodynamically and *freezes out*. In practice, freezeout is often assumed to happen as a rapid decoupling: when a specified criterion is satisfied (e.g. when a fluid cell drops below a critical energy or entropy density) then the hydrodynamic fields in the unit cell are mapped onto hadronic equilibrium Bose/Fermi distributions. This treatment assumes that hydrodynamics is valid all the way down to the kinetic freezeout temperature, below which one has noninteracting hadrons. Alternatively, at a higher temperature close to the crossover where hadrons are formed, one can map the hydrodynamic fields onto a hadron cascade which accounts for the effects of rescattering in the interacting hadronic phase without assuming that its behavior is hydrodynamic [115, 116]. Indeed, recent work suggests that hadronization may be triggered by cavitation induced by the large bulk viscosity in the vicinity of the crossover temperature [117]. As we have discussed, v_2 is insensitive to details of how the late-time evolution is treated because v_2 is generated during the epoch when the collision region is azimuthally anisotropic. Nevertheless, these late-time issues do matter when one does a global fit to v_2 and the single-particle spectra, since the latter are affected by the radial flow which is built up over the entire history of the collision. Finally, the validity of results from any hydrodynamic calculation depends on the assumption that a hydrodynamic description is applicable. This assumption can be checked at late times by checking the sensitivity to how freezeout is modelled and can be checked at early times by confirming the insensitivity of results to the values of the second order hydrodynamic coefficients and to the initialization of the higher order off-diagonal elements of the energy momentum tensor: if hydrodynamics is valid, the gradients must be small enough at all times that second order effects are small compared to first order effects.

In practice, the dependence of physics conclusions on all these model assumptions has to be established by systematically varying the initial conditions and freeze-out prescriptions within a wide physically motivated parameter range, and comparing to data on both the single particle spectra (i.e. the radial velocity) and the azimuthal flow anisotropy coefficient v_2 . At the current time, several generic observations have emerged from pursuing this program in comparison to data from RHIC:

1. Perfect fluid dynamics reproduces the size and centrality dependence of v_2

The RHIC data on single inclusive hadronic spectra $dN/p_T dp_T dy$ and their leading azimuthal dependence $v_2(p_T)$ can be reproduced in magnitude and shape by ideal hydrodynamic calculations, for particles with $p_T < 1$ GeV. The hydrodynamic picture is




Figure 2.7: The elliptic flow v_2 versus p_T for a large number of identified hadrons (pions, kaons, protons, Λ 's) showing the comparison between an *ideal* hydrodynamic calculation to data from RHIC. Figure taken from [118].

expected to break down for sufficiently small wavelength, i.e. high momenta, consistent with the observation that significant deviations occur for $p_T > 2$ GeV (see left panel of Fig. 2.6). The initialization time for these calculations is $\tau_0 = 0.6 - 1$ fm. If τ_0 is chosen larger, the agreement between ideal hydrodynamics and data is spoiled. This gives significant support to a picture in which thermalization is achieved fast. It also provides the first indication that the shear viscosity of the fluid produced at RHIC must be small.

2. The mass-ordering of identified hadron spectra

The p_T -differential azimuthal asymmetry $v_2(p_T)$ of identified single inclusive hadron spectra shows a characteristic mass ordering in the range of $p_T < 2$ GeV: at small p_T , the azimuthal asymmetry of light hadrons is significantly more pronounced than that of heavier hadrons, see Fig. 2.7. This qualitative agreement of hydrodynamic simulations with data supports the picture that all hadron species emerge from a single fluid moving with a common flow field.

3. Data support small dissipative coefficients such as shear viscosity

Above the crossover temperature, the largest dissipative correction is expected to arise from shear viscosity η , which enters the equations of motion of second order dissipative hydrodynamics in the combination η/s , where s is the entropy density. (Bulk viscosity may be sizable close to the crossover temperature [86].) As seen in Fig. 2.8, shear viscous corrections decrease collectivity, and so v_2 decreases with increasing η/s . The data are seen to favor small values $\eta/s < 0.2$, with $\eta/s > 0.5$ strongly disfavored. To give an example of the way this conclusion is stated in one recent review, it would be surprising if η/s turned out to be larger than about $(3 - 5)/(4\pi)$ once all physical effects are properly included [110]. The figure also shows that sensitivity to our lack of knowledge of the initial energy density profile precludes a precise determination of η/s

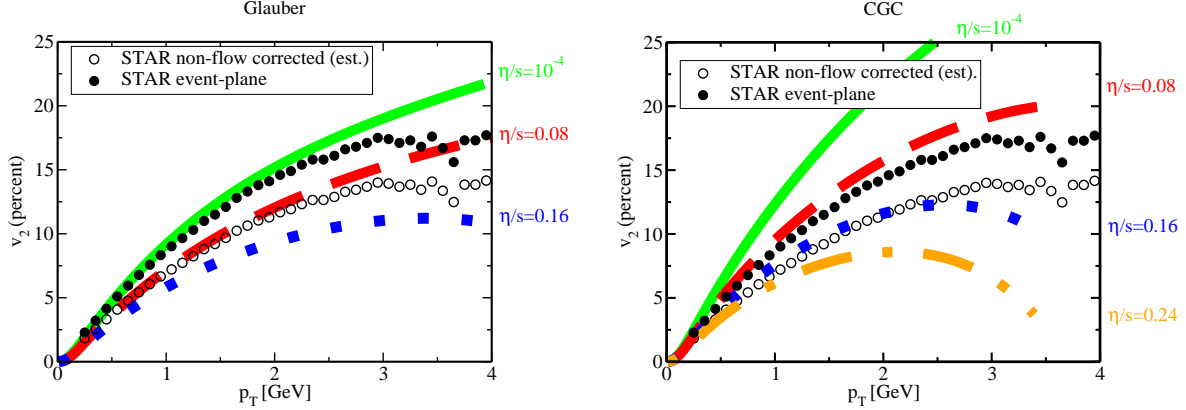


Figure 2.8: The elliptic flow v_2 versus p_T for charged hadrons showing the comparison between hydrodynamic calculations with varying shear viscosity η and data from RHIC. The two data sets show the reduction in v_2 when one uses the four-particle cumulant analysis that strips out effects of statistical two-particle correlations. The initial profile for the energy density across the almond-shaped collision region is obtained from (a) the Glauber model and (b) the color-glass condensate model. Figure taken from Ref. [103].

at present.³ The smallness of η/s is remarkable, since almost all other known liquids have $\eta/s > 1$ and most have $\eta/s \gg 1$. The one liquid that is comparably close to ideal is an ultracold gas of strongly coupled fermionic atoms, whose η/s is also $\ll 1$ and may be comparably small to that of the quark-gluon plasma produced at RHIC [120]. Both these fluids are much better described by ideal hydrodynamics than water is. Both have η/s comparable to the value $1/4\pi$ that, as we shall see in Section 6.2, characterizes any strongly coupled gauge theory plasma with a gravity dual in the large number of colors limit.

We note that while hydrodynamic calculations reproduce elliptic flow, a treatment in which the Boltzmann equation for quark and gluon (quasi)particles is solved, including all $2 \rightarrow 2$ scattering processes with the cross-sections as calculated in perturbative QCD, fails dramatically. It results in values of v_2 that are much smaller than in the data. Agreement with data can only be achieved if the parton scattering cross-sections are increased *ad hoc* by more than a factor of 10 [121]. With such large cross-sections, a Boltzmann description cannot be reliable since the mean free path of the particles becomes comparable to or smaller than the interparticle spacing. Another way of reaching the same conclusion is to note that if a perturbative description of the QGP as a gas of interacting quasiparticles is valid, the effective QCD coupling α_s describing the interaction among these quasiparticles must be small, and for small α_s perturbative calculations of η/s are controlled and yield parametrically large values $\propto 1/\alpha_s^2 \ln \alpha_s$. It is not possible to get as small a value of η/s as the data requires

³Recent analyses that include equations of state based upon parametrizations of recent lattice calculations [119] and include dissipative effects from the late-time hadron gas epoch as well as estimates of the effects of bulk viscosity [108, 109], all of which turn out to be small, have maintained these conclusions: η/s seems to lie within the range $(1 - 2.5)/(4\pi)$ with the greatest remaining source of uncertainty coming from our lack of knowledge of the initial energy density profile [114].

from the perturbative calculation without increasing α_s to the point that the calculation is invalid. In contrast, as we shall see in Section 6.2, any gauge theory with a gravity dual must have $\eta/s = 1/4\pi$ in the large- N_c and strong coupling limit and, furthermore, the plasma fluids described by these theories in this limit do not have any well-defined quasiparticles. This calculational framework thus seems to do a much better job of capturing the qualitative features needed for a successful phenomenology of collective flow in heavy ion collisions.

As we completed the preparation of this review, the first measurement of the elliptic flow v_2 in heavy ion collisions with $\sqrt{s} = 2.76$ TeV was reported by the ALICE collaboration at the LHC [122]. It is striking that the $v_2(p_T)$ that they find for charged particles in three different impact parameter bins agrees within error bars at all values of p_T out to beyond 4 GeV with that measured at $\sqrt{s} = 0.2$ TeV by the STAR collaboration at RHIC. Quantitative comparison of these data to hydrodynamic calculations is expected from various authors in the coming months, but at a qualitative level these data indicate that the quark-gluon plasma produced at the LHC is comparably strongly coupled, with comparably small η/s , to that produced and studied at RHIC.

2.3 Jet quenching

Having learned that heavy ion collisions produce a low viscosity, strongly coupled, fluid we now turn to experimental observables with which we may study properties of the fluid beyond just how it flows. There are many such observables available. In this subsection and the next we shall describe two classes of observables, selected because in both cases there is (the promise of) a substantive interplay between data from RHIC (and the LHC) and qualitative insights gained from the analysis of strongly coupled plasmas with dual gravity descriptions.

Jet quenching refers to a suite of experimental observables that together reveal what happens when a very energetic quark or gluon (with momentum much greater than the temperature) plows through the strongly coupled plasma. Some measurements focus on how rapidly the energetic parton loses its energy; other measurements give access to how the strongly coupled fluid responds to the energetic parton passing through it. These energetic partons are not external probes; they are produced within the same collision that produces the strongly coupled plasma itself.

In a small fraction of *proton-proton* collisions at $\sqrt{s} = 200$ GeV, partons from the incident protons scatter with a large momentum transfer, producing back-to-back partons in the final state with transverse momenta of order ten or a few tens of GeV. These “hard” processes are rare, but data samples are large enough that they are nevertheless well studied. The high transverse momentum partons in the final state manifest themselves in the detector as jets. Individual high- p_T hadrons in the final state come from such hard processes and are typically found within jets. In addition to copious data from proton-(anti)proton collisions, there is a highly developed quantitatively controlled calculational framework built upon perturbative QCD that is used to calculate the rates for hard processes in high energy hadron-hadron collisions. These calculations are built upon factorization theorems. Consider as an example the single inclusive charged hadron spectrum at high- p_T , Fig. 2.1 (right). That is, the production cross-section for a single charged hadron with a given high transverse momentum p_T , regardless of what else is produced in the hadron-hadron collision. This quantity is

calculated as a convolution of separate (factorized) functions that describe different aspects of the process: (i) the process-independent parton distribution function gives the probability of finding partons with a given momentum fraction in the incident hadrons; (ii) the process-dependent hard scattering cross-section gives the probability that those partons scatter into final state partons with specified momenta; and (iii) the process-independent parton fragmentation functions that describe the probability that a final state parton fragments into a jet that includes a charged hadron with transverse momentum p_T . Functions (i) and (iii) are well-measured and at high transverse momentum function (ii) is both systematically calculated and well-measured. This body of knowledge provides a firm foundation, a well-defined baseline with respect to which we can measure changes if such a hard scattering process occurs instead in an ultrarelativistic heavy ion collision.

In hard scattering processes in which the momentum transfer Q is high enough, the partonic hard scattering cross section (function (ii) above) is expected to be the same in an ultrarelativistic heavy ion collision as in a proton-proton collision. This is so because the hard interaction occurs on a timescale and length scale $\propto 1/Q$ which is too short to resolve any aspects of the hot and dense strongly interacting medium that is created in the same collision. The parton distribution functions (function (i) above) are different in nuclei than in nucleons, but they may be measured in proton-nucleus, deuteron-nucleus, and electron-nucleus collisions. The key phenomenon that is unique to ultrarelativistic nucleus-nucleus collisions is that after a very energetic parton is produced, unless it is produced at the edge of the fireball heading outwards it must propagate through as much as 5-10 fm of the hot and dense medium produced in the collision. These hard partons therefore serve as well-calibrated probes of the strongly coupled plasma whose properties we are interested in. The presence of the medium results in the hard parton losing energy and changing the direction of its momentum. The change in the direction of its momentum is often referred to as “transverse momentum broadening”, a phrase which needs explanation. “Transverse” here means perpendicular to the original direction of the hard parton. (This is different from p_T , the component of the (original) momentum of the parton that is perpendicular to the beam direction.) “Broadening” refers to the effect on a jet when the directions of the momenta of many hard partons within it are kicked; averaged over many partons in one jet, or perhaps in an ensemble of jets, there is no change in the mean momentum but the spread of the momenta of the individual partons broadens.

Because the rates for hard scattering processes drop rapidly with increasing p_T , energy loss translates into a reduction in the number of partons produced with a given p_T . (Partons with the given p_T must have been produced with a higher p_T , and are therefore rarer than they would be in proton-proton collisions.) Transverse momentum broadening is expected to lead to more subtle modifications of jet properties. Furthermore, the hard parton dumps energy into the medium, which motivates the use of observables involving correlations between soft final state hadrons and a high momentum hadron. Most generally, “jet quenching” refers to the whole suite of medium-induced modifications of high- p_T processes in heavy ion collisions and modifications of the medium in heavy ion collisions in which a high- p_T process occurs, all of which have their origin in the propagation of a highly energetic parton through the strongly coupled plasma.

The most pictorial, although not the most generic, manifestation of jet quenching in heavy

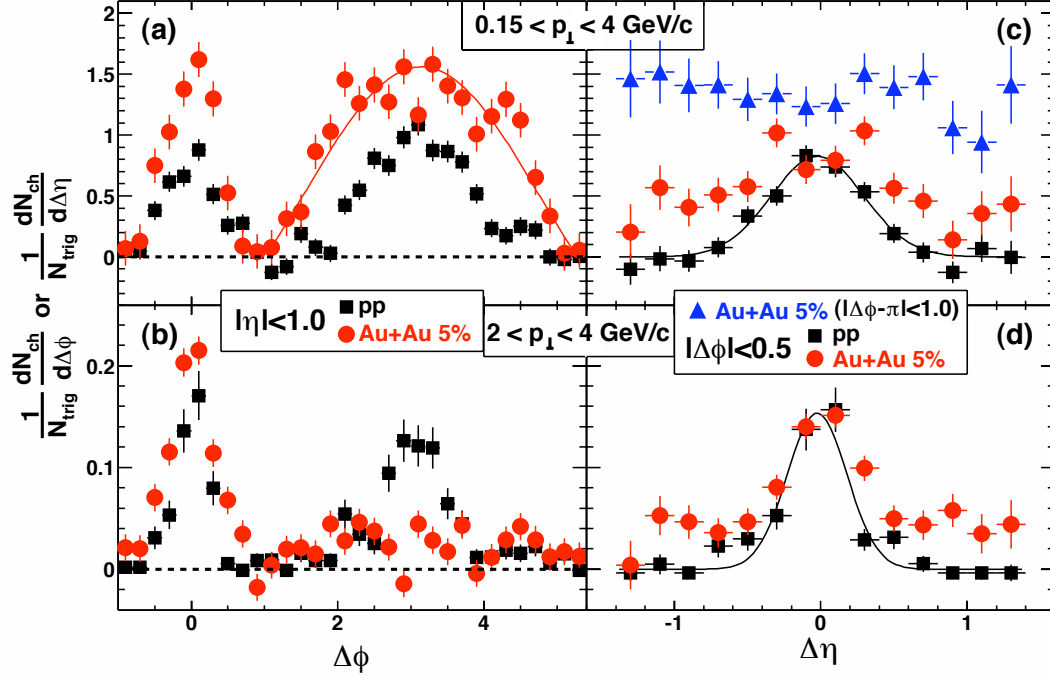


Figure 2.9: Dihadron azimuthal and longitudinal two-particle correlations in RHIC collisions with $\sqrt{s} = 200 \text{ GeV}$ in the STAR detector for high trigger p_T , $4 < p_T^{\text{trig}} < 6 \text{ GeV}$, and low (panels (a) and (c)) and intermediate (panels (b) and (d)) associated transverse momenta. In each panel, data from nearly head-on heavy ion collisions (red) are compared to data from proton-proton collisions (black). Panels (a) and (b) show the correlation in $\Delta\phi$, the difference in azimuthal angle of the two particles. Panels (c) and (d) show the correlation in pseudorapidity for those particles that are close to each other in azimuth (red and black points) or close to back to back (blue points). Figure taken from [123].

ion collisions at RHIC is provided by the data in Fig. 2.9. In this analysis, one selects events in which there is a hadron in the final state with $p_T > 4$ GeV, the trigger hadron. The distribution of all other hadrons with $p_T > 2$ GeV in azimuthal angle relative to the trigger hadron is shown in panel (b). In proton-proton collisions, we see peaks at 0 and π radians, corresponding to the jet in which the trigger hadron was found and the jet produced back to back with it. In contrast, in gold-gold collisions the away-side jet is missing. In panel (a), this analysis is repeated for *all* hadrons, not just those with $p_T > 2$ GeV. Here, we see an enhancement in gold-gold collisions on the away side. The picture that these data convey is that we have triggered on events in which one jet escapes the medium while the parton going the other way loses so much energy that it produces no hadrons with $p_T > 2$ GeV, dumping its energy instead into soft particles. The jet has been quenched and the medium has been modified. This picture is further supported by the longitudinal rapidity distribution shown in Fig. 2.9 (c) and (d). On the near-side, the rapidity distribution shows the shape of a jet-like structure. On the away side, however, one finds an enhanced multiplicity distribution whose shape over many units in pseudo-rapidity is consistent with background. This measurement is illustrative, but not generic in the sense that it depends on the choice of p_T cuts used in the analysis. A more generic consequence of jet quenching, equally direct but less pictorial than that in Fig. 2.9, is provided by simply measuring the diminution of the number of high- p_T hadrons observed in heavy ion collisions. We shall turn to this momentarily.

At the time of writing, the first results on jet quenching in heavy ion collisions at $\sqrt{s} = 2.76$ TeV at the LHC are just becoming available [124, 125]. It is too soon to offer a quantitative analysis in a review such as this one, but it is clear from these early results that jet quenching remains strong even for the much higher energy jets that are produced in hard parton-parton scattering at these much higher collision energies. And, in LHC collisions evidence of highly asymmetric dijets [124], as if one parton escaped relatively unscathed while its back-to-back partner was very significantly degraded by the presence of the medium, can now be obtained from single events rather than statistically as in Fig. 2.9.

2.3.1 Single inclusive high-pt spectra and “jet” measurements

The RHIC heavy ion program has established that the measurement of single inclusive hadronic spectra yields a generic manifestation of jet quenching. Because the spectra in hadron-hadron collisions are steeply falling functions of p_T , if the hard partons produced in a heavy ion collision lose energy as they propagate through the strongly coupled plasma shifting the spectra leftward — to lower energy — is equivalent to depressing them. This effect is quantified via the measurement of the nuclear modification factor R_{AB}^h , which characterizes how the number of hadrons h produced in a collision between nucleus A and nucleus B differs from the number produced in an equivalent number of proton-proton collisions:

$$R_{AB}^h(p_T, \eta, \text{centrality}) = \frac{\frac{dN_{\text{medium}}^{AB \rightarrow h}}{dp_T d\eta}}{\langle N_{\text{coll}}^{AB} \rangle \frac{dN_{\text{vacuum}}^{pp \rightarrow h}}{dp_T d\eta}}. \quad (2.24)$$

Here, $\langle N_{\text{coll}}^{AB} \rangle$ is the average number of inelastic nucleon-nucleon collisions in A - B collisions within a specified range of centralities. This number is typically determined by inferring the transverse density distribution of nucleons in a nucleus from the known radial density profile

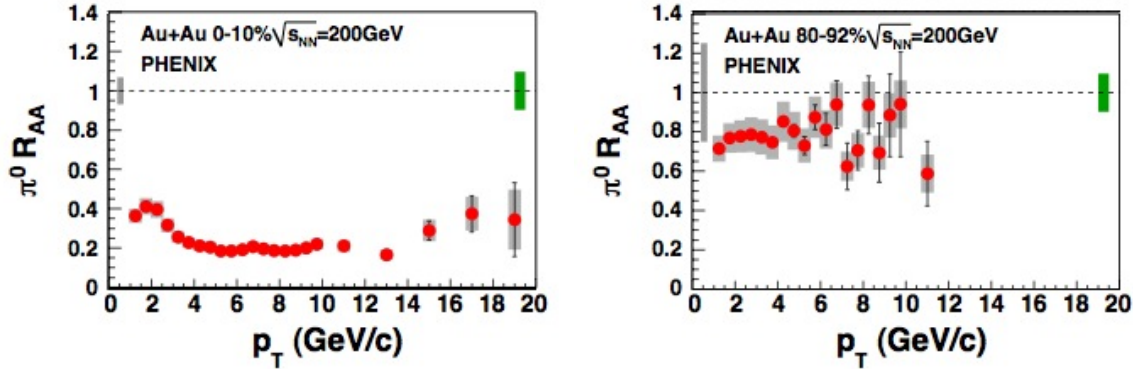


Figure 2.10: R_{AA} for neutral pions as a function of p_T for central (left) and peripheral (right) collisions. Data taken from [126].

of nuclei, and then calculating the average number of collisions with the help of the inelastic nucleon-nucleon cross section. This so-called Glauber calculation can be checked experimentally by independent means, for instance via the measurement of the nuclear modification factor for photons discussed below.

The nuclear modification factor depends in general on the transverse momentum p_T and pseudo-rapidity η of the particle, the particle identity h , the centrality of the collision and the orientation of the particle trajectory with respect to the reaction plane (which is often averaged over). If R_{AB} deviates from 1 this reflects either medium effects or initial state effects — the parton distributions in A and B need not be simply related to those in correspondingly many protons. Measurements of R_{dA} in deuteron- A collisions — which is a good proxy for R_{pA} — are used to determine whether an observed deviation of R_{AA} from 1 is due to initial state effects or the effects of parton energy loss in medium.

At mid-rapidity, RHIC data on R_{AuAu} (which is often written as R_{AA}) show the following generic features:

1. Characteristic strong centrality dependence of R_{AA} .

By varying the centrality of a heavy ion collision, one changes the typical in-medium path length over which hard partons produced in these collisions must propagate through the dense matter. For the most central head-on collisions (e.g., 0-10 % centrality), the average L is large, for a peripheral collision (e.g., 80-92 % centrality), the average L is small. RHIC data (see Fig. 2.10) show that for the most peripheral centrality bin, the nuclear modification factors are consistent with the absence of medium-effects, while R_{AA} decreases monotonically with increasing centrality and reaches about 0.2 — suppression by a factor of five! — for the most central collisions [127–130]. The left panel in Fig. 2.10 is a direct manifestation of jet quenching: 80% of the hard π^0 's that would be seen in the absence of a medium are gone.

2. Jet quenching is not observed in R_{dAu} .

In deuteron-gold collisions, R_{dAu} is consistent with or greater than 1 for all centralities

and all transverse momenta.⁴ Jet quenching is not observed. In fact, the centrality dependence is opposite to that seen in gold-gold collisions, with R_{dAu} reaching maximal values of around 1.5 for $p_T = 3 - 5$ GeV/ c in the most central collisions [128, 130]. The high- p_T hadrons are measured at or near mid-rapidity, meaning that they are well separated from the fragments of the struck gold nucleus. And, d-Au collisions do not produce a medium in the final state. In these collisions, therefore, the partons produced in hard scattering processes and tallied in R_{dAu} do not have to propagate through any matter after they are produced. The fact that R_{dAu} is consistent with or greater than 1 in these collisions therefore demonstrates that the jet quenching measured in R_{AuAu} is attributable to the propagation of the hard partons produced in heavy ion collisions through the medium that is present only in those collisions.

3. Photons are not quenched.

For single inclusive photon spectra, the nuclear modification factor shows only mild deviations from $R_{AuAu}^\gamma \approx 1$ [131]. Within errors, these are consistent with perturbative predictions that take into account the nuclear modifications of parton distribution functions (mainly the isospin difference between protons and neutrons) [132]. Since photons, unlike partons or hadrons, do not interact strongly with the medium, this gives independent support that the jet quenching observed in heavy ion collisions is a final state effect. And, it provides experimental evidence in support of the Glauber-type calculation of the factor $\langle N_{coll}^{AA} \rangle$ in (2.24) discussed above.

4. Apparently p_T -independent and species-independent suppression of R_{AA} at high p_T .

At high transverse momentum, $p_T > 7$ GeV, the suppression factor is approximately independent of p_T (see Fig. 2.10) in all centrality bins [60, 133, 134]. R_{AuAu} is suppressed all the way out to the highest p_T in Fig. 2.10, p_T that are high enough that hadronization of the parton must occur only far outside the medium.⁵ Moreover, R_{AuAu}^h is independent of the species of the hadron h [135]. Either of these observations eliminates the possibility that hadrons are formed within the medium and then lose energy upon propagating through the medium, since different hadrons would have different cross sections for interaction with the medium. This data supports the picture that the origin of the observed suppression is energy loss by a parton propagating through the medium prior to its hadronization.

5. R_{AA} for heavy-flavored and light-flavored hadrons is comparable.

The current experiments at RHIC measure charm and bottom quarks via their decay products, typically via the measurement of single electrons produced in the weak decays of those quarks [136]. These analyses cannot separate the contributions from charm and bottom quarks, since both have decay modes containing single electrons. The interest in the dependence of parton energy loss on parton mass arises from the fact that models predict a significant dependence. It is a matter of ongoing discussion to what extent the

⁴While their physical origin is not unambiguously constrained, the small deviations from unity observed in R_{dAu} reflect the presence of initial state effects in the colliding nuclei. For instance any transverse momentum broadening of partons prior to hard processes is expected to deplete R_{dAu} at low p_T and enhance it at higher p_T . Alternatively, modifications in the nuclear parton distribution functions could also lead to such enhancement.

⁵We shall see later that the fact that the suppression of R_{AA} is almost p_T -independent arises from a combination of various effects. Here we stress only the fact that it is small all the way out to the highest p_T .

uncertainties in the existing data are already small enough to put interesting constraints on models of parton energy loss. More progress can be expected in the near future, once detector upgrades at RHIC and measurements at the LHC allow for differentiation of bottom quarks via their displaced decay vertices.

In short, these observations support a picture in which highly energetic partons are produced in high momentum transfer processes within the medium as if they were in the vacuum, but where these partons subsequently lose a significant fraction of their initial energy due to interactions with the medium. Jet quenching is a partonic final state effect that depends on the length of the medium through which the parton must propagate. It is expected to have many consequences in addition to the strong suppression of single inclusive hadron spectra, which tend to be dominated by the most energetic hadronic fragments of parent partons. Rather, the entire parton fragmentation process is expected to be modified, with consequences for observables including multi-particle jet-like correlations and calorimetric jet measurements. (The early results from heavy ion collisions at the LHC [124, 125] confirm the expectation that at high enough jet energy calorimetric jet observables are modified by the presence of the medium [124].) We shall not review recent experimental and theoretical efforts to characterize multi-particle effects of jet quenching here, but we shall touch upon them briefly in Section 7.6. Furthermore, the energy deposited into the medium by the energetic parton also has interesting and potentially observable effects. We shall discuss these in Section 7.3; the interested reader is also referred to the literature [137, 138].

2.3.2 Analyzing jet quenching

For concreteness, we shall focus in this section on those aspects of the analysis of jet quenching that bear upon the calculation of the nuclear modification factor R_{AA} defined in (2.24). We shall describe other aspects of the analysis of jet quenching more briefly, as needed, in subsequent sections. The single inclusive hadron spectra which define R_{AA} are typically calculated upon assuming that the modification of the spectra in nucleus-nucleus collisions relative to that in proton-proton collisions arises due to parton energy loss. This assumption is well supported by data, as we have described above. But, from a theoretical point of view it is an assumption, not backed up by any formal factorization theorem. Upon making this assumption, we write

$$d\sigma_{(\text{med})}^{AA \rightarrow h + \text{rest}} = \sum_f d\sigma_{(\text{vac})}^{AA \rightarrow f + X} \otimes P_f(\Delta E, L, \hat{q}, \dots) \otimes D_{f \rightarrow h}^{(\text{vac})}(z, \mu_F^2). \quad (2.25)$$

Here,

$$d\sigma_{(\text{vac})}^{AA \rightarrow f + X} = \sum_{ijk} f_{i/A}(x_1, Q^2) \otimes f_{j/A}(x_2, Q^2) \otimes \hat{\sigma}_{ij \rightarrow f + k} \quad (2.26)$$

and $f_{i/A}(x, Q^2)$ are the nuclear parton distribution functions and $\sigma_{ij \rightarrow f + k}$ are the perturbatively calculable partonic cross sections. The medium dependence enters via the function $P_f(\Delta E, L, \hat{q}, \dots)$, which characterizes the probability that a parton f produced with cross section $\sigma_{ij \rightarrow f + k}$ loses energy ΔE while propagating over a path length L in a medium. This probability depends of course on properties of the medium, which are represented schematically in this formula by the symbol \hat{q} , the jet quenching parameter. We shall see below that

in the high parton energy limit, the properties of the medium enter P_f only through one parameter, and in that limit \hat{q} can be defined precisely. At nonasymptotic parton energies, \hat{q} in (2.25) is a place-holder, representing all relevant attributes of the medium. It is often conventional to refer to the combination of P_f and $D_{f \rightarrow h}^{(\text{vac})}$ together as a modified fragmentation function. It is only in the limit of high parton energy where one can be sure that the parton emerges from the medium before fragmenting into hadrons in vacuum that these two functions can be cleanly separated as we have done in (2.25). This aspect of the ansatz (2.25) is supported by the data: as we have described above, all hadrons exhibit the same suppression factor indicating that R_{AA} is due to partonic energy loss, before hadronization.

The dynamics of how parton energy is lost to the medium is specified in terms of the probability $P_f(\Delta E, L, \hat{q}, \dots)$. In the high parton energy limit, the parton loses energy dominantly by inelastic processes that are the QCD analogue of bremsstrahlung: the parton radiates gluons as it interacts with the medium. It is a familiar fact from electromagnetism that bremsstrahlung dominates the loss of energy of an electron moving through matter in the high energy limit. The same is true in calculations of QCD parton energy loss in the high-energy limit, as established first in Refs. [139–141]. The hard parton undergoes multiple inelastic interactions with the spatially extended medium, and this induces gluon bremsstrahlung. Here and throughout, by the high parton energy limit we mean the combined set of limits that can be summarized as:

$$E \gg \omega \gg |\mathbf{k}|, |\mathbf{q}| \equiv \left| \sum_i \mathbf{q}_i \right| \gg T, \Lambda_{\text{QCD}} , \quad (2.27)$$

where E is the energy of the high energy projectile parton, where ω and \mathbf{k} are the typical energy and momentum of the gluons radiated in the elementary radiative processes $q \rightarrow qg$ or $g \rightarrow gg$, and where \mathbf{q} is the transverse momentum (transverse to its initial direction) accumulated by the projectile parton due to many radiative interactions in the medium, and where T and Λ_{QCD} represent any energy scales that characterize the properties of the medium itself. This set of approximations underlies all analytical calculations of radiative parton energy loss to date [140–145]. The premise of the analysis is the assumption that QCD at scales of order $|\mathbf{k}|$ and $|\mathbf{q}|$ is weakly coupled, even if the medium (with its lower characteristic energy scales of order T and Λ_{QCD}) is strongly coupled. We shall spend most of this section on the analysis valid in this high parton energy limit, in which case all we need ask of analyses of strongly coupled gauge theories with gravity duals is insight into those properties of the strongly coupled medium that enter into the calculation of jet quenching in QCD. However, this analysis based upon the limits (2.27) may not be under quantitative control when applied to RHIC data, since at RHIC the partons in question have energies of at most a few tens of GeV, meaning that one can question whether all the scales separated by \gg in (2.27) are in fact well-separated. We shall close this section with a brief look at elastic scattering as an example of an energy loss process that is relevant when (2.27) is not satisfied. And, in Section 7.1 we shall see that at low enough energies that physics at all scales in the problem all the way up to E is strongly coupled (or in a conformal theory like $\mathcal{N} = 4$ SYM in which physics is strongly coupled at all scales) new approaches are needed. The analysis based upon (2.27) that we focus on in this section will be under better control when applied to LHC data, since in heavy ion collisions at the LHC partons with p_T all the way up to a few hundred GeV will be produced and therefore available as probes.

The analysis of radiative energy loss starts from (and extends) the eikonal formalism, so we must begin with a few ideas and some notation from this approach (for a self-contained introduction and references to earlier work, see e.g. Ref. [146,147]). As seen by a high energy parton, a target that is spatially extended but of finite thickness appears Lorentz contracted, so in the projectile rest frame the parton propagates through the target in a short period of time and the transverse position of the projectile does not change during the propagation. So, at ultra-relativistic energies, the main effect of the target on the projectile is a “rotation” of the parton’s color due to the color field of the target. These rotation phases are given by Wilson lines along the (straight line) trajectories of the propagating projectile:

$$W(\mathbf{x}) = \mathcal{P} \exp\left\{i \int dz^- T^a A_a^+(\mathbf{x}, z^-)\right\}. \quad (2.28)$$

Here, \mathbf{x} is the *transverse* position of the projectile — which does not change as the parton propagates at the speed of light along the $z^- \equiv (z - t)/\sqrt{2}$ lightlike direction. A^+ is the large component of the target color field and T^a is the generator of $SU(N)$ in the representation corresponding to the given projectile — fundamental if the hard parton is a quark and adjoint if it is a gluon. The eikonal approach to scattering treats the (unphysical, in the case of colored projectiles) setting in which the projectile impinges on the target from outside, after propagating for an arbitrarily long time and building up a fully developed Weizsäcker-Williams field proportional to $g \frac{\mathbf{x}}{\mathbf{x}^2}$ (a coherent state cloud of gluons dressing the bare projectile). The interaction of this dressed projectile with the target results in an eikonal phase (Wilson line) for the projectile itself and for each gluon in the cloud. Gluon radiation then corresponds to the decoherence of components of the dressed projectile that pick up different phases. Analysis of this problem yields a calculation of $N_{\text{prod}}(\mathbf{k})$, the number of radiated gluons with momentum \mathbf{k} , with the result: [146,147]

$$\begin{aligned} N_{\text{prod}}(\mathbf{k}) = & \\ & \frac{\alpha_s C_F}{2\pi} \int d\mathbf{x} d\mathbf{y} e^{i\mathbf{k}\cdot(\mathbf{x}-\mathbf{y})} \frac{\mathbf{x}\cdot\mathbf{y}}{\mathbf{x}^2 \mathbf{y}^2} \left[1 - \frac{1}{N^2 - 1} \langle \text{Tr} [W^{A\dagger}(\mathbf{x}) W^A(\mathbf{0})] \rangle \right. \\ & - \frac{1}{N^2 - 1} \langle \text{Tr} [W^{A\dagger}(\mathbf{y}) W^A(\mathbf{0})] \rangle \\ & \left. + \frac{1}{N^2 - 1} \langle \text{Tr} [W^{A\dagger}(\mathbf{y}) W^A(\mathbf{x})] \rangle \right], \quad (2.29) \end{aligned}$$

where the C_F prefactor is for the case where the projectile is a quark in the fundamental representation, where the projectile is located at transverse position $\mathbf{0}$, and where the $\langle \dots \rangle$ denotes averaging over the gluon fields of the target. If the target is in thermal equilibrium, these are thermal averages.

Although the simple result (2.29) is not applicable to the physically relevant case as we shall describe in detail below, we can nevertheless glean insights from it that will prove relevant. We note that the entire medium-dependence of the gluon number spectrum (2.29) is determined by target expectation values of the form $\langle \text{Tr} [W^{A\dagger}(\mathbf{x}) W^A(\mathbf{y})] \rangle$ of two eikonal Wilson lines. The jet quenching parameter \hat{q} that will appear below defines the fall-off properties of this

correlation function in the transverse direction $L \equiv |\mathbf{x} - \mathbf{y}|$:

$$\langle \text{Tr} [W^A(\mathcal{C})] \rangle \approx \exp \left[-\frac{1}{4\sqrt{2}} \hat{q} L^- L^2 \right] \quad (2.30)$$

in the limit of small L , with L^- (the extent of the target along the z^- direction) assumed large but finite [148, 149]. Here, the contour \mathcal{C} traverses a distance L^- along the light cone at transverse position \mathbf{x} , and it returns at transverse position \mathbf{y} . These two long straight lightlike lines are connected by short transverse segments located at $z^- = \pm L^-/2$, far outside the target. We see from the form of (2.29) that $|\mathbf{k}|$ and L are conjugate: the radiation of gluons with momentum $|\mathbf{k}|$ is determined by Wilson loops with transverse extent $L \sim 1/|\mathbf{k}|$. This means that in the limit (2.27), the only property of the medium that enters (2.29) is \hat{q} . Furthermore, inserting (2.30) into (2.29) yields the result that the gluons that are produced have a typical \mathbf{k}^2 that is of order $\hat{q}L^-$. This suggests that \hat{q} can be interpreted as the transverse momentum squared picked up by the hard parton per distance L^- that it travels, an interpretation that can be validated more rigorously via other calculations [150, 151].

The reason that the eikonal formalism cannot be applied verbatim to the problem of parton energy loss in heavy ion collisions is that the high energy partons we wish to study do not impinge on the target from some distant production site. They are produced within the same collision that produces the medium whose properties they subsequently probe. As a consequence, they are produced with significant virtuality. This means that even if there were no medium present, they would radiate copiously. They would fragment in what is known in QCD as a parton shower. The analysis of medium-induced parton energy loss then requires understanding the interference between radiation in vacuum and the medium-induced bremsstrahlung radiation. It turns out that the resulting interference resolves longitudinal distances in the target [140–143], meaning that its description goes beyond the eikonal approximation. The analysis of parton energy loss in the high energy limit (2.27) must include terms that are subleading in $1/E$, and therefore not present in the eikonal approximation, that describe the leading interference effects. To keep these $\mathcal{O}(1/E)$ effects, one must replace eikonal Wilson lines by retarded Green's functions that describe the propagation of a particle with energy E from position z_1^-, \mathbf{x}_1 to position z_2^-, \mathbf{x}_2 without assuming $\mathbf{x}_1 = \mathbf{x}_2$ [141, 152, 153]. (In the $E \rightarrow \infty$ limit, $\mathbf{x}_1 = \mathbf{x}_2$ and the eikonal Wilson line is recovered.) It nevertheless turns out that even after Wilson lines are replaced by Green's functions the only attribute of the medium that arises in the analysis, in the limit (2.27), is the jet quenching parameter \hat{q} defined in (2.30) that already arose in the eikonal approximation.

We shall not present the derivation, but it is worth giving the complete (albeit somewhat formal) result for the distribution of gluons with energy ω and transverse momentum \mathbf{k} that a high energy parton produced within a medium radiates:

$$\begin{aligned} \omega \frac{dI}{d\omega d\mathbf{k}} &= \frac{\alpha_s C_R}{(2\pi)^2 \omega^2} 2\text{Re} \int_{\xi_0}^{\infty} dy_l \int_{y_l}^{\infty} d\bar{y}_l \int d\mathbf{u} e^{-i\mathbf{k}\cdot\mathbf{u}} \exp \left[-\frac{1}{4} \int_{\bar{y}_l}^{\infty} d\xi \hat{q}(\xi) \mathbf{u}^2 \right] \\ &\times \frac{\partial}{\partial \mathbf{x}} \cdot \frac{\partial}{\partial \mathbf{u}} \int_{\mathbf{x}=\mathbf{r}(y_l)=\mathbf{0}}^{\mathbf{u}=\mathbf{r}(\bar{y}_l)} \mathcal{D}\mathbf{r} \exp \left[\int_{y_l}^{\bar{y}_l} d\xi \left(\frac{i\omega}{2} \dot{\mathbf{r}}^2 - \frac{1}{4} \hat{q}(\xi) \mathbf{r}^2 \right) \right]. \end{aligned} \quad (2.31)$$

We now walk through the notation in this expression. The Casimir operator C_R is in the representation of the projectile parton. The integration variables ξ , y_l and \bar{y}_l are all positions

along the z^- lightcone direction. ξ_0 is the z^- at which the projectile parton was created in a hard scattering process. Since we are not taking this to $-\infty$, the projectile is not assumed on shell. The projectile parton was created at the transverse position $\mathbf{x} = \mathbf{0}$. The integration variable \mathbf{u} is also a transverse position variable, conjugate to \mathbf{k} . The path integral is over all possible paths $\mathbf{r}(\xi)$ going from $\mathbf{r}(y_l) = \mathbf{0}$ to $\mathbf{r}(\bar{y}_l) = \mathbf{u}$. The derivation of (2.31) proceeds by writing $dI/d\omega d\mathbf{k}$ in terms of a pair of retarded Green's functions in their path-integral representations, one of which describes the radiated gluon in the amplitude, radiated at y_l , and the other of which describes the radiated gluon in the conjugate amplitude, radiated at \bar{y}_l . The expression (2.31) then follows after a lengthy but purely technical calculation [142]. The properties of the medium enter (2.31) only through the jet quenching parameter $\hat{q}(\xi)$. The compact expression (2.31) has been derived in the so-called path-integral approach [141]. For other, related, formulations of QCD parton energy loss, we refer the reader to Refs. [83, 144, 145, 150, 154–174].

The notation $\hat{q}(\xi)$ allows for the possibility that the nature of the medium and thus its \hat{q} changes with time as the hard parton propagates through it. If we approximate the medium as unchanging, $\hat{q}(\xi)$ is just a constant. In the strongly expanding medium of a heavy ion collision, however, this is not a good approximation. It turns out that the path-integral in (2.31) can be solved analytically in the saddle point approximation for quenching parameters of the form $\hat{q}(\xi) = \hat{q}_0 (\xi_0/\xi)^\alpha$ [175]. Here the free exponent α may be chosen in model calculations to take values characterizing a one-dimensional longitudinal expansion with so-called Bjorken scaling ($\alpha = 1$), or scenarios which also account for the transverse expansion $1 < \alpha < 3$. Remarkably, one finds that irrespective of the value of α , for fixed in-medium path length $L^-/\sqrt{2}$ the transverse momentum integrated gluon energy distribution (2.31) has the same ω -dependence if $\hat{q}(\xi)$ is simply replaced by a constant given by the linear line-averaged transport coefficient [176]

$$\langle \hat{q} \rangle \equiv \frac{1}{2 L^{-2}} \int_{\xi_0}^{\xi_0 + L^-} d\xi (\xi - \xi_0) \hat{q}(\xi). \quad (2.32)$$

In practice, this means that comparisons of different parton energy loss calculations to data can be performed as if the medium were static. The line-averaged transport coefficient $\langle \hat{q} \rangle$ determined in this way can then be related via (2.32) to the transport coefficient at a given time, once a model for the expansion of the medium is specified. Hence, we can continue our discussion for the case $\hat{q}(\xi) = \hat{q}$ without loss of generality.

The result (2.31) is both formal and complicated. However, its central qualitative consequences can be characterized almost by dimensional analysis. All dimensionful quantities can be scaled out of (2.31) if ω is measured in units of the so-called characteristic gluon energy

$$\omega_c \equiv \hat{q}(L^-)^2 \quad (2.33)$$

and the transverse momentum \mathbf{k}^2 in units of $\hat{q}L^-$ [177]. In a numerical analysis of (2.31), one finds that the transverse momentum distribution of radiated gluons scales indeed with $\hat{q}L^-$, as expected for the transverse momentum due to the Brownian motion in momentum space that is induced by multiple small angle scatterings. If one integrates the gluon distribution (2.31) over transverse momentum and takes the upper limit of the \mathbf{k} -integration to infinity, one recovers [177] an analytical expression first derived by Baier, Dokshitzer, Mueller, Peigné

and Schiff [140]:

$$\omega \frac{dI_{\text{BDMPS}}}{d\omega} = \frac{2\alpha_s C_R}{\pi} \ln \left| \cos \left[(1+i) \sqrt{\frac{\omega_c}{2\omega}} \right] \right|, \quad (2.34)$$

which yields the limiting cases

$$\omega \frac{dI_{\text{BDMPS}}}{d\omega} \simeq \frac{2\alpha_s C_R}{\pi} \begin{cases} \sqrt{\frac{\omega_c}{2\omega}} & \text{for } \omega \ll \omega_c, \\ \frac{1}{12} \left(\frac{\omega_c}{\omega}\right)^2 & \text{for } \omega \gg \omega_c, \end{cases} \quad (2.35)$$

for small and large gluon energies. In the soft gluon limit, the BDMPS spectrum (2.34) displays the characteristic $1/\sqrt{\omega}$ dependence, which persists up to a gluon energy of the order of the characteristic gluon energy (2.33). Hence, ω_c can be viewed as an effective energy cut-off, above which the contribution of medium-induced gluon radiation is negligible. These analytical limits provide a rather accurate characterization of the full numerical result. In particular, one expects from the above expressions that the average parton energy loss $\langle \Delta E \rangle$, obtained by integrating (2.31) over \mathbf{k} and ω , is proportional to $\propto \int_0^{\omega_c} d\omega \sqrt{\omega_c}/\omega \propto \omega_c$. One finds indeed

$$\langle \Delta E \rangle_{\text{BDMPS}} \equiv \int_0^\infty d\omega \omega \frac{dI_{\text{BDMPS}}}{d\omega} = \frac{\alpha_s C_R}{2} \omega_c. \quad (2.36)$$

This is the well-known $(L^-)^2$ -dependence of the average radiative parton energy loss [140, 141, 178]. In summary, the main qualitative properties of the medium-induced gluon energy distribution (2.31) are the scaling of \mathbf{k}^2 with $\hat{q}L^-$ dictated by Brownian motion in transverse momentum space, the $1/\sqrt{\omega}$ dependence of the \mathbf{k} -integrated distribution characteristic of the non-abelian Landau-Pomeranchuk-Migdal (LPM) effect, and the resulting $(L^-)^2$ -dependence of the average parton energy loss.

Medium-induced gluon radiation modifies the correspondence between the initial parton momentum and the final hadron momenta. We now sketch how the resulting $P(\Delta E)$ in (2.25) can be estimated. If gluons are emitted independently, $P(\Delta E)$ is the normalized sum of the emission probabilities for an arbitrary number of n gluons which carry away the total energy ΔE : [179]

$$P(\Delta E) = \exp \left[- \int_0^\infty d\omega \frac{dI}{d\omega} \right] \sum_{n=0}^\infty \frac{1}{n!} \left[\prod_{i=1}^n \int d\omega_i \frac{dI(\omega_i)}{d\omega} \right] \delta \left(\Delta E - \sum_{i=1}^n \omega_i \right). \quad (2.37)$$

Here, the factor $\exp \left[- \int_0^\infty d\omega \frac{dI}{d\omega} \right]$ denotes the probability that no energy loss occurs. This factor ensures that $P(\Delta E)$ is properly normalized, $\int d\Delta E P(\Delta E) = 1$. The mean energy loss is then

$$\langle \Delta E \rangle = \int d\Delta E (\Delta E) P(\Delta E) = \int d\omega \omega \frac{dI}{d\omega}, \quad (2.38)$$

as above, but it turns out that $P(\Delta E)$ is a very broad distribution, not peaked around its mean. In particular, as seen from Fig. 2.1, single inclusive hadron spectra are distributions which fall steeply with p_T . The modifications which parton energy loss induce on such distributions cannot be characterized by an average energy loss. Rather, what matters for a steeply falling distribution is not how much energy a parton loses on average, but rather which fraction of all the partons gets away with much less than the average energy loss [179].




Figure 2.11: Data for the nuclear modification factor of π^0 in central 0 – 5% Au+Au collisions, compared to a model calculation of jet quenching with $\langle \hat{q} \rangle$ values of 0.3, 0.9, 1.2, 1.5, 2.1, 2.9, 4.4, 5.9, 7.4, 10.3, 13.2, 17.7, 25.0, 40.5 and 101.4 GeV^2/fm . Data taken from [180].

This so-called trigger bias effect is quantitatively very important, and can be accounted for by the probability distribution (2.37).

Making contact with phenomenology requires the implementation of $P(\Delta E)$ in a model in which hard scattering events are distributed with suitable probability at locations in the transverse plane, with the hard partons produced going in random back-to-back directions and in which a hydrodynamic calculation is used to model how the medium in which the hard partons find themselves evolves subsequently, as the partons propagate through it. One such model is the PQM model [181], which calculates the hadronic spectrum (2.25) by interfacing the probability distributions $P(\Delta E)$ of Ref. [177] with a model of the geometrical distribution of matter in the collision region. Although the PQM model is built upon the gluon energy distribution (2.35), which gives the leading contribution to parton energy loss in the high energy limit (2.27), it is nevertheless a model in several senses beyond its treatment of the distribution of energy density in space and time. For example, the factorized ansatz of the inclusive cross section (2.25) is not justified from first principles but is an experimentally testable assumption. Also, the form (2.37) of the probability distribution $P(\Delta E)$ assumes that in the case of multiple gluon emissions, the consecutive energy degradation of the leading parton is negligible. For this to be a good approximation, the total energy loss must be much smaller than the projectile energy, $\Delta E \ll E$. All these aspects of the model are expected to become reliable in the high energy limit (2.27). The results of a comparison of this model to data, done by the PHENIX collaboration, are shown in Fig. 2.11. We see that for a suitable choice of the line-averaged transport coefficient $\langle \hat{q} \rangle$, defined in (2.32), the calculation does a reasonable job of describing the data. In this analysis, the approximate p_T -independence of R_{AuAu} seen in the data arises from an interplay between the rapidly dropping parton production cross-section at p_T 's exceeding $\sim \sqrt{s}/10$, the fact that there are always some partons produced near the edge of the fireball that emerge relatively unscathed and the fact that $P(\Delta E)$ is significant even for $\Delta E = 0$ even though the mean ΔE is large. The PHENIX

authors quote for this PQM model study a jet quenching parameter which is constrained by the experimental data as $13.2_{-3.2}^{+2.1}$ and $13.2_{-5.2}^{+6.3}$ GeV²/fm at the one and two standard deviation levels, respectively.

More recently, the quantitative comparison between data from RHIC and the jet quenching calculations that we have sketched in this section has been revisited in Ref. [167]. These authors use data on R_{AA} as in Fig. 2.11 but in addition they use data on dihadron correlations as well as data on heavy quark suppression. And, they include a hydrodynamical modelling of the expanding and cooling plasma. They parametrize the jet quenching parameter as

$$\hat{q} = 2K\varepsilon^{3/4} \quad (2.39)$$

where ε is the energy density of the medium through which the energetic partons are propagating and where K is a parameter to be obtained via comparing jet quenching calculations to data. (In a weakly coupled quark-gluon plasma, $K \approx 1$ [156].) ε varies with position and decreases with time according to a hydrodynamical calculation. These authors find that they are able to obtain a more stable fit to the value of K than to a time-averaged \hat{q} . Fitting to data from both PHENIX and STAR from RHIC collisions at $\sqrt{s} = 200$ GeV, they obtain

$$K = 4.1 \pm 0.6 . \quad (2.40)$$

We will compare this result to calculations done for strongly coupled plasmas in gauge theories with dual gravitational descriptions in Section 7.5.

There is currently a vigorous debate about whether hard partons produced (and then quenched) in RHIC collisions have a high enough energy E for the approximation (2.27) to be under control, and for the above-mentioned model assumptions to be valid. The question of whether the comparison of RHIC data with different model implementations of parton energy loss yield numerically consistent results is also still debated, see e.g. [180]. Here, we solely remark that while radiative energy loss is known to dominate in the high energy limit, other mechanisms and in particular elastic energy loss may contribute significantly at lower energies, and this may not be negligible at RHIC [160, 182, 183]. In radiative energy loss, in the high energy limit, the longitudinal momentum of the incident parton is shared between the longitudinal momentum of the radiated gluon and that of the outgoing high energy parton. Longitudinal momentum transfer to the medium is negligible in comparison to that carried by the radiated gluon. In elastic scattering, however, the only energy lost by the projectile parton is that due to transferring longitudinal momentum to the medium. What is clear from the outset is that if elastic mechanisms contribute, then parton energy loss must depend on more attributes of the medium than just \hat{q} , since \hat{q} knows nothing about how “bits” of the medium recoil longitudinally when struck. Note that this discussion goes through unchanged even in the context of a medium that is a strongly coupled fluid with no identifiable scattering centers. If the medium is weakly coupled, then the conclusion can be stated more precisely: \hat{q} does not differentiate between a medium made of heavy scatterers that hardly recoil and one made of light scatterers that recoil more easily.

2.4 Quarkonia in hot matter

One way of thinking about the operational meaning of the statement that quark-gluon plasma is deconfined is to ask what prevents the formation of a meson within quark-gluon plasma.

The answer is that the attractive force between a quark and an antiquark which are separated by a distance of order the size of a meson is screened by the presence of the quark-gluon plasma between them. This poses a quantitative question: how close together do the quark and antiquark have to be in order for their attraction not to be screened? How close together do they have to be in order for them to feel the same attraction that they would feel if they were in vacuum? It was first suggested by Matsui and Satz [184] in 1986 that measurements of how many quarkonia — mesons made of a heavy quark-antiquark pair — are produced in heavy ion collisions could be used as a tool with which to answer this question, because they are significantly smaller than typical mesons or baryons.

The generic term quarkonium refers to the charm-anticharm or charmonium, mesons (J/Ψ , Ψ' , χ_c , ...) and the bottom-antibottom, or bottomonium, mesons (Υ , Υ' , ...). The first quarkonium state that was discovered was the $1s$ state of the $c\bar{c}$ bound system, the J/Ψ . It is roughly half the size of a typical meson like the ρ . The bottomonium $1s$ state, the Υ is smaller again by roughly another factor of two. It is therefore expected that if one can study quark-gluon plasma in a series of experiments with steadily increasing temperature, J/ψ mesons survive as bound states in the quark-gluon plasma up to some dissociation temperature that is higher than the crossover temperature (at which generic mesons and baryons made of light quarks fall apart.) More realistically, what Matsui and Satz suggested is that if high energy heavy ion collisions create deconfined quark-gluon plasma that is hot enough, then color screening would prevent charm and anticharm quarks from binding to each other in the deconfined interior of the droplet of matter produced in the collision, and as a result the number of J/Ψ mesons produced in the collisions would be suppressed. However, Υ s should survive as bound states to even higher temperatures, until the quark-antiquark attraction is screened even on the short length scale corresponding to their size.

To study this effect, Matsui and Satz suggested comparing the temperature dependence of the screening length for the quark-antiquark force, which can be obtained from lattice QCD calculations, with the J/Ψ meson radius calculated in charmonium models. They then discussed the feasibility to detect this effect clearly in the mass spectrum of e^+e^- dilepton pairs. Between 1986, when Matsui and Satz launched this line of investigation, suggesting it as a quantitative means of characterizing the formation and properties of deconfined matter, and today we know of no other measurement that has been advocated as a more direct experimental signature for the deconfinement transition. And, there is hardly any other measurement whose phenomenological analysis has turned out to be more involved. In this subsection, we shall describe both the appeal of studying quarkonia in the hot matter produced in heavy ion collisions and the practical difficulties. The theoretical basis for the argument of Matsui and Satz has evolved considerably within the last two decades [185]. Moreover, the debate over how to interpret these measurements is by now informed by data on J/Ψ -suppression in nucleus-nucleus collisions at the CERN SPS [186, 187], at RHIC [188] and at the LHC [189]. There is also a good possibility that qualitatively novel information will become accessible in future high statistics runs at RHIC and in heavy ion collisions at the LHC.

A sketch of the basic idea of Matsui and Satz is shown in Fig. 2.12. In very general terms, one expects that the attractive interaction between the heavy quark and anti-quark in a putative bound state is sensitive to the medium in which the heavy particles are embedded, and that this attraction weakens with increasing temperature. If the distance between the




Figure 2.12: Schematic picture of the dissociation of a $Q\bar{Q}$ -pair in hot QCD matter due to color screening. Figure taken from Ref. [185]. The straight black lines attached to the heavy Q and \bar{Q} indicate, that these quarks are external probes, in contrast to the dynamical quarks within the quark-gluon plasma.

heavy quark and anti-quark is much smaller than $1/T$, there will not be much quark-gluon plasma between them. Equivalently, typical momentum scales in the medium are of order the temperature T , and so the medium cannot resolve the separation between the quark-antiquark pair if they are much closer together than $1/T$. However, if the distance is larger, then the bound state is resolved, and the color charges of the heavy quarks are screened by the medium, see Fig. 2.12. In a heavy ion collision in which quark-gluon plasma with a temperature T is created, only those quarkonia with radii that are smaller than some length scale of order $1/T$ can form. These basic arguments support the idea that quarkonium production rates are an indicator of whether quark-gluon plasma is produced, and at what temperature.

In section 3.3, we review lattice calculations of the heavy quark static free energy $F_{Q\bar{Q}}(r)$. This static potential is typically defined via how the correlation function of a pair of Polyakov loops, namely test quarks at fixed spatial positions whose worldlines wrap around the periodic Euclidean time direction, falls off as the separation between the test quarks is increased. This static potential is renormalized such that it matches the zero temperature result at small distances. Calculations of $F_{Q\bar{Q}}(r)$ were the earliest lattice results which substantiated the core idea that a quarkonium bound state, placed in hot QCD matter, dissociates (“melts”) above a critical temperature. As we now discuss, phenomenological models of quarkonium in matter are based upon interpreting $F_{Q\bar{Q}}(r)$ as the potential in a Schrödinger equation whose eigenvalues and eigenfunctions describe the heavy $Q\bar{Q}$ bound states. There is no rigorous basis for this line of reasoning, and if pushed too far it faces various conceptual challenges as we shall discuss in Section 3.3. However, these models remain valuable as a source of semi-quantitative intuition.

At zero temperature, lattice results for $F_{Q\bar{Q}}(r)$ in QCD without dynamical quarks are well approximated by the ansatz $F_{Q\bar{Q}}(r) = \sigma r - \frac{\alpha}{r}$, where the linear term that dominates

at long distance is characterized by the string tension $\sigma \simeq 0.2 \text{ GeV}^2$ and the perturbative Coulomb term α/r is dominant at short distances. In QCD with dynamical quarks, beyond some radius r_c the potential flattens because as the distance between the external Q and \bar{Q} is increased, it becomes energetically favorable to break the color flux tube connecting them by producing a light quark-antiquark pair from the vacuum which, in a sense, screens the potential. With increasing temperature, the distance r_c decreases, that is, the colors of Q and \bar{Q} are screened from each other at increasingly shorter distances. This is seen clearly in Fig. 2.12 in Section 3.3. These lattice results are well parametrized by a screened potential of the form [185, 190]

$$F_{Q\bar{Q}}(r) = -\frac{\alpha}{r} + \sigma r \left(\frac{1 - e^{-\mu r}}{\mu r} \right), \quad (2.41)$$

where $\mu \equiv \mu(T)$ can be interpreted at high temperatures as a temperature-dependent Debye screening mass. For suitably chosen $\mu(T)$, this ansatz reproduces the flattening of the potential found in lattice calculations at the finite large distance value $F_{Q\bar{Q}}(\infty) = \sigma/\mu(T)$. Taking this $Q\bar{Q}$ free energy $F_{Q\bar{Q}}(r, T)$ as the potential in a Schrödinger equation, one may try to determine which bound states in this potential remain, as the potential is weakened as the temperature increases. Such potential model studies have led to predictions of the dissociation temperatures T_d of the charmonium family, which range from $T_d(J/\Psi) \simeq 2.1 T_c$ to $T_d(\Psi') \simeq 1.1 T_c$ for the more loosely bound and therefore larger $2s$ state. The deeply bound, small, $1s$ state of the bottomonium family is estimated to have a dissociation temperature $T_d(\Upsilon(1S)) > 4 T_c$, while dissociation temperatures for the corresponding $2s$ and $3s$ states were estimated to lie at $1.6 T_c$ and $1.2 T_c$ respectively [185, 190]. Because the leap from the static quark-antiquark potential to a Schrödinger equation is not rigorously justified, the uncertainties in quantitative results obtained from these potential models are difficult to estimate. (For more details on why this is so, see Section 3.3.) However, these models with their inputs from lattice QCD calculations do provide strong qualitative support for the central idea of Matsui and Satz that quarkonia melt in hot QCD matter, and they provide strong support for the qualitative expectation that this melting proceeds *sequentially*, with smaller bound states dissociating at a higher temperature.

Experimental data on the yield of J/Ψ mesons in nucleus-nucleus collisions have been reported by the NA50 [186] and NA60 [187] collaborations at the CERN SPS and by the PHENIX and STAR collaborations at RHIC. As one increases the size of the collision system (either by varying the impact parameter selection or by changing the nuclear species used in the experiments) these yields turn out to be increasingly suppressed, relative to the yields measured in proton-proton or proton-nucleus collisions at the same center of mass energy. The measurements made in proton-nucleus collisions are particularly important to use as a baseline, since without them one would not know how much charmonium suppression arises due to the interaction between the charm-anticharm pair and ordinary confined hadronic matter [191]. The operational procedure for separating such hadronic phase effects is to measure them separately in proton-nucleus collisions [192], and to establish then to what extent the number of J/Ψ mesons produced in nucleus-nucleus collisions drops below the yield extrapolated from proton-nucleus collisions [193]. We note that for nucleus-nucleus collisions, experimental data exist so far solely for the J/Ψ . The other bound states of the charmonium family (Ψ' , χ_c , ...) have not been characterized due to the lack of statistics and resolution. Also, a characterization of bottomonium states in nuclear matter is missing. Due

to the much higher production rates for heavy ion collisions at the TeV scale, this situation is expected to change at the LHC. Moreover, the higher luminosity of future RHIC runs may give some access to the bottomonium system. Thus, measurements expected in the near future may provide evidence for the sequential quarkonium suppression pattern, which is a generic prediction of all models of quarkonium suppression and which has not yet been tested.

The analysis of data on charmonium has to take into account a significant number of important confounding effects. Here, we cannot discuss the phenomenology of these effects in detail, but we provide a list of the most important ones:

1. *Contributions from the decays of excited states*

In proton-proton collisions, a significant fraction of the observed yield of J/Ψ mesons is known to arise from the production of excited states like the Ψ' and χ_c , which subsequently decay to J/Ψ . In a nucleus-nucleus collision, the suppression of the excited states is expected to set in at a lower temperature since these states are larger in size than the ground state J/Ψ . In particular, it has been proposed that the observed suppression of the J/Ψ mesons at RHIC and at the SPS may arise solely from the dissociation of the more loosely bound Ψ' and χ_c states [190], with the J/Ψ s themselves remaining bound in the quark-gluon plasma produced in all heavy ion collisions to date. Regardless of whether this conclusion turns out to be quantitatively correct, it is certainly apparent that in the absence of separate measurements of the production of the excited states, any conclusions about the observed J/Ψ -suppression require careful modeling of, and inferences about, the contribution of the decays of excited states.

2. *Collective dynamics of the heavy ion collision: “explosive expansion”*

Lattice calculations are done for heavy quark bound states which are at rest in a hot, static, medium. In heavy ion collisions, however, even if the droplet of hot matter equilibrates rapidly, its temperature drops quickly during the subsequent explosive expansion. The observed quarkonium suppression must therefore result from a suitable time average over a dynamical medium. This is challenging in many ways. One issue that arises is the question of how long a bound state must be immersed in a sufficiently hot heat bath in order to melt. Or, phrased better, how long must the temperature be above the dissociation temperature T_d in order to prevent an heavy quark and antiquark produced at the initial moment of the collision from binding to each other and forming a quarkonium meson?

3. *Collective dynamics of the heavy ion collision: “hot wind”*

Another issue that faces any data analysis is that quarkonium mesons may be produced moving with significant transverse momentum through the hot medium. In their own reference frame, the putative quarkonium meson sees a hot wind. Phenomenologically, the question arises whether this leads to a stronger suppression since the bound state sees some kind of blue-shifted heat bath (an idea, which we will refine in Section 7.7), or whether the bound state is less suppressed since it can escape the heat bath more quickly.

4. *Formation of quarkonium bound states*

Neither quarkonia nor equilibrated quark-gluon plasma are produced at time zero in a

heavy ion collision. Quarkonia have to form, for instance by a colored $c\bar{c}$ -pair radiating a gluon to turn into a color-singlet quarkonium state. This formation process is not understood in elementary interactions or in heavy ion collisions. However, since the formation process takes time, it is *a priori* unclear whether any observed quarkonium suppression is due to the effects of the hot QCD matter on a formed quarkonium bound state or on the precursor of such a bound state, which may have different attenuation properties in the hot medium. And, it is unclear whether the suppression is due to processes occurring after the liquid-like strongly coupled plasma in approximate local thermal equilibrium is formed or earlier, before equilibration.

5. *Recombination as novel mechanism of quarkonium formation*

QCD is flavor neutral and thus charm is produced in $c\bar{c}$ pairs in primary interactions. If the average number of pairs produced per heavy ion collision is $\lesssim 1$, then all charmonium mesons produced in heavy ion collisions must be made from a c and a \bar{c} produced in the same primary interaction. At RHIC and even more so at the LHC, however, more than one $c\bar{c}$ pair is produced per collision, raising the possibility of a new charmonium production mechanism in which a c and a \bar{c} from different primary $c\bar{c}$ -pairs meet and combine to form a charmonium meson [194]. If this novel quarkonium production mechanism were to become significant, it could reduce the quarkonium suppression or even turn it into quarkonium enhancement. However, early LHC data on J/Ψ production do not show signs of any enhancement [189].

We have included this section about quarkonium in the present review, since calculations based on the AdS/CFT correspondence and reviewed in Sections 7.7 and 8 have provided complementary information for phenomenological modeling, in particular by calculating heavy quark potentials within a moving heat bath and by determining mesonic dispersion relations. The above discussion illustrates the context in which such information is useful, but it also emphasizes that such information is not sufficient. An understanding of quarkonium production in heavy ion collisions relies on phenomenological modelling as the bridge between experimental observations and the theoretical analysis of the underlying properties of hot QCD matter.

Chapter 3

Results from lattice QCD

At very high temperature, where the QCD coupling constant $g(T)$ is perturbatively small, hard thermal loop resummed perturbation theory provides a quantitatively controlled approach to QCD thermodynamics. However, in a wide temperature range around the QCD phase transition which encompasses the experimentally accessible regime, perturbative techniques become unreliable. Nonperturbative lattice-regularized calculations provide the only known, quantitatively reliable, technique for the determination of thermodynamic properties of QCD matter within this regime.

We shall not review the techniques by which lattice-regularized calculations are implemented. We merely recall that the starting point of lattice-regularized calculations at nonzero temperature is the imaginary time formalism, which allows one to write the QCD partition function in Euclidean space-time with a periodic imaginary time direction of length $1/T$. Any thermodynamic quantity can be obtained via suitable differentiation of the partition function. At zero baryon chemical potential, the QCD partition function is given by the exponent of a real action, integrated over all field configurations in the Euclidean space-time. Since the action is real, the QCD partition function can then be evaluated using standard Monte Carlo techniques, which require the discretization of the field configurations and the evaluation of the action on a finite lattice of space-time points. Physical results are obtained by extrapolating calculated results to the limit of infinite volume and vanishing lattice spacing. In principle, this is a quantitatively reliable approach. In practice, lattice-regularized calculations are CPU-expensive: the size of lattices in modern calculations does not exceed $48^3 \times 64$ [195] and these calculations nevertheless require the most powerful computing devices (currently at the multi-teraflop scale). In the continuum limit, such lattices correspond typically to small volumes of $\approx (4 \text{ fm})^3$ [195]. This means that properties of QCD matter which are dominated by long-wavelength modes are difficult to calculate with the currently available computing resources, and there are only first exploratory studies. For the same reason, it is in practice difficult to carry out calculations using light quark masses that yield realistically light pion masses. Light quarks are also challenging because of the CPU-expensive complications which arise from the formulation of fermions on the lattice.

In addition to the practical challenges above, conceptual questions arise in two important domains. First, at nonzero baryon chemical potential, the Euclidean action is no longer real, meaning that the so-called fermion sign problem precludes the use of standard Monte Carlo

techniques. Techniques have been found that evade this problem, but only in the regime where the quark chemical potential $\mu_B/3$ is sufficiently small compared to T . Second, conceptual questions arise in the calculation of any physical quantities that cannot be written as derivatives of the partition function. Calculating many such quantities that are of considerable interest requires the analytic continuation of lattice results from Euclidean to Minkowski space (see below) which is always under-constrained since the Euclidean calculations can only be done at finitely many values of the Euclidean time. This means that lattice-regularized calculations, at least as currently formulated, are not optimized for calculating transport coefficients and answering questions about, say, far-from-equilibrium dynamics or jet quenching.

We allude to these practical and conceptual difficulties to illustrate why alternative strong coupling techniques, including the use of the AdS/CFT correspondence, are and will remain of great interest for the study of QCD thermodynamics and quark-gluon plasma in heavy ion collisions, even though lattice techniques can be expected to make steady progress in the coming years. In the remainder of this Section, we discuss the current status of lattice calculations of some quantities of interest in QCD at nonzero temperature. We shall begin in Section 3.1 with quantities whose calculation does not run into any of the conceptual difficulties we have mentioned, before turning to those that do.

3.1 The QCD Equation of State from the lattice

The QCD equation of state at zero baryon chemical potential, namely the relation between the pressure and the energy density of hot QCD matter, is an example of a quantity that is well-suited to lattice-regularized calculation since, as a thermodynamic quantity, it can be obtained via suitable differentiations of the Euclidean partition function. And, the phenomenological motivation for determining this quantity from QCD from first principles is great since, as we have seen in Section 2.2, it is the most important microphysical input for hydrodynamic calculations. It illustrates the practical challenges of doing lattice-regularized calculations with light quarks that we have mentioned above that while accurate calculations of the thermodynamics of pure glue QCD ($N_f = 0$) have existed for a long time [196], the extraction of the equation of state of quark-gluon plasma with light quarks with their physical masses from calculations that have been extrapolated to the continuum limit has become possible only recently [65, 85, 197].

The current understanding of QCD thermodynamics at the physical point [85] is summarized in Fig. 3.1. In the left panel, the pressure of QCD matter (in thermal equilibrium, with zero baryon chemical potential) is plotted as a function of its temperature. In order to provide a physically meaningful reference, it is customary to compare this quantity to the Stefan-Boltzmann result

$$P_{SB} = \frac{8\pi^2}{45} \left(1 + \frac{21}{32} N_f \right) T^4, \quad (3.1)$$

for a free gas of noninteracting gluons and massless quarks. This benchmark is indicated by the arrow in the figure. As illustrated by this plot, the number of degrees of freedom rises rapidly above a temperature $T_c \sim 170$ MeV; at higher temperatures, the pressure takes an almost constant value which deviates from that of a noninteracting gas of quarks and gluons by approximately 20%. This deviation is still present at temperatures as high as 1 GeV, and




Figure 3.1: Results from a lattice calculation of QCD thermodynamics with physical quark masses ($N_f = 3$, with appropriate light and strange masses). Left panel: Temperature dependence of the pressure in units of T^4 . Right panel: The trace anomaly ($\epsilon - 3P$) in units of T^4 . Data are for lattices with the same temporal extent, meaning the same temperature, but with varying numbers of points in the Euclidean time direction N_τ . The continuum limit corresponds to taking $N_\tau \rightarrow \infty$. Figures taken from Ref. [85].

convergence to the noninteracting limit is only observed at asymptotically high temperatures ($T > 10^8$ GeV [198]) which are far from the reach of any collider experiment. The right panel shows the trace anomaly, $\epsilon - 3P$, in units of T^4 in the same range of temperatures. $\epsilon - 3P$ is often called the “interaction measure”, but this terminology is quite misleading since both noninteracting quarks and gluons on the one hand and very strongly interacting conformal matter on the other have $\epsilon - 3P = 0$, with ϵ/T^4 and P/T^4 both independent of temperature. Large values of $(\epsilon - 3P)/T^4$ necessarily indicate significant interactions among the constituents of the plasma, but small values of this quantity should in no way be seen as indicating a lack of such interactions. We see in the figure that $(\epsilon - 3P)/T^4$ rises rapidly in the vicinity of T_c . This rapid rise corresponds to the fact that ϵ/T^4 rises more rapidly than $3P/T^4$, approaching roughly 80% of its value in a noninteracting gas of quarks and gluons at a lower temperature, between 200 and 250 MeV. At higher temperatures, as $3P/T^4$ rises toward roughly 80% of its noninteracting value, $(\epsilon - 3P)/T^4$ falls off with increasing temperature and the quark-gluon plasma becomes more and more conformal. Remarkably, after a proper re-scaling of the number of degrees of freedom and T_c , all the features described above remain the same when the number of colors of the gauge group is increased and extrapolated to the large $N_c \rightarrow \infty$ limit [199–201].

The central message for us from these lattice calculations of the QCD equation of state is that at high enough temperatures the thermodynamics of the QCD plasma becomes conformal while deviations from conformality are most severe at and just above T_c . This suggests that the use of conformal theories (in which calculations can be done via gauge/gravity duality as described in much of this review) as vehicles by which to gain insights into real-world quark-gluon plasma may turn out to become (even?) more quantitatively reliable when applied to data from heavy ion collisions at the LHC than when applied to those at RHIC. In this respect, it is also quite encouraging that the charged particle elliptic flow $v_2(p_T)$ measured very recently in heavy ion collisions at $\sqrt{s} = 2.76$ TeV at the LHC [122] is, within error bars, the same as that measured at RHIC. Quantitative hydrodynamic analyses of these data will

come soon, but at a qualitative level they indicate that the quark-gluon plasma produced at the LHC is comparably strongly coupled to that at RHIC.

One of the first questions to answer with a calculation of the equation of state in hand is whether the observed rapid rise in ε/T^4 and P/T^4 corresponds to a phase transition or to a continuous crossover. In QCD without quarks, a first order deconfining phase transition is expected due to the breaking of the Z_N center symmetry. This symmetry is unbroken in the confined phase and broken above T_c by a nonzero expectation value for the Polyakov loop. The expected first order phase transition is indeed seen in lattice calculations [196]. The introduction of quarks introduces a small explicit breaking of the Z_N symmetry even at low temperatures, removing this argument for a first order phase transition. However, in QCD with massless quarks there must be a sharp phase transition (first order with three flavors of massless quarks, second order with two) since chiral symmetry is spontaneously broken at low temperatures and unbroken at high temperatures. *This* argument for the necessity of a transition vanishes for quarks with nonzero masses, which break chiral symmetry explicitly even at high temperatures. So, the question of what happens in QCD with physical quark masses, two light and one strange, cannot be answered by any symmetry argument. Since both the center and chiral symmetries are explicitly broken at all temperatures, it is possible for the transition from a hadron gas to quark-gluon plasma as a function of increasing temperature to occur with no sharp discontinuities. And, in fact, lattice calculations have shown that this is what happens: the dramatic increase in ε/T^4 and P/T^4 occurs continuously [197]. This is shown most reliably via the fact that the peaks in the chiral and Polyakov loop susceptibilities are unchanging as one increases the physical spatial volume V of the lattice on which the calculation is done. If there were a first order phase transition, the heights of the peaks of these susceptibilities should grow $\propto V$ in the large V limit; for a second order phase transition, they should grow proportional to some fractional power of V . But, for a continuous crossover no correlation length diverges at T_c and all physical quantities, including the heights of these susceptibilities, should be independent of V once $V^{1/3}$ is larger than the longest correlation length. This is indeed what is found [197]. The fact that the transition is a continuous crossover means that there is no sharp definition of T_c , and different operational definitions can give different values. However, the analysis performed in [202] indicates that the chiral susceptibility and the Polyakov loop susceptibility peak in the range of $T = 150 - 170$ MeV.

Despite the absence of a phase transition in the mathematical sense, well above T_c QCD matter is deconfined, since the Polyakov loop takes on large nonzero values. In this high temperature regime, the matter that QCD describes is best understood in terms of quarks and gluons. This does not, however, imply that the interactions amongst the plasma constituents is negligible. Indeed, we have already seen in Section 2.2 that in the temperature regime accessible in heavy ion collisions at RHIC, the quark-gluon plasma behaves like a liquid, not at all like a gas of weakly coupled quasiparticles. And, as we will discuss in Section 6, explicit calculations done via the AdS/CFT correspondence show that in the large- N_c limit in gauge theories with gravity duals which are conformal, and whose coupling can therefore be chosen, the thermodynamic quantities change by only 25% when the coupling is varied from zero (noninteracting gas) to infinite (arbitrarily strongly coupled liquid). This shows that thermodynamic quantities are rather insensitive to the strength of the interactions among the constituents (or volume elements) of quark-gluon plasma.

Finally, we note that calculations of QCD thermodynamics done via perturbative methods have been compared to the results obtained from lattice-regularized calculations. As is well known (see for example [203] and references therein), the expansion of the pressure in powers of the coupling constant g is a badly convergent series and, what is more, cannot be extended beyond order $g^6 \log(1/g)$, where nonperturbative input is required. This means that perturbative calculations must resort to resummations and indeed different resummation schemes have been developed over the years [204–210]. The effective field theory techniques developed in [204, 205], in particular, exploit a fundamental feature of any perturbative picture of the plasma: at weak coupling, $\mu_D \propto gT$ and these methods all exploit the smallness of μ_D relative to T since the basis of their formulation is that physics at these two energy scales is well separated. As we will see in Section 6.3, this characteristic is in fact essential for any description of the plasma in terms of quasiparticles. The analysis performed in Ref. [206] showed that in the region of $T = 1 - 3T_c$ these effective field theory calculations of the QCD pressure become very sensitive to the matching between the scales μ_D and T , which indicates that there is no separation of these scales. This was foreshadowed much earlier by calculations of various different correlation lengths in the plasma phase which showed that at $T = 2T_c$ some correlation lengths that are $\propto 1/(g^2T)$ at weak coupling are in fact significantly *shorter* than others that are $\propto 1/(gT)$ at weak coupling [211], and showed that the perturbative ordering of these length scales is only achieved for $T > 10^2 T_c$. Despite the success of other resummation techniques [207] in reproducing the main features of QCD thermodynamics, the absence of any separation of scales indicates that there are very significant interactions among constituents and casts doubt upon any approach based upon the existence of well-defined quasiparticles.

3.1.1 Flavor susceptibilities

The previous discussion focussed on thermodynamics in the absence of expectation values for any of the conserved (flavor) charges of QCD. As is well known, these charges are a consequence of the three flavor symmetries that QCD possesses: the U(1) symmetries generated by electric charge, Q , and baryon number, B , and a global SU(3) flavor symmetry. Within SU(3) there are two U(1) subgroups which can be chosen as those generated by Q and by strangeness S . Conservation of Q is fundamental to the standard model, since the U(1) symmetry generated by A is gauge symmetry. Conservation of S is violated explicitly by the weak interactions and conservation of B is violated by exceedingly small nonperturbative weak interactions, and perhaps by yet to be discovered beyond standard model physics. As we are interested only in physics on QCD timescales, we can safely treat S and B as conserved. (Instead of taking B , Q and S as the conserved quantities, we could just as well have chosen the linear combinations of them corresponding to the numbers of up, down and strange quarks.) With three conserved quantities, we can introduce three independent chemical potentials. In spite of the difficulties in studying QCD at nonzero chemical potential on the lattice, derivatives of the pressure with respect to these chemical potentials at zero chemical potential can be calculated. These derivatives describe moments of the distribution of these conserved quantities in an ensemble of volumes of quark-gluon plasma, and hence can be related to event-by-event fluctuations in heavy ion collision experiments.

When all three chemical potentials vanish, the lowest nonzero moments are the quadratic




Figure 3.2: Left: quadratic fluctuations of the baryon number, electric charge and strangeness normalized to their respective Stefan-Boltzmann limit. Right: off-diagonal susceptibility χ_{11}^{BS} normalized to the diagonal baryon charge fluctuations. The upper line corresponds to the Stefan-Boltzmann limits and the lower line corresponds to the hadron resonance gas. Figures taken from [213].

charge fluctuations, i.e. the diagonal and off-diagonal susceptibilities defined as

$$\chi_2^X = \frac{1}{VT} \frac{\partial^2}{\partial \mu_X \partial \mu_X} \log Z(T, \mu_X, \dots) = \frac{1}{VT^3} \langle N_X^2 \rangle, \quad (3.2)$$

$$\chi_{12}^{XY} = \frac{1}{VT} \frac{\partial^2}{\partial \mu_X \partial \mu_Y} \log Z(T, \mu_X, \mu_Y, \dots) = \frac{1}{VT^3} \langle N_X N_Y \rangle, \quad (3.3)$$

where Z is the partition function and the N_X are the numbers of B , Q or S charge present in the volume V . The diagonal susceptibilities quantify the fluctuations of the conserved quantum numbers in the plasma and the off-diagonal susceptibilities measure the correlations among the conserved quantum numbers, and are more sensitive to the nature of the charge carriers [212].

Lattice results for these quantities [213] are shown in Fig. 3.2; unlike the calculation of the pressure described in the previous section, these calculations have not been performed with physical quarks, but with larger quark masses ($m_\pi = 220$ MeV) for which $T_c = 200$ MeV is higher. In the left panel, the diagonal susceptibilities are shown as a function of temperature. The susceptibilities are divided by their values in a noninteracting gas of gluons and massless quarks. In the right panel, one of the off diagonal susceptibilities, χ_{11}^{SB} , is shown; this result has been divided by the diagonal baryon susceptibility to ensure that this quantity approaches unity in the noninteracting limit. In both cases, a rapid rise is observed and above a temperature as low as $T = 1.5 T_c$ they are compatible with a noninteracting gas.

This is a puzzling result. On the one hand, both lattice calculations of $(\varepsilon - 3P)/T^4$ and experimental measurements of elliptic flow at RHIC indicate that significant interactions among the plasma constituents must be present above T_c . On the other hand, the lattice calculations shown in Fig. 3.2 suggest that the flavor degrees of freedom appear to be uncorrelated above about $1.5 T_c$, which would appear to favor a quasiparticle interpretation of the plasma already at these rather low temperatures. This observation further illustrates the fact that thermodynamic properties alone do not resolve the structure of the QGP, since they do not yield an answer to as simple a question as whether it behaves like a liquid or like a gas

of quasiparticles. We will come back to the apparently contradictory pictures suggested by the different static properties of the plasma in Section 6.1.2 and we move on now to further determine the structure of the plasma via studying dynamical quantities rather than just static ones. The lattice calculation of dynamical quantities, which require time and therefore Minkowski spacetime in their formulation, are subject to the conceptual challenges that we described above, meaning that the lattice results that we have described in this section are at present significantly more reliable than those to which we now turn.

3.2 Transport coefficients from the lattice

Transport coefficients, such as the shear viscosity, are essential in the description of the real time dynamics of a system, since they describe how small deviations away from equilibrium relax toward equilibrium. As we have discussed in Section 2.2, the shear viscosity plays a particularly important role as it provides the connection between experimental data on elliptic flow and conclusions about the strongly coupled nature of the quark-gluon plasma produced in RHIC collisions. In this section we describe how transport coefficients can be determined via lattice gauge theory calculations.

Transport coefficients can be extracted from the low momentum and low frequency limits of the Green's functions of a suitable conserved current of the theory, see Appendix A. To illustrate this point, we concentrate on two examples: the stress tensor components T^{xy} , and the longitudinal component of some conserved $U(1)$ current $J^i(\omega, \mathbf{k})$ which can be written $J(\omega, k) \hat{k}$, with ω, k the Fourier modes. The stress tensor correlator determines the shear viscosity; the current-current correlator determines the diffusion constant for the conserved charge associated with the current. (The conserved charge could be baryon number, strangeness or electric charge in QCD or could be some R -charge in a supersymmetric theory.) The retarded correlators of these operators are defined by

$$G_R^{xy\ xy}(t, x) = -i\theta(t) \langle [T^{xy}(t, x) T^{xy}(0, 0)] \rangle, \quad (3.4)$$

$$G_R^{JJ}(t, x) = -i\theta(t) \langle [J(t, x) J(0, 0)] \rangle. \quad (3.5)$$

And, according to the Green-Kubo relation (A.9) the low momentum and low frequency limits of these correlators yield

$$\eta = - \lim_{\omega \rightarrow 0} \frac{\text{Im } G_R^{xy\ xy}(\omega, k=0)}{\omega}, \quad (3.6)$$

$$D\chi = - \lim_{\omega \rightarrow 0} \frac{\text{Im } G_R^{JJ}(\omega, k=0)}{\omega}, \quad (3.7)$$

where η is the shear viscosity, D is the diffusion constant of the conserved charge, and χ is the charge susceptibility. Note that χ is a thermodynamic quantity which can be extracted from the partition function by suitable differentiation and so is straightforward to calculate on the lattice, while η and D are transport properties which describe small deviations from equilibrium. In general, for any conserved current operator \mathcal{O} whose retarded correlator is given by

$$G_R(t, x) = -i\theta(t) \langle [\mathcal{O}(t, x) \mathcal{O}(0, 0)] \rangle, \quad (3.8)$$

if we define a quantity μ by

$$\mu = - \lim_{\omega \rightarrow 0} \frac{\text{Im } G_R(\omega, k = 0)}{\omega}, \quad (3.9)$$

then μ is a transport coefficient, possibly multiplied by a thermodynamic quantity.

Transport coefficients can be computed in perturbation theory. However, since the quark-gluon plasma not too far above T_c is strongly coupled, it is preferable to extract information about the values of the transport coefficients from lattice calculations. Doing so is, however, quite challenging. The difficulty arises from the fact that lattice quantum field theory is formulated in such a way that real time correlators cannot be calculated directly. Instead, these calculations determine the thermal or Euclidean correlator

$$G_E(\tau, x) = \langle \mathcal{O}_E(\tau, x) \mathcal{O}_E(0, 0) \rangle \quad (3.10)$$

where the Euclidean operator is defined from its Minkowski counterpart by

$$\mathcal{O}_M^{\mu_1 \dots \mu_n}{}_{\nu_1 \dots \nu_m}(-i\tau, x) = (-i)^r (i)^s \mathcal{O}_E^{\mu_1 \dots \mu_n}{}_{\nu_1 \dots \nu_m}(\tau, x) \quad (3.11)$$

where r and s are the number of time indices in $\{\mu_1 \dots \mu_n\}$ and $\{\nu_1 \dots \nu_m\}$ respectively. Using the Kubo-Martin-Schwinger relation

$$\langle \mathcal{O}(t, x) \mathcal{O}(0, 0) \rangle = \langle \mathcal{O}(0, 0) \mathcal{O}(t - i\beta, x) \rangle, \quad (3.12)$$

the Euclidean correlator G_E can be related to the imaginary part of the retarded correlator,

$$\rho(\omega, k) \equiv -2 \text{Im } G_R(\omega, k), \quad (3.13)$$

which is referred to as the spectral density. The relation between G_E (which can be calculated on the lattice) and ρ (which determines the transport coefficient) takes the form of a convolution with a known kernel:

$$G_E(\tau, k) = (-1)^{r+s} \int_0^\infty \frac{d\omega}{2\pi} \frac{\cosh\left(\omega\left(\tau - \frac{1}{2T}\right)\right)}{\sinh\left(\frac{\omega}{2T}\right)} \rho(\omega, k). \quad (3.14)$$

A typical lattice computation provides values (with errors) for the Euclidean correlator at a set of values of the Euclidean time, namely $\{\tau_i, G_E(\tau_i, k)\}$. In general, it is not possible to extract a continuous function $\rho(\omega)$ from a limited number of points on $G_E(\tau)$ without making assumptions about the functional form of either the spectral function or the Euclidean correlator. Note also that the Euclidean correlator at any one value of τ receives contributions from the spectral function at all frequencies. This makes it hard to disentangle the low frequency behavior of the spectral function from a measurement of the Euclidean correlator at a limited number of values of τ .

The extraction of the transport coefficient is also complicated by the fact that the high frequency part of the spectral function ρ typically makes a large contribution to the measured G_E . At large ω , the spectral function is the same at nonzero temperature as at zero temperature and is given by

$$\rho(\omega, k = 0) = A \omega^{2\Delta-d}, \quad (3.15)$$




Figure 3.3: Left panel: Ratio of the stress tensor Euclidean correlator calculated on the lattice in Ref. [214] to that in the free theory for QCD with three colors and zero flavors at four values of the Euclidean time $x_0 = \tau$ and two temperatures T . This theory has a first order deconfinement transition, and T is given in units of the critical temperature T_c for this transition. Right panel: Stress tensor Euclidean correlator for $\mathcal{N} = 4$ SYM from Ref. [215]. The solid line corresponds to infinite coupling and dashed lines to the free theory.

where Δ is the dimension of the operator \mathcal{O} and d is the dimension of spacetime. In QCD, the constant A can be computed in perturbation theory. For the two examples that we introduced explicitly above, the spectral functions are given at $k = 0$ to leading order in perturbation theory by

$$\rho_R^{JJ}(\omega, k = 0) = \frac{N_c}{6\pi} \omega^2, \quad (3.16)$$

$$\rho_R^{xy,xy}(\omega, k = 0) = \frac{\pi(N_c^2 - 1)}{5(4\pi)^2} \omega^4, \quad (3.17)$$

where N_c is the number of colors. These results are valid at any ω to leading order in perturbation theory; because QCD is asymptotically free, they are the dominant contribution at large ω . This asymptotic domain of the spectral function does not contain any information about the transport coefficients, but it makes a large contribution to the Euclidean correlator. This means that the extraction of the contribution of the transport coefficient, which is small in comparison and τ -independent, requires very precise lattice calculations.

The results of lattice computations for the shear correlator are shown in the left panel of Fig. 3.3. The finite temperature Euclidean correlator is normalized to the free theory correlator at the same temperature. The measured correlator deviates from the free one only by about 10 – 20%. The errors of the numerical computation illustrate that it is hard to distinguish the computed correlator from the free one, specially at the higher temperature. It is important to stress that the fact that the measured correlator is close to the free one comes from the fact that both receive a large contribution from the large ω region of the spectral function, and therefore cannot be interpreted as a signature of large viscosity. To

illustrate this point, it is illuminating to study $\mathcal{N} = 4$ SYM as a concrete example in which we can compare weak and strong coupling behavior with both determined analytically. As we will discuss in Section 6.2, the AdS/CFT correspondence allows us to compute ρ in the limit of infinite coupling, where the viscosity is small. From this AdS/CFT result, we can then compute the Euclidean correlator via eq. (3.14). The result is shown in the right panel of Fig. 3.3. In the same figure, we show the Euclidean correlator at zero coupling — noting that in the zero coupling limit the viscosity diverges as does the length scale above which hydrodynamics is valid. As in the lattice computation in the left panel of the figure, the difference between the weak coupling and strong coupling Euclidean correlator is small and is only significant around $\tau = 1/(2T)$, where G_E is smallest and the contributions from the small- ω region of ρ are most visible against the “background” from the large- ω region of ρ . For this correlator in this theory, the difference between the infinite coupling and zero coupling limits is only at most 10%. Thus, the $\mathcal{N} = 4$ SYM calculation gives us the perspective to realize that the small deviation between the lattice and free correlators in QCD must not be taken as an indication that the QGP at these temperatures behaves as a free gas. It merely reflect the lack of sensitivity of the Euclidean correlator to the low frequency part of the spectral function.

The extraction of transport information from the 4 points in the left panel of Fig. 3.3, as done in Ref. [214], requires assumptions about the spectral density. Since the high frequency behavior of the spectral function is fixed due to asymptotic freedom, a first attempt can be made by writing

$$\frac{\rho(\omega)}{\omega} = \frac{\rho_{LF}(\omega)}{\omega} + \theta(\omega - \Lambda) \frac{\rho_{HF}(\omega)}{\omega}, \quad (3.18)$$

where

$$\rho_{HF}(\omega) = \frac{\pi(N^2 - 1)}{5(4\pi)^2} \frac{\omega^4}{\tanh \omega/4T} \quad (3.19)$$

is the free theory result at the high frequencies where this result is valid. In the analysis performed in [214], the parameter Λ is always chosen to be $\geq 5T$. The functional form of the low frequency part ρ_{LF} should be chosen such that ρ_{LF} vanishes at high frequency. A Breit-Wigner ansatz

$$\rho_{LF}/\omega = \frac{\eta}{\pi(1 + b^2\omega^2)} = \rho_{BW}/\omega \quad (3.20)$$

provides a simple example with which to start (and is in fact the form that arises in perturbation theory [216].) This ansatz does not provide a good fit, but it nevertheless yields an important lesson. Fitting the parameters in this ansatz to the lattice results for G_E at four values of τ favours large values (larger than T) for the width $\Gamma = 2/b$ of the low frequency Breit-Wigner structure. This result provides quantitative motivation for the assumption that the width of any peak or other structure at low frequency must be larger than T . From this assumption, a bound may be derived on the viscosity as follows. Since a wider function than a Breit-Wigner peak of width $\Gamma = T$ would lead to larger value of ρ_{LF} for $\omega < \sqrt{2}T$ and since the spectral function is positive definite, we have

$$G_E \left(\frac{1}{2T}, k = 0 \right) \geq \frac{1}{T^5} \left[\int_0^{2T} \rho_{BW}(\omega) + \int_\Lambda^\infty \rho_{HF}(\omega) \right] \frac{d\omega}{\sinh \omega/2T}. \quad (3.21)$$

From this condition and the measured value of $G_E(\frac{1}{2T}, k = 0)$, an upper bound on the shear viscosity η can be obtained, resulting in

$$\eta/s < \begin{cases} 0.96 & (T = 1.65T_c) \\ 1.08 & (T = 1.24T_c), \end{cases} \quad (3.22)$$

with s the entropy density [214]. The idea here is: (i) we know how much the $\omega > \Lambda$ region contributes to the integral $\int d\omega \rho(\omega)/\sinh \omega/2T$ which is what the lattice calculation determines, and (ii) we make the motivated assumption that the narrowest a peak at $\omega = 0$ can be is T , and (iii) we can therefore put an upper bound on $\rho(0)$ by assuming that the entire contribution to the integral that does not come from $\omega > \Lambda$ comes from a peak at $\omega = 0$ with width T . The bound is conservative because it comes from assuming that ρ is zero at intermediate ω 's between T and Λ . Surely ρ receives some contribution from this intermediate range of ω , meaning that the bounds on η/s obtained from this analysis are conservative.

Going beyond the conservative bound (3.22) and making an estimate of η is challenging, given the finite number of points at which $G_E(\tau)$ is measured, and relies on physically motivated parameterizations of the spectral function. A sophisticated parameterization was introduced in Ref. [214] under the basic assumption that there are no narrow structures in the spectral function, which is supported by the Breit-Wigner analysis discussed above. In Ref. [214], the spectral function was expanded in an ordered basis of orthonormal functions with an increasing number of nodes, defined and ordered such that the first few functions are those that make the largest contribution to the Euclidean correlator; in other words the latter is most sensitive to the contribution of these functions. Due to the finite number of data points and their finite accuracy, the basis has to be truncated to the first few functions, which is a way of formalizing the assumption that there are no narrow structures in the spectral function. The analysis based on such parameterization leads to small values of the ratio of the shear viscosity to the entropy density. In particular,

$$\eta/s = \begin{cases} 0.134(33) & (T = 1.65T_c) \\ 0.102(56) & (T = 1.24T_c). \end{cases} \quad (3.23)$$

The errors include both statistical errors and an estimate of those systematic errors due to the truncation of the basis of functions used in the extraction. The results of this study are compelling since, as discussed in Section 2.2, they are consistent with the experimentally extracted bounds on the shear viscosity of the QGP via hydrodynamical fits to data on elliptic flow in heavy ion collisions. These results are also remarkably close to $\eta/s = 1/4\pi \approx 0.08$, which is obtained in the infinite coupling limit of $\mathcal{N} = 4$ SYM theory and which we will discuss extensively in Section 6.2.

The lattice studies to date must be taken as exploratory, given the various difficulties that we have described. As explained in Ref. [217], there are ways to do better (in addition to using finer lattices and thus obtaining G_E at more values of τ). For example, a significant improvement may be achieved by analyzing the spectral function at varying nonzero values of the momentum k . One can then exploit energy and momentum conservation to relate different Euclidean correlators to the same spectral function, in some cases constraining the same spectral function with 50-100 quantities calculated on the lattice rather than just 4.

Furthermore, the functional form of the spectral function is predicted order by order in the hydrodynamic expansion and this provides guidance in interpreting the Euclidean data. These analyses are still in progress, but results reported to date [217] are consistent with (3.23), given the error estimate therein.

Let us conclude the discussion by remarking on the main points. The Euclidean correlators calculated on the lattice are dominated by the contribution of the temperature-independent high frequency part of the spectral function, reducing their sensitivity to the transport properties that we wish to extract. This fact, together with the finite number of points on the Euclidean correlators that are available from lattice computations, complicates the extraction of the shear viscosity from the lattice. Under the mild assumption that there are no narrow structures in the spectral function, an assumption that is supported by the lattice data themselves as we discussed, current lattice computations yield a conservative upper bound $\eta/s < 1$ on the shear viscosity of the QGP at $T = (1.2 - 1.7)T_c$. A compelling but exploratory analysis of the lattice data has also been performed, yielding values of $\eta/s \approx 0.1$ for this range of temperatures. In order to determine η/s with quantitative control over all systematic errors, however, further investigation is needed — integrating information obtained from many Euclidean correlators at nonzero k as well as pushing to finer lattices.

3.3 Quarkonium Spectrum from the Lattice

Above the critical temperature, quarks and gluons are not confined. As we have discussed at length in Section 2, experiments at RHIC have taught us that in this regime QCD describes a quark-gluon plasma in which the interactions among the quarks and gluons are strong enough to yield a strongly coupled liquid. It is also possible that these interactions can result in the formation of bound states within the deconfined fluid [218]. This observation is of particular relevance for quarkonium mesons formed from heavy quarks in the plasma, namely quarks with $M \gg T$. For these quarks, $\alpha_s(M)$ is small and the zero temperature mesons are, to a first approximation, described by a Coulomb-like potential between a $Q - \bar{Q}$ pair. Thus, the typical radius of the quarkonium meson is $r_M \sim 1/\alpha_s M \ll 1/T$. As a consequence, the properties of these quarkonium mesons cannot be strongly modified in the plasma. Quarkonia are therefore expected to survive as bound states up to a temperature that is high enough that the screening length of the plasma has decreased to the point that it is of order the quarkonium radius [184].

The actual masses of the heavy quarks that can be accessed in heavy ion collisions, the charm and the bottom, are large enough that charmonium and bottomonium mesons are expected above the deconfinement transition, but they are not so large that these mesons are expected to be unmodified by the quark-gluon plasma produced in ultrarelativistic heavy ion collisions. Heavy ion collisions at RHIC (at the LHC) may reach temperatures high enough to dissociate all the charmonium (bottomonium) states, and charmonia are certainly not expected to survive in the quark-gluon plasma produced at the LHC. It is a non-trivial challenge to determine what QCD predicts for the temperatures up to which a particular quarkonium meson survives as a bound state, and above which it dissociates. In this Section, we review lattice QCD calculations done with this goal in mind. This is a subject of on-going research, and definitive results for the dissociation temperatures of various quarkonia are not




Figure 3.4: Wilson line representing the propagation of a heavy quark-antiquark pair. The line at $-r/2$ is the heavy quark propagator in imaginary time while the line at $r/2$ is the antiquark. The space links ensure gauge invariance. The singlet free energy is obtained by setting $\tau = \beta$.

yet in hand. For an example of a recent review on this subject, see Ref. [219].

Some of the earliest [184, 220] attempts to describe the in-medium heavy mesons are based on solving the Schrödinger equation for a pair of heavy quarks in a potential determined from a lattice calculation. These approaches are known generically as potential models. In this approach, it is assumed that the interactions between the quark anti-quark pairs and the medium can be expressed in the form of a temperature dependent potential. The mesons are identified as the bound states of quarks in this potential. Such an approach has been very successful at zero temperature [221] and in this context it can be put in firm theoretical grounds by means of a non-relativistic effective theory for QCD [222, 223]. However, at nonzero temperature it is not clear how to determine this potential from first principles. (For some attempts in this direction, see Ref. [224].)

If the binding energy of the quarkonium meson is small compared to the temperature and to any other energy scale that characterizes the medium, the potential can be extracted by analyzing a static (infinitely massive) $Q - \bar{Q}$ pair, in the color singlet representation, separated by a distance r . In this limit, both the quark and the antiquark remain static on the time-scale over which the medium fluctuates, and their propagators in the medium reduce to Wilson lines along the time axis. In the imaginary time formalism, these two Wilson lines wind around the periodic imaginary time direction and they are separated in space by the distance r . These quark and anti-quark Wilson lines are connected by spatial links to ensure gauge invariance. These spatial links can be thought of as arising via applying a point-splitting procedure at the point where the quark and antiquark pair are produced by a local color singlet operator. A sketch of this Wilson line is shown in Fig. 3.4.

At zero temperature, the extension of the Wilson lines in the imaginary time direction τ can be taken to infinity; this limit yields Wilson's definition of the heavy quark potential [225]. In contrast, at nonzero temperature the imaginary time direction is compact and the imaginary time τ is bounded by $1/T$. Nevertheless, inspired by the zero temperature case the early




Figure 3.5: Lattice results for the singlet free energy $F_1(\mathbf{r}, T)$ as a function of the distance r for different temperatures T , quoted as fractions of the critical temperature T_c at which the crossover from hadron gas to quark-gluon plasma occurs. The left panel shows results for QCD without quarks [228–230] and the right panel for 2+1 flavor QCD [231]. The fact that below T_c the free energy goes above the zero temperature result is a lattice artifact [224].

studies *postulated* that the potential should be obtained from the Wilson line with $\tau = \beta = 1/T$. This Wilson line can be interpreted as the singlet free energy of the heavy quark pair, i. e. the energy change in the plasma due to the presence of a pair of quarks at a fixed distance and at fixed temperature [226, 227].

Lattice results for the singlet free energy are shown in Fig. 3.5. In the left panel we show results for the gluon plasma described by QCD without any quarks [228–230]. The solid line in this figure denotes the $T = 0$ result, which rises linearly with the separation r at large r , as expected due to confinement. The potential is well approximated by the ansatz

$$F_1(r) = \sigma r - \frac{\alpha}{r}, \quad (3.24)$$

where the linear long-distance part is characterized by the string tension $\sqrt{\sigma} = 420 \text{ MeV}$ [232] and the perturbative $1/r$ piece describes the short-distance regime. Below T_c , as the temperature increases the theory remains confined but the string tension decreases. For temperatures larger than T_c , the theory is not confined and the free energy flattens to a finite value in the large- r limit. At these temperatures, the color charge in the plasma screens the interaction between the heavy Q and \bar{Q} . In QCD with light dynamical quarks, as in Fig. 3.5 b) from Ref. [231], the situation is more complicated. In this case, the free energy flattens to a finite limit at large distance even at zero temperature, since once the heavy quark and antiquark have been pulled far enough apart it becomes favorable to produce a light $q - \bar{q}$ pair from the medium (in this case the vacuum) which results in the formation of $Q\bar{q}$ and $\bar{Q}q$ mesons that can then be moved far apart without any further expenditure of energy. In vacuum this process is usually referred to as “string-breaking”. In vacuum, at distances that are small enough that string-breaking does not occur the potential can be approximated by (3.24), but with a reduced string tension $\sqrt{\sigma} \approx 200 \text{ MeV}$ [84]. Above T_c , the potential is screened at large distances by the presence of the colored fluid, with the screening length beyond which the potential flattens shrinking with increasing temperature, just as in the absence of quarks.

As a consequence, in the right panel of the figure the potential evolves relatively smoothly with increasing temperature, with string-breaking below T_c becoming screening at shorter and shorter distance scales above T_c . The decrease in the screening length with increasing temperature is a generic result, and it leads us to expect that quarkonium mesons are expected to dissociate when the temperature is high enough that the vacuum quarkonium size corresponds to a quark-antiquark separation at which the potential between the quark and antiquark is screened [184].

After precise lattice data for the singlet free energy became available, several authors have used them to solve the Schrödinger equation. Since the expectation value of the Wilson loop in Fig. 3.4 leads to the singlet *free* energy and not to the singlet *internal* energy, it has been argued that the potential to be used in the Schrödinger equation should be that obtained after first subtracting the entropy contribution to the free energy, namely

$$U(r, T) = F(r, T) - T \frac{dF(r, T)}{dT}. \quad (3.25)$$

Analyses performed with this potential indicate that the J/ψ meson survives deconfinement, existing as a bound state up to a dissociation temperature that lies in the range $T_{\text{diss}} \sim (1.5 - 2.5) T_c$ [233–237]. It is also a generic feature of these potential models that, since they are larger in size, other less bound charmonium states like the χ_c and ψ' dissociate at a lower temperature [220], typically at temperatures as low as $T = 1.1 T_c$. Let us state once more that these calculations are based on two key model assumptions: first, that the charm and bottom quarks are heavy enough for a potential model to apply and, second, that the potential is given by eq. (3.25). Neither assumption has been demonstrated from first principles

Given that potential models *are* models, there has also been a lot of effort to extract model-independent information about the properties of quarkonium mesons in the medium at nonzero temperature by using lattice techniques to calculate the Euclidean correlation function of a color singlet operator of the type

$$J_\Gamma(\tau, \mathbf{x}) = \bar{\psi}(\tau, \mathbf{x}) \Gamma \psi(\tau, \mathbf{x}), \quad (3.26)$$

where $\psi(\tau, \mathbf{x})$ is the heavy quark operator and $\Gamma = 1, \gamma_\mu, \gamma_5, \gamma_5 \gamma_\mu, \gamma_\mu \gamma_\nu$ correspond to the scalar, vector, pseudo-scalar, pseudo-vector and tensor channels. As in the case of the transport coefficients whose analysis we described in Section 3.2, in order to obtain information about the in-medium mesons we are interested in extracting the spectral functions of these operators. As in Section 3.2, the Minkowski-space spectral function cannot be calculated directly on the lattice; it must instead be inferred from lattice calculations of the Euclidean correlator

$$G_E(\tau, \mathbf{x}) = \langle J_\Gamma(\tau, \mathbf{x}) J_\Gamma(0, \mathbf{0}) \rangle. \quad (3.27)$$

which is related to the spectral function as in eq. (3.14).

The current-current correlator can be understood as describing the interaction of an external vector meson which couples only to heavy quarks with the plasma. This interaction can proceed by scattering with the heavy quarks and antiquarks present in the plasma or by mixing between the singlet quarkonium meson with (light quark) states from within the plasma that have the same quantum numbers as the external quarkonium. The first physical process leads to a large absorption of those vector mesons in which the ratio ω/q matches




Figure 3.6: Schematic view of the current-current spectral function as a function of frequency for heavy quarks. The structure at small frequency $\omega \sim 0$ is the transport peak which is due to the interaction of the external current with heavy quarks and antiquarks from the plasma. At $\omega = 2M$ there is a threshold for pair production. An in-medium bound state, like a quarkonium meson, appears as a peak below the threshold.

the velocity of heavy quarks in the medium, yielding the so-called transport peak at small ω . The second physical process populates the near-threshold region of $\omega \sim 2M$. Since the thermal distribution of the velocity of heavy quarks and anti-quarks is Maxwellian with a mean velocity $v \sim \sqrt{T/M}$, the transport peak is well-separated from the threshold region. Thus, the spectral function contains information not only about the properties of mesons in the medium, but also about the transport properties of the heavy quarks in the plasma. A sketch of the general expectation for this spectral function is shown in Fig. 3.6. Given these expectations, the extraction of properties of quarkonium mesons in the plasma from the Euclidean correlator must take into account the presence of the transport peak. It is worth mentioning that for the particular case of pseudo-scalar quarkonia, the transport peak is suppressed by mass [238]; thus, the extraction of meson properties is simplest in this channel. All other channels, including in particular the vector channel, include contributions from the transport peak.

From the relation (3.14) between the Euclidean correlator and the spectral function, it is clear that the Euclidean correlator has two sources of temperature dependence: the temperature dependence of the spectral function itself which is of interest to us and the temperature dependence of the kernel in the relation (3.14). Since the latter is a trivial kinematical factor, lattice calculations of the Euclidean correlator are often presented compared to

$$G_{\text{recon}}(\tau, T) = \int_0^\infty d\omega \frac{\cosh(\omega(\tau - 1/2T))}{\sinh(\omega/2T)} \sigma(\omega, T = 0), \quad (3.28)$$

which takes into account the modification of the heat kernel. Any further temperature dependence of G_E relative to that in G_{recon} is due to the temperature dependence of the spectral function.

In Fig. 3.7 we show the ratio of the computed lattice correlator G_E to G_{recon} defined in




Figure 3.7: The ratio of the Euclidean correlator G_E to G_{recon} defined in (3.28) in the pseudo-scalar (left) and vector (right) channels for charm quarks versus the imaginary time τ [239]. Note that the transport contribution is suppressed by the mass of the charmed quark only in the pseudo-scalar channel.

(3.28) in the pseudoscalar and vector channels for charm quarks [239]. The temperature dependence of this ratio is due only to the temperature dependence of the spectral function. The pseudoscalar correlator shows little temperature dependence up to temperatures as high as $T = 1.5 T_c$ while the vector correlator varies significantly in that range of temperatures. Since, as we have already mentioned, the transport peak is suppressed in the pseudoscalar channel, the lack of temperature dependence of the Euclidean correlator in this channel can be interpreted as a signal of the survival of pseudoscalar charmed mesons (the η_c) above deconfinement. However, the Euclidean correlator is a convolution integral over the spectral density and the thermal kernel, and in principle the spectral density can change radically while leaving the convolution integral relatively unchanged. So, the spectral function must be extracted before definitive conclusions can be drawn.

There has been a lot of effort towards extracting these spectral densities in a model-independent way directly from the Euclidean correlators. The method that has been developed the furthest is called the Maximum Entropy Method (MEM) [240]. It is an algorithm designed to find the most probable spectral function compatible with the lattice data on the Euclidean correlator. This problem is underconstrained, since one has available lattice calculations of the Euclidean correlator only at finitely many values of the Euclidean time τ , each with error bars, and one is seeking to determine a function of ω . This means that the algorithm must take advantage of prior information about the spectral function that is known absent any calculations of the value of the Euclidean correlator in the form of a default model for the spectral function. Examples of priors that are taken into account include information about asymptotic behavior and sum rules. The MEM method has been very successful in extracting the spectral functions at zero temperature, where it turns out that the extracted functions have little dependence on the details of how the priors are implemented in a default model for the spectral function. The application of the same MEM procedure at nonzero temperature is complicated by two facts: the number of data points is smaller at finite temperature than at zero temperature and the temporal extent of the correlators is limited to

$1/T$, which is reduced as the temperature increases. The first problem is a computational problem, which can be ameliorated over time as computing power grows by reducing the temporal lattice spacing and thus increasing the number of lattice points within the extent $1/T$. The second problem is intrinsic to the nonzero temperature calculation; all the structure in the Minkowski space spectral function, as in the sketch in Fig. 3.6, gets mapped onto fine details of the Euclidean correlator within a small interval of τ meaning that at nonzero temperature it takes much greater precision in the calculation of the Euclidean correlator in order to disentangle even the main features of the spectral function.

To date, extractions of the pseudoscalar spectral density at nonzero temperature via the MEM indicate, perhaps not surprisingly, that the spectral function including its η_c peak remains almost unchanged up to $T \approx 1.5 T_c$ [239, 241–243], especially when the comparison that is made is with the zero temperature spectral function extracted from only a reduced number of points on the Euclidean correlator. The application of the MEM to the vector channel also indicates survival of the J/ψ up to $T \approx 1.5 - 2 T_c$ [239, 241–244], but it fails to reproduce the transport peak that must be present in this correlator near $\omega = 0$. It has been argued that most of the temperature dependence of the vector correlator seen in Fig. 3.7 is due to the temperature dependence of the transport peak [245]. (Note that since the transport peak is a narrow structure centered at zero frequency, it corresponds to a temperature-dependent contribution to the Euclidean correlator that is approximately τ -independent.) This is supported by the fact that the τ -derivative of the ratio of correlators is much less dependent on T [246] and by the analysis of the spectral functions extracted after introducing a transport peak in the default model of the MEM [244]. When the transport peak is taken into account the MEM also shows that J/ψ may survive at least up to $T = 1.5 T_c$ [244]. However, above T_c both the vector and the pseudoscalar channels show strong dependence (much stronger than at zero temperature) on the default model via which prior information is incorporated in the MEM [239, 244], which makes it difficult to extract solid conclusions on the survival of charmonium states from this method. Despite these uncertainties, the conclusions of the MEM analyses agree with those reached via analyses of potential models in which the internal energy (3.25) is used as the potential. However, before this agreement can be taken as firm evidence for the survival of charmonium states well above the phase transition, it must be shown that the potential models and the lattice calculations are compatible in other respects. To this we now turn.

Potential models can be used for more than determining whether a temperature-dependent potential admits bound states: they provide a prediction for the entire spectral density. It is then straightforward to start with such a predicted spectral density and compute the Euclidean correlator that would be obtained in a lattice calculation if the potential model correctly described all aspects of the physics. One can then compare the Euclidean correlator predicted by the potential model with that obtained in lattice computations like those illustrated in Fig. 3.7. Following this approach, the authors of Refs. [247, 248] have shown that neither the spectral function obtained via identifying the singlet internal energy as the potential nor the one obtained via identifying the singlet free energy as the potential correctly reproduce the Euclidean correlator found in lattice calculations. This means that conclusions drawn based upon either of these potentials cannot be quantitatively reliable in all respects. These authors then proposed a more phenomenological approach, constructing a phenomenological potential (containing many of the qualitative features of the singlet free energy but

differing from it) that reproduces the Euclidean correlator obtained in lattice calculations at the percent level [247, 248]. These authors also point out that at nonzero temperature all putative bound states must have some nonzero thermal width, and states whose binding energy is smaller than this width should not be considered bound. These considerations lead the authors of Refs. [247, 248] to conclude that the J/ψ and η_c dissociate by $T \sim 1.2T_c$ while less bound states like the χ_c or ψ' do not survive the transition at all. These conclusions differ from those obtained via the MEM. Although these conclusions are dependent on the potential used, an important and lasting lesson from this work is that the spectral function above T_c can be very different from that at zero temperature even if the Euclidean correlator computed on the lattice does not show any strong temperature dependence. This lesson highlights the challenge, and the need for precision, in trying to extract the spectral function from lattice calculations of the Euclidean correlator in a model independent fashion.

Finally, we note once again that there is no argument from first principles for using the Schrödinger equation with either the phenomenological potential of Refs. [247, 248] or the internal energy or the free energy as the potential. To conclude this Section, we would like to add some remarks on why the identification of the potential with the singlet internal or singlet free energy cannot be correct [249, 250]. If the quarkonium states can be described by a Schrödinger equation, the current-current correlator must reduce to the propagation of a quark-antiquark pair at a given distance r from each other. The correlator must then satisfy

$$\left(-\partial_\tau + \frac{\nabla^2}{2M} - 2M - V(\tau, \mathbf{r})\right) G_M(\tau, \mathbf{r}) = 0 \quad (3.29)$$

where we have added the subscript M to remind the reader that this expression is only valid in the near threshold region. From this expression, it is clear that the potential can be extracted from the infinitely massive limit, where the propagation of the pair is given by the Wilson line W_M in Fig. 3.4 (up to a trivial phase factor proportional to $2Mt$). In this limit, the potential in the Euclidean equation (3.29) is then defined by

$$-\partial_\tau W_M(\tau, \mathbf{r}) = V(\tau, \mathbf{r})W_M(\tau, \mathbf{r}) , \quad (3.30)$$

where τ and r are the sides of the Wilson loop in Fig. 3.4. In principle, the correct real time potential $V(t, \mathbf{r})$ should then be obtained via analytic continuation of $V(\tau, \mathbf{r})$. And, for bound states with sufficiently low binding energy it would then suffice to consider the long time limit of the potential, $V_\infty(\mathbf{r}) \equiv V(t = \infty, \mathbf{r})$.

The difficulty of extracting the correct potential resides in the analytic continuation from $V(\tau, \mathbf{r})$ to $V(t, \mathbf{r})$. At zero temperature, τ is not periodic and we can take the $\tau \rightarrow \infty$ limit and relate what we obtain to V_∞ . At nonzero temperature, τ is periodic and so there is no $\tau \rightarrow \infty$ limit. It is also apparent that V_∞ need not coincide with the value of $V(\beta, r)$ as postulated in some potential models; in fact, due to the periodicity of τ a lot of information is lost by setting $\tau = \beta$ [250]. Explicit calculations within perturbative thermal field theory, where the analytic continuation can be performed, show that V_∞ does not coincide with the internal energy (3.25) and, what is more, the in-medium potential develops an imaginary part which can be interpreted as the collisional width of the state in the plasma [249]. If the log of the Wilson loop is a quadratic function of the gauge potential, as in QED or in QCD to leading order in perturbation theory, then it is possible to show that the real part of the

potential agrees with the singlet free energy [250]; however this is not the case in general. So, the correct potential to describe shallow bound states is not given by eq. (3.25) and care must be taken in drawing conclusions from potential model calculations.

Chapter 4

Introducing the gauge/string duality

Sections 4 and 5 together constitute a primer on gauge/string duality, written for a QCD audience.

Our goal in this section is to state what we mean by gauge/string duality, via a clear statement of the original example of such a duality [251–253], namely the conjectured equivalence between a certain conformal gauge theory and a certain gravitational theory in Anti de Sitter spacetime. We shall do this in section 4.3. In order to get there, in section 4.1 we will first motivate from a gauge theory perspective why there must be such a duality. Then, in section 4.2, we will give the reader a look at all that one needs to know about string theory itself in order to understand section 4.3, and indeed to read this review.

Since some of the contents in this section are by now standard textbook material, in some cases we will not give specific references. The reader interested in a more detailed review of string theory may consult the many textbooks available such as [254–261]. The reader interested in complementary aspects or extra details about the gauge/string duality may consult some of the many existing reviews, e.g. [1–6, 262–265].

4.1 Motivating the duality

Although the AdS/CFT correspondence was originally discovered [251–253] by studying D-branes and black holes in string theory, the fact that such an equivalence may exist can be directly motivated from certain aspects of gauge theories and gravity.¹ In this subsection we motivate such a direct path from gauge theory to string theory without going into any details about string theory and D-branes.

¹Since string theory is a quantum theory of gravity and the standard Einstein gravity arises as the low-energy limit of string theory, we will use the terms gravity and string theory interchangeably below.

4.1.1 An intuitive picture: Geometrizing the renormalization group flow

Consider a quantum field theory (more generally, a many-body system) in d -dimensional Minkowski spacetime with coordinates (t, \vec{x}) , possibly defined with a short-distance cutoff ϵ . From the work of Kadanoff, Wilson and others in the sixties, a good way to describe the system is to organize the physics in terms of length (or energy) scales, since degrees of freedom at widely separated scales are largely decoupled from each other. If one is interested in properties of the system at a length scale $z \gg \epsilon$, instead of using the bare theory defined at scale ϵ , it is more convenient to integrate out short-distance degrees of freedom and obtain an effective theory at length scale z . Similarly, for physics at an even longer length scale $z' \gg z$, it is more convenient to use the effective theory at scale z' . This procedure defines a renormalization group (RG) flow and gives rise to a continuous family of effective theories in d -dimensional Minkowski spacetime labeled by the RG scale z . One may now visualize this continuous family of d -dimensional theories as a single theory in a $(d + 1)$ -dimensional spacetime with the RG scale z now becoming a spatial coordinate.² By construction, this $(d + 1)$ -dimensional theory should have the following properties:

1. The theory should be intrinsically non-local, since an effective theory at a scale z should only describe physics at scales longer than z . However, there should still be some degree of locality in the z -direction, since degrees of freedom of the original theory at different scales are not strongly correlated with each other. For example, the renormalization group equations governing the evolution of the couplings are local with respect to length scales.
2. The theory should be invariant under reparametrizations of the z -coordinate, since the physics of the original theory is invariant under reparametrizations of the RG scale.
3. All the physics in the region below the Minkowski plane at z (see fig. 4.1) should be describable by the effective theory of the original system defined at a RG scale z . In particular, this $(d + 1)$ -dimensional description has only the number of degrees of freedom of a d -dimensional theory.

In practice, it is not yet clear how to ‘merge’ this continuous family of d -dimensional theories into a coherent description of a $(d + 1)$ -dimensional system, or whether this way of rewriting the renormalization group gives rise to something sensible or useful. Property number (3) above, however, suggests that if such a description is indeed sensible, the result may be a theory of quantum gravity. The clue comes from the holographic principle [267, 268] (for a review, see [269]), which says that a theory of quantum gravity in a region of space should be described by a *non-gravitational* theory living at the boundary of that region. In particular, one may think of the quantum field theory as living on the $z = 0$ slice, the boundary of the entire space.

We now see that the gauge/gravity duality, when interpreted as a geometrization of the RG evolution of a quantum field theory, appears to provide a specific realization of the holographic principle. An important organizing principle which follows from this discussion

²Arguments suggesting that the string dual of a Yang-Mills theory must involve an extra dimension were put forward in [266].




Figure 4.1: A geometric picture of AdS_5 .

is the UV/IR connection [270,271] between the physics of the boundary and the bulk systems. From the viewpoint of the bulk, physics near the $z = 0$ slice corresponds to physics near the boundary of the space, i.e. to large-volume or IR physics. In contrast, from the viewpoint of the quantum field theory, physics at small z corresponds to short-distance physics, i.e. UV physics.

4.1.2 The large- N_c expansion of a non-Abelian gauge theory vs. the string theory expansion

The heuristic picture of the previous section does not tell us for which many-body system such a gravity description is more likely to exist, or what kind of properties such a gravity system should have. A more concrete indication that a many-body theory may indeed have a gravitational description comes from the large- N_c expansion of a non-Abelian gauge theory.

That it ought to be possible to reformulate a non-Abelian gauge theory as a string theory can be motivated at different levels. After all, string theory was first invented to describe strong interactions. Different vibration modes of a string provided an economical way to explain many resonances discovered in the sixties which obey the so-called Regge behavior, i.e. the relation $M^2 \propto J$ between the mass and the angular momentum of a particle. After the formulation of QCD as the microscopic theory for the strong interactions, confinement provided a physical picture for possible stringy degrees of freedom in QCD. Due to confinement, gluons at low energies behave to some extent like flux tubes which can close on themselves or connect a quark-antiquark pair, which naturally suggests a possible string formulation. Such a low-energy effective description, however, does not extend to high energies if the theory becomes weakly coupled, or to non-confining gauge theories.

A strong indication that a fundamental (as opposed to effective) string theory description may exist for any non-Abelian gauge theory (confining or not) comes from 't Hooft's large- N_c

expansion [272]. Due to space limitations, here we will not give a self-contained review of the expansion (see e.g. [273–275] for reviews) and will only summarize the most important features. The basic idea of 't Hooft was to treat the number of colors N_c for a non-Abelian gauge theory as a parameter, take it to be large, and expand physical quantities in $1/N_c$. For example, consider the Euclidean partition function for a $U(N_c)$ pure gauge theory with gauge coupling g :

$$Z = \int DA_\mu \exp\left(-\frac{1}{4g^2} \int d^4x \operatorname{Tr} F^2\right) . \quad (4.1)$$

Introducing the 't Hooft coupling

$$\lambda = g^2 N_c , \quad (4.2)$$

one finds that the vacuum-to-vacuum amplitude $\log Z$ can be expanded in $1/N_c$ as

$$\log Z = \sum_{h=0}^{\infty} N_c^{2-2h} f_h(\lambda) = N_c^2 f_0(\lambda) + f_1(\lambda) + \frac{1}{N_c^2} f_2(\lambda) + \dots , \quad (4.3)$$

where $f_h(\lambda)$, $h = 0, 1, \dots$ are functions of the 't Hooft coupling λ only. What is remarkable about the large- N_c expansion (4.3) is that, at a fixed λ , Feynman diagrams are organized by their topologies. For example, diagrams which can be drawn on a plane without crossing any lines (“planar diagrams”) all have the same N_c dependence, proportional to N_c^2 , and are included in $f_0(\lambda)$. Similarly, $f_h(\lambda)$ includes the contributions of all diagrams which can be drawn on a two-dimensional surface with h holes without crossing any lines. Given that the topology a two-dimensional compact orientable surface is classified by its number of holes, the large- N_c expansion (4.3) can be considered as an expansion in terms of the topology of two-dimensional compact surfaces.

This is in remarkable parallel with the perturbative expansion of a closed string theory, which expresses physical quantities in terms of the propagation of a string in spacetime. The worldsheet of a *closed* string is a two-dimensional compact surface and the string perturbative expansion is given by a sum over the topologies of two-dimensional surfaces. For example the vacuum-to-vacuum amplitude in a string theory can be written as

$$\mathcal{A} = \sum_{h=0}^{\infty} g_s^{2h-2} F_h(\alpha') = \frac{1}{g_s^2} F_0(\alpha') + F_1(\alpha') + g_s^2 F_2(\alpha') + \dots , \quad (4.4)$$

where g_s is the string coupling, $2\pi\alpha'$ is the inverse string tension, and $F_h(\alpha')$ is the contribution of 2d surfaces with h holes.

Comparing (4.3) and (4.4), it is tempting to identify (4.3) as the perturbative expansion of some string theory with

$$g_s \sim \frac{1}{N_c} \quad (4.5)$$

and the string tension given as some function of the 't Hooft coupling λ . Note that the identification of (4.3) and (4.4) is more than just a mathematical analogy. Consider for example $f_0(\lambda)$, which is given by the sum over all Feynman diagrams which can be drawn on a plane (which is topologically a sphere). Each planar Feynman diagram can be thought of as a discrete triangulation of the sphere. Summing all planar diagrams can then be thought

of as summing over all possible discrete triangulations of a sphere, which in turn can be considered as summing over all possible embeddings of a two-dimensional surface with the topology of a sphere in some ambient spacetime. This motivates the conjecture of identifying $f_0(\lambda)$ with F_0 for some closed string theory, but leaves open what the specific string theory is.

One can also include quarks, or more generally matter in the fundamental representation. Since quarks have N_c degrees of freedom, in contrast with the N_c^2 carried by gluons, including quark loops in the Feynman diagrams will lead to $1/N_c$ suppressions. For example, in a theory with N_f flavours, the single-quark loop planar-diagram contribution to the vacuum amplitude scales as $\log Z \sim N_f N_c$ rather than as N_c^2 as in (4.3). In the large- N_c limit with finite N_f , the contribution from quark loops is thus suppressed by powers of N_f/N_c . Feynman diagrams with quark loops can also be classified using topologies of two-dimensional surfaces, now with inclusion of surfaces with boundaries. Each boundary can be identified with a quark loop. On the string side, two-dimensional surfaces with boundaries describe worldsheets of a string theory containing both closed and open strings, with boundaries corresponding to the worldlines of the endpoints of the open strings.

4.1.3 Why AdS?

Assuming that a d -dimensional field theory can be described by a $(d+1)$ -dimensional string or gravity theory, we can try to derive some properties of the $(d+1)$ -dimensional spacetime. The most general metric in $d+1$ dimensions consistent with d -dimensional Poincare symmetry can be written as

$$ds^2 = \Omega^2(z) (-dt^2 + d\vec{x}^2 + dz^2) , \quad (4.6)$$

with z is the extra spatial direction. Note that in order to have translational symmetries in the t, \vec{x} directions, the warp factor $\Omega(z)$ can depend on z only. At this stage not much can be said of the form of $\Omega(z)$ for a general quantum field theory. However if we consider field theories which are conformal (CFTs), then we can determine $\Omega(z)$ using the additional symmetry constraints! A conformally invariant theory is invariant under the scaling

$$(t, \vec{x}) \rightarrow C(t, \vec{x}) \quad (4.7)$$

with C a constant. For the gravity theory formulated in (4.6) to describe such a field theory, the metric (4.6) should respect the scaling symmetry (4.7) with the simultaneous scaling of the z coordinate $z \rightarrow Cz$, since z represents a length scale in the boundary theory. For this to be the case we need $\Omega(z)$ to scale as

$$\Omega(z) \rightarrow C^{-1}\Omega(z) \quad \text{under} \quad z \rightarrow Cz . \quad (4.8)$$

This uniquely determines

$$\Omega(z) = \frac{R}{z} , \quad (4.9)$$

where R is a constant. The metric (4.6) can now be written as

$$ds^2 = \frac{R^2}{z^2} (-dt^2 + d\vec{x}^2 + dz^2) , \quad (4.10)$$

which is precisely the line element of $(d + 1)$ -dimensional anti-de Sitter spacetime, AdS_{d+1} . This is a maximally symmetric spacetime with curvature radius R and constant negative curvature proportional to $1/R^2$. See e.g. [276] for a detailed discussion of the properties of AdS space.

In addition to Poincare symmetry and the scaling (4.7), a conformal field theory in d -dimension is also invariant under d special conformal transformations, which altogether form the d -dimensional conformal group $SO(2, d)$. It turns out that the isometry group³ of (4.10) is also $SO(2, d)$, precisely matching that of the field theory. Thus one expects that a conformal field theory should have a string theory description in AdS spacetime!

Note that it is not possible to use the discussion of this section to deduce the precise string theory dual of a given field theory, nor the precise relations between their parameters. In next section we will give a brief review of some essential aspects of string theory which will then enable us to arrive at a precise formulation of the duality, at least for some gauge theories.

4.2 All you need to know about string theory

Here we will review some basic concepts of string theory and D-branes, which will enable us to establish an equivalence between IIB string theory in $AdS_5 \times S^5$ and $\mathcal{N} = 4$ SYM theory. Although some of the contents of this section are not indispensable to understand some of the subsequent chapters, they are important for building the reader's intuition about the AdS/CFT correspondence.

4.2.1 Strings

Unlike quantum field theory, which describes the dynamics of point particles, string theory is a quantum theory of interacting, relativistic one-dimensional objects. It is characterized by the string tension, T_{str} , and by a dimensionless coupling constant, g_s , that controls the strength of interactions. It is customary to write the string tension in terms of a fundamental length scale ℓ_s , called string length, as

$$T_{\text{str}} \equiv \frac{1}{2\pi\alpha'} \quad \text{with} \quad \alpha' \equiv \ell_s^2 . \quad (4.11)$$

We now describe the conceptual steps involved in the definition of the theory, in a first-quantized formulation, i.e. we consider the dynamics of a single string propagating in a fixed spacetime. Although perhaps less familiar, an analogous first-quantized formulation also exists for point particles [277, 278], whose second-quantized formulation is a quantum field theory. In the case of string theory, the corresponding second-quantized formulation is provided by string field theory, which contains an infinite number of quantum fields, one for each of the vibration modes of a single string. In this review we will not need to consider such a formulation. For the moment we also restrict ourselves to closed strings — we will discuss open strings in the next section.

A string will sweep out a two-dimensional worldsheet which, in the case of a closed string, has no boundary. We postulate that the action that governs the dynamics of the string

³Namely the spacetime coordinate transformations which leave the metric invariant.

is simply the area of this worldsheet. This is a natural generalization of the action for a relativistic particle, which is simply the length of its worldline. In order to write down the string action explicitly, we parametrize the worldsheet with local coordinates σ^α , with $\alpha = 0, 1$. For fixed worldsheet time, $\sigma^0 = \text{const.}$, the coordinate σ^1 parametrizes the length of the string. Let x^M , with $M = 0, \dots, D - 1$, be spacetime coordinates. The trajectory of the string is then described by specifying x^M as a function of σ^α . In terms of these functions, the two-dimensional metric $g_{\alpha\beta}$ induced on the string worldsheet has components

$$g_{\alpha\beta} = \partial_\alpha x^M \partial_\beta x^N g_{MN}, \quad (4.12)$$

where g_{MN} is the spacetime metric. (If x^M are Cartesian coordinates in flat spacetime then $g_{MN} = \eta_{MN} = \text{diag}(- + \dots +)$.) The action of the string is then given by

$$S_{\text{str}} = -T_{\text{str}} \int d^2\sigma \sqrt{-\det g}. \quad (4.13)$$

In order to construct the quantum states of a single string, one needs to quantize this action. It turns out that the quantization imposes strong constraints on the spacetime one started with; not all spacetimes allow a consistent string propagation — see e.g. [254]. For example, if we start with a D -dimensional Minkowski spacetime, then a consistent string theory exists only for $D = 26$. Otherwise the spacetime Lorentz group becomes anomalous at the quantum level and the theory contains negative norm states.

Physically, different states in the spectrum of the two-dimensional worldsheet theory correspond to different vibration modes of the string. From the spacetime viewpoint, each of these modes appears as a particle of a given mass and spin. The spectrum typically contains a finite number of massless modes and an infinite tower of massive modes with masses of order $m_s \equiv \ell_s^{-1}$. A crucial fact about a closed string theory is that one of the massless modes is a particle of spin two, i.e. a graviton. This is the reason why string theory is, in particular, a theory of quantum gravity. The graviton describes small fluctuations of the spacetime metric, implying that the fixed spacetime that we started with is actually dynamical.

One can construct other string theories by adding degrees of freedom to the string worldsheet. The theory that will be of interest here is a supersymmetric theory of strings, the so-called type IIB superstring theory [279, 280], which can be obtained by adding two-dimensional worldsheet fermions to the action (4.13). Although we will of course be interested in eventually breaking supersymmetry in order to obtain a dual description of QCD, it will be important that the underlying theory be supersymmetric, since this will guarantee the stability of our constructions. For a superstring, absence of negative-norm states requires the dimension of spacetime to be $D = 10$. In addition to the graviton, the massless spectrum of IIB superstring theory includes two scalars, a number of antisymmetric tensor fields, and various fermionic partners as required by supersymmetry. One of the scalars, the so-called dilaton Φ , will play an important role here.

Interactions can be introduced geometrically by postulating that two strings can join together and that one string can split into two through a vertex of strength g_s — see Fig. 4.2. Physical observables like scattering amplitudes can be found by summing over string propagations (including all possible splittings and joinings) between initial and final states. After fixing all the gauge symmetries on the string worldsheets, such a sum reduces to a sum




Figure 4.2: ‘Geometrical’ interactions in string theory: two strings in initial states i and j can join into one string in a state k (left) or vice-versa (right).




Figure 4.3: Sum over topologies contributing to the two-to-two amplitude.

over the topologies of two-dimensional surfaces, with contributions from surfaces of h holes weighted by a factor g_s^{2h-2} . This is illustrated in Fig. 4.3 for the two-to-two amplitude.

At low energies $E \ll m_s$, one can integrate out the massive string modes and obtain a low-energy effective theory for the massless modes. Since the massless spectrum of a closed string theory always contains a graviton, to second order in derivatives, the low energy effective action has the form of Einstein gravity coupled to other (massless) matter fields, i.e.

$$S = \frac{1}{16\pi G} \int d^D x \sqrt{-g} \mathcal{R} + \dots, \quad (4.14)$$

where \mathcal{R} is the Ricci scalar for the spacetime with D spacetime dimensions and where the dots stand for additional terms associated with the rest of massless modes. For type IIB superstring theory, the full low-energy effective action at the level of two derivatives is given by the so-called IIB supergravity [281, 282], a supersymmetric generalization of (4.14) (with $D = 10$). The higher-order corrections to (4.14) take the form of a double expansion, in powers of $\alpha' E^2$ from integrating out the massive stringy modes, and in powers of the string coupling g_s from loop corrections.

We conclude this section by making two important comments. First, we note that the ten-dimensional Newton’s constant G in type IIB supergravity can be expressed in terms of the string coupling and the string length as

$$16\pi G = (2\pi)^7 g_s^2 \ell_s^8. \quad (4.15)$$

The dependence on ℓ_s follows from dimensional analysis, since in D dimensions Newton’s constant has dimension $(\text{length})^{D-2}$. The dependence on g_s follows from considering two-to-two string scattering. The leading string theory diagram, depicted in Fig. 4.4(a), is proportional




Figure 4.4: (a) Tree-level contribution, of order g_s^2 , to a two-to-two scattering process in string theory. The low-energy limit of this tree-level diagram must coincide with the corresponding field theory diagram depicted in (b), which is proportional to Newton’s constant, G .

to g_s^2 , since it is obtained by joining together the two diagrams of Fig. 4.2. The corresponding diagram in supergravity is drawn in Fig. 4.4(b), and is proportional to G . The requirement that the two amplitudes yield the same result at energies much lower than the string scale implies $G \propto g_s^2$.

Second, the string coupling constant g_s is not a free parameter, but is given by the expectation value of the dilaton field Φ as $g_s = e^\Phi$. As a result, g_s may actually vary over space and time. Under these circumstances we may still speak of the string coupling constant, e.g. in formulas like (4.15) or (4.19), meaning the asymptotic value of the dilaton at infinity, $g_s = e^{\Phi_\infty}$.

4.2.2 D-branes and gauge theories

Perturbatively, string theory is a theory of strings. Non-perturbatively, the theory also contains a variety of higher-dimensional solitonic objects. D-branes [283] are a particularly important class of solitons. To be definite, let us consider a superstring theory (e.g. type IIA or IIB theory) in a 10-dimensional flat Minkowski spacetime, labeled by time $t \equiv x_0$ and spatial coordinates x_1, \dots, x_9 . A Dp -brane is then a “defect” where closed strings can break and open strings can end that occupies a p -dimensional subspace — see Fig. 4.5, where the x -directions are parallel to the branes and the y -directions are transverse to them. When closed strings break, they become open strings. The end points of the open strings can move freely along the directions of the D-brane, but cannot leave it by moving in the transverse directions. Just like domain walls or cosmic strings in a quantum field theory, D-branes are dynamical objects which can move around. A Dp -brane then sweeps out a $(p+1)$ -dimensional worldvolume in spacetime. D0-branes are particle-like objects, D1-branes are string-like, D2-branes are membrane-like, etc. Stable Dp -branes in Type IIA superstring theory exist for $p = 0, 2, 4, 6, 8$, whereas those in Type IIB have $p = 1, 3, 5, 7$ [283].⁴ This can be seen in a variety of (not unrelated) ways, two of which are: (i) In these cases the corresponding Dp -branes preserve a fraction of the supersymmetry of the underlying theory; (ii) In these cases the corresponding Dp -branes are the lightest states that carry a conserved charge.

Introducing a D-brane adds an entirely new sector to the theory of closed strings, consisting of open strings whose endpoints must satisfy the boundary condition that they lie

⁴D9-branes also exist, but additional consistency conditions must be imposed in their presence. We will not consider them here.




Figure 4.5: Stack of D-branes.

on the D-brane. Recall that in the case of closed strings we started with a fixed spacetime and discovered, after quantization, that the close string spectrum corresponds to dynamical fluctuations of the spacetime. An analogous situation holds for open strings on a D-brane. Suppose we start with a Dp -brane extending in the $x^\mu = (x^0, x^1, \dots, x^p)$ directions, with transverse directions labelled as $y^i = (x^{p+1}, \dots, x^9)$. Then, after quantization, one obtains an open string spectrum which can be identified with fluctuations of the D-brane.

More explicitly, the open string spectrum consists of a finite number of massless modes and an infinite tower of massive modes with masses of order $m_s = 1/\ell_s$. For a single Dp -brane, the massless spectrum consists of an Abelian gauge field $A_\mu(x)$, $\mu = 0, 1, \dots, p$, $9 - p$ scalar fields $\phi^i(x)$, $i = 1, \dots, 9 - p$, and their superpartners. Since these fields are supported on the D-brane, they depend only on the x^μ coordinates along the worldvolume, but not on the transverse coordinates. The $9 - p$ scalar excitations ϕ^i describe fluctuations of the D-brane in the transverse directions y^i , including deformations of the brane's shape and linear motions. They are the exact parallel of familiar collective coordinates for a domain wall or a cosmic string in a quantum field theory, and can be understood as the Goldstone bosons associated to the subset of translational symmetries spontaneously broken by the brane. The presence of a $U(1)$ gauge field $A_\mu(x)$ as part of collective excitations lies at the origin of many fascinating properties of D-branes, which (as we will discuss below) ultimately lead to the gauge/string duality. Although this gauge field is less familiar in the context of quantum field theory solitons (see e.g. [284–286]), it can nevertheless be understood as a Goldstone mode associated to large gauge transformations spontaneously broken by the brane [287–289].

Another striking new feature of D-branes, which has no parallel in field theory, is the




Figure 4.6: Strings stretching between two D-branes.

appearance of a non-Abelian gauge theory when multiple D-branes become close to one another [290]. In addition to the degrees of freedom pertaining to each D-brane, now there are new sectors corresponding to open strings stretched between different branes. For example, consider two parallel branes separated from each other by a distance r , as shown in Fig. 4.6. Now there are four types of open strings, depending on which brane their endpoints lie on. The strings with both endpoints on the same brane give rise, as before, to two massless gauge vectors, which can be denoted by $(A_\mu)^1_1$ and $(A_\mu)^2_2$, where the upper (lower) numeric index labels the brane on which the string starts (ends). Open strings stretching between different branes give rise to two additional vector fields $(A_\mu)^1_2$ and $(A_\mu)^2_1$, which have a mass given by the tension of the string times the distance between the branes, i.e. $m = r/2\pi\alpha'$. These become massless when the branes lie on top of each other, $r = 0$. In this case there are four massless vector fields altogether, $(A_\mu)^a_b$ with $a, b = 1, 2$, which precisely correspond to the gauge fields of a non-Abelian $U(2)$ gauge group. Similarly, one finds that the $9 - p$ massless scalar fields also become 2×2 matrices $(\phi^i)^a_b$, which transform in the adjoint representation of the $U(2)$ gauge group. In the general case of N_c parallel coinciding branes one finds a $U(N_c)$ multiplet of non-Abelian gauge fields with $9 - p$ scalar fields in the adjoint representation of $U(N_c)$. The low-energy dynamics of these modes can be determined by integrating out the massive open string modes, and it turns out to be governed by a non-Abelian gauge theory [290]. To be more specific, let us consider N_c D3-branes in type IIB theory. The massless spectrum consists of a gauge field A_μ , six scalar fields ϕ^i , $i = 1, \dots, 6$ and four Weyl fermions, all of which are in the adjoint representation of $U(N_c)$ and can be written as $N_c \times N_c$ matrices. At the two-derivative level the low-energy effective action for these modes turns out to be precisely [290] the $\mathcal{N} = 4$ super-Yang-Mills theory with gauge group $U(N_c)$ in (3+1) dimensions [291, 292] (for reviews see e.g. [263, 293]), the bosonic part of whose Lagrangian

can be written as

$$\mathcal{L} = -\frac{1}{g^2} \text{Tr} \left(\frac{1}{4} F^{\mu\nu} F_{\mu\nu} + \frac{1}{2} D_\mu \phi^i D^\mu \phi^i + [\phi^i, \phi^j]^2 \right), \quad (4.16)$$

with the Yang-Mills coupling constant given by

$$g^2 = 4\pi g_s. \quad (4.17)$$

Equation (4.16) is in fact the (bosonic part of the) most general renormalizable Lagrangian consistent with $\mathcal{N} = 4$ global supersymmetry. Due to the large number of supersymmetries the theory has many interesting properties, including the fact that the beta function vanishes exactly [294–300] (see section 4.1 of [265] for a one-paragraph proof). Consequently, the coupling constant is scale-independent and the theory is conformally invariant.

Note that the $U(1)$ part of (4.16) is free and can be decoupled. Physically, the reason for this is as follows. Excitations of the overall, diagonal $U(1)$ subgroup of $U(N_c)$ describe motion of the branes' centre of mass, i.e. rigid motion of the entire system of branes as a whole. Because of the overall translation invariance, this mode decouples from the remaining $SU(N_c) \subset U(N_c)$ modes that describe motion of the branes relative to one another. This is the reason why, as we will see, IIB strings in $AdS_5 \times S^5$ are dual to $\mathcal{N} = 4$ super-Yang-Mills theory with gauge group $SU(N_c)$.

The Lagrangian (4.16) receives higher-derivative corrections suppressed by $\alpha' E^2$. The full system also contains closed string modes (e.g. gravitons) which propagate in the bulk of the ten-dimensional spacetime (see Fig. 4.7) and the full theory contains interactions between closed and open strings. The strength of interactions of closed string modes with each other is controlled by Newton's constant G , so the dimensionless coupling constant at an energy E is GE^8 . This vanishes at low energies and so in this limit closed strings become non-interacting, which is essentially the statement that gravity is infrared-free. Interactions between closed and open strings are also controlled by the same parameter, since gravity couples universally to all forms of matter. Therefore at low energies closed strings decouple from open strings. We thus conclude that at low energies the interacting sector reduces to an $SU(N_c)$ $\mathcal{N} = 4$ SYM theory in four dimensions.

Before closing this section we note that, for a single Dp -brane with constant $F_{\mu\nu}$ and $\partial_\mu \phi^i$, all higher-order α' -corrections to (4.16) (or its p -dimensional generalizations) can be resummed exactly into the so-called Dirac-Born-Infeld (DBI) action [301]

$$S_{\text{DBI}} = -T_{\text{Dp}} \int d^{p+1}x e^{-\Phi} \sqrt{-\det(g_{\mu\nu} + 2\pi\ell_s^2 F_{\mu\nu})}, \quad (4.18)$$

where

$$T_{\text{Dp}} = \frac{1}{(2\pi)^p g_s \ell_s^{p+1}} \quad (4.19)$$

is the tension of the brane, namely its mass per unit spatial volume. In this action, Φ is the dilaton and $g_{\mu\nu}$ denotes the induced metric on the brane. In flat space, the latter can be written more explicitly as

$$g_{\mu\nu} = \eta_{\mu\nu} + (2\pi\ell_s^2)^2 \partial_\mu \phi^i \partial_\nu \phi^i. \quad (4.20)$$




Figure 4.7: Open and closed string excitations of the full system.

The first term in (4.20) comes from the flat spacetime metric along the worldvolume directions, and the second term arises from fluctuations in the transverse directions. Expanding the action (4.18) to quadratic order in F and $\partial\phi$ one recovers the Abelian version of eqn. (4.16). The non-Abelian generalization of the DBI action (4.18) is not known in closed form — see e.g. [302] for a review. Corrections to (4.18) beyond the approximation of slowly-varying fields have been considered in [303–306].

4.2.3 D-branes as spacetime geometry

Due to their infinite extent along the x -directions, the total mass of a Dp -brane is infinite. However, the mass per unit p -volume, known as the tension, is finite and is given in terms of fundamental parameters by equation (4.19). The dependence of the tension on the string length is dictated by dimensional analysis. The inverse dependence on the coupling g_s is familiar from solitons in quantum field theory (see e.g. [284–286]) and signals the non-perturbative nature of D-branes, since it implies that they become infinitely massive (even per unit volume) and hence decouple from the spectrum in the perturbative limit $g_s \rightarrow 0$. The crucial difference is that the D-branes' tension scales as $1/g_s$ instead of the $1/g^2$ scaling that is typical of field theory solitons. This dependence can be anticipated based on the divergences of string perturbation theory [307] and, as we will see, is of great importance for the gauge/string duality.

In a theory with gravity, all forms of matter gravitate. D-branes are no exception, and their presence deforms the spacetime metric around them. The spacetime metric sourced by N_c Dp -branes can be found by explicitly solving the supergravity equations of motion [308–310]. For illustration we again use the example of D3-branes in type IIB theory, for which one finds:

$$ds^2 = H^{-1/2} (-dt^2 + dx_1^2 + dx_2^2 + dx_3^2) + H^{1/2} (dr^2 + r^2 d\Omega_5^2) . \quad (4.21)$$

The metric inside the parentheses in the second term is just the flat metric in the y -directions transverse to the D3-branes written in spherical coordinates, with radial coordinate $r^2 = y_1^2 + \dots + y_6^2$. The function $H(r)$ is given by

$$H = 1 + \frac{R^4}{r^4} , \quad (4.22)$$

where

$$R^4 = 4\pi g_s N_c \ell_s^4 . \quad (4.23)$$

Let us gain some physical intuition regarding this solution. Since D3-branes extend along three spatial directions, their gravitational effect is similar to that of a point particle with mass $M \propto N_c T_{D3}$ in the six transverse directions. Thus the metric (4.21) only depends on the radial coordinate r of the transverse directions. For $r \gg R$ we have $H \simeq 1$ and the metric reduces to that of flat space with a small correction proportional to

$$\frac{R^4}{r^4} \sim \frac{N_c g_s \ell_s^4}{r^4} \sim \frac{GM}{r^4} , \quad (4.24)$$

which can be interpreted as the gravitational potential due to a massive object of mass M in




Figure 4.8: Excitations of the system in the closed string description.

six spatial dimensions.⁵ Note that in the last step in eqn. (4.24) we have used the fact that $G \propto g_s^2 \ell_s^8$ and $M \propto N_c T_{D3} \propto N_c / g_s \ell_s^4$ — see (4.15) and (4.19).

The parameter R can thus be considered as the length scale characteristic of the range of the gravitational effects of N_c D3-branes. These effects are weak for $r \gg R$, but become strong for $r \ll R$. In the latter limit, we may neglect the ‘1’ in eqn. (4.23), in which case the metric (4.21) reduces to

$$ds^2 = ds_{AdS_5}^2 + R^2 d\Omega_5^2, \quad (4.25)$$

where

$$ds_{AdS_5}^2 = \frac{r^2}{R^2} (-dt^2 + dx_1^2 + dx_2^2 + dx_3^2) + \frac{R^2}{r^2} dr^2 \quad (4.26)$$

is the metric (4.10) of five-dimensional anti-de Sitter spacetime written in terms of $r = R^2/z$. We thus see that in the strong gravity region the ten-dimensional metric factorizes into $AdS_5 \times S^5$.

We conclude that the geometry sourced by the D3-branes takes the form displayed in Fig. 4.8: far away from the branes the spacetime is flat, ten-dimensional Minkowski space, whereas close to the branes a ‘throat’ geometry of the form $AdS_5 \times S^5$ develops. The size of the throat is set by the length-scale R , given by (4.23). As we will see, the spacetime geometry (4.21) can be considered as providing an *alternative* description of the D3-branes.

⁵Recall that a massive object of mass M in D spatial dimensions generates a gravitational potential GM/r^{D-2} at a distance r from its position.

4.3 The AdS/CFT conjecture

In the last two Sections we have seen two descriptions of D3-branes. In the description of Section 4.2.2, which we will refer to as the open string description, D-branes correspond to a hyperplane in a *flat* spacetime. In this description the D-branes' excitations are open strings living on the branes, and closed strings propagate outside the branes — see Fig. 4.7.

In contrast, in the description of Section 4.2.3, which we will call the closed string description, D-branes correspond to a spacetime geometry in which only *closed* strings propagate, as displayed in fig. 4.8. In this description there are no open strings. In this case the low-energy limit consists of focusing on excitations that have arbitrarily low energy with respect to an observer in the asymptotically flat Minkowski region. We have here two distinct sets of degrees of freedom, those propagating in the Minkowski region and those propagating in the throat — see fig. 4.8. In the Minkowski region the only modes that remain are those of the massless ten-dimensional graviton supermultiplet. Moreover, at low energies these modes decouple from each other, since their interactions are proportional to GE^8 . In the throat region, however, the whole tower of massive string excitations survives. This is because a mode in the throat must climb up a gravitational potential in order to reach the asymptotically flat region. Consequently, a closed string of arbitrarily high proper energy in the throat region may have an arbitrarily low energy as seen by an observer at asymptotic infinity, provided the string is located sufficiently deep down the throat. As we focus on lower and lower energies these modes become supported deeper and deeper in the throat, and so they decouple from those in the asymptotic region. We thus conclude that in the closed string description, the interacting sector of the system at low energies reduces to closed strings in $AdS_5 \times S^5$.

These two representations are tractable in different parameter regimes. For $g_s N_c \ll 1$, we see from eqn. (4.23) that $R \ll \ell_s$, i.e. the radius characterizing the gravitational effect of the D-branes becomes small in string units, and closed strings feel a flat spacetime everywhere except very close to the hyperplane where the D-branes are located. In this regime the closed string description is not useful since one would need to understand sub-string-scale geometry. In the opposite regime, $g_s N \gg 1$, we find that $R \gg \ell_s$ and the geometry becomes weakly curved. In this limit the closed string description simplifies and essentially reduces to classical gravity. Instead, the open string description becomes intractable, since $g_s N_c$ controls the loop expansion of the theory and one would need to deal with strongly coupled open strings. Note that both representations exist in both limiting regimes, and in between.

To summarize, the two descriptions of N_c D3-branes that we have discussed and their low-energy limits are:

1. A hyperplane in a *flat* spacetime with open strings attached. The low-energy limit is described by $\mathcal{N} = 4$ SYM theory (4.16) with a gauge group $SU(N_c)$.
2. A curved spacetime geometry (4.21) where only closed strings propagate. The low-energy limit is described by closed IIB string theory in $AdS_5 \times S^5$.

It is natural to conjecture that these two descriptions are equivalent. Equating in particular their low-energy limits, we are led to conjecture that

$$\left\{ \mathcal{N} = 4 SU(N_c) \text{ SYM theory} \right\} = \left\{ \text{IIB string theory in } AdS_5 \times S^5 \right\}. \quad (4.27)$$

From equations (4.17) and (4.23) we find how the parameters of the two theories are related to one another:

$$g_s = \frac{g^2}{4\pi}, \quad \frac{R}{\ell_s} = (g^2 N_c)^{1/4}. \quad (4.28)$$

One can also use the ten-dimensional Newton constant (4.15) in place of g_s in the first equation above and obtain equivalently

$$\frac{G}{R^8} = \frac{\pi^4}{2N_c^2}, \quad \frac{R}{\ell_s} = (g^2 N_c)^{1/4}. \quad (4.29)$$

Note that, in particular, the first equation in (4.28) implies that the criterion that $g_s N_c$ be large or small translates into the criterion that the gauge theory 't Hooft coupling $\lambda = g^2 N_c$ be large or small. Therefore the question of which representation of the D-branes is tractable becomes the question of whether the gauge theory is strongly or weakly coupled. We will come back to this in section 5.

The discussion above relates string theory to $\mathcal{N} = 4$ SYM theory at zero temperature, as we were considering the ground state of the N_c D3-branes. On the supergravity side this corresponds to the so-called extremal solution. The above discussion can easily be generalized to a nonzero temperature T by exciting the degrees of freedom on the D3-branes to a finite temperature, which corresponds [311, 312] to the so-called non-extremal solution [310]. It turns out that the net effect of this is solely to modify the AdS part of the metric, replacing (4.26) by

$$ds^2 = \frac{r^2}{R^2} (-f dt^2 + dx_1^2 + dx_2^2 + dx_3^2) + \frac{R^2}{r^2 f} dr^2, \quad (4.30)$$

where

$$f(r) = 1 - \frac{r_0^4}{r^4}. \quad (4.31)$$

Equivalently, in terms of the z -coordinate of section 4.1 we replace (4.10) by

$$ds^2 = \frac{R^2}{z^2} \left(-f dt^2 + dx_1^2 + dx_2^2 + dx_3^2 + \frac{dz^2}{f} \right), \quad (4.32)$$

where

$$f(z) = 1 - \frac{z^4}{z_0^4}. \quad (4.33)$$

These two metrics are related by the simple coordinate transformation $z = R^2/r$, and represent a black brane in AdS spacetime with a horizon located at $r = r_0$ or $z = z_0$ which extends in all three spatial directions of the original brane. As we will discuss in more detail in the next section, r_0 and z_0 are related to the temperature of the $\mathcal{N} = 4$ SYM theory as

$$r_0 \propto \frac{1}{z_0} \propto T. \quad (4.34)$$

Thus we conclude that $\mathcal{N} = 4$ SYM theory at finite temperature is described by string theory in an AdS black brane geometry.

To summarize this section, we have arrived at a duality (4.27) of the type anticipated in section 4.1, that is, an equivalence between a conformal field theory with zero β -function and trivial RG-flow and string theory on a scale-invariant metric that looks the same at any z . In the finite-temperature case, Eqn. (4.34) provides an example of the expected relationship between energy scale in the gauge theory, set in this case by the temperature, and position in the fifth dimension, set by the location of the horizon.

Chapter 5

General aspects of the duality

5.1 Gauge/gravity duality

In the last section we outlined the string theory reasoning behind the equivalence (4.27) between $\mathcal{N} = 4$ $SU(N_c)$ SYM theory and type IIB string theory on $AdS_5 \times S^5$. $\mathcal{N} = 4$ SYM theory is the unique maximally supersymmetric gauge theory in $(3 + 1)$ dimensions, whose field content includes a gauge field A_μ , six real scalars $\phi^i, i = 1, \dots, 6$, and four Weyl fermions $\chi_a, a = 1, \dots, 4$, all them in the adjoint representation of the gauge group. The metric of $AdS_5 \times S^5$ is given by

$$ds^2 = ds_{AdS_5}^2 + R^2 d\Omega_5^2, \quad (5.1)$$

with

$$ds_{AdS_5}^2 = \frac{r^2}{R^2} \eta_{\mu\nu} dx^\mu dx^\nu + \frac{R^2}{r^2} dr^2, \quad r \in (0, \infty). \quad (5.2)$$

In the above equation $x^\mu = (t, \vec{x})$, $\eta_{\mu\nu}$ is the Minkowski metric in four spacetime dimensions, and $d\Omega_5^2$ is the metric on a unit five-sphere. The metric (5.2) covers the so-called ‘Poincaré patch’ of a global AdS spacetime, and it is sometimes convenient to rewrite (5.2) using a new radial coordinate $z = R^2/r \in (0, \infty)$, in terms of which we have

$$ds_{AdS_5}^2 = g_{MN} dx^M dx^N = \frac{R^2}{z^2} (\eta_{\mu\nu} dx^\mu dx^\nu + dz^2), \quad x^M = (z, x^\mu), \quad (5.3)$$

as used earlier in (4.10).

In (5.3), each constant- z slice of AdS_5 is isometric to four-dimensional Minkowski spacetime with x^μ identified as the coordinates of the gauge theory (see also fig. 4.1). As $z \rightarrow 0$ we approach the ‘boundary’ of AdS_5 . This is a boundary in the conformal sense of the word but not in the topological sense, since the prefactor R^2/z^2 in (5.3) approaches infinity there. Although this concept can be given a precise mathematical meaning, we will not need these details here. As motivated in Section 4.1.1 it is natural to imagine that the Yang-Mills theory lives at the boundary of AdS_5 . For this reason, below we will often refer to it as the boundary theory. As $z \rightarrow \infty$, we approach the so-called Poincaré horizon, at which the prefactor R^2/z^2 and the determinant of the metric go to zero.

5.1.1 UV/IR connection

Due to the warp factor R^2/z^2 in front of the Minkowski metric in (5.3), energy and length scales along Minkowski directions in AdS_5 are related to those in the gauge theory by a z -dependent rescaling. More explicitly, consider an object with energy E_{YM} and size d_{YM} in the gauge theory. These are the energy and the size of the object measured in units of the coordinates t and \vec{x} . From (5.3) we see that the corresponding proper energy E and proper size d of this object in the bulk are

$$d = \frac{R}{z} d_{\text{YM}}, \quad E = \frac{z}{R} E_{\text{YM}}, \quad (5.4)$$

where the second relation follows from the fact that the energy is conjugate to time, and so it scales with the opposite scale factor than d . We thus see that physical processes in the bulk with identical proper energies but occurring at different radial positions correspond to different gauge theory processes with energies that scale as $E_{\text{YM}} \sim 1/z$. In other words, a gauge theory process with a characteristic energy E_{YM} is associated with a bulk process localized at $z \sim 1/E_{\text{YM}}$ [251, 270, 271]. This relation between the radial direction z in the bulk and the energy scale of the boundary theory makes concrete the heuristic discussion of Section 4.1.1 that led us to identify the z -direction with the direction along the renormalization group flow of the gauge theory. In particular the high-energy (UV) limit $E_{\text{YM}} \rightarrow \infty$ corresponds to $z \rightarrow 0$, i.e. to the near-boundary region, while the low-energy (IR) limit $E_{\text{YM}} \rightarrow 0$ corresponds to $z \rightarrow \infty$, i.e. to the near-horizon region.

In a conformal theory, there exist excitations of arbitrarily low energies. This is reflected in the bulk in the fact that the geometry extends all the way to $z \rightarrow \infty$. As we will see in Section 5.2.2, for a confining theory with a mass gap m , the geometry ends smoothly at a finite value $z_0 \sim 1/m$. Similarly, at a finite temperature T , which provides an effective IR cutoff, the spacetime will be cut off by an event horizon at a finite $z_0 \sim 1/T$ (see Section 5.2.1).

5.1.2 Strong coupling from gravity

$\mathcal{N} = 4$ SYM theory is a scale-invariant theory characterized by two parameters: the Yang-Mills coupling g and the number of color N_c . The theory on the right hand side of (4.27) is a quantum gravity theory in a maximally symmetric spacetime which is characterized by the Newton constant G and the string scale ℓ_s in units of the curvature radius R . The relations between these parameters are given by (4.29). Recalling that $G \sim \ell_p^8$, with ℓ_p the Planck length, these relations imply

$$\frac{\ell_p^8}{R^8} \propto \frac{1}{N_c^2}, \quad \frac{\ell_s^2}{R^2} \propto \frac{1}{\sqrt{\lambda}}, \quad (5.5)$$

where $\lambda = g^2 N_c$ is the 't Hooft coupling and we have only omitted purely numerical factors.

The full IIB string theory on $AdS_5 \times S^5$ is rather complicated and right now a systematic treatment of it is not available. However, as we will explain momentarily, in the limit

$$\frac{\ell_p^8}{R^8} \ll 1, \quad \frac{\ell_s^2}{R^2} \ll 1 \quad (5.6)$$

the theory dramatically simplifies and can be approximated by classical supergravity, which is essentially Einstein's general relativity coupled to various matter fields. An immediate consequence of the relations (5.5) is that the limit (5.6) corresponds to

$$N_c \gg 1, \quad \lambda \gg 1. \quad (5.7)$$

Equation (4.27) then implies that the planar, strongly coupled limit of the SYM theory can be described using just classical supergravity.

Let us return to why string theory simplifies in the limit (5.6). Consider first the requirement $\ell_s^2 \ll R^2$. This can be equivalently rewritten as $m_s^2 \gg \mathcal{R}$ or as $T_{\text{str}} \gg \mathcal{R}$, where $\mathcal{R} \sim 1/R^2$ is the typical curvature scale of the space where the string is propagating. The condition $m_s^2 \gg \mathcal{R}$ means that one can omit the contribution of all the massive states of *microscopic* strings in low-energy processes. In other words, only the massless modes of microscopic strings, i.e. the supergravity modes, need to be kept in this limit. This is tantamount to treating these strings as point-like particles and ignoring their extended nature, as one would expect from the fact that their typical size, ℓ_s , is much smaller than the typical size of the space where they propagate, R . The so-called α' -expansion on the string side (with $\alpha' = \ell_s^2$), which incorporates stringy effects associated with the finite length of the string in a derivative expansion, corresponds on the gauge theory side to an expansion around infinite coupling in powers of $1/\sqrt{\lambda}$.

The extended nature of the string, however, cannot be ignored in all cases. As we will see in the context of the Wilson loop calculations of Section 5.4 and in many other examples in Section 7, the description of certain physical observables requires one to consider long, *macroscopic* strings whose typical size is much larger than R — for example, this happens when the string description of such observables involves non-trivial boundary conditions on the string. In this case the full content of the second condition in (5.6) is easily understood by rewriting it as $T_{\text{str}} \gg \mathcal{R}$. This condition says that the tension of the string is very large compared to the typical curvature scale of the space where it is embedded, and therefore implies that fluctuations around the classical shape of the string can be neglected. These long strings can still break and reconnect, but in between such processes their dynamics is completely determined by the Nambu-Goto equations of motion. In these cases, the α' -expansion (that is, the expansion in powers of $1/\sqrt{\lambda}$) incorporates stringy effects associated with fluctuations of the string that are suppressed at $\lambda \rightarrow \infty$ by the tension of the string becoming infinite in this limit. From this viewpoint, the fact that the massive modes of microscopic strings can be omitted in this limit is just the statement that string fluctuations around a point-like string can be neglected.

Consider now the requirement $\ell_p^8 \ll R^8$. Since the ratio ℓ_p^8/R^8 controls the strength of quantum gravitational fluctuations, in this regime we can ignore quantum fluctuations of the spacetime metric and talk about a fixed spacetime like $AdS_5 \times S^5$. The quantum gravitational corrections can be incorporated in a power series in ℓ_p^8/R^8 , which corresponds to the $1/N_c^2$ expansion in the gauge theory. Note from (4.28) that taking the $N_c \rightarrow \infty$ limit at fixed λ corresponds to taking the string coupling $g_s \rightarrow 0$, meaning that quantum corrections corresponding to loops of string breaking off or reconnecting are suppressed in this limit.

In summary, we conclude that the strong-coupling limit in the gauge theory suppresses the stringy nature of the dual string theory, whereas the large- N_c limit suppresses its quantum

nature. When both limits are taken simultaneously the full string theory reduces to a classical gravity theory with a finite number of fields.

Given that the S^5 factor in (5.1) is compact, it is often convenient to express a 10-dimensional field in terms of a tower of fields in AdS_5 by expanding it in terms of harmonics on S^5 . For example, the expansion of a scalar field $\phi(x, \Omega)$ can be written schematically as

$$\phi(x, \Omega) = \sum_{\ell} \phi_{\ell}(x) Y_{\ell}(\Omega), \quad (5.8)$$

where x and Ω denote coordinates in AdS_5 and S^5 respectively, and $Y_{\ell}(\Omega)$ denote the spherical harmonics on S^5 . Thus, for many purposes (but not all) the original duality (4.27) can also be considered as the equivalence of $\mathcal{N} = 4$ SYM theory (at strong coupling) with a gravity theory in AdS_5 only. This perspective is very useful in two important aspects. First, it makes manifest that the duality (4.27) can be viewed as an explicit realization of the holographic principle mentioned in Section 4.1.1, with the bulk spacetime being AdS_5 and the boundary being four-dimensional Minkowski spacetime. Second, as we will mention in Section 5.2.3, this helps to give a unified treatment of many different examples of the gauge/gravity duality. In most of this review we will adopt this five-dimensional perspective and work only with fields in AdS_5 .

After dimensional reduction on S^5 , the supergravity action can be written as

$$S = \frac{1}{16\pi G_5} \int d^5x [\mathcal{L}_{\text{grav}} + \mathcal{L}_{\text{matt}}], \quad (5.9)$$

where

$$\mathcal{L}_{\text{grav}} = \sqrt{-g} \left(\mathcal{R} + \frac{12}{R^2} \right) \quad (5.10)$$

is the Einstein-Hilbert Lagrangian with a negative cosmological constant $\Lambda = -6/R^2$ and $\mathcal{L}_{\text{matt}}$ is the Lagrangian for matter fields. In the general case, the latter would include the infinite towers $\phi_{\ell}(x)$ coming from the expansion on the S^5 . The metric (5.3) is a maximally-symmetric solution of the equations of motion derived from the action (5.9) with all matter fields set to zero.

The relation between the effective five-dimensional Newton's constant G_5 and its ten-dimensional counterpart G can be read off from the reduction of the Einstein-Hilbert term,

$$\frac{1}{16\pi G} \int d^5x d^5\Omega \sqrt{-g_{10}} \mathcal{R}_{10} = \frac{R^5 \Omega_5}{16\pi G_5} \int d^5x \sqrt{-g_5} \mathcal{R}_5 + \dots, \quad (5.11)$$

where $\Omega_5 = \pi^3$ is the volume of a unit S^5 . This implies

$$G_5 = \frac{G}{\Omega_5} = \frac{G}{\pi^3 R^5}, \quad \text{i.e.} \quad \frac{G_5}{R^3} = \frac{\pi}{2N_c^2}, \quad (5.12)$$

where in the last equation we made use of (4.29).

5.1.3 Symmetries

Let us now examine the symmetries on both sides of the correspondence. The $\mathcal{N} = 4$ SYM theory is invariant not only under dilatations but under $\text{Conf}(1, 3) \times SO(6)$. The first factor is the conformal group of four-dimensional Minkowski space, which contains the Poincaré group, the dilatation symmetry generated by D , and four special conformal transformations whose generators we will denote by K_μ . The second factor is the R-symmetry of the theory under which the ϕ^i in (4.16) transform as a vector. In addition, the theory is invariant under sixteen ordinary or ‘Poincaré’ supersymmetries, the fermionic superpartners of the translation generators P_μ , as well as under sixteen special conformal supersymmetries, the fermionic superpartners of the special conformal symmetry generators K_μ .

The string side of the correspondence is of course invariant under the group of diffeomorphisms, which are gauge transformations. The subgroup of these consisting of large gauge transformations that leave the asymptotic (i.e. near the boundary) form of the metric invariant is precisely $SO(2, 4) \times SO(6)$. The first factor, which is isomorphic to $\text{Conf}(1, 3)$, corresponds to the isometry group of AdS_5 , and the second factor corresponds to the isometry group of S^5 . As usual, large gauge transformations must be thought of as global symmetries, so we see that the bosonic global symmetry groups on both sides of the correspondence agree. In more detail, the Poincaré group of four-dimensional Minkowski spacetime is realized inside $SO(2, 4)$ as transformations that act separately on each of the constant- z slices in (5.3) in an obvious manner. The dilation symmetry of Minkowski spacetime is realized in AdS_5 as the transformation $(t, \vec{x}) \rightarrow \Lambda(t, \vec{x})$, $z \rightarrow \Lambda z$ (with Λ a positive constant), which indeed leaves the metric (5.3) invariant. The four special conformal transformations of Minkowski spacetime are realized in a slightly more involved way as isometries of AdS_5 .

An analogous statement can be made for the fermionic symmetries. $AdS_5 \times S^5$ is a maximally supersymmetric solution of type IIB string theory, and so it possesses thirty-two Killing spinors which generate fermionic isometries. These can be split into two groups that match those of the gauge theory.¹

We therefore conclude that the global symmetries are the same on both sides of the duality. It is important to note, however, that on the gravity side the global symmetries arise as large gauge transformations. In this sense there is a correspondence between global symmetries in the gauge theory and gauge symmetries in the dual string theory. This is an important general feature of all known gauge/gravity dualities, to which we will return below after discussing the field/operator correspondence. It is also consistent with the general belief that the only conserved charges in a theory of quantum gravity are those associated with global symmetries that arise as large gauge transformations.

5.1.4 Matching the spectrum: the field/operator correspondence

We now consider the mapping between the spectra of two theories. To motivate the main idea, we begin by recalling that the SYM coupling constant g^2 is identified (up to a constant) with the string coupling constant g_s . As discussed below (4.15), in string theory this is given by $g_s = e^{\Phi_\infty}$, where Φ_∞ is the value of the dilaton at infinity, in this case at the AdS boundary

¹In both boundary and bulk, bosonic and fermionic symmetries combine together to form a supergroup $SU(2, 2|4)$.

(∂ AdS). This suggests that deforming the gauge theory by changing the value of a coupling constant corresponds to changing the value of a bulk field at ∂ AdS. More generally, one may imagine deforming the gauge theory action as

$$S \rightarrow S + \int d^4x \phi(x) \mathcal{O}(x), \quad (5.13)$$

where $\mathcal{O}(x)$ is a local, gauge-invariant operator and $\phi(x)$ is a possibly point-dependent coupling, namely a source. If $\phi(x)$ is constant, then the deformation above corresponds to simply changing the coupling for the operator $\mathcal{O}(x)$. The example of g suggests that to each possible source $\phi(x)$ for each possible local, gauge-invariant operator $\mathcal{O}(x)$ there must correspond a dual bulk field $\Phi(x, z)$ (and vice-versa) such that its value at the AdS boundary may be identified with the source, namely:

$$\phi(x) = \Phi|_{\partial AdS}(x) = \lim_{z \rightarrow 0} \Phi(x, z). \quad (5.14)$$

As will be discussed around Eqn. (5.26), Eqn. (5.14) is correct for a massless field Φ , but it must be generalized for a massive field.

This one-to-one map between bulk fields in AdS and local, gauge-invariant operators in the gauge theory is known as the field/operator correspondence. The field and the operator must have the same quantum numbers under the global symmetries of the theory, but there is no completely general and systematic recipe to identify the field dual to a given operator. Fortunately, an additional restriction is known for a very important set of operators in any gauge theory: conserved currents associated to global symmetries, such as the $SO(6)$ symmetry in the case of the $\mathcal{N} = 4$ SYM theory. The source $A_\mu(x)$ coupling to a conserved current $J^\mu(x)$ as

$$\int d^4x A_\mu(x) J^\mu(x) \quad (5.15)$$

may be thought of as an external background gauge field, and from (5.14) we can view it as the boundary value of a dynamical gauge field $A_\mu(x, z)$ in AdS. This identification is very natural given the discussion in Section 5.1.3 that continuous global symmetries in the boundary theory should correspond to large gauge transformations in the bulk.

An especially important set of conserved currents in any translationally-invariant theory are those encapsulated in the energy-momentum tensor operator $\mathcal{T}^{\mu\nu}(x)$. The source $g_{\mu\nu}(x)$ coupling to $\mathcal{T}^{\mu\nu}(x)$ as

$$\int d^4x g_{\mu\nu}(x) \mathcal{T}^{\mu\nu}(x) \quad (5.16)$$

can be interpreted as an external spacetime metric deformation. According to (5.14) we can then associate it with the boundary value of the bulk metric $g_{\mu\nu}(x, z)$. We thus reach the important conclusion that the dual of a translationally-invariant gauge theory, in which the energy-momentum tensor is conserved, must involve dynamical gravity.

5.1.5 Normalizable vs. non-normalizable modes and mass-dimension relation

Having motivated the field/operator correspondence, we now elaborate on two important aspects of this correspondence: how the conformal dimension of an operator is related to

properties of the dual bulk field [252, 253], and how to interpret normalizable and non-normalizable modes of a bulk field in the boundary theory [313, 314].

For illustration we will consider a massive bulk scalar field Φ , dual to some scalar operator \mathcal{O} in the boundary theory. Although our main interest is the case in which the boundary theory is four-dimensional, it is convenient to present the equations for a general boundary spacetime dimension d . For this reason we will work with a generalization of the AdS metric (5.3) in which $x^\mu = (t, \vec{x})$ denote coordinates of a d -dimensional Minkowski spacetime.

The bulk action for Φ can be written as

$$S = -\frac{1}{2} \int dz d^d x \sqrt{-g} [g^{MN} \partial_M \Phi \partial_N \Phi + m^2 \Phi^2] + \dots . \quad (5.17)$$

We have canonically normalized Φ , and the dots stand for terms of order higher than quadratic. We have omitted these terms because they are proportional to positive powers Newton's constant, and are therefore suppressed by positive powers of $1/N_c$.

Since the bulk spacetime is translationally invariant along the x^μ -directions, it is convenient to introduce a Fourier decomposition in these directions by writing²

$$\Phi(z, x^\mu) = \int \frac{d^d k}{(2\pi)^d} e^{ik \cdot x} \Phi(z, k^\mu), \quad (5.18)$$

where $k \cdot x \equiv \eta_{\mu\nu} k^\mu x^\nu$ and $k^\mu \equiv (\omega, \vec{k})$, with ω and \vec{k} the energy and the spatial momentum, respectively. In terms of these Fourier modes the equation of motion for Φ derived from the action (5.17) is

$$z^{d+1} \partial_z (z^{1-d} \partial_z \Phi) - k^2 z^2 \Phi - m^2 R^2 \Phi = 0, \quad k^2 = -\omega^2 + \vec{k}^2. \quad (5.19)$$

Near the boundary $z \rightarrow 0$, the second term in (5.19) can be neglected and the equation can be readily solved to find the general asymptotic form of the solution:

$$\Phi(z, k) \approx A(k) z^{d-\Delta} + B(k) z^\Delta \quad \text{as } z \rightarrow 0, \quad (5.20)$$

where

$$\Delta = \frac{d}{2} + \nu, \quad \nu = \sqrt{m^2 R^2 + \frac{d^2}{4}}. \quad (5.21)$$

The integration ‘constants’ A and B actually acquire a dependence on k through the requirement that the solution be regular in the interior of AdS, i.e. for all $z > 0$. Fourier-transforming (5.20) back into coordinate space, we then find

$$\Phi(z, x) \approx A(x) z^{d-\Delta} + B(x) z^\Delta \quad \text{as } z \rightarrow 0. \quad (5.22)$$

The exponents in (5.22) are real provided

$$m^2 R^2 \geq -\frac{d^2}{4}. \quad (5.23)$$

²For notational simplicity we will use the same symbol to denote a function and its Fourier transform, distinguishing them only through their arguments.

In fact, one can show that the theory is stable for any m^2 in the range (5.23), whereas for $m^2 R^2 < -d^2/4$ there exist modes that grow exponentially in time and the theory is unstable [315–317]. In other words, in AdS space a field with a negative mass-squared does not lead to an instability provided the mass-squared is not ‘too negative’. Equation (5.23) is often called the Breitenlohner-Freedman (BF) bound. In the stable region (5.23) one must still distinguish between the finite interval $-d^2/4 \leq m^2 R^2 < -d^2/4 + 1$ and the rest of the region, $m^2 R^2 \geq -d^2/4 + 1$. In the first case both terms in (5.22) are normalizable with respect to the inner product

$$(\Phi_1, \Phi_2) = -i \int_{\Sigma_t} dz d\vec{x} \sqrt{-g} g^{tt} (\Phi_1^* \partial_t \Phi_2 - \Phi_2 \partial_t \Phi_1^*), \quad (5.24)$$

where Σ_t is a constant- t slice. We will comment on this case at the end of this section. For the moment let us assume that $m^2 R^2 \geq -d^2/4 + 1$. In this case the first term in (5.22) is non-normalizable and the second term, which is normalizable, does not affect the leading boundary behavior. As motivated in the last subsection, the boundary value of a bulk field Φ should be identified with the source for the corresponding boundary operator \mathcal{O} . Since in (5.22) the boundary behavior of Φ is controlled by $A(x)$, the presence of such a non-normalizable term should correspond to a deformation of the boundary theory of the form

$$S_{\text{bdry}} \rightarrow S_{\text{bdry}} + \int d^d x \phi(x) \mathcal{O}(x), \quad \text{with } \phi(x) = A(x). \quad (5.25)$$

In other words, *the non-normalizable term determines the boundary theory Lagrangian*. In particular, we see that in order to obtain a finite source $\phi(x)$, equation (5.14) should be generalized for $\Delta \neq d$ to

$$\phi(x) = \Phi|_{\partial \text{AdS}}(x) \equiv \lim_{z \rightarrow 0} z^{\Delta-d} \Phi(z, x). \quad (5.26)$$

In contrast, the normalizable modes are elements of the bulk Hilbert space. More explicitly, in the canonical quantization one expands Φ in terms of a basis of normalizable solutions of (5.19), from which one can then build the Fock space and compute the bulk Green’s functions, etc. The equivalence between the bulk and boundary theories implies that their respective Hilbert spaces should be identified. Thus we conclude that *normalizable modes should be identified with states of the boundary theory*. This identification gives an important tool for finding the spectrum of low-energy excitations of a strongly coupled gauge theory, and we shall verify it below by showing that normalizable solutions of the wave equation for Φ with well-defined momentum k give rise to poles in the two-point function $\langle \mathcal{O}(q) \mathcal{O}(-q) \rangle$ of its dual operator at momentum $q = k$. In the particular example at hand, one can readily see from (5.19) that, for a given \vec{k} , there is a continuous spectrum of ω , consistent with the fact that the boundary theory is scale invariant.

Furthermore, as will be discussed in Section 5.3 (and in Appendix C), the coefficient $B(x)$ of the normalizable term in (5.22) can be identified with the expectation value of \mathcal{O} in the presence of the source $\phi(x) = A(x)$, namely

$$\langle \mathcal{O}(x) \rangle_\phi = 2\nu B(x). \quad (5.27)$$

In the particular case of a purely normalizable solution, i.e. one with $A(x) = 0$, this equation yields the expectation value of the operator in the undeformed theory.

Equations (5.22), (5.25) and (5.27) imply that Δ , introduced in (5.21), should be identified as the conformal dimension of the boundary operator \mathcal{O} dual to Φ . Indeed, recall that a scale transformation of the boundary coordinates $x^\mu \rightarrow \Lambda x^\mu$ corresponds to the isometry $x^\mu \rightarrow \Lambda x^\mu$, $z \rightarrow \Lambda z$ in the bulk. Since Φ is a scalar field, under such an isometry it transforms as $\Phi'(\Lambda z, \Lambda x^\mu) = \Phi(z, x^\mu)$, which implies that the corresponding functions in the asymptotic form (5.22) must transform as $A'(\Lambda x^\mu) = \Lambda^{\Delta-d} A(x^\mu)$ and $B'(\Lambda x^\mu) = \Lambda^{-\Delta} B(x^\mu)$. This means that $A(x)$ and $B(x)$ have mass scaling dimensions $d - \Delta$ and Δ , respectively. Eqs. (5.25) and (5.27) are then consistent with each other and imply that $\mathcal{O}(x)$ has mass scaling dimension Δ .

The relations (5.21) and the near-boundary behaviour (5.22) also apply to a bulk spin-two field. In particular, they apply to the five-dimensional metric³ [253], for which $m^2 = 0$. Instead, in the case of a bulk p -form the dimension Δ of the dual operator is the largest root of the equation

$$m^2 R^2 = (\Delta - p)(\Delta + p - d) \quad (5.28)$$

and the near-boundary fall-off becomes [253]

$$\Phi_{\mu_1 \dots \mu_p} \approx A_{\mu_1 \dots \mu_p} z^{d-p-\Delta} + B_{\mu_1 \dots \mu_p} z^{\Delta-p} \quad \text{as } z \rightarrow 0. \quad (5.29)$$

A gauge field A_μ corresponds to the particular case $p = 1$, $m^2 = 0$. These results show that the metric and the gauge field are dual to operators of dimensions $\Delta = d$ and $\Delta = d - 1$, as expected for their respective dual operators, the stress-energy tensor and a vector current — see Section 5.1.4. An intuitive way to understand the near-boundary behaviours above is to note that the transverse traceless part of metric fluctuations behaves like a minimally coupled massless scalar — see Section 6.2.2. Consequently, the boundary behavior of this part, and by covariance of the rest of the metric (in Fefferman-Graham coordinates), is the same as that of a massless scalar. Instead, for a p -form one finds a massless scalar but with a spacetime-dependent coupling.

Before closing this section, let us return to the range $-d^2/4 \leq m^2 < -d^2/4 + 1$. We shall be brief because this is not a case that arises in later Sections. Since in this case both terms in (5.22) are normalizable, either one can be used to build the Fock space of physical states of the theory [315, 316]. This gives rise to two different boundary CFTs in which the dimensions of the operator $\mathcal{O}(x)$ are Δ or $d - \Delta$, respectively [319]. It was later realized [320, 321] that even more general quantizations are possible in which the modes used to build the physical states have both A and B nonzero, with a relation between the two that ensures that the asymptotic symmetries of AdS are preserved [322, 323].

5.2 Generalizations

5.2.1 Nonzero temperature and nonzero chemical potential

As discussed in Section 4.3, the same string theory reasoning giving rise to the equivalence (4.27) can be generalized to nonzero temperature by replacing the pure AdS metric (5.2)

³Although in some cases additional logarithmic terms must be included — see e.g. [318].

by that of a black brane in AdS_5 [312], Eqn. (4.30), which we copy here for convenience:

$$ds^2 = \frac{r^2}{R^2} (-f dt^2 + d\vec{x}^2) + \frac{R^2}{r^2 f} dr^2, \quad f(r) = 1 - \frac{r_0^4}{r^4}. \quad (5.30)$$

Equivalently, in terms of the z -coordinate, we replace Eqn. (5.2) by Eqn. (4.32), i.e.

$$ds^2 = \frac{R^2}{z^2} (-f dt^2 + d\vec{x}^2) + \frac{R^2}{z^2 f} dz^2, \quad f(z) = 1 - \frac{z^4}{z_0^4}. \quad (5.31)$$

The metrics above have an event horizon at $r = r_0$ and $z = z_0$, respectively, and the regions outside the horizon correspond to $r \in (r_0, \infty)$ and $z \in (0, z_0)$. This generalization can also be directly deduced from (4.27) as the black brane (5.30)–(5.31) is the only metric on the gravity side that satisfies the following properties: (i) it is asymptotically AdS_5 ; (ii) it is translationally-invariant along all the boundary directions and rotationally-invariant along the boundary spatial directions; (iii) it has a temperature and satisfies all laws of thermodynamics. It is therefore natural to identify the temperature and other thermodynamical properties of (5.30)–(5.31) with those of the SYM theory at nonzero temperature.

We mention in passing that there is also a nice connection between the black brane geometry (5.30)–(5.31) and the thermal-field formulation of finite-temperature field theory in terms of real time. Indeed, the fully extended spacetime of the black brane has two boundaries. Each of them supports an identical copy of the boundary field theory which can be identified with one of the two copies of the field theory in the Schwinger-Keldysh formulation. The thermal state can also be considered as the maximally-entangled state of the two field theories. For more details see [324, 325].

The Hawking temperature of the black brane can be calculated via the standard method [326] (see Appendix B for details) of demanding that the Euclidean continuation of the metric (5.31) obtained by the replacement $t \rightarrow -it_E$,

$$ds_E^2 = \frac{R^2}{z^2} (f dt_E^2 + dx_1^2 + dx_2^2 + dx_3^2) + \frac{R^2}{z^2 f} dz^2, \quad (5.32)$$

be regular at $z = z_0$. This requires that t_E be periodically identified with a period β given by

$$\beta = \frac{1}{T} = \pi z_0. \quad (5.33)$$

The temperature T is identified with the temperature of the boundary SYM theory since, since t_E corresponds precisely to the Euclidean time coordinate of the boundary theory. We emphasize here that while the Lorentzian spacetime (5.31) can be extended beyond the horizon $z = z_0$, the Euclidean metric (5.32) only exists for $z \in (0, z_0)$ as the spacetime ends at $z = z_0$, and ends smoothly once the choice (5.33) is made.

For a boundary theory with a $U(1)$ global symmetry, like $\mathcal{N} = 4$ SYM theory, one can furthermore turn on a chemical potential μ for the corresponding $U(1)$ charge. From the discussion of Section 5.1.4, this requires that the bulk gauge field A_μ which is dual to a boundary current J_μ satisfy the boundary condition

$$\lim_{z \rightarrow 0} A_t = \mu. \quad (5.34)$$

The above condition along with the requirement that the field A_μ should be regular at the horizon implies that there should be a radial electric field in the bulk, i.e. the black hole is now charged. We will not write the metric of a charged black hole explicitly, as we will not use it in this review. For more details see [327–332]. Similarly, in the case of theories with fundamental flavour introduced as probe D-branes, a baryon number chemical potential corresponds to an electric field on the branes [333–345].

5.2.2 A confining theory

Although our main interest is the deconfined phase of QCD, in this section we will briefly describe a simple example of a duality for which the field theory possesses a confining phase [312]. For simplicity we have chosen a model in which the field theory is three-dimensional, but all the essential features of this model extend to the string duals of more realistic confining theories in four dimensions.

We start by considering $\mathcal{N} = 4$ SYM theory at finite temperature. In the Euclidean description the system lives on $\mathbb{R}^3 \times S^1$. The circle direction corresponds to the Euclidean time, which is periodically identified with period $\beta = 1/T$. As is well known, at length scales much larger than β one can effectively think of this theory as the Euclidean version of pure three-dimensional Yang-Mills theory. The reasoning is that at these scales one can perform a Kaluza-Klein reduction along the circle. Since the fermions of the $\mathcal{N} = 4$ theory obey antiperiodic boundary conditions around the circle, their zero-mode is projected out, which means that all fermionic modes acquire a tree-level mass of order $1/\beta$. The scalars of the $\mathcal{N} = 4$ theory are periodic around the circle, but they acquire masses at the quantum level through their couplings to the fermions. The only fields that cannot acquire masses are the gauge bosons of the $\mathcal{N} = 4$ theory, since masses for them are forbidden by gauge invariance. Thus, at long distances the theory reduces to a pure Yang-Mills theory in three dimensions, which is confining and has a mass gap. The Lorentzian version of the theory is simply obtained by analytically continuing one of the \mathbb{R}^3 directions into the Lorentzian time. Thus, in this construction the ‘finite temperature’ of the original four-dimensional theory is a purely theoretical device. The effective Lorentzian theory in three dimensions is at zero temperature.

In order to obtain the gravity description of this theory we just need to implement the above procedure on the gravity side. We start with the Lorentzian metric (5.1)-(5.2) dual to $\mathcal{N} = 4$ SYM at zero-temperature. Then we introduce a nonzero temperature by going to Euclidean signature via $t \rightarrow -it_E$ and periodically identifying the Euclidean time. This results in the metric (5.32). Finally, we analytically continue one of the \mathbb{R}^3 directions, say x_3 , back into the new Lorentzian time: $x_3 \rightarrow it$. The final result is the metric

$$ds^2 = \frac{R^2}{z^2} (-dt^2 + dx_1^2 + dx_2^2 + f dt_E^2) + \frac{R^2}{z^2 f} dz^2. \quad (5.35)$$

In this metric the directions t, x_1, x_2 correspond to the directions in which the effective three-dimensional YM theory lives. The direction t_E is now a compact spatial direction. Note that since the original metric (5.32) smoothly ends at $z = z_0$, so does (5.35). This leads to a dramatic difference between the gauge theory dual to (5.35) and the original $\mathcal{N} = 4$ theory: the fact that the radial direction smoothly closes off at $z = z_0$ introduces a mass scale in the

boundary theory. To see this, note that the warp factor R^2/z^2 has a lower bound. Thus, when applying the discussion of Section 5.1.1 to (5.35), E_{YM} in Eqn. (5.4) will have a lower limit of order $M \sim 1/z_0$, implying that the theory develops a mass gap of this order. This can also be explicitly verified by solving the equation of motion of a classical bulk field (which is dual to some boundary theory operator) in the metric (5.35): for any fixed \vec{k} one finds a discrete spectrum of normalizable modes with a mass gap of order M . (Note that since the size of the circle parametrized by t_E is proportional to $1/z_0$, the mass gap is in fact comparable to the energies of Kaluza-Klein excitations on the circle.) As explained in Section 5.1.5, these normalizable modes can be identified with the glueball states of the boundary theory.

The fact that the gauge theory dual to the geometry (5.35) is a confining theory is further supported by several checks, including the following two. First, analysis of the expectation value of a Wilson loop reveals an area law, as will be discussed in Section 5.4. Second, the gravitational description can be used to establish that the theory described by (5.35) undergoes a deconfinement phase transition at a temperature $T_c \sim M$ set by the mass gap, above which the theory is again described by a geometry with a black hole horizon [312] (see [3, 4] for reviews).

The above construction resulted in an effective confining theory in three dimensions because we started with the theory on the worldvolume of D3-branes, which is a four-dimensional SYM theory. By starting instead with the near horizon solution of a large number of non-extremal D4-branes, which describes a SYM theory in five dimensions, the above procedure leads to the string dual of a Lorentzian confining theory that at long distance reduces to a four-dimensional pure Yang-Mills theory [312]. This has been used as the starting point of the Sakai-Sugimoto model for QCD [346, 347], which incorporates spontaneous chiral symmetry breaking and its restoration at high temperatures [348, 349]. For reviews on some of these topics see for example [3, 4].

5.2.3 Other generalizations

In addition to (4.27), many other examples of gauge/string dualities are known in different spacetime dimensions (see e.g. [262] and references therein). These include theories with fewer supersymmetries and theories which are not scale invariant, in particular confining theories [350–352] (see e.g. [353, 354] for reviews).

For a d -dimensional conformal theory, the dual geometry on the gravity side contains a factor of AdS_{d+1} and some other compact manifold.⁴ When expanded in terms of the harmonics of the compact manifold, the duality again reduces to that between a d -dimensional conformal theory and a gravity theory in AdS_{d+1} . In particular, in the classical gravity limit, this reduces to Einstein gravity in AdS_{d+1} coupled to various matter fields with the precise spectrum of matter fields depending on the specific theory under consideration. For a non-conformal theory the dual geometry is in general more complicated. Some simple, early examples were discussed in [355]. If a theory has a mass gap it is always the case that the dual bulk geometry closes off at some finite value of z_0 as in the example of Section 5.2.2.

All known examples of gauge/string dual pairs share the following common features with (4.27): (i) the field theory is described by elementary bosons and fermions coupled to non-Abelian

⁴Not necessarily in a direct product; the product may be warped.

gauge fields whose gauge group is specified by some N_c ; (ii) the string description reduces to classical (super)gravity in the large- N_c , strong coupling limit of the field theory. In this review we will use (4.27) as our prime example for illustration purposes, but the discussion can be immediately applied to other examples including non-conformal ones.

5.3 Correlation functions of local operators

In this section we will explain how to calculate correlation functions of local gauge-invariant operators of the boundary theory in terms of the dual gravity description. In this review we are mostly concerned with the properties of strongly coupled gauge theories at nonzero temperature. For this reason we are particularly interested in real-time retarded correlators, since these are relevant for determining linear responses, transport coefficients, spectral functions, etc. Below, we will first describe the general prescription for computing n -point Euclidean correlation functions, and then turn to the computation of finite-temperature retarded two-point functions.

5.3.1 Euclidean correlators

In view of the field/operator correspondence discussed in Sections 5.1.4 and 5.1.5, it is natural to postulate that the Euclidean partition functions of the two theories must agree upon the identification (5.26), namely that [252, 253]

$$Z_{\text{CFT}}[\phi(x)] = Z_{\text{string}}[\Phi|_{\partial\text{AdS}}(x)]. \quad (5.36)$$

Note that this equation makes sense because both sides of the equation are functionals of the same variables. Indeed, the most general partition function on the CFT side would include a source for each gauge-invariant operator in the theory, so one should think of $\phi(x)$ in Eqn. (5.36) as succinctly indicating the collection of all such sources. Since AdS has a boundary, in order to define the path integral over bulk fields in AdS one needs to specify a boundary condition for each field. The collection of all such boundary conditions is indicated by $\Phi|_{\partial\text{AdS}}(x)$ in Eqn. (5.36).

The right-hand side of (5.36) is in general not easy to compute, but it simplifies dramatically in the classical gravity limit (5.6), in which it can be obtained using the saddle point approximation as

$$Z_{\text{string}}[\phi] \simeq \exp\left(S^{(\text{ren})}[\Phi_c^{(E)}]\right), \quad (5.37)$$

where we have absorbed a conventional minus sign into the definition of the Euclidean action which avoids having some additional minus signs in various equations below and in the analytic continuation to Lorentzian signature.

In this equation $S^{(\text{ren})}[\Phi_c^{(E)}]$ is the renormalized on-shell classical supergravity action [356–362], namely the classical action evaluated on a solution $\Phi_c^{(E)}$ of the equations of motion determined by the boundary condition

$$\lim_{z \rightarrow 0} z^{\Delta-d} \Phi_c^{(E)}(z, x) = \phi(x) \quad (5.38)$$

and by the requirement that it be regular everywhere in the interior of the spacetime. The on-shell action needs to be renormalized because it typically suffers from infinite-volume (i.e. IR) divergences due to the integration region near the boundary of AdS [253]. These divergences are dual to UV divergences in the gauge theory, consistent with the UV/IR correspondence. The procedure to remove these divergences on the gravity side is well understood and goes under the name of ‘holographic renormalization’. Although it is an important ingredient of the gauge/string duality, it is also somewhat technical. In Appendix C we briefly review it in the context of a two-point function calculation. The interested reader may consult some of the references in this paragraph, as well as the review [318] for details.

Corrections to Eqn. (5.37) can be included as an expansion in α' and g_s , which correspond to $1/\sqrt{\lambda}$ and $1/N_c$ corrections in the gauge theory, respectively. Note that since the classical action (5.9) on the gravity side is proportional to $1/G_5$, from Eqn. (5.12) we see that $S^{(\text{ren})}[\Phi_c^{(E)}] \sim N_c^2$, as one would expect for the generating functional of an $SU(N_c)$ SYM theory in the large- N_c limit.

We conclude from Eqs. (5.36) and (5.37) that, in the large- N_c and large- λ limit, connected correlation functions of the gauge theory are given simply by functional derivatives of the on-shell, classical gravity action:

$$\langle \mathcal{O}(x_1) \dots \mathcal{O}(x_n) \rangle = \left. \frac{\delta^n S^{(\text{ren})}[\Phi_c^{(E)}]}{\delta \phi(x_1) \dots \delta \phi(x_n)} \right|_{\phi=0}. \quad (5.39)$$

It is often convenient to consider the one point function of \mathcal{O} in the presence of the source ϕ , which is given by

$$\langle \mathcal{O}(x) \rangle_\phi = \frac{\delta S^{(\text{ren})}[\Phi_c^{(E)}]}{\delta \phi(x)} = \lim_{z \rightarrow 0} z^{d-\Delta} \frac{\delta S^{(\text{ren})}[\Phi_c^{(E)}]}{\delta \Phi_c^{(E)}(z, x)}. \quad (5.40)$$

A very important application of this formula arises when the operator in question is the gauge theory stress-energy tensor. In this case the dual five-dimensional field is the bulk metric and, as we discussed below eqn. (5.29), $\Delta = d$, so the equation above becomes

$$\langle \mathcal{T}^{\mu\nu}(x) \rangle = \lim_{z \rightarrow 0} \frac{\delta S^{(\text{ren})}[g^{(E)}]}{\delta g_{\mu\nu}^{(E)}(x, z)}. \quad (5.41)$$

At this point we must distinguish between a tensor density and a tensor. From eqn. (5.16) it is clear that $\mathcal{T}^{\mu\nu}$ is a tensor density. In order to construct a tensor from it we define

$$\langle T^{\mu\nu}(x) \rangle = \lim_{z \rightarrow 0} \frac{2}{\sqrt{-g^{(E)}(x, z)}} \frac{\delta S^{(\text{ren})}[g^{(E)}]}{\delta g_{\mu\nu}^{(E)}(x, z)}. \quad (5.42)$$

Except for the dependence on the holographic coordinate z , the reader will probably recognize the right-hand side as a familiar definition of the stress tensor in classical field theory (see e.g. [363]). The square root of the metric in the denominator turns the result into a true tensor. This and the factor of 2 above can be fixed in a variety of ways, including the requirements that $T_{\mu\nu}$: agree with the known results in simple cases such as a free scalar field; be the appropriately-normalized generator of translations; be the appropriately-normalized

stress-tensor that appears on the right-hand side of Einstein equations. Eqn. (5.42) will play an important role in Section 7.3.

In classical mechanics, it is well known that the variation of the action with respect to the boundary value of a field results in the canonical momentum Π conjugate to the field, where the boundary in that case is usually a constant-time surface (see e.g. [364]). In the present case the boundary is a constant- z surface, but it is still useful to proceed by analogy with classical mechanics and to think of the derivative in the last term in (5.40) as the renormalized canonical momentum conjugate to $\Phi_c^{(E)}$ evaluated on the classical solution:

$$\Pi_c^{(\text{ren})}(z, x) = \frac{\delta S^{(\text{ren})}[\Phi_c^{(E)}]}{\delta \Phi_c^{(E)}(z, x)}. \quad (5.43)$$

With this definition eqn. (5.40) takes the form

$$\langle \mathcal{O}(x) \rangle_\phi = \lim_{z \rightarrow 0} z^{d-\Delta} \Pi_c^{(\text{ren})}(z, x). \quad (5.44)$$

More explicitly, for a bulk scalar field Φ with action (5.17) the corresponding Π is given by

$$\Pi = -g^{zz} \sqrt{-g} \partial_z \Phi. \quad (5.45)$$

We show in Appendix C that this implies that

$$\langle \mathcal{O}(x) \rangle_\phi = 2\nu B(x) \quad (5.46)$$

when identifying $A(x)$ in (5.22) with $\phi(x)$. In linear response theory the response in momentum space, i.e the expectation value of an operator, is proportional to the corresponding source, and the constant of proportionality (for each momentum) is the two-point function of the operator:

$$\langle \mathcal{O}(\omega_E, \vec{k}) \rangle_\phi = G_E(\omega_E, \vec{k}) \phi(\omega_E, \vec{k}). \quad (5.47)$$

Thus equation (5.46) immediately yields

$$G_E(\omega_E, \vec{k}) = \frac{\langle \mathcal{O}(\omega_E, \vec{k}) \rangle_\phi}{\phi(\omega_E, \vec{k})} = \lim_{z \rightarrow 0} z^{2(d-\Delta)} \frac{\Pi_c^{(\text{ren})}}{\Phi_c^{(E)}} = 2\nu \frac{B(\omega_E, \vec{k})}{A(\omega_E, \vec{k})}, \quad (5.48)$$

where ω_E denotes Euclidean frequency. In Appendix C, we give an explicit evaluation of (5.48). For some early work on the evaluation of higher-point functions see Refs. [365–367]. Incidentally, Eqn. (5.48) shows that the Euclidean correlator possesses a pole precisely at those frequencies for which $A(\omega_E, \vec{k})$ vanishes. In other words, the poles of the two-point function are in one-to-one correspondence with normalizable solutions of the equations of motion, as anticipated in the paragraph below Eqn. (5.26).

5.3.2 Real-time thermal correlators

We now proceed to the prescription for calculating real-time correlation functions. We will focus our discussion on *retarded* two-point functions because of their important role in characterizing linear response. Also, once the retarded function is known one can then use standard relations to obtain the other Green's functions.

We start with the relation between the retarded and the Euclidean two-point functions in momentum space,

$$G_E(\omega_E, \vec{k}) = G_R(i\omega_E, \vec{k}), \quad \omega_E > 0, \quad (5.49)$$

or inversely

$$G_R(\omega, \vec{k}) = G_E(-i(\omega + i\epsilon), \vec{k}). \quad (5.50)$$

If the Euclidean correlation functions G_E are known exactly, the retarded functions G_R can then be obtained through the simple analytic continuation (5.50). In most examples of interest, however, the Euclidean correlation functions can only be found numerically and analytic continuation to Lorentzian signature becomes difficult. Thus, it is important to develop techniques to calculate real-time correlation functions directly. Based on an educated guess that passed several consistency checks, a prescription for calculating retarded two-point functions in Lorentzian signature was first proposed by Son and Starinets in Ref. [368]. Ref. [325] later justified the prescription and extended it to n -point functions. Here we will follow the treatment given in Refs. [369, 370]. For illustration we consider the retarded two-point function for a scalar operator \mathcal{O} at nonzero temperature, which can be obtained from the propagation of the dual scalar field Φ in the geometry of an AdS black hole. The action for Φ again takes the form (5.17) with g_{MN} now given by the black brane metric (5.31).

Before giving the prescription, we note that in Lorentzian signature one cannot directly apply the procedure summarized by Eqn. (5.39) to obtain retarded functions. There are two immediate complications/difficulties. First, the Lorentzian black hole spacetime contains an event horizon and one also needs to impose appropriate boundary conditions there when solving the classical equation of motion for Φ . Second, since partition functions are defined in terms of path integrals, the resulting correlation functions should be time-ordered.⁵ As we now review, both complications can be dealt with in a simple manner.

The idea is to analytically continue the Euclidean classical solution $\Phi_c^{(E)}(\omega_E, \vec{k})$, as well as Eqs. (5.44)–(5.48), to Lorentzian signature according to (5.50). Clearly the analytic continuation of $\Phi_c^{(E)}(\omega_E, \vec{k})$,

$$\Phi_c(\omega, \vec{k}) = \Phi_c^{(E)}(-i(\omega + i\epsilon), \vec{k}), \quad (5.51)$$

solves the Lorentzian equation of motion. In addition, this solution obeys the in-falling boundary condition at the future event horizon of the black-brane metric (5.31). This property is important as it ensures that the retarded correlator is causal and only propagates information forward in time. This is intuitive since we expect that, classically, information can fall into the black hole horizon but not come out, so the retarded correlator should have no out-going component. Although it is intuitive, given its importance let us briefly verify that the in-falling boundary condition is satisfied. The Lorentzian equation of motion in momentum space for Φ_c in the black brane metric (5.31) takes the form

$$z^5 \partial_z \left[z^{-3} f(z) \partial_z \Phi \right] + \frac{\omega^2 z^2}{f(z)} \Phi - \vec{k}^2 z^2 \Phi - m^2 R^2 \Phi = 0, \quad (5.52)$$

⁵While it is possible to obtain Feynman functions this way, the procedure is quite subtle, since Feynman functions require imposing different boundary conditions for positive- and negative-frequency modes at the horizon and the choices of positive-frequency modes are not unique in a black hole spacetime. The correct choice corresponds to specifying the so-called Hartle-Hawking vacuum. For details see [325]. In contrast, the retarded function does not depend on the choice of the bulk vacuum in the classical limit as the corresponding bulk retarded function is given by the commutator of the corresponding bulk field — see Eqn. (5.61).

where $\vec{k}^2 = \delta^{ij} k_i k_j$. The corresponding Euclidean equation is obtained by setting $\omega = i\omega_E$. Near the horizon $z \rightarrow z_0$, since $f \rightarrow 0$ the last two terms in (5.52) become negligible compared with the second term and can be dropped. The resulting equation (with only the first two terms of (5.52)) then takes the simple form

$$\begin{aligned} \text{Lorentzian :} \quad & \partial_\xi^2 \Phi + \omega^2 \Phi = 0, \\ \text{Euclidean :} \quad & \partial_\xi^2 \Phi - \omega_E^2 \Phi = 0, \end{aligned} \tag{5.53}$$

in terms of a new coordinate

$$\xi \equiv \int^z \frac{dz'}{f(z')}. \tag{5.54}$$

Since $\xi \rightarrow +\infty$ as $z \rightarrow z_0$, in order for the Euclidean solution to be regular at the horizon we must choose the solution with the decaying exponential, i.e. $\Phi_c^{(E)}(\omega_E, \xi) \sim e^{-\omega_E \xi}$. The prescription (5.51) then yields $\Phi_c(\omega, \xi) \sim e^{i\omega \xi}$. Going back to coordinate space we find that near the horizon

$$\Phi_c(t, \xi) \sim e^{-i\omega(t-\xi)}. \tag{5.55}$$

As anticipated, this describes a wave propagating towards the direction in which ξ increases, i.e. falling into the horizon. (Had we chosen the opposite sign in the prescription (5.51) we would have obtained an out-going wave, as appropriate for the advanced correlator which obeys an out-going boundary condition at the past event horizon of the metric (5.31), and has no in-falling component.)

Given a Lorentzian solution satisfying the in-falling boundary condition at the horizon which can be expanded near the boundary according to (5.20), Eqs. (5.44)–(5.48) can be analytically continued to Lorentzian signature to obtain an intrinsic Lorentzian prescription for the expectation value and the retarded two-point function of an operator. The former is given by

$$\langle \mathcal{O}(\omega, \vec{k}) \rangle_\phi = \lim_{z \rightarrow 0} z^{d-\Delta} \Pi_c^{(\text{ren})} = 2\nu B(\omega, \vec{k}), \tag{5.56}$$

with $\phi(\omega, \vec{k}) = A(\omega, \vec{k})$, and the latter takes the form

$$G_R(\omega, \vec{k}) = \lim_{z \rightarrow 0} z^{2(d-\Delta)} \frac{\Pi_c^{(\text{ren})}}{\Phi_c(\omega, \vec{k})} = 2\nu \frac{B(\omega, \vec{k})}{A(\omega, \vec{k})}. \tag{5.57}$$

For practical purposes, let us recapitulate here the main result of this section, namely the algorithmic procedure for computing the real-time, finite-temperature retarded two-point function of a local, gauge-invariant operator $\mathcal{O}(x)$. This consists of the following steps:

1. Identify the bulk mode $\Phi(x, z)$ dual to $\mathcal{O}(x)$.
2. Find the Lorentzian-signature bulk effective action for Φ to quadratic order, and the corresponding linearized equation of motion in momentum space.
3. Find a solution $\Phi_c(k, z)$ to this equation with the boundary conditions that the solution is in-falling at the horizon and behaves as

$$\Phi_c(z, k) \approx A(k) z^{d-\Delta} + B(k) z^\Delta \tag{5.58}$$

near the boundary ($z \rightarrow 0$), where Δ is the dimension of $\mathcal{O}(x)$, d is the spacetime dimension of the boundary theory and $A(k)$ should be thought of as an arbitrary source for $\mathcal{O}(k)$. $B(k)$ is not an independent quantity but is determined by the boundary condition at the horizon and $A(k)$.

4. The retarded Green's function for \mathcal{O} is then given by

$$G_R(k) = 2\nu \frac{B(k)}{A(k)}, \quad (5.59)$$

where ν is defined in Eqn. (5.21).

In Section 8.5.2 we will discuss in detail an example of a retarded correlator of two electromagnetic currents.

Before closing this section, we note that an alternative way to compute boundary correlation functions which works in both Euclidean and Lorentzian signature is [371]

$$\langle \mathcal{O}(x_1) \cdots \mathcal{O}(x_n) \rangle = \lim_{z_i \rightarrow 0} (2\nu z_1^\Delta) \cdots (2\nu z_n^\Delta) \langle \Phi(z_1, x_1) \cdots \Phi(z_n, x_n) \rangle \quad (5.60)$$

where the correlator on the right-hand side is a correlation function in the bulk theory. In (5.60) it should be understood that whatever ordering one wants to consider, it should be same on both sides. For example, for the retarded two-point function G_R of \mathcal{O}

$$G_R(x_1 - x_2) = \lim_{z_1, z_2 \rightarrow 0} (2\nu z_1^\Delta) (2\nu z_2^\Delta) \mathcal{G}_R(z_1, x_1; z_2, x_2), \quad (5.61)$$

where \mathcal{G}_R denotes the retarded Green's function of the bulk field Φ .

5.4 Wilson loops

The expectation values of Wilson loops

$$W^r(\mathcal{C}) = \text{Tr} \mathcal{P} \exp \left[i \int_{\mathcal{C}} dx^\mu A_\mu(x) \right], \quad (5.62)$$

are an important class of non-local observables in any gauge theory. Here, $\int_{\mathcal{C}}$ denotes a line integral along the closed path \mathcal{C} , $W^r(\mathcal{C})$ is the trace of an $SU(N)$ -matrix in the representation r (one often considers fundamental or adjoint representations, i.e. $r = F, A$), the vector potential $A_\mu(x) = A_\mu^a(x) T^a$ can be expressed in terms of the generators T^a of the corresponding representation, and \mathcal{P} denotes path ordering. The expectation values of Wilson loops contain information about the non-perturbative physics of non-Abelian gauge field theories and have applications to many physical phenomena such as confinement, thermal phase transitions, quark screening, etc. For many of these applications it is useful to think of the path \mathcal{C} as that traversed by a quark. We will discuss some of these applications in Section 7. Here, we describe how to compute expectation values of Wilson loops in a strongly coupled gauge theory using its gravity description.

We again use $\mathcal{N} = 4$ SYM theory as an example. Now recall that the field content of this theory includes six scalar fields $\vec{\phi} = (\phi^1, \cdots, \phi^6)$ in the adjoint representation of the gauge




Figure 5.1: String worldsheet associated with a Wilson loop.

group. This means that in this theory one can write down the following generalization of (5.62) [372, 373]:

$$W(\mathcal{C}) = \frac{1}{N_c} \text{Tr} P \exp \left[i \oint_{\mathcal{C}} ds \left(A_\mu \dot{x}^\mu + \vec{n} \cdot \vec{\phi} \sqrt{\dot{x}^2} \right) \right], \quad (5.63)$$

where $\vec{n}(s)$ is a unit vector in \mathbb{R}^6 that parametrizes a path in this space (or, more precisely, in S^5), just like $x^\mu(s)$ parametrizes a path in $\mathbb{R}^{(1,3)}$. The factor of $\sqrt{\dot{x}^2}$ is necessary to make $\vec{n} \cdot \vec{\phi} \sqrt{\dot{x}^2}$ a density under worldline reparametrizations. Note that the operators (5.62) and (5.63) are equivalent in the case of a light-like loop (as will be discussed in Section 7.5) for which $\dot{x}^2 = 0$.

An important difference between the operators (5.62) and (5.63) is that (5.62) breaks supersymmetry, whereas (5.63) is locally 1/2-supersymmetric, meaning that for a straight-line contour (that is time-like in Lorentzian signature) the operator is invariant under half of the supercharges of the $\mathcal{N} = 4$ theory.

We will now argue that the generalized operator (5.63) has a dual description in terms of a string worldsheet. For this purpose it is useful to think of the loop \mathcal{C} as the path traversed by a quark. Although the $\mathcal{N} = 4$ SYM theory has no quarks, we will see below that these can be simply included by introducing in the gravity description open strings attached to a D-brane sitting at some radial position proportional to the quark mass. The endpoint of the open string on the D-brane is dual to the quark, so the boundary $\partial\Sigma$ of the string worldsheet Σ must coincide with the path \mathcal{C} traversed by the quark — see Fig. 5.1. This suggests that we must identify the expectation value of the Wilson loop operator, which gives the partition function (or amplitude) of the quark traversing \mathcal{C} , with the partition function of the dual string worldsheet Σ [372, 373]:

$$\langle W(\mathcal{C}) \rangle = Z_{\text{string}}[\partial\Sigma = \mathcal{C}]. \quad (5.64)$$

For simplicity, we will focus on the case of an infinitely heavy (non-dynamical) quark. This means that we imagine that we have pushed the D-brane all the way to the AdS boundary. Under these circumstances the boundary $\partial\Sigma = \mathcal{C}$ of the string worldsheet also lies within the boundary of AdS.

The key point to recall now is that the string endpoint couples both to the gauge field and to the scalar fields on the D-brane. This is intuitive since, after all, we obtained these fields as the massless modes of a quantized open string with endpoints attached to the D-brane. Physically, the coupling to the scalar fields is just a reflection of the fact that a string ending on a D-brane ‘pulls’ on it and deforms its shape, thus exciting the scalar fields which parametrize this shape. The direction orthogonal to the D-brane in which the string pulls is specified by \vec{n} . The coupling to the gauge field reflects the fact that the string endpoint behaves as a point-like particle charged under this gauge field. We thus conclude that an open string ending on a D-brane with a fixed \vec{n} excites both the gauge and the scalar fields, which suggests that the correct Wilson loop operator dual to the string worldsheet must include both types of fields and must therefore be given by (5.63).

The dual description of the operator (5.62) is the same as that of (5.63) except that the Dirichlet boundary conditions on the string worldsheet along the S^5 directions must be replaced by Neumann boundary conditions [374] (see also [375]). One immediate consequence is that, to leading order, the strong coupling results for the Wilson loop (5.63) with constant \vec{n} and for the Wilson loop (5.62) are the same. However, the two results differ at the next order in the $1/\sqrt{\lambda}$ expansion, since in the case of (5.62) we would have to integrate over the point on the sphere where the string is sitting. More precisely, at the one-loop level in the α' -expansion one finds that the determinants for quadratic fluctuations are different in the two cases [376].

In the large- N_c , large- λ limit, the string partition function $Z_{\text{string}}[\partial\Sigma = \mathcal{C}]$ greatly simplifies and is given by the exponential of the classical string action, i.e.

$$Z_{\text{string}}[\partial\Sigma = \mathcal{C}] = e^{iS(\mathcal{C})} \quad \rightarrow \quad \langle W(\mathcal{C}) \rangle = e^{iS(\mathcal{C})} . \quad (5.65)$$

The classical action $S(\mathcal{C})$ can in turn be obtained by extremizing the Nambu-Goto action for the string worldsheet with the boundary condition that the string worldsheet ends on the curve \mathcal{C} . More explicitly, parameterizing the two-dimensional world sheet by the coordinates $\sigma^\alpha = (\tau, \sigma)$, the location of the string world sheet in the five-dimensional spacetime with coordinates x^M is given by the Nambu-Goto action (4.13). The fact that the action is invariant under coordinate changes of σ^α will allow us to pick the most convenient worldsheet coordinates (τ, σ) for each occasion.

Note that the large- N_c and large- λ limits are both crucial for (5.65) to hold. Taking $N_c \rightarrow \infty$ at fixed λ corresponds to taking the string coupling to zero, meaning that we can ignore the possibility of loops of string breaking off from the string world sheet. Additionally taking $\lambda \rightarrow \infty$ corresponds to sending the string tension to infinity, which implies that we can neglect fluctuations of the string world sheet. Under these circumstances the string worldsheet ‘hanging down’ from the contour \mathcal{C} takes on its classical configuration, without fluctuating or splitting off loops.

As a simple example let us first consider a contour \mathcal{C} given by a straightline along the time direction with length \mathcal{T} which describes an isolated static quark at rest. On the field theory

side we expect that the expectation value of the Wilson line should be given by

$$\langle W(\mathcal{C}) \rangle = e^{-iM\mathcal{T}}, \quad (5.66)$$

where M is the mass of the quark. From the symmetry of the problem, the corresponding bulk string worldsheet should be that of a straight string connecting the boundary and the Poincaré horizon and translated along the time direction by \mathcal{T} . The action of such a string worldsheet is infinite since the proper distance from the boundary to the center of AdS is infinite. This is consistent with the fact that the external quark has an infinite mass. A finite answer can nevertheless be obtained if we introduce an IR regulator in the bulk, putting the boundary at $z = \epsilon$ instead of $z = 0$. From the IR/UV connection this corresponds to introducing a short-distance (UV) cutoff in the boundary theory. Choosing $\tau = t$ and $\sigma = z$ the string worldsheet is given by $x^i(\sigma, \tau) = \text{const.}$, and the induced metric on the worldsheet is then given by

$$ds^2 = \frac{R^2}{\sigma^2}(-d\tau^2 + d\sigma^2). \quad (5.67)$$

Evaluating the Nambu-Goto action on this solution yields

$$S = S_0 \equiv -\frac{\mathcal{T}R^2}{2\pi\alpha'} \int_{\epsilon}^{\infty} \frac{dz}{z^2} = -\frac{\sqrt{\lambda}}{2\pi\epsilon} \mathcal{T}, \quad (5.68)$$

where we have used the fact that $R^2/\alpha' = \sqrt{\lambda}$. Using (5.65) and (5.66) we then find that

$$M = \frac{\sqrt{\lambda}}{2\pi\epsilon}. \quad (5.69)$$

5.4.1 Rectangular loop: vacuum

Now let us consider a rectangular loop sitting at a constant position on the S^5 [372, 373]. The long side of the loop extends along the time direction with length \mathcal{T} , and the short side extends along the x_1 -direction with length L . We will assume that $\mathcal{T} \gg L$. Such a configuration can be thought of as consisting of a static quark-antiquark pair separated by a distance L . Therefore we expect that the expectation value of the Wilson loop (with suitable renormalization) gives the potential energy between the pair, i.e. we expect that

$$\langle W(\mathcal{C}) \rangle = e^{-iE_{\text{tot}}\mathcal{T}} = e^{-i(2M+V(L))\mathcal{T}} = e^{iS(\mathcal{C})}, \quad (5.70)$$

where E_{tot} is the total energy for the whole system and $V(L)$ is the potential energy between the pair. In the last equality we have used (5.65). We will now proceed to calculate $S(\mathcal{C})$ for a rectangular loop.

It is convenient to choose the worldsheet coordinates to be

$$\tau = t, \quad \sigma = x_1. \quad (5.71)$$

Since $\mathcal{T} \gg L$, we can assume that the surface is translationally invariant along the τ -direction, i.e. the extremal surface should have nontrivial dependence only on σ . Given the symmetries of the problem we can also set

$$x_3(\sigma) = \text{const.}, \quad x_2(\sigma) = \text{const.} \quad (5.72)$$




Figure 5.2: String (red) associated with a quark-antiquark pair.

Thus the only nontrivial function to solve for is $z = z(\sigma)$ (see Fig. 5.2), subject to the boundary condition

$$z\left(\pm\frac{L}{2}\right) = 0. \quad (5.73)$$

Using the form (5.3) of the spacetime metric and Eqs. (5.71)–(5.72), the induced metric on the worldsheet is given by

$$ds_{ws}^2 = \frac{R^2}{z^2} (-d\tau^2 + (1 + z'^2)d\sigma^2), \quad (5.74)$$

giving rise to the Nambu-Goto action

$$S_{\text{NG}} = -\frac{R^2\mathcal{T}}{2\pi\alpha'} \int_{-\frac{L}{2}}^{\frac{L}{2}} d\sigma \frac{1}{z^2} \sqrt{1 + z'^2}, \quad (5.75)$$

where $z' = dz/d\sigma$. Since the action and the boundary condition are symmetric under $\sigma \rightarrow -\sigma$, $z(\sigma)$ should be an even function of σ . Introducing dimensionless coordinates via

$$\sigma = L\xi, \quad z(\sigma) = Ly(\xi) \quad (5.76)$$

we then have

$$S_{\text{NG}} = -\frac{2R^2\mathcal{T}}{2\pi\alpha' L} Q, \quad \text{with } Q = \int_0^{\frac{1}{2}} \frac{d\xi}{y^2} \sqrt{1 + y'^2}. \quad (5.77)$$

Note that Q is a numerical constant. As we will see momentarily, it is in fact divergent and therefore it should be defined more carefully. The equation of motion for y is given by

$$y'^2 = \frac{y_0^4 - y^4}{y^4} \quad (5.78)$$

with y_0 the turning point at which $y' = 0$, which by symmetry should happen at $\xi = 0$. y_0 can thus be determined by the condition

$$\frac{1}{2} = \int_0^{\frac{1}{2}} d\xi = \int_0^{y_0} \frac{dy}{y'} = \int_0^{y_0} dy \frac{y^2}{\sqrt{y_0^4 - y^4}} \quad \rightarrow \quad y_0 = \frac{\Gamma(\frac{1}{4})}{2\sqrt{\pi}\Gamma(\frac{3}{4})}. \quad (5.79)$$

It is then convenient to change integration variable in Q from ξ to y to get

$$Q = y_0^2 \int_0^{y_0} \frac{dy}{y^2 \sqrt{y_0^4 - y^4}}. \quad (5.80)$$

This is manifestly divergent at $y = 0$, but the divergence can be interpreted as coming from the infinite rest masses of the quark and the antiquark. As in the discussion after (5.66), we can obtain a finite answer by introducing an IR cutoff in the bulk by putting the boundary at $z = \epsilon$, i.e. by replacing the lower integration limit in (5.80) by ϵ . The potential $V(L)$ between the quarks is then obtained by subtracting $2M\mathcal{T}$ from (5.77) (with M given by (5.69)) and then taking $\epsilon \rightarrow 0$ at the end of the calculation. One then finds the finite answer

$$V(L) = -\frac{4\pi^2}{\Gamma^4(\frac{1}{4})} \frac{\sqrt{\lambda}}{L}, \quad (5.81)$$

where again we used the fact that $R^2/\alpha' = \sqrt{\lambda}$ to translate from gravity to gauge theory variables. Note that the $1/L$ dependence is simply a consequence of conformal invariance. The non-analytic dependence on the coupling, i.e. the $\sqrt{\lambda}$ factor, could not be obtained at any finite order in perturbation theory. From the gravity viewpoint, however, it is a rather generic result, since it is due the fact that the tension of the string is proportional to $1/\alpha'$. The above result is valid at large λ . At small λ , the potential between a quark and an anti-quark in an $\mathcal{N} = 4$ theory is given by [377]

$$E = -\frac{\pi\lambda}{L} \quad (5.82)$$

to lowest order in the weak-coupling expansion.

It is remarkable that the calculation of a Wilson loop in a strongly interacting gauge theory has been simplified to a classical mechanics problem no more difficult than finding the catenary curve describing a string suspended from two points, hanging in a gravitational field — in this case the gravitational field of the AdS spacetime.

Note that given (5.81), the boundary short-distance cutoff ϵ in (5.69) can be interpreted as the size of the external quark. One might have expected (incorrectly) that a short distance cutoff on the size of the quark should be given by the Compton wavelength $1/M \sim \epsilon/\sqrt{\lambda}$, which is much smaller than ϵ . Note that the size of a quark should be defined by either its Compton wavelength or by the distance between a quark and an anti-quark at which the potential is of the order of the quark mass, whichever is bigger. In a weakly coupled theory, the Compton wavelength is bigger, while in a strongly coupled theory with potential (5.81), the latter is bigger and is of order ϵ .

5.4.2 Rectangular loop: nonzero temperature

We now consider the expectation of the rectangular loop at nonzero temperature [378, 379]. In this case the bulk gravity geometry is given by that of the black brane (5.31). The set-up of the calculation is exactly the same as in Eqs. (5.71)–(5.73) for the vacuum. The induced worldsheet metric is now given by

$$ds_{ws}^2 = \frac{R^2}{z^2} \left(-f(z)d\tau^2 + \left(1 + \frac{z'^2}{f} \right) d\sigma^2 \right), \quad (5.83)$$




Figure 5.3: String (red) associated with a quark-antiquark pair in a plasma with temperature $T > 0$. The preferred configuration beyond a certain separation L_s consists of two independent strings.

which yields the Nambu-Goto action

$$S_{\text{NG}} = -\frac{R^2 \mathcal{T}}{2\pi\alpha'} \int_{-\frac{L}{2}}^{\frac{L}{2}} d\sigma \frac{1}{z^2} \sqrt{f(z) + z'^2}. \quad (5.84)$$

The crucial difference between the equation of motion following from (5.84) and that following from (5.75) is that in the present case there exists a maximal value $L_s \sim 1/T$ beyond which no nontrivial solutions exist [378, 379] — see Fig. 5.3. Instead, the solution beyond this maximal separation consists of two disjoint vertical strings ending at the black hole horizon. The physical reason can be easily understood qualitatively from the figure. At some separation, the lowest point on the string touches the horizon. Surely at and beyond this separation the string can minimize its energy by splitting into two independent strings, each of which falls through the horizon. The precise value of L_s is defined as the quark-antiquark separation at which the free energy of the disconnected configuration becomes smaller than that of the connected configuration. This happens at a value of L for which the lowest point of the connected configuration is close to but still somewhat above the horizon. Once $L > L_s$, the quark-antiquark separation can then be increased further at no additional energy cost, so the potential becomes constant and the quark and the antiquark are perfectly screened from each other by the plasma between them. (See, for example, Ref. [380] for a careful discussion of the corrections to this large- N_c , large- λ result.)

5.4.3 Rectangular loop: a confining theory

For comparison, let us consider the expectation value of a rectangular loop in the $2 + 1$ -dimensional confining theory [312] (for a review see [381]) whose metric is given by (5.35), which we reproduce here for convenience:

$$ds^2 = \frac{R^2}{z^2} (-dt^2 + dx_1^2 + dx_2^2 + f dt_E) + \frac{R^2}{z^2 f} dz^2, \quad f = 1 - \frac{z^4}{z_0^4}. \quad (5.85)$$




Figure 5.4: String (red) associated with a quark-antiquark pair in a confining theory.

As discussed earlier the crucial difference between (5.85) and AdS is that the spacetime (5.85) ends smoothly at a finite value $z = z_0$, which introduces a scale in the theory. The difference as compared to the finite-temperature case is that in the confining geometry the string has no place to end, so in order to minimize its energy it tends to drop down to z_0 and to run parallel there — see Fig. 5.4.

Again the set-up of the calculation is completely analogous to the cases above. The induced worldsheet metric is now given by

$$ds_{ws}^2 = \frac{R^2}{z^2} \left(-d\tau^2 + \left(1 + \frac{z'^2}{f} \right) d\sigma^2 \right), \quad (5.86)$$

and the corresponding the Nambu-Goto action is

$$S_{\text{NG}} = -\frac{R^2 \mathcal{T}}{2\pi\alpha'} \int_{-\frac{L}{2}}^{\frac{L}{2}} d\sigma \frac{1}{z^2} \sqrt{1 + \frac{z'^2}{f(z)}}. \quad (5.87)$$

When L is large, the string quickly drops to $z = z_0$ and runs parallel there. We thus find that the action can be approximated by (after subtracting the vertical parts which can be interpreted as due to the static quark masses)

$$-S(\mathcal{C}) - 2MT \approx \frac{R^2 \mathcal{T}}{2\pi\alpha'} \frac{L}{z_0^2}, \quad (5.88)$$

which gives rise to a confining potential

$$V(L) = \sigma_s L, \quad \sigma_s = \frac{\sqrt{\lambda}}{2\pi z_0^2}. \quad (5.89)$$

The constant σ_s can be interpreted as the effective string tension. As mentioned in Sec. 5.2.2 the mass gap for this theory is $M \sim 1/z_0$, so we find that $\sigma_s \sim \sqrt{\lambda} M^2$. Although we have described the calculation only for one example of a confining gauge theory, the qualitative features of Fig. 5.4 generalize. In any confining gauge theory with a dual gravity description, as a quark-antiquark pair are separated the string hanging beneath them sags down to some ‘depth’ z_0 and then as the separation is further increased it sags no further. Further increasing the separation means adding more and more string at the same depth z_0 , which costs an energy that increases linearly with separation. Clearly, any metric in which a suspended string

behaves like this cannot be conformal; it has a length scale z_0 built into it in some way. This length scale z_0 in the gravitational description corresponds via the IR/UV correspondence to the mass gap $M \sim 1/z_0$ for the gauge theory and to the size of the ‘glueballs’ in the gauge theory, which is of order z_0 .

To summarize, we note that the qualitative behavior of the Wilson loop discussed in various examples above is only determined by gross features of the bulk geometry. The $1/L$ behavior (5.82) in the conformal vacuum follows directly from the scaling symmetry of the bulk geometry; the area law (5.88) in the confining case has to do with the fact that a string has no place to end in the bulk when the geometry smoothly closes off; and the screening behavior at finite temperature is a consequence of the fact that a string can fall through the black hole horizon. The difference between Figs. 5.2 and 5.4 highlights the fact that $\mathcal{N} = 4$ SYM theory is not a good model for the vacuum of a confining theory like QCD. However, as we will discuss in Section 7.6, the potential obtained from Fig. 5.3 is not a bad caricature of what happens in the deconfined phase of QCD. This is one of many ways of seeing that $\mathcal{N} = 4$ SYM at $T \neq 0$ is more similar to QCD above T_c than $\mathcal{N} = 4$ SYM at $T = 0$ is to QCD at $T = 0$. A heuristic way of thinking about this is to note that at low temperatures the putative horizon would be at a $z_{\text{hor}} > z_0$, i.e. it is far below the bottom of Fig. 5.4, and therefore it plays no role while at large temperatures, the horizon is far above z_0 and it is z_0 that plays no role. At some intermediate temperature, the theory has undergone a phase transition from a confined phase described by Fig. 5.4 into a deconfined phase described by Fig. 5.3.⁶ Unlike in QCD, this deconfinement phase transition is a 1st order phase transition in the large- N_c , strong coupling limit under consideration, and the theory in the deconfined phase loses all memory about the confinement scale z_0 . Presumably corrections away from this limit, in particular finite- N_c corrections, could turn the transition into a higher-order phase transition or even a cross-over.

5.5 Introducing fundamental matter

All the matter degrees of freedom of $\mathcal{N} = 4$ SYM, the fermions and the scalars, transform in the adjoint representation of the gauge group. In QCD, however, the quarks transform in the fundamental representation. Moreover, most of what we know about QCD phenomenologically comes from the study of quarks and their bound states. Therefore, in order to construct holographic models more closely related to QCD, we must introduce degrees of freedom in the fundamental representation. It turns out that there is a rather simple way to do this in the limit in which the number of quark species, or flavours, is much smaller than the number of colours, i.e. when $N_f \ll N_c$. Indeed, in this limit the introduction of N_f flavours in the gauge theory corresponds to the introduction of N_f D-brane probes in the AdS geometry sourced by the D3-branes [383–385]. This is perfectly consistent with the well-known fact that the topological representation of the large- N_c expansion of a gauge theory with quarks involves Riemann surfaces with boundaries — see Section 4.1.2. In the string description, these surfaces correspond to the worldsheets of open strings whose endpoints must be attached to D-branes. In the context of the gauge/string duality, the intuitive idea is that closed strings

⁶The way we have described the transition is a crude way of thinking about the so-called Hawking-Page phase transition between a spacetime without and with a black hole [253, 382].




Figure 5.5: Excitations of the system in the open string description.

living in AdS are dual to gauge-invariant operators constructed solely out of gauge fields and adjoint matter, e.g. $\mathcal{O} = \text{Tr}F^2$, whereas open strings are dual to meson-like operators, e.g. $\mathcal{O} = \bar{q}q$. In particular, the two end-points of an open string, which are forced to lie on the D-brane probes, are dual to a quark and an antiquark, respectively.

5.5.1 The decoupling limit with fundamental matter

The fact that the introduction of gauge theory quarks corresponds to the introduction of D-brane probes in the string description can be more ‘rigorously’ motivated by repeating the arguments of Sections 4.2.2, 4.2.2 and 4.3 in the presence of N_f D p -branes, as indicated in Fig. 5.5. We shall be more precise about the value of p and the precise orientation of the branes later; for the moment we simply assume $p > 3$.

As in Section 4.2.2, when $g_s N_c \ll 1$ the excitations of this system are accurately described by interacting closed and open strings living in flat space. In this case, however, the open string sector is richer. As before, open strings with both end-points on the D3-branes give rise, at low energies, to the $\mathcal{N} = 4$ SYM multiplet in the adjoint of $SU(N_c)$. We see from Eqn. (4.17) that the coupling constant for these degrees of freedom is dimensionless, and therefore these degrees of freedom remain interacting at low energies. The coupling constant for the open strings with both end-points on the D p -branes, instead, has dimensions of $(\text{length})^{p-3}$. Therefore the effective dimensionless coupling constant at an energy E scales

as $g_{Dp} \propto E^{p-3}$. Since we assume that $p > 3$, this implies that, just like the closed strings, the p - p strings become non-interacting at low energies. Finally, consider the sector of open strings with one end-point on the D3-branes and one end-point on the Dp -branes. These degrees of freedom transform in the fundamental of the gauge group on the D3-branes and in the fundamental of the gauge group on the Dp -branes, namely in the bifundamental of $SU(N_c) \times SU(N_f)$. Consistently, these 3- p strings interact with the 3-3 and the p - p strings with strengths given by the corresponding coupling constants on the D3-branes and on the Dp -branes. At low energies, therefore, only the interactions with the 3-3 strings survive. In addition, since the effective coupling on the Dp -branes vanishes, the corresponding gauge group $SU(N_f)$ becomes a global symmetry group. This is the origin of the flavour symmetry expected in the presence of N_f (equal mass) quark species in the gauge theory.

To summarize, when $g_s N_c \ll 1$ the low-energy limit of the D3/ Dp system yields two decoupled sectors. The first sector is free and consists of closed strings in ten-dimensional flat space and p - p strings propagating on the worldvolume of N_f Dp -branes. The second sector is interacting and consists of a four-dimensional $\mathcal{N} = 4$ SYM multiplet in the adjoint of $SU(N_c)$, coupled to the light degrees of freedom coming from the 3- p strings. We will be more precise about the exact nature of these degrees of freedom later, but for the moment we emphasize that they transform in the fundamental representation of the $SU(N_c)$ gauge group, and in the fundamental representation of a global, flavour symmetry group $SU(N_f)$.

Consider now the closed string description at $g_s N_c \gg 1$. In this case, as in Section 4.2.3, the D3-branes may be replaced by their backreaction on spacetime. If we assume that $g_s N_f \ll 1$, which is consistent with $N_f \ll N_c$, we may still neglect the backreaction of the Dp -branes. In other words, we may treat the Dp -branes as probes living in the geometry sourced by the D3-branes, with the Dp -branes not modifying this geometry. The excitations of the system in this limit consist of closed strings and open p - p strings that propagate in two different regions, the asymptotically flat region and the $AdS_5 \times S^5$ throat – see Fig. 5.6. As in section 4.2.3, these two regions decouple from each other in the low-energy limit. Also as in section 4.2.3, in this limit the strings in the asymptotically flat region become non-interacting, whereas those in the throat region remain interacting because of the gravitational redshift.

Comparing the two descriptions above, we see that the low-energy limit at both small and large values of $g_s N_c$ contains a free sector of closed and open p - p strings. As in section 4.3 we identify these free sectors, and we conjecture that the interacting sectors on each side provide dual descriptions of the same physics. In other words, we conjecture that the $\mathcal{N} = 4$ SYM coupled to N_f flavours of fundamental degrees of freedom is dual to type IIB closed strings in $AdS_5 \times S^5$, coupled to open strings propagating on the worldvolume of N_f Dp -brane probes.

It is worth clarifying the following conceptual point before closing this section. It is sometimes stated that, in the 't Hooft limit in which $N_c \rightarrow \infty$ with N_f fixed, the dynamics is completely dominated by the gluons, and therefore that the quarks can be completely ignored. One may then wonder what the interest of introducing fundamental degrees of freedom in a large- N_c theory may be. There are several answers to this. First of all, in the presence of fundamental matter, it is more convenient to think of the large- N_c limit *à la* Veneziano, in which N_f/N_c is kept small but finite. Any observable can then be expanded in powers of $1/N_c^2$ and N_f/N_c . As we will see, this is precisely the limit that is captured by the dual description in terms of N_f D-brane probes in $AdS_5 \times S^5$. The leading D-brane contribution

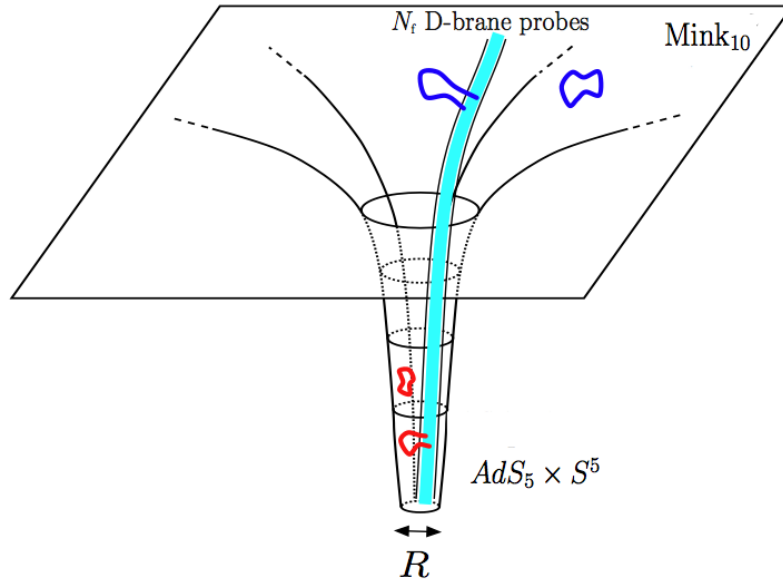


Figure 5.6: Excitations of the system in the second description.

will give us the leading contribution of the fundamental matter, of relative order N_f/N_c . The Veneziano limit is richer than the 't Hooft limit, since setting $N_f/N_c = 0$ one recovers the 't Hooft limit. The second point is that, even in the 't Hooft limit, the quarks should not be regarded as irrelevant, but rather as valuable probes of the gluon-dominated dynamics. It is their very presence in the theory that allows one to ask questions about heavy quarks in the plasma, jet quenching, meson physics, photon emission, etc. The answers to these questions are of course dominated by the gluon dynamics, but without dynamical quarks in the theory such questions cannot even be posed. There is a completely analogous statement in the dual gravity description: To leading order the geometry is not modified by the presence of D-brane probes, but one needs to introduce these probes in order to pose questions about heavy quarks in the plasma, parton energy loss, mesons, photon production, etc. In this sense, the D-brane probes allow one to decode information already contained in the geometry.

5.5.2 Models with fundamental matter

Above, we motivated the inclusion of fundamental matter via the introduction of N_f 'flavour' Dp -brane probes in the background sourced by N_c 'colour' D3-branes. However, we were deliberately vague about the value of p , about the relative orientation between the flavour and the colour branes, and about the precise nature of the flavour degrees of freedom in the gauge theory. Here we will address these points. Since we assumed $p > 3$ in order to decouple the p - p strings, and since we wish to consider stable Dp -branes in type IIB string theory, we must have $p = 5$ or $p = 7$ — see Section 4.2.2. In other words, we must consider D5- and D7-brane probes.




Figure 5.7: D3-D5 configuration (5.90) with a string stretching between them. The 12-directions common to both branes are suppressed.

Consider first adding flavour D5-branes. We will indicate the relative orientation between these and the colour D3-branes by an array like, for example,

$$\begin{array}{rcccccccc}
 \text{D3:} & 1 & 2 & 3 & - & - & - & - & - & - \\
 \text{D5:} & 1 & 2 & - & 4 & 5 & 6 & - & - & -
 \end{array}
 \tag{5.90}$$

This indicates that the D3- and the D5-branes share the 12-directions. The 3-direction is transverse to the D5-branes, the 456-directions are transverse to the D3-branes, and the 789-directions are transverse to both sets of branes. This means that the two sets of branes can be separated along the 789-directions, and therefore they do not necessarily intersect, as indicated in Fig. 5.7. It turns out that the lightest states of a D3-D5 string have a minimum mass given by what one would have expected on classical grounds, namely $M = T_{\text{str}} L = L/2\pi\ell_s^2$, where T_{str} is the string tension (4.11) and L is the minimum distance between the D3- and the D5-branes.⁷ These states can therefore be arbitrarily light, even massless, provided L is sufficiently small. Generic excited states, as usual, have an additional mass set by the string scale alone, m_s . The only exception are excitations in which the string

⁷In order to really establish this formula one must quantise the D3-D5 strings and compute the ground state energy. In the case at hand, the result coincides with the classical expectation. The underlying reason is that, because the configuration (5.90) preserves supersymmetry, corrections to the classical ground state energy coming from bosonic and fermionic quantum fluctuations cancel each other out exactly. For other brane configurations like (5.92) this does not happen.

moves rigidly with momentum \vec{p} in the 12-directions, in which case the energy squared is just $M^2 + \vec{p}^2$. This is an important observation because it means that in the decoupling limit, in which one focuses on energies $E \ll m_s$, only a finite set of modes of the D3-D5 strings survive, and moreover these modes can only propagate along the directions common to both branes. From the viewpoint of the dual gauge theory, this translates into the statement that the degrees of freedom in the fundamental representation are localised on a defect — in the example at hand, on a plane that extends along the 12-directions and lies at a constant position in the 3-direction. As an additional example, the configuration

$$\begin{array}{l} \text{D3:} \quad 1 \quad 2 \quad 3 \quad - \quad - \quad - \quad - \quad - \quad - \\ \text{D5:} \quad 1 \quad - \quad - \quad 4 \quad 5 \quad 6 \quad 7 \quad - \quad - \end{array} \quad (5.91)$$

corresponds to a dual gauge theory in which the fundamental matter is localised on a line — the 1-direction.

We thus conclude that, if we are interested in adding to the $\mathcal{N} = 4$ SYM theory fundamental matter degrees of freedom that propagate in 3+1 dimensions (just like the gluons and the adjoint matter), then we must orient the flavour D-branes so that they extend along the 123-directions. This condition leaves us with two possibilities:

$$\begin{array}{l} \text{D3:} \quad 1 \quad 2 \quad 3 \quad - \quad - \quad - \quad - \quad - \quad - \\ \text{D5:} \quad 1 \quad 2 \quad 3 \quad 4 \quad 5 \quad - \quad - \quad - \quad - \end{array} \quad (5.92)$$

and

$$\begin{array}{l} \text{D3:} \quad 1 \quad 2 \quad 3 \quad - \quad - \quad - \quad - \quad - \quad - \\ \text{D7:} \quad 1 \quad 2 \quad 3 \quad 4 \quad 5 \quad 6 \quad 7 \quad - \quad - \quad . \end{array} \quad (5.93)$$

So far we have not been specific about the precise nature of the fundamental matter degrees of freedom — for example, whether they are fermions or bosons, etc. This also depends on the relative orientation of the branes. It turns out that for the configuration (5.92), the ground state energy of the D3-D5 strings is (for sufficiently small L) negative, that is, the ground state is tachyonic, signaling an instability in the system. This conclusion is valid at weak string coupling, where the string spectrum can be calculated perturbatively. While it is possible that the instability is absent at strong coupling, we will not consider this configuration further in this review.

We are therefore left with the D3-D7 system (5.93). Quantisation of the D3-D7 strings shows that the fundamental degrees of freedom in this case consist of N_f complex scalars and N_f Dirac fermions, all of them with equal masses given by

$$M_q = \frac{L}{2\pi\alpha'} . \quad (5.94)$$

In a slight abuse of language, we will collectively refer to all these degrees of freedom as ‘quarks’. The fact that they all have exactly equal masses is a reflection of the fact that the addition of the N_f D7-branes preserves a fraction of the original supersymmetry of the SYM theory. More precisely, the original $\mathcal{N} = 4$ is broken down to $\mathcal{N} = 2$, under which the fundamental scalars and fermions transform as part of a single supermultiplet. In the rest of the review, especially in Section 8, we will focus our attention on this system as a model for gauge theories with fundamental matter.

Chapter 6

Bulk properties of strongly coupled plasma

Up to this point in this review, we have laid the groundwork needed for what is to come in two halves. In Sections 2 and 3 we have introduced the theoretical, phenomenological and experimental challenges posed by the study of the deconfined phase of QCD and in Sections 4 and 5 we have motivated and described gauge/string duality, providing the reader with most of the conceptual and computational machinery necessary to perform many calculations. Although we have foreshadowed their interplay at various points, these two long introductions have to a large degree been separately self-contained. In the next three sections, we weave these strands together. In these sections, we shall review applications of gauge/gravity duality to the study of the strongly coupled plasma of $\mathcal{N} = 4$ SYM theory at nonzero temperature, focussing on the ways in which these calculations can guide us toward the resolution of the challenges described in Sections 2 and 3.

The study of the zero temperature vacuum of strongly coupled $\mathcal{N} = 4$ SYM theory is a rich subject with numerous physical insights into the dynamics of gauge theories. Given our goal of gaining insights into the deconfined phase of QCD, we will largely concentrate on the description of strongly coupled $\mathcal{N} = 4$ SYM theory at nonzero temperature, where it describes a strongly coupled non-abelian plasma with $\mathcal{O}(N_c^2)$ degrees of freedom. The vacua of QCD and $\mathcal{N} = 4$ SYM theory have very different properties. However, when we compare $\mathcal{N} = 4$ SYM at $T \neq 0$ with QCD at a temperature above the temperature T_c of the crossover from a hadron gas to quark-gluon plasma, many of the qualitative distinctions disappear or become unimportant. In particular:

1. QCD confines, while $\mathcal{N} = 4$ does not. This is a profound difference in vacuum. But, above its T_c QCD is no longer confining. The fact that its $T = 0$ quasiparticles are hadrons within which quarks are confined is not particularly relevant at temperatures above T_c .
2. In QCD, chiral symmetry is broken by a chiral condensate which sets a scale that is certainly not present in $\mathcal{N} = 4$ SYM theory. However, in QCD above its T_c the chiral condensate melts away and this distinction between the vacua of the two theories also ceases to be relevant.

3. $\mathcal{N} = 4$ is a scale invariant theory while in QCD scale invariance is broken by the confinement scale, the chiral condensate, and just by the running of the coupling constant. Above T_c , we have already dispensed with the first two scales. Also, as we have described in Section 3, QCD thermodynamics is significantly nonconformal just above $T_c \sim 170$ MeV, but at higher temperatures the quark-gluon plasma becomes more and more scale invariant, at least in its thermodynamics. (Thermodynamic quantities converge to their values in the noninteracting limit, due to the running of the coupling towards zero, only at vastly higher temperatures which are far from the reach of any collider experiment.) So, here again, QCD above (but not asymptotically far above) its T_c is much more similar to $\mathcal{N} = 4$ SYM theory at $T \neq 0$ than the vacua of the two theories are.
4. $\mathcal{N} = 4$ SYM theory is supersymmetric. However, supersymmetry is explicitly broken at nonzero temperature. In a thermodynamic context, this can be seen by noting that fermions have antiperiodic boundary conditions along the Euclidean time circle while bosons are periodic. For this reason, supersymmetry does not play a major role in the characterization of properties of the $\mathcal{N} = 4$ SYM plasma at nonzero temperatures.
5. QCD is an asymptotically free theory and, thus, high energy processes are weakly coupled. However, as we have described in Section 2, in the regime of temperatures above T_c that are accessible to heavy ion collision experiments the QCD plasma is strongly coupled, which opens a window of applicability for strong coupling techniques.

For these and other reasons, the strongly coupled plasma of $\mathcal{N} = 4$ SYM theory has been studied by many authors with the aim of gaining insights into the dynamics of deconfined QCD plasma.

In fairness, we should also mention the significant differences between the two theories that remain at nonzero temperature:

1. $\mathcal{N} = 4$ SYM theory with $N_c = 3$ has more degrees of freedom than QCD with $N_c = 3$. To seek guidance for QCD from results in $\mathcal{N} = 4$ SYM, the challenge is to evaluate how an observable of interest depends on the number of degrees of freedom, as we do at several points in Section 7. The best case scenario is that there is no such dependence, as for example arises for the ratio η/s between the shear viscosity and the entropy density that we introduced in Section 2 and that we shall discuss in Section 6.2 below.
2. Most of the calculations that we shall report are done in the strong coupling ($\lambda \rightarrow \infty$) limit. This is of course a feature not a bug. The ability to do these calculations in the strong coupling regime is a key part of the motivation for all this work. But, although in the temperature regime of interest $g^2(T)N_c = 4\pi N_c \alpha_s(T)$ is large, it is not infinite. This motivates the calculation of corrections to various results that we shall discuss that are proportional to powers of $1/\lambda$, for the purpose of testing the robustness of conclusions drawn from calculations done with $\lambda \rightarrow \infty$.
3. QCD has $N_c = 3$ colors, while all the calculations that we shall report are done in the $N_c \rightarrow \infty$ limit. Although the large- N_c approximation is familiar in QCD, the standard way of judging whether it is reliable in a particular context is to compute

corrections suppressed by powers of $1/N_c$. And, determining the $1/N_c^2$ corrections to the calculations done via the gauge/string duality that we review remains an outstanding challenge.

4. Although we have argued above that the distinction between bosons and fermions is not important at nonzero temperature, the distinction between degrees of freedom in the adjoint or fundamental representation of $SU(N_c)$ is important. QCD has $N_f = 3$ flavors in the fundamental representation, namely $N_f = N_c$. These fundamental degrees of freedom contribute significantly to its thermodynamics at temperatures above T_c . And, the calculations that we shall report are either done with $N_f = 0$ or with $0 < N_f \ll N_c$. Extending methods based upon gauge/string duality to the regime in which $N_f \sim N_c$ remains an outstanding challenge.

The plasmas of QCD and strongly coupled $\mathcal{N} = 4$ SYM theory certainly differ. At the least, using one to gain insight into the other follows in the long tradition of modelling, in which a theoretical physicist employs the simplest instance of a theory that captures the essence of a suite of phenomena that are of interest in order to gain insights. $\mathcal{N} = 4$ SYM theory may not seem simple from a field theory perspective, but its gravitational description makes it clear that it is in fact the simplest, most symmetric, strongly coupled non-Abelian plasma. The question then becomes whether there are quantities or phenomena that are universal across many different strongly coupled plasmas. The qualitative, and in some instances even semi-quantitative, successes that we shall review that have been achieved in comparing results or insights obtained in $\mathcal{N} = 4$ SYM theory to those in QCD suggest a positive answer to this question, but no precise definition of this new kind of universality has yet been conjectured. Absent a precise understanding of such a universality, we can hope for reliable insights into QCD but not for controlled calculations.

We begin our description of the $\mathcal{N} = 4$ strongly coupled plasma in this Section by characterizing its macroscopic properties, i.e., those that involve temporal and spatial scales much larger than the microscopic scale $1/T$. In Section 6.1 we briefly review the determination of the thermodynamics of $\mathcal{N} = 4$ SYM theory. The quantities that we calculate are accessible in QCD, via lattice calculations as we have reviewed in Section 3, meaning that in Section 6.1 we will be able to compare calculations done in $\mathcal{N} = 4$ SYM theory via gauge/string duality to reliable information about QCD. In Section 6.2 we turn to transport coefficients like the shear viscosity η , which govern the relaxation of small deviations away from thermodynamic equilibrium. Lattice calculations of such quantities are in their infancy but, as we have seen in Section 2.2, phenomenological analyses of collective effects in heavy ion collisions in comparison to relativistic, viscous, hydrodynamic calculations are yielding information about η/s in QCD. Finally, Section 6.3 will be devoted to illustrating one of the most important qualitative differences between the strongly coupled $\mathcal{N} = 4$ plasma and any weakly coupled plasma: the absence of quasiparticles. As we will argue in this Section, this is a generic feature of strong coupling which, at least at a qualitative level, provides a strong motivation in the context of the physics of QCD above T_c for performing studies within the framework of gauge/string duality.

6.1 Thermodynamic properties

6.1.1 Entropy, energy and free energy

As discussed in Section 5.2.1, $\mathcal{N} = 4$ SYM theory in equilibrium at nonzero temperature is described in the gravity theory by introducing black-branes which change the AdS_5 metric to the black-brane metric (5.31) and lead to the formation of an event-horizon at position z_0 . As in standard black hole physics, the presence of the horizon allows us to compute the entropy in the gravity description, which is given by the Bekenstein-Hawking formula

$$S_{\lambda=\infty} = S_{\text{BH}} = \frac{A_3}{4G_5}, \quad (6.1)$$

where A_3 is the area of the 3-dimensional event horizon of the non-compact part of the metric and G_5 is the five-dimensional Newton constant. This entropy is identified as the entropy in the strong-coupling limit [311]. The area A_3 is determined from the horizon metric, obtained by setting $t = \text{const}$, $z = z_0$ in eqn. (5.31), i.e.

$$ds_{\text{Hor}}^2 = \frac{R^2}{z_0^2} (dx_1^2 + dx_2^2 + dx_3^2). \quad (6.2)$$

The total horizon area is, then,

$$A_3 = \frac{R^3}{z_0^3} \int dx_1 dx_2 dx_3, \quad (6.3)$$

where $\int dx_1 dx_2 dx_3$ is the volume in the gauge theory. While the total entropy is infinite, the entropy density per unit gauge-theory volume is finite and is given by

$$s_{\lambda=\infty} = \frac{S_{\text{BH}}}{\int dx_1 dx_2 dx_3} = \frac{R^3}{4G_5 z_0^3} = \frac{\pi^2}{2} N_c^2 T^3, \quad (6.4)$$

where in the last equality we have used Eqs. (5.12) and (5.33) to translate the gravity parameters z_0 , R and G_5 into the gauge-theory parameters T and N_c . Note that we would have obtained the same result if we had used the full ten-dimensional geometry, which includes the S^5 . In this case the horizon would have been 8-dimensional, of the form $A_8 = A_3 \times S^5$, and the entropy would have taken the form

$$S_{\text{BH}} = \frac{A_8}{4G} = \frac{A_3 V_{S^5}}{4G}, \quad (6.5)$$

which equals (6.1) by virtue of the relation (5.12) between the 10- and the 5-dimensional Newton constants.

Once the entropy density is known, the rest of the thermodynamic potentials are obtained through standard thermodynamic relations. In particular, the pressure P obeys $s = \partial P / \partial T$, and the energy density is given by $\varepsilon = -P + Ts$. Thus we find:

$$\varepsilon_{\lambda=\infty} = \frac{3\pi^2}{8} N_c^2 T^4, \quad P_{\lambda=\infty} = \frac{\pi^2}{8} N_c^2 T^4. \quad (6.6)$$

The N_c and temperature dependence of these results could have been anticipated. The former follows from the fact that the number of degrees of freedom in an $SU(N_c)$ gauge theory in its deconfined phase grows as N_c^2 , whereas the latter follows from dimensional analysis, since the temperature is the only scale in the $\mathcal{N} = 4$ theory. What is remarkable about these results is that they show that the prefactors in front of the N_c and temperature dependence in these thermodynamic quantities attain finite values in the limit of infinite coupling, $\lambda \rightarrow \infty$, which is the limit in which the gravity description becomes strictly applicable.

It is instructive to compare the above expressions at infinite coupling with those for the free $\mathcal{N} = 4$ SYM theory, i.e. at $\lambda = 0$. Since $\mathcal{N} = 4$ SYM has 8 bosonic and 8 fermionic adjoint degrees of freedom and since the contribution of each boson to the free entropy is $2\pi^2 T^3/45$ whereas the contribution of each fermion is $7/8$ of that of a boson, the zero coupling entropy is given by

$$s_{\lambda=0} = \left(8 + 8 \times \frac{7}{8}\right) \frac{2\pi^2}{45} 8(N_c^2 - 1)T^3 \simeq \frac{2\pi^2}{3} N_c^2 T^3, \quad (6.7)$$

where in the last equality we have used the fact that $N_c \gg 1$. As before, the N_c and T dependences are set by general arguments. The only difference between the infinite and zero coupling entropies is an overall numerical factor: comparing Eqs. (6.4) and (6.7) we find [311]

$$\frac{s_{\lambda=\infty}}{s_{\lambda=0}} = \frac{P_{\lambda=\infty}}{P_{\lambda=0}} = \frac{\varepsilon_{\lambda=\infty}}{\varepsilon_{\lambda=0}} = \frac{3}{4}. \quad (6.8)$$

This is a very interesting result: while the coupling of $\mathcal{N} = 4$ changes radically between the two limits, the thermodynamic potentials vary very mildly. This observation is, in fact, not unique to the special case of $\mathcal{N} = 4$ SYM theory, but seems to be a generic phenomenon for field theories with a gravity dual. In fact, in Ref. [386] it was found that for several different classes of theories, each encompassing infinitely many instances, the change in entropy between the infinitely strong and infinitely weak coupling limit is

$$\frac{s_{\text{strong}}}{s_{\text{free}}} = \frac{3}{4}h, \quad (6.9)$$

with h a factor of order one, $\frac{8}{9} \leq h \leq 1.09662$. These explicit calculation strongly suggest that the thermodynamic potentials of non-Abelian gauge-theory plasmas (at least for near-conformal ones) are quite insensitive to the particular value of the gauge coupling. This is particularly striking since, as we will see in Sections 6.2 and 6.3, the transport properties of these gauge theories change dramatically as a function of coupling, going from a nearly ideal gas-like plasma of quasiparticles at weak coupling to a nearly ideal liquid with no quasiparticles at strong coupling. So, we learn an important lesson from the calculations of thermodynamics at strong coupling via gauge/string duality: thermodynamic quantities are not good observables for distinguishing a weakly coupled gas of quasiparticles from a strongly coupled liquid; transport properties and the physical picture of the composition of the plasma are completely different in these two limits, but no thermodynamic quantity changes much.

Returning to the specific case of $\mathcal{N} = 4$ SYM theory, in this case the leading finite- λ correction to (6.8) has been calculated, yielding [387]

$$\frac{s_{\lambda=\infty}}{s_{\lambda=0}} = \frac{P_{\lambda=\infty}}{P_{\lambda=0}} = \frac{\varepsilon_{\lambda=\infty}}{\varepsilon_{\lambda=0}} = \frac{3}{4} + \frac{1.69}{\lambda^{3/2}} + \dots, \quad (6.10)$$

suggesting that this ratio increases from $3/4$ to $7/8$ as λ drops from infinity down to $\lambda \sim 6$, corresponding to $\alpha_{\text{SYM}} \sim 0.5/N_c$. This reminds us that the control parameter for the strong-coupling approximation is $1/\lambda$, meaning that it can be under control down to small values of α_{SYM} . The ratio (6.8) is also expected to receive corrections of order $1/N_c^2$, but these have not been computed.

It is also interesting to compare (6.8) to what we know about QCD thermodynamics from lattice calculations like those described in Section 3.1. The ratio (6.8) has the advantage that the leading dependence on the number of degrees of freedom drops out, making it meaningful to compare directly to QCD. While theories which have been analyzed in Ref. [386] are rather different from QCD, the regularity observed in these theories compel us to evaluate the ratio of the entropy density computed in the lattice calculations to that which would be obtained for free quarks and gluons. Remarkably, Fig. 3.1 shows that for $T = (2 - 3)T_c$, the coefficient defined in (6.9) is $h \simeq 1.07$, which is in the ballpark of what the calculations done via gauge/gravity duality have taught us to expect for a strongly coupled gauge theory. While this observation is interesting, by itself it is not strong evidence that the QCD plasma at these temperatures is strongly coupled. The central lesson is, in fact, that the ratio (6.8) is quite insensitive to the coupling. The proximity of the lattice results to the value for free quarks and gluons should never have been taken as indicating that the quark-gluon plasma at these temperatures is a weakly coupled gas of quasiparticles. And, now that experiments at RHIC that we described in Section 2.2 combined with calculations that we shall review in Section 6.2 have shown us a strongly coupled QCD plasma, the even closer proximity of the lattice results for QCD thermodynamics to that expected for a strongly coupled gauge theory plasma should also not be overinterpreted.

6.1.2 Holographic susceptibilities

The previous discussion focussed on a plasma at zero chemical potential μ . While gauge/gravity duality allows us to explore the phase diagram of the theory at nonzero values of μ , in order to parallel our discussion of QCD thermodynamics in Section 3, in our analysis of strongly coupled $\mathcal{N} = 4$ SYM theory here we will concentrate on the calculation of susceptibilities. As explained in Section 3.1.1, their study requires the introduction of $U(1)$ conserved charges. In $\mathcal{N}=4$ SYM, there is an $SU(4)$ global symmetry, the R -symmetry, which in the dual gravity theory corresponds to rotations in the 5-sphere. Chemical potential is introduced by studying rotating black-holes in these coordinates [329, 331, 388, 389]; these solutions demand non-vanishing values of an abelian vector potential A_μ in the gravitational theory which, in turn, lead to a non-vanishing charge density n in the gauge theory proportional to the angular momentum density of the black hole. The chemical potential can be extracted from the boundary value of the temporal component of the Maxwell field as in (5.34) and is also a function of the angular momentum of the black hole. The explicit calculation performed in Ref. [390] leads to

$$n = \frac{N_c^2 T^2}{8} \mu, \quad (6.11)$$

in the small chemical potential limit. Note that, unlike in QCD, the susceptibility $dn/d\mu$ inferred from Eq. (6.11) is proportional to N_c^2 instead of N_c . This is a trivial consequence of the fact that R -symmetry operates over adjoint degrees of freedom.

As in the case of the entropy, the different number of degrees of freedom can be taken into account by comparing the susceptibility at strong coupling to the free theory result, which yields

$$\frac{\chi_{\lambda=\infty}}{\chi_{\lambda=0}} = \frac{1}{2}, \quad (6.12)$$

where $\chi_{\lambda=0} = N_c^2 T^2/4$ [215]. Similarly to the case of the entropy density, the ratios of susceptibilities between these two extreme limits saturates into an order one constant. In spite of the radical change in the dynamics of the degrees of freedom in the two systems, all the variation in this observable is a 50% reduction, comparable to the 25% reduction of the energy density in the same limit. On the other hand, the lattice calculations in Section 3.1.1 show that above T_c the susceptibilities approach closer to their values in the free limit more quickly than in the case of the energy density. In fact, the rapid approach of the susceptibilities to the free limit has been interpreted by some as a sign that the QCD quark-gluon plasma is not strongly coupled [212, 391–393]. Another possible interpretation of this fact could be that while the gluonic (adjoint) degrees of freedom in the plasma are strongly coupled, the fermionic (fundamental) ones remain quasi-free. This interpretation is based on the observation that for fermions the lowest Matsubara frequency, $\omega_F = \pi T$, is different than that for gluons, $\omega_G = 0$, and thus the thermodynamics of fermions are insensitive to the softer components of the gluon fields. This possible interpretation of the discrepancy between the $\mathcal{N} = 4$ SYM susceptibilities and the QCD susceptibilities would hinge on the fermions in the QCD plasma remaining weakly coupled at scales $\sim \pi T$ while softer gluonic modes become strongly coupled. We note that if this interpretation is taken literally, the quarks in the QCD plasma should increase its η/s which does not seem to be indicated by the elliptic flow data. Nevertheless, it is not clear whether this deviation poses a serious challenge to the interpretation of the QCD plasma as strongly coupled, given the manifest insensitivity of thermodynamic quantities to the coupling. Furthermore, the value of this deviation is not universal and it can be different for other holographic models which are closer to QCD than $\mathcal{N} = 4$ SYM. It is in fact conceivable that if it were possible to use gauge/string duality to analyze strongly coupled theories with $N_f \sim N_c$ and compute the susceptibility for a $U(1)$ charge carried by the fundamental degrees of freedom in such a theory, it would turn out to be close to that at weak coupling as in QCD, even when all degrees of freedom are strongly coupled. Were this speculation to prove correct, it would be an example of a result from QCD leading to insight into strongly coupled gauge theories with a gravitational description, i.e. it would be an example of using the duality in the opposite direction from that in which we apply it throughout most of this review. We shall not speculate further on this discrepancy, simply acknowledging that its interpretation is an open question.

6.2 Transport properties

We now turn to the calculation of the transport coefficients of a strongly coupled plasma with a dual gravitational description, which control how such a plasma responds to small deviations from equilibrium. As we have reviewed in Section 3.2, since the relaxation of these perturbations is intrinsically a real time process, the lattice determination of trans-

port coefficients is very challenging. While initial steps toward determining them in QCD have been taken, definitive results are not in hand. As a consequence, the determination of transport coefficients via gauge/string duality is extremely valuable since it opens up their analysis in a regime which is not tractable otherwise. A remarkable consequence of this analysis, which we review in Section 6.2.2, is a universal relation between the shear viscosity and the entropy density for the plasmas in all strongly coupled large- N_c gauge theories with a gravity dual [394–397]. This finding, together with the comparison of the universal result $\eta/s = 1/(4\pi)$ with values extracted by comparing data on elliptic flow in heavy ion collisions to analyses in terms of viscous hydrodynamics as we have reviewed in Section 2.2, has been one of the most influential results obtained via the gauge/string duality.

6.2.1 A general formula for transport coefficients

The most straightforward way in which transport coefficients can be determined using the gauge/gravity correspondence is via Green-Kubo formulas, see Appendix A, which rely on the analysis of the retarded correlators in the field theory at small four-momentum. The procedure for determining these correlators using the correspondence has been outlined in Section 5.3. In this section we will try to keep our analysis as general as possible so that it can be used for the transport coefficient that describes the relaxation of any conserved current in the theory. In addition, we will not restrict ourselves to the particular form of the metric (5.31) so that our discussion can be applied to any theory with a gravity dual. Our discussion will closely follow the formalism developed in [369], which builds upon earlier analyses in Refs. [394–401].

In general, if the field theory at nonzero temperature is invariant under translations and rotations, the gravitational theory will be described by a $(4 + 1)$ -dimensional metric of the form

$$ds^2 = -g_{tt}dt^2 + g_{zz}dz^2 + g_{xx}\delta_{ij}dx^i dx^j = g_{MN}dx^M dx^N \quad (6.13)$$

with all the metric components solely dependent on z . Since a nonzero temperature is characterized in the dual theory by the presence of an event horizon, we will assume that g_{tt} has a first order zero and g_{zz} has a first order pole at a particular value $z = z_0$.

We are interested in computing the transport coefficient χ associated with some operator \mathcal{O} in this theory, namely

$$\chi = - \lim_{\omega \rightarrow 0} \lim_{\vec{k} \rightarrow 0} \frac{1}{\omega} \text{Im} G_R(\omega, \mathbf{k}). \quad (6.14)$$

(See Appendix A for the exact definition of G_R and for a derivation of this formula.) For concreteness, we assume that the quadratic effective action for the bulk mode ϕ dual to \mathcal{O} has the form of a massless scalar field¹

$$S = -\frac{1}{2} \int d^{d+1}x \sqrt{-g} \frac{1}{q(z)} g^{MN} \partial_M \phi \partial_N \phi, \quad (6.15)$$

where $q(z)$ is a function of z and can be considered a spacetime-dependent coupling constant. As we will see below, Eqs. (6.13) and (6.15) apply to various examples of interest including

¹Note that restricting to a massless mode does not result in much loss of generality, since almost all transport coefficients calculated to date are associated with operators whose gravity duals are massless fields. The only exception is the bulk viscosity.

the shear viscosity and the momentum broadening for the motion of a heavy quark in the plasma. Since transport coefficients are given by the Green-Kubo formula Eq. (6.14), the general expression for the retarded correlator (5.57) with $\Delta = d$ since $m = 0$ leads to

$$\chi = - \lim_{k_\mu \rightarrow 0} \lim_{z \rightarrow 0} \text{Im} \left\{ \frac{\Pi(z, k_\mu)}{\omega \phi(z, k_\mu)} \right\} = - \lim_{k_\mu \rightarrow 0} \lim_{z \rightarrow 0} \frac{\Pi(z, k_\mu)}{i\omega \phi(z, k_\mu)}, \quad (6.16)$$

where Π is the canonical momentum of the field ϕ :

$$\Pi = \frac{\delta S}{\delta \partial_z \phi} = - \frac{\sqrt{-g}}{q(z)} g^{zz} \partial_z \phi. \quad (6.17)$$

The last equality in (6.16) follows from the fact that the real part of $G_R(k)$ vanishes faster than linearly in ω as $k \rightarrow 0$, as is proven by the fact that the final result that we will obtain, Eq. (6.24), is finite and real.

In (6.16) both Π and ϕ must be solutions of the classical equations of motion which, in the Hamiltonian formalism, are given by (6.17) together with

$$\partial_z \Pi = - \frac{\sqrt{-g}}{q(z)} g^{\mu\nu} k_\mu k_\nu \phi. \quad (6.18)$$

The evaluation of χ , following Eq. (6.16), requires the determination of both $\omega \phi$ and Π in the small four momentum $k_\mu \rightarrow 0$ limit. Remarkably, in this limit the equations of motion (6.17) and (6.18) are trivial

$$\partial_z \Pi = 0 + \mathcal{O}(k_\mu \omega \phi), \quad \partial_z(\omega \phi) = 0 + \mathcal{O}(\omega \Pi), \quad (6.19)$$

and both quantities become independent of z , which allows their evaluation at any z . For simplicity, and since the only restriction on the general metric (6.13) is that it possesses a horizon, we will evaluate them at z_0 where the in-falling boundary condition should be imposed. Our assumptions on the metric imply that in the vicinity of the horizon $z \rightarrow z_0$

$$g_{tt} = -c_0(z_0 - z), \quad g_{zz} = \frac{c_z}{z_0 - z}, \quad (6.20)$$

and eliminating Π from (6.17) and (6.18) we find an equation for ϕ given by

$$\sqrt{\frac{c_0}{c_z}}(z_0 - z) \partial_z \left(\sqrt{\frac{c_0}{c_z}}(z_0 - z) \partial_z \phi \right) + \omega^2 \phi = 0. \quad (6.21)$$

The two general solutions for this equation are

$$\phi \propto e^{-i\omega t} (z_0 - z)^{\pm i\omega \sqrt{c_z/c_0}}. \quad (6.22)$$

Imposing in-falling boundary condition implies that we should take the negative sign in the exponent. Therefore, from Eq. (6.22) we find that at the horizon

$$\partial_z \phi = \sqrt{\frac{g_{zz}}{-g_{tt}}} (i\omega \phi), \quad (6.23)$$

and using Eqs. (6.17) and (6.22) we obtain

$$\chi = - \lim_{k_\mu \rightarrow 0} \lim_{z \rightarrow 0} \frac{\Pi(z, k_\mu)}{i\omega\phi(z, k_\mu)} = - \lim_{k_\mu \rightarrow 0} \lim_{z \rightarrow z_0} \frac{\Pi(z, k_\mu)}{i\omega\phi(z, k_\mu)} = \frac{1}{q(z_0)} \sqrt{\frac{-g}{-g_{zz}g_{tt}}} \Big|_{z_0}. \quad (6.24)$$

Note that the last equality in (6.24) can also be written as

$$\chi = \frac{1}{q(z_0)} \frac{A}{V}, \quad (6.25)$$

where A is the area of the horizon and V is the spatial volume of the boundary theory. From our analysis of the thermodynamic properties of the plasma in Section 6, the area of the event horizon is related to the entropy of the boundary theory via

$$s = \frac{A}{V} \frac{1}{4G_N}. \quad (6.26)$$

From this analysis we conclude that for any theory with a gravity dual the ratio of any transport coefficient to the entropy density depends solely on the properties of the dual fields at the horizon,

$$\frac{\chi}{s} = \frac{4G_N}{q(z_0)}. \quad (6.27)$$

In the next subsection we will use this general expression to compute the shear viscosity of the AdS plasma.

Finally, we would like to remark that the above discussion applies to more general effective actions of the form

$$S = -\frac{1}{2} \int \frac{d\omega d^{d-1}k}{(2\pi)^d} dz \sqrt{-g} \left[\frac{g^{zz}(\partial_z\phi)^2}{Q(z; \omega, k)} + P(z; \omega, k)\phi^2 \right], \quad (6.28)$$

provided that the equations of motion (6.19) remain trivial in the zero-momentum limit. This implies that Q should go to a nonzero constant at zero momentum and P must be at least quadratic in momenta. For (6.28) the corresponding transport coefficient χ is given by

$$\chi = \frac{1}{Q(z_0, k_\mu = 0)} \frac{A}{V} \quad \text{and} \quad \frac{\chi}{s} = \frac{4G_N}{Q(z_0, k_\mu = 0)}. \quad (6.29)$$

6.2.2 Universality of the Shear Viscosity

We now apply the result of last subsection to the calculation of the shear viscosity η of a strongly coupled plasma described by the metric (6.13). As in Appendix A we must compute the correlation function of the operator $\mathcal{O} = T_{xy}$, where the coordinates x and y are orthogonal to the momentum vector. The bulk field ϕ dual to \mathcal{O} should have a metric perturbation h_{xy} as its boundary value. It then follows that $\phi = (\delta g)_y^x = g_{zz}h_{xy}$, where δg is the perturbation of the bulk metric. For Einstein gravity in a geometry with no off-diagonal components in the background metric, as in (6.13), a standard analysis of the Einstein equations to linear order in the perturbation upon assuming that the momentum vector is perpendicular to the (x, y)

plane shows that the effective action for ϕ is simply that of a minimally coupled massless scalar field, namely

$$S = -\frac{1}{16\pi G_N} \int d^{d+1}x \sqrt{-g} \left[\frac{1}{2} g^{MN} \partial_M \phi \partial_N \phi \right]. \quad (6.30)$$

The prefactor $1/16\pi G_N$ descends from that of the Einstein-Hilbert action. This action has the form of Eq. (6.15) with

$$q(z) = 16\pi G_N = \text{const}, \quad (6.31)$$

which, together with Eq. (6.27), leads to the celebrated result

$$\frac{\eta_{\lambda=\infty}}{s_{\lambda=\infty}} = \frac{1}{4\pi} \quad (6.32)$$

that was first obtained in 2001 by Policastro, Son and Starinets [398]. In (6.31), we have added the subscript $\lambda = \infty$ to stress that the numerator and denominator are both computed in the strict infinite coupling limit. Remarkably, this ratio converges to a constant at strong coupling. And, this is not only a feature of $\mathcal{N} = 4$ SYM theory because this derivation applies to any gauge theory with a gravity dual given by Einstein gravity coupled to matter fields, since in Einstein gravity the coupling constant for gravity is always given by Eq. (6.31). In this sense, this result is universal [394–397] since it applies in the strong-coupling and large- N_c limits to the large class of theories with a gravity dual, regardless of whether the theories are conformal or not, confining or not, supersymmetric or not and with or without chemical potential. In particular, if large- N_c QCD has a gravity dual, its η/s should also be given by $1/(4\pi)$ up to corrections due to the finiteness of the coupling. Even if large- N_c QCD does not have a gravity dual, Eq. 6.32 may still apply since the universality of this result may be due to generic properties of strongly coupled theories (for example the absence of quasiparticles, see Section 6.3) which may not depend on whether they are dual to a gravitational theory.

The finite-coupling corrections to Eq. (6.32) in $\mathcal{N} = 4$ SYM theory are given by [402–405]

$$\frac{\eta_{\lambda \rightarrow \infty}}{s_{\lambda \rightarrow \infty}} = \frac{1}{4\pi} \left(1 + \frac{15 \zeta(3)}{\lambda^{3/2}} + \dots \right), \quad (6.33)$$

where $\zeta(3) = 1.20\dots$ is Apéry’s constant. While the coefficient of the leading finite- λ correction proportional to $1/\lambda^{3/2}$ is universal when expressed in terms of the parameters R and l_s in the gravity theory, this coefficient is different for different gauge theories [406], since the relation between R and l_s and the gauge theory parameters λ and N_c changes. Thus, expression (6.33) is only valid for $\mathcal{N} = 4$ SYM theory. It is interesting to notice that, according to Eq. (6.33), η/s increases to $\sim 2/(4\pi)$ once λ decreases to $\lambda \sim 7$, meaning $\alpha_{\text{SYM}} \sim 0.5/N_c$. This is the same range of couplings at which the finite coupling corrections (6.10) to thermodynamic quantities become significant. These results together suggest that strongly coupled theories with gravity duals may yield insight into the quark-gluon plasma in QCD even down to apparently rather small values of α_s , at which λ is still large.

To put the result (6.32) into further context, we can compare this strong coupling result to results for the same ratio η/s at weak coupling in both $\mathcal{N} = 4$ SYM theory and QCD. These have been computed at next to leading log accuracy, and take the form

$$\frac{\eta_{\lambda \rightarrow 0}}{s_{\lambda \rightarrow 0}} = \frac{A}{\lambda^2 \log(B/\sqrt{\lambda})}, \quad (6.34)$$

with $A = 6.174$ and $B = 2.36$ in $\mathcal{N} = 4$ SYM theory and $A = 34.8$ (46.1) and $B = 4.67$ (4.17) in QCD with $N_f = 0$ ($N_f = 3$) [407, 408], where we have defined $\lambda = g^2 N_c$ in QCD as in $\mathcal{N} = 4$ SYM theory. Quite unlike the strong coupling result (6.32), these weak-coupling results show a strong dependence on λ , and in fact diverge in the weak-coupling limit. The divergence reflects the fact that a weakly-coupled gauge-theory plasma is a gas of quasiparticles, with strong dissipative effects. In a gas, η/s is proportional to the ratio of the mean-free path of the quasiparticles to their average separation. A large mean-free path, and hence a large η/s , mean that momentum can easily be transported over distances that are long compared to the average spacing between particles. In the $\lambda \rightarrow 0$ limit the mean-free path diverges. The strong $1/\lambda^2$ dependence of η/s can be traced to the fact that the two-particle scattering cross-section is proportional to g^4 . It is reasonable to guess that the λ -dependence of η/s in $\mathcal{N} = 4$ SYM theory is monotonic, increasing from $1/(4\pi)$ as in (6.33) as λ decreases from ∞ and then continuing to increase until it diverges according to (6.34) as $\lambda \rightarrow 0$. The weak-coupling result (6.34) also illustrates a further important point: η/s is not universal for weakly-coupled gauge theory plasmas. The coefficients A and B can vary significantly from one theory to another, depending on their particle content. It is only in the strong coupling limit that universality emerges, with all large- N_c theories with a gravity dual having plasmas with $\eta/s = 1/(4\pi)$. And, we shall see in Section 6.3 that a strongly coupled gauge theory plasma does not have quasiparticles, which makes it less surprising that η/s at strong coupling is independent of the particle content of the theory at weak coupling.

One lesson from the calculations of η/s is that this quantity changes significantly with the coupling constant, going from infinite at zero coupling to $1/(4\pi)$ at strong coupling, at least for large- N_c theories with gravity duals. This is in marked contrast to the behavior of the thermodynamic quantities described in Section 6, which change only by 25% over the same large range of couplings. Thermodynamic observables are insensitive to the coupling, whereas η/s is a much better indicator of the strength of the coupling because it is a measure of whether the plasma is liquid-like or gaseous.

These observations prompt us to revisit the phenomenological extraction of the shear viscosity of quark-gluon plasma in QCD from measurements of elliptic flow in heavy ion collisions, described in Section 2.2. As we saw, the comparison between data and calculations done using relativistic viscous hydrodynamic yields a conservative conclusion that $\eta/s < (3-5)/(4\pi)$, with a current estimate being that η/s seems to lie within the range $(1-2.5)/(4\pi)$ in QCD, in the same ballpark as the strong-coupling result (6.32). And, as we reviewed in Section 3.2, current lattice calculations of η/s in $N_f = 0$ QCD come with caveats but also indicate a value that is in the ballpark of $1/(4\pi)$, likely somewhat above it. Given the sensitivity of η/s to the coupling, these comparisons constitute one of the main lines of evidence that, in the temperature regime accessible at RHIC, the quark-gluon plasma is a strongly coupled fluid. If we were to attempt to extrapolate the weak-coupling result (6.34) for η/s in QCD with $N_f = 3$ to the values of η/s favored by experiment, we would need $\lambda \sim (14-24)$, well beyond the regime of applicability of perturbation theory. (To make this estimate we had to set the log in (6.34) to 1 to avoid negative numbers, which reflects the fact that the perturbative result is being applied outside its regime of validity.)

A central lesson from the strong-coupling calculation of η/s via gauge/string duality, arguably even more significant than the qualitative agreement between the result (6.32) and cur-

rent extractions of η/s from heavy ion collision data, is simply the fact that values of $\eta/s \ll 1$ are possible in non-abelian gauge theories, and in particular in non-abelian gauge theories whose thermodynamic observables are not far from weak-coupling expectations. These calculations, done via gauge/string duality, provided theoretical support for considering a range of small values of η/s that had not been regarded as justified before, and inferences drawn from RHIC data have pushed η/s into this regime. The computation of the shear viscosity that we have just reviewed is one of the most influential results supporting the notion that the application of gauge/gravity duality can yield insights into the phenomenology of hot QCD matter.

It has also been conjectured [396] that the value of η/s in Eq. (6.32) is, in fact, a lower bound for all systems in nature. This conjecture is supported by the finite-coupling corrections shown in Eq. (6.33). And, all substances known in the laboratory satisfy the bound. Among conventional liquids, the lowest η/s is achieved by liquid helium, but it is about an order of magnitude above $1/(4\pi)$; water — after which hydrodynamics is named — has an η/s that is larger still, by about another order of magnitude. The best liquids known in the laboratory are the quark-gluon plasma produced in heavy ion collisions and an ultracold gas of fermionic atoms at the unitary point, at which the s -wave atom-atom scattering length has been dialed to infinity [2], both of which have η/s in the ballpark of $1/(4\pi)$ but, according to current estimates, somewhat larger.

However, in recent years the conjecture that (6.32) is a lower bound on η/s has been questioned and counter-examples have been found among theories with gravity duals. As emphasized in Section 5, Einstein gravity in the dual gravitational description corresponds to the large- λ and large- N_c limit of the boundary gauge theory. When higher order corrections to Einstein gravity are included, which correspond to $1/\sqrt{\lambda}$ or $1/N_c$ corrections in the boundary gauge theory, Eq. (6.32) will no longer be universal. In particular, as pointed out in Refs. [409, 410] and generalized in Refs. [411–419], generic higher derivative corrections to Einstein gravity can violate the proposed bound. Eq. (6.29) indicates that η/s is smaller than Eq. (6.32) if the “effective” gravitational coupling for the h_y^x polarization at the horizon is *stronger* than the universal value (6.31) for Einstein gravity. Gauss-Bonnet gravity as discussed in Refs. [409, 420] is an example in which this occurs. There, the effective action for h_x^y has the form of Eq. (6.28) with the effective coupling $Q(r)$ at the horizon satisfying [409]

$$\frac{1}{Q(r_0)} = \frac{(1 - 4\lambda_{\text{GB}})}{16\pi G_N}, \quad (6.35)$$

leading to

$$\frac{\eta}{s} = \frac{(1 - 4\lambda_{\text{GB}})}{4\pi}, \quad (6.36)$$

where λ_{GB} is the coupling for the Gauss-Bonnet higher-derivative term. Thus, for $\lambda_{\text{GB}} > 0$ the graviton in this theory is more strongly coupled than that of Einstein gravity and the value of η/s is smaller than $1/4\pi$. In Ref. [410], an explicit gauge theory has been proposed whose gravity dual corresponds to $\lambda_{\text{GB}} > 0$. (See Refs. [405, 421] for generalizations.) Despite not being a lower bound, the smallness of η/s , the qualitative agreement between Eq. (6.32) and values obtained from heavy ion collisions, and the universality of the result (6.32) which applies to any gauge theory with a gravity dual in the large- N_c and strong-coupling limits,

are responsible for the great impact that this calculation done via gauge/string duality has had on our understanding of the properties of deconfined QCD matter.

6.2.3 Bulk viscosity

As we have discussed in Section 2.2, while the bulk viscosity ζ is very small in the QCD plasma at temperatures larger than $1.5 - 2 T_c$, with ζ/s much smaller than $1/4\pi$, ζ/s rises in the vicinity of T_c , a feature which can be important for heavy ion collisions. Since the plasma of a conformal theory has zero bulk viscosity, $\mathcal{N} = 4$ SYM theory is not a useful example to study the bulk viscosity of a strongly coupled plasma. However, the bulk viscosity has been calculated both in more sophisticated examples of the gauge/string duality in which the gauge theory is not conformal [422–426], as well as in AdS/QCD models that incorporate an increase in the bulk viscosity near a deconfinement phase transition [51, 427, 428].

We will only briefly review what is possibly the simplest among the first type of examples, the so-called Dp -brane theory. This is a $(p + 1)$ -dimensional cousin of $\mathcal{N} = 4$ SYM, namely a $(p + 1)$ -dimensional SYM theory (with 16 supercharges) living at the boundary of the geometry describing a large number of non-extremal black Dp -branes [355] with $p \neq 3$. The case $p = 3$ is $\mathcal{N} = 4$ SYM, while the cases $p = 2$ and $p = 4$ correspond to non-conformal theories in $(2 + 1)$ - and $(4 + 1)$ -dimensions. We emphasize that we choose this example for its simplicity rather than because it is directly relevant for phenomenology.

The metric sourced by a stack of black Dp -branes can be written as

$$ds^2 = \alpha' \frac{(d_p \tilde{\lambda} z^{3-p})^{\frac{1}{5-p}}}{z^2} \left(-\tilde{f} dt^2 + ds_p^2 + \left(\frac{2}{5-p} \right)^2 \frac{dz^2}{\tilde{f}} + z^2 d\Omega_{8-p}^2 \right), \quad (6.37)$$

where

$$\tilde{\lambda} = g^2 N, \quad \tilde{f} = 1 - \left(\frac{z}{z_0} \right)^{\frac{14-2p}{5-p}}, \quad d_p = 2^{7-2p} \pi^{\frac{9-3p}{2}} \Gamma \left(\frac{7-p}{2} \right) \quad (6.38)$$

and

$$g^2 = (2\pi)^{p-2} g_s \alpha'^{\frac{13-p}{2}} \quad (6.39)$$

is the Yang-Mills coupling constant, which is dimensionful if $\partial \neq 3$. For $p = 2$ and 4 there is also a nontrivial profile for the dilaton field but we shall not give its explicit form here. The metric above is dual to $(p + 1)$ -dimensional SYM theory at finite temperature.

The bulk viscosity can be computed from the dual gravitational theory via the Kubo formula (A.10). However this computation is more complicated in the bulk channel than in the shear channel and we will not reproduce it here. An alternative and simpler way to compute the bulk viscosity is based on the fact that, in the hydrodynamic limit, the sound mode has the following dispersion relation:

$$\omega = c_s q - \frac{i}{\epsilon + p} \left(\frac{d-1}{d-2} \eta + \frac{\zeta}{2} \right) q^2 + \dots, \quad (6.40)$$

with c_s the speed of sound. Thus, ζ contributes to the damping of sound. In the field theory, the dispersion relation for the sound mode can be found by examining the poles of the retarded Green's function for the stress tensor in the sound channel. As discussed

in Section 5.3.1, on the gravity side these poles correspond to normalizable solutions to the equations of motion for metric perturbations. Due to the in-falling boundary conditions at the horizon, see Eq. (5.55), the frequencies of these modes have nonzero and negative imaginary parts and they are therefore called quasi-normal modes. The explicit computation of these modes for the AdS black hole metric can be used to extract the dispersion relation for sound waves in the strongly coupled $\mathcal{N} = 4$ SYM plasma, yielding an alternative derivation of its shear viscosity [429]. The explicit computation of these quasi-normal modes for the metric (6.37) performed in Ref. [425] showed that the sound mode has the dispersion relation

$$\omega = \sqrt{\frac{5-p}{9-p}}q - i \frac{2}{9-p} \frac{q^2}{2\pi T} + \dots \quad (6.41)$$

from which one finds that (after using $\eta/s = 1/(4\pi)$)

$$c_s = \sqrt{\frac{5-p}{9-p}}, \quad \frac{\zeta}{s} = \frac{(3-p)^2}{2\pi p(9-p)}. \quad (6.42)$$

The above expressions imply an interesting relation [424]

$$\frac{\zeta}{\eta} = 2 \left(\frac{1}{p} - c_s^2 \right) = 2 (c_{s,\text{CFT}}^2 - c_s^2), \quad (6.43)$$

where we have used the fact that the sound speed for a CFT in $(p+1)$ -dimension is $c_{s,\text{CFT}} = 1/\sqrt{p}$. This result might not seem surprising since the bulk viscosity of a theory which is close to being conformal can be expanded in powers of $c_{s,\text{CFT}}^2 - c_s^2$, which is a measure of deviation from conformality. The nontrivial result is that even though the Dp-brane gauge theories are *not* close to being conformal, their bulk viscosities are nevertheless linear in $c_{s,\text{CFT}}^2 - c_s^2$. While this is an interesting observation, it is not clear to what extent it is particular to the Dp-brane gauge theories or whether it is more generic.

6.2.4 Relaxation times and other 2nd order transport coefficients

As we have reviewed in Section 2.2.3, transport coefficients correspond to the leading order gradient expansion of an interacting theory which corrects the hydrodynamic description. *A priori*, there is no reason to stop the extraction of these coefficients at first order, and higher order ones can be (and have been) computed using gauge/string duality. Of particular importance is the determination of the five second order coefficients, τ_π , κ , λ_1 , λ_2 , λ_3 defined in Eq. (2.21). Unlike for the first order coefficients, the gravitational computation of these second order coefficients is quite technical and we shall not review it here. We shall only describe the main points and refer the reader to Refs. [90, 91] for details.

The strategy for determining these coefficients is complicated by the fact that the three coefficients λ_i involve only non-linear combinations of the hydrodynamic fields. Thus, even though formulae can be derived for the linear coefficients τ_π and κ [90, 430], the non-linear coefficients cannot be determined from two-point correlators, since these coefficients are invisible in the linear perturbation analysis of the background. Their determination, thus,

demands the small gradient analysis of non-linear solutions to the Einstein equations² as performed in Ref. [91] (see also Ref. [90]) which yields

$$\begin{aligned} \tau_\pi &= \frac{2-\ln 2}{2\pi T}, & \kappa &= \frac{\eta}{\pi T}, \\ \lambda_1 &= \frac{\eta}{2\pi T}, & \lambda_2 &= -\frac{\eta \ln 2}{\pi T}, & \lambda_3 &= 0. \end{aligned} \tag{6.44}$$

These results are valid in the large- N_c and strong-coupling limit. Finite coupling corrections to some of these coefficients can be found in Ref. [431]. Additionally, the first and second order coefficients have been studied in a large class of non-conformal theories with or without flavor in Refs. [432, 433].

To put these results in perspective we will compare them to those extracted in the weakly coupled limit of QCD ($\lambda \ll 1$) [434]. We shall not comment on the values of all the coefficients, since, as discussed in Section 2.2.3, the only one with any impact on current phenomenological applications to heavy ion collisions is the shear relaxation time τ_π . In the weak-coupling limit,

$$\lim_{\lambda \rightarrow 0} \tau_\pi \simeq \frac{5.9}{T} \frac{\eta}{s}, \tag{6.45}$$

where the result is expressed in such a way as to show that τ_π and η have the same leading-order dependence on the coupling λ (up to logarithmic corrections). For comparison, the strong-coupling result from (6.44) may be written as

$$\lim_{\lambda \gg 1} \tau_\pi = \frac{7.2}{T} \frac{\eta}{s}, \tag{6.46}$$

which is remarkably close to (6.45). But, of course, the value of η/s is vastly different in the weak- and strong-coupling limits. On general grounds, one may expect that relaxation and equilibration processes are more efficient in the strong coupling limit, since they rely on the interactions between different modes in the medium. This general expectation is satisfied for the shear relaxation time of the $\mathcal{N} = 4$ SYM plasma, with τ_π diverging at weak-coupling and taking on the small value

$$\lim_{\lambda \rightarrow \infty} \tau_\pi \simeq \frac{0.208}{T}. \tag{6.47}$$

in the strong-coupling limit. For the temperatures $T > 200$ MeV, which are relevant for the quark-gluon plasma produced in heavy ion collisions, this relaxation time is of the order of 0.2 fm/c, which is much smaller than perturbative expectations. We have recalled already at other places in this review that caveats enter if one seeks quantitative guidance for heavy ion phenomenology on the basis of calculations made for $\mathcal{N} = 4$ SYM plasma. However, the qualitative (and even semi-quantitative) impact of the result (6.47) on heavy ion phenomenology should not be underestimated: the computation of τ_π demonstrated for the first time that at least some excitations in a strongly coupled non-abelian plasma dissipate on timescales that are much shorter than $1/T$, i.e. on time scales much shorter than 1 fm/c. Such small values did not have a theoretical underpinning before, and they are clearly relevant for phenomenological studies based on viscous fluid dynamic simulations. As we reviewed in Section 2.2,

² Kubo-like formulas involving three-point correlators (as opposed to two) can also be used to determine the coefficients λ_i [430]. At the time of writing, this approach had not been explored within the gauge/gravity context.

the success of the comparison of such simulations to heavy ion collision data implies that a hydrodynamic description of the matter produced in these collisions is valid only ~ 1 fm/c after the collision. Although this equilibration time is related to out-of-equilibrium dynamics, whereas τ_π is related to near-equilibrium dynamics (only to second order), the smallness of τ_π makes the rapid equilibration time seem less surprising. As in the case of η/s , the gauge/gravity calculation of τ_π has made it legitimate to consider values of an important parameter that had not been considered before by showing that this regime arises in the strongly coupled plasma of a quantum field theory that happens to be accessible to reliable calculation because it possesses a gravity dual.

Let us also mention that the second order transport coefficients are known for the same non-conformal gauge theories whose bulk viscosity we discussed in Section 6.2.3. Since conformal symmetry is broken in these models, there are a total of 15 first- and second-order transport coefficients, 9 more than in the conformal case (including both shear and bulk viscosities in the counting) [92]. In addition, the velocity of sound c_s is a further independent parameter that characterizes the zeroth-order hydrodynamics of non-conformal plasmas, whose equations of state are not given simply by $P = \varepsilon/3$. As for the case of the bulk viscosity, the variable $(\frac{1}{3} - c_s^2)$ can be used to parametrize deviations from conformality, and all transport coefficients can indeed be written explicitly as functions of $(\frac{1}{3} - c_s^2)$ [435].

Let us conclude this section with a curious remark. As we saw in Section 2.2.3, the second order hydrodynamic equations are hyperbolic and, as such, they are causal. However, hydrodynamic equations are a small gradient expansion of the full theory and there is no *a priori* reason why hydrodynamic dispersion relations must provide physically sensible results for short wavelength, highly energetic, excitations. In particular, the group velocity $\frac{d\omega(k)}{dk}$ of a hydrodynamic excitation does not need to remain smaller than the velocity of light in the limit $k \rightarrow \infty$. And yet, the structure of the second order viscous hydrodynamic equations ensures that the group velocities of all hydrodynamic excitations approach finite values in the limit $k \rightarrow \infty$. For example, the hydrodynamic propagation speeds for the shear mode and the sound mode depend on $(\varepsilon + p) = Ts$, the shear and bulk viscosities, as well as the corresponding relaxation times τ_π and τ_Π . In the limit $k \rightarrow \infty$, their group velocities take the simple and general forms

$$\lim_{k \rightarrow \infty} \frac{d\omega_{\text{sound}}(k)}{dk} = \sqrt{\frac{\eta}{\tau_\pi (\varepsilon + p)}}, \quad (6.48)$$

$$\lim_{k \rightarrow \infty} \frac{d\omega_{\text{shear}}(k)}{dk} = \sqrt{c_s^2 + \frac{4}{3} \frac{\eta}{\tau_\pi (\varepsilon + p)} + \frac{\zeta}{\tau_\Pi (\varepsilon + p)}}. \quad (6.49)$$

For a derivation, see Ref. [111]. In principle, one could find a field theory with relaxation times τ_π and/or τ_Π that are so short that the limit $\lim_{k \rightarrow \infty} \frac{d\omega(k)}{dk}$ of some mode exceeds the velocity of light; in this case one would simply state that this limit is outside the range of validity of hydrodynamics. But, for these two modes and for all hydrodynamic modes studied so far, and for all field theories studied so far, when the values of all quantities are plugged into these expressions, the resulting $\lim_{k \rightarrow \infty} \frac{d\omega(k)}{dk}$ describes propagation within the forward light cone. We emphasize that the current understanding of this finding admits the possibility that it is accidental.

6.3 Quasiparticles and spectral functions

In Sections 6.1 and 6.2 we have illustrated the power of gauge/string duality by performing, in a remarkably simple way, computations that via standard field theoretical methods either take Teraflop-years of computer time or are not accessible. However, the simplicity of the calculations comes with a price. Because we do the calculations in the dual gravitational description of the theory, all we get are reliable results; we do not get the kind of intuition of what is happening in the gauge theory that we would get automatically from a field theory calculation done with Feynman diagrams or could get with effort from one done on the lattice. The gravitational calculation yields answers, and new kinds of intuition, but since by using it we are abandoning the description of the plasma in terms of quark and gluon quasi-particles interacting with each other, we are losing our prior physical intuition about how the dynamics of the gauge theory works. In particular, from the results that we extract alone it is difficult to understand whether the dynamics within a strongly coupled plasma differs in a qualitative way from those in a weakly coupled plasma, or merely differs quantitatively. We have given up the description in terms of quasiparticles, but maybe the familiar quasiparticles or some new kind of quasiparticles are in fact nevertheless present. We rule out this possibility in this Section, illustrating that a strongly coupled non-abelian gauge-theory plasma really is qualitatively different than a weakly coupled one: while in perturbation theory the degrees of freedom of the plasma are long-lived quasiparticle excitations which carry momentum, color and flavor, there are no quasiparticles in the strongly coupled plasma. The pictures that we are used to using that frame how we think about a weakly coupled plasma are simply invalid for the strongly coupled case.

Determining whether a theory possesses quasiparticles with a given set of quantum numbers is a conceptually well defined task: it suffices to analyze the spectral function of operators with that set of quantum numbers and look for narrow peaks in momentum space. In weakly coupled Yang-Mills theories, the quasiparticles (gluons and quarks in QCD) are colored and are identified by studying operators that are *not* gauge invariant. Within perturbation theory, it can be shown that the poles of these correlators, which determine the physical properties of the quasiparticles, are gauge invariant [436]. However, non-perturbative gauge-invariant operators corresponding to these excitations are not known, which complicates the search for these quasiparticles at strong coupling. Note, however, that even if such operators were known, demonstrating the absence of quasiparticles with the same quantum numbers as in the perturbative limit does not guarantee the absence of quasiparticles, since at strong coupling the system could reorganize itself into quasiparticles with different quantum numbers. Thus, proving the absence of quasiparticles along these lines would require exploring all possible spectral functions in the theory. Fortunately, there is an indirect method which can answer the question of whether any quasiparticles that carry some conserved ‘charge’ (including momentum) exist, although this method cannot determine the quantum numbers of the long-lived excitations if any are found to exist. The method involves the analysis of the small frequency structure of the spectral functions of those conserved currents of the theory which do not describe a propagating hydrodynamic mode like sound. As we will see, the presence of quasiparticles leads to a narrow structure (the transport peak) in these spectral functions [215,216]. In what follows we will use this method to demonstrate that the strongly coupled $\mathcal{N} = 4$ SYM plasma does not possess any quasiparticles that carry momentum. In

order to understand how the method works, we first apply it at weak coupling where there are quasiparticles to find.

6.3.1 Quasiparticles in perturbation theory

We start our analysis by using kinetic theory to predict the general features of the low-frequency structure of correlators of conserved currents in a weakly coupled plasma. Kinetic theory is governed by the Boltzmann equation, which describes excitations of a quasiparticle system at scales which are long compared to the inter-particle separation. The applicability of the kinetic description demands that there is a separation of scales such that the duration of interactions among particles is short compared to their mean free path (λ_{mfp}) and that multiparticle distributions are consequently determined by the single particle distributions. In Yang-Mills theories at nonzero temperature and weak coupling, kinetic theory is important since it coincides with the Hard Thermal Loop description [437–442], which is the effective field theory for physics at momentum scales of order gT , and the Boltzmann equation can be derived from first principles [442–447]. In Yang-Mills theory at weak coupling and nonzero temperature, the necessary separation of scales arises by virtue of the small coupling constant g , since $\lambda_{\text{mfp}} \sim 1/(g^4 T)$ and the time-scale of interactions is $1/\mu_D \sim 1/(gT)$, where $1/\mu_D$ is the Debye screening length of the plasma.³ The small value of the coupling constant also leads to the factorization of higher-point correlation functions.

In the kinetic description, the system is characterized by a distribution function

$$f(x, \mathbf{p}), \tag{6.50}$$

which determines the number of particles of momentum \mathbf{p} at space-time position x . Note that this position should be understood as the center of a region in space-time with a typical size much larger, at least, than the de-Broglie wave length of the particles, as demanded by the uncertainty principle. As a consequence, the Fourier transform of x , which we shall denote by $K = (\omega, \mathbf{q})$, must be much smaller than the typical momentum scale of the particles, $K \ll p \sim T$. (Here and below, when we write a criterion like $K \ll p$ we mean that both ω and $|\mathbf{q}|$ must be $\ll |\mathbf{p}|$.) Due to this separation in momentum scales, the x -dependence of the distribution functions is said to describe the soft modes of the gauge theory while the momenta \mathbf{p} are those of the hard modes. If K is sufficiently small (smaller than the inverse inter-particle separation $\sim T$), the mode describes collective excitations which involve the motion of many particles, while \mathbf{p} is the momentum of those particles. In this case, the Fourier-transformed distribution $f(K, \mathbf{p})$ can be interpreted approximately as the number of particles within the wavelength of the excitation. At the long distances at which the kinetic-theory description is valid, particles are on mass shell, as determined by the position of the peaks in the correlation functions of the relevant operators ($p^0 = E_{\mathbf{p}}$), and these hard modes describe particles that follow classical trajectories, at least between the microscopic collisions. All the properties of the system can be extracted from the distribution function.

³Strictly speaking, $\lambda_{\text{mfp}} \sim 1/(g^4 T)$ is the length-scale over which an order 1 change of the momentum-vector of the quasiparticles occurs. Over the shorter length-scale $1/\mu_D$, soft exchanges (of order gT ; not enough to change the momenta which are $\sim T$ significantly) occur. These soft exchanges are not relevant for transport.

In particular, the stress tensor is given by

$$T^{\mu\nu}(x) = \int \frac{d^3\mathbf{p}}{(2\pi)^3} \frac{p^\mu p^\nu}{E} f(x, \mathbf{p}). \quad (6.51)$$

Since all quasiparticles carry energy and momentum, we will concentrate only on the kinetic theory description of stress-tensor correlators. Our analysis is analogous to the one performed for the determination of the Green-Kubo formulae in Appendix A, and proceeds by studying the response of the system to small metric fluctuations. The dynamics are, then, governed by the Boltzmann equation which states the continuity of the distribution function f up to particle collisions [448]:

$$E \frac{d}{dt} f(x, \mathbf{p}) = p^\mu \partial_{x^\mu} f(x, \mathbf{p}) + E_p \frac{d\mathbf{p}}{dt} \frac{\partial}{\partial \mathbf{p}} f(x, \mathbf{p}) = \mathcal{C}[f], \quad (6.52)$$

where $\mathcal{C}[f]$ is the collision term which encodes the microscopic collisions among the plasma constituents and vanishes for the equilibrium distribution $f_{\text{eq}}(E_p)$ (which does not depend on x and which does not depend on the direction of \mathbf{p}). In writing (6.52), we are assuming that $\mathbf{p} = E \mathbf{v}_p$, where \mathbf{v}_p is the velocity of the particle. In curved space, in the absence of external forces, the Boltzmann equation becomes

$$p^\mu \partial_{x^\mu} f(x, \mathbf{p}) - \Gamma_{\mu\nu}^\lambda p^\mu p^\nu \partial_{p^\lambda} f(x, \mathbf{p}) = \mathcal{C}[f], \quad (6.53)$$

where $\Gamma_{\mu\nu}^\lambda$ are the Christoffel symbols of the background metric. As in Appendix A, we shall determine the stress-tensor correlator by introducing a perturbation in which the metric deviates from flat space by a small amount, $g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}$, and studying the response of the system. Even though the analysis for a generic perturbation can be performed, it will suffice for our purposes to restrict ourselves to fluctuations which in Fourier space have only one non-vanishing component $h_{xy}(K)$. We choose the directions x and y perpendicular to the wave-vector \mathbf{q} , which lies in the z direction. For this metric, the only Christoffel symbols that are non-vanishing at leading order in h_{xy} are $\Gamma_{xy}^t = \Gamma_{ty}^x = \Gamma_{tx}^y = -i\omega h_{xy}/2$ and $\Gamma_{zy}^x = \Gamma_{zy}^y = -\Gamma_{xy}^z = iq h_{xy}/2$.

We will assume that prior to the perturbation the system is in equilibrium. In response to the external disturbance the equilibrium distribution changes

$$f(x, \mathbf{p}) = f_{\text{eq}}(E_p) + \delta f(x, \mathbf{p}). \quad (6.54)$$

In the limit of a small perturbation, the modified distribution function $\delta f(x, \mathbf{p})$ is linear in the perturbation h_{xy} . We will also assume that the theory is rotationally invariant so that the energy of the particle E_p is only a function of the modulus of $p^2 = g_{ij} p^i p^j$. As a consequence, the metric perturbation also changes the on-shell relation, and the equilibrium distribution must also be expanded to first order in the perturbation, yielding

$$f_{\text{eq}} = f_0 + f'_0 p^x p^y \frac{|v_p|}{p} h_{xy} \approx f_0 + f'_0 \frac{p^x p^y}{E} h_{xy}, \quad (6.55)$$

where f_0 is the equilibrium distribution in flat space, $f'_0(E) = df_0(E)/dE$, and the velocity is given by $v_p = dE_p/dp$. In the last equality we have again approximated $v_p \approx p/E$.

The solution of the Boltzmann equation requires the computation of the collision term \mathcal{C} . In general this is a very complicated task since it takes into account the interactions among all the system constituents, which are responsible for maintaining equilibrium. However, since our only goal is to understand generic features of the spectral function, it will be sufficient to employ the relaxation time approximation

$$\mathcal{C} = -E \frac{f - f_{\text{eq}}}{\tau_R} \quad (6.56)$$

for the collision term, in which the parameter τ_R is referred to as the relaxation time.⁴ Since small perturbations away from equilibrium are driven back to equilibrium by particle collisions, the relaxation time must be of the order of the mean free path λ_{mfp} (which is long compared to the interparticle distance). The relaxation-time approximation is a very significant simplification of the full dynamics, but it will allow us to illustrate the main points that we wish to make. A complete analysis of the collision term within perturbation theory for the purpose of extracting the transport coefficients of a weakly coupled plasma can be found in Refs. [407, 434, 449, 450].

Within the approximation (6.56), upon taking into account that the distribution function does not depend independently on the energy of the particles in Eq. (6.53), the solution to the linearized Boltzmann equation is given by

$$\delta f(K, \mathbf{p}) = \frac{-i\omega p^x p^y f'_0(p)}{-i\omega + i\mathbf{v}_p \mathbf{q} + \frac{1}{\tau_R}} \frac{h_{xy}(K)}{E}. \quad (6.57)$$

Substituting this into Eq. (6.51) we learn that the perturbation of the distribution function leads to a perturbation of the stress tensor given by

$$\delta T^{\mu\nu}(K) = \int \frac{d^3p}{(2\pi)^3} \frac{p^\mu p^\nu}{E} \delta f(K, \mathbf{p}) = -G_R^{xy,xy}(K) h_{xy}(K), \quad (6.58)$$

where the retarded correlator is given by

$$G_R^{xy,xy}(K) = - \int \frac{d^3p}{(2\pi)^3} v^x v^y \frac{\omega p^x p^y f'_0(p)}{\omega - \mathbf{q}\mathbf{v}_p + \frac{i}{\tau_R}}. \quad (6.59)$$

From the definition (3.13), the spectral function associated with this correlator is

$$\rho^{xy,xy}(K) = -\omega \int \frac{d^3p}{(2\pi)^3} \frac{(p^x p^y)^2}{E^2} f'_0(p) \frac{\frac{2}{\tau_R}}{(\omega - \mathbf{q}\mathbf{v}_p)^2 + \frac{1}{\tau_R^2}}. \quad (6.60)$$

Obtaining this spectral function was our goal, because as we shall now see it has qualitative features that indicate the presence (in this weakly coupled plasma) of quasiparticles.

⁴In this approximation, this relaxation time coincides with the shear relaxation time: $\tau_R = \tau_\pi$ [97]. However, since τ_π is a property of the theory itself (defined as the appropriate coefficient in the effective field theory, aka the hydrodynamic expansion) whereas τ_R is a parameter specifying a simplified approximation to the collision kernel, which in general is not of the form eq. (6.56), we will maintain the notational distinction between τ_π and τ_R .

To clarify the structure of the spectral function (6.60) we begin by describing the free theory limit, in which $\tau_R \rightarrow \infty$ since the collision term vanishes. In this limit, the Lorentzian may be replaced by a δ -function, yielding

$$\rho^{xy,xy}(K) = -\omega \int \frac{d^3p}{(2\pi)^3} \frac{(p^x p^y)^2}{E^2} f'_0(p) 2\pi \delta(\omega - \mathbf{q} \cdot \mathbf{v}_p) . \quad (6.61)$$

The δ -function arises because, in this limit, the external perturbation (the gravity wave) interacts with free particles. The δ -function selects those components of the gravity wave whose phase velocity $\omega/|\mathbf{q}|$ coincides with the velocity of any of the free particles in the plasma. Thus, in the free-theory limit, this δ -function encodes the existence of free particles in the plasma. For an isotropic distribution of particles, such as the thermal distribution, at any $\mathbf{q} \neq 0$ the integration over angles washes out the δ -function and one is left with some function of ω that is characterized by the typical momentum scale of the particles ($\sim T$) and that is not of interest to us here.⁵ On the other hand, at $\mathbf{q} = 0$ we find that $\frac{1}{\omega} \rho^{xy,xy}(\omega, 0)$ is proportional to $\delta(\omega)$. This δ -function at $\omega = 0$ in the low-momentum spectral function is a direct consequence of the presence of free particles in the plasma. As we now discuss, the effect of weak interactions is to dress the particles into quasiparticles and to broaden the δ -function into a narrow, tall, peak at $\omega = 0$.

When the interactions do not vanish, we can proceed by relating the relaxation time to the shear viscosity. To do so, we work in the hydrodynamic limit in which all momenta must be smaller than any internal scale. This means that we can set \mathbf{q} to zero, but we must keep the relaxation time τ_R finite. The spectral density at zero momentum is then given by

$$\rho^{xy,xy}(\omega, 0) = -\omega \int \frac{d^3p}{(2\pi)^3} \frac{(p^x p^y)^2}{E^2} f'_0(p) \frac{\frac{2}{\tau_R}}{\omega^2 + \frac{1}{\tau_R^2}} . \quad (6.62)$$

Note that the spectral density at zero momentum has a peak at $\omega = 0$, and note in particular that the width in ω of this peak is $\sim 1/\tau_R \ll T$. The spectral density has vanishing strength for $\omega \gg 1/\tau_R$. This low-frequency structure in the zero-momentum spectral function is called the ‘transport peak’. It is clear that in the $\tau_R \rightarrow \infty$ limit it becomes the δ -function that characterizes the spectral density of the free theory that we described above. Here, in the presence of weak interactions, this peak at $\omega = 0$ is a direct consequence of the presence of momentum-carrying quasiparticles whose mean free time is $\sim \tau_R$.

The expression (6.62) is only valid for $\omega \ll T$ where the modes are correctly described by the Boltzmann equation. For $\omega \gg T$, since the quasiparticles can be resolved the structure of the spectral density is close to that in vacuum. The separation of scales in the spectral density is directly inherited from the separation of scales which allows the Boltzmann description. Finally, using the Green-Kubo formula for the shear viscosity (A.9), we find

$$\eta = -\tau_R \int \frac{d^3p}{(2\pi)^3} \frac{(p^x p^y)^2}{E^2} f'_0(p) . \quad (6.63)$$

⁵A distinct peak at in the spectral density at some $\omega \neq 0$ could be observed if the initial distribution were very anisotropic. This can arise if the theory has a (gauged) conserved charge and if the system is analyzed in the presence of a constant force that acts on this charge — i.e. an electric field.




Figure 6.1: Left: Sketch of the spectral function at zero momentum as a function of frequency for a weakly coupled plasma, as obtained from kinetic theory. The narrow structure at small frequency is the transport peak with a width $1/\tau_R$ that is suppressed by the coupling ($1/\tau_R \sim g^4 T$). The thermal mass is $m_{\text{th}} \propto gT$. Right: Spectral function for the shear channel in the strongly coupled plasma of $\mathcal{N} = 4$ SYM theory computed via gauge/string duality [215] (solid) and a comparison with the vacuum spectral function (dashed) which it approaches at high frequencies. The vertical axis of this figure has been scaled by the shear viscosity $\eta = s/4\pi$ of the strongly coupled plasma. Note that the definition of $\rho^{xyxy} = -\text{Im}G_R/\pi$ used in Ref. [215] is different from that in Eq. (3.13) by a factor of $\pi/2$.

Thus, since η is determined by the collisions among the quasiparticles, we can understand $1/\tau_R$ as the width that arises because the quasiparticles do not have well-defined momenta due to the collisions among them. In particular, in perturbation theory [407, 449]

$$\frac{1}{\tau_R} \sim \frac{1}{T} g^4 \ln \frac{1}{g} \sim \frac{1}{\lambda_{\text{mfp}}} . \quad (6.64)$$

Let us summarize the main points. The zero-momentum spectral densities of a plasma with quasiparticles have a completely distinctive structure: there is a separation of scales between the scale T (the typical momentum of the quasiparticles in the plasma) and the much lower scale $1/\lambda_{\text{mfp}}$. In particular, there is a narrow peak in $\rho(\omega, 0)/\omega$ around $\omega = 0$ of width $\tau_R \sim 1/\lambda_{\text{mfp}}$ and height 2η . At larger frequencies, the strength of the spectral function is very small. At the scale of the mass of the quasiparticles, the spectral function grows again. For massless particles or those with mass much smaller than any temperature-related scale, the role of the mass threshold is played by the thermal mass of the particles, gT , which is much higher than the scale $1/\lambda_{\text{mfp}} \sim g^4 T$ associated with the mean free path due to the weakness of the coupling. Finally, above the scale T the structure of the spectral function approaches what it would be in vacuum. A sketch of this behavior can be found in Fig. 6.1. These qualitative features are independent of any details of the theory, and do not even depend on its symmetries. All that matters is the existence of momentum-carrying quasiparticles. In the presence of quasiparticles, no matter what their quantum numbers are, these qualitative features must be present in the spectral density.

6.3.2 Absence of quasiparticles at strong coupling

We return now to the strongly coupled $\mathcal{N} = 4$ SYM plasma, with its dual gravitational description, in order to compare the expectation (6.62) for how the spectral density should look if the plasma contains any momentum-carrying quasiparticles to an explicit computation of the retarded correlator at strong coupling, of course done via gauge/string duality. In this subsection we will benefit from the general analyses of Sections 6.2.1 and 6.2.2. As in the kinetic-theory computation, we study the response to a metric fluctuation $h_{xy}(\omega, \mathbf{q})$ in the boundary theory, with the same conventions as before. As in Section 6.2.2, the fluctuation in the boundary leads to a metric perturbation in the bulk, $h_x^y = g^{yy}h_{yx}$ of the form

$$h_x^y(\omega, q, z) = \phi(\omega, q, z)e^{-i\omega t + iqz}. \quad (6.65)$$

The field ϕ is governed by the classical action (6.30) which yields an equation of motion for $\phi(\omega, q, z)$ that is given by

$$\phi''(\omega, q, u) - \frac{1+u^2}{uf}\phi'(\omega, q, u) + \frac{\mathfrak{w}^2 - \mathfrak{q}^2(1-u^2)}{uf^2}\phi(\omega, q, u) = 0, \quad (6.66)$$

where $u = z^2$, $\mathfrak{w} = \omega/(2\pi T)$ and $\mathfrak{q} = q/(2\pi T)$. We may now use the general program outlined in Section 5.3.2 to determine the retarded correlator. It is given by Eq. (5.59) which, together with Eqs. (6.17) and (6.31), leads to

$$G_R^{xy,xy} = -\lim_{u \rightarrow 0} \frac{1}{16\pi G_N} \frac{\sqrt{-g}g^{uu}\partial_u\phi(\mathfrak{w}, \mathfrak{q}, u)}{\phi_0(\mathfrak{w}, \mathfrak{q}, u)}, \quad (6.67)$$

where $\phi(\mathfrak{w}, \mathfrak{q}, u)$ is the solution to the equation of motion (6.66) with in-falling boundary conditions at the horizon. For arbitrary values of \mathfrak{w} and \mathfrak{q} , Eq. (6.66) must be solved numerically [215, 451]. From the correlator (6.67), the spectral function is evaluated using the definition (3.13). The result of this computation at zero spatial momentum $\mathfrak{q} = 0$ is shown in the right panel of Fig. 6.1, where we have plotted ρ/ω which should have a peak at $\omega = 0$ if there are any quasiparticles present.

In stark contrast to the kinetic-theory expectation, there is no transport peak in the spectral function at strong coupling. In fact, the spectral function has no interesting structure at all at small frequencies. The numerical computation whose results are plotted in the right panel of Fig. 6.1 also shows that there is no separation of scales in the spectral function. In the strong-coupling calculation, quite unlike in perturbation theory, the small and large frequency behaviors join smoothly and the spectral density is only a function of $\mathfrak{w} = \omega/2\pi T$. This could perhaps have been expected in a conformal theory with no small coupling constant, but note that a free massless theory is conformal and that theory does have a δ -function peak in its spectral function at zero frequency. So, having the explicit computation that gauge/string duality provides is necessary to give us confidence in the result that there is no transport peak in the strongly coupled plasma. The absence of the transport peak shows unambiguously that there are no momentum-carrying quasiparticles in the strongly coupled plasma. Thus, the physical picture of the system is completely different from that in perturbation theory.

To conclude this section we would like to argue that the absence of quasiparticles is a generic property of strong coupling and is not specific to any particular theory with any

particular symmetries or matter content. To do so, let us recall that in the kinetic theory calculation the separation of scales required for its consistency are a consequence of the weak coupling; this is so in perturbative QCD or in perturbative $\mathcal{N} = 4$ or in any weakly coupled plasma. Now, imagine increasing the coupling. According to kinetic theory, independent of the symmetries or the matter content of the theory, the width of the transport peak grows and its height decreases as the coupling increases. This reflects the fact that as the coupling grows so does the width of the quasiparticles. Extrapolating this trend to larger and larger couplings leads to the disappearance of the transport peak which, at a qualitative level, agrees nicely with the strong-coupling result for the $\mathcal{N} = 4$ SYM plasma obtained by explicit computation and shown in the right panel of Fig. 6.1. This observation is one of the most salient motivations for the phenomenological applications of AdS-based techniques. Since, as we have argued extensively in Sections 2.2 and 3 and 6.2, the quark-gluon plasma of QCD at temperatures a few times its T_c is strongly coupled, the quasiparticle picture that has conventionally been used to think about its dynamics is unlikely to be valid in this regime. This makes it very important to have new techniques at our disposal that allow us to study strongly coupled plasmas with no quasiparticles, seeking generic consequences of the absence of quasiparticles for physical observables. Gauge/gravity duality is an excellent tool for these purposes, as we have already seen in Sections 6.1 and 6.2 and as we will further see in the remaining Sections of this review. Indeed, as we use gauge/gravity duality to calculate more, and more different, physical observables we will discover that the calculations done in the dual gravitational description begin to yield a new form of physical intuition, phrased in the dual language rather than in the gauge theory language, in addition to yielding reliable results.

Chapter 7

Probing strongly coupled plasma

As discussed in Sections 2.3 and 2.4, two of the most informative probes of strongly coupled plasma that are available in heavy ion collisions are the rare highly energetic partons and quarkonium mesons produced in these collisions. In this Section and in Section 8, we review results obtained by employing the AdS/CFT correspondence that are shedding light on these classes of phenomena. In Sections 7.1 and 7.2, we review how a test quark of mass M moving through the strongly coupled $\mathcal{N} = 4$ SYM plasma loses energy and picks up transverse momentum. In Section 7.3 we consider how the strongly coupled plasma responds to the hard parton plowing through it; that is, we describe the excitations of the medium which result. In Section 7.4, we review a calculation of the stopping distance of a single light quark moving through the strongly coupled plasma. Throughout Sections 7.1, 7.2, 7.3 and 7.4 we assume that all aspects of the phenomena associated with an energetic parton moving through the plasma are strongly coupled. In Section 7.5, we review an alternative approach in which we assume that QCD is weakly coupled at the energy and momentum scales that characterize gluons radiated from the energetic parton, while the medium through which the energetic parton and the radiated gluons propagate is strongly coupled. In this case, one uses the AdS/CFT correspondence only in the calculation of those properties of the strongly coupled plasma that arise in the calculation of radiative parton energy loss and transverse momentum broadening. In Section 7.6, we describe a calculation of synchrotron radiation in strongly coupled $\mathcal{N} = 4$ SYM theory that allows the construction of a narrowly collimated beam of gluons (and adjoint scalars), opening a new path toward analyzing jet quenching.

In Section 7.7, we review those insights into the physics of quarkonium mesons in heavy ion collisions that have been obtained via AdS/CFT calculations of the temperature-dependent screening of the potential between a heavy quark and antiquark. To go farther, we need to introduce a holographic description of quarkonium-like mesons themselves. In Section 8, we first review this construction and then review the insights that it has yielded. In addition to shedding light upon the physics of quarkonia in hot matter that we have introduced in Section 2.4, as we review in Section 8.6.2 these calculations have also resulted in the discovery of a new and significant process by which a hard parton propagating through a strongly coupled plasma can lose energy: Cherenkov radiation of quarkonium mesons.

7.1 Parton energy loss via a drag on heavy quarks

When a heavy quark moves through the strongly coupled plasma of a conformal theory, it feels a drag force and consequently loses energy [452, 453]. We shall review the original calculation of this drag force in $\mathcal{N} = 4$ SYM theory [452, 453]; it has subsequently been done in many other gauge theories with dual gravitational descriptions [454–464]. In calculations of the drag on heavy quarks, one determines the energy per unit time needed to maintain the forced motion of the quark in the plasma. In these calculations one regards the quark as an external source moving at fixed velocity, v , and one performs thermal averages over the medium. This picture can be justified if the mass of the quark is assumed to be much larger than the typical momentum scale of the medium (temperature), and if the motion of the quark is studied in a time window that is large compared with the relaxation scale of the medium but short compared to the time it takes the quark to change its trajectory. In this limit the heavy quark is described by a Wilson line along the worldline of the quark.

The dual description of the Wilson line is given by a classical string hanging down from the quark on the boundary of AdS. Since we are considering a single quark, the other end of the string hangs down into the bulk of the AdS space. We consider the stationary situation, in which the quark has been moving at a fixed velocity for a long time, meaning that the shape of the string trailing down and behind it is no longer changing with time. For concreteness, we will assume that the quark moves in the x_1 direction, and we choose to parametrize the string world sheet by $\tau = t$ and $\sigma = z$. By symmetry, we can set two of the perpendicular coordinates, x_2 and x_3 to a constant. The problem of finding the string profile reduces, then, to finding a function

$$x_1(\tau, \sigma), \quad (7.1)$$

that fulfills the string equations of motion. The string solution must also satisfy the boundary condition

$$x_1(t, z \rightarrow 0) = vt. \quad (7.2)$$

Since we are interested in the stationary situation, the string solution takes the form

$$x_1(t, z) = vt + \zeta(z), \quad (7.3)$$

with $\zeta(z \rightarrow 0) = 0$. We work in an $\mathcal{N} = 4$ plasma, whose dual gravitational description is the AdS black hole with the metric $G_{\mu\nu}$ given in (5.31). The induced metric on the string worldsheet $g_{\alpha\beta} = G_{\mu\nu} \partial_\alpha x^\mu \partial_\beta x^\nu$ is then given by

$$ds_{ws}^2 = \frac{R^2}{z^2} \left(- (f(z) - v^2) d\tau^2 + \left(\frac{1}{f(z)} + \zeta'^2(z) \right) d\sigma^2 + v \zeta'(z) v (d\tau d\sigma + d\sigma d\tau) \right), \quad (7.4)$$

where, as before, $f(z) = 1 - z^4/z_0^4$ and $\zeta'(z)$ denotes differentiation with respect to z .

The Nambu-Goto action for this string reads

$$S = -\frac{R^2}{2\pi\alpha'} \mathcal{T} \int \frac{dz}{z^2} \sqrt{\frac{f(z) - v^2 + f(z)^2 \zeta'^2(z)}{f(z)}} = \mathcal{T} \int dz \mathcal{L}, \quad (7.5)$$




Figure 7.1: String solutions of eq. (7.8). The physical (unphysical) solution in which momentum flows into (out of) the horizon and the string trails behind (curves ahead) of the quark at the boundary is plotted in the left (right) panel.

with \mathcal{T} the total time traveled by the quark. Extremizing this action yields the equations of motion that must be satisfied by $\zeta(z)$. The action (7.5) has a constant of motion given by the canonical momentum

$$\Pi_z^1 = \frac{\partial \mathcal{L}}{\partial x'_1} = -\frac{R^2}{2\pi\alpha'} \frac{1}{z^2} \frac{f(z)^{3/2} \zeta'(z)}{\sqrt{f(z) - v^2 + f(z)^2 \zeta'^2(z)}}, \quad (7.6)$$

which coincides with the longitudinal momentum flux in the z direction. In terms of Π_z^1 , the equation of motion for ζ obtained from (7.5) takes the form

$$\zeta'^2(z) = \left(\frac{2\pi\alpha'}{R^2} \Pi_z^1 \right)^2 \frac{z^4}{f(z)^2} \frac{f(z) - v^2}{f(z) - \left(\frac{2\pi\alpha'}{R^2} \Pi_z^1 \right)^2 z^4}. \quad (7.7)$$

The value of Π_z^1 can be fixed by inspection of this equation, as follows: both the numerator and the denominator of (7.7) are positive at the boundary $z = 0$ and negative at the horizon $z = z_0$; since $\zeta'(z)$ is real, both the numerator and the denominator must change sign at the same z ; this is only the case if

$$\Pi_z^1 = \pm \frac{R^2}{2\pi\alpha' z_0^2} \gamma v, \quad (7.8)$$

with $\gamma = 1/\sqrt{1 - v^2}$ the Lorentz γ factor. Thus, stationary solutions can only be found for these values of the momentum flux. (Or, for $\Pi_z^1 = 0$, for which $\zeta = \text{constant}$. This solution has real action only for $v = 0$.)

The two solutions (7.8) correspond to different boundary choices of boundary conditions at the horizon. Following Refs. [452, 453], we choose the solution for which the momentum flux along the string world sheet flows from the boundary into the horizon, corresponding to the physical case in which the energy provided by the external agent that is pulling the quark through the plasma at constant speed is dissipated into the medium. This solution to (7.7) is given by

$$\zeta(z) = -\frac{v z_0}{2} \left(\operatorname{arctanh} \left(\frac{z}{z_0} \right) - \arctan \left(\frac{z}{z_0} \right) \right). \quad (7.9)$$

As illustrated in Fig. 7.1, this solution describes a string that trails behind the moving quark as it hangs down from it into the bulk spacetime.

The momentum flux flowing down from the boundary, along the string world sheet (7.9), and towards the horizon determines the amount of momentum lost by the quark in its propagation through the plasma. In terms of the field theory variables,

$$\frac{dp}{dt} = -\Pi_z^1 = -\frac{\pi T^2 \sqrt{\lambda}}{2} \gamma v. \quad (7.10)$$

Note, however, that in the stationary situation we have described, there is by construction no change in the actual momentum of the quark at the boundary; instead, in order to keep the quark moving with constant speed v against the force (7.10) there must be some external agency pushing the quark through the strongly coupled plasma. This force can be viewed as due to a constant electric field acting on the string endpoint, with the magnitude of the field given by

$$\mathcal{E} = \frac{\pi T^2 \sqrt{\lambda}}{2} \gamma v. \quad (7.11)$$

The physical setup described by the string (7.9) is thus that of forced motion of the quark through the plasma at constant speed in the presence of a constant electric field. The external force on the quark balances the backward *drag force* (7.10) on the quark exerted by the medium through which it is moving. To make it explicit that the medium exerts a drag force, we can rewrite (7.10) as

$$\frac{dp}{dt} = -\eta_D p, \quad (7.12)$$

with $p = M\gamma v$ the relativistic expression of the momentum of the quark and M the mass of the heavy particle. The drag coefficient is then

$$\eta_D = \frac{\pi \sqrt{\lambda} T^2}{2M}. \quad (7.13)$$

For test quarks with $M \rightarrow \infty$, as in the derivation above, this result is valid for motion with arbitrarily relativistic speeds v . It is remarkable that the energy loss of a heavy quark moving through the quark-gluon plasma with constant speed is described so simply, as due to a drag force. In contrast, in either a weakly coupled plasma [465] or a strongly coupled plasma that is not conformal [466], dp/dt is not proportional to p even at low velocities.

We shall see in Section 7.2 that a heavy quark moving through the strongly coupled plasma of $\mathcal{N} = 4$ SYM theory experiences transverse and longitudinal momentum broadening, in addition to losing energy via the drag that we have analyzed above. We shall review the implications of the understanding of how the presence of the strongly coupled plasma affects the motion of heavy quarks for heavy ion collision phenomenology at the end of Section 7.2.

7.1.1 Regime of validity of the drag calculation

In the derivation of the drag force above, we considered a test quark with $M \rightarrow \infty$. The result is, however, valid for quarks with finite mass M , as long as M is not too small. As we now show, the criterion that must be satisfied by M depends on the velocity of the quark

v . The closer v is to 1, the larger M must be in order for the energy loss of the quark to be correctly described via the drag force calculated above. In deriving the regime of validity of the drag calculation, we shall assume for simplicity that we are interested in large enough $\gamma = 1/\sqrt{1-v^2}$ that the M above which the calculation is valid satisfies $M \gg \sqrt{\lambda T}$. We will understand the need for this condition in Section 8.

The introduction of quarks with finite mass M in the fundamental representation of the gauge group corresponds in the dual gravitational description to the introduction of D7-branes [385], as we have reviewed in Section 5.5 and as we will further pursue in Section 8. The D7-brane extends from the boundary at $z = 0$ down to some z_q , related to the mass of the quarks it describes by

$$M = \frac{\sqrt{\lambda}}{2\pi z_q}, \quad (7.14)$$

a result that we shall explain in Section 8.1. The physical reason that the calculation of the drag force breaks down if M is too small or v is too large is that if the electric field \mathcal{E} required to keep the quark moving at constant speed v gets too large, one gets copious production of pairs of quarks and antiquarks with mass M , and the picture of dragging a single heavy quark through the medium breaks down completely [467]. The parametric dependence of the critical field \mathcal{E}_c at which pair production becomes copious can be estimated by inspection of how the Dirac-Born-Infeld action for the D7-brane depends on \mathcal{E} , namely

$$S_{DBI} \sim \sqrt{1 - \left(\frac{2\pi\alpha'}{R^2} \mathcal{E} z_q^2\right)^2}. \quad (7.15)$$

The critical maximum field strength that the D7 brane can support is the \mathcal{E}_c at which this action vanishes. This yields a criterion for the validity of the drag calculation, namely that \mathcal{E} must be less than of order

$$\mathcal{E}_c = \frac{2\pi M^2}{\sqrt{\lambda}}. \quad (7.16)$$

This maximum value of the electric field implies a maximum value of γ up to which the drag calculation can be applied for quarks with some finite mass M . From eq. (7.11) and eq. (7.16), this criterion is

$$\gamma v < \left(\frac{2M}{\sqrt{\lambda T}}\right)^2. \quad (7.17)$$

We shall assume that $M \gg \sqrt{\lambda T}$, meaning that in (7.17) we can take $\gamma v \simeq \gamma$. And, the estimate is only parametric, so the factor of two is not to be taken seriously. Thus, the result to take away is that the drag calculation is valid as long as

$$\gamma \lesssim \left(\frac{M}{\sqrt{\lambda T}}\right)^2. \quad (7.18)$$

The argument in terms of pair production for the limit (7.18) on the quark velocity gives a nice physical understanding for its origin, but this limit arises in a variety of other ways. For example, at (7.18) the velocity of the quark v becomes equal to the local speed of light in the bulk at $z = z_q$, where the trailing string joins onto the quark on the D7-brane. For example,

at (7.18) the screening length L_s (described below in Section 7.7) at which the potential between a quark and antiquark is screened becomes as short as the Compton wavelength of a quark of mass M , meaning that the calculation of Section 7.7 is also valid only in the regime (7.18) [149].

Yet further understanding of the meaning of the limit (7.18) can be gained by asking the question of what happens if the electric field is turned off, and the quark moving with speed v begins to decelerate due to the drag force on it. We would like to be able, at least initially, to calculate the energy loss of this now decelerating quark by assuming that this energy loss is due to the drag force, which from (7.10) means

$$\left. \frac{dE}{dt} \right|_{\text{drag}} = -\frac{\pi}{2} \sqrt{\lambda} T^2 \gamma v^2 = -\frac{\pi}{2} \sqrt{\lambda} T^2 \frac{pv}{M}. \quad (7.19)$$

However, once the quark is decelerating it is natural to expect that, due to its deceleration, it radiates and loses energy via this radiation also. The energy lost by a quark in strongly coupled $\mathcal{N} = 4$ SYM theory moving in vacuum along a trajectory with arbitrary acceleration has been calculated by Mikhailov [468]. For the case of a linear trajectory with deceleration a , his result takes the form

$$\left. \frac{dE}{dt} \right|_{\text{vacuum radiation}} = -\frac{\sqrt{\lambda}}{2\pi} a^2 \gamma^6 = -\frac{\sqrt{\lambda}}{2\pi} \frac{1}{M^2} \left(\frac{dp}{dt} \right)^2. \quad (7.20)$$

At least initially, dp/dt will be that due to the drag force, namely (7.10). We now see that the condition that dE/dt due to the vacuum radiation (7.20) caused by the drag-induced deceleration (7.10) be less than dE/dt due to the drag itself (7.19) simplifies considerably and becomes

$$\gamma < \left(\frac{2M}{\sqrt{\lambda} T} \right)^2, \quad (7.21)$$

the same criterion that we have seen before. This gives further physical intuition into the criterion for the validity of the drag calculation and at the same time demonstrates that this calculation cannot be used in the regime in which energy loss due to deceleration-induced radiation becomes dominant.

Motivated by the above considerations, the authors of Ref. [469] considered the (academic) case of a test quark moving in a circle of radius L with constant angular frequency ω . They showed that in this circumstance, dE/dt is given by (7.19), as if due to drag with no radiation, as long as $\omega^2 \ll (\pi T)^2 \gamma^3$, with γ the Lorentz factor for velocity $v = L\omega$. But, for $\omega^2 \gg (\pi T)^2 \gamma^3$, the energy loss of the quark moving in a circle through the plasma is precisely what it would be in vacuum according to Mikhailov's result, which becomes

$$\left. \frac{dE}{dt} \right|_{\text{vacuum radiation}} = \frac{\sqrt{\lambda}}{2\pi} v^2 \omega^2 \gamma^4 = \frac{\sqrt{\lambda}}{2\pi} a^2 \gamma^4 \quad (7.22)$$

for circular motion. Note that the radiative energy loss (7.22) is greater than that due to drag, (7.19), for

$$\omega^2 \gg (\pi T)^2 \gamma^3, \quad (7.23)$$

so the result of the calculation is that energy loss is dominated by that due to acceleration-induced radiation or that due to drag wherever each is larger. (Where they are comparable in magnitude, the actual energy loss is somewhat less than their sum [469].) This calculation shows that the *calculational method* that yields the result that a quark moving in a straight line with constant speed v in the regime (7.21) loses energy via drag can yield other results in other circumstances (see [470–472] for further examples). In the case of circular motion, the criterion for the validity of the calculational method is again (7.21), but there is a wide range of parameters for which this criterion and (7.23) are both satisfied [469]. This means that for a quark in circular motion, the calculation is reliable in a regime where energy loss is as if due to radiation in vacuum.

7.2 Momentum broadening of a heavy quark

In the same regime in which a heavy quark moving through the strongly coupled plasma of $\mathcal{N} = 4$ SYM theory loses energy via drag, as reviewed in Section 7.1, it is also possible to use gauge/gravity duality to calculate the transverse (and, in fact, longitudinal) momentum broadening induced by motion through the plasma [467, 473–475]. We shall review these calculations in this section. They have been further analyzed [476–478], and extended to study the effects of nonconformality [466, 479, 480] and acceleration [481, 482].

For non-relativistic heavy quarks, the result (7.12) is not surprising. The dynamics of this particle is that of Brownian motion which can be described by the effective equation of motion

$$\frac{dp}{dt} = -\eta_D p + \xi(t), \quad (7.24)$$

where $\xi(t)$ is a random force that encodes the interaction of the medium with the heavy probe and that causes the momentum broadening that we describe in this Section. For heavy quarks, we have seen in (7.13) that η_D is suppressed by mass. This reflects the obvious fact that the larger the mass the harder it is to change the momentum of the particle. Thus, for a heavy quark the typical time for such a change, $1/\eta_D$, is long compared to any microscopic time scale of the medium τ_{med} . This fact allows us to characterize the force distribution by the two point correlators

$$\begin{aligned} \langle \xi_T(t) \xi_T(t') \rangle &= \kappa_T \delta(t - t'), \\ \langle \xi_L(t) \xi_L(t') \rangle &= \kappa_L \delta(t - t'), \end{aligned} \quad (7.25)$$

where the subscripts L and T refer to the forces longitudinal and transverse to the direction of the particle's motion. Here, we are also assuming an isotropic plasma which leads to $\langle \xi_L(t) \rangle = \langle \xi_T(t) \rangle = 0$. In general, the force correlator would have a nontrivial dependence on the time difference (different from $\delta(t - t')$). However, since the dynamics of the heavy quark happen on timescales that are much larger than τ_{med} , we can approximate all medium correlations as happening instantaneously. It is then easy to see that the coefficient κ_T (κ_L) corresponds to the mean squared transverse (longitudinal) momentum transferred to the heavy quark per unit time. For example, the transverse momentum broadening is given by

$$\langle \mathbf{p}_\perp^2 \rangle = 2 \int dt dt' \langle \xi_T(t)_T \xi_T(t') \rangle = 2\kappa_T \mathcal{T}, \quad (7.26)$$

where \mathcal{T} is the total time duration (which should be smaller than $1/\eta_D$) and where the 2 is the number of transverse dimensions. It is clear from the correlator that κ_T is a property of the medium, independent of any details of the heavy quark probe. Our goal in this section is to calculate κ_T and κ_L . We shall do so first at low velocity, and then throughout the velocity regime in which the calculation of the drag force is valid.

Before we begin, we must show that in the limit we are considering the noise distribution is well characterized by its second moment. Odd number correlators vanish because of symmetry, so the first higher moment to consider is the fourth moment of the distribution of the transverse momentum picked up by the heavy quark moving through the plasma

$$\langle \mathbf{p}_\perp^4 \rangle = \int dt_1 dt_2 dt_3 dt_4 \langle \xi_T(t_1) \xi_T(t_2) \xi_T(t_3) \xi_T(t_4) \rangle . \quad (7.27)$$

The four-point correlator may be decomposed as

$$\begin{aligned} \langle \xi_T(t_1) \xi_T(t_2) \xi_T(t_3) \xi_T(t_4) \rangle &= \langle \xi_T(t_1) \xi_T(t_2) \xi_T(t_3) \xi_T(t_4) \rangle_c \\ &+ \langle \xi_T(t_1) \xi_T(t_2) \rangle \langle \xi_T(t_3) \xi_T(t_4) \rangle \\ &+ \langle \xi_T(t_1) \xi_T(t_3) \rangle \langle \xi_T(t_2) \xi_T(t_4) \rangle \\ &+ \langle \xi_T(t_1) \xi_T(t_4) \rangle \langle \xi_T(t_2) \xi_T(t_3) \rangle , \end{aligned} \quad (7.28)$$

which is the definition of the connected correlator. Due to time translational invariance, the connected correlator is a function

$$\langle \xi_T(t_1) \xi_T(t_2) \xi_T(t_3) \xi_T(t_4) \rangle_c = f(t_4 - t_1, t_3 - t_1, t_2 - t_1) . \quad (7.30)$$

As before, the correlator has a characteristic scale of the order of the medium scale. As a consequence, since the expectation value due to the connected part has only one free integral, we find

$$\langle \mathbf{p}_\perp^4 \rangle = \left(3 (2\kappa_T)^2 + \mathcal{O}\left(\frac{\tau_{\text{med}}}{\mathcal{T}}\right) \right) \mathcal{T}^2 , \quad (7.31)$$

where the dominant term comes from the disconnected parts in eq. (7.28). Since we are interested in times parametrically long compared to τ_{med} , we can neglect the connected part of the correlator.

7.2.1 κ_T and κ_L in the $p \rightarrow 0$ limit

The dynamical equations (7.24) together with (7.25) constitute the Langevin description of heavy quarks in a medium. In the $p \rightarrow 0$ limit, there is no distinction between transverse and longitudinal, meaning that both the fluctuations in (7.25) must be described by the same correlator with $\kappa_L = \kappa_T \equiv \kappa$. The Langevin equations (7.24) and (7.25) describe the time evolution of the probability distribution for the momentum of an ensemble of heavy quarks in a medium. A standard analysis shows that, independent of the initial probability distribution, after sufficient time any solution to the Langevin equation yields the probability distribution

$$\mathcal{P}(\mathbf{p}, t \rightarrow \infty) = \left(\frac{1}{\pi} \frac{\eta_D}{\kappa} \right)^{3/2} \exp \left\{ -\mathbf{p}^2 \frac{\eta_D}{\kappa} \right\} , \quad (7.32)$$

which coincides with the equilibrium (i.e. Boltzmann) momentum distribution for the heavy quark provided that

$$\eta_D = \frac{\kappa}{2MT}. \quad (7.33)$$

This expression is known as the Einstein relation. Thus, the Langevin dynamics of non-relativistic heavy quarks is completely determined by the momentum broadening κ , and the heavy quarks equilibrate at asymptotic times.

The Einstein relation (7.33) together with the computation of η_D in (7.13) for strongly coupled $\mathcal{N} = 4$ SYM theory allow us to infer the value of κ for this strongly coupled conformal plasma, namely

$$\kappa = \pi\sqrt{\lambda}T^3. \quad (7.34)$$

The dynamical equation (7.12) that we used in the previous section does not include the noise term simply because in that section we were describing the change in the mean heavy quark momentum in the ensemble of this section.

7.2.2 Direct calculation of the noise term

We would like to have a direct computation of the noise term in the description of a heavy quark in a strongly coupled gauge theory plasma. There are two motivations for this: 1) to explicitly check that the Einstein relation (7.33) is fulfilled and 2) to compute the momentum broadening for moving heavy quarks, which are not in equilibrium with the plasma and to which the Einstein relation therefore does not apply. This computation is somewhat technical; the reader interested only in the results for κ_T and κ_L for a moving heavy quark may skip to Section 7.2.3.

We need to express the momentum broadening in terms which are easily computed within the gauge/gravity correspondence. To do so, we prepare a state of the quark at an initial time t_0 which is moving at given velocity v in the plasma. In quantum mechanics, the state is characterized by a density matrix, which is a certain distribution of pure states

$$\rho(t_0) = \sum_n w(n) |n\rangle \langle n| \quad (7.35)$$

where the sum is performed over a complete set of states and the weight $w(n)$ is the ensemble. For a thermal distribution, the states are eigenstates of the Hamiltonian and $w(n) = \exp\{-E_n/T\}$.

In the problem we are interested in, the density matrix includes not only the quark degrees of freedom but also the gauge degrees of freedom. However, we start our discussion using a one particle system — initially for illustrative purposes. In this case, the distribution function of the particle is defined from the density matrix as

$$\hat{f}(x, x'; t_0) = \sum_n w(n) \langle x | n \rangle \langle n | x' \rangle, \quad (7.36)$$

where, as usual, $\langle x | n \rangle$ is the wave function of the particle in the state $|n\rangle$. It is also common to call $f(x, x')$ the density matrix. It is conventional to introduce the mean and relative coordinates and express the density matrix as

$$f(X, r; t_0) = \hat{f}\left(X + \frac{r}{2}, X - \frac{r}{2}; t_0\right), \quad (7.37)$$

where $X = (x + x')/2$ and $r = x - x'$. It is then easy to see that the mean position and mean momentum of the single particle with a given density matrix are given by

$$\begin{aligned}\langle x \rangle &= \text{tr} \{ \rho(t_0) x \} = \int dx x \hat{f}(x, x; t_0) = \int dX X f(X, 0; t_0) \\ \langle p \rangle &= \text{tr} \{ \rho(t_0) p \} = \int dx \frac{-i}{2} (\partial_x - \partial_{x'}) \hat{f}(x, x'; t_0) \Big|_{x'=x} \\ &= -i \int dX \partial_r f(X, r; t_0) \Big|_{r=0},\end{aligned}\tag{7.38}$$

meaning that r is the conjugate variable to the momentum and, of most interest to us, the mean squared momentum of the distribution is

$$\langle p^2 \rangle = - \int dX \partial_r^2 f(X, r; t_0) \Big|_{r=0}.\tag{7.39}$$

Returning now to the problem of interest to us, we must consider an ensemble containing the heavy quark and also the gauge field degrees of freedom. Since we assume the mass of the quark to be much larger than the temperature, we can describe the pure states of the system as

$$|A'\rangle = Q_a^\dagger(x) |A\rangle.\tag{7.40}$$

where $|A\rangle$ is a state of the gauge fields only, $|A'\rangle$ denotes a state of the heavy quark plus the gauge fields, and $Q_a^\dagger(x)$ is the creation operator (in the Schrödinger picture) of a heavy quark with color a at position x . Corrections to this expression are (exponentially) suppressed by T/M . The Heisenberg representation of the operator $Q(x)$ satisfies the equation of motion

$$(iu \cdot D - M) Q = 0,\tag{7.41}$$

where u is the four velocity of the quark and D is the covariant derivative with respect to the gauge fields of the medium. This equation realizes the physical intuition that the heavy quark trajectory is not modified by the interaction with the medium, which leads only to a modification of the quark's phase.¹

The full density matrix of the system, ρ , describes an ensemble of all the degrees of freedom of the system. Since we are only interested in the effects of the medium on the momentum of the heavy quark probe, we can define a one-body density matrix from the full density matrix by integrating over the gauge degrees of freedom

$$\begin{aligned}f(X, r; t_0) &= \left\langle Q_a^\dagger \left(X - \frac{r}{2}; \right) U_{ab} Q_b \left(X + \frac{r}{2} \right) \right\rangle \\ &= \text{Tr} \left[\rho Q_a^\dagger \left(X - \frac{r}{2} \right) U_{ab} Q_b \left(X + \frac{r}{2} \right) \right],\end{aligned}\tag{7.42}$$

where the trace is taken over a complete set of states

$$\sum_{A,a} \int dx Q_a^\dagger(x) |A\rangle \langle A| Q_a(x).\tag{7.43}$$

¹The expression (7.41) can also be derived from the Dirac equation by performing a Foldy-Wouthuysen transformation, which in the heavy quark rest frame is given by $Q = \exp\{\gamma \cdot D/2M\} \psi$, where $\gamma = 1/\sqrt{1-v^2}$.




Figure 7.2: Time contour C in the complex time plane for the path integral (7.47). Here, $\Delta t \equiv t - t_0$ and the $i\epsilon$ prescription in time is translated to the longitudinal coordinate x since the quark trajectory is $x = vt$. The two-point functions computed from the partition function (7.47) are evaluated at two arbitrary points t_C and t'_C on the contour.

Note that the inclusion of the operators in the trace in (7.42) plays the same role as the projectors $|x\rangle$ in (7.36). The gauge link U_{ab} in (7.42) joins the points $X + r/2$ and $X - r/2$ to ensure gauge invariance. In the long time limit, the precise path is not important, and we will assume that U_{ab} is a straight link. To simplify our presentation, we shall explicitly treat only transverse momentum broadening, which means taking the separation r to be in a direction perpendicular to the direction of motion of the heavy quark, \mathbf{r}_\perp .

At a later time t , after the heavy quark has propagated through the plasma for a time $t - t_0$, the one-body density matrix has evolved from (7.42) to

$$f(X, \mathbf{r}_\perp; t) = \text{Tr} \left[\rho e^{iH(t-t_0)} Q_a^\dagger \left(X - \frac{\mathbf{r}_\perp}{2} \right) e^{-iH(t-t_0)} e^{iH(t-t_0)} U_{ab} e^{-iH(t-t_0)} e^{iH(t-t_0)} Q_b \left(X + \frac{\mathbf{r}_\perp}{2} \right) e^{-iH(t-t_0)} \right], \quad (7.44)$$

where we have introduced evolution operators to express the result in the Heisenberg picture. We then introduce a complete set of states, obtaining

$$f(X, \mathbf{r}_\perp; t) = \int d\mathbf{x}_1 d\mathbf{x}_2 \sum_{A_1, A_2, A_3, A_4} \rho_{a_1 a_2}[\mathbf{x}_1, \mathbf{x}_2; A_1, A_2] \langle A_2 | Q_{a_2}(\mathbf{x}_2) Q_a^\dagger \left(X - \frac{\mathbf{r}_\perp}{2}; t \right) | A_3 \rangle \langle A_3 | U_{ab}(t) | A_4 \rangle \langle A_4 | Q_b \left(X + \frac{\mathbf{r}_\perp}{2}; t \right) Q_{a_1}^\dagger(x_1) | A_1 \rangle, \quad (7.45)$$

where we have defined

$$\rho_{a_1 a_2}[\mathbf{x}_1, \mathbf{x}_2; A_1, A_2] \equiv \langle A_1 | Q_{a_1}(\mathbf{x}_1) \rho Q_{a_2}^\dagger(\mathbf{x}_2) | A_2 \rangle. \quad (7.46)$$

The expression (7.45) can be expressed as a path integral. Note that the second line in (7.45) is an anti-time ordered correlator; thus, its path integral representation involves a time reversal of the usual path integral. Instead of introducing two separate path integrals corresponding to the second and fourth lines of (7.45), we introduce the time contour shown in Fig. 7.2 and use this contour to define a single path integral. The fields A_1 and A_2 are the




Figure 7.3: Graphical representation of eq. (7.48). The Wilson line indicated by the black line is denoted $W_C[\mathbf{r}_\perp/2, -\mathbf{r}_\perp/2]$. This Wilson line is traced with the initial density matrix, $\rho_{a_1 a_2}^o$. The horizontal axis is along the time direction and the vertical axis is along one of the transverse coordinates, x_\perp . $\Delta t \equiv t - t_0$.

values at the endpoints of the contour. The one-body density matrix then reads

$$\begin{aligned}
f(X, \mathbf{r}_\perp; t) &= \sum_{A_1, A_2} \int d\mathbf{x}_1 d\mathbf{x}_2 \int [DA] \mathcal{D}Q \mathcal{D}Q^\dagger e^{i \int_C d^4x \{ \mathcal{L}_{YM} + Q^\dagger(iu \cdot D - M)Q \}} \\
&\quad \rho_{a_1 a_2}[\mathbf{x}_1, \mathbf{x}_2; A_1, A_2] U_{ab}(t) \\
&\quad Q_{a_2}(\mathbf{x}_2, t_0 - i\epsilon) Q_a^\dagger\left(X - \frac{\mathbf{r}_\perp}{2}, t - i\epsilon\right) \\
&\quad Q_b\left(X + \frac{\mathbf{r}_\perp}{2}, t\right) Q_{a_1}^\dagger(\mathbf{x}_1, t_0) .
\end{aligned} \tag{7.47}$$

Integrating out the quark field and working to leading order in T/M (neglecting the fermionic determinant) yields

$$f(X, \mathbf{r}_\perp; t) = \left\langle \text{tr} \left[\rho \left[X + \frac{\mathbf{r}_\perp}{2}, X - \frac{\mathbf{r}_\perp}{2}; A_1, A_2 \right] W_C \left[\frac{\mathbf{r}_\perp}{2}, -\frac{\mathbf{r}_\perp}{2} \right] \right] \right\rangle_A , \tag{7.48}$$

where the subscript A indicates averaging with respect to the gauge fields, and where the Wilson line $W_C[\mathbf{r}_\perp/2, -\mathbf{r}_\perp/2]$ is defined in Fig. 7.3. We have used the fact that the Green's function of eq. (7.41) is the (contour ordered) Wilson line.

Next, we perform a Taylor expansion of the time-evolved density matrix (7.48) about $\mathbf{r}_\perp = 0$, obtaining

$$\begin{aligned}
f(X, \mathbf{r}_\perp; t) &= f(X, 0; t) + \\
&\quad \frac{\mathbf{r}_\perp^2}{2} \left\langle \text{tr} \left[\frac{\partial^2}{\partial \mathbf{r}_\perp^2} \rho \left[X + \frac{\mathbf{r}_\perp}{2}, X - \frac{\mathbf{r}_\perp}{2}; A_1, A_2 \right] W_C[0] \right] \right\rangle_A + \\
&\quad \frac{\mathbf{r}_\perp^2}{2} \kappa_T \Delta t \langle \text{tr} [\rho[X, X; A_1, A_2] W_C[0]] \rangle_A .
\end{aligned} \tag{7.49}$$

The second term in this expression involves only derivatives of the initial density matrix; thus, it is the mean transverse momentum squared of the initial distribution (which may be supposed to be small). In the last term, which scales with the elapsed time Δt , we have defined

$$\kappa_T \Delta t = \frac{1}{4} \frac{1}{\langle \text{tr} \rho W_C[0] \rangle_A} \int_C dt_C dt'_C \left\langle \text{tr} \rho[X, X; A_1, A_2] \frac{\delta^2 W_C[\delta y]}{\delta y(t_C) \delta y(t'_C)} \right\rangle_A , \tag{7.50}$$

where t_C denotes time along the contour depicted in Fig. 7.2. $\kappa_T \Delta t$ is the mean transverse momentum squared picked up by the heavy quark during the time Δt . We have expressed the transverse derivatives of the Wilson line as functional derivatives with respect to the path of the Wilson line. The path δy denotes a small transverse fluctuation $\delta y(t)$ away from the path $X_1 = vt$.

The contour δy may be split into two pieces, δy_1 and δy_2 , which run along the time ordered and anti-time ordered part of the path. Thus, the fluctuation calculation defines four correlation functions

$$iG_{11}(t, t') = \frac{1}{\langle \text{tr} \rho^o W_C[0, 0] \rangle_A} \left\langle \text{tr} \rho^o \frac{\delta^2 W_C[\delta y_1, 0]}{\delta y_1(t) \delta y_1(t')} \right\rangle_A, \quad (7.51)$$

$$iG_{22}(t, t') = \frac{1}{\langle \text{tr} \rho^o W_C[0, 0] \rangle_A} \left\langle \text{tr} \rho^o \frac{\delta^2 W_C[0, \delta y_2]}{\delta y_1(t) \delta y_2(t')} \right\rangle_A, \quad (7.52)$$

$$iG_{12}(t, t') = \frac{1}{\langle \text{tr} \rho^o W_C[0, 0] \rangle_A} \left\langle \text{tr} \rho^o \frac{\delta^2 W_C[\delta y_1, \delta y_2]}{\delta y_1(t) \delta y_2(t')} \right\rangle_A, \quad (7.53)$$

$$iG_{21}(t, t') = \frac{1}{\langle \text{tr} \rho^o W_C[0, 0] \rangle_A} \left\langle \text{tr} \rho^o \frac{\delta^2 W_C[\delta y_2, \delta y_2]}{\delta y_1(t') \delta y_2(t)} \right\rangle_A. \quad (7.54)$$

Note that the first two correlators correspond to time ordered and anti-time ordered correlators, while the last two are unordered. We can then divide the integration over t_C and $t_{C'}$ in (7.50) into four parts corresponding to the cases where each of t_C and $t_{C'}$ is on the upper or lower half of the contour in Fig. 7.2. In the large Δt limit we can then use time translational invariance to cast (7.50) as

$$\kappa_T = \lim_{\omega \rightarrow 0} \frac{1}{4} \int dt e^{+i\omega t} [iG_{11}(t, 0) + iG_{22}(t, 0) + iG_{12}(t, 0) + iG_{21}(t, 0)]. \quad (7.55)$$

This admittedly rather formal expression for κ_T is as far as we can go in general. In Section 7.2.3 we evaluate κ_T (and κ_L) in the strongly coupled plasma of $\mathcal{N} = 4$ SYM theory.

Although our purpose in deriving the expression (7.55) is to use it to analyze the case $v \neq 0$, it can be further simplified in the case that $v = 0$. On the time scales under consideration, the static quark is in equilibrium with the plasma, and the Kubo-Martin-Schwinger relation which takes the form

$$i[G_{11}(\omega) + G_{22}(\omega) + G_{12}(\omega) + G_{21}(\omega)] = -4 \coth\left(\frac{\omega}{2T}\right) \text{Im} G_R(\omega) \quad (7.56)$$

for $\epsilon \rightarrow 0$ applies [483]. Here, G_R is the retarded correlator. Thus, we find

$$\kappa_T(v = 0) = \lim_{\omega \rightarrow 0} \left(-\frac{2T}{\omega}\right) \text{Im} G_R(\omega). \quad (7.57)$$

If $v \neq 0$, however, we must evaluate the four correlators in the expression (7.55).

7.2.3 κ_T and κ_L for a moving heavy quark

We see from the expression (7.50) that the transverse momentum broadening coefficient κ_T is extracted by analyzing small fluctuations in the path of the Wilson line depicted in Fig. 7.3.

In the strongly coupled plasma of $\mathcal{N} = 4$ SYM theory, we can use gauge/gravity duality to evaluate κ_T starting from (7.50). In the dual gravitational description, the small fluctuations in the path of the Wilson line amount to perturbing the location on the boundary at which the classical string (whose unperturbed shape is given by (7.9)) terminates according to

$$(x_1(t, z), 0, 0) \rightarrow (x_1(t, z), y(t, z), 0) . \quad (7.58)$$

The perturbations of the Wilson line at the boundary yield fluctuations on the string world sheet dragging behind the quark. Because we wish to calculate κ_T , in (7.58) we have only introduced perturbations transverse to the direction of motion of the quark. We shall quote the result for κ_L at the end; calculating it requires extending (7.58) to include perturbations to the function $x_1(t, z)$.

In order to analyze fluctuations of the string world sheet, we begin by casting the metric induced on the string world sheet in the absence of any perturbations

$$ds_{\text{ws}}^2 = \frac{R^2}{z^2} \left(- (f(z) - v^2) d\tau^2 + \frac{\hat{f}(z)}{f^2(z)} d\sigma^2 - v^2 \frac{z^2/z_0^2}{f(z)} (d\tau d\sigma + d\sigma d\tau) \right) , \quad (7.59)$$

in a simpler form. In (7.59), we have defined $\hat{f}(z) \equiv 1 - z^4/(z_0^4 \gamma^2)$. The induced metric (7.59) is diagonalized by the change of world sheet coordinates

$$\begin{aligned} \hat{t} &= \frac{t}{\sqrt{\gamma}} + \frac{z_0}{2\sqrt{\gamma}} \left(\arctan\left(\frac{z}{z_0}\right) - \operatorname{arctanh}\left(\frac{z}{z_0}\right) \right. \\ &\quad \left. - \sqrt{\gamma} \arctan\left(\frac{\sqrt{\gamma}z}{z_0}\right) + \sqrt{\gamma} \operatorname{arctanh}\left(\frac{\sqrt{\gamma}z}{z_0}\right) \right) , \\ \hat{z} &= \sqrt{\gamma}z , \end{aligned} \quad (7.60)$$

in terms of which the induced metric takes the simple form

$$ds_{\text{ws}}^2 = \frac{R^2}{\hat{z}^2} \left(-f(\hat{z}) d\hat{t}^2 + \frac{1}{f(\hat{z})} d\hat{z}^2 \right) . \quad (7.61)$$

Note that this has the same form as the induced metric for the world sheet hanging below a motionless quark, upon making the replacement $(\hat{t}, \hat{z}) \rightarrow (t, z)$. In particular, the metric (7.61) has a horizon at $\hat{z} = z_0$, which means that the metric describing the world sheet of the string trailing behind the moving quark has a world sheet horizon at $z = z_{\text{ws}} \equiv z_0/\sqrt{\gamma}$. For $v \rightarrow 0$, the location of the world sheet horizon drops down toward the spacetime horizon at $z = z_0$. But, for $v \rightarrow 1$, the world sheet horizon moves closer and closer to the boundary at $z = 0$, i.e. towards the ultraviolet. As at any horizon, the singularity at $z = z_{\text{ws}}$ (i.e. at $\hat{z} = z_0$) in (7.61) is just a coordinate singularity. In the present case, this is manifest since (7.61) was obtained from (7.59) which is regular at $z = z_{\text{ws}}$ by a coordinate transformation (7.60). Nevertheless, the world sheet horizon has clear physical significance: at $z = z_{\text{ws}}$ the local speed of light at this depth in the bulk matches the speed v with which the quark at the boundary is moving. Furthermore, and of direct relevance to us here, because of the world sheet horizon at $z = z_{\text{ws}}$ fluctuations of the string world sheet at $z > z_{\text{ws}}$, below — to the

infrared of — the world sheet horizon, are causally disconnected from fluctuations at $z < z_{\text{ws}}$ above the world sheet horizon and in particular are causally disconnected from the boundary at $z = 0$.

The remarkable consequence of the picture that emerges from the above analysis of the unperturbed string world sheet trailing behind the quark at the boundary moving with speed v is that the momentum fluctuations of this quark can be thought of as due to the Hawking radiation on the string world sheet, originating from the world sheet horizon at $z = z_{\text{ws}}$ [467, 474]. It is as if the force fluctuations that the quark in the boundary gauge theory feels are due to the fluctuations of the string world sheet to which it is attached, with these fluctuations arising due to the Hawking radiation originating from the world sheet horizon. It will therefore prove useful to calculate the Hawking temperature of the world sheet horizon, which we denote T_{ws} . As detailed in Appendix B this can be done in the standard fashion, upon using a further coordinate transformation to write the metric (7.61) in the vicinity of the world sheet horizon in the form $ds_{\text{ws}}^2 = -b^2\rho^2 d\hat{t}^2 + d\rho^2$ for some constant b , where the world sheet horizon is at $\rho = 0$. Then, it is a standard argument that in order to avoid having a conical singularity at $\rho = 0$ in the Euclidean version of this metric, namely $ds_{\text{ws}}^2 = b^2\rho^2 d\hat{\theta}^2 + d\rho^2$, $b\hat{\theta}$ must be periodic with period 2π . The periodicity of the variable $\hat{\theta}$, namely $2\pi/b$, is $1/T$. Since at the boundary, where $z = 0$, Eq. (7.60) becomes $\hat{t} = t/\sqrt{\gamma}$, this argument yields

$$T_{\text{ws}} = \frac{T}{\sqrt{\gamma}}, \quad (7.62)$$

a result that we shall use below.

We have gained significant physical intuition by analyzing the unperturbed string world sheet, but in order to obtain a quantitative result for κ_T we must introduce the transverse fluctuations $y(t, z)$ defined in (7.58) explicitly. We write the Nambu-Goto action for the string world sheet with $y(t, z) \neq 0$, and expand it to second order in y , obtaining the zeroth order action (7.5) plus a second order contribution

$$S_T^{(2)}[y] = \frac{\gamma R^2}{2\pi\alpha'} \int \frac{d\hat{t}d\hat{z}}{\hat{z}^2} \frac{1}{2} \left(\frac{\dot{y}^2}{f(\hat{z})} - f(\hat{z})y'^2 \right) \quad (7.63)$$

where $\dot{}$ and \prime represent differentiation with respect to \hat{t} and \hat{z} respectively. This action is conveniently expressed as

$$S_T^{(2)}[y] = -\frac{\gamma R^2}{2\pi\alpha'} \int \frac{d\hat{t}d\hat{z}}{\hat{z}^2} \frac{1}{2} \sqrt{-h} h^{ab} \partial_a y \partial_b y \quad (7.64)$$

with h_{ab} the induced metric on the unperturbed world sheet that we have analyzed above. The existence of the world sheet horizon means that we are only interested in solutions to the equations of motion for the transverse fluctuations y obtained from (7.64) that satisfy infalling boundary conditions at the world sheet horizon. This constraint in turn implies a relation among the correlators analogous to those in (7.55) that describe the transverse fluctuations of the world sheet, and in fact the relation turns out to be analogous to the Kubo-Martin-Schwinger relation (7.56) among the gauge theory correlators [467]. Consequently, for a quark moving with velocity v the transverse momentum broadening coefficient $\kappa_T(v)$ is given by the

same expression (7.57) that is valid at $v = 0$ with T replaced by the world sheet temperature T_{ws} of (7.62) [467, 474]. That is,

$$\kappa_T(v) = \lim_{\omega \rightarrow 0} \left(-\frac{2T_{\text{ws}}}{\omega} \text{Im} \hat{G}_R(\omega) \right), \quad (7.65)$$

where \hat{G}_R denotes the retarded correlator at the world-sheet horizon. The fact that in the strongly coupled theory there is a KMS-like relation at $v \neq 0$ after all is a nontrivial consequence of the development of the world-sheet horizon.

The computation of the retarded correlator follows the general procedure of Ref. [368] described in Section 5.3. Since the action (7.63) is a function of \hat{t} which is given by $t/\sqrt{\gamma}$ at the boundary, the retarded correlator is a function of $\hat{\omega} = \sqrt{\gamma}\omega$ (with ω the frequency of oscillations at the boundary). To avoid this complication, and in particular in order to be able to apply the general results for $\text{Im} G_R$ that we derived in Section 6.2, it is convenient to define

$$\tilde{t} = \sqrt{\gamma} \hat{t}, \quad (7.66)$$

so that $\tilde{t} = t$ at the boundary. We now wish to apply the general expressions (A.10), (6.16) and (6.17). In order to do so, we identify the world sheet metric h^{ab} and the field y in the action (7.64) with the metric g^{MN} and the field ϕ in the action (6.15), meaning that in our problem the function q in (6.15) takes the specific form

$$\frac{1}{q(z)} = \frac{\gamma R^2}{2\pi\alpha'} \frac{1}{\hat{z}^2} = \frac{\sqrt{\lambda}}{2\pi z^2}. \quad (7.67)$$

Furthermore, for the two-dimensional world sheet metric we have $-h = h_{\tilde{t}\tilde{t}} h_{zz}$, meaning that from the general result (6.24) we find

$$-\lim_{\omega \rightarrow 0} \frac{\text{Im} \hat{G}_R(\omega)}{\omega} = \frac{1}{q(z_{\text{ws}})} = \frac{\gamma\sqrt{\lambda}}{2\pi} (\pi T)^2, \quad (7.68)$$

and thus

$$\kappa_T = \sqrt{\lambda\gamma}\pi T^3, \quad (7.69)$$

which is our final result for the transverse momentum broadening coefficient.

The analysis of longitudinal fluctuations and the extraction of κ_L proceed analogously to the analysis we have just reviewed, except that in (7.58) we introduce a perturbation to the function x_1 instead of a transverse perturbation y . At quadratic order, there is no coupling between the transverse and longitudinal perturbations. Remarkably, the action for longitudinal fluctuations of the string is the same as that for transverse fluctuations eq. (7.63) up to a constant:

$$S_L^{(2)}[x] = \gamma^2 S_T^{(2)}[x], \quad (7.70)$$

with γ the Lorentz factor. Following the analogous derivation through, we conclude that

$$\kappa_L = \gamma^2 \kappa_T = \gamma^{5/2} \sqrt{\lambda} \pi T^3. \quad (7.71)$$

This result shows that κ_L depends very strongly on the velocity of the heavy quark. Indeed, κ_L grows faster with increasing velocity than the energy squared of the heavy quark, $\gamma^2 M^2$.

Thus, the longitudinal momentum acquired by a quark moving through a region of strongly coupled $\mathcal{N} = 4$ SYM plasma of finite extent does not become a negligible fraction of the energy of the quark in the high energy limit. This is very different from the behavior of a quark moving through a weakly coupled QCD plasma, in which the longitudinal momentum transferred to the quark can be neglected in the high energy limit. However, we should keep in mind that due to the bound (7.18), for a given value of the mass M and the coupling $\sqrt{\lambda}$ the calculation of κ_L (and of κ_T) is only valid for finite energy quarks, with γ limited by (7.18).

We see from the expressions (7.69) and (7.71) for κ_T and κ_L derived by explicit analysis of the fluctuations that in the $v \rightarrow 0$ limit we have $\kappa_T = \kappa_L = \kappa$ with κ given by (7.34), as we obtained previously from the drag coefficient η_D via the use of the Einstein relation (7.33). This is an example of the fluctuation-dissipation theorem.

In the gauge theory, momentum broadening is due to the fluctuating force exerted on the heavy quark by the fluctuating plasma through which it is moving. In the dual gravitational description, the quark at the boundary feels a fluctuating force due to the fluctuations of the world sheet that describes the profile of the string to which the quark is attached. These fluctuations have their origin in the Hawking radiation of fluctuations of the string world sheet originating from the world sheet horizon. The explicit computation of this world sheet Hawking radiation for a quark at rest was performed in Refs. [477, 484], and these results nicely reproduce those we have obtained within a Langevin formalism. This computation was extended to quarks moving at nonzero velocity in Refs. [475, 478].

7.2.4 Implications for heavy quarks in heavy ion collisions

Models based upon Langevin dynamics have been used by a number of authors to describe the motion of heavy quarks within the hot expanding fluid produced in heavy ion collisions [7, 465, 485–493], with the goal of using data from RHIC [136, 494, 495] to constrain the diffusion constant D that describes the Brownian motion of a heavy quark in the hot fluid. These analyses typically assume that the relative velocity of the heavy quark and the hot fluid is not relativistic, meaning that $\kappa_T = \kappa_L \equiv \kappa$. In this regime, the diffusion constant is given by $D = 2T^2/\kappa$, meaning that the result (7.34) translates into the statement that a heavy quark in the strongly coupled $\mathcal{N} = 4$ SYM theory plasma obeys a Langevin equation with

$$D = \frac{4}{\sqrt{\lambda}} \frac{1}{2\pi T} \approx \frac{1.1}{2\pi T} \sqrt{\frac{1}{\alpha_{\text{SYM}} N_c}} . \quad (7.72)$$

The diffusion constant D parametrizes how strongly the heavy quark couples to the medium, with smaller D corresponding to stronger coupling and shorter mean free path. D is well defined even if it is so small that it does not make sense to define a mean free path, as is the case in the plasma of strongly coupled $\mathcal{N} = 4$ SYM theory, given (7.72). In a weakly coupled QCD plasma, perturbative calculations suggest [465]

$$D_{\text{weakly coupled}} \approx \frac{14}{2\pi T} \left(\frac{.33}{\alpha_s} \right)^2 \quad (7.73)$$

or larger by a factor of more than three [487], although it should be noted that the perturbative expansion converges poorly meaning that such calculations can only become quantita-

tively reliable at values of α_s that are much smaller than 0.33 [496, 497]. Note that the result (7.72) is valid for $\mathcal{N} = 4$ SYM theory when $\lambda = 4\pi\alpha_{\text{SYM}}N_c$ is large, and $2\pi TD$ is small, while the result (7.73) is valid for QCD when α_s is small, and $2\pi TD$ is large.

At present, the central lessons from RHIC data on heavy quarks are qualitative. This is because, so far, the best experimental constraints on the spectra of c and b quarks come mainly from single-electron spectra that are known to be dominated by the semileptonic decays of heavy quarks. The experimental pointing resolution is not yet sufficient to separate b from c quark decays via their displaced vertices. Upgrades at PHENIX and STAR will soon allow one to make this separation. Moreover, with the onset of heavy ion collisions at the LHC, it will soon be possible to measure B- and D-meson spectra up to $p_T \sim 20$ GeV, e.g. via the decay channel $D^0 \rightarrow K^+ \pi^-$. Also, heavy quark energy loss can be characterized via muon spectra at the LHC. The experimental understanding of how heavy quarks behave in the matter produced in relativistic heavy ion collisions can be expected to become much more quantitative with these developments.

At present, there are two qualitative lessons from the RHIC data [136, 494, 495]: (i) R_{AA} for isolated electrons (some linear combination of c and b quarks) is small, not that much bigger than R_{AA} for pions; and (ii) v_2 for isolated electrons is significant, not that much smaller than that for pions. At a qualitative level, both these observations suggest that c and b quarks are strongly coupled to the medium, as they must initially be slowed by their relative motion through it and must then flow along with the collective hydrodynamic expansion. Theoretical calculations involve calculating b and c quark production, using a Langevin model to describe their diffusion in the hot fluid produced in the collision, which is described hydrodynamically, and then modelling the freezeout and decay of the heavy B and D mesons, before finally comparing to data on R_{AA} and v_2 for isolated electrons. Two recent calculations along these lines are qualitatively consistent with RHIC data as long as the diffusion constant is in the range [488]

$$D_{\text{RHIC}} \approx \frac{2-6}{2\pi T} \ , \quad (7.74)$$

or $(3-5)/(2\pi T)$ [487]. Certainly the extraction of D , or equivalently κ , will become more quantitative once b 's and c 's can be measured separately. On the theoretical side, it will be interesting to redo these analyses with $\kappa_T \neq \kappa_L$ and both velocity-dependent as in (7.69) and (7.71). It will be interesting to see whether the explicit velocity dependence of κ_T and κ_L serves to improve the fit to R_{AA} and v_2 for b and c quarks. There is some tension in the present analyses of isolated electrons with a single fit parameter D , with the R_{AA} data favoring somewhat larger values of D while the v_2 data favors somewhat smaller values [136].

It is striking that D for strongly coupled $\mathcal{N} = 4$ SYM theory, given in (7.72), is even in the same ballpark as that extracted by fitting to present RHIC data, given in (7.74). To understand whether this is a coincidence, we need to ask how D would change if we could deform $\mathcal{N} = 4$ SYM theory so as to turn it into QCD. This is not a question to which the answer is known, but we can make two observations. First, in a large class of conformal theories, at a given value of T , N_c and λ both κ and the drag coefficient η_D scale with the square root of the entropy density. (The argument is the same as that for the jet quenching parameter \hat{q} of Section 7.5, and is given in Ref. [149].) The number of degrees of freedom in QCD is smaller than that in $\mathcal{N} = 4$ SYM theory by a factor of 47.5/120 for $N_c = 3$, suggesting

that κ and η_D should be smaller in QCD by a factor of $\sqrt{47.5/120} = 0.63$, making D larger by a factor of $\sqrt{120/47.5} = 1.59$.² Second, $\mathcal{N} = 4$ SYM theory is of course conformal, while QCD is not. Analysis of a toy model in which nonconformality can be introduced by hand suggests that turning on nonconformality to a degree suggested by lattice calculations of QCD thermodynamics reduces D , perhaps by as much as a factor of two [466]. Turning on nonconformality in $\mathcal{N} = 2^*$ theory also reduces D . So, D in a strongly coupled QCD plasma is not known reliably but the reduction in degrees of freedom and the nonconformality seem to push in opposite directions, suggesting that D_{QCD} may not be far from that given in (7.72). It will be very interesting to watch how this story evolves once it is possible to measure b and c quarks separately in heavy ion collisions.

7.3 Disturbance of the plasma induced by an energetic heavy quark

In Sections 7.1 and 7.2 we have analyzed the effects of the strongly coupled plasma of $\mathcal{N} = 4$ SYM theory on an energetic heavy quark moving through it, focussing on how the heavy quark loses energy in Section 7.1 and on the momentum broadening that it experiences in Section 7.2. In this Section, we turn the tables and analyze the effects of the energetic heavy quark on the medium through which it is propagating [7, 137, 138, 498–525]. From the point of view of QCD calculations and heavy ion collision phenomenology, the problem of understanding the response of the medium to an energetic probe is quite complicated. An energetic particle passing through the medium can excite the medium on many different wavelengths. And, even if the medium was thermalized prior to the interaction with the probe, the disturbance caused by the probe must drive the medium out of equilibrium, at least close to the probe. And, non-equilibrium processes are difficult to treat, especially at strong coupling.

It is natural to attempt to describe the disturbance of the medium using hydrodynamics, with the energetic particle treated as a source for the hydrodynamic equations. This approach is based on two assumptions. First, one must assume that the medium itself can be described hydrodynamically. Second, one has to assume that the non-equilibrium disturbance in the vicinity of the energetic particle relaxes to some locally equilibrated (but still excited) state after the energetic particle has passed on a timescale that is short compared to the lifetime of the hydrodynamic medium itself. The first assumption is clearly supported by data from heavy ion collisions at RHIC, as discussed in Section 2.2. The second assumption is stronger, and less well justified. Even though, as we saw in Section 2.2, there is evidence from the data that in heavy ion collisions at RHIC a hydrodynamic medium in local thermal equilibrium forms rapidly, after only a short initial thermalization time, it is not clear *a priori* that

²Gubser has suggested that the difference in the numbers of degrees of freedom between QCD and $\mathcal{N} = 4$ SYM theory be taken into account by comparing results in the two theories at differing temperatures, with $T_{\text{QCD}} = (47.5/120)^{1/4} T_{\text{SYM}}$ [7, 460]. This prescription is inconsistent with known results for a large class of conformal theories [149], in which \hat{q} , κ , η_D and D all scale with the square root of the number of degrees of freedom when compared at fixed T , N_c and λ . Perhaps coincidentally, since the drag coefficient η_D is proportional to T^2 scaling it by the square root of the number of degrees of freedom is equivalent to scaling T by the one-fourth power of the number of degrees of freedom as prescribed in Refs. [7, 460]. But, this prescription is not correct when applied to $D \propto 1/T$ or to $\kappa \propto T^3$ or $\hat{q} \propto T^3$.

the relaxation time for the disturbance caused by an energetic quark plowing through this medium is comparably short, particularly since the density of the medium drops with time. Finally, even if a hydrodynamic approach to the dynamics of these disturbances is valid, the details of the functional form of this hydrodynamic source are unknown, since the relaxation process is not under theoretical control.

Keeping the above difficulties in mind, it is still possible to use the symmetries of the problem and some physical considerations to make some progress toward understanding the source for the hydrodynamic equations corresponding to the disturbance caused by an energetic quark. If the propagating parton is sufficiently energetic, we may assume that it moves at a fixed velocity; this ansatz forces the source to be a function of $x - vt$, with the parton moving in the x -direction. We may also assume that the source has cylindrical symmetry around the parton direction. We may also constrain the source by the amount of energy and momentum that is fed into the plasma, which for the case of the plasma of strongly coupled $\mathcal{N} = 4$ SYM theory we calculated in Section 7.1. In an infinite medium, at late enough times, all the energy lost by the probe must thermalize and be incorporated into heating and/or hydrodynamic motion. (This may not be a good approximation for a very energetic parton propagating through weakly coupled plasma of finite extent since, as we have discussed in Section 2.3, in this setting the parton loses energy by the radiation of gluons whose energy and momentum are large relative to the temperature of the medium, which may escape from the medium without being thermalized.)

Although the caveats above caution against attempting to draw quantitative conclusions without further physical inputs, the success of the hydrodynamical description of the medium itself support the conclusion that there must be some hydrodynamic response to the passage of the energetic particle through it. In particular, there must be some coupling of the energetic particle to sound waves in the medium. Since an energetic particle is moving through the plasma with a velocity greater than the velocity of sound c_s , the coupling of the energetic particle to the sound mode must lead to the formation of a Mach cone, namely a sound front moving away from the trajectory of the energetic particle at the Mach angle

$$\cos \Theta_M = \frac{c_s}{v}, \quad (7.75)$$

with v the velocity of the energetic particle. However, absent further physical inputs it is hard to estimate how strong the Mach cone produced by a *point-like* particle shooting through the medium will be.

Several years ago, the interest in these phenomena intensified considerably when the formation of Mach cones was suggested as a candidate to explain certain nontrivial particle correlations seen in heavy ion collisions at RHIC, illustrated in the top-left panel of Fig. 2.9. Recall that in these experiments one triggers on (that is, selects events containing) a high energy hadron in the final state, and then measures the distribution of all the other soft hadrons in azimuthal angle $\Delta\phi$ relative to the direction of the trigger particle. After subtracting the effects of elliptic flow, the correlation functions for nucleus-nucleus collisions show a near-side jet of hadrons at small $\Delta\phi$ and a very broad distribution of soft particles on the away-side, at angles around $\Delta\phi \sim \pi$, which may have low peaks at $\Delta\phi = \pi \pm \phi_v$ with $\phi_v \approx 1 - 1.2$ radians, on the shoulders of the away-side distribution rather than at its center at $\Delta\phi = \pi$. This correlation pattern is very different from that seen in proton-proton

collisions, in which the away-side peak is broader than the near-side peak but is approximately Gaussian in shape, centered about a single peak at $\Delta\phi = \pi$. Many authors have argued that this nontrivial structure is the debris resulting from the near-complete quenching (thermalization) of the energetic particle that was initially produced back-to-back with the hard parton that became the trigger hadron, with the Mach cone excited in the medium by this energetic particle leading to the peaks on the shoulders of the away-side distribution at $\Delta\phi = \pi \pm \phi_v$ [137, 138, 501, 521]. This qualitative picture motivates the need for a good quantitative model for the coupling between an energetic parton and the collective modes of the strongly coupled plasma.

More recently, some doubt has been cast on the interpretation of the qualitative features seen in the top-left panel of Fig. 2.9 in terms of a Mach cone. Alver and Roland realized that even though $v_3 = 0$ on average, it is nonzero in individual heavy ion collisions [526]. On average the shape of the overlap between two nuclei in an ensemble of collisions all with the same nonzero impact parameter is almond-shaped, meaning that $v_n = 0$ for all odd n . But, individual nuclei are made of 200 nucleons, meaning that when they are Lorentz-contracted into pancakes these pancakes are not perfectly circular. This means that the overlap between two individual nuclei will not be perfectly almond shaped. In general, the shape of this overlap zone will have some “triangularity”, which raises the possibility of a nonzero v_3 in the final state, since the hydrodynamic expansion of an initial state that is initially triangular in position space will result in a nonzero $v_3 = \langle \exp[i3(\phi - \Phi_R)] \rangle$ in momentum space, and hence in the observed hadrons. (Recall that Φ_R is the angle of the reaction plane.) In the analysis of data that goes into producing Fig. 2.9, v_2 has been subtracted but v_3 has not. If v_3 is nonzero in individual events, the consequence in Fig. 2.9 would be an excess at angles $\Delta\phi = 0$ and $\Delta\phi = \pi \pm \frac{\pi}{3}$. And, Alver and Roland noted, whenever low peaks are seen at some ϕ_v as in Fig. 2.9, these peaks always seem to occur at $\phi_v \approx \pi/3$, independent of any cuts that are made that should change the velocity of the energetic particle on the away side and consequently should change the putative Mach angle. The preliminary model calculations in Ref. [526] suggest that the incident nuclei are sufficiently asymmetric that the “triangular flow”, which is to say v_3 , that is produced is sufficient to explain the features in Fig. 2.9 that had previously been attributed to the excitation of a Mach cone. As of this writing, it remains to be seen whether once the event-by-event v_3 is quantified and subtracted, there will be any remaining evidence for a Mach cone [526].

Notwithstanding these recent developments, it remains an interesting question of principle to understand in quantitative terms how strong a Mach cone is induced by the passage of an energetic point-like quark through strongly coupled plasma. Remarkably, every one of the difficulties associated with answering this problem in QCD, sketched above, can be addressed for the case of an energetic heavy quark propagating through the strongly coupled plasma of $\mathcal{N} = 4$ SYM theory. In this section we shall review this calculation, which is done via the dual gravitational description of the phenomenon. As in Sections 7.1 and 7.2, we shall assume that the relevant physics is strongly coupled at all length scales, treating the problem entirely within strongly coupled $\mathcal{N} = 4$ SYM theory. In this calculation, the AdS/CFT correspondence is used to determine the stress tensor of the medium, excited by the passing energetic quark, at all length scales. This dynamical computation will allow us to quantify to what extent hydrodynamics can be used to describe the response of the strongly coupled plasma of this theory to the disturbance produced by the energetic quark, as well as to study

the relaxation of the initially far-from-equilibrium disturbance. This calculation applies to quarks with mass M whose velocity respects the bound (7.18).

7.3.1 Hydrodynamic preliminaries

From the point of view of hydrodynamics, the disturbance of the medium induced by the passage of an energetic probe must be described by adding some source to the conservation equation:

$$\partial_\mu T^{\mu\nu}(x) = J^\nu(x) . \quad (7.76)$$

As we have stressed above, we do not know the functional form of the source, since it not only involves the way in which energy is lost by the energetic particle but also how this energy is thermalized and how it is incorporated into the medium. The source will in general depend not only on the position of the quark but also on its velocity. In this subsection, we will use general considerations valid in any hydrodynamic medium to constrain the functional form of the source. From eq. (7.76) it is clear that the amount of energy-momentum deposited in the plasma is given by.³

$$\frac{dP^\nu}{dt} = \int d^3x J^\nu(x) . \quad (7.77)$$

We now attempt to characterize the hydrodynamic modes that can be excited in the plasma due to the deposition of the energy (7.77). We will assume, for simplicity, that the perturbation on the background plasma is small. We will also assume than the background plasma is static. The modification of the stress tensor

$$\delta T^{\mu\nu} \equiv T^{\mu\nu} - T_{\text{background}}^{\mu\nu} \quad (7.78)$$

satisfies a linear equation.

Since in the hydrodynamic limit the stress tensor is characterized by the local energy density, ϵ , and the three components of the fluid spatial velocity, u^i , there are only 4 independent fields, which can be chosen to be

$$\mathcal{E} \equiv \delta T^{00} \quad \text{and} \quad S^i \equiv \delta T^{0i} . \quad (7.79)$$

Using the hydrodynamic form of the stress tensor, (2.10), all other stress tensor components can be expressed as a function of these variables. Since we have assumed that these perturbations are small, all the stress-tensor components can be expanded to first order in the four independent fields (7.79).

In Fourier space, keeping the shear viscosity correction, the linearized form of the equations (7.76) for the mode with a wave vector \mathbf{q} that has the magnitude $q \equiv |\mathbf{q}|$ take the form

$$\begin{aligned} \partial_t \mathcal{E} + iq S_L &= J^0 , \\ \partial_t S_L + ic_s^2 q \mathcal{E} + \frac{4}{3} \frac{\eta}{\varepsilon_0 + p_0} q^2 S_L &= J_L , \\ \partial_t \mathbf{S}_T + \frac{\eta}{\varepsilon_0 + p_0} q^2 \mathbf{S}_T &= \mathbf{J}_T , \end{aligned} \quad (7.80)$$

³ We note as an aside that if the source moves supersonically, one component of its energy loss is due to the emission of sound waves. This is conventionally known as sonic drag, and is a part of the energy loss computed in Section 7.1.

where $\mathbf{S} = S_L \mathbf{q}/q + \mathbf{S}_T$, $\mathbf{J} = J_L \mathbf{q}/q + \mathbf{J}_T$, L and T refer to longitudinal and transverse relative to the hydrodynamical wave-vector \mathbf{q} and ε_0 , p_0 , $c_s = \sqrt{dp/d\varepsilon}$ and η are the energy density, pressure, speed of sound and shear viscosity of the unperturbed background plasma. We observe that the longitudinal and transverse modes are independent. This decomposition is possible since the homogeneous equations have a $SO(2)$ symmetry corresponding to rotation around the wave vector \mathbf{q} . The spin zero (longitudinal) and spin one (transverse) modes correspond to the sound and diffusion mode respectively. (The spin two mode is a subleading perturbation in the gradient expansion, since its leading contribution is proportional to velocity gradients.) After combining the first two equations of Eq. (7.80) and doing a Fourier transformation, we find

$$\begin{aligned} \left(\omega^2 - c_s^2 q^2 + i \frac{4}{3} \frac{\eta}{\varepsilon_0 + p_0} q^2 \omega \right) S_L &= i c_s^2 q J^0 + i \omega J_L, \\ \left(i \omega - \frac{\eta}{\varepsilon_0 + p_0} q^2 \right) \mathbf{S}_T &= -\mathbf{J}_T. \end{aligned} \quad (7.81)$$

The sound mode (S_L) satisfies a wave equation and propagates with the speed of sound while the diffusion mode (\mathbf{S}_T), which does not propagate, describes the diffusion of transverse momentum as opposed to wave propagation. We also note that only the sound mode results in fluctuations of the energy density, while the diffusion mode involves only momentum densities (the S^i of Eq. (7.79)). In the linear approximation that we are using, the excitation of the diffusion mode produces fluid motion but does not affect the energy density. This result can be further illustrated by expressing the energy fluctuations in terms of the velocity fields

$$\delta T^{00} = \delta \varepsilon + \frac{1}{2} (\varepsilon + P) (\delta v)^2 + \dots \quad (7.82)$$

The second term in this expression corresponds to the kinetic energy contribution of the fluid motion which takes a non-relativistic form due to the small perturbation approximation. This expression is quadratic in the velocity fluctuation, and thus is not described in the linearized approximation. The sound mode corresponds to both compression/rarefaction of the fluid and motion of the fluid; sound waves result in fluctuations of the energy density as a consequence of the associated compression and rarefaction. But, the diffusion mode corresponds to fluid motion only and, to this order, does not affect the energy density.

Solving the linearized hydrodynamic equations (7.80) yields hydrodynamic fields given by

$$\mathcal{E}(t, \mathbf{x}) = \int \frac{d\omega}{2\pi} \frac{d^3 q}{(2\pi)^3} \frac{i q J_L + i \omega J^0 - \Gamma_s q^2 J^0}{\omega^2 - c_s^2 q^2 + i \Gamma_s q^2 \omega} e^{-i\omega t + i \mathbf{q} \cdot \mathbf{x}} \quad (7.83)$$

$$\mathbf{S}_L(t, \mathbf{x}) = \int \frac{d\omega}{2\pi} \frac{d^3 q}{(2\pi)^3} \frac{\mathbf{q}}{q} \frac{c_s^2 i q J^0 + i \omega J_L}{\omega^2 - c_s^2 q^2 + i \Gamma_s q^2 \omega} e^{-i\omega t + i \mathbf{q} \cdot \mathbf{x}} \quad (7.84)$$

$$\mathbf{S}_T(t, \mathbf{x}) = \int \frac{d\omega}{2\pi} \frac{d^3 q}{(2\pi)^3} \frac{-\mathbf{J}_T}{i \omega - D q^2} e^{-i\omega t + i \mathbf{q} \cdot \mathbf{x}} \quad (7.85)$$

where the sound attenuation length and the diffusion constant are

$$\Gamma_s \equiv \frac{4}{3} \frac{\eta}{\varepsilon_0 + p_0} \quad (7.86)$$

$$D \equiv \frac{\eta}{\varepsilon_0 + p_0}. \quad (7.87)$$

We note in passing that the integral of the longitudinal momentum density over all space vanishes.

The hydrodynamic solutions (7.83), (7.84) and (7.85) are only of formal value without any information about the source. And, as we have stressed above, a lot of nonlinear, non-equilibrium physics goes into determining the source as a function of the coordinates. Still, we can make some further progress. If we assume that the energetic quark moves at a constant velocity v for a long time (as would be the case if the quark is either ultrarelativistic or very heavy) then we expect

$$J^\mu(\omega, k) = 2\pi\delta(\omega - \mathbf{v} \cdot \mathbf{q})J_v^\mu(\mathbf{q}) , \quad (7.88)$$

where the factor $\delta(\omega - \mathbf{v} \cdot \mathbf{q})$ comes from Fourier transforming $\delta(\mathbf{x} - \mathbf{v}t)$. We also note that far away from the source, and at sufficiently small q that we can neglect any energy scales characteristic of the medium and any internal structure of the particle moving through the medium, the only possible vectors from which to construct the source are \mathbf{v} and \mathbf{q} . In this regime, we may decompose the source as

$$\begin{aligned} J_v^0(\mathbf{q}) &= e_0(\mathbf{q}) \\ \mathbf{J}_v(\mathbf{q}) &= \mathbf{v} g_0(\mathbf{q}) + \mathbf{q} g_1(\mathbf{q}) . \end{aligned} \quad (7.89)$$

Then, inspection of the solutions (7.83), (7.84) and (7.85) together with the observation that a particle moving with a velocity close to the speed of light loses similar amounts of energy and momentum, shows that, at least for an ultrarelativistic probe, nonvanishing values of $e_0(\mathbf{q})$ must be linked to nonvanishing values of $g_0(\mathbf{q})$. We call this case scenario 1. However, if the interaction of the probe with the plasma is such that both g_0 and e_0 are zero (or parametrically small compared to g_1), from Eqs. (7.77) and (7.89) and since $\mathbf{q} g_1(\mathbf{q})$ is a total derivative, one may mistakenly conclude that the energetic probe has created a disturbance carrying zero energy and momentum. In this scenario, which we shall call scenario 2, the energy and momentum loss are actually quadratic in the fluctuations. These two scenarios lead to disturbances with different characteristics. In scenario 2, only the sound mode is excited while in scenario 1, both the sound and diffusion mode are excited. The correct answer for a given energetic probe may lie in between these two extreme cases.

The phenomenological implications of this analysis depend critically on the degree to which the diffusion mode is excited. This mode leads to an excess of momentum density along the direction of the source which does not propagate out of the region of deposition, but only diffuses away. Therefore, the diffusion mode excited by an energetic quark moving through the plasma corresponds to a wake of moving fluid, trailing behind the quark and moving in the same direction as the quark. In a heavy ion collision, therefore, the diffusion wake excited by the away-side energetic quark will become hadrons at $\Delta\phi \sim \pi$, whereas the Mach cone will become a cone of hadrons with moment at some angle away from $\Delta\phi = \pi$. If most of the energy dumped into the medium goes into the diffusion wake, even if a Mach cone were produced it would be overwhelmed in the final state, and invisible in the data. Only in the case in which the diffusion mode is absent (or sufficiently small) is the formation of a Mach cone reflected in a visible nontrivial correlation of the type seen in Fig. 2.9.

7.3.2 AdS computation

In Section 7.1 we have computed the amount of energy lost by a heavy quark as it plows through the strongly coupled $\mathcal{N} = 4$ SYM theory plasma. In order to address the fate of this energy, we must determine the stress tensor of the gauge theory fluid at the boundary that corresponds to the string (7.9) trailing behind the quark in the bulk. In the dual gravitational theory, this string modifies the metric of the (4+1)-dimensional geometry. That is, it produces gravitational waves. The stress energy tensor of the gauge theory plasma at the boundary is determined by the asymptotic behavior of the bulk metric perturbations as they approach the boundary [505, 509, 510, 513].

The modifications of the 4 + 1-dimensional metric due to the presence of the trailing string are obtained by solving the Einstein equations

$$\mathcal{R}_{\mu\nu} - \frac{1}{2}G_{\mu\nu}(\mathcal{R} + 2\Lambda) = \kappa_5^2 t_{\mu\nu} , \quad (7.90)$$

where $\kappa_5^2 = 4\pi^2 R^3 / N_c^2$ and $\Lambda = 6/R^2$ with R the AdS radius and where $t_{\mu\nu}$ is the five-dimensional string stress tensor, which can be computed from the Nambu-Goto action:

$$t^{\mu\nu} = -\frac{1}{2\pi\alpha'} \int d\tau d\sigma \frac{\sqrt{-h}}{\sqrt{-G}} h^{ab} \partial_a X^\mu(\tau, \sigma) \partial_b X^\nu(\tau, \sigma) \delta^{(5)}(x - X(\tau, \sigma)) , \quad (7.91)$$

where h^{ab} is the induced metric on the string and $X(\tau, \sigma)$ is the string profile. For the case of a trailing string (7.9), the stress tensor is given by

$$\begin{aligned} t_{00} &= s(f + v^2 z^4 / z_0^4) , \\ t_{0i} &= -s v_i , \\ t_{0z} &= -s v^2 z^2 / z_0^2 f , \\ t_{zz} &= s(f - v^2) / f^2 , \\ t_{ij} &= s v_i v_j , \\ t_{iz} &= s v_i z^2 / z_0^2 f , \end{aligned} \quad (7.92)$$

where

$$s = \frac{z\gamma\sqrt{\lambda}}{2\pi R^3} \delta^3(\mathbf{x} - vt - \zeta(z)) , \quad (7.93)$$

with $\zeta(z)$ the string profile (7.9). After solving the Einstein equations (7.90) with the string stress tensor (7.92), the expectation value of the boundary stress tensor is then obtained by following the prescription (5.42), namely by performing functional derivatives of the Hilbert action evaluated on the classical solution with respect to the boundary metric.

We need to analyze the small fluctuations on top of the background AdS black hole metric. Denoting these fluctuations by $h_{\mu\nu}$ and the background metric by $g_{\mu\nu}$, the left-hand side of the Einstein equations (7.90) are given to leading order in $h_{\mu\nu}$ by

$$\begin{aligned} &-D^2 h_{\mu\nu} + 2D^\sigma D_{(\mu} h_{\nu)\sigma} - D_\mu D_\nu h + \frac{8}{R^2} h_{\mu\nu} \\ &+ \left(D^2 h - D^\sigma D^\delta h_{\sigma\delta} - \frac{4}{R^2} h \right) g_{\mu\nu} = 0 , \end{aligned} \quad (7.94)$$

with D_μ the covariant derivative with respect to the full metric, namely $g_{\mu\nu} + h_{\mu\nu}$. This equation has a gauge symmetry

$$h_{\mu\nu} \rightarrow h_{\mu\nu} + D_\mu \xi_\nu + D_\nu \xi_\mu, \quad (7.95)$$

inherited from reparameterization invariance, that, together with 5 constraints from the linearized Einstein equations (7.94), reduces the number of degrees of freedom from 15 to 5. It is therefore convenient to introduce gauge invariant combinations which describe the independent degrees of freedom. These can be found after Fourier transforming the $(3+1)$ -dimensional coordinates. The gauge invariants can be classified by how they transform under $SO(2)$ rotations around the wave vector \mathbf{q} . Upon introducing $H_{\mu\nu} = z^2 h_{\mu\nu}/R^2$, one possible choice of gauge invariants is given by [509, 513]

$$\begin{aligned} Z_{(0)} &= q^2 H_{00} + 2\omega q H_{0q} + \omega^2 H_{qq} + \frac{1}{2} [(2-f)q^2 - \omega^2] H \\ Z_{(1)\alpha} &= (H'_{0\alpha} - i\omega H_{\alpha 5}) \\ Z_{(2)\alpha\beta} &= \left(H_{\alpha\beta} - \frac{1}{2} H \delta_{\alpha\beta} \right) \end{aligned} \quad (7.96)$$

where $q \equiv |\mathbf{q}|$, $\hat{q} \equiv \mathbf{q}/q$, $H_{0q} \equiv H_{0i}\hat{q}^i$, $H_{qq} \equiv H_{ij}\hat{q}^i\hat{q}^j$, α and β (which are each either 1 or 2) are space coordinates transverse to \hat{q} , $'$ means ∂_z , and $H \equiv H_{\alpha\alpha}$. When written in terms of these gauge invariants, the Einstein equations (7.94) become three independent equations for $Z_{(0)}$, $Z_{(1)\alpha}$ and $Z_{(2)\alpha\beta}$, which correspond to the spin zero, one and two fluctuations of the stress tensor. We focus on the spin zero and spin one fluctuations, since these are the relevant modes in the hydrodynamic limit. Their equations of motion are given by

$$Z''_{(1)\alpha} + \frac{zf' - 3f}{zf} Z'_{(1)\alpha} + \frac{3f^2 - z(zq^2 + 3f')f + z^2\omega^2}{z^2 f^2} Z_{(1)\alpha} = S_{(1)\alpha} \quad (7.97)$$

$$\begin{aligned} Z''_{(0)} + \frac{1}{u} \left[1 + \frac{uf'}{f} + \frac{24(q^2 f - \omega^2)}{q^2(uf' - 6f) + 6\omega^2} \right] Z'_{(0)} \\ + \frac{1}{f} \left[-q^2 + \frac{\omega^2}{f} - \frac{32q^2 z^6 / z_0^8}{q^2(uf' - 6f) + 6\omega^2} \right] Z_{(0)} = S_{(0)}, \end{aligned} \quad (7.98)$$

where the sources are combinations of the string stress tensor and its derivatives. Choosing one of the transverse directions (which we shall denote by $\alpha = 1$) to lie in the (\mathbf{v}, \mathbf{q}) plane, the source for the trailing string is given explicitly by

$$\begin{aligned} S_{(1)1} &= \frac{2\kappa_5^2 \gamma \sqrt{\lambda}}{R^3} \frac{v q_\perp}{qf} \delta(\omega - \mathbf{v} \cdot \mathbf{q}) e^{-i\mathbf{q} \cdot \zeta}, \\ S_{(1)2} &= 0, \end{aligned} \quad (7.99)$$

$$\begin{aligned} S_{(0)} &= \frac{\kappa_5^2 \gamma \sqrt{\lambda}}{3R^3} \frac{q^2(v^2 + 2) - 3\omega^2}{q^2} \frac{z [q^4 z^8 + 48iq^2 z_0^2 z^5 - 9(q^2 - \omega^2)^2 z_0^8]}{f(fq^2 + 2q^2 - 3\omega^2)z_0^8} \\ &\quad \times \delta(\omega - \mathbf{v} \cdot \mathbf{q}) e^{-i\mathbf{q} \cdot \zeta}, \end{aligned} \quad (7.100)$$

where q_\perp is the magnitude of the component of \mathbf{q} perpendicular to \mathbf{v} . The boundary action can be expressed in terms of the gauge invariants $Z_{(1)\alpha}$ and $Z_{(0)}$ plus certain counterterms

(terms evaluated at the boundary). This procedure, which can be found in Ref. [513], is somewhat cumbersome but straightforward, and we shall not repeat it here. Once this is achieved, the stress tensor components can be obtained from the classical solution to (7.97)-(7.98), following the prescription (5.42).

To find the classical solution to (7.97)-(7.98) we must specify boundary conditions. Since the quark propagates in flat space, the metric fluctuations must vanish at the boundary. Also, since we are interested in the response of the medium, the solution must satisfy retarded boundary condition, meaning that at the horizon it must be composed only of infalling modes. Thus, we may construct the Green's function

$$G_s(z, z') = \frac{1}{W_s(z')} (\theta(z' - z)g_s^n(z)g_s^i(z') + \theta(z - z')g_s^i(z)g_s^n(z')) \quad (7.101)$$

where the subscript s which can be 0 or 1 denotes the spin component and g_s^n and g_s^i denote the normalizable and infalling solutions to the homogeneous equations obtained by setting the left-hand side of (7.97) equal to zero. W_s is the Wronskian of the two homogeneous solutions. The full solution to (7.97) may be then written as

$$Z_s(z) = \int_0^{z_h} dz' G_s(z, z') S_s(z') \quad (7.102)$$

Close to the boundary, these solutions behave as

$$\begin{aligned} Z_{(0)} &= z^3 Z_{(0)}^{[3]} + z^4 Z_{(0)}^{[4]} + \dots \\ \vec{Z}_{(1)} &= z^2 \vec{Z}_{(1)}^{[2]} + z^3 \vec{Z}_{(1)}^{[3]} + \dots \end{aligned} \quad (7.103)$$

The components $Z_{(0)}^{[3]}$ and $Z_{(1)}^{[2]}$ can be computed analytically and are temperature independent. They yield a divergent contribution to the boundary stress tensor. However, this contribution is analytic in q and, thus, has δ -function support at the position of the heavy quark. This divergent contribution is the contribution of the heavy quark mass to the boundary theory stress tensor. The response of the boundary theory gauge fields to the disturbance induced by the passing energetic quark is encoded in the components $Z_{(0)}^{[4]}$ and $Z_{(1)}^{[3]}$, which must be computed numerically. After expressing the boundary actions in terms of gauge invariants, the nondivergent spin zero and one components of the boundary stress tensor are given by

$$\mathcal{T}_0 = \frac{4q^2}{3\kappa_5^2(q^2 - \omega^2)^2} Z_{(0)}^{[4]} + D + \varepsilon_0, \quad (7.104)$$

$$\vec{\mathcal{T}}_1 = -\frac{L^3}{2\kappa_5} \vec{Z}_{(1)}^{[3]}. \quad (7.105)$$

where $\mathcal{T}_0 = T^{00}$ and $\vec{\mathcal{T}}_1 = T^{0a}\hat{e}_a$, with \hat{e}_a the spatial unit vectors orthogonal to the spatial momentum \mathbf{q} , and where the counterterm D is a complicated function of ω and q that depends on the quark velocity and the plasma temperature and that is given in Ref. [513].

Results from Ref. [513] on the numerical computation of the disturbance in the gauge theory plasma created by a supersonic quark moving with speed $v = 0.75$ are shown in

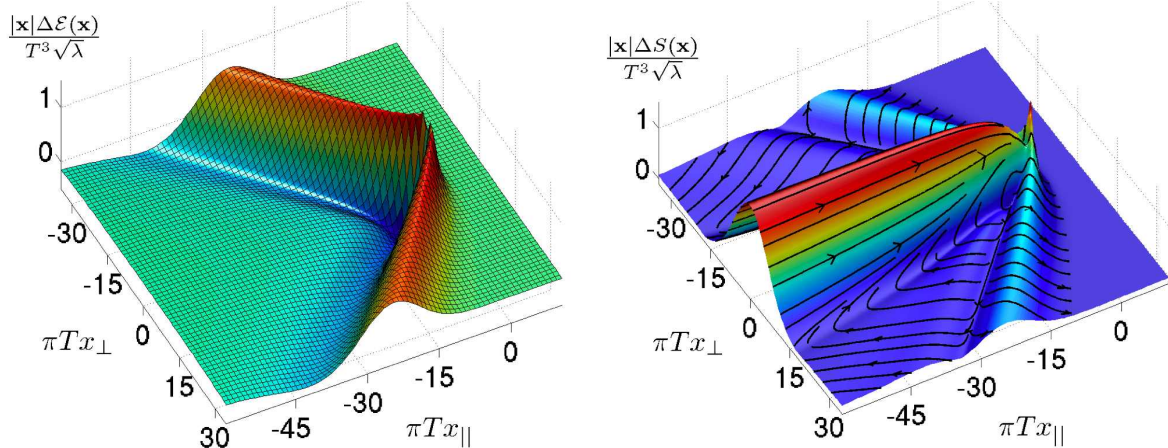


Figure 7.4: Energy density (left) and momentum flux (right) induced by the passage of a supersonic heavy quark moving through the strongly coupled $\mathcal{N} = 4$ SYM theory plasma in the x_{\parallel} direction with speed $v = 0.75$. ($\Delta\varepsilon(\mathbf{x})$ is the difference between $\varepsilon(\mathbf{x})$ and the equilibrium energy density; since $\mathbf{S} = 0$ in equilibrium, $\Delta\mathbf{S}(\mathbf{x})$ is simply $\mathbf{S}(\mathbf{x})$.) The flow lines on the surface are flow lines of $\Delta\mathbf{S}(\mathbf{x})$. These disturbances are small compared to the background energy density and pressure of the plasma (both of which are $\propto N_c^2$). The perturbation is small and it is well described by linearized hydrodynamics everywhere except within a distance $R \approx 1.6/T$ from the quark. Since the perturbation is small, the kinetic energy contribution of the diffusion mode to the energy density is suppressed by N_c^2 and, thus, it does not contribute in the left panel.

Fig. 7.4. The left panel shows the energy density of the disturbance and clearly demonstrates that a Mach cone has been excited by the supersonic quark. The front is moving outwards at the Mach angle Θ_M , where $\cos \Theta_M = c_s/v = 4/(3\sqrt{3})$. Recall from our general discussion above that fluid motion is invisible in the energy density, to the linear order at which we are working; the energy density is nonzero wherever the fluid is compressed. Thus, the Mach cone is made up of sound modes, as expected. In the right panel of Fig. 7.4, we see the density of fluid momentum induced by the supersonic quark. This figure reveals the presence of a sizable wake of moving fluid behind the quark, a wake that is invisible in the energy density and is therefore made up of moving fluid without any associated compression, meaning that it is made up of diffusion modes. We conclude that the supersonic quark passing through the strongly coupled plasma excites both the sound mode and the diffusion mode, meaning that the interaction of the quark with the plasma is as in what we called scenario 1 above. Quantitatively, it turns out that the momentum carried by the sound waves is greater than that carried by the diffusion wake, but only by a factor of $1 + v^2$ [511].

Since hydrodynamics describes the long wavelength limit of the stress tensor excitation, it is reasonable to find a Mach cone at long distances. And, since the gravitational equations whose solution we have reviewed are linear, the long distance behavior of the gauge theory fluid must be described by linearized hydrodynamics. It is easy to justify the linearization from the point of view of the field theory: the background plasma has an energy density that is proportional to N_c^2 while that of the perturbation is proportional to the number of flavors, which is just $N_f = 1$ in the present case since we are considering only one quark. The strong coupling computation leads to a perturbation of magnitude $N_f\sqrt{\lambda}$. Thus, the energy density

of the fluctuations are suppressed by $\sqrt{\lambda}/N_c^2$ with respect to that of the background plasma, justifying the linearized treatment. Remarkably, it turns out that disturbances like those in Fig. 7.4 are well described by hydrodynamics everywhere except within $\approx 1.6/T$ of the position of the quark [513]. So, the calculation that we have reviewed in this section is important for two reasons. First, it demonstrates that a *point-like* probe passing through the strongly coupled plasma does indeed excite hydrodynamic modes. And, second, it demonstrates that in the strongly coupled plasma, the resulting disturbance relaxes to a hydrodynamic excitation in local thermal equilibrium surprisingly close to the probe.

The observation that point particles moving through the strongly coupled fluid excite sound waves, which are collective excitations, is at odds with intuition based upon the interaction of, say, electrons with water. In this example, most of the energy lost by the electron is transferred to photons and not to the medium. These photons, in turn, have long mean free paths and dissipate their energy far away from the electron (or escape the medium entirely). Thus, the effective size of the region where energy is dissipated is very large, given by the photon mean free path. Hydrodynamics will only describe the physics on longer length scales than this. The reason that no Mach cone is formed is that the length scale over which the energy is deposited is long compared to the length scale over which the electron slows and stops. The situation is similar in weakly coupled gauge theory plasmas; even though the gauge modes in these theories do interact, they still have long mean free paths proportional to $1/g^4$. In sharp contrast, in the strongly coupled plasma of $\mathcal{N} = 4$ SYM theory there are no long-lived quasiparticle excitations (leave apart photons) that could transport the energy deposited by the point-like particle over long distances. Instead, all the energy lost by the point-like probe is dumped into collective hydrodynamic modes over a characteristic length scale $\sim 1/T$, which is the only length scale in this conformal plasma.

7.3.3 Implications for heavy ion collisions

The calculation that we have reviewed in this section suggests that a high energy quark plowing through the strongly coupled plasma produced in heavy ion collisions at RHIC should excite a Mach cone. As we argued just above, this phenomenon is not expected in a weakly coupled plasma. The Mach cone should have consequences that are observable in the soft particles on the away-side of a high energy trigger hadron. However, for a hydrodynamic solution like that in Fig. 7.4, it turns out that the diffusion wake contains enough momentum flux along the direction of the energetic particle to “fill in” the center of the Mach cone, meaning that the Mach cone is not sufficiently prominent as to result in peaks in the particle distribution at $\Delta\phi = \pi \pm \Theta_M$ [138, 518]. As we discussed above, the observed peaks at $\Delta\phi = \pi \pm \phi_v$ receive a significant contribution from the event-by-event v_3 due to event-by-event fluctuations that introduce “triangularity”. Detecting evidence for Mach cones in heavy ion collisions will require careful subtraction of these effects from the data on the one hand, and careful theoretical analysis of the effects of the rapid expansion of the fluid produced in heavy ion collisions on the putative Mach cones on the other.

7.3.4 Disturbance excited by a moving quarkonium meson

Strong coupling calculations like that of the disturbance excited by an energetic quark moving through the plasma of $\mathcal{N} = 4$ SYM theory can help guide the construction of more phenomenological models of the coupling of energetic particles to hydrodynamic modes. To further that end, we close with an example which shows that not all probes behave in the same way.

As we shall describe in Section 7.7, a simple way of modelling a “quarkonium” meson made from a heavy quark and antiquark embedded in the strongly coupled plasma of $\mathcal{N} = 4$ SYM theory is to consider a string with both ends at the boundary — the ends representing the quark and antiquark. We shall see in Section 7.7 that even when this string is moving through the plasma, it hangs straight downward into the AdS black hole metric, rather than trailing behind as happens for the string hanging downward from a single moving quark. The fact that the “U” of string hangs straight down and does not trail behind the moving quark and antiquark implies that the heavy quarkonium meson moving through the strongly coupled plasma does not lose any energy, at least at leading order. The energy loss of such a meson has been computed and is in fact nonzero but is suppressed by $1/N_c^2$ [527].

Despite the fact that the leading order quarkonium energy loss vanishes, the leading order disturbance of the fluid through which the meson is moving does not vanish [528]. Instead, the meson excites a Mach cone with no diffusion wake, providing an example of what we called scenario 2 at the end of Section 7.3.1. It is as though the moving meson “dresses itself” with a Mach cone, and then the meson and its Mach cone propagate through the fluid without dissipation, to leading order. To illustrate this point, the metric fluctuation and consequent boundary stress tensor induced by a semiclassical string with both ends on the boundary moving with a velocity \mathbf{v} has been calculated [528]. For a string with the two endpoints aligned along the direction of motion and separated by a distance l the long distance part (low momentum) part of the associated stress tensor is given by

$$\delta T^{00} = \frac{\Pi}{q^2 - 3(q \cdot \mathbf{v})^2} \left(-q^2 (1 + 2v^2\sigma) - 3v^2 (q \cdot \mathbf{v})^2 (1 - 2\sigma) \right) \quad (7.106)$$

$$\delta T^{0i} = \frac{\Pi}{k^2 - 3(q \cdot \mathbf{v})^2} 2q \cdot \mathbf{v} (1 - (1 - v^2)\sigma) q^i + 2\Pi \sigma v^i \quad (7.107)$$

$$\delta T^{ij} = \frac{\Pi}{k^2 - 3(q \cdot \mathbf{v})^2} 2(q \cdot \mathbf{v})^2 (-1 + (1 - v^2)\sigma) \delta^{ij} - \frac{\Pi}{v^2} (1 + 2v^2\sigma) \frac{v^i v^j}{v^2} \quad (7.108)$$

where $\sigma = \sigma(l, T)$ is a dimensionless function of the length of the meson and the temperature and the prefactor takes the form

$$\Pi = \sqrt{\lambda} \frac{F(lT)}{l}, \quad (7.109)$$

with F a dimensional function.

The expression (7.106) clearly shows that all the spin 0,1 and 2 components of the stress tensor are excited. The spin 0 components are multiplied by the sound propagator, signaling the emission of sound waves. (Note that in the low- q limit the width of the sound pole

vanishes.) The spin 1 component corresponds to the terms proportional to the velocity of the particle v^i . These terms are analytic in q ; in particular it seems that there is no pole contribution from the diffusive mode. More careful analysis shows that the diffusive mode decays faster than that excited by a quark probe.

The magnitude of the disturbance in the strongly coupled plasma that is excited by a passing quarkonium meson is no smaller than that excited by a passing quark. However, the total integral of the energy and momentum deposited is zero, as can be seen by multiplying the momentum densities by $\omega = \mathbf{v} \cdot \mathbf{q}$ and taking the limit $q \rightarrow 0$. This is consistent with the fact that, to the order at which this calculation has been done, the meson does not lose any energy. This is an interesting example since it indicates that the loss of energy and the excitation of hydrodynamic modes are distinct phenomena, controlled by different physics. This example also illustrates the value of computations done at strong coupling in opening ones eyes to new possibilities: without these calculations it would have been very hard to guess or justify that such a separation in magnitude between the strength of the hydrodynamic fields excited by a probe and the energy lost by that probe could be possible. It would be interesting to analyze the soft particles in heavy ion collisions in which a high transverse momentum quarkonium meson is detected, to see whether there is any hint of a Mach cone around the meson — in this case without the complication of soft particles from a diffusion wake filling in the cone.

7.4 Stopping light quarks

As we have discussed extensively in Section 2.3, the dominant energy loss process for a parton moving through the QCD plasma with energy E in the limit in which $E \rightarrow \infty$ is gluon radiation, and in this limit much (but not all; see Section 7.5) of the calculation can be done at weak coupling. However, since it is not clear how quantitatively reliable the $E \rightarrow \infty$ approximation is for the jets produced in RHIC collisions, it is also worth analyzing the entire problem of parton energy loss and jet quenching at strong coupling to the degree that is possible. For the case of a heavy quark propagating through the strongly coupled plasma of $\mathcal{N} = 4$ SYM theory, this approach has been pursued extensively, yielding the many results that we have reviewed in the previous three subsections. Less work has been done on the energy loss of an energetic light quark or gluon in the $\mathcal{N} = 4$ SYM plasma, in particular since they do not fragment into anything like a QCD jet. (As we shall review in Section 7.6, there are no true jets in $\mathcal{N} = 4$ SYM theory, although it is possible to cook up a collimated beam of radiation.) It is nevertheless worth asking how a light quark or gluon loses energy in the $\mathcal{N} = 4$ SYM plasma in the hope that even if this is not a good model for jets and their quenching in QCD some qualitative strong-coupling benchmarks against which to compare experimental results may be obtained. This program has been pursued in Refs. [529–531].

As introduced in Section 5.5 and described extensively in Section 8, dynamical quarks can be introduced into $\mathcal{N} = 4$ SYM theory by introducing a D7-brane that fills the 3 + 1 Minkowski dimensions and fills the fifth dimension from the boundary at $z = 0$ down to $z = z_q$. The mass of the (heavy) quarks that this procedure introduces in the gauge theory is $\sqrt{\lambda}/(2\pi z_q)$. Light quarks are obtained by taking $z_q \rightarrow \infty$, meaning that the D7-brane fills all of the z -dimension. At $T \neq 0$, what matters is that the D7-brane fills the z -dimension all

the way down to, and below, the horizon. In this set up, a light quark-antiquark pair moving back-to-back each with some initial high energy can be modelled as a quark and anti-quark located at some depth z that are moving apart from each other in, say, the x direction and that are connected by a string [530, 531]. The quark and antiquark must be within the D7 brane, but since this D7 brane fills all of z there is nothing stopping them from falling to larger z as they fly apart from each other, and ultimately there is nothing stopping them from falling into the horizon. It should be evident from this description that there is an arbitrariness to the initial condition: At what z should the quark and antiquark be located initially? What should the string profile be initially? What should the initial profile of the velocity of the string be? These choices correspond in the gauge theory to choices about the initial quantum state of the quark-antiquark pair and the gauge fields surrounding them. And, there is no known way to choose these initial conditions so as to obtain a QCD-like jet so the choices made end up being arbitrary. (The analogous set up for a back-to-back pair of high energy gluons [529] involves a doubled loop of string, rather than an open string with a quark and antiquark at its ends.)

Ambiguities about the initial conditions notwithstanding, several robust qualitative insights have been obtained from these calculations. First, the quark and the antiquark always fall into the horizon after travelling some finite distance x_{stopping} . (The string between them falls into the horizon also.) x_{stopping} corresponds in the gauge theory to the stopping distance for the initially energetic quark, namely the distance that it takes this quark to slow down, thermalize, and equilibrate with the bulk plasma — the gauge theory analogue of falling into the horizon. This is qualitatively reminiscent of the discovery at RHIC of events with a single jet (manifest as a high p_T hadron that is triggered on) but no jet (no high p_T hadrons) back-to-back with it.

Second, although x_{stopping} does depend on details of the initial conditions, the dominant dependence is that it scales like $E^{1/3}$, where E is the initial energy of the quark [529, 531–533]. More precisely, upon analyzing varied initial conditions the maximum possible stopping distance is given by [531]

$$x_{\text{stopping}} = \frac{C}{T} \left(\frac{E}{T\sqrt{\lambda}} \right)^{1/3}, \quad (7.110)$$

with $C \approx 0.5$. If there is a regime of E and T in which it is reasonable to treat the entire problem of jet quenching at strong coupling, and if in this regime the droplet of plasma produced in a heavy ion collision is large enough and lives long enough that it can stop and thermalize an initial parton with energy E that would in vacuum have become a jet, then the scaling (7.110) has interesting qualitative consequences. For example, if this scaling applies to collisions with two different collision energies $\sqrt{s_1}$ and $\sqrt{s_2}$, yielding plasmas that form at different temperatures T_1 and T_2 , then jets in these two experiments whose energies satisfy $E_1/E_2 \sim (T_1/T_2)^4$ should have similar observed phenomenology. Turning this speculation into semi-quantitative expectations for experimental observables requires careful study of jet stopping in a realistic model of the dynamics in space and time of the expanding droplet of plasma produced in a heavy ion collision.

Third, a light quark with initial energy E that loses this energy over a distance x_{stopping} loses most of its energy near the end of its trajectory, where it thermalizes (falls into the horizon) [531]. This pattern of energy loss is reminiscent of the ‘Bragg peak’ that characterizes

the energy loss of a fast charged particle in ordinary matter, where the energy loss has a pronounced peak near the stopping point. It is quite different from the behavior of a heavy quark in strongly coupled plasma which, as we saw in Section 7.1, loses energy at a rate proportional to its momentum, making it reasonable to expect that a heavy quark that slows from a high velocity to a stop loses more energy earlier in its trajectory than later.

It will be very interesting to see how these insights fare when compared with results on jet quenching in heavy ion collisions at the LHC, and in particular to comparisons between such results and results at RHIC energies.

7.5 Calculating the jet quenching parameter

As we have described in Section 2.3, when a parton with large transverse momentum is produced in a hard scattering that occurs within a heavy ion collision, the presence of the medium in which the energetic parton finds itself has two significant effects: it causes the parton to lose energy and it changes the direction of the parton’s momentum. The latter effect is referred to as “transverse momentum broadening”. In the high parton energy limit, as established first in Refs. [139–141], the parton loses energy dominantly by inelastic processes that are the QCD analogue of bremsstrahlung: the parton radiates gluons as it interacts with the medium. It is crucial to the calculation of this radiative energy loss process that the incident hard parton, the outgoing parton, and the radiated gluons are all continually being jostled by the medium in which they find themselves: they are *all* subject to transverse momentum broadening. The transverse momentum broadening of a hard parton is described by $P(k_\perp)$, defined as the probability that after propagating through the medium for a distance L the hard parton has acquired transverse momentum k_\perp . For later convenience, we shall choose to normalize $P(k_\perp)$ as follows:

$$\int \frac{d^2 k_\perp}{(2\pi)^2} P(k_\perp) = 1 . \quad (7.111)$$

From the probability density $P(k_\perp)$, it is straightforward to obtain the mean transverse momentum picked up by the hard parton per unit distance travelled (or, equivalently in the high parton energy limit, per unit time):

$$\hat{q} \equiv \frac{\langle k_\perp^2 \rangle}{L} = \frac{1}{L} \int \frac{d^2 k_\perp}{(2\pi)^2} k_\perp^2 P(k_\perp) . \quad (7.112)$$

$P(k_\perp)$, and consequently \hat{q} , can be evaluated for a hard quark or a hard gluon. In the calculation of radiative parton energy loss [140–145,155] that we have reviewed in Section 2.3 and that is also reviewed in Refs. [150,154,157–159,168,169,173], \hat{q} for the radiated gluon plays a central role, and this quantity is referred to as the “jet quenching parameter”. Consequently, \hat{q} should be thought of as a (or even the) property of the strongly coupled medium that is “measured” (perhaps constrained is a better phrase) by radiative parton energy loss and hence jet quenching. But, it is important to note that \hat{q} is *defined* via transverse momentum broadening only. Radiation and energy loss do not arise in its definition, although they are central to its importance.

The BDMPS calculation of parton energy loss in QCD involves a number of scales which must be well-separated in order for this calculation to be relevant. The radiated gluons have energy of order $\omega_c \sim \hat{q}L^2$ and transverse momenta of order $\sqrt{\hat{q}L}$. Both these scales must be much less than E and much greater than T . And, α_S evaluated at both these scales must be small enough that physics at these scales is weakly coupled, even if physics at scales of order T is strongly coupled. In heavy ion collisions at RHIC, with the highest energy partons having E only of order many tens of GeV, this separation of scales can be questioned. It is expected that in heavy ion collisions at the LHC it will be possible to study the interaction of partons with energies of order a few hundred GeV, improving the reliability of the calculations reviewed in this subsection. In the RHIC regime, and particularly for heavy quarks, the alternative approach of the previous subsection in which the entire problem is assumed to be strongly coupled and parton energy loss occurs via drag is just as plausibly relevant. In this subsection, we shall review the calculation of \hat{q} — the property of the plasma that describes transverse momentum broadening directly and, in the high parton energy limit in which the relevant scales are well separated, controls the transverse momentum and the energy of the gluon radiation that dominates parton energy loss.

It has been shown via several different calculations done via conventional field theoretical methods [150, 151, 534] that the transverse momentum broadening of a hard parton in the \mathcal{R} representation of $SU(N)$ takes the form

$$P(k_\perp) = \int d^2x_\perp e^{-ik_\perp \cdot x_\perp} \mathcal{W}_\mathcal{R}(x_\perp) \quad (7.113)$$

with

$$\mathcal{W}_\mathcal{R}(x_\perp) = \frac{1}{d(\mathcal{R})} \left\langle \text{Tr} \left[W_\mathcal{R}^\dagger[0, x_\perp] W_\mathcal{R}[0, 0] \right] \right\rangle \quad (7.114)$$

where

$$W_\mathcal{R}[x^+, x_\perp] \equiv P \left\{ \exp \left[ig \int_0^{L^-} dx^- A_\mathcal{R}^+(x^+, x^-, x_\perp) \right] \right\} \quad (7.115)$$

is the representation- \mathcal{R} Wilson line along the lightcone, $L^- = \sqrt{2}L$ is the distance along the lightcone corresponding to travelling a distance L through the medium, and where $d(\mathcal{R})$ is the dimension of the representation \mathcal{R} . The result (7.113) is similar to (7.48) although the physical context in which it arises is different as is the path followed by the Wilson line, and one of the derivations [150] of (7.113) is analogous to the derivation of (7.48) that we reviewed in Section 7.2. Another derivation [151] proceeds via the use of the optical theorem to relate $P(k_\perp)$ to an appropriate forward scattering matrix element which can then be calculated explicitly via formulating the calculation of transverse momentum broadening in the language of Soft Collinear Effective Theory [535–539]. This derivation in particular makes it clear that (7.113) is valid whether the plasma through which the energetic quark is propagating, i.e. the plasma which is causing the transverse momentum broadening, is weakly coupled or strongly coupled. The result (7.113) is an elegant expression saying that the probability for the quark to obtain transverse momentum k_\perp is simply given by the two-dimensional Fourier transform in x_\perp of the expectation value (7.114) of two light-like Wilson lines separated in the transverse plane by the vector x_\perp . Note that the requirement (7.111) that the probability distribution $P(k_\perp)$ be normalized is equivalent to the requirement that $\mathcal{W}_\mathcal{R}(0) = 1$.




Figure 7.5: The Schwinger-Keldysh contour that must be used in the evaluation of $\mathcal{W}_{\mathcal{R}}(x_{\perp})$. It is similar to that in Fig. 7.2.

It is important to notice that the expectation value of the trace of the product of two light-like Wilson lines that arises in $P(k_{\perp})$ and hence in \hat{q} , namely $\mathcal{W}_{\mathcal{R}}(x_{\perp})$ of (7.114), has a different operator ordering from that in a standard Wilson loop. Upon expanding the exponential, each of the A^+ that arise can be written as the product of an operator and a group matrix: $A^+ = (A^+)^a t^a$. It is clear (for example either by analogy with our discussion around (7.45) in the analysis of momentum broadening of heavy quarks or from the explicit derivation in Ref. [151]) that in $\mathcal{W}_{\mathcal{R}}(x_{\perp})$ both the operators and the group matrices are path ordered. In contrast, in a conventional Wilson loop the group matrices are path ordered but the operators are time ordered. Because the operators in (7.114) are path ordered, the expectation value in (7.114) should be described using the Schwinger-Keldysh contour in Fig. 7.5 with one of the light-like Wilson lines on the $\text{Im } t = 0$ segment of the contour and the other light-like Wilson line on the $\text{Im } t = -i\epsilon$ segment of the contour. The infinitesimal displacement of one Wilson line with respect to the other in Fig. 7.5 ensures that the operators from the two lines are ordered such that all operators from one line come before any operators from the other. In contrast, the loop \mathcal{C} for a standard Wilson loop operator lies entirely at $\text{Im } t = 0$, and the operators for a standard Wilson loop are time ordered.

The transverse momentum broadening of a hard parton with energy E is due to repeated interactions with gluons from the medium which, if the medium is in equilibrium at temperature T , carry transverse momenta of order T and lightcone momenta of order T^2/E [151, 540]. The relation (7.113) between $P(k_{\perp})$ (and hence \hat{q}) and the expectation value \mathcal{W} of (7.114) is valid as long as $E \gg \hat{q}L^2$ (which is to say E must be much greater than the characteristic energy of the radiated gluons) even if $\alpha_S(T)$ is in no way small, i.e. it is valid in the large- E limit even if the hard parton is interacting with a strongly coupled plasma and even if the soft interactions that generate transverse momentum broadening are not suppressed by any weak coupling either [151]. However, in this circumstance even though (7.113) is valid it was not particularly useful until recently because there is no known conventional field theoretical *evaluation* of \mathcal{W} for a strongly coupled plasma. (Since lattice quantum field theory is formulated in Euclidean space, it is not well-suited for the evaluation of the expectation value of lightlike Wilson lines.) In this subsection we review the evaluation of \mathcal{W} , and hence \hat{q} , in the strongly coupled plasma of $\mathcal{N} = 4$ SYM theory with gauge group $SU(N_c)$ in the large N

and strong coupling limit using its gravitational dual, namely the AdS Schwarzschild black hole at nonzero temperature [148, 149, 151, 456, 458, 461, 462, 541–544]. The calculation is not simply an application of results reviewed in Section 5.4 both because the operators are path ordered and because the Wilson lines are light-like.

We begin by sketching how the standard AdS/CFT procedure for computing a Wilson loop in the fundamental representation in the large N_c and strong coupling limit, reviewed in Section 5.4, applies to a light-like Wilson loop with standard operator ordering [148, 149], and then below describe how the calculation (but not the result) changes when the operator ordering is as in (7.114). Consider a Wilson loop operator $W(\mathcal{C})$ specified by a closed loop \mathcal{C} in the $(3+1)$ -dimensional field theory, and thus on the boundary of the $(4+1)$ -dimensional AdS space. $\langle W(\mathcal{C}) \rangle$ is then given by the exponential of the classical action of an extremized string worldsheet Σ in AdS which ends on \mathcal{C} . The contour \mathcal{C} lives within the $(3+1)$ -dimensional Minkowski space boundary, but the string world sheet Σ attached to it hangs “down” into the bulk of the curved five-dimensional AdS₅ spacetime. More explicitly, consider two long parallel light-like Wilson lines separated by a distance x_\perp in a transverse direction.⁴ (The string world sheet hanging down into the bulk from these two Wilson lines can be visualized as in Fig. 7.7 below if one keeps everything in that figure at $\text{Im } t = 0$, i.e. if one ignores the issue of operator ordering.) Upon parameterizing the two-dimensional world sheet by the coordinates $\sigma^\alpha = (\tau, \sigma)$, the location of the string world sheet in the five-dimensional spacetime with coordinates x^μ is

$$x^\mu = x^\mu(\tau, \sigma) \tag{7.116}$$

and the Nambu-Goto action for the string world sheet is given by

$$S = -\frac{1}{2\pi\alpha'} \int d\sigma d\tau \sqrt{-\det g_{\alpha\beta}}. \tag{7.117}$$

Here,

$$g_{\alpha\beta} = G_{\mu\nu} \partial_\alpha x^\mu \partial_\beta x^\nu \tag{7.118}$$

is the induced metric on the world sheet and $G_{\mu\nu}$ is the metric of the $(4+1)$ -dimensional AdS₅ spacetime. Denoting by $S(\mathcal{C})$ the classical action which extremizes the Nambu-Goto action (7.117) for the string worldsheet with the boundary condition that it ends on the curve \mathcal{C} , the expectation value of the Wilson loop operator is then given by

$$\langle W(\mathcal{C}) \rangle = \exp [i \{ S(\mathcal{C}) - S_0 \}], \tag{7.119}$$

where the subtraction S_0 is the action of two disjoint strings hanging straight down from the two Wilson lines. In order to obtain the thermal expectation value at nonzero temperature, one takes the metric $G_{\mu\nu}$ in (7.118) to be that of an AdS Schwarzschild black hole (5.30) with a horizon at $r = r_0$ and Hawking temperature $T = r_0/(\pi R^2)$. The AdS curvature radius R and the string tension $1/(2\pi\alpha')$ are related to the 't Hooft coupling in the Yang-Mills theory $\lambda \equiv g^2 N_c$ by $\sqrt{\lambda} = R^2/\alpha'$.

We shall assume that the length of the two light-like lines $L^- = \sqrt{2}L$ is much greater than their transverse separation x_\perp , which can be justified after the fact by using the result for

⁴Note that for a light-like contour \mathcal{C} , the Wilson line (5.63) of $\mathcal{N} = 4$ SYM theory reduces to the familiar (7.115).

$\mathcal{W}(x_\perp)$ to show that the x_\perp -integral in (7.113) is dominated by values of x_\perp that satisfy $x_\perp \ll 1/\sqrt{\hat{q}L} \sim 1/\sqrt{\sqrt{\lambda}LT^3}$. As long as we are interested in $L \gg 1/T$, then $x_\perp \ll 1/(T\lambda^{1/4}) \ll 1/T \ll L$. With $L^- \gg x_\perp$, we can ignore the ends of the light-like Wilson lines and assume that the shape of the surface Σ is translationally invariant along the light-like direction. The action (7.117) now takes the form

$$S = i \frac{\sqrt{2}r_0^2\sqrt{\lambda}L^-}{2\pi R^4} \int_0^{x_\perp/2} d\sigma \sqrt{1 + \frac{r'^2 R^4}{r^4 - r_0^4}} \quad (7.120)$$

where the shape of the worldsheet Σ is described by the function $r(\sigma)$ that satisfies $r(\pm \frac{x_\perp}{2}) = \infty$, which preserves the symmetry $r(\sigma) = r(-\sigma)$, and where $r' = \partial_\sigma r$. The equation of motion for $r(\sigma)$ is then

$$r'^2 = \frac{\gamma^2}{R^4} (r^4 - r_0^4) \quad (7.121)$$

with γ an integration constant. Eq. (7.121) has two solutions. One has $\gamma = 0$ and hence $r' = 0$, meaning that $r(\sigma) = \infty$ for all σ : the surface Σ stays at infinity. Generalizations of this solution have also been studied [545, 546]. We shall see below that such solutions are not relevant. The other solution has $\gamma > 0$. It “descends” from $r(\pm \frac{x_\perp}{2}) = \infty$ and has a turning point where $r' = 0$ which, by symmetry, must occur at $\sigma = 0$. From (7.121), the turning point must occur at the horizon $r = r_0$. Integrating (7.121) gives the condition that specifies the value of γ :

$$\frac{x_\perp}{2} = \frac{R^2}{\gamma} \int_{r_0}^{\infty} \frac{dr}{\sqrt{r^4 - r_0^4}} = \frac{aR^2}{\gamma r_0} \quad (7.122)$$

where we have defined

$$a \equiv \sqrt{\pi} \Gamma(\frac{5}{4}) / \Gamma(\frac{3}{4}) \approx 1.311. \quad (7.123)$$

Putting all the pieces together, we find [148, 149]

$$S = \frac{ia\sqrt{\lambda}TL^-}{\sqrt{2}} \sqrt{1 + \frac{\pi^2 T^2 x_\perp^2}{4a^2}}. \quad (7.124)$$

We see that S is imaginary, because when the contour \mathcal{C} at the boundary is lightlike the surface Σ hanging down from it is spacelike. It is worth noting that S had to turn out to be imaginary, in order for $\langle W \rangle$ in (7.119) to be real and the transverse momentum broadening $P(k_\perp)$ to be real, as it must be since it is a probability distribution. The surface Σ that we have used in this calculation descends from infinity, skims the horizon, and returns to infinity. Note that the surface descends all the way to the horizon regardless of how small x_\perp is. This is reasonable on physical grounds, as we expect $P(k_\perp)$ to depend on the physics of the thermal medium [148, 149]. We shall see below that it is also required on mathematical grounds: when we complete the calculation by taking into account the nonstandard operator ordering in (7.113), we shall see that only a worldsheet that touches the horizon is relevant [151].

We now consider the computation of (7.114), with its nonstandard operator ordering corresponding to putting one of the two light-like Wilson lines on the $\text{Im } t = 0$ contour in Fig. 7.5 and the other on the $\text{Im } t = -i\epsilon$ contour. The procedure we shall describe is a specific example of the more general discussion of Lorentzian AdS/CFT given recently in Refs. [547–550].




Figure 7.6: Penrose diagrams for Lorentzian ($\text{Im } t = 0$) and Euclidean ($\text{Re } t = 0$) sections of an AdS black hole. In the right panel, the two light-like Wilson lines are points at $r = \infty$, indicated by colored dots. These dots are the boundaries of a string world sheet that extends inward to $r = r_0$, which is at the origin of the Euclidean section of the black hole. In the left panel, the string world sheet and its endpoints at $r = \infty$ are shown at $\text{Re } t = 0$; as $\text{Re } t$ runs from $-\infty$ to ∞ , the string worldsheet sweeps out the whole of quadrant I .

In order to compute (7.114) we first need to construct the bulk geometry corresponding to the $\text{Im } t = -i\epsilon$ segment of the Schwinger-Keldysh contour in Fig. 7.5. For this purpose it is natural to consider the black hole geometry with complex time. In Fig. 7.6, we show two slices of this complexified geometry. The left plot is the Penrose diagram for the fully extended black hole spacetime with quadrant I and III corresponding to the slice $\text{Im } t = 0$ and $\text{Im } t = -\frac{\beta}{2}$ respectively, while the right plot is for the Euclidean black hole geometry, i.e. corresponding to the slice $\text{Re } t = 0$. Note that because the black hole has a nonzero temperature, the imaginary part of t is periodic with the period given by the inverse temperature β . In the left plot the imaginary time direction can be considered as a circular direction coming out of the paper at quadrant I , going a half circle to reach quadrant III and then going into the paper for a half circle to end back at I . In the right plot the real time direction can be visualized as the direction perpendicular to the paper.

The first segment of the Schwinger-Keldysh contour in Fig. 7.5, with $\text{Im } t = 0$, lies at the boundary ($r = \infty$) of quadrant I in Fig. 7.6. The second segment of the Schwinger-Keldysh contour, with $\text{Im } t = -i\epsilon$, lies at the $r = \infty$ boundary of a copy of I that in the left plot of Fig. 7.6 lies infinitesimally outside the paper and in the right plot of Fig. 7.6 lies at an infinitesimally different angle. We shall denote this copy of I by I' . The geometry and metric in I' are identical to those of I . Note that I' and I are joined together at the horizon $r = r_0$, namely at the origin in the right plot of Fig. 7.6. Now, the thermal expectation value (7.114) can be computed by putting the two parallel light-like Wilson lines at the boundaries of I and I' , and finding the extremized string world sheet which ends on both of them. Note that since I and I' meet only at the horizon, the only way for there to be a nontrivial (i.e. connected) string world sheet whose boundary is the two Wilson lines in (7.114) is for such a string world sheet to touch the horizon. Happily, this is precisely the feature of the string world sheet




Figure 7.7: String configuration for the thermal expectation value of (7.114).

found in the explicit calculation that we reviewed above. So, we can use that string world sheet in the present analysis, with the only difference being that half the string world sheet now lies on I and half on I' , as illustrated in Fig. 7.7.⁵

We conclude that the result for the expectation value (7.114), with its nonstandard path ordering of operators, is identical to that obtained in Refs. [148, 149] for a light-like Wilson loop with standard time ordering of operators [151]. That is, in strongly coupled $\mathcal{N} = 4$ SYM theory $\mathcal{W}(x_\perp)$ in the adjoint representation is given by

$$\mathcal{W}_A(x_\perp) = \exp \left[-\sqrt{2}a\sqrt{\lambda}L^{-1}T \left(\sqrt{1 + \frac{\pi^2 T^2 x_\perp^2}{4a^2}} - 1 \right) \right]. \quad (7.125)$$

We have quoted the result for $\mathcal{W}_A(x_\perp)$, which is given by $\mathcal{W}_{\mathcal{F}}^2(x_\perp)$ in the large- N_c limit,

⁵ The calculation of \hat{q} in $\mathcal{N} = 4$ SYM theory via (7.114) nicely resolves a subtlety. As we saw above, in addition to the extremized string configuration which touches the horizon, the string action also has another trivial solution which lies solely at the boundary, at $r = \infty$. Based on the connection between position in the r dimension in the gravitational theory and energy scale in the quantum field theory, the authors of Ref. [148, 149] argued that physical considerations (namely the fact that \hat{q} should reflect thermal physics at energy scales of order T) require selecting the extremized string configuration that touches the horizon. Although this physical argument remains valid, we now see that it is not necessary. In (7.114), the two Wilson lines are at the boundaries of I and I' , with different values of $\text{Im} t$. That means that there are no string world sheets that connect the two Wilson lines without touching the horizon. So, once we have understood how the nonstandard operator ordering in (7.114) modifies the boundary conditions for the string world sheet, we see that the trivial world sheet of Refs. [148, 149] and all of its generalizations in Refs. [545, 546] do not satisfy the correct boundary conditions. The nontrivial world sheet illustrated in Fig. 7.7, which is sensitive to thermal physics [148, 149, 466], is the only extremized world sheet bounded by the two lightlike Wilson lines in (7.114) [151].

because that is what arises in the analysis of jet quenching, see Section 2.3.2. (Radiative parton energy loss depends on the medium through the transverse momentum broadening of the radiated gluons, which are of course in the adjoint representation.) The x_\perp -independent term in the exponent in (7.125), namely “the -1”, is the finite subtraction of S_0 , which was identified in Ref. [148] as the action of two disjoint strings hanging straight down from the two Wilson lines to the horizon of the AdS black hole. Our calculation serves as a check of the value of S_0 , since only with the correct S_0 do we obtain $\mathcal{W}_A(0) = 1$ and a correctly normalized probability distribution $P(k_\perp)$. Note that our field theory set-up requires $L^-T \gg 1$, and our supergravity calculation requires $\lambda \gg 1$, meaning that our result (7.125) is only valid for

$$\sqrt{\lambda} L^-T \gg 1. \quad (7.126)$$

In this regime, (7.125) is very small unless $\pi x_\perp T / (2a)$ is small. This means that when we take the Fourier transform of (7.125) to obtain the probability distribution $P(k_\perp)$, in the regime (7.126) where the calculation is valid the Fourier transform is dominated by small values of x_\perp , for which

$$\mathcal{W}_A(x_\perp) \simeq \exp \left[-\frac{\pi^2}{4\sqrt{2}a} \sqrt{\lambda} L^-T^3 x_\perp^2 \right], \quad (7.127)$$

and we therefore obtain

$$P(k_\perp) = \frac{4\sqrt{2}a}{\pi\sqrt{\lambda}T^3L^-} \exp \left[-\frac{\sqrt{2}ak_\perp^2}{\pi^2\sqrt{\lambda}T^3L^-} \right]. \quad (7.128)$$

Thus, the probability distribution $P(k_\perp)$ is a Gaussian and the jet-quenching parameter (7.112) can easily be evaluated, yielding [148]

$$\hat{q} = \frac{\pi^{3/2}\Gamma(\frac{3}{4})}{\Gamma(\frac{5}{4})} \sqrt{\lambda}T^3. \quad (7.129)$$

The probability distribution (7.128) has a simple physical interpretation: the probability that the quark has gained transverse momentum k_\perp is given by diffusion in transverse momentum space with a diffusion constant given by $\hat{q}L$. This is indeed consistent with the physical expectation that transverse momentum broadening in a strongly coupled plasma is due to the accumulated effect of many small kicks by gluons from the medium: the quark performs Brownian motion in momentum space even though in coordinate space it remains on a light-like trajectory.

If we attempt to plug RHIC-motivated numbers into the result (7.129), taking $T = 300$ MeV, $N_c = 3$, $\alpha_{\text{SYM}} = \frac{1}{2}$ and therefore $\lambda = 6\pi$ yields $\hat{q} = 4.5$ GeV²/fm, which turns out to be in the same ballpark as the values of \hat{q} inferred from RHIC data on the suppression of high momentum partons in heavy ion collisions [148, 149], as we reviewed around Fig. 2.11 in Section 2.3. We can also write the result (7.129) as

$$\hat{q} \simeq 57 \sqrt{\alpha_{\text{SYM}} \frac{N_c}{3}} T^3, \quad (7.130)$$

which can be compared to the result (2.39,2.40) extracted via comparison to RHIC data in Ref. [167]. To make the comparison, we need to relate the QCD energy density ε appearing

in (2.39) to T . Lattice calculations of QCD thermodynamics indicate $\varepsilon \sim (9 - 11)T^4$ in the temperature regime that is relevant at RHIC [85]. This then means that if in (7.130) we take α_{SYM} within the range $\alpha_{\text{SYM}} = .66_{-.25}^{+.34}$, the result (7.130) for the strongly coupled $\mathcal{N} = 4$ SYM plasma is consistent with the result (2.40) obtained via comparing QCD jet quenching calculations to RHIC data. The extraction of \hat{q} by inference from LHC data on several-hundred-GeV jets should be under better control, since then the separation of scales upon which the QCD calculations (and the basic assumption that energy loss is dominated by gluon radiation) will be more quantitatively reliable.

We have reviewed the $\mathcal{N} = 4$ SYM calculation, but the jet quenching parameter can be calculated in any conformal theory with a gravity dual [149]. In a large class of such theories in which the spacetime for the gravity dual is $\text{AdS}_5 \times M_5$ for some internal manifold M_5 other than the five-sphere S_5 which gives $\mathcal{N} = 4$ SYM theory [149],

$$\frac{\hat{q}_{\text{CFT}}}{\hat{q}_{\mathcal{N}=4}} = \sqrt{\frac{s_{\text{CFT}}}{s_{\mathcal{N}=4}}}, \quad (7.131)$$

with s the entropy density. This result makes a central qualitative lesson from (7.129) clear: in a strongly coupled plasma, the jet quenching parameter is not proportional to the entropy density or to some number density of distinct scatterers. This qualitative lesson is more robust than any attempt to make a quantitative comparison to QCD. But, we note that if QCD were conformal, (7.131) would suggest

$$\frac{\hat{q}_{\text{QCD}}}{\hat{q}_{\mathcal{N}=4}} \approx 0.63. \quad (7.132)$$

And, analysis of how \hat{q} changes in a particular toy model in which nonconformality can be introduced by hand then suggests that introducing the degree of nonconformality seen in QCD thermodynamics may increase \hat{q} by a few tens of percent [466].⁶ Putting these two observations together, perhaps it is not surprising that the \hat{q} for the strongly coupled plasma of $\mathcal{N} = 4$ SYM theory is in the same ballpark as that extracted by comparison with RHIC data.

7.6 Quenching a beam of strongly coupled synchrotron radiation

In Sections 7.1, 7.2 and 7.3 we have gained insights into parton energy loss and jet quenching via studying how a single heavy quark slows and gets kicked as it moves through the strongly coupled plasma of $\mathcal{N} = 4$ SYM theory, and how that single heavy quark excites the plasma through which it moves. In Section 7.4, we have seen how a single light quark or gluon is stopped and thermalized by the strongly coupled $\mathcal{N} = 4$ SYM plasma, assuming that all aspects of this physics are described at strong coupling. In Section 7.5 we have analyzed jet quenching per se, but have done so via the strategy of working as far as possible within weakly coupled QCD and only using a holographic calculation within $\mathcal{N} = 4$ SYM theory for one small part of the story, namely the calculation of the jet quenching parameter \hat{q} through

⁶ \hat{q} also increases with increasing nonconformality in strongly coupled $\mathcal{N} = 2^*$ gauge theory [461, 462].

which the physics of the strongly coupled medium enters the calculation. This approach is justified in the high jet energy limit, where there is a clean separation of scales and where jet quenching is due to medium-induced gluon radiation. But, the jets being studied at present at RHIC are not sufficiently energetic as to make one confident in the quantitative reliability of these approximations. In addition to motivating the study of jet quenching in heavy ion collisions at the LHC, this raises an obvious theoretical question: can we study jets themselves in $\mathcal{N} = 4$ SYM theory? Suppose we assume that the physics at all relevant scales is strongly coupled, an assumption that is in a sense the opposite of that we made in Section 7.5. But, instead of considering just a single quark, can we analyze how an actual jet is modified by the strongly coupled plasma of $\mathcal{N} = 4$ SYM theory, relative to how it would have developed in the vacuum of that theory? This question is pressing since one of the main goals of the ongoing high- p_T measurements in heavy ion collisions at RHIC is full jet reconstruction, with experimentalists aiming to answer questions like how much wider in angle are jets in heavy ion collisions than in proton-proton collisions? For example, the STAR collaboration is measuring the ratio of the energy within a cone of radius 0.2 radians to that within a cone of radius 0.4 radians, and comparing that ratio in gold-gold and proton-proton collisions. It would be of value to get some insights into what to expect for such an observable in a strongly coupled plasma, via a holographic calculation in $\mathcal{N} = 4$ SYM theory.

Unfortunately, the answer to the question as we have just posed it is negative. In Ref. [551], Hofman and Maldacena considered the following thought experiment. Suppose you did electron-positron scattering in a world in which the electron and positron coupled to $\mathcal{N} = 4$ SYM theory through a virtual photon, just as in the real world they couple to QCD. What would happen in high energy scattering? Would there be any “jetty” events? They showed that the answer is no. Instead, the final state produced by a virtual photon in the conformal $\mathcal{N} = 4$ SYM theory is a spherically symmetric outflow of energy. Similar conclusions were also reached in Refs. [530, 552]. The bottom line is that there are no jets in strongly coupled $\mathcal{N} = 4$ SYM theory, which would seem to rule out using this theory to study how jets are modified by propagating through the strongly coupled plasma of this theory.

Recent work [553] may provide a path forward. These authors have found a way of creating a beam of gluons (and color-adjoint scalars) that is tightly collimated in angle and that propagates outwards forever in the $\mathcal{N} = 4$ SYM theory vacuum without spreading in angle. This beam is not literally a jet, since it is not produced far off shell. But, we know from Hofman and Maldacena that a far off shell “photon” does not result in jets in this theory. And, this beam of nonabelian radiation may yield a better cartoon of a jet than a single quark, heavy or light, or a single gluon, all of which have been analyzed previously. This beam is created by synchrotron radiation from a test quark that is moving in a circle of radius L with constant angular frequency ω , as in Ref. [469] but now, in Ref. [553], in vacuum instead of in medium. The way the calculation is done is sketched in Fig. 7.8. The logic of the calculation is as we described in Section 7.3, and so we shall not present it in detail. It turns out, however, that in vacuum the shape of the rotating string in the bulk and the corresponding form of the energy density of the outward-propagating radiation on the boundary can both be determined analytically [553].

Because the vacuum of $\mathcal{N} = 4$ SYM theory is not similar to that of our world, studying the radiation of a test quark in circular motion does not have direct phenomenological mo-




Figure 7.8: Cartoon of the gravitational description of synchrotron radiation at strong coupling: the quark rotating at the boundary trails a rotating string behind it which hangs down into the bulk AdS_5 space. This string acts as a source of gravitational waves in the bulk, and this gravitational radiation induces a stress tensor on the four-dimensional boundary. By computing the bulk-to-boundary propagator one obtains the boundary-theory stress-tensor that describes the radiated energy. The entire calculation can be done analytically [553].

tivation. However, the angular distribution of the radiation turns out to be *very* similar to that of classic synchrotron radiation [553]: when the quark is moving along its circle with an ultrarelativistic velocity v , the radiation is produced in a tightly collimated beam with an angular width that is proportional to $1/\gamma$, with γ the Lorentz boost factor associated with v , see Fig. 7.9. The spiral pulse of radiation propagates outward without spreading, and the radial thickness (or, equivalently, the duration in time) of the pulse is proportional to $1/\gamma^3$. So, as a function of increasing γ the beam, if Fourier analyzed, contains radiation at increasingly short wavelengths and high frequencies.

Although it has recently been explained from the gravitational point of view in beautifully geometric terms [554], from the point of view of the non-abelian gauge theory it is surprising that the angular distribution of the radiation at strong coupling is so similar (see Ref. [553] for quantitative comparisons) to that at weak coupling, where what is radiated is a mixture of colored — and therefore interacting — gluons and scalars. The fact that, even when the coupling is arbitrarily strong, as the pulses of radiation propagate outwards they do not spread at all and never isotropize indicates that intuition based upon parton branching [552] (namely that the non-abelian character of the radiation should result in energy flowing from short to long wavelengths as the pulses propagate outwards and should yield isotropization




Figure 7.9: Cutaway plots of $r^2\mathcal{E}/P$ for a test quark in circular motion with $v = 1/2$ and $v = 3/4$. Here, \mathcal{E} is the energy density and P is the total power radiated per unit time. We see a spiral of radiation, propagating radially outwards at the speed of light, without any spreading. The spiral is localized about $\theta = \pi/2$ with a characteristic width $\delta\theta \propto 1/\gamma$. The radial thickness of the spiral is proportional to $1/\gamma^3$.

at large distances) is invalid in this context.⁷

For our purposes, the result that the “beam” of radiated gluons (and scalars) propagates outward with a fixed angular width that we can select by picking γ is fortuitous. It means that via this roundabout method we have found a way of making a state that looks something like a jet. It is not the same in all respects as a jet in QCD in that it is not produced via the fragmentation of an initially far offshell parton. But, it is a collimated beam of gluons of known, and controllable, angular width. For this reason, the results obtained in the formal setting of a test quark moving in a circle open the way to new means of modeling jet quenching in heavy ion collisions. The challenge is to repeat the calculation of Ref. [553] at nonzero temperature, in the presence of a horizon in the bulk spacetime. Upon doing so, one could watch the initially tightly collimated beam of synchrotron radiation interact with the strongly coupled plasma that would then be present. As long as $\omega^2\gamma^3 \gg \pi^2T^2$, we know from Ref. [469] that the total power radiated is as if in vacuum [468]. In this regime, we therefore expect to see strongly coupled synchrotron radiation as in Fig. 7.9 on length scales small compared to $1/T$. But, as in Section 7.3, we expect that this beam of radiation should eventually thermalize locally, likely converting into an outgoing hydrodynamic wave moving

⁷Very recent work [555] shows that isotropization via some analogue of parton branching is also not the correct picture for the radiation studied in Ref. [551] by Hofman and Maldacena: this radiation also propagates outward as a pulse without any spreading, but this pulse is spherically symmetric at all radii. So, in the case studied by Hofman and Maldacena there is no process of isotropization because the radiation is isotropic at all times while in the case illustrated in Fig. 7.9 there is no isotropization because the radiation never becomes isotropic.

at the speed of sound before ultimately isotropizing, dissipating and thermalizing completely. One could see by how much a beam prepared such that in vacuum it has a known, small, angular extent spreads in angle after propagating through $1/T$ (or $2/T$, or $3/T$) of strongly coupled plasma. Such an analysis should provide interesting benchmarks against which to compare data from RHIC on how much wider in angle jets of a certain energy are in heavy ion collisions as compared to in proton-proton collisions. Early results on jets in heavy ion collisions at the LHC [124] indicate that, for the high energy jets available in these collisions, it will be possible to make calorimetric measurements of how the jets initiated by partons propagating in the medium produced in these collisions are broadened in angle.

7.7 Velocity-scaling of the screening length and quarkonium suppression in a hot wind

We saw in Section 2.4 that, because they are smaller than typical hadrons in QCD, heavy quarkonium mesons survive as bound states even at temperatures above the crossover from a hadron gas to quark-gluon plasma. However, if the temperature of the quark-gluon plasma is high enough, they eventually dissociate. An important physical mechanism underlying the dissociation is the weakening attraction between the heavy quark and anti-quark in the bound state because the force between their color charges is screened by the medium. The dissociation of charmonium and bottomonium bound states has been proposed as a signal for the formation of a hot and deconfined quark-gluon plasma in heavy ion collisions [184], and as a means of gauging the temperatures reached during the collisions.

In the limit of large quark mass, the interaction between the quark and the anti-quark in a bound state in the thermal medium can be extracted from the thermal expectation value of the Wilson loop operator $\langle W^F(\mathcal{C}_{\text{static}}) \rangle$, with $\mathcal{C}_{\text{static}}$ a rectangular loop with a short side of length L in a spatial direction (say x_1) and a long side of length \mathcal{T} along the time direction. This expectation value takes the form

$$\langle W^F(\mathcal{C}_{\text{static}}) \rangle = \exp[-i\mathcal{T}E(L)] . \quad (7.133)$$

where $E(L)$ is the (renormalized) free energy of the quark-antiquark pair with the self-energy of each quark subtracted. $E(L)$ defines an effective potential between the quark anti-quark pair. The screening of the force between color charges due to the presence of the medium manifests itself in the flattening of $E(L)$ for L greater than some characteristic length scale L_s called the screening length. In QCD, the flattening of the potential occurs smoothly, as seen in the lattice calculations illustrated in Fig. 3.5 in Section 3.3, and one must make an operational definition of L_s . For example, in the parametrization of (2.41), L_s can be set equal to $1/\mu$. L_s decreases with increasing temperature and can be used to estimate the scale of the dissociation temperature T_{diss} as

$$L_s(T_{\text{diss}}) \sim d \quad (7.134)$$

where d is the size of a particular mesonic bound state at zero temperature. The idea here is that once the temperature of the quark-gluon plasma is high enough that the potential between a quark and an antiquark separated by a distance corresponding to the size of a

particular meson has been fully screened, that meson can no longer exist as a bound state in the plasma. This means that larger quarkonium states dissociate at lower temperatures, and means that the ground-state bottomonium meson survives to the highest temperatures of all. As we discussed at length in Section 2.4, there are many important confounding effects that must be taken into account in order to realize the goal of using data on charmonium production in heavy ion collisions to provide evidence for this sequential pattern of quarkonium dissociation as a function of increasing temperature. In this Section, we shall focus only on one of these physical effects, one on which calculations done via gauge/gravity duality have shed some light [149, 459, 466, 556–571].

In heavy ion collisions, quarkonium mesons are produced moving with some velocity \vec{v} with respect to the medium. It is thus important to understand the effects of nonzero quarkonium velocity on the screening length and consequent dissociation of bound states. To describe the interaction between a quark–anti-quark pair that is moving relative to the medium, it is convenient to boost into a frame in which the quark anti-quark pair is at rest, but feels a hot wind of QGP blowing past them. The effective quark potential can again be extracted from (7.133) evaluated in the boosted frame with \mathcal{T} now interpreted as the proper time of the dipole. While much progress has been made in using lattice QCD calculations to extract the effective potential between a quark–anti-quark pair at rest in the QGP, there are significant difficulties in using Euclidean lattice techniques to address the (dynamical as opposed to thermodynamic) problem of a quark–antiquark pair in a hot wind. In the strongly coupled plasma of $\mathcal{N} = 4$ SYM theory with large N_c , however, the calculation can be done using gauge/gravity duality [149, 557, 558], and requires only a modest extension of the standard methods reviewed in Section 5.4. Here, we sketch the derivation from Ref. [149].

We start with a rectangular Wilson loop whose short transverse space-like side

$$\sigma = x_1 \in \left[-\frac{L}{2}, \frac{L}{2}\right] \tag{7.135}$$

defines the separation L between the quark–antiquark pair and whose long time-like sides extend along the $x_3 = vt$ direction, describing a pair moving with speed v in the x_3 direction. In this frame, the plasma is at rest and the spacetime metric in the gravitational description is the familiar AdS black hole (5.30). We then apply a Lorentz boost that rotates this Wilson loop into the rest frame (t', x'_3) of the quark–anti-quark pair:

$$dt = dt' \cosh \eta - dx'_3 \sinh \eta, \tag{7.136}$$

$$dx_3 = -dt' \sinh \eta + dx'_3 \cosh \eta, \tag{7.137}$$

where the rapidity η is given by $\tanh \eta = v$, meaning that $\cosh \eta = \gamma$. After the AdS black hole metric has been transformed according to this boost, it describes the moving hot wind of plasma felt by the quark–anti-quark pair in its rest frame.

In order to extract $E(L)$ it suffices to work in the limit in which the time-like extent of the Wilson loop \mathcal{T} is much greater than its transverse extent L , meaning that the corresponding string worldsheet “suspended” from this Wilson loop and “hanging down” into the bulk is invariant under translations along the long direction of the Wilson loop. Parametrizing the two-dimensional world sheet with the coordinates σ and $\tau = t$, the dependence on τ is then trivial. The task is reduced to calculating the curve $r(\sigma)$ along which the worldsheet descends

into the bulk from positions on the boundary brane which we take to be located at $r = r_0\Lambda$, with Λ a dimensionless UV cutoff that we shall take to infinity at the end of the calculation. That is, the boundary conditions on $r(\sigma)$ are

$$r\left(\pm\frac{L}{2}\right) = r_0\Lambda. \quad (7.138)$$

It is then helpful to introduce dimensionless variables

$$r = r_0y, \quad \tilde{\sigma} = \sigma\frac{r_0}{R^2}, \quad l = \frac{Lr_0}{R^2} = \pi LT, \quad (7.139)$$

where $T = \frac{r_0}{\pi R^2}$ is the temperature. Upon dropping the tilde, one is then seeking to determine the shape $y(\sigma)$ of the string world sheet satisfying the boundary conditions $y(\pm\frac{l}{2}) = \Lambda$. From the boosted AdS black hole metric, one finds that the Nambu-Goto action which must be extremized takes the form

$$S(\mathcal{C}) = -\sqrt{\lambda}\mathcal{T}T \int_0^{l/2} d\sigma \mathcal{L}, \quad (7.140)$$

with a Lagrangian that reads ($y' = \partial_\sigma y$)

$$\mathcal{L} = \sqrt{(y^4 - \cosh^2 \eta) \left(1 + \frac{y'^2}{y^4 - 1}\right)}. \quad (7.141)$$

We must now determine $y(\sigma)$ by extremizing (7.141). This can be thought of as a classical mechanics problem, with σ the analogue of time. Since \mathcal{L} does not depend on σ explicitly, the corresponding Hamiltonian

$$\mathcal{H} \equiv \mathcal{L} - y' \frac{\partial \mathcal{L}}{\partial y'} = \frac{y^4 - \cosh^2 \eta}{\mathcal{L}} = q \quad (7.142)$$

is a constant of the motion, which we denote by q . In the calculation we are reviewing in this section, we take $\Lambda \rightarrow \infty$ at fixed, finite, rapidity η . In this limit, the string worldsheet in the bulk is time-like, and $E(L)$ turns out to be real. (The string world-sheet bounded by the rectangular Wilson loop that we are considering becomes space-like if $\sqrt{\cosh \eta} > \Lambda$. In order to recover the light-like Wilson loop used in the calculation of the jet quenching parameter in Section 7.5, one must first take $\eta \rightarrow \infty$ and only then take $\Lambda \rightarrow \infty$.)

It follows from the Hamiltonian (7.142) that solutions $y(\sigma)$ with $\Lambda > \sqrt{\cosh \eta}$ satisfy the equation of motion

$$y' = \frac{1}{q} \sqrt{(y^4 - 1)(y^4 - y_c^4)} \quad (7.143)$$

with

$$y_c^4 \equiv \cosh^2 \eta + q^2. \quad (7.144)$$

Note that $y_c^4 > \cosh^2 \eta \geq 1$. The extremal string world sheet begins at $\sigma = -\ell/2$ where $y = \Lambda$, and “descends” in y until it reaches a turning point, namely the largest value of y at which $y' = 0$. It then “ascends” from the turning point to its end point at $\sigma = +\ell/2$ where $y = \Lambda$. By symmetry, the turning point must occur at $\sigma = 0$. We see from (7.143) that

in this case, the turning point occurs at $y = y_c$ meaning that the extremal surface stretches between y_c and Λ . The integration constant q can then be determined from the equation $\frac{l}{2} = \int_0^{\frac{l}{2}} d\sigma$ which, upon using (7.143), becomes

$$l = 2q \int_{y_c}^{\Lambda} dy \frac{1}{\sqrt{(y^4 - y_c^4)(y^4 - 1)}} . \quad (7.145)$$

The action for the extremal surface can be found by substituting (7.143) into (7.140) and (7.141), yielding

$$S(l) = -\sqrt{\lambda} \mathcal{T} T \int_{y_c}^{\Lambda} dy \frac{y^4 - \cosh^2 \eta}{\sqrt{(y^4 - 1)(y^4 - y_c^4)}} . \quad (7.146)$$

Equation (7.146) contains not only the potential between the quark-antiquark pair but also the static mass of the quark and antiquark considered separately in the moving medium. (Recall that we have boosted to the rest frame of the quark and antiquark, meaning that the quark-gluon plasma is moving.) Since we are only interested in the quark-antiquark potential, we need to subtract the action S_0 of two independent quarks from (7.146) in order to obtain the quark-anti-quark potential in the dipole rest frame:

$$E(L)\mathcal{T} = -S(l) + S_0 . \quad (7.147)$$

The string configuration corresponding to a single quark at rest in a moving $\mathcal{N} = 4$ SYM plasma was obtained in Refs. [452, 453], as we have reviewed in Section 7.1. From this configuration one finds that

$$S_0 = -\sqrt{\lambda} \mathcal{T} T \int_1^{\Lambda} dy . \quad (7.148)$$

To extract the quark-antiquark potential, we use (7.145) to solve for q in terms of l and then plug the corresponding $q(l)$ into (7.146) and (7.147) to obtain $E(L)$. Note that (7.145) is manifestly finite as $\Lambda \rightarrow \infty$ and the limit can be taken directly. (7.146) and (7.148) are divergent separately when taking $\Lambda \rightarrow \infty$, but the difference (7.147) is finite.

We now describe general features of (7.145) and (7.147). Denoting the RHS of (7.145) (with $\Lambda = \infty$) as function $l(q)$, one finds that for a given η , $l(q)$ has a maximum $l_{\max}(\eta)$ and equation (7.145) has no solution when $l > l_{\max}(\eta)$. Thus for $l > l_{\max}$, the only string worldsheet configuration is two disjoint strings and from (7.147), $E(L) = 0$, i.e. the quark and anti-quark are completely screened to the order of the approximation we are considering. We can define the screening length as

$$L_s \equiv \frac{l_{\max}(\eta)}{\pi T} . \quad (7.149)$$

At $\eta = 0$, i.e. the dipole at rest with the medium, one finds that

$$L_s(0) \approx \frac{0.87}{\pi T} \approx \frac{0.28}{T} . \quad (7.150)$$

Similar criteria are used in the definition of screening length in QCD [572], although in QCD there is no sharply defined length scale at which screening sets in. Lattice calculations of the




Figure 7.10: The screening length l_{\max} times its leading large- η dependence $\sqrt{\cosh(\eta)}$. The exact results are given for dipoles oriented perpendicular to the wind ($\theta = \pi/2$) and parallel to the wind ($\theta = 0$). The $\theta = \pi/2$ curve is compared to the analytical large- η approximation (7.151). Keeping only the first term in this analytical expression corresponds to a horizontal line on the figure; including the term proportional to $(\cosh \eta)^{-5/2}$ improves the agreement with the exact result.

static potential between a heavy quark and antiquark in QCD indicate a screening length $L_s \sim 0.5/T$ in hot QCD with two flavors of light quarks [573] and $L_s \sim 0.7/T$ in hot QCD with no dynamical quarks [230]. The fact that there *is* a sharply defined L_s is an artifact of the limit in which we are working, in which $E(L) = 0$ for $L > L_s$.⁸

$L_s(\eta)$ can be obtained numerically, with results as illustrated in Fig. 7.10. One finds that $L_s(\eta)$ decreases with increasing velocity, indicating that quarkonia dissociate at a lower temperature when they are moving. An analytical expression for L_s can also be obtained in the limit of high velocity. Expanding (7.145) in powers of $1/y_c^4$ and truncating the corresponding expression at order $1/y_c^4$, one finds

$$\begin{aligned}
 l_{\max} &= \frac{\sqrt{2\pi} \Gamma(\frac{3}{4})}{3^{3/4} \Gamma(\frac{1}{4})} \left(\frac{2}{\cosh^{1/2} \eta} + \frac{1}{5 \cosh^{5/2} \eta} + \dots \right) \\
 &= 0.74333 \left(\frac{1}{\cosh^{1/2} \eta} + \frac{1}{10 \cosh^{5/2} \eta} + \dots \right). \tag{7.151}
 \end{aligned}$$

In Fig. 7.10, exact results for $l_{\max} = \pi T L_s(\eta)$ as a function of η , obtained numerically, are compared to the expression (7.151) that is valid in the high velocity limit. One sees that the

⁸We are considering the contribution to $E(L)$ that is proportional to $\sqrt{\lambda}$. For $l \gg l_{\max}$, the leading contribution to $E(L)$ is proportional to λ^0 and is determined by the exchange of the lightest supergravity mode between the two disjoint strings [380].

screening length decreases with increasing velocity to a good approximation according to the scaling [556–558]

$$L_s(v) \simeq \frac{L_s(0)}{\cosh^{1/2} \eta} = \frac{L_s(0)}{\sqrt{\gamma}}, \quad (7.152)$$

with $\gamma = 1/\sqrt{1-v^2}$. This velocity dependence suggests that L_s should be thought of as, to a good approximation, proportional to (energy density) $^{-1/4}$, since the energy density increases like γ^2 as the wind velocity is boosted. The velocity-scaling of L_s has proved robust in the sense that it applies in various strongly coupled plasmas other than $\mathcal{N} = 4$ SYM [466, 559, 561, 566] and in the sense that it applies to baryons made of heavy quarks also [569].

We have only described the calculation for the case in which the direction of the hot wind is perpendicular to the dipole. With a little more effort, the calculation can be extended to general angles and has been analyzed in detail in [149], to which we refer the readers for more details. The conclusion of [149] was that the dependences of L_s and $E(L)$ on the angle between the dipole and the wind are very weak. For example, the black line in Fig. 7.10 gives the η dependence of the screening length when the wind direction is parallel to the dipole. We see the difference from the perpendicular case is only about 12%.

If the velocity-scaling of L_s (7.152) holds for QCD, it will have qualitative consequences for quarkonium suppression in heavy ion collisions [149, 557]. From (7.134), the dissociation temperature $T_{\text{diss}}(v)$, defined as the temperature above which J/Ψ or Υ mesons with a given velocity do not exist, should scale with velocity as

$$T_{\text{diss}}(v) \sim T_{\text{diss}}(v=0)(1-v^2)^{1/4}, \quad (7.153)$$

since $T_{\text{diss}}(v)$ should be the temperature at which the screening length $L_s(v)$ is comparable to the size of the meson bound state. The scaling (7.153) indicates that slower mesons can exist up to higher temperatures than faster ones. As illustrated schematically in Fig. 7.11, this scaling indicates that J/Ψ suppression at RHIC (and Υ suppression at the LHC) may increase markedly for J/Ψ 's (Υ 's) with transverse momentum p_T above some threshold, on the assumption that the temperatures reached in RHIC (LHC) collisions do not reach the dissociation temperature of J/Ψ (Υ) mesons at zero velocity [185, 190]. Modelling this effect requires embedding results for quarkonium production in hard scatterings in nuclear collisions into a hydrodynamic code that describes the motion of the quark-gluon fluid produced in the collision, in order to evaluate the velocity of the hot wind felt by each putative quarkonium meson. Such an analysis indicates that once p_T is above the threshold at which $T_{\text{diss}}(v)$ has dropped below the temperature reached in the collision, the decline in the J/Ψ survival probability is significant, by more than a factor of four (two) in central (peripheral) collisions [574, 575]. We should caution that, as we discussed in Section 2.4, in modelling quarkonium production and suppression versus p_T in heavy ion collisions, various other effects like secondary production or formation of J/Ψ mesons outside the hot medium at high p_T [576] remain to be quantified. The quantitative importance of these and other effects may vary significantly, depending on details of their model implementation. In contrast, Eq. (7.153) was obtained directly from a field-theoretic calculation and its implementation will not introduce additional model-dependent uncertainties.




Figure 7.11: A $1/\sqrt{\gamma}$ -velocity scaling of the screening length in QCD would imply a J/Ψ dissociation temperature $T_{\text{diss}}(p_T)$ that decreases significantly with p_T , while that for the heavier Υ is affected less at a given p_T . The curves are schematic, in that we have arbitrarily taken $T_{\text{diss}}(0)$ for the J/Ψ to be $2.1 T_c$ and we have increased $T_{\text{diss}}(0)$ for the Υ over that for the J/ψ by a factor corresponding to its smaller size in vacuum. At a qualitative level, we expect to see fewer J/Ψ (Υ) mesons at p_T 's above that at which their dissociation temperature is comparable to the temperatures reached in heavy ion collisions at RHIC (at the LHC).

The threshold p_T above which the production of J/Ψ mesons falls off due to their motion through the quark-gluon plasma depends sensitively on the difference between $T(\text{diss}(v=0))$ and the temperature reached in the collision [575], but should be somewhere around 5 GeV; perhaps a conservative range is 2 to 10 GeV. Present data on the suppression of J/Ψ production in CuCu collisions at RHIC show no signs of increased suppression above some p_T threshold [577, 578], but the error bars are large above $p_T \sim 4$ GeV. The kinematical range in which this novel quarkonium suppression mechanism is operational lies within experimental reach of near-future high-luminosity AuAu runs at RHIC and will be studied thoroughly at the LHC in both the J/Ψ and Υ channels.

The analysis of this section is built upon the calculation of the potential between a test quark and antiquark in the strongly coupled plasma of $\mathcal{N} = 4$ SYM theory, a theory which in and of itself has no mesons. Gaining insight into the physics of quarkonium mesons from calculations of the screening of the static quark-antiquark potential has a long history in QCD, as we have seen in Section 3.3. But, we have also seen in that section that these approaches are gradually being superseded as lattice QCD calculations of quarkonium spectral functions themselves are becoming available. In the present context also, we would like to

move beyond drawing inferences about mesons from analyses of the potential $E(L)$ and the screening length L_s to analyses of mesons themselves. This is the subject of Section 8, in which we shall carefully describe how once we have added heavy quarks to $\mathcal{N} = 4$ SYM by adding a D7-brane in the gravity dual [385], as in Section 5.5, the fluctuations of the D7-brane then describe the quarkonium mesons of this theory. We shall review the construction first in vacuum and then in the presence of the strongly coupled plasma at nonzero temperature. We shall find that the results of this section prove robust, in that the velocity-scaling (7.153) has also been obtained [568] by direct analysis of the dispersion relations of mesons in the plasma [565, 568]. These mesons have a limiting velocity that is less than the speed of light and that decreases with increasing temperature [565], and whose temperature dependence is equivalent to (7.153) up to few percent corrections that have been computed [568] and that we shall show. This is a key part of the story, with the velocity-dependent dissociation temperature of this section becoming a temperature dependent limiting velocity for explicitly constructed quarkonium mesons in Section 8. However, this cannot be the whole story since the dispersion relations seem to allow for mesons with arbitrarily large momentum even though they limit their velocity. The final piece of the story is described in Section 8.4.2, where we review the calculation of the leading contribution to the widths of these mesons [571], which was neglected in the earlier calculations of their dispersion relations. Above some momentum, the width grows rapidly, increasing like p_1^2 . And, the momentum above which this rapid growth of the meson width sets in is just the momentum at which the meson velocity first approaches its limiting value. The physical picture that emerges is that at the momentum at which the mesons reach a velocity such that the hot wind they are feeling has a temperature sufficient to dissociate them, according to the analysis of this Section built upon the calculation of L_s , their widths in fact grow rapidly [571].

Chapter 8

Quarkonium mesons

As discussed in section 2.4, heavy quarks and quarkonium mesons, with masses such that $M/T \gg 1$, constitute valuable probes of the QGP. Since dynamical questions about these probes are very hard to answer from first principles, here we will study analogous questions in the strongly coupled $\mathcal{N} = 4$ SYM plasma. In this case the gauge/string duality provides the tool that makes a theoretical treatment possible. Although for concreteness we will focus on the $\mathcal{N} = 4$ plasma, many of the results that we will obtain are rather universal in the sense that, at least qualitatively, they hold for any strongly coupled gauge theory with a string dual. Such results may give us insights relevant for the QCD quark-gluon plasma at temperatures at which it is reasonably strongly coupled.

The QGP only exists at temperatures $T > T_c$, so in QCD the condition $M/T \gg 1$ can only be realised by taking M to be large. In contrast, $\mathcal{N} = 4$ SYM is a conformal theory with no confining phase, so all temperatures are equivalent. In the presence of an additional scale, namely the quark or the meson mass, the physics only depends on the ratio M/T . This means that in the $\mathcal{N} = 4$ theory the condition $M/T \gg 1$ can be realised by fixing T and sending M to infinity, or by fixing M and sending $T \rightarrow 0$; both limits are completely equivalent. In particular, the leading-order approximation to the heavy quark or quarkonium meson physics, in an expansion in T/M , may be obtained by setting $T = 0$. For this reason, this is the limit that we will study first.

We will follow the nomenclature common in the QCD literature and refer to mesons made of two heavy quarks as ‘quarkonium mesons’ or ‘quarkonia’, as opposed to using the term ‘heavy mesons’, which commonly encompasses mesons made of one heavy and one light quark.

8.1 Adding quarks to $\mathcal{N} = 4$ SYM

In section 5.5 we saw that N_f flavours of fundamental matter can be added to $\mathcal{N} = 4$ SYM by introducing N_f D7-brane probes into the geometry sourced by the D3-branes, as indicated by the array (5.93), which we reproduce here (with the time direction included) for convenience:

$$\begin{array}{cccccccccccc} \text{D3:} & 0 & 1 & 2 & 3 & - & - & - & - & - & - & - \\ \text{D7:} & 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & - & - & . \end{array} \tag{8.1}$$

Before we proceed, let us clarify an important point. $\mathcal{N} = 4$ SYM is a conformal theory,

i.e. its β -function vanishes exactly. Adding matter to it, even if the matter is massless, makes the quantum-mechanical β -function positive, at least perturbatively. This means that the theory develops a Landau pole in the UV and is therefore not well defined at arbitrarily-high energy scales.¹ However, since the beta-function (for the 't Hooft coupling λ) is proportional to N_f/N_c , the Landau pole occurs at a scale of order e^{N_c/N_f} . This is exponentially large in the limit of interest here, $N_f/N_c \ll 1$, and in fact the Landau pole disappears altogether in the strict probe limit $N_f/N_c \rightarrow 0$. On the string side, the potential pathology associated with a Landau pole manifests itself in the fact that a completely smooth solution that incorporates the backreaction of the D7-branes may not exist [383, 579–583]. In any case, the possible existence of a Landau pole at high energies will not be of concern for the applications reviewed here. In the gauge theory, it will not prevent us from extracting interesting infrared physics, just as the existence of a Landau pole in QED does not prevent one from calculating the conductivity of an electromagnetic plasma. In the string description, we will not go beyond the probe approximation, so the backreaction of the D7-branes will not be an issue.² And finally, we note that we will work with the D3/D7 model because of its simplicity. We could work with a more sophisticated model with better UV properties, but this would make the calculations more involved while leaving the physics we are interested in essentially unchanged.

As illustrated in Fig. 8.1, the D3-branes and the D7-branes can be separated a distance L in the 89-directions. This distance times the string tension, Eqn. (4.11), is the minimum energy of a string stretching between the D3- and the D7-branes. Since the quarks arise as the lightest modes of these 3-7 strings, this energy is precisely the bare quark mass:

$$M_q = \frac{L}{2\pi\alpha'}. \quad (8.2)$$

An important remark here is the fact that the branes in Fig. 8.1 are implicitly assumed to be embedded in flat spacetime. In Section 5.5 this was referred to as the ‘first’ or ‘open-string’ description of the D3/D7 system, which is reliable in the regime $g_s N_c \ll 1$, in which the backreaction of the D3-branes on spacetime can be ignored. One of our main tasks in the following sections will be to understand how this picture is modified in the opposite regime, $g_s N_c \gg 1$, when the D3-branes are replaced by their backreaction on spacetime. In this regime the shape of the D7-branes may or may not be modified, but Eqn. (8.2) will remain true provided the appropriate definition of L , to be given below, is used.

Although $\mathcal{N} = 4$ SYM is a conformal theory, the addition of quarks with a nonzero mass introduces a scale and gives rise to a rich spectrum of quark-antiquark bound states, i.e. mesons. In the following section we will study the meson spectrum in this theory at zero temperature in the regime of strong ‘t Hooft coupling, $g_s N_c \gg 1$. On the gauge theory side this is inaccessible to conventional methods such as perturbation theory, but on the string side a classical description in terms of D7-brane probes in a weakly curved $AdS_5 \times S^5$ applies. Our first task is thus to understand in more detail the way in which the D7-branes are embedded in this geometry. Since this is crucial for subsequent sections, we will in fact provide a fair amount of detail here.

¹Non-perturbatively, the possibility that a strongly coupled fixed point exists must be ruled out before reaching this conclusion. See [579] for an argument in this direction based on supersymmetry.

²For a review of ‘unquenched’ models, i.e. those in which the flavour backreaction is included, see [584].




Figure 8.1: D3/D7 system at weak coupling.

8.2 Zero temperature

8.2.1 D7-brane embeddings

We begin by recalling that the coordinates in the $AdS_5 \times S^5$ metric (5.1), (5.2) can be understood as follows. The four directions t, x_i correspond to the 0123-directions in (8.1). The 456789-directions in the space transverse to the D3-branes give rise to the radial coordinate r in AdS_5 , defined through

$$r^2 = x_4^2 + \dots + x_9^2, \quad (8.3)$$

as well as five angles that parametrise the S^5 . We emphasize that, once the gravitational effect of the D3-branes is taken into account, the 6-dimensional space transverse to the D3-branes is not flat, so the x^4, \dots, x^9 coordinates are not Cartesian coordinates. However, they are still useful to label the different directions in this space.

The D7-branes share the 0123-directions with the D3-branes, so from now on we will mainly focus on the remaining directions. In the 6-dimensional space transverse to the D3-branes, the D7-branes span only a 4-dimensional subspace parametrised by x_4, \dots, x_7 . Since the D7-branes preserve the $SO(4)$ rotational symmetry in this space, it is convenient to introduce a radial coordinate u such that

$$u^2 = x_4^2 + \dots + x_7^2, \quad (8.4)$$

as well as three spherical coordinates, denoted collectively by Ω_3 , that parametrise an S^3 .




Figure 8.2: Coordinates in the 6-dimensional space transverse to the D3-branes. Each axis actually represents two directions, *i.e.* a plane (or, equivalently, the radial direction in that plane). The asymptotic distance $L = U(u = \infty)$ is proportional to the quark mass, Eqn. (8.2). We emphasize that the directions parallel to the D3-branes (the gauge theory directions t, x_i) are suppressed in this picture, and they should not be confused with the D7 directions shown in the figure, which lie entirely in the space transverse to the D3-branes.

Similarly, it is useful to introduce a radial coordinate U in the 89-plane through

$$U = x_8^2 + x_9^2, \quad (8.5)$$

as well as a polar angle α . In terms of these coordinates one has

$$dx_4^2 + \dots + dx_9^2 = du^2 + u^2 d\Omega_3^2 + dU^2 + U^2 d\alpha^2. \quad (8.6)$$

Obviously, the overall radial coordinate r satisfies $r^2 = u^2 + U^2$.

Since the D7-branes only span the 4567-directions, they only wrap an S^3 inside the S^5 . The D7-brane worldvolume may thus be parametrized by the coordinates $\{t, x_i, u, \Omega_3\}$. In order to specify the D7-branes' embedding one must then specify the remaining spacetime coordinates, U and α , as functions of, in principle, all the worldvolume coordinates. However, translational symmetry in the $\{t, x_i\}$ -directions and rotational symmetry in the $\{\Omega_3\}$ -directions allow U and α to depend only on u .

In order to understand this dependence, consider first the case in which the spacetime curvature generated by the D3-branes is ignored. In this case, the D7-branes lie at a constant position in the 89-plane, see Fig. 8.2. In other words, their embedding is given by $\alpha(u) = \alpha_0$ and $U(u) = L$, where α_0 and L are constants. The first equation can be understood as saying that, because of the $U(1)$ rotational symmetry in the 89-plane, the D7-branes can sit at any




Figure 8.3: Possible bending of the D7-branes at nonzero temperature. The asymptotic distance $L \equiv U(\infty)$ is proportional to the bare quark mass M_q , whereas the minimum distance $U(0)$ is related (albeit in a way more complicated than simple proportionality) to the quark thermal mass.

constant angular position; choosing α_0 then breaks the symmetry. Since this $U(1)$ symmetry is respected by the D3-branes' backreaction (i.e. since the $AdS_5 \times S^5$ metric is $U(1)$ -invariant), it is easy to guess (correctly) that $\alpha(u) = \alpha_0$ is still a solution of the D7-branes' equation of motion in the presence of the D3-branes' backreaction.

The second equation, $U(u) = L$, says that the D7-branes lie at a constant distance from the D3-branes. In the absence of the D3-branes' backreaction this is easily understood: there is no force on the D7-branes and therefore they span a perfect 4-plane. In the presence of backreaction, one should generically expect that the spacetime curvature deforms the D7-branes as in Fig. 8.3, bending them towards the D3-branes at the origin. The reason that this does not happen for the D3/D7 system at zero temperature is that the underlying supersymmetry of the system guarantees an exact cancellation of forces on the D7-branes. In fact, it is easy to verify directly that $U(u) = L$ is still an exact solution of the D7-branes' equations of motion in the presence of the D3-branes' backreaction. The constant L then determines the quark mass through Eqn. (8.2). We will see below that the introduction of nonzero temperature breaks supersymmetry completely, and that consequently $U(u)$ becomes a non-constant function that one must solve for, and that this function contains information about the ground state of the theory in the presence of quarks. For example, its asymptotic behaviour encodes the value of the bare quark mass M_q and the quark condensate $\langle \bar{\psi}\psi \rangle$, whereas its value at $u = 0$ is related to the quark thermal mass M_{th} . Since in this section we work at $T = 0$, any nonzero quark mass corresponds to $M_q/T \rightarrow \infty$. In this sense one must think of the quarks in question as the analogue of heavy quarks in QCD, and of the quark condensate as the analogue of $\langle \bar{c}c \rangle$ or $\langle \bar{b}b \rangle$. However, when we consider a nonzero temperature




Figure 8.4: D7-branes' embedding in $AdS_5 \times S^5$. At nonzero temperature this picture is slightly modified: First, a horizon appears at $r = r_0 > 0$, and second, the D7-branes terminate at $r = U(0) < L$. This 'termination point' corresponds to the tip of the branes in Fig. 8.3.

in subsequent sections, whether the holographic quarks described by the D7-branes are the analogues of heavy or light quarks in QCD will depend on how their mass (or, more precisely, the mass of the corresponding mesons) compares to the temperature.

We have concluded that, at zero temperature, the D7-branes lie at $U = L$ and are parametrised by $\{t, x_i, u, \Omega_3\}$. In terms of these coordinates, the induced metric on the D7-branes takes the form

$$ds^2 = \frac{u^2 + L^2}{R^2} (-dt^2 + dx_i^2) + \frac{R^2}{u^2 + L^2} du^2 + \frac{R^2 u^2}{u^2 + L^2} d\Omega_3^2. \quad (8.7)$$

We see that if $L = 0$ then this metric is exactly that of $AdS_5 \times S^3$. The AdS_5 factor suggests that the dual gauge theory should still be conformally invariant. This is indeed the case in the limit under consideration: If $L = 0$ the quarks are massless and the theory is classically conformal, and in the probe limit $N_f/N_c \rightarrow 0$ the quantum mechanical β -function, which is proportional to N_f/N_c , vanishes. If $L \neq 0$ then the metric above becomes $AdS_5 \times S^3$ only asymptotically, i.e. for $u \gg L$, reflecting the fact that in the gauge theory conformal invariance is explicitly broken by the quark mass $M_q \propto L$, but is restored asymptotically at energies $E \gg M_q$. We also note that, if $L \neq 0$, then the radius of the three-sphere is not constant, as displayed in fig. 8.4; in particular, it shrinks to zero at $u = 0$ (corresponding to $r = L$), at which point the D7-branes 'terminate' from the viewpoint of the projection on

AdS_5 [385]. In other words, if $L \neq 0$ then the D7-branes fill the AdS_5 factor of the metric only down to a minimum value of the radial direction proportional to the quark mass. As we anticipated above, at nonzero temperature one must distinguish between the bare and the thermal quark masses, related to $U(\infty)$ and $U(0)$ respectively. In this case the position in AdS at which the D7-branes terminate is $r = U(0) < L$, and therefore they fill the AdS space down to a radial position related to the thermal mass. Note also that at finite temperature a horizon is present at $r = r_0 > 0$.

8.2.2 Meson spectrum

We are now ready to compute the spectrum of low-spin mesons in the D3/D7 system following Ref. [585]. The spectrum for more general Dp/Dq systems was computed in [586–588]. Recall that mesons are described by open strings attached to the D7-branes. In particular, spin-zero and spin-one mesons correspond to the scalar and vector fields on the D7-branes. Large-spin mesons can be described as long, semi-classical strings [556, 585, 589–596], but we will not review them here.

For simplicity, we will focus on scalar mesons. Following Section 5.1.5, we know that in order to determine the spectrum of scalar mesons, we need to determine the spectrum of normalisable modes of small fluctuations of the scalar fields on the D7-branes. At this point we restrict ourselves to a single D7-brane, i.e. we set $N_f = 1$, in which case the dynamics is described by the DBI action (4.18). At leading order in the large- N_c expansion, the spectrum for $N_f > 1$ consists of N_f^2 identical copies of the single-flavour spectrum [597].

We use the coordinates in Eqn. (8.7) as worldvolume coordinates for the D7-brane, which we collectively denote by σ^μ . The physical scalar fields on the D7-brane are then $x^8(\sigma^\mu)$, $x^9(\sigma^\mu)$. By a rotation in the 89-plane we can assume that, in the absence of fluctuations, the D7-brane lies at $x^8 = 0$, $x^9 = L$. Then the fluctuations can be parametrised as

$$x^8 = 0 + \varphi(\sigma^\mu) , \quad x^9 = L + \tilde{\varphi}(\sigma^\mu) , \quad (8.8)$$

with φ and $\tilde{\varphi}$ the scalar fluctuations around the fiducial embedding. In order to determine the normalisable modes, it suffices to work to quadratic order in $\varphi, \tilde{\varphi}$. Substituting (8.8) in the DBI action (4.18) and expanding in $\varphi, \tilde{\varphi}$ leads to a quadratic Lagrangian whose corresponding equation of motion is

$$\frac{R^4}{(u^2 + L^2)^2} \square \varphi + \frac{1}{u^3} \partial_u (u^3 \partial_u \varphi) + \frac{1}{u^2} \nabla^2 \varphi = 0 , \quad (8.9)$$

where \square is the four-dimensional d'Alembertian associated with the Cartesian coordinates t, x_i , and ∇^2 is the Laplacian on the three-sphere. The equation for $\tilde{\varphi}$ takes exactly the same form. Modes that transform non-trivially under rotations on the sphere correspond to mesons that carry nonzero R-charge. Since QCD does not possess an R-symmetry, we will restrict ourselves to R-neutral mesons, meaning that we will assume that φ does not depend on the coordinates of the sphere. We can use separation of variables to write these modes as

$$\varphi = \phi(u) e^{iq \cdot x} , \quad (8.10)$$

where $x = (t, x_i)$. Each of these modes then corresponds to a physical meson state in the gauge theory with a well defined four-dimensional mass given by its eigenvalue under \square , that

is, $M^2 = -q^2$. For each of these modes, Eqn. (8.9) results in an equation for $\phi(u)$ that, after introducing dimensionless variables through

$$\bar{u} = \frac{u}{L}, \quad \bar{M}^2 = -\frac{k^2 R^4}{L^2}, \quad (8.11)$$

becomes

$$\partial_{\bar{u}}^2 \phi + \frac{3}{\bar{u}} \partial_{\bar{u}} \phi + \frac{\bar{M}^2}{(1 + \bar{u}^2)^2} \phi = 0. \quad (8.12)$$

This equation can be solved in terms of hypergeometric functions. The details can be found in Ref. [585], but we will not give them here because most of the relevant physics can be extracted as follows.

Eqn. (8.12) is a second-order, ordinary differential equation with two independent solutions. The combination we seek must satisfy two conditions: It must be normalisable as $\bar{u} \rightarrow \infty$, and it must be regular as $\bar{u} \rightarrow 0$. For arbitrary values of \bar{M} , both conditions cannot be simultaneously satisfied. In other words, the values of \bar{M} for which physically acceptable solutions exist are quantised. Since Eqn. (8.12) contains no dimensionful parameters, the values of \bar{M} must be pure numbers. These can be explicitly determined from the solutions of (8.12) and they take the form [585]

$$\bar{M}^2 = 4(n+1)(n+2), \quad n = 0, 1, 2, \dots \quad (8.13)$$

Using this, and $M^2 = -q^2 = \bar{M}^2 L^2 / R^4$, we derive the result that the four-dimensional mass spectrum of scalar mesons is

$$M(n) = \frac{2L}{R^2} \sqrt{(n+1)(n+2)} = \frac{4\pi M_q}{\sqrt{\lambda}} \sqrt{(n+1)(n+2)}, \quad (8.14)$$

where in the last equality we have used the expressions $R^2/\alpha' = \sqrt{\lambda}$ and (8.2) to write R and L in terms of gauge theory parameters. We thus conclude that the spectrum consists of a discrete set of mesons with a mass gap given by the mass of the lightest meson:³

$$M_{\text{mes}} = 4\pi\sqrt{2} \frac{M_q}{\sqrt{\lambda}}. \quad (8.15)$$

Since this result is valid at large 't Hooft coupling, $\lambda \gg 1$, the mass of these mesons is much smaller than the mass of two constituent quarks. In other words, the mesons in this theory are very deeply bound. In fact, the binding energy

$$E_B \equiv 2M_q - M_{\text{mes}} \lesssim 2M_q \sim \sqrt{\lambda} M_{\text{mes}} \quad (8.16)$$

is so large that it almost cancels the rest energy of the quarks. This is clear from the gravity picture of ‘meson formation’ (see Fig. 8.5), in which two strings of opposite orientation stretching from the D7-brane to $r = 0$ (the quark-antiquark pair) join together to form an open string with both ends on the D7-brane (the meson). This resulting string is much shorter than the initial ones, and hence corresponds to a configuration with much lower energy. This feature is an important difference with quarkonium mesons in QCD, such as

³In order to compare this and subsequent formulas with ref. [585] and others, note that our definition (4.17) of g^2 differs from the definition in some of those references by a factor of 2, for example $g_{\text{[here]}}^2 = 2g_{\text{[585]}}^2$.




Figure 8.5: String description of a quark, an antiquark and a meson. The string that describes the meson can be much shorter than those describing the quark and the antiquark.

charmonium or bottomonium, which are not deeply bound. Although this certainly means that caution must be exercised when trying to compare the physics of quarkonium mesons in holographic theories with the physics of quarkonium mesons in QCD, the success or failure of these comparisons cannot be assessed at this point. We will discuss this assessment in detail below, once we have learned more about the physics of holographic mesons. Suffice it to say here that some of this physics, such as the temperature or the velocity dependence of certain meson properties, turns out to be quite general and may yield insights into some of the challenges related to understanding the physics of quarkonia within the QCD quark-gluon plasma.

We close this section with a consistency check. The behaviour of the fluctuation modes at infinity is related to the high-energy properties of the theory. At high energy, we can ignore the effect of the mass of the quarks and the theory becomes conformal. The $u \rightarrow \infty$ behaviour is then related to the UV operator of the lowest conformal dimension, Δ , that has the same quantum numbers as the meson [252, 253]. Analysis of this behaviour for the solutions of Eqs. (8.9), (8.12) shows that $\Delta = 3$ [585], as expected for a quark-bilinear operator.

8.3 Nonzero temperature

8.3.1 D7-brane embeddings

We now turn to the case of nonzero temperature, $T \neq 0$. This means that we must study the physics of a D7-brane in the black brane metric (cf. Eqn. (5.30))

$$ds^2 = \frac{r^2}{R^2} (-f dt^2 + dx_1^2 + dx_2^2 + dx_3^2) + \frac{R^2}{r^2 f} dr^2 + R^2 d\Omega_5^2, \quad (8.17)$$

where

$$f(r) = 1 - \frac{r_0^4}{r^4}, \quad r_0 = \pi R^2 T. \quad (8.18)$$

The study we must perform is conceptually analogous to that of the last few sections, but the equations are more involved and most of them must be solved numerically. These technical details are not very illuminating, and for this reason we will not dwell into them. Instead, we will focus on describing in detail the main conceptual points and results, as well as the physics behind them, which in fact can be understood in very simple and intuitive terms.

As mentioned above, at $T \neq 0$ all supersymmetry is broken. We therefore expect that the D7-branes will be deformed by the non-trivial geometry. In particular, the introduction of nonzero temperature corresponds, in the string description, to the introduction of a black brane in the background. Intuitively, we expect that the extra gravitational attraction will bend the D7-branes towards the black hole. This simple conclusion, which was anticipated in previous sections, has far-reaching consequences. At a qualitative level, most of the holographic physics of mesons in a strongly coupled plasma follows from this conclusion. An example of the D7-branes' embedding for a small value of T/M_q is depicted in two slightly different ways in Fig. 8.6.

The qualitative physics of the D3/D7 system as a function of the dimensionless ratio T/M_q is now easy to guess, and is captured by Fig. 8.7. At zero temperature the horizon has zero size and the D7-branes span an exact hyperplane. At nonzero but sufficiently small T/M_q , the gravitational attraction from the black hole pulls the branes down but the branes' tension can still compensate for this. The embedding of the branes is thus deformed, but the branes remain entirely outside the horizon. Since in this case the induced metric on the D7-branes has no horizon, we will call this type of configuration a 'Minkowski embedding'. In contrast, above a critical temperature T_{diss} ,⁴ the gravitational force overcomes the tension of the branes and these are pulled into the horizon. In this case the induced metric on the branes possesses an event horizon, inherited from that of the spacetime metric. For this reason we will refer to such configurations as 'black hole embeddings'. Between these two types of embeddings there exists a so-called critical embedding in which the branes just 'touch the horizon at a point'. The existence of such an interpolating solution might lead one to suspect that the phase transition between Minkowski and black hole embeddings is continuous, i.e. of second or higher order. However, as we will see in the next section, thermodynamic considerations reveal that a first-order phase transition occurs between a Minkowski and a black hole embedding. In other words, the critical embedding is skipped over by the phase transition, and near-critical embeddings turn out to be metastable or unstable.

As illustrated by the figures above, the fact that the branes bend towards the horizon implies that the asymptotic distance between the two differs from their minimal distance. As we will see in Section 8.3.2, the asymptotic distance is proportional to the microscopic or 'bare' quark mass, since it is determined by the non-normalizable mode of the field that describes the branes' bending. In contrast, the minimal distance between the branes and the horizon includes thermal (and quantum) effects, and for this reason we will refer to the mass of a string stretching between the bottom of the branes and the horizon (shown as a wiggly red curve in the figures) as the 'thermal' quark mass. Note that this vanishes in the black

⁴The reason for the subscript will become clear shortly.




Figure 8.6: D7-branes' embedding for small T/M_q . The branes bend towards the horizon, shown in dark grey. The radius of the horizon is proportional to its Hawking temperature, which is identified with the gauge theory temperature T — see Eqn. (8.18). The asymptotic position of the D7-branes is proportional to the bare quark mass, M_q . The minimum distance between the branes and the horizon is related to the thermal quark mass, because this is the minimum length of a string (shown as a red wiggly line) stretching between the branes and the horizon. The top figure shows the two relevant radial directions in the space transverse to the D3-branes, U and u (introduced in Eqs. (8.4) and (8.5)), together with the gauge theory directions x_i (time is suppressed). The horizon has topology $\mathbb{R}^3 \times S^5$, where the first factor corresponds to the gauge-theory directions. This 'cylinder-like' topology is manifest in the top figure. Instead, in the bottom figure the gauge theory directions are suppressed and the S^3 wrapped by the D7-branes in the space transverse to the D3-branes is shown, as in Figs. 8.2 and 8.3. In this figure only the S^5 factor of the horizon is shown.




Figure 8.7: Various D7-brane configurations in a black D3-brane background with increasing temperature from left to right. At low temperatures, the probe branes close off smoothly above the horizon. At high temperatures, the branes fall through the event horizon. In between, a critical solution exists in which the branes just ‘touch’ the horizon at a point. The critical configuration is never realized: a first-order phase transition occurs from a Minkowski to a black hole embedding (or vice versa) before the critical solution is reached.

hole phase.

Although we will come back to this important point below, we wish to emphasize right from the start that *the phase transition under discussion is not a confinement-deconfinement phase transition*, since the presence of a black hole implies that both phases are deconfined. Instead, we will see that *the branes’ phase transition corresponds to the dissociation of heavy quarkonium mesons*. In order to illustrate the difference most clearly, consider first a holographic model of a confining theory, as described in Section 5.2.2; below we will come back to the case of $\mathcal{N} = 4$ SYM. For all such confining models, the difference between the deconfinement and the dissociation phase transitions is illustrated in Fig. 8.8. Below T_c , the theory is in a confining phase and therefore no black hole is present. At some T_c , a deconfinement transition takes place, which in the string description corresponds to the appearance of a black hole whose size is proportional to T_c . If the quark mass is sufficiently large compared to T_c then the branes remain outside the horizon (top part of the figure); otherwise they fall through the horizon (bottom part of the figure). The first case corresponds to heavy quarkonium mesons that remain bound in the deconfined phase, and that eventually dissociate at some higher $T_{\text{diss}} > T_c$. The second case describes light mesons that dissociate as soon as the deconfinement transition takes place.

Fig. 8.8 also applies to $\mathcal{N} = 4$ SYM theory with $T_c = 0$ in the sense that, although the vacuum of the theory is not confining, there is no black hole at $T = 0$. Note also that mesons only exist provided $M_q > 0$, since otherwise the theory is conformal and there is no particle spectrum. This means that in $\mathcal{N} = 4$ SYM theory any meson is a heavy quarkonium meson that remains bound for some range of temperatures above $T_c = 0$, as described by the top part of Fig. 8.8. In the case $M_q = 0$ we cannot properly speak of mesons, but we see that the situation is still described by the bottom part of the figure in the sense that in this case the branes fall through the horizon as soon as T is raised above $T_c = 0$.

The universal character of the meson dissociation transition was emphasized in Refs. [565,




Figure 8.8: Top: Sufficiently heavy quarkonium mesons remain bound in the deconfined phase (above T_c) and dissociate at $T_{\text{diss}} > T_c$. Bottom: In contrast, light mesons dissociate as soon as the deconfinement phase transition at $T = T_c$ takes place. This picture also applies to $\mathcal{N} = 4$ SYM theory with $T_c = 0$, as described in the main text. In $\mathcal{N} = 4$ SYM theory, the top (bottom) panel applies when $M_q > 0$ ($M_q = 0$).

[598], which we will follow in our presentation. Specific examples were originally seen in [597, 599, 600], and aspects of these transitions in the D3/D7 system were studied independently in [601–604]. Similar holographic transitions appeared in a slightly different framework in [348, 605–607]. The D3/D7 system at nonzero temperature has been studied upon including the backreaction of the D7-branes in Ref. [464].

8.3.2 Thermodynamics of D7-branes

In this subsection we shall show that the phase transition between Minkowski and black hole embeddings is a discontinuous, first-order phase transition. The reader willing to accept this without proof can safely skip to Section 8.3.3. Since we are working in the canonical ensemble (i.e. at fixed temperature) we must compute the free energy density of the system per unit gauge theory three-volume, F , and determine the configuration that minimizes it. In the gauge theory we know that this takes the form

$$F = F_{\mathcal{N}=4} + F_{\text{flavour}}, \quad (8.19)$$

where the first term is the $O(N_c^2)$ free energy of the $\mathcal{N} = 4$ SYM theory in the absence of quarks, and the second term is the $O(N_c N_f)$ contribution due to the presence of quarks in the fundamental representation. Since the SYM theory without quarks is conformal, dimensional analysis completely fixes the first factor to be of the form $F_{\mathcal{N}=4} = C(\lambda)T^4$, where C is a possibly coupling-dependent coefficient of order N_c^2 . In contrast, in the presence of quarks of mass M_q there is a dimensionless ratio T/M_q on which the flavour contribution can depend non-trivially. Our purpose is to determine this contribution to leading order in the large- N_c , strong coupling limit.

Our tool is of course the dual description of the $\mathcal{N} = 4$ SYM theory with flavour as a system of N_f D7-brane probes in the gravitational background of N_c black D3-branes. As usual in finite-temperature physics, the free energy of the system may be computed through the identification

$$\beta F = S_E, \quad (8.20)$$

where $\beta = 1/T$ and S_E is the Euclidean action of the system. In our case this takes the form

$$S_E = S_{\text{sugra}} + S_{\text{D7}}. \quad (8.21)$$

The first term is the contribution from the black hole gravitational background sourced by the D3-branes, and is computed by evaluating the Euclideanized supergravity action on this background. The second term is the contribution from the D7-branes, and is computed by evaluating the Euclidean version of the DBI action (4.18) on a particular D7-brane configuration. The decomposition (8.21) is the dual version of that in (8.19). The supergravity action scales as $1/g_s^2 \sim N_c^2$, and thus yields the free energy of the $\mathcal{N} = 4$ SYM theory in the absence of quarks, i.e. we identify

$$S_{\text{sugra}} = \beta F_{\mathcal{N}=4}. \quad (8.22)$$

Similarly, the D7-brane action scales as $N_f/g_s \sim N_c N_f$, and represents the flavour contribution to the free energy:

$$S_{\text{D7}} = \beta F_{\text{D7}} = \beta F_{\text{flavour}}. \quad (8.23)$$

We therefore conclude that we must first find the solutions of the equations of motion of the D7-branes for any given values of T and M_q , then evaluate their Euclidean actions, and finally use the identification above to compare their free energies and determine the thermodynamically preferred configuration.

As explained above, in our case solving the D7-brane equations of motion just means finding the function $U(u)$, which is determined by the condition that the D7-brane action be extremized. This leads to an ordinary, second-order, non-linear differential equation for $U(u)$. Its precise form can be found in e.g. Ref. [565], but is not very illuminating. However, it is easy to see that it implies the asymptotic, large- u behaviour

$$U(u) = \frac{m r_0}{\sqrt{2}} + \frac{c r_0^3}{2\sqrt{2} u^2} + \dots, \quad (8.24)$$

where m and c are constants. The factors of r_0 have been introduced to make these constants dimensionless, whereas the numerical factors have been chosen to facilitate comparison with the literature. As usual (and, in particular, as in Section 5.1.5), the leading and subleading terms correspond to the non-normalizable and to the normalizable modes, respectively. Their coefficients are therefore proportional to the source and the expectation values of the corresponding dual operator in the gauge theory. In this case, the position of the brane $U(u)$ is dual to the quark-mass operator $\mathcal{O}_m \sim \bar{\psi}\psi$, so m and c are proportional to the quark mass and the quark condensate, respectively. The precise form of \mathcal{O}_m can be found in Ref. [335], where it is shown that the exact relation between m, c and $M_q, \langle \mathcal{O}_m \rangle$ is

$$M_q = \frac{r_0 m}{2^{3/2} \pi \ell_s^2} = \frac{1}{2\sqrt{2}} \sqrt{\lambda} T m, \quad (8.25)$$

$$\langle \mathcal{O}_m \rangle = -2^{3/2} \pi^3 \ell_s^2 N_f T_{D7} r_0^3 c = -\frac{1}{8\sqrt{2}} \sqrt{\lambda} N_f N_c T^3 c. \quad (8.26)$$

In particular, we recover the fact that the asymptotic value

$$L = \lim_{u \rightarrow \infty} U(u) = \frac{m r_0}{\sqrt{2}} \quad (8.27)$$

is related to the quark mass through Eqn. (8.2), as anticipated in previous sections.

It is interesting to note that the dimensionless mass m is given by the simple ratio

$$m = \frac{\bar{M}}{T}, \quad (8.28)$$

where

$$\bar{M} = \frac{2\sqrt{2} M_q}{\sqrt{\lambda}} = \frac{M_{\text{mes}}}{2\pi} \quad (8.29)$$

is (up to a constant) precisely the meson gap at zero temperature, given in Eqn. (8.15). As mentioned in Section 8.2.1, and as we will elaborate upon in Section 8.4.3, \mathcal{O}_m must be thought of as the analogue of a heavy- or light-quark bilinear operator in QCD depending on whether the ratio $M_{\text{mes}}/T \sim m$ is large or small, respectively.

The constants m and c can be understood as the two integration constants that completely determine a solution of the second-order differential equation obeyed by $U(u)$. Mathematically, these two constants are independent, but the physical requirement that the solution




Figure 8.9: Quark condensate c versus $T/\bar{M} = 1/m$. The blue dashed (red continuous) curves correspond to the Minkowski (black hole) embeddings. The dotted vertical line indicates the precise temperature of the phase transition. The point where the two branches meet corresponds to the critical embedding.




Figure 8.10: Some representative D7-brane embeddings from the region in which c is multi-valued. The three profiles correspond to the same value of m but differ in their value of c . Two of them, represented by blue, dashed curves, are of Minkowski type. The third one, represented by a red, continuous curve, is a black hole embedding.

be regular in the interior relates them to one another. The equation for $U(u)$ can be solved numerically (see, e.g. Ref. [565]), and the resulting possible values of c for each value of m are plotted in Fig. 8.9. We see from the ‘large-scale’ plot on the left that c is a single-valued function of m for most values of the latter. However, the zoom-in plot on the right shows that, in a small region around $1/m = T/\bar{M} \simeq 0.766$, three values of c are possible for a given value of m ; a pictorial representation of a situation of this type is shown in Fig. 8.10. This multivaluedness is related to the existence of the phase transition which, as we will see, proceeds between points A and B through a discontinuous jump in the quark condensate and other physical quantities. The point in Fig. 8.9 where the Minkowski and the black hole branches meet corresponds to the critical embedding.

Having determined the regular D7-brane configurations, one must now compute their free energies and compare them in order to determine which one is preferred in the multivalued region. The result is shown in Fig. 8.11(top), where the normalization constant is given

by [565, 598]

$$\mathcal{N} = \frac{2\pi^2 N_f T_{D7} r_0^4}{4T} = \frac{\lambda N_f N_c}{64} T^3. \quad (8.30)$$

The plot on the right shows the classic ‘swallow tail’ form, typically associated with a first-order phase transition. As anticipated, Minkowski embeddings have the lowest free energy for temperatures $T < T_{\text{diss}}$, whereas the free energy is minimized by black hole embeddings for $T > T_{\text{diss}}$, with $T_{\text{diss}} \simeq 0.77\bar{M}$ (i.e. $m \simeq 1.3$). At $T = T_{\text{diss}}$ the Minkowski and the black hole branches meet and the thermodynamically-preferred embedding changes from one type to the other. The first-order nature of the phase transition follows from the fact that several physical quantities jump discontinuously across the transition. An example is provided by the quark condensate which, as illustrated in Fig. 8.9, makes a finite jump between the points labelled A and B. Similar discontinuities also appear in other physical quantities, like the entropy and energy density. These are easily obtained from the free energy through the usual thermodynamic relations

$$S = -\frac{\partial F}{\partial T}, \quad E = F + TS, \quad (8.31)$$

and the results are shown in Fig. 8.11. From the plots of the energy density one can immediately read off the qualitative behaviour of the specific heat $c_V = \partial E / \partial T$. In particular, note that this slope must become negative as the curves approach the critical solution, indicating that the corresponding embeddings are thermodynamically unstable. Examining the fluctuation spectrum of the branes, we will show that a corresponding dynamical instability, manifested by a meson state becoming tachyonic, is present exactly for the same embeddings for which $c_V < 0$. One may have thought that the phases near the critical point were metastable and thus accessible by ‘super-cooling’ the system, but instead it turns out that over much of the relevant regime such phases are unstable.

We see from (8.30) that $\mathcal{N} \sim \lambda N_c N_f T^3$, which means that the leading contribution of the D7-branes to all the various thermodynamic quantities will be order $\lambda N_c N_f$, in comparison to N_c^2 for the usual bulk gravitational contributions. The $N_c N_f$ dependence, anticipated below Eqn. (8.19), follows from large- N_c counting. In contrast, as noted in Ref. [598, 608], the factor of λ represents a strong-coupling enhancement over the contribution of a simple free-field estimate for the $N_c N_f$ fundamental degrees of freedom. From the viewpoint of the string description, this enhancement is easy to understand by reexamining the relative normalization of the two terms in Eqn. (8.21) more carefully than we did above. Ignoring only order-one, purely numerical factors, the supergravity action scales as $1/G$, with $G \sim g_s^2 \ell_s^8$ the ten-dimensional Newton’s constant, whereas the D7-brane action scales as $N_f T_{D7} \sim N_f / g_s \ell_s^8$. The ratio between the two normalizations is therefore

$$GN_f T_{D7} \sim g_s N_f \sim g^2 N_f \sim \frac{\lambda N_f}{N_c}. \quad (8.32)$$

Thus the flavour contribution is suppressed with respect to the leading $O(N_c^2)$ contribution by $\lambda N_f / N_c$, i.e. it is of order $\lambda N_c N_f$. We will come back to this point in the next subsection.

As the calculations above were all performed in the limit $N_c, \lambda \rightarrow \infty$ (with N_f fixed), it is natural to ask how the detailed results depend on this approximation. Since the phase transition is first order, we expect that its qualitative features will remain unchanged within a finite radius of the $1/N_c, 1/\lambda$ expansions. Of course, finite- N_c and finite- λ corrections may

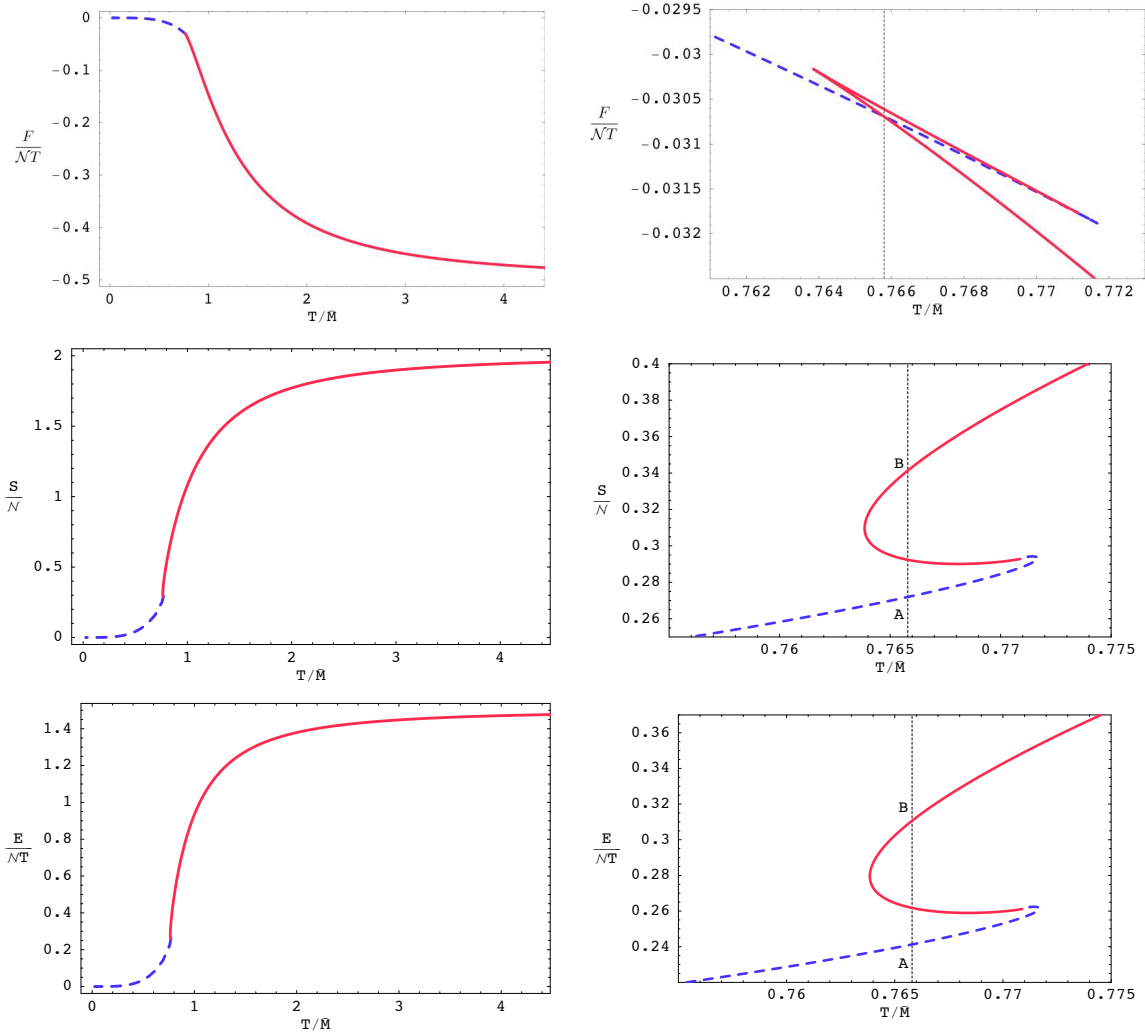


Figure 8.11: Free energy, entropy and energy densities for a D7-brane in a black D3-brane background; note that $\mathcal{N} \propto T^3$. The blue dashed (red continuous) curves correspond to the Minkowski (black hole) embeddings. The dotted vertical line indicates the precise temperature of the phase transition.

eventually modify the behaviour described above. For example, at large but finite N_c the black hole will emit Hawking radiation and each bit of the probe branes will experience a thermal bath at a temperature determined by the local acceleration. Similarly, finite 't Hooft coupling corrections, which correspond to higher-derivative corrections both to the supergravity action and the D-brane action, will become important if the spacetime or the brane curvatures become large. It is certainly clear that both types of corrections will become more and more important as the lower part of a Minkowski brane approaches the horizon, since as this happens the local temperature and the branes (intrinsic) curvature at their tip increase. However, at the phase transition the minimum separation between the branes and the horizon is not parametrically small, and therefore the corrections above can be made arbitrarily small by taking N_c and λ sufficiently large but still finite. This confirms our expectation on general grounds that the qualitative aspects of the phase transition should be robust within a finite radius around the $1/N_c = 0$, $1/\lambda = 0$ point. Of course, these considerations do not tell us whether the dissociation transition is first order or a crossover at $N_c = 3$.

8.3.3 Quarkonium thermodynamics

We have seen above that, in a large class of strongly coupled gauge theories with fundamental matter, this matter undergoes a first-order phase transition described on the gravity side by a change in the geometry of the probe D-branes. In this section we will elaborate on thermodynamical aspects of this transition from the gauge-theory viewpoint. Once we have learned more about the dynamics of holographic mesons in subsequent sections, in Section 8.4.3 we will return to the gauge-theory viewpoint and discuss possible implications for the dynamics of quarkonium mesons in the QCD plasma.

The temperature scale at which the phase transition takes place is set by the meson gap at zero temperature, $T_{\text{diss}} \sim M_{\text{mes}}$. As well as giving the mass gap in the meson spectrum, $1/M_{\text{mes}}$ is roughly the characteristic size of these bound states [587, 609]. The gluons and other adjoint fields are already in a deconfined phase at T_{diss} , so this new transition is not a confinement/deconfinement transition. Rather, the most striking feature of the new phase transition is the change in the meson spectrum, and so we refer to it as a ‘dissociation’ or ‘melting’ transition.

In the low-temperature phase, below the transition, stable mesons exist and their spectrum is discrete and gapped. This follows from the same general principles as in the zero-temperature case. The meson spectrum corresponds to the spectrum of normalizable fluctuations of the D7-branes around their fiducial embedding. For Minkowski embeddings the branes close off smoothly outside the black hole horizon and the admissible modes must also satisfy a regularity condition at the tip of the branes. On general grounds, we expect that the regular solution at the tip of the branes evolves precisely into the normalizable solution at the boundary only for a certain set of discrete values of the meson mass. We will study the meson spectrum in detail in Section 8.4.1, and in Section 8.4.2 we will see that mesons acquire finite decay widths at finite N_c or finite coupling. Since the phase under consideration is not a confining phase, we can also introduce deconfined quarks into the system, represented by fundamental strings stretching between the D7-branes and the horizon. At a figurative level, in this phase we might describe quarks in the adjoint plasma as a ‘suspension’. That is, when

quarks are added to this phase, they retain their individual identities. More technically, we may just say that quarks are well defined quasiparticles in the Minkowski phase.

In the high-temperature phase, at $T > T_{\text{diss}}$, no stable mesons exist. Instead, as we will discuss in more detail in Section 8.5, the excitations of the fundamental fields in this phase are characterised by a discrete spectrum of quasinormal modes on the black hole embeddings [610, 611]. The spectral function of some two-point meson correlators in the holographic theory, of which we will see an example in Section 8.5.2, still exhibits some broad peaks in a regime just above T_{diss} , which suggests that a few broad bound states persist just above the dissociation phase transition [611, 612]. This is quite analogous to the lattice approach, where similar spectral functions are examined to verify the presence or absence of bound states. Hence, identifying T_{diss} with the dissociation temperature should be seen as a (small) underestimate of the temperature at which mesons completely cease to exist. An appropriate figurative characterization of the quarks in this high-temperature phase would be as a ‘solution’. If one attempts to inject a localised quark charge into the system, it quickly falls through the horizon, i.e. it spreads out across the entire plasma and its presence is reduced to diffuse disturbances of the supergravity and worldvolume fields, which are soon damped out [610, 611]. Technically speaking, we may just state that quarks are not well defined quasiparticles in the black hole phase.

The physics above is potentially interesting in connection with QCD, since evidence from several sources indicates that heavy quarkonium mesons remain bound in a range of temperatures above T_c . We will analyze this connection in more detail in Section 8.4.3, once we have learned more about the properties of holographic mesons in subsequent sections. Here we would just like to point out one simple physical parallel. The question of quarkonium bound states surviving in the quark-gluon plasma was first addressed by comparing the size of the bound states to the screening length in the plasma [184]. In the D3/D7 system, the size of the mesons can be inferred, for example, from the structure functions, and the relevant length scale that emerges is $d_{\text{mes}} \sim \sqrt{\lambda}/M_q$ [609]. This can also be heuristically motivated as follows. As discussed in Section 5.4 (see Eqn. (5.81)) at zero temperature the potential between a quark-antiquark pair separated by a distance ℓ is given by

$$V \sim -\frac{\sqrt{\lambda}}{\ell} . \quad (8.33)$$

We can then estimate the size d_{mes} of a meson by requiring $E_B \sim |V(d_{\text{mes}})|$, where E_B is the binding energy (8.16).⁵ This gives

$$d_{\text{mes}} \sim \frac{\sqrt{\lambda}}{E_B} \sim \frac{\sqrt{\lambda}}{M_q} \sim \frac{1}{M_{\text{mes}}} \sim \frac{R^2}{L} . \quad (8.34)$$

The last equality follows from Eqn. (8.14) and is consistent with expectations based on the UV/IR correspondence [587], since on the gravity side mesons are excitations near $r = L$. Just for comparison, we remind the reader that the weak-coupling formula for the size of quarkonium is $d_{\text{weak}} \sim 1/(g^2 M_q)$.

⁵Eqn. (8.16) was derived at zero temperature, but as we will see in Section 8.4.1 it is also parametrically correct at nonzero temperature.

One intuitive way to understand why a meson has a very large size compared to its inverse binding energy or to the inverse quark mass is that, due to strong-coupling effects, the quarks themselves have an effective size of order d_{mes} . The effective size of a quark is defined as the largest of the following two scales: (i) its Compton wavelength, or (ii) the distance between a quark-antiquark pair at which their potential energy is large enough to pair-produce additional quarks and antiquarks. In a weakly coupled theory (i) is larger, whereas in a strongly coupled theory (ii) is larger. From Eqn. (8.33) we see that this criterion gives an effective quark size of order $\sqrt{\lambda}/M_q$ instead of $1/M_q$. This heuristic estimate is supported by an explicit calculation of the size of the gluon cloud that dresses a quark [613]. These authors computed the expectation value $\langle \text{Tr} F^2(x) \rangle$ sourced by a quark of mass M_q and found that the characteristic size of the region in which this expectation value is nonzero is precisely $\sqrt{\lambda}/M_q$.

As reviewed in Section 5.4.2, holographic studies of Wilson lines at nonzero temperature [378, 379] reveal that the relevant screening length of the SYM plasma is of order $L_s \sim 1/T$ — see Eqn. (7.150). The argument that the mesons should dissociate when the screening length is shorter than the size of these bound states then yields $T_{\text{diss}} \sim M_q/\sqrt{\lambda} \sim M_{\text{mes}}$, in agreement with the results of the detailed calculations explained in previous sections. We thus see that the same physical reasoning which is used in QCD to estimate the dissociation temperature of, e.g., the J/ψ meson can also be used to understand the dissociation of mesons in the the $\mathcal{N} = 4$ SYM theory. This may still seem counterintuitive in view of the fact that the binding energy of these mesons is much larger than T_{diss} . In other words, one might have expected that the temperature required to break apart a meson would be of the order of the binding energy, $E_B \sim M_q$, instead of being parametrically smaller,

$$T_{\text{diss}} \sim M_{\text{mes}} \sim E_B/\sqrt{\lambda}. \quad (8.35)$$

However, this intuition relies on the expectation that the result of dissociating a meson is a quark-antiquark pair of mass $2M_q$. The gravity description makes it clear that this is not the case at strong coupling, since above T_{diss} the branes fall through the horizon. Heuristically, one may say that this means that the ‘constituent’ or ‘thermal’ mass of the quarks becomes effectively zero. However, a more precise statement is simply that in the black hole phase quark-like quasiparticles simply do not exist, and therefore for the purpose of the present discussion it becomes meaningless to attribute a mass to them.

One point worth emphasizing is that there are two distinct processes that are occurring at $T \sim M_{\text{mes}}$. If we consider e.g. the entropy density in Fig. 8.11, we see that the phase transition occurs in the midst of a crossover signalled by a rise in S/T^3 . We may write the contribution of the fundamental matter to the entropy density as

$$S_{\text{flavour}} = \frac{1}{8} \lambda N_f N_c T^3 H(x), \quad (8.36)$$

where $x = \lambda T^2/M_q$ and $H(x)$ is the function shown in the plot of the free energy density in the top panels of Fig. 8.11. H rises from 0 at $x = 0$ to 2 as $x \rightarrow \infty$, but the most dramatic part of this rise occurs in the vicinity of $x = 1$. Hence it seems that new degrees of freedom, i.e. the fundamental quarks, are becoming ‘thermally activated’ at $T \sim M_{\text{mes}}$. We note that the phase transition produces a discontinuous jump in which H only increases by

about 0.07, i.e. the jump at the phase transition only accounts for about 3.5% of the total entropy increase. Thus the phase transition seems to play a small role in this crossover and produces relatively small changes in the thermal properties of the fundamental matter, such as the energy and entropy densities.

As M_{mes} sets the scale of the mass gap in the meson spectrum, it is tempting to associate the crossover above with the thermal excitation of mesonic degrees of freedom. However, the pre-factor $\lambda N_f N_c$ in (8.36) indicates that this reasoning is incorrect: if mesons provided the relevant degrees of freedom, we should have $S_{\text{flavour}} \propto N_f^2$. Such a contribution can be obtained either by a one-loop calculation of the fluctuation determinant around the classical D7-brane configuration, or by taking into consideration the D7-branes backreaction to second order in the N_f/N_c expansion as in [464, 614–616]. One can make an analogy here with the entropy of a confining theory (cf. Section 5.2.2). In the low-temperature, confining phase the absence of a black hole horizon implies that the classical-gravity saddle point yields zero entropy, which means that the entropy is zero at order N_c^2 . One must look at the fluctuation determinant to see the entropy contributed by the supergravity modes, i.e. by the gauge-singlet glueballs, which is of order N_c^0 .

We thus see that the factor of $N_f N_c$ in S_{flavour} is naturally interpreted as counting the number of degrees of freedom associated with deconfined quarks, with the factor of λ demonstrating that the contribution of the quarks is enhanced at strong coupling. A complementary interpretation of (8.36) comes from reorganizing the pre-factor as

$$\lambda N_f N_c = (g^2 N_f) N_c^2. \quad (8.37)$$

The latter expression suggests that the result corresponds to the first-order correction of the adjoint entropy due to quark loops. As explained at the end of Section 5.5.1, we are working in a ‘not quite’ quenched approximation, in that contributions of the D7-branes represent the leading order contribution in an expansion in N_f/N_c , and so quark loops are suppressed but not completely. In view of the discussion below Eqn. (8.32), it is clear that the expansion for the classical gravitational back-reaction of the D7-branes is controlled by $\lambda N_f/N_c = g^2 N_f$. Hence this expansion corresponds to precisely the expansion in quark loops on the gauge theory side.

We conclude that the strongly coupled theory brings together these two otherwise distinct processes. That is, because the $\mathcal{N} = 4$ SYM theory is strongly coupled at all energy scales, the dissociation of the quarkonium bound states and the thermal activation of the quarks happen at essentially the same temperature. Note that this implies that the phase transition should not be thought of as exclusively associated with a discontinuous change in the properties of mesons — despite the fact that this is the aspect that is more commonly emphasized. The phase transition is also associated with a discontinuous change in the properties of quarks since, as explained above, these exist as well defined quasiparticles in the Minkowski phase but not in the black hole phase. In fact, as the discussion around Eqn. (8.37) makes clear, in the $\mathcal{O}(N_f N_c)$ -approximation considered here the observed discontinuous jump in the thermodynamic functions comes entirely from the discontinuous change in the properties of quarks. In this approximation, the discontinuous jump in the thermodynamic functions associated with the discontinuous change in the properties of mesons simply cannot be detected, since it is of order N_f^2 and its determination would require a one-loop calculation. Fortunately,




Figure 8.12: A qualitative representation of the simplest possibility interpolating between the weak- and the strong-coupling regimes in $\mathcal{N} = 4$ SYM theory. The solid and the dotted black curves correspond to $T = T_{\text{diss}}$. At strong coupling this corresponds to a first-order phase transition (solid black curve), whereas at weak coupling it corresponds to a crossover (dotted black curve). The dashed red curve corresponds to $T = T_{\text{activ}}$. At strong coupling this takes place immediately after the phase transition, whereas at weak coupling it is widely separated from T_{diss} .

however, the change in the mesons' properties can be inferred, e.g. from the comparison of their spectra above and below T_{diss} .

It is instructive to contrast this behavior with that which is expected to occur at weak coupling. In this regime, one expects that the dissociation of the quarkonium mesons may well be just a crossover rather than a (first-order) transition. Moreover, since the weakly bound mesons are much larger than $1/M_{\text{mes}} \sim 1/(2M_q)$, their dissociation transition will occur at a T_{diss} that is much lower than M_q . On the other hand, the quarks would not be thermally activated until the temperature $T_{\text{activ}} \sim M_q$, above which the number densities of unbound quarks and antiquarks are no longer Boltzmann-suppressed. Presumably, the thermal activation would again correspond to a crossover rather than a phase transition. The key point is that these two temperatures are widely separated at weak coupling. Fig. 8.12 is an ‘artistic’ representation of the simplest behaviour which would interpolate between strong and weak coupling. One might expect that the dissociation point and the thermal activation are very close for $\lambda \gg 1$. The line of first-order phase transitions must end somewhere and so one might expect that it terminates at a critical point around $\lambda \sim 1$. Below this point, both processes would only represent crossovers and their respective temperatures would diverge from one another, approaching the weak-coupling behaviour described above.

We close with a comment about a possible comparison to QCD. Although it would be interesting to look for signs of a crossover or a phase transition associated with quarkonium dissociation, for example in lattice QCD, the above discussion makes it clear that much caution must be exercised in trying to compare with the holographic results described here. The differences can be traced back to the fact that, unlike the holographic theory considered

here, QCD is not strongly coupled at the scale set by the mass of the heavy quark or of the corresponding heavy quarkonium meson. For this reason, in QCD the binding energy of a quarkonium meson is $E_B \ll M_{\text{mes}} \lesssim 2M_q$ and, since one expects that $T_{\text{diss}} \sim E_B$, this implies that at the dissociation temperature the quarkonium contribution to (say) the total entropy density would be Boltzmann suppressed, i.e. it would be of order $S_{\text{flavour}} \sim N_f^2 \exp(-M_{\text{mes}}/T_{\text{diss}}) \ll 1$. In contrast, in the holographic setup there is no exponential suppression because $T_{\text{diss}} \sim M_{\text{mes}}$. Note also that the quarkonium contribution should scale as N_f^2 , and therefore the exponential suppression is a further suppression on top of the already small one-loop contribution discussed in the paragraph above Eqn. (8.37). That is, there are two sources of suppression relative to the leading $O(N_f N_c)$ -contribution in the holographic theory. Although N_f/N_c is not small in QCD, the Boltzmann suppression is substantial and will likely make the thermodynamic effects of any quarkonium dissociation transition quite a challenge to identify.

8.4 Quarkonium mesons in motion and in decay

In previous sections, we examined the thermodynamics of the phase transition between Minkowski and black hole embeddings, and we argued that from the gauge-theory viewpoint it corresponds to a meson-dissociation transition. In particular, we argued that quarkonium bound states exist on Minkowski embeddings, i.e. at $T < T_{\text{diss}}$, that they are absolutely stable in the large- N_c , strong-coupling limit, and that their spectrum is discrete and gapped. We will begin Section 8.4.1 by studying this spectrum quantitatively, which will allow us to understand how the meson spectrum is modified with respect to that at zero temperature, described in Section 8.2.2. The spectrum on black hole embeddings will be considered in Section 8.5.

After describing the spectrum of quarkonium mesons at rest, we will determine their dispersion relations. This will allow us to study mesons in motion with respect to the plasma and, in particular, to determine how the dissociation temperature depends on the meson velocity. As discussed in Section 2.4, one of the hallmarks of a quark-gluon plasma is the screening of colored objects. Heavy quarkonia provide an important probe of this effect since the existence (or absence) of quark-antiquark bound states and their properties are sensitive to the screening properties of the medium in which they are embedded. In Section 7.7 we studied this issue via computing the potential between an external quark-antiquark pair, at rest in the plasma or moving through it with velocity v . In particular, we found that the dissociation temperature scales with v as

$$T_{\text{diss}}(v) \simeq T_{\text{diss}}(v=0)(1-v^2)^{1/4}, \quad (8.38)$$

which could have important implications for the phenomenon of quarkonium suppression in heavy ion collisions. By studying dynamical mesons in a thermal medium, we will be able to reexamine this issue in a more ‘realistic’ context.

We will show in Section 8.4.2 that both finite- N_c and finite-coupling corrections generate nonzero meson decay widths, as one would expect in a thermal medium. We shall find that the dependence of the widths on the meson momentum yields further understanding of how (8.38) arises.

We will close in Section 8.4.3 with a discussion of the potential connections between the properties of quarkonium mesons in motion in a holographic plasma and those of quarkonium mesons in motion in the QCD plasma.

8.4.1 Spectrum and dispersion relations

In order to determine the meson spectrum on Minkowski embeddings, we proceed as in Section 8.2.2. For simplicity we will focus on fluctuations of the position of the branes $U(u)$ with no angular momentum on the S^3 , i.e. we write

$$\delta U = \mathcal{U}(u) e^{-i\omega t} e^{i\mathbf{q}\cdot\mathbf{x}}. \quad (8.39)$$

The main difference between this equation and its zero-temperature counterpart (8.10) is that in the latter case Lorentz invariance implies the usual relation $\omega^2 - q^2 = M^2$ between the energy ω , the spatial three-momentum q , and the mass M of the meson. At nonzero temperature, boost invariance is broken because the plasma defines a preferred frame in which it is at rest and the mesons develop a non-trivial dispersion relation $\omega(q)$. In the string description this is determined by requiring normalizability and regularity of $\mathcal{U}(u)$: For each value of q , these two requirements are mutually compatible only for a discrete set of values $\omega_n(q)$, where different values of n label different excitation levels of the meson. We define the ‘rest mass’ of a meson as the energy $\omega(0)$ at vanishing three-momentum, $q = 0$, in the rest-frame of the plasma.

Fig. 8.13, taken from Ref. [565], shows the rest mass of the mesons as a function of temperature and quark mass. Note that in the zero-temperature limit, the spectrum coincides with the zero-temperature spectrum (8.14). In particular, the lightest meson has a mass squared matching Eqn. (8.29): $M_{\text{mes}}^2 = 4\pi^2 \bar{M}^2 \simeq 39.5 \bar{M}^2$.

The meson masses decrease as the temperature increases. Heuristically, this can be understood in geometrical terms from Fig. 8.6, which shows that the thermal quark mass M_{th} decreases as the temperature increases and the tip of the D7-branes gets closer to the black hole horizon. The thermal shift in the meson masses becomes more significant at the phase transition, and slightly beyond this point some modes actually become tachyonic. This happens precisely in the same region in which Minkowski embeddings become thermodynamically unstable because $c_V < 0$. In other words, Minkowski embeddings develop thermodynamic and dynamic instabilities at exactly the same T/\bar{M} , just beyond that at which the first order dissociation transition occurs.

We now turn to quarkonium mesons moving through the plasma, that is to modes with $q \neq 0$. The dispersion relation for scalar mesons was first computed in Ref. [565] and then revisited in Ref. [568]. The dispersion relation for (transverse) vector mesons appeared in Ref. [617]. An exhaustive discussion of the dispersion relations for all these cases can be found in Ref. [618]. The result for the lowest-lying ($n = 0$) vector, scalar and pseudoscalar quarkonia is shown in Fig. 8.14, taken from Ref. [618]. Fig. 8.15, taken from Ref. [568], shows the group velocity $v_g = d\omega/dq$ for the $n = 0$ scalar mesons at three different temperatures.

An important feature of these plots is their behavior at large momentum. In this regime we find that ω grows linearly with q . Naively, one might expect that the constant of propor-




Figure 8.13: Meson mass spectrum $M^2 = \omega^2|_{q=0}$ for Minkowski embeddings in the D3/D7 system. Continuous curves correspond to radially excited mesons with radial quantum number $n = 0, 1, 2$ from bottom to top, respectively. Dashed lines correspond to mesons with angular momentum on the S^3 . The dashed vertical line indicates the temperature of the phase transition. Note that modes become tachyonic slightly beyond this temperature.




Figure 8.14: Left: Dispersion relation for the transverse (black, continuous curve) and longitudinal (red, dashed curve) $n = 0$ modes of a heavy vector meson with $v_{\text{lim}} = 0.35$ in the $\mathcal{N} = 4$ SYM plasma. The dual D7-brane has $m = 1.3$, corresponding to a temperature just below T_{diss} . Right: Analogous curves for a scalar (black, continuous curve) and pseudoscalar (red, dashed curve) meson. In both plots the blue, continuous straight lines correspond to $\omega = vq$ for some v such that $v_{\text{lim}} < v \leq 1$. The black, dotted, vertical lines mark the crossing points between the meson dispersion relations and the blue lines.

tionality should be one. However, one finds instead that

$$\omega = v_{\text{lim}} q, \quad (8.40)$$

where $v_{\text{lim}} < 1$ and where v_{lim} depends on $m = \bar{M}/T$ but at a given temperature is the same for all quarkonium modes. In the particular case of $m = 1.3$, illustrated in Fig. 8.14, one has $v_{\text{lim}} \simeq 0.35$. In other words, there is a subluminal limiting velocity for quarkonium mesons moving through the plasma. And, as illustrated in Fig. 8.15, one finds that the limiting velocity decreases with increasing temperature. Fig. 8.15 also illustrates another generic feature of the dispersion relations, namely that the maximal group velocity is attained at some $q_m < \infty$ and as q is increased further the group velocity approaches v_{lim} from above. Since v_g at q_m is not much greater than v_{lim} , we will not always distinguish between these two velocities. We will come back to the physical interpretation of q_m at the end of this subsection.

The existence of a subluminal limiting velocity, which was discovered in [565] and subsequently elaborated upon in [568], is easily understood from the perspective of the dual gravity description [565, 568]. Recall that mesonic states have wave functions supported on the D7-branes. Since highly energetic mesons are strongly attracted by the gravitational pull of the black hole, their wave-function is very concentrated at the bottom of the branes (see Fig. 8.6). Consequently, their velocity is limited by the local speed of light at that point. As seen by an observer at the boundary, this limiting velocity is

$$v_{\text{lim}} = \sqrt{-g_{tt}/g_{xx}} \Big|_{\text{tip}}, \quad (8.41)$$

where g is the induced metric on the D7-branes. Because of the black hole redshift, v_{lim} is lower than the speed of light at infinity (i.e. at the boundary), which is normalized to unity.




Figure 8.15: Group velocities v_g for $n = 0$ scalar meson modes with $T/T_{\text{diss}} \approx 0.65, 0.92$ and 1 from top to bottom. We see that $v_{\text{lim}} < 1$ decreases with increasing temperature. (v_{lim} would approach zero if we included the unstable Minkowski embeddings with $T > T_{\text{diss}}$.) The group velocity approaches its large- q value v_{lim} from above, i.e. v_g reaches a maximum before settling into the limiting velocity v_{lim} . The maximum exists also for the top curve even though it is less clearly visible. We will refer to the momentum at which v_g reaches the maximum as q_m . Clearly q_m decreases with temperature.

Note that, as the temperature increases, the bottom of the brane gets closer to the horizon and the redshift becomes larger, thus further reducing v_{lim} ; this explains the temperature dependence in Fig. 8.15. In the gauge theory, the above translates into the statement that v_{lim} is lower than the speed of light in the vacuum. The reason for this interpretation is that the absence of a medium in the gauge theory corresponds to the absence of a black hole on the gravity side, in which case $v_{\text{lim}} = 1$ everywhere. Eqn. (8.41) yields $v_{\text{lim}} \simeq 0.35$ at $m = 1.3$, in agreement with the numerical results displayed in Fig. 8.14.

It is also instructive to plot v_{lim} as a function of T/T_{diss} , as done in Fig. 8.16, taken from [568]. Although this curve was derived as a limiting meson velocity at a given temperature, it can also be read (by asking where it cuts horizontal lines rather than vertical ones) as giving $T_{\text{diss}}(v)$, the temperature below which mesons with a given velocity v are found and above which no mesons with that velocity exist. In order to compare this result for T_{diss} at all velocities to (8.38), one can parametrize the curve in Fig. 8.16 as

$$T_{\text{diss}}(v) = f(v)(1 - v^2)^{1/4} T_{\text{diss}}(0). \quad (8.42)$$

In the left panel of Fig. 8.16, the dashed line is obtained by setting $f(v) = 1$, which is of course just (8.38). In the right panel, $f(v)$ is shown to be close to 1 for all velocities, varying between 1.021 at its maximum and 0.924 at $v = 1$. Recall that the scaling (8.38) was first obtained via the analysis of the potential between a moving test quark and antiquark, as described in Section 7.7. The weakness of the dependence of $f(v)$ on v is a measure of the robustness with which that simple scaling describes the velocity dependence of the




Figure 8.16: Left panel: The solid curve is the limiting velocity v_{lim} as a function of T/T_{diss} , where T_{diss} is the temperature of the dissociation transition at zero velocity. The dissociation transition occurs at the dot, where $v_{\text{lim}} = 0.27$. The dashed curve is the approximation obtained by setting $f(v) = 1$ in Eqn. (8.42). Right panel: $f(v)$, the ratio of the solid and dashed curves in the left panel at a given v . We see that $f(v)$ is within a few percent of 1 at all velocities.

dissociation temperature for quarkonium mesons in a fully dynamical calculation. In other words, to a good approximation $v_{\text{lim}}(T)$ can be determined by setting $v = v_{\text{lim}}$ on the right-hand side of (8.38), yielding

$$v_{\text{lim}}(T) \simeq \sqrt{1 - \left(\frac{T}{T_{\text{diss}}(v=0)} \right)^4}. \quad (8.43)$$

Thus we reach a rather satisfactory picture that the subluminal limiting velocity (8.40) is in fact a manifestation in the physics of dynamical mesons of the velocity-enhanced screening of Section 7.7. However, in the case of the low-spin mesons whose dynamics we are considering in this Section, there is an important addition to our earlier picture: Although the quarkonium mesons have a limiting velocity, they can nevertheless manage to remain bound at arbitrarily large momenta thanks to their modified dispersion relations. The latter allow the group velocity to remain less than v_{lim} , and consequently $T_{\text{diss}}(v)$ as given in (8.38) to remains higher than T , all the way out to arbitrarily large momenta. In other words, there exist meson bound states of arbitrarily large spatial momentum, but no matter how large the momentum the group velocity never exceeds v_{lim} . In this sense, low-spin mesons realize the first of two simple possibilities by which mesons may avoid exceeding v_{lim} . A second possibility, more closely related to the analysis of Section 7.7, is that meson states with momentum larger than a certain value simply cease to exist. This possibility is realized in the case of high-spin mesons. Provided $J \gg 1$, these mesons can be reliably described as long, semiclassical strings whose ends are attached to the bottom of the D7-branes. The fact that the endpoints do not fall on top of one another is of course due to the fact that they are rotating around one another in such a way that the total angular momentum of the string is J . These type of mesons were first studied [585] at zero temperature, and subsequently considered at nonzero temperature in Ref. [556]. These authors also studied the possibility that, at the same time that the endpoints of the string rotate around one another in a given plane, they also move with a certain velocity in the direction orthogonal to that plane. The result of the analysis was that, for a fixed spin J , string solutions exist only up to a maximum velocity $v_{\text{lim}} < 1$.

As we saw in Fig. 8.15, the group velocity of quarkonium mesons reaches a *maximum* at some value of the momentum $q = q_m$ before approaching the limiting value v_{lim} . There is a simple intuitive explanation for the location of q_m : it can be checked numerically that q_m is always close to the ‘limiting momentum’ q_{lim} that would follow from (8.38) if one assumes the *standard* dispersion relation for the meson. Thus, q_m can be thought of as a characteristic momentum scale where the velocity-enhanced screening effect starts to be important. For the curves in Fig. 8.15, to the left of the maximum one finds approximately standard dispersion relations with a thermally corrected meson mass. To the right of the maximum, the dispersion relations approach the limiting behavior (8.40), with v_g approaching v_{lim} , as a consequence of the enhanced screening.

8.4.2 Decay widths

We saw above that at $T < T_{\text{diss}}$ (Minkowski embeddings) there is a discrete and gapped spectrum of absolutely stable quarkonium mesons, i.e. the mesons have zero width. The reason is that in this phase the D-branes do not touch the black hole horizon. Since the mesons’ wave-functions are supported on the branes, this means that the mesons cannot fall into the black hole. In the gauge theory this translates into the statement that the mesons cannot disappear into the plasma, which implies that the meson widths are strictly zero in the limit $N_c, \lambda \rightarrow \infty$. This conclusion only depends on the topology of the Minkowski embedding. In particular, it applies even when higher-order perturbative corrections in α' are included, which implies that the widths of mesons should remain zero to all orders in the perturbative $1/\sqrt{\lambda}$ expansion. In contrast, in the black hole phase the D-branes fall into the black hole and a meson has a nonzero probability of disappearing through the horizon, that is, into the plasma. As a consequence, we expect the mesons to develop thermal widths in the black hole phase, even in the limit $N_c, \lambda \rightarrow \infty$. In fact, as we will see in Section 8.5, the widths are generically comparable to the energies of the mesons, and hence the mesons can no longer be interpreted as quasiparticles.

We thus encounter a somewhat unusual situation: the quarkonium mesons are absolutely stable for $T < T_{\text{diss}}$, but completely disappear for $T > T_{\text{diss}}$. The former is counterintuitive because, on general grounds, we expect that any bound states should always have a nonzero width when immersed in a medium with $T > 0$. In the case of these mesons, we expect that they can decay and acquire a width through the following channels:

1. Decay to gauge singlets such as glueballs, lighter mesons, etc;
2. Breakup by high-energy gluons (right diagram in Fig. 8.17);
3. Breakup by thermal medium quarks (left diagram in Fig. 8.17).

Process (1) is suppressed by $1/N_c^2$ (glueballs) or $1/N_c$ (mesons), while (2) and (3) are unsuppressed in the large- N_c limit. Since (1) is also present in the vacuum, we will focus on (2) and (3), which are medium effects. They are shown schematically in Fig. 8.17.

The width due to (2) is proportional to a Boltzmann factor $e^{-\beta E_B}$ for creating a gluon that is energetic enough to break up the bound state, while that due to (3) is proportional to a Boltzmann factor $e^{-\beta M_{\text{th}}}$ for creating a thermal quark, where M_{th} is the thermal mass




Figure 8.17: Sketches taken from Ref. [571] showing the relevant thermal processes contributing to the meson width. q (\bar{q}) denotes a quark (antiquark). The left diagram corresponds to the breakup of a meson by a quark from the thermal medium, while the right diagram corresponds to breakup of a meson by an energetic gluon. For large λ the first process is dominant, coming from the single instanton sector.

of the quark — see Fig. 8.6. Given Eqn. (8.35) and the fact that in the Minkowski phase $T < T_{\text{diss}}$, both Boltzmann factors are suppressed by $e^{-\sqrt{\lambda}} \sim e^{-R^2/\alpha'}$, so we recover the result that these mesons are stable in the infinite- λ limit. In particular, there is no width at any perturbative order in the $1/\sqrt{\lambda}$ -expansion, consistent with the conclusion from the string theory side. Furthermore, since the binding energy is $E_B \approx 2M_{\text{th}}$, in the regime where λ is large (but not infinite), the width from process (3) will dominate over that from process (2).

We now review the result from the string theory calculation of the meson widths in Ref. [571]. As discussed above, the width is non-perturbative in $1/\sqrt{\lambda} \sim \alpha'/R^2$, and thus should correspond to some instanton effect on the string worldsheet. The basic idea is very simple: even though in a Minkowski embedding the brane is separated from the black hole horizon and classically a meson living on the brane cannot fall into the black hole, quantum mechanically (from the viewpoint of the string worldsheet) it has a nonzero probability of tunneling into the black hole and the meson therefore develops a width. At leading order, the instanton describing this tunneling process is given by a (Euclidean) string worldsheet stretching between the tip of the D7-brane to the black hole horizon (see Fig. 8.6) and winding around the Euclidean time direction. Heuristically, such a worldsheet creates a small tunnel between the brane and the black hole through which mesons can fall into the black hole. The instanton action is βM_{th} , as can be read off immediately from the geometric picture just described, and its exponential gives rise to the Boltzmann factor expected from process (3). From the gauge theory perspective, such an instanton can be interpreted as creating a thermal quark from the medium, and a meson can disappear into the medium via interaction with it as shown in the left diagram of Fig. 8.17.

The explicit expression for the meson width due to such instantons is somewhat complicated, so we refer the reader to the original literature [571]. Although the width appears to be highly model dependent and is exponentially small in the regime of a large but finite λ under consideration, remarkably its momentum dependence has some universal features at

large momentum in [571]. Specifically, one finds that

$$\frac{\Gamma(q)}{\Gamma(0)} = \frac{|\psi(\text{tip}; \vec{q})|^2}{|\psi(\text{tip}; \vec{q} = 0)|^2}, \quad (8.44)$$

where $\Gamma(q)$ denotes the width of a meson with spatial momentum q and $\psi(\text{tip}; q)$ its wave-function evaluated at the tip of the D7-branes (i.e. where it is closest to the black hole). This result is intuitively obvious because a meson tunnels into the black hole from the tip of the branes. In particular, as discussed in detail in Ref. [568], at large momentum q the wave-function becomes localized around the tip of the brane and can be approximated by that of a spherical harmonic oscillator with a potential proportional to $q^2 z^2$, where z is the proper distance from the tip of the branes.⁶ It then immediately follows that for large q the width (8.44) scales as q^2 . Furthermore for temperatures $T \ll M_{\text{mes}}$ and $q \gg M_{\text{mes}}^3/T^2$, one finds the closed-form expression

$$\frac{\Gamma_n(q)}{\Gamma_n(0)} \approx \frac{2(4\pi)^4}{(n+2)(n+3/2)} \frac{T^4 q^2}{M_{\text{mes}}^6}, \quad (8.45)$$

where n labels different mesonic excitations (see (8.13)).

It is also instructive to plot the full q -dependence of (8.44) obtained numerically, as done in Fig. 8.18 (taken from [571]) for $n = 0$ mesons at various temperatures. Fig. 8.18 has the interesting feature that the width is roughly constant for small q , but turns up quadratically around $q/M_{\text{mes}} \approx 0.52(T_{\text{diss}}/T)^2$. This is roughly the momentum $q_m \sim q_{\text{lim}}$ at which the group velocity of a meson achieves its maximum in Fig. 8.15 which, as discussed in Section 8.4.1, can be considered as the characteristic momentum scale where velocity-enhanced screening becomes significant. This dramatic increase of meson widths beyond q_m can also be understood intuitively: when velocity-enhanced screening becomes significant, interaction between the quark and antiquark in a meson becomes further weakened, which makes it easier for a thermal medium quark or gluon to break it apart.

We now briefly comment on the gravity description of process (2) mentioned earlier, i.e. the right diagram in Fig. 8.17. For such a process to happen the gluon should have an energy above the binding energy of the meson. The density of such gluons is thus suppressed by $e^{-2\beta M_{\text{th}}}$ and should be described by an instanton and anti-instanton. We expect that contributions from such processes are also controlled by the value of the meson wave-function at the tip of the branes, and thus likely have similar growth with momentum.

Finally, we note that, as T increases, M_{th} decreases and thus the meson width increases quickly with temperature, but remains exponentially suppressed until T_{diss} is reached, after which we are in the black hole phase. As will be discussed in Section 8.5, in this phase quarkonium quasiparticles no longer exist.

8.4.3 Connection with the quark-gluon plasma

Let us now recapitulate the main qualitative features regarding heavy mesons in a strongly coupled plasma:

⁶Note that there are four transverse directions along the D7-brane as we move away from the tip (not including the other $(3+1)$ dimensions parallel to the boundary). Thus this is a four-dimensional harmonic oscillator.




Figure 8.18: The behavior of the width as a function of q for $T/T_{\text{diss}} = 0.99, 0.71, 0.3, 0.13$ from left to right. The dashed lines are analytic results for large momenta.

1. They survive deconfinement.
2. Their dispersion relations have a subluminal limiting velocity at large momentum. The limiting velocity decreases with increasing temperature and as a result the motion of a meson with large momentum dramatically slows down near T_{diss} .
3. At large momenta, meson widths increase dramatically with momentum.
4. The limiting velocity is reached and the increase in widths applies when $q \gg q_{\text{lim}}$, where q_{lim} is the ‘limiting’ momentum following from (8.38) if one assumes the *standard* dispersion relation.

Properties (1)–(3) are universal in the sense that they apply to the deconfined phase of any gauge theory with a string dual in the large- N_c , strong-coupling limit. The reason for this is that they are simple consequences of general geometric features following from two universal aspects of the gauge/string duality: (i) the fact that the deconfined phase of the gauge theory is described on the gravity side by a black hole geometry [312], and (ii) the fact that a finite number N_f of quark flavours is described by N_f D-brane probes [384, 385]. Property (4) was established by explicit numerical calculations in specific models. However, given that q_{lim} can be motivated in a model-independent way from (8.38), it is likely to also be universal even though this was not manifest in our discussion above.

We have seen that properties (2) and (3) can be considered direct consequences of velocity-enhanced screening, which as discussed in Section 7.7 can have important implications for quarkonium suppression in heavy ion collisions.

It is interesting that properties (1) and to some degree (2) can be independently motivated in QCD whether or not a string dual of QCD exists. The original argument [184] for (1) is

simply that the heavier the quarkonium meson, the smaller its size. And, it is reasonable to expect a meson to remain bound until the screening length in the plasma becomes comparable to the meson size, and for sufficiently heavy quarkonia this happens at $T_{\text{diss}} > T_c$. As we have discussed in Sections 2.4 and 3.3, this conclusion is supported by calculations of both the static quark-antiquark potential and of Minkowski-space spectral functions in lattice-regularized QCD. The ballpark estimate for the dissociation temperature of heavy mesons suggested by the above studies roughly agrees with that from the D3/D7 system. For example, for the J/ψ meson the former estimate is $T_c \lesssim T_{\text{diss}} \lesssim 2T_c$. Allowing for a certain range in the precise value of $150 \text{ MeV} \lesssim T_c \lesssim 190 \text{ MeV}$, this translates into $300 \text{ MeV} \lesssim T_{\text{diss}} \lesssim 380 \text{ MeV}$. In the D3/D7 model, we see from Fig. 8.11 that meson states melt at $T_{\text{diss}} \simeq 0.766\bar{M}$. The scale \bar{M} is related to the mass M_{mes} of the lightest meson in the theory at zero temperature through Eqn. (8.29). Therefore we have $T_{\text{diss}}(M_{\text{mes}}) \simeq 0.122M_{\text{mes}}$. For the J/ψ , taking $M_{\text{mes}} \simeq 3 \text{ GeV}$ gives $T_{\text{diss}}(J/\psi) \simeq 366 \text{ MeV}$. Although it is gratifying that this comparison leads to qualitative agreement, it must be taken with some caution because meson bound states in the D3/D7 system are deeply bound, i.e. $M_{\text{mes}} \ll 2M_q$, whereas the binding energy of charmonium states in QCD is a small fraction of the charm mass, i.e. $M_{c\bar{c}} \simeq 2M_c$. An additional difference comes from the fact that in QCD the dissociation of charmonium states is expected to happen sequentially, with excited states (that are larger) dissociating first, whereas in the D3/D7 system all meson states are comparable in size and dissociate at the same temperature. Presumably, in the D3/D7 system this is an artifact of the large- N_c , strong-coupling approximation under consideration, and thus corrections away from this limit should make holographic mesons dissociate sequentially too.

There is a simple (but incomplete) argument for property (2) that applies to QCD just as well as to $\mathcal{N} = 4$ SYM theory [556–558, 568]: a meson moving through the plasma with velocity v experiences a higher energy density, boosted by a factor of γ^2 . Since energy density is proportional to T^4 , this can be thought of as if the meson sees an effective temperature that is boosted by a factor of $\sqrt{\gamma}$, meaning $T_{\text{eff}}(v) = (1 - v^2)^{-1/4}T$. A velocity-dependent dissociation temperature scaling like (8.38) follows immediately and from this a subluminal limiting velocity (8.43) can be inferred. Although this argument is seductive, it can be seen in several ways that it is incomplete. For example, we would have reached a different $T_{\text{eff}}(v)$ had we started by observing that the entropy density s is boosted by a factor of γ and is proportional to T^3 . And, furthermore, there really is no single effective temperature seen by the moving quarkonium meson. The earliest analysis of quarkonia moving through a weakly coupled QCD plasma with some velocity v showed that the meson sees a blue-shifted temperature in some directions and a red-shifted temperature in others [619]. Although the simple argument does not work by itself, it does mean that all we need from the calculations done via gauge/string duality is the result that $T_{\text{diss}}(v)$ behaves as if it is controlled by the boosted energy density — i.e. we need the full calculation only for the purpose of justifying the use of the particular simple argument that works. This suggests that property (2), and in particular the scaling in Eqs. (8.38) and (8.43), are general enough that they may apply to the quark-gluon plasma of QCD whether or not it has a gravity dual.

As explained towards the end of Section 8.4.1, there are at least two simple ways in which a limiting velocity for quarkonia may be implemented. It may happen that meson states with momentum above a certain q_{lim} simply do not exist, in which case one expects that $v_{\text{lim}} = v(q_{\text{lim}})$. The second possibility is that the dispersion relation of mesons may

become dramatically modified beyond a certain q_{lim} in such a way that, although meson states of arbitrarily high momentum exist, their group velocity never exceeds a certain v_{lim} . It is remarkable that both possibilities are realized in gauge theories with a string dual, the former by high-spin mesons and the latter by low-spin mesons. However, note that even in the case of low-spin mesons, q_{lim} remains the important momentum scale beyond which we expect more significant quarkonium suppression for two reasons. First, meson widths increase significantly for $q > q_{\text{lim}}$. Therefore, although it is an overstatement to say that these mesons also cease to exist above q_{lim} , their existence becomes more and more transient at higher and higher q . Second, due to the modified dispersion relation mesons with $q > q_{\text{lim}}$ slow down and they spend a longer time in the medium, giving the absorptive imaginary part more time to cause dissociation. It will be very interesting to see whether future measurements at RHIC or the LHC will show the suppression of J/Ψ or Υ production increasing markedly above some threshold transverse momentum p_T , as we described in Section 7.7.

In practice, our ability to rigorously verify the properties (1)–(4) in QCD is limited due to the lack of tools that are well-suited for this purpose. It is therefore reassuring that they hold for all strongly coupled, large- N_c plasmas with a gravity dual, for which the gravity description provides just such a tool.

8.5 Black hole embeddings

We now consider the phase $T > T_{\text{diss}}$, which is described by a D7-brane with a black hole embedding. We will give a qualitative argument that in this regime the system *generically* contains no quarkonium quasiparticles. We have emphasized the word ‘generically’ because exceptions arise when certain large ratios of physical scales are introduced ‘by hand’, as we will see later. We will illustrate the absence of quasiparticles in detail by computing a spectral function of two electromagnetic currents in the next subsection.

8.5.1 Absence of quasiparticles

In the gravity description of physics at $T > T_{\text{diss}}$, the meson widths may be seen by studying so-called quasi-normal modes. These are analogues of the fluctuations we studied in the case of Minkowski embeddings in that a normalizable fall-off is imposed at the boundary. However, the regularity condition at the tip of the branes is replaced by the so-called in-falling boundary condition at the horizon. Physically, this is the requirement that energy can flow into the horizon but cannot come out of it (classically). Mathematically, it is easy to see that this boundary condition forces the frequency of the mode to acquire a negative imaginary part, and thus corresponds to a nonzero meson width. The meson in question may then be considered a quasiparticle if and only if this width is much smaller than the real part of the frequency. In the case at hand, the meson widths increase as the area of the induced horizon on the branes increases, and go to zero only when the horizon shrinks to zero size. This is of course to be expected, since it is the presence of the induced horizon that causes the widths to be nonzero in the first place. We are thus led to the suggestion that meson-like quasiparticles will be present in the black hole phase only when the size of the induced horizon on the branes can be made parametrically small. This expectation can be directly

verified by explicit calculation of the quasi-normal modes on the branes [610, 620–622], and we will confirm it indirectly below by examining the spectral function of two electromagnetic currents. For the moment, let us just note that this condition is not met in the system under consideration because, as soon as the phase transition at $T = T_{\text{diss}}$ takes place, the area of the induced horizon on the brane is an order-one fraction of the area of the background black hole. This can be easily seen from Fig. 8.11 by comparing the entropy density (which is a measure of the horizon area) at the phase transition to the entropy density at asymptotically high temperatures:

$$\frac{s_{\text{transition}}}{s_{\text{high } T}} \approx \frac{0.3}{2} \approx 15\%. \quad (8.46)$$

This indicates that there is no parametric reason to expect quasiparticles with narrow widths above the transition. We shall confirm by explicit calculation in the next subsection that there are no quasiparticle excitations in the black hole phase.

8.5.2 Meson spectrum from a spectral function

Here we will illustrate some of the general expectations discussed above by examining the in-medium spectral function of two electromagnetic currents in the black hole phase. We choose this particular correlator because it is related to thermal photon production, which we will discuss in the next Section. We will see that no narrow peaks exist for stable black hole embeddings, indicating the absence of long-lived quasiparticles. These peaks will appear, however, as we artificially push the system into the unstable region close to the critical embedding (see Fig. 8.7), thus confirming our expectation that quasiparticles should appear as the area of the induced horizon on the branes shrinks to zero size.

$\mathcal{N} = 4$ SYM coupled to N_f flavours of equal-mass quarks is an $SU(N_c)$ gauge theory with a global $U(N_f)$ symmetry. In order to couple this theory to electromagnetism we should gauge a $U(1)_{\text{EM}}$ subgroup of $U(N_f)$ by adding a dynamical photon A_μ to the theory; for simplicity we will assume that all quarks have equal electric charge, in which case $U(1)_{\text{EM}}$ is the diagonal subgroup of $U(N_f)$. In this extended theory we could then compute correlation functions of the conserved current J_μ^{EM} that couples to the $U(1)_{\text{EM}}$ gauge field. The string dual of this $SU(N_c) \times U(1)_{\text{EM}}$ gauge theory is unknown, so we cannot perform this calculation holographically. However, as noted in [623], we can perform it to leading order in the electromagnetic coupling constant e , because at this order correlation functions of electromagnetic currents in the gauged and in the ungauged theories are identical. This is very simple to understand diagrammatically, as illustrated for the two-point function in Fig. 8.19. In the ungauged theory only $SU(N_c)$ fields ‘run’ in the loops, represented by the shaded blob. The gauged theory contains additional diagrams in which the photon also runs in the loops, but these necessarily involve more photon vertices and therefore contribute at higher orders in e . Thus one can use the holographic description to compute the ‘ $SU(N_c)$ blob’ and obtain the result for the correlator to leading order in e .

Using this observation, the authors of Ref. [623] first did a holographic computation of the spectral density of two R-symmetry currents in $\mathcal{N} = 4$ SYM theory, to which finite-coupling corrections were computed in [624, 625]. The result for the R-charge spectral density is identical, up to an overall constant, with the spectral density of two electromagnetic currents




Figure 8.19: Diagrams contributing to the two-point function of electromagnetic currents. The external line corresponds to a photon of momentum q . As explained in the text, to leading order in the electromagnetic coupling constant only $SU(N_c)$ fields ‘run’ in the loops represented by the shaded blob.

in $\mathcal{N} = 4$ SYM theory coupled to massless quarks. This, and the extension to nonzero quark mass, were obtained in Ref. [612], which we now follow.

The relevant spectral function is defined as

$$\chi_{\mu\nu}(k) = 2 \operatorname{Im} G_{\mu\nu}^{\text{R}}(k), \quad (8.47)$$

where $k^\mu = (\omega, q)$ is the photon null momentum (i.e. $\omega^2 = q^2$) and

$$G_{\mu\nu}^{\text{R}}(k) = i \int d^{d+1}x e^{-ik_\mu x^\mu} \Theta(t) \langle [J_\mu^{\text{EM}}(x), J_\nu^{\text{EM}}(0)] \rangle \quad (8.48)$$

is the retarded correlator of two electromagnetic currents. The key point in this calculation is to identify the field in the string description that is dual to the operator of interest here, namely the conserved current J_μ^{EM} . We know from the discussion in Section 5 that conserved currents are dual to gauge fields on the string side. Moreover, since J_μ^{EM} is constructed out of fields in the fundamental representation, we expect its dual field to live on the D7-branes. The natural (and correct) candidate turns out to be the $U(1)$ gauge field associated with the diagonal subgroup of the $U(N_f)$ gauge group living on the worldvolume of the N_f D7-branes. Once this is established, one must just follow the general prescription explained in Section 5. The technical details of the calculation can be found in Ref. [612], so here we will only describe the results and their interpretation. In addition, we will concentrate on the trace of the spectral function, $\chi^\mu{}_\mu(k) \equiv \eta^{\mu\nu} \chi_{\mu\nu}(k)$, since this is the quantity that determines the thermal photon production by the plasma (see next section).

The trace of the spectral function for stable black hole embeddings is shown in Fig. 8.20 for several values of the quark mass m . Note that this is a function of only one variable, since for an on-shell photon $\omega = q$. The normalisation constant that sets the scale on the vertical axis is $\tilde{\mathcal{N}}_{\text{D7}} = N_f N_c T^2 / 4$. The $N_f N_c$ scaling of the spectral function reflects the number of electrically charged degrees of freedom in the plasma; in the case of two R-symmetry currents, $N_f N_c$ would be replaced by N_c^2 [623]. All curves decay as $\omega^{-1/3}$ for large frequencies. Note that $\chi \sim \omega$ as $\omega \rightarrow 0$. This is consistent with the fact that the value at the origin of each of the curves yields the electric conductivity of the plasma at the corresponding quark mass,




Figure 8.20: Trace of the spectral function as a function of the dimensionless frequency $\bar{\omega} = \omega/2\pi T$ for (from top to bottom on the left-hand side) $m = \{0, 0.6, 0.85, 0.93, 1.15, 1.25, 1.32\}$. The last value corresponds to that at which the phase transition from a black hole to a Minkowski embedding takes place. Recall that $\tilde{N}_{D7} = N_f N_c T^2/4$.

namely

$$\sigma = \frac{e^2}{4} \lim_{\omega \rightarrow 0} \frac{1}{\omega} \eta^{\mu\nu} \chi_{\mu\nu}(\omega = q). \quad (8.49)$$

This formula is equivalent to the perhaps-more-familiar expression in terms of the zero-frequency limit of the spectral function at vanishing spatial momentum (see Appendix A):

$$\sigma = \frac{e^2}{6} \lim_{\omega \rightarrow 0} \frac{1}{\omega} \delta^{ij} \chi_{ij}(\omega, q = 0) = \frac{e^2}{6} \lim_{\omega \rightarrow 0} \frac{1}{\omega} \eta^{\mu\nu} \chi_{\mu\nu}(\omega, q = 0), \quad (8.50)$$

where in the last equality we used the fact that $\chi_{00}(\omega \neq 0, q = 0) = 0$, as implied by the Ward identity $k^\mu \chi_{\mu\nu}(k) = 0$. To see that the two expressions (8.49) and (8.50) are equivalent, suppose that q points along the 1-direction. Then the Ward identity, together with the symmetry of the spectral function under the exchange of its spacetime indices, imply that $\omega^2 \chi_{00} = q^2 \chi_{11}$. For null momentum this yields $-\chi_{00} + \chi_{11} = 0$, so we see that Eqn. (8.49) reduces to

$$\sigma = \frac{e^2}{4} \lim_{\omega \rightarrow 0} \frac{1}{\omega} \left[\chi_{22}(\omega = q) + \chi_{33}(\omega = q) \right]. \quad (8.51)$$

The diffusive nature of the hydrodynamic pole of the correlator implies that at low frequency and momentum the spatial part of the spectral function behaves as

$$\chi_{ij}(\omega, q) \sim \frac{\omega^3}{\omega^2 + D^2 q^4}, \quad (8.52)$$

where D is the diffusion constant for electric charge. This means that we can replace $q = \omega$ by $q = 0$ in Eqn. (8.51), thus arriving at expression (8.50).




Figure 8.21: Trace of the spectral function as a function of the dimensionless frequency $\bar{\omega} = \omega/2\pi T$ for non-stable black hole embeddings. Curves with higher, narrower peaks correspond to embeddings that are closer to the critical embedding.

For the purpose of our discussion, the most remarkable feature of the spectral functions displayed in Fig. 8.20 is the absence of any kind of high, narrow peaks that may be associated with a quasiparticle excitation in the plasma. This feature is shared by thermal spectral functions of other operators on stable black hole embeddings. We thus confirm our expectation that no quasiparticles exist in this phase. In order to make contact with the physics of the Minkowski phase, in which we do expect the presence of quarkonium quasiparticles, the authors of Ref. [612] computed the spectral function for black hole embeddings beyond the phase transition, i.e. in the region below T_{diss} in which these embeddings are metastable or unstable. The results for the spectral function are shown in Fig. 8.21. The most important feature of these plots is the appearance of well defined peaks in the spectral function, which become higher and narrower, seemingly approaching delta-functions, as the embedding approaches the critical embedding (see Fig. 8.7). Thus the form of the spectral function appears to approach the form we expect for Minkowski embeddings,⁷ namely an infinite sum of delta-functions supported at a discrete set of energies $\omega^2 = q^2$. (However, a precise map between the peaks in Fig. 8.21 and the meson spectrum in a Minkowski embedding is not easy to establish [626].) Each of these delta-functions is associated with a meson mode on the D7-branes with *null* four-momentum. The fact that the momentum is null may seem surprising in view of the fact that, as explained above, the meson spectrum in the Minkowski phase possesses a mass gap, but in fact it follows from the dispersion relation for these mesons displayed in Fig. 8.14. To see this, consider the dispersion relation $\omega(q)$ for a given meson in the Minkowski phase. The fact that there is a mass gap means that $\omega > 0$ at $q = 0$. On the other hand, in the limit of infinite spatial momentum, $q \rightarrow \infty$, the dispersion relation takes the form $\omega \simeq v_{\text{lim}} q$ with $v_{\text{lim}} < 1$. Continuity then implies that there must exist a value of q

⁷An analogous result was found in [611] for time-like momenta.

such that $\omega(q) = q$. This is illustrated in Fig. 8.14 by the fact that the dispersion relations intersect the blue lines. Since in the Minkowski phase the mesons are absolutely stable in the large- N_c , strong-coupling limit under consideration, we see that each of them gives rise to a delta-function-like (i.e. zero-width) peak in the spectral function of electromagnetic currents at null momentum. Below we will see some potential implications of this result for heavy-ion collisions.

8.6 Two universal predictions

We have just seen that the fact that heavy mesons remain bound in the plasma, and the fact that their limiting velocity is subluminal, imply that the dispersion relation of a heavy meson must cross the light-cone, defined by $\omega = q$, at some energy $\omega = \omega_{\text{peak}}$ indicated by the vertical line in Fig. 8.14. In this section we will see that this simple observation leads to two universal consequences. Implications for deep inelastic scattering have been studied in [627] but will not be reviewed here.

8.6.1 A meson peak in the thermal photon spectrum

At the crossing point between the meson dispersion relation and the light-cone, the meson four-momentum is null, that is $\omega_{\text{meson}}^2 = q_{\text{meson}}^2$. If the meson is flavourless and has spin one, then at this point its quantum numbers are the same as those of a photon. Such a meson can then decay into an on-shell photon, as depicted in Fig. 8.22. Note that, in the vacuum, only the decay into a virtual photon would be allowed by kinematics. In the medium, the decay can take place because of the modified dispersion relation of the meson. Also, note that the decay will take place unless the photon-meson coupling vanishes for some reason (e.g. a symmetry). No such reason is known in QCD.

The decay process of Fig. 8.22 contributes a resonance peak, at a position $\omega = \omega_{\text{peak}}$, to the in-medium spectral function of two electromagnetic currents (8.48) evaluated at null-momentum $\omega = q$. This in turn produces a peak in the spectrum of thermal photons emitted by the plasma,

$$\frac{dN_\gamma}{d\omega} \sim e^{-\omega/T} \chi_\mu^\mu(\omega, T). \quad (8.53)$$

The width of this peak is the width of the meson in the plasma.

The analysis above applies to an infinitely-extended plasma at constant temperature. Assuming that these results can be extrapolated to QCD, a crucial question is whether a peak in the photon spectrum could be observed in a heavy ion collision experiment. Natural heavy vector mesons to consider are the J/ψ and the Υ , since these are expected to survive deconfinement. We wish to compare the number of photons coming from these mesons to the number of photons coming from other sources. Accurately calculating the meson contribution would require a precise theoretical understanding of the dynamics of these mesons in the quark-gluon plasma, which at present is not available. Our goal will therefore be to estimate the order of magnitude of this effect with a simple recombination model. The details can be found in Ref. [628], so here we will only describe the result for heavy ion collisions at LHC energies.




Figure 8.22: In-medium vector meson–photon mixing. The imaginary part of this diagram yields the meson decay width into photons.




Figure 8.23: Thermal photon spectrum for LHC energies, $T_{\text{diss}} = 1.25 T_c$ and $M_{\text{charm}} = 1.7$ GeV. The (arbitrary) normalization is the same for all curves. The continuous, monotonically decreasing, blue curve is the background from light quarks. The continuous, red curve is the signal from J/ψ mesons. The dashed, black curve is the sum of the two. See Ref. [628] for details.

The result is summarized in Fig. 8.23, which shows the thermal photon spectrum coming from light quarks, the contribution from J/ψ mesons, and the sum of the two, for a thermal charm mass $M_{\text{charm}} = 1.7$ GeV and a J/ψ dissociation temperature $T_{\text{diss}} = 1.25 T_c$. Although the value of M_{charm} is relatively high, the values of M_{charm} and T_{diss} are within the range commonly considered in the literature. For the charm mass a typical range is $1.3 \leq M_{\text{charm}} \leq 1.7$ GeV because of a substantial thermal contribution — see e.g. Ref. [236] and Refs. therein. The value of T_{diss} is far from settled, but a typical range is $T_c \leq T_{\text{diss}} \leq 2T_c$ [190, 230, 241, 242, 247, 573, 629–633]. We have chosen these values for illustrative purposes, since they lead to an order-one enhancement in the spectrum. We emphasize, however, that whether this photon excess manifests itself as a peak, or only as an enhancement smoothly distributed over a broader range of frequencies, depends sensitively on these and other parameters. Qualitatively, the dependence on the main ones is as follows. Decreasing the quark mass decreases the magnitude of the J/ψ contribution. Perhaps surprisingly, higher values of T_{diss} make the peak less sharp. The in-medium width of the J/ψ used in Fig. 8.23 was 100 MeV. Increasing this by a factor of two turns the peak into an enhancement. Crucially, the J/ψ contribution depends quadratically on the $c\bar{c}$ cross-section. Since at RHIC energies this is




Figure 8.24: D-branes and open string in a black brane geometry.

believed to be ten times smaller than at LHC energies, the enhancement discussed above is presumably unobservable at RHIC.

These considerations show that a precise determination of the enhancement is not possible without a very detailed understanding of the in-medium dynamics of the J/ψ . On the other hand, they also illustrate that there exist reasonable values of the parameters for which this effect yields an order-one enhancement, or even a peak, in the spectrum of thermal photons produced by the quark-gluon plasma. This thermal excess is concentrated at photon energies roughly between 3 and 5 GeV. In this range the number of thermal photons in heavy ion collisions at the LHC is expected to be comparable to or larger than that of photons produced in initial partonic collisions that can be described using perturbative QCD [634]. Thus, we expect the thermal excess above to be observable even in the presence of the pQCD background.

The authors of Ref. [628] also examined the possibility of an analogous effect associated with the Υ meson, in which case $\omega_{\text{peak}} \sim 10$ GeV. At these energies the number of thermal photons is very much smaller than that coming from initial partonic collisions, so an observable effect is not expected.

8.6.2 A new mechanism of quark energy loss: Cherenkov emission of mesons

We now turn to another universal prediction that follows from the existence of a subluminal limiting velocity for mesons in the plasma. Consider a highly energetic quark moving through the plasma. In order to model this we consider a string whose endpoint moves with an arbitrary velocity v at an arbitrary radial position r_q — see Fig. 8.24. Roughly speaking, the interpretation of r_q in the gauge theory is that of the inverse size of the gluon cloud that dresses the quark. This can be seen, for example, by holographically computing the profile of $\langle \text{Tr} F^2(x) \rangle$ around a static quark source dual to a string whose endpoint sits at $r = r_q$ [613].

Two simple observations now lead to the effect that we are interested in. The first one is that the string endpoint is charged under the scalar and vector fields on the branes. In the gauge theory, this corresponds to an effective quark-meson coupling (see Fig. 8.25) of order $e \sim 1/\sqrt{N_c}$. Physically, this can be understood very simply. The fields on the branes describe fluctuations around the equilibrium configuration. The string endpoint pulls on the branes and therefore excites (i.e. it is charged under) these fields. The branes' tension is




Figure 8.25: Effective quark-meson coupling.

of order $1/g_s \sim N_c$, where g_s is the string coupling constant, whereas the string tension is N_c -independent. This means that the deformation of the branes caused by the string is of order $e^2 \sim 1/N_c$. We thus conclude that the dynamics of the ‘branes+string endpoint’ system is (a generalization of) that of classical electrodynamics in a medium in the presence of a fast-moving charge.

The second observation is that the velocity of the quark may exceed the limiting velocity of the mesons, since the redshift at the position of the string endpoint is smaller than that at the bottom of the branes. As in ordinary electrodynamics, if this happens then the string endpoint loses energy by Cherenkov-radiating into the fields on the branes. In the gauge theory, this translates into the quark losing energy by Cherenkov-radiating scalar and vector quarkonium mesons. The rate of energy loss is set by the square of the coupling, and is therefore of order $1/N_c$.

The quantitative details of the energy lost to Cherenkov radiation of quarkonium mesons by a quark propagating through the $\mathcal{N} = 4$ plasma can be found in [617, 618], so here we will only describe the result. For simplicity, we will assume that the quark moves with constant velocity along a straight line at a constant radial position. In reality, r_q and v will of course decrease with time because of the black hole gravitational pull and the energy loss. However, we will concentrate on the initial part of the trajectory (which is long provided the initial quark energy is large), for which r_q and v are approximately constant [531] — see Fig. 8.24. Finally, for illustrative purposes we will focus on the energy radiated into the transverse modes of vector mesons. The result is depicted in Fig. 8.26, and its main qualitative features are as follows.

As expected, we see that the quark only radiates into meson modes with phase velocity lower than v — those to the right of the dashed, vertical lines in Fig. 8.14. For fixed r_q , the energy loss increases monotonically with v up to the maximum allowed value of v — the local speed of light at r_q . As r_q decreases, the characteristic momentum q_{char} of the modes into which the energy is deposited increases. These modes become increasingly peaked near the bottom of the branes, and the energy loss diverges. However, this mathematical divergence is removed by physical effects we have not taken into account. For example, for sufficiently large q the mesons’ wave functions become concentrated on a region whose size is of order the string length, and hence stringy effects become important [568]. Also, as we saw in Section 8.4.2, mesons acquire widths $\Gamma \propto q^2$ at large q [571] and can no longer be treated as well defined quasiparticles. Finally, the approximation of a constant- v , constant- r_q trajectory ceases to be valid whenever the energy loss rate becomes large.

The Cherenkov radiation of quarkonium mesons by quarks depends only on the qualitative features of the dispersion relation of Fig. 8.14, which are universal for all gauge theory plasmas with a dual gravity description. Moreover, as we explained in Section 8.4.3, it is conceivable




Figure 8.26: Cherenkov energy loss into the transverse mode of vector quarkonium mesons. The continuous curves correspond to increasing values of r_q from left to right. The dotted curve is defined by the endpoints of the constant- r_q curves. See Refs. [617,618] for details.

that they may also hold for QCD mesons such as the J/ψ or the Υ whether or not a string dual of QCD exists. Here we will examine some qualitative consequences of this assumption for heavy ion collision experiments. Since the heavier the meson the more perturbative its properties become, we expect that our conclusions are more likely to be applicable to the charmonium sector than to the bottomonium sector.

An interesting feature of energy loss by Cherenkov radiation of quarkonia is that, unlike other energy-loss mechanisms, it is largely independent of the details of the quark excited state, such as the precise features of the gluon cloud around the quark, etc. In the gravity description these details would be encoded in the precise profile of the entire string, but the Cherenkov emission only depends on the trajectory of the string endpoint. This leads to a dramatic simplification which, with the further approximation of rectilinear uniform motion, reduces the parameters controlling the energy loss to two simple ones: the string endpoint velocity v and its radial position r_q . In order to obtain a ballpark estimate of the magnitude of the energy loss, we will assume that in a typical collision quarks are produced with order-one values of r_q (in units of R^2T). Under these circumstances the energy loss is of order unity in units of $(2\pi T)^2/N_c$, which for a temperature range of $T = 200 - 400$ MeV and $N_c = 3$ leads to $dE/dx \approx 2 - 8$ GeV/fm. This is of the same order of magnitude as other mechanisms of energy loss in the plasma; for example, the BDMPs radiative energy loss $dE/dx = \alpha_s C_F \hat{q} L/2$ yields values of $dE/dx = 7 - 40$ GeV/fm for $\hat{q} = 1 - 5$ GeV²/fm, $\alpha_s = 0.3$ and $L \approx 6$ fm. Since our gravity calculation is strictly valid only in the infinite-quark-energy limit (because of the linear trajectory approximation), we expect that our estimate is more likely to be applicable to highly energetic quarks at the LHC than to those at RHIC.

Even if in the quark-gluon plasma the magnitude of Cherenkov energy loss turns out to be subdominant with respect to other mechanisms, its velocity dependence and its geometric features may still make it identifiable. Indeed, this mechanism would only operate for quarks moving at velocities $v > v_{\text{lim}}$, with v_{lim} the limiting velocity of the corresponding quarkonium meson in the plasma. The presence of such a velocity threshold is the defining characteristic

of Cherenkov energy loss. The precise velocity at which the mechanism starts to operate may actually be higher than v_{lim} in some cases, since the additional requirement that the energy of the quark be equal or larger than the in-medium mass of the quarkonium meson must also be met.

Cherenkov mesons would be radiated at a characteristic angle $\cos \theta_c = v_{\text{lim}}/v$ with respect to the emitting quark, where v is the velocity of the quark. Taking the gravity result as guidance, v_{lim} could be as low as $v_{\text{lim}} = 0.35$ at the quarkonium dissociation temperature [565], corresponding to an angle as large as $\theta_c \approx 1.21$ rad. This would result in an excess of heavy quarkonium associated with high-energy quarks passing through the plasma. Our estimate of the energy loss suggests that the number of emitted J/ψ 's, for example, could range from one to three per fm. This emission pattern is similar to the emission of sound waves by an energetic parton [138] that we have reviewed in Section 7.3, in that both effects lead to a non-trivial angular structure. One important difference, however, is that the radiated quarkonium mesons would not thermalize and hence would not be part of a hydrodynamic shock wave. As in the Mach cone case, the meson emission pattern could be reflected in azimuthal dihadron correlations triggered by a high- p_T hadron. Due to surface bias, the energetic parton in the triggered direction is hardly modified, while the one propagating in the opposite direction moves through a significant amount of medium, emitting quarkonium mesons. Thus, under the above assumptions, the dihadron distribution with an associated J/ψ would have a ring-like structure peaked at an angle $\theta \approx \pi - \theta_c$. Even if this angular structure were to prove hard to discern, the simpler correlation that in events with a high- p_T hadron there are more J/Ψ mesons than in typical events may suffice as a distinctive signature, although further phenomenological modelling is required to establish this.

A final observation is that Cherenkov energy loss also has a non-trivial temperature dependence, since it requires that there are meson-like states in the plasma, and therefore it does not take place at temperatures above the meson dissociation temperature. Similarly, it is reasonable to assume that it does not occur at temperatures below T_c , since in this case we do not expect the meson dispersion relations to become spacelike.⁸ Under these circumstances, the Cherenkov mechanism is only effective over a limited range of temperatures $T_c < T < T_{\text{diss}}$ which, if $T_{\text{diss}} \gtrsim 1.2T_c$ as in Ref. [247], is a narrow interval. As was pointed out in Ref. [635], a mechanism of energy loss which is confined to a narrow range of temperatures in the vicinity of T_c concentrates the emission of energetic partons to a narrow layer within the collision geometry and is able to explain v_2 -data at high p_T at RHIC [636, 637]. Provided that the meson dissociation temperature T_{diss} is not much larger than T_c , the Cherenkov radiation of quarkonium mesons is one such mechanism.

⁸This assumption is certainly correct for plasmas with a gravity dual, since the corresponding geometry does not include a black hole horizon if $T < T_c$.

Chapter 9

Concluding remarks and outlook

Since the purpose of heavy ion collisions is to study the properties of Quantum Chromodynamics at extreme temperature and energy density, any successful phenomenology must ultimately be based on QCD. However, as discussed in Section 2, heavy ion phenomenology requires strong-coupling techniques not only for bulk thermodynamic quantities like the QCD equation of state, but also for many dynamical quantities at nonzero temperature, such as transport coefficients, relaxation times and quantities accessed by probes propagating through a plasma. By now, lattice-regularized QCD calculations provide well controlled results for the former, but progress on all the latter quantities is likely to be slow since one needs to overcome both conceptual limitations and limitations in computing power. Alternative strong-coupling tools are therefore desirable. The gauge/string duality provides one such tool for performing non-perturbative calculations for a wide class of non-abelian plasmas.

In this review we have mostly focussed on results obtained within one particular example of a gauge/string duality, namely the case in which the gauge theory is $\mathcal{N} = 4$ SYM or a small deformation thereof. One reason for this is pedagogical: $\mathcal{N} = 4$ SYM is arguably the simplest and best understood case of a gauge/string duality. By now many examples are known of more sophisticated string duals of non-supersymmetric, non-conformal QCD-like theories that exhibit confinement, spontaneous chiral symmetry breaking, thermal phase transitions, etc. However, many of these features become unimportant in the deconfined phase. For this reason, for the purpose of studying the QCD quark-gluon plasma, the restriction to $\mathcal{N} = 4$ SYM not only yields a gain in simplicity, but also does not imply a significant loss of generality, at least at the qualitative level. Moreover, none of these ‘more realistic’ theories can be considered in any sense a controlled approximation to QCD. Indeed, many differences remain including the presence of adjoint fermions and scalar fields, the lack of asymptotic freedom, the large- N_c approximation, etc. Some of these differences may be overcome if string theory in asymptotically AdS spacetimes becomes better understood. However, in the supergravity (plus classical strings and branes) limit currently accessible these caveats remain important to bear in mind when comparing to QCD.

In this context, it is clearly questionable to assess the interplay between heavy ion phenomenology and the gauge/string duality correspondence solely on the basis of testing the numerical agreement between theory and experiment. Rather, one should view this interplay in light of the standard scientific strategy that to gain significant insight into problems

that cannot be addressed within the current state of the art, it is useful to find a closely related theory within which such problems can be addressed with known technology and which encompasses the essential features of interest. For many dynamical features of phenomenological interest in heavy ion physics, controlled strong-coupling calculations in QCD are indeed not in immediate reach with the current state of the art. In contrast, within the gauge/string correspondence, it has been possible to formulate and solve the same problems in the strongly coupled plasmas of a large class of non-abelian quantum field theories. Among these, strongly coupled $\mathcal{N} = 4$ SYM theory at large- N_c turns out to provide the simplest model for the strongly coupled plasma being produced and probed in heavy ion collisions. Very often in the past, when theoretical physicists have introduced some model for the purpose of capturing the essence of some phenomenon or phenomena involving strongly coupled dynamics, the analysis of that model has then required further uncontrolled approximations. (Examples abound: Nambu–Jona-Lasinio models in which the QCD interaction is first replaced by a four-fermi coupling but one then still has to make a mean-field approximation; linear sigma models, again followed by a mean-field approximation; bag models; . . .) A great advantage of using a quantum field theory with a gravity dual as a model is that once we have picked such a theory, the calculations needed to address the problems of interest can be done rigorously at strong coupling, without requiring any further compromise. In some cases, the mere formulation of the problem in a way suitable for a gravitational dual calculation can lead to new results within QCD [90, 638, 639]. In many others, as we have seen, the existence of these solutions allows one to examine and understand the physics responsible for the processes of interest. The most important output of a successful model is understanding. Controlled quantitative calculations come later. Understanding how the dynamics works, what is important and what is extraneous, what the right picture is that helps you to think about the physics in a way that is both insightful and predictive, these must all come first.

At the least, the successes to date of the applications of gauge/string duality to problems arising from heavy ion collisions indicate that it provides us with a successful model, in the sense of the previous paragraph. However, there are many indications that it provides more. In solving these problems, some regularities have emerged in the form of universal properties, by which we mean properties common to all strongly coupled theories with gravity duals in the large- N_c , strong-coupling limit. These include both quantitative observables, such as the ratio of the thermodynamic potentials at strong and weak coupling (Section 6) and the value $\eta/s = 1/4\pi$ at strong coupling (Section 6.2), and qualitative features, such as the familiar fact that heavy quarkonium mesons remain bound in the plasma as well as the discovery that the dissociation temperature for quarkonium mesons drops with increasing meson velocity v like $(1-v^2)^{1/4}$ (Sections 7.7 and 8.4) and that high-momentum dispersion relations become space-like (Section 8.4). The discovery of these generic properties is important in order to extract lessons for QCD. Indeed, the fact that some properties are valid in a class of gauge-theory plasmas so broad as to include theories in different numbers of dimensions, with different field content, with or without chemical potentials, with or without confinement and chiral symmetry breaking, etc. leads one to suspect that such properties might be universal across the plasmas in a class of theories that includes the QCD — whether or not a string dual of QCD itself exists. The domain of applicability of this putative universality is at present unknown, both in the sense that we do not know to what theories it may apply and in the sense that we cannot say *a priori* which observables and phenomena are universal and which

others are theory-specific details. One guess as to a possible characterization in theory space could be that these universal features may be common across all gauge-theory plasmas that have no quasiparticle description (Section 6.3).

Even results obtained via the gauge/string duality that are not universal may provide guidance for our understanding of QCD at nonzero temperature and for heavy ion phenomenology. In many cases, these are strong-coupling results that differ parametrically from the corresponding weak-coupling results and therefore deliver valuable qualitative messages for the modeling of heavy ion collisions. In particular, the small values of the ratio η/s , of the relaxation time τ_π (Section 6.2) and of the heavy-quark diffusion constant (Section 7.2) showed that such small values can be realized in a gauge-theory plasma. In addition, the result for η/s in $\mathcal{N} = 4$ theory, combined with the results for the entropy density, pressure and energy density (Section 6.1) have taught us that a theory can have almost identical thermodynamics at zero and infinite coupling and yet have radically different hydrodynamics. The lesson that this provides for QCD is that the thermodynamic observables, although they are available from lattice simulations, are not good indicators of whether the quark-gluon plasma is weakly or strongly coupled, whereas the transport coefficients are.

Important lessons have also been extracted from the strong-coupling calculations of the jet quenching parameter \hat{q} and the heavy-quark drag coefficient η_D and momentum broadening κ (Sections 7.1, 7.2 and 7.5). These showed not only that these quantities can attain values significantly larger than indicated by perturbative estimates but also that while in perturbation theory both \hat{q} and κ are proportional to the entropy density, this is not the case at strong coupling, where both these quantities and η_D scale with the square root of the number of degrees of freedom. This result, which is valid for a large class of theories, corrected a naïve physical expectation that was supported by perturbation theory.

Perhaps even more fundamentally, the availability of rigorous, reliable results for any strongly coupled plasma (leave apart for a large class of them) can alter the very intuition we use to think about the dynamics of the quark-gluon plasma. In perturbation theory, one thinks of the plasma as being made of quark and gluon quasiparticles. However, gravity calculations of correlation functions valid at strong coupling show no evidence of the existence of any quasiparticle excitations composed from gluons and light quarks (Section 6.3). (Heavy, small, quarkonium mesons do survive as quasiparticles up to some dissociation temperature (Sections 8.5.2 and 8.6).) The presence or absence of quasiparticles is a major qualitative difference between the weak- and strong-coupling pictures of the plasma which is largely independent of the caveats associated with the use of gauge/string duality that we have described above.

Finally, two important roles played by the gauge/string duality are that of a testing ground of existing ideas and models relevant for heavy ion collisions, and that of a source of new ones in a regime in which guidance and inspiration from perturbation theory may be inapplicable or misleading. In its first role, the duality provides a rigorous field-theoretical framework within which to verify our intuition about the plasma. For example, explicit calculations gave support to previously suggested ideas about the hydrodynamical response of the plasma to high-energy particles (Section 7.3) or the possibility that heavy quarkonium mesons survive deconfinement (Section 8.3.1). In its second role, the duality has generated qualitatively new ideas which could not have been guessed from perturbation theory. Examples of these are the

non-trivial velocity dependence of screening lengths (Section 7.7), the in-medium conversion of mesons into photons (Section 8.6), the energy loss of heavy quarks via Cherenkov emission of mesons (Section 8.6.2), and the appearance of a phase transition associated with the dissociation of heavy quarkonium bound states (Section 8.3.2).

In summary, while it is true that caution and a critical mind must be exercised when trying to extract lessons from any gauge/gravity calculation, paying particular attention to its limitations and range of applicability, it is also undeniable that over the last few years a broad suite of qualitatively novel insights relevant for heavy ion phenomenology have emerged from detailed and quantitative calculations in the gravity duals of non-abelian field theories. As the phenomenology of heavy ion collisions moves to new, more quantitative, and more detailed, studies in the RHIC program, and as it moves to novel challenges at the LHC, we have every reason to expect that experimental information about additional properties of hot QCD matter will come into theoretical focus. Understanding properties of the QCD plasma, as well as its response to and its effects on probes, at strong coupling will therefore remain key issues in future analyses. We expect that the gauge/string correspondence will continue to play an important role in making progress on these issues.

Acknowledgements

We gratefully acknowledge our many collaborators with whom and from whom we have learned the subjects that we have reviewed. We also acknowledge T. Faulkner for helping us with some of the figures in this review. The work of HL was supported in part by a DOE Outstanding Junior Investigator grant. This research was supported in part by the DOE Offices of Nuclear and High Energy Physics under grants #DE-FG02-94ER40818 and #DE-FG02-05ER41360. DM is supported in part by grants 2009-SGR-168, MEC FPA 2007-66665-C02, MEC FPA 2007-66665-C01 and CPAN CSD2007-00042 Consolider-Ingenio 2010. JCS is supported by a Marie Curie Intra-European Fellowship PIEF-GA-2008-220207.

Appendix A

Green-Kubo formula for transport coefficients

Transport coefficients of a gauge theory plasma, such as the shear viscosity η , can be extracted from correlation functions of the gauge theory via a relation known as the Green-Kubo formula. Here, we derive this relation for the case of the shear viscosity. Let us consider a system in equilibrium and let us work in the fluid rest frame, meaning that $u^\mu = (1, \mathbf{0})$. Deviations from equilibrium are studied by introducing a small external source of the type

$$S = S_0 + \frac{1}{2} \int d^4x T^{\mu\nu} h_{\mu\nu}, \quad (\text{A.1})$$

where S_0 (S) is the action of the theory in the absence (presence) of the perturbation $h_{\mu\nu}$. To leading order in the perturbation, the expectation value of the stress tensor is

$$\langle T^{\mu\nu}(x) \rangle = \langle T^{\mu\nu}(x) \rangle_0 - \frac{1}{2} \int d^4y G_R^{\mu\nu, \alpha\beta}(x-y) h_{\alpha\beta}(y) \quad (\text{A.2})$$

where the subscript 0 indicates the unperturbed expectation value and the retarded correlator is given by

$$iG_R^{\mu\nu, \alpha\beta}(x-y) \equiv \theta(x^0 - y^0) \left\langle \left[T^{\mu\nu}(x), T^{\alpha\beta}(y) \right] \right\rangle. \quad (\text{A.3})$$

To extract the shear viscosity, we concentrate on an external perturbation of the form

$$h_{xy}(t, z). \quad (\text{A.4})$$

Upon Fourier transforming, this is equivalent to using rotational invariance to choose the wave vector of the perturbation, \mathbf{k} , along the \hat{z} direction. The off-diagonal components of the stress tensor are then given by

$$\langle T^{xy}(\omega, k) \rangle = -G_R^{xy, xy}(\omega, k) h_{xy}(\omega, k), \quad (\text{A.5})$$

to linear order in the perturbation. In the long wavelength limit, in which the typical variation of the perturbed metric is large compared with any correlation length, we obtain

$$\langle T^{xy} \rangle(t, z) = - \int \frac{d\omega}{2\pi} e^{-i\omega t} G_R^{xy, xy}(\omega, k=0) h_{xy}(\omega, z). \quad (\text{A.6})$$

This long wavelength expression may be compared to the hydrodynamic approximation by studying the reaction of the system to the source within the effective theory. The source $h_{\mu\nu}$ can be interpreted as a modification on the metric,

$$g_{\mu\nu} \rightarrow g_{\mu\nu} + h_{\mu\nu} . \quad (\text{A.7})$$

To leading order in the perturbation, the shear tensor defined in eq. (2.13) is given by

$$\sigma_{xy} = 2\Gamma_{xy}^0 = \partial_0 h_{xy} , \quad (\text{A.8})$$

where $\Gamma_{\nu\rho}^\mu$ are the Christoffel symbols. The hydrodynamic approximation is valid in the long time limit, when all microscopic processes have relaxed. In this limit, we can compare the linear response expression (A.5) to the expression obtained upon making the hydrodynamic approximation, namely (2.11). We conclude that

$$\eta = - \lim_{\omega \rightarrow 0} \frac{1}{\omega} \lim_{\vec{k} \rightarrow 0} \text{Im} G_R^{xy,xy}(\omega, \mathbf{k}) . \quad (\text{A.9})$$

This result is known as the Green-Kubo formula for the shear viscosity.

The above discussion for the stress tensor can also be generalized to other conserved currents. In general, the low frequency limit of $G_R(\omega, \vec{k})$ for a conserved current operator O defines a transport coefficient χ

$$\chi = - \lim_{\omega \rightarrow 0} \lim_{\vec{k} \rightarrow 0} \frac{1}{\omega} \text{Im} G_R(\omega, \mathbf{k}) , \quad (\text{A.10})$$

where the retarded correlator is defined analogously to eq. (A.3)

$$iG_R(x-y) \equiv \theta(x^0 - y^0) \langle [O(x), O(y)] \rangle . \quad (\text{A.11})$$

Appendix B

Hawking temperature of a general black brane metric

Here we calculate the Hawking temperature for a general class of black brane metrics of the form

$$ds^2 = g(r) \left[-f(r) dt^2 + d\vec{x}^2 \right] + \frac{1}{h(r)} dr^2. \quad (\text{B.1})$$

where we assume that $f(r)$ and $h(r)$ have a first-order zero at the horizon $r = r_0$, whereas $g(r)$ is non-vanishing there. We follow the standard method [326] and demand that the Euclidean continuation of the metric (B.1),

$$ds^2 = g(r) \left[f(r) dt_{\text{E}}^2 + d\vec{x}^2 \right] + \frac{1}{h(r)} dr^2, \quad (\text{B.2})$$

obtained by the replacement $t \rightarrow -it_{\text{E}}$, be regular at the horizon. Expanding (B.2) near $r = r_0$ one finds

$$ds^2 \approx \rho^2 d\theta^2 + d\rho^2 + g(r_0) d\vec{x}^2, \quad (\text{B.3})$$

where we have introduced new variables ρ, θ defined as

$$\rho = 2\sqrt{\frac{r - r_0}{h'(r_0)}}, \quad \theta = \frac{t_{\text{E}}}{2} \sqrt{g(r_0) f'(r_0) h'(r_0)}. \quad (\text{B.4})$$

The first two terms in the metric (B.3) describe a plane in polar coordinates, so in order to avoid a conical singularity at $\rho = 0$ we must require θ to have period 2π . From (B.4) we then see that the period $\beta = 1/T$ of the Euclidean time must be

$$\beta = \frac{1}{T} = \frac{4\pi}{\sqrt{g(r_0) f'(r_0) h'(r_0)}}. \quad (\text{B.5})$$

Appendix C

Holographic renormalization, one-point functions and an example of a Euclidean two-point function

Here we will illustrate the general prescription of Section 5.3.1 for computing Euclidean correlators. We will begin by giving a derivation of Eq. (5.46) at linear order in the external source, although we note that Eq. (5.46) is in fact valid at the nonlinear level [319], as follows from a simple generalization of the discussion that we shall present. Then, we will calculate the two-point function of a scalar operator $\mathcal{O}(x)$ in $\mathcal{N} = 4$ SYM at zero temperature. Although our main interest is in four-dimensional boundary theories, for the sake of generality we will present the formulas for a general dimension d .

Let Φ be the scalar field in AdS dual to \mathcal{O} . The Euclidean two-point function of \mathcal{O} is then given by the right-hand side of Eq. (5.39) with $n = 2$. In order to evaluate this, we first need to solve the classical equation of motion for Φ subject to the boundary condition (5.38), and then evaluate the action on that solution. Since in order to obtain the two-point function we only need to take two functional derivatives of the action, it suffices to keep only the terms in the action that are quadratic in Φ , ignoring all interaction terms. At this level, the action is given by Eq. (5.17), except without the minus sign inside $\sqrt{-g}$, as appropriate for Euclidean signature:

$$S = -\frac{1}{2} \int dz d^d x \sqrt{g} [g^{MN} \partial_M \Phi \partial_N \Phi + m^2 \Phi^2] + \dots \quad (\text{C.1})$$

Note that we have adopted an overall sign convention for the Euclidean action appropriate for (5.37). The metric is that of pure Euclidean AdS and takes the form

$$ds^2 = \frac{R^2}{z^2} (dz^2 + \delta_{\mu\nu} dx^\mu dx^\nu) . \quad (\text{C.2})$$

We will work in momentum space along the boundary directions. The equation of motion for $\Phi(z, k)$ then takes the form (5.19), which we reproduce here for convenience

$$z^{d+1} \partial_z \left(z^{1-d} \partial_z \phi \right) - k^2 z^2 \Phi - m^2 R^2 \Phi = 0 , \quad (\text{C.3})$$

with $k^2 = \delta_{\mu\nu} k^\mu k^\nu$ as appropriate in Euclidean signature.

A simple integration by parts shows that, when evaluated on a solution Φ_c , S reduces to the boundary term

$$S[\Phi_c] = -\frac{1}{2} \lim_{\epsilon \rightarrow 0} \int_{z=\epsilon} \frac{d^d k}{(2\pi)^d} \Pi_c(-k) \Phi_c(k), \quad (\text{C.4})$$

where Π_c is the canonical momentum associated with the z -foliation,

$$\Pi = -\sqrt{-g} g^{zz} \partial_z \Phi, \quad (\text{C.5})$$

evaluated at the solution Φ_c . Since $z = 0$ is a regular singular point of eq. (C.3), it is possible to choose a basis for $\Phi_{1,2}$ given by

$$\Phi_1 \rightarrow z^{d-\Delta}, \quad \Phi_2 \rightarrow z^\Delta, \quad \text{as } z \rightarrow 0, \quad (\text{C.6})$$

with the corresponding canonical momentum $\Pi_{1,2}(z, k)$ behaving as

$$\Pi_1 \rightarrow -(d-\Delta) z^{-\Delta}, \quad \Pi_2 \rightarrow -\Delta z^{-(d-\Delta)}, \quad \text{as } z \rightarrow 0, \quad (\text{C.7})$$

where

$$\Delta = \frac{d}{2} + \nu, \quad \nu = \sqrt{\frac{d^2}{4} + m^2 R^2}. \quad (\text{C.8})$$

Then Φ_c and its canonical momentum can be expanded as

$$\begin{aligned} \Phi_c(z, k) &= A(k) \Phi_1(z, k) + B(k) \Phi_2(z, k), \\ \Pi_c(z, k) &= A(k) \Pi_1(z, k) + B(k) \Pi_2(z, k), \end{aligned} \quad (\text{C.9})$$

as in (5.20), and the classical on-shell action becomes

$$\begin{aligned} S[\Phi_c] &= -\frac{1}{2} \lim_{\epsilon \rightarrow 0} \int_{z=\epsilon} \frac{d^d k}{(2\pi)^d} [(A(-k)A(k)\Pi_1(-k)\Phi_1(k) + B(-k)B(k)\Pi_2(-k)\Phi_2(k) \\ &\quad + A(-k)B(k)(\Pi_1(-k)\Phi_2(-k) + \Phi_1(-k)\Pi_2(k))]. \end{aligned} \quad (\text{C.10})$$

Note that given $\nu > 0$, in the $\epsilon \rightarrow 0$ limit the first term on the RHS of (C.10) contains divergences and thus S requires renormalization. These divergences can be interpreted as dual to UV divergences of the boundary gauge theory. A local counter-term action S_{ct} defined on the cutoff surface $z = \epsilon$ can be introduced to cancel the divergences. From (C.10) we need to choose ¹

$$S_{ct} = \frac{1}{2} \int_{z=\epsilon} \frac{d^d k}{(2\pi)^d} \frac{\Pi_1(-k)}{\Phi_1(k)} \Phi(-k) \Phi(k). \quad (\text{C.11})$$

Note that in this expression $\Pi_1(-k)/\Phi_1(k)$ is a function of z which in the $z \rightarrow \epsilon$ limit is independent of k . The renormalized on-shell action is then given by

$$S^{(\text{ren})}[\Phi_c] \equiv S[\Phi_c] + S_{ct}[\Phi_c] = \frac{1}{2} \int \frac{d^d k}{(2\pi)^d} 2\nu A(-k)B(k), \quad (\text{C.12})$$

¹We assume that 2ν is not an integer. If that is the case, then extra logarithmic terms arise. See discussion in [318].

where we have dropped terms which vanish in the $\epsilon \rightarrow 0$ limit; the action is now finite.

We now impose the (momentum space version of the) boundary condition (5.38) on Φ_c , i.e.

$$\Phi_c(\epsilon, k) = \epsilon^{d-\Delta} \phi(k) \quad \text{as } \epsilon \rightarrow 0, \quad (\text{C.13})$$

which from equation (C.9) gives

$$A(k) = \phi(k) + \text{terms that vanish as } \epsilon \rightarrow 0. \quad (\text{C.14})$$

We also need to impose the condition that Φ_c be regular everywhere in the interior. This extra condition then fixes the solution of (C.3) completely, which in turn determines the ratio $\chi \equiv B/A$ in terms of which $B = \chi\phi$. The renormalized action can now be written as

$$S^{(\text{ren})}[\Phi_c] = \frac{1}{2} \int \frac{d^d k}{(2\pi)^d} 2\nu \chi \phi(-k) \phi(k). \quad (\text{C.15})$$

It follows that the one-point function is given by

$$\langle \mathcal{O}(k) \rangle_\phi = \frac{\delta S^{(\text{ren})}[\Phi_c]}{\delta \phi(-k)} = 2\nu \chi \phi(k) = 2\nu B(k), \quad (\text{C.16})$$

which is the momentum space version of (5.46). The two point function is then

$$G_E(k) = \frac{\langle \mathcal{O}(k) \rangle_\phi}{\phi(k)} = 2\nu \frac{B(k)}{A(k)}. \quad (\text{C.17})$$

Note that the above discussion only uses the form of equation (C.3) near $z = 0$ and thus applies to generic asymptotically-AdS geometries.

For pure AdS, equation (C.3) can in fact be solved exactly with $\Phi_{1,2}$ given by

$$\Phi_2 = \Gamma(1 + \nu) \left(\frac{k}{2}\right)^{-\nu} z^{\frac{d}{2}} I_\nu(kz), \quad \Phi_1 = \Gamma(1 - \nu) \left(\frac{k}{2}\right)^\nu z^{\frac{d}{2}} I_{-\nu}(kz), \quad (\text{C.18})$$

where $I(x)$ is the modified Bessel function of the first kind. Requiring Φ_c to be regular at $z \rightarrow \infty$ determines the solution up to an overall multiplicative constant:

$$\Phi_c = z^{\frac{d}{2}} K_\nu(kz), \quad (\text{C.19})$$

where $K(x)$ is the modified Bessel function of the second kind. From (C.9), (C.18) and (C.19) we then find that

$$\frac{B}{A} = \frac{\Gamma(-\nu)}{\Gamma(\nu)} \left(\frac{k}{2}\right)^{2\nu} \quad (\text{C.20})$$

and thus

$$G_E(k) = 2\nu \frac{\Gamma(-\nu)}{\Gamma(\nu)} \left(\frac{k}{2}\right)^{2\nu}. \quad (\text{C.21})$$

Bibliography

- [1] D. T. Son and A. O. Starinets, “Viscosity, Black Holes, and Quantum Field Theory,” [Ann. Rev. Nucl. Part. Sci. **57** \(2007\) 95–118](#), [arXiv:0704.0240 \[hep-th\]](#).
- [2] T. Schafer and D. Teaney, “Nearly Perfect Fluidity: From Cold Atomic Gases to Hot Quark Gluon Plasmas,” [Rept. Prog. Phys. **72** \(2009\) 126001](#), [arXiv:0904.3107 \[hep-ph\]](#).
- [3] K. Peeters and M. Zamaklar, “The string/gauge theory correspondence in QCD,” [Eur. Phys. J. ST **152** \(2007\) 113–138](#), [arXiv:0708.1502 \[hep-ph\]](#).
- [4] D. Mateos, “String Theory and Quantum Chromodynamics,” [Class. Quant. Grav. **24** \(2007\) S713–S740](#), [arXiv:0709.1523 \[hep-th\]](#).
- [5] J. Erdmenger, N. Evans, I. Kirsch, and E. Threlfall, “Mesons in Gauge/Gravity Duals - A Review,” [Eur. Phys. J. **A35** \(2008\) 81–133](#), [arXiv:0711.4467 \[hep-th\]](#).
- [6] S. S. Gubser and A. Karch, “From gauge-string duality to strong interactions: a Pedestrian’s Guide,” [Ann. Rev. Nucl. Part. Sci. **59** \(2009\) 145–168](#), [arXiv:0901.0935 \[hep-th\]](#).
- [7] S. S. Gubser, S. S. Pufu, F. D. Rocha, and A. Yarom, “Energy loss in a strongly coupled thermal medium and the gauge-string duality,” [arXiv:0902.4041 \[hep-th\]](#).
- [8] F. Gelis, E. Iancu, J. Jalilian-Marian, and R. Venugopalan, “The Color Glass Condensate,” [arXiv:1002.0333 \[hep-ph\]](#).
- [9] Y. Hatta, E. Iancu, and A. Mueller, “Deep inelastic scattering at strong coupling from gauge/string duality: The Saturation line,” [JHEP **0801** \(2008\) 026](#), [arXiv:0710.2148 \[hep-th\]](#).
- [10] Y. Hatta, E. Iancu, and A. Mueller, “Deep inelastic scattering off a N=4 SYM plasma at strong coupling,” [JHEP **0801** \(2008\) 063](#), [arXiv:0710.5297 \[hep-th\]](#).
- [11] A. Mueller, A. Shoshi, and B.-W. Xiao, “Deep inelastic and dipole scattering on finite length hot N=4 SYM matter,” [Nucl.Phys. **A822** \(2009\) 20–40](#), [arXiv:0812.2897 \[hep-th\]](#).
- [12] C. Bayona, H. Boschi-Filho, and N. R. Braga, “Deep inelastic scattering off a plasma with flavour from D3-D7 brane model,” [Phys.Rev. **D81** \(2010\) 086003](#), [arXiv:0912.0231 \[hep-th\]](#).

- [13] J. L. Albacete, Y. V. Kovchegov, and A. Taliotis, “DIS on a Large Nucleus in AdS/CFT,” [JHEP](#) **0807** (2008) 074, [arXiv:0806.1484 \[hep-th\]](#).
- [14] L. Cornalba and M. S. Costa, “Saturation in Deep Inelastic Scattering from AdS/CFT,” [Phys.Rev.](#) **D78** (2008) 096010, [arXiv:0804.1562 \[hep-ph\]](#).
- [15] R. Baier, A. H. Mueller, D. Schiff, and D. Son, “‘Bottom up’ thermalization in heavy ion collisions,” [Phys.Lett.](#) **B502** (2001) 51–58, [arXiv:hep-ph/0009237 \[hep-ph\]](#).
- [16] P. Romatschke and M. Strickland, “Collective modes of an anisotropic quark gluon plasma,” [Phys.Rev.](#) **D68** (2003) 036004, [arXiv:hep-ph/0304092 \[hep-ph\]](#).
- [17] S. Mrowczynski, “Instabilities driven equilibration of the quark-gluon plasma,” [Acta Phys.Polon.](#) **B37** (2006) 427–454, [arXiv:hep-ph/0511052 \[hep-ph\]](#).
- [18] P. B. Arnold, J. Lenaghan, G. D. Moore, and L. G. Yaffe, “Apparent thermalization due to plasma instabilities in quark-gluon plasma,” [Phys.Rev.Lett.](#) **94** (2005) 072302, [arXiv:nucl-th/0409068 \[nucl-th\]](#).
- [19] P. B. Arnold, “Quark-Gluon Plasmas and Thermalization,” [Int.J.Mod.Phys.](#) **E16** (2007) 2555–2594, [arXiv:0708.0812 \[hep-ph\]](#).
- [20] K. Kang and H. Nastase, “High energy QCD from Planckian scattering in AdS and the Froissart bound,” [Phys.Rev.](#) **D72** (2005) 106003, [arXiv:hep-th/0410173 \[hep-th\]](#).
- [21] H. Nastase, “The RHIC fireball as a dual black hole,” [arXiv:hep-th/0501068 \[hep-th\]](#).
- [22] H. Nastase, “More on the RHIC fireball and dual black holes,” [arXiv:hep-th/0603176 \[hep-th\]](#).
- [23] H. Nastase, “AdS-CFT and the RHIC fireball,” [Prog.Theor.Phys.Suppl.](#) **174** (2008) 274–285, [arXiv:0805.3579 \[hep-th\]](#).
- [24] R. A. Janik and R. B. Peshanski, “Asymptotic perfect fluid dynamics as a consequence of AdS/CFT,” [Phys. Rev.](#) **D73** (2006) 045013, [arXiv:hep-th/0512162](#).
- [25] P. M. Chesler and L. G. Yaffe, “Holography and colliding gravitational shock waves in asymptotically AdS₅ spacetime,” [arXiv:1011.3562 \[hep-th\]](#).
- [26] S. Lin and E. Shuryak, “Toward the AdS/CFT gravity dual for High Energy Collisions: I.Falling into the AdS,” [Phys.Rev.](#) **D77** (2008) 085013, [arXiv:hep-ph/0610168 \[hep-ph\]](#).
- [27] Y. V. Kovchegov and A. Taliotis, “Early Time Dynamics in Heavy Ion Collisions from AdS/CFT Correspondence,” [Phys.Rev.](#) **C76** (2007) 014905, [arXiv:0705.1234 \[hep-ph\]](#).

- [28] S. Lin and E. Shuryak, “Toward the AdS/CFT gravity dual for high energy collisions. 2. The stress tensor on the boundary,” [Phys.Rev. D77 \(2008\) 085014](#), [arXiv:0711.0736 \[hep-th\]](#).
- [29] K. Kajantie, J. Louko, and T. Tahkokallio, “Gravity dual of conformal matter collisions in 1+1 dimensions,” [Phys.Rev. D77 \(2008\) 066001](#), [arXiv:0801.0198 \[hep-th\]](#).
- [30] D. Grumiller and P. Romatschke, “On the collision of two shock waves in AdS(5),” [JHEP 0808 \(2008\) 027](#), [arXiv:0803.3226 \[hep-th\]](#).
- [31] S. S. Gubser, S. S. Pufu, and A. Yarom, “Entropy production in collisions of gravitational shock waves and of heavy ions,” [Phys.Rev. D78 \(2008\) 066014](#), [arXiv:0805.1551 \[hep-th\]](#).
- [32] J. L. Albacete, Y. V. Kovchegov, and A. Taliotis, “Modeling Heavy Ion Collisions in AdS/CFT,” [JHEP 0807 \(2008\) 100](#), [arXiv:0805.2927 \[hep-th\]](#).
- [33] S. Lin and E. Shuryak, “Toward the AdS/CFT Gravity Dual for High Energy Collisions. 3. Gravitationally Collapsing Shell and Quasiequilibrium,” [Phys.Rev. D78 \(2008\) 125018](#), [arXiv:0808.0910 \[hep-th\]](#).
- [34] L. Alvarez-Gaume, C. Gomez, A. Sabio Vera, A. Tavanfar, and M. A. Vazquez-Mozo, “Critical formation of trapped surfaces in the collision of gravitational shock waves,” [JHEP 0902 \(2009\) 009](#), [arXiv:0811.3969 \[hep-th\]](#).
- [35] P. M. Chesler and L. G. Yaffe, “Horizon formation and far-from-equilibrium isotropization in supersymmetric Yang-Mills plasma,” [Phys. Rev. Lett. 102 \(2009\) 211601](#), [arXiv:0812.2053 \[hep-th\]](#).
- [36] V. Balasubramanian, A. Bernamonti, J. de Boer, N. Copland, B. Craps, et al., “Thermalization of Strongly Coupled Field Theories,” [Phys.Rev.Lett. 106 \(2011\) 191601](#), [arXiv:1012.4753 \[hep-th\]](#).
- [37] S. Lin and E. Shuryak, “Grazing Collisions of Gravitational Shock Waves and Entropy Production in Heavy Ion Collision,” [Phys.Rev. D79 \(2009\) 124015](#), [arXiv:0902.1508 \[hep-th\]](#).
- [38] S. S. Gubser, S. S. Pufu, and A. Yarom, “Off-center collisions in AdS(5) with applications to multiplicity estimates in heavy-ion collisions,” [JHEP 0911 \(2009\) 050](#), [arXiv:0902.4062 \[hep-th\]](#).
- [39] G. Beuf, M. P. Heller, R. A. Janik, and R. Peschanski, “Boost-invariant early time dynamics from AdS/CFT,” [JHEP 0910 \(2009\) 043](#), [arXiv:0906.4423 \[hep-th\]](#).
- [40] E. Avsar, E. Iancu, L. McLerran, and D. Triantafyllopoulos, “Shockwaves and deep inelastic scattering within the gauge/gravity duality,” [JHEP 0911 \(2009\) 105](#), [arXiv:0907.4604 \[hep-th\]](#).

- [41] P. M. Chesler and L. G. Yaffe, “Boost invariant flow, black hole formation, and far-from-equilibrium dynamics in $N = 4$ supersymmetric Yang-Mills theory,” [Phys. Rev.](#) **D82** (2010) 026006, [arXiv:0906.4426](#) [hep-th].
- [42] S. Lin and E. Shuryak, “On the critical condition in gravitational shock wave collision and heavy ion collisions,” [arXiv:1011.1918](#) [hep-th].
- [43] I. Aref’eva, A. Bagrov, and L. Joukovskaya, “Critical Trapped Surfaces Formation in the Collision of Ultrarelativistic Charges in (A)dS,” [JHEP](#) **1003** (2010) 002, [arXiv:0909.1294](#) [hep-th].
- [44] J. Erlich, E. Katz, D. T. Son, and M. A. Stephanov, “QCD and a holographic model of hadrons,” [Phys.Rev.Lett.](#) **95** (2005) 261602, [arXiv:hep-ph/0501128](#) [hep-ph].
- [45] L. Da Rold and A. Pomarol, “Chiral symmetry breaking from five dimensional spaces,” [Nucl. Phys.](#) **B721** (2005) 79–97, [arXiv:hep-ph/0501218](#).
- [46] A. Karch, E. Katz, D. T. Son, and M. A. Stephanov, “Linear confinement and AdS/QCD,” [Phys.Rev.](#) **D74** (2006) 015005, [arXiv:hep-ph/0602229](#) [hep-ph].
- [47] O. Andreev and V. I. Zakharov, “Heavy-quark potentials and AdS/QCD,” [Phys.Rev.](#) **D74** (2006) 025023, [arXiv:hep-ph/0604204](#) [hep-ph].
- [48] S. J. Brodsky and G. F. de Teramond, “Light-Front Dynamics and AdS/QCD Correspondence: The Pion Form Factor in the Space- and Time-Like Regions,” [Phys.Rev.](#) **D77** (2008) 056007, [arXiv:0707.3859](#) [hep-ph].
- [49] U. Gursoy and E. Kiritsis, “Exploring improved holographic theories for QCD: Part I,” [JHEP](#) **0802** (2008) 032, [arXiv:0707.1324](#) [hep-th].
- [50] U. Gursoy, E. Kiritsis, and F. Nitti, “Exploring improved holographic theories for QCD: Part II,” [JHEP](#) **0802** (2008) 019, [arXiv:0707.1349](#) [hep-th].
- [51] U. Gursoy, E. Kiritsis, L. Mazzanti, and F. Nitti, “Deconfinement and Gluon Plasma Dynamics in Improved Holographic QCD,” [Phys. Rev. Lett.](#) **101** (2008) 181601, [arXiv:0804.0899](#) [hep-th].
- [52] K. Kajantie, T. Tahkokallio, and J.-T. Yee, “Thermodynamics of AdS/QCD,” [JHEP](#) **0701** (2007) 019, [arXiv:hep-ph/0609254](#) [hep-ph].
- [53] C. Ballon Bayona, H. Boschi-Filho, N. R. Braga, and L. A. Pando Zayas, “On a Holographic Model for Confinement/Deconfinement,” [Phys.Rev.](#) **D77** (2008) 046002, [arXiv:0705.1529](#) [hep-th].
- [54] J. Polchinski and M. J. Strassler, “Hard scattering and gauge / string duality,” [Phys.Rev.Lett.](#) **88** (2002) 031601, [arXiv:hep-th/0109174](#) [hep-th].
- [55] H. Boschi-Filho and N. R. Braga, “Gauge / string duality and scalar glueball mass ratios,” [JHEP](#) **0305** (2003) 009, [arXiv:hep-th/0212207](#) [hep-th].

- [56] H. Boschi-Filho and N. R. Braga, “QCD / string holographic mapping and glueball mass spectrum,” [Eur.Phys.J. C32 \(2004\) 529–533](#), [arXiv:hep-th/0209080 \[hep-th\]](#).
- [57] P. Anderson, “More is Different,” [Science, New Series 177 \(1972\) 393–396](#).
- [58] **PHOBOS** Collaboration, B. B. Back *et al.*, “Charged-particle pseudorapidity distributions in Au + Au collisions at $s(\text{NN})^{1/2} = 62.4\text{-GeV}$,” [Phys. Rev. C74 \(2006\) 021901](#), [nucl-ex/0509034](#).
- [59] **The ALICE** Collaboration, K. Aamodt *et al.*, “Charged-particle multiplicity density at mid-rapidity in central Pb-Pb collisions at $\sqrt{s(\text{NN})} = 2.76\text{ TeV}$,” [arXiv:1011.3916 \[nucl-ex\]](#).
- [60] **STAR** Collaboration, J. Adams *et al.*, “Transverse momentum and collision energy dependence of high $p(\text{T})$ hadron suppression in Au + Au collisions at ultrarelativistic energies,” [Phys. Rev. Lett. 91 \(2003\) 172302](#), [arXiv:nucl-ex/0305015](#).
- [61] B. B. Back *et al.*, “The significance of the fragmentation region in ultrarelativistic heavy ion collisions,” [Phys. Rev. Lett. 91 \(2003\) 052303](#), [arXiv:nucl-ex/0210015](#).
- [62] B. B. Back *et al.*, “The PHOBOS perspective on discoveries at RHIC,” [Nucl. Phys. A757 \(2005\) 28–101](#), [arXiv:nucl-ex/0410022](#).
- [63] N. Armesto *et al.*, “Heavy Ion Collisions at the LHC - Last Call for Predictions,” [J. Phys. G35 \(2008\) 054001](#), [arXiv:0711.0974 \[hep-ph\]](#).
- [64] J. D. Bjorken, “Highly Relativistic Nucleus-Nucleus Collisions: The Central Rapidity Region,” [Phys. Rev. D27 \(1983\) 140–151](#).
- [65] A. Bazavov *et al.*, “Equation of state and QCD transition at finite temperature,” [Phys. Rev. D80 \(2009\) 014504](#), [arXiv:0903.4379 \[hep-lat\]](#).
- [66] **PHENIX** Collaboration, S. S. Adler *et al.*, “Identified charged particle spectra and yields in Au + Au collisions at $s(\text{NN})^{1/2} = 200\text{-GeV}$,” [Phys. Rev. C69 \(2004\) 034909](#), [arXiv:nucl-ex/0307022](#).
- [67] **STAR** Collaboration, J. Adams *et al.*, “Pion, kaon, proton and anti-proton transverse momentum distributions from p + p and d + Au collisions at $s(\text{NN})^{1/2} = 200\text{-GeV}$,” [Phys. Lett. B616 \(2005\) 8–16](#), [arXiv:nucl-ex/0309012](#).
- [68] **STAR** Collaboration, J. Adams *et al.*, “Experimental and theoretical challenges in the search for the quark gluon plasma: The STAR collaboration’s critical assessment of the evidence from RHIC collisions,” [Nucl. Phys. A757 \(2005\) 102–183](#), [arXiv:nucl-ex/0501009](#).
- [69] A. Andronic, P. Braun-Munzinger, and J. Stachel, “Hadron production in central nucleus nucleus collisions at chemical freeze-out,” [Nucl. Phys. A772 \(2006\) 167–199](#), [arXiv:nucl-th/0511071](#).

- [70] **BRAHMS** Collaboration, I. G. Bearden et al., “Nuclear stopping in Au + Au collisions at $s(\text{NN})^{1/2} = 200\text{-GeV}$,” Phys. Rev. Lett. **93** (2004) 102301, [nucl-ex/0312023](#).
- [71] P. Braun-Munzinger, K. Redlich, and J. Stachel, “Particle production in heavy ion collisions,” [arXiv:nucl-th/0304013](#).
- [72] J. Y. Ollitrault, “Nucleus nucleus collisions at RHIC: A review,” Pramana **67** (2006) 899–914.
- [73] P. F. Kolb and U. W. Heinz, “Hydrodynamic description of ultrarelativistic heavy-ion collisions,” [arXiv:nucl-th/0305084](#).
- [74] **STAR** Collaboration, J. Adams et al., “Azimuthal anisotropy in Au + Au collisions at $s(\text{NN})^{1/2} = 200\text{-GeV}$,” Phys. Rev. **C72** (2005) 014904, [nucl-ex/0409033](#).
- [75] N. Borghini, P. M. Dinh, and J.-Y. Ollitrault, “A new method for measuring azimuthal distributions in nucleus nucleus collisions,” Phys. Rev. **C63** (2001) 054906, [arXiv:nucl-th/0007063](#).
- [76] P. F. Kolb, U. W. Heinz, P. Huovinen, K. J. Eskola, and K. Tuominen, “Centrality dependence of multiplicity, transverse energy, and elliptic flow from hydrodynamics,” Nucl. Phys. **A696** (2001) 197–215, [arXiv:hep-ph/0103234](#).
- [77] T. Hirano, U. W. Heinz, D. Kharzeev, R. Lacey, and Y. Nara, “Hadronic dissipative effects on elliptic flow in ultrarelativistic heavy-ion collisions,” Phys. Lett. **B636** (2006) 299–304, [arXiv:nucl-th/0511046](#).
- [78] D. Forster, Hydrodynamic Fluctuations, Broken Symmetry, and Correlation Functions. W. A. Benjamin, Inc., Advanced Book Program Reading, Massachusetts, 1975.
- [79] D. T. Son, “Hydrodynamics of nuclear matter in the chiral limit,” Phys. Rev. Lett. **84** (2000) 3771–3774, [hep-ph/9912267](#).
- [80] T. Hirano and K. Tsuda, “Collective flow and two pion correlations from a relativistic hydrodynamic model with early chemical freeze out,” Phys. Rev. **C66** (2002) 054905, [arXiv:nucl-th/0205043](#).
- [81] Y. Hama et al., “3D relativistic hydrodynamic computations using lattice- QCD inspired equations of state,” Nucl. Phys. **A774** (2006) 169–178, [arXiv:hep-ph/0510096](#).
- [82] C. Nonaka and S. A. Bass, “Space-time evolution of bulk QCD matter,” Phys. Rev. **C75** (2007) 014902, [arXiv:nucl-th/0607018](#).
- [83] T. Renk, J. Ruppert, C. Nonaka, and S. A. Bass, “Jet-quenching in a 3D hydrodynamic medium,” Phys. Rev. **C75** (2007) 031902, [arXiv:nucl-th/0611027](#).
- [84] M. Cheng et al., “The QCD Equation of State with almost Physical Quark Masses,” [arXiv:0710.0354 \[hep-lat\]](#).

- [85] S. Borsanyi et al., “The QCD equation of state with dynamical quarks,” [arXiv:1007.2580 \[hep-lat\]](#).
- [86] H. B. Meyer, “A calculation of the bulk viscosity in SU(3) gluodynamics,” [Phys. Rev. Lett.](#) **100** (2008) 162001, [arXiv:0710.3717 \[hep-lat\]](#).
- [87] I. Mueller, “Zum Paradoxon der Waermeleitungstheorie,” [Z. Phys.](#) **198** (1967) 329.
- [88] W. Israel, “Nonstationary irreversible thermodynamics: A Causal relativistic theory,” [Ann. Phys.](#) **100** (1976) 310–331.
- [89] W. Israel and J. M. Stewart, “Transient relativistic thermodynamics and kinetic theory,” [Ann. Phys.](#) **118** (1979) 341–372.
- [90] R. Baier, P. Romatschke, D. T. Son, A. O. Starinets, and M. A. Stephanov, “Relativistic viscous hydrodynamics, conformal invariance, and holography,” [JHEP](#) **0804** (2008) 100, [arXiv:0712.2451 \[hep-th\]](#).
- [91] S. Bhattacharyya, V. E. Hubeny, S. Minwalla, and M. Rangamani, “Nonlinear Fluid Dynamics from Gravity,” [JHEP](#) **0802** (2008) 045, [arXiv:0712.2456 \[hep-th\]](#).
- [92] P. Romatschke, “Relativistic Viscous Fluid Dynamics and Non-Equilibrium Entropy,” [Class. Quant. Grav.](#) **27** (2010) 025006, [arXiv:0906.4787 \[hep-th\]](#).
- [93] A. Muronga, “Second order dissipative fluid dynamics for ultra- relativistic nuclear collisions,” [Phys. Rev. Lett.](#) **88** (2002) 062302, [arXiv:nucl-th/0104064](#).
- [94] D. Teaney, “Effect of shear viscosity on spectra, elliptic flow, and Hanbury Brown-Twiss radii,” [Phys. Rev.](#) **C68** (2003) 034913, [arXiv:nucl-th/0301099](#).
- [95] A. Muronga, “Causal Theories of Dissipative Relativistic Fluid Dynamics for Nuclear Collisions,” [Phys. Rev.](#) **C69** (2004) 034903, [arXiv:nucl-th/0309055](#).
- [96] A. Muronga and D. H. Rischke, “Evolution of hot, dissipative quark matter in relativistic nuclear collisions,” [arXiv:nucl-th/0407114](#).
- [97] R. Baier, P. Romatschke, and U. A. Wiedemann, “Dissipative hydrodynamics and heavy ion collisions,” [Phys.Rev.](#) **C73** (2006) 064903, [arXiv:hep-ph/0602249 \[hep-ph\]](#).
- [98] R. Baier and P. Romatschke, “Causal viscous hydrodynamics for central heavy-ion collisions,” [Eur. Phys. J.](#) **C51** (2007) 677–687, [arXiv:nucl-th/0610108](#).
- [99] P. Romatschke and U. Romatschke, “Viscosity Information from Relativistic Nuclear Collisions: How Perfect is the Fluid Observed at RHIC?,” [Phys. Rev. Lett.](#) **99** (2007) 172301, [arXiv:0706.1522 \[nucl-th\]](#).
- [100] H. Song and U. W. Heinz, “Suppression of elliptic flow in a minimally viscous quark-gluon plasma,” [Phys. Lett.](#) **B658** (2008) 279–283, [arXiv:0709.0742 \[nucl-th\]](#).

- [101] K. Dusling and D. Teaney, “Simulating elliptic flow with viscous hydrodynamics,” [Phys. Rev. C77 \(2008\) 034905](#), [arXiv:0710.5932 \[nucl-th\]](#).
- [102] H. Song and U. W. Heinz, “Causal viscous hydrodynamics in 2+1 dimensions for relativistic heavy-ion collisions,” [Phys. Rev. C77 \(2008\) 064901](#), [arXiv:0712.3715 \[nucl-th\]](#).
- [103] M. Luzum and P. Romatschke, “Conformal Relativistic Viscous Hydrodynamics: Applications to RHIC results at $\sqrt{(s_{NN})} = 200$ GeV,” [Phys. Rev. C78 \(2008\) 034915](#), [arXiv:0804.4015 \[nucl-th\]](#).
- [104] H. Song and U. W. Heinz, “Multiplicity scaling in ideal and viscous hydrodynamics,” [Phys. Rev. C78 \(2008\) 024902](#), [arXiv:0805.1756 \[nucl-th\]](#).
- [105] D. Molnar and P. Huovinen, “Dissipative effects from transport and viscous hydrodynamics,” [J. Phys. G35 \(2008\) 104125](#), [arXiv:0806.1367 \[nucl-th\]](#).
- [106] H. Song and U. W. Heinz, “Extracting the QGP viscosity from RHIC data – a status report from viscous hydrodynamics,” [arXiv:0812.4274 \[nucl-th\]](#).
- [107] M. Luzum and P. Romatschke, “Viscous Hydrodynamic Predictions for Nuclear Collisions at the LHC,” [arXiv:0901.4588 \[nucl-th\]](#).
- [108] H. Song and U. W. Heinz, “Viscous hydrodynamics with bulk viscosity – uncertainties from relaxation time and initial conditions,” [arXiv:0907.2262 \[nucl-th\]](#).
- [109] H. Song and U. W. Heinz, “Interplay of shear and bulk viscosity in generating flow in heavy-ion collisions,” [Phys. Rev. C81 \(2010\) 024905](#), [arXiv:0909.1549 \[nucl-th\]](#).
- [110] U. W. Heinz, “Early collective expansion: Relativistic hydrodynamics and the transport properties of QCD matter,” [arXiv:0901.4355 \[nucl-th\]](#).
- [111] P. Romatschke, “New Developments in Relativistic Viscous Hydrodynamics,” [arXiv:0902.3663 \[hep-ph\]](#).
- [112] D. A. Teaney, “Viscous Hydrodynamics and the Quark Gluon Plasma,” [arXiv:0905.2433 \[nucl-th\]](#).
- [113] C. Shen, U. W. Heinz, P. Huovinen, and H. Song, “Systematic parameter study of hadron spectra and elliptic flow from viscous hydrodynamic simulations of Au+Au collisions at $\sqrt{(s_{NN})} = 200$ GeV,” [arXiv:1010.1856 \[nucl-th\]](#).
- [114] H. Song, S. A. Bass, U. W. Heinz, T. Hirano, and C. Shen, “200 A GeV Au+Au collisions serve a nearly perfect quark- gluon liquid,” [arXiv:1011.2783 \[nucl-th\]](#).
- [115] S. A. Bass and A. Dumitru, “Dynamics of hot bulk QCD matter: From the quark-gluon plasma to hadronic freeze-out,” [Phys. Rev. C61 \(2000\) 064909](#), [arXiv:nucl-th/0001033](#).
- [116] D. Teaney, J. Lauret, and E. V. Shuryak, “A hydrodynamic description of heavy ion collisions at the SPS and RHIC,” [arXiv:nucl-th/0110037](#).

- [117] K. Rajagopal and N. Tripuraneni, “Bulk Viscosity and Cavitation in Boost-Invariant Hydrodynamic Expansion,” [arXiv:0908.1785 \[hep-ph\]](#).
- [118] U. W. Heinz, “‘RHIC serves the perfect fluid’: Hydrodynamic flow of the QGP,” [arXiv:nucl-th/0512051](#).
- [119] P. Huovinen and P. Petreczky, “QCD Equation of State and Hadron Resonance Gas,” [Nucl. Phys. **A837** \(2010\) 26–53](#), [arXiv:0912.2541 \[hep-ph\]](#).
- [120] J. E. Thomas, “Is an Ultra-Cold Strongly Interacting Fermi Gas a Perfect Fluid?,” [arXiv:0907.1040](#).
- [121] D. Molnar and M. Gyulassy, “Saturation of elliptic flow at RHIC: Results from the covariant elastic parton cascade model MPC,” [Nucl. Phys. **A697** \(2002\) 495–520](#), [arXiv:nucl-th/0104073](#).
- [122] **The ALICE** Collaboration, K. Aamodt *et al.*, “Elliptic flow of charged particles in Pb-Pb collisions at 2.76 TeV,” [arXiv:1011.3914 \[nucl-ex\]](#).
- [123] **STAR** Collaboration, J. Adams *et al.*, “Distributions of charged hadrons associated with high transverse momentum particles in p p and Au + Au collisions at $\sqrt{s(NN)} = 200$ -GeV,” [Phys. Rev. Lett. **95** \(2005\) 152301](#), [arXiv:nucl-ex/0501016](#).
- [124] T. A. Collaboration, “Observation of a Centrality-Dependent Dijet Asymmetry in Lead-Lead Collisions at $\sqrt{s(NN)} = 2.76$ TeV with the ATLAS Detector at the LHC,” [arXiv:1011.6182 \[hep-ex\]](#).
- [125] **ALICE** Collaboration, K. Aamodt *et al.*, “Suppression of Charged Particle Production at Large Transverse Momentum in Central Pb–Pb Collisions at $\sqrt{s_{NN}} = 2.76$ TeV,” [arXiv:1012.1004 \[nucl-ex\]](#).
- [126] **PHENIX** Collaboration, A. Adare *et al.*, “Suppression pattern of neutral pions at high transverse momentum in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV and constraints on medium transport coefficients,” [arXiv:0801.4020 \[nucl-ex\]](#).
- [127] **BRAHMS** Collaboration, I. Arsene *et al.*, “Transverse momentum spectra in Au + Au and d + Au collisions at $\sqrt{s(NN)} = 200$ -GeV and the pseudorapidity dependence of high $p(T)$ suppression,” [Phys. Rev. Lett. **91** \(2003\) 072305](#), [arXiv:nucl-ex/0307003](#).
- [128] **PHENIX** Collaboration, S. S. Adler *et al.*, “Absence of suppression in particle production at large transverse momentum in $\sqrt{s(NN)} = 200$ -GeV d + Au collisions,” [Phys. Rev. Lett. **91** \(2003\) 072303](#), [nucl-ex/0306021](#).
- [129] **PHOBOS** Collaboration, B. B. Back *et al.*, “Centrality dependence of charged hadron transverse momentum spectra in d + Au collisions at $\sqrt{s(NN)} = 200$ -GeV,” [Phys. Rev. Lett. **91** \(2003\) 072302](#), [arXiv:nucl-ex/0306025](#).

- [130] **STAR** Collaboration, J. Adams *et al.*, “Evidence from $d + Au$ measurements for final-state suppression of high $p(T)$ hadrons in $Au + Au$ collisions at RHIC,” *Phys. Rev. Lett.* **91** (2003) 072304, [nucl-ex/0306024](#).
- [131] **PHENIX** Collaboration, T. Isobe, “Systematic study of high- $p(T)$ direct photon production with the PHENIX experiment at RHIC,” *J. Phys.* **G34** (2007) S1015–1018, [arXiv:nucl-ex/0701040](#).
- [132] F. Arleo, “Hard pion and prompt photon at RHIC, from single to double inclusive production,” *JHEP* **09** (2006) 015, [arXiv:hep-ph/0601075](#).
- [133] **PHOBOS** Collaboration, B. B. Back *et al.*, “Centrality dependence of charged hadron transverse momentum spectra in $Au + Au$ collisions from $s(NN)^{1/2} = 62.4\text{-GeV}$ to 200-GeV ,” *Phys. Rev. Lett.* **94** (2005) 082304, [arXiv:nucl-ex/0405003](#).
- [134] **PHENIX** Collaboration, S. S. Adler *et al.*, “Common suppression pattern of η and π^0 mesons at high transverse momentum in $Au + Au$ collisions at $s(NN)^{1/2} = 200\text{-GeV}$,” *Phys. Rev. Lett.* **96** (2006) 202301, [arXiv:nucl-ex/0601037](#).
- [135] **PHENIX** Collaboration, S. S. Adler *et al.*, “High- $p(T)$ charged hadron suppression in $Au + Au$ collisions at $s(NN)^{1/2} = 200\text{-GeV}$,” *Phys. Rev.* **C69** (2004) 034910, [nucl-ex/0308006](#).
- [136] **PHENIX** Collaboration, A. Adare *et al.*, “Energy Loss and Flow of Heavy Quarks in $Au+Au$ Collisions at $\sqrt{s_{NN}} = 200\text{ GeV}$,” *Phys. Rev. Lett.* **98** (2007) 172301, [nucl-ex/0611018](#).
- [137] H. Stoecker, “Collective Flow signals the Quark Gluon Plasma,” *Nucl. Phys.* **A750** (2005) 121–147, [arXiv:nucl-th/0406018](#).
- [138] J. Casalderrey-Solana, E. V. Shuryak, and D. Teaney, “Conical flow induced by quenched QCD jets,” *J. Phys. Conf. Ser.* **27** (2005) 22–31, [arXiv:hep-ph/0411315](#).
- [139] M. Gyulassy and X.-n. Wang, “Multiple collisions and induced gluon Bremsstrahlung in QCD,” *Nucl. Phys.* **B420** (1994) 583–614, [arXiv:nucl-th/9306003](#).
- [140] R. Baier, Y. L. Dokshitzer, A. H. Mueller, S. Peigne, and D. Schiff, “Radiative energy loss and $p(T)$ -broadening of high energy partons in nuclei,” *Nucl. Phys.* **B484** (1997) 265–282, [arXiv:hep-ph/9608322](#).
- [141] B. G. Zakharov, “Radiative energy loss of high energy quarks in finite-size nuclear matter and quark-gluon plasma,” *JETP Lett.* **65** (1997) 615–620, [arXiv:hep-ph/9704255](#).
- [142] U. A. Wiedemann, “Gluon radiation off hard quarks in a nuclear environment: Opacity expansion,” *Nucl. Phys.* **B588** (2000) 303–344, [arXiv:hep-ph/0005129](#).
- [143] M. Gyulassy, P. Levai, and I. Vitev, “Reaction operator approach to non-Abelian energy loss,” *Nucl. Phys.* **B594** (2001) 371–419, [arXiv:nucl-th/0006010](#).

- [144] X.-f. Guo and X.-N. Wang, “Multiple Scattering, Parton Energy Loss and Modified Fragmentation Functions in Deeply Inelastic eA Scattering,” [Phys. Rev. Lett.](#) **85** (2000) 3591–3594, [arXiv:hep-ph/0005044](#).
- [145] X.-N. Wang and X.-f. Guo, “Multiple parton scattering in nuclei: Parton energy loss,” [Nucl. Phys.](#) **A696** (2001) 788–832, [arXiv:hep-ph/0102230](#).
- [146] Y. V. Kovchegov and A. H. Mueller, “Gluon production in current nucleus and nucleon nucleus collisions in a quasi-classical approximation,” [Nucl. Phys.](#) **B529** (1998) 451–479, [arXiv:hep-ph/9802440](#).
- [147] A. Kovner and U. A. Wiedemann, “Eikonal evolution and gluon radiation,” [Phys. Rev.](#) **D64** (2001) 114002, [arXiv:hep-ph/0106240](#).
- [148] H. Liu, K. Rajagopal, and U. A. Wiedemann, “Calculating the jet quenching parameter from AdS/CFT,” [Phys. Rev. Lett.](#) **97** (2006) 182301, [arXiv:hep-ph/0605178](#).
- [149] H. Liu, K. Rajagopal, and U. A. Wiedemann, “Wilson loops in heavy ion collisions and their calculation in AdS/CFT,” [JHEP](#) **03** (2007) 066, [arXiv:hep-ph/0612168](#).
- [150] J. Casalderrey-Solana and C. A. Salgado, “Introductory lectures on jet quenching in heavy ion collisions,” [Acta Phys. Polon.](#) **B38** (2007) 3731–3794, [arXiv:0712.3443 \[hep-ph\]](#).
- [151] F. D’Eramo, H. Liu, and K. Rajagopal, “Transverse Momentum Broadening and the Jet Quenching Parameter, Redux,” [Phys.Rev.](#) **D84** (2011) 065015, [arXiv:1006.1367 \[hep-ph\]](#).
- [152] B. Z. Kopeliovich, A. V. Tarasov, and A. Schafer, “Bremsstrahlung of a quark propagating through a nucleus,” [Phys. Rev.](#) **C59** (1999) 1609–1619, [arXiv:hep-ph/9808378](#).
- [153] U. A. Wiedemann, “Transverse dynamics of hard partons in nuclear media and the QCD dipole,” [Nucl. Phys.](#) **B582** (2000) 409–450, [arXiv:hep-ph/0003021](#).
- [154] R. Baier, D. Schiff, and B. G. Zakharov, “Energy loss in perturbative QCD,” [Ann. Rev. Nucl. Part. Sci.](#) **50** (2000) 37–69, [arXiv:hep-ph/0002198](#).
- [155] P. B. Arnold, G. D. Moore, and L. G. Yaffe, “Photon and Gluon Emission in Relativistic Plasmas,” [JHEP](#) **06** (2002) 030, [arXiv:hep-ph/0204343](#).
- [156] R. Baier, “Jet quenching,” [Nucl. Phys.](#) **A715** (2003) 209–218, [arXiv:hep-ph/0209038](#).
- [157] A. Kovner and U. A. Wiedemann, “Gluon radiation and parton energy loss,” [arXiv:hep-ph/0304151](#).
- [158] M. Gyulassy, I. Vitev, X.-N. Wang, and B.-W. Zhang, “Jet quenching and radiative energy loss in dense nuclear matter,” [arXiv:nucl-th/0302077](#).

- [159] P. Jacobs and X.-N. Wang, “Matter in extremis: Ultrarelativistic nuclear collisions at RHIC,” *Prog. Part. Nucl. Phys.* **54** (2005) 443–534, [hep-ph/0405125](#).
- [160] S. Wicks, W. Horowitz, M. Djordjevic, and M. Gyulassy, “Elastic, Inelastic, and Path Length Fluctuations in Jet Tomography,” *Nucl. Phys.* **A784** (2007) 426–442, [arXiv:nucl-th/0512076](#).
- [161] H. Zhang, J. F. Owens, E. Wang, and X.-N. Wang, “Dihadron Tomography of High-Energy Nuclear Collisions in NLO pQCD,” *Phys. Rev. Lett.* **98** (2007) 212301, [arXiv:nucl-th/0701045](#).
- [162] A. Majumder, C. Nonaka, and S. A. Bass, “Jet modification in three dimensional fluid dynamics at next-to-leading twist,” *Phys. Rev.* **C76** (2007) 041902, [arXiv:nucl-th/0703019](#).
- [163] G.-Y. Qin et al., “Radiative and Collisional Jet Energy Loss in the Quark- Gluon Plasma at RHIC,” *Phys. Rev. Lett.* **100** (2008) 072301, [arXiv:0710.0605](#) [[hep-ph](#)].
- [164] K. Zapp, G. Ingelman, J. Rathsman, J. Stachel, and U. A. Wiedemann, “A Monte Carlo Model for ‘Jet Quenching’,” *Eur. Phys. J.* **C60** (2009) 617–632, [arXiv:0804.3568](#) [[hep-ph](#)].
- [165] S. A. Bass et al., “Systematic Comparison of Jet Energy-Loss Schemes in a realistic hydrodynamic medium,” *Phys. Rev.* **C79** (2009) 024901, [arXiv:0808.0908](#) [[nucl-th](#)].
- [166] D. d’Enterria, “Jet quenching,” [arXiv:0902.2011](#) [[nucl-ex](#)].
- [167] N. Armesto, M. Cacciari, T. Hirano, J. L. Nagle, and C. A. Salgado, “Constraint fitting of experimental data with a jet quenching model embedded in a hydrodynamical bulk medium,” *J. Phys.* **G37** (2010) 025104, [arXiv:0907.0667](#) [[hep-ph](#)].
- [168] A. Accardi, F. Arleo, W. K. Brooks, D. D’Enterria, and V. Muccifora, “Parton Propagation and Fragmentation in QCD Matter,” *Riv. Nuovo Cim.* **32** (2010) 439–553, [arXiv:0907.3534](#) [[nucl-th](#)].
- [169] U. A. Wiedemann, “Jet Quenching in Heavy Ion Collisions,” [arXiv:0908.2306](#) [[hep-ph](#)].
- [170] I. Vitev and B.-W. Zhang, “Jet tomography of high-energy nucleus-nucleus collisions at next-to-leading order,” *Phys. Rev. Lett.* **104** (2010) 132001, [arXiv:0910.1090](#) [[hep-ph](#)].
- [171] W. A. Horowitz and B. A. Cole, “Systematic theoretical uncertainties in jet quenching due to gluon kinematics,” *Phys. Rev.* **C81** (2010) 024909, [arXiv:0910.1823](#) [[hep-ph](#)].
- [172] X.-F. Chen, C. Greiner, E. Wang, X.-N. Wang, and Z. Xu, “Bulk matter evolution and extraction of jet transport parameter in heavy-ion collisions at RHIC,” *Phys. Rev.* **C81** (2010) 064908, [arXiv:1002.1165](#) [[nucl-th](#)].

- [173] A. Majumder and M. Van Leeuwen, “The theory and phenomenology of perturbative QCD based jet quenching,” [arXiv:1002.2206 \[hep-ph\]](#).
- [174] S. Caron-Huot, “Loops and trees,” [arXiv:1007.3224 \[hep-ph\]](#).
- [175] R. Baier, Y. L. Dokshitzer, A. H. Mueller, and D. Schiff, “Radiative energy loss of high energy partons traversing an expanding QCD plasma,” [Phys. Rev. C58 \(1998\) 1706–1713](#), [arXiv:hep-ph/9803473](#).
- [176] C. A. Salgado and U. A. Wiedemann, “A dynamical scaling law for jet tomography,” [Phys. Rev. Lett. 89 \(2002\) 092303](#), [arXiv:hep-ph/0204221](#).
- [177] C. A. Salgado and U. A. Wiedemann, “Calculating quenching weights,” [Phys. Rev. D68 \(2003\) 014008](#), [arXiv:hep-ph/0302184](#).
- [178] R. Baier, Y. L. Dokshitzer, A. H. Mueller, S. Peigne, and D. Schiff, “Radiative energy loss of high energy quarks and gluons in a finite-volume quark-gluon plasma,” [Nucl. Phys. B483 \(1997\) 291–320](#), [arXiv:hep-ph/9607355](#).
- [179] R. Baier, Y. L. Dokshitzer, A. H. Mueller, and D. Schiff, “Quenching of hadron spectra in media,” [JHEP 09 \(2001\) 033](#), [arXiv:hep-ph/0106347](#).
- [180] **PHENIX** Collaboration, A. Adare et al., “Quantitative Constraints on the Opacity of Hot Partonic Matter from Semi-Inclusive Single High Transverse Momentum Pion Suppression in Au+Au collisions at $\sqrt{(s_{NN})} = 200$ GeV,” [Phys. Rev. C77 \(2008\) 064907](#), [arXiv:0801.1665 \[nucl-ex\]](#).
- [181] A. Dainese, C. Loizides, and G. Paic, “Leading-particle suppression in high energy nucleus nucleus collisions,” [Eur. Phys. J. C38 \(2005\) 461–474](#), [arXiv:hep-ph/0406201](#).
- [182] A. Peshier, “Collisional jet quenching becomes probable,” [Phys. Rev. C75 \(2007\) 034906](#), [arXiv:hep-ph/0607299](#).
- [183] R. Kolevatov and U. A. Wiedemann, “Varying and inverting the mass hierarchy in collisional energy loss,” [arXiv:0812.0270 \[hep-ph\]](#).
- [184] T. Matsui and H. Satz, “J/psi Suppression by Quark-Gluon Plasma Formation,” [Phys. Lett. B178 \(1986\) 416](#).
- [185] H. Satz, “Colour deconfinement and quarkonium binding,” [J. Phys. G32 \(2006\) R25](#), [arXiv:hep-ph/0512217](#).
- [186] **NA50** Collaboration, B. Alessandro et al., “A new measurement of J/psi suppression in Pb-Pb collisions at 158 GeV per nucleon,” [Eur. Phys. J. C39 \(2005\) 335–345](#), [arXiv:hep-ex/0412036](#).
- [187] **NA60** Collaboration, R. Arnaldi et al., “J/psi production in indium-indium collisions at 158- GeV/nucleon,” [Phys. Rev. Lett. 99 \(2007\) 132302](#).

- [188] **PHENIX** Collaboration, A. Adare *et al.*, “J/psi production vs centrality, transverse momentum, and rapidity in Au + Au collisions at $s(NN)^{1/2} = 200$ - GeV,” [Phys. Rev. Lett.](#) **98** (2007) 232301, [arXiv:nucl-ex/0611020](#).
- [189] T. A. Collaboration, “Measurement of the centrality dependence of J/ψ yields and observation of Z production in lead-lead collisions with the ATLAS detector at the LHC,” [arXiv:1012.5419](#) [hep-ex].
- [190] F. Karsch, D. Kharzeev, and H. Satz, “Sequential charmonium dissociation,” [Phys. Lett.](#) **B637** (2006) 75–80, [arXiv:hep-ph/0512239](#).
- [191] S. J. Brodsky and A. H. Mueller, “Using Nuclei to Probe Hadronization in QCD,” [Phys. Lett.](#) **B206** (1988) 685.
- [192] D. Kharzeev and H. Satz, “Charmonium interaction in nuclear matter,” [Phys. Lett.](#) **B356** (1995) 365–372, [arXiv:hep-ph/9504397](#).
- [193] **PHENIX** Collaboration, S. S. Adler *et al.*, “J/psi production and nuclear effects for d + Au and p + p collisions at $s(NN)^{1/2} = 200$ -GeV,” [Phys. Rev. Lett.](#) **96** (2006) 012304, [arXiv:nucl-ex/0507032](#).
- [194] R. L. Thews, M. Schroedter, and J. Rafelski, “Enhanced J/psi production in deconfined quark matter,” [Phys. Rev.](#) **C63** (2001) 054905, [arXiv:hep-ph/0007323](#).
- [195] S. Durr *et al.*, “Ab-Initio Determination of Light Hadron Masses,” [Science](#) **322** (2008) 1224–1227, [arXiv:0906.3599](#) [hep-lat].
- [196] G. Boyd *et al.*, “Thermodynamics of SU(3) Lattice Gauge Theory,” [Nucl. Phys.](#) **B469** (1996) 419–444, [arXiv:hep-lat/9602007](#).
- [197] Y. Aoki, G. Endrodi, Z. Fodor, S. D. Katz, and K. K. Szabo, “The order of the quantum chromodynamics transition predicted by the standard model of particle physics,” [Nature](#) **443** (2006) 675–678, [arXiv:hep-lat/0611014](#).
- [198] Z. Fodor, “QCD Thermodynamics,” [PoS LAT2007](#) (2007) 011, [arXiv:0711.0336](#) [hep-lat].
- [199] B. Bringoltz and M. Teper, “The Pressure of the SU(N) lattice gauge theory at large-N,” [Phys.Lett.](#) **B628** (2005) 113–124, [arXiv:hep-lat/0506034](#) [hep-lat].
- [200] S. Datta and S. Gupta, “Exploring the gluon(c) plasma,” [Nucl.Phys.](#) **A830** (2009) 749C–752C, [arXiv:0906.3929](#) [hep-lat].
- [201] M. Panero, “Thermodynamics of the QCD plasma and the large-N limit,” [Phys.Rev.Lett.](#) **103** (2009) 232001, [arXiv:0907.3719](#) [hep-lat].
- [202] **Wuppertal-Budapest** Collaboration, S. Borsanyi *et al.*, “Is there still any T_c mystery in lattice QCD? Results with physical masses in the continuum limit III,” [JHEP](#) **09** (2010) 073, [arXiv:1005.3508](#) [hep-lat].

- [203] U. Kraemmer and A. Rebhan, “Advances in perturbative thermal field theory,” [Rept. Prog. Phys.](#) **67** (2004) 351, [arXiv:hep-ph/0310337](#).
- [204] E. Braaten and A. Nieto, “Free Energy of QCD at High Temperature,” [Phys. Rev.](#) **D53** (1996) 3421–3437, [arXiv:hep-ph/9510408](#).
- [205] K. Kajantie, M. Laine, K. Rummukainen, and Y. Schroder, “The pressure of hot QCD up to $g^6 \ln(1/g)$,” [Phys. Rev.](#) **D67** (2003) 105008, [arXiv:hep-ph/0211321](#).
- [206] A. Hietanen, K. Kajantie, M. Laine, K. Rummukainen, and Y. Schroder, “Three-dimensional physics and the pressure of hot QCD,” [Phys. Rev.](#) **D79** (2009) 045018, [arXiv:0811.4664 \[hep-lat\]](#).
- [207] J. P. Blaizot, E. Iancu, and A. Rebhan, “The entropy of the QCD plasma,” [Phys. Rev. Lett.](#) **83** (1999) 2906–2909, [arXiv:hep-ph/9906340](#).
- [208] J. O. Andersen, E. Braaten, and M. Strickland, “Hard thermal loop resummation of the free energy of a hot gluon plasma,” [Phys.Rev.Lett.](#) **83** (1999) 2139–2142, [arXiv:hep-ph/9902327 \[hep-ph\]](#).
- [209] J. O. Andersen, E. Petitgirard, and M. Strickland, “Two loop HTL thermodynamics with quarks,” [Phys.Rev.](#) **D70** (2004) 045001, [arXiv:hep-ph/0302069 \[hep-ph\]](#).
- [210] J. O. Andersen, L. E. Leganger, M. Strickland, and N. Su, “NNLO hard-thermal-loop thermodynamics for QCD,” [arXiv:1009.4644 \[hep-ph\]](#).
- [211] A. Hart, M. Laine, and O. Philipsen, “Static correlation lengths in QCD at high temperatures and finite densities,” [Nucl. Phys.](#) **B586** (2000) 443–474, [arXiv:hep-ph/0004060](#).
- [212] V. Koch, A. Majumder, and J. Randrup, “Baryon-strangeness correlations: A diagnostic of strongly interacting matter,” [Phys. Rev. Lett.](#) **95** (2005) 182301, [arXiv:nucl-th/0505052](#).
- [213] M. Cheng et al., “Baryon Number, Strangeness and Electric Charge Fluctuations in QCD at High Temperature,” [Phys. Rev.](#) **D79** (2009) 074505, [arXiv:0811.1006 \[hep-lat\]](#).
- [214] H. B. Meyer, “A calculation of the shear viscosity in SU(3) gluodynamics,” [Phys. Rev.](#) **D76** (2007) 101701, [arXiv:0704.1801 \[hep-lat\]](#).
- [215] D. Teaney, “Finite temperature spectral densities of momentum and R- charge correlators in $N = 4$ Yang Mills theory,” [Phys. Rev.](#) **D74** (2006) 045025, [arXiv:hep-ph/0602044](#).
- [216] G. Aarts and J. M. Martinez Resco, “Transport coefficients, spectral functions and the lattice,” [JHEP](#) **04** (2002) 053, [arXiv:hep-ph/0203177](#).
- [217] H. B. Meyer, “Transport properties of the quark-gluon plasma from lattice QCD,” [Nucl. Phys.](#) **A830** (2009) 641c–648c, [arXiv:0907.4095 \[hep-lat\]](#).

- [218] E. V. Shuryak and I. Zahed, “Rethinking the properties of the quark gluon plasma at T approx. T(c),” [Phys. Rev. **C70** \(2004\) 021901](#), [arXiv:hep-ph/0307267](#).
- [219] A. Bazavov, P. Petreczky, and A. Velytsky, “Quarkonium at Finite Temperature,” [arXiv:0904.1748 \[hep-ph\]](#).
- [220] F. Karsch, M. T. Mehr, and H. Satz, “Color Screening and Deconfinement for Bound States of Heavy Quarks,” [Z. Phys. **C37** \(1988\) 617](#).
- [221] E. Eichten, K. Gottfried, T. Kinoshita, K. D. Lane, and T.-M. Yan, “Charmonium: Comparison with Experiment,” [Phys. Rev. **D21** \(1980\) 203](#).
- [222] A. Pineda and J. Soto, “Effective field theory for ultrasoft momenta in NRQCD and NRQED,” [Nucl. Phys. Proc. Suppl. **64** \(1998\) 428–432](#), [arXiv:hep-ph/9707481](#).
- [223] N. Brambilla, A. Pineda, J. Soto, and A. Vairo, “Potential NRQCD: An effective theory for heavy quarkonium,” [Nucl. Phys. **B566** \(2000\) 275](#), [arXiv:hep-ph/9907240](#).
- [224] N. Brambilla, J. Ghiglieri, A. Vairo, and P. Petreczky, “Static quark-antiquark pairs at finite temperature,” [Phys. Rev. **D78** \(2008\) 014017](#), [arXiv:0804.0993 \[hep-ph\]](#).
- [225] K. G. Wilson, “CONFINEMENT OF QUARKS,” [Phys. Rev. **D10** \(1974\) 2445–2459](#).
- [226] O. Philipsen, “Non-perturbative formulation of the static color octet potential,” [Phys. Lett. **B535** \(2002\) 138–144](#), [arXiv:hep-lat/0203018](#).
- [227] S. Nadkarni, “Nonabelian Debye Screening. 1. The Color Averaged Potential,” [Phys. Rev. **D33** \(1986\) 3738](#).
- [228] O. Kaczmarek, F. Karsch, P. Petreczky, and F. Zantow, “Heavy Quark Anti-Quark Free Energy and the Renormalized Polyakov Loop,” [Phys. Lett. **B543** \(2002\) 41–47](#), [arXiv:hep-lat/0207002](#).
- [229] O. Kaczmarek, F. Karsch, P. Petreczky, and F. Zantow, “Heavy quark free energies, potentials and the renormalized Polyakov loop,” [Nucl. Phys. Proc. Suppl. **129** \(2004\) 560–562](#), [arXiv:hep-lat/0309121](#).
- [230] O. Kaczmarek, F. Karsch, F. Zantow, and P. Petreczky, “Static quark anti-quark free energy and the running coupling at finite temperature,” [Phys. Rev. **D70** \(2004\) 074505](#), [arXiv:hep-lat/0406036](#).
- [231] O. Kaczmarek, “Screening at finite temperature and density,” [PoS **CPOD07** \(2007\) 043](#), [arXiv:0710.0498 \[hep-lat\]](#).
- [232] S. Necco and R. Sommer, “The N(f) = 0 heavy quark potential from short to intermediate distances,” [Nucl. Phys. **B622** \(2002\) 328–346](#), [arXiv:hep-lat/0108008](#).
- [233] C.-Y. Wong, “Heavy quarkonia in quark gluon plasma,” [Phys. Rev. **C72** \(2005\) 034906](#), [arXiv:hep-ph/0408020](#).

- [234] W. M. Alberico, A. Beraudo, A. De Pace, and A. Molinari, “Heavy quark bound states above $T(c)$,” [*Phys. Rev.* **D72** \(2005\) 114011](#), [arXiv:hep-ph/0507084](#).
- [235] E. V. Shuryak and I. Zahed, “Towards a theory of binary bound states in the quark gluon plasma,” [*Phys. Rev.* **D70** \(2004\) 054507](#), [arXiv:hep-ph/0403127](#).
- [236] A. Mocsy and P. Petreczky, “Quarkonia correlators above deconfinement,” [*Phys. Rev.* **D73** \(2006\) 074007](#), [arXiv:hep-ph/0512156](#).
- [237] W. M. Alberico, A. Beraudo, A. De Pace, and A. Molinari, “Potential models and lattice correlators for quarkonia at finite temperature,” [*Phys. Rev.* **D77** \(2008\) 017502](#), [arXiv:0706.2846 \[hep-ph\]](#).
- [238] G. Aarts and J. M. Martinez Resco, “Continuum and lattice meson spectral functions at nonzero momentum and high temperature,” [*Nucl. Phys.* **B726** \(2005\) 93–108](#), [arXiv:hep-lat/0507004](#).
- [239] A. Jakovac, P. Petreczky, K. Petrov, and A. Velytsky, “Quarkonium correlators and spectral functions at zero and finite temperature,” [*Phys. Rev.* **D75** \(2007\) 014506](#), [arXiv:hep-lat/0611017](#).
- [240] M. Asakawa, T. Hatsuda, and Y. Nakahara, “Maximum entropy analysis of the spectral functions in lattice QCD,” [*Prog. Part. Nucl. Phys.* **46** \(2001\) 459–508](#), [arXiv:hep-lat/0011040](#).
- [241] M. Asakawa and T. Hatsuda, “ J/ψ and η/c in the deconfined plasma from lattice QCD,” [*Phys. Rev. Lett.* **92** \(2004\) 012001](#), [arXiv:hep-lat/0308034](#).
- [242] S. Datta, F. Karsch, P. Petreczky, and I. Wetzorke, “Behavior of charmonium systems after deconfinement,” [*Phys. Rev.* **D69** \(2004\) 094507](#), [arXiv:hep-lat/0312037](#).
- [243] G. Aarts, C. Allton, M. B. Oktay, M. Peardon, and J.-I. Skullerud, “Charmonium at high temperature in two-flavor QCD,” [*Phys. Rev.* **D76** \(2007\) 094513](#), [arXiv:0705.2198 \[hep-lat\]](#).
- [244] H. T. Ding, O. Kaczmarek, F. Karsch, H. Satz, and W. Soeldner, “Charmonium correlators and spectral functions at finite temperature,” [arXiv:0910.3098 \[hep-lat\]](#).
- [245] T. Umeda, “A constant contribution in meson correlators at finite temperature,” [*Phys. Rev.* **D75** \(2007\) 094502](#), [arXiv:hep-lat/0701005](#).
- [246] S. Datta and P. Petreczky, “Zero mode contribution in quarkonium correlators and in-medium properties of heavy quarks,” [*J. Phys.* **G35** \(2008\) 104114](#), [arXiv:0805.1174 \[hep-lat\]](#).
- [247] A. Mocsy and P. Petreczky, “Can quarkonia survive deconfinement ?,” [*Phys. Rev.* **D77** \(2008\) 014501](#), [arXiv:0705.2559 \[hep-ph\]](#).
- [248] A. Mocsy and P. Petreczky, “Color Screening Melts Quarkonium,” [*Phys. Rev. Lett.* **99** \(2007\) 211602](#), [arXiv:0706.2183 \[hep-ph\]](#).

- [249] M. Laine, O. Philipsen, P. Romatschke, and M. Tassler, “Real-time static potential in hot QCD,” *JHEP* **03** (2007) 054, [arXiv:hep-ph/0611300](#).
- [250] A. Beraudo, J. P. Blaizot, and C. Ratti, “Real and imaginary-time $Q\bar{Q}$ correlators in a thermal medium,” *Nucl. Phys.* **A806** (2008) 312–338, [arXiv:0712.4394 \[nucl-th\]](#).
- [251] J. M. Maldacena, “The large N limit of superconformal field theories and supergravity,” *Adv. Theor. Math. Phys.* **2** (1998) 231–252, [arXiv:hep-th/9711200](#).
- [252] S. S. Gubser, I. R. Klebanov, and A. M. Polyakov, “Gauge theory correlators from non-critical string theory,” *Phys. Lett.* **B428** (1998) 105–114, [arXiv:hep-th/9802109](#).
- [253] E. Witten, “Anti-de Sitter space and holography,” *Adv. Theor. Math. Phys.* **2** (1998) 253–291, [arXiv:hep-th/9802150](#).
- [254] M. B. Green, J. H. Schwarz, and E. Witten, “Superstring Theory. Vol. 1 & 2,” . , Cambridge, Uk: Univ. Pr. (1987), Cambridge Monographs On Mathematical Physics.
- [255] D. Lust and S. Theisen, “Lectures on string theory,” *Lect. Notes Phys.* **346** (1989) 1–346.
- [256] J. Polchinski, “String theory. Vol. 1 & 2,” . , Cambridge, UK: Univ. Pr. (1998).
- [257] C. V. Johnson, “D-branes,” . , Cambridge, USA: Univ. Pr. (2003) 548 p.
- [258] B. Zwiebach, “A first course in string theory,” . , Cambridge, UK: Univ. Pr. (2004) 558 p.
- [259] E. Kiritsis, “String theory in a nutshell,” . , Princeton, USA: Univ. Pr. (2007) 588 p.
- [260] K. Becker, M. Becker, and J. H. Schwarz, “String theory and M-theory: A modern introduction,” . , Cambridge, UK: Cambridge Univ. Pr. (2007) 739 p.
- [261] M. Dine, “Supersymmetry and string theory: Beyond the standard model,” . , Cambridge, UK: Cambridge Univ. Pr. (2007) 515 p.
- [262] O. Aharony, S. S. Gubser, J. M. Maldacena, H. Ooguri, and Y. Oz, “Large N field theories, string theory and gravity,” *Phys. Rept.* **323** (2000) 183–386, [arXiv:hep-th/9905111](#).
- [263] E. D’Hoker and D. Z. Freedman, “Supersymmetric gauge theories and the AdS/CFT correspondence,” [arXiv:hep-th/0201253](#).
- [264] J. M. Maldacena, “Lectures on AdS/CFT,” [arXiv:hep-th/0309246](#).
- [265] J. Polchinski, “Introduction to Gauge/Gravity Duality,” [arXiv:1010.6134 \[hep-th\]](#).
- [266] A. M. Polyakov, “The wall of the cave,” *Int. J. Mod. Phys.* **A14** (1999) 645–658, [arXiv:hep-th/9809057](#).

- [267] G. 't Hooft, “Dimensional reduction in quantum gravity,” [arXiv:gr-qc/9310026](#) , gr-qc/9310026.
- [268] L. Susskind, “The World as a hologram,” [J. Math. Phys.](#) **36** (1995) 6377–6396, [arXiv:hep-th/9409089](#).
- [269] R. Bousso, “The holographic principle,” [Rev. Mod. Phys.](#) **74** (2002) 825–874, [arXiv:hep-th/0203101](#).
- [270] L. Susskind and E. Witten, “The holographic bound in anti-de Sitter space,” [arXiv:hep-th/9805114](#) , hep-th/9805114.
- [271] A. W. Peet and J. Polchinski, “UV/IR relations in AdS dynamics,” [Phys. Rev.](#) **D59** (1999) 065011, [arXiv:hep-th/9809022](#).
- [272] G. 't Hooft, “A Planar Diagram Theory for Strong Interactions,” [Nucl. Phys.](#) **B72** (1974) 461.
- [273] E. Witten, “Baryons in the $1/n$ Expansion,” [Nucl. Phys.](#) **B160** (1979) 57.
- [274] S. Coleman, “Aspects of symmetry,”. Cambridge, UK: Univ. Pr. (1985) 402 p.
- [275] A. V. Manohar, “Large N QCD,” [arXiv:hep-ph/9802419](#).
- [276] S. W. Hawking and G. F. R. Ellis, “The Large scale structure of space-time,”. , Cambridge University Press, Cambridge, 1973.
- [277] A. M. Polyakov, “Gauge Fields and Strings,”. Chur, Switzerland: Harwood (1987) 301 P. (Contemporary Concepts in Physics, 3).
- [278] R. Casalbuoni, J. Gomis, and G. Longhi, “The Relativistic Point Revisited in the Light of the String Model,” [Nuovo Cim.](#) **A24** (1974) 249.
- [279] M. B. Green and J. H. Schwarz, “Supersymmetrical String Theories,” [Phys. Lett.](#) **B109** (1982) 444–448.
- [280] J. H. Schwarz, “Superstring Theory,” [Phys. Rept.](#) **89** (1982) 223–322.
- [281] J. H. Schwarz and P. C. West, “Symmetries and Transformations of Chiral N=2 D=10 Supergravity,” [Phys. Lett.](#) **B126** (1983) 301.
- [282] J. H. Schwarz, “Covariant Field Equations of Chiral N=2 D=10 Supergravity,” [Nucl. Phys.](#) **B226** (1983) 269.
- [283] J. Polchinski, “Dirichlet-Branes and Ramond-Ramond Charges,” [Phys. Rev. Lett.](#) **75** (1995) 4724–4727, [arXiv:hep-th/9510017](#).
- [284] S. R. Coleman, “Classical Lumps and their Quantum Descendants,” [Subnucl. Ser.](#) **13** (1977) 297.
- [285] R. Jackiw, “Quantum Meaning of Classical Field Theory,” [Rev. Mod. Phys.](#) **49** (1977) 681–706.

- [286] R. Rajaraman, “Solitons and Instantons. An Introduction to Solitons and Instantons in Quantum Field Theory,” . , Amsterdam, Netherlands: North-holland (1982) 409p.
- [287] C. G. Callan, Jr., J. A. Harvey, and A. Strominger, “Worldbrane actions for string solitons,” [Nucl. Phys. **B367** \(1991\) 60–82.](#)
- [288] D. M. Kaplan and J. Michelson, “Zero Modes for the D=11 Membrane and Five-Brane,” [Phys. Rev. **D53** \(1996\) 3474–3476, arXiv:hep-th/9510053.](#)
- [289] T. Adawi, M. Cederwall, U. Gran, B. E. W. Nilsson, and B. Razaznejad, “Goldstone tensor modes,” [JHEP **02** \(1999\) 001, arXiv:hep-th/9811145.](#)
- [290] E. Witten, “Bound states of strings and p-branes,” [Nucl. Phys. **B460** \(1996\) 335–350, arXiv:hep-th/9510135.](#)
- [291] L. Brink, J. H. Schwarz, and J. Scherk, “Supersymmetric Yang-Mills Theories,” [Nucl. Phys. **B121** \(1977\) 77.](#)
- [292] F. Gliozzi, J. Scherk, and D. I. Olive, “Supersymmetry, Supergravity Theories and the Dual Spinor Model,” [Nucl. Phys. **B122** \(1977\) 253–290.](#)
- [293] S. Kovacs, “N=4 supersymmetric Yang-Mills theory and the AdS/SCFT correspondence,” [arXiv:hep-th/9908171.](#)
- [294] L. V. Avdeev, O. V. Tarasov, and A. A. Vladimirov, “Vanishing of the Three Loop Charge Renormalization Function in a Supersymmetric Gauge Theory,” [Phys. Lett. **B96** \(1980\) 94–96.](#)
- [295] M. T. Grisaru, M. Rocek, and W. Siegel, “Zero Three Loop beta Function in N=4 Superyang-Mills Theory,” [Phys. Rev. Lett. **45** \(1980\) 1063–1066.](#)
- [296] M. F. Sohnius and P. C. West, “Conformal Invariance in N=4 Supersymmetric Yang-Mills Theory,” [Phys. Lett. **B100** \(1981\) 245.](#)
- [297] W. E. Caswell and D. Zanon, “Zero Three Loop Beta Function in the N=4 Supersymmetric Yang-Mills Theory,” [Nucl. Phys. **B182** \(1981\) 125.](#)
- [298] P. S. Howe, K. S. Stelle, and P. K. Townsend, “Miraculous Ultraviolet Cancellations in Supersymmetry Made Manifest,” [Nucl. Phys. **B236** \(1984\) 125.](#)
- [299] S. Mandelstam, “Light Cone Superspace and the Ultraviolet Finiteness of the N=4 Model,” [Nucl. Phys. **B213** \(1983\) 149–168.](#)
- [300] L. Brink, O. Lindgren, and B. E. W. Nilsson, “The Ultraviolet Finiteness of the N=4 Yang-Mills Theory,” [Phys. Lett. **B123** \(1983\) 323.](#)
- [301] R. G. Leigh, “Dirac-Born-Infeld Action from Dirichlet Sigma Model,” [Mod. Phys. Lett. **A4** \(1989\) 2767.](#)
- [302] A. A. Tseytlin, “Born-Infeld action, supersymmetry and string theory,” [arXiv:hep-th/9908105.](#)

- [303] O. D. Andreev and A. A. Tseytlin, “Partition Function Representation for the Open Superstring Effective Action: Cancellation of Mobius Infinities and Derivative Corrections to Born-Infeld Lagrangian,” [Nucl. Phys. **B311** \(1988\) 205.](#)
- [304] Y. Kitazawa, “Effective Lagrangian for Open Superstring from Five Point Function,” [Nucl. Phys. **B289** \(1987\) 599.](#)
- [305] C. P. Bachas, P. Bain, and M. B. Green, “Curvature terms in D-brane actions and their M-theory origin,” [JHEP **05** \(1999\) 011,](#) [arXiv:hep-th/9903210.](#)
- [306] M. B. Green and M. Gutperle, “D-instanton induced interactions on a D3-brane,” [JHEP **02** \(2000\) 014,](#) [arXiv:hep-th/0002011.](#)
- [307] S. H. Shenker, “The Strength of nonperturbative effects in string theory,”. Presented at the Cargese Workshop on Random Surfaces, Quantum Gravity and Strings, Cargese, France, May 28 - Jun 1, 1990.
- [308] G. W. Gibbons and K.-i. Maeda, “Black Holes and Membranes in Higher Dimensional Theories with Dilaton Fields,” [Nucl. Phys. **B298** \(1988\) 741.](#)
- [309] D. Garfinkle, G. T. Horowitz, and A. Strominger, “Charged black holes in string theory,” [Phys. Rev. **D43** \(1991\) 3140.](#)
- [310] G. T. Horowitz and A. Strominger, “Black strings and P-branes,” [Nucl. Phys. **B360** \(1991\) 197–209.](#)
- [311] S. S. Gubser, I. R. Klebanov, and A. W. Peet, “Entropy and Temperature of Black 3-Branes,” [Phys. Rev. **D54** \(1996\) 3915–3919,](#) [arXiv:hep-th/9602135.](#)
- [312] E. Witten, “Anti-de Sitter space, thermal phase transition, and confinement in gauge theories,” [Adv. Theor. Math. Phys. **2** \(1998\) 505–532,](#) [arXiv:hep-th/9803131.](#)
- [313] V. Balasubramanian, P. Kraus, and A. E. Lawrence, “Bulk vs. boundary dynamics in anti-de Sitter spacetime,” [Phys. Rev. **D59** \(1999\) 046003,](#) [arXiv:hep-th/9805171.](#)
- [314] V. Balasubramanian, P. Kraus, A. E. Lawrence, and S. P. Trivedi, “Holographic probes of anti-de Sitter space-times,” [Phys. Rev. **D59** \(1999\) 104021,](#) [arXiv:hep-th/9808017.](#)
- [315] P. Breitenlohner and D. Z. Freedman, “Positive Energy in anti-De Sitter Backgrounds and Gauged Extended Supergravity,” [Phys. Lett. **B115** \(1982\) 197.](#)
- [316] P. Breitenlohner and D. Z. Freedman, “Stability in Gauged Extended Supergravity,” [Ann. Phys. **144** \(1982\) 249.](#)
- [317] L. Mezincescu and P. K. Townsend, “Stability at a Local Maximum in Higher Dimensional Anti-de Sitter Space and Applications to Supergravity,” [Ann. Phys. **160** \(1985\) 406.](#)
- [318] K. Skenderis, “Lecture notes on holographic renormalization,” [Class. Quant. Grav. **19** \(2002\) 5849–5876,](#) [arXiv:hep-th/0209067.](#)

- [319] I. R. Klebanov and E. Witten, “AdS/CFT correspondence and symmetry breaking,” [Nucl. Phys. **B556** \(1999\) 89–114](#), [arXiv:hep-th/9905104](#).
- [320] E. Witten, “Multi-trace operators, boundary conditions, and AdS/CFT correspondence,” [arXiv:hep-th/0112258](#).
- [321] M. Berkooz, A. Sever, and A. Shomer, “Double-trace deformations, boundary conditions and spacetime singularities,” [JHEP **05** \(2002\) 034](#), [arXiv:hep-th/0112264](#).
- [322] T. Hertog and K. Maeda, “Black holes with scalar hair and asymptotics in $N = 8$ supergravity,” [JHEP **07** \(2004\) 051](#), [arXiv:hep-th/0404261](#).
- [323] M. Henneaux, C. Martinez, R. Troncoso, and J. Zanelli, “Asymptotically anti-de Sitter spacetimes and scalar fields with a logarithmic branch,” [Phys. Rev. **D70** \(2004\) 044034](#), [arXiv:hep-th/0404236](#).
- [324] J. M. Maldacena, “Eternal black holes in Anti-de-Sitter,” [JHEP **04** \(2003\) 021](#), [arXiv:hep-th/0106112](#).
- [325] C. P. Herzog and D. T. Son, “Schwinger-Keldysh propagators from AdS/CFT correspondence,” [JHEP **03** \(2003\) 046](#), [arXiv:hep-th/0212072](#).
- [326] G. W. Gibbons and S. W. Hawking, “Action Integrals and Partition Functions in Quantum Gravity,” [Phys. Rev. **D15** \(1977\) 2752–2756](#).
- [327] A. Chamblin, R. Emparan, C. V. Johnson, and R. C. Myers, “Charged AdS black holes and catastrophic holography,” [Phys. Rev. **D60** \(1999\) 064018](#), [arXiv:hep-th/9902170](#).
- [328] A. Chamblin, R. Emparan, C. V. Johnson, and R. C. Myers, “Holography, thermodynamics and fluctuations of charged AdS black holes,” [Phys. Rev. **D60** \(1999\) 104026](#), [arXiv:hep-th/9904197](#).
- [329] S. S. Gubser, “Thermodynamics of spinning D3-branes,” [Nucl. Phys. **B551** \(1999\) 667–684](#), [arXiv:hep-th/9810225](#).
- [330] R.-G. Cai and K.-S. Soh, “Critical behavior in the rotating D-branes,” [Mod. Phys. Lett. **A14** \(1999\) 1895–1908](#), [arXiv:hep-th/9812121](#).
- [331] M. Cvetič and S. S. Gubser, “Phases of R-charged black holes, spinning branes and strongly coupled gauge theories,” [JHEP **04** \(1999\) 024](#), [arXiv:hep-th/9902195](#).
- [332] M. Cvetič and S. S. Gubser, “Thermodynamic Stability and Phases of General Spinning Branes,” [JHEP **07** \(1999\) 010](#), [arXiv:hep-th/9903132](#).
- [333] K.-Y. Kim, S.-J. Sin, and I. Zahed, “Dense hadronic matter in holographic QCD,” [arXiv:hep-th/0608046](#).
- [334] N. Horigome and Y. Tanii, “Holographic chiral phase transition with chemical potential,” [JHEP **01** \(2007\) 072](#), [arXiv:hep-th/0608198](#).

- [335] S. Kobayashi, D. Mateos, S. Matsuura, R. C. Myers, and R. M. Thomson, “Holographic phase transitions at finite baryon density,” *JHEP* **02** (2007) 016, [arXiv:hep-th/0611099](#).
- [336] D. Mateos, S. Matsuura, R. C. Myers, and R. M. Thomson, “Holographic phase transitions at finite chemical potential,” *JHEP* **11** (2007) 085, [arXiv:0709.1225 \[hep-th\]](#).
- [337] S. Nakamura, Y. Seo, S.-J. Sin, and K. P. Yogendran, “A new phase at finite quark density from AdS/CFT,” *J. Korean Phys. Soc.* **52** (2008) 1734–1739, [arXiv:hep-th/0611021](#).
- [338] A. Karch and A. O’Bannon, “Metallic AdS/CFT,” *JHEP* **09** (2007) 024, [arXiv:0705.3870 \[hep-th\]](#).
- [339] D. Yamada, “Sakai-Sugimoto Model at High Density,” *JHEP* **10** (2008) 020, [arXiv:0707.0101 \[hep-th\]](#).
- [340] O. Bergman, G. Lifschytz, and M. Lippert, “Holographic Nuclear Physics,” *JHEP* **11** (2007) 056, [arXiv:0708.0326 \[hep-th\]](#).
- [341] J. L. Davis, M. Gutperle, P. Kraus, and I. Sachs, “Stringy NJL and Gross-Neveu models at finite density and temperature,” *JHEP* **10** (2007) 049, [arXiv:0708.0589 \[hep-th\]](#).
- [342] M. Rozali, H.-H. Shieh, M. Van Raamsdonk, and J. Wu, “Cold Nuclear Matter In Holographic QCD,” *JHEP* **01** (2008) 053, [arXiv:0708.1322 \[hep-th\]](#).
- [343] S. Nakamura, Y. Seo, S.-J. Sin, and K. P. Yogendran, “Baryon-charge Chemical Potential in AdS/CFT,” *Prog. Theor. Phys.* **120** (2008) 51–76, [arXiv:0708.2818 \[hep-th\]](#).
- [344] K. Ghoroku, M. Ishihara, and A. Nakamura, “D3/D7 holographic Gauge theory and Chemical potential,” *Phys. Rev.* **D76** (2007) 124006, [arXiv:0708.3706 \[hep-th\]](#).
- [345] A. Karch and A. O’Bannon, “Holographic Thermodynamics at Finite Baryon Density: Some Exact Results,” *JHEP* **11** (2007) 074, [arXiv:0709.0570 \[hep-th\]](#).
- [346] T. Sakai and S. Sugimoto, “Low energy hadron physics in holographic QCD,” *Prog. Theor. Phys.* **113** (2005) 843–882, [arXiv:hep-th/0412141](#).
- [347] T. Sakai and S. Sugimoto, “More on a holographic dual of QCD,” *Prog. Theor. Phys.* **114** (2005) 1083–1118, [arXiv:hep-th/0507073](#).
- [348] O. Aharony, J. Sonnenschein, and S. Yankielowicz, “A holographic model of deconfinement and chiral symmetry restoration,” *Annals Phys.* **322** (2007) 1420–1443, [arXiv:hep-th/0604161](#).
- [349] A. Parnachev and D. A. Sahakyan, “Photoemission with Chemical Potential from QCD Gravity Dual,” *Nucl. Phys.* **B768** (2007) 177–192, [arXiv:hep-th/0610247](#).

- [350] J. Polchinski and M. J. Strassler, “The string dual of a confining four-dimensional gauge theory,” [arXiv:hep-th/0003136](#).
- [351] I. R. Klebanov and M. J. Strassler, “Supergravity and a confining gauge theory: Duality cascades and chiSB-resolution of naked singularities,” *JHEP* **08** (2000) 052, [arXiv:hep-th/0007191](#).
- [352] J. M. Maldacena and C. Nunez, “Towards the large N limit of pure N = 1 super Yang Mills,” *Phys. Rev. Lett.* **86** (2001) 588–591, [arXiv:hep-th/0008001](#).
- [353] O. Aharony, “The non-AdS/non-CFT correspondence, or three different paths to QCD,” [hep-th/0212193](#) (2002) , [arXiv:hep-th/0212193](#).
- [354] M. J. Strassler, “The duality cascade,” [hep-th/0505153](#) (2005) , [arXiv:hep-th/0505153](#).
- [355] N. Itzhaki, J. M. Maldacena, J. Sonnenschein, and S. Yankielowicz, “Supergravity and the large N limit of theories with sixteen supercharges,” *Phys. Rev.* **D58** (1998) 046004, [arXiv:hep-th/9802042](#).
- [356] M. Henningson and K. Skenderis, “The holographic Weyl anomaly,” *JHEP* **07** (1998) 023, [arXiv:hep-th/9806087](#).
- [357] V. Balasubramanian and P. Kraus, “A stress tensor for anti-de Sitter gravity,” *Commun. Math. Phys.* **208** (1999) 413–428, [arXiv:hep-th/9902121](#).
- [358] R. C. Myers, “Stress tensors and Casimir energies in the AdS/CFT correspondence,” *Phys. Rev.* **D60** (1999) 046002, [arXiv:hep-th/9903203](#).
- [359] R. Emparan, C. V. Johnson, and R. C. Myers, “Surface terms as counterterms in the AdS/CFT correspondence,” *Phys. Rev.* **D60** (1999) 104001, [arXiv:hep-th/9903238](#).
- [360] P. Kraus, F. Larsen, and R. Siebelink, “The gravitational action in asymptotically AdS and flat spacetimes,” *Nucl. Phys.* **B563** (1999) 259–278, [arXiv:hep-th/9906127](#).
- [361] S. de Haro, S. N. Solodukhin, and K. Skenderis, “Holographic reconstruction of spacetime and renormalization in the AdS/CFT correspondence,” *Commun. Math. Phys.* **217** (2001) 595–622, [arXiv:hep-th/0002230](#).
- [362] M. Bianchi, D. Z. Freedman, and K. Skenderis, “Holographic Renormalization,” *Nucl. Phys.* **B631** (2002) 159–194, [arXiv:hep-th/0112119](#).
- [363] L. D. Landau and E. M. Lifshitz, “Course Theoretical Physics. Vol. 2: The Classical Theory of Fields,”. , Elsevier (1975).
- [364] L. D. Landau and E. M. Lifshitz, “Textbook on Theoretical Physics. Vol. 1: Mechanics,”. , Elsevier (1976).
- [365] D. Z. Freedman, S. D. Mathur, A. Matusis, and L. Rastelli, “Correlation functions in the CFT(d)/AdS($d + 1$) correspondence,” *Nucl. Phys.* **B546** (1999) 96–118, [arXiv:hep-th/9804058](#).

- [366] H. Liu and A. A. Tseytlin, “D = 4 super Yang-Mills, D = 5 gauged supergravity, and D = 4 conformal supergravity,” [Nucl. Phys. **B533** \(1998\) 88–108](#), [arXiv:hep-th/9804083](#).
- [367] G. Chalmers, H. Nastase, K. Schalm, and R. Siebelink, “R-current correlators in N = 4 super Yang-Mills theory from anti-de Sitter supergravity,” [Nucl. Phys. **B540** \(1999\) 247–270](#), [arXiv:hep-th/9805105](#).
- [368] D. T. Son and A. O. Starinets, “Minkowski-space correlators in AdS/CFT correspondence: Recipe and applications,” [JHEP **09** \(2002\) 042](#), [arXiv:hep-th/0205051](#).
- [369] N. Iqbal and H. Liu, “Universality of the hydrodynamic limit in AdS/CFT and the membrane paradigm,” [Phys.Rev. **D79** \(2009\) 025023](#), [arXiv:0809.3808 \[hep-th\]](#).
- [370] N. Iqbal and H. Liu, “Real-time response in AdS/CFT with application to spinors,” [Fortsch. Phys. **57** \(2009\) 367–384](#), [arXiv:0903.2596 \[hep-th\]](#).
- [371] T. Banks, M. R. Douglas, G. T. Horowitz, and E. J. Martinec, “AdS dynamics from conformal field theory,” [arXiv:hep-th/9808016](#).
- [372] J. M. Maldacena, “Wilson loops in large N field theories,” [Phys. Rev. Lett. **80** \(1998\) 4859–4862](#), [arXiv:hep-th/9803002](#).
- [373] S.-J. Rey and J.-T. Yee, “Macroscopic strings as heavy quarks in large N gauge theory and anti-de Sitter supergravity,” [Eur. Phys. J. **C22** \(2001\) 379–394](#), [arXiv:hep-th/9803001](#).
- [374] L. F. Alday and J. Maldacena, “Comments on gluon scattering amplitudes via AdS/CFT,” [JHEP **0711** \(2007\) 068](#), [arXiv:0710.1060 \[hep-th\]](#).
- [375] N. Drukker, D. J. Gross, and H. Ooguri, “Wilson loops and minimal surfaces,” [Phys.Rev. **D60** \(1999\) 125006](#), [arXiv:hep-th/9904191 \[hep-th\]](#).
- [376] N. Drukker, D. J. Gross, and A. A. Tseytlin, “Green-Schwarz string in AdS(5) x S**5: Semiclassical partition function,” [JHEP **0004** \(2000\) 021](#), [arXiv:hep-th/0001204 \[hep-th\]](#).
- [377] J. K. Erickson, G. W. Semenoff, R. J. Szabo, and K. Zarembo, “Static potential in N = 4 supersymmetric Yang-Mills theory,” [Phys. Rev. **D61** \(2000\) 105006](#), [arXiv:hep-th/9911088](#).
- [378] S.-J. Rey, S. Theisen, and J.-T. Yee, “Wilson-Polyakov loop at finite temperature in large N gauge theory and anti-de Sitter supergravity,” [Nucl. Phys. **B527** \(1998\) 171–186](#), [arXiv:hep-th/9803135](#).
- [379] A. Brandhuber, N. Itzhaki, J. Sonnenschein, and S. Yankielowicz, “Wilson loops in the large N limit at finite temperature,” [Phys. Lett. **B434** \(1998\) 36–40](#), [arXiv:hep-th/9803137](#).

- [380] D. Bak, A. Karch, and L. G. Yaffe, “Debye screening in strongly coupled $N=4$ supersymmetric Yang-Mills plasma,” [JHEP **08** \(2007\) 049](#), [arXiv:0705.0994 \[hep-th\]](#).
- [381] J. Sonnenschein, “Stringy confining Wilson loops,” [arXiv:hep-th/0009146 \[hep-th\]](#).
- [382] S. W. Hawking and D. N. Page, “Thermodynamics of Black Holes in anti-De Sitter Space,” [Commun. Math. Phys. **87** \(1983\) 577](#).
- [383] O. Aharony, A. Fayyazuddin, and J. M. Maldacena, “The large N limit of $N = 2,1$ field theories from three- branes in F-theory,” [JHEP **07** \(1998\) 013](#), [arXiv:hep-th/9806159](#).
- [384] A. Karch and L. Randall, “Open and closed string interpretation of SUSY CFT’s on branes with boundaries,” [JHEP **06** \(2001\) 063](#), [arXiv:hep-th/0105132](#).
- [385] A. Karch and E. Katz, “Adding flavor to AdS/CFT,” [JHEP **06** \(2002\) 043](#), [arXiv:hep-th/0205236](#).
- [386] T. Nishioka and T. Takayanagi, “Free Yang-Mills vs. Toric Sasaki-Einstein,” [Phys. Rev. **D76** \(2007\) 044004](#), [arXiv:hep-th/0702194](#).
- [387] S. S. Gubser, I. R. Klebanov, and A. A. Tseytlin, “Coupling constant dependence in the thermodynamics of $N = 4$ supersymmetric Yang-Mills theory,” [Nucl. Phys. **B534** \(1998\) 202–222](#), [arXiv:hep-th/9805156](#).
- [388] J. G. Russo, “New compactifications of supergravities and large N QCD,” [Nucl. Phys. **B543** \(1999\) 183–197](#), [arXiv:hep-th/9808117](#).
- [389] P. Kraus, F. Larsen, and S. P. Trivedi, “The Coulomb branch of gauge theory from rotating branes,” [JHEP **03** \(1999\) 003](#), [arXiv:hep-th/9811120](#).
- [390] D. T. Son and A. O. Starinets, “Hydrodynamics of R-charged black holes,” [JHEP **03** \(2006\) 052](#), [arXiv:hep-th/0601157](#).
- [391] C. Sasaki, B. Friman, and K. Redlich, “Susceptibilities and the Phase Structure of a Chiral Model with Polyakov Loops,” [Phys.Rev. **D75** \(2007\) 074013](#), [arXiv:hep-ph/0611147 \[hep-ph\]](#).
- [392] M. Bluhm and B. Kampfer, “Flavor Diagonal and Off-Diagonal Susceptibilities in a Quasiparticle Model of the Quark-Gluon Plasma,” [Phys.Rev. **D77** \(2008\) 114016](#), [arXiv:0801.4147 \[hep-ph\]](#).
- [393] P. Petreczky, “Lattice QCD at finite temperature : Present status,” [Nucl.Phys. **A830** \(2009\) 11C–18C](#), [arXiv:0908.1917 \[hep-ph\]](#).
- [394] P. Kovtun, D. T. Son, and A. O. Starinets, “Holography and hydrodynamics: Diffusion on stretched horizons,” [JHEP **10** \(2003\) 064](#), [arXiv:hep-th/0309213](#).

- [395] A. Buchel and J. T. Liu, “Universality of the shear viscosity in supergravity,” [*Phys.Rev.Lett.* **93** \(2004\) 090602](#), [arXiv:hep-th/0311175 \[hep-th\]](#).
- [396] P. Kovtun, D. Son, and A. Starinets, “Viscosity in strongly interacting quantum field theories from black hole physics,” [*Phys.Rev.Lett.* **94** \(2005\) 111601](#), [arXiv:hep-th/0405231 \[hep-th\]](#). An Essay submitted to 2004 Gravity Research Foundation competition.
- [397] A. Buchel, “On universality of stress-energy tensor correlation functions in supergravity,” [*Phys.Lett.* **B609** \(2005\) 392–401](#), [arXiv:hep-th/0408095 \[hep-th\]](#).
- [398] G. Policastro, D. Son, and A. Starinets, “The Shear viscosity of strongly coupled N=4 supersymmetric Yang-Mills plasma,” [*Phys.Rev.Lett.* **87** \(2001\) 081601](#), [arXiv:hep-th/0104066 \[hep-th\]](#).
- [399] O. Saremi, “Shear waves, sound waves on a shimmering horizon,” [arXiv:hep-th/0703170](#).
- [400] A. O. Starinets, “Quasinormal spectrum and the black hole membrane paradigm,” [*Phys. Lett.* **B670** \(2009\) 442–445](#), [arXiv:0806.3797 \[hep-th\]](#).
- [401] M. Fujita, “Non-equilibrium thermodynamics near the horizon and holography,” [*JHEP* **10** \(2008\) 031](#), [arXiv:0712.2289 \[hep-th\]](#).
- [402] A. Buchel, J. T. Liu, and A. O. Starinets, “Coupling constant dependence of the shear viscosity in N=4 supersymmetric Yang-Mills theory,” [*Nucl. Phys.* **B707** \(2005\) 56–68](#), [arXiv:hep-th/0406264](#).
- [403] A. Buchel, “Shear viscosity of boost invariant plasma at finite coupling,” [*Nucl. Phys.* **B802** \(2008\) 281–306](#), [arXiv:0801.4421 \[hep-th\]](#).
- [404] A. Buchel, “Resolving disagreement for η/s in a CFT plasma at finite coupling,” [*Nucl. Phys.* **B803** \(2008\) 166–170](#), [arXiv:0805.2683 \[hep-th\]](#).
- [405] R. C. Myers, M. F. Paulos, and A. Sinha, “Quantum corrections to η/s ,” [*Phys. Rev.* **D79** \(2009\) 041901](#), [arXiv:0806.2156 \[hep-th\]](#).
- [406] A. Buchel, R. C. Myers, M. F. Paulos, and A. Sinha, “Universal holographic hydrodynamics at finite coupling,” [*Phys. Lett.* **B669** \(2008\) 364–370](#), [arXiv:0808.1837 \[hep-th\]](#).
- [407] P. B. Arnold, G. D. Moore, and L. G. Yaffe, “Transport coefficients in high temperature gauge theories. II: Beyond leading log,” [*JHEP* **05** \(2003\) 051](#), [arXiv:hep-ph/0302165](#).
- [408] S. C. Huot, S. Jeon, and G. D. Moore, “Shear viscosity in weakly coupled N = 4 super Yang-Mills theory compared to QCD,” [*Phys.Rev.Lett.* **98** \(2007\) 172303](#), [arXiv:hep-ph/0608062 \[hep-ph\]](#).

- [409] M. Brigante, H. Liu, R. C. Myers, S. Shenker, and S. Yaida, “Viscosity Bound Violation in Higher Derivative Gravity,” [*Phys. Rev.* **D77** \(2008\) 126006](#), [arXiv:0712.0805 \[hep-th\]](#).
- [410] Y. Kats and P. Petrov, “Effect of curvature squared corrections in AdS on the viscosity of the dual gauge theory,” [*JHEP* **01** \(2009\) 044](#), [arXiv:0712.0743 \[hep-th\]](#).
- [411] R. Brustein and A. J. M. Medved, “The ratio of shear viscosity to entropy density in generalized theories of gravity,” [*Phys. Rev.* **D79** \(2009\) 021901](#), [arXiv:0808.3498 \[hep-th\]](#).
- [412] R. Brustein and A. J. M. Medved, “The shear diffusion coefficient for generalized theories of gravity,” [*Phys. Lett.* **B671** \(2009\) 119–122](#), [arXiv:0810.2193 \[hep-th\]](#).
- [413] R.-G. Cai, Z.-Y. Nie, and Y.-W. Sun, “Shear Viscosity from Effective Couplings of Gravitons,” [*Phys. Rev.* **D78** \(2008\) 126007](#), [arXiv:0811.1665 \[hep-th\]](#).
- [414] X.-H. Ge, S.-J. Sin, S.-F. Wu, and G.-H. Yang, “Shear viscosity and instability from third order Lovelock gravity,” [*Phys. Rev.* **D80** \(2009\) 104019](#), [arXiv:0905.2675 \[hep-th\]](#).
- [415] S. S. Pal, “ η/s at finite coupling,” [*Phys. Rev.* **D81** \(2010\) 045005](#), [arXiv:0910.0101 \[hep-th\]](#).
- [416] F.-W. Shu, “The Quantum Viscosity Bound In Lovelock Gravity,” [*Phys. Lett.* **B685** \(2010\) 325–328](#), [arXiv:0910.0607 \[hep-th\]](#).
- [417] A. J. Amsel and D. Gorbonos, “The Weak Gravity Conjecture and the Viscosity Bound with Six-Derivative Corrections,” [*JHEP* **11** \(2010\) 033](#), [arXiv:1005.4718 \[hep-th\]](#).
- [418] R. Brustein and A. J. M. Medved, “Unitarity constraints on the ratio of shear viscosity to entropy density in higher derivative gravity,” [arXiv:1005.5274 \[hep-th\]](#).
- [419] X. O. Camanho, J. D. Edelstein, and M. F. Paulos, “Lovelock theories, holography and the fate of the viscosity bound,” [arXiv:1010.1682 \[hep-th\]](#).
- [420] M. Brigante, H. Liu, R. C. Myers, S. Shenker, and S. Yaida, “The Viscosity Bound and Causality Violation,” [*Phys.Rev.Lett.* **100** \(2008\) 191601](#), [arXiv:0802.3318 \[hep-th\]](#).
- [421] A. Buchel, R. C. Myers, and A. Sinha, “Beyond $\eta/s = 1/4\pi$,” [*JHEP* **03** \(2009\) 084](#), [arXiv:0812.2521 \[hep-th\]](#).
- [422] P. Benincasa, A. Buchel, and A. O. Starinets, “Sound waves in strongly coupled non-conformal gauge theory plasma,” [*Nucl. Phys.* **B733** \(2006\) 160–187](#), [arXiv:hep-th/0507026](#).

- [423] A. Buchel, “Transport properties of cascading gauge theories,” [Phys. Rev.](#) **D72** (2005) 106002, [arXiv:hep-th/0509083](#).
- [424] A. Buchel, “Bulk viscosity of gauge theory plasma at strong coupling,” [Phys. Lett.](#) **B663** (2008) 286–289, [arXiv:0708.3459](#) [[hep-th](#)].
- [425] J. Mas and J. Tarrío, “Hydrodynamics from the Dp-brane,” [JHEP](#) **05** (2007) 036, [arXiv:hep-th/0703093](#).
- [426] A. Buchel and C. Pagnutti, “Bulk viscosity of N=2* plasma,” [Nucl. Phys.](#) **B816** (2009) 62–72, [arXiv:0812.3623](#) [[hep-th](#)].
- [427] S. S. Gubser, A. Nellore, S. S. Pufu, and F. D. Rocha, “Thermodynamics and bulk viscosity of approximate black hole duals to finite temperature quantum chromodynamics,” [Phys. Rev. Lett.](#) **101** (2008) 131601, [arXiv:0804.1950](#) [[hep-th](#)].
- [428] S. S. Gubser, S. S. Pufu, and F. D. Rocha, “Bulk viscosity of strongly coupled plasmas with holographic duals,” [JHEP](#) **08** (2008) 085, [arXiv:0806.0407](#) [[hep-th](#)].
- [429] G. Policastro, D. T. Son, and A. O. Starinets, “From AdS / CFT correspondence to hydrodynamics,” [JHEP](#) **0209** (2002) 043, [arXiv:hep-th/0205052](#) [[hep-th](#)].
- [430] G. D. Moore and K. A. Sohrawi, “Kubo Formulae for Second-Order Hydrodynamic Coefficients,” [arXiv:1007.5333](#) [[hep-ph](#)].
- [431] A. Buchel and M. Paulos, “Relaxation time of a CFT plasma at finite coupling,” [Nucl. Phys.](#) **B805** (2008) 59–71, [arXiv:0806.0788](#) [[hep-th](#)].
- [432] F. Bigazzi, A. L. Cotrone, and J. Tarrío, “Hydrodynamics of fundamental matter,” [JHEP](#) **1002** (2010) 083, [arXiv:0912.3256](#) [[hep-th](#)].
- [433] F. Bigazzi and A. L. Cotrone, “An elementary stringy estimate of transport coefficients of large temperature QCD,” [JHEP](#) **1008** (2010) 128, [arXiv:1006.4634](#) [[hep-ph](#)].
- [434] M. A. York and G. D. Moore, “Second order hydrodynamic coefficients from kinetic theory,” [Phys. Rev.](#) **D79** (2009) 054011, [arXiv:0811.0729](#) [[hep-ph](#)].
- [435] I. Kanitscheider and K. Skenderis, “Universal hydrodynamics of non-conformal branes,” [JHEP](#) **04** (2009) 062, [arXiv:0901.1487](#) [[hep-th](#)].
- [436] E. Braaten and R. D. Pisarski, “Resummation and Gauge Invariance of the Gluon Damping Rate in Hot QCD,” [Phys. Rev. Lett.](#) **64** (1990) 1338.
- [437] U. W. Heinz, “Kinetic Theory for Nonabelian Plasmas,” [Phys. Rev. Lett.](#) **51** (1983) 351.
- [438] E. Braaten and R. D. Pisarski, “Soft Amplitudes in Hot Gauge Theories: A General Analysis,” [Nucl. Phys.](#) **B337** (1990) 569.

- [439] J. Frenkel and J. C. Taylor, “High Temperature Limit of Thermal QCD,” [Nucl. Phys. B334 \(1990\) 199](#).
- [440] J. C. Taylor and S. M. H. Wong, “IN QCD The Effective Action of Hard Thermal Loops in QCD,” [Nucl. Phys. B346 \(1990\) 115–128](#).
- [441] P. F. Kelly, Q. Liu, C. Lucchesi, and C. Manuel, “Classical transport theory and hard thermal loops in the quark - gluon plasma,” [Phys. Rev. D50 \(1994\) 4209–4218](#), [arXiv:hep-ph/9406285](#).
- [442] J. P. Blaizot and E. Iancu, “Kinetic equations for long wavelength excitations of the quark - gluon plasma,” [Phys. Rev. Lett. 70 \(1993\) 3376–3379](#), [arXiv:hep-ph/9301236](#).
- [443] E. Calzetta and B. L. Hu, “Nonequilibrium Quantum Fields: Closed Time Path Effective Action, Wigner Function and Boltzmann Equation,” [Phys. Rev. D37 \(1988\) 2878](#).
- [444] E. A. Calzetta, B. L. Hu, and S. A. Ramsey, “Hydrodynamic transport functions from quantum kinetic field theory,” [Phys. Rev. D61 \(2000\) 125013](#), [arXiv:hep-ph/9910334](#).
- [445] J.-P. Blaizot and E. Iancu, “A Boltzmann equation for the QCD plasma,” [Nucl.Phys. B557 \(1999\) 183–236](#), [arXiv:hep-ph/9903389 \[hep-ph\]](#).
- [446] J.-P. Blaizot and E. Iancu, “The quark-gluon plasma: Collective dynamics and hard thermal loops,” [Phys. Rept. 359 \(2002\) 355–528](#), [arXiv:hep-ph/0101103](#).
- [447] P. B. Arnold, G. D. Moore, and L. G. Yaffe, “Effective kinetic theory for high temperature gauge theories,” [JHEP 01 \(2003\) 030](#), [arXiv:hep-ph/0209353](#).
- [448] E. M. Lifshitz and L. P. Pitaevskii, [Physical Kinetics](#). Butterworth Heinemann, 1981.
- [449] P. B. Arnold, G. D. Moore, and L. G. Yaffe, “Transport coefficients in high temperature gauge theories: (I) Leading-log results,” [JHEP 11 \(2000\) 001](#), [arXiv:hep-ph/0010177](#).
- [450] J. Hong and D. Teaney, “Spectral densities for hot QCD plasmas in a leading log approximation,” [Phys. Rev. C82 \(2010\) 044908](#), [arXiv:1003.0699 \[nucl-th\]](#).
- [451] P. Kovtun and A. Starinets, “Thermal spectral functions of strongly coupled $N = 4$ supersymmetric Yang-Mills theory,” [Phys. Rev. Lett. 96 \(2006\) 131601](#), [arXiv:hep-th/0602059](#).
- [452] C. P. Herzog, A. Karch, P. Kovtun, C. Kozcaz, and L. G. Yaffe, “Energy loss of a heavy quark moving through $N = 4$ supersymmetric Yang-Mills plasma,” [JHEP 07 \(2006\) 013](#), [arXiv:hep-th/0605158](#).
- [453] S. S. Gubser, “Drag force in AdS/CFT,” [Phys. Rev. D74 \(2006\) 126005](#), [arXiv:hep-th/0605182](#).

- [454] C. P. Herzog, “Energy loss of heavy quarks from asymptotically AdS geometries,” [JHEP 09 \(2006\) 032](#), [arXiv:hep-th/0605191](#).
- [455] E. Caceres and A. Guijosa, “Drag force in charged $N = 4$ SYM plasma,” [JHEP 11 \(2006\) 077](#), [arXiv:hep-th/0605235](#).
- [456] E. Caceres and A. Guijosa, “On Drag Forces and Jet Quenching in Strongly Coupled Plasmas,” [JHEP 12 \(2006\) 068](#), [arXiv:hep-th/0606134](#).
- [457] T. Matsuo, D. Tomino, and W.-Y. Wen, “Drag force in SYM plasma with B field from AdS/CFT,” [JHEP 10 \(2006\) 055](#), [arXiv:hep-th/0607178](#).
- [458] E. Nakano, S. Teraguchi, and W.-Y. Wen, “Drag Force, Jet Quenching, and AdS/QCD,” [Phys. Rev. D75 \(2007\) 085016](#), [arXiv:hep-ph/0608274](#).
- [459] P. Talavera, “Drag force in a string model dual to large- N QCD,” [JHEP 01 \(2007\) 086](#), [arXiv:hep-th/0610179](#).
- [460] S. S. Gubser, “Comparing the drag force on heavy quarks in $N = 4$ super- Yang-Mills theory and QCD,” [Phys. Rev. D76 \(2007\) 126003](#), [arXiv:hep-th/0611272](#).
- [461] U. Gursoy, E. Kiritsis, G. Michalogiorgakis, and F. Nitti, “Thermal Transport and Drag Force in Improved Holographic QCD,” [JHEP 12 \(2009\) 056](#), [arXiv:0906.1890 \[hep-ph\]](#).
- [462] C. Hoyos-Badajoz, “Drag and jet quenching of heavy quarks in a strongly coupled $N=2^*$ plasma,” [JHEP 09 \(2009\) 068](#), [arXiv:0907.5036 \[hep-th\]](#).
- [463] G. Bertoldi, F. Bigazzi, A. Cotrone, and J. D. Edelstein, “Holography and unquenched quark-gluon plasmas,” [Phys.Rev. D76 \(2007\) 065007](#), [arXiv:hep-th/0702225 \[hep-th\]](#).
- [464] F. Bigazzi, A. L. Cotrone, J. Mas, A. Paredes, A. V. Ramallo, et al., “D3-D7 Quark-Gluon Plasmas,” [JHEP 0911 \(2009\) 117](#), [arXiv:0909.2865 \[hep-th\]](#).
- [465] G. D. Moore and D. Teaney, “How much do heavy quarks thermalize in a heavy ion collision?,” [Phys. Rev. C71 \(2005\) 064904](#), [arXiv:hep-ph/0412346](#).
- [466] H. Liu, K. Rajagopal, and Y. Shi, “Robustness and Infrared Sensitivity of Various Observables in the Application of AdS/CFT to Heavy Ion Collisions,” [JHEP 08 \(2008\) 048](#), [arXiv:0803.3214 \[hep-ph\]](#).
- [467] J. Casalderrey-Solana and D. Teaney, “Transverse momentum broadening of a fast quark in a $N = 4$ Yang Mills plasma,” [JHEP 04 \(2007\) 039](#), [arXiv:hep-th/0701123](#).
- [468] A. Mikhailov, “Nonlinear waves in AdS/CFT correspondence,” [arXiv:hep-th/0305196](#).
- [469] K. B. Fadafan, H. Liu, K. Rajagopal, and U. A. Wiedemann, “Stirring Strongly Coupled Plasma,” [Eur. Phys. J. C61 \(2009\) 553–567](#), [arXiv:0809.2869 \[hep-ph\]](#).

- [470] M. Chernicoff and A. Guijosa, “Acceleration, Energy Loss and Screening in Strongly-Coupled Gauge Theories,” [JHEP](#) **0806** (2008) 005, [arXiv:0803.3070 \[hep-th\]](#).
- [471] M. Chernicoff, J. Garcia, and A. Guijosa, “Generalized Lorentz-Dirac Equation for a Strongly-Coupled Gauge Theory,” [Phys.Rev.Lett.](#) **102** (2009) 241601, [arXiv:0903.2047 \[hep-th\]](#).
- [472] M. Chernicoff, J. Garcia, and A. Guijosa, “A Tail of a Quark in N=4 SYM,” [JHEP](#) **0909** (2009) 080, [arXiv:0906.1592 \[hep-th\]](#).
- [473] J. Casalderrey-Solana and D. Teaney, “Heavy quark diffusion in strongly coupled N = 4 Yang-Mills,” [Phys. Rev.](#) **D74** (2006) 085012, [arXiv:hep-ph/0605199](#).
- [474] S. S. Gubser, “Momentum fluctuations of heavy quarks in the gauge-string duality,” [Nucl. Phys.](#) **B790** (2008) 175–199, [arXiv:hep-th/0612143](#).
- [475] J. Casalderrey-Solana, K.-Y. Kim, and D. Teaney, “Stochastic String Motion Above and Below the World Sheet Horizon,” [JHEP](#) **12** (2009) 066, [arXiv:0908.1470 \[hep-th\]](#).
- [476] F. Dominguez, C. Marquet, A. H. Mueller, B. Wu, and B.-W. Xiao, “Comparing energy loss and p_{\perp} -broadening in perturbative QCD with strong coupling $\mathcal{N} = 4$ SYM theory,” [Nucl. Phys.](#) **A811** (2008) 197–222, [arXiv:0803.3234 \[nucl-th\]](#).
- [477] J. de Boer, V. E. Hubeny, M. Rangamani, and M. Shigemori, “Brownian motion in AdS/CFT,” [JHEP](#) **07** (2009) 094, [arXiv:0812.5112 \[hep-th\]](#).
- [478] G. C. Giecold, E. Iancu, and A. H. Mueller, “Stochastic trailing string and Langevin dynamics from AdS/CFT,” [JHEP](#) **07** (2009) 033, [arXiv:0903.1840 \[hep-th\]](#).
- [479] J.-J. Peng and S.-Q. Wu, “Covariant anomalies and Hawking radiation from charged rotating black strings in anti-de Sitter spacetimes,” [Phys. Lett.](#) **B661** (2008) 300–306, [arXiv:0801.0185 \[hep-th\]](#).
- [480] U. Gursoy, E. Kiritsis, L. Mazzanti, and F. Nitti, “Langevin diffusion of heavy quarks in non-conformal holographic backgrounds,” [arXiv:1006.3261 \[hep-th\]](#).
- [481] B.-W. Xiao, “On the exact solution of the accelerating string in AdS_5 space,” [Phys. Lett.](#) **B665** (2008) 173–177, [arXiv:0804.1343 \[hep-th\]](#).
- [482] E. Caceres, M. Chernicoff, A. Guijosa, and J. F. Pedraza, “Quantum Fluctuations and the Unruh Effect in Strongly- Coupled Conformal Field Theories,” [JHEP](#) **06** (2010) 078, [arXiv:1003.5332 \[hep-th\]](#).
- [483] M. Le Bellac, [Thermal Field Theory](#). Cambridge Monographs on Mathematical Physics, 1996.
- [484] D. T. Son and D. Teaney, “Thermal Noise and Stochastic Strings in AdS/CFT,” [JHEP](#) **07** (2009) 021, [arXiv:0901.2338 \[hep-th\]](#).

- [485] H. van Hees, V. Greco, and R. Rapp, “Heavy-quark probes of the quark-gluon plasma at RHIC,” [*Phys. Rev.* **C73** \(2006\) 034913](#), [arXiv:nucl-th/0508055](#).
- [486] W. A. Horowitz and M. Gyulassy, “Heavy quark jet tomography of Pb + Pb at LHC: AdS/CFT drag or pQCD energy loss?,” [*Phys. Lett.* **B666** \(2008\) 320–323](#), [arXiv:0706.2336 \[nucl-th\]](#).
- [487] H. van Hees, M. Mannarelli, V. Greco, and R. Rapp, “Nonperturbative Heavy-Quark Diffusion in the Quark-Gluon Plasma,” [*Phys. Rev. Lett.* **100** \(2008\) 192301](#), [arXiv:0709.2884 \[hep-ph\]](#).
- [488] Y. Akamatsu, T. Hatsuda, and T. Hirano, “Heavy Quark Diffusion with Relativistic Langevin Dynamics in the Quark-Gluon Fluid,” [*Phys. Rev.* **C79** \(2009\) 054907](#), [arXiv:0809.1499 \[hep-ph\]](#).
- [489] A. Beraudo, A. De Pace, W. M. Alberico, and A. Molinari, “Transport properties and Langevin dynamics of heavy quarks and quarkonia in the Quark Gluon Plasma,” [*Nucl. Phys.* **A831** \(2009\) 59–90](#), [arXiv:0902.0741 \[hep-ph\]](#).
- [490] R. Rapp and H. van Hees, “Heavy Quarks in the Quark-Gluon Plasma,” [arXiv:0903.1096 \[hep-ph\]](#).
- [491] J. Noronha, M. Gyulassy, and G. Torrieri, “Constraints on AdS/CFT Gravity Dual Models of Heavy Ion Collisions,” [arXiv:0906.4099 \[hep-ph\]](#).
- [492] J. Noronha, M. Gyulassy, and G. Torrieri, “Conformal Holography of Bulk Elliptic Flow and Heavy Quark Quenching in Relativistic Heavy Ion Collisions,” [arXiv:1009.2286 \[nucl-th\]](#).
- [493] W. M. Alberico *et al.*, “Heavy quark dynamics in the QGP: R_{AA} and v_2 from RHIC to LHC,” [arXiv:1011.0400 \[hep-ph\]](#).
- [494] **PHENIX** Collaboration, S. S. Adler *et al.*, “Nuclear modification of electron spectra and implications for heavy quark energy loss in Au + Au collisions at $s(NN)^{1/2} = 200$ -GeV,” [*Phys. Rev. Lett.* **96** \(2006\) 032301](#), [arXiv:nucl-ex/0510047](#).
- [495] **STAR** Collaboration, B. I. Abelev *et al.*, “Transverse momentum and centrality dependence of high- p_T non-photonically electron suppression in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV,” [*Phys. Rev. Lett.* **98** \(2007\) 192301](#), [arXiv:nucl-ex/0607012](#).
- [496] S. Caron-Huot and G. D. Moore, “Heavy quark diffusion in perturbative QCD at next-to-leading order,” [*Phys. Rev. Lett.* **100** \(2008\) 052301](#), [arXiv:0708.4232 \[hep-ph\]](#).
- [497] S. Caron-Huot and G. D. Moore, “Heavy quark diffusion in QCD and $N = 4$ SYM at next-to-leading order,” [*JHEP* **02** \(2008\) 081](#), [arXiv:0801.2173 \[hep-ph\]](#).
- [498] D. H. Rischke, H. Stoecker, and W. Greiner, “Flow in Conical Shock Waves: a Signal for the Deconfinement Transition?,” [*Phys. Rev.* **D42** \(1990\) 2283–2292](#).

- [499] J. Ruppert and B. Muller, “Waking the colored plasma,” [Phys. Lett. **B618** \(2005\) 123–130](#), [arXiv:hep-ph/0503158](#).
- [500] L. M. Satarov, H. Stoecker, and I. N. Mishustin, “Mach shocks induced by partonic jets in expanding quark- gluon plasma,” [Phys. Lett. **B627** \(2005\) 64–70](#), [arXiv:hep-ph/0505245](#).
- [501] T. Renk and J. Ruppert, “Mach cones in an evolving medium,” [Phys. Rev. **C73** \(2006\) 011901](#), [arXiv:hep-ph/0509036](#).
- [502] J. Casalderrey-Solana and E. V. Shuryak, “Conical flow in a medium with variable speed of sound,” [arXiv:hep-ph/0511263](#).
- [503] J. Casalderrey-Solana, E. V. Shuryak, and D. Teaney, “Hydrodynamic flow from fast particles,” [arXiv:hep-ph/0602183](#).
- [504] T. Renk and J. Ruppert, “The rapidity structure of Mach cones and other large angle correlations in heavy-ion collisions,” [Phys. Lett. **B646** \(2007\) 19–23](#), [arXiv:hep-ph/0605330](#).
- [505] J. J. Friess, S. S. Gubser, G. Michalogiorgakis, and S. S. Pufu, “The stress tensor of a quark moving through $N = 4$ thermal plasma,” [Phys. Rev. **D75** \(2007\) 106003](#), [arXiv:hep-th/0607022](#).
- [506] A. Yarom, “The high momentum behavior of a quark wake,” [Phys. Rev. **D75** \(2007\) 125010](#), [arXiv:hep-th/0702164](#).
- [507] S. S. Gubser and S. S. Pufu, “Master field treatment of metric perturbations sourced by the trailing string,” [Nucl. Phys. **B790** \(2008\) 42–71](#), [arXiv:hep-th/0703090](#).
- [508] A. Yarom, “On the energy deposited by a quark moving in an $N=4$ SYM plasma,” [Phys. Rev. **D75** \(2007\) 105023](#), [arXiv:hep-th/0703095](#).
- [509] P. M. Chesler and L. G. Yaffe, “The wake of a quark moving through a strongly-coupled $\mathcal{N} = 4$ supersymmetric Yang-Mills plasma,” [Phys. Rev. Lett. **99** \(2007\) 152001](#), [arXiv:0706.0368 \[hep-th\]](#).
- [510] S. S. Gubser, S. S. Pufu, and A. Yarom, “Energy disturbances due to a moving quark from gauge- string duality,” [JHEP **09** \(2007\) 108](#), [arXiv:0706.0213 \[hep-th\]](#).
- [511] S. S. Gubser, S. S. Pufu, and A. Yarom, “Sonic booms and diffusion wakes generated by a heavy quark in thermal AdS/CFT,” [Phys. Rev. Lett. **100** \(2008\) 012301](#), [arXiv:0706.4307 \[hep-th\]](#).
- [512] S. S. Gubser and A. Yarom, “Universality of the diffusion wake in the gauge-string duality,” [Phys. Rev. **D77** \(2008\) 066007](#), [arXiv:0709.1089 \[hep-th\]](#).
- [513] P. M. Chesler and L. G. Yaffe, “The stress-energy tensor of a quark moving through a strongly-coupled $N=4$ supersymmetric Yang-Mills plasma: comparing hydrodynamics and AdS/CFT,” [Phys. Rev. **D78** \(2008\) 045013](#), [arXiv:0712.0050 \[hep-th\]](#).

- [514] J. Noronha, G. Torrieri, and M. Gyulassy, “Near Zone Navier-Stokes Analysis of Heavy Quark Jet Quenching in an $\mathcal{N} = 4$ SYM Plasma,” [Phys. Rev. **C78** \(2008\) 024903](#), [arXiv:0712.1053 \[hep-ph\]](#).
- [515] R. B. Neufeld, B. Muller, and J. Ruppert, “Sonic Mach Cones Induced by Fast Partons in a Perturbative Quark-Gluon Plasma,” [Phys. Rev. **C78** \(2008\) 041901](#), [arXiv:0802.2254 \[hep-ph\]](#).
- [516] S. S. Gubser and A. Yarom, “Linearized hydrodynamics from probe-sources in the gauge-string duality,” [Nucl. Phys. **B813** \(2009\) 188–219](#), [arXiv:0803.0081 \[hep-th\]](#).
- [517] R. B. Neufeld, “Fast Partons as a Source of Energy and Momentum in a Perturbative Quark-Gluon Plasma,” [Phys. Rev. **D78** \(2008\) 085015](#), [arXiv:0805.0385 \[hep-ph\]](#).
- [518] J. Noronha, M. Gyulassy, and G. Torrieri, “Di-Jet Conical Correlations Associated with Heavy Quark Jets in anti-de Sitter Space/Conformal Field Theory Correspondence,” [Phys. Rev. Lett. **102** \(2009\) 102301](#), [arXiv:0807.1038 \[hep-ph\]](#).
- [519] R. B. Neufeld, “Mach cones in the quark-gluon plasma: Viscosity, speed of sound, and effects of finite source structure,” [Phys. Rev. **C79** \(2009\) 054909](#), [arXiv:0807.2996 \[nucl-th\]](#).
- [520] M. Gyulassy, J. Noronha, and G. Torrieri, “Conical Di-jet Correlations from a Chromo-Viscous Neck in AdS/CFT,” [arXiv:0807.2235 \[hep-ph\]](#).
- [521] B. Betz, M. Gyulassy, J. Noronha, and G. Torrieri, “Anomalous Conical Di-jet Correlations in pQCD vs AdS/CFT,” [Phys. Lett. **B675** \(2009\) 340–346](#), [arXiv:0807.4526 \[hep-ph\]](#).
- [522] B. Betz et al., “Universality of the Diffusion Wake from Stopped and Punch- Through Jets in Heavy-Ion Collisions,” [Phys. Rev. **C79** \(2009\) 034902](#), [arXiv:0812.4401 \[nucl-th\]](#).
- [523] R. B. Neufeld and B. Muller, “The sound produced by a fast parton in the quark-gluon plasma is a ‘crescendo’,” [Phys. Rev. Lett. **103** \(2009\) 042301](#), [arXiv:0902.2950 \[nucl-th\]](#).
- [524] R. B. Neufeld and T. Renk, “The Mach cone signal and energy deposition scenarios in linearized hydrodynamics,” [Phys. Rev. **C82** \(2010\) 044903](#), [arXiv:1001.5068 \[nucl-th\]](#).
- [525] B. Betz, J. Noronha, G. Torrieri, M. Gyulassy, and D. H. Rischke, “Universal Flow-Driven Conical Emission in Ultrarelativistic Heavy-Ion Collisions,” [arXiv:1005.5461 \[nucl-th\]](#).
- [526] B. Alver and G. Roland, “Collision geometry fluctuations and triangular flow in heavy-ion collisions,” [Phys. Rev. **C81** \(2010\) 054905](#), [arXiv:1003.0194 \[nucl-th\]](#).

- [527] K. Dusling et al., “Quarkonium transport in thermal AdS/CFT,” [JHEP 10 \(2008\) 098](#), [arXiv:0808.0957 \[hep-th\]](#).
- [528] S. S. Gubser, S. S. Pufu, and A. Yarom, “Shock waves from heavy-quark mesons in AdS/CFT,” [JHEP 07 \(2008\) 108](#), [arXiv:0711.1415 \[hep-th\]](#).
- [529] S. S. Gubser, D. R. Gulotta, S. S. Pufu, and F. D. Rocha, “Gluon energy loss in the gauge-string duality,” [JHEP 10 \(2008\) 052](#), [arXiv:0803.1470 \[hep-th\]](#).
- [530] P. M. Chesler, K. Jensen, and A. Karch, “Jets in strongly-coupled $N = 4$ super Yang-Mills theory,” [Phys. Rev. D79 \(2009\) 025021](#), [arXiv:0804.3110 \[hep-th\]](#).
- [531] P. M. Chesler, K. Jensen, A. Karch, and L. G. Yaffe, “Light quark energy loss in strongly-coupled $N = 4$ supersymmetric Yang-Mills plasma,” [Phys. Rev. D79 \(2009\) 125015](#), [arXiv:0810.1985 \[hep-th\]](#).
- [532] P. Arnold and D. Vaman, “Jet quenching in hot strongly coupled gauge theories revisited: 3-point correlators with gauge-gravity duality,” [JHEP 1010 \(2010\) 099](#), [arXiv:1008.4023 \[hep-th\]](#).
- [533] P. Arnold and D. Vaman, “Jet quenching in hot strongly coupled gauge theories simplified,” [JHEP 1104 \(2011\) 027](#), [arXiv:1101.2689 \[hep-th\]](#).
- [534] Z.-t. Liang, X.-N. Wang, and J. Zhou, “The Transverse-momentum-dependent Parton Distribution Function and Jet Transport in Medium,” [Phys. Rev. D77 \(2008\) 125010](#), [arXiv:0801.0434 \[hep-ph\]](#).
- [535] C. W. Bauer, S. Fleming, and M. E. Luke, “Summing Sudakov logarithms in $B \rightarrow X \gamma$ in effective field theory,” [Phys. Rev. D63 \(2000\) 014006](#), [arXiv:hep-ph/0005275](#).
- [536] C. W. Bauer, S. Fleming, D. Pirjol, and I. W. Stewart, “An effective field theory for collinear and soft gluons: Heavy to light decays,” [Phys. Rev. D63 \(2001\) 114020](#), [arXiv:hep-ph/0011336](#).
- [537] C. W. Bauer and I. W. Stewart, “Invariant operators in collinear effective theory,” [Phys. Lett. B516 \(2001\) 134–142](#), [arXiv:hep-ph/0107001](#).
- [538] C. W. Bauer, D. Pirjol, and I. W. Stewart, “Soft-Collinear Factorization in Effective Field Theory,” [Phys. Rev. D65 \(2002\) 054022](#), [arXiv:hep-ph/0109045](#).
- [539] C. W. Bauer, S. Fleming, D. Pirjol, I. Z. Rothstein, and I. W. Stewart, “Hard scattering factorization from effective field theory,” [Phys. Rev. D66 \(2002\) 014017](#), [arXiv:hep-ph/0202088](#).
- [540] A. Idilbi and A. Majumder, “Extending Soft-Collinear-Effective-Theory to describe hard jets in dense QCD media,” [Phys. Rev. D80 \(2009\) 054022](#), [arXiv:0808.1087 \[hep-ph\]](#).
- [541] A. Buchel, “On jet quenching parameters in strongly coupled non-conformal gauge theories,” [Phys. Rev. D74 \(2006\) 046006](#), [arXiv:hep-th/0605178](#).

- [542] F.-L. Lin and T. Matsuo, “Jet quenching parameter in medium with chemical potential from AdS/CFT,” [Phys. Lett. B641 \(2006\) 45–49](#), [arXiv:hep-th/0606136](#).
- [543] S. D. Avramis and K. Sfetsos, “Supergravity and the jet quenching parameter in the presence of R-charge densities,” [JHEP 01 \(2007\) 065](#), [arXiv:hep-th/0606190](#).
- [544] N. Armesto, J. D. Edelstein, and J. Mas, “Jet quenching at finite ’t Hooft coupling and chemical potential from AdS/CFT,” [JHEP 09 \(2006\) 039](#), [arXiv:hep-ph/0606245](#).
- [545] P. C. Argyres, M. Edalati, and J. F. Vazquez-Poritz, “Spacelike strings and jet quenching from a Wilson loop,” [JHEP 04 \(2007\) 049](#), [arXiv:hep-th/0612157](#).
- [546] P. C. Argyres, M. Edalati, and J. F. Vazquez-Poritz, “Lightlike Wilson loops from AdS/CFT,” [JHEP 03 \(2008\) 071](#), [arXiv:0801.4594 \[hep-th\]](#).
- [547] K. Skenderis and B. C. van Rees, “Real-time gauge/gravity duality,” [Phys. Rev. Lett. 101 \(2008\) 081601](#), [arXiv:0805.0150 \[hep-th\]](#).
- [548] K. Skenderis and B. C. van Rees, “Real-time gauge/gravity duality: Prescription, Renormalization and Examples,” [JHEP 05 \(2009\) 085](#), [arXiv:0812.2909 \[hep-th\]](#).
- [549] B. C. van Rees, “Real-time gauge/gravity duality and ingoing boundary conditions,” [Nucl. Phys. Proc. Suppl. 192-193 \(2009\) 193–196](#), [arXiv:0902.4010 \[hep-th\]](#).
- [550] E. Barnes, D. Vaman, C. Wu, and P. Arnold, “Real-time finite-temperature correlators from AdS/CFT,” [Phys.Rev. D82 \(2010\) 025019](#), [arXiv:1004.1179 \[hep-th\]](#).
- [551] D. M. Hofman and J. Maldacena, “Conformal collider physics: Energy and charge correlations,” [JHEP 05 \(2008\) 012](#), [arXiv:0803.1467 \[hep-th\]](#).
- [552] Y. Hatta, E. Iancu, and A. H. Mueller, “Jet evolution in the N=4 SYM plasma at strong coupling,” [JHEP 05 \(2008\) 037](#), [arXiv:0803.2481 \[hep-th\]](#).
- [553] C. Athanasiou, P. M. Chesler, H. Liu, D. Nickel, and K. Rajagopal, “Synchrotron radiation in strongly coupled conformal field theories,” [Phys. Rev. D81 \(2010\) 126001](#), [arXiv:1001.3880 \[hep-th\]](#).
- [554] V. E. Hubeny, “Relativistic Beaming in AdS/CFT,” [arXiv:1011.1270 \[hep-th\]](#).
- [555] Y. Hatta, E. Iancu, A. H. Mueller, and D. N. Triantafyllopoulos, “Aspects of the UV/IR correspondence : energy broadening and string fluctuations,” [arXiv:1011.3763 \[hep-th\]](#).
- [556] K. Peeters, J. Sonnenschein, and M. Zamaklar, “Holographic melting and related properties of mesons in a quark gluon plasma,” [Phys. Rev. D74 \(2006\) 106008](#), [arXiv:hep-th/0606195](#).
- [557] H. Liu, K. Rajagopal, and U. A. Wiedemann, “An AdS/CFT calculation of screening in a hot wind,” [Phys. Rev. Lett. 98 \(2007\) 182301](#), [arXiv:hep-ph/0607062](#).

- [558] M. Chernicoff, J. A. Garcia, and A. Guijosa, “The energy of a moving quark-antiquark pair in an $N = 4$ SYM plasma,” *JHEP* **09** (2006) 068, [arXiv:hep-th/0607089](#).
- [559] E. Caceres, M. Natsuume, and T. Okamura, “Screening length in plasma winds,” *JHEP* **10** (2006) 011, [arXiv:hep-th/0607233](#).
- [560] P. C. Argyres, M. Edalati, and J. F. Vazquez-Poritz, “No-drag string configurations for steadily moving quark- antiquark pairs in a thermal bath,” *JHEP* **01** (2007) 105, [arXiv:hep-th/0608118](#).
- [561] S. D. Avramis, K. Sfetsos, and D. Zoakos, “On the velocity and chemical-potential dependence of the heavy-quark interaction in $N = 4$ SYM plasmas,” *Phys. Rev.* **D75** (2007) 025009, [arXiv:hep-th/0609079](#).
- [562] J. J. Friess, S. S. Gubser, G. Michalogiorgakis, and S. S. Pufu, “Stability of strings binding heavy-quark mesons,” *JHEP* **04** (2007) 079, [arXiv:hep-th/0609137](#).
- [563] M. Chernicoff and A. Guijosa, “Energy Loss of Gluons, Baryons and k-Quarks in an $N=4$ SYM Plasma,” *JHEP* **02** (2007) 084, [arXiv:hep-th/0611155](#).
- [564] S. D. Avramis, K. Sfetsos, and K. Siampos, “Stability of strings dual to flux tubes between static quarks in $N=4$ SYM,” *Nucl. Phys.* **B769** (2007) 44–78, [arXiv:hep-th/0612139](#).
- [565] D. Mateos, R. C. Myers, and R. M. Thomson, “Thermodynamics of the brane,” *JHEP* **05** (2007) 067, [arXiv:hep-th/0701132](#).
- [566] M. Natsuume and T. Okamura, “Screening length and the direction of plasma winds,” *JHEP* **09** (2007) 039, [arXiv:0706.0086 \[hep-th\]](#).
- [567] S. D. Avramis, K. Sfetsos, and K. Siampos, “Stability of string configurations dual to quarkonium states in AdS/CFT,” *Nucl. Phys.* **B793** (2008) 1–33, [arXiv:0706.2655 \[hep-th\]](#).
- [568] Q. J. Ejaz, T. Faulkner, H. Liu, K. Rajagopal, and U. A. Wiedemann, “A limiting velocity for quarkonium propagation in a strongly coupled plasma via AdS/CFT,” *JHEP* **04** (2008) 089, [arXiv:0712.0590 \[hep-th\]](#).
- [569] C. Athanasiou, H. Liu, and K. Rajagopal, “Velocity Dependence of Baryon Screening in a Hot Strongly Coupled Plasma,” *JHEP* **05** (2008) 083, [arXiv:0801.1117 \[hep-th\]](#).
- [570] R. C. Myers and A. Sinha, “The fast life of holographic mesons,” *JHEP* **06** (2008) 052, [arXiv:0804.2168 \[hep-th\]](#).
- [571] T. Faulkner and H. Liu, “Meson widths from string worldsheet instantons,” *Phys. Lett.* **B673** (2009) 161–165, [arXiv:0807.0063 \[hep-th\]](#).
- [572] F. Karsch, “Properties of the Quark Gluon Plasma: A lattice perspective,” *Nucl. Phys.* **A783** (2007) 13–22, [arXiv:hep-ph/0610024](#).

- [573] O. Kaczmarek and F. Zantow, “Static quark anti-quark interactions in zero and finite temperature QCD. I: Heavy quark free energies, running coupling and quarkonium binding,” [Phys. Rev. D71 \(2005\) 114510](#), [arXiv:hep-lat/0503017](#).
- [574] T. Gunji, H. Hamagaki, T. Hatsuda, and T. Hirano, “Onset of J/psi melting in quark-gluon fluid at RHIC,” [Phys. Rev. C76 \(2007\) 051901](#), [arXiv:hep-ph/0703061](#).
- [575] T. Gunji, H. Hamagaki, T. Hatsuda, and T. Hirano, “Onset of J/psi melting in quark-gluon fluid at RHIC,” [J. Phys. G: Nucl. Part. Phys. 35 \(2008\) 104137](#).
- [576] F. Karsch and R. Petronzio, “ χ and J/ψ Suppression in Heavy Ion Collisions and a Model for its Momentum Dependence,” [Z. Phys. C37 \(1988\) 627](#).
- [577] **STAR** Collaboration, B. I. Abelev et al., “J/psi production at high transverse momentum in p+p and Cu+Cu collisions at $\sqrt{s_{NN}} = 200$ GeV,” [Phys. Rev. C80 \(2009\) 041902](#), [arXiv:0904.0439 \[nucl-ex\]](#).
- [578] **PHENIX** Collaboration, E. T. Atomssa, “J/psi Elliptic Flow, High pT Suppression and Upsilon Measurements in A+A Collisions by the PHENIX Experiment,” [Nucl. Phys. A830 \(2009\) 331c–334c](#), [arXiv:0907.4787 \[nucl-ex\]](#).
- [579] E. D’Hoker and Y. Guo, “Rigidity of SU(2,2—2)-symmetric solutions in Type IIB,” [JHEP 05 \(2010\) 088](#), [arXiv:1001.4808 \[hep-th\]](#).
- [580] M. Grana and J. Polchinski, “Gauge / gravity duals with holomorphic dilaton,” [Phys. Rev. D65 \(2002\) 126005](#), [arXiv:hep-th/0106014](#).
- [581] B. A. Burrington, J. T. Liu, L. A. Pando Zayas, and D. Vaman, “Holographic duals of flavored N = 1 super Yang-Mills: Beyond the probe approximation,” [JHEP 02 \(2005\) 022](#), [arXiv:hep-th/0406207](#).
- [582] I. Kirsch and D. Vaman, “The D3/D7 background and flavor dependence of Regge trajectories,” [Phys. Rev. D72 \(2005\) 026007](#), [arXiv:hep-th/0505164](#).
- [583] F. Benini, F. Canoura, S. Cremonesi, C. Nunez, and A. V. Ramallo, “Unquenched flavors in the Klebanov-Witten model,” [JHEP 02 \(2007\) 090](#), [arXiv:hep-th/0612118](#).
- [584] C. Nunez, A. Paredes, and A. V. Ramallo, “Unquenched flavor in the gauge/gravity correspondence,” [Adv. High Energy Phys. 2010 \(2010\) 196714](#), [arXiv:1002.1088 \[hep-th\]](#).
- [585] M. Kruczenski, D. Mateos, R. C. Myers, and D. J. Winters, “Meson spectroscopy in AdS/CFT with flavour,” [JHEP 07 \(2003\) 049](#), [arXiv:hep-th/0304032](#).
- [586] D. Arean and A. V. Ramallo, “Open string modes at brane intersections,” [JHEP 04 \(2006\) 037](#), [arXiv:hep-th/0602174](#).
- [587] R. C. Myers and R. M. Thomson, “Holographic mesons in various dimensions,” [JHEP 09 \(2006\) 066](#), [arXiv:hep-th/0605017](#).

- [588] A. V. Ramallo, “Adding open string modes to the gauge / gravity correspondence,” [Mod. Phys. Lett. **A21** \(2006\) 1481–1494](#), [arXiv:hep-th/0605261](#).
- [589] M. Kruczenski, L. A. P. Zayas, J. Sonnenschein, and D. Vaman, “Regge trajectories for mesons in the holographic dual of large- $N(c)$ QCD,” [JHEP **06** \(2005\) 046](#), [arXiv:hep-th/0410035](#).
- [590] A. Paredes and P. Talavera, “Multiflavour excited mesons from the fifth dimension,” [Nucl. Phys. **B713** \(2005\) 438–464](#), [arXiv:hep-th/0412260](#).
- [591] K. Peeters, J. Sonnenschein, and M. Zamaklar, “Holographic decays of large-spin mesons,” [JHEP **02** \(2006\) 009](#), [arXiv:hep-th/0511044](#).
- [592] A. L. Cotrone, L. Martucci, and W. Troost, “String splitting and strong coupling meson decay,” [Phys. Rev. Lett. **96** \(2006\) 141601](#), [arXiv:hep-th/0511045](#).
- [593] M. Bando, A. Sugamoto, and S. Terunuma, “Meson strings and flavor branes,” [Prog. Theor. Phys. **115** \(2006\) 1111–1127](#), [arXiv:hep-ph/0602203](#).
- [594] F. Bigazzi and A. L. Cotrone, “New predictions on meson decays from string splitting,” [JHEP **11** \(2006\) 066](#), [arXiv:hep-th/0606059](#).
- [595] C. Park, “The dissociation of a heavy meson in the quark medium,” [Phys. Rev. **D81** \(2010\) 045009](#), [arXiv:0907.0064 \[hep-ph\]](#).
- [596] M. Ali-Akbari and K. Bitaghsir Fadafan, “Rotating mesons in the presence of higher derivative corrections from gauge-string duality,” [Nucl. Phys. **B835** \(2010\) 221–237](#), [arXiv:0908.3921 \[hep-th\]](#).
- [597] M. Kruczenski, D. Mateos, R. C. Myers, and D. J. Winters, “Towards a holographic dual of large- $N(c)$ QCD,” [JHEP **05** \(2004\) 041](#), [arXiv:hep-th/0311270](#).
- [598] D. Mateos, R. C. Myers, and R. M. Thomson, “Holographic phase transitions with fundamental matter,” [Phys. Rev. Lett. **97** \(2006\) 091601](#), [arXiv:hep-th/0605046](#).
- [599] J. Babington, J. Erdmenger, N. J. Evans, Z. Guralnik, and I. Kirsch, “Chiral symmetry breaking and pions in non-supersymmetric gauge / gravity duals,” [Phys. Rev. **D69** \(2004\) 066007](#), [arXiv:hep-th/0306018](#).
- [600] I. Kirsch, “Generalizations of the AdS/CFT correspondence,” [Fortsch. Phys. **52** \(2004\) 727–826](#), [arXiv:hep-th/0406274](#).
- [601] T. Albash, V. G. Filev, C. V. Johnson, and A. Kundu, “A topology-changing phase transition and the dynamics of flavour,” [Phys. Rev. **D77** \(2008\) 066004](#), [arXiv:hep-th/0605088](#).
- [602] T. Albash, V. G. Filev, C. V. Johnson, and A. Kundu, “Global Currents, Phase Transitions, and Chiral Symmetry Breaking in Large N_c Gauge Theory,” [JHEP **12** \(2008\) 033](#), [arXiv:hep-th/0605175](#).

- [603] V. G. Filev, C. V. Johnson, R. C. Rashkov, and K. S. Viswanathan, “Flavoured large N gauge theory in an external magnetic field,” [JHEP **10** \(2007\) 019](#), [arXiv:hep-th/0701001](#).
- [604] A. Karch and A. O’Bannon, “Chiral transition of $N = 4$ super Yang-Mills with flavor on a 3-sphere,” [Phys. Rev. **D74** \(2006\) 085033](#), [arXiv:hep-th/0605120](#).
- [605] A. Parnachev and D. A. Sahakyan, “Chiral phase transition from string theory,” [Phys. Rev. Lett. **97** \(2006\) 111601](#), [arXiv:hep-th/0604173](#).
- [606] Y.-h. Gao, W.-s. Xu, and D.-f. Zeng, “NGN, QCD(2) and chiral phase transition from string theory,” [JHEP **08** \(2006\) 018](#), [arXiv:hep-th/0605138](#).
- [607] E. Antonyan, J. A. Harvey, and D. Kutasov, “The Gross-Neveu model from string theory,” [Nucl. Phys. **B776** \(2007\) 93–117](#), [arXiv:hep-th/0608149](#).
- [608] D. Mateos, R. C. Myers, and R. M. Thomson, “Holographic viscosity of fundamental matter,” [Phys. Rev. Lett. **98** \(2007\) 101601](#), [arXiv:hep-th/0610184](#).
- [609] S. Hong, S. Yoon, and M. J. Strassler, “Quarkonium from the fifth dimension,” [JHEP **04** \(2004\) 046](#), [arXiv:hep-th/0312071](#).
- [610] C. Hoyos-Badajoz, K. Landsteiner, and S. Montero, “Holographic Meson Melting,” [JHEP **04** \(2007\) 031](#), [arXiv:hep-th/0612169](#).
- [611] R. C. Myers, A. O. Starinets, and R. M. Thomson, “Holographic spectral functions and diffusion constants for fundamental matter,” [JHEP **11** \(2007\) 091](#), [arXiv:0706.0162 \[hep-th\]](#).
- [612] D. Mateos and L. Patino, “Bright branes for strongly coupled plasmas,” [JHEP **11** \(2007\) 025](#), [arXiv:0709.2168 \[hep-th\]](#).
- [613] J. L. Hovdebo, M. Kruczenski, D. Mateos, R. C. Myers, and D. J. Winters, “Holographic mesons: Adding flavor to the AdS/CFT duality,” [Int. J. Mod. Phys. **A20** \(2005\) 3428–3433](#).
- [614] S. A. Cherkis and A. Hashimoto, “Supergravity solution of intersecting branes and AdS/CFT with flavor,” [JHEP **11** \(2002\) 036](#), [arXiv:hep-th/0210105](#).
- [615] M. Gomez-Reino, S. G. Naculich, and H. Schnitzer, “Thermodynamics of the localized D2-D6 system,” [Nucl. Phys. **B713** \(2005\) 263–277](#), [arXiv:hep-th/0412015](#).
- [616] J. Erdmenger and I. Kirsch, “Mesons in gauge / gravity dual with large number of fundamental fields,” [JHEP **12** \(2004\) 025](#), [arXiv:hep-th/0408113](#).
- [617] J. Casalderrey-Solana, D. Fernandez, and D. Mateos, “A New Mechanism of Quark Energy Loss,” [Phys. Rev. Lett. **104** \(2010\) 172301](#), [arXiv:0912.3717 \[hep-ph\]](#).
- [618] J. Casalderrey-Solana, D. Fernandez, and D. Mateos, “Cherenkov mesons as in-medium quark energy loss,” [arXiv:1009.5937 \[hep-th\]](#).

- [619] M. C. Chu and T. Matsui, “Dynamic Debye Screening for a Heavy Anti-Quark Pair Traversing a Quark-Gluon Plasma,” [*Phys. Rev.* **D39** \(1989\) 1892](#).
- [620] N. Evans and E. Threlfall, “Mesonic quasinormal modes of the Sakai-Sugimoto model at high temperature,” [*Phys. Rev.* **D77** \(2008\) 126008](#), [arXiv:0802.0775 \[hep-th\]](#).
- [621] M. Kaminski *et al.*, “Quasinormal modes of massive charged flavor branes,” [*JHEP* **03** \(2010\) 117](#), [arXiv:0911.3544 \[hep-th\]](#).
- [622] M. Kaminski, K. Landsteiner, J. Mas, J. P. Shock, and J. Tarrio, “Holographic Operator Mixing and Quasinormal Modes on the Brane,” [*JHEP* **02** \(2010\) 021](#), [arXiv:0911.3610 \[hep-th\]](#).
- [623] S. Caron-Huot, P. Kovtun, G. D. Moore, A. Starinets, and L. G. Yaffe, “Photon and dilepton production in supersymmetric Yang- Mills plasma,” [*JHEP* **12** \(2006\) 015](#), [arXiv:hep-th/0607237](#).
- [624] B. Hassanain and M. Schvellinger, “Holographic current correlators at finite coupling and scattering off a supersymmetric plasma,” [*JHEP* **1004** \(2010\) 012](#), [arXiv:0912.4704 \[hep-th\]](#).
- [625] B. Hassanain and M. Schvellinger, “Towards ’t Hooft parameter corrections to charge transport in strongly-coupled plasma,” [*JHEP* **1010** \(2010\) 068](#), [arXiv:1006.5480 \[hep-th\]](#).
- [626] A. Paredes, K. Peeters, and M. Zamaklar, “Mesons versus quasi-normal modes: undercooling and overheating,” [*JHEP* **05** \(2008\) 027](#), [arXiv:0803.0759 \[hep-th\]](#).
- [627] E. Iancu and A. H. Mueller, “Light-like mesons and deep inelastic scattering in finite-temperature AdS/CFT with flavor,” [*JHEP* **02** \(2010\) 023](#), [arXiv:0912.2238 \[hep-th\]](#).
- [628] J. Casalderrey-Solana and D. Mateos, “Prediction of a Photon Peak in Relativistic Heavy Ion Collisions,” [*Phys. Rev. Lett.* **102** \(2009\) 192302](#), [arXiv:0806.4172 \[hep-ph\]](#).
- [629] P. Petreczky and K. Petrov, “Free energy of a static quark anti-quark pair and the renormalized Polyakov loop in three flavor QCD,” [*Phys. Rev.* **D70** \(2004\) 054503](#), [arXiv:hep-lat/0405009](#).
- [630] A. Mocsy, “Potential Models for Quarkonia,” [*Eur. Phys. J.* **C61** \(2009\) 705–710](#), [arXiv:0811.0337 \[hep-ph\]](#).
- [631] T. Hatsuda, “Hadrons above $T(c)$,” [*Int. J. Mod. Phys.* **A21** \(2006\) 688–693](#), [arXiv:hep-ph/0509306](#).
- [632] H. Satz, “Quarkonium Binding and Dissociation: The Spectral Analysis of the QGP,” [*Nucl. Phys.* **A783** \(2007\) 249–260](#), [arXiv:hep-ph/0609197](#).
- [633] M. Asakawa, T. Hatsuda, and Y. Nakahara, “Hadronic spectral functions above the QCD phase transition,” [*Nucl. Phys.* **A715** \(2003\) 863–866](#), [arXiv:hep-lat/0208059](#).

- [634] F. Arleo, D. G. d’Enterria, and D. Peressounko, “Direct photon spectra in Pb-Pb at $\sqrt{s} = 5.5$ TeV: hydrodynamics+pQCD predictions,” [arXiv:0707.2357](#) [[nucl-th](#)].
- [635] J. Liao and E. Shuryak, “Angular Dependence of Jet Quenching Indicates Its Strong Enhancement Near the QCD Phase Transition,” [Phys. Rev. Lett.](#) **102** (2009) 202302, [arXiv:0810.4116](#) [[nucl-th](#)].
- [636] **PHENIX** Collaboration, S. S. Adler et al., “A detailed study of high-p(T) neutral pion suppression and azimuthal anisotropy in Au + Au collisions at $\sqrt{s(NN)}^{1/2} = 200$ -GeV,” [Phys. Rev.](#) **C76** (2007) 034904, [arXiv:nucl-ex/0611007](#).
- [637] **STAR** Collaboration, J. Adams et al., “Azimuthal anisotropy and correlations at large transverse momenta in p + p and Au + Au collisions at $\sqrt{s(NN)}^{1/2} = 200$ -GeV,” [Phys. Rev. Lett.](#) **93** (2004) 252301, [arXiv:nucl-ex/0407007](#).
- [638] M. Laine, G. D. Moore, O. Philipsen, and M. Tassler, “Heavy Quark Thermalization in Classical Lattice Gauge Theory: Lessons for Strongly-Coupled QCD,” [JHEP](#) **0905** (2009) 014, [arXiv:0902.2856](#) [[hep-ph](#)].
- [639] S. Caron-Huot, M. Laine, and G. D. Moore, “A Way to estimate the heavy quark thermalization rate from the lattice,” [JHEP](#) **0904** (2009) 053, [arXiv:0901.1195](#) [[hep-lat](#)].