

From hydrodynamics to far-from-equilibrium dynamics

In Chapter 6 we have described ways in which holographic calculations have yielded insight into various properties of strongly coupled plasma that is at rest and in thermal equilibrium or, in our discussion of transport properties, infinitesimally close to being at rest and in thermal equilibrium. In this chapter, we release these restrictions. By the end of the chapter, we will be analyzing violent dynamical processes that are initially very far from thermal equilibrium and that may provide a caricature of the dynamics in the earliest moments of a heavy ion collision, or at least what that dynamics would be if the physics then was already strongly coupled. Before the development of gauge/gravity duality, there were no reliable quantum field theoretical calculations valid in far-from-equilibrium, highly time-dependent, strongly coupled matter. We shall build up the holographic tools that can now provide such calculations in stages over the course of the chapter. In doing so we shall make the connections to heavy ion physics manifest throughout, but it is important to realize that these tools are of considerable value in any other quantum field theoretical context in which the physics is strongly coupled and the questions of interest include far-from-equilibrium dynamics and thermalization.

We begin by letting the strongly coupled plasma move. In Sections 7.1 and 7.2 we show how to construct the gravitational description of solutions to the hydrodynamic equations for strongly coupled plasma in motion. That is, we continue to assume local thermal equilibrium, but we let the plasma move and flow. Upon making the standard assumptions that in a conventional description lead to hydrodynamics, namely upon assuming that the length scales that describe gradients of the flow velocity are long compared to all microscopic length scales that describe the plasma itself, we show that the flowing plasma has an equivalent gravitational description as a black brane whose metric, including its horizon, is undulating. This correspondence between hydrodynamic flow and the gravitational dynamics of an undulating metric, which we make explicit in Section 7.2, is basic to the relevance of holography to heavy ion collisions. In a heavy ion collision, the strongly

coupled plasma that is produced explodes outwards and is never at rest. However, as we have seen in Section 2.2, for much of the time that it is in existence the way that it expands and flows is well described by the laws of hydrodynamics.

In Section 7.2, we use the correspondence between hydrodynamics and gravity to recast the holographic determination of the constitutive relations that describe the strongly coupled plasma itself. The explicit examples that we focus on include its equation of state and the relation between its shear viscosity and its entropy density, both of which we have already calculated in Chapter 6. As we saw in that chapter, the results obtained from these holographic calculations resonate in many ways with what we are learning from lattice QCD calculations and from comparisons of hydrodynamic calculations to data from heavy ion collision experiments.

In Section 7.3, we further relax our assumptions and begin our discussion of the dynamics of strongly coupled matter that is far from equilibrium, and in particular the processes by which such matter equilibrates. Because QCD is asymptotically free, it is unlikely that the matter in the far-from-equilibrium conditions that characterize the very earliest moments of a sufficiently high energy heavy ion collision is itself strongly coupled. This means that, to the degree that realizable heavy ion collisions approach this high energy regime, we should not expect holographic calculations to provide as good models of, or as reliable insights into, the pre-equilibrium dynamics in heavy ion collisions as is possible for the hydrodynamics of the expanding strongly coupled plasma that emerges a little later or for the probes of this plasma that we shall discuss in Chapter 8. However, before the development of holographic approaches far-from-equilibrium dynamics in strongly coupled many-body systems was famously difficult to understand by any means. This makes the holographic analyses that we shall present beginning in Section 7.3 of considerable interest from a perspective that goes well beyond heavy ion collision physics. In the context of heavy ion collisions, understanding how the matter produced at early times, whatever its nature, isotropizes, how its motion comes to be governed by the laws of hydrodynamics, and how it reaches local thermal equilibrium have long been seen as central puzzles. If we can find quantities that seem to robustly characterize the equilibration of strongly coupled plasma starting from a large variety of widely varying initial conditions, perhaps we can gain insights into these questions even if holographic calculations are not able to capture the details of the initial conditions specific to heavy ion collisions. By the end of this Chapter we shall see that the equilibration timescale itself may be just such a quantity.

In Section 7.4 we describe how to prepare a far-from-equilibrium initial state whose subsequent evolution we wish to follow. In Sections 7.5 and 7.6 we present a complete analysis of the equilibration of a particularly simple class of far-from-equilibrium initial states, namely states which are initially spatially homogeneous.

For example, perhaps the equation of state is initially very far from that in thermal equilibrium. And, perhaps the matter is initially far from isotropic, with the pressures acting in different directions far from equal. In all these cases, the final state after enough time passes is strongly coupled plasma, at rest, in thermal equilibrium. In Section 7.7 we generalize our analysis to circumstances in which the final state is strongly coupled plasma in local thermal equilibrium that continues to flow, in a boost invariant expansion. Finally, in Section 7.8 we describe the holographic analysis of the collision of two sheets of energy density, infinite in their transverse extent, finite in thickness along the “beam direction”, slamming into each other at the speed of light. The final state in this case is strongly coupled plasma that continues to expand, hydrodynamically, but without boost invariance. This calculation represents the best holographic caricature to date of the collision of two large, highly Lorentz-contracted, nuclei.

7.1 Hydrodynamics and gauge/gravity duality

At distance and time scales much larger than the inverse temperature and any other microscopic dynamical distance and time scales, a quantum many-body system in *local thermal equilibrium* should be described by hydrodynamics. In the context of gauge/gravity duality, we thus expect hydrodynamics to emerge from the gravity description at large distance and time scales. More specifically, any solution to the equations of hydrodynamics that describes some flowing strongly coupled plasma in the boundary theory should have a corresponding bulk gravity solution. In Section 5.2.1 we discovered that in the boundary theory a fluid that is at rest and in thermal equilibrium with temperature T is described in the bulk gravity theory by a black brane solution whose Hawking temperature is the same T . Now, a typical hydrodynamic flow can be thought of as long-wavelength “ripples” on top of an equilibrium state. Accordingly, the corresponding dual gravity solution can be heuristically visualized as long-wavelength “ripples” on top of a static black brane.

The holographic correspondence between hydrodynamics and gravity was pioneered in a series of works by Policastro, Son and Starinets who were the first to work out the bulk gravity solutions for various hydrodynamic phenomena including the diffusion of momentum and the propagation and attenuation of sound waves [692, 691, 554, 748, 749]. In this section we describe bulk gravity solutions dual to boundary hydrodynamic flows which eventually approach thermal equilibrium. We shall follow the approach of Ref. [155]. (See also Refs. [705, 473].) Later in the chapter we shall consider more general situations, including both boost invariant hydrodynamic flows in which the fluid is in local thermal equilibrium but never becomes static and far-from-equilibrium dynamics. To keep our presentation manageable in scope, we shall consider only fluids that are conformal and that are

neutral with respect to all conserved charges. It is typically relatively straightforward to extend the construction that we discuss here to more general situations. We will briefly comment on these generalizations, many of which are of considerable current interest, at the end of Section 7.2.

We reviewed the formulation of hydrodynamics in Section 2.2.3. To set the stage, let us first establish notation and highlight some of the salient aspects of the standard formulation. In many applications (as in Section 2.2.3) it is helpful to use curvilinear coordinates to describe the hydrodynamic flow even when the boundary theory is flat or to consider boundary theories with curved geometry, as in some examples in subsequent sections. However, in this introductory presentation in which our goal is to illustrate general ideas in a simple context we shall consider only a flat boundary described with Cartesian coordinates. The generalization to curvilinear coordinates or a curved boundary can be made by replacing ∂_μ by the covariant derivative throughout. For a neutral fluid, the equations for hydrodynamics are simply the conservation of the stress tensor:

$$\partial_\mu T^{\mu\nu} = 0, \quad (7.1)$$

where $T^{\mu\nu}$ denotes the expectation value of the quantum stress tensor operator. $T^{\mu\nu}$ is in turn expressed via constitutive relations in terms of a derivative expansion of four hydrodynamic fields which we will choose to be the temperature $T(x)$ in the local fluid rest frame and the local fluid four-velocity $u^\mu(x)$, normalized according to $u^\mu u_\mu = -1$. Up to first order in derivatives, $T^{\mu\nu}$ can be written as

$$T^{\mu\nu} = \varepsilon(T) u^\mu u^\nu + P(T) \Delta^{\mu\nu} - \eta(T) \sigma^{\mu\nu} - \zeta(T) \partial_\lambda u^\lambda \Delta^{\mu\nu} + \dots, \quad (7.2)$$

where

$$\Delta_{\mu\nu} \equiv \eta_{\mu\nu} + u_\mu u_\nu \quad \text{and} \quad \sigma_{\mu\nu} \equiv \left(\Delta_\mu^\alpha \Delta_\nu^\beta - \frac{1}{3} \Delta_{\mu\nu} \Delta^{\alpha\beta} \right) (\partial_\alpha u_\beta + \partial_\beta u_\alpha). \quad (7.3)$$

The indices are raised and lowered using the Minkowski metric $\eta_{\mu\nu}$. The coefficients $\varepsilon(T)$, $P(T)$, $\eta(T)$, and $\zeta(T)$ are the energy density, pressure, shear and bulk viscosities respectively. It is possible to continue the derivative expansion (7.2) to any desired order by enumerating all possible terms allowed by symmetries and the local second law of thermodynamics. For example, the expansion of $T^{\mu\nu}$ up to second order in derivatives was given earlier in Eq. (2.24).

The question we would like to answer is how to derive (7.1) and (7.2) from the bulk gravity theory. In particular, we should be able to use the dual gravitational description to obtain precise expressions for ε , P , η , ζ and the coefficients of all higher order terms (such as those in (2.24)) that specify the hydrodynamics in the corresponding boundary theory. To achieve this goal, we need to find the most general solution to the bulk Einstein equations which describes the moving fluid

in local thermal equilibrium in the boundary theory and then obtain the boundary stress tensor corresponding to such a solution. We shall first explain how the conservation of the stress tensor (7.1) arises from the Einstein equations. Then, in Section 7.2, we describe a systematic procedure for deriving the constitutive relations (7.2) from solutions of Einstein equations and we provide an explicit calculation of the relations that are needed up to first order.

7.1.1 Conservation of the stress tensor and the Einstein equations

We now show that the conservation of stress tensor (7.1) can be obtained from a subset of the Einstein equations in an asymptotically AdS₅ spacetime. The full Einstein equations are given by

$$E_{MN} \equiv \mathcal{R}_{MN} - \frac{1}{2}g_{MN}\mathcal{R} + \Lambda g_{MN} = 0, \quad \text{where} \quad \Lambda \equiv -\frac{6}{R^2} \quad (7.4)$$

and R is the curvature of the asymptotic AdS₅. We follow the index convention of previous chapters with $x^M = (z, x^\mu)$, where z is the radial direction and $x^\mu = (t, \vec{x})$ are the spacetime directions along the boundary. We will be mainly interested in the evolution of the bulk metric along the radial direction. For this purpose, it is convenient to visualize the bulk spacetime as foliated by constant- z hypersurfaces Σ_z , which are spanned by the boundary coordinates x^μ , and treat the z -direction as a Euclidean “time”. We can then apply well-developed techniques for analyzing the time evolution of Einstein equations to the radial evolution that is of interest in the present context.

The Einstein equations can be separated into three groups depending on whether E_{MN} has zero, one or two indices along boundary directions:

$$\mathcal{H}_z \equiv E_{zM}n^M = 0, \quad (7.5)$$

$$\mathcal{H}_\mu \equiv E_{\mu M}n^M = 0, \quad (7.6)$$

$$E_{\mu\nu} = 0, \quad (7.7)$$

where n^M is the unit vector normal to Σ_z . Equations (7.7) contain second derivatives in z and are often called dynamical equations, while (7.5) and (7.6) contain only first derivatives in z and are often called the Hamiltonian constraint and the momentum constraint, respectively. A discussion of the implications of these constraint equations for the Ward identities of the boundary theory can be found in Refs. [664, 665]. Via the Bianchi identity, the structure of the Einstein equations is such that if the constraint equations (7.5) and (7.6) are satisfied on a single z -slice, the dynamical equations (7.7) will ensure that they are satisfied everywhere. Thus we need only impose (7.5) and (7.6) at a single value of z , for example at the boundary $z = 0$.

For a metric which does not have cross terms between z and x^μ ,

$$ds^2 = g_{zz}dz^2 + g_{\mu\nu}dx^\mu dx^\nu, \quad (7.8)$$

Eqs. (7.5) and (7.6) reduce simply to $E_{zz} = 0$ and $E_{\mu z} = 0$. For more general metrics that include cross terms,

$$ds^2 = g_{zz}dz^2 + 2g_{z\mu}dx^\mu dz + g_{\mu\nu}dx^\mu dx^\nu, \quad (7.9)$$

the unit normal n^M no longer lies along the z -direction. Equations (7.5) and (7.6) are then the appropriate constraint equations. In (7.9), $g_{\mu\nu}$ is the induced metric for Σ_z and we denote its inverse by $g^{\mu\nu}$.

It is standard textbook result (see e.g. Ref. [793]) that the momentum constraint (7.6) can be written explicitly as

$$\mathcal{H}_\mu = D_\mu (K^{\mu\nu} - g^{\mu\nu} K) = 0, \quad (7.10)$$

where $K_{\mu\nu}$ is the extrinsic curvature for a constant- z hypersurface Σ_z , $K \equiv g^{\mu\nu} K_{\mu\nu}$ is its trace, and D_μ is the intrinsic covariant derivative on Σ_z associated with $g_{\mu\nu}$. Now, according to the standard AdS dictionary, the boundary stress tensor can be obtained from the bulk metric via [111]

$$T^{\mu\nu} = \lim_{z \rightarrow 0} \frac{1}{8\pi G_5} \frac{R^6}{z^6} \left(K^{\mu\nu} - g^{\mu\nu} K - \frac{3}{R} g^{\mu\nu} \right), \quad (7.11)$$

with G_5 the five-dimensional Newton constant and R the AdS radius. We reviewed the derivation of (7.11) in Section 5.3.2, where it was Eq. (5.51). We have denoted $\langle T^{\mu\nu} \rangle$ by just $T^{\mu\nu}$ on the left-hand side of (7.11) as is standard in the hydrodynamic literature. At the boundary $z = 0$, $g_{\mu\nu}$ is proportional to the boundary Minkowski metric $\eta_{\mu\nu}$ and D_μ becomes the ordinary derivative ∂_μ . From (7.11) we therefore conclude that when the constraint (7.10) is imposed at the boundary it implies that

$$\partial_\mu T^{\mu\nu} = 0. \quad (7.12)$$

This then establishes that the constraint equations (7.6) correspond precisely to the conservation of the stress tensor in the boundary theory.

A similar analysis of the Hamiltonian constraint (7.5) at the boundary (slightly more involved than the analysis above because the Hamiltonian constraint is quadratic in the extrinsic curvature) shows that it implies that the boundary stress tensor is traceless [664, 665],

$$T^\mu{}_\mu = 0, \quad (7.13)$$

expressing the fact that the boundary theory is conformal.

Equations (7.11)–(7.13) are valid for the case of pure gravity in AdS₅ with a flat boundary metric $\eta_{\mu\nu}$. If the boundary metric is not flat then the expression in

parentheses in (7.11) contains an additional term $-RG^{\mu\nu}/2$, with $G^{\mu\nu}$ the Einstein tensor of the boundary metric. Since this term is divergence free with respect to the covariant derivative defined by the boundary metric, Eq. (7.12) then still holds with the replacement of ∂_μ by the covariant derivative. Finally, the right-hand side of (7.13) is replaced by [450]

$$\mathcal{A} = \frac{R^3}{8\pi G_5} \left(\frac{1}{8} \mathfrak{R}_{\mu\nu} \mathfrak{R}^{\mu\nu} - \frac{1}{24} \mathfrak{R}^2 \right), \tag{7.14}$$

where $\mathfrak{R}_{\mu\nu}$ and \mathfrak{R} are the Ricci tensor and the Ricci scalar of the boundary metric, respectively. This modification expresses the fact that the boundary theory may possess a conformal anomaly when placed in a curved manifold.

7.2 Constitutive relations from gravity

Having identified how Eqs. (7.6), a subset of the Einstein equations governing the gravitational physics in the bulk, imply the conservation of the stress tensor (7.1) in the boundary theory, in this section we outline a systematic procedure for deriving the constitutive relations (7.2) from gravity to all orders. This is achieved by finding the gravitational realization of any solution to the hydrodynamic equations order by order in a derivative expansion, and in particular by making the heuristic picture that boundary hydrodynamic flows correspond to “ripples” on a static black brane explicit. So, the procedure boils down to developing an iterative procedure with which to find a general solution of the bulk Einstein equations which describes the flowing boundary fluid in local thermal equilibrium.

Our starting point is the black brane metric discussed in Section 5.2.1 which we copy here for convenience:

$$ds^2 = \frac{R^2}{z^2} \left(-f dt^2 + d\vec{x}^2 + \frac{dz^2}{f} \right), \quad \text{with} \quad f(z) \equiv 1 - (\pi T z)^4. \tag{7.15}$$

This metric has an event horizon at

$$z = \frac{1}{\pi T} \tag{7.16}$$

and describes a system in thermal equilibrium with temperature T . Instead of using the coordinate t which becomes singular at the horizon where $f = 0$, it is more convenient for our purposes to use the so-called Eddington–Finkelstein coordinate v defined by

$$dv = dt - \frac{dz}{f}, \quad v = t - \int_0^z \frac{dz'}{f(z')}. \tag{7.17}$$

When written in terms of v , the metric (7.15) becomes

$$ds^2 = \frac{R^2}{z^2} (-2dv dz - f dv^2 + d\vec{x}^2), \quad (7.18)$$

from which we can see that lines of constant v and \vec{x} are light-like. The coordinate v has the nice feature that it reduces to t at the boundary $z = 0$, but remains non-singular at the horizon as can be seen from (7.18) and as is further illustrated in the Penrose diagram of Fig. 7.1. As will become clear later, the use of a coordinate like this that is regular at the horizon helps significantly in simplifying the analysis. Henceforth, we will denote $x^\mu \equiv (v, \vec{x})$.

The metric (7.18) (or (7.15)) describes a system at rest. The metric corresponding to a system moving with a constant four-velocity u_μ can be obtained by a Lorentz boost in x^μ , resulting in

$$ds^2 = \frac{R^2}{z^2} [2u_\mu dx^\mu dz + (-f u_\mu u_\nu + \Delta_{\mu\nu}) dx^\mu dx^\nu], \quad (7.19)$$

where $\Delta_{\mu\nu}$ is the projector introduced earlier in (7.3). Here and below, we will always use $u^\mu = \eta^{\mu\nu} u_\nu$ and $u_\mu u^\mu = -1$.

Now let us consider a system which is only in *local* thermal equilibrium, described by slowly varying local temperature $T(x^\mu)$ and flow velocity $u_\mu(x^\nu)$. The corresponding bulk metric describing such a non-equilibrium state is in general not known precisely. Nevertheless, the metric

$$ds_0^2 = g_{MN}^{(0)} dx^M dx^N = \frac{R^2}{z^2} [2u_\mu(x^\lambda) dx^\mu dz + h_{\mu\nu}^{(0)}(x^\lambda, z) dx^\mu dx^\nu] \quad (7.20)$$

with

$$h_{\mu\nu}^{(0)}(x^\lambda, z) = -f(T(x^\lambda)z) u_\mu(x^\lambda) u_\nu(x^\lambda) + \Delta_{\mu\nu}(x^\lambda), \quad (7.21)$$

obtained by replacing the constant parameters u_μ and T in (7.19) by $T(x^\lambda)$ and $u_\mu(x^\lambda)$ should provide a reasonable approximation. Owing to the coordinate dependence of those parameters, (7.20) no longer solves the Einstein equations. But, as $T(x^\lambda)$ and $u_\mu(x^\lambda)$ become more and more slowly varying, it should provide a better and better approximation. In particular, since the failure of (7.20) to solve the Einstein equations is solely due to gradients of $T(x^\lambda)$, $u_\mu(x^\lambda)$ along the boundary space and time directions, as these quantities do not depend on z , one expects to be able to correct the metric (7.20) order by order in an expansion in the number of derivatives of $T(x^\lambda)$ and $u_\mu(x^\lambda)$. That is, the full metric can be written in an expansion of the form

$$g = g^{(0)}(T(x^\mu), u_\mu(x^\nu); z) + \epsilon g^{(1)}(T(x^\mu), u_\mu(x^\nu); z) + \epsilon^2 g^{(2)}(T(x^\mu), u_\mu(x^\nu); z) + \dots, \quad (7.22)$$

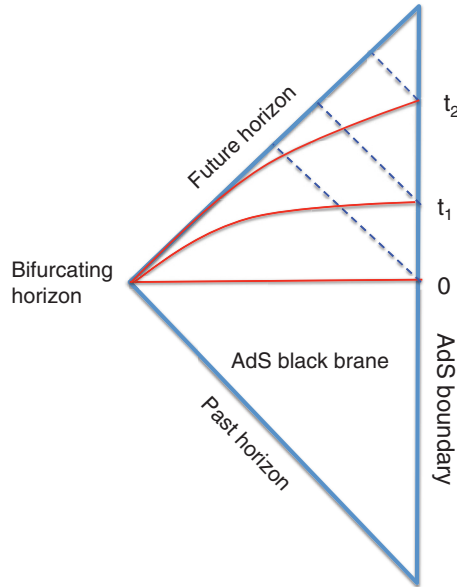


Figure 7.1 Penrose diagram for the region outside the horizon of an AdS black brane. The metric on a Penrose diagram is related to the actual spacetime metric by an overall (spacetime-dependent) scale factor chosen such that the entire infinite spacetime is transformed into a diagram of finite size. Since an overall scale factor in a metric does not change the causal structure, the Penrose diagram can be used to visualize the causal structure of the actual spacetime. In particular, in a Penrose diagram, light travels along 45 degree lines. (For textbook discussions of Penrose diagrams for black hole spacetimes, see Refs. [619, 683], for example.) In the figure, the vertical line on the right denotes the AdS boundary $z = 0$ with time running in the vertical direction. The boundary spatial directions \vec{x} are suppressed. (Each point in the diagram should be considered as an \mathbb{R}^3 .) The red lines ending at the boundary points $t = 0, t_1, t_2$ are lines of constant t , the time coordinate in the metric (7.15). The dashed blue lines originating at the same boundary points are lines of constant v , the Eddington–Finkelstein coordinate of (7.18). Along each such slice z increases from 0 at the boundary to the value (7.16) at the horizon. Notice that all constant- t slices meet at the “bifurcating horizon” which is the point where the past and future horizons meet. This is one way to see that the t -coordinate becomes singular at the horizon. In contrast, constant v -slices are infalling null geodesics from the boundary to the horizon. They provide a one-to-one map from points on the boundary to points on the future horizon.

where the zeroth metric $g^{(0)}$ is given by (7.20) and $g^{(n)}$ are *local* functions of T , u_μ and their derivatives along boundary directions, with n being the number of boundary derivatives. Here, ϵ is a book-keeping device to keep track of the total number of boundary derivatives and will be set to 1 at the end.

We will see below that the structure of the Einstein equations does admit a derivative expansion (7.22) of the metric in terms of hydrodynamic fields $T(x^\mu)$

and $u_\mu(x^\nu)$. With this expansion in hand, the boundary stress tensor obtained from (7.22) via (7.11) will then yield the constitutive relations. Combined with the result from Section 7.1 that the constraint equations (7.6) from among the Einstein equations imply the conservation of the boundary stress tensor, we will then have a full description of hydrodynamics on the boundary emerging from the Einstein equations that describe gravity in the bulk.

It is worth pausing to stress the picture that is intrinsic to the way we have written the metric in (7.20) and (7.21). The functions appearing in this metric that are varying as a function of the boundary coordinates x^λ are the standard variables of hydrodynamics, the flow velocity u_μ and the temperature T , which describe how the boundary theory fluid is flowing. If the procedure we have outlined can be carried out, it implies that the same information is encoded in the undulations of the bulk spacetime. In particular, since the T in (7.21) is undulating, the location at which the black brane apparent horizon sits is moving “up and down” in z as the fluid moves. At places and times where the boundary theory fluid is compressed, T increases, and the apparent horizon in the bulk moves to smaller values of z , closer to the boundary. Where the boundary fluid expands, T decreases, and the horizon moves to larger z farther away from the boundary. So, there is a direct relationship between the undulation of the horizon and the metric in the bulk gravitational theory on the one hand and the motion of the hydrodynamic fluid on the other.

Recalling that the horizon area of a stationary black hole in the bulk spacetime corresponds to the entropy of the boundary system, once we discover that a quintessential feature of the bulk metric dual to a hydrodynamic flow is that the event horizon is undulating and evolving dynamically it is natural to propose that the local area element on the horizon in this dynamical context corresponds to the local entropy current of the hydrodynamic flow [154]. More explicitly, writing the area form of a spatial section of the horizon¹ as

$$\mathcal{A} = \frac{1}{3!} a_{\mu_1\mu_2\mu_3} dx^{\mu_1} dx^{\mu_2} dx^{\mu_3}, \quad (7.23)$$

one can define the entropy current J_μ^S as the dual of \mathcal{A} divided by the familiar $4G_5$, i.e.

$$J_S^\mu = \frac{1}{4G_5} \frac{1}{3!} \epsilon^{\mu\mu_1\mu_2\mu_3} a_{\mu_1\mu_2\mu_3}. \quad (7.24)$$

In the static case, this current has a nonzero component only in the time direction and J_S^0 reduces to the standard area formula for the black hole entropy, as in Eq. (6.1). As a further consistency check, it is possible to show that even though

¹ At any point on the horizon the horizon is spanned by three spatial directions and one null direction. We can use the affine parameter along the family of null geodesics to define a foliation of the horizon. A spatial section of the horizon is a slice of the horizon at constant affine parameter, spanned by three spatial directions. For the case of a static AdS black brane, the spatial section of the horizon is spanned by the \bar{x} -directions.

J_S^μ in (7.24) is defined in terms of quantities at the horizon, when it is interpreted as a boundary theory current its divergence $\partial_\mu J_S^\mu$ is a Lorentz scalar [154], as it should be. And, finally, one can show that the fact that the area of the horizon in the bulk spacetime does not decrease corresponds to the fact that in the boundary theory $\partial_\mu J_S^\mu$ is non-negative [154], as expected for the entropy current.

Before turning to the construction of (7.22) in detail in the next subsection, let us mention here some of its key features. While for a bulk metric to yield (7.2) via (7.11) it is only necessary that a derivative expansion (7.22) exists near the boundary $z \sim 0$, we will see using the Eddington–Finkelstein coordinate (7.17) that such an expansion in fact exists for all z outside the event horizon (i.e. everywhere along the dashed lines in Fig. 7.1). Furthermore, we shall see that the problem of solving the Einstein equations for (7.22) in a derivative expansion factorizes into two separate problems:

- (1) Solving hydrodynamic equations for the hydrodynamic fields $T(x^\lambda)$ and $u_\mu(x^\lambda)$ in a derivative expansion.
- (2) Finding the radial evolution of the Einstein equations at a given boundary point x^μ . This is a one-dimensional problem that reduces to ordinary differential equations with sources, and can easily be solved.

7.2.1 Constitutive relations from gravity: explicit construction

In this subsection, we describe in detail how to solve the Einstein equations order-by-order in derivatives to obtain the derivative expansion (7.22). We will first discuss some general aspects of the construction and then carry out the explicit calculation up to first order.

General aspects

The Einstein equations are invariant under bulk diffeomorphisms. We can use such a diffeomorphism to impose that the metric satisfies

$$g_{zz} = 0, \quad g_{z\mu} = \frac{R^2}{z^2} u_\mu \tag{7.25}$$

to all orders. In other words, we fix a gauge in which the full metric can be written as

$$ds^2 \equiv g_{MN} dx^M dx^N = \frac{R^2}{z^2} [2u_\mu dx^\mu dz + h_{\mu\nu}(x^\lambda, z) dx^\mu dx^\nu] \tag{7.26}$$

and the expansion (7.22) becomes an expansion for $h_{\mu\nu}$ which can be written

$$h_{\mu\nu} = h_{\mu\nu}^{(0)}(T(x^\lambda), u_\mu(x^\lambda), z) + \epsilon h_{\mu\nu}^{(1)}(T(x^\lambda), u_\mu(x^\lambda), z) + \epsilon^2 h_{\mu\nu}^{(2)}(T(x^\lambda), u_\mu(x^\lambda), z) + \dots, \tag{7.27}$$

with $h^{(0)}$ given by (7.21). Once we go beyond zeroth order, the heuristic picture of a direct relationship between undulations of the bulk metric and the motion of the boundary hydrodynamic fluid remains, but the specific relation between the location of the horizon and the local T in the fluid can become more complicated.

Given that the constraint equations (7.6) in the bulk imply the conservation of the boundary stress tensor, the boundary hydrodynamic equations must be a part of the Einstein equations which we shall solve order by order in a derivative expansion. We thus expect that the hydrodynamic fields entering (7.27) should also have a derivative expansion

$$T(x^\mu) = T^{(0)}(x^\mu) + \epsilon T^{(1)}(x^\mu) + \dots, \quad u_\mu(x^\nu) = u_\mu^{(0)}(x^\nu) + \epsilon u_\mu^{(1)}(x^\nu) + \dots \tag{7.28}$$

Our task is then to substitute (7.26)–(7.28) into the Einstein equations (7.5)–(7.7), and solve the resulting equations at each order in ϵ ,

$$\mathcal{H}_z^{(n)} = 0, \quad n = 0, 1, \dots \tag{7.29}$$

$$\mathcal{H}_\mu^{(n)} = 0, \quad n = 0, 1, \dots \tag{7.30}$$

$$E_{\mu\nu}^{(n)} = 0, \quad n = 0, 1, \dots \tag{7.31}$$

where $\mathcal{H}_z^{(n)}$ is the coefficient of ϵ^n in the expansion of \mathcal{H}_z and similarly for the others. Note that, when obtaining (7.29)–(7.31), in order for ϵ to keep track of the total number of boundary derivatives each boundary derivative in the Einstein equations should give rise to a factor of ϵ , in addition to those in (7.27) and (7.28). Thus for each term in the n th order equations (7.29)–(7.31), the sum of the number of x^μ -derivatives and all the upper indices in $h^{(m)}$, $T^{(k)}$, $u_\mu^{(l)}$ should be exactly n . As we will see below this power counting rule has important consequences.

By construction, the zeroth order equations, which do not contain any boundary derivatives, are solved by the zeroth order metric $g^{(0)}(T^{(0)}, u_\mu^{(0)}; z)$ given in (7.20) since the black brane metric satisfies the Einstein equations with a constant T and u_μ .

At any order $n \geq 1$, from the general structure of the perturbative expansion of differential equations and using the power counting that we have defined we can deduce the following regarding Eqs. (7.29)–(7.31).

- (1) $u_\mu^{(n)}$ and $T^{(n)}$ cannot appear in the n th order equations at all. (Because they are already of order n , $u_\mu^{(n)}$ and $T^{(n)}$ certainly cannot appear with any boundary derivatives acting on them. Thus, if they appear they effectively behave as constant shifts of T and u_μ , which solve all the equations but are trivial.)
- (2) The constraint equations (7.30), which are the n th order terms in the expansion of (7.6), become precisely the hydrodynamic equations (7.12) at n th order, i.e.

$$\partial_\mu T_{(n-1)}^{\mu\nu} = 0, \quad n = 1, 2, \dots, \tag{7.32}$$

where $T_{(n-1)}^{\mu\nu}$ is the boundary stress tensor expanded to order $n - 1$ and is obtained from $g^{(n-1)}$ via (7.11).

- (3) Equations (7.29) and (7.31) can be used to solve for $h^{(n)}$. There are altogether eleven equations to solve for the ten components of $h_{\mu\nu}^{(n)}$. There is one redundancy, as once the constraint equation (7.29) is imposed on a single constant- z hypersurface, the other dynamical equations will ensure that it is satisfied everywhere. The power counting rule implies that the differential equations for $h^{(n)}$ with $n \geq 1$ that are obtained from (7.29) and (7.31) must take the form

$$\mathbb{H} [T^{(0)}, u_{\mu}^{(0)}] h^{(n)}(z, x^{\mu}) = s_n, \quad n \geq 1 \tag{7.33}$$

where the differential operator \mathbb{H} must have the following properties.

- (a) It can not contain any boundary derivatives.
- (b) It can depend on the zeroth order quantities $T^{(0)}$ and $u_{\mu}^{(0)}$ only. Neither boundary derivatives of $T^{(0)}$ or $u_{\mu}^{(0)}$ nor higher order quantities like $h^{(k)}$, $T^{(k)}$ or $u_{\mu}^{(k)}$ with $k \geq 1$ can appear. In other words, \mathbb{H} evaluated at a boundary point x^{μ} only depends on the values of $T^{(0)}$ and $u_{\mu}^{(0)}$ at that same boundary point.
- (c) It is independent of n .
- (d) It is a linear differential operator in z with at most two derivatives.

The first three properties of \mathbb{H} follow immediately from the power counting rule, and can be described by saying that \mathbb{H} is “ultra-local” along the boundary directions. As a result, one can integrate (7.33) in z point-by-point in x^{μ} . The equations at different x^{μ} ’s do not interfere with one another. The last property (d) follows from the fact that the Einstein equations are second order differential equations.

- (4) The s_n on the right-hand side of (7.33) are source terms (i.e. terms with no z derivatives) which “measure” the failure of the metric g to satisfy the Einstein equations at each order due to the dependence of its parameters on location in the boundary spacetime. The s_n are local functions of $h^{(n)}$, $T^{(k)}$ and $u_{\mu}^{(k)}$ for $k < n$ and their derivatives (subject to the power counting rule), and must contain at least one boundary derivative in each term.

Given that the differential operator \mathbb{H} is ultra-local, meaning that it can be integrated along the radial direction at any location in the boundary directions with no dependence on any other locations, solving the full Einstein equations (7.29)–(7.31) factorizes into two separate problems: solving the boundary hydrodynamic equations (7.32) and solving the radial evolution equation (7.33). We can then integrate the radial equation at a single boundary point, say $x^{\mu} = 0$, and easily generalize the results to all points. It is also convenient to work in the local rest frame, meaning that in (7.26) we choose coordinates such that $u^{\mu}(x^{\lambda} = 0) = (1, 0, 0, 0)$.

With the above general aspects of the derivative expansion in hand, we are now ready to work out the solutions of the Einstein equations (7.29)–(7.31) explicitly up to first order. We shall only push our explicit results to this order as doing so is already enough to illustrate the general procedure and to obtain the explicit form of \mathbb{H} in (7.33), whose integration could then be carried out to higher order.

Explicit solutions up to first order

At first order $n = 1$, the constraint equations (7.30), following from the discussion of (7.10)–(7.12), reduce to

$$\partial_\mu T_{(0)}^{\mu\nu} = 0. \quad (7.34)$$

$T_{(0)}^{\mu\nu}$ is the zeroth order boundary stress tensor obtained via (7.11) from $g^{(0)}$, given by (7.20). Not surprisingly, one finds a perfect fluid form (see Appendix D for a derivation)

$$T_{(0)}^{\mu\nu} = (\varepsilon(T^{(0)}) + P(T^{(0)})) u_{(0)}^\mu u_{(0)}^\nu + P(T^{(0)}) \eta^{\mu\nu} \quad (7.35)$$

with

$$\varepsilon(T) = \frac{3R^3(\pi T)^4}{16\pi G_5}, \quad P(T) = \frac{R^3(\pi T)^4}{16\pi G_5}. \quad (7.36)$$

For $\mathcal{N} = 4$ super-Yang–Mills theory, from (5.12) which we reproduce here for convenience,

$$\frac{R^3}{8\pi G_5} = \frac{N_c^2}{4\pi^2}, \quad (7.37)$$

we then have

$$\varepsilon(T) = \frac{3N_c^2\pi^2 T^4}{8}, \quad P(T) = \frac{N_c^2\pi^2 T^4}{8}, \quad (7.38)$$

which agree with the expressions that we found for the energy density and pressure previously in Section 6.1. So, we have reproduced the equation of state relating ε and P , which is the only constitutive relation that arises in the zeroth order stress tensor, for the strongly coupled plasma of $\mathcal{N} = 4$ SYM theory.

Equations (7.34) are simply the hydrodynamic equations for a perfect fluid. They can be solved for $T^{(0)}(x)$ and $u_\mu^{(0)}$. Expanding around $x^\mu = 0$ in the local rest frame, working to first order we find that $\partial_\mu u_0 = 0$ and that Eqs. (7.34) reduce to

$$\frac{1}{3} \nabla \cdot u = -\frac{\partial_\nu T}{T}, \quad (7.39)$$

$$\partial_\nu u_i = -\frac{\partial_i T}{T}, \quad (7.40)$$

where

$$\nabla \cdot u \equiv \sum_{i=1}^3 \partial_i u_i . \tag{7.41}$$

These results can also be obtained directly from (7.30). Recalling how T appears in the bulk metric specified by (7.20) and (7.21), we see that Eq. (7.39) is an explicit illustration (to first order in the derivative expansion) of the fact that as the boundary theory fluid expands or compresses the apparent horizon in the bulk spacetime moves to larger or smaller z .

So far, we have pursued the general algorithm that we have laid out previously through step (2) above, i.e. through Eq. (7.32). We have obtained the spacetime-varying quantities T and u^μ that specify the fluid motion and the bulk metric to zeroth order in the derivative expansion, and have obtained $T^{\mu\nu}$ to zeroth order and hence the zeroth order constitutive relation (7.38), namely the equation of state. To get the first order constitutive relations, we need $T^{\mu\nu}$ to first order, which means that we need to obtain the $h_{\mu\nu}^{(1)}$ by completing steps (3) and (4) in the general algorithm, which is to say solving (7.33). Once we have $h_{\mu\nu}$ to first order, we will have $g_{\mu\nu}$ to first order and from (7.11) we can then obtain $T^{\mu\nu}$ to first order, and the first order constitutive relations.

Following this algorithm, let us look at (7.29) and (7.31) with $n = 1$. At $x^\mu = 0$, the various components of $h^{(1)}$ can be classified according to their transformation properties under spatial $SO(3)$ rotations in the local rest frame:

$$\text{scalar : } h_{vv}^{(1)}, \quad \mathfrak{h}^{(1)} \equiv \frac{1}{3} \sum_{i=1}^3 h_{ii}^{(1)} \tag{7.42}$$

$$\text{vector : } h_{vi}^{(1)} \tag{7.43}$$

$$\text{tensor : } \alpha_{ij}^{(1)} \equiv h_{ij}^{(1)} - \delta_{ij} \mathfrak{h}^{(1)} \tag{7.44}$$

and the Einstein equations for the three different sectors decouple from one another. The equations for these components are of the form (7.33) and are given at $x^\mu = 0$ in the local rest frame by the explicit expressions below, with primes below denoting z -derivatives.

(1) The scalar sector:

$$z^4 (z^{-4} h_{vv})' - (2 + f) \mathfrak{h}' = -2 \nabla \cdot u, \tag{7.45}$$

$$\mathfrak{h}'' = 0. \tag{7.46}$$

(2) The vector sector:

$$f \left[z^3 \partial_z (z^{-3} h'_{vi}) + \frac{3}{z} \partial_v u_i \right] + 4\pi^4 T^3 z^3 (\partial_i T + T \partial_v u_i) = 0. \tag{7.47}$$

(3) The tensor sector:

$$z^3 (z^{-3} f \alpha'_{ij})' = -\frac{3}{z} \sigma_{ij}, \quad (7.48)$$

where

$$\sigma_{ij} \equiv \partial_i u_j + \partial_j u_i - \frac{2}{3} \delta_{ij} \nabla \cdot u. \quad (7.49)$$

For notational simplicity, in the equations above we have suppressed various superscripts. One should understand h as $h^{(1)}$ while T and u_i should be understood as $T^{(0)}$ and $u_i^{(0)}$. Note that the left-hand sides of Eqs. (7.45)–(7.48) can be understood as the explicit definitions of the operator \mathbb{H} , that we introduced in general terms in (7.33), acting on fields in the different sectors. Equations (7.45)–(7.48) are all first order ordinary differential equations in z with sources. Therefore, they can all be integrated easily. We require the solutions to be normalizable at the boundary and regular at the black brane horizon (7.16) of the zeroth order solution where the function f is zero and a potential singularity could arise. These conditions fix some of the integration constants, but not all. We will set all other integration constants to zero as they arise from solving the homogeneous parts of Eqs. (7.45)–(7.48) and simply correspond to shifting the parameters of the zeroth order solutions.

With these considerations in mind, the scalar sector equations (7.45)–(7.46) then have solutions given by

$$\mathfrak{h} = 0, \quad h_{vv} = \frac{2\nabla \cdot u}{3} z. \quad (7.50)$$

The vector equation (7.47) has a singularity at the event horizon coming from the factor f on its left-hand side. Note, however, that Eq. (7.47) can be further simplified using the constraint equation (7.40). In particular, the potential troublesome factor f cancels on both sides, yielding

$$z^3 (z^{-3} h'_{vi})' = -\frac{3}{z} \partial_v u_i, \quad (7.51)$$

which now has a regular solution

$$h_{vi} = z \partial_v u_i. \quad (7.52)$$

Turning now to Eq. (7.48) for the tensor sector equation, upon integrating it once we get

$$z^{-3} f \alpha'_{ij} = \sigma_{ij} (z^{-3} - (\pi T)^3), \quad (7.53)$$

where we have chosen the integration constant to ensure that α_{ij} is regular at the horizon in the next integration, namely

$$\alpha_{ij} = F(z) \sigma_{ij} \quad (7.54)$$

with

$$F(z) = \int_0^z dz' \frac{z'^3}{f} (z'^{-3} - (\pi T)^3) . \tag{7.55}$$

As mentioned earlier in our general discussion, $u^{(1)}$ and $T^{(1)}$ do not appear and are unconstrained at this order.

Collecting (7.50), (7.52) and (7.54), we find that the bulk metric specified as in (7.26) and (7.27) receives a first order contribution given by

$$\frac{R^2}{z^2} h_{\mu\nu}^{(1)} dx^\mu dx^\nu = \frac{R^2}{z^2} \left(\frac{2}{3} (\nabla \cdot u) z dv^2 + 2z \partial_\nu u_i dv dx^i + F(z) \sigma_{ij} dx^i dx^j \right) . \tag{7.56}$$

We have derived this expression in the fluid rest frame at $x^\mu = 0$. We can immediately generalize it to obtain the first order correction to the metric at a generic point in spacetime where the fluid velocity four-vector is some generic u_μ by making the substitutions

$$u_i \rightarrow u_\mu, \quad \partial_\nu \rightarrow u^\lambda \partial_\lambda, \quad \partial_i \rightarrow \Delta_\mu{}^\nu \partial_\nu, \quad dv \rightarrow -u_\mu dx^\mu, \quad dx^i \rightarrow \Delta^\mu{}_\nu dx^\nu, \tag{7.57}$$

where $\Delta_{\mu\nu}$ is the projector introduced in (7.3). Upon making these substitutions, σ_{ij} becomes $\sigma_{\mu\nu}$ defined in (7.3) and from (7.56) we find that the general expression for the first order contribution to the bulk metric is given by

$$h_{\mu\nu}^{(1)} = \frac{2}{3} z (\partial_\lambda u^\lambda) u_\mu u_\nu - z u^\lambda \partial_\lambda (u_\mu u_\nu) + F(z) \sigma_{\mu\nu} . \tag{7.58}$$

With the metric for the bulk spacetime now determined up to first order in the derivative expansion, we take this metric and use (7.11) to obtain the first order contribution to the stress tensor, finding

$$T_{\mu\nu}^{(1)} = -\frac{R^3 (\pi T)^3}{16\pi G_5} \sigma_{\mu\nu} . \tag{7.59}$$

(For details, see Appendix D.) Comparing (7.59) with (7.2), we finally conclude that the shear and bulk viscosities of the strongly coupled plasma are given by

$$\eta = \frac{R^3 (\pi T)^3}{16\pi G_5} = \frac{s}{4\pi} \quad \text{and} \quad \zeta = 0 , \tag{7.60}$$

respectively, where s is the entropy density of the system. In this way we recover the results for these transport coefficients that were derived via a different method in Section 6.2.

Finally, with the first order correction to the metric $h_{\mu\nu}^{(1)}$ in hand we can evaluate the entropy current (7.24) explicitly to first order in the derivative expansion, obtaining [154]

$$J_S^\mu = s u^\mu , \tag{7.61}$$

where s is the entropy density given by the area of the horizon in the zeroth order black brane metric. Note that (7.61) is the standard zeroth order expression; the explicit evaluation of (7.24) shows that there is no first order contribution. We can then take the divergence of (7.61) and use the conservation equation (7.12) (which relates terms that are first and second order in the derivative expansion since $T_{\mu\nu}$ contains zeroth and first order terms) to show that for a neutral conformal fluid, in which the bulk viscosity vanishes,

$$\partial_\mu J_S^\mu = \frac{\eta}{2T} \sigma_{\mu\nu} \sigma^{\mu\nu} . \quad (7.62)$$

As expected at leading order, we see that the shear viscosity η controls the production of entropy as such a fluid flows. In a fluid with conserved currents like those we have discussed in Section 6.2.5, both Eqs. (7.61) and (7.62) contain additional terms proportional to the conductivity or conductivities and, if any of the currents are anomalous, further terms introduced by the anomalies [750].

7.2.2 Generalizations

The first order calculation above can be extended to higher orders. At n -th order, one first solves the constraint equations (7.30), or equivalently the hydrodynamic equations (7.32), and then solves the equations (7.33) for $h_{\mu\nu}^{(n)}$ which arise from (7.29) and (7.31). Around a single point, (7.32) become algebraic equations at each order. The integration of (7.33) is very similar to that at the first order except that the sources are different. Thus these equations can be solved straightforwardly to all orders although the number of terms in s_n and $T_{n-1}^{\mu\nu}$ increase quickly.

As discussed around Eq. (2.24), in a conformal theory five additional transport parameters arise at second order in the derivative expansion. The values of these quantities can be found by extending the derivative expansion of the Einstein equations to second order [155], as we have described. In our discussion of transport coefficients in Section 6.2.4, we have already quoted the values of these five second order parameters for the strongly coupled plasma of $\mathcal{N} = 4$ SYM theory in Eq. (6.45), where their physical implications were also discussed.

The iterative procedure outlined here can also be straightforwardly generalized to other fluid systems, including charged fluids, fluids with spontaneous symmetry breaking and Goldstone bosons, including superfluids, fluids driven by external forces, and nonconformal fluids. We refer readers to Refs. [705, 473] for reviews and a more extensive reference list. A particularly interesting result among these generalizations is the modification of charged fluid hydrodynamics that is induced by quantum anomalies [787, 789, 788, 791, 792, 341, 117, 750] that we have discussed in Section 6.2.5.

7.3 Introduction to far-from-equilibrium dynamics

We have seen in previous sections that the equilibrium and near-equilibrium properties of strongly coupled plasma with a gravity dual are encoded in the equilibrium and near-equilibrium properties of the dual black brane. In particular, we saw in Section 7.1 that the effective theory that governs long-wavelength fluctuations around an equilibrium black brane is hydrodynamics itself.

In this section we turn to the formation of strongly coupled plasma starting from some initial far-from-equilibrium state in a theory that is strongly coupled at all length scales. Without taking advantage of holography, there are no known methods for doing reliable calculations of far-from-equilibrium strongly coupled dynamics in quantum field theory. Within conventional quantum field theoretical methods, the adjectives “strong coupling” on the one hand and “time-dependent” or “far-from-equilibrium” on the other represent major outstanding challenges, separately and even more so in concert. When holography can be applied, the dual gravitational description of the far-from-equilibrium strongly coupled dynamics consists of the formation of a highly disturbed, far-from-equilibrium black hole horizon, and its subsequent relaxation towards an equilibrium state. This is in general not easy to analyze, but much progress has occurred in recent years. The study of this type of dynamics on the gravity side requires solving Einstein’s equations in the presence of strong time (and possibly space) dependence. This can be done analytically for certain highly fine-tuned initial conditions [361, 350] or if some approximations are made [495, 642, 493, 152, 110, 109], but generically it can only be done numerically [696, 477, 163, 476, 164, 478, 368, 807, 492, 370, 290, 291, 292, 448, 119, 14, 287, 286, 218]. Thus we will focus on what studies using numerical relativity methods have taught us about black hole formation in 4+1-dimensional spacetimes that are asymptotically AdS and, consequently, what these studies have taught us about equilibration of hot strongly coupled matter in non-Abelian gauge theories. Since an important goal is to learn generic lessons that provide insight into dynamics during the initial stages of a heavy ion collision, we will concentrate on studies in which the gauge theory lives on Minkowski space, as opposed to, for example, a three-dimensional sphere.

As we discussed at the beginning of this chapter, because QCD is asymptotically free the dominant dynamics at the earliest moments of a sufficiently energetic heavy ion collision is expected to be weakly coupled, with the relevant (weak) coupling being α_{QCD} evaluated at the (short) distance scale corresponding to the mean spacing between gluons in the transverse plane at the moment when the two highly Lorentz-contracted nuclei collide. As we have seen in Section 2.2, though, after

a time that is of order or perhaps even less than 1 fm we find that the collision has produced a strongly coupled nearly perfect liquid. A complete account of the dynamics that starts with a far-from-equilibrium weakly coupled state and results in a strongly coupled hydrodynamic liquid must involve both weakly and strongly coupled dynamics. One motivation for the investigations in this chapter, where we will watch the formation of a strongly coupled hydrodynamic liquid starting from a wide variety of far-from-equilibrium states that are themselves also strongly coupled, is the hope that by understanding far-from-equilibrium dynamics and equilibration in both the strong and weak coupling limits we can bracket the real world physics. We shall see that the equilibration timescale itself is an example of a quantity where these investigations have indeed yielded insights along these lines.

At present it is too ambitious to envision literally simulating the gravity dual of a full collision that starts with widely separated nuclei heading towards each other in $3 + 1$ -dimensional Minkowski space, since the weakly coupled aspects of the nuclear physics present great challenges on the gravity side. Simulating the ultra-relativistic collision of two Lorentz-contracted spheres of strongly coupled matter is certainly conceivable, however. In Section 7.8 we will study a toy model of such a collision in $\mathcal{N} = 4$ super-Yang–Mills in an approximation in which the “nuclei” are taken to be infinitely big and have been replaced by sheets of energy density with a finite thickness in the “beam” direction that are translation-invariant in the transverse directions. Before we get to this model, however, we will introduce the construction of far-from-equilibrium states in general terms in Section 7.4 and will then study several somewhat less physical but simpler systems. In Sections 7.5 and 7.6 we shall treat the formation of static strongly coupled plasma from initial conditions that are homogeneous but strongly anisotropic, and in Section 7.7 we shall consider the formation of an expanding, boost-invariant, volume of strongly coupled plasma. Not only will these prior studies lay the ground work that will allow us to understand Section 7.8, but they will also teach us interesting lessons in their own right about the far-from-equilibrium dynamics of strongly coupled matter and its equilibration. Because these settings are simpler, also, it has been possible in these contexts to investigate a wide variety of initial conditions which makes it possible to get a sense of what features of the non-equilibrium dynamics are generic.

Useful references where holographic calculations of far-from-equilibrium dynamics are reviewed from viewpoints complementary to those adopted here include Refs. [143, 472, 494, 738].

Finally, a comment on notation. Throughout the rest of this chapter we will set the AdS radius R to unity.

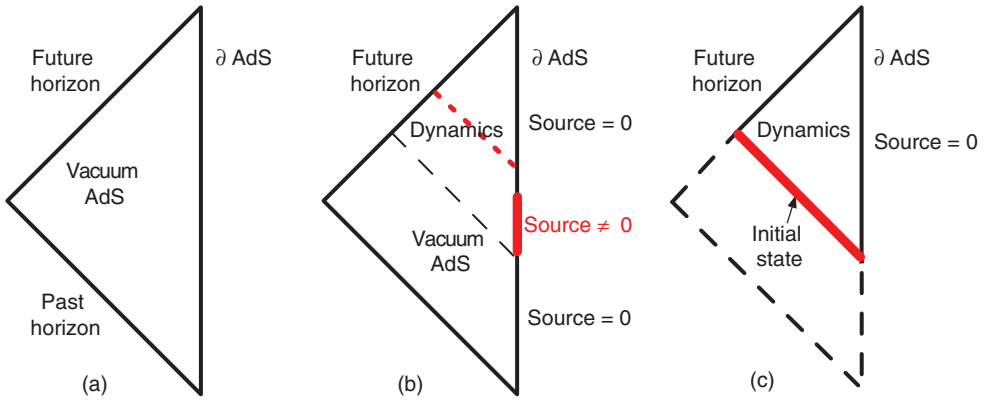


Figure 7.2 Penrose diagram of the spacetime outside the horizon of: (a) vacuum AdS; (b) the spacetime created when a source at the boundary of AdS is turned on; (c) the spacetime associated with the evolution of an initial state specified on an initial-time hypersurface in the bulk spacetime. Figure taken from Ref. [447].

7.4 Constructing far-from-equilibrium states

In order to study the far-from-equilibrium dynamics of plasma formation we must first prepare a far-from-equilibrium initial state whose subsequent evolution and thermalization we will study. This can be done in (at least) two ways. The initial state can be defined implicitly as the state that results from acting on the ground state of the theory with an external source [290, 291], or it can be defined explicitly by specifying initial conditions [292, 448, 446, 447].

In the first approach one turns on a time-dependent source in the boundary gauge theory with compact support in time. Before the source is turned on the system is in its ground state. The work done on the system by the external source takes it from its ground state to an excited state. This excited state evolves in time and, after the source is turned off, eventually relaxes to a thermal state, in equilibrium.

The description of this process on the gravity side is encoded in the Penrose diagram depicted in Fig. 7.2b. The physics is easy to understand by recalling that, through the gauge/string duality, a source in the gauge theory is identified with the value at the boundary of AdS of an appropriate supergravity (or string) field. Turning on the source at some time t_i therefore changes the boundary condition for this field in a time-dependent way. As a result, a wave of radiation is sent into the bulk at $t = t_i$. Since this cannot propagate faster than the speed of light, causality implies that this process cannot affect the geometry below the 45° black dashed line in Fig. 7.2b. The part of the geometry below this line is therefore dual to the CFT ground state, namely it is a piece of anti-de Sitter space with no excitations in it, to which we will simply refer as “vacuum AdS”. The work done by the source

in the gauge theory translates into the energy carried into the bulk by the wave. If the injected energy is such that it results in a finite energy density (which means that the total injected energy is infinite), then we expect the formation of a planar, regular, future horizon which at late times will be in equilibrium with some Hawking temperature $T > 0$. This is easy to understand on the gauge theory side, since we expect that an interacting system with finite energy density will eventually reach thermal equilibrium. On the gravity side this corresponds to the formation of a horizon with a nonzero Hawking temperature.

The region of interest is that labeled “Dynamics”, above the black dashed line in Fig. 7.2b. As is clear from the causal structure of the Penrose diagram, the metric and the other supergravity fields in this region are completely determined by the fact that the spacetime is vacuum AdS below the dashed line, together with the knowledge of the sources at the boundary of AdS for $t \geq t_i$. In other words, the problem of determining the supergravity fields in the region of interest is well posed. Once this problem is solved, which typically must be done numerically, the entire evolution of the bulk gravitational state and therefore of the boundary gauge theory state is known, and any observable can be computed. For example, the time-dependent expectation values of operators can be read off from the form of the corresponding dual fields near the AdS boundary at a given time. In subsequent sections we will illustrate this procedure by computing the expectation value of the stress tensor.

The second, explicit, approach to preparing a far-from-equilibrium initial state [292, 447] is illustrated by the Penrose diagram in Fig. 7.2c. In this case the initial state is specified explicitly on the gravity side in terms of the metric and other supergravity fields on some initial-time slice, depicted as a 45° thick red line in the figure. The only restriction on these fields is that they must satisfy the constraints associated with Einstein’s equations, as well as appropriate boundary conditions near the AdS boundary. Once these initial values are specified along the thick line, the problem of their future time evolution is again well posed, as illustrated by the causal structure of the diagram. The initial-time slice shown in Fig. 7.2c is null, but one can equally well choose a spacelike hypersurface, as in Refs. [448, 446]. In all these explicit approaches, the region below the initial-time slice is not of interest; only the future evolution of the system is. Moreover, no external sources are turned on, which immediately defines the boundary conditions near the AdS boundary mentioned above. One simply specifies an initial state, which generically will be far from equilibrium, lets it go, and watches it evolve towards an equilibrium state.

The two approaches described above are related. For example, once the source has been turned off in Fig. 7.2b and the evolution has been determined, one could read off the values of the supergravity fields at an appropriate initial-time slice such as that shown as a dotted red line in the figure. Obviously these values satisfy

the supergravity constraints and evolve to the future in the absence of boundary sources. Therefore any combination of sources *à la* Fig. 7.2b defines an initial state *à la* Fig. 7.2c. Presumably the converse is also true, namely any admissible initial state in Fig. 7.2c can also be obtained by turning on and off an appropriate combination of sources. However, these sources may be highly non-local. For this reason it is useful to explore both approaches rather than restricting attention to either one of them.

7.5 Isotropization of homogeneous plasma

We will now illustrate the general discussion of the previous section with the simplest possible case: the evolution of a homogeneous, but initially far-from-equilibrium, CFT state towards an equilibrium plasma state. The homogeneity assumption means that we work at strictly zero spatial momentum in Fourier space. Since by definition hydrodynamic modes have dispersion relations $\omega(\vec{q})$ such that $\omega \rightarrow 0$ as $q \rightarrow 0$, this implies that no hydrodynamic modes will get excited. To further simplify the physics, we will restrict ourselves to studying pure gravity in AdS_5 , which is a consistent truncation of type IIB supergravity on $\text{AdS}_5 \times S^5$. On the CFT side this simplification amounts to focusing on a sector of the dynamics in which the stress tensor is the only operator that acquires a nonzero expectation value. All other operators have vanishing one-point functions in this sector. Higher order correlation functions may be non-trivial and they could be computed in principle, but we will focus on the one-point function of the stress tensor.

In order to “create” a far-from-equilibrium state, the authors of Ref. [290] turn on a time-dependent, anisotropic source for the stress tensor of the boundary gauge theory, following the general strategy that we have illustrated in Fig. 7.2b. In other words, they turn on a non-normalizable mode of the bulk metric such that the metric of the gauge theory takes the form

$$ds^2 = -dt^2 + e^{B_0(t)} d\mathbf{x}_T^2 + e^{-2B_0(t)} dx_L^2, \quad (7.63)$$

where $\mathbf{x}_T = \{x_1, x_2\}$ are referred to as the transverse directions and x_L is referred to as the longitudinal direction. There is rotational symmetry only within the transverse plane. The function $B_0(t)$ describes a time-dependent shear in the gauge theory metric and can be chosen at will. Ref. [290] chooses

$$B_0(t) = \frac{1}{2} c \left[1 - \tanh\left(\frac{t}{\Delta}\right) \right], \quad (7.64)$$

where c is a nonzero constant and Δ is the characteristic time scale of the source. At asymptotic early and late times, $t \rightarrow \mp\infty$, B_0 becomes constant and has no physical effect, since it can be simply absorbed in the metric (7.63) by a rescaling

of the x -coordinates. Over a period of time of order Δ , however, B_0 induces a time-dependent rescaling of the transverse coordinates with respect to the longitudinal one. This way of creating a far-from-equilibrium state can be thought of as analogous to placing the gauge theory in a “cosmological” background or subjecting it to a strong gravitational wave, for a finite period of time $\sim \Delta$. Once the background metric becomes flat again, one is left with the gauge theory in Minkowski space in a highly excited state which then relaxes to equilibrium in the absence of external forces. As we will see, the excited state produced in this way possesses an anisotropic stress tensor with different transverse and longitudinal pressures, $P_T \neq P_L$. The process we are interested in is the evolution of these pressures towards a common value at asymptotically late times, namely the isotropization of the plasma once the source has been turned off.

The causal structure of the diagram in Fig. 7.2b suggests that it will prove convenient to write the bulk metric in the Eddington–Finkelstein (EF) coordinates that we introduced in (7.17) and (7.18). In the present context, the bulk metric takes the form

$$ds^2 = 2dvdr - A dv^2 + \Sigma^2 (e^{-2B} dx_L^2 + e^B d\mathbf{x}_T^2) \quad (7.65)$$

in EF coordinates, with A , B and Σ in general being functions of r and v that must be chosen such that Einstein’s equations are satisfied. The coordinate v is the EF time in the bulk. As we have seen, it coincides with the gauge theory time at the boundary, which lies at $r = \infty$. Curves of constant v are infalling null geodesics from the boundary, for which r is an affine parameter. Outgoing null geodesics obey

$$\frac{dr}{dv} = \frac{A}{2}. \quad (7.66)$$

In these coordinates the equilibrium black brane solution is given by

$$A = r^2 f(r), \quad f(r) = \left(1 - \frac{r_0^4}{r^4}\right), \quad \Sigma = r, \quad B = 0. \quad (7.67)$$

This can be seen by setting

$$v = t - g(r), \quad g'(r) = -\frac{1}{A(r)}, \quad (7.68)$$

which brings the metric (7.67) to the form

$$ds^2 = r^2 (-f(r)dt^2 + dx_L^2 + d\mathbf{x}_T^2) + \frac{dr^2}{r^2 f(r)}, \quad (7.69)$$

which is familiar from (7.18), in which $z = R^2/r$.

The formulation of general relativity in which a spacetime is constructed by means of a foliation by null hypersurfaces is called “the characteristic formulation”.

The utility of this formulation lies in the fact that Einstein’s equations are integrated in from the boundary along infalling null radial geodesics. Therefore, any numerical error made at the boundary, where Einstein’s equations are singular because of the diverging conformal factor in the metric (7.69), instantaneously falls to finite r away from the singular point in Einstein’s equations. While this tames the singular point in Einstein’s equations at $r = \infty$, it does not completely ameliorate it. One must still solve Einstein’s equations very well near the boundary. Two successful approaches thus far are (i) to solve Einstein’s equations semi-analytically for r greater than some UV cut-off r_{\max} and match the semi-analytic solution onto the numerical solution at $r = r_{\max}$, as in e.g. Ref. [290], or (ii) to discretize Einstein’s equations using pseudospectral methods, as in e.g. Refs. [292, 447]. In the latter approach one can directly impose boundary conditions at $r = \infty$, as the exponential convergence of pseudospectral methods outpaces the power-law singularities in Einstein’s equations.

In the coordinates (7.65) Einstein’s equations take the nested form

$$0 = \Sigma (\dot{\Sigma})' + 2\Sigma' \dot{\Sigma} - 2\Sigma^2, \tag{7.70}$$

$$0 = \Sigma (\dot{B})' + \frac{3}{2}(\Sigma' \dot{B} + B' \dot{\Sigma}), \tag{7.71}$$

$$0 = A'' + 3B' \dot{B} - 12\Sigma' \dot{\Sigma} / \Sigma^2 + 4, \tag{7.72}$$

$$0 = \ddot{\Sigma} + \frac{1}{2}(\dot{B}^2 \Sigma - A' \dot{\Sigma}), \tag{7.73}$$

$$0 = \Sigma'' + \frac{1}{2}B'^2 \Sigma, \tag{7.74}$$

where $h' \equiv \partial_r h$ and $\dot{h} \equiv \partial_v h + \frac{1}{2}A \partial_r h$ are derivatives along ingoing and outgoing null geodesics, respectively. Equations (7.70)–(7.72) are dynamical equations, whereas Eqs. (7.73) and (7.74) are constraints. Equation (7.74) is a constraint in the familiar sense of general relativity: if it holds on a given constant-time slice then it holds at any other time by virtue of the dynamical equations. This equation therefore will constrain the possible states that we are allowed to specify on the initial-time slice in Fig. 7.2c. Equation (7.73) is a constraint in a perhaps less familiar but analogous sense: if it satisfied on a given constant- r slice then it is satisfied everywhere by virtue of the dynamical equations. In our case we will impose this constraint on the $r = \infty$ slice by imposing the boundary conditions

$$A(r, v) \simeq r^2 + \dots, \tag{7.75}$$

$$\Sigma(r, v) \simeq r + \dots, \tag{7.76}$$

$$B(r, v) \simeq B_0(v) + \dots, \tag{7.77}$$

where the dots stand for subleading terms in the large- r expansion. Substituting these expressions in the bulk metric (7.65) and dividing by the conformal factor r^2 , as usual, we see that we indeed reproduce the boundary metric (7.63).




Figure 7.3 Energy density, longitudinal pressure and transverse pressure, all divided by $N_c^2/2\pi^2$ and all in units of $1/\Delta^4$, as a function of boundary time for $c = 2$. Recall that Δ is the characteristic timescale of the source (7.64). Figure taken from Ref. [290].

The problem now reduces to integrating Eqs. (7.70)–(7.72) numerically. Once a solution is found, the boundary stress tensor can then be read off from the normalizable mode of the metric near the AdS boundary. Details on the numerical integration can be found in the original references and we will not dwell into them here. Instead, we will concentrate on describing the physical results.

The combination of homogeneity in three spatial dimensions and rotational invariance in the two-dimensional transverse plane implies that the stress tensor can be written as

$$\langle T_\nu^\mu \rangle = \frac{N_c^2}{2\pi^2} \text{diag}[\mathcal{E}(t), \mathcal{P}_L(t), \mathcal{P}_T(t), \mathcal{P}_T(t)]. \quad (7.78)$$

(Throughout the remainder of this chapter we will use \mathcal{E} and \mathcal{P} for energy densities and pressures rescaled by a factor of $N_c^2/2\pi^2$. We denote the longitudinal and transverse pressure by \mathcal{P}_L and \mathcal{P}_T ; it is also common to refer to them as \mathcal{P}_\parallel and \mathcal{P}_\perp .) Figure 7.3 shows a plot of the energy density and transverse and longitudinal pressures produced by the changing boundary geometry (7.63), with the parameter c in (7.64) chosen as $c = 2$. The energy density and pressures all begin at zero in the distant past when the system is in its vacuum state, and at late times approach thermal equilibrium values given by

$$T_\nu^\mu = \frac{\pi^2}{8} N_c^2 T^4 \text{diag}(3, 1, 1, 1), \quad (7.79)$$

where T is the final equilibrium temperature. Non-monotonic behavior is seen when the boundary geometry changes most rapidly around time zero.

Figure 7.4 displays a congruence of outgoing radial null geodesics, again for $c = 2$. The surface shading shows A/r^2 . In the SYM vacuum (i.e., at early times) this quantity equals 1, while at late times




Figure 7.4 The congruence of outgoing radial null geodesics. The boundary is at $1/r = 0$, at the top of the figure. The surface shading displays A/r^2 . The excised region is beyond the apparent horizon, which is shown by the dashed green line. The geodesic shown as a heavier black line is the event horizon; it separates geodesics which escape to the boundary from those which cannot escape. Figure taken from Ref. [290].

$$\frac{A}{r^2} = 1 - \left(\frac{r_0}{r}\right)^4. \quad (7.80)$$

In the SYM vacuum, outgoing geodesics are given by

$$\frac{1}{r} + \frac{v}{2} = \text{const.}, \quad (7.81)$$

and appear as straight lines in the early part of Fig. 7.4. In the vicinity of $t = 0$, when the boundary geometry is changing rapidly and producing infalling gravitational radiation, the geodesic congruence changes dramatically from the zero-temperature form to a finite-temperature form. As is evident from the figure, at late times some outgoing geodesics do escape to the boundary, while others fall into the bulk and never escape. Separating the “escaping” and “plunging” geodesics is one geodesic that does neither – this geodesic, shown as the black line in Fig. 7.4, defines the true event horizon of the geometry.

Excised from the plot is a region of the geometry behind the apparent horizon, which is shown by the dashed line. This excision is necessary since somewhere inside the apparent horizon there must be a singularity, and if the region of the spacetime near the singularity were included in the calculation the numerics would break down. This excision can have no consequences for physics outside the event horizon, including in particular for the near-boundary behavior of the metric that determines the boundary theory stress tensor. From the boundary point of view, it would be safe to excise the entire region of the spacetime that lies inside the event horizon. The event horizon is the null hypersurface of spacetime that separates




Figure 7.5 We plot the area of the apparent horizon and the event horizon as a function of boundary time, again for $c = 2$. Nearly all the growth of the apparent horizon area occurs in the interval $-2 < t < 0$, during which the boundary geometry is changing rapidly. In contrast, the area of the event horizon grows in the distant past long before the boundary geometry is significantly perturbed, reflecting the global nature of event horizons discussed in the text. Figure taken from Ref. [290].

those points in the spacetime that are causally connected with the boundary from those that are not. As such, the event horizon is a global, non-local, concept whose determination requires knowledge of the entire history of the spacetime. In this sense it is teleological in nature. One extreme manifestation of this in the case of interest here is that the event horizon of Fig. 7.4 extends to the infinite past beyond the time at which the boundary theory was first perturbed, as shown more clearly in Fig. 7.5. Because the location of the event horizon can only be determined after the entire calculation has been completed, it is not possible to excise the entire region of spacetime inside the event horizon as the calculation is being done. In contrast, if an apparent horizon (defined below) can be found its location can be determined at any time and it is always inside the event horizon. So, excising a region of spacetime inside the apparent horizon that includes the singularity is guaranteed to be safe, and this is what has been done in Fig. 7.4.

The event horizon is coordinate independent but is defined only globally. In contrast, the apparent horizon is a local but coordinate dependent concept. Technically it is defined as the outer-most marginally trapped surface. The reader can consult Refs. [436, 793, 83, 172] for a general technical discussion, and Refs. [173, 174, 175] for a discussion in the context of the fluid/gravity correspondence. Here we will only give a heuristic explanation. Consider a spacelike slice in a given spacetime, and a closed surface within this slice. For example, in Minkowski spacetime one may consider a constant- t three-dimensional slice, and a two-sphere within it. Now imagine constructing two new surfaces by following light rays shot both inwards and outwards from each point on the original surface. In the presence of weak or no gravity, e.g. in the case of a two-sphere

Table 7.1 *Final equilibrium temperature T and isotropization time t_{iso} (in units of $1/T$ or Δ), for various values of c . The isotropization time t_{iso} is the time after which the pressures deviate from their equilibrium values by less than 10%. Table taken from Ref. [290].*

$ c $	1	1.5	2	2.5	3	3.5	4
$T \Delta$	0.23	0.31	0.41	0.52	0.65	0.79	0.94
$t_{\text{iso}} T$	0.67	0.68	0.71	0.92	1.2	1.5	1.8
t_{iso}/Δ	3.0	2.2	1.7	1.8	1.8	1.9	1.9

in a Minkowski spacetime, the area of the surface increases along the outgoing light rays and it decreases along the ingoing ones. In contrast, if the spacetime curvature is sufficiently strong then the area may decrease along both sets of rays. In this case the surface is called a “trapped surface”. A “marginally trapped surface” corresponds to the limiting case in which the area remains constant along the outgoing direction. The apparent horizon is a local concept, but it is not coordinate independent because it depends on a choice of a specific spacelike hypersurface. The importance of the apparent horizon lies in the fact that, under certain conditions, it can be shown that it must always lie inside an event horizon, as is the case in Fig. 7.4. This means that one can safely excise all the region (or part of it, as in Fig. 7.4) inside the apparent horizon, since this will be causally disconnected from the region outside the event horizon, and in particular from the boundary.

It is worth pointing out that in a fully dynamical, far-from-equilibrium, setting neither the area density of the event horizon nor the area density of the apparent horizon correspond to an entropy density of the far-from-equilibrium matter in the boundary quantum field theory. The acausal nature of the event horizon illustrated in Figs. 7.4 and 7.5 makes it clear that its area cannot be proportional to an entropy density since if it were there would be entropy present in the quantum field theory vacuum long before the process that excites it begins. The coordinate dependence of the apparent horizon makes it clear that its area also cannot correspond to any physical observable in the boundary theory. None of this should come as a surprise, because in the quantum field theory there is in fact no notion of entropy density that is well defined far from equilibrium. The standard thermodynamic relations that determine the entropy density from the energy density and the pressure, which are well defined, are (approximately) valid only in (near) equilibrium.

Table 7.1 shows, for various values of c , the final equilibrium temperature T and a measure of the isotropization time t_{iso} . (These quantities only depend on $|c|$, not on the sign of c .) Let us define t_{iso} as the earliest time after which the transverse

and longitudinal pressures are always both equal to their final values to within 10%. When $|c| \gtrsim 2$, we find that $t_{\text{iso}} \approx 2\Delta$, while for $|c| \lesssim 2$, $t_{\text{iso}} \approx 0.7/T$. So, t_{iso} is always comparable to either 2Δ or $0.7/T$, whichever is larger. (When $|c| \sim 2$, the two quantities 2Δ and $0.7/T$ are comparable in magnitude.)

The results in Table 7.1 can be understood qualitatively with intuitive arguments. If $\Delta \ll 1/T$, the external source pumps energy into the system only during a very brief time that, in the $\Delta \rightarrow 0$ limit, cannot control t_{iso} . In this regime, both t_{iso} and the final equilibrium temperature must be determined only by (the appropriate power of) the energy density that is pumped into the system, which is controlled by $|c|$. So, on dimensional grounds, t_{iso} must be proportional to $1/T$. From a gravitational perspective, as we shall discuss further in the next section the relaxation is controlled by the lowest quasinormal mode, whose damping rate is proportional to $1/T$. In the opposite regime, where $\Delta \gg 1/T$, even though the total amount of energy density that is pumped into the system is large (because this regime is achieved when $|c|$ is large) the energy is pumped in slowly and the system can respond adiabatically to the deformation in the geometry. In this regime, the stress tensor of the boundary fluid is never far from that of an equilibrium fluid, albeit one whose temperature is changing with time. Once the source turns off, which happens after a time of order Δ , the fluid is already close to its final equilibrium state. Perhaps it takes a time of order $1/T$ to get there, but that time is much shorter than Δ . So, t_{iso} is proportional to Δ . From a gravitational perspective, this adiabatic behavior arises because as we have discussed in Section 6.4 the relaxation times of non-hydrodynamic quasinormal modes are proportional to the inverse of the one-fourth power of the local energy density and hence vanish when $|c| \rightarrow \infty$. These qualitative considerations provide a complete understanding of the physics behind the result of the full calculation, namely that t_{iso} goes from $\sim 0.7/T$ to $\sim 2\Delta$ as a function of increasing $|c|$, but of course they do not give us the factors of 0.7 or 2.

Although the setting we have analyzed here is quite far from that in a heavy ion collision, it is interesting to note that $t_{\text{iso}} \approx 0.7/T$ corresponds to a time $\sim 0.3 \text{ fm}/c$ when $T = 500 \text{ MeV}$. This is about a factor of two faster than the upper bounds on the thermalization times inferred from hydrodynamic modeling of RHIC collisions [543, 441]. Reference [735] provides a recent example of such modeling that indicates that if the equilibration time in a RHIC collision were as short as 0.4 fm, the equilibration temperature would be just above 500 MeV. We shall return to such comparisons later, after we have seen holographic calculations of the equilibration process starting from many more, different, far-from-equilibrium initial states. Drawing conclusions from the results in this section alone would be hard to justify, but these results do already hint that equilibration times in heavy ion collisions may be longer than they would be if the physics were strongly coupled from start to finish.




Figure 7.6 Close-limit approximation for the collision of two black holes.

7.6 Isotropization of homogeneous plasma, simplified

In the previous section, we studied the isotropization of homogeneous plasma by solving the non-linear Einstein's equations in the presence of an external force, as depicted in Fig. 7.2b. The purpose of this section, in which we follow Ref. [447] closely, is two-fold. First, to perform a similar analysis in the absence of external forces, as depicted in Fig. 7.2c. Second, to show that the problem can be dramatically simplified by linearizing Einstein's equations. This simplification will allow us to analyze many possible initially far-from-equilibrium states.

Inspiration for this simplification comes from the so-called “close-limit approximation” [697] in the context of black hole mergers in four-dimensional general relativity in asymptotically flat spacetime – see Fig. 7.6 (left). If the impact parameter is small enough, then a single common horizon forms around the two incident black holes when they are close enough. At much later times the system will settle down to a single black hole in equilibrium. The close-limit approximation is the statement that the evolution of the initial horizon, from the moment it forms until the system settles down to its final state, is described well by Einstein's equations linearized around the final equilibrium black hole – see Fig. 7.6 (center). This is quite surprising because *a priori* one might have expected that the initial common horizon could not in general be viewed as a small perturbation of the horizon of the final equilibrium black hole. Yet, the close-limit approximation predicts with good accuracy, in particular, the form of the gravitational radiation emitted to infinity in the merger and ring-down phases of the collision [57], depicted in Fig. 7.6 (right). The direct analogue of this radiation in our case will be the holographic stress tensor determined from the metric near the boundary of AdS.

We will study the isotropization of a large number of anisotropic initial states in the absence of external sources. Each state will be specified on the gravity side by an entire function on the initial-time slice shown in Fig. 7.2c, and hence it will be characterized by an arbitrary number of scales. Conservation of the stress tensor for homogeneous plasma in the absence of external sources implies that the energy density \mathcal{E} (but not the entropy density) must be constant in time. Since in a homogeneous situation an equilibrium state is completely characterized by its

energy density, this means that the final state is known without solving for the dynamical evolution. On the CFT side, it is the homogeneous, isotropic plasma with an energy density \mathcal{E} equal to the initial energy density, with a pressure given by $\mathcal{E}/3$ and with a temperature proportional to $\mathcal{E}^{1/4}$. On the gravity side it is a static, isotropic black brane with the same temperature. This *a priori* knowledge of the final state makes the linear approximation particularly simple: we will linearize Einstein's equations around the static black brane (7.67) and use them to evolve each initial state. As expected on general grounds, the dynamical evolution shows that an event horizon (but not necessarily an apparent horizon) is already present on the initial-time slice for each of the states we consider. By comparing the full numerical evolution on the gravity side with its linear approximation, we will see that the latter predicts the time evolution of the CFT stress tensor with surprising accuracy (see [119] for related observations), in analogy with the prediction of the gravitational radiation at infinity by the close-limit approximation. As in that case, we emphasize that the applicability of the linear approximation is not guaranteed *a priori*, since in general our initial states will not be near-equilibrium states.

Let us now be more precise about the specification of the initial states. In the absence of sources the asymptotic form of the metric functions takes the form²

$$A = r^2 + \frac{a_4}{r^2} - \frac{2b_4(v)^2}{7r^6} + \dots, \quad (7.82)$$

$$B = \frac{b_4(v)}{r^4} + \frac{b'_4(v)}{r^5} + \dots, \quad (7.83)$$

$$\Sigma = r - \frac{b_4(v)^2}{7r^7} + \dots. \quad (7.84)$$

As usual, the normalizable modes a_4 and $b_4(v)$ are not determined by the boundary conditions but must be read off from a full bulk solution that is regular in the interior. These modes are dual to the expectation value of the stress tensor (7.78). In the absence of external sources, the energy density is constant and conservation and tracelessness of the stress tensor imply that the two pressures in (7.78) may be written as

$$\mathcal{P}_L(t) = \frac{1}{3}\mathcal{E} - \frac{2}{3}\Delta\mathcal{P}(t), \quad (7.85)$$

$$\mathcal{P}_T(t) = \frac{1}{3}\mathcal{E} + \frac{1}{3}\Delta\mathcal{P}(t), \quad (7.86)$$

in terms of \mathcal{E} and a single function $\Delta\mathcal{P} \equiv \mathcal{P}_T - \mathcal{P}_L$ that measures the degree of anisotropy. For the specific case of strongly coupled $SU(N_c)$ $\mathcal{N} = 4$ super-Yang–Mills theory at large N_c , this relation is

² The case with sources is discussed in Ref. [290].

$$\mathcal{E} = -\frac{3a_4}{4}, \quad \Delta\mathcal{P}(t) = 3b_4(t). \quad (7.87)$$

Note that, although \mathcal{E} is constant in time, a physical temperature can only be assigned to the system once (near) equilibrium is reached, in which case $\mathcal{E} = 3\pi^4 T^4/4$.

As we mentioned above, Eq. (7.74) is a constraint on the possible initial states because it relates two of the metric functions on the initial-time slice. We choose B as the independent variable because it is directly related to the CFT anisotropy through Eq. (7.87). Thus each initial state is specified by a constant a_4 and a function of the radial coordinate $B(v = 0, r)$. Note that for positive Σ the constraint (7.74) implies $\Sigma'' \leq 0$, which in combination with the asymptotic behavior $\Sigma \simeq r$ means that Σ will vanish at some $r \geq 0$ on the initial-time slice. Generically this corresponds to a curvature singularity. However, for all the initial states which the numerical code of Ref. [447] was able to evolve in a stable manner, the region where $\Sigma = 0$ is hidden behind an event horizon and hence it has no effect on the physics.

Upon considering small fluctuations around the equilibrium black brane solution (7.67), one finds that A and Σ are unmodified at linear order whereas the B fluctuation obeys Eq. (7.71) with Σ and A as in (7.67). Thus, in order to determine the evolution of the boundary stress tensor in the linear approximation we only need to solve a linear equation for B . At this order the position and the area of the horizon are unmodified. The leading correction to these quantities is obtained by solving the respective linear equations for A and Σ with a source that is quadratic in the leading solution for B .

The authors of Ref. [447] considered around 1000 initial states, for all of which the numerical code converged nicely. Most of these states were generated by taking the ratio of two tenth-degree polynomials in r with randomly generated coefficients and subtracting from them the appropriate powers of r to ensure the correct near-boundary asymptotics. A few other states were constructed “by hand” to ensure qualitative differences between them by requiring that the initial B be localized at different positions along the radial direction, that it be quickly oscillating, etc. For some profiles, an apparent horizon was present on the initial-time slice. For others, it was not. In order to evaluate the accuracy of the linearized analysis, the authors of Ref. [447] first determined the time evolution of each state by solving the full, nonlinear, Einstein’s equations. They then solved the linear equation for B , again for each of the ~ 1000 initial states. In each case, the pressure anisotropy was read off by extracting $b_4(t)$ from the near-boundary behavior (7.82). The upper panel of Fig. 7.7 shows the result obtained by solving the full Einstein’s equations for a representative initial state. The lower panel in Fig. 7.7 shows the difference between the full solution and the result obtained via the linear approximation for

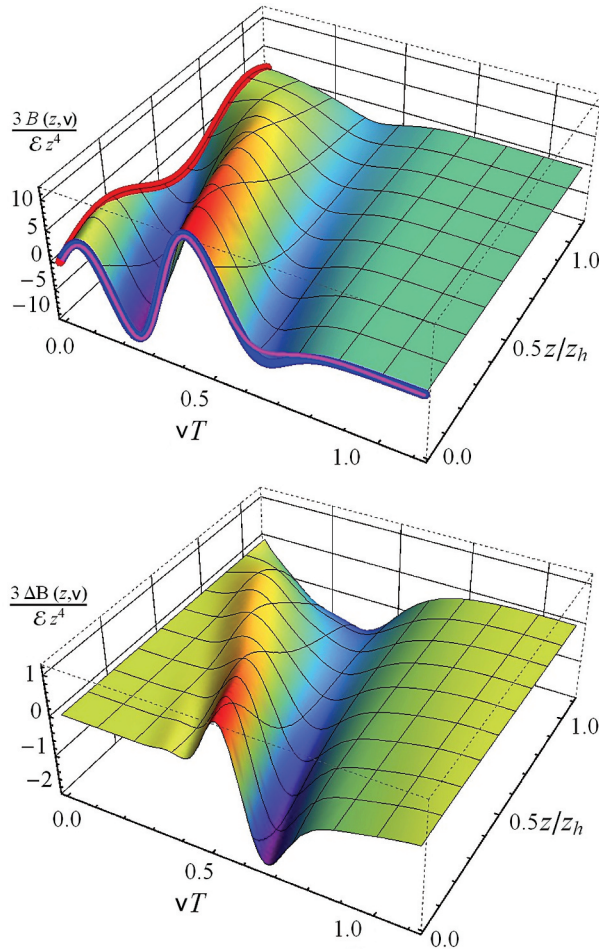


Figure 7.7 Upper panel: solution $B(v, z)$ (with $z \equiv 1/r$) obtained from the full Einstein's equations. The intersection of the surface with the $v = 0$ plane is the initial profile $B(v = 0, z) = \frac{4}{5}(z/z_h)^4 \sin(8z/z_h)$, and is shown there as a thick red curve. The intersection of the surface with the $z = 0$ plane corresponds to $\Delta\mathcal{P}(t)/\mathcal{E}$ as obtained from the full Einstein's equations, and is shown as a thick blue curve. The thin purple curve on the same plane shows the value of $\Delta\mathcal{P}(t)/\mathcal{E}$ as obtained from the linear approximation. Lower panel: difference between the full solution and the linear approximation. As evidenced by the thick and thin curves at $z = 0$ in the upper panel, this difference is small, so the scale on the vertical axis has been stretched in order to make it visible. Figure taken from Ref. [447].

this state. The ratio in the overall scales of the vertical axes in the plots, $2/10$, gives a rough estimate of the accuracy of the linear approximation, namely 20%, which is remarkable given that the evolution is definitely far-from-equilibrium. This feature is illustrated by the thick blue curve at $z = 0$ in the upper panel of Fig. 7.7, which




Figure 7.8 Time evolution of the areas of the event (top curve) and apparent (red curve) horizons for the initial state of Fig. 7.7. The dot at the origin signifies that there is no apparent horizon for this state at the initial time. From that time until the start of the lower curve there is no apparent horizon within the range of the radial coordinate covered by our grid, but there could be one at a deeper position. Figure taken from Ref. [447].

shows that the pressure anisotropy is almost an order of magnitude larger than the energy density at some points during the evolution.

As in the previous section, it is interesting to examine the time evolution of the area densities of the event and apparent horizons, since these coincide with the entropy density once the system has reached equilibrium. Figure 7.8 shows that both of these quantities are larger at the end of the evolution than at the beginning, suggesting that entropy is indeed generated during the out-of-equilibrium evolution of the system in the absence of sources that we are describing in this section.

As in the previous section, we define the isotropization time t_{iso} as the time beyond which $\Delta\mathcal{P}(t)/\mathcal{E} \leq 0.1$. Figure 7.9 is a histogram that summarizes the isotropization times of the 1000 initial states. One of the horizontal axes shows the isotropization times obtained from the full evolution, Tt_{iso} , measured in units of the final temperature. The other horizontal axis shows the relative error in the determination of this quantity that is made by using the linear approximation, namely the difference between t_{iso} as determined by the full Einstein's equations and by the linear approximation. The height of each bar indicates the number of states in each bin. We see that isotropization times are typically $t_{\text{iso}} \lesssim 1/T$, with T the final temperature, although they can of course be shorter for some initial states that happen to start closer to equilibrium. A large majority of the isotropization times found in this study lie in the range $(0.6 - 1)/T$, indicating that the result $t_{\text{iso}} \sim 0.7/T$ that we found in Section 7.5 by analyzing how the system isotropizes in response to a single family of sources is representative.

Figure 7.9 also shows that the linear approximation works with an accuracy of 20% or better for most states. Inspection “by hand” of the cases where our criterion suggests that the approximation is working less well indicates that in fact it works remarkably well even for these initial states. This is illustrated in Fig. 7.10




Figure 7.9 Results for the isotropization times obtained from the full evolution of 1000 initial states, and for the differences between the full and the linearized evolution (normalized by the full isotropization time). The height of each bar indicates the number of states in each bin. Figure taken from Ref. [447].

which shows the time evolution of the pressure anisotropy for a state in the bin marked with an arrow in the histogram of Fig. 7.9, where the linear approximation seems to be making a relatively large (around 65%) error in the final t_{iso} . We see that the linear approximation (thin curve) follows the exact evolution (thick curve) very closely indeed on the scale of the initial anisotropy. However, the fact that our isotropization criterion makes no reference to this scale means that a late-time deviation that is tiny ($\sim 1/30$) on this scale translates into an error that our isotropization criterion counts as large. One could develop an improved criterion, but Fig. 7.9 already makes the points we need to make.

The fact that the linear approximation works fairly well for such a large number of far-from-equilibrium states is surprising. Of course, it is well known that small perturbations around equilibrated plasma can be described in linear-response theory. Equivalently, small perturbations around the dual horizon can be described by linearizing Einstein's equations around the equilibrium black hole solution. This means that, for a homogeneous but anisotropic perturbation, one may expect the linear approximation to be applicable whenever $\Delta\mathcal{P}/\mathcal{E} \ll 1$. What is remarkable is that, in a strongly coupled CFT with a gravity dual, the linear approximation actually works fairly accurately for perturbations that are far larger, even with $\Delta\mathcal{P}$ an order of magnitude larger than \mathcal{E} .

We have focused on predicting the expectation value of the holographic stress tensor. Since this is read off from the normalizable mode of the metric near the boundary, it is the direct analog of the wave-form computed in the close-limit




Figure 7.10 Time evolution of the pressure anisotropy for a state in the bin marked with an arrow in the histogram of Fig. 7.9. The lower plot zooms in around the isotropization time. The thick blue curve corresponds to the exact evolution. The thin red curve corresponds to linear approximation. The horizontal green lines lie at $\Delta\mathcal{P}(t)/\mathcal{E} = \pm 0.1$. Figure courtesy of the authors of Ref. [449].

approximation in calculations of black hole mergers in asymptotically flat space-time, as in Fig. 7.6. Note, though, that our results indicate that the linear approximation in AdS works not only asymptotically but also deep within the bulk, as illustrated by Fig. 7.7.

As we have seen in Section 6.4, in Fourier space one may distinguish between hydrodynamic quasinormal modes with dispersion relations $\omega(q)$ such that $\omega \rightarrow 0$ as $q \rightarrow 0$, and nonhydrodynamic quasinormal modes (QNMs), for which $\omega(0) \neq 0$. If a perturbation is anisotropic but homogeneous then the relaxation back to equilibrium involves exclusively the non-hydrodynamic QNMs. In this sense the dynamics that we have studied in this section can be thought of as the far-from-equilibrium dynamics of the non-hydrodynamic QNMs. If the perturbation is small then these modes evolve towards equilibrium linearly, independently of each other and on a time scale set by the imaginary parts of their frequencies. One could

imagine extending the description to not-so-small perturbations by including nonlinearities in the form of interactions between the QNMs, but naively one would expect this effective description to break down for order-one anisotropies. Instead, the results of this section imply that, for homogeneous, strongly coupled plasma with a gravity dual, the isotropization process is still reasonably well described by nonhydrodynamic QNMs that evolve approximately linearly and independently of each other, even in the presence of large anisotropies (see Ref. [119] for related observations). This can be verified explicitly by expanding and evolving B in terms of a sufficient number of QNMs. Figure 7.11 shows a comparison for several initial states between the time evolution of the stress tensor as determined by the full nonlinear evolution (thick blue curves), by the full linear evolution (dotted red curves) and by the linear evolution truncated to a few QNMs (thin purple curves). For each plot, we have indicated the factor by which the area density of the event horizon increases throughout the evolution. The fact that in some cases this factor can be as large as $A_{\text{fin}}/A_{\text{ini}} \sim 25$ is another indication that we are considering initial states that are far from equilibrium.

The fact that the evolution is well described by QNMs means that, just as in the near-equilibrium case, the relaxation towards equilibrium is characterized by a few non-hydrodynamic quasinormal frequencies. In particular, a naive (under)estimate of the isotropization time can be obtained from the imaginary part of the lowest non-hydrodynamic quasinormal frequency, as in the top panel of Fig. 6.3 in Section 6.4, and is given by $\text{Im } \omega_0 \simeq -8.5T$. Since our initial states typically have anisotropies of the order of $1 \lesssim (\Delta\mathcal{P}/\mathcal{E})_{\text{ini}} \lesssim 20$, requiring

$$\left(\frac{\Delta\mathcal{P}}{\mathcal{E}}\right)_{\text{ini}} \exp\left(\text{Im } \omega_0 t_{\text{iso}}\right) \lesssim 0.1 \quad (7.88)$$

gives $0.27 \lesssim T t_{\text{iso}} \lesssim 0.62$. The reason why this may be an underestimate is that the degree of anisotropy carried by each individual QNM can be much larger, typically as large as $\Delta\mathcal{P}/\mathcal{E} \sim 500$, with the total anisotropy being much smaller due to cancellations among different modes. Assuming $(\Delta\mathcal{P}/\mathcal{E})_{\text{ini}} \simeq 500$ one gets $T t_{\text{iso}} \sim 1$.

Intuitively (and very crudely) the applicability of the linear approximation seems to be related to the fact that any nonlinearities generated by the Einstein equations are quickly absorbed by the horizon. This suggests that the linear approximation may be applicable to more general situations than the very simple one considered here, in particular to situations in which the final state is not known *a priori* and/or in which hydrodynamic modes become excited. The same intuition also suggests that the linear approximation should not be applicable to the description of strong gravitational dynamics in the absence of horizons. In particular, it is not expected to describe the formation of a horizon. Yet, as we have seen, it can be very useful indeed in describing its subsequent evolution. From a practical viewpoint this

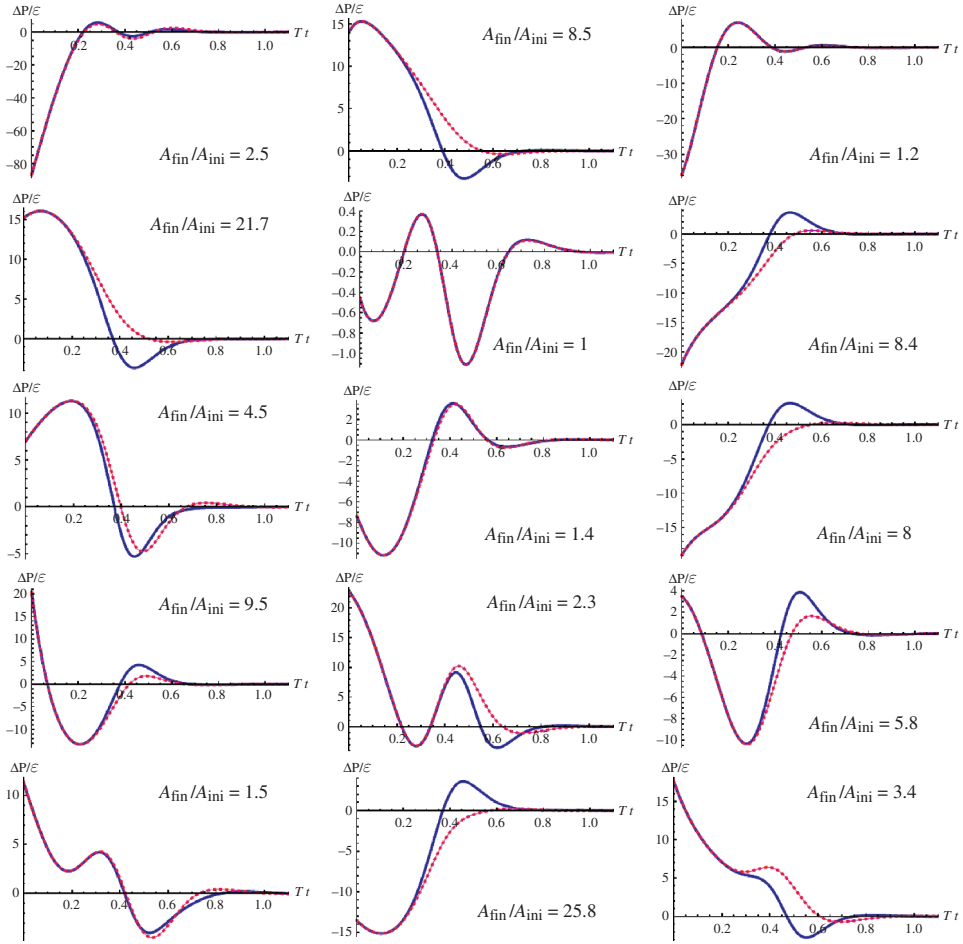


Figure 7.11 Time evolution of the pressure anisotropy for several initial states as determined by the full nonlinear evolution (thick blue curves), by the full linear evolution (dotted red curves) and by the linear evolution truncated to a few QNMs (thin purple curves). The latter two are so similar that in all the panels they appear on top of each other, differing by less than the width of the dotted curves. For each plot we have indicated the factor by which the area density of the event horizon increases throughout the evolution. Figure courtesy of the authors of Ref. [449].

is because of the technical simplification at the level of solving Einstein’s equations. However, the real power of this approximation lies on a more conceptual level, since it implies that the superposition principle applies. In our case, for example, this means that the evolution of an initial profile that takes the form $B(r) = \sum_n B_n(r)$ is given by the sum of the evolutions of each of the $B_n(r)$.

We close this section with a comment on the accuracy of the linear approximation. If one is interested in doing precision physics in a specific CFT with a known gravity dual, then 20% accuracy may not be good enough. However, if the goal

is to learn robust lessons that might be extrapolated to the far-from-equilibrium dynamics of real world QGP produced in heavy ion collisions, then this accuracy is quite likely sufficient. The more important question is how to apply it in a setting in which two zero-temperature objects collide and plasma results, meaning that on the gravity side a horizon forms during the far-from-equilibrium evolution. Even if the linear approximation can only be applied starting after the horizon has formed, because of the insights that it yields and because the simplification that it brings makes it feasible to analyze and compare very many “initial” states with horizons it is a very important tool in our toolbox.

7.7 Hydrodynamization of boost-invariant plasma

The homogeneity assumed in previous sections provides a dramatic simplification of Einstein’s equations in AdS₅, since it reduces the generic problem of a 4 + 1-dimensional evolution to one in 1 + 1 dimensions because the problem in the boundary theory is 0 + 1-dimensional. The principal drawback of this assumption is that it freezes the hydrodynamic modes, since at $q = 0$ (i.e. in homogeneous plasma) they must have $\omega = 0$. In any more generic setting, it is the hydrodynamic modes with small but nonzero q and ω that would actually dominate the late-time dynamics.

The first step in relaxing the assumption of homogeneity is to allow the dynamics of the fluid to depend on one spatial coordinate – that we shall call the longitudinal direction – while maintaining translation invariance in the other two – transverse – coordinates. Generically, this makes the hydrodynamic problem in the boundary theory 1 + 1-dimensional and so makes the gravitational calculation in AdS₅ 2 + 1 dimensional. We shall describe such a calculation in Section 7.8. In this section we shall make the further simplifying assumption that the longitudinal expansion is boost invariant, a simplification that has been used in many hydrodynamic analyses since it was introduced by Bjorken more than 30 years ago [165]. This has the great technical advantage of keeping the boundary theory problem effectively 0 + 1-dimensional, and therefore keeping the dual gravitational problem that must be solved in AdS₅ 1 + 1-dimensional, while nevertheless introducing hydrodynamic expansion. This is the simplest possible way of unfreezing the hydrodynamic modes, permitting the study of their far-from-equilibrium evolution and equilibration. In this set-up we will be able to see the transition between an early, far-from-equilibrium, phase of the dynamics when the boost invariant expansion does not satisfy hydrodynamic equations of motion and a late phase when the expansion becomes a conventional Bjorken flow solution of hydrodynamics. In other words, we will see the fluid “hydrodynamize” as it expands. The gravity solution that describes boost-invariant hydrodynamic expansion can be understood analytically [495, 642, 493, 445, 530], so we will begin our presentation with this late-time dynamics.

7.7.1 Boost-invariant hydrodynamics

In order to describe boost-invariant dynamics, it is convenient to introduce proper time τ and spacetime rapidity y coordinates through

$$t = \tau \cosh y, \quad x_L = \tau \sinh y. \quad (7.89)$$

In these coordinates the Minkowski metric takes the form

$$ds^2 = -d\tau^2 + \tau^2 dy^2 + dx_T^2 \quad (7.90)$$

and boost transformations act as a shift of y , so boost invariance is simply the statement that physical quantities are independent of y . Since we will retain homogeneity in the transverse plane, physical observables will only depend on τ . This dependence leads to nonzero gradients with respect to the Cartesian coordinate x_L and therefore to hydrodynamic behavior.

Before we begin the calculation, it is important to provide some context for the assumption of boost invariance in heavy ion collisions. If the hydrodynamic expansion of the fluid produced in a heavy ion collision were independent of the spacetime rapidity y , then after this fluid hadronizes the distribution of the momenta of the final state hadrons would be independent of the momentum space rapidity $\frac{1}{2} \ln((E + p_L)/(E - p_L))$, introduced in Section 2.1 and also conventionally denoted by y . This can be argued for on symmetry grounds, and therefore arises in any standard algorithm for relating the stress tensor of the hydrodynamic fluid to the momenta of final state particles, the simplest of which can be found in Ref. [299]. Of course, no finite energy collision can yield a flat rapidity distribution extending from $y = -\infty$ to $y = +\infty$, but if the stress tensor of the hydrodynamic fluid is independent of the spacetime y over some wide but finite range of y then the distribution of final state hadrons in momentum-space y will feature a broad flat plateau. As we mentioned in Section 2.1, however, we now know from data that in heavy ion collisions at RHIC the charged particle rapidity distribution dN_{ch}/dy does not have this shape: it looks roughly Gaussian [762], meaning that if there is a plateau around $y \sim 0$ it is relatively narrow.

Because the data require it, nowadays the state of the art hydrodynamic calculations that seek to describe the matter produced in heavy ion collisions as in Section 2.2 describe droplets that expand with non-trivial dependence on all spatial coordinates. Because these 3 + 1-dimensional relativistic viscous hydrodynamics calculations are challenging, though, many authors begin by assuming boost invariant longitudinal expansion and focusing on the hydrodynamic expansion in the transverse directions. And, as we shall see below, if the initial conditions for hydrodynamics are boost invariant then the hydrodynamic evolution remains boost invariant, regardless of how the fluid flows in the transverse dimensions. We shall

initially simplify even further, assuming translation invariance in the transverse plane. We do so just for the purpose of making our calculations tractable and in fact, as we shall see in Section 7.7.5, the first calculations in which the fluid is allowed to move in the transverse plane are now being done.

Early motivations for assuming boost-invariant hydrodynamics, going back to Ref. [165], were based upon a simplified picture of the dynamics of nucleus–nucleus collisions in the limit of infinite collision energy. In this limit, in the center-of-mass frame the incident nuclei are pancakes with zero thickness colliding at the speed of light. After the collision, the fragments of the nuclei themselves are assumed to stay arbitrarily close to the lightcones, at arbitrarily large positive and negative y , while the future lightcone at finite y is seen as containing particles that had been newly created in the collision. And, inspired by early data, it was assumed that particle creation would be boost invariant in the high collision energy limit.

A large body of more recent data indicates that the dynamical assumptions we have just described are not valid: dN_{ch}/dy does not feature a broad flat plateau and although the fragments of the incident nuclei, which can for example be tracked by the net proton density as in Fig. 2.6b, do end up on average at higher $|y|$ in higher energy collisions, they are present also at $y \sim 0$. These data have motivated the investigation of other simplifying assumptions for the longitudinal dynamics (going back to Landau’s assumption [565] that the incident nuclei initially stop at $y = 0$, which is in many ways the antithesis of boost invariance, and including assumptions that span the space from Landau to Bjorken [157]). For our purposes, though, what we need is a simplifying assumption within which we can study how matter that is initially far-from-equilibrium hydrodynamizes as it expands, and the best choice for this specific purpose is the assumption of boost invariance.

Using the coordinates (7.90), it is straightforward to derive the equations of motion for boost invariant hydrodynamics in any conformal plasma. Assuming for simplicity $y \rightarrow -y$ symmetry, the stress tensor has only three nonzero components $T_{\tau\tau}$, T_{yy} and $T_{x_2x_2} = T_{x_3x_3} \equiv T_{xx}$. Since we are dealing with a conformal gauge theory, $T_{\mu\nu}$ is traceless:

$$-T_{\tau\tau} + \frac{1}{\tau^2}T_{yy} + 2T_{xx} = 0. \quad (7.91)$$

Energy–momentum conservation $\nabla_\mu T^{\mu\nu} = 0$ gives a second relation among the components:

$$\tau \frac{d}{d\tau} T_{\tau\tau} + T_{\tau\tau} + \frac{1}{\tau^2} T_{yy} = 0. \quad (7.92)$$

Using the relations (7.91) and (7.92), all components of the energy momentum tensor can be expressed in terms of the time-dependent energy density $\mathcal{E}(\tau)$:

$$T_{\tau\tau} \equiv \mathcal{E}, \quad T_{yy} = -\tau^3 \mathcal{E}' - \tau^2 \mathcal{E}, \quad T_{xx} = \mathcal{E} + \frac{1}{2} \tau \mathcal{E}', \quad (7.93)$$

where $\mathcal{E}' \equiv d\mathcal{E}/d\tau$ and where we have now scaled a factor of $N_c^2/(2\pi^2)$ out of all components of the stress tensor so that \mathcal{E} here is defined as in (7.78). Note that all these conditions are purely kinematical in nature. The dynamics of the theory will choose a specific $\mathcal{E}(\tau)$.

At sufficiently late times we expect (and shall confirm below) that, owing to the continued expansion of the fluid, spatial gradients will decrease in magnitude, making all viscous effects less and less important. We therefore expect the dynamics to approach that of a perfect, inviscid, fluid. In particular, it should become locally isotropic in the local rest frame. In terms of the pressures this means that $P_L = T_y^y = P_T = T_x^x$. Using (7.93), this translates into the differential equation

$$-\tau \mathcal{E}' - \mathcal{E} = \mathcal{E} + \frac{1}{2} \tau \mathcal{E}' \quad (7.94)$$

for the energy density. This equation has the simple and well-known ‘‘Bjorken solution’’ $\mathcal{E} \propto \tau^{-4/3}$ [165] where, for later convenience, we shall write the proportionality constant as

$$\mathcal{E} = \mathcal{E}_0 \frac{\Lambda^4}{(\Lambda\tau)^{4/3}}, \quad (7.95)$$

where \mathcal{E}_0 is defined by the relationship between the energy density of the conformal plasma in local thermal equilibrium and its temperature, $\mathcal{E} = \mathcal{E}_0 T^4$, meaning that $\mathcal{E}_0 = 3\pi^4/4$ in strongly coupled $\mathcal{N} = 4$ SYM theory. (Recall from (7.78) that the energy density is given by the scaled energy density \mathcal{E} multiplied by a factor of $N_c^2/(2\pi^2)$.) We have introduced the integration constant Λ , with dimensions of energy, that specifies a particular solution. We shall further interpret Λ below. Substituting into the expressions for the pressures we see that this leads to the conformal equation of state $\mathcal{E} = 3P$. Since $\mathcal{E} = \mathcal{E}_0 T^4$, the temperature decreases at late times when the expansion is described by ideal hydrodynamics according to

$$T = \frac{\Lambda}{(\Lambda\tau)^{1/3}}. \quad (7.96)$$

We now see that we have defined Λ such that when $\tau = 1/\Lambda$ the temperature is given by Λ . Equivalently, $T = 1/\tau$ at the time when $\tau = 1/\Lambda$. The boost invariant hydrodynamic expansion of a perfect liquid is thus fully specified by the value of Λ , with solutions with larger Λ being those in which $T(\tau)\tau$ reaches 1 at earlier τ . Finally, the entropy density scales as

$$s \sim T^3 \sim \frac{\Lambda^3}{\Lambda\tau}. \quad (7.97)$$

Since the volume element of the metric (7.90) grows as τ , it follows that the total entropy remains constant, as expected for a fluid with zero viscosity.

Recall that the full solution to the equations describing boost-invariant expansion starting from any arbitrary boost-invariant initial state is expected to behave as an ideal fluid, as above, as $\tau \rightarrow \infty$, since at asymptotically late times gradient corrections should become negligible and the system should approach local equilibrium. At late but finite times, however, viscous effects produce corrections to the leading behavior (7.95). In this way the hydrodynamic expansion becomes a late-time expansion in powers of $\tau^{-2/3}$. This power can be understood from the fact that the hydrodynamic expansion is controlled by the product $T^{-1}\nabla$, with $T \sim \tau^{-1/3}$ and the size of gradient corrections being $\nabla \sim 1/\tau$. For example, including the first and second order hydrodynamic corrections the stress-energy tensor that describes the boost-invariant hydrodynamic expansion of a conformal fluid takes the form [107, 291]

$$\begin{aligned}\mathcal{E} &= \frac{\mathcal{E}_0\Lambda^4}{(\Lambda\tau)^{4/3}} \left[1 - \frac{2\eta_0}{(\Lambda\tau)^{2/3}} + \frac{C}{(\Lambda\tau)^{4/3}} + \dots \right], \\ \mathcal{P}_T &= \frac{\mathcal{E}_0\Lambda^4}{3(\Lambda\tau)^{4/3}} \left[1 - \frac{C}{(\Lambda\tau)^{4/3}} + \dots \right], \\ \mathcal{P}_L &= \frac{\mathcal{E}_0\Lambda^4}{3(\Lambda\tau)^{4/3}} \left[1 - \frac{6\eta_0}{(\Lambda\tau)^{2/3}} + \frac{5C}{(\Lambda\tau)^{4/3}} + \dots \right].\end{aligned}\quad (7.98)$$

The constant η_0 is related to the shear viscosity of the plasma through $\eta = \eta_0\mathcal{E}_0T^3$. The constant C is related to second order hydrodynamic relaxation times. In the plasma of strongly coupled $\mathcal{N} = 4$ SYM theory [107, 530],

$$\eta_0 = \frac{1}{3\pi}, \quad C = \frac{1 + 2\ln 2}{18\pi^2}.\quad (7.99)$$

We see from (7.98) that, as expected, at later and later times the gradient terms through which the effects of viscosity and higher order corrections to ideal hydrodynamics enter become less and less important.

If the expansion (7.98) were extended to include terms that are higher and higher order in $\tau^{-2/3}$, more and more coefficients that characterize the static plasma (like η_0 and C in (7.98)) would appear but the solution itself would still be specified just by the single parameter Λ . The late-time hydrodynamic behavior of a boost invariant expansion starting from any arbitrary boost invariant initial condition must be of the form (7.98) for some value of Λ . This means that (7.98) is a scaling solution, in the sense that for any value of Λ there are many different far-from-equilibrium initial conditions that will evolve into the same form (7.98) at late time. As the fluid expands, it loses the memory of the details of its initial conditions. And, if all

one knows is the late-time expansion (7.98) it is impossible to run the clock backwards and reproduce the initial conditions from which the late-time state (7.98) was obtained. In all these ways, the hydrodynamization of an expanding boost invariant fluid is analogous to the equilibration that we described in Sections 7.5 and 7.6.

7.7.2 Late-time gravity solution

In this section, we follow Ref. [495]. These authors addressed the question of whether the ideal-fluid behavior that is expected at late times on general grounds arises dynamically from Einstein's equations. For this purpose they considered the most general 4 + 1-dimensional bulk metric allowed by the symmetries of the problem:

$$ds^2 = \frac{1}{z^2} \left[-e^{a(\tau,z)} d\tau^2 + \tau^2 e^{b(\tau,z)} dy^2 + e^{c(\tau,z)} dx_{\perp}^2 \right] + \frac{dz^2}{z^2}, \quad (7.100)$$

where we are now using the radial coordinate z for which the boundary lies at $z = 0$. They began their analysis by allowing for a general form of the energy density $\mathcal{E} \sim 1/\tau^\alpha$, although α is constrained to lie in the range $0 < \alpha < 4$ just by the requirement that the energy density be non-negative in any frame. The question they posed was for what values of α a regular solution of the form (7.100) exists. In principle, Einstein's equations for this ansatz yield a system of coupled, nonlinear partial differential equations in two variables, which in general is intractable analytically. The insight of Ref. [495] was the realization that at late times the solution can be written in terms of a single scaling variable $z/\tau^{\alpha/4}$. This reduces Einstein's equations to ordinary differential equations for which an analytic solution valid at asymptotically late times can be found. This solution exists for any α in the range above, but the solution is free of naked curvature singularities only for $\alpha = 4/3$. For this particular value of α the solution takes the form

$$ds^2 = \frac{1}{z^2} \left[- \left(1 - \frac{z^4}{z_0^4}\right)^2 \left(1 + \frac{z^4}{z_0^4}\right)^{-1} d\tau^2 + \left(1 + \frac{z^4}{z_0^4}\right) (\tau^2 dy^2 + dx_{\perp}^2) \right] + \frac{dz^2}{z^2}, \quad (7.101)$$

with $z_0 \sim \tau^{1/3}$. This metric is boost invariant and it possesses a receding horizon³ at $z = z_0$, suggesting that it describes boost-invariant, cooling plasma with temperature $T \sim 1/z_0 \sim 1/\tau^{1/3}$, as in (7.96). Similarly, the total entropy, which is proportional to the total horizon area, scales as $s \sim \tau/z_0^3 \sim \text{const.}$, again in agreement with ideal-fluid hydrodynamics. We therefore conclude that Einstein's equations, together with the physical requirement of regularity of the gravity solution, reproduce the late-time ideal-fluid dynamics expected on general grounds.

³ Note that at the asymptotically late times at which the solution (7.101) is valid the event and the apparent horizons are expected to coincide.

The leading correction to the metric (7.101) has been computed [642], and the first correction to the energy density agrees with that in (7.98). Furthermore, the corrected metric is regular at and outside the horizon if and only if the viscosity coefficient η_0 in (7.98) is given by (7.99) [493], which is to say if and only if $\eta/s = 1/(4\pi)$ as in static plasma [690]. By computing the second order correction, the authors of Ref. [444] then determined a coefficient of second order hydrodynamics, the relaxation time. Further examination of this correction indicated the existence of a subleading logarithmic divergence that could not be cancelled regardless of the choice of transport coefficients [135], but this apparent singularity turned out to be simply an artifact of the coordinates chosen to write the metric (7.101), which are problematic for the discussion of regularity issues at the horizon [154, 153] (see Section 7.2, in particular Fig. 7.1). In Refs. [445, 530, 531] the gravity solution was constructed in the Eddington–Finkelstein coordinates, and it was found that in these coordinates the geometry is indeed regular.

In summary, not only do the Einstein equations in a dynamical setting predict the correct late-time behavior as a function of proper time, but they also predict the correct transport coefficients at any order. Note that no information about the initial conditions of the plasma is necessary to derive these late-time features. The reason is that this information is dissipated along the flow, so the behavior at late times is universal. All the information about the initial conditions is encoded in the single dimensionful constant Λ . This loss of memory of details about the initial conditions is behind the existence of a scaling solution at late times. At early times the dynamics is strongly dependent on the initial conditions and no universal solution exists [151]. In order to connect the early-time dynamics and the late-time behavior, a solution valid at all times is needed. This can be constructed by solving Einstein's equations numerically, as we now describe.

7.7.3 Full gravity solution

In this section we follow Ref. [291] closely. The strategy is the same as in Section 7.5, namely to create a far-from-equilibrium state by acting on the CFT vacuum with an external source during a finite period of time. In the present case we are interested in replacing homogeneity along the longitudinal direction by boost invariance, so the boundary metric (7.63) gets replaced by

$$ds^2 = -d\tau^2 + e^{\gamma(\tau)} dx_T^2 + \tau^2 e^{-2\gamma(\tau)} dy^2. \quad (7.102)$$

The function $\gamma(\tau)$ characterizes a time-dependent shear in the boundary metric which serves to excite the CFT from its vacuum to some far-from-equilibrium state. The authors of Ref. [291] choose




Figure 7.12 A spacetime diagram depicting several stages of the evolution of the field theory state in response to the changing spatial geometry. Figure taken from Ref. [291].

$$\gamma(\tau) = c \Theta \left(1 - \frac{(\tau - \tau_0)^2}{\Delta^2} \right) \left[1 - \frac{(\tau - \tau_0)^2}{\Delta^2} \right]^6 \times \exp \left[-\frac{1}{1 - \frac{(\tau - \tau_0)^2}{\Delta^2}} \right], \quad (7.103)$$

where Θ is the unit-step function. Inclusion of the $[1 - (\tau - \tau_0)^2/\Delta^2]^6$ factor makes the first few derivatives of $\gamma(\tau)$ better behaved as $\tau - \tau_0 \rightarrow \pm\Delta$. The function $\gamma(\tau)$ has compact support and is infinitely differentiable; $\gamma(\tau)$ and all its derivatives vanish at and outside the endpoints of the interval (τ_i, τ_f) , with $\tau_i \equiv \tau_0 - \Delta$ and $\tau_f \equiv \tau_0 + \Delta$. We choose $\tau_0 \equiv \frac{5}{4}\Delta$ so the geometry is flat at $\tau = 0$. Choosing $\tau_0 \geq \Delta$ is convenient for numerics as our coordinate system becomes singular on the $\tau = 0$ lightcone. The particular choice $\tau_0 = \frac{5}{4}\Delta$ is made so that the numerical results (which begin at $\tau = 0$) contain a small interval of unmodified geometry before the deformation turns on. We choose to measure all dimensionful quantities in units where $\Delta = 1$, so $\tau_i = 1/4$ and $\tau_f = 9/4$.

Figure 7.12 shows a spacetime diagram depicting several stages in the evolution of the SYM state schematically. Hyperbolae inside the forward lightcone are constant- τ surfaces. Prior to $\tau = \tau_i$, the system is in the ground state. The region of spacetime where the geometry is deformed from flat space by the external source specified by (7.103) is shown as the red region labeled I in Fig. 7.12. At coordinate time $t = \tau_i$ the geometry of spacetime begins to deform in the vicinity of $x_L = 0$. As time progresses, the deformation splits into two localized regions centered about $x_L \sim \pm t$, which subsequently separate and move in the $\pm x_L$ directions at speeds asymptotically approaching the speed of light. After the “pulse” of spacetime deformation passes, the system will be left at $\tau = \tau_f$ in an excited, anisotropic, non-equilibrium state. That is, the deformation in the geometry will have done work on the field theory state. As the excited far-from-equilibrium state then evolves in time it is boost invariant but not hydrodynamic. This region is shown as the yellow region labeled II in Fig. 7.12. It is in this region that we can study the relaxation of a far-from-equilibrium non-equilibrium state. After some later proper time τ_{hydro} , the system will have relaxed to a point where a

hydrodynamic description of the continuing evolution is accurate. This final hydrodynamic regime, in green and labeled III in Fig. 7.12, is the regime whose dynamics is described by (7.98). As the late-time hydrodynamic solution to boost invariant flow is known analytically, we choose to define τ_{hydro} as the time after which the stress tensor coincides with the hydrodynamic approximation to better than 10%.

Our task, then, is to find τ_{hydro} and, in particular, to see how it correlates with quantities such as the energy density, from which an effective temperature can be defined through

$$\mathcal{E}(\tau) = \mathcal{E}_0 T^4(\tau), \quad (7.104)$$

where the purely numerical factor is $\mathcal{E}_0 = 3\pi^4/4$ in $\mathcal{N} = 4$ SYM theory after the rescaling (7.78). If the system is far from equilibrium there is no sense in which a temperature can be defined and the quantity $T(\tau)$ should simply be thought of as an alternative measure of the energy density. At late times, though, $T(\tau)$ approaches the local temperature in the hydrodynamic regime. We denote the effective temperature at time τ_{hydro} by $T_{\text{hydro}} \equiv T(\tau_{\text{hydro}})$. As explained at the end of Section 7.5, in the $c \rightarrow \infty$ limit the energy pumped into the system by the source (here, the source (7.103)) is large and so is T_{hydro} , meaning that $1/T_{\text{hydro}} \ll \tau_f$. In this regime, we expect that $\tau_{\text{hydro}} - \tau_f \ll \tau_f$. And, we expect the evolution throughout region II of Fig. 7.12 to be adiabatic in the sense that the non-hydrodynamic degrees of freedom remain close to equilibrium and the description of the dynamics is close to hydrodynamic (with a changing energy density) at all times. In particular, a hydrodynamic description without driving terms will be accurate the moment the geometry stops changing. Hence, in this regime one learns little about the dynamics associated with the relaxation of non-hydrodynamic modes.

More interesting is the case where the effective temperature at τ_{hydro} satisfies $1/T_{\text{hydro}} \gtrsim \tau_f$. This is the regime we will study. In this regime, the system can be significantly out of equilibrium after the source turns off at τ_f and the boundary geometry becomes flat. We will see that when this is the case the entire process of hydrodynamization occurs over a time which is less than or comparable to $1/T_{\text{hydro}}$.

Given the symmetries of the physical situation that we wish to study we can write the metric on the gravity side in the form

$$ds^2 = 2dr dv - A dv^2 + \Sigma^2 \left[e^B d\mathbf{x}_T^2 + e^{-2B} dy^2 \right], \quad (7.105)$$

where A , B , and Σ are all functions of r and Eddington–Finkelstein time v only. Note that we are back to using the radial coordinate r , with the boundary at $r \rightarrow \infty$. In these coordinates, this metric is analogous to (7.65), except that the EF time v at the boundary coincides in this case with the proper time τ of Eq. (7.102). Similarly, our task is to solve Einstein's equations (7.74) but with the boundary conditions (7.77) replaced by




Figure 7.13 The congruence of outgoing radial null geodesics. The surface shading displays A/r^2 . Before time $v_i = 1/4$, this quantity equals one. The excised region lies inside the apparent horizon, which is shown by the dashed purple line. The event horizon is shown as a solid blue curve which separates geodesics which escape to the boundary from those which cannot escape. At late times $v \gtrsim 1.5$ the event horizon coincides with the apparent horizon, and both slowly fall deeper into the bulk. This is the gravity dual of the hydrodynamic expansion of a boost invariant fluid at late times. Figure taken from Ref. [291].

$$A(r, v) \simeq r^2 + \dots, \tag{7.106}$$

$$\Sigma(r, v) \simeq r v^{1/3} + \dots, \tag{7.107}$$

$$B(r, v) \simeq \gamma(v) - \frac{2}{3} \log v + \dots, \tag{7.108}$$

where the dots stand for subleading terms in the large- r expansion. Technical details can be found in Ref. [291]. Here we will just describe the results.

On the gravity side, the results that we now describe are qualitatively the same for any c . Figure 7.13 shows a congruence of outgoing radial null geodesics for $c = 1$. The geodesics are obtained by integrating

$$\frac{dr}{dv} = \frac{1}{2} A(r, v). \tag{7.109}$$

The shaded surface in the plot displays the value of A/r^2 . Excised from the plot is a region of the geometry behind the apparent horizon, whose location is shown by the dotted line.

At times $v < v_i = 1/4$, the boundary geometry is static and $A/r^2 = 1$. The outgoing geodesic congruence at early times therefore satisfies

$$v + \frac{2}{r} = \text{const.}, \quad (7.110)$$

and hence appears as parallel straight lines on the left side of Fig. 7.13. These are just radial geodesics in AdS_5 , which is the geometry dual to the initial zero-temperature ground state. After time v_i the boundary geometry starts to change, A/r^2 deviates from unity, and the congruence departs from the zero-temperature form (7.110).

As is evident from Fig. 7.13, and just as we saw in Fig. 7.4 for the case of equilibration without expansion, at late times some geodesics escape up to the boundary and some plunge deep into the bulk. Separating escaping from plunging geodesics is precisely one geodesic that does neither. This geodesic, shown as the thick solid curve in the figure, defines the location of a null surface inside which all events are causally disconnected from observers on the boundary. This surface is the event horizon of the geometry.

After the time $v_f = 9/4$, the boundary geometry becomes flat and unchanging, no additional gravitational radiation is produced, and the bulk geometry approaches a slowly evolving form. The rapid relaxation of high frequency modes can clearly be seen in the behavior of A/r^2 shown in Fig. 7.13 – all of the high frequency structure in the plot appears only during the time interval where the boundary geometry is changing and creating gravitational radiation. Physically, the rapid relaxation of high frequency modes occurs because the horizon acts as an absorber of gravitational radiation and low frequency modes simply take more time to fall into the horizon than high frequency modes. Therefore, as time progresses the geometry relaxes onto a smooth universal form whose temporal variations become slower and slower as $v \rightarrow \infty$.

As we saw in the previous section, one can systematically construct a boost-invariant late-time solution to Einstein's equations. At leading order the solution takes the form (7.101) and is characterized by a receding horizon with approximate position $1/r_0 = z_0 \sim v^{1/3}$. As time progresses, the horizon slowly falls deeper into the bulk, and the temperature of the black hole decreases as $v^{-1/3}$. The falling of the horizon into the bulk, as an inverse power of v , is clearly visible in the calculation presented in Fig. 7.13.

We now turn to a discussion of the results for boundary field theory observables. Figure 7.14 shows plots of the energy density and transverse and longitudinal pressures sourced by the changing boundary geometry (7.102) with $c = \pm 1$ between $\tau_i = 1/4$ and $\tau_f = 9/4$ and evolving subsequently. These quantities begin at zero before time τ_i , when the system is in the vacuum state, and deviate




Figure 7.14 Energy density, longitudinal pressure and transverse pressure, all divided by $N_c^2/2\pi^2$ and all in units of $1/\Delta^4$, as functions of time for $c = -1$ (left) and $c = +1$ (right), where c and Δ were defined in (7.103). The energy density and pressures start off at zero before time $\tau_i = 1/4$, when the system is in the vacuum state. During the interval of time $\tau \in (\tau_i, \tau_f) = (1/4, 9/4)$, the boundary geometry is changing and doing work on the field theory state. After time τ_f , the deformation in the geometry turns off, the field theory state evolves, and subsequently relaxes onto a hydrodynamic description. The smooth tails at late times in both plots occur during the hydrodynamic regime. At late times, from top to bottom, the three curves (in both plots) correspond to the energy density \mathcal{E} , transverse pressure \mathcal{P}_\perp , and longitudinal pressure \mathcal{P}_\parallel . Figure taken from Ref. [291].

from zero once the boundary geometry starts to vary. During the interval of time where the boundary geometry is changing, the energy density generally grows and the pressures oscillate rapidly: work is being done by the source on the field theory state. After time τ_f , the boundary geometry becomes flat and no longer does any work on the system. As time progresses, non-hydrodynamic degrees of freedom relax and at late times the evolution of the system is governed by hydrodynamics. The late time hydrodynamic behavior manifests itself as the smooth tails appearing at late times in Fig. 7.14.

The two sets of plots in Fig. 7.14, contrasting $c = +1$ and $c = -1$, are qualitatively similar, with the main difference being the phase of the oscillations in the pressures. For example, for $c = -1$ the transverse pressure is negative at τ_f whereas for $c = +1$ the transverse pressure is positive and larger than the longitudinal pressure, which is nearly zero at τ_f . Furthermore, from the figure one sees that, for either sign of c , the transverse pressure approaches the longitudinal pressure from above, in agreement with the hydrodynamic prediction (7.98).

To facilitate a quantitative comparison between the numerical results for the stress tensor and the late-time hydrodynamic expansions, Fig. 7.15 shows the energy density and pressures for $c = 1/4, 1$ and $3/2$, with the corresponding hydrodynamic forms (7.98) plotted on top of the numerical data. The single




Figure 7.15 Energy density, longitudinal pressure and transverse pressure, all divided by $N_c^2/2\pi^2$ and all in units of $1/\Delta^4$, as functions of time for $c = 1/4$, $c = 1$ and $c = 3/2$. From top to bottom, the continuous curves are the energy density, the transverse pressure, and the longitudinal pressure. The dashed curves in each plot show the second order viscous hydrodynamic approximation (7.98) to the different stress tensor components. Figure taken from Ref. [291].

parameter Λ that specifies the hydrodynamic solution is obtained by fitting to the late time results. The plots start at time $\tau = \tau_f$. In all three plots, one clearly sees the stress-energy components approach their hydrodynamic approximations. Moreover, in all plots one sees a substantial anisotropy even at the late times where a hydrodynamic treatment is applicable. From (7.98) we see that this means that the effect of viscosity is very evident in these results. On the time scales depicted in Fig. 7.15, the boost invariant hydrodynamic expansion is described by viscous hydrodynamics, not by ideal hydrodynamics.

Looking at the right-hand panel of Fig. 7.15, we see that for $c = 3/2$ the energy density and both pressures are already quite close to their hydrodynamic values at $\tau = \tau_f$. Indeed, in this panel the time τ_{hydro} after which the full results are within 10% of their hydrodynamic values is 2.3 [291], meaning that $\tau_{\text{hydro}} - \tau_f = 0.05$. The system is almost hydrodynamic at τ_f and hydrodynamizes very soon thereafter. This reflects the fact that, as we have discussed, when $|c|$ is large the energy density, and the effective temperature $T(\tau)$ defined from it, are pumped up high enough that during the time between τ_i and τ_f when the system is being driven its nonhydrodynamic modes evolve almost adiabatically and so are hardly excited. Hence the system is already almost hydrodynamic when the source turns off.

Looking at the left-hand panel of Fig. 7.15, with $c = 1/4$, we see that at the time $\tau = \tau_f$ the transverse and longitudinal pressures are almost equal and opposite in magnitude meaning that at this time the system is very far from hydrodynamic. The curves show the system hydrodynamizing, and it turns out that $\tau_{\text{hydro}} - \tau_f = 0.85$ [291]. In the units of the figure, the effective temperature at the

time τ_{hydro} is given by $T_{\text{hydro}} = 0.27$ [291], and $\tau_{\text{hydro}} - \tau_f = 0.24/T_{\text{hydro}}$. We see that when measured in units of $1/T_{\text{hydro}}$ the time that it takes the system to hydrodynamize after the source is turned off is still very short, even though at time τ_f the system is manifestly far from hydrodynamic. It would be premature based upon this single example to conclude that $\tau_{\text{hydro}}T_{\text{hydro}}$, with τ_{hydro} the hydrodynamization time for a system whose final state is a boost invariant expanding plasma, is shorter than the $\tau_{\text{iso}}T$ that we found in our analysis in Section 7.6 of the equilibration (i.e. in that case isotropization) of a system whose final state is a static plasma with temperature T . It is true that in Section 7.6 we found $\tau_{\text{iso}}T = (0.6 - 1)$ and here we have found $(\tau_{\text{hydro}} - \tau_f)T_{\text{hydro}} = 0.24$. But, first of all, it is not clear whether in making this comparison we should or should not “count” some or all of the time between τ_i and τ_f when the boost invariant system is being driven. After all, the system is presumably already hydrodynamizing while it is being driven. And, second, we see from Fig. 7.15 that in a system with an expanding final state isotropization happens long after hydrodynamization: if we were to define a τ_{iso} from the difference between the longitudinal and transverse pressures in this section as we did in Section 7.6 it would be significantly greater than τ_{hydro} . In the present context, however, and in fact in heavy ion collisions, it is the hydrodynamization time τ_{hydro} that is of interest.

We can also use the results that we have presented here to investigate whether hydrodynamic behavior sets in when higher order terms in the hydrodynamic expansion (7.98) become comparable to lower order terms or, instead, sets in at a time determined by when it is that non-hydrodynamic quasinormal modes have damped away leaving the longer-lived hydrodynamic modes dominant. In the three panels of Fig. 7.15, from left to right, hydrodynamization occurs when $\Lambda\tau_{\text{hydro}}$ is given by 0.89, 1.9 and 2.6 [291], meaning that in all three panels $\Lambda\tau_{\text{hydro}} \gtrsim 1$. Examining the size of the coefficients in the series (7.98) shows that the second-order $(\Lambda\tau)^{-4/3}$ terms are quite small compared to the leading $(\Lambda\tau)^{-2/3}$ viscous terms when $\Lambda\tau \gtrsim 1$; they only become comparable when $\Lambda\tau \lesssim 0.05$. In other words, in the period $0.05 \lesssim \Lambda\tau \lesssim 1$ higher order terms in the hydrodynamic expansion (7.98) are small and yet hydrodynamics is still not yet applicable. This indicates that the physics which determines the onset of hydrodynamic behavior is not associated with higher order terms in the hydrodynamic expansion becoming comparable to lower order terms. Rather, it must be the case that the expansion hydrodynamizes at $\tau \sim \tau_{\text{hydro}}$ because it is at that time that the non-hydrodynamic modes that are not described at all by (7.98) and that damp away exponentially in time are becoming insignificant relative to the slowly relaxing hydrodynamic modes. If $|c|$ is not large, when $\tau < \tau_{\text{hydro}}$ the nonhydrodynamic modes dominate and the expansion is far from equilibrium. We can draw two (related) conclusions from this. First, if all we know is the late-time gradient expansion (7.98), asking

before what time this expansion breaks down is *not* an accurate way of identifying the hydrodynamization time and the domain of validity of hydrodynamics. A similar conclusion was also reached in Ref. [48] by analyzing small perturbations on top of infinite static plasma. Second, even if we know many terms in the late-time gradient expansion (7.98), knowing only this it is impossible to run the equations of hydrodynamics backwards in time and reconstruct the far-from-equilibrium initial conditions. This follows from the fact that many very different far-from-equilibrium initial states can end up at late times in a boost-invariant hydrodynamic expansion (7.98) with the same Λ . This loss of memory of the initial conditions is characteristic of any equilibration process, and the hydrodynamization of a plasma that in boost-invariant expansion is no exception.

It is also interesting to notice that in all three panels of Fig. 7.15 at the time of hydrodynamization the fluid is markedly anisotropic, which means that in the hydrodynamic expansion (7.98) the first derivative terms are not much smaller than the zeroth derivative terms. The most extreme example is \mathcal{P}_L in the left-hand panel of Fig. 7.15, where the first derivative term in (7.98) is almost 70% as large as the zeroth order term at $t = t_{\text{hydro}}$. And yet, as we have described above, in all cases the second order terms are very small at the time of hydrodynamization. What this suggests is that although ideal, zeroth order, hydrodynamics of course becomes valid at asymptotically late times, at the times shown in Fig. 7.15 it does rather badly because it includes no dissipation and at these times dissipation is important. However, once the lowest order term that includes dissipation (i.e. the first order terms in the derivative expansion (7.98)) are included, the important physics that was missed at zeroth order is incorporated. With no further qualitatively new physics being missed, the second order terms are small. So, what we learn is that at the time of hydrodynamization the hydrodynamics is dissipative hydrodynamics, with first order terms behaving like additional leading terms since they are the leading dissipative terms.

The results illustrated in Fig. 7.15 clearly show the system hydrodynamizing before it isotropizes. Isotropization, like hydrodynamization, happens continuously with strict isotropy only being achieved in the infinite time limit. So, to make a quantitative comparison between τ_{hydro} and an isotropization time τ_{iso} we need to introduce a criterion for isotropization, just as we did for hydrodynamization. If we define τ_{iso} as the proper time when $(\mathcal{P}_T - \mathcal{P}_L)/\mathcal{P}_T = 0.1$, meaning that in the local fluid rest frame the pressure in the fluid is within 10% of being isotropic, we see from (7.98) that $\tau_{\text{iso}} = 15.6/\Lambda$ ($\tau_{\text{iso}} = 16.1/\Lambda$) for an expanding boost invariant hydrodynamic flow if we work to second (first) order in gradients. And, if we use (7.96) as an operational definition of a temperature $T(\tau)$ then $\tau_{\text{iso}}T(\tau_{\text{iso}}) = 6.2$ ($\tau_{\text{iso}}T(\tau_{\text{iso}}) = 6.4$). We saw above that in all three cases illustrated in Fig. 7.15 the hydrodynamization time τ_{hydro} is substantially less than this τ_{iso} meaning that, as

anticipated in Section 2.2, the strongly coupled plasma hydrodynamizes at a time when it is still significantly anisotropic. It then expands and cools according to the laws of viscous, i.e. dissipative, hydrodynamics with entropy being produced until $\tau \sim \tau_{\text{iso}}$. After $\tau \sim \tau_{\text{iso}}$, the boost-invariant expansion continues but the fluid is now close to isotropic meaning that gradient terms are unimportant and the expansion is well described by ideal, inviscid, hydrodynamics with no further production of entropy.⁴ We shall see via many examples throughout the remainder of this chapter that this ordering of events, hydrodynamization before isotropization, is generic when an expanding strongly coupled plasma forms, whether boost invariant or not, starting from varied far-from-equilibrium initial states. Although we know of no proof that hydrodynamization always happens first at strong coupling, we also know of no counterexamples.

Although hydrodynamization before isotropization should always have been seen as a logical possibility, in fact before the holographic calculations for strongly coupled fluids that we are describing in this chapter were done this possibility was not much considered. The expectation, based partly upon weak coupling intuition and partly upon not anticipating that first order terms in the derivative expansion of hydrodynamics could be significant at a time when second and higher order terms have already become insignificant, was that hydrodynamization would occur after the τ_{iso} defined above from (7.98). If this were the case, at the time of hydrodynamization the fluid would already be close to isotropic and the subsequent hydrodynamic expansion would be close to ideal with almost no entropy production. In such a setting, it was shown in Ref. [79] that isotropization could in fact occur before hydrodynamization. In this case, isotropization cannot be described hydrodynamically and τ_{iso} cannot be obtained from (7.98). It was even shown that in the epoch between early isotropization and later hydrodynamization the expansion of the weakly coupled matter can be described with equations that take the same form as the equations of hydrodynamics [79] albeit with constitutive relations, including in particular the equation of state, that can differ from those of hydrodynamics. This is possible because the processes that change the constitutive relations back to the equilibrium ones take a time that is much longer than τ_{iso} ; in the perturbative analysis of Ref. [79], these two different time scales are controlled by different powers of the small coupling constant. We now understand, therefore, that all these considerations are relevant only if the plasma that forms is weakly (possibly very weakly) coupled at the time that it hydrodynamizes. We have learned from holographic calculations, like those illustrated in Fig. 7.15 and like

⁴ No further production of entropy unless or until there is some large increase in the viscosity of the fluid as it cools, as can happen after a phase transition. For example, attempts to describe the late-time hadron gas phase of QCD via hydrodynamics require viscosities that increase rapidly with decreasing temperature [695, 274, 275, 316]. However, with the methods of this chapter, and indeed of this book, we are not seeking to gain insights into the physics of heavy ion collisions at late times, during or after the transition to hadronic matter.

those described throughout the rest of this chapter, that when a strongly coupled plasma is formed it hydrodynamizes first, then expands anisotropically and hydrodynamically, and only later isotropizes – with the isotropization and the associated cessation of entropy production being described well by hydrodynamics.

It is important to note that if $\tau_{\text{hydro}} < \tau_{\text{iso}}$ then in the analysis of flow observables in heavy ion collisions that we described in Section 2.2 the important time scale is τ_{hydro} . Following the conventions of the literature about flow in heavy ion collisions, in that section we referred to an “equilibration time”. What is meant by this phrase, in that context, is the time scale after which the expanding fluid can be described by the equations of viscous hydrodynamics, which is to say τ_{hydro} . So, the conclusion from that section that is of interest to us in this chapter should be phrased as saying that the agreement between the data on single particle spectra and azimuthally anisotropic flow in heavy ion collisions and hydrodynamic calculations implies that $\tau_{\text{hydro}} \leq (0.6\text{--}1)$ fm.

7.7.4 An all-order criterion for boost invariant hydrodynamization

In the previous subsection we introduced a deformation of the four-dimensional boundary theory metric in order to pump energy and momentum into the vacuum at early times and in this way create a boost-invariant far-from-equilibrium state. As in Fig. 7.2b, and as in the analysis of a nonexpanding plasma in Section 7.5, we had a source acting in the boundary for some duration in time. This made the identification of the hydrodynamization time a little ambiguous, since presumably the system was already beginning to hydrodynamize while the source was still acting. We now want to analyze the boost invariant expansion of an initial state created as in Fig. 7.2c, as we did for a nonexpanding plasma in Section 7.6. This has been accomplished by the authors of Refs. [448, 446], who found a way to impose boost-invariant, far-from-equilibrium, initial conditions in the bulk at $\tau = 0$. This removes the ambiguity of the previous subsection, although from the perspective of heavy ion physics it is also unrealistic: heavy ion collisions cannot be boost invariant at $\tau = 0$ since the colliding nuclei have a nonzero Lorentz-contracted thickness. The authors of Refs. [448, 446] developed a new numerical framework for solving the numerical relativity problem that, in the bulk, describes boost-invariant expansion in the boundary theory. We shall not describe their formalism but, very loosely speaking, it is a boost-invariant generalization of the formalism of Appendix D. They were then able to analyze the evolution and hydrodynamization of the expanding plasma resulting from a wide range of initial conditions at $\tau = 0$. These correspond, in their set-up, to specifying a single metric coefficient function (the “initial profile”) for the initial geometry on the hypersurface $\tau = 0$. We shall describe the results they obtain for 29 different initial profiles below.

Before we present results, we shall derive an all-order criterion for the validity of boost-invariant hydrodynamics, i.e. for the determination of the hydrodynamization time τ_{hydro} , introduced in Refs. [448, 446]. Doing so requires recasting the equations of boost-invariant hydrodynamics in terms of the effective temperature $T(\tau)$ that we introduced in (7.104). Recall that $T(\tau)$ is simply an alternative measure of the energy density if the system is far from equilibrium. All-order viscous hydrodynamics, namely the extension of (7.98) to arbitrarily high order, amounts to presenting the stress tensor as a series of terms expressed in terms of flow velocities u^μ and their derivatives with coefficients being proportional to appropriate powers of T , the proportionality constants being the transport coefficients. Hydrodynamic equations are just the conservation equations $\nabla_\mu T^{\mu\nu} = 0$, which are then by construction first order differential equations for T . In the case of a conformal fluid, in boost invariant expansion, in the hydrodynamic regime the effective temperature must take the form

$$T = \Lambda f(\Lambda\tau) \quad (7.111)$$

for some function f . This can be seen by taking the fourth root of the first equation in (7.98), and relies upon the fact that if the fluid is conformal there can be no dimensionful parameters present other than Λ . The only information about the initial conditions that the fluid “remembers” after it hydrodynamizes is contained in the constant Λ . Multiplying (7.111) by τ and inverting we arrive at

$$\Lambda = \frac{h(\tau T)}{\tau}, \quad (7.112)$$

where h is the function defined such that $h[xf(x)] = x$. Differentiating (7.112), we obtain

$$\tau \frac{d}{d\tau} \tau T = F_{\text{hydro}}(\tau T), \quad (7.113)$$

where $F_{\text{hydro}}(x) \equiv h(x)/h'(x)$. We conclude that in a conformal plasma in boost-invariant expansion, the hydrodynamic equation for the scale invariant quantity

$$w \equiv \tau T(\tau) \quad (7.114)$$

takes on the simple form

$$\frac{\tau}{w} \frac{d}{d\tau} w = \frac{F_{\text{hydro}}(w)}{w}, \quad (7.115)$$

where $F_{\text{hydro}}(w)$ is completely determined in terms of the transport coefficients of the theory [448, 446], much in the spirit of [586]. Intuitively, the reason that (7.115) does not hold outside the hydrodynamic regime is that before the system hydrodynamizes it has not yet lost memory of the initial conditions, and so the evolution depends on the physical scales that characterize those initial conditions.

For the plasma of strongly coupled $\mathcal{N} = 4$ SYM theory, $F_{\text{hydro}}(w)$ is known explicitly up to terms corresponding to third order hydrodynamics [173]:

$$\frac{F_{\text{hydro}}(w)}{w} = \frac{2}{3} + \frac{1}{9\pi w} + \frac{1 - \log 2}{27\pi^2 w^2} + \frac{15 - 2\pi^2 - 45 \log 2 + 24 \log^2 2}{972\pi^3 w^3} + \dots \quad (7.116)$$

The advantage of the result (7.115) over, for example, a seemingly simpler expression like (7.111) lies in the fact that, if the boost-invariant expansion of the fluid is governed entirely by hydrodynamics, including dissipative terms up to any high order or even resummed, then, on a plot of $\frac{\tau}{w} \frac{d}{d\tau} w$ as a function of w , trajectories for all initial conditions must lie on a single curve given by $F_{\text{hydro}}(w)/w$. If, on the other hand, genuine non-equilibrium processes intervene, i.e. if non-hydrodynamic modes have been excited, then we should observe a wide range of curves which all merge for sufficiently large w , after the system hydrodynamizes. Thus Eq. (7.115) can be used to test whether the stress tensor is of hydrodynamic form even without knowing the specific form of $F_{\text{hydro}}(w)$. Thus, it provides an all-order criterion for the hydrodynamization of a conformal plasma in boost invariant expansion.

In the top panel of Fig. 7.16 we plot $\frac{\tau}{w} \frac{d}{d\tau} w$ as a function of w for trajectories corresponding to 29 different initial states. It is clear from the plot that non-hydrodynamic modes are very important in the initial stage of plasma evolution. Yet, for all the sets of initial data, the curves merge into a single curve characteristic of hydrodynamics for $w > 0.7$. When plotted on this scale, hydrodynamization appears to be occurring even earlier but the vertical scale in the top panel has been extended in order to show all the very far-from-equilibrium dynamics at early times. In the bottom panel of Fig. 7.16 we show a plot of the evolution of the pressure anisotropy for a single initial state. Using (7.93) and (7.115),

$$1 - \frac{3\mathcal{P}_L}{\mathcal{E}} = 12 \frac{F(w)}{w} - 8, \quad (7.117)$$

so the quantity plotted in the lower panel of Fig. 7.16 is almost the same as that plotted in the upper panel, except now with a vertical scale chosen such that the details of the approach to hydrodynamization are visible. We compare the result to the corresponding curves for first, second and third order hydrodynamics. We observe, on the one hand, excellent agreement with hydrodynamics for $w > 0.63$ and, on the other hand, a quite sizable pressure anisotropy after hydrodynamization, as in the previous subsection, meaning that the anisotropic fluid is described well by viscous hydrodynamics. We can also see that, again as we discussed in the previous subsection, the hydrodynamic expansion itself is well under control at the hydrodynamization time. Non-hydrodynamic modes are important before $w = 0.63$, which cannot be inferred if all we know is the late time hydrodynamic behavior.




Figure 7.16 Top: plot of $\frac{\tau}{w} \frac{d}{d\tau} w$ versus w for 29 initial profiles. Bottom: the solid curve shows $1 - \frac{3\mathcal{P}_L}{\mathcal{E}} = \frac{2(\mathcal{P}_T - \mathcal{P}_L)}{\mathcal{E}}$, a measure of the pressure anisotropy, for a selected profile. The three dotted curves (top, middle and bottom) represent first, second and third order hydrodynamic fits. Figure taken from Ref. [448].

For all 29 initial conditions analyzed in Refs. [448, 446], hydrodynamization occurs at the latest by $w \equiv \tau T(\tau) = 0.7$. In the example that we analyzed in the previous subsection, when we pumped energy into the system over some period of time ending at τ_f , hydrodynamization occurred at τ_{hydro} with $(\tau_{\text{hydro}} - \tau_f)T(\tau_{\text{hydro}}) = 0.24$. Our estimate here that when the system is initialized far from equilibrium at $\tau = 0$ it hydrodynamizes by a time τ_{hydro} with $\tau_{\text{hydro}}T(\tau_{\text{hydro}}) = 0.7$, or perhaps somewhat smaller, confirms our speculation that in the analysis of the

previous section the hydrodynamization process has already begun before τ_f . It is also interesting to notice how similar the criterion $\tau_{\text{hydro}}T(\tau_{\text{hydro}}) = 0.7$ is to the results we obtained in Section 7.6 for the thermalization times of 1000 different initial states that result in a static plasma with final temperature T . The analyses of the equilibration of a homogeneous plasma in Sections 7.5 and 7.6 and the hydrodynamization of an expanding boost invariant plasma in this section all point toward the same conclusion: when strongly coupled plasma is formed by starting with some far-from-equilibrium state and letting it equilibrate or hydrodynamize, the time that this process takes is of order, or maybe even slightly less than, the inverse of the temperature at which the process concludes.

7.7.5 Boost-invariant hydrodynamization with radial flow

Throughout this chapter, we have assumed homogeneity in the transverse plane. It is of obvious interest to lift this restriction, since colliding ions are finite in transverse extent and result in a distribution of energy density that varies nontrivially in the transverse plane. The first step away from homogeneity is to assume only rotational symmetry in the transverse plane, meaning that initially (and during the subsequent expansion) the energy density profile is independent of the azimuthal angle. Quite recently, analytic solutions to the equations of viscous boost invariant hydrodynamics with this geometry have been found for the first time by Gubser [397]. Even more recently, the hydrodynamization of initially far-from-equilibrium states with this geometry which subsequently expand both longitudinally (in a boost-invariant fashion) and radially in the transverse plane has been analyzed for the first time by van der Schee [783]. The assumption of boost invariance in the longitudinal direction together with rotational symmetry in the transverse plane makes the gravitational problem 2 + 1-dimensional, and it can be solved using pseudo-spectral methods [783]. The author of Ref. [783] analyzes the boost invariant expansion of what is initially a “blob” of energy about 14 fm in diameter in the transverse plane. He chooses initial conditions in which the longitudinal pressure vanishes and the transverse pressure is half the energy density. He watches these initial conditions evolve and hydrodynamize. The result of the calculation is that hydrodynamization occurs at a τ_{hydro} at which $\tau_{\text{hydro}}T(\tau_{\text{hydro}}) \approx 0.8\text{--}0.9$, quite comparable to the hydrodynamization times found above for boost invariant expansion with no transverse dynamics. With the choice of initial transverse pressure profile in Ref. [783], gradients in the transverse pressure are such that by the time of hydrodynamization the fluid is expanding in the radial direction in the transverse plane with a velocity that reaches about $0.1c$ near the edge of the blob, where the gradient in the pressure profile is highest.

7.8 Colliding sheets of energy

In the previous section we studied the hydrodynamization of expanding plasma under the strong assumption of boost invariance. In this section we will relax this condition and we will study the dynamics of a collision of two sheets of energy, finite in thickness but infinite in transverse extent, in $\mathcal{N} = 4$ SYM theory. Perhaps this can be viewed as an instructive caricature of the collision of two large, highly Lorentz-contracted, nuclei. Introducing nontrivial dynamics in the transverse plane would yield an even better caricature with which to study the hydrodynamization of strongly coupled plasma as in a heavy ion collision. We saw in Section 7.7.5 that the first steps in this direction are just now being taken, albeit in a boost invariant setting. As we show in this section, the assumption of boost invariance can also be dispensed with.

Multiple authors have discussed collisions of infinitely extended planar shock waves in SYM, which in the dual description becomes a problem of colliding gravitational shock waves in asymptotically AdS spacetime. Existing work has examined qualitative properties and trapped surfaces [390, 405, 35, 36, 546, 547, 411, 580, 533], possible early-time behavior [549, 390, 151, 773], and expected late time asymptotics [495, 496]. As no analytic solution is known for this gravitational problem, solving the gravitational initial-value problem numerically is the only way to obtain quantitative results which properly connect early- and late-time behavior. This was done in Ref. [292], whose results we shall describe here. Although much of the earlier work concerned singular shocks with vanishing thickness, in Ref. [292] Chesler and Yaffe were able to analyze the collision and subsequent evolution and hydrodynamization of planar sheets whose energy density is everywhere finite, with a Gaussian profile in the “beam” direction, incident at the speed of light.

Diffeomorphism invariance plus translation invariance in two spatial directions allows one to write the bulk metric in the form

$$ds^2 = 2dv dr - A dv^2 + \Sigma^2 (e^B dx_{\perp}^2 + e^{-2B} dz^2) + 2F dv dz, \quad (7.118)$$

where A , B , Σ , and F are all functions of the bulk radial coordinate r , of the time v , and of the longitudinal coordinate z along which the waves will collide. As usual, we use generalized infalling Eddington–Finkelstein (EF) coordinates, and the EF time v coincides with the Minkowski time at the boundary, which lies at $r \rightarrow \infty$. Note that the crossed $dv dz$ term in the metric is necessary to describe the expected energy flux

$$S \equiv \frac{2\pi^2}{N_c^2} T^{tz} \quad (7.119)$$

in the z direction.

We want our initial gravitational data to be dual to two well-separated sheets of energy in the gauge theory, with finite thickness and energy density, moving towards each other at the speed of light. For a single sheet moving in the $\mp z$ direction, one possible choice on the gravity side is a planar shock of the form [495]

$$ds^2 = r^2 (-dx_+ dx_- + d\mathbf{x}_\perp^2) + \frac{1}{r^2} \left[dr^2 + h(x_\pm) dx_\pm^2 \right], \quad (7.120)$$

with $x_\pm \equiv t \pm z$, and h an arbitrary function – see e.g. Refs. [150, 91] for detailed discussions of these type of solutions. Note that the time t in this form of the metric is not an EF time but rather the analog of the time t shown in Fig. 7.1. The function h is chosen to be a Gaussian with width w and amplitude μ^3 :

$$h(x_\pm) \equiv \mu^3 (2\pi w^2)^{-1/2} e^{-\frac{1}{2}x_\pm^2/w^2}. \quad (7.121)$$

Note that the term $h(x_\pm) dx_\pm^2/r^2$ in the metric has precisely the correct fall-off to correspond to a vacuum expectation value of the T_{++} component of the boundary stress tensor, i.e. $\langle T_{++} \rangle \propto h(x_\pm)$, as corresponds to an excitation that propagates at the speed of light. The energy density per unit area of the shock is $\mu^3(N_c^2/2\pi^2)$. If the shock profile h has compact support, then a superposition of right- and left-moving shocks solves Einstein's equations at early times when the incoming shocks have disjoint support. Although this is not exactly true for our Gaussian profiles, the residual error in Einstein's equations is negligible when the separation of the incoming shocks is more than a few times the shock width. Following Ref. [292], we choose a width $w = 0.75/\mu$ and an initial separation of the shocks $\Delta z = 6.2/\mu$. We evolve the system for a total time $\Delta t = 9.1/\mu$.

Because of its light-like nature, all the curvature invariants of the metric (7.120) are finite. Nevertheless, as pointed out in [91], this metric possesses a naked curvature singularity, since tidal forces diverge in the region $h(x_\pm) \neq 0$, $r \rightarrow 0$ [682]. Recall that $r \rightarrow 0$ can be reached in finite affine parameter along future-directed causal curves, in particular along geodesics, and that when $h = 0$ it corresponds to the horizon of AdS. Presumably this singularity is related to the observation [292] that, when expressed in terms of EF coordinates, the functions A and F in the metric (7.118) seem to grow without bound as $r \rightarrow 0$. Following Ref. [292], we regulate this problem by adding to the metric (7.120) an additional piece representing a static, infinite thermal bath within which the shocks propagate. On the gravity side this corresponds to a static, infinite horizon at a small value of r that cloaks the singularity present for $r \rightarrow 0$. The bath acts as a regulator in the sense that its energy density is smaller by a factor of 50 than the peak energy in the shocks. This large separation implies that the effect of the bath on the propagation of the shocks is small and also implies that the temperature of the plasma at the time of hydrodynamization is much larger than the regulator temperature. However, the presence of the singularity means that the regulator cannot be strictly removed.




Figure 7.17 Energy density \mathcal{E}/μ^4 as a function of the boundary time t and longitudinal coordinate z . Figure taken from Ref. [292].

With the solution (7.120) plus the regulator in hand, the remaining tasks are (i) to transform these initial data to the EF coordinates of the ansatz (7.118), which is done numerically, (ii) to evolve Einstein's equations with these initial data, and (iii) to read off the boundary stress tensor from the near-boundary fall-off of the resulting metric. Figure 7.17 shows the energy density \mathcal{E} as a function of time and longitudinal position obtained from this calculation. On the left, one sees two incoming shocks propagating toward each other at the speed of light. After the collision, centered on $t=0$, energy is deposited throughout the region between the two receding energy density maxima. As the dynamics of the collision is strongly coupled, the energy density after the collision does not at all resemble two outgoing sheets of energy that have passed through each other. For example, after the collision in Fig. 7.17 the two receding maxima are moving outwards at less than the speed of light. To elaborate on this point, Fig. 7.18 shows a contour plot of the energy flux \mathcal{S} for positive t and z . The dashed curve shows the location of the maximum of the energy flux. The inverse slope of this curve, equal to the speed with which the receding maxima in the energy density are moving, is $v = 0.86$ at late times. This is the most dramatic manifestation of the fact that the dynamics is not boost invariant. The solid line shows the point beyond which $\mathcal{S}/\mu^4 < 10^{-4}$, and has slope 1. Evidently, the leading disturbance from the collision moves outwards at the speed of light, but the maxima in \mathcal{E} and \mathcal{S} move significantly slower. The collision has substantially slowed down the sheets of




Figure 7.18 Energy flux \mathcal{S}/μ^4 as a function of the boundary time t and longitudinal coordinate z . Figure taken from Ref. [292].

energy and, furthermore, it has resulted in the deposition of energy density between the receding sheets, in the vicinity of $z = 0$. This energy density, seen expanding and cooling in Fig. 7.17, is the plasma whose hydrodynamization we wish to quantify.

In Fig. 7.19 we plot the transverse and longitudinal pressures at $z = 0$ as a function of time. At $z = 0$, the pressures increase dramatically during the collision, resulting in a system which is very anisotropic and far from equilibrium. At $t = -0.23/\mu$, where \mathcal{P}_L has its maximum, \mathcal{P}_L is roughly 5 times larger than \mathcal{P}_T . At late times, the pressures asymptotically approach each other.

We expect that at sufficiently late times the evolution of $T^{\mu\nu}$ will be described by hydrodynamics. To test the validity of hydrodynamics, in Fig. 7.19 we also plot (as dashed lines) the pressures $\mathcal{P}_T^{\text{hydro}}$ and $\mathcal{P}_L^{\text{hydro}}$ predicted from the energy density by the first order viscous hydrodynamic constitutive relations [107]. At $z = 0$ the hydrodynamic constitutive relations hold within 15% at time $t_{\text{hydro}} = 2.4/\mu$, with improving accuracy thereafter.

At $z = 0$, where the flux $\mathcal{S} = 0$, the constitutive relations imply that the difference between $\mathcal{P}_T^{\text{hydro}}$ and $\mathcal{P}_L^{\text{hydro}}$ is purely due to viscous effects. Figure 7.19 shows that there is a large difference between \mathcal{P}_T and \mathcal{P}_L when hydrodynamics first becomes applicable, implying that viscous effects are substantial. As in all




Figure 7.19 Longitudinal (continuous green curve) and transverse pressure (continuous blue curve) in units of μ^4 as a function of time t , at $z = 0$. Also shown for comparison (dashed black curves) are the pressures predicted by the viscous hydrodynamic constitutive relations. Figure taken from Ref. [292].

the boost-invariant examples that we analyzed earlier, hydrodynamization of the strongly coupled fluid produced by the collision between the sheets of energy illustrated in Fig. 7.17 occurs before isotropization. Hydrodynamization is followed by an epoch of anisotropic hydrodynamic expansion during which entropy is produced. Only later, at an isotropization time that comes well after the times visible in Figs. 7.17 and 7.19, the pressures in the fluid become locally isotropic, entropy production ceases, and the subsequent expansion is described by ideal hydrodynamics.

At $z = 0$, when $t = t_{\text{hydro}} = 2.4/\mu$ the effective temperature, defined from the energy density according to (7.104) is $T = 0.27\mu$, meaning that $t_{\text{hydro}}T(t_{\text{hydro}}) = 0.65$. We see that the hydrodynamization time for the plasma produced in the collisions of this section is quite comparable to the estimates that we obtained in the context of boost-invariant expansion in Section 7.7.

As we conclude this chapter, we see a common conclusion emerging from a wide variety of calculations. We saw in Sections 7.5 and 7.6 that the thermalization of a static strongly coupled plasma with equilibrium temperature T takes a time in the range $(0.6-1)/T$, for a wide variety of far-from-equilibrium initial conditions. We then saw in Section 7.7 that when a strongly coupled plasma forms in a boost-invariant expansion, it hydrodynamizes in a time $\tau_{\text{hydro}} \lesssim 0.7/T(\tau_{\text{hydro}})$, again for a wide variety of far-from-equilibrium initial states. Here, $T(\tau_{\text{hydro}})$ is the effective temperature of the plasma, defined from the fourth root of its energy density, at the time that it hydrodynamizes. Finally, in this section where we have

dispensed with boost invariance, at least in the one instance of colliding sheets of energy density that we have analyzed we find $t_{\text{hydro}}T(t_{\text{hydro}}) = 0.65$. The conclusion that the hydrodynamization of strongly coupled plasma takes a time satisfying $t_{\text{hydro}}T(t_{\text{hydro}}) = 0.6\text{--}1$ seems robust indeed.

As we have discussed in Section 2.2, the comparison of data on identified particle spectra and elliptic flow in heavy ion collisions at RHIC to hydrodynamic calculations indicates that the quark–gluon plasma produced in RHIC collisions hydrodynamizes before a time of order $0.6\text{--}1$ fm/c [443, 543]. Before the investigations described in this chapter, this was always seen as rapid thermalization since analyses of thermalization that are based upon weakly coupled physics point to significantly longer thermalization times. How does our strong coupling conclusion compare? In doing hydrodynamic calculations to compare with data from heavy ion collisions, it has become conventional to initialize the hydrodynamic calculation $\tau = 0.6$ fm after the collision. For example, in Ref. [734] we find recent viscous hydrodynamic calculations that fit RHIC data well if the temperature at the middle of the fireball is initially $T(0.6 \text{ fm}) = 347\text{--}379$ MeV, with the variation coming from uncertainty in the shape of the energy density profile at that time. We see that $\tau = 0.6$ fm after a RHIC collision, $\tau T(\tau) \sim 1.1\text{--}1.2$. Although we do not have full-fledged hydrodynamic calculations that start this early to call upon, we see that $\tau T(\tau) = 0.65$ corresponds to a τ that is somewhere around 0.3 fm. In heavy ion collisions at the LHC, $T(0.6 \text{ fm}) = 444\text{--}485$ MeV [734], meaning that $\tau T(\tau) = 0.65$ corresponds to a τ that is somewhere around 0.2 fm. Following the discussion with which we began this chapter, it would be inappropriate to take these as estimates for the hydrodynamization times of heavy ion collisions *per se*. We expect that at the very beginning of a heavy ion collision, the dominant physics is not yet strongly coupled. The impact of these strong coupling estimates is that they teach us that the more than ten-year-old result that the matter produced in RHIC collisions takes at most $0.6\text{--}1$ fm to hydrodynamize should no longer be seen as “rapid thermalization”, since this time scale is comfortably longer than what we now know we should expect if the physics of heavy ion collisions were strongly coupled from the start.

We close by noting that there are two other ways to connect the results from the calculation of colliding sheets in this section to heavy ion collisions. The authors of Ref. [292] noted that if we equate the energy per unit transverse area of the sheets with the mean energy per unit transverse area of the Lorentz-contracted nuclei incident at RHIC, this sets the scale for μ without appealing to the hydrodynamization temperature in the final state. This estimate yields $\mu \sim 2.3$ GeV, which would then mean that $t_{\text{hydro}} = 2.4/\mu$ corresponds to $t_{\text{hydro}} \sim 0.2$ fm. The reason that this yields a larger value of μ than the one that we obtained implicitly above when we used the hydrodynamization temperature is that, in the collision of the two sheets of

strongly coupled matter, the energy density near $z \sim 0$ in the final state is greater than in a collision between heavy ions with the same energy per unit transverse area. The other option is to use the width of the sheets to set the scale. In the collision that we have described, $w = 0.75/\mu$ meaning that $t_{\text{hydro}} = 2.4/\mu = 3.2 w$. It is a little hard to compare the Gaussian profiles of the colliding sheets of energy with the profiles of Lorentz-contracted nuclei in a quantitative way, but this suggests $t_{\text{hydro}} \sim 0.3$ fm, consistent with what we obtained via the hydrodynamization temperature above. Based on experience in simpler contexts earlier in this chapter, we expect that the hydrodynamization time is controlled by the inverse of the hydrodynamization temperature when that time scale is much longer than w , and by w when that is the longer time scale. For the specific value of $w\mu$ that we have used, the two estimates are comparable. It would be interesting to investigate the collisions of sheets of energy density with Gaussian profiles having varying values of $w\mu$, and then with profiles having other shapes. And, it will be interesting to investigate colliding disks with a non-trivial profile in the transverse plane as, we have seen in Section 7.7.5, is already now possible in a boost-invariant setting.