IOPscience iopscience.iop.org

Home Search Collections Journals About Contact us My IOPscience

Strongly correlated quantum fluids: ultracold quantum gases, quantum chromodynamic plasmas and holographic duality

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

2012 New J. Phys. 14 115009

(http://iopscience.iop.org/1367-2630/14/11/115009)

View the table of contents for this issue, or go to the journal homepage for more

Download details:
IP Address: 18.51.1.228
The article was downloaded on 28/03/2013 at 14:24

Please note that terms and conditions apply.

New Journal of Physics
The open access journal for physics

Strongly correlated quantum fluids: ultracold quantum gases, quantum chromodynamic plasmas and holographic duality

Allan Adams[1], **Lincoln D Carr**[2,3,6], **Thomas Schäfer**[4], **Peter Steinberg**[5] **and John E Thomas**[4]
[1] Center for Theoretical Physics, MIT, Cambridge, MA 02139, USA
[2] Physics Institute, University of Heidelberg, D-69120 Heidelberg, Germany
[3] Department of Physics, Colorado School of Mines, Golden, CO 80401, USA
[4] Department of Physics, North Carolina State University, Raleigh, NC 27695, USA
[5] Brookhaven National Laboratory, Upton, NY 11973, USA
E-mail: lcarr@mines.edu

New Journal of Physics **14** (2012) 115009 (121pp)
Received 23 May 2012
Published 19 November 2012
Online at http://www.njp.org/
doi:10.1088/1367-2630/14/11/115009

Abstract. Strongly correlated quantum fluids are phases of matter that are intrinsically quantum mechanical and that do not have a simple description in terms of weakly interacting quasiparticles. Two systems that have recently attracted a great deal of interest are the quark–gluon plasma, a plasma of strongly interacting quarks and gluons produced in relativistic heavy ion collisions, and ultracold atomic Fermi gases, very dilute clouds of atomic gases confined in optical or magnetic traps. These systems differ by 19 orders of magnitude in temperature, but were shown to exhibit very similar hydrodynamic flows. In particular, both fluids exhibit a robustly low shear viscosity to entropy density ratio, which is characteristic of quantum fluids described by holographic duality, a mapping from strongly correlated quantum field theories to weakly curved higher dimensional classical gravity. This review explores the connection

---

[6] Author to whom any correspondence should be addressed.

Content from this work may be used under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike 3.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

© IOP Publishing Ltd and Deutsche Physikalische Gesellschaft

between these fields, and also serves as an introduction to the focus issue of *New Journal of Physics* on 'Strongly Correlated Quantum Fluids: From Ultracold Quantum Gases to Quantum Chromodynamic Plasmas'. The presentation is accessible to the general physics reader and includes discussions of the latest research developments in all three areas.

Contents

1. Introduction **2**
2. Ultracold quantum gases **6**
  2.1. Ultracold Fermi gas experiments 9
  2.2. Scale invariance and universality 13
  2.3. Experimental determination of the equation of state 16
  2.4. Experimental studies of the phase transition 20
  2.5. Universal hydrodynamics and transport 22
  2.6. The Bardeen–Cooper–Schrieffer to Bose–Einstein condensate crossover in lattices 25
  2.7. Recent and new directions in crossover physics 27
3. Quantum chromodynamics, the quark–gluon plasma and heavy-ion collisions **30**
  3.1. Quantum chromodynamics and the phase diagram 30
  3.2. Weakly versus strongly coupled plasmas 35
  3.3. Nuclear collisions: initial conditions 37
  3.4. Particle multiplicities 41
  3.5. Hydrodynamic flow 43
  3.6. Jet quenching 47
  3.7. Heavy quarks 50
4. Holographic duality **53**
  4.1. Why should holography be true? Two heuristic pictures 55
  4.2. Essential holography 62
  4.3. Applied holography 80
5. Conclusions **96**
  5.1. Open problems and questions in ultracold quantum gases 96
  5.2. Open problems and questions in quantum chromodynamic plasmas 98
  5.3. Open problems and questions in holographic duality 101
Acknowledgments **104**
References **104**

1. Introduction

This review covers the convergence between three at first sight disparate fields: ultracold quantum gases, quantum chromodynamic (QCD) plasmas and holographic duality. *Ultracold quantum gases* have opened up new vistas in many-body physics, from novel quantum states of matter to quantum computing applications. There are over 100 experiments on ultracold quantum gases around the world on every continent except Antarctica. *The QCD plasma*, also called the *quark–gluon plasma* (QGP), has been the subject of an intensive experimental investigation for more than two decades, continuing now at the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory (BNL) and the Large Hadron Collider (LHC) of the European Organization for Nuclear Research (CERN). A QGP is predicted to have occurred in the first microsecond after the Big Bang, and re-creation of the QGP on the Earth at present allows us to probe the physics of the early universe. *Holographic duality* is a powerful mapping from strongly interacting quantum field theories, where the very concept of a quasiparticle can lose meaning, to weakly curved higher dimensional classical gravitational theories. This duality provides a new approach for modeling strongly interacting quantum systems, yielding fresh insights into previously intractable quantum many-body problems key to understanding experiments such as ultracold quantum gases and the QGP.

What do these three fields have in common? All treat *strongly correlated quantum fluids*. Generically, strong interactions give rise to strong correlations. By *strongly correlated* we mean that we cannot describe a system by working perturbatively from non-interacting particles or quasiparticles. In the case of electrons in condensed matter systems, this means that theories constructed from single-particle properties, such as the Hartree–Fock approximation, cannot describe a material. In the case of fluids, we mean that kinetic theories based on quasiparticle degrees of freedom, in particular the Boltzmann equation, fail[7]. The natural candidates for building quasiparticles are quarks and gluons in the case of the QGP, neutrons and protons in the case of nuclear matter and atoms in the case of ultracold atomic gases. In strongly interacting systems the mean free path for these excitations is comparable to the interparticle spacing, and quasiparticles lose their identity. Even though kinetic theory fails, nearly ideal, low-viscosity hydrodynamics is a very good description of these systems. This is a central prediction of holographic duality, and has been verified experimentally for both ultracold quantum gases and the QGP, as we will explore in this review.

As shown in figure 1, strongly correlated quantum fluids cover a wide range of scales in temperature and pressure[8]. We remind the reader that temperature T and energy E are equivalent up to a factor of Boltzmann's constant, $k_{\mathrm{B}} = 1.3806503 \times 10^{-23}\,\mathrm{J\,K^{-1}}$, with $E = k_{\mathrm{B}}T$. We focus on fluids that can be studied in bulk, as opposed to quantum liquids that exist on lattices. We show ultracold Fermi gases, liquid helium, neutron matter in proto-neutron stars and the QGP. For comparison we also show a classical fluid, water, and a classical plasma, the Coulomb plasma in the sun.

Figure 2 shows that despite the large range in scale there is a remarkable universality in the transport behavior of strongly correlated quantum fluids. Transport properties of the fluid can be

---

[7] Fluids are materials that obey the equations of hydrodynamics. The word *liquid* refers to a phase of matter that cannot be distinguished from a gas in terms of symmetry, but exhibits short-range correlations similar to those in a solid, and is separated from the gas phase by a line of first-order transitions that terminates at a critical endpoint. A plasma is a gas of charged particles. Gases, liquids and plasmas behave as fluids if probed on very long length scales. Weakly coupled systems exhibit a single-particle behavior if probed on microscopic scales, but strongly coupled systems behave like fluids also on short scales. Liquids are typically more strongly correlated than gases, and are more likely to behave like a fluid.

[8] The points in figure 1 correspond to the range of temperatures for which the transport measurements shown in figure 2 have been made. For the ultracold atomic Fermi gas experiments described in section 2.1 the critical temperature is roughly 500 nK (the exact value depends on the trap geometry and the number of particles; Bose gases have been cooled to temperatures below 1 nK). The data points for helium and water are centered around the critical endpoint of the liquid gas transition. The point for the solar plasma corresponds to the geometric mean of the temperatures in the core and the corona. The neutron matter point is at $T = 1\,\mathrm{MeV}/k_{\mathrm{B}} = 1.2 \times 10^{10}\,\mathrm{K}$ and at a density $n = 0.1\,n_0$, where $n_0 = 0.14\,\mathrm{fm}^{-3}$ is the nuclear matter saturation density. Neutron stars are born at $T \simeq 10\,\mathrm{MeV}/k_{\mathrm{B}}$, and they can cool to temperatures below $1\,\mathrm{keV}/k_{\mathrm{B}}$. The critical temperature of the QGP is $T_{\mathrm{c}} \simeq 150\,\mathrm{MeV}/k_{\mathrm{B}} = 1.75 \times 10^{12}\,\mathrm{K}$. Experiments with heavy ions explore temperatures up to $\sim 3T_{\mathrm{c}}$.

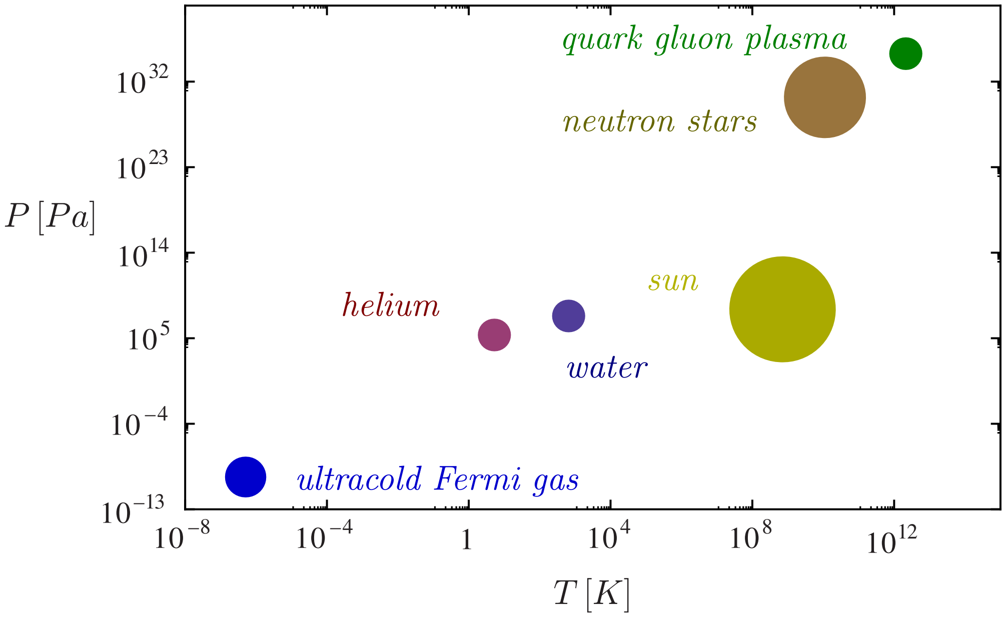

Figure 1. Temperature and pressure scales of extreme quantum matter. Ultracold quantum gases are the coldest matter produced to date, while the QGP is the hottest, together spanning about 19 orders of magnitude in temperature and about 44 orders of magnitude in pressure. However, these systems exhibit very similar hydrodynamic behavior, as characterized by the shear viscosity to entropy density ratio shown in figure 2. We include two other well-known quantum fluids, liquid helium and hot proto-neutron star matter, as well as a classical fluid, water and a classical plasma, the Coulomb plasma in the sun.

characterized in terms of its shear viscosity η, which governs dissipation due to internal friction. A dimensionless measure of dissipative effects is the ratio η/s of shear viscosity to entropy density in units of $\hbar/k_B$. Near the critical point, where the role of correlations is expected to be strongest, the ratio η/s has a minimum. For classical fluids the minimum value is much bigger than $\hbar/k_B$, but for strongly correlated quantum fluids η/s is of order $\hbar/k_B$, indicating that dissipation is governed by quantum effects. We observe, in particular, that η/s for the QGP and ultracold Fermi gases is quite similar, even though the absolute values of η and s differ by many orders of magnitude[9].

Remarkably, these values of η/s lie near a lower bound, $\eta/s \geqslant \hbar/(4\pi k_B)$, which arises in the study of $4+1$ dimensional black holes in classical Einstein gravity. These gravitational theories are conjectured to be dual to certain $3+1$ dimensional quantum field theories; see section 4. This lower bound is known to be non-universal; it can be violated in a more general class of theories dual to a gravitational theory known as the Gauss–Bonnet gravity. Imposing

---

[9] The theoretical curves, as well as the data for helium and water, correspond to systems in the thermodynamic limit. The lattice data for the QGP and the ultracold Fermi gas have finite volume corrections that have not been fully quantified. The experimental data point for the QGP is based on an analysis that assumes η/s to be temperature independent. The data points for the ultracold Fermi gas show the ratio of trap averages of η and s. The local value of η/s at the center of the trap is likely to be smaller than the ratio of the averages; see section 5.1.

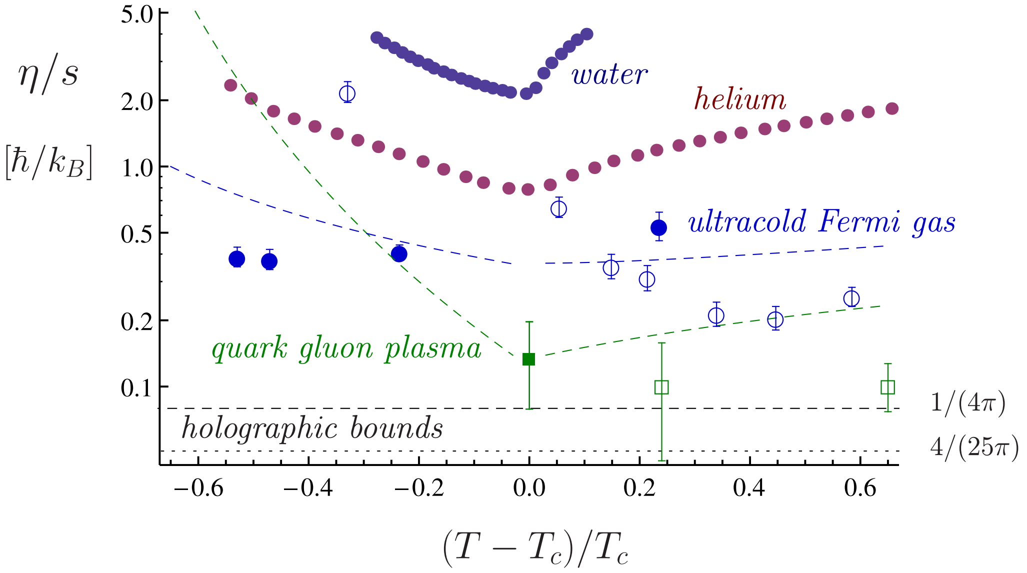

Figure 2. Transport properties of strongly correlated fluids. The ratio of shear viscosity η to entropy density s as a function of $(T-T_c)/T_c$, where T_c is the superfluid transition temperature in the case of ultracold Fermi gases, the deconfinement temperature in the case of QCD and the critical temperature at the endpoint of the liquid gas transition in the case of water and helium. The data for water and helium are from [1], the ultracold Fermi gas data are from [2], the QGP point (square) is taken from the analysis of [3], the lattice QCD data (open squares) are from [4] and the lattice data for the ultracold Fermi gas (open circles) are the 8^3 data from [5]. The dashed curves are theory curves from [6–9]. The theories are scaled by overall factors to match the data near T_c. The lines labeled 'holographic bounds' correspond to the Kovtun, Son and Starinets (KSS) bound $\hbar/(4\pi k_B)$ [11] and the Gauss–Bonnet bound $(16/25)\hbar/(4\pi k_B)$ [10]. Similar compilations can be found in [11–13].

basic physical requirements such as causality and positivity leads to a slightly smaller bound[10] $\eta/s \geqslant (16/25)\hbar/(4\pi k_B)$. The main feature of the results obtained using holographic dualities is that, at strong coupling, η/s is both unusually small and relatively insensitive to the precise strength of the interactions, as long as they are strong. This is in sharp contrast to the predictions of kinetic theory for a weakly interacting gas. As a result, η/s serves as a probe of the strength of correlations in a quantum fluid.

We have chosen to focus on the fields of ultracold quantum gases and the QGP not only for their range of energy and density scales, but also because of their broad intrinsic interest. Ultracold fermions are connected to a wide variety of exotic, strongly interacting systems in nature, ranging from high-temperature superconductors to nuclear matter. They are incredibly flexible many-body systems that allow nearly arbitrary tuning of interactions, symmetries, spin structure, effective mass and imposed lattice structures. The QGP, on the other hand, explores a very different regime from other particle physics experiments: thousands of particles are

[10] Whether this value represents a true lower bound, or whether more general classes of fluids with even smaller values of η/s can exist, is an active area of research; see section 4.3.

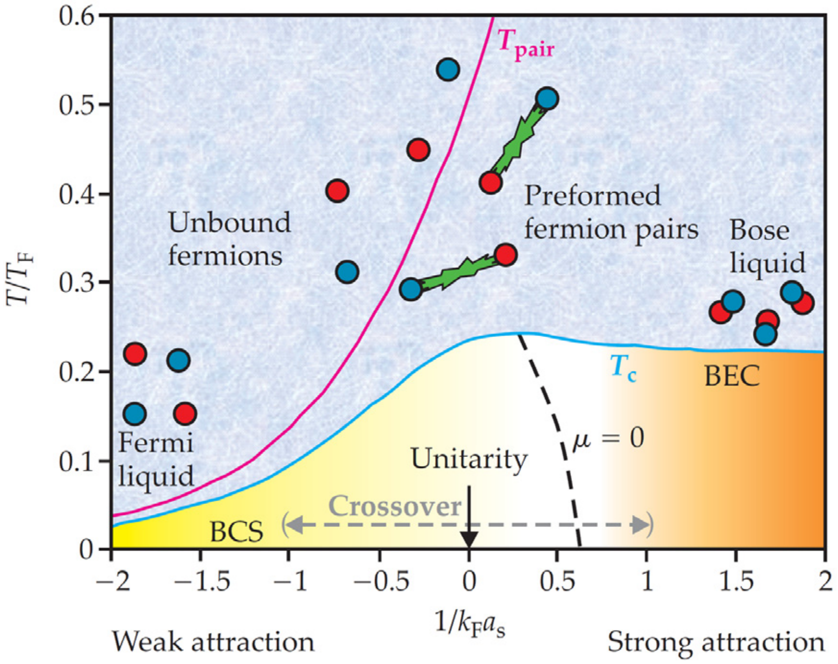

Figure 3. Ultracold Fermi gas phase diagram. Sketch of the BCS to BEC crossover for ultracold Fermi gases. When the scattering length a_s passes through a pole, so that $1/(k_F a_s) \to 0$, one obtains a strongly correlated fluid, the *unitary gas*. The critical temperature T_c for the phase transition only approaches the pairing temperature T_{pair} in the limit $1/(k_F a) \to -\infty$. The *crossover region* is the strongly interacting regime, loosely defined as $|1/(k_F a_s)| < 1$. Note that we denote the scattering length by a in the text. Figure reproduced with permission from [33].

produced, and these particles form a hot and dense environment recreating the conditions of the early universe.

Our review is outlined as follows. A key introductory or heuristic plot is presented in each of our three fields for the general physics reader. In section 2, figure 3 shows the phase diagram for the Bardeen–Cooper–Schrieffer (BCS) to Bose–Einstein condensate (BEC) crossover, midway through which the *unitary quantum gas*, a strongly correlated quantum fluid, is obtained. In section 3, the QCD phase diagram is shown in figure 11, including the strongly and weakly interacting QGP being explored currently at RHIC and the LHC. In section 4, not only is an extensive heuristic description of holographic duality provided in section 4.1, but also a *holographic dictionary* is presented in table 1, mapping quantities in strongly interacting quantum fields similar to the unitary quantum gas or the QGP to their gravitational duals. The three main sections also cover all 39 papers from three fields forming this focus issue. Finally, in section 5, rather than a summary, we instead conclude with open questions in each field.

2. Ultracold quantum gases

Ultracold quantum gases provide a unique table-top paradigm for exploring the properties of quantum many-body systems in nature [14–20], from the thermodynamics of high-temperature superconductors to the hydrodynamics of QGPs. These gases are mainly made of alkali metal atoms but also more recently other atoms as well as diatomic molecules. They can be fermionic or bosonic with a wide variety of internal hyperfine spin structures. They can be made strongly or weakly interacting and both attractive and repulsive. They are contained in a variety of magnetic and optical traps in one, two and three dimensions, including optical lattices. The latter gives rise to arbitrary lattice structures. Because these gases are dilute and very cold, they are described by first-principles theories developed from low-energy binary scattering between atoms, and well-known interactions with magnetic and optical fields. The collection of atoms can be probed and manipulated by external laser beams and pulses, as well as external magnetic fields.

In this review we focus on atomic Fermi gases [17, 18, 21, 22], in particular strongly interacting Fermi gases. Several have been cooled to degeneracy using evaporative cooling methods. The most widely studied atoms are ^{40}K [23] and ^{6}Li [24–27].[11] Experiments are carried out at temperatures in the nanokelvin to microkelvin range, with typical densities from 10^{11} to 10^{14} atoms cm^{-3}. For an ^{6}Li atom at nanokelvin temperatures, the thermal de Broglie wavelength

$$\lambda_{\mathrm{dB}} = \hbar\sqrt{\frac{2\pi}{mk_{\mathrm{B}}T}} \tag{1}$$

is of the order of several μm. Quantum degeneracy occurs when the de Broglie wavelength is greater than or of the order of the particle spacing, $\lambda_{\mathrm{dB}}n^{-1/3} \gtrsim 1$, where n is the density; this condition is equivalent to $T/T_{\mathrm{F}} \geqslant 1$, where T_{F} is the Fermi temperature[12]. Our interest is in ultracold Fermi gases that are quantum degenerate: $T/T_{\mathrm{F}} \leqslant 1$.

A cloud of trapped dilute fermions is typically about 100 μm in size, and is deformed by harmonic trapping fields into prolate or oblate forms, commonly called a cigar or a pancake. In the degenerate regime the cloud is stabilized against collapse by Pauli pressure [23, 24, 26]. The size of the cloud depends on the interplay between the trapping potential, the Pauli pressure and interactions between the atoms. Because of the low density and ultracold temperatures, these interactions are dominated by an effective s-wave contact interaction. The scattering amplitude is of the form

$$f(k) = \frac{1}{-1/a + r_0k^2/2 - \mathrm{i}k}, \tag{2}$$

where a is the s-wave scattering length and r_0 is the effective range. Higher partial waves as well as short-range corrections are suppressed by powers of $r_0/\lambda_{\mathrm{dB}}$ and $r_0n^{1/3}$.[13] The scattering length is widely tunable by a Feshbach resonance [32], an external magnetic field that brings a weakly bound excited molecular state into resonance with the unbound atomic scattering state.

Each of the different trapped atomic elements used in ultracold quantum gas experiments has an internal spin structure due to the hyperfine structure of the atom, that is, the combination of the nuclear spin and, in the case of the alkali metals, the electron outside the closed shell. For instance, ^{6}Li has a nuclear spin of 1 and one unpaired electron. The two lowest hyperfine

[11] Degeneracy has also been achieved in metastable ^{3}He* [28], in ^{171}Yb and ^{173}Yb [29], and recently, in ^{87}Sr [30] and ^{161}Dy [31].

[12] In this review, T_{F} is always defined with respect to the non-interacting Fermi gas, $T_{\mathrm{F}} = \hbar^2k_{\mathrm{F}}^2/(2mk_{\mathrm{B}})$ with $k_{\mathrm{F}} = (3\pi^2n)^{1/3}$ for a two-component gas.

[13] The range of the atomic potential is of the order of the van der Waals length $l = (mC_6/\hbar^2)^{1/4}$, where C_6 controls the van der Waals tail of the atomic potential, $V \sim C_6/r^6$. We assume that the p-wave scattering length is natural, meaning that $a_{\mathrm{p}} \sim r_0$.

states have a total spin of 1/2, and the remaining four have a total spin of 3/2. By selecting out hyperfine states through appropriate laser-induced transitions and trapping and cooling methods, experiments can thus work with a variety of spin structures. The case in which two hyperfine states are trapped is effectively equivalent to a spin-1/2 atom. Tuning the scattering length by using a Feshbach resonance, one obtains three distinct regimes, shown in figure 3. The first is for weak attractive scattering, $-k_{\mathrm{F}}a \ll 1$, where k_{F} is the Fermi momentum. Then for temperatures well below the Fermi temperature T_{F}, one obtains a BCS state [34] or s-wave superconductivity. We call this an *atomic Fermi superfluid*, since our systems are in fact neutral. In such a state, fermions of opposite spin join to make Cooper pairs, but their average pair size ξ_{c} is greater than the interparticle spacing $n^{-1/3}$, so that they are overlapping: $\xi_{\mathrm{c}}n^{1/3} \gg 1$. Tuning a as in the phase diagram, we observe that the scattering length passes through a pole; note that the figure shows temperature as a function of the inverse scattering length, $1/(k_{\mathrm{F}}a)$. Thus as $a \to \pm\infty$, $1/(k_{\mathrm{F}}a) \to 0$, and one obtains a second regime, called the *unitary gas*. This middle regime is a strongly correlated fluid, and one finds that $\xi_{\mathrm{c}}n^{1/3} \simeq 1$, i.e. the pair size is approximately equal to the interparticle spacing. Finally, for large positive scattering lengths, the paired fermions make much more tightly bound molecules, and one obtains a molecular BEC, similar to the well-known atomic BECs. This regime is depicted on the far right of figure 3. In practice these molecules are still quite large, thousands of Bohr or more, but still much smaller than the interparticle spacing, so that $\xi_{\mathrm{c}}n^{1/3} \ll 1$.

The upper curve in the figure depicts the pair formation temperature T_{pair}, which in general is distinct from the critical temperature for superfluidity, T_{c} [35]. Note that superfluidity is associated with the spontaneous breakdown of a global symmetry, the $U(1)$ phase symmetry of the wavefunction, and T_{c} is therefore always well defined. T_{pair}, on the other hand, is not associated with a symmetry or a local order parameter and may not be well defined. This remark is particularly relevant in the BCS regime, where the size of the pairs is large compared to the interparticle spacing. Physically, we expect that in the BCS regime there are no pre-formed pairs, and T_{c} and T_{pair} are very close together.

Although we refer to these systems as ultracold, in terms of the dimensionless ratio T/T_{F}, and in comparison with solid-state systems, they are quite hot. In the unitary regime the phase transition occurs at $T_{\mathrm{c}}/T_{\mathrm{F}} = 0.167(23)$ [36], compared to typical solid state superconductors in which $T_{\mathrm{c}}/T_{\mathrm{F}} \lesssim 0.01$. The unitary Fermi gas is a very high T_{c} superfluid. As indicated in figure 3, in the BCS regime the temperature required for achieving a phase transition to an atomic Fermi superfluid is quite low. In this regime, the critical temperature is given by the weak coupling expression [37]

$$\frac{T_{\mathrm{c}}}{T_{\mathrm{F}}} \simeq \frac{4 \times 2^{1/3} \mathrm{e}^{\gamma}}{\pi \mathrm{e}^{7/3}} \exp\left(-\frac{\pi}{k_{\mathrm{F}}|a|}\right), \tag{3}$$

where $\gamma \simeq 0.577$ is the Euler constant. The numerical value of the factor in front of the exponent is 0.277. We observe that even though equation (3) is formally valid only in the limit $k_{\mathrm{F}}|a| \ll 1$, it also provides a reasonable estimate for T_{c} at unitarity. This is somewhat of an accident, because higher order corrections in $k_{\mathrm{F}}|a|$ are divergent in the unitary limit.

In the following, we explore the unitary regime of the BCS–BEC crossover for ultracold Fermi gases. In section 2.1 we present an overview of experiments on these systems. In section 2.2 we focus on universal aspects of unitary gases. In the strongly correlated regime the scattering length diverges, and the remaining length scales in the problem are the Fermi length $1/k_{\mathrm{F}}$ and the de Broglie wavelength λ_{dB}, given in equation (1). Thus many theoretical statements can be made despite the lack of a small parameter or perturbative calculations. In sections 2.3 and 2.4, we treat the thermodynamics and the structure of the phase diagram for unitary gases, and in section 2.5 we focus on transport properties. Section 2.6 presents an overview of ultracold Fermi gases in optical lattices. Finally, in section 2.7 we treat new directions in unitary gases as presented in this focus issue, including novel experimental probes, solitons, imbalanced systems and polarons, disorder, quantum phase transitions, Efimov physics and the use of three hyperfine states to explore $SU(3)$ physics and connections to the QGP.

2.1. Ultracold Fermi gas experiments

Historically, atomic Fermi gases were first brought to degeneracy at JILA in 1999 [23], using a mixture of two hyperfine states in ^{40}K to enable s-wave scattering in a magnetic trap. Dual species radio-frequency (RF)-induced evaporation produced a degenerate, weakly interacting sample, with $T/T_F \simeq 0.5$. Later, degeneracy was achieved in fermionic ^{6}Li by direct evaporation from a magneto-optical trap (MOT)-loaded optical trap [26] and by sympathetic cooling with another species [24, 25, 27], producing a lower T/T_F. However, the minimum reduced temperature was initially limited to $T/T_F \simeq 0.2$, which may have been a consequence of trap-noise-induced heating [40] or, at the lowest temperatures, Fermi hole heating [38] in combination with evaporative cooling [39].

Optical traps enabled a dramatic improvement in the efficiency of evaporation and the creation of strongly interacting Fermi gases, through the use of magnetically tunable collisional resonances or Feshbach resonances [41]. Feshbach resonances in fermionic atoms were initially characterized in 2002 by several groups [42–44]. For a recent review of Feshbach resonances see [32]. In a Feshbach resonance, a bias magnetic field tunes the total energy of a pair of colliding atoms in the incoming open (triplet) channel into resonance with a molecular bound state in an energetically closed (singlet) channel. At resonance, the zero-energy s-wave scattering length a diverges and the collision cross section attains the *unitary* limit, proportional to the square of the de Broglie wavelength, i.e. $\sigma = 4\pi/k^2$, where $\hbar k$ is the relative momentum. The collision cross section therefore *increases* with decreasing temperature, enabling highly efficient evaporative cooling in optical traps and much lower reduced temperatures $T/T_F \simeq 0.05$.

An optical trap is formed by a focused laser beam. Atoms are polarized by the field and attracted to the high-intensity region when the laser is detuned below resonance with the resonant optical transition, so that the induced dipole moment is in phase with the field. For large detunings, obtained using infrared lasers, the trapping potential is independent of the atomic hyperfine state, enabling several species to be trapped, which is ideal for Fermi gases [45]. Forced evaporation is accomplished by slowly lowering the intensity of the optical trap laser beam. Near a Feshbach resonance, a highly degenerate sample can be produced in a few seconds [46].

A milestone in the Fermi gas field was the observation in 2002 of a strongly interacting, degenerate Fermi gas [46], in the so-called BEC–BCS crossover regime, using this method. In contrast to Bose gases, which are unstable and undergo three-body loss on millisecond time scales near Feshbach resonances, the two-component Fermi gas was found to be stable, as the Pauli principle suppresses three-body s-wave scattering [47]. Released from the cigar-shaped trapping potential of the focused beam, the cloud was observed to expand much more rapidly in the initially narrow direction, compared to the long direction, as a consequence of the much larger pressure gradient along the narrow axis. Consequently, the aspect ratio inverts from a cigar

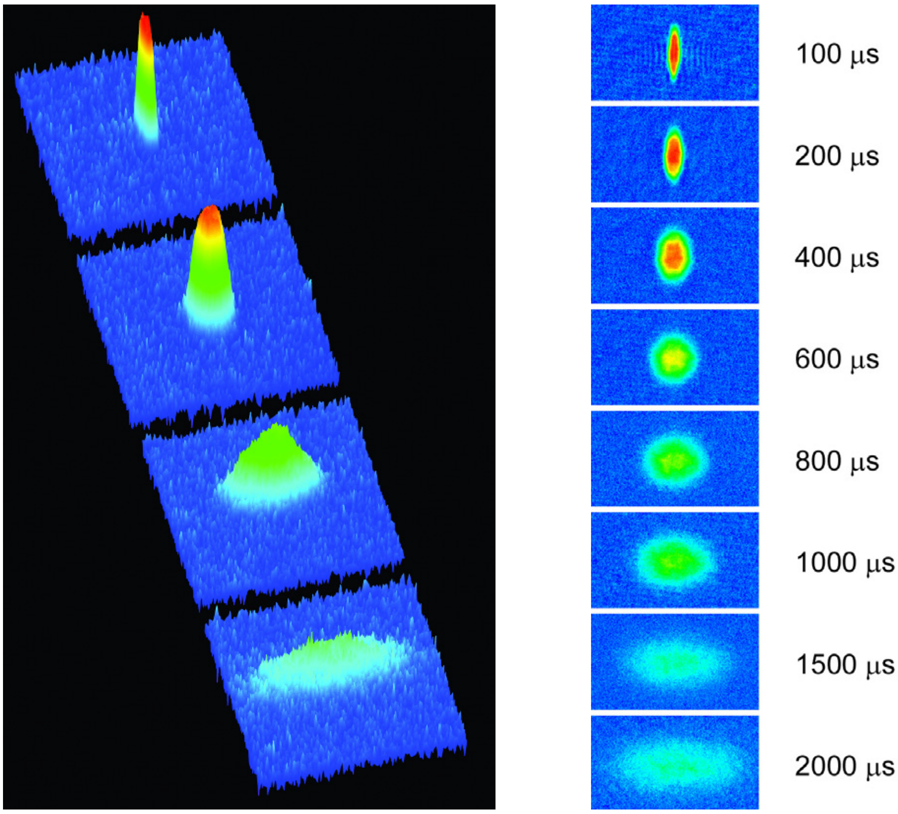

Figure 4. Experimental images. Elliptic flow of a strongly interacting Fermi gas as a function of time after release from a cigar-shaped optical trap: from top to bottom, 100 μs to 2 ms after release. The pressure gradient is much larger in the initially narrow directions of the cloud than in the long direction, causing the gas to expand much more rapidly along the initially narrow directions, inverting the aspect ratio. Achieving a nearly perfect elliptic flow requires extremely low shear viscosity, as is the case for a QGP. The sequence of images is created by recreating similar initial conditions and destructively imaging the cloud at different times after release [46].

to an ellipse, as shown in figure 4. Remarkably, the same type of *elliptic flow* is also observed in the momentum distribution of an expanding QGP produced in an off-center collision of two heavy ions; see section 3.5. There, the temperature is 19 orders of magnitude hotter and the particle density is 25 orders of magnitude greater than that of the Fermi gas. In both systems, however, this nearly perfect ‘elliptic’ flow, figure 4, is a consequence of extremely low-viscosity hydrodynamics, which persists in the normal, non-superfluid unitary gas and deeply connects these two apparently disparate fields.

The creation of a degenerate Fermi gas near a Feshbach resonance was followed in 2003 by the first measurements of the interaction energy [48, 49] and the creation of Bose-condensed dimer molecules [50–53]. In 2004, condensed fermionic atom pairs were observed using a fast magnetic field sweep to project the pairs onto stable molecular dimers [54, 55]. Using this method, the first phase diagram in the crossover region was obtained (see figure 3) as a function of magnetic field and temperature, albeit using the temperature of the ideal gas obtained by an adiabatic sweep to a non-interacting regime [54]. Evidence for superfluidity in a Fermi gas was provided by the measurement of collective mode frequencies and damping rates versus temperature [56] and magnetic field [57], and by the measurement of the pairing gap by RF spectroscopy [58]. The observation of a vortex lattice in 2005 provided a direct verification of Fermi superfluid behavior [59].

Also in 2005, initial thermodynamic measurements were made by adding a controlled amount of energy to the cloud and measuring an empirical temperature from the corresponding cloud spatial profile [60]. However, the results were model dependent, as the calibration of the empirical temperature relied on comparing the measured cloud profiles with theoretical predictions. Model-independent measurements were soon to follow, based on universal behavior in the unitary regime, where the local thermodynamic quantities, such as pressure, are functions only of the density n and the temperature T [61].

Model-independent measurements of the total energy E of a resonantly interacting Fermi gas are based on the Virial theorem, which holds for a unitary gas as a consequence of universality and yields the energy directly from the cloud profile [62]. Using entropic cooling [63, 64], a model-independent measurement of the total entropy S was accomplished by an adiabatic sweep of the bias magnetic field from resonance to the weakly interacting regime, where the entropy can be calculated from the cloud profile [65]. As $T = \partial E/\partial S$, these measurements enabled the first model-independent temperature calibration and estimates of the critical parameters in the strongly interacting regime [66]. A refined temperature calibration is used in the measurement of universal quantum viscosity, as described in this focus issue [2].

Measurements of global thermodynamic quantities from the cloud profiles in the strongly interacting regime are now superseded by model-independent measurements of local quantities [67, 68]. Using the Gibbs–Duhem relation

$$\mathrm{d}P = n\,\mathrm{d}\mu \tag{4}$$

at constant temperature, the local pressure is obtainable from the integrated column density, where the local chemical potential μ is determined by the known trap potential [69, 70]. Combined with a temperature measurement, the local equation of state $P(\mu, T)$ or $P(n, T)$ is determined. The most precise local measurements avoid temperature measurement, which introduces the most uncertainty, by determining the pressure, density and compressibility from the cloud profiles. The resulting equation of state reveals clearly a lambda transition, and provides the best measurement of the critical parameters for a unitary Fermi gas [36], as described in detail in section 2.3. Measurements of equilibrium thermodynamic quantities are now used as stringent tests of predictions, as described in this focus issue by Hu and co-workers [71]. These thermodynamic measurements are connected to universal hydrodynamics and transport measurements, as described in [2].

We proceed to describe the all-optical methods developed at Duke in 2002 [26, 46], as one specific example of experimental techniques, which are closely tied to the theme of this focus issue, namely viscosity and transport measurements on Fermi gases in the universal regime [2]. A degenerate, strongly interacting Fermi gas is readily made by all-optical methods [46]. As sketched in the left panel of figure 5, an MOT is used to pre-cool a 50 : 50 mixture of spin-up and spin-down ^{6}Li atoms, which is loaded into a CO_2 laser optical trap and magnetically tuned to an s-wave Feshbach resonance. Atoms from the source (lower right, green cylinder) form an atomic beam (blue arrow) that is slowed by radiation pressure forces from a resonant laser beam (top left, opposing red arrow). For ^{6}Li, the deceleration is $2 \times 10^6\ \mathrm{m\,s^{-2}}$, slowing the atoms from oven thermal speeds of $1\ \mathrm{km\,s^{-1}}$ to tens of $\mathrm{m\,s^{-1}}$ over a distance of a fraction of a meter. Six laser beams (three thick red lines) then propagate toward the center of the MOT

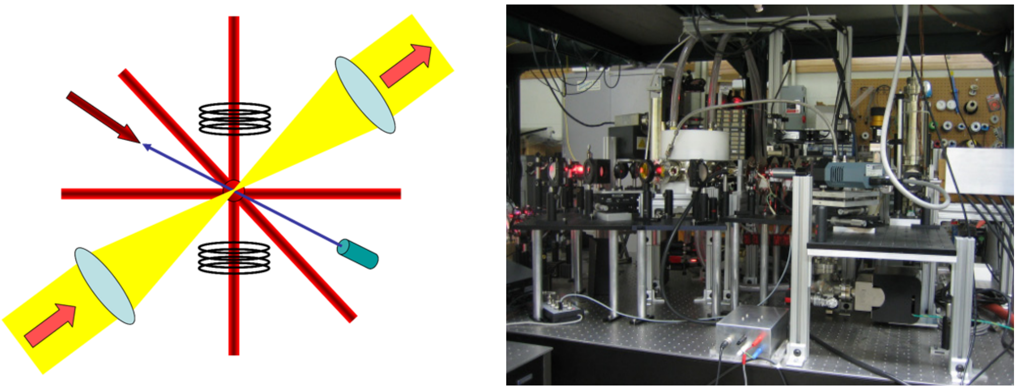

Figure 5. Ultracold quantum gas experimental apparatus. Left: sketch of the experimental apparatus for ultracold fermions. Right: apparatus for the Duke experiments (currently at North Carolina State University). Compare to a sketch of the QGP experiment at the LHC in figure 12: the quantum gas experiment is about ten times smaller (2.5 versus 26 m), but the size of the trapped ultracold gas is 11 orders of magnitude larger (a few hundreds of micrometers versus a few femtometers). The ultracold quantum gas is at nanokelvin temperatures, or pico-eV, compared to the deconfinement temperature of $\simeq 2 \times 10^{12}$ K in the QGP, or 200 MeV, created by colliding gold nuclei at energies of 100 GeV per nucleon.

(point of intersection of three thick red lines), creating inward damping forces that cool the atoms. Opposing magnetic fields, created by two coils (stacked black circles, top and bottom), spatially tune the atomic resonance frequency, causing the six beams to produce a harmonic restoring force at the MOT center. Typical MOT clouds are a few millimeters across and contain several hundreds of million atoms. A trapping laser beam (shown in yellow) is focused (lenses indicated by two light blue ovals) at the MOT center and loaded. After turning off the MOT beams and the MOT magnetic field, an additional bias magnetic field tunes the atoms to a collisional (Feshbach) resonance. Forced evaporation near resonance, by lowering the trap depth, rapidly cools the cloud to quantum degeneracy, i.e. $T/T_{\mathrm{F}} \ll 1$, producing a cigar-shaped cloud that is typically a few microns in diameter and several hundreds of microns long, containing several hundred thousands of atoms. In the right panel of figure 5 is shown the experimental apparatus from the Duke laboratory, currently located at North Carolina State University. From right to left in the photo: the oven assembly where hot fermions are produced (aluminum housing); the camera to produce density images (blue device in the foreground); a Zeeman slower to bring atoms into MOT (middle, behind camera); bias field magnets containing MOT in ultrahigh vacuum (white plastic housings); ZnSe lens for the CO_2 laser trapping beam and optical table (left).

In the simplest case, the optical trap consists of a single laser beam, focused into the center of the MOT and detuned well below the atomic resonance frequency to suppress spontaneous scattering, which would otherwise heat the atoms. For an optical trapping laser detuned below the atomic resonance, the induced atomic dipole moment is in-phase with the trapping laser field, so that the atoms are attracted to the high-field region at the trap focus, i.e. the effective trapping potential is $U = -\alpha\langle E^2\rangle/2$, where the polarizability $\alpha > 0$ and $\langle E^2\rangle$ is proportional to the trap laser intensity, time averaged over a few optical cycles. The effective potential then has the same spatial profile as the intensity of the focused trap laser. For ultracold atoms, the energy per particle is typically quite small compared to the depth of the optical trap, so that the atoms vibrate in a nearly harmonic confining potential. The vibration frequencies of the atoms in each direction ω_i, $i = x, y, z$, of the trap are readily determined by parametric resonance: the trap laser intensity is modulated and the size of the cloud is measured as a function of modulation frequency. When the modulation frequency is twice the harmonic oscillation frequency, the energy of the atoms increases, causing the density profile to increase in size. This method permits a precise characterization of the trap parameters.

All information about the cloud is obtained by absorption imaging: a spatially uniform short (several μs) low-intensity laser pulse is transmitted through the atom cloud, which partially absorbs the light. The shadow cast by the absorption is imaged onto a CCD (charge coupled device) array to record the image, which is analyzed to extract the column density, integrated along the line of sight. This 'laser flash photography' method provides real space images with a resolution of a few microns, on a time scale short compared to the time scale over which the atoms move significantly compared to the spatial resolution. Both non-destructive and destructive imaging techniques are possible. In the latter case the entire cloud is destroyed by the laser pulse in order to make the most complete possible image. One then runs the same experiment multiple times, with nearly identical initial conditions, to obtain an average picture of temporal evolution. While the CCD measures just the density or $g^{(1)}$ correlations, in fact it is possible to extract density–density or $g^{(2)}$ correlations from the noise on an image [72, 73].

What do experimental measurements actually look like? In figure 4 are shown absorption images from the 2002 Duke experiment on elliptic flow [46]. In the experiments, $N = 7.5 \times 10^4$ atoms in each of the two lowest hyperfine states were cooled to degeneracy in a CO_2 laser trap, with a reduced temperature T/T_{FI} between 0.08 and 0.2. Here, $T_{\mathrm{FI}} = \hbar\bar{\omega}(6N)^{1/3}/k_{\mathrm{B}}$ is the Fermi temperature for an ideal gas at the center of a harmonic trap[14], where $\bar{\omega} = 2\pi \times 2160(65)$ Hz is the geometric mean of the trap oscillation frequencies, measured by parametric resonance as described above. For these parameters, the Fermi temperature is $T_{\mathrm{FI}} = 7.9\,\mu$K. For an ideal gas in the trap, the Fermi radii are $\sigma_x = 3.6\,\mu$m in the narrow directions and $\sigma_z = 103\,\mu$m in the long direction.

2.2. *Scale invariance and universality*

Studies of trapped ultracold Fermi gases have provided important information about the phase diagram, the equation of state, transport properties and quasiparticle properties of strongly correlated Fermi gases. This is possible because under the conditions typically encountered in the experiments, local properties of the trapped gas directly correspond to equilibrium properties of the homogeneous Fermi gas. Consider the ground state of N harmonically trapped fermions. Hohenberg and Kohn [74] showed that the solution of the N-body Schrödinger equation is equivalent to the minimum of the energy functional

$$E[n(\mathbf{x})] = \int \mathrm{d}\mathbf{x}\, (\mathcal{E}(n(\mathbf{x})) + n(\mathbf{x})U(\mathbf{x}))\,, \tag{5}$$

where $n(\mathbf{x})$ is the density, subject to the condition $\int \mathrm{d}\mathbf{x}\, n(\mathbf{x}) = N$, $\mathcal{E}(n)$ is the energy density functional and $U(\mathbf{x})$ is the trap potential. If the density is sufficiently smooth we can write $\mathcal{E}(n)$

---

[14] T_{FI} is defined by $T_{\mathrm{FI}} = T_{\mathrm{F}}(n_0(0))$, where $T_{\mathrm{F}}(n) = \hbar^2 k_{\mathrm{F}}(n)^2/(2mk_{\mathrm{B}})$ is the local Fermi temperature of a non-interacting gas evaluated at the center of the trap. An equivalent definition is that $k_{\mathrm{B}}T_{\mathrm{FI}} = E_{\mathrm{FI}}$, where $E_{\mathrm{FI}} = \hbar\bar{\omega}(3N)^{1/3}$ is the Fermi energy of N non-interacting fermions in a harmonic trap. The advantage of T_{FI} is that it depends only on N and $\bar{\omega}$.

as a function of the local density and its gradients. On dimensional grounds we have

$$\mathcal{E}(n(\mathbf{x})) = \frac{c_0 \hbar^2}{m} n(\mathbf{x})^{5/3} + \frac{c_1 \hbar^2}{m} \frac{(\vec{\nabla} n(\mathbf{x}))^2}{n(\mathbf{x})} + O\left(\nabla^4 n(\mathbf{x})\right), \tag{6}$$

where $c_0, c_1, \ldots$ are dimensionless constants. At unitarity the coefficients c_i are pure numbers, but for a finite scattering length they become functions of na^3. To first approximation we can neglect the gradient terms. Then the density is given by $n(\mathbf{x}) = n_{\mathrm{eq}}(\mu - U(\mathbf{x}))$, where $n_{\mathrm{eq}}(\mu)$ is the equilibrium density at the chemical potential μ and zero temperature. This approximation is known as the local density approximation. Gradient terms are suppressed by $(\omega_\perp/\mu)^2 \sim 1/(\lambda_z N)^{2/3}$, where $\lambda_z = \omega_z/\omega_\perp$ is the trap deformation. Typical experiments involve $\lambda_z \simeq 0.025$–0.1 and $N \geqslant 10^5$, so corrections beyond the local density approximation are quite small. These arguments generalize to systems at non-zero temperature. In this case the density of the trapped system is $n(\mathbf{x}) = n_{\mathrm{eq}}(\mu - U(\mathbf{x}), T)$.

The equilibrium density can be determined from the equation of state, $P = P(\mu, T)$, through the thermodynamic relation[15] $n = \partial P/\partial\mu$. In the following we also frequently refer to the relation $P = P(n, T)$ as the equation of state. At unitarity the interaction is scale invariant and the only scales in the many-body system are the interparticle distance $n^{-1/3}$ and the de Broglie wavelength, given in equation (1). Dimensional analysis implies that the equation of state must be of the form

$$P(n, T) = \frac{\hbar^2 n^{5/3}}{m} f(n\lambda_{\mathrm{dB}}^3), \tag{7}$$

where $f(x)$ is a universal function. At zero temperature the pressure is proportional to $n^{5/3}/m$. This implies, in particular, that the pressure is given by a numerical constant times the pressure of a free Fermi gas. The same is true of the energy per particle and the chemical potential. It has become standard to denote the ratio of the energy per particle of the unitary gas and the free Fermi gas as the Bertsch parameter ξ,

$$\frac{E}{N} = \xi \left(\frac{E}{N}\right)_0. \tag{8}$$

Bertsch posed the calculation of the parameter ξ as a challenge problem to the many-body physics community in 1999 [75]. At the time, the problem was stated in the context of a toy model for dilute neutron matter; see section 3.1.

Using thermodynamic identities we can show that equation (7) implies that $P = \frac{2}{3}\epsilon$, where ϵ is the energy density. This relation is analogous to the equation of state of a scale-invariant relativistic gas, $P = \frac{1}{3}\epsilon$, as discussed in section 3.2. For a trapped gas the relation between pressure and energy density implies a Virial theorem: in a harmonic trap, the internal energy of the system is equal to the potential energy of the trapping potential [76, 77],

$$\int \mathrm{d}\mathbf{x}\, \epsilon(\mathbf{x}) = \int \mathrm{d}\mathbf{x}\, n(\mathbf{x}) U(\mathbf{x}). \tag{9}$$

These universal relations have been extended in many ways; see [78] for a review. An important class of relations, discovered by Tan, connects the derivative of thermodynamic quantities with respect to $1/a$ to short-range correlation in the gas. Tan defined the contact density $\mathcal{C}$ via [79–81]

$$\frac{\mathrm{d}\epsilon}{\mathrm{d}(a^{-1})} = -\frac{\hbar^2}{4\pi m}\mathcal{C}, \tag{10}$$

---

[15] Here and in the remainder of this review we have dropped the subscript 'eq'.

where the derivative is taken at constant entropy density. The contact density appears in a number of thermodynamic relations. The universal equation of state, for example, is given by

$$P = \frac{2}{3}\epsilon + \frac{\hbar^2}{12\pi m a}\mathcal{C}. \tag{11}$$

More remarkable is the fact that $\mathcal{C}$ controls short-distance correlations in the dilute Fermi gas. The tail of the momentum distribution is given by

$$n_\sigma(k) \to \frac{C}{k^4} \quad \left(|a|^{-1} \ll k \ll r_0^{-1}\right), \tag{12}$$

where $C = \int \mathrm{d}^3x\, \mathcal{C}(x)$ is the integrated contact, $n_\sigma(k)$ is the momentum distribution[16] in the spin state σ and r_0 is the range of the interaction. There are similar expressions for the asymptotic behavior of other correlation observables such as the static and dynamic structure factors, and the dynamic shear viscosity; see [82] for a review. In this focus issue, Kuhnle *et al* [83] present a comprehensive set of measurements of the contact as a function of interaction strength and reduced temperature. These results can be compared to new theoretical predictions discussed by Hu and co-workers [71].

Below the critical temperature for superfluidity, the superfluid flow velocity v_s can be viewed as an additional thermodynamic variable. The response of the pressure to the superfluid velocity defines the superfluid mass density

$$\rho_\mathrm{s} = m n_\mathrm{s} = -\left.\frac{\partial^2 P}{\partial v_\mathrm{s}^2}\right|_{v_n=0}, \tag{13}$$

where the derivative is taken in the rest frame of the normal fluid, meaning that $v_n = 0$. The superfluid mass density can be measured using rotating clouds [84]. The second moment of the trap integrated value of the superfluid mass density determines the quenching of the moment of inertia. New measurements of the moment of inertia can be found in [85].

For small values of $n|a|^3$ the equation of state $P(n, T)$ can be computed in perturbation theory. This program was initiated by Lee and Yang [86] and Huang and Yang [87]. At unitarity, weak coupling methods can be used in the limit of high temperature. This is based on the observation that the binary cross section at unitarity is $\sigma = 4\pi/k^2$. At high temperature the mean momentum is large and the average thermal cross section is small. The equation of state can be written as an expansion in $n\lambda_\mathrm{dB}^3$, which is the well-known Virial expansion. We have

$$P = nk_\mathrm{B}T\left\{1 + b_2(n\lambda_\mathrm{dB}^3) + O((n\lambda_\mathrm{dB}^3)^2)\right\}, \tag{14}$$

with $b_2 = -1/(2\sqrt{2})$ at unitarity [88, 89]. Analytic approaches in the non-perturbative regime $n\lambda_\mathrm{dB}^3 \sim 1$ are based on extrapolating to the unitary limit from different regimes in the phase diagram. For this purpose the phase diagram has been studied as a function of the strength of the interaction, the number of species and the number of spatial dimensions. The oldest theory of this type is the Noziéres–Schmitt–Rink (NSR) theory [90–92], which is based on a set of many-body diagrams that correctly describe both the BCS and BEC limits. NSR theory works surprisingly well, despite the formal lack of a small parameter at unitarity. For example, the basic form of the critical temperature sketched in figure 3 is correctly reproduced. Modern theories of this type are typically based on self-consistent T-matrix approximations; see [93, 94]. Another idea is to generalize the unitary Fermi gas to $2N$ spin states [95]. Mean field theory is reliable in the limit $N \to \infty$, and the interesting case $N = 1$ can be studied by expanding in $1/N$. This method is of interest in connection with holographic dualities, because the gravitational dual is expected to be classical in the limit that the number of degrees of freedom is large. Finally, it was proposed to use the number of dimensions as a control parameter. The unitary limit is perturbative in both $D = 2$ and 4 spatial dimensions [96]. The interesting case $D = 3$ can be studied as an expansion around $D = 2 + \epsilon$ or $4 - \epsilon$ dimensions [97].

---

[16] The momentum distribution is normalized as $\int \mathrm{d}\mathbf{k}/(2\pi)^3 n_\sigma(k) = N_\sigma$, where N_σ is the total number of atoms in the state σ.

These methods are promising, but currently the only techniques that provide reliable and systematically improvable results in the regime $n\lambda_{\mathrm{dB}}^3 \simeq 1$ are quantum Monte Carlo calculations. At zero temperature the standard technique is Green function Monte Carlo [98, 99]. This method relies on a variational initial wavefunction, which is used as the initial condition for an imaginary time diffusion process. The Monte Carlo method suffers from a fermion sign problem which is addressed using the fixed node approximation. At finite temperature a number of groups have performed imaginary time path integral Monte Carlo calculations [100–103]. These calculations do not rely on variational input, and they do not suffer from a sign problem, but they are formulated on a space–time lattice and require an extrapolation to zero lattice spacing. One new technique is bold diagrammatic Monte Carlo, which is based on sampling the sum of all Feynman diagrams. The method suffers from a sign problem, but convergence in the regime above the critical temperature for superfluidity was found to be very good [104].

2.3. *Experimental determination of the equation of state*

The equation of state describes a functional relation between key thermodynamic variables, such as the pressure $P(\mu, T)$ as a function of chemical potential μ and temperature T. In Fermi gases, as stated in section 2.1, what we can actually measure is the density profile of a trapped cloud. There are various techniques for translating density measurements into thermodynamic quantities, all relying on the local density approximation.

The first thermodynamic measurements took place at Duke in 2005, where global thermodynamic variables were measured. In the initial experiments [60], a controlled amount of energy was added to the trapped cloud by abruptly releasing it from the optical trap, allowing it to expand hydrodynamically by a known factor, and then recapturing the cloud after a selected expansion time, thereby increasing the potential energy. After allowing the gas to equilibrate, the cloud profile was measured to determine an empirical temperature, which was later calibrated by comparing to theoretical cloud profiles, predicted as a function of reduced temperature. The resulting energy versus temperature curve was compared to that measured for an ideal Fermi gas in the same trap, and showed a departure from ideal gas behavior at a certain temperature, which yielded an estimate of the superfluid-to-normal fluid transition temperature. However, the results suffered from being model dependent, as calibration of the empirical temperature relied on a comparison of the measured cloud profiles with theoretical predictions. To avoid this model dependence, in 2006 the JILA group measured the potential energy of the strongly interacting cloud of ^{40}K as a function of the ideal Fermi gas temperature that was obtained after an adiabatic sweep of the bias field to the non-interacting regime above resonance [105].

Model-independent determination of thermodynamic quantities was done by the Duke group in 2007 [65], where the total energy E and the total entropy S of a trapped cloud were measured from cloud profiles, by exploiting the universal behavior of a unitary Fermi gas at a Feshbach resonance. From equation (9) we know that for a harmonic trapping potential the total energy is twice the average potential energy, $E = 2\langle U\rangle = 3\, m\omega_z^2\langle z^2\rangle$. Hence, by measuring the harmonic oscillator frequency and mean square cloud size, the total energy is readily determined

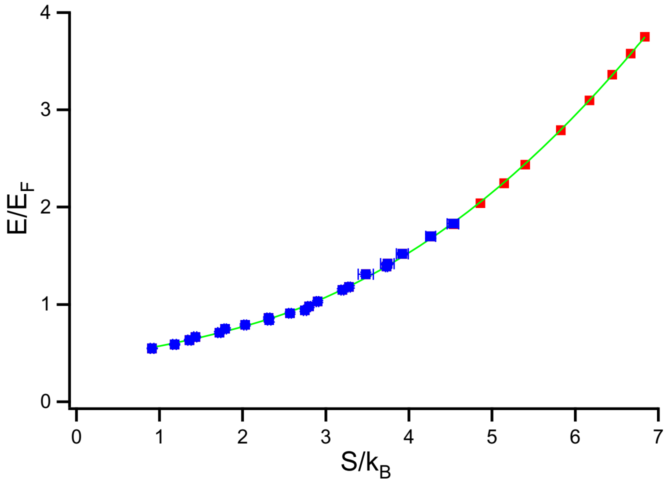

Figure 6. Total energy per particle of a strongly interacting Fermi gas in the universal regime versus the entropy per particle. The blue dots show the entropy obtained by adiabatically sweeping the magnetic field from 840 to 1200 G, where the gas is weakly interacting. The red dots are the theoretical calculations using a second virial coefficient approximation. The green curve is a fit using two power laws, which determines the temperature $T = \partial E/\partial S$. Figure reproduced from [2] with permission.

from cloud images. This method was demonstrated experimentally in [62]. With the energy in the strongly interacting regime measured from cloud images, the entropy is determined by adiabatically sweeping the magnetic field to a weakly interacting regime. There, the entropy and mean square cloud size are readily calculated as a function of the temperature, yielding the entropy as a function of the mean square cloud size of the weakly interacting gas. Adiabatic behavior is verified by a round-trip sweep. Figure 6 shows the energy per particle as a function of the entropy per particle, in the universal regime. The temperature is determined by fitting a smooth curve [2, 66]. Hu *et al* [106] combined the demonstrated universal behavior of the global thermodynamic quantities by reanalyzing the measurements in ^{40}K and showing that the ^{40}K and ^{6}Li data fit on a single thermodynamic curve.

As already mentioned briefly in section 2.1, model-independent measurement of global thermodynamic variables was superseded in 2010 by model-independent measurement of local thermodynamic quantities, which can be directly compared to predictions, within the local density approximation. Equation (4) determines the local pressure P from the local density n and the local chemical potential, $\mu = \mu_g - U$, where μ_g is the global chemical potential and U is the known trapping potential. Absorption imaging directly yields the column density $\tilde{n}(x, z) = \int \mathrm{d}y\, n(x, y, z)$ for an imaging beam propagating along y. In a cylindrically symmetric harmonic trap $U = m\omega_\perp^2(x^2 + y^2)/2 + m\omega_z^2 z^2/2$, where $\omega_\perp$ and ω_z are the radial and axial trapping frequencies, respectively, we can write $-\mathrm{d}\mu = \mathrm{d}U = m\omega_\perp^2 \rho\, \mathrm{d}\rho = m\omega_\perp^2\, \mathrm{d}x\, \mathrm{d}y/(2\pi)$ from which we see that the pressure is determined from the doubly integrated density, i.e. the integrated column density $\bar{n}(z) = \int \mathrm{d}x\, \mathrm{d}y\, n(x, y, z)$. For a $50:50$ mixture of spin-up and spin-down fermions with n the total density, the pressure is determined as a function of z, $P(z, T) = m\omega_\perp^2\, \bar{n}(z)/(2\pi)$. For $x = y = 0$, the corresponding chemical potential is $\mu(z) = \mu_g - m\omega_z^2 z^2/2$, so that $P(\mu, T)$ is determined if the temperature and global chemical potential can be determined.

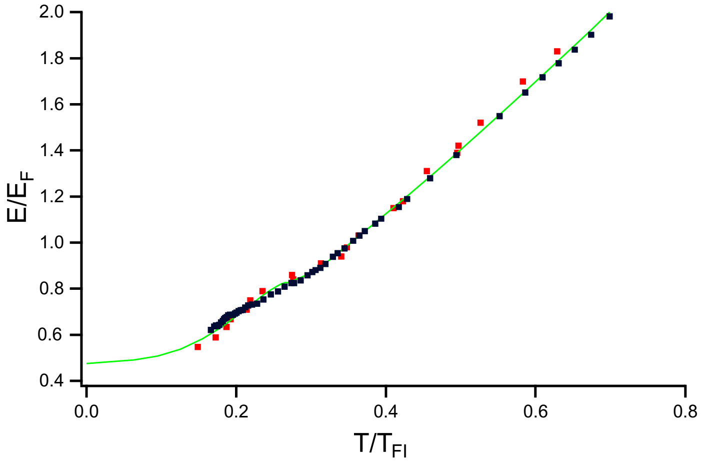

Figure 7. Measured energy versus the temperature obtained from the calibration of [2] (red dots); for comparison, we show the data obtained by the ENS group [68] (black dots) and the theory of Hu and co-workers [71] (green curve). Figure reproduced from [2], with permission.

To determine the temperature, the Tokyo group [67] used the temperature calibration by the Duke group to obtain T from the total energy and hence from the mean square cloud size [66]. An improved version of this calibration is described in this focus issue [2]. The ENS group [68] directly determined the temperature by using a weakly interacting ^{7}Li impurity to measure the temperature of a strongly interacting ^{6}Li gas. The global chemical potential was determined from the wings of the density profile using a fit based on a Virial expansion, yielding $P(\mu, T)$, from which all other local thermodynamic quantities, such as the energy density and entropy density, were determined. These results enable a direct comparison with predictions. Further, the total entropy S and the total energy E of the trapped cloud are readily determined by integration, and can be compared with the corresponding global quantities measured by the Duke group. The total energy versus temperature is shown in figure 7, where the improved temperature calibration of [2] from this focus issue is displayed for the Duke data. As quite different methods were employed to make the measurements, their close agreement indicates the correctness of the data.

The latest studies of local thermodynamics by the MIT group [36] use refined methods, where measurements of the isothermal compressibility κ directly from the density profiles replace temperature measurements, yielding an equation of state $n(P, \kappa)$. This eliminates the determination of the temperature and local chemical potential from fits at the edges of the cloud, which produced the most uncertainty in previous work. For this method, the three-dimensional (3D) density $n(x, y, z)$ is determined by tomographic imaging, using an inverse Abel transform [107] to determine n from the measured column density. The trap potential is carefully characterized by determining the surfaces of constant density and the hence constant chemical potential, in a very shallow trap where the axial z trap potential is almost perfectly harmonic and therefore known. Equation (4) yields the pressure

$$P(U, T) = \int_{-\infty}^{\mu} \mathrm{d}\mu \, n(\mu', T) = \int_{U}^{\infty} \mathrm{d}U \, n(U, T), \tag{15}$$

where the unknown global potential μ_g is not needed, since the integral is over the known trap potential [36]. The compressibility is the change in density with respect to a change in the local

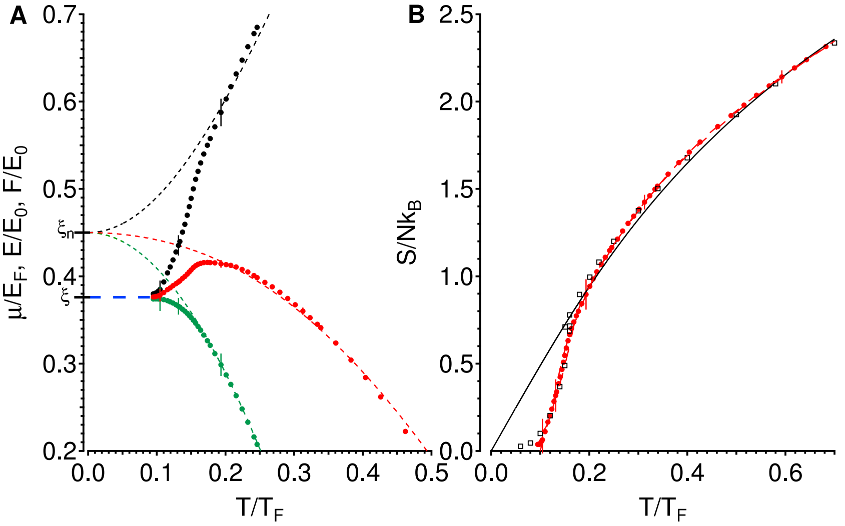

Figure 8. Experimentally measured thermodynamics for a unitary Fermi gas at MIT. Left panel: energy, free energy and chemical potential as a function of temperature. Right panel: entropy as a function of temperature. See the text for an explanation of the curves. Figure reproduced with permission from [36].

trapping potential U, and is determined from the density profile by

$$\kappa = \frac{1}{n}\frac{\partial n}{\partial P}\bigg|_T = -\frac{1}{n^2}\frac{\partial n}{\partial U}\bigg|_T, \tag{16}$$

since $\mathrm{d}P = n\,\mathrm{d}\mu = -n\,\mathrm{d}U$ at constant T. Thus, the actual observed equation of state is the functional relation $n(\kappa, P)$, measured directly from the density distribution, the clear and direct experimental observable for ultracold quantum gases as discussed in section 2.1. From $n(P, \kappa)$, all other local thermodynamic quantities are determined [36], as shown in figure 8.

These experiments provide the best current value for the Bertsch parameter of $\xi = 0.376(5)$, which is consistent with the value obtained in measurements of global quantities, $\xi = 1 + \beta = 0.39(2)$ [66]. These measurements can also be compared with theory predictions. The two most recent quantum Monte Carlo calculations give $\xi \leqslant 0.383(1)$ [108] and $\xi = 0.3968^{+0.0076}_{-0.0077}$ [109]. See [109] for an extensive compilation of analytic results and earlier Monte Carlo calculations.

Figure 8 shows several representations of the equation of state, and provides a glimpse of the level of precision that can be achieved in present experiments. In the left panel are shown the chemical potential μ, energy E and free energy F. The right panel shows the entropy per particle $S/(Nk_\mathrm{B})$ versus T/T_F. The chemical potential (red solid circles) is normalized by the Fermi energy; energy (black solid circles) and free energy (green solid circles) are normalized by $E_0 = \frac{3}{5}NE_\mathrm{F}$, which is the energy per particle in a uniform Fermi gas. At high temperatures all quantities approximately track those for a non-interacting Fermi gas, shifted by $\xi_n - 1$ with $\xi_n \simeq 0.45$ (dashed curves). The peak in the chemical potential roughly coincides with the onset of superfluidity. In the very low-temperature regime, μ/E_F, E/E_0 and F/F_0 all approach ξ (blue dashed line). At high temperatures, the entropy closely tracks that of a non-interacting Fermi gas (black solid curve). The open squares are from the self-consistent T-matrix calculation [94]. A few representative error bars are shown, representing means ± standard deviation.

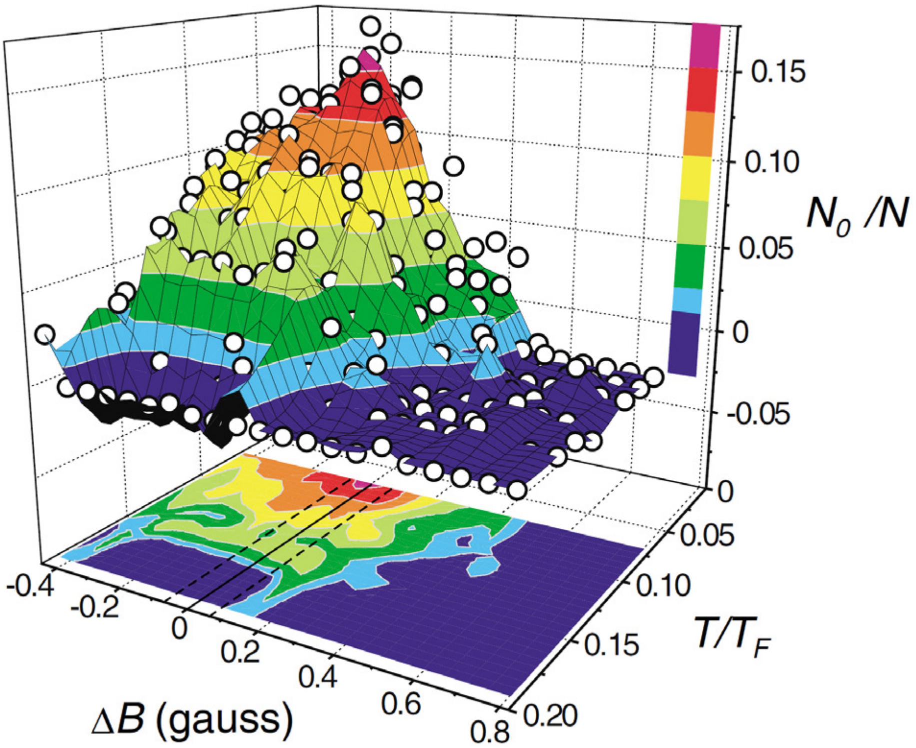

Figure 9. Experimental Fermi gas phase diagram. Experimental measurements of the BCS–BEC crossover for ultracold Fermi gases (compare to the sketch in figure 3). Since a magnetic Feshbach resonance is used to tune the interaction strength, one axis is magnetic field. In this measurement, the BCS side is on the $B > 0$ side, so figure 3 should be reversed for comparison; it is not possible to go deep into the BEC side because the lifetimes become too short in this experiment. $\Delta B < 0.6$ contains the strongly interacting region $k_F|a| > 1$, and the dashed lines indicate uncertainty in the precise position of the Feshbach resonance. Figure reproduced with permission from [54].

2.4. Experimental studies of the phase transition

The first measurements of the phase transition at unitarity were made at JILA, shown in figure 9, by a pair projection technique. After the pairs were created in the BCS and unitary regimes, a rapid magnetic field sweep was used to pair-wise project the fermions onto molecules to protect them during subsequent expansion measurements, where the molecular momentum distribution was measured to determine the condensed pair fraction. The fraction of near-zero momentum molecules is interpreted as the percentage of condensed Fermi unitary or BCS pairs. In this early experiment, ^{40}K was used, and the measured temperature was that of an ideal Fermi gas, rather than the temperature of the interacting gas. This ideal Fermi gas temperature was obtained by ballistic expansion after an adiabatic sweep to the weakly interacting regime above resonance, yielding the condensed pair fraction versus ideal gas temperature and magnetic field. Despite the uncertainties regarding the temperature and the pair conversion efficiency we observe that the shape of the transition line is qualitatively consistent with theoretical expectations, as summarized in figure 3.

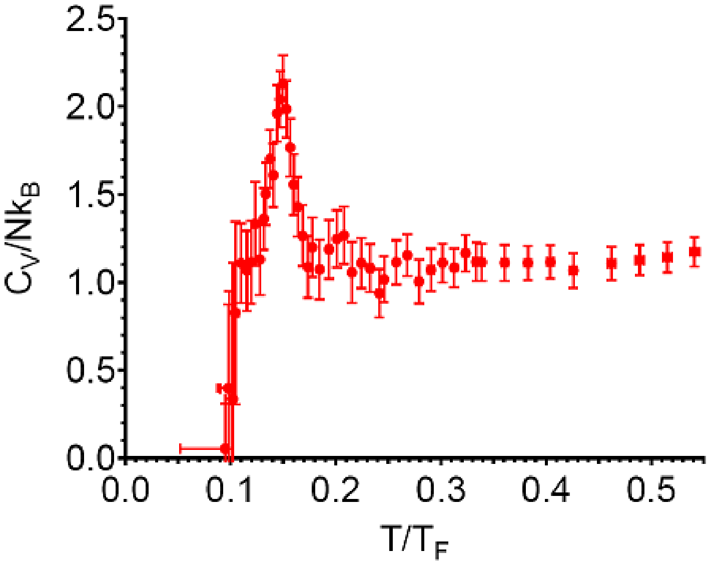

Figure 10. Superfluid phase transition of a unitary Fermi gas. Specific heat per particle, $C_V/(Nk_B)$, as a function of quantum degeneracy, T/T_F. The phase transition is clearly evident at $T_c/T_F = 0.167(13)$. Figure adapted with permission from [36].

Initial estimates of the critical parameters of the trapped gas were done by the Duke group, first based on model-dependent measurements of the energy versus temperature [60] and later based on model-independent measurements of the total energy E and entropy S, obtained as described above. Assuming that a phase transition would be manifested in a change in the scaling of E with S, two power laws were used to fit the $E(S)$ data, one to fit the high-temperature data and the other to fit the low-temperature data, which joined at a point S_c. The continuity of the energy and temperature was used as a constraint, and the critical entropy S_c was estimated from the joining point that minimized the χ^2 for the fit [66].

This fit method is useful for calibrating the temperature, as described in [2] of this focus issue. However, the fitted critical parameters are dependent on the form of the fit function, making the uncertainty difficult to quantify. Further, the global observables suffer from the trap averaging that masks an abrupt local phase transition, initially near the trap center. Hence, the curve of total energy versus entropy is too smooth to obtain reliable estimates of the critical parameters.

Using the refined local measurements described in section 2.3, the MIT group has traced out the phase transition in the unitary regime, without any need for rapid magnetic sweeps, fitting parameters, thermometry, or indeed any kind of theory besides elementary thermodynamic considerations. The main goal of these experiments was to obtain a cusp-like signal of the phase transition to superfluidity in a unitary Fermi gas, by focusing on second-order derivatives of the pressure, where a cusp appears. Although superfluidity was established by the creation of vortex lattices [59], the actual phase transition itself had been only indirectly observed.

In figure 10 is shown the specific heat per particle, clearly displaying the superfluid phase transition. Experimentally, the specific heat is derived from the compressibility κ and the pressure P [36],

$$\frac{C_V}{k_B N} = \frac{5}{2}\frac{T_F}{T}\left(\tilde{p} - \frac{1}{\tilde{\kappa}}\right), \tag{17}$$

where $\tilde{\kappa} = \kappa/\kappa_0$ and $\tilde{p} \equiv P/P_0$ are normalized to the non-interacting Fermi gas compressibility $\kappa_0 = \frac{3}{2}\frac{1}{nE_F}$ and pressure $P_0 = \frac{2}{5}nE_F$, respectively. Using $n = (\partial P/\partial\mu)_T$, the compressibility can be written as $\kappa = (1/n^2)(\partial n/\partial\mu)_T = (1/n^2)(\partial^2 P/\partial^2\mu)_T$. As κ is a second derivative of the pressure, the specific heat shows a clear cusp-like signature (figure 10). Qualitatively, the behavior of C_V can be understood as follows: as one approaches the phase transition from above, $T/T_c > 1$, the compressibility increases due to the attraction between fermions; below the phase transition, $T/T_c < 1$, the compressibility decreases because fermions are bound into pairs, and it becomes more difficult to squeeze the gas, i.e. to change the single-particle density.

2.5. *Universal hydrodynamics and transport*

Transport properties of the unitary Fermi gas are of interest for several reasons. The first reason is related to the main theme of this review: holographic dualities suggest a new kind of universality in the transport properties of strongly interacting quantum fluids. We expect, in particular, that the shear viscosity to entropy density ratio is close to the value $\eta/s \sim \hbar/(4\pi k_B)$ originally discovered in the QGP, and first obtained theoretically using the AdS /CFT correspondence [11], where AdS is a special maximally symmetric space–time described in detail in section 4.2, and CFT stands for *conformal field theory*. The second reason is that transport properties are very sensitive to the strength of the interaction, and the types of quasiparticles present in the system. The Bertsch parameter, which characterizes the effect of interactions on the energy per particle, varies by about a factor of two between the weak coupling (BCS) and strong coupling (unitarity) limits. The shear viscosity, on the other hand, changes by many orders of magnitude. Finally, quantum-limited transport has also been observed in systems that are of great practical significance, in particular in the strange metal phase of high-T_c compounds; see the contribution by Guo *et al* to this focus issue [110].

Transport properties have been studied experimentally by exciting hydrodynamic modes, such as collective oscillations [56, 111–114], collective flow [46, 115], sound [116] and rotational modes [117]. In a system that can be described in terms of quasiparticles the hydrodynamic description is valid if the Knudsen number $Kn = l_{\mathrm{mfp}}/L$, the ratio of the mean free path l_{mfp} to the system size L, is small[17]. In the unitary gas the mean free path is $l_{\mathrm{mfp}} = 1/(n\sigma)$, where n is the density and $\sigma = 4\pi/k^2$ is the universal cross section. In the high-temperature limit the thermal average cross section is $\sigma = 4\lambda_{\mathrm{dB}}^2$. The Knudsen number of a unitary gas confined in a cigar-shaped harmonic trap is

$$Kn = \frac{3\pi^{1/2}}{4(3\lambda_z N)^{1/3}}\left(\frac{T}{T_{\mathrm{FI}}}\right)^2, \tag{18}$$

where we have taken L as the radius in the narrow or z direction. Here, N is the number of particles, λ_z was defined previously as the aspect ratio of the trap and T_{FI} is the global

---

[17] Criteria for the validity of hydrodynamics can also be formulated if there is no underlying quasiparticle description, a situation that is of great interest in connection with holographic dualities. In this case, hydrodynamics is based on a gradient expansion of the conserved currents. The ratio of the $O(v)$ to $O(\partial v)$ terms in the stress tensor is known as the Reynolds number, $Re = vLmn/\eta$. Validity of the gradient expansion requires that the Reynolds number be large.

Fermi temperature for a harmonically trapped ideal gas; see section 2.1. Using $N = 2 \times 10^5$ and $\lambda_z = 0.045$ as in [113], we conclude that hydrodynamics is expected to be valid for $T \lesssim 5T_{\mathrm{FI}}$. For the Fermi gas viscosity measurements described in this focus issue [2, 115], the maximum temperature is $T \simeq 1.5\, T_{\mathrm{FI}}$ and $Kn \leqslant 0.09$.

Nearly ideal hydrodynamic behavior was first observed in the expansion of a unitary Fermi gas after release from a deformed trap [46]; see figure 4. For a ballistic gas the expansion reflects the isotropic local momentum distribution in the trap. As a result the gas expands in all directions and the cloud slowly becomes spherical. For a hydrodynamic system the expansion is driven by gradients in the pressure. In the case of a deformed cloud the gradients are largest in the short direction of the trap, and the expansion takes place mostly in the transverse direction. As a result, the cloud eventually becomes elongated along what was originally the short direction. This phenomenon is analogous to the elliptic flow observed in heavy-ion collisions, as described in section 3.5. What is also remarkable is the fact that even though the gas becomes more dilute as it expands, this effect is compensated for by the growth in the mean cross section. As a consequence, the gas remains hydrodynamic throughout the expansion. Ballistic behavior sets in eventually only because of imperfections, such as the fact that the scattering length is not truly infinite.

The role of dissipative effects, in particular shear viscosity, was first studied in collective modes. The radial breathing mode can be excited by removing the confining potential, letting the gas expand for a short period of time, and then restoring the potential. Hydrodynamic behavior can be established by measuring the frequency of the breathing mode. For an ideal fluid $\omega = \sqrt{10/3}\,\omega_0$, whereas in a ballistic system $\omega = 2\omega_0$ [118, 119]. The transition from ballistic behavior in the BCS limit to hydrodynamics in the unitary limit was observed experimentally in [111, 112]. In the hydrodynamic regime damping is expected to be dominated by dissipative terms in the equations of fluid dynamics. The energy dissipation is given by

$$\dot{E} = -\int d\mathbf{x} \left\{ \frac{\eta(\mathbf{x})}{2} \left(\nabla_i v_j + \nabla_j v_i - \frac{2}{3}\delta_{ij} (\nabla \cdot \mathbf{v}) \right)^2 + \zeta(\mathbf{x})\, (\nabla \cdot \mathbf{v})^2 + \frac{\kappa(\mathbf{x})}{T} (\nabla T)^2 \right\}, \tag{19}$$

where v_i is the fluid velocity, η is the shear viscosity, ζ is the bulk viscosity, κ is the thermal conductivity, and all derivatives, divergences and gradients are spatial. At unitarity the system is scale invariant and the bulk viscosity is expected to vanish [78, 120]. This prediction was experimentally checked by Cao *et al* [121] and Dusling and Schäfer [121]. Thermal conductivity is not important because the system remains isothermal in ideal hydrodynamics. Temperature gradients only appear due to shear viscosity, and their contribution to dissipation is higher order in the gradient expansion. This means that damping is dominated by shear viscosity.

Collective modes in ideal fluid dynamics are described by scaling solutions of Euler's equation. This means that the shape of the density profile does not change during the evolution, and that the velocity is linear in the coordinates. The solution is analogous to Hubble flows in cosmology and the Bjorken expansion of a QGP, as discussed in section 3.5. In the case of a scaling solution, the shear stresses $\partial_i v_j$ are spatially constant, and the energy dissipated only depends on the spatial integral of η. On dimensional grounds we can write $\eta = \hbar n \alpha_n$. For a scale invariant system α_n is only a function of the dimensionless variable $n^{2/3}\hbar^2/(mk_{\mathrm{B}}T)$, i.e. a function of the reduced temperature T/T_{F}. The reduced temperature varies across the trap, but for a given fluid element it remains approximately constant during the hydrodynamic evolution of the system. This implies that the damping constant of a collective mode is related to the spatial average of the shear viscosity in the initial equilibrium state

$$\langle\alpha_n\rangle = \frac{1}{N}\int d\mathbf{x}\,\eta(x). \tag{20}$$

The extracted values of $\langle\alpha_n\rangle$ can be converted into the trap averaged shear viscosity to entropy density ratio by using the measured entropy per particle. This type of analysis was originally carried out in [122, 123], where it was observed that $\eta/s \lesssim 0.5\hbar/k_B$ in the vicinity of T_c. More recently, Cao *et al* [115] showed that in the high-temperature regime α_n exhibits the scaling behavior expected from the solution of the linearized Boltzmann equation. The shear viscosity due to elastic two-body scattering has the form $\eta \simeq npl_{\rm mfp}$, where $p \sim \lambda_{\rm dB}^{-1}$ is the mean quasiparticle momentum. As we saw above, $l_{\rm mfp} \sim 1/(n\sigma) \sim 1/(n\lambda_{\rm dB}^2)$. This implies that at high temperature $\eta \sim \lambda_{\rm dB}^{-3}$. The coefficient of proportionality was determined by Bruun and co-workers [8] and Brun and Smith [124]. They find that

$$\eta = \frac{15}{32\sqrt{\pi}}\frac{(mk_BT)^{3/2}}{\hbar^2}. \tag{21}$$

There is an important problem related to equation (21) that affects the extraction of the shear viscosity from experiments with scaling flows. Equation (21), which is reliable in the high-temperature or low-density part of the cloud, is independent of the density[18]. As a consequence the integral in equation (20) diverges in the low-density region. A solution to this problem was proposed in [126, 127]: in the low-density regime, the viscous relaxation time $\tau_R \simeq \eta/(nk_BT)$, which is the time it takes for the dissipative stresses to relax to the Navier–Stokes form $\eta(\nabla_i v_j + \nabla_j v_i - \frac{2}{3}\delta_{ij}\nabla_k v_k)$, becomes very large. Since the dissipative stresses are initially zero, taking relaxation into account suppresses the contribution from the dilute corona. A simplified version of this approach was used in Cao *et al* [2, 115]. The relaxation time and its relation to the spectral function of the shear tensor are discussed in the contribution by Braby *et al* [128].

The shear viscosity drops with temperature and is expected to reach a minimum near T_c. In this regime quantum effects are important. T-matrix calculations can be found in [129] and in the contributions by Guo *et al* and LeClair in this focus issue [110, 130]. It is also possible that in this regime quasiparticle descriptions break down completely, and the most efficient description of the unitary gas near T_c is in terms of a suitable weakly coupled holographic dual. Progress toward constructing holographic duals of non-relativistic quantum fluids is summarized in section 4.3.4.

In the low-temperature superfluid regime the appropriate description is superfluid (two-fluid) hydrodynamics [131]. Superfluid hydrodynamics predicts the existence of additional hydrodynamic modes, in particular second sound, and contains additional transport coefficients that come into play if there is relative motion between the superfluid and normal components of the fluid. There are proposals for exciting second sound modes in the literature [132], but these ideas have not yet been confirmed. At very low temperature the shear viscosity is expected to be dominated by phonons, similar to liquid helium or dilute Bose gases. Elastic phonon scattering gives $\eta \sim 1/T^5$ [133], whereas inelastic processes can give a slower increase at low temperature, $\eta \sim 1/T$ [9]. These predictions are difficult to verify experimentally because the phonon free path is quite large.

---

[18] This is a general property of the viscosity of dilute gases, and was first noticed by Maxwell. The result was experimentally confirmed by Maxwell himself, who measured the damping of oscillating discs in a partially evacuated container [125].

The short mean free path in the strongly interacting normal fluid suggests that not only the viscosity, but also other transport coefficients may exhibit universal behavior. The spin diffusion constant was recently studied by Sommer *et al* [134, 135]. In the first paper, Sommer *et al* observed collisions between polarized Fermi gas clouds. The colliding clouds are initially very far from equilibrium, but at late times the system relaxes diffusively. The corresponding relaxation time can be used to measure the spin drag and the spin diffusion constant D_s. The spin diffusion constant in the homogeneous system is defined by Fick's law,

$$\jmath_s = -D_s \nabla M, \tag{22}$$

where $\jmath_s$ is the spin current, and $M = n_\uparrow - n_\downarrow$ is the polarization. In this focus issue, Sommer *et al* follow up on these studies by measuring the damping of the spin dipole mode in strongly polarized gases [135]. A theoretical study of the spin drag relaxation rate for a repulsive gas is presented by Duine *et al* [136] in this focus issue. They show that spin fluctuations enhance the spin drag in the vicinity of the Stoner ferromagnetic transition. Similarities in the spin transport in the unitary gas and graphene are studied by Müller and Nguyen in this focus issue [137].

A calculation of the diffusion constant based on the two-body Boltzmann equation can be found in [134, 138]. They find that

$$D_s = \frac{9\pi^{3/2}}{32\sqrt{2}} \frac{\hbar}{m} \left(\frac{T}{T_F}\right)^{3/2}. \tag{23}$$

Similar to the shear viscosity, the spin diffusion constant drops with decreasing temperature. Near the critical temperature D_s is expected to approach the universal value $D_s \sim \hbar/m$. This behavior is indeed observed in the experiment [134]; see also the recent analysis of Bruun and Pethick [139]. It is interesting to compare this result with the observed minimum of the shear viscosity. Shear viscosity governs the rate of momentum diffusion. The associated diffusion constant is $D_\eta = \eta/(mn)$. Near T_c we have $\eta/s \simeq 0.5\hbar/k_B$ and $s/n \simeq k_B$. This implies that $D_\eta \simeq 0.5\hbar/m$, comparable to what is seen in spin diffusion. A similar correlation between shear viscosity and diffusion was observed in the QGP, as discussed in section 3.7.

2.6. The Bardeen–Cooper–Schrieffer to Bose–Einstein condensate crossover in lattices

A question quite distinct from that of the BCS–BEC crossover in ultracold Fermi gases in the continuum is how fermions pair, and reach the unitary regime, in a discrete context. Optical lattices, or crystals of light, are created from interfering laser beams. They make a sinusoidal standing wave that traps fermions via the ac Stark effect. The strongly discretized regime, i.e. in the lowest band(s) and tight binding approximations, is quite distinct from that explored by both the QGP and continuum ultracold quantum gas experiments. The essential features of weakly interacting fermions in this context have already been explored and understood [140]; however, the strongly interacting regime remains in question. Moreover, the lattice regime presents an additional challenge for holographic duality, as it too has a unitary point where $1/(k_F a) \to 0$. The question of how to correctly model this problem in the context of a lattice is subtle, as simply attaching a band index to fermionic fields leads to hundreds of bands deep in the BEC regime, and is therefore numerically and practically intractable.

In the continuum, a single-channel model treating only fermions qualitatively reproduces the phase diagram of figures 3 and 9. This qualitative approach can be taken, for instance, by using the generalized BCS ansatz, technically valid only for weak interaction and high density, at the mean-field level for arbitrary interaction and finding the chemical potential self-consistently by fixing the average number of particles [141]. This method was first used by Eagles in the context of superconductivity in low carrier concentration systems [142] and later by Leggett [143] and Noziéres and Schmitt-Rink [90] explicitly for the BCS–BEC crossover at zero temperature and finite temperature, respectively, together called NSR theory. In NSR one begins with a lattice, but then quickly takes a continuum or low-momentum limit. Interestingly, if this limit is not taken one obtains a completely incorrect prediction: instead of the BEC critical temperature tending to a constant non-zero value as $1/(k_{\mathrm{F}}a) \to \infty$, as in figure 3, the critical temperature tends to zero algebraically. This unphysical tendency is because in order to hop or tunnel between lattice sites, a pair of fermions must be broken and tunnel one by one. Then for very strong pair binding energy the process is prohibitively expensive. In fact, even on the BCS side such models utilizing a single band have been shown to fail quantitatively for quite small values of $k_{\mathrm{F}}a < 0$ [144, 145]. This model is called the *Hubbard Hamiltonian* [146] and is also a proposed model for high-temperature superconductivity [147].

A solution to this quandry is to introduce a *two-channel model* incorporating both fermions and bosons [148–151]. The fermions then represent the unbound atoms scattering at threshold (open channel), while the bosons represent the weakly bound molecule brought into resonance (closed channel). Such two-channel models are mentioned briefly in the condensed matter context as long ago as 1985 [152] and explored seriously in partial form starting in 1995 [153]; they continue to be explored as the *Cooperon model*, in current research on high-temperature superconductivity [154]. However, only in ultracold quantum gases have all hopping, interaction and interconversion terms been included to make a *Fermi–Bose Hubbard Hamiltonian* in the lattice [155]. A series of papers attempted this approach with steady improvement over time, starting from simply attaching a band index to each channel [155, 156] up to performing a renormalization procedure to produce effective molecules in the lattice and minimize the number of bands that must be included [157–161]. The latter method has led to a Hamiltonian that is so much far from the original Hubbard Hamiltonian that it has been given a new name, the *Fermi resonance Hamiltonian*.

The Fermi resonance Hamiltonian is

$$\hat{H}_{\mathrm{eff}} = -t_f \sum_{\sigma\in\{\uparrow,\downarrow\}} \sum_{\langle i,j\rangle} \hat{a}^{\dagger}_{i\sigma}\hat{a}_{j\sigma} + E_0 \sum_{\sigma\in\{\uparrow,\downarrow\}} \sum_{i} \hat{n}^{(f)}_{i\sigma} - \sum_{\alpha\in\mathcal{M}} \sum_{i,j} t^{\alpha}_{i,j}\hat{d}^{\dagger}_{i,\alpha}\hat{d}_{j,\alpha}$$
$$+ \sum_{\alpha\in\mathcal{M}} \bar{\nu}_{\alpha} \sum_{i} \hat{n}^{(b)}_{i\alpha} + \sum_{\alpha\in\mathcal{M}} \sum_{ijk} g^{\alpha}_{i-j,i-k} \left[\hat{d}^{\dagger}_{i,\alpha}\hat{a}_{j,\uparrow}\hat{a}_{k,\downarrow} + \mathrm{h.c.}\right], \tag{24}$$

where $\hat{a}^{\dagger}_{i\sigma}$ creates a particle with spin σ in the lowest open channel band Wannier state centered at lattice site i; $\hat{d}^{\dagger}_{i,\alpha}$ creates a particle in the αth dressed molecule Wannier state centered at site i; $\hat{n}^{(f)}_{i\sigma}$ is the number operator for fermions in the lowest Bloch band; and $\hat{n}^{(b)}_{i\alpha}$ is the number operator for the αth dressed molecule state. The set of dressed molecules $\mathcal{M}$ which are included dynamically can be determined on energetic and symmetry grounds from the two-particle solution. In order, the terms in equation (24) represent tunneling of atoms in the lowest Bloch band between neighboring lattice sites i and j; the energy $E_0 = \sum_{\mathbf{q}} E_{1,\mathbf{q}}/N^3$ of a fermion in the lowest band with respect to the zero of energy; tunneling of the dressed molecular center of mass between two lattice sites i and j, not necessarily nearest neighbors; detunings of the dressed molecules from the lowest band two-particle scattering continuum; and resonant coupling between the lowest band fermions at sites j and k in different internal states and a dressed molecule at site i. The Fermi resonance Hamiltonian is a two-channel resonance model, between unpaired fermions in the lowest band, and dressed molecules nearby in energy [161]. Among other unusual predictions it makes, in contrast to usual Hubbard physics, is a significant diagonal hopping, pairing between atoms which do not lie along a principal axis of the lattice to form a dressed molecule and multiple molecular bound states induced by the lattice. Thus, the lattice problem is indeed quite different from both the continuum crossover problem and the well-known Hubbard physics.

In practice, the solution of the crossover problem in the lattice must occur in three steps. Firstly, for a given lattice strength and interacting strength the two-body problem must be solved exactly, including a renormalization procedure to get effective Wannier states for the molecules. Secondly, the coefficients in equation (24) are calculated. Thirdly, the Hamiltonian itself must be solved. Potential solution methods range from matrix product state methods in one dimensional (1D) to mean field, quantum Monte Carlo and dynamical mean field theory methods in higher dimensions. Holographic duality may offer new approaches to this Hamiltonian in the future. The crossover problem on the lattice remains very much an open problem and therefore we return to it in section 5.

2.7. Recent and new directions in crossover physics

In this subsection, we cover some of the new directions in unitary Fermi gases and related systems not discussed in previous sections, as explored in this focus issue.

2.7.1. New experimental probes. We require new experimental probes into unitary gases in order to better understand their behavior. Thus nearly all papers on ultracold fermions in this focus issue suggest an experimentally measurable effect. However, three papers realize new probes directly in experiments. First, Riedl *et al* [85] measure the quenching of the moment of inertia in the unitary regime. This quenching is a well-known signal of superfluidity, as starting with liquid ^{4}He it has been demonstrated that the moment of inertia drops below its classical value as one decreases a system through its critical temperature for the superfluid transition. Unlike other methods of measurement in ultracold Fermi gases such as frequency and damping rate of collective excitations in response to small perturbations, quenching of the moment of inertia zeros in specifically on superfluid effects. Thus one can distinguish between nearly ideal fluid dynamics in the normal phase, which is characterized by very low viscosity but allows rotational flow and irrotational superfluid hydrodynamics. A second new experimental probe is spin diffusion measurement, covered by three sets of contributors, one experimental [135] and two theoretical [136, 138], as we already touched on in section 2.5. The spin diffusion constant D_s and the momentum diffusion constant D_η exhibit a similar temperature dependence, and an analogous relation between heavy quark and momentum diffusion is expected in the QGP. Overall, diffusion constants offer another point of comparison for bounds predicted by holographic duality. Finally, a third new experimental probe is Tan's contact density, given in its original form in equation (10), and covered in this focus issue in both an experimental [83] and a theoretical contribution [71]. Tan's contact density connects microscopic scattering properties to macroscopic thermodynamic observables throughout the BCS–BEC crossover. Finally, as shown in figure 7, an improved temperature calibration is realized in this focus issue to better characterize the key ratio of viscosity to entropy [2].

2.7.2. Solitons to polarons. We now turn to new directions in theory. The simplest and oldest approach to describing the BCS–BEC crossover is BCS mean field theory, in the form of the

Boguliubov–de Gennes (BdG) equations describing quasiparticles in the BCS phase. Despite the fact that these equations are technically only valid for weak interactions $-k_{\mathrm{F}}a \ll 1$ and are based on a quasiparticle picture, one can attempt an extrapolation through the crossover by self-consistently solving the gap equation, BdG equations, and density and normalization of quasiparticles. Both Spuntarelli *et al* [162] and Baksmaty *et al* [163] describe this approach in detail in this focus issue; the method is numerically challenging as it requires the solution of simultaneous nonlinear partial differential equations. Spuntarelli predicts the features of dark solitons through the crossover. Solitons appear as robust local minima in the gap with a strong indicator in the density as well, even at unitarity. Solitons are also described by holographic duality by Keränen *et al* in this focus issue [164], providing a point of comparison to mean field theory. Further recent work on solitons through the BCS–BEC crossover can be found in [165–167]. Baksmaty *et al* use the BdG approach to shed light on experimental data in imbalanced Fermi gases, i.e. those with more spin-up than spin-down fermions. Imbalanced Fermi gases in elongated traps are shown to display strong violation of the local density approximation in the form of phase separation, quite a different perspective from the thermodynamic one.

Imbalanced Fermi gases, sometimes also called spin-polarized Fermi gases, open up a new territory for observing quantum phase transitions and have recently been used to explore a Fermi liquid description of the strongly interacting normal phase [168] and to observe a Fulde–Ferrel–Larkin–Ovchinnikov (FFLO) state in 1D Fermi gases [169]. The FFLO state is a pairing between two different Fermi surfaces, resulting in a finite center-of-mass momentum of the pairs. Spin-dependent scattering is suggested as a novel experimental probe in a theoretical contribution by Sheehy [170], to measure the pairing spin gap in locally imbalanced Fermi gases, where the imbalance is induced by coupling to a spin-dependent potential. Imbalanced gases are also used in spin diffusion measurements, as described in section 2.5. In the limit in which a Fermi gas is strongly imbalanced, a polaron model can be used, where the minority species is dressed by the Fermi sea of the majority species and the resulting quasiparticle is called a polaron; in the extreme case a single atom of the minority species is considered, leading to a single polaron. Exactly how such concepts apply at unitarity where quasiparticle concepts tend to break down is an open question. In this focus issue, Sadeghzadeh *et al* [171] predict the existence of a new metastable state consisting of a Fermi sea of polarons.

2.7.3. Disorder and quantum phase transitions. As discussed in section 2.6, lattice physics is only barely beginning to be treated in Fermi gases through the BCS–BEC crossover. Two new directions are treated for ultracold Fermions in optical lattices in this focus issue. Firstly, Han and Sá de Melo [172] treat the BCS–BEC crossover in the presence of disorder. They find the superfluid near unitarity to be much more robust against disorder than in the BCS and BEC regimes, and give some useful observations on how the different practical realizations of disorder in optical lattices can lead to different considerations for crossover physics. Secondly, lattices give one access to strongly correlated systems which have nothing to do with the BCS–BEC crossover. That is, near quantum phase transitions one generally obtains scale-invariant phenomena that can be potentially studied with holographic duality. Zhang *et al* [173] cover experiments on one such system in this focus issue, ultracold bosons near the Mott insulator-to-superfluid transition in the Hubbard model in two dimensions. Then the same kinds of issues as we have discussed throughout section 2 reappear, from a universal equation of state to transport. The first half of this paper forms a very readable overview for the general reader, while the second half gives a status report on experiments.

The massive ground state degeneracy induced by lattices is the principle underlying the achievement of strongly correlated states in bosons. One can obtain a similar degeneracy either by rotating a system to achieve such degeneracy in the lowest Landau level and get quantum Hall physics or by taking advantage of the spin degree of freedom. The latter strategy is pursued by Shlyapnikov and Tsvelik [174]. Chromium, a Bose-condensed atom with spin $S = 3$, provides a practical working example, and they find a commensurate–incommensurate quantum phase transition, among other possibilities in a rich phase diagram.

2.7.4. Trimers: from Efimov states to SU(3). Although we briefly mentioned three-body losses in section 2.1, in fact three-body losses are the dominant loss mechanism at unitarity, whether for fermions or for bosons. Three-body physics shows some real surprises in ultracold fermions: even when there is no two-body bound state at all, one finds a denumerably infinite and universal[19] set of three-body bound states in the strongly interacting regime, called the *Efimov* states. Efimov states can occur for various combinations of bosons and fermions of the same or different masses. The evidence for such a series of states is found in three-body losses. Braaten and Hammer wrote a review back in 2006 [175]. However, this area continues to develop rapidly, motivated strongly by ultracold experiments. In this focus issue, Wang and Esry cover this important area [176], extending our knowledge of Efimov physics beyond the low-energy threshold collisions typical of ultracold gases to the higher-energy regime relevant to colliding clouds. They treat both positive and negative scattering length and both broad and narrow Feshbach resonances [32]: in narrow Feshbach resonances Efimov effects are less universal in nature, requiring a second scattering potential parameter beyond the scattering length, namely the effective range. Losses in fermions also show a puzzling feature, a maximum at a magnetic field below the pole of the Feshbach resonance where the scattering length diverges. Zhang and Ho suggest an explanation for this anomaly in terms of an interplay between atom–atom, atom–dimer and dimer–dimer interactions. Based on these considerations and incorporating temperature and trapping parameters, they develop rate equations which match all experimental loss features, covering data in four different laboratories taken over seven years.

Finally, so far we have discussed solely $SU(2)$ physics. However, the hyperfine spin manifold of ultracold fermions allows for a variety of spin structures to be created. In particular, the preparation of degenerate three-state mixtures with three-fold ($SU(3)$) symmetric attractive interactions opens up new territory in the connection between cold atoms and nuclear matter. Such a three-state gas can be used to explore pairwise superfluidity and spontaneous symmetry breaking, color superconductivity and a superfluid-to-trimer gas transition that mimics the deconfinement-to-hadronization transition in quark matter [177]. In symmetric three-component gases, it is possible that more than one pairing field can simultaneously become non-zero [178]. Magnetism and domain formation can also be studied [179]. A key issue is to perform experiments at low density in order to reduce the three-body decay rates $\propto n^3$, which are not Pauli-suppressed in three-state systems. Three-body decay arises from recombination, which is enhanced and modified near a Feshbach resonance, where Efimov states exist with binding energies well beyond the threshold region [176]. By tuning away from resonance and using RF dressing to precisely symmetrize the pair-wise interactions, and confining the atoms in a lattice to increase the critical reduced temperature, the goal of creating cold-atom analogues of quark matter appears to be within reach [177].

---

[19] By universal we mean that the details of the scattering potential do not matter for Efimov states.

3. Quantum chromodynamics, the quark–gluon plasma and heavy-ion collisions

In this section we will give an introduction to the physics of the QGP. This is a vast subject and we will not attempt to provide a comprehensive review here. Recent reviews of various issues related to quark matter and the QGP can be found in [20, 180–184] and the results of the experimental program at the RHIC at BNL are summarized in [185]. Some results of the LHC program at the European Organization for Nuclear Research (CERN) can be found in the contributions [186, 187, 189] and in the recent review [190].

We will focus on observables and experimental results that are important in establishing the presence of strong correlations in the QGP and that provide direct connections to the physics of ultracold atomic gases and holographic dualities. In section 3.1 we introduce quantum chromodynamics (QCD), the theory of the strong interaction, and discuss the QCD phase diagram. The differences between weakly coupled QCD plasmas and strongly coupled QCD fluids are discussed in section 3.2. The experimental study of high-temperature QCD matter is based on heavy-ion collisions. We provide an overview of the experimental program in section 3.3 and discuss specific observables (multiplicities, flow, energy loss and heavy quarks) in sections 3.4–3.7. Finally, a discussion of open issues can be found in section 5.2.

3.1. Quantum chromodynamics and the phase diagram

QCD describes the behavior of strongly interacting matter, such as protons, neutrons and nuclei as well as the hot and dense matter created in heavy-ion collisions, in terms of quarks and gluons and their interactions. The complicated phenomenology of the strong interaction is encoded in a deceptively simple Lagrangian. The Lagrangian is formulated in terms of quark fields $q^c_{\alpha\, f}$ and gluon fields A^a_μ. Here, $\alpha = 1, \ldots, 4$ is a Dirac spinor index (corresponding, in the Dirac representation, to quarks and anti-quarks with spin up and down), $c = 1, \ldots, N_c$ is a color index and $f = up, down, strange, charm, bottom, top$ is a flavor index. In our world the number of colors is $N_c = 3$, but as a theoretical laboratory it is useful to consider theories with $N_c \neq 3$. The dynamics of the theory is governed by the color degrees of freedom. The gluon field A^a_μ is a vector field (such as the photon) labeled by an adjoint color index $a = 1, \ldots, 8$. The octet of gluon fields can be used to construct a matrix valued field $A_\mu = A^a_\mu \frac{\lambda^a}{2}$, where λ^a is a set of traceless, Hermitian, 3×3 matrices. Repeated indices are assumed to be summed throughout our treatment. The QCD Lagrangian is

$$\mathcal{L} = -\frac{1}{4} G^a_{\mu\nu} G^a_{\mu\nu} + \sum_f^{N_f} \bar{q}_f (\mathrm{i}\gamma^\mu D_\mu - m_f) q_f , \tag{25}$$

where $G^a_{\mu\nu}$ is the QCD field strength tensor:

$$G^a_{\mu\nu} = \partial_\mu A^a_\nu - \partial_\nu A^a_\mu + g f^{abc} A^b_\mu A^c_\nu , \tag{26}$$

and $f^{abc} = 4\mathrm{i}\, \mathrm{Tr}([\lambda^a, \lambda^b]\lambda^c)$ is a set of numbers called the $SU(3)$ structure constants. The covariant derivative acting on the quark fields is

$$\mathrm{i}D_\mu q = \left(\mathrm{i}\partial_\mu + g A^a_\mu \frac{\lambda^a}{2} \right) q , \tag{27}$$

and m_f is the mass of the quarks. The different terms in equation (25) describe the interaction between quarks and gluons, as well as nonlinear three- and four-gluon interactions. It is important to note that, except for the number of flavors and their masses, the structure of the QCD Lagrangian is completely fixed by a local $SU(3)$ color symmetry.

For many purposes, we can consider the light flavors (up, down and strange) to be approximately massless, and the heavy flavors (charm, bottom and top) to be infinitely massive. In this limit the QCD Lagrangian contains a single dimensionless parameter, the coupling constant g. If quantum effects are taken into account the coupling becomes scale dependent. At leading order the running coupling constant is

$$g^2(q^2) = \frac{16\pi^2}{b_0 \log(q^2/\Lambda^2_{\mathrm{QCD}})}, \qquad b_0 = \frac{11}{3}N_c - \frac{2}{3}N_f, \tag{28}$$

where q is a characteristic momentum and N_f is the number of active flavors ($N_f = 3$ if we consider charm, bottom and top to be infinitely heavy). The result in equation (28) implies that, as a quantum theory, QCD is not characterized by a dimensionless coupling but by a dimensionful scale, the QCD scale parameter Λ_{QCD}. This effect is called dimensional transmutation [191]. We also observe that the coupling decreases with increasing momentum. This is the phenomenon of asymptotic freedom [192, 193]. The flip side of asymptotic freedom is anti-screening or confinement: the effective interaction between quarks increases with distance.

In massless QCD the scale parameter is an arbitrary parameter (a QCD ‘standard kilogram’), but once QCD is embedded in the standard model and quarks acquire masses by electroweak symmetry breaking the QCD scale is fixed by the choice of units in the standard model. The numerical value of Λ_{QCD} depends on the renormalization scheme used to derive equation (28). Physical masses, as well as the value of b_0, are independent of this choice. In the *modified minimal subtraction* ($\overline{MS}$), one finds that $\Lambda_{\mathrm{QCD}} \simeq 200\,\mathrm{MeV}$; see the QCD section in [194].

Asymptotic freedom and the symmetries of QCD determine the main phases of strongly interacting matter that appear in the QCD phase diagram shown in figure 11. In this figure we show the phases of QCD as a function of the temperature T and the baryon chemical potential μ. The chemical potential μ controls the baryon density ρ, defined as 1/3 times the number density of quarks minus the number density of anti-quarks.

At zero temperature and chemical potential the interaction between quarks is dominated by large distances and the effective coupling is large. As a consequence, quarks and gluons are permanently confined in color singlet hadrons, with masses of order Λ_{QCD}. For example, the proton has a mass of $m_{\mathrm{p}} = 935\,\mathrm{MeV}$.[20]

If we think of the proton as composed of three *constituent quarks* this implies that quarks have effective masses $m_Q \simeq m_{\mathrm{p}}/3 \simeq \Lambda_{\mathrm{QCD}}$. This should be compared to the bare up and down quark masses which are of the order of 10 MeV.

Strong interactions between quarks and anti-quarks lead to a vacuum condensate of $\bar{q}q$ pairs, $\langle \bar{q}q \rangle \simeq -\Lambda^3_{\mathrm{QCD}}$ [195–197]. This vacuum condensate spontaneously breaks the approximate chiral $SU(3)_{\mathrm{L}} \times SU(3)_{\mathrm{R}}$ flavor symmetry of the QCD Lagrangian. Chiral symmetry breaking implies the existence of Goldstone bosons, particles with masses smaller than Λ_{QCD}. These particles are known as pions, kaons and etas[21].

[20] Note that we are quoting the mass in units of energy, setting the speed of light equal to unity.

[21] The $SU(3)_{\mathrm{L}} \times SU(3)_{\mathrm{R}}$ is explicitly broken by quark masses and the mass of the charged pion is $m_\pi = 139\,\mathrm{MeV}$, which is not much smaller than Λ_{QCD}. The lightest non-Goldstone particle is the rho meson, which has a mass $m_\rho = 770\,\mathrm{MeV}$.

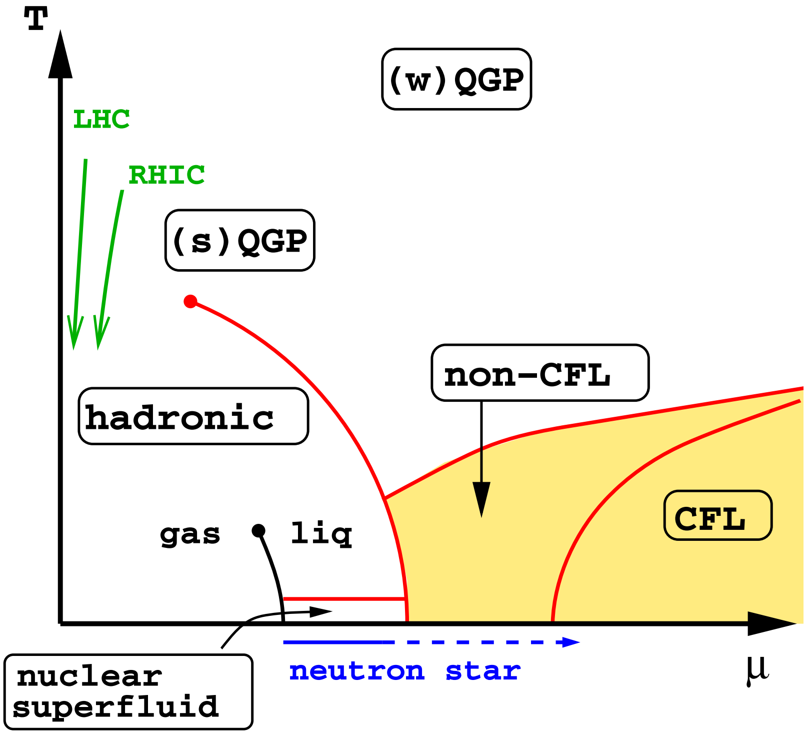

Figure 11. Schematic phase diagram of QCD as a function of temperature T and baryon chemical potential μ. QGP refers to the QGP, and sQGP/wQGP denote the regions of the phase diagram where the plasma is strongly/weakly coupled. The CFL (color–flavor locked) phase is the color superconducting phase that occurs at asymptotically large chemical potential. LHC/RHIC denote the regions of the phase diagram that are being explored by the experimental heavy-ion programs at the LHC and RHIC. The red and black points denote the critical endpoints of the chiral and nuclear liquid–gas phase transitions, respectively.

The quark–anti-quark condensate $\langle\bar{q}q\rangle$ is analogous to the di-fermion condensate $\langle qq\rangle$ that characterizes BCS superconductors [34]. BCS pairing involves particles with opposite momenta near the Fermi surface. The chiral condensate is an order parameter for the pairing of fermions and anti-fermions near the surface of the *Dirac sea*. The most important difference between BCS pairing and chiral condensation is that because of the finite density of states near the Fermi surface, BCS pairing can take place at weak coupling. Chiral condensation, on the other hand, only happens at strong coupling[22]. We will see below that in QCD at large baryon density there is a transition between $\bar{q}q$ and qq pairing. This transition is driven by the competition between the stronger coupling in the $\bar{q}q$ channel and the growing density of states for qq pairing.

At very high temperature quarks and gluons have thermal momenta $p \sim T \gg \Lambda_{\mathrm{QCD}}$ (note that we are writing momenta in units of energy, again setting $c = 1$.). Asymptotic freedom implies that these particles are weakly interacting and that they form a plasma of mobile color charges, the QGP [199, 200]. We note that the argument that the QGP at asymptotically high temperature is weakly coupled is somewhat more subtle than it might appear at first sight. If two particles in the plasma interact via large-angle scattering, then the momentum transfer is large, and the effective coupling is weak because of asymptotic freedom. However, the color Coulomb interaction is dominated by small-angle scattering, and it is not immediately clear why the effective interaction that governs small-angle scattering is weak. The important point is that in a high-temperature plasma there is a large thermal population ($n \sim T^3$) of mobile charges that screen the interaction at distances beyond the Debye length $r_D \sim 1/(gT)$. We also note that even in the limit $T \gg \Lambda_{\rm QCD}$ the QGP contains a non-perturbative sector of static magnetic color fields [201]. This sector is strongly coupled, but it does not contribute to thermodynamic or transport properties of the plasma in the limit $T \to \infty$.

---

[22] In ultracold atomic gases in the BEC regime, pair condensation involves pre-formed pairs. Whether or not pre-formed $\bar{q}q$ pairs exist in the QGP above T_c is still being investigated; see, for example, [198]. We should note that there is no completely rigorous criterion for the existence of a bound state embedded in a plasma. In practice, researchers have looked for peaks in the spectral function associated with the correlator $\langle\bar{q}q(x)\bar{q}q(0)\rangle$.

The plasma phase exhibits neither color confinement nor chiral symmetry breaking. This means that the high-temperature QGP phase must be separated from the low-temperature hadronic phase by a phase transition. The nature of this transition is very sensitive to the values of the quark masses. In QCD with massless u,d and infinitely massive s,c,b,t quarks the transition is second order [202]. In the case of massless (or sufficiently light) u,d,s quarks the transition is first order. Lattice simulations show that for realistic quark masses, $m_u \simeq m_d \simeq 10\,\text{MeV}$ and $m_s \simeq 120\,\text{MeV}$, the phase transition is a rapid crossover [203]. The transition temperature, defined in terms of the chiral susceptibility, is $T_c \simeq 151 \pm 3 \pm 3\,\text{MeV}$ [204, 205].

The transition is believed to strengthen as a function of chemical potential, so that there is a critical μ at which the crossover turns into a first-order phase transition [206]. This point is the critical endpoint of the chiral phase transition. Due to the fermion sign problem it is very difficult to locate the critical endpoint using simulations on the lattice. A number of exploratory calculations have been performed [207–210], but at this point it is not even clear whether the idea that the transition strengthens as the baryon chemical potential increases is correct [211]. The critical endpoint is interesting because it is the only point on the phase transition line at which the correlation length diverges (there is a similar endpoint on the nuclear liquid–gas transition line). This means that the critical point may manifest itself in heavy-ion collisions in terms of enhanced fluctuations [212]. The idea is that one can tune the baryon chemical potential by changing the collision energy, since lower beam energy allows for more stopping of the baryons in the initial state and therefore leads to higher chemical potential. As the collision energy is varied one looks for non-monotonic behavior of fluctuation and correlation observables. A typical observable is the variance (divided by the mean) of the net number of protons (protons minus anti-protons) in a finite sub-volume, $\langle (\Delta N_p)^2 \rangle / \langle N_p \rangle$ with $\Delta N_p = N_p - \langle N_p \rangle$. A beam energy scan is now under way at RHIC, and similar scans have been performed at lower energy as part of the CERN fixed target program [213, 214].

At low temperature the first phase one encounters when the chemical potential is increased from zero is nuclear matter, a strongly correlated superfluid composed of approximately non-relativistic neutrons and protons. It is interesting to note that nuclear matter in weak $\mathrm{n} \leftrightarrow \mathrm{p} + \mathrm{e} + \bar{\nu}_e$ equilibrium is neutron rich[23] and that the neutron–neutron scattering length[24]

---

[23] Isolated neutrons are unstable with regard to the decay into protons, electrons and anti-neutrinos. If we ignore the role of electrons, then dense matter is composed of equal numbers of protons and neutrons because this configuration optimizes the total Fermi energy. Taking electrons into account we observe that for a given density of protons and electrons the Fermi energy of the electrons is much bigger than that of protons, and the lowest energy state is neutron rich.

[24] 1 fm (Fermi)= 10^{-15} m is the typical unit of length in nuclear and particle physics. 1 fm is approximately equal to the radius of the proton and neutron.

$a_{\mathrm{nn}} \simeq 20\,\mathrm{fm}$ is much larger than the average inter-particle spacing, $r_{\mathrm{nn}} \simeq 1.9\,\mathrm{fm}$ at nuclear matter saturation density. As a consequence, dilute nuclear matter is closely related to the unitary atomic Fermi gases discussed in section 2.

At very large chemical potential, we can use arguments similar to those in the high-temperature limit to establish that quarks and gluons are weakly coupled. The main difference between cold quark matter and the hot QGP is that because of the large density of states near the quark Fermi surface even weak interactions can cause qualitative changes in the ground state of dense matter. In particular, attractive interactions between quark pairs lead to color superconductivity and the formation of a $\langle qq \rangle$ condensate. Since quarks carry color, flavor and spin labels, many superconducting phases are possible. The most symmetric of these, known as the color–flavor locked (CFL) phase, is predicted to exist at very high density [215, 216]. In the CFL phase the diquark order parameter is $\langle q^A_{\alpha f} q^B_{\beta g} \rangle \sim \epsilon_{\alpha\beta}\epsilon^{ABC}\epsilon_{fgC}$. This order parameter has a number of interesting properties. It breaks the $U(1)$ symmetry associated with baryon number, leading to superfluidity, and it breaks the chiral $SU(3)_{\mathrm{L}} \times SU(3)_{\mathrm{R}}$ symmetry. Except for Goldstone modes the spectrum is fully gapped; fermions acquire a BCS-pairing gap, and gauge fields are screened by the Meissner effect. This implies that the CFL phase, even though it arises from a superdense liquid of quarks, shares many properties of superfluid nuclear matter.

The CFL phase involves equal pair condensates $\langle ud \rangle = \langle us \rangle = \langle ds \rangle$ of all three light quark flavors. As the density is lowered, effects of the non-zero strange quark mass become more important, and less symmetric phases are likely to appear [180]. Calculations based on weak coupling suggest that the first non-CFL phase is a CFL-like phase with a Bose condensate of kaons, and the second phase involves a standing meson wave superimposed on the kaon condensate [217, 218]. Other possibilities which may appear in strong coupling include a phase with up–down pairing only (2SC), or separate spin-one condensates of up, down and strange quarks.

Some guidance an analyzing the competition between these phases may come from studying analogue models of the quark–hadron phase transition based on ultracold atomic gases. A number of authors have pointed out the possibility of a BCS–BEC crossover in dense quark matter [219]. The basic idea is that as the density is lowered, quark Cooper pairs become more strongly bound and form a diquark Bose condensate. At lower densities one may also encounter a mixture of condensed diquarks and unpaired quarks. The physics of this system can be studied using ultracold boson–fermion mixtures [24, 25, 220, 221]. Finally, it may be possible to study the transition from nucleons, bound states of three quarks, to color superconducting diquarks using three-species systems of ultracold fermions [177–179, 222].

Color superconductivity affects both the equation of state and the transport properties of dense quark matter. These effects may manifest themselves in the structure and evolution of neutron stars. Possible signatures of pairing appear in the mass–radius relation, the cooling curve and the relation between the spin-down rate and the magnetic field. There are two contributions in this focus issue that touch on these issues. Shovkovy and Wang [223] study the bulk viscosity of normal quark matter in the large-amplitude regime. Anglani *et al* [224] investigate the interaction of collective modes in the CFL phase. This calculation is a first step toward more accurate calculations of transport properties of the CFL phase [225, 226].

3.2. Weakly versus strongly coupled plasmas

At asymptotically high temperature the coupling is weak and properties of the QGP can be systematically computed. One of the most basic properties is the equation of state. In the following we will consider a related quantity, the entropy density. The perturbative expansion of the entropy density is [227, 228]

$$s = T^3 \left\{ c_0 + c_2 \, g^2 + c_3 \, g^3 + \cdots \right\}, \tag{29}$$

where g is the QCD coupling constant from equation (28) and we have chosen units such that $\hbar = k_{\mathrm{B}} = 1$. The expansion is performed with the coupling $g = g(\bar{\mu})$ evaluated at a fixed scale $\bar{\mu}$. Higher order terms contain logarithms of the form $\log(2\pi T/\bar{\mu})$, which combine with lower order terms to eliminate the dependence on the arbitrary scale $\bar{\mu}$, and lead to an expansion in terms of the running coupling $g(q)$ evaluated at a scale $q \simeq 2\pi T$. The first term in equation (29) corresponds to the Stefan–Boltzmann law, and c_0 is proportional to the number of degrees of freedom,

$$c_0 = \frac{2\pi^2}{45} \left(2 \left(N_c^2 - 1 \right) + 4 N_c N_f \frac{7}{8} \right), \tag{30}$$

where $2(N_c^2 - 1)$ is the number of degrees of freedom from the gluons and $4N_cN_f$ from the quarks. The naive perturbative expansion is an expansion in powers of g^2. Odd powers of g appear because of infrared divergences. As we approach the phase transition the coupling becomes large and higher order terms are no longer small. Indeed, the convergence properties of the expansion in equation (29) are extremely poor: the series shows no signs of converging unless the coupling is taken to be much smaller than one, $g \ll 1$, corresponding to completely unrealistic temperatures of the order of 1 TeV. The convergence can be improved significantly by using a self-consistent quasiparticle expansion [229]. This means that the perturbative expansion is formulated not in terms of free quarks and gluons, but in terms of quasiquarks and quasigluons which have effective masses and effective interactions.

Quasiparticle expansions fit lattice data quantitatively down to temperatures $T \sim 2T_{\mathrm{c}}$. In this regime $s/s_0 \simeq 0.85$, where s_0 is the entropy density of a non-interacting plasma. In the past this was frequently taken as evidence that quasiparticles are not strongly coupled at temperatures relevant to the early stages of heavy-ion collisions at RHIC or the LHC. A new perspective on this question is provided by holographic duality; see section 4. Holographic duality maps the strongly interacting quantum field theories onto one higher dimensional classical gravity, in particular anti-de Sitter space (AdS). In the original version of the mapping, the quantum field theory (QFT) is a QCD-like theory known as the $N = 4$ supersymmetric (SUSY) Yang–Mills theory [230]. This theory is conformal. This means that its coupling constant does not run, $g^2(q^2) = \text{const}$, and there is no analogue of Λ_{QCD}. It also does not have a phase transition, and so the theory is in the plasma phase for all values of the coupling. The particle content of the $N = 4$ SUSY Yang–Mills theory differs from that of QCD. The theory has no quarks, and in addition to gluons it contains supersymmetric fermionic partners of gluons called gluinos as well as additional colored scalar fields. Nevertheless, the plasma phase shares many features of the QGP, and by focusing on a ratio such as s/s_0 we can remove the difference in the number of degrees of freedom. Later versions of holographic duality relax conformal and supersymmetric requirements, but still do not map precisely onto QCD.

Using holographic duality it was found that in the limit of strong coupling and a large number of colors the entropy density of the SUSY Yang–Mills theory is $s/s_0 = 0.75$ [231].

This value is remarkably close to the result observed in lattice QCD calculations near the phase transition, casting doubt on an interpretation of the data in terms of weakly coupled quasiparticles. More generally, holographic duality demonstrates that thermodynamic properties of the plasma need not be very sensitive to the strength of the interaction.

Holographic duality also shows that, in contrast to equilibrium properties, transport properties of the plasma are very sensitive probes of the strength of the interaction. In perturbative QCD the shear viscosity of three-flavor QCD is [6, 232, 233]

$$\eta = \frac{kT^3}{g^4 \log(\mu^*/m_{\rm D})}, \tag{31}$$

where $k = 106.67$, $\mu^* = 2.96T$ and $m_{\rm D} \sim gT$ is the screening mass. Shear viscosity is related to the rate of momentum diffusion. A simple estimate of the shear viscosity is [234]

$$\eta \simeq \frac{1}{3} np\, l_{\rm mfp} \simeq \frac{1}{3}\frac{p}{\sigma_{\rm T}}, \tag{32}$$

where n is the density, p is the mean momentum, $l_{\rm mfp}$ is the mean free path and $\sigma_{\rm T}$ is the transport cross section. In a perturbative QCD plasma the typical momentum is $p \sim T$, and the transport cross section is $\sigma_{\rm T} \sim g^4 \log(g) T^{-2}$. The time scale for momentum diffusion is $\eta/(sT) \sim 1/(g^4 \log(g) T)$. We note that in the weak coupling limit this number is parametrically large. In a QGP at $T = 200\,\text{MeV}$, just above the phase transition, we have $T^{-1} \simeq 1\,\text{fm}/c$ and $\eta/s \simeq 9.2/g^4$. A typical value of the coupling is $g \simeq 2$ (corresponding to $\alpha_{\rm s} \simeq 0.3$), which implies that $\eta/(sT) \simeq 0.6\,\text{fm}/c$.

The strong coupling limit of η/s in the SUSY Yang–Mills plasma was studied by Policastro *et al* [235]. They find that

$$\frac{\eta}{s} = \frac{1}{4\pi}\left(1 + O\left(\lambda^{-3/2}\right)\right), \tag{33}$$

where $\lambda = g^2 N_c$ is the 't Hooft coupling. This result implies that at strong coupling momentum equilibration is almost an order of magnitude more rapid than it is in weak coupling. The ratio η/s is also close to a bound that was argued to arise from the uncertainty relation [234]. The idea is that, because of quantum mechanics, the product $p\, l_{\rm mfp}$ in equation (32) cannot become smaller than Planck's constant $\hbar$ (note that we are using units in which $\hbar = 1$). This implies that $\eta/s \lesssim \frac{1}{3}(\frac{n}{s}) \simeq 0.09$, where we have used the entropy per particle of a free gas [234]. The uncertainty argument has never been made precise, since the mean free path estimate is not valid at strong coupling. However, the holographic duality calculation was discovered to be quite general. It was shown that the strong coupling limit of η/s is universal in a large class of theories that have gravitational duals, and that the $O(\lambda^{-3/2})$ corrections are positive [236, 237]. These observations led KSS to make the conjecture that

$$\frac{\eta}{s} \geqslant \frac{1}{4\pi} \tag{34}$$

is a universal bound [11] that applies to all fluids. The status of this conjecture is discussed in more detail in section 4 and sketched in figure 2. There are some known counterexamples involving theories with gravitational duals described by higher derivative gravity. However, it seems clear that equation (34) applies to a large class of theories that include generalizations of QCD.

Experiments on collective flow in heavy-ion collisions at RHIC and the LHC (discussed in section 3.4) indicate that the viscosity of the QGP near the critical temperature is indeed close to the proposed bound and that equilibration must be very rapid. Together with the observation of a large energy loss of highly energetic probes of the plasma, these results led to the conclusion that the QGP produced at RHIC must be strongly coupled [185, 238].

In the case of the SUSY Yang–Mills plasma, one can show that strong coupling implies the absence of well-defined quasiparticles. This can be seen most clearly by studying the spectral function $\rho(\omega)$ of the stress tensor correlation function. Kubo's formula relates the value of the spectral function at zero energy to the shear viscosity. The spectral function at finite energy carries information about the physical excitations that contribute to momentum relaxation. In weak coupling the spectral function has a peak at zero energy. The width of the peak, $\Gamma \sim g^4 \log(g) T$, is related to the quasiparticle lifetime. The spectral function in the strong coupling limit was computed in [239, 240]. It was found that the spectral density is completely featureless: the intercept at zero energy smoothly connects to the continuum contribution $\rho(\omega) \sim \omega^4$.

In the case of the QGP, the evidence regarding the existence of quasiparticles is ambiguous. The stress tensor correlation function in lattice QCD was studied by Meyer [4]. The reconstructed spectral function is smooth, but the resolution was not sufficient to exclude the presence of quasiparticles. There are a number of results that have been interpreted as favoring the existence of quasiparticles. One is the fact that lattice calculations of fluctuations of conserved charges, such as baryon number and strangeness, are compatible with the behavior of a free gas, even near T_c [241]. An experimental observation that has been cited as evidence for the presence of quasiparticles is the approximate quark-number scaling of the elliptic flow parameter v_2 [242]. Another experimental observable that may shed some light on the quasiparticle structure is the heavy quark diffusion constant D, which we will discuss in section 3.7. The main observation is that kinetic theory and holographic duality make very different predictions about the relation between the momentum relaxation time $\eta/(sT)$ and the heavy quark relaxation time $m_Q D/T$ [181].

3.3. Nuclear collisions: initial conditions

The experimental study of heavy-ion collisions at relativistic energies, which were long seen as the best environments for the production of the QGP, began in the late 1980s with fixed target programs both at CERN near Geneva, and at BNL, outside of New York City; see [243–245] for historical overviews. These two programs were essential in building up a vital experimental community, formed of larger and larger collaborations as the experiments became more and more sophisticated. The programs at CERN and BNL demonstrated that heavy-ion collisions produce strongly interacting, approximately equilibrated matter. The next step was taken with the turn-on of the RHIC at Brookhaven in 2000, which brought heavy-ion physics into the collider era, with accessibility to processes calculatable with perturbative QCD. Colliders are accelerators with counter-circulating beams in which the entire collision energy is available for particle production. RHIC collides gold ions with a beam energy such that every nucleon in each of the beams has an energy of 100 GeV. Since the rest mass energy of a nucleon is about 1 GeV, this implies a relativistic γ factor of about 100. When two bunches of counter-rotating beams overlap in each of the RHIC experimental halls, at a rate of thousands of times per second, the 1 billion ions in each bunch typically induce less than one collision. In a fraction of these collisions, particularly the ones where the nuclei are head-on, thousands of particles are produced by the conversion of kinetic energy into mass energy.

The particles are recorded in detectors which were originally located in four experimental areas located around the ring. The early RHIC program had two large multipurpose detectors, STAR and PHENIX, with roughly 500 collaborators each, and two smaller detectors, BRAHMS and PHOBOS, with about 50 collaborators each [185]. STAR and PHENIX are both large spectrometers. STAR focuses primarily on charged hadrons using a *time projection chamber* of 4 m diameter that creates a 3D image of collision events at a relatively low rate. PHENIX is a combination of drift chambers, particle identification counters and calorimeters that make precise measurements of both charged particles and photons. Both of these detectors were inspired by the collider detectors of the previous generation and make measurements mainly at large angles with respect to the beam directions. To make sure physics was not missed in other regions of phase space, BRAHMS and PHOBOS were both designed to make more limited measurements near the beam axis, BRAHMS with a narrow-band spectrometer and PHOBOS with a large-coverage single-layer silicon detector, sensitive to roughly 75% of the total number of charged particles produced per event. A range of the data from these experiments will be shown in later sections, and in several of the contributions to this focus issue [186, 187, 189]. As of 2011, RHIC has completed its 11th experimental run, after colliding ions with a wide range of energies (from 7.7 to 200 GeV) and nuclear species (protons, deuterons, copper and gold).

As RHIC was taking data, the LHC at CERN was being built, along with its three large detector systems ALICE, ATLAS and CMS. The LHC provides heavy-ion collisions of primarily lead, slightly larger than gold, at center of mass energies up to 28 times that reached by RHIC. This translates into higher particle densities and temperatures and, even more significantly, into much higher rates for large momentum-transfer processes, such as jets, photons and heavy flavor, discussed below. The ALICE detector was designed around another large time projection chamber, similar to but larger than the one in STAR at RHIC, with a set of additional detectors to identify different hadron species as well as photons and electrons in limited angular regions. The layout of the ALICE detector is shown in figure 12. The two larger experiments, ATLAS and CMS, were general-purpose detectors deploying charged particle tracking and hermetic calorimetry over a large solid angle. However, the stringent design requirements to search for new high-mass particles led to detectors that were quite capable for the higher multiplicities expected in heavy-ion collisions at the LHC and which have excellent capabilities for high-energy processes, similar to that needed for Higgs and searches beyond the Standard Model. As of late 2011, the LHC has completed its second lead ion run, as well as a first feasibility study for future proton lead collisions.

In the following, we will concentrate on the results from Au + Au at the top RHIC energy of 100 GeV per nucleon. The transverse radius of an Au nucleus is approximately 6 fm, and the duration of a heavy-ion event is $\tau \sim 6$–$10\,\mathrm{fm}/c$. The estimate for the lifetime comes from hydrodynamic simulations which we will describe in section 3.5. The simplest observable in a heavy-ion experiment, typically published very soon after a new machine becomes operational, is the total multiplicity of the produced particles. In Au + Au collisions at 100 GeV per nucleon the total multiplicity is about 7000; see section 3.4. Somewhat more detailed information is provided by the spectra $\mathrm{d}N/\mathrm{d}^3p$ of produced particles (e.g. [246, 247]). The momenta can be

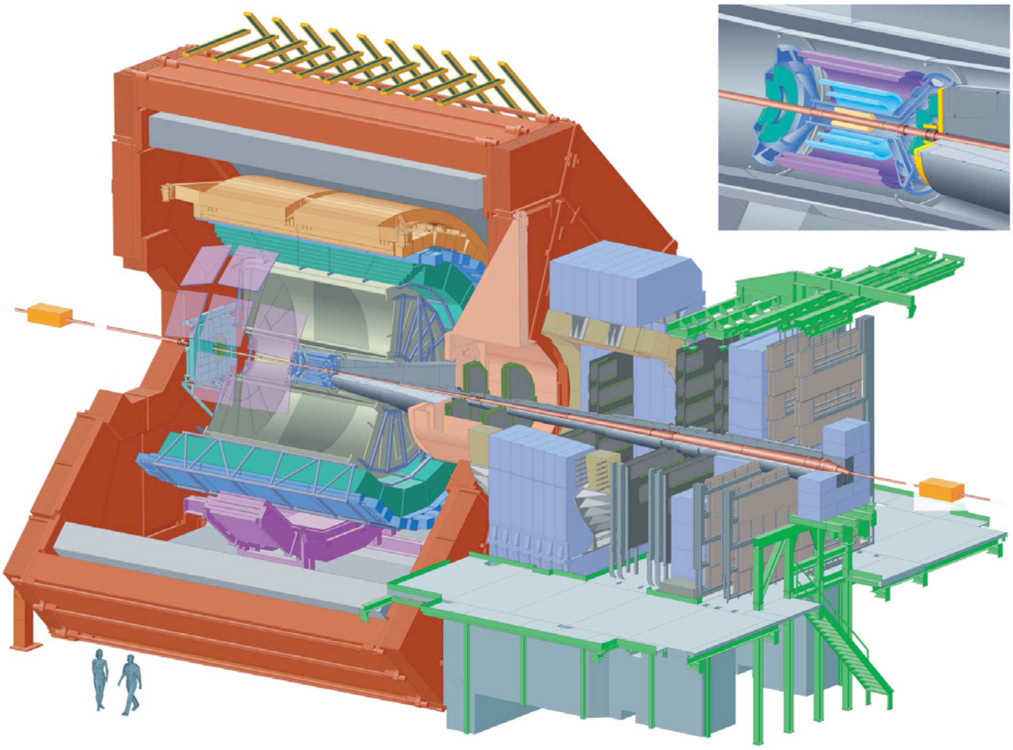

Figure 12. Schematic layout of the ALICE detector at the LHC. The beam lines are the thin yellow tubes entering from both sides. The outermost (red) layer is a magnet. The magnetic field helps to discriminate between particles of different charge-to-mass ratios. The central gray barrel is a time projection chamber and the outer layers inside the magnet contain calorimeters, transition radiation and time-of-flight detectors. The innermost layers shown in the blow-up contain silicon detectors for a precise determination of the initial production vertices.

decomposed into a transverse momentum $p_T^2 = p_x^2 + p_y^2$ and a longitudinal momentum p_z; see figure 13. In the relativistic regime the natural variable to describe the motion in the z direction is the rapidity,

$$y = \frac{1}{2}\log\left(\frac{E+p_z}{E-p_z}\right). \tag{35}$$

At RHIC the energy of the colliding nuclei is $100+100\,\mathrm{GeV}$ per nucleon, and the separation in rapidity is $\Delta y = 10.6$.

A simple picture of the initial state of the fireball created in the collision was suggested by Bjorken [249]. He proposed that the two highly Lorentz contracted nuclei pass through each other and create a longitudinally expanding fireball in which particles are produced. In the original model the number of produced particles is independent of rapidity, and the subsequent evolution is invariant under boosts along the z-axis. A simple model of the initial energy density in the transverse plane is the Glauber model [248, 250]. The Glauber model is based on the observation that high-energy scattering can be described in the eikonal (geometric optics)

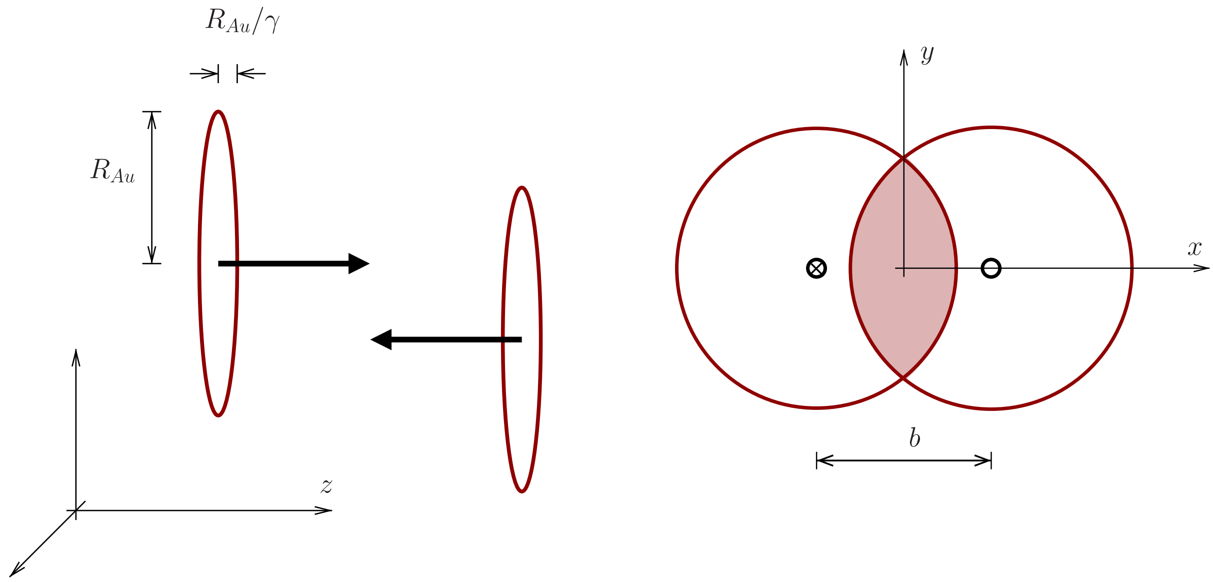

Figure 13. Geometry of a high-energy heavy-ion collision. Left: collision of two Lorentz contracted gold nuclei. The beam direction is the z-axis. Right: the same collision in the transverse plane. The impact parameter is along the x-axis, and the remaining transverse direction is the y-axis.

approximation. The initial entropy density in the transverse plane is

$$s(\mathbf{x}_\perp, b) \propto T_A(\mathbf{x}_\perp + \mathbf{b}/2)\Big[1 - \exp(-\sigma_{\rm NN}\, T_A(\mathbf{x}_\perp - \mathbf{b}/2))\Big] \\ + T_A(\mathbf{x}_\perp - \mathbf{b}/2)\Big[1 - \exp(-\sigma_{\rm NN}\, T_A(\mathbf{x}_\perp + \mathbf{b}/2))\Big], \qquad (36)$$

where $\mathbf{b}$ is the impact parameter,

$$T_A(\mathbf{x}_\perp) = \int \mathrm{d}z\, \rho_A(\mathbf{x}) \qquad (37)$$

is the thickness function, $\mathbf{x} \equiv (x, y, z)$ and $\mathbf{x}_\perp \equiv (x, y)$, and $\sigma_{\rm NN}(\sqrt{s})$ is the nucleon–nucleon cross section. We also define the nuclear density as $\rho_A(\mathbf{x})$. The idea behind the Glauber model is that the initial entropy density is proportional to the number of nucleons per unit area which experience an inelastic collision, which we call the *number of participants*, $N_{\rm part}$. Other variants exist. For instance, one can distribute the energy density according to the number of binary nucleon–nucleon collisions, $N_{\rm coll}$; see [251] for a comparison. A more sophisticated theory of the initial energy density is provided by the *color glass condensate* [252, 253]. This model leads to somewhat steeper initial transverse energy density distributions.

It should also be noted that the simple Glauber model is typically not used by modern calculations, since it does not account for the large fluctuations which are currently attributed to the event-wise variations in the nucleon configuration coming from each nucleus and the collision process itself. Figure 14 illustrates this by means of a single event simulated by a *Glauber Monte Carlo* code, which counts participants and collisions by means of a simple prescription that collisions occur when two nucleons are within $d < \sqrt{\sigma_{\rm NN}/\pi}$, where $\sigma_{\rm NN}$ is the total nucleon–nucleon inelastic cross section [248].

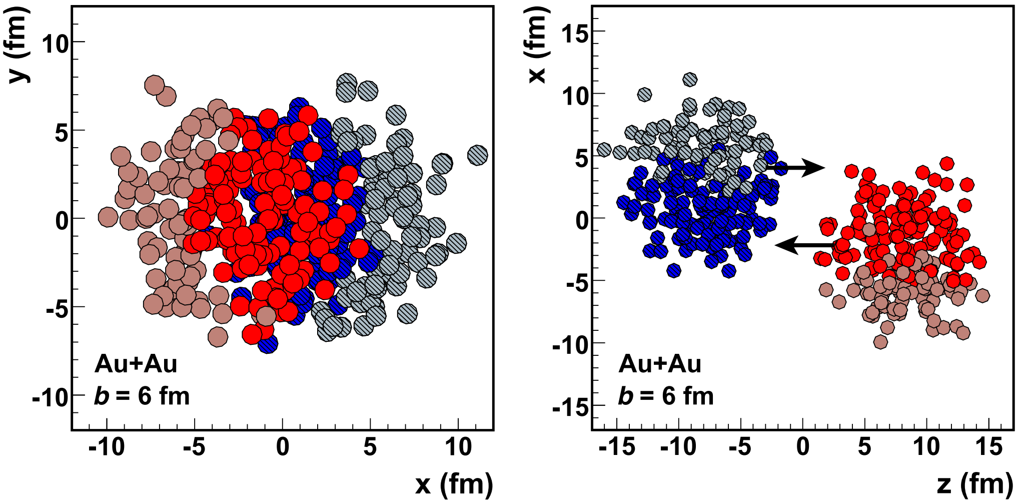

Figure 14. Glauber Monte Carlo calculation showing the collision of two gold nuclei at an impact parameter $b = 6\,\text{fm}$ head-on (left panel) and from a side view (right panel). The Glauber model is a geometric model for the high-energy scattering of composite quantum mechanical states. Spectator nucleons from the two nuclei are shown in pale red and blue and participants are shown as dark red and blue circles. Figure reproduced from [248] with permission.

3.4. Particle multiplicities

Particle multiplicities are important both as first measurements at new machines, but also as a means of testing the current theoretical approaches in their ability to predict the degrees of freedom available at higher energies.

From the theoretical perspective, multiplicities are argued to be a good proxy for the initial gluon density. This relies on one essential non-trivial argument that subsequent scatterings in the QGP, after the initial phase where hard collisions occur and energy is deposited in the interaction region, do not increase the overall entropy, and thus do not increase the overall multiplicity. This is certainly not exactly true, as both thermalization and viscous effects during the hydrodynamic evolution will produce some amount of entropy, but there is evidence that the total entropy is dominated by initial particle production [255]. Assuming that entropy is conserved one can derive the *Bjorken estimate* for the initial entropy density,

$$s_0 = \frac{3.6}{\pi R^2 \tau_0} \left(\frac{\mathrm{d}N}{\mathrm{d}y} \right), \tag{38}$$

where τ_0 is the thermalization time, R is the nuclear radius and $\mathrm{d}N/\mathrm{d}y$ is the total number of particles per unit rapidity. This estimate is based on a model of the hydrodynamic evolution which we will discuss in more detail in section 3.5. The left panel in figure 15 shows experimental results for $\mathrm{d}N/\mathrm{d}y$ of charged particles at RHIC. Assuming that $N(\text{all}) = 1.5N$ (charged) (most particles are pions) we obtain $\mathrm{d}N/\mathrm{d}y|_{y=0} \simeq 975$ for Au + Au collisions as 100 GeV per nucleon. A conservative estimate for the equilibration time is $\tau_0 = 1\,\text{fm}/c$; see section 5.2. This corresponds to an initial entropy density $s_0 \simeq 30\,\text{fm}^{-3}$. For a weakly interacting QGP this implies an initial temperature $T_0 \simeq 230\,\text{MeV}$, and an initial energy density

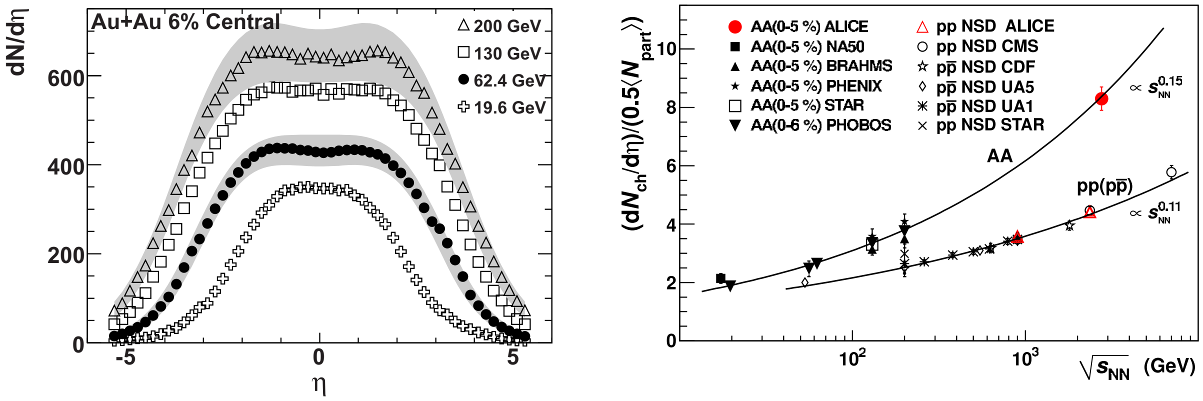

Figure 15. Left: the number of charged particles per unit pseudorapidity η as a function of η for Au + Au collisions at several different energies. Pseudorapidity is defined like rapidity, but using the approximation $E \simeq |\vec{p}|$, and is more easily measured experimentally than the rapidity y. For light particles, such as pions, $\eta \simeq y$; note, however, that corrections to this relation are biggest at $\eta = 0$. Data from the PHOBOS collaboration at RHIC [185]. Right: multiplicity $\mathrm{d}N/\mathrm{d}\eta$ per participant pair at midrapidity for both nucleon–nucleon and nucleus–nucleus collisions as a function of the collision energy. Data and compilation of earlier results from the ALICE collaboration at the LHC; see [254] for original references.

$\epsilon_0 = 5\,\mathrm{GeV\,fm^{-3}}$. This number is significantly larger than the critical energy density for forming a QGP, $\epsilon_{\mathrm{crit}} \simeq 1\,\mathrm{GeV\,fm^{-3}}$ [256].

Experiments typically focus on studying how multiplicities scale with beam energy and collision geometry. Empirically, the multiplicity per unit rapidity scales as a small power of the beam energy; see the right panel of figure 15. In nucleus–nucleus collisions, data ranging from the fixed target SPS program all the way to the LHC are described by $\mathrm{d}N/\mathrm{d}y|_{y=0} \sim s^{\alpha}$, where s is the square of the center of mass energy in the nucleon–nucleon system and $\alpha \simeq 0.15$. In nucleon–nucleon collisions the exponent is somewhat smaller, $\alpha \simeq 0.11$.

As a first approximation, it is natural to expect that heavy-ion collisions can be built by a linear superposition of proton–proton or, more accurately, nucleon–nucleon collisions. This is the basis of the *wounded nucleon model* of the 1970s which was invoked to explain how multiplicities in proton–nucleus collisions tended to scale as $N_{\mathrm{ch}}^{\mathrm{p+A}} = N_{\mathrm{ch}}^{\mathrm{p+p}} \times N_{\mathrm{part}}/2$ [257, 258]. In this expression, N_{part} is the number of participating nucleons estimated either using a Glauber model or by counting the number of slow proton tracks, e.g. in emulsion. In this picture, soft particle production was expected to result from the excitation and subsequent independent decay of individual nucleons. Hard processes, in which single *partons* (quarks and gluons) from one nucleon scattered off partons of another nucleon, were also discovered in the 1970s. As these were not simply decay products of a soft excitation but were short-range elastic scatterings, they were expected to scale linearly with the number of binary collisions, also predictable using a Glauber model.

In order to compare to experimental results, calculations typically scale the produced multiplicity by the number of participants, since the multiplicity is dominated by soft particles, with transverse momenta under 2 GeV. The experiments use the observed distribution of multiplicities at a fixed beam energy to estimate the *centrality* of each collision, usually expressed as a percentage of the total inelastic cross section. The 10% highest multiplicity events are the 0–10% centrality bin, the next highest are the 10–20% and so on. These bins are then compared to the events with the 10% smallest impact parameter, or 10% highest multiplicity events, in either a simple Glauber model or one which includes some model of the experimental fluctuations [248].

The first comparisons of theory and experiment at RHIC using the multiplicity in the most central events were surprising. While many models predicted a combination of soft and semi-hard physics processes, it was the models that were predominantly composed of soft processes which came closest to the experimental data. This included particularly models based on parton saturation, which were preferred by the early data. The situation is a bit different at the LHC, as described briefly in the contribution by Steinberg [186]. At the higher energies, models which include a combination of hard and soft processes seem to do better at predicting the multiplicities in the most central collisions. However, the *centrality dependence*, i.e. the relationship between the number of measured charged particles and the number of participating nucleons (or effective nucleon–nucleon collisions), is essentially the same at RHIC and the LHC [259]. At present it is not clear how to reconcile this observation with the apparent increase in the fraction of bulk particle production stemming from hard processes.

3.5. Hydrodynamic flow

The particle density measured by the early experiments is not only valuable to compare with different production models. It is also an essential empirical input that can be used to determine the initial conditions for hydrodynamical models of the subsequent evolution. The hydrodynamic description is valid if the system is in local thermal equilibrium and if the thermodynamic variables are varying smoothly. If the system has a kinetic description in terms of quasiparticles this implies that the ratio of the mean free path over the characteristic length scale of the system, known as the *Knudsen number*, must be small. There is evidence that this condition is satisfied during the early stages of the evolution, but it must eventually break down during the late, dilute, stage of the evolution, since the mean free path increases with diluteness. The time when the mean free path becomes larger than the system size, more precisely, the expansion time multiplied by the mean velocity, is called *freezeout*.

There are several experimental observations that provide evidence for the assumption that heavy-ion collisions create a thermalized state. The first observation is that the overall abundances of produced particles, including very rare ones, is described by a simple thermal model that contains only two parameters, the freezeout temperature T and the baryon chemical potential μ [183]. The second observation is that for momenta less than about 2 GeV the spectra $\mathrm{d}N/\mathrm{d}^3p$ of produced particles follow a Boltzmann distribution characterized by the freezeout temperature and a collective radial expansion velocity [182]. This *radial flow* is clearly seen from the fact that the spectra of heavy particles, which receive a larger momentum boost from the collective velocity field, have a larger apparent temperature than the spectra of light particles.

The third piece of evidence in support of not only thermalization, but early thermalization, which means equilibration significantly before freezeout, is the observation of strong azimuthal anisotropies, typically called *elliptic flow*, in non-central heavy-ion collisions. Elliptic flow represents the collective response of the system to pressure gradients in the initial state. The analogous phenomenon in ultracold atomic gases is discussed in section 2.1. At finite impact parameter the initial state has the shape of an ellipse, with the short axis along the x-direction, and the long axis along the y-direction; see figure 13. This implies that pressure gradients along the x-axis are larger than along the y-axis. Hydrodynamic evolution converts the initial pressure gradients into velocity gradients in the final state. Elliptic flow is sensitive to early thermalization because the initial anisotropy that drives elliptic flow disappears with time. Elliptic flow is quantified in terms of the second Fourier coefficient v_2 of the particle distribution in the transverse plane,

$$p_0 \frac{\mathrm{d}N}{\mathrm{d}^3 p}\bigg|_{p_z=0} = v_0(p_\mathrm{T})(1 + 2v_1(p_\mathrm{T})\cos(\phi) + 2v_2(p_\mathrm{T})\cos(2\phi) + \cdots), \tag{39}$$

where ϕ is the angle between the momentum vector and the x-axis. Odd Fourier moments such as v_1 and v_3 arise mainly from fluctuations in the initial state [260, 261]. Techniques for measuring v_2 and systematic trends in the results are described in some detail in the contribution by Snellings [187]. Snellings emphasizes that measurements of v_2 constrain the equation of state and the transport properties of the QGP. A similar conclusion is reached in the contribution of Lisa *et al* [188] to this focus issue. These authors focus on direct measurements of the shape of the final state using HBT (Hanbury–Brown–Twiss) interferometry.

In this section we will explain how elliptic flow can be used to constrain the ratio η/s introduced in section 3.2. In a relativistic fluid the equations of energy and momentum conservation can be written as a single equation

$$\partial_\mu T^{\mu\nu} = 0, \tag{40}$$

where $T^{\mu\nu}$ is the energy momentum tensor. In ideal fluid dynamics the form of $T_{\mu\nu}$ is completely fixed by the Lorentz invariance,

$$T^{\mu\nu} = (\epsilon + P)u^\mu u^\nu + P\eta^{\mu\nu}, \tag{41}$$

where u^μ is the fluid velocity ($u^2 = -1$) and $\eta^{\mu\nu} = \mathrm{diag}(-1, 1, 1, 1)$ is the metric tensor. The hydrodynamic equations have to be supplemented by an equation of state $P = P(\epsilon)$. The four equations given in equation (40) can be split into longitudinal and transverse components relative to the fluid velocity. The longitudinal equation is equivalent to entropy conservation

$$\partial_\mu (s u^\mu) = 0, \tag{42}$$

and the transverse equation is the relativistic Euler equation

$$D u_\mu = -\frac{1}{\epsilon + P} \nabla^\perp_\mu P, \tag{43}$$

where $D = u^\mu \partial_\mu$ and $\nabla^\perp_\mu = \partial_\mu - u_\mu D$. We observe that the inertia of a relativistic fluid is governed by $\epsilon + P$. In the non-relativistic limit we can have $u^\mu \simeq (1, \vec{v})$ and $\epsilon + P \simeq \rho$ (energy density is dominated by rest mass energy). We also find that $D \simeq \partial_t + v \cdot \partial$ and $\vec{\nabla}^\perp \simeq \vec{\partial}$. These approximations led to the usual Euler equation. We also note that $\nabla^\perp_\mu P = c_\mathrm{s}^2 \nabla^\perp_\mu \epsilon$, where c_s is the speed of sound. This implies that for a given initial energy density gradient the resulting acceleration is determined by the speed of sound. Ideal fluid dynamics corresponds to the limit that variations in the hydrodynamic variables occur on scales much larger than the mean free path. Dissipation arises from the leading gradient terms in the energy momentum tensor. In the rest frame of the fluid these corrections have the same form as in non-relativistic fluids. We have

$$\delta T^{ij} = -\eta \left(\partial^i u^j + \partial^j u^i - \frac{2}{3}\delta^{ij} \partial \cdot u \right) - \zeta \delta^{ij} \partial \cdot u, \tag{44}$$

where η and ζ are the shear and bulk viscosity.

The application of hydrodynamics to relativistic heavy-ion collisions goes back to the work of Landau [262] and Bjorken [249]. Bjorken discussed a simple scaling solution of the equations of fluid dynamics that corresponds to the space–time picture discussed in section 3.3. This solution provides a natural starting point for more detailed studies in the ultra-relativistic domain. In the Bjorken solution the initial entropy density is independent of rapidity, and the subsequent evolution is invariant under boosts along the z-axis. The evolution in proper time is the same for all comoving observers. The flow velocity is

$$u_\mu = \gamma(1, 0, 0, v_z) = (t/\tau, 0, 0, z/\tau), \tag{45}$$

where $\gamma = \sqrt{1 - v_z^2}$ is the boost factor and $\tau = \sqrt{t^2 - z^2}$ is the proper time. The velocity field in equation (45) solves the relativistic Euler equation (43). In particular, there is no longitudinal acceleration. The remaining hydrodynamic variables are determined by entropy conservation. Equation (42) gives

$$\frac{\mathrm{d}}{\mathrm{d}\tau}\left[\tau s(\tau)\right] = 0 \tag{46}$$

and $s(\tau) = s_0 \tau_0/\tau$. For an ideal relativistic gas $s \sim T^3$ and $T \sim 1/\tau^{1/3}$. We saw in the previous section that the initial entropy density is constrained by the final state multiplicity. Typical parameters at RHIC are $\tau_0 \simeq 0.6$–$1.6\,\mathrm{fm}$ and $T_0 \simeq 300$–425 MeV. We note that the initial temperature is significantly larger than the critical temperature for the QCD phase transition. The temperature drops as a function of τ and eventually the system becomes too dilute for the hydrodynamic evolution to make sense. At this point, the hydrodynamic description is matched to kinetic theory, generally using the formalism described by Cooper and Frye in the early 1970s [263], and the spectra of produced particles are computed.

In order to quantitatively describe the observed particle distributions several improvements of the simple Bjorken model are necessary. Firstly, one has to include the transverse expansion of the system [265]. Secondly, one has to include deviations from boost invariance in the longitudinal directions. Rapidity distributions at RHIC most likely result from physics somewhere in between the Bjorken scenario, which assumes boost invariance, and the Landau picture, which assumes complete stopping of the initial nuclei [266]. One also has to include realistic equations of state and take into account the geometry of the initial state. Results for $v_2(p_\mathrm{T})$ obtained from a calculation in ideal hydrodynamics are shown in figure 16. This calculation focuses on the central rapidity regime and maintains the assumption of boost invariance. The figure shows the result for a given centrality class, corresponding to a specific range of impact parameters. Hydrodynamic calculations show that the elliptic flow response v_2 is approximately linear in the spatial anisotropy $\langle y^2 - x^2 \rangle / \langle x^2 + y^2 \rangle$ of the initial state. We observe that ideal hydrodynamics provides an excellent fit to the RHIC data for $p_\mathrm{T} \lesssim 1.0$–$1.5\,\mathrm{GeV}$, depending on the particle species. This includes the observed hierarchy in the p_T dependence of v_2 for different species at low p_T. The mass splitting of $v_2(p_\mathrm{T})$ reflects an approximate transverse energy $E_\mathrm{T} = (p_\mathrm{T}^2 + m^2)$ scaling of the particle spectra in hydrodynamics.

Having established a baseline description of the spectra using ideal hydrodynamics, we can now discuss the role of dissipative effects. We begin with the effect of shear and bulk viscosity

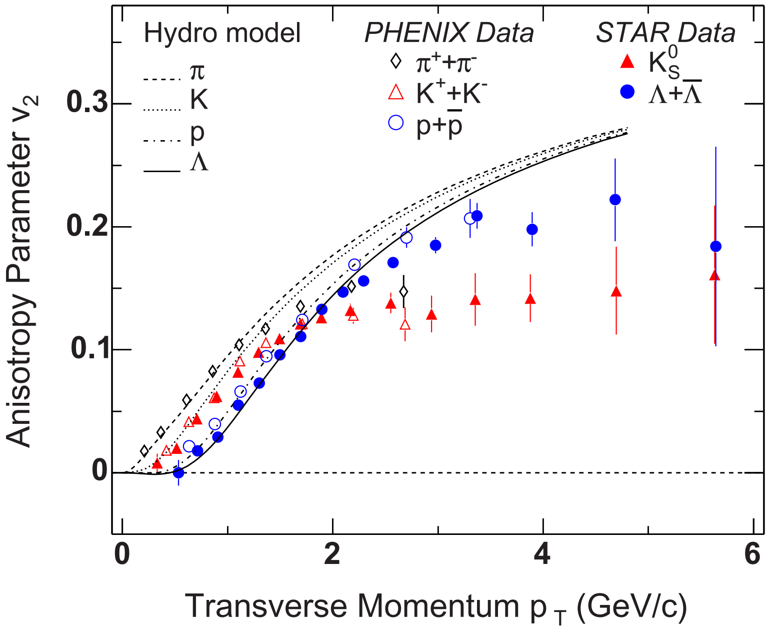

Figure 16. Differential elliptic flow $v_2(p_\mathrm{T})$ of identified hadrons from minimum bias Au + Au collisions at 200 GeV per nucleon compared to predictions from ideal (non-dissipative) hydrodynamics. Figure reproduced from [264], with permission.

on the Bjorken solution. The scaling flow given in equation (45) is a solution of the relativistic Navier–Stokes equation. If the transverse expansion of the system is neglected, viscosity does not affect the flow profile but it does generate entropy. We find that

$$\frac{1}{s}\frac{\mathrm{d}s}{\mathrm{d}\tau} = -\frac{1}{\tau}\left(1 - \frac{\frac{4}{3}\eta + \zeta}{sT\tau}\right). \tag{47}$$

The applicability of the Navier–Stokes equation requires that the viscous correction be small [234],

$$\frac{\eta}{s} + \frac{3}{4}\frac{\zeta}{s} \ll \frac{3}{4}(T\tau). \tag{48}$$

For the Bjorken solution, $T\tau \sim \tau^{2/3}$ grows with time, and this condition is most restrictive during the early stages of the evolution. Using $\tau_0 = 1$ fm and $T_0 = 300$ MeV gives $\eta/s < 0.6$. This result implies that hydrodynamics cannot be used in relativistic heavy-ions collisions unless the QGP is strongly coupled and the shear viscosity is small.

It is instructive to study in more detail the viscous contribution to the stress tensor. Neglecting bulk viscosity the stresses in the central rapidity slice are given by

$$T_{zz} = P - \frac{4}{3}\frac{\eta}{\tau}, \qquad T_{xx} = T_{yy} = P + \frac{2}{3}\frac{\eta}{\tau}. \tag{49}$$

This means that shear viscosity decreases the longitudinal pressure and increases the transverse one. In the Bjorken scenario there is no acceleration, but if pressure gradients are taken into account shear viscosity will tend to increase radial flow. At finite impact parameter shear viscosity reduces the pressure along the x-direction and increases the pressure in the y-direction. As a consequence there is less acceleration in the x-direction, and elliptic flow is suppressed.

Viscosity modifies the stress tensor, and via the matching to kinetic theory at freeze-out, this modification changes the distribution functions f_p of the produced particles. In [267] a simple quadratic ansatz for the leading correction δf to the distribution function was proposed,

$$\delta f_p = \frac{1}{2T^3}\frac{\eta}{s} f_0(1 \pm f_0) p_\alpha p_\beta \partial^{\langle\alpha} u^{\beta\rangle}, \tag{50}$$

where f_0 is the Bose–Einstein/Fermi–Dirac distribution and $\partial^{\langle\alpha} u^{\beta\rangle}$ is the symmetric traceless tensor that appears in equation (44). This ansatz is a very good approximation to the result of a more involved calculation using kinetic theory [6]. The modified distribution function leads to a modification of the single-particle spectrum. For a simple Bjorken expansion and at large p_T, we find that

$$\frac{\delta(\mathrm{d}N)}{\mathrm{d}N_0} = \frac{1}{3\tau_f T_f}\frac{\eta}{s}\left(\frac{p_\mathrm{T}}{T_f}\right)^2, \tag{51}$$

where $\mathrm{d}N_0$ is the number of particles produced in ideal hydrodynamics, $\delta(\mathrm{d}N)$ is the dissipative correction and τ_f is the freeze-out time. There is an analogous formula for the second Fourier moment of the spectrum, related to v_2 [267]. We observe that the dissipative correction to the spectrum is controlled by the same parameter $\eta/(s\tau T)$ that appeared in the entropy equation. We also note that the viscous term grows with p_T. These results are in agreement with experiment: deviations from ideal hydrodynamics grow with p_T, and they are larger in smaller systems (which freeze out earlier). More detailed analyses can be found in the contributions by Snellings [187] and Nagle *et al* [189]. A conservative bound for η/s at RHIC is $\eta/s < 0.4$, but the best fits tend to give values that are even smaller, $\eta/s \simeq 0.1$–0.2.

3.6. Jet quenching

Another way to directly probe the density of gluons present in the initial state is to scatter fast partons, formed in the initial collisions, which can be detected in the final state as hadronic jets. This idea was proposed as long ago as 1982 by Bjorken, placed on a firmer theoretical footing in the early 1990s, and finally discovered experimentally by RHIC experiments in the early part of the 2000s [268]. Figure 17 shows the *nuclear suppression factor*, defined as

$$R_\mathrm{AA} = \frac{1}{N_\mathrm{coll}}\frac{\mathrm{d}N/\mathrm{d}p_\mathrm{T}^\mathrm{AA}}{\mathrm{d}N/\mathrm{d}p_\mathrm{T}^\mathrm{pp}}, \tag{52}$$

i.e. the yield of a particle species in heavy-ion collisions, typically measured as a function of its transverse momentum, divided by a similar yield in proton–proton collisions, divided by the number of binary collisions calculated from a Glauber model. This ratio essentially tests the hypothesis that each individual binary collision has an equal probability to induce a hard process identical to that found in proton–proton collisions. The latter acts as a reference system in which hard processes can usually be calculated perturbatively [269]. As can be seen in figure 17, only the production of direct photons appears to be unmodified, and even then only below 12–13 GeV, above which isospin effects are expected [270]. Conversely, light hadrons such as π^0 and η are suppressed by about a factor of five above 5 GeV and perhaps slowly rise at high p_T. Electrons from charmed hadron decays are unsuppressed at low p_T but quickly fall to a level similar to that found for light hadrons, as discussed in more detail in the next section.

Considerable effort has been devoted to the development of a perturbative QCD-based formalism that describes the energy loss of fast partons moving through a hot, dense medium.

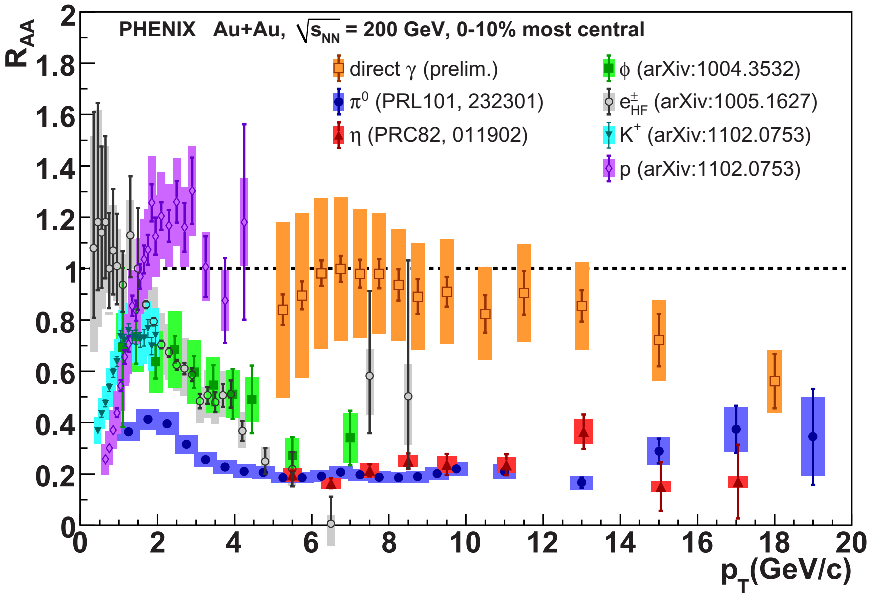

Figure 17. Nuclear suppression factor R_{AA} for a variety of hadron species, photons and electrons measured in PHENIX. Note that the dominant source of electrons is the decay of heavy quarks (bottom and charm). References to the data are indicated on the figure itself. The definition of R_{AA} is discussed in the text. The value $R_{\mathrm{AA}} = 1$ implies that for this particular probe a nucleus–nucleus collision behaves like a simple superposition of nucleon–nucleon collisions. Figure provided courtesy of Carla Vale, BNL.

The main transport parameter that appears in the theory is the transverse momentum diffusion constant $\hat{q}$ [274],

$$\hat{q} = \rho \int \mathrm{d}q_T^2\, q_T^2\, \frac{\mathrm{d}\sigma}{\mathrm{d}q_T^2}, \tag{53}$$

which determines the mean transverse momentum kick of the fast parton per unit length traveled. Here, $\mathrm{d}\sigma/\mathrm{d}q_T^2$ is the differential cross section for the scattering of the parton off the constituents of the dense medium, and ρ is the density of the medium. Recent work has focused on non-perturbative definitions of $\hat{q}$ that make no reference to the structure of the dense medium; see [184] for a review. The energy loss per unit length scales as $\mathrm{d}E/\mathrm{d}z \sim \hat{q}L$ for short path lengths ($L < L_c$) and as $\mathrm{d}E/\mathrm{d}z \sim \sqrt{\hat{q}E}$ for long path lengths. The characteristic length scale is $L_c \sim \sqrt{E/\hat{q}}$.

Extracting $\hat{q}$ from experimental data on jet quenching at RHIC has proven to be difficult. Based on an analysis of R_{AA} using a Monte-Carlo implementation of the theory of perturbative energy loss [274, 275], the PHENIX collaboration has reported that $\hat{q} = 13.2^{+2.1}_{-3.2}\,\mathrm{GeV}^2\,\mathrm{fm}^{-1}$ (the 2σ errors are $\hat{q} = 13.2^{+6.3}_{-5.2}\,\mathrm{GeV}^2\,\mathrm{fm}^{-1}$) [276]. This number is large compared to expectations for a perturbative QCD plasma, $\hat{q} \simeq 1$–$2\,\mathrm{GeV}^2\,\mathrm{fm}^{-1}$ [277]. However, there are significant uncertainties associated with different implementations of energy loss. Bass *et al* [278] consider three different methods and find that equally good descriptions for R_{AA} can be obtained for values of $\hat{q}$ ranging from $\hat{q} = 2.5$ to $10\,\mathrm{GeV}^2\,\mathrm{fm}^{-1}$.

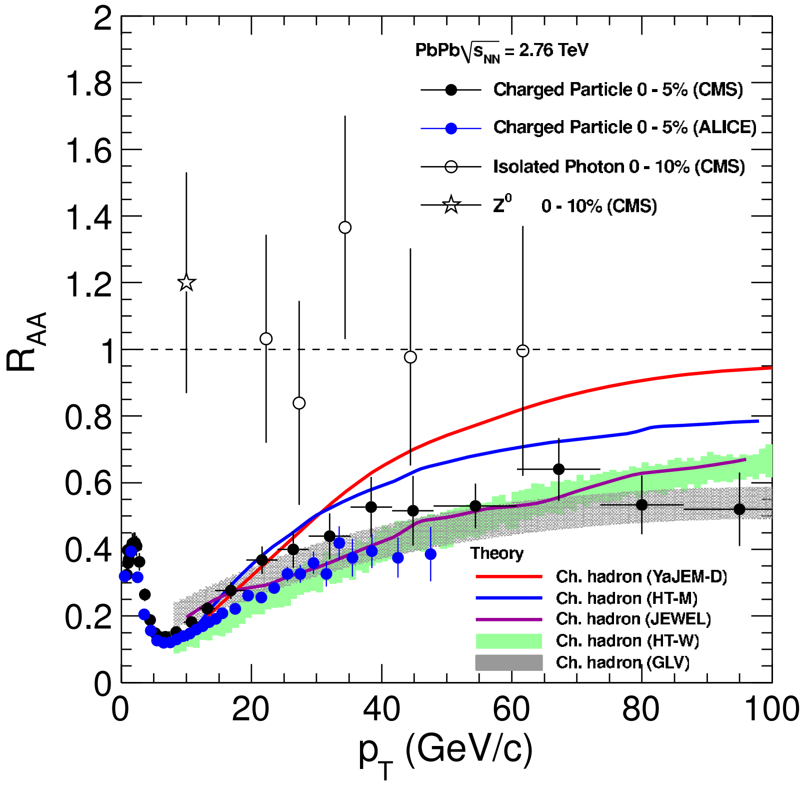

Figure 18. Nuclear suppression factor R_{AA} for a variety of hadron species, photons and Z^0 bosons measured in Pb–Pb collisions at the LHC, from [190]. The plot shows the results of the CMS and the ALICE collaboration compared to various theoretical models. See [271, 272] for the original data and [190] for references to the theory predictions.

The LHC offers significantly larger jet production cross sections, a much increased p_T range and the possibility for detailed studies of not only leading particles but also identified jets and jet shapes. First results from the LHC are shown in figures 18 and 19. Figure 18 shows the observable R_{AA} introduced above. The low-p_T behavior agrees with the results at RHIC, despite the much larger collision energy. At high p_T the suppression is smaller, $R_{AA} \sim 0.5$, but there is no hint of R_{AA} approaching 1. This is seen even more dramatically in figure 19, which shows the suppression factor R_{CP} for identified jets and not just high-p_T hadrons. The ratio R_{CP} is defined relative to the yield in very peripheral nucleus–nucleus collisions. We observe that R_{CP} is approximately equal to 0.5 for jets with transverse energies as high as 300 GeV. These results have important implications for the dependence of energy loss on the energy density of the medium, but the theoretical analysis of the LHC data is still very much in progress. A discussion of the first jet data from the LHC can be found in the contribution by Steinberg to this focus issue [186].

There are some important connections between $\hat{q}$ and other transport parameters. In a quasiparticle description the cross section in equation (53) is the same cross section that governs momentum diffusion of approximately thermal particles, i.e. shear viscosity. This implies that $\eta/s \propto T^3/\hat{q}$, where the constant of proportionality is roughly 1. In the context of kinetic theory, a small shear viscosity therefore implies strong jet quenching [279]. In detail, the situation

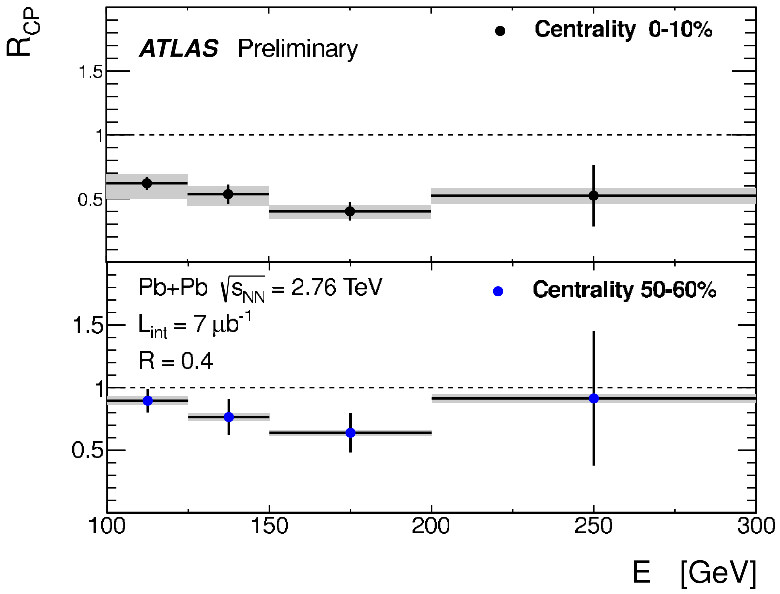

Figure 19. The jet central-to-peripheral ratio R_{CP} for different centralities in Pb–Pb collisions at the LHC, measured by the ATLAS collaboration [273]. The plots show the ratio of the jet yield in central 0–10% and semi-peripheral 50–60% events relative to the yield in very peripheral collisions as a function of the transverse energy E_{T} of the jet.

is more complicated. One can show that the viscous correction to $v_2(p_{\mathrm{T}})$ at moderate p_{T} is proportional to $1/\sqrt{\hat{q}}$ and not $1/\hat{q}$ [280]. Also, elliptic flow at large p_{T} is mainly a probe of the path length dependence of energy loss [281]. In the next section we will see that there are important connections between energy loss of light quarks and energy loss of heavy quarks.

3.7. Heavy quarks

An important diagnostic of the properties of the QGP is the drag force on a heavy quark initially produced in a hard collision during the pre-thermal-equilibrium stage of the collision. We will see below that the drag force can be related to the heavy quark diffusion constant. Combining experimental constraints on heavy quark diffusion and shear viscosity, which is related to momentum diffusion, provides additional information about transport properties of the nearly perfect fluid created in heavy-ion collisions.

The diffusion of heavy quarks in a QCD plasma was first studied by Svetitsky [288]. In the following we will follow the arguments presented by Teaney and Moore [286]. Consider a small density n_{Q} of heavy quarks inside a QGP. In a heavy-ion collision, heavy charm and bottom quarks are produced in hard collisions between quarks and gluons during the pre-equilibrium stage of the reaction. If there is enough time, and if the interaction is sufficiently strong, then the heavy quarks can reach thermal equilibrium[25]. In this case the time evolution of the density can be described by a diffusion equation

---

[25] The charm and bottom quark masses are $m_{\mathrm{c}} \simeq 1.3\,\mathrm{GeV}$ and $m_{\mathrm{b}} \simeq 4.2\,\mathrm{GeV}$, respectively. This implies that even if thermal equilibrium is reached, the density of heavy quarks is suppressed by large Boltzmann factors $\exp(m_Q/T)$. We will see below that the expected thermalization times are larger than those of light quarks and gluons by a factor m_Q/T.

$$\frac{\partial n_Q}{\partial t} = D\nabla^2 n_Q, \tag{54}$$

where D is the diffusion constant. We can relate D to the drag force on the quark by using a stochastic (Langevin) equation:

$$\frac{d\mathbf{p}}{dt} = -\eta_D \mathbf{p} + \xi(t), \qquad \langle \xi_i(t)\xi_j(t')\rangle = \kappa\delta_{ij}\delta(t-t'). \tag{55}$$

Here, $\mathbf{p}$ is the momentum of the particle, η_D is the drag coefficient and $\xi(t)$ is a stochastic force. In a kinetic theory picture, the stochastic force models collisions with quarks and gluons in the plasma. The coefficient κ is related to the mean square momentum change per unit time, $3\kappa = \langle(\Delta\mathbf{p})^2\rangle/(\Delta t)$. The Langevin equation can be integrated to determine the mean squared momentum. In the long time limit ($t \gg \eta_D^{-1}$) the particle thermalizes and we expect that $\langle \mathbf{p}^2\rangle = 3mT$. This requirement leads to the Einstein relation

$$\eta_D = \frac{\kappa}{2mT}. \tag{56}$$

The relation between η_D and the diffusion constant can be determined from the mean square displacement. At late times $\langle[\Delta\mathbf{x}(t)]^2\rangle = 6D|t|$ and

$$D = \frac{T}{m\eta_D} = \frac{2T^2}{\kappa}. \tag{57}$$

The diffusion constant for heavy quarks in a QGP can be determined by computing the mean square momentum transfer per unit time. At weak coupling and for approximately thermal heavy quarks, the diffusion constant is dominated by heavy quark scattering on light quarks and gluons, $qQ \to qQ$ and $gQ \to gQ$. These processes are similar to the processes that determine the shear viscosity, except that heavy quarks move slowly and the interaction only involves the color Coulomb interactions, whereas shear viscosity is sensitive to both electric and magnetic interactions. The leading order result in QCD with three light flavors is [286, 288]

$$D = \frac{6\pi}{g^4 T \log(2T/m_D)}. \tag{58}$$

Comparing this result with equation (31), we observe that heavy quark and momentum diffusion are related to each other. In the relevant range of coupling constants and keeping terms beyond the leading logarithm, one finds that $DT \simeq 6(\eta/s)$. We note that the heavy quark relaxation time η_D^{-1} contains an extra factor m_Q/T compared to the hydrodynamic relaxation time $\eta/(sT)$. For charm quarks at $T = 200\,\text{MeV}$ this factor is $m_c/T \simeq 7$. This implies that even if hydrodynamic behavior is reached very quickly, $\eta/(sT) \sim 0.2\,\text{fm}$, charm quarks have barely enough time to equilibrate, $m_c D/T \sim 8\,\text{fm}$.

This simple relation between DT and η/s is broken in the strong coupling limit of the SUSY Yang–Mills plasma. Using holographic duality, one finds that [289–291]

$$D = \frac{2}{\pi T}\frac{1}{\sqrt{\lambda}}. \tag{59}$$

This result implies that in the strong coupling limit there is no bound on the diffusion constant and that even very heavy quarks can possibly equilibrate. It is not clear how λ should be chosen

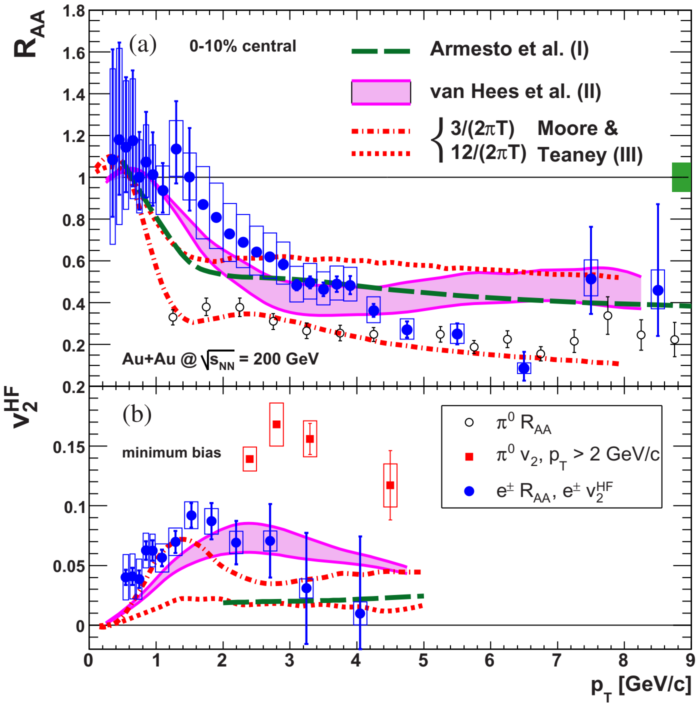

Figure 20. (a) R_{AA} of heavy-flavor electrons in 0–10% central collisions [282] compared with π^0 data [283] and model calculations (curves I [284], II [285] and III [286]). Model III is the Langevin calculation discussed in the text for two different values of the diffusion constant. Model II contains some extra processes, in particular final state coalescence. Model I is based on energy loss. The box to the far right at $R_{AA} = 1$ shows the uncertainty in T_{AA}. (b) v_2^{HF} of heavy-flavor electrons in minimum bias collisions compared with the π^0 data [287] and the same models.

in order to compare equation (59) to experiments in QCD. Gubser has advocated a value $\lambda \simeq 6\pi$ and concludes that $D \simeq 1/(2\pi T)$ [292].

Experimental information on the diffusion constant comes from the observation of single electrons from charm and bottom decays, as shown in figure 20 [282]. There are data on both the nuclear modification factor $R_{AA} = \mathrm{d}N_{\mathrm{Au+Au}}/(\langle T_{AA}\rangle \mathrm{d}\sigma_{\mathrm{p+p}})$, where $\mathrm{d}N_{\mathrm{Au+Au}}$ is the differential yield in Au + Au, $\langle T_{AA}\rangle$ is the nuclear overlap and $\mathrm{d}\sigma_{\mathrm{p+p}}$ is the differential cross section in p + p, as well as on the elliptic flow parameter v_2. The nuclear modification factor is sensitive to energy loss, which in turn is governed by the drag force. Drag also implies that the elliptic flow of light quarks and gluons induces a non-vanishing v_2 parameter for heavy quarks. Different models for the drag force are shown as dashed and dash-dotted lines in figure 20. We observe that these models can account qualitatively for the behavior of R_{AA} but fail to reproduce the observed flow. This indicates that charm quarks show some degree of equilibration. A Langevin simulation with $(2\pi T)D = (4–6)$ provides a qualitative description of both R_{AA} and v_2. It is interesting to note that this result, combined with the estimate $DT \simeq 6(\eta/s)$, gives $(4\pi)\eta/s \simeq 1.3–2.0$, which is close to the result obtained from the elliptic flow of light particles [282].

There are a number of uncertainties in this analysis that will be addressed in the future. Experiments have not yet been able to determine the flavor of the heavy quark, and there are significant theoretical uncertainties in predictions of the charm and bottom spectra. Future experiments will be able to identify the flavor of the heavy quark. Two contributions in this focus issue address transport properties of heavy quarks. Meyer [293] presents a lattice study of the Euclidean chromo-electric field correlation function. This correlation function is related by a Kubo formula to the diffusion constant. Rapp and Riek [294] perform a non-perturbative T-matrix analysis of the charm quark diffusion constant and the charmonium spectral function. These and other studies, together with ongoing efforts to measure the spectra of identified heavy quarks, will shed light on the relation between the different transport properties of the plasma such as diffusion, energy loss and shear viscosity and help to characterize the degree to which the initial state in a heavy-ion collision thermalizes.

4. Holographic duality

Holography is a duality relating QFT and gravity. Roughly speaking, holographic duality maps the quantum physics of strongly correlated many-body systems to the classical dynamics of black hole horizons in one higher dimension, replacing quasiparticles with geometry as the salient degrees of freedom. As such, holography is very much in the tradition of emergent critical phenomena: when the system is strongly coupled, new weakly coupled degrees of freedom dynamically emerge. The novelty is that the emergent fields live in a dynamical space–time with an extra spatial dimension. This extra dimension plays the role of an energy scale in the QFT, with motion along the extra dimension representing a change of scale, or renormalization group (RG) flow, in the QFT. Holography thus translates problems in quantum many-body physics, such as thermodynamics and transport physics, into equivalent problems in classical gravity.

The power of this reorganization is that various phenomena which are easy to see, or perhaps even universal, in one presentation may be very surprising in the dual presentation, but just as universal. For example, it is a classic result that the physics of black hole horizons in general relativity (GR) is largely independent of the details of the black hole [295–299], with small fluctuations of the horizon obeying the equations of viscous hydrodynamics [300]. The dissipation of waves in this fluid encodes the absorption of energy by the black hole, with the shear viscosity taking a universal value determined by basic properties of the Einstein–Hilbert action [300, 301]. Holographically, this is a remarkable fact: it tells us that any generic strongly interacting quantum many-body system at finite temperature and density and with a sufficiently large number of degrees of freedom per unit volume (necessary for the application of holographic duality) can be expected to behave, at low energies, like a nearly perfect liquid and not like a gas of long-lived quasiparticles as one might naively imagine.

More generally, holography gives us an entirely new way to define quantum field theories. Explicitly, holography provides a recipe [302–304] for using classical gravity in $(d+1)$ dimensions to compute quantum amplitudes that manifestly satisfy the consistency conditions of a d-dimensional QFT (locality, causality, etc)[26]. In general, we do not know how to identify the corresponding QFT in terms of more familiar tools, for example by specifying a Hamiltonian governing the interactions of a set of well-defined quasiparticles. Moreover, as we shall explore in some detail below, the holographic description is generally reliable precisely in situations

[26] This recipe is colloquially referred to as the *holographic dictionary*, a sketch of which is presented in table 1.

where most traditional techniques are not: when λ, the typical coupling in the QFT, is strong, $\lambda \gg 1$, and when N, the number of degrees of freedom per unit volume, is large, $N \gg 1$.[27] In such cases we can simply take the gravitational description as a constructive definition of the QFT. Such *holographic QFTs* thus define a special subset of the space of well-defined QFTs which does not depend on any quasiparticle picture or conventional perturbation theory. Importantly, some of the most theoretically and experimentally interesting real-world systems manifestly do not have any well-defined quasiparticles upon which to base a standard QFT. Holography thus provides an entirely new way to construct consistent models of these strongly correlated quantum many-body systems, replacing quasiparticles with geometry as the central organizing principle.

An important corollary is that it is typically more fruitful to use holographic QFTs as windows on a general class of phenomena, or to study general concepts in the space of QFTs, than to try to exactly reproduce or solve a specific QFT of previous interest. Indeed, building 'the holographic dual' of one's favorite QFT is generally quixotic, as the regimes of validity of holographic QFTs typically exclude the theories of previous interest. The classic example is $SU(N)$ gauge theory: QCD is given by $N = 3$, while the holographic dual is classical only when $N \gg 1$ and weakly coupled only when the 't Hooft coupling is also large, $\lambda \gg 1$. It is thus futile to try to reproduce QCD exactly. Where holography has proven useful, rather, is in studying general properties of $SU(N)$ gauge theories at high temperature and densities where many interesting phenomena arise which appear to be relatively insensitive to the precise value of N. For example, holographic models suggest that any strongly coupled QGP should behave like a liquid whose viscosity is very low and relatively insensitive to the precise value of the coupling, in sharp contrast to the large viscosity and strong coupling dependence predicted by weakly coupled QCD. The low viscosity of the QGP observed at RHIC (see section 3.5) thus suggests that the RHIC fireball is indeed an extremely strongly coupled quantum liquid.

To be sure, the simplest holographic models are in many ways very different from QCD—in the most well-understood example there are no quarks and no mass gap! What is remarkable from this point of view is that many of the striking features of the simplest holographic models persist even after the inclusion of quarks, a mass gap and other phenomenologically important ingredients. For example, in more realistic holographic models of QCD which contain fundamental quarks and display confinement, the physics at low energy is again governed by hydrodynamics with an exceptionally low and coupling-insensitive viscosity. Such holographic models may thus be treated of as computationally tractable toy models which exhibit a rich set of behaviors analogous to those observed in the laboratory, the study of which can reveal qualitative, and sometimes even quantitative, general properties of the larger class of strongly coupled QFTs.

It is important to stress that while the original discovery of holographic duality [302] required esoteric tools such as supersymmetry and conformal invariance[28], the duality itself does not depend on supersymmetry or conformality[29]. In fact, most current work on holographic duality involves non-supersymmetric, non-conformal QFTs. Indeed, we now have a long list of examples in which each of these constraints is weakened or removed (see, e.g., [305–308] and references therein). To date there is no example of a violation of a sharp holographic duality. From the current point of view, holographic duality is simply a true, if as yet unproven, fact about quantum field theories and quantum gravity.

---

[27] More generally, holographic duality relates QFT in d-dimensions to quantum gravity in $(d+1)$ dimensions, with the gravitational description becoming classical when the QFT is strongly coupled, as discussed in section 4.2.3. We will focus on regimes where the QFT is strongly coupled and the gravitational dual classical, but we emphasize that holography remains true even when quantum gravitational (string theoretic) effects become important.

[28] Indeed, holographic duality is often referred to as the *AdS-CFT correspondence* and also as *gauge–gravity duality*, among many other names.

While our understanding of the duality remains in various ways incomplete, holographic model building has already generated novel insights, including predictions for the anomalously low viscosity of cold atom gases, which follow from the same universal horizon physics as that of strongly coupled $SU(N)$ plasmas, and lessons about jet quenching and rapid thermalization in the RHIC fireball [184]. Beyond providing new tools with which to model specific phenomena, holography has sometimes suggested new organizing principles for strongly quantum dynamics, as seen, for example, in semi-holographic models of non-Fermi liquids [309–311], or in the application of classical numerical relativity to study far-from-equilibrium dynamics in extreme quantum liquids [312].

The goal of this section is to provide a brief introduction to the basic structure of holography, first heuristically and then more precisely, as well as its origins and a few of its applications. We will also review a few major themes in the holographic study of extreme quantum matter and give a broad, if selective view of some of the key open questions in the field. Of necessity, we omit many important topics and references. For a more detailed introduction, several excellent reviews are available, including [20, 184, 310, 313–321], from which we have drawn heavily in preparing various parts of this section.

4.1. *Why should holography be true? Two heuristic pictures*

Why should a QFT in d-dimensions have anything to do with gravity in $d+1$, or vice versa? To give some intuition we describe two heuristic arguments that motivate such a holographic correspondence. We first consider a system that includes gravity, i.e., objects falling into a black hole, and argue that it should admit an equivalent description without gravity in one lower dimension. We then examine a strictly non-gravitational system, a QFT on a lattice, and argue that it should be related to an equivalent theory with gravity in one higher dimension.

4.1.1. Starting with gravity: falling into a black hole. Imagine standing far from a black hole[30] in $(d+1)$ dimensions and sending a robotic probe on a straight line trajectory directly toward the black hole. To follow the trajectory of the robot, we have attached to it a strobe light which flashes once a second. According to the robot, it will itself reach the horizon in finite proper time, as measured for example by a clock onboard the robot, and will thus have flashed its strobe only a finite number of times before crossing the horizon into the black hole interior.

---

[29] The role of all this structure is simple: given enough symmetry, it becomes possible to compute quantities in an interacting QFT at both weak and strong coupling. Holographic duality was discovered when Maldacena pointed out [302] that scattering amplitudes in the *maximally* supersymmetric four-dimensional QFT were, in fact, identical to amplitudes for the low-energy scattering of closed strings off a maximally symmetric 5D black hole in string theory. Conformal symmetry still plays an important role in holographic QFTs, but no more so than it does in any interacting QFT: conformal symmetry arises dynamically at fixed points of the RG.

[30] For example, in a spaceship following a distant circular orbit of the black hole. We emphasize that everything in section 4.1.1 follows from standard results about black holes in GR. For a clear and concise review of the physics of black holes in GR, see [322] and references therein. A more complete reference list can be found in [323].

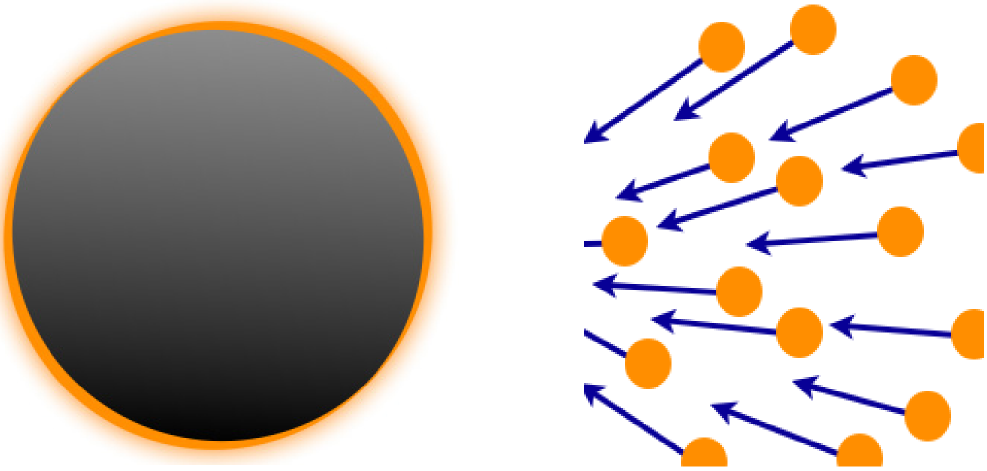

Figure 21. Two views of a swarm of robots falling toward a black hole in an isotropic shell formation. (left) A distant observer sees the robots slow down and spread over the horizon. According to her, the proper description of the robots is as a non-gravitating fluid stretched over the horizon. (right) An observer falling along with the robots would see them continue steadily towards and beyond the horizon, responding as one would expect to the black hole's gravitational pull. According to the in-falling observer, the appropriate description is GR in the bulk.

However, according to you, the distant observer, the story looks very different. As the robot falls into the gravitational well of the black hole, the light emitted by the flashing strobe must do work to escape the gravitational potential, and is thus red-shifted. By the constancy of the speed of light, this means that the flashes arrive to you at ever greater intervals. As the robot approaches the horizon, the observed red-shift in fact diverges, so that the final flash emitted at the moment of crossing the horizon takes an infinite time to arrive at the distant observer; light from flashes after the robot has crossed the horizon will never make it out to you. Thus, what you actually see is not the robot falling into the black hole, but rather the robot approaching the black hole and gradually slowing down and compressing into a thin membrane on the surface of the black hole. Indeed, to you, it appears as if the robot has stopped falling toward the massive black hole entirely, as if it had stopped responding to gravity at all (figure 21).

Now imagine sending in a swarm of robots in an isotropic shell around the black hole, so as to increase the mass of the black hole by a small amount corresponding to the matter and energy in the collapsing shell. Again, according to you, the distant observer, the shell will never actually appear to fall into the black hole. Rather, you will see it approach the black hole, slow down and effectively stop, forming a non-gravitating membrane which stretches over the true horizon. As each new bit of matter reaches this *stretched horizon*, it generates a disturbance in the membrane which spreads in waves through the stretched horizon as if through a fluid. Remarkably, all of the observable behavior of this system can be precisely captured by a non-gravitational viscous hydrodynamical model for a fluid on the stretched horizon which lives in only d rather than $(d+1)$ space–time dimensions.

If we describe this same process from the perspective of one of the robotic probes, the physics looks very different. In this frame, nothing special happens as the robots approach and pass the horizon. In contrast, the swarm of robots continues to fall toward the black hole, interacting with each other and with the gravity of the black hole in the full volume of $(d+1)$ dimensions. These dynamics can be modeled in a straighforward manner with GR in $(d+1)$ dimensions coupled to the matter and energy of the robots and the black hole.

This leaves us with two very different descriptions of our probes as they approach the horizon, one involving the gravitational dynamics of probes falling into a $(d+1)$-dimensional black hole and the other involving the hydrodynamic response to our probes of a d-dimensional fluid. Both descriptions accurately capture the behavior of our probes as viewed by two different observers. But there cannot be two facts of the matter about the physics of our probes before they cross the horizon: two observers cannot accurately observe contradictory events unfold[31]. The gravitational dynamics of a $(d+1)$-dimensional black hole must thus be, in some deeply non-local way, equivalent to the hydrodynamics of a d-dimensional fluid.

4.1.2. Starting without gravity: taking the renormalization group literally. Rather than start with gravity, let us start with a familiar non-gravitational field theory. Consider a system on a lattice with lattice spacing a and Hamiltonian,

$$H=\sum_{x,i} J_i(x)\mathcal{O}^i(x). \tag{60}$$

Here, x labels the sites in the lattice, i labels the various operators $\mathcal{O}^i(x)$ defined at each site and the $J_i(x)$ are coupling constants/sources for the operators[32] $\mathcal{O}^i$. Note that the sources will, in general, depend on both space and time; we take x to stand for both as is the typical space–time notation in relativity. Given this setup, what we generally want to compute is the physics of the ground state and the low-energy excitations over this ground state as a function of the microscopic coupling constants at the lattice scale. In general, however, exactly diagonalizing the Hamiltonian is intractably difficult.

Kadanoff and Wilson taught us a beautiful approach to this problem: the RG [324–326].

The basic move is to iteratively coarse-grain the lattice, making the lattice spacing larger with each step, so that at each step a single site represents the average of multiple sites in the previous lattice. We then tune the couplings so as to preserve the physics of the ground state and the low-energy excitations over it, replacing the fixed coupling $J_i(x)$ with a scale-dependent coupling $J_i(x,u)$, where u denotes the length scale at which we probe the system. By iterating this procedure, we eventually arrive at an effective description of long-wavelength modes of the system. As explained by Kadanoff and Wilson, the resulting flow of the couplings with scale u can be encoded in a β-function which is, remarkably, local in energy scale,

$$u\frac{\partial}{\partial u}J_i(x,u)=\beta_i(J_j(x,u),u). \tag{61}$$

When the beta function can be determined, for example in perturbation theory, this renormalization group approach is extremely powerful. Indeed, even finding the fixed points of the RG flow, corresponding to scale-invariant or conformal points governed by CFTs, can be enormously revealing, so much so that we often *define* scale-dependent QFTs with non-trivial RG flows by starting with a well-understood CFT and turning on a relevant deformation to generate the desired RG flow.

---

[31] Explicitly, so long as both observers remain outside the black hole, it remains physically possible for them to exchange information and compare their observations.

[32] We write sources with index down and operators with index up, for later notational convenience.

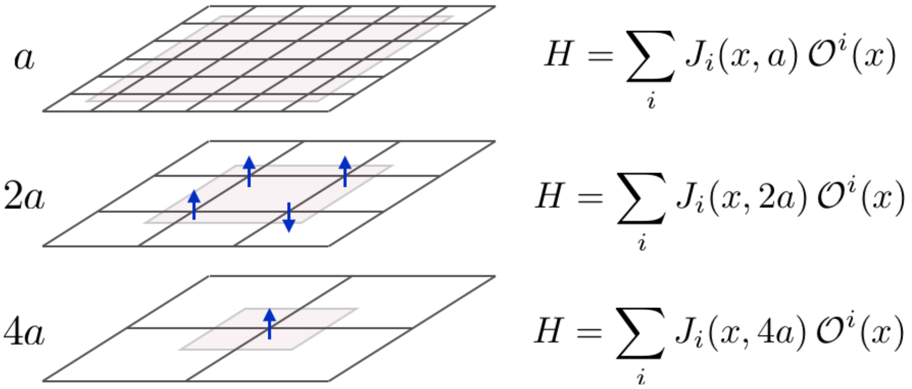
a
$H = \sum_i J_i(x, a)\, \mathcal{O}^i(x)$
$2a$
$H = \sum_i J_i(x, 2a)\, \mathcal{O}^i(x)$
$4a$
$H = \sum_i J_i(x, 4a)\, \mathcal{O}^i(x)$

Figure 22. Coarse-graining spins on a lattice á la Kadanoff and Wilson. At each step we coarse-grain by replacing the degrees of freedom on a block of sites by an average value on a single site, rescaling the lattice spacing a and tuning the couplings $J_i(x, a)$ in the Hamiltonian H so that the physics of the ground state and low-lying excitations remain invariant under the scaling operation.

However, in many complex, strongly coupled systems, including many strongly correlated systems governed by interesting conformal fixed points, the β-functions cannot be straightforwardly derived. It is thus tempting to search for a reorganization of the RG which determines the correct RG flow of the couplings without requiring an explicit calculation of the β-functions.

To this end, consider the following recasting of the Kadanoff–Wilson picture (figure 22). As before, consider our series of coarse-grained lattices with coarse-grained Hamiltonians and suitably tuned couplings, $J_i(x, u)$. Now imagine arranging each coarse-graining of the lattice into a stack ordered by scale, so that at each subsequent level of the stack the lattice spacing grows, $u = \{a, 2a, 4a, \ldots\}$. By construction, motion down the stack reproduces RG flow in the original lattice theory. In this one-higher-dimensional hyper-lattice, however, the couplings at each scale, $J_i(x, u)$, look a lot like fluctuating fields in a one-higher-dimensional lattice, varying both in space–time and in the scale direction (figure 23).

This raises an interesting question. We know that the true solution $J_i(x, u)$ solves the β-function equations for our original lattice theory, but we do not know what the correct β-functions for our lattice model are. Might there be some simple equation of motion on this one-higher-dimensional lattice whose solution is also $J_i(x, u)$? More generally, might there be a conventional local field theory[33] whose dynamics encode the full, and unknown, β-functions of the original lattice theory?

Consider the full RG stack of lattices labeled by a new RG coordinate, r, which runs from the original UV lattice cutoff, $r = a$, to the deep infrared (IR), $r \to \infty$. We seek a simple QFT defined on this one-higher-dimensional space whose $(d + 1)$-dimensional bulk fields, $\Phi_i(x, r)$, are in one-to-one correspondence with the couplings, $J_i(x)$, of the underlying lattice theory, with the values of the bulk fields at the UV end of the stack, $r = a$, determined by the microscopic couplings,

$$\Phi_i(x, a) = J_i(x). \tag{62}$$

---

[33] By a conventional local QFT we mean a QFT with a finite number of local fields governed by a Lagrangian with canonical quadratic kinetic terms and local interactions. One can certainly imagine other possibilities, but this structure turns out to be particularly useful in what follows.

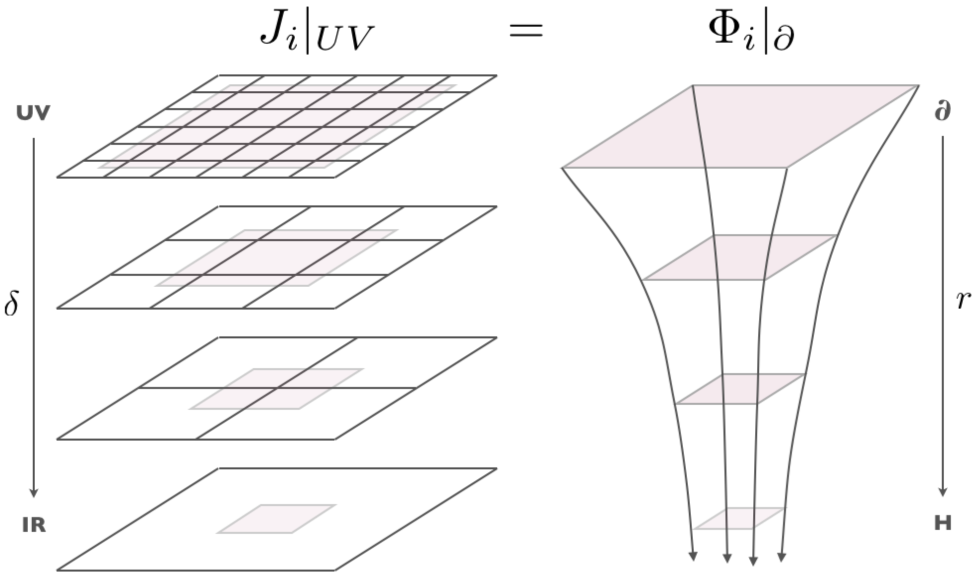

Figure 23. Holography promotes the tower of coarse-grained lattices to a one-higher-dimensional lattice, with the RG scale recast as a spatial dimension and the running couplings J_i replaced by dynamical fields Φ_i which asymptotically approach the ultraviolet (UV) couplings at the UV top of the stack and which satisfy a gravitational equation of motion in the bulk of the stack.

As a consequence, the bulk field $\Phi(x, r)$ must have the same charges, tensor structure and other quantum numbers as the corresponding coupling[34], $J_i(x)$. Thus with each scalar operator $\mathcal{O}(x)$ is associated a bulk scalar field $\Phi(x, r)$ such that $\Phi(x, a)\mathcal{O}(x)$ is a scalar operator we can add to our Hamiltonian, with each current operator $\mathcal{J}^\mu(x)$ is associated a bulk vector field, $A_\mu(x, r)$, such that $A_\mu(x, a)\mathcal{J}^\mu(x)$ is a scalar and so forth. In particular, with the canonical stress-tensor $T^{\mu\nu}(x)$ is associated a canonical spin-2 field in the bulk, $g_{\mu\nu}(x, r)$.

What can the Lagrangian of this bulk QFT be? It is not the original lattice Hamiltonian—the fields of the bulk QFT encode the dynamics of the *couplings* in RG space and not of the operators on the original lattice. *A priori*, the natural thing to do is to write down the most general effective field theory allowed by the fields and symmetries of the system. However, a simple argument greatly constrains the possibilities. Any QFT defined by perturbation from a fixed point CFT as discussed above includes among its operators a canonical stress-energy tensor, $T_{\mu\nu}$, arising as the Noether current corresponding to the space–time symmetry group. The bulk QFT should thus include a corresponding canonical spin-2 field, $g_{\mu\nu}$. But by the Weinberg and Weinberg–Witten theorems[35] and assuming the existence of a Lorentz-invariant continuum limit, any spin-2 field either must decouple at low energy, which would imply that sources of momentum or energy in the original lattice theory do not affect the system, or must couple universally according to the equivalence principle. In other words, the bulk field theory must either be topological or be a theory of gravity that reduces, at long wavelength and low energy, to GR.

---

[34] Recall that our couplings are defined as the coefficients of the associated operators in the Hamiltonian, and so have definite quantum numbers. Note too that, at fixed points of the β-functions where the system is scale invariant, the dimension of an operator is also a quantum number. This is also true in perturbations around such conformal fixed points.

[35] Weinberg [327] argued that any Lorentz-invariant theory of a spin-2 field must respect the equivalence principle or suffer IR pathologies; the Weinberg–Witten theorem [328] further states that any massless spin-2 field must either be the graviton or suffer one of a host of pathologies, for example not having a conserved stress tensor. For an illuminating discussion of these constraints, see [329] and references therein.

If all this can be done consistently, we would appear to have two distinct descriptions of our system: the original lattice theory with lattice spacing a and sources $J_i(x)$ which satisfy first-order β-function equations; and a gravitational description living in a one-higher-dimensional bulk whose fields, $\Phi_i(x, r)$, are in one-to-one correspondence with the sources $J_i(x)$ but which satisfy conventional second-order local field equations. In other words, we have two ways to compute correlation functions in our system: either we may work in the lattice theory and solve the QFT problem; or we may work in the gravitational description and solve a GR problem.

In order for these two descriptions to have any chance of being equivalent[36], the bulk theory must satisfy two among many rather unusual properties, both of which turn out to be basic properties of black holes. Firstly, the entropy as computed in the two descriptions must be the same. In the original lattice QFT, the entropy will generically be extensive, with the entropy S_{QFT} of the d-dimensional QFT in a region R_d scaling linearly with its spatial volume, $S_{\mathrm{QFT}}(R_d) \propto V(R_d)$. As a result, the entropy as measured in the one-greater-dimensional bulk gravity, S_{GR}, must be sub-extensive, scaling as the area bounding a region, $S_{\mathrm{GR}}(R_{d+1}) \propto A(R_{d+1}) \propto V(\partial R_{d+1})$. This is a basic feature of any theory of gravity[37]: the entropy in a volume is bounded by the entropy of a black hole that fits inside that volume [330]. Since the entropy of a black hole is proportional to the surface area of its horizon, $S_{\mathrm{BH}} = \frac{1}{4}A_{\mathrm{H}}$, this tells us that the entropy of any region in a theory of gravity is bounded by its surface area, $S_{\mathrm{GR}}(R_d) \leqslant \frac{1}{4}A(R_d)$. Indeed, the fact that the entropy of a black hole can be entirely associated with its surface, and that as a result the entropy in gravity is necessarily sub-extensive, is the origin of the term *holography* [331, 332][38].

Secondly, the β-function equations in the boundary lattice theory are first order in the scale r, while the gravitational equations of motion in the bulk are second order in r. Physically, this means that specifying the sources in the lattice QFT at one scale completely determines the coupling at all other scales, i.e. completely fixes the RG trajectory. By contrast, in the gravitational system we have two solutions to our second-order equation, so we must further specify the derivative of the couplings with scale to uniquely fix the RG trajectory. How can these be equivalent? Somehow it must be the case that the gravitational system automatically selects one of the two solutions of the bulk equations as physical.

It is again the physics of black holes which explains how this can happen. Imagine solving a simple wave equation in the neighborhood of a black hole horizon. In principle, there are two solutions to the second-order wave equation. However, since things can fall into a black hole horizon but nothing can escape, it is natural to organize the solutions near the horizon into in-going and out-going waves. Moreover, when working in the Euclidean signature as is done, for example, in computing thermodynamic data, only the in-going solution is regular at the horizon; the out-going solution is always irregular at the horizon. Thus, while the gravitational field equations are indeed second order, the presence of a black hole horizon effectively adds a second boundary condition, so that we again need only specify a single boundary condition to completely determine the solution. Since the location of the horizon encodes the thermodynamic variables (temperature, entropy) of the system, the physical meaning of this second boundary condition is to specify the state of the QFT.

---

[36] By *equivalent*, we mean that their partition functions are exactly equal, so that both descriptions contain precisely the same information, as we shall spell out in detail in section 4.2.3. Note that this implies their free energies and other thermodynamic data must similarly, and identically, match.

[37] Various other theories, including topological theories, share the property of sub-extensivity. However, by Weinberg–Witten, to have both an area law and an interacting spin-2 field, gravity appears to be the only option.

[38] The term *holography* is used by analogy with familiar holograms, and is meant to convey the surprising fact that the information about the $(d+1)$-dimensional bulk space–time is being reliably encoded in the data of the QFT living on the d-dimensional boundary. The analogy is in various ways quite apt. For example, information that is local in the $(d+1)$-dimensional bulk is encoded non-locally in the d-dimensional boundary. However, one should not take the analogy too literally; for example, the encoding here involves quantum mechanics and gravity.

These heuristics suggest a connection between QFT in d space–time dimensions and quantum gravity in $d+1$ space–time dimensions, with the fields in the gravitational system exactly paired with sources in the dual QFT and with the extra spatial coordinate in the gravitational bulk playing the role of an RG scale. The next sections will describe a precise and computationally effective theory of this connection known as holographic duality and review some of the lessons gleaned from its study to date.

4.1.3. Coda: the surprise of locality. It is worth pausing to contemplate just how remarkable this proposal truly is. Consider a lattice theory in d-dimensions which has, in its Hamiltonian, a real parameter q. You might think of q as a global charge, or as a momentum, or something else entirely, which can in principle be varied adiabatically. All correlation functions in the theory are thus functions of d-dimensional space–time, (x, t), as well as of our parameter, q.

We might be tempted to take a stack of such lattices, each with a different value of q, call the 'space' labeled by (x, t, q) a $(d+1)$-dimensional space–time, and claim that the resulting tensor product of lattice QFTs is well described by a $(d+1)$-dimensional QFT living on this $(d+1)$-dimensional space–time.

This would be wrong. Being $(d+1)$-dimensional means that all observables should transform appropriately under translations, rotations, accelerations and general coordinate transformations among x, t and q. They should also respect $(d+1)$-dimensional causality and locality. But that is absurd! What does it mean to rotate between the x direction and the q parameter? Meanwhile, there is absolutely no reason for correlators of two operators which live on lattices with very different values of q to vanish outside some $q-t$ light cone, or for operators from distant q-slices to have vanishing equal-time commutators in whatever state you choose to call the ground state. Locality and causality are almost sure to be violated, and badly. This is not a $(d+1)$-dimensional system.

To summarize, saying that the system is well described by a local, causal, unitary $(d+1)$-dimensional QFT implies a long list of constraints on the $(d+1)$-dimensional correlation functions. These translate into constraints on the original lattice theory and on the q-dependence of its correlation functions. It is by no means obvious that any such stack of lattices can ever have a causal, local $(d+1)$-dimensional description. From this point of view, holographic duality, in which the role of q is played by the energy scale and in which the $(d+1)$-dimensional description is in fact local, causal and even classical when the d-dimensional QFT is very strongly interacting, is nothing short of astonishing. Precisely why, when and how this locality arises is perhaps the deepest question in holography[39].

---

[39] See, for example, [333] for a discussion of locality in holographic duality, and [334] for an interesting spin of the problem.

4.2. Essential holography

Holographic duality [302–304] is a precise equivalence between certain d-dimensional quantum field theories and $(d+1)$-dimensional gravitational theories which provides a sharp realization of the heuristics described in section 4.1. In particular, all the basic features noted above will reappear below: the fields in the bulk correspond to the couplings in the QFT; the RG flow of the QFT is encoded in the radial evolution of the gravitational theory along an extra dimension; and black hole horizons play a key role. Of particular importance is the precise geometry of the emergent space–time, as well as the precise relationship between observables in the QFT and those of the gravitational dual. We will explore a few simple examples along the way.

For clarity of presentation, we will focus our discussion on the simplest possible QFTs from the point of view of the RG, i.e. Lorentz-invariant CFTs which are fixed points of the RG. Since the RG maps to radial evolution in the gravitational dual, QFTs corresponding to RG fixed points must be dual to geometries which are translationally invariant along the emergent radial direction up to overall rescalings; adding Lorentz invariance then entirely fixes the dual geometry to be AdS, whose geometry we will now describe[40].

4.2.1. Gravity and matter in anti-de Sitter space. AdS_{d+1} is a homogeneous, isotropic space–time with $d+1$ space–time dimensions and constant negative curvature whose metric can be conveniently written as[41]

$$\mathrm{d}s^2 = \frac{L^2}{r^2}\left[-\mathrm{d}t^2 + \mathrm{d}\vec{x}^2 + \mathrm{d}r^2\right], \tag{63}$$

where $\vec{x}$ are $(d-1)$ spatial coordinates, t is a time-like coordinate and r is a final 'radial' spatial coordinate. The constant parameter L, known as the AdS-radius, sets the radius of curvature, with the Ricci curvature scalar given by the constant

$$R = -\frac{d(d+1)}{L^2}. \tag{64}$$

The space is thus weakly curved when L is large and strongly curved when L is small. At fixed values of the radial coordinate r_*, the metric reduces to the d-dimensional Minkowski metric on flat space–time rescaled by L^2/r_*^2,

$$\mathrm{d}s^2_{r_*} = \frac{L^2}{r_*^2}\left[-\mathrm{d}t^2 + \mathrm{d}\vec{x}^2\right]. \tag{65}$$

It is thus useful to think of AdS as a stack of slices of flat space, with r controlling the physical scale on that slice. The slice at $r \to 0$ is referred to as the *boundary*. Note that the volume element of the boundary slice at $r \to 0$ diverges, while the proper distance to the boundary also diverges. Thus the boundary at coordinate $r \to 0$ should be understood as the set of points at *spatial infinity*[42].

---

[40] The first examples of holographic duality involved such AdS/CFT dual pairs, and hence holographic duality is often referred to as AdS/CFT. Just as we can always construct a general QFT by perturbing away from a CFT fixed point, we can extend such AdS/CFT dualities to more general QFTs and geometries by perturbing both sides in corresponding ways, so we do not lose too much by focusing on these examples.

[41] For the remainder of our discussion of holography we will work in *natural units* in which $\hbar = 1$ and $c = 1$.

[42] In computations, it will often be useful to regulate this divergence by introducing a cutoff at $r = a$, with $a \ll L$, removing the regulator, $a \to 0$, only at the end of all computations.

Importantly, the full geometry is invariant under not just translations, rotations and Lorentz boosts along each slice, but also under general scaling transformations

$$(r, \vec{x}, t) \to (\alpha r, \alpha \vec{x}, \alpha t) \tag{66}$$

as well as rotations and boosts that mix r with the other dimensions. The full isometry group of AdS_{d+1} is, in fact, identical to the conformal group in d space–time dimensions[43].

It is often useful to think of AdS as a harmonic trap for gravity. More precisely, the constant negative curvature of AdS acts as a harmonic trap such that if you sit inside AdS at an arbitrary point and throw a ball, it will inevitably fall back to you in finite time. Indeed, if you fire out a photon, that photon will run away to the boundary at $r = 0$ and return, again in finite observed time[44]. And yet while nothing, not even gravity, gets out of the AdS box, the space–time remains completely homogeneous and isotropic. AdS thus acts as a peculiarly graceful IR regulator for gravity which breaks none of the symmetries of flat space.

All of this implies an intimate connection between the radial coordinate in the bulk of AdS and spatial scales along the boundary: probing short distances (or high energies) along the boundary corresponds to probing the bulk only near the boundary at $r = 0$, while probing long distances (or low energies) along the boundary corresponds to data deep in the interior of AdS. Roughly, then, we can think of the region near the boundary of AdS as associated with the UV physics of the boundary, and of the deep interior of AdS as associated with the IR physics of the boundary.

To visualize this association, consider a fixed tension rope whose ends are glued to the boundary at $r = 0$. Due to the constant negative curvature of AdS, the bulk of the rope is drawn into the bulk of the AdS space–time, forming an arc drooping away from the boundary. If the ends of the rope are held close together—probing short distances or high energy in the four-dimensional (4D) boundary—the rope barely hangs into the bulk (figure 24). If we instead pull the ends of the rope far apart—probing long distances or low energy in the 4D boundary—the rope dips much further into the bulk. In this way, probing the IR of the boundary corresponds to probing the deep interior of the AdS, $r \to \infty$, while probing the UV of the boundary corresponds to focusing on the near-boundary region of the geometry, $r \to 0$. Note that the near-boundary cutoff at $r = a \ll 1$ thus acts as a UV regulator, while the horizon at $r = r_{\mathrm{H}}$ provides a natural IR regulator.

Importantly, the AdS is a solution to the equations of motion of a generic Wilsonian action for the metric

$$I_{\mathrm{Gravity}} = \frac{1}{16\pi G_{\mathrm{N}}} \int \mathrm{d}^{d+1}x \, \sqrt{-g} \left(-2\Lambda + R + c_2 R^2 + c_3 R^3 + \cdots \right). \tag{67}$$

Here, G_{N} is the Newton constant, g is the determinant of the space–time metric, $g = \det(g_{\mu\nu})$, R is the Ricci curvature scalar built out of two derivatives of the metric, $R \sim \partial\partial g$, Λ is a cosmological constant (also known as the tension of the vacuum), and the $\cdots$ represent all other scalars one can build out of the metric and its derivatives. The R term in the action plays the role of a two-derivative kinetic term for the metric, $g_{\mu\nu}$. GR is defined by retaining only the

---

[43] Note that our regulator at $r = a$ explicitly breaks conformal symmetry but does not break space–time rotational or translational invariance in x and t, as a more conventional UV regulator such as a hard momentum cutoff would. Matching regulators between bulk and boundary turns out to be quite subtle, cf [335].

[44] Unlike the electromagnetic harmonic traps used in trapping cold atoms, AdS is homogeneous and isotropic: there is no center to which all objects return. Rather, the curvature ensures that any two initially parallel trajectories always curve toward each other, no matter where they begin, and then oscillate.

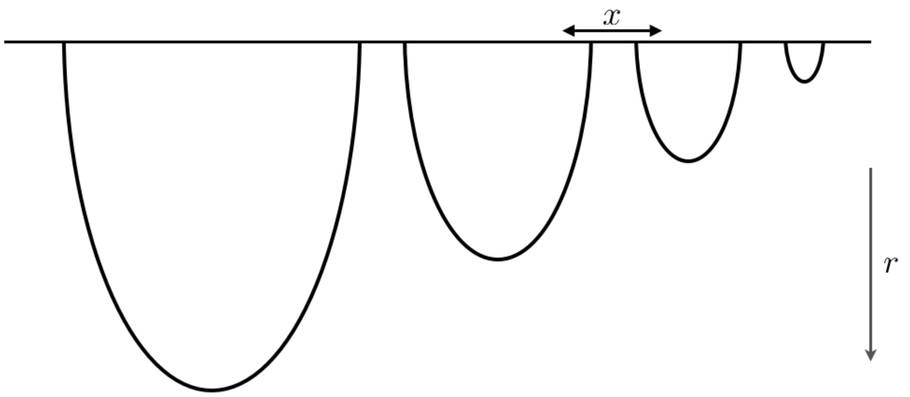

Figure 24. Hanging rope in AdS. When the ends are close together in the spatial direction, x, so that the rope probes short distances and high energies in the QFT, the rope stretches only slightly into the bulk of the AdS space–time. When the ends are far apart, probing long distances and low energies in the QFT, the rope stretches deep into the AdS bulk. Thus the boundary of AdS is associated with the UV of the QFT, while the deep bulk is associated with the IR.

lowest dimension kinetic term, R, giving the Einstein–Hilbert action whose equation of motion is the Einstein equation,

$$R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R = -\Lambda g_{\mu\nu}. \tag{68}$$

The $c_n R^n + \cdots$ terms then represent higher-derivative corrections to GR. The reader can verify that the AdS metric solves the Einstein equation (68) with the AdS radius L determined by the cosmological constant, Λ, as $\Lambda = -\frac{\mathrm{d}(d-2)}{2L^2}$. One can then show [336] that AdS continues to be a solution even when we include generic higher-curvature corrections: all that changes is the precise relationship between Λ and L.

Importantly, the Ricci scalar in AdS scales as $R \sim \frac{1}{L^2}$ (cf (64)). As a result, we may neglect higher-curvature corrections in AdS when L^2 is large compared to the appropriate powers of the dimensionful couplings c_n. For example, if the higher-derivative operators are generated by quantum gravity effects, the dimensional couplings are controlled by the Planck length, ℓ_{p}. Quantum corrections to the metric can thus be neglected in AdS when L is large compared to the Planck length, $\frac{L}{\ell_{\mathrm{p}}} \gg 1$. Similarly, GR receives corrections from string theory even at the classical level. Such corrections become important at a characteristic length scale known as the string scale, ℓ_{s}. In AdS, these stringy corrections can be neglected so long as $\frac{L}{\ell_{\mathrm{s}}} \gg 1$. Our generic Wilsonian action for quantum gravity in AdS is thus well approximated by classical GR when the AdS radius is large compared to all scales in the problem,

$$\frac{L}{\ell_{\mathrm{p}}} \gg 1, \qquad \frac{L}{\ell_{\mathrm{s}}} \gg 1, \ldots, \tag{69}$$

where the ... denotes other sources of corrections to GR which may kick in at other length scales.

Black holes in AdS. If pure empty AdS is a ground state for gravity, finite-temperature states correspond to black holes inside AdS. The simplest such asymptotically AdS black hole is the

AdS–Schwarzchild black brane,

$$\mathrm{d}s^2 = \frac{L^2}{r^2}\left[-f(r)\,\mathrm{d}t^2 + \mathrm{d}\vec{x}^2 + \frac{1}{f(r)}\mathrm{d}r^2\right], \tag{70}$$

with emblackening factor

$$f(r) = 1 - \frac{r^d}{r_\mathrm{H}^d}. \tag{71}$$

Near the asymptotic boundary at $r \to 0$, $f \to 1$, so this metric is asymptotically AdS_{d+1}. At $r = r_\mathrm{H}$, however, $f \to 0$, signaling the presence of a black hole horizon. This horizon in turn shields us from a physical singularity at $r \to \infty$ by ensuring that nothing which is inside the horizon, and thus sensitive to the singularity, can ever escape to influence events in the rest of the space–time. Since the entire solution is translationally invariant in the $(d-1)$ spatial directions, $\vec{x}$, this black brane is not a compact object, but rather extended in all directions other than r.

Importantly, a classical black hole is a thermodynamic object with a definite temperature, energy and entropy, as shown by Bekenstein and Hawking (see [330] and references therein). If this were not the case, we could violate the third law of thermodynamics by throwing all of our waste heat into a black hole and thus build perfectly efficient Carnot engines. Black holes in AdS similarly carry definite temperatures, energies and entropies, although the details are a bit different from the flat space case [337]. In the case at hand, the Hawking temperature T, energy density ϵ and entropy density s of the black brane are then

$$T = \frac{d}{4\pi\, r_\mathrm{H}}, \qquad \epsilon = \frac{d-1}{16\pi\, r_\mathrm{H}^d}\left(\frac{L}{\ell_\mathrm{p}}\right)^{d-1}, \qquad s = \frac{1}{4\, r_\mathrm{H}^{d-1}}\left(\frac{L}{\ell_\mathrm{p}}\right)^{d-1}, \tag{72}$$

where ℓ_p is the Planck length defined by $G_\mathrm{N} = \ell_\mathrm{p}^{d-1}$. The ratio $\frac{L}{\ell_\mathrm{p}}$ thus measures the AdS scale in Planck units.

Physically, these thermodynamic relations are telling us that if we increase ϵ by throwing additional mass or energy into the black brane, the horizon swells outward toward the asymptotic boundary ($r_\mathrm{H} \to 0$) and the black brane heats up. The specific heat of this black brane is thus positive. This means that black holes in AdS can come to equilibrium. Again, this can be traced to AdS playing the role of a harmonic trap for gravity: any radiation that the black hole evaporates away will, in finite time, fall back into the horizon[45].

Charged black holes in AdS. More generally, we can add charge to our black brane by adding a non-trivial Maxwell field

$$\mathrm{d}s^2 = L^2\frac{-f(r)\mathrm{d}t^2 + \mathrm{d}\vec{x}^2 + \frac{1}{\mathrm{f(r)}}\mathrm{d}r^2}{r^2}, \qquad A = A_t(r)\mathrm{d}t, \tag{73}$$

with emblackening factor

$$f = 1 - M\,r^d + Q^2\,r^{2(d-1)}, \tag{74}$$

---

[45] This is very different from black holes in asymptotically flat space, whose specific heat is negative and which never come to equilibrium in the absence of external forces. Black holes in AdS thus behave like familiar systems in the thermodynamic limit, while black holes in flat space behave more like finite-volume sub-systems.

and electromagnetic scalar potential[46]

$$A_t(r) = \mu \left(1 - \left(\frac{r}{r_{\rm H}}\right)^{d-2}\right), \tag{75}$$

where $\mu = \frac{2Q}{\mathcal{C}}\, r_{\rm H}^{d-2}$ and $\mathcal{C} = \sqrt{\frac{2(d-2)}{d-1}}$.

This metric and gauge field together extremize the Einstein–Maxwell action[47],

$$I_{\rm ME} = \frac{1}{16\pi G_{\rm N}} \int {\rm d}^d x\, \sqrt{-g} \left[-2\Lambda + R - \frac{L^2}{4e^2} F^2\right], \tag{76}$$

again with $\Lambda = \frac{-d(d-2)}{2L^2}$. In this geometry, the horizon lies at the radial position $r = r_{\rm H}$ implicitly defined as the value of r where $f(r)$ vanishes. M and Q then determine the Hawking temperature of the horizon,

$$T = \frac{d}{4\pi r_{\rm H}} \left(1 - \frac{d-2}{d} Q^2 r_{\rm H}^{2d-2}\right), \tag{77}$$

as well as its energy, entropy and charge densities,

$$\epsilon = M\, \frac{d-1}{16\pi} \left(\frac{L}{\ell_{\rm p}}\right)^{d-1} \qquad s = \frac{1}{4\, r_{\rm H}^{d-1}} \left(\frac{L}{\ell_{\rm p}}\right)^{d-1}, \qquad \rho = Q\, \frac{d-1}{8\pi \mathcal{C}} \left(\frac{L}{\ell_{\rm p}}\right)^{d-1}. \tag{78}$$

It is then straightforward to check that these variables satisfy the first law of thermodynamics, ${\rm d}\epsilon = T\,{\rm d}s + \mu\,{\rm d}\rho$, with μ playing the role of the chemical potential.

A curious feature of these charged black holes is that the zero-temperature limit has finite entropy density. To see this, note that we can tune Q to make T vanish, $Q \to Q_* = \sqrt{\frac{d}{d-2}}\frac{1}{r_{\rm H}^{d-1}}$. In this limit, $r_{\rm H}$, and thus the entropy density, remains non-zero. This finite zero-temperature entropy represents an enormous degeneracy of the ground state of this black hole[48] and correspondingly a large number of potential instabilities of these black hole solutions; these turn out to play an interesting role in holographic descriptions of quantum phase transitions. For our present purposes, let it suffice to say that this is an interesting feature of these black branes which will crop up from time to time in the following.

Matter in AdS: the effective action as functional of boundary conditions. The fact that AdS behaves like a homogeneous harmonic trap has important consequences when we consider matter fields in AdS. Consider, for simplicity, a scalar field Φ with mass m^2 in the background of an uncharged AdS black brane (70) with classical action

---

[46] Note that we are working with units in which $\hbar = 1$ and $c = 1$. For example, M and Q scale as energy and charge densities on the (non-compact) horizon, $M \sim E/V_{d-1} \sim L^{-d}$ and $Q \sim 1/V_{d-1} \sim L^{-(d-1)}$, as required by equation (74). Similarly, both the $d+1$-dimensional scalar potential, A_t, and the d-dimensional chemical potential, μ, scale as $\frac{1}{L}$.

[47] Here $F_{\mu\nu} = \partial_{[\mu} A_{\nu]}$ is the totally antisymmetric electromagnetic field strength tensor ($E^i = -F^{0,i}$ and $B^i = -\epsilon^{ijk} F_{jk}$) and $F^2 = F^{\mu\nu} F_{\mu\nu} = (\vec{E}^2 - \vec{B}^2)$ is the Maxwell kinetic term.

[48] The precise microcanonical counting of this ground-state entropy was a long-standing puzzle in quantum gravity which was first solved, for a specific set of black holes in string theory, by Strominger and Vafa in 1996 [338]. A similar counting can now be performed for a much larger set of black holes, and in fact played a very important role in the discovery of holographic duality.

$$I_\Phi \propto \int \mathrm{d}^{d+1}x\, \sqrt{-g}\left(-\frac{1}{2}\,(\partial\Phi)^2 - \frac{m^2}{2}\Phi^2 + \cdots\right), \tag{79}$$

where the $\cdots$ denotes interactions with other matter fields which we will, for the moment, neglect. Working for a moment with a simple plane wave with momentum k in the boundary dimensions, $\Phi(x, r) = \Phi(r)\mathrm{e}^{\mathrm{i}k\cdot x}$, the wave equation following from the above action becomes

$$r^2 f\, \Phi''(r) - r\left[rf' - (d-1)f\right]\Phi'(r) - \left[k^2r^2 + m^2L^2\right]\Phi(r) = 0. \tag{80}$$

This equation has two simple degenerate points, one at the boundary where $r = 0$ and the other at the horizon where $f = 0$. Let us study a general solution in the neighborhood of each.

Near the boundary at $r \to 0$, where $f \to 1$, the general solution of this equation reduces to

$$\Phi \sim \phi_{d-\Delta}(k)\, r^{d-\Delta} + \phi_\Delta(k)\, r^\Delta + \cdots, \tag{81}$$

where the scaling dimension Δ is determined by[49] $\Delta(\Delta - d) = m^2L^2$. So long as $m^2 \geqslant -d^2/4L^2$, the scaling dimension is real. This Breitenloner–Freedman (BF) bound tells us that a small negative m^2 does not lead to an instability in AdS as it would in flat space—instead, the would-be instability is lifted by the harmonic potential generated by the AdS curvature. Note that in the flat space limit, $L \to \infty$, this window of ‘allowed tachyons’ disappears. So long as the BF bound is satisfied, specifying a solution reduces to specifying the two radial integration constants, $\phi_{d-\Delta}$ and ϕ_Δ.

A key fact about this asymptotic behavior is that, so long as $m^2L^2 > 1 - \frac{d^2}{4}$, the mode $\phi_{d-\Delta}$ is *non-normalizable* according to the natural norm on a constant-time spatial slice, Σ_t,

$$(\Phi_1, \Phi_2) = -\mathrm{i}\int_{\Sigma_t} \mathrm{d}z\, \mathrm{d}\vec{x}\sqrt{-g}g^{tt}\left(\Phi_1^*\partial_t\Phi_2 - \Phi_2^*\partial_t\Phi_1\right). \tag{82}$$

Varying this non-normalizable mode thus corresponds to a large (divergent) change in the action. We must thus fix the non-normalizable mode to make the variational problem in the bulk well posed, for example by specifying a boundary condition for the bulk field Φ near the boundary, $r \to 0$. This leaves us with a one-parameter family of solutions labeled by the remaining integration constant, ϕ_Δ, which can be chosen arbitrarily.

We thus draw a general conclusion for studying matter fields in AdS: in order for the least-action principle to be well defined in AdS, we must fix the values of the non-normalizable modes of all fields. This is typically done by fixing boundary conditions for the non-normalizable modes of all fields at the AdS boundary, $r \to 0$.

Near the horizon at r_{H}, where $f \sim f_H(1 - r) \to 0$, the equation of motion again degenerates. Here, however, the physics is rather different. Physically, this equation has two kinds of solution: one which is regular, so that the degenerating $f\Phi''$ term is negligible; and one which is irregular at the horizon, so that the $f\Phi''$ term is not negligible. But why are half the solutions irregular? The reason was alluded to above: near the horizon, all waves can be expressed as a superposition of in-going and out-going waves. To see that this is precisely what is going on, let us study our scalar equation at non-zero frequency. Recalling that the magnitude of the d-momentum k in the metric (70) is $k^2 = -\frac{r^2}{f}\omega^2 + r^2\vec{k}^2$, the dominant terms in the scalar equation near the horizon are

---

[49] A similar scaling behavior is obtained for any set of matter fields, with their resulting scaling dimensions Δ determined by their precise spins, masses and interactions.

$$f\,\partial_r(f\,\partial_r\Phi) + \omega^2\Phi \sim 0. \tag{83}$$

It is useful to switch to the near-horizon coordinate ρ defined by $f\partial_r = \partial_\rho$, i.e. $\mathrm{e}^{-\rho} \sim (1-r)^{1/f_{\mathrm{H}}}$, in terms of which the horizon lies at $\rho \to \infty$. In terms of this coordinate, the general solution can be expressed as a superposition of in-going and out-going waves as

$$\Phi = \Phi_{\mathrm{in}}(\rho)\,\mathrm{e}^{-\mathrm{i}\omega(t-\rho)} + \Phi_{\mathrm{out}}(\rho)\,\mathrm{e}^{-\mathrm{i}\omega(t+\rho)}. \tag{84}$$

For future reference, note that the in(out)-going solutions satisfy the first-order constraint at r_{H},

$$f\,\partial_r\Phi_{\mathrm{in}} = \mathrm{i}\omega\Phi_{\mathrm{in}}, \qquad f\,\partial_r\Phi_{\mathrm{out}} = -\mathrm{i}\omega\Phi_{\mathrm{out}}. \tag{85}$$

Note that, in terms of the original radial coordinate, r, the phases of both of these time-dependent solutions accumulate near the horizon. One can check that, in the limit of zero frequency, the in-going wave is regular at the horizon while the out-going wave is irregular, matching our zero-frequency analysis above. Note, too, that this suggests a natural prescription for analytic continuation to Euclidean time: we should choose the continuation so that the in-going wave is regular at the Euclidean horizon, $\mathrm{e}^{\mathrm{i}\omega\rho} \to \mathrm{e}^{-\omega_{\mathrm{E}}\rho}$, i.e. we should define $\omega = \mathrm{i}\omega_{\mathrm{E}}$, where $\omega_{\mathrm{E}} > 0$.

Again, this turns out to be a general result: in the presence of a black hole horizon, the Laplacian in any tensor representation takes the form $\Box \sim r^2 f\partial_r^2 + \frac{r^2\omega^2}{f} + \ldots$, so we must impose a boundary condition on our solutions such that they correspond to regular, in-falling modes at the horizon.

Thus, while our equation of motion is second order, there is only one linearly independent solution which is regular everywhere in the bulk and in-going at the horizon. Our gravity problem in AdS, in the presence of a black hole horizon, becomes effectively a first-order problem[50]. Explicitly, fixing the non-normalizable mode $\phi_{d-\Delta}$ at the boundary and imposing in-falling boundary conditions at the horizon completely determines our solution for Φ throughout the bulk of AdS and as a result determines ϕ_Δ. However, in order to actually compute ϕ_Δ given $\phi_{d-\Delta}$, we must solve our problem not just at the boundary but all the way through the bulk to the horizon, where we impose the appropriate boundary conditions. The relation between $\phi_{d-\Delta}$ and ϕ_Δ is thus also determined by IR physics and not just UV physics.

It is illuminating to see this work in detail in an analytically solvable example. Consider again our scalar field Φ but now in pure AdS, i.e. with no black hole. This corresponds to studying our dual QFT at zero temperature and zero density. Expanding in plane waves, $\mathrm{e}^{-\mathrm{i}(\omega t - kx)}$, the general solution to the bulk equation of motion can be found in closed form

$$\Phi(t,x,r) = \left[\phi_{\mathrm{reg}}\, r^{\frac{d}{2}}\, K_\nu(\kappa r) + \phi_{\mathrm{irreg}}\, r^{\frac{d}{2}}\, I_\nu(\kappa r)\right]\mathrm{e}^{-\mathrm{i}(\omega t - kx)}, \tag{86}$$

where $\nu = \Delta - \frac{d}{2}$, $\kappa = \sqrt{\omega^2 + \vec{k}^2}$, K_ν and I_ν are modified Bessel functions and ϕ_{reg} and ϕ_{irreg} are the two integration constants. Since $I_\nu(\kappa r) \sim \mathrm{e}^{\kappa r}$ near the 'AdS horizon' at $r \to \infty$, regularity requires that $\phi_{\mathrm{irreg}} = 0$. We thus have just a single integration constant, ϕ_{reg}, which fixes the overall normalization of the in-falling mode in Φ. Near the boundary, the two asymptotic integration constants $\phi_{d-\Delta}$ and ϕ_Δ in (81) are not independent. We can see how they are related by setting $\phi_{\rm irreg} = 0$ in (86) and expanding the solution near the boundary. After a little algebra, this gives

---

[50] A similar story holds for fermions, although the details differ because the Dirac equation is naturally first order [339].

$$\phi_\Delta = \frac{\Gamma(\frac{d}{2} - \Delta)}{2^{\Delta-\frac{d}{2}}\Gamma(\Delta - \frac{d}{2})}\left(\omega^2 + k^2\right)^{\Delta-\frac{d}{2}} \phi_{d-\Delta}. \tag{87}$$

Thus, once we impose regularity on our bulk solutions, we need only one boundary condition in the UV at $r \to 0$ to fully specify a solution of the second-order bulk equations of motion. Meanwhile, we know that we need to fix the non-normalizable mode $\phi_{d-\Delta}$ at the boundary so that our variational problem is well posed. Once we impose regularity in the bulk, i.e. in-falling at the horizon, the value of ϕ_Δ is completely determined.

This teaches us an important lesson. Suppose that we want to study the partition function for our gravitational theory in AdS, integrating over all bulk fields, $\Phi(x, r)$. To make the variational problem well posed, we must specify a boundary condition for each of the bulk fields at the boundary of AdS, as discussed above,

$$\phi_{d-\Delta}(x) = \lim_{r\to 0} r^{\Delta-d}\Phi(x, r), \tag{88}$$

where the factor of $r^{\Delta-d}$ should be thought of as the appropriate wavefunction renormalization of the bulk field $\Phi(x, r)$. This means that the partition function in AdS, and thus the effective action $\Gamma_{\rm AdS} = -ln[Z_{\rm AdS}]$, is a functional of the boundary conditions for all bulk fields,

$$Z_{\rm AdS}[\phi_{d-\Delta}(x)] \equiv Z_{\rm AdS}[\Phi[\phi_{d-\Delta}(x)]]. \tag{89}$$

What one means by 'a theory of gravity in AdS'. In discussing holography, we will regularly refer to classical gravity, quantum gravity and perturbative gravity in AdS, so it is worth taking a moment to define our terms here.

By *quantum gravity in AdS* we mean a complete quantum description of the gravitational system, i.e. string theory in AdS. In general, string theory is a rich and complicated quantum theory which has no simple perturbative classical description. However, in some cases closed strings have a tractable semiclassical low-energy expansion involving a metric, $g_{\mu\nu}$, minimally coupled to a host of dynamical matter fields, $\{\Phi_i\}$, governed by an effective action of the form

$$I_{\rm GR} = \frac{1}{16\pi G_{\rm N}}\int {\rm d}^{d+1}x\,\sqrt{-g}\,(-2\Lambda + R + \cdots + \mathcal{L}_{\rm matter}(\Phi_i))\,. \tag{90}$$

Here $\mathcal{L}_{\rm matter}(\Phi_i)$ is the Lagrangian density for the matter fields and R is the Ricci curvature scalar, which is roughly two derivatives of the metric. $\sqrt{g}R$ plays the role of a kinetic term for the metric in GR. The $\cdots$ then represents an infinite tower of higher-derivative terms involving the metric, for example R^2, R^4, etc, which represent closed string corrections to GR. Since these higher-derivative terms are also higher-dimension irrelevant operators, they must appear with dimensionful couplings. In general, these couplings are controlled by two independent length scales: $\ell_{\rm p}$, the Planck length, which controls quantum corrections to the dynamics of classical strings; and $\ell_{\rm s}$, which controls classical 'stringy' corrections to the dynamics of point particles.

When expanding the theory around AdS with AdS-length L (and thus with Ricci curvature $R \sim L^{-2}$), we thus have two dimensionless parameters controlling the various possible corrections to GR: $\frac{\ell_{\rm p}}{L}$, which tells us whether quantum corrections are important; and $\frac{\ell_{\rm s}}{L}$, which tells us whether stringy corrections are important. Note that ensuring that the curvature is weak in Planck units, $\frac{\ell_{\rm p}}{L} \ll 1$, does ensure that the system can be treated classically, but does not tell us that the gravity must be exactly GR: the curvature of our AdS may still be large compared to the string scale, so higher-curvature corrections to the Lagrangian of GR may not be negligible. For pure classical GR to be a good approximation, two independent conditions must thus hold,

$$\frac{\ell_{\mathrm{p}}}{L} \ll 1 \quad \text{and} \quad \frac{\ell_{\mathrm{s}}}{L} \ll 1, \tag{91}$$

i.e. the curvature must be small compared both to the Planck scale and to the string scale.

4.2.2. Field theories in flat space. Our next job is to define the d-dimensional QFT we want to study. A canonical way to do so begins by specifying a list of local operators $\mathcal{O}_i$ labeled by their Lorentz structure, their charges q_i and their scaling dimensions Δ_i, all defined at some UV fixed point. To generate an RG flow of interest, we perturb away from this fixed point by turning on appropriate sources J_i. The basic observables of the theory are then correlation functions of products of local operators

$$\langle \mathcal{O}_1(x_1) \ldots \mathcal{O}_n(x_n) \rangle. \tag{92}$$

These can be conveniently encoded in terms of the quantum generating functional,

$$Z_{\mathrm{QFT}}[J_i(x)] = \langle \mathrm{e}^{\int \mathrm{d}x^d J_i(x) \mathcal{O}_i(x)} \rangle, \tag{93}$$

where the $J_i(x)$ represent a set of sources and couplings for the operators $\mathcal{O}_i$. Note that the scaling dimensions and tensor structure of the sources J_i are completely determined by those of $\mathcal{O}_i$. This allows us to express correlation functions as derivatives of the partition function

$$\langle \mathcal{O}_1(x_1) \ldots \mathcal{O}_n(x_n) \rangle = \frac{\delta^n \ln Z_{\mathrm{QFT}}[J(x)]}{\delta J_1(x_1) \ldots \delta J_n(x_n)} \bigg|_{J_i=0}, \tag{94}$$

where the restriction $|_{J_i=0}$ means that evaluate the final result with all remaining sources turned off. Solving the theory then requires one to compute the partition function, $Z_{\mathrm{QFT}}[J_i(x)]$.

Canonical example: $SU(N)$ Yang–Mills at large N. The classic example of a field theory with a well-understood and controlled holographic dual, to which we will appeal below, is a special version of the $SU(N)$ Yang–Mills gauge theory in flat 4D space–time called the $\mathcal{N} = 4$ theory. Here $\mathcal{N}$ refers to the amount of supersymmetry enjoyed by the theory; $\mathcal{N} = 4$ is the most one can have in 4D without including gravity. For our purposes, the role of supersymmetry is nothing more than a way of turning off quantum mechanics without totally trivializing the theory: with $\mathcal{N} = 4$ supersymmetry, some quantities do not receive quantum corrections beyond one loop, and so can be computed at weak coupling and reliably extrapolated to strong coupling. In particular, the b-function of the theory can be computed exactly and is identically zero for all values of the coupling[51].

The basic ingredients of the $\mathcal{N} = 4$ theory are a gauge group, $G = SU(N)$, a gauge field, $A_\mu(x)$, transforming in the adjoint of G, six scalars $\phi^I(x)$ also transforming in the adjoint of

---

[51] Our ability to compute various quantities at both strong and weak coupling in the $\mathcal{N} = 4$ theory was key to the original discovery of holographic duality, and is the reason why this example is both illuminating and canonical. However, that is the extent of the role of supersymmetry, to provide unusually simple and tractable examples. It is no more a necessary feature of holography than spherical symmetry is a necessary property of hydrogen. The maximally supersymmetric $\mathcal{N} = 4$ theory is an illuminating example to study, so we turn to it now, but we stress that the general structure of holographic duality thus revealed is more general, and does not depend on supersymmetry.

G and further enjoying a global $SO(6)$ symmetry, and a large set of fermions. We can safely neglect the fermions for now, as their only role is to ensure supersymmetry. The Lagrangian is

$$\mathcal{L}_{\mathrm{YM}} = -\frac{1}{4g^2}\left(\operatorname{Tr} F^2 + \operatorname{Tr}|D\phi_I|^2 + \cdots\right), \tag{95}$$

where g^2_{YM} is the gauge coupling, D_μ is the gauge-covariant derivative, $\cdots$ denotes fermionic terms we neglect and all fields are canonically normalized with an overall factor of $-\frac{1}{4g^2}$ in front of their kinetic terms.

What are the good operators of this theory? We may suppose that any good gauge-invariant combination of the fundamental fields is a perfectly good observable. However, since we have a global symmetry, it is useful to write this list in terms of $SO(6)$ irreducible representations. For example, the following three Lorentz scalars,

$$\mathcal{O} = \sum_I \operatorname{Tr}\phi_I^2, \quad \mathcal{O}_{IJ} = \operatorname{Tr}\phi_{(I}\phi_{J)}, \quad \mathcal{O}^\mu = \operatorname{Tr}\mathcal{J}^\mu, \tag{96}$$

transform as a scalar, a symmetric two-tensor and a vector, respectively, under the global $SO(6)$, with the latter corresponding to a charge current operator. In the weak-coupling limit, their scaling dimensions Δ_i are just their engineering dimensions which can be read off the Lagrangian[52].

Now, as pointed out by 't Hooft [341] (see also [342]), anytime you have a gauge theory of $N \times N$ matrices, there is a natural way to organize the Feynman diagrams in terms of underlying genus-g surfaces[53]. Explicitly, since every field comes with two gauge indices, and since all indices in a gauge invariant observable must be contracted, every Feynman diagram can be written as a *ribbon diagram* in which each propagator is a ribbon with a gauge running along each side of the ribbon. One can then show that any given Feynman diagram, with any number of underlying loops, can be drawn without the index lines crossing only on a genus-g Riemann surface, where the specific genus g depends on precisely how the index lines are contracted to make the given Feynman diagram. Explicitly, a diagram with genus $g = 0$ can be drawn on a piece of paper with no lines crossing, a diagram of genus $g = 1$ must be drawn on a torus to have no lines cross, etc.

A remarkable result is that if we reorganize the loop expansion in g^2_{YM} as a double expansion in N and $\lambda = g^2_{\mathrm{YM}}N$, then every observable can be expressed as a power series expansion in N where the powers of N for a given diagram is determined only by the genus g on which that diagram can be drawn with no crossings. For example, the free energy $\mathcal{F} = \ln Z$ takes the form

$$\mathcal{F} = N^2 f_0(\lambda) + f_1(\lambda) + \frac{1}{N^2} f_2(\lambda) + \cdots = N^2 \sum_g \frac{f_g(\lambda)}{N^{2g}}, \tag{97}$$

where $f_g(\lambda)$ is the sum of all diagrams arising at genus g and is a function of only λ, independent of N. A similar result is obtained for every possible observable, with the only difference being the overall factor of N,

---

[52] Here, *engineering dimension* refers to the naive dimension of an operator as determined by dimensional analysis on the kinetic terms in the classical Lagrangian, while the true scaling dimension is determined by the scale dependence of correlation functions containing the operator in the full interacting quantum theory. For a review of engineering and general scaling dimensions, see, e.g., the text by Sachdev [340].

[53] Here the *genus* refers to the number of handles on a two-dimensional surface: a sphere or plane has genus zero, a donut has genus 1, a donut with two holes has genus 2, etc.

$$\mathcal{A} = N^m \left(\mathcal{A}_0(\lambda) + \frac{1}{N^2}\mathcal{A}_1(\lambda) + \frac{1}{N^4}\mathcal{A}_2(\lambda) + \cdots \right) = N^m \sum_g \frac{\mathcal{A}_g(\lambda)}{N^{2g}}. \tag{98}$$

This double expansion implies a remarkable simplification at large N: if $N \gg 1$, the only term that matters in the genus expansion is the leading planar term; any diagram that cannot be drawn on a piece of paper without crossings, while possibly large in numerical value, is dwarfed by the much larger terms coming from the planar diagrams.

Another way of expressing this simplification at large N is on terms of *large N factorization*. Consider a set of single-trace operators[54] of the form $\mathcal{O}_i = \mathrm{Tr}\,(\ldots)$. At large N, correlation functions of single trace operators factorize

$$\langle \mathcal{O}_1(x_1)\mathcal{O}_2(x_2)\rangle = \langle \mathcal{O}_1(x_1)\rangle\langle \mathcal{O}_2(x_2)\rangle + O\left(\frac{1}{N^2}\right). \tag{99}$$

This follows because the disconnected parts of the diagrams necessarily involve more closed index loops than connected parts and thus additional powers of N. This does not imply that the large N theory is free, for while these correlators do factorize, the individual one-point functions remain non-trivial. For example, the anomalous dimensions of all operators are controlled by the 't Hooft coupling, λ.

We thus see that large N gauge theories have two natural coupling constants: $\frac{1}{N}$, which controls the genus expansion and factorization; and λ, which controls the perturbative corrections to each term in the genus expansion and the anomalous dimensions of large N factorized operators. Suggestively, this is precisely the structure of closed string perturbation theory, where every amplitude is expressed as a sum of contributions from various genera, with the loop-counting parameter given by the string coupling g_{s} and the amplitude of a given genus determined by an auxiliary QFT whose interactions are determined by the string tension α' via the relation

$$\mathcal{A} = g_{\mathrm{s}}^n \left(\mathcal{A}_0(\alpha') + g_{\mathrm{s}}^2\mathcal{A}_1(\alpha') + g_{\mathrm{s}}^4\mathcal{A}_2(\alpha') + \cdots \right) = g_{\mathrm{s}}^n \sum_g g_{\mathrm{s}}^{2g}\mathcal{A}_g(\alpha'). \tag{100}$$

This analogy led various people to speculate that large N gauge theories should be captured by some theory of closed strings. This turned out to be wrong, but only just barely: at large N, these theories *are* dual to closed string theories, but the closed strings live in one higher dimension!

As an aside, while not every theory is a gauge theory of $N \times N$ matrices, it is useful to keep in mind this example when thinking about more general examples. What N^2 is really measuring is the number of degrees of freedom per unit volume, or per lattice site if we put the theory on a lattice. Similarly, λ is measuring the strength of the dominant interactions when the number of degrees of freedom grows large. Many theories that are not simple gauge theories, nonetheless, have a useful notion of N and λ. Exactly when one obtains such a controlled double expansion remains, in general, a key open question in holography.

4.2.3. The holographic dictionary. The basic claim of holography is that every d-dimensional QFT defined as above can be exactly reorganized into a $(d+1)$-dimensional quantum theory of gravity and matter propagating in AdS_{d+1}. The precise relationship between these two theories is known as the *holographic dictionary*, a sketch of which is outlined in table 1. The rest of this

---

[54] The meaning of single-trace operators is slightly subtle, but can roughly be understood as the set of operators which have local, extensive classical limits.

Table 1. Elements of the holographic dictionary.

Boundary QFT				Bulk gravity
Operator	$\mathcal{O}(x)$	$\longleftrightarrow$	$\Phi(x,r)$	Field
Spin	$s_{\mathcal{O}}$	$\longleftrightarrow$	s_{Φ}	Spin
Global charge	$q_{\mathcal{O}}$	$\longleftrightarrow$	q_{Φ}	Gauge charge
Scaling dimension	$\Delta_{\mathcal{O}}$	$\longleftrightarrow$	m_{Φ}	Mass
Source	$J(x)$	$\longleftrightarrow$	$\Phi(x,r)\rvert_{\partial}$	Boundary value (BV)
Expectation value	$\langle\mathcal{O}(x)\rangle$	$\longleftrightarrow$	$\Pi_{\Phi}(x,r)\rvert_{\partial}$	BV of radial momentum
Global symmetry group	G	$\longleftrightarrow$	G	Gauge symmetry group
Source for global current	$\mathcal{A}_{\mu}(x)$	$\longleftrightarrow$	$A_{\mu}(x,r)\rvert_{\partial}$	BV of gauge field
Expectation of current	$\langle\mathcal{J}^{\mu}(x)\rangle$	$\longleftrightarrow$	$\Pi^{\mu}_{A}(x,r)\rvert_{\partial}$	BV of momentum
Stress tensor	$T^{\mu\nu}(x)$	$\longleftrightarrow$	$g_{\mu\nu}(x,r)$	Space–time metric
Source for stress energy	$h_{\mu\nu}(x)$	$\longleftrightarrow$	$g_{\mu\nu}(x,r)\rvert_{\partial}$	BV of metric
Expected stress energy	$\langle T^{\mu\nu}(x)\rangle$	$\longleftrightarrow$	$\Pi^{\mu\nu}_{g}(x,r)\rvert_{\partial}$	BV of momentum
No. of degrees of freedom per space–time point	N^2	$\longleftrightarrow$	$\left(\frac{L}{\ell_{\mathrm{p}}}\right)^{d-1}$	Radius of curvature In Planck units
Characteristic strength of interactions	λ	$\longleftrightarrow$	$\left(\frac{L}{\ell_{\mathrm{s}}}\right)^{d}$	Radius of curvature In string units
QFT partition function with sources $J_i(x)$	$Z_{\mathrm{QFT}_d}[J_i]$	$\longleftrightarrow$	$Z_{\mathrm{QG}_{d+1}}[\Phi_i[J_i]]$	QG partition function in AdS w/ $\Phi_i\rvert_{\partial}=J_i$
QFT partition function at strong coupling	$Z^{\lambda,N\gg 1}_{\mathrm{QFT}_d}[J_i]$	$\longleftrightarrow$	$\mathrm{e}^{-I_{\mathrm{GR}_{d+1}}[\Phi[J_i]]}$	Classical GR action in AdS w/ $\Phi_i\rvert_{\partial}=J_i$
QFT n-point functions at strong coupling	$\langle\mathcal{O}_1(x_1)\ldots\mathcal{O}_n(x_n)\rangle$	$\longleftrightarrow$	$\left.\frac{\delta^n I_{\mathrm{GR}_{d+1}}[\Phi[J_i]]}{\delta J_1(x_1)\ldots\delta J_n(x_n)}\right\rvert_{J_i=0}$	Classical derivatives of the on-shell classical gravitational action
Thermodynamic state		$\longleftrightarrow$		Black hole
temperature	T	$\longleftrightarrow$	T_{H}	Hawking temperature ∼ mass
Chemical potential	μ	$\longleftrightarrow$	Q	Charge of black hole
Free energy	F	$\longleftrightarrow$	$I_{\mathrm{GR}}\rvert_{(\mathrm{on-shell})}$	On-shell bulk action
Entropy	S	$\longleftrightarrow$	A_{H}	Area of horizon

subsection will develop various entries in the dictionary, including a few canonical examples of computations performed via the holographic dictionary.

Operators and fields. The first entry in the holographic dictionary relates the operators of the QFT to the matter fields in the bulk gravity: with every local operator $\mathcal{O}_i(x)$ with quantum numbers (or *charges*) q_i and scaling dimension Δ_i in our QFT is associated a field $\Phi_i(x,r)$ in AdS which carries the Lorentz structure and quantum numbers required to act as a coupling for $\mathcal{O}_i(x)$ in the Hamiltonian. Note that the dimension of the operator, Δ_i, is determined by the mass of the bulk field, as we saw above in (81).[55] This means that the mass and other couplings of the bulk gravity theory generally do not correspond to interactions of the boundary theory in any simple way, but rather determine which boundary QFT we are studying.

---

[55] More generally, Δ depends on the Lorentz structure of the field Φ_i and on the full set of its bulk interactions. As a simple example, the relation for a scalar operator corresponding to a scalar field in the bulk is $m^2L^2=\Delta(\Delta-d)$.

An equivalence of partition functions. Given this map, the central claim of holographic duality [302–304] can be succinctly expressed as

$$Z_{\mathrm{QFT}}[J_i] = Z_{\mathrm{QG}}[\Phi[J_i]], \tag{101}$$

where $Z_{\mathrm{QFT}}[J_i]$ is the partition function of the QFT as a function of the sources J_i for each operator $\mathcal{O}_i$, while $Z_{\mathrm{QG}}[\Phi[J_i]]$ is the quantum partition function of the gravitational theory described in section 4.2.1 computed in an AdS space–time background[56]. As discussed, the quantum gravity partition function must be evaluated on field configurations Φ_i which asymptote at the boundary to the sources J_i of the QFT, hence the notation $\Phi[J]$. As in the heuristic derivation above, the bulk fields are precisely the coupling constants of the QFT promoted to dynamical fields on the full RG-extended space–time in which the RG scale becomes a physical coordinate.

An important ingredient in this recipe is a proper definition of the boundary value of the bulk field. As we saw above for a bulk scalar Φ of mass m, dual to a scalar operator $\mathcal{O}$ of dimension Δ given by $\Delta(d-\Delta) = m^2L^2$, the boundary value of the bulk field is in general divergent. To specify the boundary condition, we must rescale the bulk field by an appropriate wavefunction renormalization which picks off the leading divergence of the bulk field near the boundary

$$J(x) \equiv \phi_{d-\Delta}(x) = \lim_{r\to 0} r^{\Delta-d}\Phi(x,r). \tag{102}$$

We will return to this renormalization when we discuss the computation of one- and two-point functions.

Strong coupling and perturbative gravity. In general, both sides of equation (101) are complicated and computationally intractable objects. However, when one side or the other can be evaluated semi-classically, we can use equation (101) to generate a controlled strong-coupling expansion via the dual description. Precisely whether, and when, such a limit is obtained is a delicate question that must be studied in detail on a case-by-case basis. To get a sense of the general structure, let us consider the canonical example, the large N limit of the 4D $\mathcal{N}=4$ gauge theory with gauge group $G = SU(N)$ and 't Hooft coupling λ. The dual theory turns out to be a closed string theory on AdS$_5$ with AdS-radius L. The holographic dictionary then tells us that

$$N^2 = \left(\frac{L}{\ell_{\mathrm{p}}}\right)^{d-1}, \quad \lambda = \left(\frac{L}{\ell_{\mathrm{s}}}\right)^{d}, \tag{103}$$

where ℓ_{p} is the Planck length determining the scale at which quantum effects occur in the bulk, and ℓ_{s} is the string length controlling the scale of stringy higher-curvature corrections to the bulk gravitational action, as discussed in section 4.2.1. The first relation in equation (103) thus tells us that the bulk description becomes classical ($L \gg \ell_{\mathrm{p}}$) only in the large-N limit[57] and that the gravitational description is deeply quantum mechanical when N is small. In other words, the loop counting parameter in the bulk is $\frac{1}{N}$.

---

[56] In general, Z_{QG} means the partition function of closed string theory expanded around the AdS geometry. In practice, we can only compute this explicitly in special cases and then only in regimes in which the string coupling and curvature are sufficiently weak. Thus, we may use this duality either to deduce the lhs of equation (101) when the gravity is classical and weakly curved or to define the rhs of equation (101) when the QFT is computable.

[57] This provides an out from the Weinberg–Witten argument, which assumes that there are a finite number of local degrees of freedom.

With an eye toward the general class of holographic theories, it is worthwhile unpacking equation (103) in more detail. In the case at hand, i.e. $SU(N)$ SYM, N^2 measures the number of degrees of freedom per point in the QFT. In particular, all extensive thermodynamic quantities (entropy, energy, etc) must scale as N^2. In this light, what equation (103) shows is that the quantum corrections in the bulk are negligible when the number of degrees of freedom per point in the dual QFT is large, and vice versa. Indeed, in a host of controlled models in string theory, the appropriate version of the first equation in equation (103) takes the form, $N^b = \left(L/\ell_{\mathrm{p}}\right)^{d-1}$, where N is a conserved charge and b is some real parameter [343]. Similarly, it is not always possible to interpret λ as a 't Hooft parameter in the dual QFT. More generally, the role of λ is played by the typical scale of anomalous dimensions in the QFT [333]. We can understand this as follows: if the QFT is weakly interacting, the effects of renormalization will generally be weak and most observables will be well approximated by their classical cousins, i.e. quantum anomalous dimensions will be small; if interactions are strong, quantum effects will generally drive anomalous dimensions to be large.

While $N \gg 1$ ensures that the Planck scale is negligible, so that the gravitational interactions can be treated with classical effective field theory, it does not tell us that the system is well approximated by pure GR. Indeed, in principle, there will be a host of higher-curvature corrections to the classical GR action suppressed by powers of $\frac{\ell_{\mathrm{s}}}{L}$. To ensure that these corrections are in fact negligible, i.e. that $L \gg \ell_{\mathrm{s}}$, we must furthermore take the 't Hooft coupling to be large, $\lambda \gg 1$. Conversely, when the QFT is weakly coupled, $\lambda \ll 1$, the bulk geometry is sufficiently strongly curved that GR is swamped by higher curvature corrections to the stringy effective action.

The main lesson here is that when the QFT is perturbative, the bulk gravity is out of control, and that when the QFT is strongly interacting and has a large entropy density, the dual gravitational description is simple semi-classical gravity. To see why such a weakly curved limit is interesting, let us consider our canonical example at large N, $N \gg 1$, and with strong 't Hooft coupling, $\lambda \gg 1$. Let us also analytically continue to imaginary time, $\omega = \mathrm{i}\omega_{\mathrm{E}}$ with $\omega_{\mathrm{E}} > 0$. In this limit, the gravity sector reduces to classical Euclidean GR. In terms of our fundamental relation, the rhs can thus be expressed as a sum over saddles. Focusing on the dominant saddle, we obtain

$$Z_{\mathrm{QFT}}[J] \simeq \mathrm{e}^{-I_{\mathrm{GR}}[\Phi[J]]}, \tag{104}$$

where

$$I_{\mathrm{GR}}[\Phi[J]] = \frac{1}{16\pi G_{\mathrm{N}}} \int \mathrm{d}^{d+1}x \sqrt{-g}\,(R - 2\Lambda + \mathcal{L}_m(\Phi)) \tag{105}$$

is the classical gravitational action expanded about the dominant classical saddle, subject again to the condition that the normalizable modes of all matter fields are determined by the sources of the dual QFT. Equation (105) thus defines a simple classical functional of the sources from which we can compute quantum correlation functions of the dual QFT via the above expressions as

$$\langle \mathcal{O}_1(x_1) \ldots \mathcal{O}_n(x_n) \rangle = \left. \frac{\delta^n I_{\mathrm{GR}}[\Phi[J_i]]}{\delta J_1(x_1) \ldots \delta J_n(x_n)} \right|_{J_i=0}. \tag{106}$$

The upshot is that, when such a dual semi-classical limit exists and is reliable, we can compute the quantum correlation functions of our QFT by finding solutions to the classical field equations of the dual gravitational system, evaluating the action on-shell as a functional of the boundary values of the fields, and taking appropriate derivatives.

Example: one- and two-point functions for a scalar operator. As an example of this machinery at work, let us compute the one- and two-point functions for a boundary operator $\mathcal{O}$ dual to a free bulk scalar field ϕ. For simplicity, we will work in Euclidean momentum space throughout.

For the one-point function, the basic form of the computation is straightforward. Intuitively, our prescription (106) tells us that the one-point function corresponds to a variation of the classical gravitational action with respect to the boundary value of the bulk field. But a variation of the action with respect to the boundary value of the field is precisely the conjugate momentum, Π, of the field Φ in the direction normal to the boundary,

$$\Pi = -\sqrt{-g}\,g^{rr}\partial_r\Phi. \tag{107}$$

Thus we should expect to find, with $k \equiv \left(\omega_{\mathrm{E}}, \vec{k}\right)$,

$$\langle\mathcal{O}(k)\rangle = \frac{\delta I_{\mathrm{GR}}[\Phi[J]]}{\delta J(k)} \sim \lim_{r\to 0}\Pi(k,r). \tag{108}$$

This is almost correct. The trouble is that the on-shell classical action, $I_{\mathrm{GR}}[\phi]$, and the bulk field, Φ, both generically diverge as we approach the boundary, cf equation (81). To absorb these divergences we must (a) regulate all quantities by evaluating them not at the boundary at $r = 0$ but at $r = \epsilon \ll 1$, (b) rescale the bulk field by an overall wavefunction normalization as in (88) and (c) add boundary counterterms to the action evaluated at $r = \epsilon$ so that the renormalized on-shell action remains finite as we remove the cutoff. In retrospect, this should not be surprising: the radial coordinate r is playing the role of a lattice cutoff, so we must work with properly renormalized quantities to avoid confusion. This *holographic renormalization* [335, 344–348] is simply the bulk realization of the UV renormalization of the boundary QFT.

In fact, we have already seen this effect above. Recall that the relationship between the amplitude of the non-normalizable mode $\phi_{d-\Delta}$ and the boundary value of the bulk field Φ also required a subtle wavefunction renormalization (88), so that the source is given by

$$J(k) = \phi_{d-\Delta}(k) = \lim_{r\to 0} r^{\Delta-d}\,\Phi(k,r). \tag{109}$$

Upon performing a similar renormalization of the bulk Euclidean action, one finds that the correct expression for the one-point function in terms of the renormalized radial momentum is[58]

$$\langle\mathcal{O}(k)\rangle = \lim_{r\to 0} r^{d-\Delta}\,\Pi(k,r). \tag{110}$$

Together with (107) and (81), this gives

$$\langle\mathcal{O}(k)\rangle = \frac{2\Delta - d}{L}\phi_{\Delta}(k), \tag{111}$$

where ϕ_Δ is the coefficient of the sub-leading term in (81). Thus, just as the non-normalizable mode of the bulk field $\phi_{d-\Delta}$ determines the source, J, for the boundary operator $\mathcal{O}_\Delta$, the subleading normalizable term ϕ_Δ determines the response, $\langle\mathcal{O}\rangle$.

---

[58] For a detailed discussion of this calculation, see appendix C of [184]. The basic strategy can be found in [237, 339, 349–351].

This makes computing the linear-response Green functions easier. In linear response theory, an infinitesimal source $J(x)$ generates a response which is linearly proportional to the source, with the ratio defining the linear-response Green function

$$G_{\mathrm{E}}(k) \equiv \frac{\langle \mathcal{O}(k)\rangle}{J(k)} = \lim_{r\to 0} r^{2(\Delta-d)} \frac{\Pi(k,r)}{\Phi(k,r)}. \tag{112}$$

Using our above expressions for the source, J, and response, $\langle \mathcal{O}\rangle$, we thus find that

$$G_{\mathrm{E}}(k) = \frac{2\Delta-d}{L} \frac{\phi_\Delta(k)}{\phi_{d-\Delta}(k)}. \tag{113}$$

Note that we could also have derived this result by taking two derivatives of the partition function.

Computing the one- and two-point functions thus follows from determining ϕ_Δ in terms of $\phi_{d-\Delta}$. As we have seen, in AdS these two modes are not independent: requiring the bulk solution to be regular at the Euclidean horizon imposes a relation which we can use to determine the response ϕ_Δ in terms of the source $\phi_{d-\Delta}$. In practice, finding the precise relation involves solving the bulk equation of motion subject to the boundary conditions, i.e. to solving a set of second-order elliptic partial differential equations. For simple examples this can be done analytically; more generally, one is forced to use some form of matched asymptotic expansion or numerical integration to determine ϕ_Δ in terms of $\phi_{d-\Delta}$.

The key point is that the computation of the quantum correlation function of the boundary QFT has been reduced to solving a set of classical partial differential equations with fixed boundary conditions.

An example where this can be done analytically is our example of a massive scalar in pure AdS, for which we have in fact already determined this relation in momentum space, equation (87). Plugging this in and again working in momentum space gives, for the Euclidean Green function,

$$G_{\mathrm{E}}(k) = \frac{2\Delta-d}{L} \frac{\Gamma\left(\frac{d}{2}-\Delta\right)}{2^{\Delta-\frac{d}{2}}\Gamma\left(\Delta-\frac{d}{2}\right)} \left(k^2\right)^{\Delta-\frac{d}{2}}. \tag{114}$$

This is exactly the form for the two-point function of a scalar operator with scaling dimension Δ in a CFT.

Real-time response and retarded Green functions. So far we have focused on Euclidean correlation functions. In fact, these holographic techniques can be readily extended to intrinsically Lorentzian computations needed for real-time response.

The simplest thing to do would be to simply compute the Euclidean Green function and analytically continue. In practice, this is very often not tractable. For example, we will sometimes be interested in particles moving on light-like trajectories which are not easily described in Euclidean continuation. Instead, one can construct an intrinsically real-time holographic prescription, as was first proposed by Son and Starinets [352] by essentially analytically continuing the Euclidean prescription. Their results have since been justified by developing a full holographic Schwinger–Keldysh formalism [353–355]. Here we will forego the formalities. We assume that such a justification can be made and proceed to the prescription.

The prescription requires repeating our Euclidean prescription step by step in the real-time Lorentzian geometry. The key difference is that we must choose an appropriate boundary condition at the horizon. The appropriate choice depends on which Green function we wish to compute[59]. Intuitively, one might expect that the retarded Green function, G_{R}, should correspond to imposing causal in-falling boundary conditions at the horizon, while the advanced Green function, G_{A}, should involve acausal out-going boundary conditions. This turns out to be precisely correct. Once we have a solution satisfying the appropriate boundary conditions, we again compute the properly renormalized on-shell Lorentzian action, identify source and response from the asymptotic behavior of the solution near the boundary and compute the Green function via a linear response. The final result takes an analogous form to equation (113),

$$G_{\mathrm{R}}(\omega,\vec{k}) = \lim_{r\to 0} r^{2(\Delta-d)}\frac{\Pi_{r,\mathrm{in}}(\omega,\vec{k},r)}{\Phi_{\mathrm{in}}(\omega,\vec{k},r)} = \frac{2\Delta-d}{L}\,\frac{\phi_{\Delta,\mathrm{in}}(\omega,\vec{k})}{\phi_{d-\Delta,\mathrm{in}}(\omega,\vec{k})}, \tag{115}$$

where ω is the Lorentzian frequency and ϕ_{Δ}, $\phi_{d-\Delta}$ are found by solving the linearized Lorentzian equations in the bulk with in-falling boundary conditions at the horizon. The advanced Green function is computed analogously to out-going boundary conditions enforced at the horizon.

Cautionary notes on the semi–classical limit. While it is true that we have simplified our job by focusing on limits where the gravitational partition function Z_{GR} can be computed by saddle point approximation, this is not the same as expanding Z_{QFT} semi-classically; in the present saddle, the classical description is of classical gravity in one higher dimension. Thus, rather than studying a semi-classical limit of the QFT governed by quasiparticle perturbation theory, we are focusing on a very different emergent limit in which quasiparticles are replaced by geometry.

Note, too, a practical consistency condition: it cannot be the case that both Z_{QFT} and Z_{QG} are simultaneously perturbative, for if that were the case we would learn that every classical field theory can be reorganized into classical gravity—but this is explicitly forbidden by the Weinberg–Witten theorem. It is thus clear that this holographic duality had to be a strong–weak duality, with one side becoming strongly quantum mechanical whenever the dual becomes semi-classical.

Finally, we could just as easily have inverted our logic and studied a limit in which the QFT partition function became semi-classical. While this is unlikely to teach us much about the QFT, since perturbation theory is already well understood, it may teach us a great deal about the strongly quantum gravitational dual. In particular, this approach provides the only non-perturbative definition of 4D quantum gravity currently known, and can be used to quickly and convincingly argue for the preservation of unitarity by the formation and evaporation of black holes.

Thermodynamics encoded by black holes. The final entry in the holographic dictionary sketched in table 1 relates the thermodynamic state and ensemble in which we place our QFT to the precise geometry in which we study the bulk gravity [356]. When the QFT is exactly conformal, the bulk geometry is pure AdS. Recalling that the UV physics of the boundary is associated with the asymptotic near-boundary region of AdS. Deforming the theory away from this conformal UV fixed point to generate a QFT with non-trivial RG flow corresponds to studying a geometry which is asymptotically AdS near the boundary but flows away from pure AdS as we run into the bulk. For example, turning on a finite temperature, T, and chemical potentials, μ_i, in this CFT corresponds to studying asymptotically AdS black hole space–times with Hawking temperature $T_{\mathrm{H}} = T$ and charges $Q_i = \mu_i$. More generally, the thermodynamic data of the QFT are entirely encoded in the thermodynamics of the black hole in the dual geometry (see table 1 for the mapping between the thermodynamic parameters of the QFT and the quantum numbers of the dual black hole).

---

[59] Since we are only studying infinitesimal sources and responses, there is no reason to restrict to finite action modes, i.e. to solutions that are regular at the horizon. This is a simplification afforded by linear response. More generally, when considering non-trivial bulk field configurations, we should again impose regularity and in-falling boundary conditions.

The classic example of the power of these thermodynamic relations is the computation of the thermodynamic properties of a large-N 4D Yang–Mills gauge theory at the strong 't Hooft coupling. The holographic dual of this system is a planar AdS black brane, as discussed in section 4.2.1. In the large-N, large-λ limit, the gravitational description reduces to GR and the black brane entropy becomes

$$S_{\mathrm{BH}} = \frac{A_{\mathrm{H}}}{4G_{\mathrm{N}}}, \tag{116}$$

where A_{H} is the area of the horizon. Since the system is translationally invariant, this diverges, so it is convenient to consider instead the entropy per unit area in the field theory

$$s = \frac{S_{\mathrm{BH}}}{A_\partial} = \frac{A_{\mathrm{H}}}{4G_{\mathrm{N}}A_\partial}, \tag{117}$$

where A_∂ is the area of the AdS boundary. Plugging in factors, this gives a strong-coupling computation of the entropy density and energy density as

$$s_{\mathrm{strong}} = \frac{\pi^2}{2}N^2T^3, \quad \epsilon_{\mathrm{strong}} = \frac{3\pi^2}{8}N^2T^4. \tag{118}$$

These strong coupling results are extremely close to the results in the free theory as computed in perturbation theory:

$$s_{\mathrm{free}} = \frac{2\pi^2}{3}N^2T^3, \quad \epsilon_{\mathrm{free}} = \frac{\pi^2}{2}N^2T^4. \tag{119}$$

Remarkably, despite running from weak to strong coupling, $\lambda = 0$ to $\lambda \to \infty$, all that happened to the thermodynamics is a mild renormalization by a factor of $\frac{3}{4}$:

$$\frac{s_{\mathrm{strong}}}{s_{\mathrm{free}}} = \frac{\epsilon_{\mathrm{strong}}}{\epsilon_{\mathrm{free}}} = \frac{3}{4}. \tag{120}$$

Studying the full λ-dependence of this ratio reveals that the energy and entropy densities flow smoothly and monotonically as we run from weak to strong coupling. Analogous results are obtained for a wide variety of gauge groups, matter contents and even dimensions, with the ratio $\frac{s_{\mathrm{strong}}}{s_{\mathrm{free}}}$ typically within 10% or so of the $\mathcal{N} = 4$ value, $\frac{3}{4}$.

This holographic result already tells us an important fact: thermodynamic quantities such as the free energy and the entropy density are not good probes of the strength of interaction of a quantum liquid. Similar results apply to other thermodynamic quantities, for which the strong/weak ratio is typically $\mathcal{O}(1)$ as well[60].

4.2.4. A bit of history. Holography is far more general than the refined context in which it was discovered. In particular, using holography as a tool does not require supersymmetry, string theory, d-branes or five-dimensional (5D) AdS, although all those ideas were central to its discovery and can provide useful guidance to the applied holographer. It is useful to address this directly.

---

[60] This can, of course, change if there is a quantum phase transition as parameters in the system are varied.

In its modern form, holography was first discovered in string theory while studying very special toy systems with as much symmetry as possible, in particular, maximally supersymmetric 4D $\mathcal{N} = 4$ $SU(N)$ Yang–Mills theory at large N. The reason supersymmetry was important in the discovery of holography is that SUSY turns off many quantum effects so that a semi-classical analysis can give, for at least some special quantities, exact results. This allowed the computation of a host of correlation functions at weak coupling which could then be reliably extrapolated to strong coupling. For example, thermodynamic quantities such as the free energy and entropy density, as well as transport properties such as conductivities, viscosities, susceptibilities and sound mode dispersion relations can all be calculated analytically, something one would never expect in a less constrained theory such as QCD. Of particular interest were *gray-body factors* of supersymmetric black holes, corresponding to the probability with which a particle would be absorbed by a black hole. These could be computed at weak coupling via classical gravitational perturbation theory and then reliably extrapolated to strong coupling, giving a precise computation of a small set of very special absorption cross-sections of deeply quantum mechanical black holes. Thus, in the late 1990s, a number of researchers intensively computed strong-coupling results for the string theorist's spherical cow.

The results were astonishing. As first pointed out by Juan Maldacena in the fall of 1997 [302], while the weakly coupled $\mathcal{N} = 4$ field theory behaved mostly as familiar 4D gauge theories, at strong coupling the correlation functions of the theory reproduced, rather miraculously, classical scattering off black holes in one higher dimension! More precisely, Maldacena [302] (and, shortly thereafter, Gubser *et al* [303] and Witten [304]) gave a precise relationship between the 4D $\mathcal{N} = 4$ theory and a particular theory of gravity in 5D AdS. Crucially, the 't Hooft coupling λ and the number of colors N of the 4D supersymmetric Yang–Mills theory were related to the AdS curvature in units of the Planck and string lengths precisely as reviewed above in equation (103). Thus, when the 4D theory is weakly coupled, $\lambda \ll 1$, the 5D theory is strongly curved and cannot be treated as a theory of perturbative gravity. Conversely, when the 4D theory is strongly coupled, $\lambda \gg 1$, and at large N, the 5D theory is weakly curved, so while the system cannot be treated as a weakly interacting gauge theory in 4D, it can be treated as a weakly coupled theory of gravity in 5D.

The point is this: while the original discovery of holographic duality required supersymmetry and conformal invariance, the duality itself does not depend on supersymmetry or conformality. Indeed, we now have a long list of examples in which each of these constraints is weakened or removed, with every consistency check passed. To date there is no example of a violation of a sharp holographic duality. From the stringy point of view, holographic duality is simply a true, if as yet unproven, fact about quantum field theories and quantum gravity.

4.3. Applied holography

Applied holography is the application of holographic dualities to construct computationally tractable toy models of behavior which has proven difficult to model with traditional tools. Key to this endeavor is the fact that various features of the physics which appear universal, or even trivial, on one side of the duality may be significant, or simply obscure, in the dual. Most current work in applied holography has used universal features of gravitational models, and in particular, the universality of black hole horizon physics, to make new discoveries about QFTs, although using universal features of weakly coupled field theories to learn about extreme phases of gravity is of interest as well. Given the aim of this focus issue, we will emphasize the first direction in what follows.

The simplest application of holography involves studying QFTs for which the precise gravitational dual is explicitly known. In such cases we can use the perturbative QFT to compute any quantities of interest at weak coupling, $\lambda \to 0$, and then use the weakly coupled gravitational dual to compute the corresponding quantities at strong coupling, $\lambda \to \infty$. This strategy has very limited applications because the precise duals of most QFTs are not yet known. At the moment we only know the precise duality for a special set of non-generic quantum field theories, although they are fairly close cousins of the theories studied in particle physics. For example, we know the precise gravitational dual of the $\mathcal{N} = 4$ theory discussed in detail above and have used it to learn a great deal about the structure and strong-coupling physics of this intricate and exotic theory (which we emphasize does not appear to be realized anywhere in nature). By contrast, we do not know the dual of the Hubbard model.

This is not to say progress has not been made, but just that it is difficult. An example of this point is the recent conjecture of a duality between the 1D Ising model and pure Einstein gravity in AdS_3. For this relation to hold the gravitational side must be treated not only quantum mechanically, but fully non-perturbatively, summing over an infinite number of topologically inequivalent locally AdS_3 space–times, including a series of quantum corrections around each classical saddle, something that can only be done in closed form in the absolutely simplest possible gravitational systems, namely pure Einstein gravity in $2+1$ dimensions. This conjecture underscores the difficulty of constructing the precise holographic dual of a garden-variety QFT.

If the precise dual is not known, a second approach to applied holography is to make use of a known exact duality with a class of QFTs that are structurally similar to the QFT of interest. Universal properties of these sibling QFTs can then be inferred to apply to the original QFT, too. An example of this strategy is the construction by Kachru *et al* of holographic dimer models [357, 358]. In this construction, the degrees of freedom of interest to the dimer problem form a small subset of the full set of holographic degrees of freedom. At low energy, this subset forms a relatively generic dimer model. By exploiting the holographic description, much can be derived about the low-energy physics, including a novel phase transition between Fermi-liquid and non-Fermi liquid behaviors. The most powerful example of this approach so far has been the study of strongly coupled quark–gluon liquids in large N toy models of QCD, for which a menagerie of models have been considered, each emphasizing and reproducing various features of the observed phenomena. As we shall see in detail in sections 4.3.1 and 4.3.2, this approach has been used to motivate the anomalously small viscosity of the RHIC fireball, to explain why the rapid thermalization of the RHIC droplet is not so surprising, and to estimate the jet quenching parameter.

However, perhaps the most interesting use of holography is to provide an entirely new way of defining consistent QFTs. Traditionally, QFTs in dimensions greater than 2 are defined by specifying some Lagrangian or Hamiltonian governing the interactions of a set of well-defined quasiparticles. However, many systems, including some of the most theoretically and experimentally interesting systems, manifestly do not have any well-defined quasiparticles upon which to base a standard QFT. From this point of view, holography looks interesting because it provides a recipe for computing quantum amplitudes that manifestly satisfy the conditions of a good QFT (local, causal, etc) but which is defined without any appeal to a quasiparticle picture or any sort of conventional perturbation theory. Holography thus provides an entirely new way of constructing consistent models of many-body physics, replacing quasiparticles with geometry as the central organizing principle.

4.3.1. Viscous hydrodynamics from gravity. Suppose we are given a QFT at finite temperature and density and want to know whether it is weakly or strongly interacting. A natural observable to query is the shear viscosity, η, which measures the efficiency of momentum transport across a velocity gradient. Explicitly, the shear viscosity η of a fluid determines the frictional force F induced on a plate of area A moving at a velocity v a distance L over a fixed surface,

$$F = \eta \frac{Av}{L}. \tag{121}$$

This relation follows from the hydrodynamic stress tensor given in equation (44). As discussed in section 3.2 shear viscosity has dimensions $\hbar n$, where n is a density. The natural thermodynamic quantity to which one can compare the viscosity is thus the entropy density, s, with dimensions of $k_{\mathrm{B}} n$. η/s thus gives a simple dimensionless measure of the efficiency of transverse momentum transport.

From kinetic theory, the transport of momentum proceeds by scattering quasiparticles between layers of the fluid, and is thus proportional to the mean free path between scattering events and fluid quasiparticles. Since the mean free path decreases as the scattering cross section increases, the shear viscosity should be inversely related and thus strongly sensitive to the strength of interactions. However, if we drive up the strength of interaction until the mean free path becomes of the order of the Compton wavelength of the would-be quasiparticle, the kinetic-theory relation between interaction and viscosity should no longer apply. Meanwhile, in this post-quasiparticle setting, the mean free path is much smaller than any typical velocity gradient, so a hydrodynamic description should set in. Thus, as the interactions become strong, we expect the shear viscosity to become small, and furthermore expect coupling dependence of the viscosity to qualitatively change, since the kinetic behavior must eventually cease. But what values should η/s take and how will it vary as the coupling grows large? *A priori*, the answers to these questions would appear to be strongly system dependent.

To address this question, let us compute η/s holographically[61]. This will give us a controlled strong-coupling estimate of η/s. However, before proceeding, let us pause to consider, from the holographic perspective, why there must be any viscosity in the first place. Recall that the holographic dual of a strongly coupled relativistic large N gauge theory at finite temperature is a black hole with non-zero mass. But anytime we have a black hole, we will have dissipation: if we scatter a wave off the black hole, some fraction of its amplitude, and thus of the momentum it carried, will be absorbed into the black hole. So the fact that there is always a generic source of dissipation is very much built into holographic duality.

A useful way to compute η/s begins with the Kubo formula for η,

$$\eta = -\lim_{\omega,k\to 0} \frac{\mathrm{Im} G_{\mathrm{R}}(\omega, k)}{\omega}, \tag{122}$$

---

[61] See [359] for a summary of the current state of such computations and [237, 339] for a more detailed exposition of the computation we outline and for references to the original literature.

which relates η to the retarded Green function G_{R} for shear modes of the stress tensor,

$$G_{\mathrm{R}}(t,x) = \langle T_{xy}(t,x) T_{xy}(0,0)\rangle\, \Theta(t). \tag{123}$$

Note that we take $\omega \to 0$ only after sending $k \to 0$. The Kubo formula follows from matching linear response theory to the expected low-frequency, low-momentum limit of the correlation function in hydrodynamics; see, for example, [317]. To obtain η we need to compute $G_{\mathrm{R}}(t,x)$. In general, this is difficult in a strongly coupled QFT. However, holographically this calculation turns out to be relatively straightforward. According to the holographic dictionary, the operator T_{xy} is dual to the corresponding shear mode of the bulk metric, g_{xy}. We thus need the response of the classical gravitational action, $I_{\mathrm{AdS}}[g]$, to variations of this mode, $\delta g_{xy} = h_{xy}$. So long as we are expanding around pure AdS, h_{xy} is governed by a simple massless scalar action of the form in equation (79), but with the normalization of the kinetic term determined by the normalization of the gravitational kinetic term, $\frac{1}{16\pi G_{\mathrm{N}}}$,

$$I_\phi = -\frac{1}{16\pi G_{\mathrm{N}}} \int \mathrm{d}^{d+1}x\, \sqrt{g}\, \frac{1}{2}\, (\partial\phi)^2\, . \tag{124}$$

Since $m = 0$, we have $\Delta(d-\Delta) = 0$, so the dimension of our dual operator is $\Delta = d$. Let us use this action and the formalism developed above to compute G_{R}. With $\Delta = d$, the retarded Green function takes the form

$$G_{\mathrm{R}}(k) = \lim_{r\to 0} \frac{\Pi_{\mathrm{in}}(\omega,\vec{k},r)}{\Phi_{\mathrm{in}}(\omega,\vec{k},r)}. \tag{125}$$

It is readily checked (cf (114)) that the real part of the ratio above scales, after sending $\vec{k} \to 0$, as ω^3, and thus vanishes in our limit. We can thus replace $\mathrm{Im} G_{\mathrm{R}}(\omega,k)$ with $\frac{1}{\mathrm{i}}\, G_{\mathrm{R}}(\omega,k)$ to obtain

$$\eta = -\lim_{\omega,k\to 0}\lim_{r\to 0} \frac{\Pi_{\mathrm{in}}(\omega,\vec{k},r)}{\mathrm{i}\omega\Phi_{\mathrm{in}}(\omega,\vec{k},r)}. \tag{126}$$

At this point, something remarkable happens. In the limit that $\omega, \vec{k} \to 0$, the radial evolution of Π and $\omega\Phi$ trivialize. This can be seen by inspection of the definition of Π and the Hamiltonian equation of motion, $\partial_r\,(\omega\Phi) \sim \omega\Pi$ and $\partial_r \Pi \sim k^\mu k^\nu \partial_\mu \partial_n \Phi$. Thus, in the limit $\omega, \vec{k} \to 0$, both Π and $\omega\Phi$ become independent of r.

We can thus evaluate the ratio appearing in our Kubo formula at any radial coordinate we like. A convenient choice is the horizon, where the solution must satisfy in-falling boundary conditions. From our earlier discussion of the in-falling solutions (85), which implied that $f\partial_r\Phi_{\mathrm{in}} = \mathrm{i}\omega\Phi_{\mathrm{in}}$ at the horizon, and using the definition $\Pi_{\mathrm{in}} = -\frac{1}{16\pi G_{\mathrm{N}}}\sqrt{-g}g^{rr}\partial_r\Phi_{\mathrm{in}}$, we have that

$$\lim_{r\to r_{\mathrm{H}}} \Pi_{\mathrm{in}} = -\frac{1}{16\pi G_{\mathrm{N}} r_{\mathrm{H}}^3} \mathrm{i}\omega\Phi_{\mathrm{in}}(r_{\mathrm{H}}). \tag{127}$$

The factor of r_{H}^{-3} can be nicely interpreted as the area of the horizon in boundary units, $\frac{A_{\mathrm{H}}}{A_\partial}$. Plugging all of this into our Kubo formula, equation (126) then becomes

$$\eta = \frac{A_{\mathrm{H}}}{16\pi G_{\mathrm{N}} A_\partial}, \tag{128}$$

where A_{H} is the area of the horizon and A_∂ is the area of the boundary. From equation (117), $\frac{A_{\mathrm{H}}}{A_\partial} = 4G_{\mathrm{N}}s$. Putting these together, we see that the viscosity-to-entropy-density ratio takes the simple form now known as the KSS form after the authors of [11] (however see also, e.g., [235, 236, 301]),

$$\frac{\eta}{s}=\frac{1}{4\pi}, \tag{129}$$

which is orders of magnitude smaller than that of typical liquids. In fact, the liquids that are believed to come close to this level of perfection are the QGP at RHIC and fermions at unitarity, as shown in figure 2.

Many features of this result are remarkable. Firstly, as motivated heuristically above, η/s is indeed a good diagnostic of strong coupling in any many-body system. Near weak coupling, QFT perturbation theory gives

$$\frac{\eta}{s}\sim\frac{A}{\lambda^2\ln(\sqrt{\lambda})}, \tag{130}$$

where the constant of proportionality is system dependent. In particular, weakly coupled QFTs generally have $\frac{\eta}{s}\gg 1$. Near strong coupling, by contrast, a more careful gravitational analysis gives

$$\frac{\eta}{s}=\frac{1}{4\pi}+O\left(\frac{1}{\lambda^{3/2}}\right), \tag{131}$$

almost independently of the details of the system. Thus, in any such system, the behavior of $\frac{\eta}{s}$ changes dramatically as we run from weak to strong coupling.

Secondly, equation (129) is remarkably generic in holographic systems. Indeed, in any QFT which is holographically dual to GR, whether the gauge group is $SO(N)$ or $SP(2N)$ or something more complicated, whether enjoying maximal supersymmetry or no supersymmetry, we obtain precisely $\frac{\eta}{s}=\frac{1}{4\pi}$. To see why this result is so general, note that the radial evolution dropped out of our computation of η, so that all that we needed was the local physics near the horizon. The rest of the holographic description (including the UV physics near the boundary) was moot. Meanwhile, s is explicitly proportional to the horizon area, so again it is completely determined by the physics of the horizon. The generality of this holographic result then follows from the exceptional universality of the physics of horizons in GR.

Observations such as these led KSS to make the conjecture that $\eta\geqslant\frac{1}{4\pi}$ is a universal lower bound on the shear viscosity of any quantum many-body system, with saturation occurring only for those strongly coupled theories which are holographically dual to pure Einstein gravity. We now understand that this result is not universal. Explicitly, equation (129) holds when the space–time geometry is AdS, the gravitational action is Einstein–Hilbert and the fluid is time independent, homogeneous and isotropic. However, when higher-derivative terms are present in the gravitational action (corresponding to finite-N and finite-λ effects in the dual QFT) or when the geometry is deformed away from AdS such that the radial evolution does not factor out (as occurs, for example, when the geometry is time dependent, spatially disordered or sufficiently anisotropic), $\frac{\eta}{s}$ can depart from the KSS value, and indeed can be lower than $\frac{1}{4\pi}$.

The cleanest example of such an effect involves [10, 360, 361] adding the leading irrelevant operator to the bulk gravitational action, the 5D Gauss–Bonnet term, a quadratic scalar built out of the curvature tensors,

$$I_{\mathrm{GB}}=\frac{1}{16\pi G_{\mathrm{N}}}\int \mathrm{d}^{d+1}x\,\sqrt{-g}\left(-2\Lambda+R+\lambda_{\mathrm{GB}}R^2_{\mathrm{GB}}\right). \tag{132}$$

Such terms arise in known string theories where we know and have control over both sides of the duality[62]. As it turns out, the holographic analysis including $\lambda_{\rm GB} \neq 0$ is straightforward. All that materially changes is the normalization of the action for the scalar mode h_{xy}, which picks an extra factor of $(1–4\lambda_{\rm GB})$, yielding

$$\frac{\eta}{s} = \frac{1-4\lambda_{\rm GB}}{4\pi}. \tag{133}$$

When $\lambda_{\rm GB} > 0$, the KSS conjecture is violated[63].

It is worth understanding why this happens. The value of η/s is controlled by the normalization of the kinetic term in the scalar action for our tensor mode h_{xy}. Typically, such coefficients are best understood as coupling constants, since scaling them away to make the kinetic term canonical shifts the coefficient into the interactions[64]. When we make the coupling stronger, as in the case of Gauss–Bonnet with $\lambda_{\rm GB} > 0$, the viscosity goes down. This is reminiscent of the dependence of viscosity on coupling in kinetic theory: increasing the strength of quasiparticle scattering inhibits transport and decreases the shear viscosity. At strong coupling, there are no longer any well-defined quasiparticles to which to apply kinetic theory; instead, the effective degrees of freedom mediating low-energy momentum transport are weakly coupled gravitons in the holographic space–time.

Interestingly, while this evidence rules out the KSS conjecture, there exist compelling arguments for a weakened version of the KSS bound. Physically, the arguments for such a refined bound involve demonstrating that violations of these bounds would lead, in a wide class of theories, either to violations of causality in the field theory [10, 362] or to violations of the positivity of energy [363]; in fact, these two pathologies are intimately related. In the case of Gauss–Bonnet gravity which is holographically dual to 4D QFTs with $\mathcal{N} = 1$ supersymmetry, both arguments lead to the same constraint, $\frac{-7}{36} \leqslant \lambda_{\rm GB} \leqslant \frac{9}{100}$, leading to the refined lower bound

$$\frac{\eta}{s} \geqslant \frac{1}{4\pi}\frac{16}{25}. \tag{134}$$

In more general models, similar bounds obtain, but the precise values differ. The meaning and reliability of such refined bounds remain a matter of active debate. For example, it has been shown [364, 365] that sufficiently contrived examples may be concocted which violate any such bound and in fact push $\frac{\eta}{s}$ arbitrarily close to zero, and even negative! However, it is not at all clear that these models are themselves well defined; indeed, one can immediately show [366] that all such models with $\frac{\eta}{s}$ negative or vanishing are pathological, with causality and self-consistency enforcing a minimal value for $\frac{\eta}{s}$. Meanwhile, whether any such bound should be expected when we give up time independence, homogeneity or isotropy is a topic of much current research. To summarize, there are well-defined lower bounds on $\frac{\eta}{s}$ in various classes of theories, but each such bound has been violated by considering a more general class of models. It remains possible that there is a simple, completely universal lower bound on $\frac{\eta}{s}$, but the evidence for such a bound is increasingly tenuous.

---

[62] In the QFT, the Gauss–Bonnet coupling $\lambda_{\rm GB}$ is related to the Euler and Weyl anomaly coefficients. While it cannot be tuned arbitrarily in these examples while remaining within the regime of validity of the classical gravitational description, it can certainly be taken away from zero in a controlled fashion [361].

[63] Note that $\lambda_{\rm GB} < \frac{1}{4}$ is a consistency condition for GB gravity. When $\lambda_{\rm GB} = \frac{1}{4}$, the kinetic term for the graviton degenerates and the theory becomes ill-defined.

[64] The universality of η/s in GR thus physically derives from the equivalence principle: in Einstein gravity, the metric couples universally with a single overall coefficient to all forms of stress energy. When we muck with this universality, for example by breaking symmetries or by modifying the way the metric interacts with itself via higher-curvature corrections, we make our chosen tensor mode of the metric couple more, or less, strongly.

Worrying about universality of a particular bound, however, misses the point. The lasting import of equation (129) is that it establishes a link between ultra-low shear viscosity and strongly interacting many-body systems. Said differently, if one finds a fluid with an extremely small shear viscosity, one has very strong evidence that it enjoys strong quantum correlations. Conversely, extreme quantum fluids should be expected to have ultra-low viscosity of the order of $\frac{1}{4\pi}$. Whatever the final story about a hard lower limit, this connection is an important lesson from holography.

4.3.2. Diffusion, jet quenching and dynamics in the holographic QGP. A related transport coefficient which can be fruitfully studied holographically is the diffusion constant for a heavy quark in a strongly coupled holographic QGP. It is an interesting example for us for several reasons. Firstly, the structure of the holographic computation is very different from the computations described above, and in fact owes a great deal to the string-theoretic origin of holographic duality. Secondly, the result is very different from the simple universality of the viscosity computation. Specifically, the diffusion constant for a heavy quark scales as

$$D = \frac{2}{\pi T}\frac{1}{\sqrt{\lambda}} \tag{135}$$

and is thus very strongly sensitive to the 't Hooft coupling, λ. In particular, this suggests that even very heavy quarks should rapidly equilibrate in the QGP, in sharp contrast to perturbative estimates. Note that this also differs sharply from the perturbative YM result[65]

$$D = \frac{6}{2\pi T}\left(\frac{1}{2\alpha_s}\right)^2, \tag{136}$$

for while equations (135) and (136) look similar, the latter is only valid for $\alpha_s \ll 1$, while the former is only valid for $\lambda \gg 1$.

The basic structure of the holographic computation is described in [289, 291]; see [286] for a perturbative analysis. Suppose that we place a heavy quark in the QGP and treat it in a Born–Oppenheimer approximation as a slow, heavy, point-like variable which sources the gauge field. As we run the RG, this charge is conserved. Holographically, this means that we should have a point-like defect not just on the UV boundary, but also on every slice of the holographic bulk. The heavy quark on the boundary is thus the endpoint of a string in the bulk which hangs down from the boundary into the AdS horizon, hanging along a straight line according to the gravitational tug of AdS. The string can be thought of as the holographic image of the color flux of the fundamental quark in the $3+1$ boundary. Explicitly, the dynamics of the string are determined by a simple least action principle, where the action is the area of the surface in space–time swept out by the string in AdS. For a static string, the extremal string is simply the string hanging straight down from the boundary quark.

To measure the diffusion constant, we need to give our quark a kick. The easiest way to do this is to drag the quark through the QGP with a fixed velocity and then determine the drag force. In equilibrium, we must have

$$f = -\eta_D p, \tag{137}$$

---

[65] This result corresponds to a specific numerical estimate of the logarithmic terms in equation (58); see equation (6.1) of [286].

where η_D is the drag coefficient. Holographically, the string hanging off the quark into the bulk will now dangle behind the moving quark, imparting a drag force to the quark. Using the action principle we just defined, the drag force is computed to be

$$f = -\frac{\pi}{2}\sqrt{\lambda}T^2\frac{v}{\sqrt{1-v^2}}. \tag{138}$$

For heavy quarks, for which the thermal corrections to the kinetic mass are negligible, this gives

$$\eta_D = \frac{\pi}{2m}\sqrt{\lambda}T^2. \tag{139}$$

Using this in the relation $D = T/\eta_D m$ derived in equation (57) yields equation (135).

It is worth commenting on how this computation is similar to the viscosity calculation and how it differs. Perhaps the most important similarity is that in both cases we computed deeply quantum-mechanical transport quantities in our strongly coupled QGP by solving simple, classical differential equations in the dual gravitational system. This is a basic hallmark of holographic duality. Secondly, the resulting computation was surprisingly independent of the details of the holographic setup, depending only on the gravitational dynamics of a string hanging into AdS from a moving point on the boundary. This suggests that a heavy fundamental in any strongly coupled QGP should have a similar qualitative behavior, namely efficient diffusion and rapid equilibration. By contrast, the final results differ in significant ways. For example, while η/s was remarkably insensitive to the coupling, D depends sensitively and is thus a sensitive probe of precisely how strongly coupled the theory is.

A similar computation leads to a holographic estimate of the jet-quenching parameter, $\hat{q}$, introduced in section 3.6. Holographically, the role of the QGP fireball is played by the dual black hole, while the hard jet can be modeled as a heavy quark at very high energy, and thus moving on an effectively light-like trajectory. Our job is thus to compute the transverse momentum gained by a light-like quark moving through the QGP. This can be reduced [367–369] to computing a light-like Wilson loop in the QGP. Holographically, the expectation value of this Wilson loop can again be computed by studying the surface swept out by the trailing string, determined as before by extremizing the string action. Since the Wilson loop is light-like, the computation must be performed in Lorentzian signature. The resulting real-time holographic analysis, while somewhat intricate, leads to a simple prediction for the jet-quenching parameter in $\mathcal{N}=4$ SYM [369, 370],

$$\hat{q}_{\mathcal{N}=4} = \frac{\pi^{3/2}\Gamma\left(\frac{3}{4}\right)}{\Gamma\left(\frac{5}{4}\right)}\sqrt{\lambda}\,T^3. \tag{140}$$

Recalling that $\lambda = g^2 N_c$ and $g^2 = 4\pi\alpha_s$, and plugging in $T \sim 300\,\mathrm{MeV}$ and $\alpha_s = \frac{1}{2}$, which are reasonable estimates for the values at RHIC, this gives $\hat{q}_{\mathcal{N}=4} \sim 4.5\,\mathrm{GeV}^2\,\mathrm{fm}^{-1}$, which is comparable with the results from RHIC; see section 3.6. Note that this holographic result leads to a semi-strongly coupled treatment of energy loss, treating transverse momentum diffusion in strong coupling via a holographic estimate of $\hat{q}$, while the actual energy loss (longitudinal drag) is done in perturbation theory, as discussed in section 3.6. For fully strongly coupled treatments, see, e.g., [371, 372].

An important lesson of this holographic computation, however, is not the specific value of $\hat{q}$ we find but rather how $\hat{q}$ depends on the physical parameters of the system. In particular, it

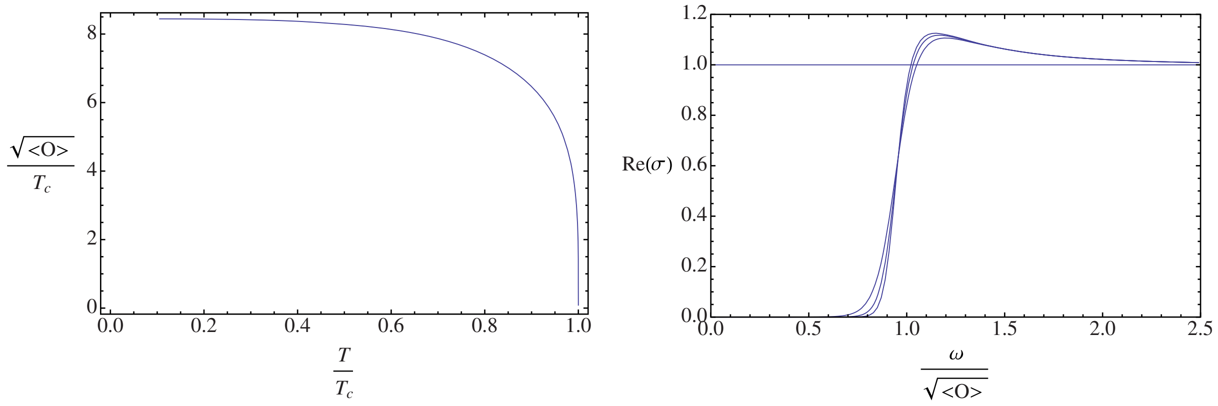

Figure 25. Condensates and conductivity of a holographic superconductor. The results are adapted from the code by Herzog [377], presented in the original work [375].

turns out that in any holographic QGP dual to pure AdS, and thus governed by a dual CFT, one finds that [370]

$$\frac{\hat{q}_{\text{CFT}}}{\hat{q}_{\mathcal{N}=4}} = \sqrt{\frac{s_{\text{CFT}}}{s_{\mathcal{N}=4}}}, \tag{141}$$

where s is the entropy density of the QGP. The energy lost by a hard parton plowing through the QGP is thus not proportional to the entropy density and so cannot be well modeled by scattering off a gas of persistent quasiparticles.

This holographic approach to studying the physics of a generic $SU(N)$ QGP to glean insight into the dynamics of e.g. the physical QGP at RHIC has been followed in many directions, including several of the papers in this focus issue. For example, Hubeny [373] uses a simple but remarkably illuminating gravitational model for an accelerated heavy quark in large-N QCD to study the collimation of synchrotron radiation emitted by a circling quark. Like the diffusion and jet-quenching computations sketched above, Hubeny's calculation involves tracing the path of the bulk string which trails behind the accelerating quark into the bulk of AdS. The result of the acceleration is that this bulk string becomes coiled, leading to the radiation emitted by the string remaining tightly collimated. Curiously, this result at first appears to be in tension with standard holographic intuition. Hubeny resolves this apparent conflict through a careful study of the gravitational backreaction induced by the coiled string.

4.3.3. Gravitational engineering and holographic superconductors. An interesting example of defining the field theory by its holographic dual is the holographic superconductor (figure 25). The idea, first outlined by Gubser [374] and developed in detail by Hartnoll *et al* [375, 376], goes roughly as follows. Suppose we have a QFT which enters a superconducting phase at low temperatures by developing a non-zero condensate for some scalar operator, $\mathcal{O}$, spontaneously breaking a global $U(1)$ symmetry. What would that look like in the holographic context? By appeal to the holographic dictionary in table 1, the global current in the boundary is dual to a gauge field in the bulk, while the charged scalar order parameter in the boundary is dual to a scalar field in the bulk, which is charged under the bulk gauge field. Spontaneously breaking (or *Higgsing*) the global $U(1)$ symmetry on the boundary by giving a vacuum expectation value to

$\mathcal{O}$ is then dual to spontaneously breaking the bulk gauge $U(1)$ via a non-trivial profile for the bulk scalar field, $\Phi(r)$.

Our goal is thus to build a gravitational system which is dual to some QFT which exhibits this symmetry breaking pattern—not a specific QFT, but any QFT with this structure. By adjusting the gravitational theory we will then be able to explore a whole space of strongly coupled QFTs with low-energy superfluid phases. Procedurally, we will take the gravitational system as a definition of the dual QFT, computing the correlation functions, etc, of the QFT entirely by appeal to the holographic dual.

The holographic dictionary thus tells us that, whatever else exists in the gravitational dual of a superconducting QFT, we will need a vector field A_μ dual to the boundary current $\mathcal{J}_m$ and a scalar field Φ dual to the scalar order parameter $\mathcal{O}$. The dictionary, however, does not tell us the Lagrangian, so we must choose it. A minimal first guess is

$$I_{\mathrm{grav}} = -\frac{1}{16\pi G_{\mathrm{N}}} \int \sqrt{g} \left(-2\Lambda + R - \frac{1}{e^2} \left[\frac{1}{4} F^2 + |D\Phi|^2 + m^2 |\Phi|^2 \right] \right). \qquad (142)$$

We can further simplify our lives by considering the *probe limit*, $g \to \infty$, in which case the matter field stress tensor (which scales as $\frac{1}{e^2}$) is negligible so that the Einstein equation, and thus the metric, is not modified from its original AdS form.

Our task now is to find non-trivial solutions of the bulk equations of motion which follow from our trial action expanded around an AdS black brane geometry, so as to put the system at finite temperature.

Non-trivial means that the scalar field Φ acquires a non-vanishing profile which spontaneously breaks the gauge symmetry in the bulk. From our dictionary, this corresponds to the boundary operator $\mathcal{O}$ acquiring a vacuum expectation value which spontaneously breaks the global symmetry on the boundary. Once we find such a solution, we can compute correlation functions in the resulting system by evaluating the bulk action on-shell as a functional of the non-normalizable modes and taking derivatives with respect to those modes, as discussed in section 4.2.3.

For simplicity, let us focus on homogeneous configurations of the form $A = A_t(r)\mathrm{d}t$ and $\Phi = \Phi(r)$. The bulk equations of motion in our AdS black brane take the form

$$r^2 f A_t'' - 2|\Phi|^2 A_t = 0, \quad r^2 f\, \Phi'' + (rf' - 2f) r \Phi' + \left(\frac{r^2 A_t^2}{f} - m^2 \right) \Phi = 0. \qquad (143)$$

Near the boundary, the general solution to the bulk field equations takes the form

$$A_t \sim \mu + \rho\, r + \cdots, \quad \Phi \sim \phi_{d-\Delta} r^{d-\Delta} + \phi_\Delta r^\Delta + \cdots, \qquad (144)$$

where $m^2 L^2 = \Delta(\Delta - d)$. A convenient choice for the mass is $m^2 L^2 = -2$, which gives $\Delta = 1$ or 2. Let us focus on $\Delta = 2$. The holographic dictionary tells us to interpret μ as the source of the $\mathcal{J}_t$ component of the current, so μ plays the role of the chemical potential for the boundary charge density. ρ, being the subleading piece of A_t near the boundary, then represents the expectation value of the operator conjugate to μ, $\rho = \langle \mathcal{J}_t \rangle$, so we can identify ρ as the charge density induced by μ. Similarly, $\phi_{d-\Delta}$, as the leading term in Φ near the boundary, is interpreted as the source for the boundary operator $\mathcal{O}$, and ϕ_Δ as the vacuum expectation value, $\phi_\Delta = \langle \mathcal{O} \rangle$.

To construct a spontaneous condensate, we should turn off the source, $\phi_{d-\Delta} = 0$, fix a reference value for the chemical potential, $\mu = \mu_{\mathrm{o}}$, and look for a solution of the bulk equations of motion which have, below some critical temperature T_c, a non-zero value for $\langle \mathcal{O} \rangle = \phi_\Delta$. This requires that we solve equations (143) subject to the boundary conditions that $\Phi(0) = 0$ (no source) and $A_t(0) = \mu_o$. Generally, this must be done numerically, as equations (143) are nonlinearly coupled. By evaluating the on-shell action and using the holographic dictionary in table 1, we can compute physical quantities. For example, we find the critical temperature by searching for the minimum T of the background black brane such that a non-trivial solution with these boundary conditions exists. We find the ac conductivities by looking at linear response for fluctuations of A_x linearized around a non-trivial solution for A_t and Φ, etc. The precise behavior we find depends on the details of the gravitational model chosen. By adding additional fields and interactions, we can use these techniques to engineer a wide range of dynamical phenomena in the dual QFT.

The point to emphasize here is that while we are explicitly computing correlation functions in a strongly coupled QFT, and indeed studying the detailed physics of transport, the QFT under investigation is defined purely through its gravitational dual. We thus call these superconductors *holographic superconductors*, and more generally refer to such QFTs as *holographic QFTs*.

Importantly, these models can be used to study much more than just the translationally invariant ground states of these strongly interacting superfluids. For example, they provide a simple and powerful framework for studying spatially ordered phases and solitons. For example, the paper by Keränen *et al* in this focus issue [164] uses solitons in these holographic superconductors as precision probes of the superfluid, including, in particular, the excitation spectrum [378, 379] around a dark soliton at unitarity in the BCS–BEC crossover [162, 165–167]. As the solution of this problem reduces to a detailed study of coupled PDEs, considerable progress can be made with common computational techniques toward a concrete prediction observable in the present experiments on ultracold quantum gases.

Notably, this holographic mechanism for spontaneously breaking a symmetry can be applied to any boundary QFT with symmetry breaking. For example, the paper by Basu *et al* in this focus issue [380] uses the holographic superconductor paradigm to study color superconductivity in holographic models of the QGP. As good effective field theorists, Basu *et al* begin by writing down a gravitational effective theory with the minimal ingredients to mock up a color superconductor and then examine what constraints they must impose on the parameters of this effective description to reproduce the expected physics of a true color superconductor.

4.3.4. Non-relativistic holography and cold atoms. Strongly correlated quantum liquids of cold atoms provide an excellent target for applied holography. From a holographic point of view, cold atoms at unitarity are very similar to $\mathcal{N} = 4$ models of the RHIC fireball: they provide examples of strongly correlated liquids at finite temperature and chemical potential that are governed by a conformal field theory. The key difference is that these systems are non-relativistic, with dispersion relations scaling as[66]

$$\omega \sim k^z, \tag{145}$$

so that the RG fixed points that govern their dynamics are non-relativistic conformal field theories (NRCFTs) [381–383]. Furthermore, in the case $z = 2$ arising in cold fermions at

[66] The scaling exponent z is called the *dynamical exponent*, and can be extracted either from the dispersion relation for low-lying modes or from the scaling of the characteristic time scale with the correlation length, $\tau \sim \xi^z$, as we approach a critical point [340].

unitarity, the symmetry algebra is enlarged [383] and includes a central *number operator*, $\hat{N}$, such that every operator has, in addition to a dimension, Δ, a conserved *particle number* eigenvalue, N.

Holographically, the non-relativistic scaling has a dramatic effect. In the case of $\mathcal{N}=4$ theory, relativistic conformal invariance was enough to determine the dual space–time, with the conformal symmetry group of the QFT, $SO(4,2)$, mapping to the isometry group of the dual space–time; this fixes the geometry dual to the $\mathcal{N}=4$ theory to be AdS$_5$. More generally, the fact that relativistic QFTs flow to Lorentz-invariant RG fixed points in the UV implies that their holographic duals should asymptotically approach pure AdS near the boundary. If our QFT does not possess Lorentz invariance, the dual space–time cannot be simple AdS. To build the holographic dual of an NRCFT, then, we must start from the beginning and identify the right space–time.

The simplest geometry with these properties is the *Lifschitz metric* [308]

$$\mathrm{d}s^2 = L^2\left(-\frac{\mathrm{d}t^2}{r^{2z}} + \frac{\mathrm{d}\vec{x}^2 + \mathrm{d}r^2}{r^2}\right). \tag{146}$$

In this geometry, every slice at a fixed r is again a copy of flat space, but the scaling of the time-like coordinate and the spatial coordinates under a shift in r is inhomogeneous, with scaling $\omega \sim k^z$. At $z=1$, this reduces precisely to the AdS case, a good check. However, at $z=2$ there is no enhancement of the isometry group of the geometry, so it is natural to look for another geometry which realizes the full enhanced $z=2$ non-relativistic conformal symmetry group. Following this logic leads to the *Schrödinger metric* [306, 307]

$$\mathrm{d}s^2 = L^2\left(-\frac{\mathrm{d}t^2}{r^{2z}} + \frac{2\mathrm{d}t\,\mathrm{d}\xi + \mathrm{d}\vec{x}^2 + \mathrm{d}r^2}{r^2}\right), \tag{147}$$

where ξ is a new periodic variable. Again we find that scaling implies that $\omega \sim k^z$. Now, however, the isometry group of the manifold is in fact enhanced at $z=2$. Moreover, since ξ is compact, we can expand every field in eigenmodes of $\hat{N} = \mathrm{i}\partial_\xi$, so that bulk fields, and thus all boundary operators to which they are dual, are labeled by both a dimension, Δ, and also a ξ-momentum, N. We thus have a geometry which enjoys both of the peculiar features of fermions at unitarity identified above.

Repeating the full construction in section 4.2 in this non-relativistic setting remains an open problem. However, considerable progress has been made in the context of both Lifschitz and Schrödinger systems, including the construction of charged black holes corresponding to NRCFTs at finite temperature and chemical potential [384–387] and the construction of some toy non-relativistic superfluids (see, e.g., [388] in this focus issue). Many of the basic universal results from the relativistic setting have been replicated. For example, in the simplest models we again find $\frac{\eta}{s} = \frac{1}{4\pi}$. Note that, while we have suppressed factors of c throughout, there are no such factors in this result for the viscosity ratio so we get a non-trivial prediction for non-relativistic systems. Furthermore, while $\frac{\eta}{s}$ is a natural observable in relativistic QFTs even if we do not know about holography (it is the dimensionless momentum diffusion constant in linearized hydrodynamics), the fact that $\frac{\eta}{s}$ is relevant in non-relativistic systems is an interesting feature of holography; the momentum diffusion constant in non-relativistic hydrodynamics is

$\frac{\eta}{n}$. In the context of Lifschitz scaling, considerably more progress has been made, including interesting results for models of strange metals [389].[67]

For models of cold atoms, on the other hand, our non-relativistic holographic models remain rather distant from the systems studied in actual experiments. Indeed, it seems likely that fermions at unitarity, like QCD, do not have weakly coupled holographic duals. However, unlike QCD, where we have already built holographic toy models which capture a great deal of the strong-coupling phenomena, non-relativistic CFTs with dynamical exponent $z = 2$ and Schrödinger symmetry have proven difficult to model holographically. In particular, the thermodynamics and phase structure of current holographic models are not obviously related to any known systems, if not outright pathological [384–386]. For example, we expect the finite-density ground state of these theories to form a superfluid which spontaneously breaks the $U(1)$ symmetry generated by the number operator, $\hat{N}$. Holographically, $\hat{N}$ is mapped to ∂_ξ, the generator of translations in the second holographic direction. Such a superfluid ground state must thus break translation invariance in the ξ-direction. Unfortunately, all our current holographic models of $z = 2$ Schrödinger CFTs feature manifest unbroken ξ-translation invariance.

This problem can be traced [390–392] to the fact that all our current models were derived from highly symmetric stringy systems which enforce this symmetry. As a result, these limitations do not appear to be intrinsic to the holographic framework but rather to the very specific models which have so far been studied. It would be of great interest to develop more realistic models (see, e.g., [392] for a recent work in this direction). A step in this direction is described in the paper by Adams and Wang in this focus issue [388], in which a superconducting ground state for such a non-relativistic holographic CFT is constructed by extending the 'holographic superconductor' strategy from AdS to these non-relativistic geometries. This leads to various interesting effects, including a surprising quantum phase transition as the density of the fluid is varied and a multicritical point where the transition to the superfluid state switches from second to first order. This does not solve the underlying problem, as the geometry remains unmodified, but is at least a proof of principle that such holographic ground states do exist.

4.3.5. Holographic non-Fermi liquids and critical phenomena.

Fermi surfaces are familiar beasts, and are generally well described by standard quasiparticle perturbation theory. This is true despite the fact that the electrons living on the Fermi surface may interact through the infinite-range Coulomb potential. The key point is that the phase space available for two electrons to scatter is extremely limited, as both the initial electron states and the final electron states must all lie on the Fermi surface. Folding in this phase space constraint, the effective interaction for (dressed) electrons scattering on the spherical Fermi surface mediated by the electromagnetic interaction is small, and hence the applicability of perturbation theory. The corresponding state is called a Fermi liquid[68].

---

[67] It is worth emphasizing that it is straightforward to build holographic models in which the dynamical exponent changes as the theory flows from the UV to IR. A particularly illuminating example involves the fate of charged black holes with large ground-state degeneracies. In at least some holographic models it has been shown that the ground state degeneracy is lifted by a dynamical instability which changes the near-horizon geometry from AdS to Lifschitz geometries which have no, or at least vastly diminished, ground state degeneracy. Since the near-horizon region encodes the IR physics of the dual QFT, such holographic models describe QFTs where the dynamical exponent of the low-energy degrees of freedom arises due to strong interactions among constituents with no simple quasiparticle description, and is thus effectively invisible in the UV regime. The dynamical scaling is in this sense *emergent*.

Experimentally, many systems with a Fermi surface are manifestly not well described by quasiparticle perturbation theory and are thus referred to as non-Fermi liquids. Perhaps the most famous such examples are high-T_c superconductors, whose spectral function (as measured e.g. by ARPES experiments) does not contain any sharp quasiparticle peaks. While it is relatively straightforward to write down phenomenological models of such non-Fermi liquids, it has proven somewhat non-trivial to build controlled microscopic models which generate non-Fermi-liquid behavior dynamically, and famously intractable to derive such non-Fermi-liquid behavior from realistic models of the crystal lattices which lead, experimentally, to such a behavior.

Since holographic QFTs neither require nor generally admit any quasiparticle description, it is natural to ask what happens when we put a holographic QFT at finite fermion number density [394–397]. The simplest way to do so is study a QFT with a fermion carrying a conserved charge under some global $U(1)$ for which we turn on a chemical potential. Holographically, this corresponds to studying a fermionic field in AdS in the presence of a charged black hole. The charge of the black hole induces a chemical potential in the QFT, while the boundary value of the bulk fermion specifies a source for the fermion in the QFT with which to probe the spectral function. To compute the spectral function, which is equal to the imaginary part of the retarded Green function for the fermion in the QFT, we turn on a known source and measure the response. Holographically, this maps to solving the Dirac equation in the background of a charged black hole in AdS subject to a Dirichlet boundary condition at the boundary and infalling boundary conditions at the horizon. This turns out to be a straightforward numerical computation. The resulting spectral functions show clear evidence of a sharp Fermi surface, but, as expected, enjoy no quasiparticle poles and so the holographic QFT describes a non-Fermi liquid.

Much work has been done to clarify the physics of these holographic Fermi liquids. For example, it has been shown that incorporating gravitational effects can change the results dramatically at low energy, with the black hole in some cases vomiting forth its charge and disappearing, leaving behind an *electron star* which is extremely dense but completely smooth, much like a neutron star but carrying a net charge [398, 399]. Meanwhile, completely different models of Fermi surfaces in holographic QFTs have been concocted [389]. Despite using very different ingredients, and having in some ways very different physics, all such models share the important property of having no simple quasiparticle poles in their spectral functions.

One of the key lessons of this study has been the role of emergent conformal symmetry in the near-horizon region [397, 400]. Recall that the near-horizon region corresponds to the deep IR of the dual QFT. Saying that this region has an emergent conformal symmetry then implies that the QFT has, at very low energy, an emergent scaling symmetry. Importantly, the dynamical exponent of this conformal symmetry is generally larger than 1 and indeed often infinite, corresponding to ultra-local physics with a finite ground-state degeneracy. As usual, this enormous ground-state degeneracy is a marker of instability, and these systems are quick to decay into a host of other phases. As such, these holographical critical points correspond to quantum critical points of the dual QFT.

---

[68] We discuss the electron gas as an example of a Fermi liquid, but we emphasize that the arguments given here are only true for purely Coulombic forces. Current–current interactions, for example, are unscreened and destroy quasiparticle behavior [393].

All of these models of holographic Fermi surfaces and holographic critical phenomena share a generic two-component structure: firstly, a set of degrees of freedom, χ, which correspond to the naive fermions living on the Fermi surface; secondly, operators ψ in an emergent IR CFT which couple to the χ fields. Integrating over the ψ-fields then generates new, strong interactions among the original χ fields, which is the origin of the strong interactions on the Fermi surface in these holographic models. The natural question then becomes, if all we care about is the low-energy physics, which is entirely encoded in the near-horizon region, why bother about the entire holographic framework? Why not just build a truncated model which only includes the fermions of interest coupled to the emergent CFT which controls the interactions? Doing so has come to be known as *semi-holography* [309, 311], and has proven quite useful (see, e.g., [400–402]).

An example of this semi-holographic approach is the paper by Nickel and Son [311] in this focus issue. In this paper, the authors argue that the transport of conserved charges is governed by an effective theory for a set of Goldstone modes coupled to an emergent near-horizon CFT. The fact that the physical modes are Goldstone modes powerfully constrains the dynamics and enforces much of the structure of holographic transport. The origin of the Goldstone bosons is quite beautiful. From the holographic dictionary, we know that the conserved $U(1)$ of the boundary QFT maps into a gauged $U(1)$ in the bulk, and thus to a dynamical emergent vector field in the bulk. Importantly, this vector has two boundary values, one at the boundary of AdS and the other as we approach the horizon. Now imagine we run the RG, integrating out the high-energy modes (which live near the boundary) shell by shell. If we integrate out the entire bulk, we are eventually left with a thin shell just outside the horizon, on the boundaries of which we have a dynamical $U(1)$ gauge field, and thus a total symmetry group which is $U(1) \times U(1)$. However, this is a fake—there is only one $U(1)$ in the system—so integrating out the degrees of freedom in the bulk between these two boundaries must spontaneously break $U(1) \times U(1)$ to a single $U(1)$. As Nickel and Son show, this indeed happens, leaving behind a Goldstone mode which remains coupled to the surviving emergent CFT. A Wilsonian analysis of the most general action allowed by the symmetries of the system then efficiently reproduces many classic results about transport in holographic models.

These types of holographic approaches to quantum critical phenomena have also been applied to myriad problems in the physics of condensed matter broadly construed. For example, the paper by Bayntun *et al* [403] in this focus issue studies plateau transitions in quantum-Hall-like holographic models. Such phenomenological models are becoming increasingly sophisticated and, as they improve, increasingly compelling. In this case, the authors construct a model for which the scaling exponent for plateau transitions is $\frac{2}{5}$, in rough agreement with the experimental values of ~ 0.42.

4.3.6. Far-from-equilibrium physics in strongly coupled quantum field theories. Perhaps the most exciting application of holographic duality is the study of far-from-equilibrium physics in strongly interacting quantum many-body systems. The basic idea here is extremely simple. As we have seen, the equilibrium thermal physics of a holographic QFT is encoded in the thermodynamics of a stationary[69] black hole horizon in the dual space–time. To study non-equilibrium physics in the QFT, then, we should study non-equilibrium physics in the dual gravity—i.e. time-dependent space–times whose black hole horizons evolve in time.

[69] In GR, *stationary* means, roughly, time independent, while *static* means, again roughly, non-rotating.

It is easy to visualize this connection. Imagine the collision of two nuclei at RHIC or LHC. At some initial time we have two approximately non-thermal objects in a head-on collision. At some special moment they collide. Quite rapidly—in fact, amazingly rapidly according to experimental results—the ensuing fireball thermalizes and then slowly evolves according to near-equilibrium hydrodynamics. Now what would this process look like in a holographic dual? Since the initial temperature is zero, the initial geometry has no black hole horizon. Meanwhile, the initial non-thermal nuclei are basically just moving sources of stress energy. According to the holographic dictionary, the source for stress energy in the boundary QFT is equal to the boundary value of the bulk metric. Our moving packets of mass are thus represented holographically by traveling gravitational waves which extend from the boundary into the deep interior of AdS. The collision of the nuclei then corresponds to a collision of the dual gravitational waves. But we know what happens when you collide gravitational waves at sufficiently high energy—they form a black hole.

However, since the waves are not perfectly thin and the process of black hole formation is not perfectly efficient, the details of black hole formation can be quite messy. In particular, the horizon does not form instantaneously but grows and expands rapidly during the collision—this black hole is far from equilibrium. Meanwhile, while the initial collision, and thus the initial horizon, will generically be inhomogeneous, anisotropic and just generally messy, it is a theorem of black hole mechanics that the horizon always settles down to a completely uniform spherical surface with a uniform radius and a uniform temperature—and indeed that it does so relatively rapidly, with fluctuations of the horizon damped by the viscosity we studied above. Thus our collision of nuclei and thermalization of the QGP fireball at RHIC, in the holographic dual, look like a collision of gravitational waves which form an initially far-from-equilibrium black hole, which then rapidly rings down to a simple uniform sphere.

Note the simplification afforded by this dual description: while the QFT description involves real-time quantum dynamics in a strongly coupled many-body system, which are extremely difficult to simulate, the holographic model involves solving a set of PDEs—the Einstein equations, together with a well-defined set of initial data and boundary conditions—which are considerably easier to simulate.

This dramatic simplification was exploited by Chesler and Yaffe [312], who simulated the formation of a holographic QGP by numerically colliding gravitational shock waves to form a black hole in AdS. An important lesson of the resulting numerical simulations is that these holographic QGPs thermalize exceedingly rapidly. The rapidity can be quantified in several ways. One simple measure is that the thermalization time is of the order of the time required for light to cross the horizon of the final equilibrium black hole. Alternatively, we can take the numerical results and rescale variables to reproduce the real scales of QCD. Doing so [312] gives a very rough prediction of the RHIC thermalization time of $0.35\,\mathrm{fm}/c$, which is safely below the upper bound estimated via hydrodynamic simulations from the RHIC data; see section 5.2. Another important lesson of this and related computations is that hydrodynamics become an excellent approximation very rapidly, with the pressure and energy tracking hydrodynamic predictions as fast as and sometimes apparently even faster than allowed by thermalization.

Importantly, the holographic models studied so far do not do nearly as good a job in other regards. For example, the most simple models fail to reproduce the observed energy dependence of the total multiplicities. This is perhaps not surprising given the simplicity of the models studied to this end, but is an important point to rectify. Further, while these models give us useful insight into the phenomena, they do not give us any guidance on how to understand the relation between such rapid thermalization and asymptotic freedom or deep inelastic scattering. Nonetheless, these are remarkable results.

5. Conclusions

We conclude this review by providing a list of open questions in the areas of ultracold quantum gases, QCD plasmas and holographic dualities. This list is far from complete, and the selection of problems is guided by the topics discussed in the main body of the review. Many of these problems will benefit from the connections between the physics of ultracold quantum gases, the QGP and holographic duality that are being explored in this focus issue. For example, improved studies of the hydrodynamic behavior of ultracold gases will benefit from what has been learned about second-order hydrodynamics in the context of the QGP and the AdS/CFT correspondence. Holographic dualities will also continue to play an important role in studying the approach from the initial, far-from-equilibrium, stage of heavy-ion collisions to the hydrodynamic regime. Far-from-equilibrium states have also been explored in ultracold atomic gases. Finally, new experimental results on the properties of the QGP obtained in the ongoing programs at RHIC and LHC, as well as new results from experiments with trapped gases, may point to new, unexpected connections between these fascinating systems.

5.1. Open problems and questions in ultracold quantum gases

(i) What is the structure of the unitary Fermi gas in both two and three dimensions? In particular, is it possible to understand the properties of the system in terms of well-defined quasiparticles?
At high temperature the gas is always weakly correlated, independent of the strength of the interaction. In the weak coupling regime we know that the quasiparticles near T_c are BCS-like gapped fermions, and that at very low temperature the quasiparticles are phonons. At unitarity it has been argued that the regime $T_c \lesssim T \lesssim 1.5T_c$ can be understood in terms of a pseudogap [35, 404, 405]. This means that there is a gap in the single-particle spectrum, but no long-range coherence. Pseudogap behavior in three dimensions has been observed using RF spectroscopy [406, 407], and it has also been seen in quantum Monte Carlo calculations [408]. Pseudogap behavior was also reported in experiments with two-dimensional systems [409]. However, it is not entirely clear to what extent thermodynamic and transport properties can be understood in terms of gapped quasiparticles. One quasiparticle that may turn out to be relevant is the polaron, a single spin-up particle immersed in a sea of spin-down fermions [410]. In highly polarized gases, RF spectroscopy has been used to measure the dispersion relation of polarons. It was observed that the polaron description can be extended to more spin balanced systems. In particular, a simple estimate based on a gas of polarons provides a very good estimate for the critical polarization at the Chandrasekhar–Clogston point, the first-order transition from an unpolarized superfluid to a polarized normal gas [411]. Recently, polarons have also been observed in two-dimensional Fermi gases [412, 413].

(ii) How do we make local measurements of transport coefficients such as shear and bulk viscosity, thermal conductivity and the spin diffusion constant?
Measurements of elliptic flow and collective mode damping provide trap averaged values of the shear viscosity (see section 2.5) but there is no direct measurement of the local

value of the shear viscosity $\eta(n, T)$. The status is roughly analogous to the situation after the first experiments aimed at measuring the equation of state. These experiments only determined the total energy and entropy [66]. Local measurements were made possible by new ideas such as the Gibbs–Duhem method employed by the ENS group [68] and the compressibility thermometer introduced by the MIT group [36]. It is possible, in principle, to unfold the currently available data for trap averaged viscosities, but this method requires a very careful treatment of the dilute corona, and a more direct method would be desirable. One possible direction is the development of new experimental techniques that generate local shear flows and study their relaxation. Another option is to measure the frequency-dependent shear viscosity via spectroscopic methods as suggested in [414]. Similar questions apply to other transport properties, such as bulk viscosity, thermal conductivity and spin diffusion. Current methods are mostly sensitive to average transport coefficients, but more local measurements are needed.

(iii) Are there reliable approaches to transport theory beyond the Boltzmann equation?
In the unitary limit, transport properties can be computed reliably using the Boltzmann equation in the limit $T \gg T_{\mathrm{F}}$. For $T \sim T_{\mathrm{F}}$ the quasiparticles are strongly interacting, and the kinetic description in terms of binary scattering between atoms breaks down. Enss *et al* [129] have computed the shear viscosity using a T-matrix resummation. A similar approach is discussed by Guo *et al* [110] in this focus issue. These methods are quite successful in describing the thermodynamics, but there is no expansion parameter, and it is not clear if non-equilibrium properties are predicted as well as equilibrium features. In section 2.2 we briefly discuss the application of $1/N$ and ϵ expansions to equilibrium properties of the unitary gas, but except for the recent work of [415] these methods have not been extended to non-equilibrium properties. Similar remarks apply to other methods as well: there are detailed studies of equilibrium properties using the exact RG [416], but no corresponding studies of transport properties. Finally, there now exist very accurate quantum Monte Carlo calculations of the equation of state, but there is only a single, pioneering, study of transport properties [5].

(iv) Are there other universal quantum gases that can be created by manipulating the dimensionality, the spin structure, the type of interaction or the dispersion relation?
The unitary gas exhibits a very high degree of universality: in the limit of infinite scattering length and zero effective range the many-body system is completely characterized by the fermion mass, the temperature and the chemical potential. The dependence on dimensionful combinations of these parameters is completely governed by dimensional analysis. Universality can be exploited to use the unitary gas as a model system for dilute neutron matter; see section 3.1. The question is whether this idea can be extended to more complicated systems, for example three-species gases that can serve as a model for the quark matter/nuclear matter transition (see section 2.7) or four-species gases that correspond to self-bound nuclear matter, which is a liquid of equal numbers of spin-up and -down protons and neutrons. Universality in systems with more than two degrees of freedom is more complicated, because the interaction depends also on at least one three-body parameter [417], and the system may require repulsive short-range forces in order to prevent collapse. New universal states can also be created by combining Bose and Fermi gases in different dimensions [418] or by manipulating the dispersion relation of the atoms. Steps in this direction were recently taken by Salger *et al* [419], who demonstrated linear dispersion and Klein tunneling of a Bose–Einstein condensate in a 1D optical lattice,

and by Taruell *et al* [420], who created a two-dimensional optical lattice with the same honeycomb lattice structure as graphene and therefore obtained low-energy fermions with a linear dispersion relation.

(v) What is the physics of the crossover in the discrete context of optical lattices?
Will the predictions of holographic duality of a lower limit on the ratio of viscosity to entropy hold true in the lattice context or not? What class of theories from holographic duality might map onto the Fermi resonance Hamiltonian or other effective two-channel models? How can we extend such models to treat the imbalanced Fermi gas? How will a polaron gas be described in the lattice context?

(vi) What will be the interplay between disorder and interactions for the unitary gas?
The relative effects of disorder and interactions on both thermal and quantum phase transitions in quantum many-body physics are an open question generally (see, e.g., [421]). An essential feature of disorder is Anderson localization, in which scattering off of many small defects conspires to localize all states above a certain threshold in three dimensions. This concept is, in fact, a wave physics or single-particle quantum concept: it has been so far unclear where and how interactions destroy or allow for Anderson localization in a generic sense. Disorder was treated in this focus issue for ultracold Fermi gases specifically by Han and Sá de Melo [172]. Here, too, we find potential connections to holographic duality, as explored in recent forays into disordered systems by Adams and Yaida [422, 423]. Although these explorations have not treated lattice physics in particular, disorder is an essential feature of lattice systems arising in nature, and can be induced artificially in various ways for optical lattices [424]. Can holographic duality capture the underlying lattice causes of disorder? How will experiments in ultracold Fermi gases combine disordered optical lattices with unitary fermions at sufficiently low temperatures? Might studies in this direction help us gain some fundamental insight into the interplay between disorder and interactions in strongly correlated systems?

5.2. *Open problems and questions in quantum chromodynamic plasmas*

(i) Can we determine all the transport properties of the QGP?
We would like to determine, with fully quantified experimental and theoretical uncertainties, the value of the shear viscosity to entropy density ratio η/s, the heavy quark diffusion constant D and the jet quenching parameter $\hat{q}$. This program is probably closest to completion in the case of η/s. In this case the basic method, second-order relativistic viscous hydrodynamics, is now fairly well understood. Final state corrections can be handled by coupling to a hadronic cascade. There is a large set of data, including the energy, impact parameter and system size dependence of elliptic flow, as well as the observation of higher harmonics that points to the validity of the hydrodynamic description. The internal consistency of hydrodynamics already implies a fairly strong bound $\eta/s \lesssim 0.5$–1.0. The main problem is that we have no direct method for establishing the initial conditions for the hydrodynamic description. This problem may be addressed by focusing on a larger set of observables, in particular higher harmonics of the flow [425, 426]. There also remain theoretical uncertainties regarding the coupling of viscous fluid dynamics to a Boltzmann description, the role of bulk viscosity and the equation of state. A discussion of the known uncertainties in η/s can be found in [427].

The main strategy for determining the heavy quark diffusion constant is to use a Langevin description of the diffusion process where the background collective expansion is taken from hydrodynamics. The most important problem at present is the fact that there are only data for single leptons, and not for identified heavy quarks. This problem will be solved by improved detection capabilities at RHIC and LHC. The main theoretical issue is how to handle the interplay between elastic and inelastic scattering. For very heavy quarks elastic scattering is dominant, but for charm quarks inelastic processes may be important. This includes hadronization effects such as coalescence.

In the case of jet energy loss there are still a number of theoretical issues that need to be resolved. The radiative energy loss of an energetic parton ($E \gg T$) in a partonic medium can be expressed in terms of a single medium parameter, the transverse diffusion constant $\hat{q}$. It is not clear if this carries over to the spectra of jets or leading particles produced in the medium. Another source of uncertainty is that different implementations of the in-medium radiation in perturbative QCD lead to significantly different values of $\hat{q}$ when applied to the RHIC data [278]. There is a significant ongoing effort to understand these differences[70]. On the experimental side, the LHC represents a large step toward higher jet energies $E \sim 100\,\mathrm{GeV}$, and significantly improved statistics in the range $E \lesssim 20\,\mathrm{GeV}$ already explored at RHIC. Both at the LHC and at RHIC, experimentalists are studying not only spectra of hard particles, but also identified jets, jet shapes and correlations between hard particles and the associated soft emission.

(ii) Can we understand early thermalization?
The success of the hydrodynamic description implies that the system must thermalize early. The question is: how to quantify this statement experimentally, and how to understand early equilibration theoretically? The most reliable constraint on the thermalization time arises from the observation of elliptic flow. Elliptic flow is driven by the anisotropy of the initial state. Since ballistic expansion dilutes the initial anisotropy, local equilibration must happen early for elliptic flow to develop. A very conservative constraint, discussed in [182], is $\tau_0 < 2.5\,\mathrm{fm}/c$. Typical values used in hydrodynamic fits are $\tau_0 \lesssim 1\,\mathrm{fm}/c^{-1}$. Another hint of early thermalization comes from the observation of thermal photons with temperatures significantly larger than those seen in hadronic spectra. The PHENIX collaboration has recently reported a photon transverse momentum spectrum with a slope parameter $T = 221 \pm 19 \pm 19\,\mathrm{MeV}$ [428]. The spectra can be reproduced in hydrodynamic calculations with thermalization times ranging from $\tau_0 \lesssim 0.15\,\mathrm{fm}/c$ to $\tau_0 \lesssim 0.6\,\mathrm{fm}/c$, corresponding to initial temperatures $T_0 \sim 300$–$600\,\mathrm{MeV}$. Establishing experimental constraints on the equilibration time is further complicated by the fact that not all observables require full thermalization. For example, for the development of radial and elliptic flow it is not necessary that longitudinal momentum distributions be fully equilibrated [429].

Understanding thermalization theoretically is difficult even if the early evolution is governed by weak coupling. The first attempt to understand thermalization starting from an initial color-glass state using $2 \to 2$ and $2 \to 3$ scattering is the bottom-up scenario described in [430]. It was later realized that this picture is modified by collective effects that lead to plasma instabilities [431–433]; see also the contribution by Dusling in this focus issue [434]. In all these approaches it is hard to understand how thermalization at RHIC can occur on a time scale $\tau_0 \lesssim 1\,\mathrm{fm}/c$.

Attempts to understand thermalization at strong coupling have focused on holographic duality; see section 4.3.6 and [435, 436] for recent overviews. Holographic duality provides an interesting geometric picture of thermalization: the initial state corresponds to two colliding shock waves, and thermalization is signaled by the formation of an event horizon. The hydrodynamic stage is described by the relaxation or *ringdown* of the black hole. In holographic duality one can achieve fast equilibration [312], but it is hard to make contact with asymptotic freedom and the well-established theory and phenomenology of nuclear deep inelastic scattering, and it is difficult to understand the observed energy and impact parameter dependence of the total multiplicity [437].

---

[70] Topical Collaboration on Jet and Electromagnetic Tomography, http://jet.lbl.gov/.

(iii) Does the QGP have a quasiparticle description?
We would like to understand how the properties of the hot and dense matter produced at RHIC and the LHC are related to its structure, in particular whether the matter can indeed be described as a plasma made of quarks and gluons. Holographic duality seem to indicate that a strongly interacting quantum fluid near the viscosity bound does not have a quasiparticle description; see sections 3.2 and 4. In section 3.2 we discussed a number of observables that have been suggested as probes of the quasiparticle structure: fluctuations, the observed constituent quark scaling of flow, and the relation between the heavy quark diffusion constant and the shear viscosity. None of these observables provides a conclusive test of the quasiparticle picture, and new ideas would certainly be helpful. We also discussed the possibility to study quasiparticles by extracting the spectral functions of the correlators of conserved charges in lattice QCD, such as energy, momentum, flavor, etc. These studies, too, are somewhat indirect, because on the lattice we only have direct access to imaginary time correlation functions. This means that the spectral functions have to be extracted via analytic continuation. Analytic continuation of numerical data is difficult, but there has been significant progress over the last couple of years [438].

(iv) Can we experimentally locate the phase transition?
There is a continuing effort to establish the presence of a critical endpoint in the QCD phase diagram. The basic idea is to look for the non-monotonic behavior of fluctuations or correlations as a function of the energy of the colliding system. The first data from the RHIC beam energy scan have recently been released [439]; see also the contribution by Mohanty [440] in this focus issue. The data contain new information about the beam energy dependence of hydrodynamic flow which will help to establish the onset of a nearly perfect fluidity. Gupta *et al* [214] have argued that data on fluctuations can be directly compared with lattice QCD and that it can be used to set the scale for the phase transition at zero chemical potential. There is not yet conclusive evidence for the existence of a critical point, but data for additional beam energies and with better statistics will become available in the near future. In heavy-ion collisions the growth of the correlation length is limited by finite size and finite time effects, and it is important to identify observables that provide the best sensitivity to critical behavior. Observables that have been considered include fluctuations of conserved charges, such as baryon number, other fluctuation observables, such as the pion-to-kaon ratio, the mean p_T, or elliptic flow, and higher moments such as skewness and curtosis [441]. There is a parallel effort to locate the critical endpoint in lattice QCD. This effort is complicated by the sign problem that affects lattice calculations for fermions at finite chemical potential [442].

The discovery of a critical point would establish the presence of a phase transition between hadronic matter and the QGP. At small chemical potential this transition is only a crossover. One might hope that it would be possible to detect the crossover through hydrodynamic effects associated with the minimum in the speed of sound near the crossover temperature [443]. However, explicit calculations seem to show that there is no direct link between the presence of a minimum in the speed of sound and the beam energy dependence of multiplicity and flow [182].

Finally, we would like to demonstrate the existence of quark matter at high density and low temperature. It is not known whether the cores of compact stars are sufficiently dense to contain quark matter or other exotic phases. The presence of quark matter can be potentially detected through its effect on the equation of state, which is reflected in the mass–radius relationship. There is a very active program to pin down this relationship using observations of compact stars [444, 445]. Once the presence of a high-density phase has been established, more detailed information regarding its properties can be derived using the cooling and spin-down behavior of compact stars.

5.3. *Open problems and questions in holographic duality*

(i) What does it take to derive a holographic duality?
We have no first-principles derivation of holographic duality. This is perhaps not surprising, since it is a strong–weak duality, but it is frustrating. We do have a number of heuristic arguments, such as those that appeared in section 4.1. We also know how to derive a long but finite list of dual pairs as decoupling limits of particular string theories, e.g. the original construction, $\mathcal{N}=4$ Yang–Mills in 4D. But we do not know how to derive the duality directly within the bulk or boundary theories[71].

This is important for two main reasons. Firstly, we do not have a thorough understanding of the consistency conditions or regimes of validity of our holographic dual descriptions. We do have general rules of thumb, such as the scalings discussed in section 4.2.3, and we can check for self-consistency in specific examples. But particularly when we are defining a QFT via the holographic prescription, it is not always obvious what the precise regime of validity is and to what extent we should trust the chosen gravitational truncation.

Secondly, if we want to generalize holographic duality, it would be very helpful to know why it is true in the more familiar examples. For instance, we would like to generalize to intrinsically anisotropic systems, to asymptotically flat or positive-curvature (dS) space–times or even to understand why this is not possible.

(ii) Given a QFT, what is the gravitational dual?
Even if a general derivation is not possible, it would be useful to have a constructive map which allowed us to take an arbitrary QFT and identify the appropriate gravitational dual. For example, what is the exact dual of QCD, not QCD plus a collection of exotic fields, not of large N QCD, not of some toy model, but pure QCD? Similarly, with an eye toward ultracold atoms, what is the gravitational dual of the Hubbard model, or of the Kondo problem?

Again, the fact that holography is a strong–weak duality should give us pause. As we have seen, when the gravitational system is weakly coupled, the dual QFT generically does not have any simple quasiparticle interpretation. If we want to define the QFT using conventional Lagrangian methods, the best we will, in general, be able to do is say that it is the strong-coupling limit of some explicit perturbative theory. In the models we understand best, the coupling that is getting strong is tightly constrained by symmetries, e.g. gauge symmetries, flavor symmetries or supersymmetry, so we know how to follow the RG flow well beyond the perturbative, quasiparticle regime. We do not know how to solve the more general case without these extraordinary symmetries and the constraints they imply.

---

[71] This is not for lack of trying; see [333, 446, 447] for recent approaches.

(iii) Can we derive gravity from field theory?
The holographic graviton looks very much like an emergent vector boson in more familiar condensed matter systems. It arises as an emergent dynamical gauge field which reorganizes the strongly coupled physics into a set of perturbative local fluctuations. However, in holography the graviton (a) lives in one higher dimension and (b) couples universally. Point (a) fits naturally with a basic feature of gravity: since gauge transformations are coordinate transformations, there can be no truly local degrees of freedom, so any emergent graviton must come along with an emergent space–time dimension with which to fix its gauge redundancy, much as any emergent vector boson must come along with an extra Hubbard–Stratanovich scalar to fix its gauge redundancy. This was, in fact, the origin of the term *holography* in gravitational physics. Point (b) is more spectacular from the point of view of emergence, but is in fact a necessary feature of any theory with a fluctuating spin-2 gauge field.

It is thus tempting to wonder whether we can, in fact, derive the graviton in a general space–time as an emergent gauge boson in a strongly interacting QFT. Some progress has been made on this front (cf [446]), but a detailed understanding remains elusive.

(iv) Do we always need the complete holographic description?
In many of the examples discussed above, and indeed almost generically in the holographic literature, most of the work is being done by a very small fraction of the total degrees of freedom in the holographic description. For example, in the computation of η/s in $\mathcal{N}=4$ SYM, whose holographic dual, string theory in $\mathrm{AdS}_5\times S^5$, includes an infinite number of dynamical massive fields, the only fields that matter in the computation are the 5D metric and a single scalar mode encoding one particular fluctuation of this metric. In practice, keeping track of the full holographic description often seems to be overkill[72].

This raises two natural questions. Firstly, when does it make sense to truncate the full holographic description to a small subset of modes? Secondly, might this truncated holographic model have a more straightforward QFT representation than the full holographic description? Said differently, might there be a 'semi-holographic' formalism defined entirely within field theory which allows us to build conventional QFTs that reproduce the interesting phenomenology of fully holographic QFTs without using higher dimensions?

Recent work on holographic Fermi liquids, holographic quantum critical points and holographic quantum liquids (see [309] and especially the work of Nickel and Son in this focus issue [311]) demonstrates that at least some of the interesting holographic effects can indeed be realized in a purely quantum field theoretic formalism. On the other hand, for questions where the nonlinear dynamics of the geometry such as far-from-equilibrium physics, as opposed to kinematics or linear response, are important, it seems unlikely that any semi-holographic description can usefully capture the salient physics. Just how far a semi-holographic approach can be pushed, then, is an extremely interesting question.

---

[72] Note that this is analogously true in many-body physics: to describe ground states and low-energy dynamics, one generally only needs a tiny corner of Hilbert space. However, in the holographic description the physics of a single bulk field represents an enormously complicated and non-local set of interactions among the boundary degrees of freedom. So, observing that the bulk description can be effectively simplified to a small subset of bulk fields does not immediately imply that the same is true of the boundary QFT. For example, in holographic models of non-Fermi liquids, the bulk physics can be truncated rather dramatically, but the dual QFT has no quasiparticle description and so has not proven amenable to any simple truncation; concretely, the spectral function does not have quasiparticle poles. With that said, it seems worthwhile to search for relations between known effective truncations in quantum many-body and holographic models.

(v) What predictions do holography make for generic strongly coupled systems?
Once we let go of the idea that our models must be defined by traditional quasiparticle effective QFTs, we can start asking general questions about the strongly coupled QFTs illuminated by our holographic examples. What common properties do these theories share? What general lessons can we draw about their kinematics, dynamics, phase structures, critical behavior, equilibration, etc? For example, we already know to expect nearly inviscid hydrodynamic behavior at finite density and low temperature. What else should we expect in general? Perhaps most importantly, what lessons can we draw for systems we really care about?

(vi) What is the holographic dual of a non-relativistic QFT?
As discussed in section 4.3.4, non-trivial progress has been made in the study of non-relativistic QFTs and many-body systems. This has led to some simple heuristic lessons, for example that we should expect the viscosity to entropy ratio to again be extremely low and that the *Bertsch parameter*[73] should be expected to be of the order of unity and remain relatively insensitive to the coupling constant near strong coupling. However, the state of current models is rather poor, with awkward thermodynamics and bizarre phase structure. What are we missing? Can we build a holographic dual in the same universality class as fermions at unitarity?

(vii) How (and why) is quantum information encoded in the bulk geometry?
Holographically, information about the thermodynamic state of the boundary QFT is encoded in the detailed geometry and matter field backgrounds of the bulk theory. For example, the temperature and chemical potential are encoded in the mass and charge of the bulk black hole. Similarly, when the QFT is in a pure state, the bulk geometry should encode the quantum correlations of the boundary QFT at zero temperature. But how?

One approach to this question is the Ryu–Takanayagi conjecture [448, 449], which states that the entanglement entropy associated with a region R in a holographic QFT is given by the area of the minimal-area surface in the bulk which ends at the boundary of the region R. Holographically, this conjecture matches the relation between the thermodynamic entropy of a black hole and the area of its horizon. Interestingly, it suggests an intimate connection between quantum and thermal entropy—and more generally, between quantum and thermal fluctuations—all of which are encoded by the physics of horizons in the holographic dual.

While a general proof is still lacking, considerable evidence has been found for this conjecture, and conversely, numerous predictions have been made for the entanglement entropy of strongly coupled many-body systems by exploiting this conjecture (see, e.g., [450–455] and references therein). For example, in this focus issue, Albash and Johnson [456] have studied the evolution of the Ryu–Takanayagi entropy under a thermal quench, leading both to evidence supporting the conjecture and to predictions on the evolution of the entanglement entropy.

---

[73] We defined the Bertsch parameter of the unitary Fermi gas in section 2.2. Here we use the term to refer to the ratio of the ground-state energy at strong coupling to that at weak coupling in a general non-relativistic CFT.

(viii) How does spatial inhomogeneity change the story?
As we have seen, holography translates quantum computations in the boundary QFT (i.e. Feynman diagrams and path integrals) into classical GR calculations in the bulk (i.e. solving nonlinear partial differential equation boundary value problems). As such, it is clear why most research to date has focused on homogeneous, isotropic problems. However, on the quantum many-body side, we know that inhomogeneity can radically alter the dynamics, and that the precise way they do so provides a detailed probe of the physics driving those dynamics. The same is true of anisotropy. Moreover, much of what we know about inhomogeneity and anisotropy in QFT comes from some form of perturbation theory. It would thus be particularly interesting to study the dynamics of holographic systems in the presence of inhomogeneity and anisotropy. Will we find evidence for many-body localization? Do $(2+1)$D QFTs at strong coupling and in the hydrodynamic regime display a canonical turbulent cascade, or an inverse cascade as is usually, but not universally, expected in two spatial dimensions?

Many of the issues discussed in sections 5.1–5.3 are addressed in the contributions to this focus issue, and we hope that this summary, as well as the insights contained in these contributions, will spur further advances in research on the intriguing subject of strongly correlated quantum fluids.

Acknowledgments

This work was supported in part by the US Department of Energy under cooperative research agreement numbers DE-FG02-03ER41260 (TS) and DE-FG02-05ER-41360 (AA) and by the Division of Materials Science and Engineering, the Office of Basic Energy Sciences, Office of Science under grant number DE-SC-0002712 (JET); by the US National Science Foundation under grant numbers PHY-0547845, PHY-0903457 and PHY-1067973 (LDC) and PHY-1067873 (JET); by the Army Research Office under grant number W911NF-11-1-0420 (JET); and by the US Air Force Office of Scientific Research under grant numbers FA9550-10-1-0072 (JET) and FA9550-11-1-0224 (LDC). We developed and wrote this review in part during the Aspen Center for Physics workshops on Quantum Simulation/Computation with Cold Atoms and Molecules (LDC), String Duals of Finite Temperature and Low-Dimensional Systems (AA) and Critical Behavior of Lattice Models in Condensed Matter and Particle Physics (LDC); at the Kavli Institute for Theoretical Physics workshops on Beyond Standard Optical Lattices (LDC), Disentangling Quantum Many-Body Systems: Computational and Conceptual Approaches (LDC) and Holographic Duality and Condensed Matter Physics (AA); and at the Heidelberg Center for Quantum Dynamics (LDC). We thank Joaquin Drut for providing some of the data in figure 2. We thank Aurel Bulgac, Chenglin Cao, Oliver DeWolfe, Ethan Dyer, Ethan Elliot, James Joseph, Kathy Levin, John McGreevy, Jessie Petricka, Michael Wall, Haibin Wu and Martin Zwierlein for useful discussions.

References

[1] Linstrom P J and Mallard W G 2011 *NIST Chemistry WebBook* (*NIST Standard Reference Database* vol 69) (Gaithersburg, MD: National Institute of Standards and Technology)

[2] Cao C, Elliott E, Wu H and Thomas J E 2011 Searching for perfect fluids: quantum viscosity in a universal Fermi gas *New J. Phys.* **13** 075007

[3] Song H, Bass S A, Heinz U, Hirano T and Shen C 2011 Hadron spectra and elliptic flow for 200 A GeV Au + Au collisions from viscous hydrodynamics coupled to a Boltzmann cascade *Phys. Rev.* C **83** 054910

[4] Meyer H B 2007 A calculation of the shear viscosity in SU(3) gluodynamics *Phys. Rev.* D **76** 101701

[5] Wlazlowski G, Magierski P and Drut J E 2012 Shear viscosity of a unitary Fermi gas *Phys. Rev. Lett.* **109** 020406

[6] Arnold P B, Moore G D and Yaffe L G 2000 Transport coefficients in high temperature gauge theories 1. Leading log results *J. High Energy Phys.* JHEP11(2000)001

[7] Prakash M, Prakash M, Venugopalan R and Welke G 1993 Nonequilibrium properties of hadronic mixtures *Phys. Rep.* **227** 321

[8] Massignan P, Bruun G M and Smith H 2005 Viscous relaxation and collective oscillations in a trapped Fermi gas near the unitarity limit *Phys. Rev.* A **71** 033607

[9] Mannarelli M, Manuel C and Tolos L 2012 Shear viscosity in a superfluid cold Fermi gas at unitarity arXiv:1201.4006 [cond-mat.quant-gas]

[10] Brigante M, Liu H, Myers R C, Shenker S and Yaida S 2008 The viscosity bound and causality violation *Phys. Rev. Lett.* **100** 191601

[11] Kovtun P, Son D T and Starinets A O 2005 Viscosity in strongly interacting quantum field theories from black hole physics *Phys. Rev. Lett.* **94** 111601

[12] Csernai L P, Kapusta J I and McLerran L D 2006 On the strongly-interacting low-viscosity matter created in relativistic nuclear collisions *Phys. Rev. Lett.* **97** 152303

[13] Lacey R A, Ajitanand N N, Alexander J M, Chung P, Holzmann W G, Issah M, Taranenko A, Danielewicz P and Stöcker H 2007 Has the QCD critical point been signaled by observations at RHIC? *Phys. Rev. Lett.* **98** 092301

[14] Dalfovo F, Giorgini S, Pitaevskii L P and Stringari S 1999 Theory of Bose–Einstein condensation in trapped gases *Rev. Mod. Phys.* **71** 463

[15] Leggett A J 2001 Bose–Einstein condensation in the alkali gases: some fundamental concepts *Rev. Mod. Phys.* **73** 307

[16] Lewenstein A, Sanpera M, Ahufinger V, Damski B, Sen De A and Sen U 2007 Ultracold atomic gases in optical lattices: mimicking condensed matter physics and beyond *Adv. Phys.* **56** 243

[17] Bloch I, Dalibard J and Zwerger W 2008 Many-body physics with ultracold gases *Rev. Mod. Phys.* **80** 885

[18] Giorgini S, Pitaevskii L P and Stringari S 2008 Theory of ultracold atomic Fermi gases *Rev. Mod. Phys.* **80** 1215

[19] Carr L D, Demille D, Krems R V and Ye J 2009 Cold and ultracold molecules: science, technology, and applications *New J. Phys.* **11** 055049

[20] Schäfer T and Teaney D 2009 Nearly perfect fluidity: from cold atomic gases to hot quark gluon plasmas *Rep. Prog. Phys.* **72** 126001

[21] Ketterle W and Zwierlein M W 2008 Making, probing and understanding ultracold Fermi gases *Nuovo Cimento* **31** 247–422

[22] Zwerger W (ed) 2011 *The BCS–BEC Crossover and the Unitary Fermi Gas* (Berlin: Springer)

[23] DeMarco B and Jin D S 1999 Onset of Fermi degeneracy in a trapped atomic gas *Science* **285** 1703

[24] Truscott A G, Strecker K E, McAlexander W I, Partridge G and Hulet R G 2001 Observation of Fermi pressure in a gas of trapped atoms *Science* **291** 2570

[25] Schreck F, Khaykovich L, Corwin K L, Ferrari G, Bourdel T, Cubizolles J and Salomon C 2001 Quasipure Bose–Einstein condensate immersed in a Fermi sea *Phys. Rev. Lett.* **87** 080403

[26] Granade S R, Gehm M E, O'Hara K M and Thomas J E 2002 All-optical production of a degenerate Fermi gas *Phys. Rev. Lett.* **88** 120405

[27] Hadzibabic Z, Gupta S, Stan C A, Schunck C H, Zwierlein M W, Dieckmann K and Ketterle W 2003 Fifty-fold improvement in the number of quantum degenerate fermionic atoms *Phys. Rev. Lett.* **91** 160401

[28] McNamara J M, Jeltes T, Tychkov A S, Hogervorst W and Vassen W 2006 Degenerate Bose–Fermi mixture of metastable atoms *Phys. Rev. Lett.* **97** 080404

[29] Fukuhara T, Takasu Y, Kumakura M and Takahashi Y 2007 Degenerate Fermi gases of ytterbium *Phys. Rev. Lett.* **98** 030401

[30] DeSalvo B J, Yan M, Mickelson P G, Martinez de Escobar Y N and Killian T C 2010 Degenerate Fermi gas of ^{87}Sr *Phys. Rev. Lett.* **105** 030402

[31] Lu M, Burdick N Q and Lev B L 2012 Quantum degenerate dipolar Fermi gas *Phys. Rev. Lett.* **108** 215301

[32] Chin C, Grimm R, Julienne P and Tiesinga E 2010 Feshbach resonances in ultracold gases *Rev. Mod. Phys.* **82** 1225–86

[33] Sá de Melo C A R 2008 When fermions become bosons: pairing in ultracold gases *Phys. Today* **61** 45–51

[34] Leggett A J 1975 A theoretical description of the new phases of liquid ^{3}He *Rev. Mod. Phys.* **47** 331–414

[35] Chen Q, Stajic J, Tan S and Levin K 2005 BCS–BEC crossover: from high temperature superconductors to ultracold superfluids *Phys. Rep.* **412** 1–88

[36] Ku M J H, Sommer A T, Cheuk L W and Zwierlein M W 2012 Revealing the superfluid lambda transition in the universal thermodynamics of a unitary Fermi gas *Science* **335** 563

[37] Gorkov L P and Melik-Barkhudarov T K 1962 Contribution to the theory of superfluidity in an imperfect Fermi gas *Sov. Phys.—JETP* **13** 1018–22

[38] Timmermans E 2001 Degenerate fermion gas heating by hole creation *Phys. Rev. Lett.* **87** 240403

[39] Carr L D, Bourdel T and Castin Y 2004 Limits of sympathetic cooling of fermions by zero-temperature bosons due to particle losses *Phys. Rev.* A **69** 033603

[40] Savard T A, O'Hara K M and Thomas J E 1997 Laser-noise-induced heating in far-off resonance optical traps *Phys. Rev.* A **56** 1095

[41] Tiesinga E, Verhaar B J and Stoof H T C 1993 Threshold and resonance phenomena in ultracold ground-state collisions *Phys. Rev.* A **47** 4114

[42] O'Hara K M, Hemmer S L, Granade S R, Gehm M E, Thomas J E, Venturi V, Tiesinga E and Williams C J 2002 Measurement of the zero crossing in a Feshbach resonance of fermionic ^{6}Li *Phys. Rev.* A **66** 041401

[43] Jochim S, Bartenstein M, Hendl G, Denschlag J H and Grimm R 2002 Magnetic field control of elastic scattering in a cold gas of fermionic lithium atoms *Phys. Rev. Lett.* **89** 273202

[44] Regal C A and Jin D S 2003 Measurement of positive and negative scattering lengths in a Fermi gas of atoms *Phys. Rev. Lett.* **90** 230404

[45] O'Hara K M, Granade S R, Gehm M E, Savard T A, Bali S, Freed C and Thomas J E 1999 Ultrastable CO_2 laser trapping of lithium Fermions *Phys. Rev. Lett.* **82** 4204

[46] O'Hara K M, Hemmer S L, Gehm M E, Granade S R and Thomas J E 2002 Observation of a strongly interacting degenerate Fermi gas of atoms *Science* **298** 2179

[47] Petrov D S, Salomon C and Shlyapnikov G V 2004 Weakly bound dimers of fermionic atoms *Phys. Rev. Lett.* **93** 090404

[48] Gehm M E, Hemmer S L, Granade S R, O'Hara K M and Thomas J E 2003 Mechanical stability of a strongly interacting Fermi gas of atoms *Phys. Rev.* A **68** 011401

[49] Bourdel T, Cubizolles J, Khaykovich L, Magālhaes K M F, Kokkelmans S J J M F, Shlyapnikov G V and Salomon C 2003 Measurement of the interaction energy near a Feshbach resonance in a 6Li Fermi gas *Phys. Rev. Lett.* **91** 020402

[50] Cubizolles J, Bourdel T, Kokkelmans S J J M F, Shlyapnikov G V and Salomon C 2003 Production of long-lived ultracold Li_2 molecules from a Fermi gas *Phys. Rev. Lett.* **91** 240401

[51] Jochim S, Bartenstein M, Altmeyer A, Hendl G, Riedl S, Chin C, Denschlag J H and Grimm R 2003 Bose–Einstein condensation of molecules *Science* **302** 2101

[52] Greiner M, Regal C A and Jin D S 2003 Emergence of a molecular Bose–Einstein condensate from a Fermi gas *Nature* **426** 537

[53] Zwierlein M W, Stan C A, Schunck C H, Raupach S M F, Gupta S, Hadzibabic Z and Ketterle W 2003 Observation of Bose–Einstein condensation of molecules *Phys. Rev. Lett.* **91** 250401

[54] Regal C A, Greiner M and Jin D S 2004 Observation of resonance condensation of fermionic atom pairs *Phys. Rev. Lett.* **92** 040403

[55] Zwierlein M W, Stan C A, Schunck C H, Raupach S M F, Kerman A J and Ketterle W 2004 Condensation of pairs of fermionic atoms near a Feshbach resonance *Phys. Rev. Lett.* **92** 120403

[56] Kinast J, Hemmer S L, Gehm M E, Turlapov A and Thomas J E 2004 Evidence for superfluidity in a resonantly interacting Fermi gas *Phys. Rev. Lett.* **92** 150402

[57] Bartenstein M, Altmeyer A, Riedl S, Jochim S, Chin C, Denschlag J H and Grimm R 2004 Crossover from a molecular Bose–Einstein condensate to a degenerate Fermi gas *Phys. Rev. Lett.* **92** 120401

[58] Chin C, Bartenstein M, Altmeyer A, Riedl S, Jochim S, Denschlag J H and Grimm R 2004 Observation of the pairing gap in a strongly interacting Fermi gas *Science* **305** 1128

[59] Zwierlein M W, Abo-Shaeer J R, Schirotzek A, Schunck C H and Ketterle W 2005 Vortices and superfluidity in a strongly interacting Fermi gas *Nature* **435** 1047–51

[60] Kinast J, Turlapov A, Thomas J E, Chen Q, Stajic J and Levin K 2005 Heat capacity of a strongly interacting Fermi gas *Science* **307** 1296

[61] Ho T-L 2004 Universal thermodynamics of degenerate quantum gases in the unitarity limit *Phys. Rev. Lett.* **92** 090402

[62] Thomas J E, Turlapov A and Kinast J 2005 Virial theorem and universality in a unitary Fermi gas *Phys. Rev. Lett.* **95** 120402

[63] Carr L D, Shlyapnikov G V and Castin Y 2004 Achieving a BCS transition in an atomic Fermi gas *Phys. Rev. Lett.* **92** 150404

[64] Carr L D, Chiaramonte R and Holland M J 2004 Thermodynamics of an atom–molecule mixture *Phys. Rev.* A **70** 043609

[65] Luo L, Clancy B, Joseph J, Kinast J and Thomas J E 2007 Measurement of the entropy and critical temperature of a strongly interacting Fermi gas *Phys. Rev. Lett.* **98** 080402

[66] Luo L and Thomas J E 2009 Thermodynamic measurements in a strongly interacting Fermi gas *J. Low Temp. Phys.* **154** 1

[67] Horikoshi M, Nakajima S, Ueda M and Mukaiyama T 2010 Measurement of universal thermodynamic functions for a unitary Fermi gas *Science* **327** 442

[68] Navon N, Nascimbène S, Chevy F and Salomon C 2010 The equation of state of a low-temperature Fermi gas with tunable interactions *Science* **328** 5979

[69] Ho T L and Zhou Q 2010 Obtaining the phase diagram and thermodynamic quantities of bulk systems from the densities of trapped gases *Nature Phys.* **6** 131

[70] Hazzard K R A and Mueller E J 2011 Techniques to measure quantum criticality in cold atoms *Phys. Rev.* A **84** 013604

[71] Liu X-J, Hu H and Drummond P D 2011 Universal contact of strongly interacting fermions at finite temperatures *New J. Phys.* **13** 035007

[72] Altman E, Demler E and Lukin M D 2004 Probing many-body states of ultra-cold atoms via noise correlations *Phys. Rev.* A **70** 013603

[73] Greiner M, Regal C A, Stewart J T and Jin D S 2005 Probing pair-correlated fermionic atoms through correlations in atom shot noise *Phys. Rev. Lett.* **94** 110401

[74] Hohenberg P and Kohn W 1964 Inhomogeneous electron gas *Phys. Rev.* B **136** 864

[75] Baker M A 2001 The MBX challenge competition: a neutron matter model *Int. J. Mod. Phys.* B **15** 1314–20

[76] Thomas J E, Kinast J and Turlapov A 2005 Virial theorem and universality in a unitary Fermi gas *Phys. Rev. Lett.* **95** 120402

[77] Werner F 2008 Virial theorems for trapped cold atoms *Phys. Rev.* A **78** 025601

[78] Castin Y and Werner F 2011 The unitary gas and its symmetry properties *Crossover BEC–BCS and the Unitary Fermi Gas* (*Springer Lecture Notes in Physics*) ed W Zwerger (Berlin: Springer) pp 127–91

[79] Tan S 2008 Energetics of a strongly correlated Fermi gas *Ann. Phys.* **323** 2952

[80] Tan S 2008 Large momentum part of a strongly correlated Fermi gas *Ann. Phys.* **323** 2971

[81] Tan S 2008 Generalized virial theorem and pressure relation for a strongly correlated Fermi gas *Ann. Phys.* **323** 2987

[82] Braaten E 2012 Universal relations for Fermions with large scattering length *Lect. Notes Phys.* **836** 193

[83] Kuhnle E D, Hoinka S, Hu H, Dyke P, Hannaford P and Vale C J 2011 Studies of the universal contact in a strongly interacting Fermi gas using Bragg spectroscopy *New J. Phys.* **13** 055010
[84] Ho T-L and Zhou Q 2010 Obtaining the phase diagram and thermodynamic quantities of bulk systems from the densities of trapped gases *Nature Phys.* **6** 131
[85] Riedl S, Sánchez Guajardo E R, Kohstall C, Denschlag J H and Grimm R 2011 Superfluid quenching of the moment of inertia in a strongly interacting Fermi gas *New J. Phys.* **13** 035003
[86] Lee T D and Yang C N 1957 Many-body problem in quantum mechanics and quantum statistical mechanics *Phys. Rev.* **105** 1119
[87] Huang K and Yang C N 1957 Quantum-mechanical many-body problem with hard-sphere interaction *Phys. Rev.* **105** 767
[88] Uhlenbeck G E and Beth E 1936 The quantum theory of the non-ideal gas I. Deviations from the classical theory *Physica* **3** 729
[89] Ho T-L and Mueller E J 2004 High temperature expansion applied to fermions near Feshbach resonance *Phys. Rev. Lett.* **92** 160404
[90] Nozières P and Schmitt-Rink S 1985 Bose–Einstein condensation in an attractive fermion gas: from weak to strong coupling superconductivity *J. Low Temp. Phys.* **59** 195–212
[91] Perali A, Pieri P, Pisani L and Strinati G C 2004 BCS–BEC crossover at finite temperature for superfluid trapped Fermi atoms *Phys. Rev. Lett.* **92** 220404
[92] Fukushima N, Ohashi Y, Taylor E and Griffin A 2007 Superfluid density and condensate fraction in the BCS–BEC crossover regime at finite temperatures *Phys. Rev.* A **75** 033609
[93] Chena Q, Stajic J, Tan S and Levin K 2005 BCS–BEC crossover: from high temperature superconductors to ultracold superfluids *Phys. Rep.* **412** 1
[94] Haussmann R, Rantner W, Cerrito S and Zwerger W 2007 Thermodynamics of the BCS–BEC crossover *Phys. Rev.* A **75** 023610
[95] Nikolic P and Sachdev S 2007 Renormalization group fixed points, universal phase diagram, and $1/N$ expansion for quantum liquids with interactions near the unitarity limit *Phys. Rev.* A **75** 033608
[96] Nussinov Z and Nussinov S 2006 Triviality of the BCS–BEC crossover in extended dimensions: implications for the ground state energy *Phys. Rev.* A **74** 053622
[97] Nishida Y and Son D T 2006 An epsilon expansion for Fermi gas at infinite scattering length *Phys. Rev. Lett.* **97** 050403
[98] Carlson J, Chang S-Y, Pandharipande V R and Schmidt K E 2003 Superfluid Fermi gases with large scattering length *Phys. Rev. Lett.* **91** 050401
[99] Astrakharchik G E, Boronat J, Casulleras J and Giorgini S 2004 Equation of state of a Fermi gas in the BEC–BCS crossover: a quantum Monte Carlo study *Phys. Rev. Lett.* **93** 200404
[100] Bulgac A, Drut J E and Magierski P 2006 Spin 1/2 fermions on a 3D-lattice in the unitary regime at finite temperatures *Phys. Rev. Lett.* **96** 090404
[101] Lee D and Schäfer T 2006 Cold dilute neutron matter on the lattice II. Results in the unitary limit *Phys. Rev.* C **73** 015202
[102] Goulko O and Wingate M 2010 Thermodynamics of balanced and slightly spin-imbalanced Fermi gases at unitarity *Phys. Rev.* A **82** 053621
[103] Drut J E, Lahde T A, Wlazlowski G and Magierski P 2012 The equation of state of the unitary Fermi gas: an update on lattice calculations *Phys. Rev.* A **85** 051601(R)
[104] Van Houcke K, Werner F, Kozik E, Prokofev N, Svistunov B, Ku M, Sommer A, Cheuk L W, Schirotzek A and Zwierlein M W 2012 Feynman diagrams versus Feynman quantum emulator *Nature Phys.* **8** 366
[105] Stewart J T, Gaebler J P, Regal C A and Jin D S 2006 The potential energy of a ^{40}K Fermi gas in the BCS-BEC crossover *Phys. Rev. Lett.* **97** 220406
[106] Hu H, Drummond P D and Liu X-J 2007 Universal thermodynamics of strongly interacting Fermi gases *Nature Phys.* **3** 469
[107] Dribinski V, Ossadtchi A, Mandelshtam V A and Reisler H 2002 Reconstruction of Abel-transformable images: the Gaussian basis-set expansion Abel transform method *Rev. Sci. Instrum.* **73** 2634

[108] Forbes M M, Gandolfi S and Gezerlis A 2011 Resonantly interacting fermions in a box *Phys. Rev. Lett.* **106** 235303

[109] Endres M G, Kaplan D B, Lee J-W and Nicholson A N 2012 Lattice Monte Carlo calculations for unitary fermions in a finite box arXiv:1203.3169 [hep-lat]

[110] Chien C-C, Guo H, Wulin D and Levin K 2011 Perfect fluids and bad metals: insights from ultracold Fermi gases *New J. Phys.* **13** 075011

[111] Kinast J, Turlapov A and Thomas J E 2004 Breakdown of hydrodynamics in the radial breathing mode of a strongly-interacting Fermi gas *Phys. Rev.* A **70** 051401

[112] Bartenstein M, Altmeyer A, Riedl S, Jochim S, Chin C, Denschlag J H and Grimm R 2004 Collective excitations of a degenerate gas at the BEC–BCS crossover *Phys. Rev. Lett.* **92** 203201

[113] Kinast J, Turlapov A and Thomas J E 2005 Damping of a unitary Fermi gas *Phys. Rev. Lett.* **94** 170404

[114] Altmeyer A, Riedl S, Kohstall C, Wright M, Geursen R, Bartenstein M, Chin C, Denschlag J H and Grimm R 2007 Precision measurements of collective oscillations in the BEC–BCS crossover *Phys. Rev. Lett.* **98** 040401

[115] Cao C, Elliott E, Joseph J, Wu H, Petricka J, Schäfer T and Thomas J E 2011 Universal quantum viscosity in a unitary Fermi gas *Science* **331** 58

[116] Joseph J, Clancy B, Luo L, Kinast J, Turlapov A and Thomas J E 2007 Sound propagation in a Fermi gas near a Feshbach resonance *Phys. Rev. Lett.* **98** 170401

[117] Clancy B, Luo L and Thomas J E 2007 Observation of nearly perfect irrotational flow in normal and superfluid strongly interacting Fermi gases *Phys. Rev. Lett.* **99** 140401

[118] Stringari S 2004 Collective oscillations of a trapped superfluid Fermi gas near a Feshbach resonance *Europhys. Lett.* **65** 749

[119] Bulgac A and Bertsch G F 2005 Collective oscillations of a trapped Fermi gas near the unitary limit *Phys. Rev. Lett.* **94** 070401

[120] Son D T 2007 Vanishing bulk viscosities and conformal invariance of unitary Fermi gas *Phys. Rev. Lett.* **98** 020604

[121] Dusling K and Schäfer T 2011 Elliptic flow of the dilute Fermi gas: from kinetics to hydrodynamics *Phys. Rev.* A **84** 013622

[122] Schäfer T 2007 The shear viscosity to entropy density ratio of trapped fermions in the unitarity limit *Phys. Rev.* A **76** 063618

[123] Turlapov A, Kinast J, Clancy B, Luo L, Joseph J and Thomas J E 2008 Is a gas of strongly interacting atomic fermions a nearly perfect fluid? *J. Low Temp. Phys.* **150** 567

[124] Bruun G M and Smith H 2005 Viscosity and thermal relaxation for a resonantly interacting Fermi gas *Phys. Rev.* A **72** 043605

[125] Maxwell J C 1866 On the viscosity or internal friction of air and other gases *Proc. R. Soc. Lond.* **XV** 14

[126] Bruun G M and Smith H 2007 Frequency and damping of the scissors mode of a Fermi gas *Phys. Rev.* A **76** 045602

[127] Schäfer T and Chafin C 2012 Scaling flows and dissipation in the dilute Fermi gas at unitarity *Lect. Notes Phys.* **836** 375

[128] Chao J, Braby M and Schäfer T 2011 Viscosity spectral functions of the dilute Fermi gas in kinetic theory *New J. Phys.* **13** 035014

[129] Enss T, Haussmann R and Zwerger W 2011 Viscosity and scale invariance in the unitary Fermi gas *Ann. Phys.* **326** 770–96

[130] LeClair A 2011 On the viscosity-to-entropy density ratio for unitary Bose and Fermi gases *New J. Phys.* **13** 055015

[131] Khalatnikov I M 1965 *Introduction to the Theory of Superfluidity* (New York: Benjamin)

[132] Taylor E, Hu H, Liu X-J, Pitaevskii L P, Griffin A and Stringari S 2009 First and second sound in a strongly interacting Fermi gas *Phys. Rev.* A **80** 053601

[133] Rupak G and Schäfer T 2007 Shear viscosity of a superfluid Fermi gas in the unitarity limit *Phys. Rev.* A **76** 053607

[134] Sommer A, Ku M, Roati G and Zwierlein M W 2011 Universal spin transport in a strongly interacting Fermi gas *Nature* **472** 201

[135] Ku M, Sommer A and Zwierlein M W 2011 Spin transport in polaronic and superfluid Fermi gases *New J. Phys.* **13** 055009

[136] Duine R A, Polini M, Raoux A, Stoof H T C and Vignale G 2011 Spin drag in ultracold Fermi mixtures with repulsive interactions *New J. Phys.* **13** 045010

[137] Müller M and Nguyen H C 2011 Collision-dominated spin transport in graphene and Fermi liquids *New J. Phys.* **13** 035009

[138] Bruun G M 2011 Spin diffusion in Ffermi gases *New J. Phys.* **13** 035005

[139] Bruun G M and Pethick C J 2011 Spin diffusion in trapped clouds of strongly interacting cold atoms *Phys. Rev. Lett.* **107** 255302

[140] Köhl M, Moritz H, Stöferle T, Günter K and Esslinger T 2005 Fermionic atoms in a three dimensional optical lattice: observing Fermi surfaces, dynamics and interactions *Phys. Rev. Lett.* **94** 080403

[141] Leggett A J 2006 What do we know about high T_c? *Nature Phys.* **2** 134

[142] Eagles D M 1969 Possible pairing without superconductivity at low carrier concentrations in bulk and thin-film superconducting semiconductors *Phys. Rev.* **186** 456–63

[143] Leggett A J 1980 *Modern Trends in the theory of Condensed Matter* (Berlin: Springer)

[144] Peter H 2010 Büchler microscopic derivation of Hubbard parameters for cold atomic gases *Phys. Rev. Lett.* **104** 090402

[145] von Stecher J, Gurarie V, Radzihovsky L and Rey A M 2011 Lattice-induced resonances in one-dimensional bosonic systems *Phys. Rev. Lett.* **106** 235301

[146] Hubbard J 1963 Electron correlations in narrow energy bands *Proc. R. Soc. Lond.* A **276** 238

[147] Anderson P W 2002 Superconductivity in high T_c cuprates: the cause is no longer a mystery *Phys. Scr.* **T102** 10

[148] Timmermans E, Tommasini P, Hussein M and Kerman A 1999 Feshbach resonances in atomic Bose–Einstein condensates *Phys. Rep.* **315** 199–230

[149] Holland M, Kokkelmans S J J M F, Chiofalo M L and Walser R 2001 Resonance superfluidity in a quantum degenerate Fermi gas *Phys. Rev. Lett.* **87** 120406

[150] Kokkelmans S J J M F and Holland M J 2002 Ramsey fringes in a Bose–Einstein condensate between atoms and molecules *Phys. Rev. Lett.* **89** 180401

[151] Milstein J N, Kokkelmans S J J M F and Holland M J 2002 Resonance theory of the crossover from Bardeen–Cooper–Schrieffer superfluidity to Bose–Einstein condensation in a dilute Fermi gas *Phys. Rev.* A **66** 043604

[152] Ranninger J and Robaszkiewicz S 1985 Superconductivity of locally paired electrons *Physica* B **135** 468

[153] Ranninger J and Robin J M 1995 The boson–fermion model of high-T_c superconductivity doping dependence *Physica* C **253** 279

[154] Yang K-Y, Kozik E, Wang X and Troyer M 2011 Diagrammatic quantum Monte Carlo solution of the two-dimensional cooperon–fermion model *Phys. Rev.* B **83** 214516

[155] Carr L D and Holland M J 2005 Quantum phase transitions in the Fermi–Bose Hubbard model *Phys. Rev.* A **72** 031604(R)

[156] Zhou F 2005 Mott states under the influence of fermion–boson conversion *Phys. Rev.* B **72** 220501(R)

[157] Dickerscheid D B M, van Oosten D, Tillema E J and Stoof H T C 2005 Quantum phases in a resonantly interacting boson–fermion mixture *Phys. Rev. Lett.* **94** 230404

[158] Diener R B and Ho T-L 2006 Fermions in optical lattices swept across Feshbach resonances *Phys. Rev. Lett.* **96** 010402

[159] Duan L-M 2005 Effective Hamiltonian for fermions in an optical lattice across a Feshbach resonance *Phys. Rev. Lett.* **95** 243202

[160] Kestner J P and Duan L-M 2010 Effective single-band models for strongly interacting fermions in an optical lattice *Phys. Rev.* A **81** 043618

[161] Wall M L and Carr L D 2012 Microscopic model for Feshbach interacting fermions in an optical lattice with arbitrary scattering length and resonance width *Phys. Rev. Lett.* **109** 055302

[162] Spuntarelli P P A, Carr L D and Strinati G C 2011 Gray solitons in a strongly interacting superfluid Fermi gas *New J. Phys.* **13** 035010

[163] Bolech C J, Baksmaty L O, Lu H and Pu H 2011 A Bogoliubov–de Gennes study of trapped spin-imbalanced unitary Fermi gases *New J. Phys.* **13** 055014

[164] Nowling S, Keränen V, Keski-Vakkuri E and Yogendran K P 2011 Solitons as probes of the structure of holographic superfluids *New J. Phys.* **13** 065003

[165] Scott R G, Dalfovo F, Pitaevskii L P and Stringari S 2011 Dynamics of dark solitons in a trapped superfluid Fermi gas *Phys. Rev. Lett.* **106** 185301

[166] Liao R and Brand J 2011 Traveling dark solitons in superfluid Fermi gases *Phys. Rev.* A **83** 041604

[167] Scott R G, Dalfovo F, Pitaevskii L P, Stringari S, Fialko O, Liao R and Brand J 2012 The decay and collisions of dark solitons in superfluid Fermi gases *New J. Phys.* **14** 023044

[168] Nascimbène S, Navon K J, Jiang N, Chevy F and Salomon C 2010 Exploring the thermodynamics of a universal Fermi gas *Nature* **463** 1057

[169] Liao Y-A, Rittner A S C, Paprotta T, Li W, Partridge G B, Hulet R G, Baur S K and Mueller E J 2010 Spin-imbalance in a one-dimensional Fermi gas *Nature* **467** 567

[170] Sheehy D E 2011 Spin-selective scatterers as a probe of pairing in a one-dimensional interacting fermion gas *New J. Phys.* **13** 035015

[171] Sadeghzadeh K, Bruun G M, Lobo C, Massignan P and Recati A 2011 Metastability in spin-polarized Fermi gases and quasiparticle decays *New J. Phys.* **13** 055011

[172] Han L and Sá de Melo C A R 2011 Evolution from Bardeen–Cooper–Schrieffer to Bose–Einstein condensate superfluidity in the presence of disorder *New J. Phys.* **13** 055012

[173] Zhang X, Hung C-L, Tung S-K, Gemelke N and Chin C 2011 Exploring quantum criticality based on ultracold atoms in optical lattices *New J. Phys.* **13** 045011

[174] Shlyapnikov G V and Tsvelik A M 2011 Polar phase of one-dimensional bosons with large spin *New J. Phys.* **13** 065012

[175] Braaten E and Hammer H W 2006 Universality in few-body systems with large scattering length *Phys. Rep.* **428** 259–390

[176] Wang Y and Esry B D 2011 Universal three-body physics at finite energy near Feshbach resonances *New J. Phys.* **13** 035025

[177] O'Hara K M 2011 Realizing analogues of color superconductivity with ultracold alkali atoms *New J. Phys.* **13** 065011

[178] Kinnunen J J, Nummi O H T and Törmä P 2011 Coexistence of pairing gaps in three-component Fermi gases *New J. Phys.* **13** 055013

[179] Titvinidze I, Privitera A, Chang S-Y, Diehl S, Baranov M A, Daley A and Hofstetter W 2011 Magnetism and domain formation in SU(3)-symmetric multi-species Fermi mixtures *New J. Phys.* **13** 035013

[180] Alford M G, Schmitt A, Rajagopal K and Schäfer T 2008 Color superconductivity in dense quark matter *Rev. Mod. Phys.* **80** 1455

[181] Shuryak E 2009 Physics of strongly coupled quark–gluon plasma *Prog. Part. Nucl. Phys.* **62** 48

[182] Heinz U W 2009 Early collective expansion: relativistic hydrodynamics and the transport properties of QCD matter *Relativistic Heavy Ion Physics* (*Landolt–Boernstein New Series* vol 23) ed R Stock (New York: Springer) chapter 5

[183] Braun-Munzinger P and Wambach J 2009 Colloquium: phase diagram of strongly interacting matter *Rev. Mod. Phys.* **81** 1031–50

[184] Casalderrey-Solana J, Liu H, Mateos D, Rajagopal K and Wiedemann U A 2011 Gauge/string duality, hot and heavy QCD ion collisions arXiv:1101.0618 [hep-th]

[185] Arsenne I *et al* (BRAHMS), Back B *et al* (PHOBOS), Adcox K *et al* (PHENIX) and Adams J *et al* (STAR) 2005 First three years of operation of RHIC *Nucl. Phys.* A **757** 1–283

[186] Steinberg P 2012 First results from the LHC heavy ion program *New J. Phys.* **14** 035006

[187] Snellings R 2011 Elliptic flow: brief a review *New J. Phys.* **13** 055008

[188] Lisa M A, Frodermann E, Graef G, Mitrovski M, Mount E, Petersen H and Bleicher M 2011 Shape analysis of strongly interacting systems: the heavy ion case *New J. Phys.* **13** 065006

[189] Nagle J L, Bearden I G and Zajc W A 2011 Quark–gluon plasma at RHIC and the LHC: perfect fluid too perfect? *New J. Phys.* **13** 075004

[190] Müller B, Schukraft J and Wyslouch B 2012 First results from Pb + Pb collisions at the LHC *Annu. Rev. Nucl. Part. Sci.* **62** 361–86 (arXiv:1202.3233 [hep-ex])

[191] Coleman S R and Weinberg E J 1973 Radiative corrections as the origin of spontaneous symmetry breaking *Phys. Rev.* D **7** 1888–910

[192] Gross D J and Wilczek F 1973 Ultraviolet behavior of nonabelian gauge theories *Phys. Rev. Lett.* **30** 1343–6

[193] Politzer H D 1973 Reliable perturbative results for strong interactions? *Phys. Rev. Lett.* **30** 1346–9

[194] Nakamura K *et al* (Particle Data Group Collaboration) 2010 Review of particle physics *J. Phys. G: Nucl. Part. Phys.* **37** 075021

[195] Gell-Mann M, Oakes R J and Renner B 1968 Behavior of current divergences under $SU(3) \times SU(3)$ *Phys. Rev.* **175** 2195–9

[196] Coleman S R and Witten E 1980 Chiral symmetry breakdown in large N chromodynamics *Phys. Rev. Lett.* **45** 100

[197] 't Hooft G 1980 *Recent Developments in Gauge Theories* (New York: Plenum) p 135
Reprinted in *Unity of Forces in the Universe 1982* vol II ed A Zee (Singapore: World Scientific) 1004

[198] Asakawa M, Hatsuda T and Nakahara Y 2001 Maximum entropy analysis of the spectral functions in lattice QCD *Prog. Part. Nucl. Phys.* **46** 459

[199] Shuryak E V 1978 Theory of hadronic plasma *Sov. Phys.—JETP* **47** 212–9

[200] Shuryak E V 1978 quark–gluon plasma and hadronic production of leptons, photons and psions *Phys. Lett.* B **78** 150

[201] Linde A D 1980 Infrared problem in thermodynamics of the Yang–Mills gas *Phys. Lett.* B **96** 289

[202] Pisarski R D and Wilczek F 1984 Remarks on the chiral phase transition in chromodynamics *Phys. Rev.* D **29** 338–41

[203] Aoki Y, Endrodi G, Fodor Z, Katz S D and Szabo K K 2006 The order of the quantum chromodynamics transition predicted by the standard model of particle physics *Nature* **443** 675–8

[204] Aoki Y, Fodor Z, Katz S D and Szabo K K 2006 The QCD transition temperature: results with physical masses in the continuum limit *Phys. Lett.* B **643** 46

[205] Aoki Y, Borsanyi S, Durr S, Fodor Z, Katz S D, Krieg S and Szabo K K 2009 The QCD transition temperature: results with physical masses in the continuum limit II *J. High Energy Phys.* JHEP06(2009)088

[206] Stephanov M A 2004 QCD phase diagram and the critical point *Prog. Theor. Phys. Suppl.* **153** 139–56

[207] Fodor Z and Katz S D 2002 Lattice determination of the critical point of QCD at finite T and mu *J. High Energy Phys.* JHEP03(2002)014

[208] Allton C R, Ejiri S, Hands S J, Kaczmarek O, Karsch F, Laermann E, Schmidt C and Scorzato L 2002 The QCD thermal phase transition in the presence of a small chemical potential *Phys. Rev.* D **66** 074507

[209] Karsch F, Allton C R, Ejiri S, Hands S J, Kaczmarek O, Laermann E and Schmidt C 2004 Where is the chiral critical point in three flavor QCD? *Nucl. Phys. Proc. Suppl.* **129** 614–6

[210] Fodor S D and Katz Z 2004 Critical point of at QCD finite T and mu, lattice results for physical quark masses *J. High Energy Phys.* JHEP04(2004)050

[211] de Forcrand P and Philipsen O 2010 Constraining the QCD phase diagram by tricritical lines at imaginary chemical potential *Phys. Rev. Lett.* **105** 152001

[212] Stephanov M A, Rajagopal K and Shuryak E V 1998 Signatures of the tricritical point in QCD *Phys. Rev. Lett.* **81** 4816–9

[213] Gazdzicki M (NA49 and NA61/SHINE Collaborations) 2011 NA49/NA61: results and plans on beam energy and system size scan at the CERN SPS *J. Phys. G: Nucl. Part. Phys.* **38** 124024

[214] Gupta S, Luo X, Mohanty B, Ritter H G and Xu N 2011 Scale for the phase diagram of quantum chromodynamics *Science* **332** 1525–8

[215] Alford M G, Rajagopal K and Wilczek F 1999 Color flavor locking and chiral symmetry breaking in high density QCD *Nucl. Phys.* B **537** 443–58

[216] Schäfer T 2000 Patterns of symmetry breaking in at QCD high baryon density *Nucl. Phys.* B **575** 269–84

[217] Bedaque P F and Schäfer T 2002 High density quark matter under stress *Nucl. Phys.* A **697** 802

[218] Schäfer T 2006 P-wave meson condensation in high density QCD *Phys. Rev. Lett.* **96** 012305

[219] Abuki H, Baym G, Hatsuda T and Yamamoto N 2010 The NJL model of dense three-flavor matter with axial anomaly: the low temperature critical point and BEC–BCS diquark crossover *Phys. Rev.* D **81** 125010

[220] Maeda K, Baym G and Hatsuda T 2009 Simulating dense QCD matter with ultracold atomic boson–fermion mixtures *Phys. Rev. Lett.* **103** 085301

[221] Maeda K 2011 Large expansion N for strongly-coupled Boson Fermion mixtures *Ann. Phys.* **326** 1032–52

[222] Rapp A, Zarand G, Honerkamp C and Hofstetter W 2007 Color superfluidity and 'baryon' formation in ultracold fermions *Phys. Rev. Lett.* **98** 160405

[223] Shovkovy I A and Wang X 2011 Bulk viscosity in nonlinear and anharmonic regime of strange quark matter *New J. Phys.* **13** 045018

[224] Anglani R, Mannarelli M and Ruggieri M 2011 Collective modes in the color flavor locked phase *New J. Phys.* **13** 055002

[225] Mannarelli M, Colucci G and Manuel C 2011 Dissipative processes in superfluid quark matter *AIP Conf. Proc.* **1317** 366–71

[226] Schäfer T 2010 Phase structure and transport properties of dense quark matter arXiv:1001.2917 [astro-ph.SR]

[227] Kajantie K, Laine M, Rummukainen K and Schroder Y 2003 The pressure of hot QCD up to $g^6\log(1/g)$ *Phys. Rev.* D **67** 105008

[228] Kraemmer U and Rebhan A 2004 Advances in perturbative thermal field theory *Rep. Prog. Phys.* **67** 351

[229] Blaizot J P, Iancu E and Rebhan A 2003 Thermodynamics of the high-temperature quark gluon plasma *Quark Gluon Plasma* vol 3 ed R Hwa and X-N Wang pp 60–122 (arXiv:hep-ph/0303185)

[230] D'Hoker E and Freedman D Z 2001 Supersymmetric gauge theories and the AdS/CFT correspondence *Strings 2004 Branes and Extra Dimensions* (*Proc. TASI 2001*) ed S S Gubser and J D Lykken (Singapore: World Scientific) pp 3–160 (arXiv:hep-th/0201253)

[231] Gubser S S, Klebanov I R and Tseytlin A A 1998 Coupling constant dependence in the thermodynamics of $N = 4$ supersymmetric Yang–Mills theory *Nucl. Phys.* B **534** 202

[232] Baym G, Monien H, Pethick C J and Ravenhall D G 1990 Transverse interactions and transport in relativistic quark–gluon and electromagnetic plasmas *Phys. Rev. Lett.* **64** 1867

[233] Arnold P, Moore G D and Yaffe L G 2003 Transport coefficients in high temperature gauge theories. II: beyond leading log *J. High Energy Phys.* JHEP05(2003)051

[234] Danielewicz P and Gyulassy M 1985 Dissipative phenomena in quark gluon plasmas *Phys. Rev.* D **31** 53

[235] Policastro G, Son D T and Starinets A O 2001 The shear viscosity of strongly coupled $N = 4$ supersymmetric Yang–Mills plasma *Phys. Rev. Lett.* **87** 081601

[236] Buchel A and Liu J T 2004 Universality of the shear viscosity in supergravity *Phys. Rev. Lett.* **93** 090602

[237] Iqbal N and Liu H 2009 Universality of the hydrodynamic limit in AdS/ and the CFT membrane paradigm *Phys. Rev.* D **79** 025023

[238] Gyulassy M and McLerran L 2005 New forms of QCD matter discovered at RHIC *Nucl. Phys.* A **750** 30

[239] Teaney D 2006 Finite temperature spectral densities of momentum and R-charge correlators in $N = 4$ Yang Mills theory *Phys. Rev.* D **74** 045025

[240] Kovtun P and Starinets A 2006 Thermal spectral functions of strongly coupled $N = 4$ supersymmetric Yang–Mills theory *Phys. Rev. Lett.* **96** 131601

[241] Koch V, Majumder A and Randrup J 2005 Baryon-strangeness correlations: a diagnostic of strongly interacting matter *Phys. Rev. Lett.* **95** 182301

[242] Voloshin S A 2003 Anisotropic flow *Nucl. Phys.* A **715** 379–88

[243] Stock R 2004 Relativistic nucleus–nucleus collisions: from the BEVALAC to RHIC *J. Phys. G: Nucl. Part. Phys.* **30** S633

[244] Satz H 2008 Maurice Jacob: making quark matter *Phys. Scr.* **78** 028004

[245] Baym G 2002 RHIC: from dreams to beams in two decades *Nucl. Phys.* A **698** XXIII

[246] Adams J *et al* (STAR Collaboration) 2004 Identified particle distributions in pp and Au + Au collisions at $\sqrt{s_{NN}} = 200$ GeV *Phys. Rev. Lett.* **92** 112301

[247] Adler S S *et al* (PHENIX Collaboration) 2004 Identified charged particle spectra and yields in Au + Au collisions at $\sqrt{s_{NN}} = 200$ GeV *Phys. Rev.* C **69** 034909

[248] Miller M L, Reygers K, Sanders S J and Steinberg P 2007 Glauber modeling in high energy nuclear collisions *Annu. Rev. Nucl. Part. Sci.* **57** 205–43

[249] Bjorken J D 1983 Highly relativistic nucleus-nucleus collisions: the central rapidity region *Phys. Rev.* D **27** 140

[250] Glauber R J 1959 *High-Energy Collisions Theory* (*Lectures in Theoretical Physics* vol 1) ed W E Brittin and L G Dunham (New York: Interscience) p 315

[251] Kolb P F, Heinz U W, Huovinen P, Eskola K J and Tuominen K 2001 Centrality dependence of multiplicity, transverse energy, and elliptic flow from hydrodynamics *Nucl. Phys.* A **696** 197

[252] McLerran L D and Venugopalan R 1994 Computing quark and gluon distribution functions for very large nuclei *Phys. Rev.* D **49** 2233

[253] Hirano T, Heinz U W, Kharzeev D, Lacey R and Nara Y 2006 Hadronic dissipative effects on elliptic flow in ultrarelativistic heavy-ion collisions *Phys. Lett.* B **636** 299

[254] Abelev B *et al* (The ALICE Collaboration) 2010 Charged-particle multiplicity density at mid-rapidity in central Pb–Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV *Phys. Rev. Lett.* **105** 252301

[255] Fries R J, Kunihiro T, Muller B, Ohnishi A and Schäfer A 2009 From 0 to 5000 in $2 \cdot 10^{-24}$ seconds: entropy production in relativistic heavy-ion collisions *Nucl. Phys.* A **830** 519C–22C

[256] Bazavov A *et al* 2009 Equation of state and QCD transition at finite temperature *Phys. Rev.* D **80** 014504

[257] Bialas A, Bleszynski M and Czyz W 1977 Relation between the Glauber model and classical probability calculus *Acta Phys. Pol.* B **8** 389–92

[258] Bialas A, Bleszynski M and Czyz W 1976 Multiplicity distributions in nucleus–nucleus collisions at high-energies *Nucl. Phys.* B **111** 461

[259] Aamodt K *et al* (ALICE Collaboration) 2011 Centrality dependence of the charged-particle multiplicity density at mid-rapidity in Pb–Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV *Phys. Rev. Lett.* **106** 032301

[260] Alver B and Roland G 2010 Collision geometry fluctuations and triangular flow in heavy-ion collisions *Phys. Rev.* C **81** 054905

[261] Teaney D and Yan L 2011 Triangularity and dipole asymmetry in heavy ion collisions *Phys. Rev.* C **83** 064904

[262] Landau L D 1953 On the multiparticle production in high-energy collisions *Izv. Akad. Nauk Ser. Fiz.* **17** 51

[263] Cooper F and Frye G 1974 Comment on the single particle distribution in the hydrodynamic and statistical thermodynamic models of multiparticle production *Phys. Rev.* D **10** 186

[264] Heinz U W 2005 RHIC serves the perfect fluid: hydrodynamic flow of the QGP arXiv:nucl-th/0512051

[265] Baym G, Friman B L, Blaizot J P, Soyeur M and Czyz W 1983 Hydrodynamics of ultrarelativistic heavy ion collisions *Nucl. Phys.* A **407** 541

[266] Steinberg P 2005 Landau hydrodynamics and RHIC phenomena *Acta Phys. Hung.* A **24** 51–7

[267] Teaney D 2003 Effect of shear viscosity on spectra, elliptic flow, and Hanbury Brown–Twiss radii *Phys. Rev.* C **68** 034913

[268] Adcox K *et al* (PHENIX Collaboration) 2002 Suppression of hadrons with large transverse momentum in central Au + Au collisions at $\sqrt{s_{NN}} = 130$ GeV *Phys. Rev. Lett.* **88** 022301

[269] Ellis S D, Kunszt Z and Soper D E 1990 The one jet inclusive cross-section at order α_s^3 quarks and gluons *Phys. Rev. Lett.* **64** 2121

[270] Arleo F 2006 Hard pion and prompt photon at RHIC, from single to double inclusive production *J. High Energy Phys.* JHEP09(2006)015

[271] Aamodt K *et al* (ALICE Collaboration) 2011 Suppression of charged particle production at large transverse momentum in central Pb–Pb collisions at $\sqrt{s_{NN}} = 2.76\,\mathrm{TeV}$ *Phys. Lett.* B **696** 30

[272] Chatrchyan S *et al* (CMS Collaboration) 2012 Study of high-pT charged particle suppression in PbPb compared to pp collisions at $\sqrt{s_{NN}} = 2.76\,\mathrm{TeV}$ *Eur. Phys. J.* C **72** 1945

[273] Cole B A (ATLAS Collaboration) 2011 Jet probes of $\sqrt{s_{NN}} = 2.76\,\mathrm{TeV}$ Pb + Pb collisions with the ATLAS detector *J. Phys. G: Nucl. Part. Phys.* **38** 124021

[274] Baier R, Dokshitzer Y L, Mueller A H, Peigne S and Schiff D 1997 Radiative energy loss of high-energy quarks and gluons in a finite volume quark–gluon plasma *Nucl. Phys.* B **483** 291–320

[275] Salgado C A and Wiedemann U A 2003 Calculating quenching weights *Phys. Rev.* D **68** 014008

[276] Adare A *et al* (PHENIX Collaboration) 2008 Quantitative constraints on the opacity of hot partonic matter from semi-inclusive single high transverse momentum pion suppression in Au + Au collisions at $\sqrt{s_{NN}} = 200\,\mathrm{GeV}$ *Phys. Rev.* C **77** 064907

[277] Baier R 2003 Jet quenching *Nucl. Phys.* A **715** 209

[278] Bass S A, Gale C, Majumder A, Nonaka C, Qin G-Y, Renk T and Ruppert J 2009 Systematic comparison of jet energy-loss schemes in a realistic hydrodynamic medium *Phys. Rev.* C **79** 024901

[279] Majumder A, Muller B and Wang X-N 2007 Small shear viscosity of a quark–gluon plasma implies strong jet quenching *Phys. Rev. Lett.* **99** 192301

[280] Dusling K, Moore G D and Teaney D 2010 Radiative energy loss and v_2 spectra for viscous hydrodynamics *Phys. Rev.* C **81** 034907

[281] Adare A *et al* (PHENIX Collaboration) 2010 Azimuthal anisotropy of neutral pion production in Au + Au collisions at $\sqrt{(s_N N)} = 200\,\mathrm{GeV}$: path-length dependence of jet quenching and the role of initial geometry *Phys. Rev. Lett.* **105** 142301

[282] Adare A *et al* (PHENIX Collaboration) 2007 Energy loss and flow of heavy quarks in Au + Au collisions at $\sqrt{s_{NN}} = 200\,\mathrm{GeV}$ *Phys. Rev. Lett.* **98** 172301

[283] Adler S S *et al* (PHENIX Collaboration) 2003 Suppressed pi^0 production at large transverse momentum in central Au + Au collisions at $\sqrt{s_{NN}} = 200\,\mathrm{GeV}$ *Phys. Rev. Lett.* **91** 072301

[284] Armesto N, Cacciari M, Dainese A, Salgado C A and Wiedemann U A 2006 How sensitive are high-p_T electron spectra at RHIC to heavy quark energy loss? *Phys. Lett.* B **637** 362–6

[285] van Hees H, Mannarelli M, Greco V and Rapp R 2008 Nonperturbative heavy-quark diffusion in the quark–gluon plasma *Phys. Rev. Lett.* **100** 192301

[286] Moore G D and Teaney D 2005 How much do heavy quarks thermalize in a heavy ion collision? *Phys. Rev.* C **71** 064904

[287] Adler S S *et al* (PHENIX Collaboration) 2006 Measurement of identified pi0 and inclusive photon v_2 and implication to the direct photon production in $\sqrt{s_{NN}} = 200\,\mathrm{GeV}$ Au + Au collisions *Phys. Rev. Lett.* **96** 032302

[288] Svetitsky B 1988 Diffusion of charmed quark in the quark gluon plasma *Phys. Rev.* D **37** 2484

[289] Herzog C P, Karch A, Kovtun P, Kozcaz C and Yaffe L G 2006 Energy loss of a heavy quark moving through $N = 4$ supersymmetric Yang–Mills plasma *J. High Energy Phys.* JHEP07(2006)013

[290] Casalderrey-Solana J and Teaney D 2006 Heavy quark diffusion in strongly coupled $N = 4$ Yang Mills *Phys. Rev.* D **74** 085012

[291] Gubser S S 2006 Drag force in AdS/CFT *Phys. Rev.* D **74** 126005

[292] Gubser S S 2007 Comparing the drag force on heavy quarks in $N = 4$ super-Yang–Mills theory and QCD *Phys. Rev.* D **76** 126003

[293] Meyer H B 2011 The errant life of a heavy quark in the quark–gluon plasma *New J. Phys.* **13** 035008

[294] Riek F and Rapp R 2011 Selfconsistent evaluation of charm and charmonium in the quark–gluon plasma *New J. Phys.* **13** 045007

[295] Israel W 1967 Event horizons in static vacuum space–times *Phys. Rev.* **164** 1776–9

[296] Israel W 1968 Event horizons in static electrovac space–times *Commun. Math. Phys.* **8** 245–60

[297] Carter B 1971 Axisymmetric black hole has only two degrees of freedom *Phys. Rev. Lett.* **26** 331–3

[298] Hawking S W 1971 Gravitational radiation from colliding black holes *Phys. Rev. Lett.* **26** 1344–6

[299] Hawking S W 1972 Black holes in general relativity *Commun. Math. Phys.* **25** 152–66

[300] Price R H and Thorne K S 1986 Membrane viewpoint on black holes: properties and evolution of the stretched horizon *Phys. Rev.* D **33** 915–41

[301] Kovtun P, Son D T and Starinets A O 2003 Holography and hydrodynamics: diffusion on stretched horizons *J. High Energy Phys.* JHEP10(2003)064

[302] Maldacena J M 1999 The large-N limit of superconformal field theories and supergravity *Int. J. Theor. Phys.* **38** 1113–33 (arXiv:hep-th/9711200)

Maldacena J M 1998 The large-N limit of superconformal field theories and supergravity *Adv. Theor. Math. Phys.* **2** 231–52

[303] Gubser S S, Klebanov I R and Polyakov A M 1998 Gauge theory correlators from noncritical string theory *Phys. Lett.* B **428** 105–14 (arXiv:hep-th/9802109)

[304] Witten E 1998 Anti-de Sitter space and holography *Adv. Theor. Math. Phys.* **2** 253–91 (arXiv:hep-th/9802150)

[305] Aharony O 2003 The NonAdS/nonCFT correspondence, or three different paths to QCD *Progress in String, Field and Particle Theory* (*Nato Science Series II: Mathematics, Physics and Chemistry* vol 104) ed L Baulieu, E Rabinovici, J Harvey, B Pioline and P Windey (Dordrecht: Kluwer) pp 3–24

[306] Son D T 2008 Toward an AdS/cold atoms correspondence: a geometric realization of the Schrödinger symmetry *Phys. Rev.* D **78** 046003

[307] Balasubramanian K and McGreevy J 2008 Gravity duals for non-relativistic CFTs *Phys. Rev. Lett.* **101** 061601

[308] Kachru S, Liu X and Mulligan M 2008 Gravity duals of Lifshitz-like fixed points *Phys. Rev.* D **78** 106005

[309] Faulkner T and Polchinski J 2011 Semi-holographic Fermi liquids *J. High Energy Phys.* JHEP06(2011)012

[310] Iqbal N, Liu H and Mezei M 2012 Lectures on holographic non-Fermi liquids and quantum phase transitions *String Theory and its Applications (TASI 2010: From MeV to the Planck Scale)* ed M Dine, T Banks and S Sachdev (Singapore: World Scientific) pp 707–815

[311] Nickel D and Son D T 2011 Deconstructing holographic liquids *New J. Phys.* **13** 075010

[312] Chesler P M and Yaffe L G 2011 Holography and colliding gravitational shock waves in asymptotically AdS_5 spacetime *Phys. Rev. Lett.* **106** 021601

[313] Hartnoll S A 2009 Lectures on holographic methods for condensed matter physics *Class. Quantum Grav.* **26** 224002

[314] Hartnoll S A 2009 Quantum critical dynamics from black holes *Understanding Quantum Phase Transitions* ed L Carr (Boca Raton FL: CRC) pp 701–23

[315] McGreevy J 2010 Holographic duality with a view toward many-body physics *Adv. High Energy Phys.* **2010** 723105

[316] Aharony O, Gubser S S, Maldacena J M, Ooguri H and Oz Y 2000 Large N field theories, string theory and gravity *Phys. Rep.* **323** 183–386

[317] Son D T and Starinets A O 2007 Viscosity, black holes and quantum field theory *Annu. Rev. Nucl. Part. Sci.* **57** 95–118

[318] Erdmenger J, Evans N, Kirsch I and Threlfall E 2008 Mesons in gauge/gravity duals—a review *Eur. Phys. J.* A **35** 81–133

[319] Herzog C P 2009 Lectures on holographic superfluidity and superconductivity *J. Phys. A: Math. Theor.* **42** 343001

[320] Gubser S S and Karch A 2009 From gauge-string duality to strong interactions: a pedestrian's guide *Annu. Rev. Nucl. Part. Sci.* **59** 145–68

[321] Lee S-S 2010 TASI lectures on emergence of supersymmetry, gauge theory and string in condensed matter systems arXiv:1009.5127 [hep-th]

[322] Hughes S A 2005 Trust but verify: the case for astrophysical black holes *eConf* C0507252(2005)L006

[323] Gallo E and Marolf D 2009 Resource letter BH-2: black holes *Am. J. Phys.* **77** 294–307

[324] Kadanoff L P 1966 Scaling laws for Ising models near T_c *Physics* **2** 263–72

[325] Wilson K G and Kogut J B 1974 The renormalization group and the epsilon expansion *Phys. Rep.* **12** 75–200

[326] Wilson K G 1975 The renormalization group: critical phenomena and the kondo problem *Rev. Mod. Phys.* **47** 773

[327] Weinberg S 1965 Infrared photons and gravitons *Phys. Rev.* **140** B516–24

[328] Weinberg S and Witten E 1980 Limits on massless particles *Phys. Lett.* B **96** 59

[329] Jenkins A 2009 Constraints on emergent gravity *Int. J. Mod. Phys.* D **18** 2249–55

[330] Bousso R 2002 The holographic principle *Rev. Mod. Phys.* **74** 825–74

[331] Hooft G 1993 Dimensional reduction in quantum gravity (essay dedicated to Abdus Salam) arXiv:gr-qc/9310026

[332] Susskind L 1995 The World as a hologram *J. Math. Phys.* **36** 6377–96

[333] Heemskerk I, Penedones J, Polchinski J and Sully J 2009 Holography from conformal field theory *J. High Energy Phys.* JHEP10(2009)079

[334] Fitzpatrick L A and Kaplan J 2011 Analyticity and the holographic S-matrix arXiv:1111.6972 [hep-th]

[335] Heemskerk I and Polchinski J 2011 Holographic and Wilsonian renormalization groups *J. High Energy Phys.* JHEP06(2011)031

[336] Adams A, Maloney A, Sinha A and Vazquez S E 2009 $1/N$ effects in non-relativistic gauge–gravity duality *J. High Energy Phys.* JHEP03(2009)097

[337] Hawking S W and Page D N 1983 Thermodynamics of black holes in anti-de Sitter space *Commun. Math. Phys.* **87** 577

[338] Strominger A and Vafa C 1996 Microscopic origin of the Bekenstein–Hawking entropy *Phys. Lett.* B **379** 99–104

[339] Iqbal N and Liu H 2009 Real-time response in AdS/CFT with application to spinors *Fortsch. Phys.* **57** 367–84

[340] Sachdev S 2001 *Quantum Phase Transitions* (Cambridge: Cambridge University Press)

[341] Hooft G 1974 A planar diagram theory for strong interactions *Nucl. Phys.* B **72** 461–73

[342] Coleman S R 1988 *$1/N$ Aspects of Symmetry* (Cambridge: Cambridge University Press)

[343] Silverstein E 2003 AdS and dS entropy from string junctions: or, the function of junction conjunctions *From Fields to Strings* vol 3 ed M Shifman *et al* (Singapore: World Scientific) pp 1848–63 (arXiv:hep-th/0308175)

[344] Balasubramanian V and Kraus P 1999 A stress tensor for anti-de Sitter gravity *Commun. Math. Phys.* **208** 413–28

[345] Bianchi M, Freedman D Z and Skenderis K 2001 How to go with an RG flow *J. High Energy Phys.* JHEP08(2001)041

[346] Bianchi M, Freedman D Z and Skenderis K 2002 Holographic renormalization *Nucl. Phys.* B **631** 159–94

[347] Skenderis K 2002 Lecture notes on holographic renormalization *Class. Quantum Grav.* **19** 5849–76

[348] Faulkner T, Liu H and Rangamani M 2011 Integrating out geometry: holographic Wilsonian RG and the membrane paradigm *J. High Energy Phys.* JHEP08(2011)051

[349] Balasubramanian V, Kraus P, Lawrence A E and Trivedi S P 1999 Holographic probes of anti-de Sitter space–times *Phys. Rev.* D **59** 104021

[350] Klebanov I R and Witten E 1999 AdS/CFT correspondence and symmetry breaking *Nucl. Phys.* B **556** 89–114

[351] Papadimitriou I and Skenderis K 2004 Correlation functions in holographic RG flows *J. High Energy Phys.* JHEP10(2004)075

[352] Son D T and Starinets A O 2002 Minkowski space correlators in AdS/CFT correspondence: recipe and applications *J. High Energy Phys.* JHEP09(2002)042

[353] Maldacena J M 2003 Eternal black holes in anti-de Sitter *J. High Energy Phys.* JHEP04(2003)021

[354] Herzog C P and Son D T 2003 Schwinger–Keldysh propagators from AdS/CFT correspondence *J. High Energy Phys.* JHEP03(2003)046

[355] Skenderis K and van Rees B C 2008 Real-time gauge/gravity duality *Phys. Rev. Lett.* **101** 081601

[356] Witten E 1998 Anti-de Sitter space, thermal phase transition and confinement in gauge theories *Adv. Theor. Math. Phys.* **2** 505–32 (arXiv:hep-th/9803131)

[357] Kachru S, Karch A and Yaida S 2010 Holographic lattices, dimers and glasses *Phys. Rev.* D **81** 026007

[358] Karch A, Kachru S and Yaida S 2011 Adventures in holographic dimer models *New J. Phys.* **13** 035004

[359] Cremonini S 2011 The shear viscosity to entropy ratio: a status report *Mod. Phys. Lett.* B **25** 1867–88

[360] Brigante M, Liu H, Myers R C, Shenker S and Yaida S 2008 Viscosity bound violation in higher derivative gravity *Phys. Rev.* D **77** 126006

[361] Kats Y and Petrov P 2009 Effect of curvature squared corrections in AdS on the viscosity of the dual gauge theory *J. High Energy Phys.* JHEP01(2009)044

[362] Buchel A and Myers R C 2009 Causality of holographic hydrodynamics *J. High Energy Phys.* JHEP08(2009)016

[363] Hofman D M 2009 Higher derivative gravity, causality and positivity of energy in a UV complete QFT *Nucl. Phys.* B **823** 174–94

[364] de Boer J, Kulaxizi M and Parnachev A 2010 Holographic Lovelock gravities and black holes *J. High Energy Phys.* JHEP06(2010)008

[365] Camanho X O and Edelstein J D 2010 Causality in AdS/CFT and Lovelock theory *J. High Energy Phys.* JHEP06(2010)099

[366] Camanho X O, Edelstein J D and Paulos M F 2011 Lovelock theories, holography and the fate of the viscosity bound *J. High Energy Phys.* JHEP05(2011)127

[367] Casalderrey-Solana J and Salgado C A 2007 Introductory lectures on jet quenching in heavy ion collisions *Acta Phys. Pol.* B **38** 3731–94

[368] D'Eramo F, Liu H and Rajagopal K 2011 Jet quenching parameter via soft collinear effective theory (SCET) *Int. J. Mod. Phys.* E **20** 1610–5

[369] Liu H, Rajagopal K and Wiedemann U A 2006 Calculating the jet quenching parameter from AdS/CFT *Phys. Rev. Lett.* **97** 182301

[370] Liu H, Rajagopal K and Wiedemann U A 2007 Wilson loops in heavy ion collisions and their calculation in AdS/CFT *J. High Energy Phys.* JHEP03(2007)066

[371] Chesler P M, Jensen K, Karch A and Yaffe L G 2009 Light quark energy loss in strongly-coupled $N = 4$ supersymmetric Yang–Mills plasma *Phys. Rev.* D **79** 125015

[372] Arnold P and Vaman D 2011 Jet quenching in hot strongly coupled gauge theories simplified *J. High Energy Phys.* JHEP04(2011)027

[373] Hubeny V E 2011 Holographic dual of collimated radiation *New J Phys.* **13** 035006

[374] Gubser S S 2008 Breaking an Abelian gauge symmetry near a black hole horizon *Phys. Rev.* D **78** 065034

[375] Hartnoll S A, Herzog C P and Horowitz G T 2008 Building a holographic superconductor *Phys. Rev. Lett.* **101** 031601

[376] Hartnoll S A, Herzog C P and Horowitz G T 2008 Holographic superconductors *J. High Energy Phys.* JHEP12(2008)015

[377] Herzog C 2011 Solitons as probes of the structure of holographic superfluids *New J. Phys.* **13** 065003

[378] Keränen V, Keski-Vakkuri E, Nowling S and Yogendran K P 2009 Dark solitons in holographic superfluids *Phys. Rev.* D **80** 121901

[379] Keränen V, Keski-Vakkuri E, Nowling S and Yogendran K P 2010 Inhomogeneous structures in holographic superfluids I. Dark solitons *Phys. Rev.* D **81** 126011

[380] Basu P, Nogueira F, Rozali M, Stang J B and Van Raamsdonk M 2011 Towards a holographic model of color superconductivity *New J. Phys.* **13** 055001

[381] Mehen T, Stewart I W and Wise M B 2000 Conformal invariance for nonrelativistic field theory *Phys. Lett.* B **474** 145–52

[382] Son D T and Wingate M 2006 General coordinate invariance and conformal invariance in nonrelativistic physics: unitary Fermi gas *Ann. Phys.* **321** 197–224

[383] Nishida Y and Son D T 2007 Nonrelativistic conformal field theories *Phys. Rev.* D **76** 086004

[384] Herzog C P, Rangamani M and Ross S F 2008 Heating up Galilean holography *J. High Energy Phys.* JHEP11(2008)080

[385] Maldacena J, Martelli D and Tachikawa Y 2008 Comments on string theory backgrounds with non-relativistic conformal symmetry *J. High Energy Phys.* JHEP10(2008)072

[386] Adams A, Balasubramanian K and McGreevy J 2008 Hot spacetimes for cold atoms *J. High Energy Phys.* JHEP11(2008)059

[387] Balasubramanian K and McGreevy J 2009 An analytic Lifshitz black hole *Phys. Rev.* D **80** 104039

[388] Adams A and Wang J 2011 Towards a non-relativistic holographic superfluid *New J. Phys.* **13** 115008

[389] Hartnoll S A, Polchinski J, Silverstein E and Tong D 2010 Towards strange metallic holography *J. High Energy Phys.* JHEP04(2010)120

[390] Kovtun P and Nickel D 2009 Black holes and non-relativistic quantum systems *Phys. Rev. Lett.* **102** 011602

[391] Barbon J L F and Fuertes C A 2009 Ideal gas matching for thermal Galilean holography *Phys. Rev.* D **80** 026006

[392] Balasubramanian K and McGreevy J 2011 The particle number in Galilean holography *J. High Energy Phys.* JHEP01(2011)137

[393] Holstein T, Norton R E and Pincus P 1973 De Haas–van Alphen effect and the specific heat of an electron gas *Phys. Rev.* B **8** 2649–56

[394] Lee S-S 2009 A non-Fermi liquid from a charged black hole: critical a Fermi ball *Phys. Rev.* D **79** 086006

[395] Liu H, McGreevy J and Vegh D 2011 Non-Fermi liquids from holography *Phys. Rev.* D **83** 065029

[396] Cubrovic M, Zaanen J and Schalm K 2009 String theory, quantum phase transitions and the emergent Fermi-liquid *Science* **325** 439–44

[397] Faulkner T, Liu H, McGreevy J and Vegh D 2011 Emergent quantum criticality, Fermi surfaces and AdS(2) *Phys. Rev.* D **83** 125002

[398] Hartnoll S A and Tavanfar A 2011 Electron stars for holographic metallic criticality *Phys. Rev.* D **83** 046003

[399] Hartnoll S A, Hofman D M and Tavanfar A 2011 Holographically smeared Fermi surface: quantum oscillations and Luttinger count in electron stars *Europhys. Lett.* **95** 31002

[400] Iqbal N, Liu H and Mezei M 2012 Semi-local quantum liquids *J. High Energy Phys.* JHEP04(2012)086

[401] Jensen K 2011 Semi-holographic quantum criticality *Phys. Rev. Lett.* **107** 231601

[402] Iqbal N, Liu H and Mezei M 2011 Quantum phase transitions in semi-local quantum liquids arXiv:1108.0425 [hep-th]

[403] Bayntun A, Burgess C P, Dolan B P and Lee S-S 2011 AdS/QHE: towards a holographic description of quantum hall experiments *New J. Phys.* **13** 035012

[404] Janko B, Maly J and Levin K 1997 Pseudogap effects induced by resonant pair scattering *Phys. Rev.* B **56** R11407

[405] Randeria M 1998 *Proc. Int. School of Physics 'Enrico Fermi'* (*Course CXXXVI*) ed G Iadonisi, J R Schrieffer and M L Chiafalo (Amsterdam: IOS Press) p 115

[406] Chin C, Bartenstein M, Altmeyer A, Riedl S, Jochim S, Hecker Denschlag J and Grimm R 2004 Observation of the pairing gap in a strongly interacting Fermi gas *Science* **305** 1128

[407] Gaebler J P, Stewart J T, Drake T E, Jin D S, Perali A, Pieri P and Strinati G C 2010 Observation of pseudogap behaviour in a strongly interacting Fermi gas *Nature Phys.* **6** 569

[408] Magierski P, Wlazlowski G and Bulgac A 2011 Onset of a pseudogap regime in ultracold Fermi gases *Phys. Rev. Lett.* **107** 145304

[409] Feld M, Fröhlich B, Vogt E, Koschorreck M and Köhl M 2011 Observation of a pairing pseudogap in a two-dimensional Fermi gas *Nature* **480** 75

[410] Schirotzek A, Wu C-H, Sommer A and Zwierlein M W 2009 Observation of Fermi polarons in a tunable Fermi liquid of ultracold atoms *Phys. Rev. Lett.* **102** 230402

[411] Lobo C, Recati A, Giorgini S and Stringari S 2006 Normal state of a polarized Fermi gas at unitarity *Phys. Rev. Lett.* **97** 200403

[412] Koschorreck M, Pertot D, Vogt E, Fröhlich B, Feld M and Köhl M 2012 Attractive and repulsive Fermi polarons in two dimensions *Nature* **485** 619

[413] Zhang Y, Ong W, Arakelyan I and Thomas J E 2012 Polarons in the radio-frequency spectrum of a quasi-two-dimensional Fermi gas *Phys. Rev. Lett.* **108** 235302

[414] Taylor E and Randeria M 2010 Viscosity of strongly interacting quantum fluids: spectral functions and sum rules *Phys. Rev.* A **81** 053610

[415] Enss T 2012 Quantum critical transport in the unitary Fermi gas *Phys. Rev.* A **86** 013616

[416] Diehl S, Floerchinger S, Gies H, Pawlowski J M and Wetterich C 2010 Functional renormalization group approach to the BCS–BEC crossover *Ann. Phys.* **522** 615

[417] Braaten E and Hammer H-W 2007 Efimov physics in cold atoms *Ann. Phys.* **322** 120

[418] Nishida Y and Tan S 2011 Liberating Efimov physics from three dimensions *Few Body Syst.* **51** 191

[419] Salger T, Grassert C, Kling S and Weitz M 2011 Klein-tunneling of a quasi-relativistic Bose–Einstein condensate in an optical lattice *Phys. Rev. Lett.* **107** 240401

[420] Tarruell L, Greif D, Uehlinger T, Jotzu G and Esslinger T 2012 Creating, moving and merging Dirac points with a Fermi gas in a tunable honeycomb lattice *Nature* **483** 302

[421] Pilati S, Giorgini S, Modugno M and Prokof'ev N 2010 Dilute Bose gas with correlated disorder: a path integral Monte Carlo study *New J. Phys.* **12** 073003

[422] Adams A and Yaida S 2011 Disordered holographic systems. I: Functional renormalization arXiv:1102.2892 [hep-th]

[423] Adams A and Yaida S 2012 Disordered holographic systems II: marginal relevance of imperfection arXiv:1201.6366 [hep-th]

[424] Jendrzejewski F, Bernard A, Mueller K, Cheinet P, Josse V, Piraud M, Pezze L, Sanchez-Palencia L, Aspect A and Bouyer P 2012 Three-dimensional localization of ultracold atoms in an optical disordered potential *Nature Phys.* **8** 398–403

[425] Alver B H, Gombeaud C, Luzum M and Ollitrault J-Y 2010 Triangular flow in hydrodynamics and transport theory *Phys. Rev.* C **82** 034913

[426] Schenke B, Jeon S and Gale C 2012 Higher flow harmonics from $(3+1)$D event-by-event viscous hydrodynamics *Phys. Rev.* C **85** 024901

[427] Song H and Heinz U W 2009 Extracting the QGP viscosity from RHIC data—a status report from viscous hydrodynamics *J. Phys. G: Nucl. Part. Phys.* **36** 064033

[428] Adare A *et al* (PHENIX Collaboration) 2010 Enhanced production of direct photons in Au + Au collisions at $\sqrt{s_{NN}} = 200$ GeV and implications for the initial temperature *Phys. Rev. Lett.* **104** 132301

[429] Martinez M, Ryblewski R and Strickland M 2012 Boost-invariant (2+1)-dimensional anisotropic hydrodynamics *Phys. Rev.* C **85** 064913

[430] Baier R, Mueller A H, Schiff D and Son D T 2001 'Bottom up' thermalization in heavy ion collisions *Phys. Lett.* B **502** 51–8

[431] Arnold P B, Lenaghan J and Moore G D 2003 QCD plasma instabilities and bottom up thermalization *J. High Energy Phys.* JHEP08(2003)002

[432] Kurkela A and Moore G D 2011 Thermalization in weakly coupled nonabelian plasmas *J. High Energy Phys.* JHEP12(2011)044

[433] Mueller A H, Shoshi A I and Wong S M H 2006 A possible modified 'bottom-up' thermalization in heavy ion collisions *Phys. Lett.* B **632** 257–60

[434] Dusling K 2011 Role of quantum fluctuations in a system with strong fields *New J. Phys.* **13** 035011

[435] Kovchegov Y V 2009 Early time dynamics in heavy ion collisions from and from CGC AdS/CFT *Nucl. Phys.* A **830** 395C–402C

[436] Gubser S S 2010 TASI lectures: collisions in anti-de Sitter space, conformal symmetry and holographic superconductors arXiv:1012.5312 [hep-th]

[437] Gubser S S, Pufu S S and Yarom A 2009 Off-center collisions in AdS(5) with applications to multiplicity estimates in heavy-ion collisions *J. High Energy Phys.* JHEP11(2009)050

[438] Meyer H B 2011 Transport properties of the quark–gluon plasma: lattice a perspective QCD *Eur. Phys. J.* A **47** 86

[439] Mohanty B (STAR Collaboration) 2011 STAR experiment results from the beam energy scan program at RHIC *J. Phys. G: Nucl. Part. Phys.* **38** 124023

[440] Mohanty B 2011 Exploring the quantum chromodynamics landscape with high-energy nuclear collisions *New J. Phys.* **13** 065031

[441] Stephanov M A 2009 Non-Gaussian fluctuations near the QCD critical point *Phys. Rev. Lett.* **102** 032301

[442] de Forcrand P 2009 Simulating QCD at finite density PoS(LAT2009)010 (arXiv:1005.0539 [hep-lat])

[443] Van Hove L 1982 Multiplicity dependence of $p(T)$ spectrum as a possible signal for a phase transition in hadronic collisions *Phys. Lett.* B **118** 138

[444] Steiner A W, Lattimer J M and Brown E F 2010 The equation of state from observed masses and radii of neutron stars *Astrophys. J.* **722** 33–54

[445] Ozel F, Baym G and Guver T 2010 Astrophysical measurement of the equation of state of neutron star matter *Phys. Rev.* D **82** 101301

[446] Lee S-S 2011 Holographic description of large N gauge theory *Nucl. Phys.* B **851** 143–60

[447] Douglas M R, Mazzucato L and Razamat S S 2011 Holographic dual of free field theory *Phys. Rev.* D **83** 071701

[448] Ryu S and Takayanagi T 2006 Holographic derivation of entanglement entropy from AdS/CFT *Phys. Rev. Lett.* **96** 181602

[449] Nishioka T, Ryu S and Takayanagi T 2009 Holographic entanglement entropy: an overview *J. Phys. A: Math. Theor.* **42** 504008

[450] Casini H and Huerta M 2012 Infinite divisibility, entanglement entropy and minimal surfaces arXiv:1203.4007 [hep-th]

[451] Callan R, He J and Headrick M 2012 Strong subadditivity and the covariant holographic entanglement entropy formula *J. High Energy Phys.* JHEP06(2012)081

[452] Headrick M and Takayanagi T 2007 A holographic proof of the strong subadditivity of entanglement entropy *Phys. Rev.* D **76** 106013

[453] Hubeny V E, Rangamani M and Takayanagi T 2007 A covariant holographic entanglement entropy proposal *J. High Energy Phys.* JHEP07(2007)062

[454] Hayden P, Headrick M and Maloney A 2011 Holographic mutual information is monogamous arXiv:1107.2940 [hep-th]

[455] Casini H, Huerta M and Myers R C 2011 Towards a derivation of holographic entanglement entropy *J. High Energy Phys.* JHEP05(2011)036

[456] Albash T and Johnson C V 2011 Evolution of holographic entanglement entropy after thermal and electromagnetic quenches *New J. Phys.* **13** 045017