

3

Results from lattice QCD at nonzero temperature

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

We shall not review the techniques by which lattice-regularized calculations are implemented. We merely recall that the starting point of lattice-regularized calculations at nonzero temperature is the imaginary time formalism, which allows one to write the QCD partition function in Euclidean spacetime with a periodic imaginary time direction of length $1/T$ [541]. Any thermodynamic quantity can be obtained via suitable differentiation of the partition function. At zero baryon chemical potential, the QCD partition function is given by the exponent of a real action, integrated over all field configurations in the Euclidean spacetime. Since the action is real, the QCD partition function can then be evaluated using standard Monte Carlo techniques, which require the discretization of the field configurations and the evaluation of the action on a finite lattice of spacetime points. Physical results are obtained by extrapolating calculated results to the limit of infinite volume and vanishing lattice spacing. In principle, this is a quantitatively reliable approach. In practice, lattice-regularized calculations are CPU-expensive: the size of lattices in modern calculations does not exceed $48^3 \times 64$ [332] and these calculations nevertheless require the most powerful computing devices (currently at the multi-teraflop scale). In the continuum limit, such lattices correspond typically to small volumes of $\approx (4 \text{ fm})^3$ [332]. This means that properties of QCD matter which are dominated by long-wavelength modes are difficult to calculate with the currently available computing resources and there are only first exploratory studies.

For the same reason, it is in practice difficult to carry out calculations using light quark masses that yield realistically light pion masses at zero temperature. Light quarks are also challenging because of the CPU-expensive complications which arise from the formulation of fermions on the lattice.

In addition to the practical challenges above, conceptual questions arise in two important domains. First, at nonzero baryon chemical potential the Euclidean action is no longer real, meaning that the so-called fermion sign problem precludes the use of standard Monte Carlo techniques. Techniques have been found that evade this problem, but only in the regime where the quark chemical potential $\mu_B/3$ is sufficiently small compared to T [352, 312, 45, 315, 353, 46, 371, 372, 313]. Second, conceptual questions arise in the calculation of any physical quantities that cannot be written as derivatives of the partition function. Many such quantities are of considerable interest. Calculating them requires the analytic continuation of lattice results from Euclidean to Minkowski space (see below) which is always under-constrained since the Euclidean calculations can only be done at finitely many values of the Euclidean time. This means that lattice-regularized calculations, at least as currently formulated, are not optimized for calculating transport coefficients and answering questions about, say, far-from-equilibrium dynamics or jet quenching.

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

3.1 The QCD equation of state from the lattice

The QCD equation of state at zero baryon chemical potential, namely the relation between the pressure and the energy density of hot QCD matter, is an example of a quantity that is well-suited to lattice-regularized calculation since, as a thermodynamic quantity, it can be obtained via suitable differentiations of the Euclidean partition function. And, the phenomenological motivation for determining this quantity from first principles in QCD is strong since, as we have seen in Section 2.2, it is the most important microphysical input for hydrodynamic calculations. Accurate calculations of the thermodynamics of pure glue QCD ($N_f = 0$) have existed for a long time [184], but the extraction of the equation of state of quark–gluon plasma with light quarks having their physical masses, and with the continuum limit taken, has become possible only recently [59, 129, 179]. This illustrates the

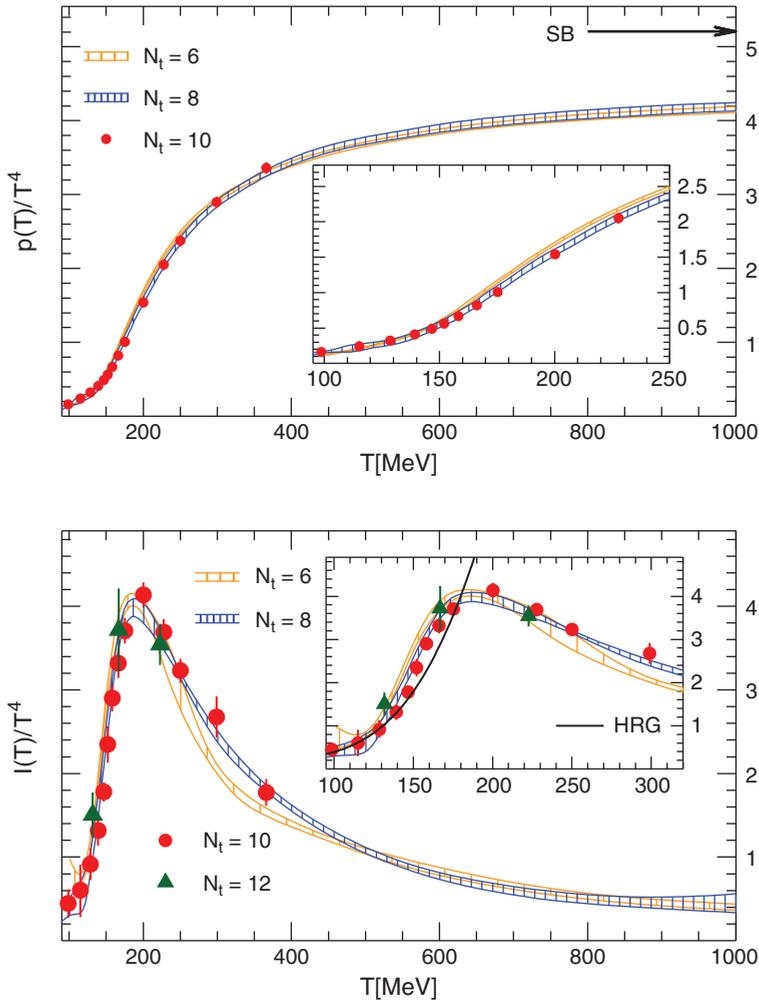


Figure 3.1 Results from a lattice calculation of QCD thermodynamics with physical quark masses ($N_f = 3$, with appropriate light and strange masses). Upper panel: temperature dependence of the pressure in units of T^4 . Lower panel: the trace anomaly ($\varepsilon - 3P$) in units of T^4 . Data are for lattices with the same temporal extent, meaning the same temperature, but with varying numbers of points in the Euclidean time direction N_τ . The continuum limit corresponds to taking $N_\tau \rightarrow \infty$. Figures taken from Ref. [179].

practical challenges of doing lattice-regularized calculations with light quarks that we have mentioned above.

The current understanding of QCD thermodynamics at the physical point [179] is summarized in Fig. 3.1. In the upper panel, the pressure of QCD matter (in thermal equilibrium, with zero baryon chemical potential) is plotted as a function of its temperature. In order to provide a physically meaningful reference, it is customary to compare this quantity to the Stefan–Boltzmann result

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

for a free gas of noninteracting gluons and massless quarks. This benchmark is indicated by the arrow in the figure. As illustrated by this plot, the number of degrees of freedom rises rapidly above a temperature $T_c \sim 170$ MeV; at higher temperatures, the pressure takes an almost constant value which deviates from that of a noninteracting gas of quarks and gluons by approximately 20%. This deviation is still present at temperatures as high as 1 GeV, and convergence to the noninteracting limit is only observed at asymptotically high temperatures ($T > 10^8$ GeV [354]) which are far from the reach of any collider experiment. The lower panel shows the trace anomaly, $\varepsilon - 3P$, in units of T^4 in the same range of temperatures. $\varepsilon - 3P$ is often called the “interaction measure”, but this terminology is quite misleading since both noninteracting quarks and gluons on the one hand and very strongly interacting conformal matter on the other have $\varepsilon - 3P = 0$, with ε/T^4 and P/T^4 both independent of temperature. Large values of $(\varepsilon - 3P)/T^4$ necessarily indicate significant interactions among the constituents of the plasma, but small values of this quantity should in no way be seen as indicating a lack of such interactions. We see in the figure that $(\varepsilon - 3P)/T^4$ rises rapidly in the vicinity of T_c . This rapid rise corresponds to the fact that ε/T^4 rises more rapidly than $3P/T^4$, approaching roughly 80% of its value in a noninteracting gas of quarks and gluons at a lower temperature, between 200 and 250 MeV. At higher temperatures, as $3P/T^4$ rises toward roughly 80% of its noninteracting value, $(\varepsilon - 3P)/T^4$ falls off with increasing temperature and the quark–gluon plasma becomes more and more conformal. Remarkably, after a proper rescaling of the number of degrees of freedom and T_c , all the features described above remain the same when the number of colors of the gauge group is increased and extrapolated to the $N_c \rightarrow \infty$ limit [198, 306, 663].

The central message for us from these lattice calculations of the QCD equation of state is that at high enough temperatures the thermodynamics of the QCD plasma becomes conformal while deviations from conformality are most severe at and just above T_c . This suggests that the use of conformal theories (in which calculations can be done via gauge/gravity duality as described in much of this book) as vehicles by which to gain insights into real-world quark–gluon plasma may be more quantitatively reliable when applied to data from heavy ion collisions at the LHC than when applied to those at RHIC. In this respect, it is also quite encouraging that the charged particle elliptic flow $v_2(p_T)$ measured very recently in heavy ion collisions at $\sqrt{s} = 2.76$ TeV at the LHC [5] is, within error bars, the same as that measured at RHIC. On a qualitative level, this indicates that the quark–gluon plasma produced at the LHC is comparably strongly coupled to that at RHIC.

One of the first questions to answer with a calculation of the equation of state in hand is whether the observed rapid rise in ε/T^4 and P/T^4 corresponds to a phase transition or to a continuous crossover. In QCD without quarks, a first order deconfining phase transition is expected due to the breaking of the Z_N center symmetry. This symmetry is unbroken in the confined phase and broken above T_c by a nonzero expectation value for the Polyakov loop [388, 769]. The expected first order phase transition is indeed seen in lattice calculations [184]. The introduction of quarks introduces a small explicit breaking of the Z_N symmetry even at low temperatures, removing this argument for a first order phase transition. However, in QCD with massless quarks there must be a sharp phase transition (first order with three flavors of massless quarks, second order with two) since chiral symmetry is spontaneously broken at low temperatures and unbroken at high temperatures. This argument for the necessity of a transition vanishes for quarks with nonzero masses, which break chiral symmetry explicitly even at high temperatures. So, the question of what happens in QCD with physical quark masses, two light and one strange, cannot be answered by any symmetry argument. Since both the center and chiral symmetries are explicitly broken at all temperatures, it is possible for the transition from a hadron gas to quark–gluon plasma as a function of increasing temperature to occur with no sharp discontinuities. And, in fact, lattice calculations have shown that this is what happens: the dramatic increase in ε/T^4 and P/T^4 occurs continuously [59]. This is shown most reliably via the fact that the peaks in the chiral and Polyakov loop susceptibilities are unchanging as one increases the physical spatial volume V of the lattice on which the calculation is done. If there were a first order phase transition, the heights of the peaks of these susceptibilities should grow $\propto V$ in the large V limit; for a second order phase transition, they should grow proportional to some fractional power of V . But, for a continuous crossover no correlation length diverges at T_c and all physical quantities, including the heights of these susceptibilities, should be independent of V once $V^{1/3}$ is larger than the longest correlation length. This is indeed what is found [59]. The fact that the transition is a continuous crossover means that there is no sharp definition of T_c , and different operational definitions can give different values. However, the analysis performed in [178] indicates that the chiral susceptibility and the Polyakov loop susceptibility peak in the range of $T = 150\text{--}170$ MeV.

Despite the absence of a phase transition in the mathematical sense, well above T_c QCD matter is deconfined, since the Polyakov loop takes on large nonzero values. In this high temperature regime, the matter that QCD describes is best understood in terms of quarks and gluons. This does not, however, imply that the interactions amongst the plasma constituents is negligible. Indeed, we have already seen in Section 2.2 that in the temperature regime accessible in heavy ion collisions at RHIC, the quark–gluon plasma behaves like a liquid, not at all like a gas

of weakly coupled quasiparticles. And, as we will discuss in Section 6.1, explicit calculations done via the AdS/CFT correspondence show that in the large- N_c limit in gauge theories with gravity duals which are conformal, and whose coupling can therefore be chosen, the thermodynamic quantities change by only 25% when the coupling is varied from zero (noninteracting gas) to infinite (arbitrarily strongly coupled liquid). This shows that thermodynamic quantities are rather insensitive to the strength of the interactions among the constituents (or volume elements) of quark–gluon plasma.

Finally, we note that calculations of QCD thermodynamics done via perturbative methods have been compared to the results obtained from lattice-regularized calculations. As is well known (see for example [556] and references therein), the expansion of the pressure in powers of the coupling constant g is a badly convergent series and, what is more, cannot be extended beyond order $g^6 \log(1/g)$, where nonperturbative input is required. This means that perturbative calculations must resort to resummations and indeed different resummation schemes have been developed over the years [185, 507, 454, 166, 50, 51, 52]. The effective field theory techniques developed in [185, 507], in particular, exploit a fundamental feature of any perturbative picture of the plasma: at weak coupling, the Debye screening mass $\mu_D \propto gT$ and these methods all exploit the smallness of μ_D relative to T since the basis of their formulation is that physics at these two energy scales is well separated. As we will see in Section 6.3, this characteristic is in fact essential for any description of the plasma in terms of quasiparticles. The analysis performed in Ref. [454] showed that in the region of $T \sim (1-3) T_c$ these effective field theory calculations of the QCD pressure become very sensitive to the matching between the scales μ_D and T , which indicates that there is no separation of these scales. This was foreshadowed much earlier by calculations of various different correlation lengths in the plasma phase which showed that at $T = 2T_c$ some correlation lengths that are $\propto 1/(g^2T)$ at weak coupling are in fact significantly *shorter* than others that are $\propto 1/(gT)$ at weak coupling [425], and showed that the perturbative ordering of these length scales is only achieved for $T > 10^2 T_c$. Despite the success of other resummation techniques [166] in reproducing the main features of QCD thermodynamics, the absence of any separation of scales indicates that there are very significant interactions among constituents and casts doubt upon any approach based upon the existence of quasiparticles.

3.1.1 Flavor susceptibilities

The previous discussion focused on thermodynamics in the absence of expectation values for any of the conserved (flavor) charges of QCD. As is well known, these charges are a consequence of the three flavor symmetries that QCD possesses: the $U(1)$ symmetries generated by electric charge, Q , and baryon number, B , and a

global $SU(3)$ flavor symmetry. Within $SU(3)$ there are two $U(1)$ subgroups which can be chosen as those generated by Q and by strangeness S . Conservation of Q is fundamental to the standard model, since the $U(1)$ symmetry is a gauge symmetry. Conservation of S is violated explicitly by the weak interactions and conservation of B is violated by exceedingly small nonperturbative weak interactions, and perhaps by yet to be discovered beyond standard model physics. As we are interested only in physics on QCD time scales, we can safely treat S and B as conserved. Instead of taking B , Q and S as the conserved quantities, we can just as well choose the linear combinations of them corresponding to the numbers of up, u , down, d , and strange quarks, s . With three conserved quantities, we can introduce three independent chemical potentials. In spite of the difficulties in studying QCD at nonzero chemical potential on the lattice, derivatives of the pressure with respect to these chemical potentials at zero chemical potential can be calculated. These derivatives describe moments of the distributions of these conserved quantities in an ensemble of volumes of quark–gluon plasma, and hence can be related to event-by-event fluctuations in heavy ion collision experiments.

When all three chemical potentials vanish, the lowest nonzero moments are the quadratic charge fluctuations, i.e. the diagonal and off-diagonal susceptibilities defined as

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

$$\chi_{11}^{XY} = \frac{1}{VT} \frac{\partial^2}{\partial \mu_X \partial \mu_Y} \log Z(T, \mu_X, \mu_Y, \dots) = \frac{1}{VT^3} \langle N_X N_Y \rangle, \quad (3.3)$$

where Z is the partition function and the N_X are the numbers of u , d or s quarks (or, equivalently, B , Q or S charge) present in the volume V . The diagonal susceptibilities quantify the fluctuations of the conserved quantum numbers in the plasma and the off-diagonal susceptibilities measure the correlations among the conserved quantum numbers, and are more sensitive to the nature of the charge carriers [540].

Lattice results for these quantities [180, 130] are shown in Fig. 3.2. In the top panel, the diagonal strange quark number susceptibility is shown as a function of temperature, at different lattice spacings and extrapolated to the continuum limit. The susceptibility is compared to its value in a noninteracting gas of gluons and quarks (dashed line) and to the expectation from a hadron resonance gas, which correctly describes the results of the lattice calculation at low temperature. Similarly to the case of the pressure, there is rapid rise in the susceptibility above T_c followed by saturation at high temperatures to a constant value that is below what it would be in a noninteracting gas. This rise, which reflects the liberation of s-quarks from hadrons, occurs over a similar range of temperatures as the rise in the pressure. The high temperature value of the susceptibility is about 90% of the Stefan–Boltzmann value, closer but not significantly closer to the noninteracting

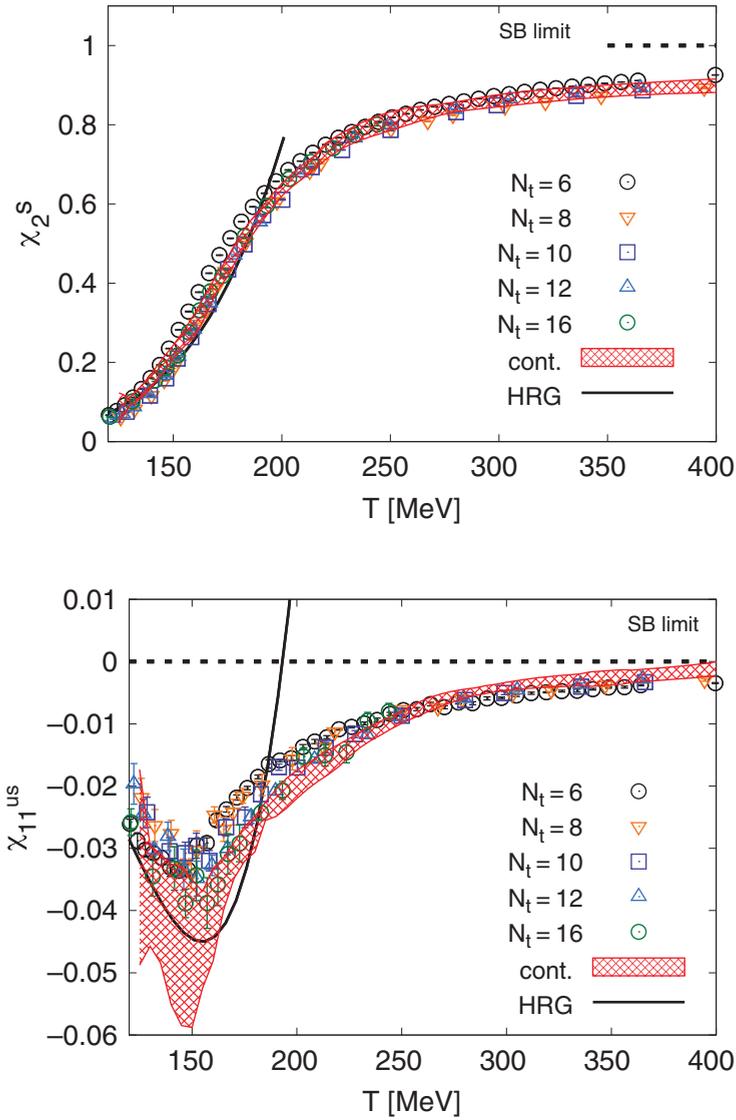


Figure 3.2 Top: quadratic fluctuations of the strange quark number. Bottom: off-diagonal susceptibility χ_{11}^{us} . In both panels the different symbols correspond to different lattice spacings. The red band is the continuum extrapolation. The continuous black line is the expectation from a hadron resonance gas and the dashed black line corresponds to the Stefan–Boltzmann (i.e. noninteracting) limit. Figures taken from [180].

limit than is the case for the pressure. Similar results are obtained for the u and d quark number susceptibilities [180]. In the bottom panel, one of the off-diagonal susceptibilities, χ_{11}^{us} , is shown. In a noninteracting plasma, the off-diagonal susceptibilities would all vanish. The results of these lattice calculations show significant

deviation from zero, with the off-diagonal value comparable to its value in a hadron resonance gas up to temperatures as high as $1.5 T_c$ and gradually approaching zero at higher temperatures.

The lattice calculations of susceptibilities further illustrate the fact that thermodynamic properties alone do not resolve the structure of the QGP, since they do not yield an answer to as simple a question as whether it behaves like a liquid or like a gas of quasiparticles. On the one hand, the diagonal susceptibilities seem close to the noninteracting limit and are, in fact, not incompatible with hard thermal loop computations [170], which supports a quasiparticle interpretation of the plasma already at these rather low temperatures. On the other hand, the off-diagonal susceptibilities are too large to be accommodated in perturbative calculations and it has even been suggested that they point towards the presence of meson-like states above deconfinement [707] (see also [740]). We will come back to the apparently contradictory pictures suggested by the different static properties of the plasma in Section 6.1.2.

3.2 Transport coefficients from the lattice

We turn now to lattice calculations that further determine the structure of the plasma via studying dynamical quantities rather than just static ones. The lattice calculation of dynamical quantities, which require time and therefore Minkowski spacetime in their formulation, are subject to the conceptual challenges that we described at the beginning of this section, meaning that the lattice results that we are going to discuss now come with caveats that we shall describe.

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

Transport coefficients can be extracted from the low momentum and low frequency limits of the Green's functions of a suitable conserved current of the theory, see Appendix A. To illustrate this point, we concentrate on two examples: the stress tensor components T^{xy} , and the longitudinal component of some conserved $U(1)$ current $J^i(\omega, \mathbf{k})$ which can be written $J(\omega, k) \hat{k}$, with ω, k the Fourier modes, and \hat{k} a vector of unit length. The stress tensor correlator determines the shear viscosity; the current–current correlator determines the diffusion constant for the conserved charge associated with the current. (The conserved charge could be baryon

number, strangeness or electric charge in QCD or could be some R -charge in a supersymmetric theory.) The retarded correlators of these operators are defined by

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

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

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

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

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

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

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

if we define a quantity μ by

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

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

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

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

where the Euclidean operator is defined from its Minkowski counterpart by

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

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

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

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

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

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

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

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

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

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

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

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

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

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

The results of lattice computations for the shear correlator are shown in the top panel of Fig. 3.3. The finite temperature Euclidean correlator is normalized to the free theory correlator at the same temperature. The measured correlator deviates

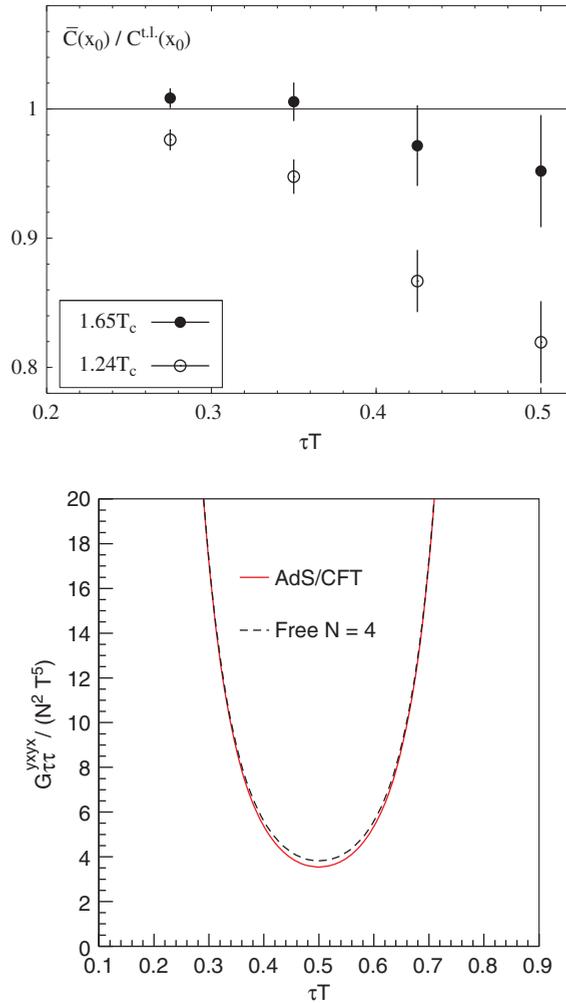


Figure 3.3 Top panel: ratio of the stress tensor Euclidean correlator calculated on the lattice in Ref. [614] to that in the free theory for QCD with three colors and zero flavors at four values of the Euclidean time $x_0 = \tau$ and two temperatures T . This theory has a first order deconfinement transition, and T is given in units of the critical temperature T_c for this transition. Bottom panel: stress tensor Euclidean correlator for $\mathcal{N} = 4$ SYM from Ref. [777]. The solid red line corresponds to infinite coupling and the dashed black line corresponds to the free theory. The solid curves are the zero temperature potential.

from the free one only by about 10%–20%. The statistical errors in the numerical computation illustrate that it is hard to distinguish the computed correlator from the free one, specially at the higher temperature. It is important to stress that the fact that the measured correlator is close to the free one comes from the fact that both receive a large contribution from the large ω region of the spectral function,

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

The extraction of transport information from the four points in the upper panel of Fig. 3.3, as done in Ref. [614], requires assumptions about the spectral density. Since the high frequency behavior of the spectral function is fixed due to asymptotic freedom, a first attempt can be made by writing

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

where

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

is the free theory result at the high frequencies where this result is valid. In the analysis performed in [614], the parameter Λ is always chosen to be $\geq 5T$. The functional form of the low frequency part ρ_{LF} should be chosen such that ρ_{LF} vanishes at high frequency. A Breit–Wigner ansatz

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

provides a simple example with which to start (and is in fact the form that arises in perturbation theory [10]). This ansatz does not provide a good fit, but it nevertheless

yields an important lesson. Fitting the parameters in this ansatz to the lattice results for G_E at four values of τ favors large values (larger than T) for the width $\Gamma = 2/b$ of the low frequency Breit–Wigner structure. This result motivates the assumption that the width of any peak or other structure at low frequency must be larger than T . From this assumption, a bound may be derived on the viscosity as follows. Since a wider function than a Breit–Wigner peak of width $\Gamma = T$ would lead to larger value of ρ_{LF} for $\omega < \sqrt{2}T$ and since the spectral function is positive definite, we have

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

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

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

with s the entropy density [614]. The idea here is: (i) we know how much the $\omega > \Lambda$ region contributes to the integral $\int d\omega \rho(\omega) / \sinh \omega/2T$ which is what the lattice calculation determines, and (ii) we make the motivated assumption that the narrowest a peak at $\omega = 0$ can be is T , and (iii) we can therefore put an upper bound on $\rho(0)$ by assuming that the entire contribution to the integral that does not come from $\omega > \Lambda$ comes from a peak at $\omega = 0$ with width T . The bound is conservative because it comes from assuming that ρ is zero at intermediate ω between T and Λ . Surely ρ receives some contribution from this intermediate range of ω , meaning that the bounds on η/s obtained from this analysis are conservative.

Going beyond the conservative bound (3.22) and making an estimate of η is challenging, given the finite number of points at which $G_E(\tau)$ is measured, and relies on physically motivated parameterizations of the spectral function. A sophisticated parameterization was introduced in Ref. [614] under the basic assumption that there are no narrow structures in the spectral function, which is supported by the Breit–Wigner analysis discussed above. In Ref. [614], the spectral function was expanded in an ordered basis of orthonormal functions with an increasing number of nodes, defined and ordered such that the first few functions are those that make the largest contribution to the Euclidean correlator; in other words the latter is most sensitive to the contribution of these functions. Owing to the finite number of data points and their finite accuracy, the basis has to be truncated to the first few functions, which is a way of formalizing the assumption that there are no narrow structures in the spectral function. The analysis based on such parameterization leads to small values of the ratio of the shear viscosity to the entropy density. In particular,

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

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

The lattice studies to date must be taken as exploratory, given the various difficulties that we have described. As explained in Ref. [616], there are ways to do better (in addition to using finer lattices and thus obtaining G_E at more values of τ). For example, a significant improvement may be achieved by analyzing the spectral function at varying nonzero values of the momentum k . One can then exploit energy and momentum conservation to relate different Euclidean correlators to the same spectral function, in some cases constraining the same spectral function with 50–100 quantities calculated on the lattice rather than just four. Furthermore, the functional form of the spectral function is predicted order by order in the hydrodynamic expansion and this provides guidance in interpreting the Euclidean data. These analyses are still in progress, but results reported to date [616] are consistent with (3.23), given the error estimate therein.

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

3.3 Quarkonium spectrum from the lattice

Above the critical temperature, quarks and gluons are not confined. As we have discussed at length in Chapter 2, experiments at RHIC have taught us that in this

regime QCD describes a quark–gluon plasma in which the interactions among the quarks and gluons are strong enough to yield a strongly coupled liquid. It is also possible that these interactions can result in the formation of bound states within the deconfined fluid [739]. This observation is of particular relevance for quarkonium mesons formed from heavy quarks in the plasma, namely quarks with $M \gg T$. For these quarks, $\alpha_s(M)$ is small and the zero temperature mesons are, to a first approximation, described by a Coulomb-like potential between a $Q-\bar{Q}$ pair. Thus, the typical radius of the quarkonium meson is $r_M \sim 1/\alpha_s M \ll 1/T$. As a consequence, the properties of these quarkonium mesons cannot be strongly modified in the plasma. Quarkonia are therefore expected to survive as bound states up to a temperature that is high enough that the screening length of the plasma has decreased to the point that it is of order the quarkonium radius [609].

The actual masses of the heavy quarks that can be accessed in heavy ion collisions, the charm and the bottom, are large enough that charmonium and bottomonium mesons are expected above the deconfinement transition, but they are not so large that these mesons are expected to be unmodified by the quark–gluon plasma produced in ultra-relativistic heavy ion collisions. As discussed in detail in Section 2.4, data indicate that heavy ion collisions at RHIC (at the LHC) reach temperatures high enough to dissociate all but the lowest lying 1s charmonium (bottomonium) states, see also Fig. 2.16. Moreover, while charmonia are not expected to survive in the quark–gluon plasma produced at the LHC, they may be regenerated when the plasma hadronizes since several dozen charm and anticharm quarks are expected in each LHC collision. It is a non-trivial challenge to determine what QCD predicts for the temperatures up to which a particular quarkonium meson survives as a bound state, and above which it dissociates. In this section, we describe the results of lattice QCD calculations done with this goal in mind. This is a subject of ongoing research, and definitive results for the dissociation temperatures of various quarkonia are not yet in hand. For an example of a recent review on this subject, see Ref. [131].

Some of the earliest [609, 520] attempts to describe the in-medium heavy mesons are based on solving the Schrödinger equation for a pair of heavy quarks in a potential determined from a lattice calculation. These approaches are known generically as potential models. In this approach, it is assumed that the interactions between the quark–antiquark pairs and the medium can be expressed in the form of a temperature-dependent potential. The mesons are identified as the bound states of quarks in this potential. Such an approach has been very successful at zero temperature [335] and in this context it can be put on firm theoretical grounds by means of a non-relativistic effective theory for QCD [681, 188]. However, at nonzero temperature it is not clear how to determine this potential from first principles. (For some attempts in this direction, see Ref. [189].)

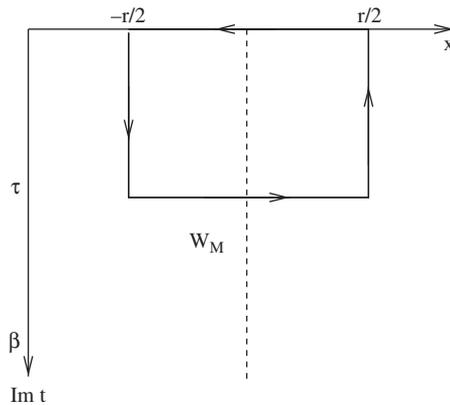


Figure 3.4 Wilson line representing the propagation of a heavy quark–antiquark pair. The line at $-r/2$ is the heavy quark propagator in imaginary time while the line at $r/2$ is the antiquark. The space links ensure gauge invariance. The singlet free energy is obtained by setting $\tau = \beta$.

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

At zero temperature, the extension of the Wilson lines in the imaginary time direction τ can be taken to infinity; this limit yields Wilson’s definition of the heavy quark potential [800]. In contrast, at nonzero temperature the imaginary time direction is compact and the imaginary time τ is bounded by $1/T$. Nevertheless, inspired by the zero temperature case the early studies *postulated* that the potential should be obtained from the Wilson line with $\tau = \beta = 1/T$. This Wilson line can be interpreted as the singlet free energy of the heavy quark pair, i.e. the energy change in the plasma due to the presence of a pair of quarks at a fixed distance and at fixed temperature [680, 641].

Lattice results for the singlet free energy are shown in Fig. 3.5. In the upper panel we show results for the gluon plasma described by QCD without any quarks

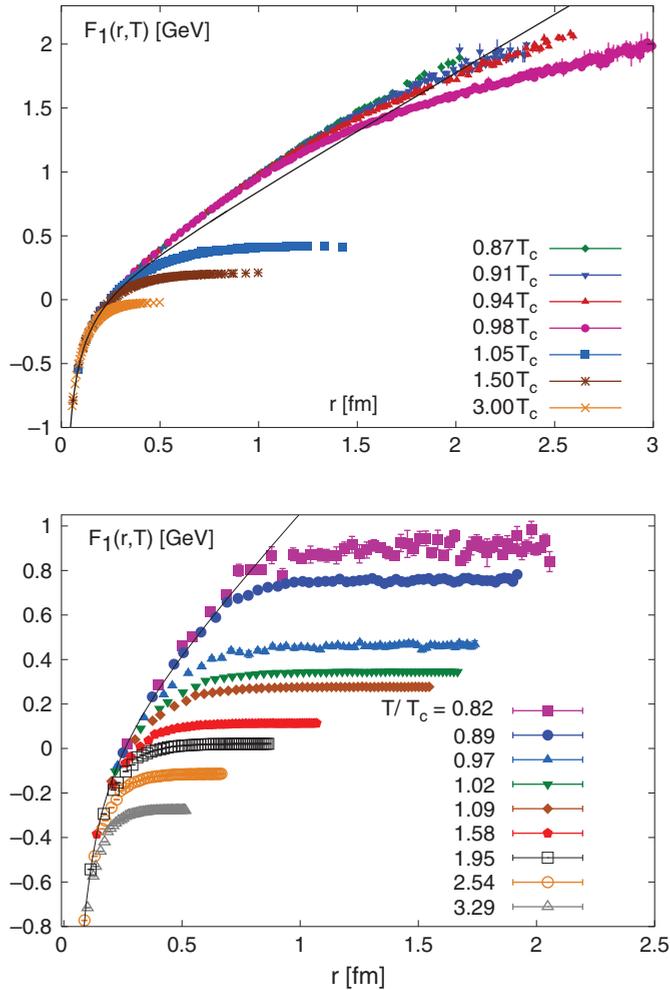


Figure 3.5 Lattice results for the singlet free energy $F_1(\mathbf{r}, T)$ as a function of the distance r for different temperatures T , quoted as fractions of the critical temperature T_c at which the crossover from hadron gas to quark–gluon plasma occurs. The solid curves are the zero temperature potential. The upper panel shows results for QCD without quarks [502, 503, 504] and the lower panel for 2+1 flavor QCD [505]. The fact that below T_c the free energy goes above the zero temperature result is a lattice artifact [189]. Figures taken from Refs. [131, 678].

[502, 503, 504]. The solid black line in this figure denotes the $T = 0$ result, which rises linearly with the separation r at large r , as expected due to confinement. The potential is well approximated by the ansatz

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

where the linear long-distance part is characterized by the string tension $\sqrt{\sigma} = 420 \text{ MeV}$ [647] and the perturbative $1/r$ piece describes the short-distance regime. Below T_c , as the temperature increases the theory remains confined but the string tension decreases. For temperatures larger than T_c , the theory is not confined and the free energy flattens to a finite value in the large- r limit. At these temperatures, the color charge in the plasma screens the interaction between the heavy Q and \bar{Q} . In QCD with light dynamical quarks, as in the lower panel of Fig. 3.5 from Ref. [505], the situation is more complicated. In this case, the free energy flattens to a finite limit at large distance even at zero temperature, since once the heavy quark and antiquark have been pulled far enough apart it becomes favorable to produce a light $q\bar{q}$ pair from the medium (in this case the vacuum) which results in the formation of $Q\bar{q}$ and $\bar{Q}q$ mesons that can then be moved far apart without any further expenditure of energy. In vacuum this process is usually referred to as “string-breaking”. In vacuum, at distances that are small enough that string-breaking does not occur the potential can be approximated by (3.24), but with a reduced string tension $\sqrt{\sigma} \approx 200 \text{ MeV}$ [277]. Above T_c , the potential is screened at large distances by the presence of the colored fluid, with the screening length beyond which the potential flattens shrinking with increasing temperature, just as in the absence of quarks. As a consequence, in the lower panel of the figure the potential evolves relatively smoothly with increasing temperature, with string-breaking below T_c becoming screening at shorter distance scales above T_c . The decrease in the screening length with increasing temperature is a generic result, and it leads us to expect that quarkonium mesons dissociate when the temperature is high enough that the vacuum quarkonium size corresponds to a quark–antiquark separation at which the potential between the quark and antiquark is screened [609].

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

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

Analyses performed with this potential indicate that the J/ψ meson survives deconfinement, existing as a bound state up to a dissociation temperature that lies in the range $T_{\text{diss}} \sim (1.5\text{--}2.5) T_c$ [806, 39, 740, 621, 40]. It is also a generic feature of these potential models that, since they are larger in size, other less bound charmonium states like the χ_c and ψ' dissociate at a lower temperature [520], typically at temperatures as low as $T = 1.1 T_c$. Let us state once more that these

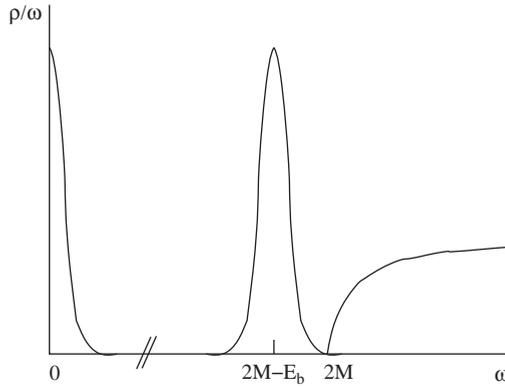


Figure 3.6 Schematic view of the current–current spectral function as a function of frequency for heavy quarks. The structure at small frequency $\omega \sim 0$ is the transport peak which is due to the interaction of the external current with heavy quarks and antiquarks from the plasma. At $\omega =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

calculations are based on two key model assumptions: first, that the charm and bottom quarks are heavy enough for a potential model to apply and, second, that the potential is given by Eq. (3.25). Neither assumption has been demonstrated from first principles.

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

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

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

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

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

The current–current correlator can be understood as describing the interaction of an external vector meson which couples only to heavy quarks in the plasma. This interaction can proceed by scattering with the heavy quarks and antiquarks

present in the plasma or by mixing the singlet quarkonium meson with (light quark) states from within the plasma that have the same quantum numbers as the external quarkonium. The first physical process leads to a large absorption of those vector mesons in which the ratio ω/q matches the velocity of heavy quarks in the medium, yielding the so-called transport peak at small ω . The second physical process populates the near-threshold region of $\omega \sim 2M$. Since the thermal distribution of the velocity of heavy quarks and antiquarks is Maxwellian with a mean velocity $v \sim \sqrt{T/M}$, the transport peak is well-separated from the threshold region. Thus, the spectral function contains information not only about the properties of mesons in the medium, but also about the transport properties of the heavy quarks in the plasma. A sketch of the general expectation for this spectral function is shown in Fig. 3.6. Given these expectations, the extraction of properties of quarkonium mesons in the plasma from the Euclidean correlator must take into account the presence of the transport peak. It is worth mentioning that for the particular case of pseudoscalar quarkonia, the transport peak is suppressed by mass [9]; thus, the extraction of meson properties is simplest in this channel. All other channels, including in particular the vector channel, include contributions from the transport peak.

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

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

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

In Fig. 3.7 we show the ratio of the computed lattice correlator G_E to G_{recon} defined in (3.28) in the pseudoscalar and vector channels for charm quarks [491]. The temperature dependence of this ratio is due only to the temperature dependence of the spectral function. The pseudoscalar correlator shows little temperature dependence up to temperatures as high as $T = 1.5 T_c$ while the vector correlator varies significantly in that range of temperatures. Since, as we have already mentioned, the transport peak is suppressed in the pseudoscalar channel, the lack of temperature dependence of the Euclidean correlator in this channel can be interpreted as a signal of the survival of pseudoscalar charmed mesons (the η_c) above deconfinement. However, the Euclidean correlator is a convolution integral over the spectral density and the thermal kernel, and in principle the spectral density

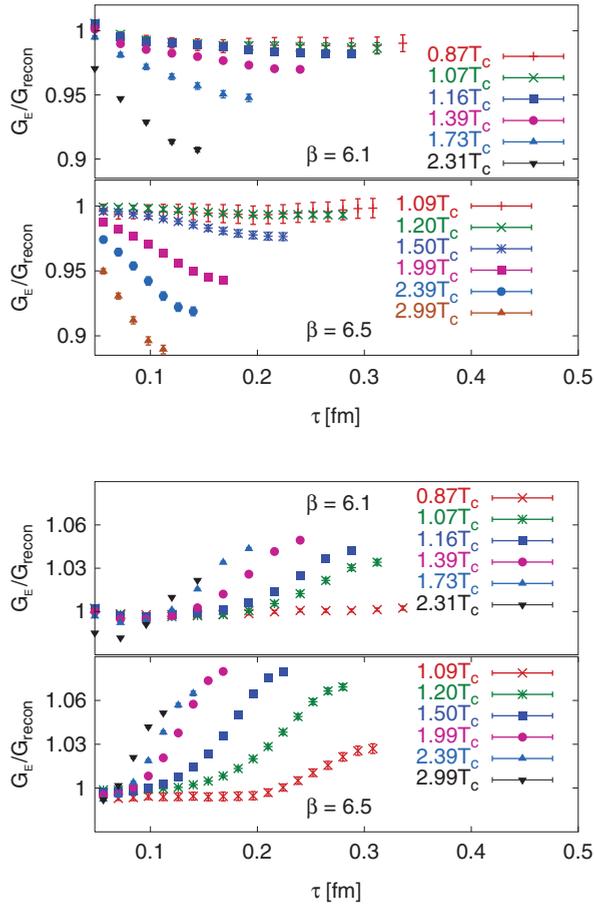


Figure 3.7 The ratio of the Euclidean correlator G_E to G_{recon} defined in (3.28) in the pseudoscalar (top) and vector (bottom) channels for charm quarks versus the imaginary time τ [491]. Note that the transport contribution is suppressed by the mass of the charmed quark only in the pseudoscalar channel. Figure taken from Ref. [491].

can change radically while leaving the convolution integral relatively unchanged. So, the spectral function must be extracted before definitive conclusions can be drawn.

There has been a lot of effort towards extracting these spectral densities in a model-independent way directly from the Euclidean correlators. The method that has been developed the furthest is called the Maximum Entropy Method (MEM) [81]. It is an algorithm designed to find the most probable spectral function compatible with the lattice data on the Euclidean correlator. This problem is underconstrained, since one has available lattice calculations of the Euclidean correlator only at finitely many values of the Euclidean time τ , each with error bars,

and one is seeking to determine a function of ω . This means that the algorithm must take advantage of prior information in the form of a default model for the spectral function. Examples of priors that are taken into account include information about asymptotic behavior and sum rules. The MEM method has been very successful in extracting the spectral functions at zero temperature, where it turns out that the extracted functions have little dependence on the details of how the priors are implemented in a default model for the spectral function. The application of the same MEM procedure at nonzero temperature is complicated by two facts: the number of data points is smaller at finite temperature than at zero temperature and the temporal extent of the correlators is limited to $1/T$, which is reduced as the temperature increases. The first problem is a computational problem, which can be ameliorated over time as computing power grows by reducing the temporal lattice spacing and thus increasing the number of lattice points within the extent $1/T$. The second problem is intrinsic to the nonzero temperature calculation; all the structure in the Minkowski space spectral function, as in the sketch in Fig. 3.6, gets mapped onto fine details of the Euclidean correlator within a small interval of τ meaning that at nonzero temperature it takes much greater precision in the calculation of the Euclidean correlator in order to disentangle even the main features of the spectral function.

To date, extractions of the pseudoscalar spectral density at nonzero temperature via the MEM indicate, perhaps not surprisingly, that the spectral function including its η_c peak remains almost unchanged up to $T \approx 1.5 T_c$ [80, 308, 491, 11], especially when the comparison that is made is with the zero temperature spectral function extracted from only a reduced number of points on the Euclidean correlator. The application of the MEM to the vector channel also indicates survival of the J/ψ up to $T \approx 1.5-2 T_c$ [80, 308, 491, 11, 322], but it fails to reproduce the transport peak that must be present in this correlator near $\omega = 0$. It has been argued that most of the temperature dependence of the vector correlator seen in Fig. 3.7 is due to the temperature dependence of the transport peak [782]. (Note that since the transport peak is a narrow structure centered at zero frequency, it corresponds to a temperature-dependent contribution to the Euclidean correlator that is approximately τ -independent.) This is supported by the fact that the τ -derivative of the ratio of correlators is much less dependent on T [307] and by the analysis of the spectral functions extracted after introducing a transport peak in the default model of the MEM [322, 325]. When the transport peak is taken into account the MEM also shows that J/ψ may survive at least up to $T = 1.5 T_c$ [322]. However, above T_c both the vector and the pseudoscalar channels show strong dependence (much stronger than at zero temperature) on the default model via which prior information is incorporated in the MEM [491, 322], which makes it difficult to extract solid conclusions on the survival of charmonium states from this method. For the

bottomonium family, the difficulties which arise from the presence of a transport peak remain. However, their influence on the lattice determination of the survival of the different states is reduced as a result of the larger mass of the b-quark, which suppresses this contribution. The large b-quark mass also makes it possible to perform a different set of lattice studies in which these uncertainties are reduced by starting from a non-relativistic effective field theory for the bottomonia states in which the transport and bound state regions of the spectral function are explicitly decoupled. This method provides access not only to the ground states of the different bottomonium channels but also to their excited states. In the vector channel, in particular, the calculations in Ref. [8] indicate that the Υ survives up to temperatures higher than $2 T_c$ while the first excited state, the $\Upsilon(2S)$, disappears below $1.5 T_c$. Notwithstanding the uncertainties of the method, the conclusions of the MEM analyses agree with those reached via analyses of potential models in which the internal energy (3.25) is used as the potential. However, before this agreement can be taken as firm evidence for the survival of the different quarkonia states well above the phase transition, it must be shown that the potential models and the lattice calculations are compatible in other respects. To this we now turn.

Potential models can be used for more than determining whether a temperature-dependent potential admits bound states: they provide a prediction for the entire spectral density. It is then straightforward to start with such a predicted spectral density and compute the Euclidean correlator that would be obtained in a lattice calculation if the potential model correctly described all aspects of the physics. One can then compare the Euclidean correlator predicted by the potential model with that obtained in lattice computations like those illustrated in Fig. 3.7. Following this approach, the authors of Refs. [623, 622] have shown that neither the spectral function obtained via identifying the singlet internal energy as the potential nor the one obtained via identifying the singlet free energy as the potential correctly reproduce the Euclidean correlator found in lattice calculations. This means that conclusions drawn based upon either of these potentials cannot be quantitatively reliable in all respects. These authors then proposed a more phenomenological approach, constructing a phenomenological potential (containing many of the qualitative features of the singlet free energy but differing from it) that reproduces the Euclidean correlator obtained in lattice calculations at the percent level [623, 622]. These authors also point out that at nonzero temperature all putative bound states must have some nonzero thermal width, and states whose binding energy is smaller than this width should not be considered bound. These considerations lead the authors of Refs. [623, 622] to conclude that the J/ψ and η_c dissociate by $T \sim 1.2T_c$ while less bound states like the χ_c or ψ' do not survive the transition at all. These conclusions differ from those obtained via the MEM. Although these conclusions are dependent on the potential used, an important and lasting lesson from this work is

that the spectral function above T_c can be very different from that at zero temperature even if the Euclidean correlator computed on the lattice does not show any strong temperature dependence. This lesson highlights the challenge, and the need for precision, in trying to extract the spectral function from lattice calculations of the Euclidean correlator in a model independent fashion.

Finally, we note once again that there is no argument from first principles for using the Schrödinger equation with either the phenomenological potential of Refs. [623, 622] or the internal energy or the free energy as the potential. To conclude this section, we would like to add some remarks on why the identification of the potential with the singlet internal or singlet free energy cannot be correct [563, 137]. If the quarkonium states can be described by a Schrödinger equation, the current–current correlator must reduce to the propagation of a quark–antiquark pair at a given distance r from each other. The correlator must then satisfy

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

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

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

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

The difficulty of extracting the correct potential resides in the analytic continuation from $V(\tau, \mathbf{r})$ to $V(t, \mathbf{r})$. At zero temperature, τ is not periodic and we can take the $\tau \rightarrow \infty$ limit and relate what we obtain to V_∞ . At nonzero temperature, τ is periodic and so there is no $\tau \rightarrow \infty$ limit. It is also apparent that V_∞ need not coincide with the value of $V(\beta, r)$ as postulated in some potential models; in fact, due to the periodicity of τ a lot of information is lost by setting $\tau = \beta$ [137]. Explicit calculations within perturbative thermal field theory, where the analytic continuation can be performed, show that V_∞ does not coincide with the internal energy (3.25) and, what is more, the in-medium potential develops a r -dependent imaginary part which can be interpreted as the collisional dissociation of the state in the plasma via processes in which momentum is exchanged with the plasma

constituents [563]. Potential model analyses, analogous to those described above but which use complex potentials inspired by the perturbative quantum field theoretical calculations, show that the imaginary part of the potential dominates the melting process of the different quarkonia states, since they become broad well before their binding energy vanishes [563] and before their size becomes comparable to the Debye screening length of the plasma [219].¹ Nevertheless, the Euclidean correlators extracted from these potentials deviate from those obtained in lattice calculation [679], as expected since the calculation of these potentials assumes a weakly coupled medium while the lattice calculations include the full thermodynamics of the strongly coupled plasma in QCD.

If the log of the Wilson loop W_M is a quadratic function of the gauge potential, as in QED or in QCD to leading order in perturbation theory, then it is possible to show that the real part of the potential agrees with the singlet free energy [137]; however, this is not the case in general and an analysis which goes beyond perturbation theory is needed. Complementary to the analysis of current correlators, there have been some attempts to extract the heavy quark potential from lattice studies [717, 221]. These analyses are based on an analytical continuation of the numerical computation of the potential (3.30) to real time via a spectral analysis of the lattice Wilson loop

$$W_M(\tau, \mathbf{r}) = \int_{-\infty}^{\infty} d\omega e^{-\omega\tau} \rho_M(\omega, \mathbf{r}). \quad (3.31)$$

The numerical inversion of this integral poses the same challenges as the extraction of the spectral density from the meson correlators we have described above. Nevertheless, current attempts in determining V_∞ from this spectral analysis show that the real part of the potential does indeed deviate from the singlet free energy above T_c , and also show a very strong increase in the imaginary part of the potential at large distances, r . While further studies are needed before V_∞ can be extracted accurately from the lattice, these calculations show that the many-body effects that lead to complex potentials are indeed of relevance for the dynamics of quarkonium mesons in the plasma. For this reason, in recent years a new approach to these dynamics based on describing the heavy quark pair as an open quantum system

¹ If the plasma is weakly coupled, and if the quark mass M is heavy enough that the quarkonium mesons are small enough that physics at the scales given by both their size and their binding energy is weakly coupled, then all the relevant scales are distinct. In this case, the Debye screening length of the plasma would become comparable to the size of the quarkonium meson at a temperature that is of order gM but the imaginary part of the potential becomes comparable to the binding energy of the meson first, at a lower temperature that is of order $g^{4/3} \left(\log \frac{1}{g}\right)^{-1/3} M$ [344, 562]. In a strongly coupled plasma in which g is large, these scales need not be separated. Indeed, the calculations that we shall describe in Section 9.4.2 indicate that in the strongly coupled plasma in $\mathcal{N} = 4$ SYM theory in which the physics of “quarkonium” mesons can be investigated using gauge/gravity duality there is no parametric difference between the scales where the imaginary part of the potential gets large and where the screening length of the plasma becomes comparable to the meson size.

has been pursued. This approach has great potential for the description of the real-time evolution of heavy states in plasma since it is based on a stochastic approach. A detailed description of this method [814, 813, 176, 31] goes beyond the scope of this book.

Finally, the collisional dissociation processes that lead to imaginary potentials have also impacted the modern description of quarkonium production in heavy ion collisions. In current modeling, these dynamics have been introduced either by including collisional dissociation processes into the traditional rate equations for the different states [337, 759], or via new attempts to describe the entire quarkonium evolution via a potential approach [766, 765, 248]. The latter approach, which is particularly suitable for bottomonium states due to their larger masses, has been very successful in describing the suppression pattern of the Υ family in heavy ion collisions at the LHC, which we have described in Section 2.4.