

PARTICLE-THEORY INPUT
FOR NEUTRON-STAR PHYSICS* **

ALEKSI VUORINEN

Department of Physics and Helsinki Institute of Physics
P.O. Box 64, 00014, University of Helsinki, Finland*Received 7 May 2024, accepted 9 May 2024,
published online 22 May 2024*

Understanding the properties and physical phase of the dense strongly interacting matter present in the cores of neutron stars or created in their binary mergers remains one of the most prominent open problems in nuclear astrophysics. While most microscopic analyses have historically relied on solvable phenomenological models of nuclear and quark matter, in recent years a model-independent approach utilizing only controlled *ab-initio* calculations and astrophysical observations has emerged as a viable alternative.

In these lecture notes, I review recent progress in first-principles weak-coupling calculations within high-density quark matter, shedding light on its thermodynamic and transport properties. I cover the most important technical tools used in such calculations, introduce selected highlight results, and explain how this information can be used in phenomenological studies of neutron-star physics. The notes do not offer a self-consistent treatment of the topics covered, but rather aim at filling gaps in existing textbooks on thermal field theory and at connecting the dots in a story developed in several recent research articles, to which the interested reader is directed for further technical details.

DOI:10.5506/APhysPolB.55.4-A4

1. Introduction

The qualitative idea that a competition between gravity and the degeneracy pressure of matter composed of nucleons could give rise to compact astrophysical objects dates back to more than 90 years [1], and the first direct observation of rapidly rotating pulsars will soon reach the ripe age of

* Presented at the LXIII Cracow School of Theoretical Physics *Nuclear Matter at Extreme Densities and High Temperatures*, Zakopane, Poland, 17–23 September, 2023.

** Funded by SCOAP³ under Creative Commons License, CC-BY 4.0.

60 [2]. Yet, it is only during the past 10 years that neutron stars (NSs) have truly become a functioning laboratory of dense nuclear and quark matter (NM and QM), owing largely to recent dramatic advances in observational astrophysics. Milestone results from the 2010s and 2020s include accurate mass measurements of several individual high-mass NSs [3–5], often taking advantage of general relativistic effects such as the Shapiro Delay [6]; increasingly precise radius measurements utilizing X-ray emission from NSs [7–12]; and most famously the first-ever detection of a gravitational-wave (GW) signal from a binary NS merger by the LIGO and Virgo collaborations in 2017 [13–15]. In addition to the first GW signal, GW170817, the same merger event gave rise to an associated electromagnetic signal across a wide spectrum, recorded by altogether 70 different observatories [16, 17]. This marked the dawn of an era of multimessenger astronomy, the future of which looks bright, with several new GW observatories capable of recording a postmerger signal, including the Einstein Telescope [18] and the Cosmic Explorer [19], being planned at the moment.

In order to draw robust microphysics lessons from the increasing amount of high-quality observational data on NS properties, it is imperative to simultaneously develop our theoretical understanding of matter at supernuclear densities. Here, the dominant physical effects are described by the theory of the strong nuclear interaction, Quantum Chromodynamics (QCD), with sub-leading but important roles played by the electromagnetic and weak sectors of the Standard Model. Due to the sign problem of lattice QCD at nonzero baryon chemical potentials μ_B [20], a nonperturbative first-principles approach available at all relevant densities is unfortunately unavailable, leaving us with effective-theory frameworks and weak-coupling expansions to work with. Indeed, the *ab-initio* tools available in the cold and dense part of the QCD phase diagram include Chiral Effective Theory (CET), valid for moderate-density NM up to somewhat above the nuclear saturation density $n_s = 0.16/\text{fm}^3$ (see, *e.g.*, [21]); perturbative QCD (pQCD), available beyond some tens of saturation densities [22, 23]; and with some reservations, methods such as Functional Renormalization Group [24] the AdS/CFT conjecture that allows access to the strongly coupled regime of many QCD-like theories albeit not QCD itself [25, 26].

In the lecture notes at hand, our focus will be on the physics of ultra-dense QM. More specifically, we aim at providing an introduction to the methods used in recent perturbative determinations of the thermodynamic and transport properties of dense unpaired QM and briefly review the application of these results to phenomenological studies of NS matter. While we will at times perform brief example calculations, these notes are not meant to serve as a stand-alone textbook. On the contrary, we will rely on existing treatments of the basics of thermal field theory and supplement them

with brief introductions to computational methods specific to finite density, pointing out the most useful references where they have been developed or applied at the state-of-the-order level.

The remaining seven sections of the notes cover the following topics:

2. Basics of perturbative thermal field theory at finite density;
3. Soft effective theories at finite temperature and density;
4. The Hard Thermal Loop effective theory;
5. High-order perturbative calculations at finite density;
6. Transport properties of dense QM;
7. Model-independent studies of the NS-matter equation of state;
8. Outlook towards future developments.

In Sections 2–6, we will pay particular attention to thermal-field-theory techniques specific to finite density that have received limited attention in existing literature, while in Section 7, we will focus on the role of QM thermodynamics in recent model-independent determinations of the NS-matter equation of state (EoS).

In the perturbative calculations performed in these lecture notes, we shall mostly employ the imaginary-time formalism of thermal field theory, implying that our four-dimensional metric is Euclidean $\eta_{\mu\nu} = \delta_{\mu\nu}$. For all other conventions, we use the definitions specified in the preface of [22].

2. Basics of perturbative thermal field theory in dense QCD

Perturbation theory is without doubt the most popular and versatile computational tool in quantum field theory (QFT), with applications ranging from collider problems in vacuum (see, *e.g.*, [27]) to thermal-field-theory challenges in hot and/or dense environments such as the early Universe or the fireball created in a heavy-ion collision [23]. Within the latter context, the weak-coupling approach can be used to determine not only Euclidean “bulk thermodynamic” quantities but also transport coefficients [28, 29] and even time-dependent quantities related to, *e.g.*, thermalization dynamics [30].

In the context of NS physics, of primary interest in these lecture notes, the phenomenologically most important quantity is the EoS of dense beta-equilibrated QCD matter, *i.e.* the relationship between its pressure and energy density. This quantity can be shown to be in a rough one-to-one correspondence with the so-called NS mass–radius relation (see Section 7), implying that both NS observations and microphysical calculations can be

used in a model-independent inference of the properties of NS matter. Such calculations form the topic of Section 7 and serve as the primary motivation for the high-order perturbative determinations of QM thermodynamics that we will be studying in this and the following three sections.

In the realm of high temperatures, perturbative thermal field theory is a well-developed and mature field, with multiple textbooks devoted to the subject. For the imaginary-time formalism, we employ here two widely used sources include [22, 31], while for the real-time alternative, we can recommend [23, 32]. Given that these references touch upon cold and dense fermionic matter only very briefly (two exceptions being Section 7 of [22] and Section 6.4 of [23]), we will here start from a brief review of the basic equilibrium thermodynamic properties of such systems at low perturbative orders.

Any weak-coupling calculation in cold and dense QM proceeds as an expansion around a system of noninteracting quarks at sizable chemical potentials and small or vanishing temperature. The physical properties of such a system resemble those of the conduction electrons in a metal — a canonical example system in undergraduate statistical mechanics — and can be qualitatively understood based on the small- T behavior of the Fermi–Dirac distribution function

$$n_F(\epsilon) = \frac{1}{e^{(\epsilon-\mu)/T} + 1} \xrightarrow{T \rightarrow 0} \Theta(\mu - \epsilon), \quad (1)$$

with $\epsilon = \epsilon(p)$ denoting the free one-particle dispersion relation. The step-function form of this quantity implies that in the strict $T = 0$ limit, all quantum-mechanical states are filled inside a momentum-space Fermi sphere with a radius determined by the relation $\epsilon(p) = \mu$ and unoccupied outside it. This reflects the effect of the Pauli exclusion principle and is in stark contrast with the low-temperature behavior of bosonic systems, characterized by a condensation of particles to the lowest quantum state available [33].

Interaction corrections to the thermodynamic properties of the free system can be obtained by expanding the path-integral representation of the grand canonical partition function in powers of the coupling constant. For a relativistic gauge field theory coupled to Dirac fermions with separately conserved number densities (thus allowing the introduction of the corresponding chemical potentials)¹, this quantity takes the form [22]

¹ In the case of QCD, this implies neglecting the effects of flavor-changing weak interactions, which is a good approximation in heavy-ion collisions, but not inside neutron stars. We will return to this issue when discussing beta equilibrium below.

$$\begin{aligned}
Z_{\text{QCD}} = & \int_{\text{periodic}} \mathcal{D}A_0^a \mathcal{D}A_k^a \int_{\text{periodic}} \mathcal{D}\bar{c}^a \mathcal{D}c^a \int_{\text{anti-periodic}} \mathcal{D}\bar{\psi}_f \mathcal{D}\psi_f \\
& \times \exp \left\{ - \int_0^\beta d\tau \int_{\mathbf{x}} \left[\frac{1}{4} F_{\mu\nu}^a F_{\mu\nu}^a + \frac{1}{2\xi} G^a G^a + \bar{c}^a \left(\frac{\delta G^a}{\delta \theta^b} \right) c^b \right. \right. \\
& \left. \left. + \bar{\psi}_f (\gamma_\mu D_\mu + m_f - \mu_f \gamma_0) \psi_f \right] \right\}. \tag{2}
\end{aligned}$$

Here, we denote gluon fields in the adjoint representation of the SU(3) gauge group by A_μ^a , $a = 1, 2, \dots, 8$ being the corresponding color index; adjoint-representation ghosts by c^a ; fundamental-representation quarks of flavor f by ψ_f ; the covariant gauge parameter by ξ ; and the function defining the covariant gauge by $G^a \equiv -\partial_\mu A_\mu^a$. Finally, the field strength tensor and the fundamental-representation covariant derivative appearing above read

$$F_{\mu\nu}^a = \partial_\mu A_\nu^a - \partial_\nu A_\mu^a + g f^{abc} A_\mu^b A_\nu^c, \quad D_\mu = \partial_\mu - ig A_\mu^a T^a, \tag{3}$$

where f^{abc} and T^a stand for the structure constants and Hermitian generators of the gauge group. Boundary conditions of the fields in the imaginary time direction τ , running from 0 to $\beta \equiv 1/T$, are indicated in the functional integral above, and the Euclidean γ matrices and other necessary quantities are properly defined and listed in [22].

Just as at high temperatures, the weak-coupling expansion of the pressure, or the logarithm of the partition function, of a cold and dense system is organized in terms of a loop expansion in connected vacuum or bubble Feynman diagrams. The first two orders of the expansion proceed without complications, with the subtraction of μ -independent vacuum terms and the renormalization of quark masses and the gauge coupling g successfully removing all $1/\epsilon$ ultraviolet (UV) divergences encountered in dimensional regularization. Analogously to the high- T case, uncanceled infrared (IR) divergences, however, appear at the three-loop order, which necessitates the use of either explicit diagrammatic resummations or low-energy effective field theories (EFTs). Given that this will be the main topic of the next two sections of these notes, we will here leave this issue aside and simply point out some of the most important technical differences between loop calculations performed in the high- and low- T realms, relying on the reader's familiarity with the former context (based on textbooks such as [22]):

- While the use of the Matsubara formalism is required at all nonzero values of temperature in the imaginary-time formalism, both bosonic and fermionic sum-integrals become continuous integrals in $D = 4 - 2\epsilon$ dimensions in the strict $T \rightarrow 0$ limit. This implies that we may write

$$\begin{aligned}
\rlap{\not\int}_P &\rightarrow \int_{-\infty}^{\infty} \frac{dp_0}{2\pi} \int \frac{d^d \mathbf{p}}{(2\pi)^d} \equiv \int_{-\infty}^{\infty} \frac{dp_0}{2\pi} \int_{\mathbf{p}} \equiv \int_P \quad (\text{bosons}), \\
\rlap{\not\int}_{\{P\}} &\rightarrow \int_{-\infty+i\mu}^{\infty+i\mu} \frac{dp_0}{2\pi} \int \frac{d^d \mathbf{p}}{(2\pi)^d} \equiv \int_{\tilde{P}} \quad (\text{fermions}), \tag{4}
\end{aligned}$$

where the tilde in the latter integration measure reminds us of the fermionic nature of the corresponding momentum. In practical calculations, it is often advantageous to start the evaluation of the momentum integrals from the temporal ones, for which so-called cutting rules (see Section 5) offer a convenient book-keeping tool. Note also that many references use slightly differing conventions for the integration measures, with angular brackets often used to signify the fermionic nature of four-momenta.

- A convenient simplification occurring in the strict $T = 0$ limit is the exact vanishing of vacuum graphs containing no quark loops — a direct consequence of the massless nature of gluons and the vanishing of scalefree integrals in dimensional regularization. At higher loop orders, the same mechanism leads to the vanishing of vacuum diagrams that contain quark loops but also a factorized purely bosonic sub-diagram, such as the first 11 diagrams of Fig. 3.
- A subtlety not present in studies of the short-lived quark–gluon plasma (QGP) created in heavy-ion collisions that needs to be taken into account when considering the high-density matter inside NSs is related to maintaining chemical (beta) equilibrium and local charge neutrality². To achieve both limits, one typically needs to add electrons to the system and implement a number of constraints between the chemical potentials of the relevant particle species.

For the three lightest quark flavors present at NS densities, the requirement of beta equilibrium implies the relations $\mu_s = \mu_d$ and $\mu_u = \mu_d - \mu_e$ between the four chemical potentials present. This allows parametrizing the system in terms of only two chemical potentials, typically taken to be μ_d and μ_e , while local charge neutrality adds one more constraint

$$\frac{2}{3}n_u - \frac{1}{3}n_d - \frac{1}{3}n_s = n_e. \tag{5}$$

The last relation allows us to further solve μ_e as a function of μ_d .

² The reason one can neglect these effects in heavy-ion physics is related to the fact that the electromagnetic and weak interactions operate on much longer timescales than the strong-interaction processes.

In the simplified limit of three massless quark flavors, typically a good approximation for bulk thermodynamic quantities at perturbative densities, the above equations admit the simple solution of $\mu_u = \mu_d = \mu_s$ and $\mu_e = 0$. This means that both local charge neutrality and beta equilibrium can be maintained without the presence of electrons and with equal number densities for the three lightest quark flavors. Corrections to this limit due to a nonzero strange-quark mass can be conveniently obtained using the quark-mass expansion scheme recently introduced in [34].

To get some hands-on experience in concrete perturbative calculations at high density, let us next consider the first two orders in the weak-coupling expansion of the pressure of cold and dense QM. Up to a μ -independent vacuum part, the leading-order (LO) pressure of a system of noninteracting quarks at nonzero flavor-dependent quark chemical potentials μ_f but vanishing temperature can be written in the simple form (*cf.* Eq. (7.43) of [22])

$$p_{\text{QCD}}^{\text{LO}}(\{\mu_f\}) = 2N_c \sum_f \int_{-\infty}^{\infty} \frac{dp_0}{2\pi} \ln \left[(p_0 + i\mu_f)^2 + p^2 + m_f^2 \right], \quad (6)$$

where the number of colors N_c is kept unspecified for the sake of generality.

It is convenient to begin the calculation with the integral over p_0 . We could in principle perform it directly in the above logarithmic form, but a more straightforward strategy is to first carry out a differentiation with respect to $E_p^2 \equiv p^2 + m_f^2$ and later integrate the result over the same parameter. To this end, we consider the integral

$$\begin{aligned} \int_{-\infty}^{\infty} \frac{dp_0}{2\pi} \frac{1}{(p_0 + i\mu_f)^2 + E_p^2} &= \int_{-\infty}^{\infty} \frac{dp_0}{2\pi} \frac{1}{(p_0 + i\mu_f - iE_p)(p_0 + i\mu_f + iE_p)} \\ &= \frac{\Theta(E_p - \mu_f)}{2E_p}, \end{aligned} \quad (7)$$

where we closed the integration contour over the upper halfplane. Integrating this expression with respect to E_p^2 from $E_p^2 = \mu^2$, where the contributions to the pressure vanish on physical grounds (this represents the quark mass threshold), we easily obtain

$$\int_{-\infty}^{\infty} \frac{dp_0}{2\pi} \ln \left[(p_0 + i\mu_f)^2 + p^2 + m_f^2 \right] = (E_p - \mu_f) \Theta(E_p - \mu_f) \quad (8)$$

and further

$$\begin{aligned}
p_{\text{QCD}}^{\text{LO}}(\{\mu_f\}) &= 2N_c \sum_f \int_{\mathbf{p}} (E_p - \mu_f) \Theta(E_p - \mu_f) \\
&= 2N_c \sum_f \int_{\mathbf{p}} (\mu_f - E_p) \Theta(\mu_f - E_p) \\
&\xrightarrow{m_f \rightarrow 0} \frac{N_c}{12\pi^2} \sum_f \mu_f^4.
\end{aligned} \tag{9}$$

Here, we have at the second equal sign discarded two terms — proportional to the 1 in $\Theta(E_p - \mu_f) = 1 - \Theta(\mu_f - E_p)$ — that either correspond to a $\mu = T = 0$ vacuum contribution or vanish in dimensional regularization, and at the last stage, finally proceeded to the massless limit where the spatial momentum integral trivializes. It is worth comparing the simplicity of this computation to the complications one encounters with its finite- T counterpart, considered, *e.g.*, in Section 7, Appendix A of [22].

Proceeding next to the Next-to-Leading Order (NLO) in the massless limit, we encounter one two-loop diagram with a quark loop, first considered in [35]. Carrying out the Lorentz and color contractions, we straightforwardly obtain

$$\begin{aligned}
-\frac{1}{2} \text{ (quark loop diagram) } &= -d_A g^2 \frac{d-1}{2} \int_{\tilde{P}} \int_Q \frac{1}{P^2(P-Q)^2} \\
&= -d_A g^2 \frac{d-1}{2} \left(\int_{\tilde{P}} \frac{1}{P^2} \right)^2,
\end{aligned} \tag{10}$$

where $d_A = N_c^2 - 1$ is the dimension of the adjoint representation of the gauge group $\text{SU}(N_c)$ and we have explicitly included the negative sign and symmetry coefficient of the diagram on the left-hand side of the equation. A clean factorization of the result into a product of two one-loop (fermionic) integrals can clearly be observed from the final form of the result.

The sole one-loop master integral appearing in the above expression is clearly identical to the one in Eq. (7), from where we can immediately take the result for the p_0 integration. A straightforward calculation utilizing this intermediate result produces now

$$\int_{\tilde{P}} \frac{1}{P^2} = -\frac{1}{4\pi^2} \int_0^\infty dp p \Theta(\mu - p) = -\frac{\mu^2}{8\pi^2}, \tag{11}$$

where we have set $d = 3$ and dropped a scalefree integral that vanishes in dimensional regularization. With this result, we obtain the pressure of cold and dense QM up to NLO,

$$p_{\text{QCD}}^{\text{NLO}}(\{\mu_f\}) = \left(\frac{N_c}{12\pi^2} - \frac{d_A g^2}{64\pi^4} \right) \sum_f \mu_f^4, \quad (12)$$

a result first derived close to 50 years ago [35].

3. Soft effective theories for deconfined QCD matter

At the three-loop order, performing the integrations for individual vacuum diagrams becomes technically considerably more demanding, with very few graphs displaying factorization to lower-order integrals (see [36] for recent general results concerning factorization). In addition, explicit computations performed both at nonzero and vanishing temperatures display uncanceled IR divergences in specific vacuum diagrams featuring single gluon propagators raised to powers higher than unity due to the presence of self-energy-type insertions. Unlike the UV divergences that cancel upon renormalization, these divergences are of a physical origin, related to contributions from long-distance interactions mediated by low-momentum gluons. To cancel them, something more is required, and here the discussion naturally separates into two different realms: nonzero and vanishing temperatures.

We begin the story from the perhaps more familiar limit of high temperatures, where the boundary conditions of the fields contributing to Eq. (2) allow us to express their dependence on the imaginary time coordinate $x_0 \equiv \tau$ through discrete Fourier series. In practice, we write

$$\phi(\tau, \mathbf{x}) = T \sum_{n=-\infty}^{\infty} \tilde{\phi}_n(\mathbf{x}) e^{i\omega_n^{\text{bos}} \tau}, \quad (13)$$

$$\psi(\tau, \mathbf{x}) = T \sum_{n=-\infty}^{\infty} \tilde{\psi}_n(\mathbf{x}) e^{i\omega_n^{\text{fer}} \tau}, \quad (14)$$

for bosonic and fermionic fields ϕ and ψ , respectively, with the Matsubara frequencies $\omega_n^{\text{bos}} = 2n\pi T$ and $\omega_n^{\text{fer}} = (2n + 1)\pi T$ acting as thermal mass terms for the three-dimensional field components carrying the index n . At high T , all three-dimensional fields except for the massless $n = 0$ modes of bosons are protected against IR problems by these nonvanishing masses, so that the minimal effective theory capable of describing the IR physics of hot QCD becomes a three-dimensional theory for the $n = 0$ components of the A_0 and A_i fields. What sets the Lorentz components of the original gauge field apart here is the fact that the gauge field of a three-dimensional theory

only has three (spatial) components. This makes the $n = 0$ component of A_0 an adjoint scalar field in the effective theory, capable of acquiring a nonzero mass term in the corresponding effective Lagrangian.

The process of building three-dimensional effective field theories for the description of high-temperature QFTs by integrating out the nonzero Matsubara modes of four-dimensional fields is known as *dimensional reduction*. It possesses a long history dating back to the original work of Appelquist and Pisarski in 1981 [37] and subsequent refinements in the 1990s by Kajantie *et al.* [38] as well as Braaten and Nieto [39]. For QCD, a particularly important step was taken in [40], where the widely-used terms Electrostatic and Magnetostatic QCD (EQCD and MQCD for short) were coined for the three-dimensional theories capable of describing the equilibrium physics of length scales $x \gtrsim 1/(gT)$ and $1/(g^2T)$, respectively. While the former of these theories takes the form of a three-dimensional Yang–Mills theory with an adjoint scalar (see, *e.g.*, [22] for more details),

$$\mathcal{L}_{\text{EQCD}} = \frac{1}{4} F_{ij}^a F_{ij}^a + \text{tr}[D_i, A_0]^2 + m_E^2 \text{tr} A_0^2 + \lambda_1 (\text{tr} A_0^2)^2 + \lambda_2 \text{tr} A_0^4, \quad (15)$$

in MQCD, the A_0 field with an $O(gT)$ mass is integrated out as well. These theories can be used to account for the soft contributions to equilibrium thermodynamic quantities such as the equation of state, but at lower temperatures, their use is limited not only by the lack of a scale hierarchy between the hard (πT) and soft (gT) scales, but also by the explicit breaking of the $Z(N_c)$ center symmetry of the Yang–Mills part of QCD. This can be further remedied by supplementing the effective theory with a more versatile field content, thus extending the applicability of dimensional reduction to somewhat lower temperatures [41, 42].

The effects of quark chemical potentials can be introduced to the dimensionally reduced EFTs in a fairly straightforward manner [43, 44], with the main effects seen in small shifts of the EFT parameters from their $\mu = 0$ values and in the introduction of one new EQCD operator of the form $\text{tr} A_0^3$. Upon increasing the parameter μ/T to values greatly exceeding unity, one, however, eventually exits the regime of validity of dimensional reduction. As demonstrated in [45], this takes place when the temperature is reduced below the scale $T \sim m_E$, where m_E stands for the leading-order electric screening mass

$$m_E^2 = g^2 \left[\frac{C_A + T_F}{3} T^2 + \frac{1}{2\pi^2} \sum_f \mu_f^2 \right] \quad (16)$$

and $T_F = N_f/2$. The physical effect that takes place when T is lowered well below m_E is that more and more bosonic Matsubara modes become soft and

need to be resummed, a process depicted in Fig. 1. In other words, the field content of EQCD is no longer sufficient to capture all of the soft physics in need of a resummed treatment.

The minimal effective theory replacing EQCD at low temperatures turns out to be a four-dimensional EFT constructed for all soft gluon modes satisfying $P^2 \lesssim m_E^2$ and thus requiring some form of resummed treatment³. This EFT, the form of which can be most straightforwardly derived by considering the soft-external-momentum limit of gluonic amplitudes, is dubbed Hard Thermal Loops (HTL) and was first introduced by Braaten and Pisarski already in 1989 [46, 47]. Since then, the HTL framework has been successfully applied to the derivation of numerous physical quantities including the bulk thermodynamic properties of high-temperature QGP at vanishing [48] and nonzero chemical potentials [49], the EoS of cold and dense QCD and QED [50–54], dynamical observables such as production and decay rates [28, 55], and even transport coefficients at various orders of perturbation theory [29, 56].




Fig. 1. An illustration of how more and more bosonic Matsubara frequencies fit inside a circle of constant radius m_E as the temperature is lowered first to a small number times $m_E \sim g\mu$ and then all the way to the $T = 0$ limit. Note that for fermions, the result would be very different due to the imaginary offset of the Matsubara frequency by $i\mu$.

4. Hard Thermal Loops: basic properties and simple applications

The development of the HTL effective theory has been documented in multiple review articles and even textbooks, including *e.g.* [23, 32, 57], and the theory itself has been extended to new physical realms and higher perturbative orders even relatively recently [58, 59]. To keep the present discussion

³ Note that fermions continue to be protected against IR problems also in the cold and dense limit through their nonvanishing chemical potentials. This can be understood in physical terms from the fact that all quantum-mechanical states are filled within the Fermi sphere, implying that only hard fermions with momenta of the order of μ can propagate.

at a tractable level, we will limit its treatment here to the parts of the theory necessary for determining the leading soft contributions to the pressure of cold and dense QM. This implies we will be mainly interested in the one-loop HTL gluon self-energy, which characterizes how the propagation of soft gluons is modified by the medium and that can be obtained from a particular (soft-external-momentum) limit of the same quantity in full QCD. While the result can be shown to take a universal form, being equally applicable in the high-temperature regime, we will for simplicity perform the computation in the strict $T = 0$ limit, allowing some technical simplifications to be implemented. After the derivation of the leading-order HTL self-energies, we will apply the result to the determination of the leading nonanalyticity in α_s of the weak-coupling expansion of the pressure of cold and dense QM.

The gluon self-energy, or the gauge field self-energy tensor, is defined by the Schwinger–Dyson equation, which relates it to the difference of the full and bare inverse propagators

$$\Pi_{\mu\nu}^{ab}(P) = (D^{-1})_{\mu\nu}^{ab}(P) - (D_{(0)}^{-1})_{\mu\nu}^{ab}(P). \quad (17)$$

Here, the free propagator takes the usual form

$$(D_{(0)})_{\mu\nu}^{ab}(P) = \frac{\delta_{\mu\mu} - (1 - \xi)P_\mu P_\nu / P^2}{P^2} \delta^{ab} \quad (18)$$

in the covariant gauges, with ξ being the corresponding gauge parameter. Simply put, $\Pi_{\mu\nu}^{ab}(P)$ is thus gauge-theory generalization of the scalar-field self-energy, which dresses the massless scalar propagator as $\frac{1}{P^2} \rightarrow \frac{1}{P^2 + \Pi(P)}$.

Although a symmetric rank-two tensor in four dimensions, in principle, contains 10 independent components, various symmetries of the system can be seen to significantly reduce this number. First, while four-dimensional Lorentz invariance is broken by the existence of a preferred frame in a system in thermal equilibrium — the rest frame of the heat bath of the medium — three-dimensional rotational invariance is enough to strongly limit the possible structures appearing in the result. In practice, it implies that the self-energy tensor can be given as a linear combination of at most four independent elements: the metric $\delta_{\mu\nu}$ and three symmetric rank-two tensors composed with the external momentum P_μ and the four-vector singling out the temporal direction, $n_\mu \equiv \delta_{\mu 0}$. Current conservation, implemented via the Slavnov–Taylor identities, further restricts the possible form of the tensor by requiring it to be transverse with respect to the external four-momentum in vacuum. As discussed in detail in [59, 60], the situation becomes more complicated in a thermal setting, where transversality depends on whether non-Abelian interactions are present, what loop order is being considered,

and even which (covariant) gauge has been chosen. Here, it suffices to note that for QCD at nonzero μ but vanishing temperature, the one-loop self-energy must be transverse irrespective of the gauge one is working in.

The transversality of the LO self-energy implies that the tensor can possess at most two independent components, typically labeled according to their transversality properties with respect to the external three-momentum. Omitting the Kronecker delta function in the adjoint color indices, the result of this decomposition reads

$$\Pi_{\mu\nu}(P) = \Pi_{\mu\nu}^{\text{vac}}(P) + \Pi_{\text{T}}(P)\mathbb{P}_{\mu\nu}^{\text{T}}(P) + \Pi_{\text{E}}(P)\mathbb{P}_{\mu\nu}^{\text{E}}(P), \quad (19)$$

where the UV divergent vacuum part $\Pi_{\mu\nu}^{\text{vac}}(P)$ represents the $T \rightarrow 0, \mu \rightarrow 0$ limit of the polarization tensor and can be seen to be proportional to $P^2\delta_{\mu\nu} - P_\mu P_\nu$, being thus subdominant for soft external momenta. The matter part of the self-energy is, on the other hand, divided into a transverse (T) and a longitudinal or Euclidean (E) part, with the corresponding projection operators reading

$$\mathbb{P}_{\mu\nu}^{\text{T}}(P) \equiv \delta_{\mu i} \delta_{\nu j} \left(\delta_{ij} - \frac{p_i p_j}{p^2} \right), \quad (20)$$

$$\mathbb{P}_{\mu\nu}^{\text{E}}(P) \equiv \delta_{\mu\nu} - \frac{P_\mu P_\nu}{P^2} - \mathbb{P}_{\mu\nu}^{\text{T}}(P). \quad (21)$$

It is straightforward to verify that apart from the obvious transversality of $\mathbb{P}_{\mu\nu}^{\text{T}}(P)$ with respect to the external three-momentum, these tensors satisfy the following properties (see, *e.g.*, [31]):

$$\mathbb{P}_{\alpha\beta}^{\text{X}}(P)\mathbb{P}_{\beta\gamma}^{\text{X}}(P) = \mathbb{P}_{\alpha\gamma}^{\text{X}}(P), \quad \text{X} = \text{T, E}, \quad (22)$$

$$\mathbb{P}_{\alpha\beta}^{\text{X}}(P)\mathbb{P}_{\beta\gamma}^{\text{Y}}(P) = 0, \quad \text{X} \neq \text{Y}, \quad (23)$$

$$\mathbb{P}_{\mu\mu}^{\text{T}}(P) = d - 1, \quad \mathbb{P}_{\mu\mu}^{\text{E}}(P) = 1. \quad (24)$$

Moving on to the actual evaluation of the gluon self-energy, we first note that the one-loop gluon polarization tensor can be obtained from the sum of all amputated one-particle-irreducible Feynman diagrams with two external gluon lines. Its general expression at nonzero T and μ is rather lengthy, and we refer the interested reader to Eq. (C1) of [45], where the Feynman gauge version of this quantity is displayed. In the strict $T = 0$ limit, some simplifications occur, though, and *e.g.*, the purely bosonic term only contributes to the vacuum limit of the tensor. If we are moreover interested in the soft (or HTL) limit of the result, where $|P| \ll \mu$ and the components of the loop momentum Q can subsequently be assumed to be dominant over those of P , we observe part of the fermionic contribution dropping out as well. This leaves us with the much more manageable $T = 0$ expression

$$\Pi_{\mu\nu}(P)|_{T=0} = -g^2 \sum_f \left\{ \int_{\tilde{Q}} \frac{2\delta_{\mu\nu}}{Q^2} - \int_{\tilde{Q}} \frac{(2Q-P)_\mu(2Q-P)_\nu}{Q^2(Q-P)^2} \right\} \quad (25)$$

to work with, where the fermionic integration measures remind us of the presence of μ_f in the corresponding fermion propagators.

To derive results for the two functions $\Pi_T(P)$ and $\Pi_E(P)$, it suffices to consider two independent components or contractions of the self-energy tensor, for which we choose $\Pi_{\mu\mu}$ and Π_{00} . Omitting terms that are clearly subdominant for soft external momentum P , we obtain

$$\Pi_{\mu\mu}(P)|_{T=0} \approx -2(D-2)g^2 \sum_f \int_{\tilde{Q}} \frac{1}{Q^2}, \quad (26)$$

$$\Pi_{00}(P)|_{T=0} \approx -g^2 \sum_f \left\{ \int_{\tilde{Q}} \frac{2}{Q^2} - \int_{\tilde{Q}} \frac{(2q_0 - p_0)^2}{Q^2(Q-P)^2} \right\}, \quad (27)$$

of which we have already evaluated the tadpole integral in Eq. (11). To this end, let us here concentrate on the latter term in the above expression for Π_{00} , *i.e.*, the only integral with nontrivial dependence on P .

A key observation simplifying the evaluation of the q_0 integral in

$$\tilde{\Pi}(P) \equiv \int_{\tilde{Q}} \frac{(2q_0 - p_0)^2}{Q^2(Q-P)^2} \quad (28)$$

is that two of the four poles, including $q_0 = -i\mu - iq$, always reside on the lower halfplane, while two, including $q_0 = -i\mu + iq$, can reside on either halfplane depending on the magnitude of the three-momentum q in comparison with the chemical potential μ . Writing the resulting step function in the form of $\Theta(q - \mu) = 1 - \Theta(\mu - q)$ and recognizing the appearance of a μ -independent integral that can be dropped, we straightforwardly obtain for the vacuum-subtracted “matter” part of the function

$$\begin{aligned} \tilde{\Pi}(P)_{\text{mat}} &= -\frac{1}{2} \int_q \frac{\Theta(\mu - q)}{q} \left\{ \frac{(2iq - p_0)^2}{(iq - p_0)^2 + (\mathbf{q} - \mathbf{p})^2} + \frac{(2iq + p_0)^2}{(iq + p_0)^2 + (\mathbf{q} - \mathbf{p})^2} \right\} \\ &= -\frac{1}{8\pi^2} \int_0^\mu dq \, q \int_0^\pi d\theta \sin \theta \left\{ \frac{(2iq - p_0)^2}{P^2 - 2iqp_0 - 2\mathbf{q} \cdot \mathbf{p}} + \text{c.c.} \right\}, \end{aligned} \quad (29)$$

where we have on the latter row set $d = 3$ and abbreviated the complex conjugate of the first expression (for real P_μ) by c.c.

The two integrations remaining in Eq. (29) can now be straightforwardly carried out. Upon an expansion of the result to leading order in P/q , the final expression becomes

$$\tilde{\Pi}(P)_{\text{mat}}^{\text{soft}} = \frac{\mu^2}{4\pi^2} \left\{ 1 - \frac{ip_0}{p} \ln \frac{ip_0 + p}{ip_0 - p} \right\}, \quad (30)$$

implying that we have obtained a remarkably simple analytic form for the soft-external-momentum limit of the original integral. In particular, all dependence on the hard scale μ has neatly factorized from the rest, leaving behind a nontrivial function of the ratio of the two independent components of the external momentum, p_0 and $p = |\mathbf{p}|$.

Combining the above result with that of Eq. (11), we are now ready to write down final expressions for the HTL self-energy components $\Pi_{\text{T}}(P)$ and $\Pi_{\text{E}}(P)$. Setting $d = 3$ everywhere, a simple calculation produces

$$\Pi_{\text{E}}(P) = \frac{P^2}{p^2} \Pi_{00}^{\text{mat}}(P), \quad \Pi_{\text{T}}(P) = \frac{1}{2} \Pi_{\mu\mu}^{\text{mat}}(P) - \frac{P^2}{2p^2} \Pi_{00}^{\text{mat}}(P), \quad (31)$$

using which [and taking into account the sum over flavors in Eqs. (26) and (27)] we quickly reach the final gauge-invariant result

$$\Pi_{\text{T}}^{\text{HTL}}(P) = -\frac{m_{\text{E}}^2}{2} \frac{P^2}{p^2} \left\{ \frac{p_0^2}{P^2} - \frac{ip_0}{2p} \ln \frac{ip_0 + p}{ip_0 - p} \right\}, \quad (32)$$

$$\Pi_{\text{E}}^{\text{HTL}}(P) = m_{\text{E}}^2 \frac{P^2}{p^2} \left[1 - \frac{ip_0}{2p} \ln \frac{ip_0 + p}{ip_0 - p} \right], \quad (33)$$

with m_{E} denoting the leading-order electric screening mass from Eq. (16). With these self-energies, the full HTL-resummed gluon propagator finally becomes

$$D_{\mu\nu}^{ab}(P) = \delta^{ab} \left\{ \frac{\mathbb{P}_{\mu\nu}^{\text{T}}(P)}{P^2 + \Pi_{\text{T}}(P)} + \frac{\mathbb{P}_{\mu\nu}^{\text{E}}(P)}{P^2 + \Pi_{\text{E}}(P)} + \frac{\xi P_{\mu} P_{\nu}}{(P^2)^2} \right\}, \quad (34)$$

where we have again reinstated the gauge-parameter term from the free propagator.

Interestingly, the final result for the HTL self-energies and the HTL-resummed gluon propagator take precisely the same forms as in the more general case of nonzero temperatures and chemical potentials; see, *e.g.*, Section 8 of [22]. It is also worth noting that both self-energy components depend on the external momentum only through the ratio $\frac{p_0}{p}$, so that defining an angle ϕ via $\tan \phi \equiv p/p_0$, we can parameterize the entire LO HTL self-energy as a function of this single dimensionless variable. This result suffices

for typical low-order computations utilizing the HTL effective theory, while at higher orders, one may, depending on the quantity being determined, also require the use of the HTL vertex functions, the HTL fermion self-energy, or subleading corrections to the above result from higher loop orders or a power expansion in P/μ . For detailed derivations of these additional features of the HTL theory, we recommend a recent review by Haque and Mustafa [57] as well as the original research article [59], where the real-time formalism of perturbation theory is applied to the gluon polarization tensor at nonzero T and μ .

As a straightforward application of the LO HTL self-energy derived above, let us now briefly return to the weak-coupling expansion of the pressure of dense zero-temperature QM. Here, we first note that the only three-loop vacuum diagram displaying an IR divergence in the $T = 0$ limit can be written in the form

$$\begin{aligned} \frac{1}{4} \text{ (diagram) } &= \frac{d_A}{4} \int_P \left\{ (d-1) \frac{\Pi_T^2(P)}{(P^2)^2} + \frac{\Pi_E^2(P)}{(P^2)^2} + \dots \right\} \\ &\rightarrow \frac{d_A}{4} \int_P \left\{ (d-1) \frac{(\Pi_T^{\text{HTL}}(P))^2}{(P^2)^2} + \frac{(\Pi_E^{\text{HTL}}(P))^2}{(P^2)^2} \right\}, \quad (35) \end{aligned}$$

where we first utilized the expansion (19) for the gluon self-energy, left out IR-convergent terms generated by the vacuum part, and finally retained only the LO IR behavior of the integrand by replacing the full self-energies by their HTL limits.

To cure the IR divergence in the final form of Eq. (35), the simplest way to proceed is to explicitly resum all ring diagrams of the same type, each taking the form

$$I_{\text{ring}}^N \equiv \frac{(-1)^N d_A}{2N} \int_P \left\{ (d-1) \frac{(\Pi_T^{\text{HTL}}(P))^N}{(P^2)^N} + \frac{(\Pi_E^{\text{HTL}}(P))^N}{(P^2)^N} \right\}, \quad (36)$$

with N denoting the number of one-loop self-energy insertions in the graph. Recognizing here the coefficients of the Taylor expansion of the function $\ln(1+x)$, we may explicitly perform the sum $\sum_{N=1}^{\infty} I_{\text{ring}}^N$, obtaining a quantity identifiable as the LO pressure of the HTL effective theory

$$\begin{aligned} p_{\text{HTL}}^{\text{LO}} &= -\frac{d_A}{2} \int_P \left\{ (d-1) \ln [K^2 + \Pi_T^{\text{HTL}}(P)] + \ln [K^2 + \Pi_E^{\text{HTL}}(P)] \right\} \\ &= -\frac{d_A m_E^D}{2} \int_{\hat{P}} \left\{ (d-1) \ln [\hat{P}^2 + \hat{\Pi}_T(\phi)] + \ln [\hat{P}^2 + \hat{\Pi}_E(\phi)] \right\}. \quad (37) \end{aligned}$$

Here, we have on the lower line defined the scaled quantities $P \equiv m_E \hat{P}$ and $\Pi_X^{\text{HTL}}(P) \equiv m_E^2 \hat{\Pi}_X(\phi)$, with X standing for either T or E.

To proceed from here, we first make use of the standard integral

$$\int \frac{d^D K}{(2\pi)^D} \ln(K^2 + m^2) = -\frac{m^D \Gamma(-D/2)}{(4\pi)^{\frac{D}{2}}} \quad (38)$$

that can be used to take care of the radial integration in Eq. (37). This quickly produces (see [22] for details)

$$p_{\text{HTL}}^{\text{LO}} = \frac{d_A m_E^D \Gamma(D/2) \Gamma(-D/2)}{(4\pi)^{\frac{D+1}{2}} \Gamma((D-1)/2)} \times \int_0^\pi d\phi \sin^{2-2\epsilon} \phi \left[(d-1) \hat{\Pi}_T^{D/2}(\phi) + \hat{\Pi}_E^{D/2}(\phi) \right], \quad (39)$$

where the prefactor is of the order of $m_E^{4-2\epsilon}/\epsilon$. It is precisely this product of a D -dependent power of m_E and $1/\epsilon$ that upon a power expansion in ϵ gives rise to a term of type $\alpha_s^2 \ln \alpha_s$ in the weak-coupling expansion of the QM pressure.

Setting now $\epsilon = 0$ in the angular integral of Eq. (39), a straightforward calculation produces

$$\int_0^\pi d\phi \sin^2 \phi \left[(d-1) \hat{\Pi}_T^2(\phi) + \hat{\Pi}_E^2(\phi) \right] = \frac{\pi}{4}. \quad (40)$$

With this result, we may finally read off the coefficient of the leading non-analyticity in the weak-coupling expansion of the pressure,

$$p_{\text{HTL}}^{\text{LO}} = \frac{d_A m_E^4 \Lambda^{-2\epsilon}}{8(4\pi)^2} \left(\frac{\Lambda^2}{m_E^2} \right)^\epsilon \frac{1}{\epsilon} + O(\epsilon^0) = -\frac{d_A m_E^4}{8(4\pi)^2} \ln \alpha_s + \dots, \quad (41)$$

a contribution first derived by Freedman and McLerran in 1977 using considerably more involved computational techniques [35]. Note that the corresponding $1/\epsilon$ divergence that we have discarded here is of UV type (in the HTL effective theory) and cancels against a corresponding IR divergence in the three-loop full theory vacuum diagram considered in Eq. (35).

5. QM thermodynamics at three and four loops

Next, we move on to the full three-loop order and beyond, *i.e.* towards the state-of-the-art in perturbative calculations for QCD thermodynamics.

Up to and including $O(g^5)$ contributions, the pressure of deconfined QCD matter at nonzero temperature and density is most conveniently expressed in a form first introduced in [51],

$$\begin{aligned}
 p_{\text{QCD}} &= p_{\text{QCD}} - p_{\text{soft}} + p_{\text{soft}} \\
 &= p_{\text{QCD}}^{\text{naive}} - p_{\text{soft}}^{\text{naive}} + p_{\text{soft}}^{\text{res}} \\
 &= p_{\text{QCD}}^{\text{naive}} + p_{\text{DR}}^{\text{res}} - p_{\text{DR}}^{\text{naive}} + p_{\text{HTL}}^{\text{res}} - p_{\text{HTL}}^{\text{naive}}.
 \end{aligned} \tag{42}$$

Here, we have first added and subtracted the pressure of a properly defined soft effective theory from the full QCD pressure, then noticed that the difference $p_{\text{QCD}} - p_{\text{soft}}$ is an IR-safe quantity that can be evaluated in a naive loop expansion in the respective theories, and finally introduced a dimensionally reduced (DR) effective theory as the minimal EFT for the $n = 0$ Matsubara modes of gluons and HTL for all the nonstatic modes.

The five terms appearing in Eq. (42) are defined as follows:

- $p_{\text{QCD}}^{\text{naive}}$ denotes the naive loop expansion of the pressure of full QCD, evaluated up to and including the three-loop order. Dimensional regularization is used to regulate both UV and IR divergences, of which the former cancel upon renormalization. IR divergences of the form $1/\epsilon$, however, remain, in addition to which the $O(\alpha_s^2)$ part of the result contains terms that diverge as $\mu^4 \ln \frac{T}{\mu}$ in the small- T limit.
- $p_{\text{DR}}^{\text{res}}$ denotes the pressure of the dimensionally reduced effective theory EQCD, evaluated in a weak-coupling expansion within this theory featuring a massive A_0 propagator. The resulting contribution to the pressure is IR safe by construction but contains a UV divergence.
- $p_{\text{DR}}^{\text{naive}}$ denotes a version of the EQCD pressure evaluated by treating the A_0 mass as an interaction. This leads to an expression that identically vanishes in dimensional regularization due to the integrals becoming scalefree, but the calculation can be seen to contain a cancellation between equal but opposite UV and IR $1/\epsilon$ poles that convert the UV divergence of $p_{\text{DR}}^{\text{res}}$ into an IR one in the difference $p_{\text{DR}}^{\text{res}} - p_{\text{DR}}^{\text{naive}}$. This remaining IR divergence cancels against a similar term in $p_{\text{QCD}}^{\text{naive}}$.
- $p_{\text{HTL}}^{\text{res}}$ and $p_{\text{HTL}}^{\text{naive}}$ denote the logarithmic HTL ring sum [cf. Eq. (37)] and its expansion in self-energies up to and including quadratic order, but with the contribution of the Matsubara zero mode left out. Both expressions are IR finite at nonzero T , but contain UV divergences that cancel in the difference of the two terms. In the $T \rightarrow 0$ limit, the difference $p_{\text{HTL}}^{\text{res}} - p_{\text{HTL}}^{\text{naive}}$ develops an IR divergence proportional to $\ln T$, which is seen to cancel against the corresponding term in the $T \rightarrow 0$ limit of $p_{\text{QCD}}^{\text{naive}}$.

All in all, combining the five terms in the final form of Eq. (42), we recover an expression that is free of both UV and IR divergences and agrees with all previously known limits of the pressure up to and including order $\alpha_s^{5/2}$. This was the main result of [51] and continues to represent the state-of-the-art result for the pressure up to and including all fully known perturbative orders (*i.e.*, not counting logarithms). The only exception to this is the large- N_f limit of QCD, where an all-orders result for the high-temperature pressure has been determined thanks to simplifications that occur in precisely this limit (see, *e.g.*, [61–64]).

Proceeding beyond the three-loop order, the derivation of new terms in the weak-coupling expansion of the pressure becomes considerably more involved and the methods used in the low- and high-temperature regimes become increasingly disjoint. Here, we restrict our discussion to the strict $T = 0$ limit, which was at the center of our attention already in the previous sections of these notes and which is importantly free of the famous Linde problem [65], plaguing high-order perturbative calculations at finite temperature. Given the complexity of the problem and the fact that parts of it remain under active work, we will keep our discussion at a mostly qualitative level, concentrating on the structure of the weak-coupling expansion, the physical origins of various contributions entering at the α_s^3 order, and the computational tools required in these calculations. For a reader interested in further computational details, we will provide numerous references to the original research articles below.

Up to and including the four lowest orders, the weak-coupling expansion of the pressure of cold ($T = 0$) and dense unpaired QM matter obtains the schematic form [50, 66]

$$\begin{aligned}
 p = p_{\text{FD}} + \alpha_s p_1^{\text{h}} + \alpha_s^2 p_2^{\text{h}} + \alpha_s^3 p_3^{\text{h}} \\
 + \alpha_s^2 p_2^{\text{s}} + \alpha_s^3 p_3^{\text{s}} \\
 + \alpha_s^3 p_3^{\text{m}},
 \end{aligned}
 \tag{43}$$

where the coefficients p_2^{s} and p_3^{m} contain linear logarithms of the coupling α_s and the coefficient p_3^{s} both linear and quadratic $(\ln^2 \alpha_s) \log$ s [67]⁴. The letters “h”, “s”, and “m” refer here to the “hard”, “soft”, and “mixed” sectors, corresponding to different regions of momentum space that contribute to the pressure at these orders. A detailed account of the different contributions can be found in [66], the results of which we summarize here, paying particular attention to the state-of-the-art α_s^3 order.

⁴ Note that many terms here also depend on logs of the renormalization scale $\bar{\Lambda}$ that cancel the scale dependence from the running of α_s order by order. We will mostly skip this topic here, but a reader interested in the precise mechanism, in which these cancellations occur, is referred to the Supplemental Material of [52].

Hard sector: Similarly to the case of the finite-temperature pressure in Eq. (42), the contribution of the hard momentum scale μ_B to the pressure of cold and dense QM originates through a naive loop expansion of the full theory pressure. This necessitates the evaluation of all four-loop vacuum, or bubble, diagrams of QCD at nonzero chemical potentials but vanishing temperature, using dimensional regularization to regulate all divergences. The UV divergences of the result are expected to be removed upon renormalization, while — similarly to the finite- T calculation described above — IR divergences will be seen to cancel against corresponding UV divergences from the EFT calculations corresponding to the soft and mixed contributions to the pressure.

At one-, two-, and three-loop orders, the evaluation of vacuum diagrams can be performed analytically for vanishing quark masses, for which it was, in fact, completed for arbitrary values of T and μ in [44]. When a nonzero strange-quark mass is implemented, numerical methods are, on the other hand, required even in the strict $T = 0$ limit, where the three-loop pressure was first determined in [68]. This computation relied on the use of the so-called cutting rules, where the underlying idea is to first perform all fermionic (but not bosonic) temporal momentum integrals, reducing the evaluation of the original Feynman integral to a sum of three-dimensional “phase-space” integrations over vacuum ($T = \mu = 0$) on-shell amplitudes. Below, we will briefly specify the cutting rules and apply them to a simple example calculation at the two-loop level, referring a reader interested in details of their formal derivation to the original article [69].

The cutting rules act on a scalarized $T = 0$ Feynman integral corresponding to either a vacuum diagram or a Euclidean N -point function with real-valued external momenta. The integral is assumed to be composed of two types of (massive or massless) propagators, bosonic and fermionic, of which the former are composed of real-valued momenta, while the temporal components of the latter are shifted by $i\mu$ as discussed in Section 2. We furthermore assume that no fermionic propagator is raised to a power higher than unity, although this restriction can be relaxed if proper care is taken in the calculation [70, 71].

With these assumptions, the cutting rules amount to the following set of five operations:

- (i) Graphically perform all possible cuts of *independent* internal fermion lines in the diagram, ranging from zero to the number of loops in the graph. Here, the independence of a given set of internal lines is equivalent to being able to choose the corresponding momenta as the loop momenta in the Feynman integral in question.

- (ii) For each cut fermion line with momentum P , remove the corresponding propagator from the integral and place the same momentum on shell in the thus generated amplitude by writing $p_0 = iE_p$, $E_p \equiv p^2 + m^2$.
- (iii) Set $\mu = 0$ in each of the uncut fermion propagators and evaluate the corresponding integrals while assuming all components of the external momenta to be real-valued.
- (iv) Integrate over the three-momenta of the cut lines with integration measure $-\Theta(\mu - E_k)/(2E_k)$ (times the usual $1/(2\pi)^3$ associated with momentum integrals), with the integrand being the vacuum ($\mu = 0$) on-shell amplitude generated by the cutting.
- (v) Sum all the individual contributions together.

This procedure reduces the evaluation of the original finite- μ diagram to a number of separate terms, in which the μ -dependence resides only in the theta functions of the phase-space integration measures. This is a significant simplification, as the values of the vacuum amplitudes can typically be taken from collider-physics literature and the phase-space integrals are by construction UV finite.

As a simple example case, we inspect next the $T = 0$, $\mu \neq 0$ integral

$$I_2(\mu) = \int_{\tilde{P}} \int_{\tilde{Q}} \frac{1}{P^2 + m^2} \frac{1}{Q^2 + m^2} \frac{1}{(P - Q)^2}, \quad (44)$$

where the momenta P and Q correspond to massive fermions at nonzero chemical potential and $P - Q$ to a massless boson. Identifying two as the number of independent fermionic momenta, the cutting rules amount to writing the integral in the form (see Fig. 2)

$$I_2(\mu) = I_2^{0\text{-cut}} + I_2^{1\text{-cut}}(\mu) + I_2^{2\text{-cut}}(\mu), \quad (45)$$

where the 0-cut part refers to the (uninteresting) $\mu = 0$ version of the original integral, while the 1- and 2-cut parts carry all the μ -dependence of the original graph

$$\begin{aligned} I_2^{1\text{-cut}}(\mu) &= -2 \int_p \frac{\Theta(\mu - E_p)}{2E_p} \left[\int_Q \frac{1}{Q^2 + m^2} \frac{1}{(P - Q)^2} \right]_{p_0 \rightarrow iE_p}, \\ I_2^{2\text{-cut}}(\mu) &= \int_p \frac{\Theta(\mu - E_p)}{2E_p} \int_q \frac{\Theta(\mu - E_q)}{2E_q} \left[\frac{1}{(P - Q)^2} \right]_{p_0 \rightarrow iE_p, q_0 \rightarrow iE_q}. \end{aligned} \quad (46)$$

$$\begin{aligned}
& \text{Solid circle with dotted line} \rightarrow \text{Dashed circle with dotted line} - 2 \int_p \frac{\theta(\mu - E_p)}{2E_p} \text{Dashed arc with dotted line} \Big|_{p_0 \rightarrow iE_p} \\
& + \int_p \frac{\theta(\mu - E_p)}{2E_p} \int_q \frac{\theta(\mu - E_q)}{2E_q} \text{Dashed lines meeting at dotted line} \Big|_{p_0 \rightarrow iE_p, q_0 \rightarrow iE_q}
\end{aligned}$$

Fig. 2. A graphical illustration of the cutting of the two-loop graph $I_2(\mu)$ discussed in the main text, with solid lines corresponding to a massive scalarized fermion propagator with nonzero μ , dashed lines to its $\mu = 0$ version, and dotted lines to a massless bosonic propagator. The figure is taken from [23].

The resulting integrals can be performed numerically at ease, or even analytically for $m = 0$ (see Appendix D of [71] for details), to reach the final result for the two-loop diagram. Note that the factor of two in the one-cut expression above results from the two separate one-cut terms producing identical results due to a symmetry of the original integral. It should be noted that such symmetries are not always present in higher-loop cases, and care must be taken not to automatically equate two terms that differ, *e.g.*, by the relative direction of fermion flow in two separate fermion loops.

The cutting rules were first used in [68] to evaluate all massive three-loop vacuum diagrams of QCD in a calculation that would have been considerably more challenging without this crucial computational aid. Proceeding all the way to four loops, the complexity of the calculations becomes considerably more challenging, however, and only two out of the altogether 52 individual graphs depicted in Fig. 3 have been fully evaluated so far [53, 72]. From the rest, the first 12 of Fig. 3 are found to vanish because they either contain a factorized scalefree integral or their color trace gives zero. Active work is currently underway to evaluate the remaining diagrams, with promising first steps taken recently in a yet unpublished article [73]. The steps involved in this process include performing the color traces and Lorentz algebra of all graphs (in a general covariant gauge), then dividing the result into group-theory-invariant sectors, and finally systematically implementing momentum shifts and other standard manipulations to achieve the cancellation of the covariant gauge parameter ξ within each group theory sector.




Fig. 3. A list of all four-loop vacuum diagrams of QCD containing at least one fermion loop, shown together with the associated signs and symmetry coefficients. The first 12 diagrams can be shown to exactly vanish due to either containing a factorized scalefree integral (first 11 graphs) or being proportional to a vanishing color algebra (graph 12). The figure has been composed by Pablo Navarrete.

After the above process, we are left with a relatively high number of independent master integrals (of the order of 100) to evaluate, which one may try to reduce using various novel methods, including but not limited to four-dimensional IBP identities recently generalized to the $T = 0$, $\mu \neq 0$ realm in [71]. Finally, once the seemingly irreducible masters have been identified, the remaining task is to evaluate them starting from the ones that factorize into lower-loop-order entities.

The evaluation of the genuine four-loop master integrals is a very challenging task that involves a mix of analytic and numerical methods, with only a handful of cases completed so far [73]. An important complementary tool, which holds for the simultaneous evaluation of all IR convergent four-loop integrals, is based on the loop-tree duality (LTD) method of vacuum perturbative QFT [74]. This method was recently generalized to nonzero chemical potentials and successfully applied to the evaluation of the infamous “bugblatter” diagram [72]⁵, which represents the LO difference between the pressures of cold and dense QCD and its phase-quenched version (see, *e.g.*, [76]). The main difference between the LTD and cutting rules methods is that in the former, one first evaluates all temporal momentum integrals instead of just the fermionic ones, which avoids the generation of artificial IR divergences and, at least in some cases, leads to more manageable numerical calculations.

⁵ For the etymology of the name of this diagram, see [75].

Soft and mixed sectors: The only soft scale present in cold and dense QM originates from the long-distance screening of gluon fields and can be identified with the $T \rightarrow 0$ limit of the chromoelectric screening mass $m_E \sim g\mu_B$ in Eq. (16), *i.e.*, the same parameter that appears in the LO HTL self-energy derived in the previous section. At the $O(\alpha_s^3)$ order for the pressure, it no longer suffices to only work with this quantity, but we need to additionally consider other elements of the effective theory, including the HTL vertex function, and derive corrections to the LO self-energy both in the form of two-loop contributions and power corrections in the soft external momentum. The necessary calculations were first described in [66] and subsequently carried out in [50, 54], to which we refer the interested reader for technical details. Here, we instead restrict our attention to a qualitative account of the origins of the soft and mixed $O(\alpha_s^3)$ contributions to the QM pressure, see Eq. (43).

The physical nature of the mixed and soft contributions to the $O(\alpha_s^3)$ pressure can be neatly illustrated through a practical example depicted in Fig. 4, borrowed from [50]. Here, we begin with a single unresummed four-loop vacuum diagram of full QCD that consists of two quark loops and four additional gluon propagators. The fermion propagators are by construction “hard”, being characterized by the scale μ present in their temporal momentum components, while the gluon momenta can be either hard, with $P \sim \mu$, or soft, with $P \sim m_E$. Should one or several of the gluonic momenta in the graph — say the vertical ones flowing into the one-loop self-energy insertion in the first diagram of Fig. 4 — become soft, a simple power-counting exercise shows that similar diagrams containing an arbitrary number of one-loop gluon self-energy insertions contribute at the same order in the weak-coupling expansion [step (i) in Fig. 4]. Summing all such diagrams using the low-momentum limit of the gluon self-energy [step (ii)] leads to the emergence of a three-loop graph containing an HTL resummed gluon propagator, which represents a typical mixed contribution to the $O(\alpha_s^3)$ pressure, denoted by p_3^m in Eq. (43). Repeating, finally, the same exercise for the two remaining unresummed gluon lines, which clearly become both soft at the same time [step (iii)], leads to a fully soft two-loop contribution computable within the HTL effective theory [step (iv)]. The final HTL graph features also two resummed HTL vertex functions, one of which arises from the lower quark loop of the original diagram, and contributes to the p_3^s term of Eq. (43).

The successful evaluation of the soft and mixed contributions to the $O(\alpha_s^3)$ pressure presents a technically very challenging task that was completed only recently in [50, 54]. In these calculations, it was crucial to ensure the proper cancellation of all IR divergences in the hard four-loop diagrams of QCD against UV divergences originating from the soft two-loop diagrams




Fig. 4. An illustration of how a single four-loop vacuum diagram in dense QCD gives rise to a three-loop mixed diagram and a two-loop HTL graph when one or several of the internal gluon lines become soft. The thin wavy lines correspond to unresummed and the thick wavy lines to HTL-resummed gluon propagators, respectively, while the gray blobs stand for HTL-resummed vertex functions.

in the HTL effective theory as well as UV and IR divergences originating from the mixed three-loop graphs. It is also worth repeating that, as noted above, the one-loop HTL self-energies needed to be generalized both by two-loop corrections, evaluated in [59], and by power corrections in the soft external momentum, derived in [54].

Convergence of the weak-coupling expansion: After the recent determination of the mixed $O(\alpha_s^3)$ contributions to the QM pressure in [54], the current state-of-the-art result for the EoS of cold and dense QM contains all terms appearing in Eq. (43) except for the finite, renormalization-scale-independent part of p_3^h , which amounts to one single number. As demonstrated in Fig. 5, taken from [54], the effect of this new order on the convergence of the weak-coupling expansion is dramatic. For typical values of the missing coefficient, currently under intense scrutiny, the new N3LO band improves the convergence of the pressure at all relevant densities. Such a result would importantly enable lowering the density, at which the high-density constraint is applied in model-independent determinations of the NS-matter EoS, from the present approx. $40 n_s$ down to perhaps $20 n_s$. As we will discuss in detail in Section 7, this would have a dramatic effect on the precision, at which we currently understand the properties of the ultradense matter present inside NS cores.




Fig. 5. The pressure of cold and dense massless QM in beta equilibrium [54]. Shown are the NLO (blue band), NNLO (yellow), and NNNLO (green) perturbative orders, of which the last one is new and contains one single undetermined coefficient c_0 set to a value determined by a Bayesian estimation performed in [54] (see also [77]).

6. Transport properties of dense QM

In addition to the bulk thermodynamic properties of NS matter, first-principles particle theory methods may provide valuable input for its transport coefficients, describing physical phenomena such as diffusion, energy and momentum dissipation, and the conduction of heat and electric currents. While in some cases also relevant for the properties of quiescent NSs, they play a particularly important role in more dynamical settings, such as supernovae explosions and the postmerger dynamics of binary NS mergers (see, *e.g.*, [78, 79]). In [80], it is argued that the most important transport coefficient for NS mergers is the bulk viscosity ζ , which has subsequently received considerable attention in recent literature, including two very recent papers [81, 82] concentrating on its determination in QM. Here, we will, however, begin the story with some general remarks followed by a brief look into other transport coefficients in (mostly unpaired) dense QM, returning to the somewhat special evaluation of ζ only at the very end of this section.

Similarly to the more widely studied case of hot QGP, the transport coefficients of cold and dense QM are technically considerably more challenging to determine than its equilibrium thermodynamic properties, which is mainly due to the inherently Minkowskian nature of the former quantities. This is reflected in the existence of only a handful of low-order perturbative results for the most important transport coefficients, many of which moreover date back tens of years [56, 83]. These quantities are typically defined

through the response of the system to an imbalance in some local quantity, and their evaluation often begins with the Boltzmann equation, describing the evolution of an off-equilibrium statistical system. Here, our treatment of the subject will be rather superficial, and the interested reader is referred to the excellent review article [84] for more details on the transport phenomena of dense QCD matter, including its confined phase.

The study of transport phenomena in dense QM often begins with the Boltzmann equation describing the behavior of the quark distribution function $f(\mathbf{r}, \mathbf{p}, t)$

$$\left(\frac{\partial}{\partial t} + \mathbf{v} \cdot \nabla_{\mathbf{r}} + \mathbf{F} \cdot \nabla_{\mathbf{p}} \right) f(\epsilon_{\mathbf{p}}) = \left(\frac{\partial f(\epsilon_{\mathbf{p}})}{\partial t} \right)_{\text{coll}}. \quad (47)$$

Here, \mathbf{F} is an external force, \mathbf{r} and \mathbf{p} denote the position and three-momentum of the quark being tracked, and the gluon-mediated scattering between quarks is described by the highly nonlinear collision term on the right-hand side of the equation (see, *e.g.*, Eq. (1) of [56]). The near-equilibrium solutions of this equation under various boundary conditions, say with a specific type of external force or a fixed flow velocity for one quark flavor, inform us about the response of the system to the perturbation in question and often allow for solving a particular transport coefficient.

A very crude approximation to solving the Boltzmann equation amounts to linearizing the collision term in the difference between the off-equilibrium and equilibrium distribution functions. This amounts to the so-called relaxation time approximation, which was applied in the first works on transport in dense QM or hot QGP, such as [85]. A major step forward was taken somewhat later in [56], where the first self-consistent perturbative study of transport in dense QM was performed with the full nonlinear collision term. In this work, a major challenge was related to the long-range interactions mediated by soft gluons, which — analogously to the $O(\alpha_s^2)$ pressure discussed above — would lead to divergent results if screening effects are not properly taken into account.

While present also at high temperatures, the effects of dynamical screening, such as Landau damping, are particularly pronounced in transport calculations at large chemical potentials and small temperatures, where the system is characterized by three distinct momentum scales, μ_B , T , and m_E . Of the three, the first can typically be assumed to be the largest, but the ordering of the latter two is *a priori* unclear in the physical systems of interest, such as binary NS mergers. This leads to technical complications in the Boltzmann description, where the quark–quark scattering matrix element within the collision term needs to be dressed with HTL-resummed gluon propagators, featuring the self-energies of Eqs. (32) and (33).

From this point onwards, the steps needed to obtain LO results for various transport coefficients depend strongly on the quantity in question, and we refer the interested reader to [56] for computational details. Here, we merely note that two key quantities that aid the determination of a number of transport coefficients from the diffusion constant to various conductivities and the shear viscosity include the momentum transfer between quarks and the so-called momentum stopping time. Since the publication of these LO results in 1993 [56], no perturbative improvements have, however, been reported, which is largely due to the technical complexity of the NLO calculations, reflected also in the first appearance of NLO transport results for high-temperature QGP as late as in 2018 [29, 86]. Promising advances have, however, been recently made in the holographic description of the same transport coefficients as well as neutrino transport in the strongly-coupled regime of QCD-like theories [87–89].

Returning finally to the transport coefficient of highest relevance for NS mergers, the bulk viscosity ζ of dense QM, we encounter a quantity whose evaluation differs dramatically from the computations discussed above. Interestingly, it turns out that in an NS setting, the dominant contribution to ζ does not originate from QCD alone, but emerges when rapid density oscillations in the macroscopic system force matter to depart from chemical equilibrium due to weak interactions not being able to keep up with the compression rate. In rotating quiescent pulsars, these oscillations are related to the different oscillation modes of the star (see, *e.g.*, [90]), while in binary NS mergers, they are often related to the complicated merger dynamics. A straightforward calculation reviewed, *e.g.*, in Appendix A of [81] shows that in the so-called neutrino-transparent regime, where the dominant contribution to the bulk viscosity comes from the nonleptonic W -boson exchange process $u + d \leftrightarrow u + s$, the bulk viscosity obtains the schematic form

$$\zeta = \frac{\lambda_1 A_1^2}{\omega^2 + (\lambda_1 C_1)^2}. \quad (48)$$

Here, ω is the angular velocity of the density oscillations, λ_1 an electroweak rate defined by $\frac{dn_d}{dt} = -\frac{dn_s}{dt} = \lambda_1(\mu_s - \mu_d)$, and the coefficients A_1 and C_1 are linear combinations of quark densities and susceptibilities, reproduced in Eqs. (4) and (5) of [81]. Interestingly, A_1 turns out to identically vanish for mass-degenerate quarks, reflecting the vanishing of the bulk viscosity for conformal systems and implying that it is necessary to include a nonzero strange-quark mass in the evaluation of the coefficients A_1 and C_1 .

Barring loop corrections to the electroweak data λ_1 (that may well be sizable, see, *e.g.*, [95]), a crucial implication of the form of Eq. (48) is that the QCD contribution to the bulk viscosity of dense QM enters solely through the thermodynamic quantities appearing in the coefficients A_1 and C_1 . As

noted above, in their evaluation it is imperative to employ a nonzero strange-quark mass, which may, however, be treated as an interaction following, *e.g.*, the mass-expansion scheme recently introduced in [34]. Precisely, this was recently done in [81], where the thermodynamic functions contributing to the bulk viscosity of unpaired QM were evaluated to an unprecedented accuracy using both pQCD and two holographic models. The result of this computation is displayed in Fig. 6, where the bulk viscosity is plotted as a function of T for various fixed baryon densities. We observe a qualitative agreement between the weak- and strong-coupling calculations within QM — even across different densities — but a sharp contrast to earlier nuclear-matter results from [94], for which the bulk viscosity peaks at a considerably higher temperature. It is tempting to speculate, whether this difference in the peak temperatures might lead to observable effects in the postmerger GW spectrum of a binary NS merger, should the effects of bulk viscous dissipation on the merger dynamics be sizable enough.

One general remark about transport coefficients in dense QM is finally in order. As the alert reader has surely noted, all of the discussion above concerned the transport properties of unpaired QM, which is a natural starting point, but very likely not the physical ground state of QCD at the low temperatures and moderate and high densities realized in NSs and their binary mergers. Instead, it is widely believed that some type of quark pairing takes place in physical QM, but due to the strongly-coupled nature of the




Fig. 6. The T dependence of the bulk viscosity in three-flavor unpaired QM [81], evaluated in pQCD at the baryon density of $40n_s$ and in two holographic models V-QCD [91, 92] and D3-D7 [93] at 5 and $10n_s$. For reference, the quantity is also shown for dense nuclear matter at three different densities [94] and at $n_B = 10n_s$ for a system of noninteracting quarks.

system, the details of the pairing channel remain unclear at nonasymptotic densities. Unlike in the case of equilibrium thermodynamics, where pairing contributions to the EoS are expected to be strongly suppressed at all densities where pQCD is applicable, many transport coefficients are moreover thought to be highly sensitive to pairing, so calculations performed in the unpaired phase will need to be generalized to color-superconducting phases in the future. Given the technical nature of this topic, we refrain from a more extensive discussion here and instead refer the interested reader to two excellent review articles by Schmitt and collaborators [84, 96].

7. Model-independent inference of NS-matter properties

After describing the determination of the thermodynamic and transport properties of unpaired QM at some length above, it is natural to finally ask, how these results can be used in a more phenomenological setting within NS physics. Here, the primary interest lies in a model-independent constraining of the EoS and other properties of strongly interacting matter up to the maximal $O(5\text{--}10\,n_s)$ densities reached inside physical NSs and the maximal $O(50\text{--}100\text{ MeV})$ temperatures reached during NS mergers. *A priori*, it is far from clear that perturbative results of the kind described in these notes contain any information of practical use in such settings, given that weak-coupling expansions in pQCD typically start converging only at densities many times higher. Somewhat counterintuitively, the perturbative high-density constraint, however, turns out very useful especially in EoS inference, and NS observations can, in turn, be shown to constrain the thermodynamic properties of QCD matter up to relatively high densities. Below, we will shed light on these developments, paying particular attention to the use of the pQCD constraint in the model-independent inference of NS-matter properties. For a reader interested in NS physics from a more astrophysical point of view, we can recommend, *e.g.*, the recent review [97].

The key microscopic quantity one typically tracks in the description of quiescent NSs is the EoS of dense QCD matter in the limits of vanishingly small temperature T , local charge neutrality⁶, and beta equilibrium. This is because it is precisely the relationship between the pressure and energy density of this type of matter that closes the famous Tolman–Oppenheimer–Volkov (TOV) equations governing hydrostatic equilibrium inside a nonrotating star [99, 100]

⁶ The requirement of local charge neutrality can, in principle, be relaxed to a global one, allowing for the presence of two phases in coexistence, relevant for the case of a first-order (deconfinement) phase transition. This is referred to as the Gibbs construction, explained in more detail in, *e.g.*, [98].

$$\begin{aligned}\frac{dM(r)}{dr} &= 4\pi r^2 \epsilon(r), \\ \frac{dp(r)}{dr} &= -\frac{(\epsilon(r) + p(r))}{r^2} \frac{(M(r) + 4\pi r^3 p(r))}{(1 - 2M(r)/r)},\end{aligned}\tag{49}$$

allowing one to solve the corresponding unique mass–radius (MR) relation that can be compared to observations. Up to small deviations caused by rotational frequencies, magnetic fields, and other similar quantities varying from star to star, the masses and radii of every NS in existence should, in principle, fall on the same MR-curve, dictated by the low-temperature EoS alone. Conversely, should an accurate simultaneous determination of the masses and radii of several individual NSs become feasible one day, one may reverse engineer the TOV equations to obtain the EoS of cold and dense beta-equilibrated QCD matter all the way from vanishing density to the maximal central densities realized in stable NSs.

The most straightforward way, in which one could, in principle, use the perturbative QM EoS in NS physics, is by extrapolating the result to much lower chemical potentials, where it would be matched to a nuclear matter EoS, similarly extrapolated from below. The simplest implementation of such a setup would typically involve a first-order phase transition at the baryon chemical potential where the two pressures are equal, while a more refined setting would allow for the presence of a mixed phase of both quark and nuclear matter. Precisely this was done in [68, 101], where the three-loop pressure of cold QM was first evaluated with a nonzero strange-quark mass and then matched with phenomenological EoSs for high-density NM (see Fig. 7). This represented a significant step forward from earlier calculations employing simplistic model EoSs for QM, such as that of the MIT bag model [102], but the approach nevertheless suffered from multiple problematic issues. First, the 3–5 n_s densities, where the matching of the NM and QM EoSs was performed, are well outside the realms of controlled first-principles calculations in both phases, so that predictions for quantities such as the transition density or the associated latent heat can be considered indicative at best. Also, there is no robust way to quantitatively assess the systematic uncertainty involved in such a setup, and potentially sizable contributions from important physical phenomena such as the onset of hyperons [103] or quark pairing [96] are altogether ignored.

The most important shortcoming of the above setup clearly originates from the assumption that the low- and high-density EoSs and their respective uncertainty ranges are applicable also in the problematic region of NS core densities. To improve from here, a natural alternative is to divide the density interval from zero to infinity into not two but three parts: a low-density regime where *ab-initio* nuclear-physics methods such as CET produce reli-

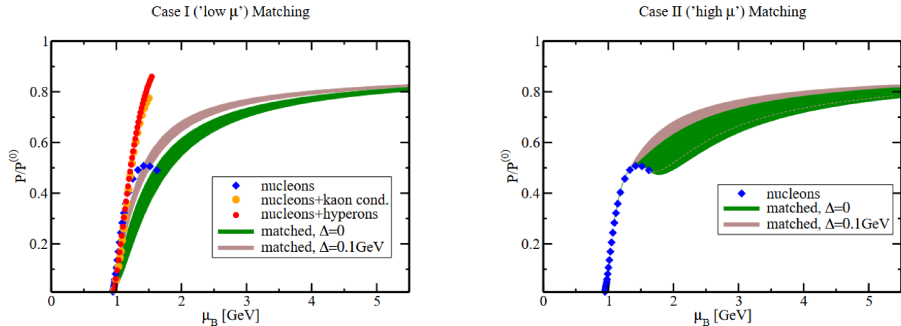


Fig. 7. Results from a naive matching of phenomenological hadronic EoSs (blue, red, and yellow dots and diamonds) to the perturbative QM EoS both with (brown band) and without (green) a phenomenological pairing contribution [68].

able results, a high-density one where the pQCD EoS for QM converges in a satisfactory manner, and an intermediate region where no first-principles results are available. In the last density interval, the EoS can be allowed to behave in any physically consistent way, with only subluminality ($c_s < 1$), thermodynamic consistency as well as a smooth matching to the low- and high-density EoSs required.

The qualitative idea described above and illustrated in Fig. 8 was first implemented by Kurkela *et al.* in 2014 [104], following a similar extrapolation of the CET EoS by Hebeler *et al.* just one year before [105]. The exploratory study of [104] involved the use of two- and three-part polytropes in the intermediate density regime, *i.e.* functions composed in a piecewise manner from ansätze of the form of $p = \kappa n_B^\gamma$. After discarding EoS specimens that do not support the existence of the most massive NSs observed, $M \approx 2M_\odot$ with M_\odot standing for the solar mass, the resulting EoS band was seen to considerably tighten from that obtained in [105] (see Fig. 8). In addition, a set of somewhat *ad-hoc* limits imposed on the polytropic γ indices in the extrapolated EoSs of [105] were now set on a more firm footing and shown to originate from the high-density pQCD constraint.

Without significant advances in astrophysical measurements that can be applied to constrain the interpolated EoS band at intermediate densities, tightening the EoS bounds derived in [104] would have been a very slow process, driven solely by advances in the low- and high-density microphysical calculations. In late 2017, a dramatic new observation was, however, announced by the LIGO and Virgo collaborations, who reported the first-ever detection of a GW signal from the binary NS merger GW170817 [13–15]. The GW signal was soon accompanied by associated gamma-ray and kilonova observations [16, 17], suggesting that the merger remnant very likely underwent a delayed gravitational collapse into a black hole (BH).

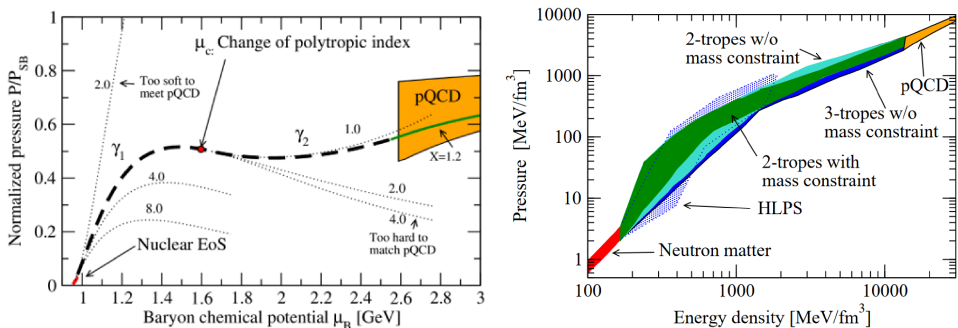


Fig. 8. Left: an illustration of the schematic idea behind the first model-independent interpolation study of the NS-matter EoS [104]. The red line on the left illustrates the low-density EoS from CET [105], the orange band on the right the high-density EoS from pQCD [68, 106], and the dashed black line in between represents a two-part polytrope. Right: the resulting EoS bands from the study presented in [104]. Shown here are bands generated with two- (green) and three-part (solid blue) polytropes as well as results from an earlier extrapolation study [105] (dotted blue) and a two-part polytrope calculation without a constraint requiring that two-solar-mass NSs be supported by all viable EoSs (turquoise).

From the point of view of EoS inference, the most useful constraint from the first LIGO/Virgo papers was a nontrivial upper bound for the tidal deformability of an approx. 1.4-solar-mass NS — a quantity that characterizes the extent to which the shapes of two inspiraling stars get deformed under each others' gravitational fields [14]. With accurate knowledge of the so-called chirp mass of the merger, this bound was quickly translated to a constraint for the NS-matter EoS, whereby a large fraction of otherwise viable candidate EoSs could be discarded for predicting too large radii and thus tidal deformabilities for light pulsars. Early papers implementing this new bound in either model-independent EoS studies or specific model EoSs included, *e.g.*, [109–125], of which the one by Annala *et al.* [109] was the first one to systematically utilize the pQCD high-density constraint.

Another important insight that emerged from the GW170817 merger was related to the likely formation of a black hole through either a so-called supramassive or hypermassive NS (see, *e.g.*, [126] for an extensive discussion of these scenarios)⁷. Depending on the scenario considered, this leads to either more stringent (the more likely hypermassive scenario) or weaker

⁷ In short, in the hypermassive NS scenario, the remnant experiences a rapid (milliseconds to a few seconds) gravitational collapse while still undergoing differential rotation, whereas in the longer-lived supramassive NS scenario, the remnant has time to turn from differential to uniform rotation, then slowly loses angular momentum, and finally collapses to a BH in a timeframe from seconds to hours.

(supramassive case) constraints that were implemented in an interpolation setup featuring the pQCD EoS in [108] (see also [110–112], which derived the first constraints on the maximal TOV mass of a nonrotating NS and [127] for a later refinement). The corresponding interpolation routine started by expressing the speed of sound squared as a piecewise-linear function of the quark chemical potential, illustrated in Fig. 9, and resulted in the EoS and MR bands of Fig. 10. In the latter results, radius constraints published by the NICER Collaboration in 2021 were implemented as well, amounting to an 11.1 km lower bound for the radius of a $2M_\odot$ NS at a 90% credence [10–12].

While the results of Fig. 10 still represent the state-of-the-art of hard-cut-type interpolation analyses, there are questions that studies of this kind are unable to answer or that would at least require a further layer of analysis:

1. What is the quantitative impact of the high-density constraint on the results, *i.e.* would similar results have been obtained also without the pQCD EoS that we have worked so hard to derive in these notes?
2. Is there a way to assign likelihoods for various EoS behaviors and take measurement uncertainties into account? The use of hard cuts in EoS inference is clearly suboptimal, as it forces one to make a binary decision (accept or reject a given EoS) based on minuscule differences, such as whether the most massive NS an EoS is able to support is $1.99M_\odot$ or $2.00M_\odot$.
3. Is it possible to transform the EoS constraints to statements about the phase of matter inside NS cores? In particular, can we somehow estimate the likelihood of NS cores containing deconfined QM?

Below, we will address each of these questions in turn, covering both the technical tools needed in this work and the results from the said analyses.




Fig. 9. A schematic illustration of the speed-of-sound interpolation routine, originally developed for [107]. The figure is taken from [108].




Fig. 10. An example of the EoS and MR results from the analysis of [108], where information from the NS radius constraints from the NICER Collaboration and the likely delayed gravitational collapse of the GW170817 merger remnant were taken into account, assuming here the more conservative supramassive scenario. The color coding refers to the maximal value the speed of sound squared attains at any density, made feasible by the c_s^2 interpolation routine of Fig. 9. The observed behavior clearly suggests that a hypothetical upper limit for c_s^2 (see, *e.g.*, the recent [128]) would carry significant additional constraining power for EoS inference.

Starting from the first point, one may, in principle, ask two slightly different questions: (1) Do the piecewise-defined interpolation functions used in calculations such as [107–109] allow for sufficient complexity for the EoS, so that the results are not biased by the EoS ansätze? and (2) How would the results change if the corresponding analyses were performed without the high-density constraint? To address both issues at one go, Komoltsev, Kurkela, and collaborators decided to inspect the impact of the pQCD constraint with an entirely different set of tools. In [129], they first demonstrated that the high-density constraint implemented at $40n_s$ limits the NS-matter EoS down to below $3n_s$ with no assumptions except for causality and standard thermodynamic relations. Following this, they employed a very conservative nonparametric Gaussian Process (GP) interpolation to pinpoint the effects of the pQCD constraint on the NS-matter EoS [130, 131]. The most important takeaway from these studies, where the pQCD limit was treated in a similar fashion to astrophysical constraints and could be turned on and off at will, was that it significantly softens the EoS at intermediate and high densities, leading to a lower maximum mass for stable NSs and a higher likelihood for the formation of a BH in binary NS mergers. In [131], the authors, in addition, introduced a less conservative method for implementing the pQCD constraint to EoS inference calculations, with the corresponding marginalized QCD likelihood function available at <https://zenodo.org/records/10592568>

As to the second question concerning likelihoods and measurement uncertainties, a very natural step forward is to extend the hard-cut analyses described above towards Bayesian statistical inference, built on the famous Bayes' theorem

$$P(\text{EoS}|\text{data}) = \frac{P(\text{data}|\text{EoS})P(\text{EoS})}{P(\text{data})}. \quad (50)$$

In short, the idea is to turn the (difficult) question of determining the relative likelihoods of various EoSs based on given observational data around and instead solve the (substantially easier) problem of determining the likelihood of the observational data based on a given EoS. Such a setup has by now been implemented to the NS-matter EoS inference by multiple groups, with early adaptations including, *e.g.*, [132–136]. A key finding from these studies is the very likely presence of a bump, *i.e.* a maximum, in the speed of sound of NS matter at densities realized within physical NSs.

Finally, successfully answering the third and last of the above questions requires knowledge of the expected properties of both nuclear and quark matter in the strongly-coupled regime, where the deconfinement transition is expected to occur but no controlled first-principles tools are available. The unavailability of *ab-initio* results makes the question somewhat ill-posed, but at least some guidance can be sought from a more tractable regime of high-energy-density QCD matter, high-temperature QGP. For this system, lattice simulations carried out by multiple collaborations over the past two decades (see, *e.g.*, [137, 138]) have convincingly demonstrated that the transition from a hot hadron gas to QGP takes place as a crossover transition at a temperature of approx. 155 MeV and an energy density of slightly below 400 MeV/fm³, although the precise numbers slightly depend on the quantities being tracked. The most striking difference between the confined and deconfined phases is related to the approximative conformal symmetry of the latter: while the properties of hadronic matter are characterized by the $O(\text{GeV})$ hadron masses, in QM, the only scaleful parameters are the $O(100 \text{ MeV})$ mass of the strange quark and the dynamically generated $\Lambda_{\text{QCD}} \sim 300 \text{ MeV}$ scale parameter. This is reflected in the rapid conformalization of lattice results for multiple quantities, when the temperature of the system is increased past the transition region, *cf.*, *e.g.*, Figs. 7 and 11 of [138].

For the cold and dense QCD matter inside NSs, the lack of first-principles predictions to compare against means that phase identification must at least partially rely on a comparison of the inferred properties of NS matter with the known properties of QM and NM at considerably higher and lower densities, respectively. Fortunately, there exist a large number of physical quantities with clear predictions in various limits to which we can compare the

results of our model-independent EoS inference. Such quantities include, *e.g.*, the speed of sound squared c_s^2 , the normalized trace anomaly and its logarithmic derivative with respect to energy density, $\Delta = (\epsilon - 3p)/(3\epsilon)$ and $\Delta' \equiv d\Delta/d\ln\epsilon$, the “conformal distance” $d_c \equiv \sqrt{\Delta^2 + (\Delta')^2}$, the polytropic index $\gamma = d\ln p/d\ln\epsilon$, as well as the normalized pressure p/p_{free} . In Table 1, taken from [139], we summarize the values these quantities take in various limits: sub-saturation-density CET calculations, model results for NM at NS-core densities, pQCD calculations around $n_B = 40 n_s$, field theories exhibiting exact conformal symmetry (CFTs), and systems undergoing a discontinuous first-order phase transition (FOPTs). A key takeaway from this table is that, as expected, high-density QM is considerably closer to the conformal limit than low- or high-density NM, implying that a possible conformalization of NS matter would be a strong indication of the presence of deconfined matter within NS cores.

The first serious attempt to inspect the possible conformalization of matter inside NS cores in a model-independent fashion was made in [107], where the primary quantities studied included the polytropic index γ , the speed of sound squared c_s^2 , and the normalized pressure p/p_{free} . The results that emerged from this hard-cut-type study are summarized in Fig. 11. In short, they indicate that while the properties of matter in the inner cores of light $1.4M_\odot$ NSs (blue and cyan diamonds in the figure) are well in line with those expected based on model calculations for dense NM, things look dramatically different in the cores of maximally massive NSs (red and magenta circles). In the latter case, the inferred properties of NS matter lie considerably closer to those of QM at perturbative densities, indicating the likely presence of deconfined matter inside the most massive stable NSs.

Table 1. A summary of the values that several physical quantities take in various limits defined in the main text. The table is taken from [139].

	CET	Dense NM	Pert. QM	CFTs	FOPT
c_s^2	$\ll 1$	$[0.25, 0.6]$	$\lesssim 1/3$	$1/3$	0
Δ	$\approx 1/3$	$[0.05, 0.25]$	$[0, 0.15]$	0	$1/3 - p_{\text{PT}}/\epsilon$
Δ'	≈ 0	$[-0.4, -0.1]$	$[-0.15, 0]$	0	$1/3 - \Delta$
d_c	$\approx 1/3$	$[0.25, 0.4]$	$\lesssim 0.2$	0	$\geq 1/(3\sqrt{2})$
γ	≈ 2.5	$[1.95, 3.0]$	$[1, 1.7]$	1	0
p/p_{free}	$\ll 1$	$[0.25, 0.35]$	$[0.5, 1]$	—	$p_{\text{PT}}/p_{\text{free}}$

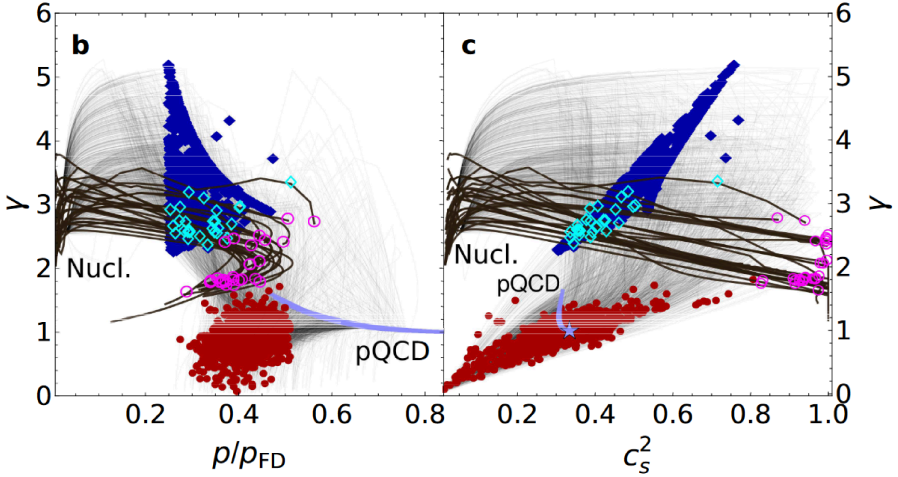


Fig. 11. The model-independent NS-matter EoSs from [107] (thin gray lines) plotted as functions of normalized pressure, polytropic index, and speed of sound squared, and compared to the high-density pQCD limit (blue band) and viable models of high-density NM (thick black lines). Shown are also blue and cyan diamonds corresponding to the centers of $1.4M_{\odot}$ NSs according to the interpolated and model EoSs, respectively, and the same for maximally massive NSs (red and magenta circles).

In the recent work [139], the argument for QM cores was put on a considerably firmer footing by tracking the conformalization of a number of new physical quantities, utilizing the largest set of NS observations to date, and most importantly, taking advantage of Bayesian inference techniques in the analysis. As an optimal quantity for tracking the conformalization of NS matter, the authors of [139] chose the conformal distance $d_c \equiv \sqrt{\Delta^2 + (\Delta')^2}$, which extends the use of the normalized conformal anomaly, first proposed in the NS context by Fujimoto *et al.* [140]. The reasoning used was simple: while Δ may take on small values due to the quantity changing its sign around a given density, d_c tends to zero only when both Δ and its rate of change are very small. In addition, when the triplet Δ , Δ' , and d_c all tend to zero, then the closely related quantities c_s^2 and γ will also approach their conformal values of $1/3$ and 1 due to the relations

$$\Delta = \frac{1}{3} - \frac{c_s^2}{\gamma}, \quad \Delta' = c_s^2 \left(\frac{1}{\gamma} - 1 \right). \quad (51)$$

In Fig. 12 (left) from [139], we indeed witness the rapid conformalization of d_c between the central densities of $1.4M_\odot$ and maximal-mass NSs. The corresponding likelihood for the existence of conformalized matter in the cores of the most massive stable NSs was found to be around 80–90%, with the remaining 10–20% likelihood corresponding to EoSs exhibiting FOPT-like behavior at the central densities of maximal-mass NSs, indicating the likely presence of a destabilizing phase transition (see also [141] for a related study with similar conclusions). Whether this observation can be interpreted as an indication of the existence of QM cores constitutes, in principle, a separate question, but some further reaffirmation can be obtained from the normalized pressure, shown in Fig. 12, right, obtaining a value indicative of a QM-like number of effective degrees of freedom at the same densities. It is indeed worth noting that at high temperatures, the deconfinement transition has been shown to occur when $p/p_{\text{free}} \lesssim 0.2$ [137, 138], while in the cold and dense case, the same quantity obtains values around 0.4 in the inner cores of maximal-mass stars.

As noted already in [107, 139] and alluded to above, one viable scenario still remains that would altogether forbid the existence of QM cores: a strong FOPT destabilizing the NS mass–radius sequence as soon as the central density of the star reaches the transition density. While it is nontrivial to make likelihood comparisons between two disjoint EoS ensembles — *e.g.*, the non-QM EoS in studies like [139] corresponded to arbitrarily rapid crossovers — this issue has been studied in recent literature to some extent. In the hard-cut study [142], it was demonstrated that if a strong first-order transition is inserted in a simple polytropic NS-matter EoS, it is possible

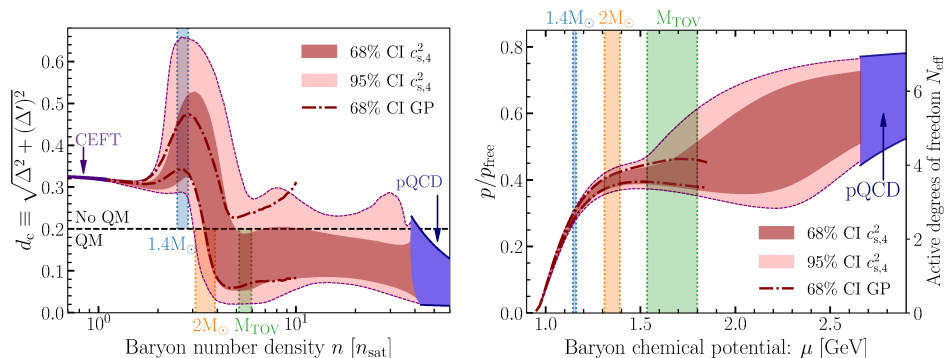


Fig. 12. The behavior of the conformal distance d_c and the normalized pressure p/p_{free} in a recent Bayesian study of the NS-matter EoS [139]. Shown are 68 and 95% credible intervals from a speed-of-sound interpolation and the 68% credible interval from a GP calculation, while the vertical bars indicate the central densities of $1.4M_\odot$, $2M_\odot$, and maximal-mass (M_{TOV}) NSs at the 68% level.

to construct EoSs that extend beyond the limits derived in papers such as [107, 108]. More recently, Komoltsev implemented FOPTs to a Bayesian framework involving GP regression [143], finding no evidence either favoring or disfavoring destabilizing solutions. Other features seen in previous EoS inferences without explicit phase transitions were seen to stay intact, though, including most importantly a clear peak in the speed of sound. More work is clearly needed to confirm or rule out the presence of a destabilizing FOPT, possibly using a future postmerger GW signal that may carry traces of the presence of a discontinuous transition [144, 145].

8. Outlook to future developments

In the lecture notes at hand, we have reviewed in some detail both the methods used in first-principles perturbative thermal field theory calculations within high-density quark matter, and the applications of these results to the model-independent determination of the neutron-star-matter equation of state. In this discussion, we have not aimed at a self-consistent text-book-style presentation, but have rather tried to aid the interested reader in their journey of self-study. We have done so by filling in a number of gaps left uncovered in existing textbooks on thermal field theory and by providing numerous references to original research articles, including both older classics of the field and more recent studies.

The topic of these lecture notes belongs to a growing subfield of nuclear and particle physics, aimed at understanding the properties of strongly interacting matter inside NSs. This is a rapidly evolving research topic, not least due to the pace, at which neutron-star observations have progressed during the past 10–15 years. Many recent advances have been made possible by an efficient interplay between microscopic theoretical calculations and new observational insights, of which the quest to discover a new phase of QCD matter inside NS cores represents a prime example. Without *ab-initio* limits for the properties of low-density nuclear matter and high-density QM, NS observations would at best provide ballpark estimates for the properties of NS interiors, and without new observational constraints, the accuracy of EoS inference would not have progressed much during the past decade.

Within the next couple of years, several important advances relevant for the physics of NSs and their potential QM cores can be expected to emerge from the microscopic side. CET and pQCD calculations of the thermodynamics of low-temperature QCD matter are advancing at a rapid pace, with the completion of the $O(\alpha_s^3)$ pressure of cold and dense QM being finally in sight. Once complete, this result is expected to dramatically improve the precision to which we know the properties of deconfined matter at densities of the order of $15\text{--}40 n_s$. This will immediately have a significant impact

on the model-independent inference of the NS-matter EoS, which continues to be an active topic of research [146–148]. Alongside such developments, the transport properties of dense QM are presently under intense scrutiny, and model-independent bounds for the temperature dependence of the NS-matter EoS, improving previous estimates such as those of [149, 150], are upcoming. As reviewed in these lecture notes, a combination of such theoretical results and future observational advances, including the potential detection of a postmerger GW signal from a binary NS merger, are expected to hold the key to resolving long-standing puzzles in nuclear astrophysics.

I would like to thank Niko Jokela, Aapeli Kärkkäinen, Mika Nurmela, Risto Paatelainen, and Tomi Ruosteoja for useful comments on early versions of these lecture notes.

REFERENCES

- [1] L.D. Landau, «On the theory of stars», in: D. Ter Haar (Ed.) «Collected Papers of L.D. Landau», *Pergamon*, 1965, pp. 60–62.
- [2] A. Hewish *et al.*, «Observation of a Rapidly Pulsating Radio Source», *Nature* **217**, 709 (1968).
- [3] J. Antoniadis *et al.*, «A Massive Pulsar in a Compact Relativistic Binary», *Science* **340**, 1233232 (2013).
- [4] NANOGrav Collaboration (H.T. Cromartie *et al.*), «Relativistic Shapiro delay measurements of an extremely massive millisecond pulsar», *Nat. Astron.* **4**, 72 (2019).
- [5] E. Fonseca *et al.*, «Refined Mass and Geometric Measurements of the High-mass PSR J0740+6620», *Astrophys. J. Lett.* **915**, L12 (2021).
- [6] I.I. Shapiro, «Fourth Test of General Relativity», *Phys. Rev. Lett.* **13**, 789 (1964).
- [7] A.W. Shaw *et al.*, «The radius of the quiescent neutron star in the globular cluster M13», *Mon. Not. R. Astron. Soc.* **476**, 4713 (2018).
- [8] J. Nättilä *et al.*, «Neutron star mass and radius measurements from atmospheric model fits to X-ray burst cooling tail spectra», *Astron. Astrophys.* **608**, A31 (2017).
- [9] M.C. Miller *et al.*, «PSR J0030+0451 Mass and Radius from NICER Data and Implications for the Properties of Neutron Star Matter», *Astrophys. J. Lett.* **887**, L24 (2019).
- [10] M.C. Miller *et al.*, «The Radius of PSR J0740+6620 from NICER and XMM-Newton Data», *Astrophys. J. Lett.* **918**, L28 (2021).
- [11] T.E. Riley *et al.*, «A NICER View of PSR J0030+0451: Millisecond Pulsar Parameter Estimation», *Astrophys. J. Lett.* **887**, L21 (2019).
- [12] T.E. Riley *et al.*, «A NICER View of the Massive Pulsar PSR J0740+6620

- Informed by Radio Timing and XMM-Newton Spectroscopy», *Astrophys. J. Lett.* **918**, L27 (2021).
- [13] LIGO Scientific, Virgo collaborations (B.P. Abbott *et al.*), «GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral», *Phys. Rev. Lett.* **119**, 161101 (2017).
 - [14] LIGO Scientific and Virgo collaborations (B. Abbott *et al.*), «GW170817: Measurements of Neutron Star Radii and Equation of State», *Phys. Rev. Lett.* **121**, 161101 (2018).
 - [15] LIGO Scientific and Virgo collaborations (B.P. Abbott *et al.*), «Properties of the Binary Neutron Star Merger GW170817», *Phys. Rev. X* **9**, 011001 (2019).
 - [16] B.P. Abbott *et al.*, «Multi-messenger Observations of a Binary Neutron Star Merger», *Astrophys. J. Lett.* **848**, L12 (2017).
 - [17] B.P. Abbott *et al.*, «Gravitational Waves and Gamma-rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A», *Astrophys. J. Lett.* **848**, L13 (2017).
 - [18] M. Maggiore *et al.*, «Science Case for the Einstein Telescope», *J. Cosmol. Astropart. Phys.* **03**, 050 (2020).
 - [19] D. Reitze *et al.*, «Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO», *Bull. Am. Astron. Soc.* **51**, 035 (2019), [arXiv:1907.04833 \[astro-ph.IM\]](#).
 - [20] P. de Forcrand, «Simulating QCD at finite density», *PoS (LAT2009)*, 010 (2009).
 - [21] R. Machleidt, D.R. Entem, «Chiral effective field theory and nuclear forces», *Phys. Rep.* **503**, 1 (2011).
 - [22] M. Laine, A. Vuorinen, «Basics of Thermal Field Theory. A Tutorial on Perturbative Computations», *Lecture Notes in Physics*, Vol. 925, *Springer*, 2016.
 - [23] J. Ghiglieri, A. Kurkela, M. Strickland, A. Vuorinen, «Perturbative Thermal QCD: Formalism and Applications», *Phys. Rep.* **880**, 1 (2020).
 - [24] J.M. Pawłowski, «Aspects of the functional renormalisation group», *Ann. Phys.* **322**, 2831 (2007).
 - [25] J. Casalderrey-Solana *et al.*, «Gauge/String Duality, Hot QCD and Heavy Ion Collisions», *Cambridge University Press*, 2014.
 - [26] C. Hoyos, N. Jokela, A. Vuorinen, «Holographic approach to compact stars and their binary mergers», *Prog. Part. Nucl. Phys.* **126**, 103972 (2022).
 - [27] J. Collins, «Foundations of Perturbative QCD», *Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology*, Vol. 32, *Cambridge University Press*, 2023.
 - [28] J. Ghiglieri *et al.*, «Next-to-leading order thermal photon production in a weakly coupled quark–gluon plasma», *J. High Energy Phys.* **2013**, 010 (2013).

- [29] J. Ghiglieri, G.D. Moore, D. Teaney, «QCD shear viscosity at (almost) NLO», *J. High Energy Phys.* **2018**, 179 (2018).
- [30] A. Kurkela, G.D. Moore, «Bjorken flow, plasma instabilities, and thermalization», *J. High Energy Phys.* **2011**, 120 (2011).
- [31] J.I. Kapusta, C. Gale, «Finite-temperature Field Theory: Principles and Applications», Cambridge Monographs on Mathematical Physics, *Cambridge University Press*, 2011.
- [32] M.L. Bellac, «Thermal Field Theory», Cambridge Monographs on Mathematical Physics, *Cambridge University Press*, 2011.
- [33] M.H. Anderson *et al.*, «Observation of Bose–Einstein Condensation in a Dilute Atomic Vapor», *Science* **269**, 198 (1995).
- [34] T. Gorda, S. Säppi, «Cool quark matter with perturbative quark masses», *Phys. Rev. D* **105**, 114005 (2022).
- [35] B.A. Freedman, L.D. McLerran, «Fermions and gauge vector mesons at finite temperature and density. III. The ground-state energy of a relativistic quark gas», *Phys. Rev. D* **16**, 1169 (1977).
- [36] A.I. Davydychev, P. Navarrete, Y. Schröder, «Factorizing two-loop vacuum sum-integrals», *J. High Energy Phys.* **2024**, 104 (2024).
- [37] T. Appelquist, R.D. Pisarski, «High-temperature Yang–Mills Theories and three-dimensional quantum chromodynamics», *Phys. Rev. D* **23**, 2305 (1981).
- [38] K. Kajantie, M. Laine, K. Rummukainen, M.E. Shaposhnikov, «Generic rules for high temperature dimensional reduction and their application to the standard model», *Nucl. Phys. B* **458**, 90 (1996).
- [39] E. Braaten, A. Nieto, «Effective field theory approach to high temperature thermodynamics», *Phys. Rev. D* **51**, 6990 (1995).
- [40] E. Braaten, A. Nieto, «Free energy of QCD at high temperature», *Phys. Rev. D* **53**, 3421 (1996).
- [41] A. Vuorinen, L.G. Yaffe, «Z(3)-symmetric effective theory for SU(3) Yang–Mills theory at high temperature», *Phys. Rev. D* **74**, 025011 (2006).
- [42] P. de Forcrand, A. Kurkela, A. Vuorinen, «Center-symmetric effective theory for high-temperature SU(2) Yang–Mills theory», *Phys. Rev. D* **77**, 125014 (2008).
- [43] A. Hart, M. Laine, O. Philipsen, «Static correlation lengths in QCD at high temperatures and finite densities», *Nucl. Phys. B* **586**, 443 (2000).
- [44] A. Vuorinen, «Pressure of QCD at finite temperatures and chemical potentials», *Phys. Rev. D* **68**, 054017 (2003).
- [45] A. Ipp, K. Kajantie, A. Rebhan, A. Vuorinen, «Pressure of deconfined QCD for all temperatures and quark chemical potentials», *Phys. Rev. D* **74**, 045016 (2006).
- [46] E. Braaten, R.D. Pisarski, «Soft amplitudes in hot gauge theories: A general analysis», *Nucl. Phys. B* **337**, 569 (1990).

- [47] E. Braaten, R.D. Pisarski, «Simple effective Lagrangian for hard thermal loops», *Phys. Rev. D* **45**, R1827 (1992).
- [48] J.O. Andersen, E. Braaten, M. Strickland, «Hard-Thermal-Loop Resummation of the Free Energy of a Hot Gluon Plasma», *Phys. Rev. Lett.* **83**, 2139 (1999).
- [49] N. Haque *et al.*, «Three-loop HTLpt thermodynamics at finite temperature and chemical potential», *J. High Energy Phys.* **2014**, 027 (2014).
- [50] T. Gorda *et al.*, «Soft Interactions in Cold Quark Matter», *Phys. Rev. Lett.* **127**, 162003 (2021).
- [51] A. Kurkela, A. Vuorinen, «Cool Quark Matter», *Phys. Rev. Lett.* **117**, 042501 (2016).
- [52] T. Gorda *et al.*, «Degenerate fermionic matter at N³LO: Quantum electrodynamics», *Phys. Rev. D* **107**, L031501 (2023).
- [53] T. Gorda *et al.*, «Soft photon propagation in a hot and dense medium to next-to-leading order», *Phys. Rev. D* **107**, 036012 (2023).
- [54] T. Gorda, R. Paatelainen, S. Säppi, K. Seppänen, «Equation of State of Cold Quark Matter to $O(\alpha_s^3 \ln \alpha_s)$ », *Phys. Rev. Lett.* **131**, 181902 (2023).
- [55] E. Braaten, R.D. Pisarski, T.-C. Yuan, «Production of Soft Dileptons in the Quark–Gluon Plasma», *Phys. Rev. Lett.* **64**, 2242 (1990).
- [56] H. Heiselberg, C.J. Pethick, «Transport and relaxation in degenerate quark plasmas», *Phys. Rev. D* **48**, 2916 (1993).
- [57] N. Haque, M.G. Mustafa, «Hard Thermal Loop — theory and applications», [arXiv:2404.08734 \[hep-ph\]](https://arxiv.org/abs/2404.08734).
- [58] S. Caron-Huot, «Hard thermal loops in the real-time formalism», *J. High Energy Phys.* **2009**, 004 (2009).
- [59] T. Gorda, R. Paatelainen, S. Säppi, K. Seppänen, «Soft gluon self-energy at finite temperature and density: hard NLO corrections in general covariant gauge», *J. High Energy Phys.* **2023**, 021 (2023).
- [60] H.A. Weldon, «Structure of the Gluon Propagator at Finite Temperature», *Ann. Phys.* **271**, 141 (1999).
- [61] G.D. Moore, «Pressure of Hot QCD at Large N_f », *J. High Energy Phys.* **2002**, 055 (2002).
- [62] A. Ipp, G.D. Moore, A. Rebhan, «Comment on and erratum to ‘Pressure of Hot QCD at Large N_f ’», *J. High Energy Phys.* **2003**, 037 (2003).
- [63] A. Ipp, A. Rebhan, A. Vuorinen, «Perturbative QCD at nonzero chemical potential: Comparison with the large N_f limit and apparent convergence», *Phys. Rev. D* **69**, 077901 (2004).
- [64] A. Gynther, A. Kurkela, A. Vuorinen, « $N_f^3 g^6$ term in the pressure of hot QCD», *Phys. Rev. D* **80**, 096002 (2009).
- [65] A.D. Linde, «Infrared problem in thermodynamics of the Yang–Mills gas», *Phys. Lett. B* **96**, 289 (1980).
- [66] T. Gorda *et al.*, «Cold quark matter at N³LO: Soft contributions», *Phys. Rev. D* **104**, 074015 (2021).

- [67] T. Gorda *et al.*, «Next-to-Next-to-Next-to-Leading Order Pressure of Cold Quark Matter: Leading Logarithm», *Phys. Rev. Lett.* **121**, 202701 (2018).
- [68] A. Kurkela, P. Romatschke, A. Vuorinen, «Cold quark matter», *Phys. Rev. D* **81**, 105021 (2010).
- [69] I. Ghisoiu *et al.*, «On high-order perturbative calculations at finite density», *Nucl. Phys. B* **915**, 102 (2017).
- [70] T. Gorda, J. Österman, S. Säppi, «Augmenting the residue theorem with boundary terms in finite-density calculations», *Phys. Rev. D* **106**, 105026 (2022).
- [71] J. Österman, P. Schicho, A. Vuorinen, «Integrating by parts at finite density», *J. High Energy Phys.* **2023**, 212 (2023).
- [72] P. Navarrete, R. Paatelainen, K. Seppänen, «Perturbative QCD meets phase quenching: The pressure of cold Quark Matter», [arXiv:2403.02180 \[hep-ph\]](#).
- [73] A. Kärkkäinen *et al.*, in preparation.
- [74] Z. Capatti, V. Hirschi, D. Kermanschah, B. Ruijl, «Loop-Tree Duality for Multiloop Numerical Integration», *Phys. Rev. Lett.* **123**, 151602 (2019).
- [75] J.P. Blaizot, E. Iancu, A. Rebhan, «Quark number susceptibilities from HTL-resummed thermodynamics», *Phys. Lett. B* **523**, 143 (2001).
- [76] G.D. Moore, T. Gorda, «Bounding the QCD Equation of State with the Lattice», *J. High Energy Phys.* **2023**, 133 (2023).
- [77] T. Gorda, O. Komoltsev, A. Kurkela, A. Mazeliauskas, «Bayesian uncertainty quantification of perturbative QCD input to the neutron-star equation of state», *J. High Energy Phys.* **2023**, 002 (2023).
- [78] D. Radice, S. Bernuzzi, A. Perego, R. Haas, «A new moment-based general-relativistic neutrino-radiation transport code: Methods and first applications to neutron star mergers», *Mon. Not. R. Astron. Soc.* **512**, 1499 (2022).
- [79] M. Chabanov, L. Rezzolla, «Impact of bulk viscosity on the post-merger gravitational-wave signal from merging neutron stars», [arXiv:2307.10464 \[gr-qc\]](#).
- [80] M.G. Alford *et al.*, «Viscous Dissipation and Heat Conduction in Binary Neutron-Star Mergers», *Phys. Rev. Lett.* **120**, 041101 (2018).
- [81] J. Cruz Rojas *et al.*, «Estimate for the bulk viscosity of strongly coupled quark matter», [arXiv:2402.00621 \[hep-ph\]](#).
- [82] J.L. Hernandez, C. Manuel, L. Tolos, «Damping of density oscillations from bulk viscosity in quark matter», [arXiv:2402.06595 \[hep-ph\]](#).
- [83] J. Madsen, «Bulk viscosity of strange quark matter, damping of quark star vibration, and the maximum rotation rate of pulsars», *Phys. Rev. D* **46**, 3290 (1992).
- [84] A. Schmitt, P. Shternin, «Reaction Rates and Transport in Neutron Stars», *Astrophys. Space Sci. Libr.* **457**, 455 (2018).

- [85] A. Hosoya, K. Kajantie, «Transport coefficients of QCD matter», *Nucl. Phys. B* **250**, 666 (1985).
- [86] J. Ghiglieri, G.D. Moore, D. Teaney, «Second-Order Hydrodynamics in Next-to-Leading-Order QCD», *Phys. Rev. Lett.* **121**, 052302 (2018).
- [87] C. Hoyos *et al.*, «Transport in Strongly Coupled Quark Matter», *Phys. Rev. Lett.* **125**, 241601 (2020).
- [88] C. Hoyos *et al.*, «Holographic approach to transport in dense QCD matter», *Phys. Rev. D* **105**, 066014 (2022).
- [89] M. Järvinen, E. Kiritsis, F. Nitti, E. Préau, «Holographic neutrino transport in dense strongly-coupled matter», *J. High Energy Phys.* **2023**, 139 (2023).
- [90] C.J. Krüger, K.D. Kokkotas, P. Manoharan, S.H. Völkel, «Fast Rotating Neutron Stars: Oscillations and Instabilities», *Front. Astron. Space Sci.* **8**, 736918 (2021).
- [91] M. Jarvinen, E. Kiritsis, «Holographic models for QCD in the Veneziano limit», *J. High Energy Phys.* **2012**, 002 (2012).
- [92] N. Jokela, M. Järvinen, J. Remes, «Holographic QCD in the Veneziano limit and neutron stars», *J. High Energy Phys.* **2019**, 041 (2019).
- [93] A. Karch, E. Katz, «Adding flavor to AdS/CFT», *J. High Energy Phys.* **2002**, 043 (2002).
- [94] M. Alford, A. Harutyunyan, A. Sedrakian, «Bulk Viscosity of Relativistic $npe\mu$ Matter in Neutron-Star Mergers», *Particles* **5**, 361 (2022).
- [95] K. Schwenzer, «How long-range interactions tune the damping in compact stars», [arXiv:1212.5242 \[nucl-th\]](#).
- [96] M.G. Alford, A. Schmitt, K. Rajagopal, T. Schäfer, «Color superconductivity in dense quark matter», *Rev. Mod. Phys.* **80**, 1455 (2008).
- [97] J. Nättilä, J.J.E. Kajava, «Fundamental Physics with Neutron Stars», in: C. Bambi, A. Santangelo (Eds.) «Handbook of X-ray and Gamma-ray Astrophysics», *Springer, Singapore* 2022.
- [98] S. Han *et al.*, «Treating quarks within neutron stars», *Phys. Rev. D* **100**, 103022 (2019).
- [99] R.C. Tolman, «Static Solutions of Einstein's Field Equations for Spheres of Fluid», *Phys. Rev.* **55**, 364 (1939).
- [100] J.R. Oppenheimer, G.M. Volkoff, «On Massive Neutron Cores», *Phys. Rev.* **55**, 374 (1939).
- [101] A. Kurkela, P. Romatschke, A. Vuorinen, B. Wu, «Looking inside neutron stars: Microscopic calculations confront observations», [arXiv:1006.4062 \[astro-ph.HE\]](#).
- [102] C. Alcock, E. Farhi, A. Olinto, «Strange Stars», *Astrophys. J.* **310**, 261 (1986).
- [103] S. Weissenborn, D. Chatterjee, J. Schaffner-Bielich, «Hyperons and massive neutron stars: Vector repulsion and SU(3) symmetry», *Phys. Rev. C* **85**, 065802 (2012); *Erratum ibid.* **90**, 019904 (2014).

- [104] A. Kurkela, E.S. Fraga, J. Schaffner-Bielich, A. Vuorinen, «Constraining Neutron Star Matter with Quantum Chromodynamics», *Astrophys. J.* **789**, 127 (2014).
- [105] K. Hebeler, J.M. Lattimer, C.J. Pethick, A. Schwenk, «Equation of State and Neutron Star Properties Constrained by Nuclear Physics and Observation», *Astrophys. J.* **773**, 11 (2013).
- [106] E.S. Fraga, A. Kurkela, A. Vuorinen, «Interacting Quark Matter Equation of State for Compact Stars», *Astrophys. J. Lett.* **781**, L25 (2014).
- [107] E. Annala *et al.*, «Evidence for quark-matter cores in massive neutron stars», *Nat. Phys.* **16**, 907 (2020).
- [108] E. Annala *et al.*, «Multimessenger Constraints for Ultradense Matter», *Phys. Rev. X* **12**, 011058 (2022).
- [109] E. Annala, T. Gorda, A. Kurkela, A. Vuorinen, «Gravitational-Wave Constraints on the Neutron-Star-Matter Equation of State», *Phys. Rev. Lett.* **120**, 172703 (2018).
- [110] B. Margalit, B.D. Metzger, «Constraining the Maximum Mass of Neutron Stars from Multi-messenger Observations of GW170817», *Astrophys. J. Lett.* **850**, L19 (2017).
- [111] L. Rezzolla, E.R. Most, L.R. Weih, «Using Gravitational-wave Observations and Quasi-universal Relations to Constrain the Maximum Mass of Neutron Stars», *Astrophys. J. Lett.* **852**, L25 (2018).
- [112] M. Ruiz, S.L. Shapiro, A. Tsokaros, «GW170817, general relativistic magnetohydrodynamic simulations, and the neutron star maximum mass», *Phys. Rev. D* **97**, 021501(R) (2018).
- [113] A. Bauswein, O. Just, H.-T. Janka, N. Stergioulas, «Neutron-star Radius Constraints from GW170817 and Future Detections», *Astrophys. J. Lett.* **850**, L34 (2017).
- [114] D. Radice, A. Perego, F. Zappa, S. Bernuzzi, «GW170817: Joint Constraint on the Neutron Star Equation of State from Multimessenger Observations», *Astrophys. J. Lett.* **852**, L29 (2018).
- [115] E.R. Most, L.R. Weih, L. Rezzolla, J. Schaffner-Bielich, «New Constraints on Radii and Tidal Deformabilities of Neutron Stars from GW170817», *Phys. Rev. Lett.* **120**, 261103 (2018).
- [116] T. Dietrich *et al.*, «Multimessenger constraints on the neutron-star equation of state and the Hubble constant», *Science* **370**, 1450 (2020).
- [117] C.D. Capano *et al.*, «Stringent constraints on neutron-star radii from multimessenger observations and nuclear theory», *Nat. Astron.* **4**, 625 (2020).
- [118] P. Landry, R. Essick, «Nonparametric inference of the neutron star equation of state from gravitational wave observations», *Phys. Rev. D* **99**, 084049 (2019).
- [119] C.A. Raithel, F. Özel, D. Psaltis, «Tidal Deformability from GW170817 as a Direct Probe of the Neutron Star Radius», *Astrophys. J. Lett.* **857**, L23 (2018).

- [120] C.A. Raithel, F. Özel, «Measurement of the nuclear symmetry energy parameters from gravitational-wave events», *Astrophys. J.* **885**, 121 (2019).
- [121] G. Raaijmakers *et al.*, «Constraining the Dense Matter Equation of State with Joint Analysis of NICER and LIGO/Virgo Measurements», *Astrophys. J. Lett.* **893**, L21 (2020).
- [122] R. Essick, P. Landry, D.E. Holz, «Nonparametric inference of neutron star composition, equation of state, and maximum mass with GW170817», *Phys. Rev. D* **101**, 063007 (2020).
- [123] N. Jokela, M. Järvinen, G. Nijs, J. Remes, «Unified weak and strong coupling framework for nuclear matter and neutron stars», *Phys. Rev. D* **103**, 086004 (2021).
- [124] M. Al-Mamun *et al.*, «Combining Electromagnetic and Gravitational-Wave Constraints on Neutron-Star Masses and Radii», *Phys. Rev. Lett.* **126**, 061101 (2021).
- [125] T. Demircik, C. Ecker, M. Järvinen, «Rapidly Spinning Compact Stars with Deconfinement Phase Transition», *Astrophys. J. Lett.* **907**, L37 (2021).
- [126] L. Baiotti, L. Rezzolla, «Binary neutron star mergers: a review of Einstein's richest laboratory», *Rep. Prog. Phys.* **80**, 096901 (2017).
- [127] Y.-Z. Fan *et al.*, «Maximum gravitational mass $M_{\text{TOV}} = 2.25^{+0.08}_{-0.07} M_{\odot}$ inferred at about 3% precision with multimessenger data of neutron stars», *Phys. Rev. D* **109**, 043052 (2024).
- [128] M. Hippert, J. Noronha, P. Romatschke, «Upper Bound on the Speed of Sound in Nuclear Matter from Transport», [arXiv:2402.14085 \[nucl-th\]](#).
- [129] O. Komoltsev, A. Kurkela, «How Perturbative QCD Constrains the Equation of State at Neutron-Star Densities», *Phys. Rev. Lett.* **128**, 202701 (2022).
- [130] T. Gorda, O. Komoltsev, A. Kurkela, «Ab-initio QCD Calculations Impact the Inference of the Neutron-star-matter Equation of State», *Astrophys. J.* **950**, 107 (2023).
- [131] O. Komoltsev *et al.*, «Equation of state at neutron-star densities and beyond from perturbative QCD», [arXiv:2312.14127 \[nucl-th\]](#).
- [132] C.A. Raithel, F. Özel, D. Psaltis, «From Neutron Star Observables to the Equation of State. II. Bayesian Inference of Equation of State Pressures», *Astrophys. J.* **844**, 156 (2017).
- [133] M.W. Coughlin, T. Dietrich, B. Margalit, B.D. Metzger, «Multimessenger Bayesian parameter inference of a binary neutron star merger», *Mon. Not. R. Astron. Soc.* **489**, L91 (2019).
- [134] H. Güven, K. Bozkurt, E. Khan, J. Margueron, «Multimessenger and multiphysics Bayesian inference for the GW170817 binary neutron star merger», *Phys. Rev. C* **102**, 015805 (2020).
- [135] T. Malik, M. Ferreira, B.K. Agrawal, C. Providência, «Relativistic Description of Dense Matter Equation of State and Compatibility with Neutron Star Observables: A Bayesian Approach», *Astrophys. J.* **930**, 17 (2022).

- [136] J.-L. Jiang, C. Ecker, L. Rezzolla, «Bayesian Analysis of Neutron-star Properties with Parameterized Equations of State: The Role of the Likelihood Functions», *Astrophys. J.* **949**, 11 (2023).
- [137] S. Borsanyi *et al.*, «The QCD equation of state with dynamical quarks», *J. High Energy Phys.* **2010**, 077 (2010).
- [138] HotQCD Collaboration (A. Bazavov *et al.*), «Equation of state in (2+1)-flavor QCD», *Phys. Rev. D* **90**, 094503 (2014).
- [139] E. Annala *et al.*, «Strongly interacting matter exhibits deconfined behavior in massive neutron stars», *Nat. Commun.* **14**, 8451 (2023).
- [140] Y. Fujimoto, K. Fukushima, L.D. McLerran, M. Praszalowicz, «Trace Anomaly as Signature of Conformality in Neutron Stars», *Phys. Rev. Lett.* **129**, 252702 (2022).
- [141] M.-Z. Han, Y.-J. Huang, S.-P. Tang, Y.-Z. Fan, «Plausible presence of new state in neutron stars with masses above $0.98 M_{\text{TOV}}$ », *Sci. Bull.* **68**, 913 (2023).
- [142] T. Gorda *et al.*, «Constraints on Strong Phase Transitions in Neutron Stars», *Astrophys. J.* **955**, 100 (2023).
- [143] O. Komoltsev, «First-order phase transitions in the cores of neutron stars», [arXiv:2404.05637 \[nucl-th\]](#).
- [144] E.R. Most *et al.*, «Signatures of Quark–Hadron Phase Transitions in General-Relativistic Neutron-Star Mergers», *Phys. Rev. Lett.* **122**, 061101 (2019).
- [145] C. Ecker, K. Topolski, M. Järvinen, A. Stehr, «Prompt Black Hole Formation in Binary Neutron Star Mergers», [arXiv:2402.11013 \[astro-ph.HE\]](#).
- [146] A.C. Sempowski *et al.*, «From chiral EFT to perturbative QCD: a Bayesian model mixing approach to symmetric nuclear matter», [arXiv:2404.06323 \[nucl-th\]](#).
- [147] T. Malik, V. Dexheimer, C. Providência, «Astrophysics and Nuclear Physics Informed Interactions in Dense Matter: Insights from PSR J0437-4715», [arXiv:2404.07936 \[nucl-th\]](#).
- [148] S.-P. Tang, Y.-J. Huang, M.-Z. Han, Y.-Z. Fan, «Upper Limit of Sound Speed in Nuclear Matter: A Harmonious Interplay of Transport Calculation and Perturbative QCD Constraint», [arXiv:2404.09563 \[astro-ph.HE\]](#).
- [149] P.M. Chesler, N. Jokela, A. Loeb, A. Vuorinen, «Finite-temperature equations of state for neutron star mergers», *Phys. Rev. D* **100**, 066027 (2019).
- [150] D. Mroczek *et al.*, «Finite-temp. expansion of the dense-matter equation of state», [arXiv:2404.01658 \[astro-ph.HE\]](#).