

Computational Challenges in QCD Thermodynamics

Carleton DeTar and Frithjof Karsch

(USQCD Collaboration)

(Dated: February 6, 2013)

CONTENTS

Preamble	ii
I. Introduction	1
II. Current results and future challenges in QCD thermodynamics	
II.A. The QCD transition and equation of state	2
II.B. The critical endpoint	6
II.C. The properties of the quark-gluon plasma	8
II.D. The chiral limit	11
III. Outlook	
III.A. Scientific challenges	14
III.B. Computational challenges	15
III.C. Computational priorities	18
IV. Summary of accomplishments and challenges	18
IV.A. Accomplishments	18
IV.B. Challenges	20
References	21

Preamble

To a large extent complex many body interactions control the various phases of strong interaction matter that are of relevance for our understanding of the nuclear force and its role in determining the structure of nuclear matter. Accounting for their effects quantitatively requires nonperturbative techniques, such as the numerical simulation of lattice-regularized Quantum Chromodynamics (QCD). Such simulations are of particular importance in the temperature range close to phase changes, where properties of the matter change rapidly. This temperature range is currently probed experimentally in relativistic heavy ion experiments.

The decadal research program in QCD thermodynamics, using numerical simulations of lattice QCD, has to a large extent been described in a Scientific Grand Challenge report issued by the U.S. Department of Energy in 2009 [1]. We will discuss here the achievements of recent years and will outline the next steps that need to be taken to reach our goal of *understanding the phases of strongly interacting matter and the role they play in the cosmos* (NSAC 2007).

This white paper describes the opportunities for lattice QCD to play a key role in supporting the future experimental relativistic heavy ion program, and discusses the plans of the U.S. lattice QCD Collaboration (USQCD) for the next five years in this area. Companion documents on *Lattice QCD for Cold Nuclear Physics*, *Lattice Gauge Theory for Physics on the Intensity Frontier* and *Lattice Gauge Theory for Physics on the Energy Frontier* address prospects in other areas of nuclear physics and high-energy physics for which lattice calculations are essential. They can be found on the homepage of USQCD: <http://www.usqcd.org>

I. INTRODUCTION

During the last decade our understanding of properties of strong interaction matter at high temperature and vanishing as well as nonvanishing baryon number density has witnessed tremendous progress. Experiments performed at the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory (BNL) have made it evident that matter at temperatures close to, but above the QCD transition temperature exhibits properties far more complex than could have been expected for a thermal medium described by perturbative QCD in the asymptotically free regime [2, 3]. The system seems to be strongly coupled with a surprisingly small ratio of shear viscosity over entropy density. Nonperturbative calculations within the framework of lattice QCD had predicted this [4] as the equation of state close to, but above the transition temperature is far from being that of an ideal gas. Instead, it has a large trace anomaly and screening of heavy quark free energies as well as thermal masses still exhibit confining features in the sense that the heavy quark potential still allows for bound hadronic states even at temperatures above the crossover transition. However, these statements still suffer, to some extent, from the artifacts of contemporary lattice calculations that arise from discretization errors or a still only approximate implementation of symmetries of the continuum theory. Refining them will be a major challenge for lattice QCD calculations in the coming years.

During the next years we will see a large number of new experimental results from heavy ion experiments at RHIC as well as the Large Hadron Collider (LHC) at CERN. The latter will probe the high temperature phase of QCD at (almost) vanishing net baryon number in a wider temperature range, providing new information about thermal dilepton and photon emission from the quark-gluon plasma, heavy quark bound states, the equilibration and diffusion of light and heavy quarks in dense matter, as well as information about other transport coefficients that characterize the *perfect fluid*.




FIG. 1. A schematic phase diagram of QCD, including a still hypothetical critical point, and the parameter range covered by the current beam energy scan at RHIC [5].

Furthermore, the Beam Energy Scan (BES) [5, 6], recently performed at RHIC, and, we hope, to be continued in upcoming years, will provide much information about fluctuations in net proton and net electric charge numbers. This information will allow us to explore the phase diagram at nonzero baryon chemical potential; detector upgrades of STAR and PHENIX will lead to new high precision data on properties of strong interaction matter close to the QCD transition.

Providing theoretical input from nonperturbative lattice QCD calculations at vanishing as well as nonvanishing values of the quark chemical potentials, extrapolated to the continuum limit and performed with physical values of the quark masses, thus remains of great importance. In order to confront experimental results on higher order cumulants of charge fluctuations with QCD calculations that go beyond

the predictions of model calculations, *e.g.*, the hadron resonance gas model, accurate numerical results are needed.

Aside from these experiment-driven motivations for future numerical calculations in QCD thermodynamics, there also exist fundamental theoretical questions related to the phase structure of quantum chromodynamics at nonzero temperature and/or nonzero density that should be addressed, as QCD is the only component of the Standard Model that stands on its own as a well-defined quantum field theory. The spontaneously broken chiral symmetry of QCD, the axial anomaly and the topological structure of the QCD vacuum have played an important role in the development of our theoretical picture of the QCD phase diagram and the mechanisms leading to phase transitions in a quantum field theory. With the increase in computing resources, we have now reached a point at which numerical simulations not only with physical values of quark masses but even with mass parameters smaller than physical have become possible. This will allow us to come close to the chiral limit of QCD and directly probe the universal properties of QCD thermodynamics, including the existence of genuine phase transitions at vanishing quark mass values and vanishing as well as nonvanishing baryon chemical potential.

These are pressing physics questions that need and can be answered during this decade. They can be answered with the anticipated development of computer hardware achieving exascale levels of performance together with the development of new algorithms and software that exploit exascale hardware for QCD calculations. The decadal research program in QCD thermodynamics has, to a large extent, been described in a Scientific Grand Challenge report issued by the U.S. Department of Energy in 2009 [1]. We will discuss here how we proceeded in recent years and will outline the next steps that need to be taken to reach our goal of *understanding the phases of strongly interacting matter and the role they play in the cosmos* (NSAC 2007).

II. CURRENT RESULTS AND FUTURE CHALLENGES IN QCD THERMODYNAMICS

A. The QCD transition and the equation of state

The QCD transition temperature and the equation of state are among the most basic quantities of QCD thermodynamics that are of obvious importance for the description of the hydrodynamic expansion of hot and dense matter created in heavy ion experiments as well as the early universe. The transition in finite temperature QCD has aspects related to deconfinement and chiral symmetry restoration. For the value of the light (u and d) quark and strange quark masses realized in nature only the chiral aspects of the transition allow for a definition of the transition temperature. In a strict sense, the true chiral phase transition occurs only at vanishing values of the light quark masses, where we expect universal critical scaling. Using the theory of critical universality, we can then relate this transition to the crossover observed at physical quark masses, resulting in a reasonably well-defined pseudocritical transition temperature and various crossover phenomena associated with the nearby chiral transition.

The properties of the chiral transition as a function of the light quark masses have been studied using improved staggered fermion formulations. It was shown that, for masses and temperatures close to the zero-mass chiral phase transition, the behavior of key chiral observables is governed by the $O(4)$ universality class in three dimensions [7] (see also section II.D).

An analysis of the pseudo-critical temperature at the crossover transition, extrapolated to the continuum limit gave the value $T_c = 154(9)$ MeV [8]. This value is in good agreement with the value for the crossover transition temperature obtained by using another staggered fermion discretization scheme, the so-called stout staggered action [9].

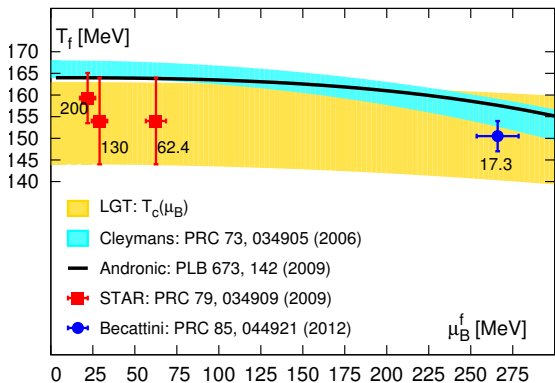


FIG. 2. The QCD transition compared with freeze-out temperatures determined in heavy ion collisions at RHIC and the SPS. The broad band shows the transition temperature $T_c(\mu_B) = T_c(0) + \kappa_B(\mu_B/T)^2$ with $T_c(0) = 154(9)$ MeV [8] and $\kappa_B = 0.0066(7)$ [11] obtained in lattice QCD calculations.

Some results for the QCD transition temperature and a comparison with freeze-out parameters are shown in Fig. 2. In order to address this issue in lattice QCD calculations based on a systematic expansion in baryon chemical potentials, one has to calculate higher order expansion coefficients in Taylor series of, *e.g.*, the chiral condensate and its susceptibility (see also section II.B).

At present simulations with improved staggered fermion actions are most advanced in determining the QCD transition temperature for light, almost physical values of the quark masses. However, also simulations with chiral fermions, *e.g.*, domain wall fermions, are now reaching a point where calculations with physical values of the quark masses are becoming feasible. Simulations with pion masses of about 200 MeV have been performed [12, 13], and it has been demonstrated that the calculation of charge fluctuations is also possible in this context [12], which should make it feasible to calculate the transition temperature as a function of the baryon chemical potential, *i.e.*, the determination of the curvature of the transition line in the QCD phase diagram, also in a chiral fermion formulation.

Calculations of the equation of state (EoS) [14, 15], including its extension to nonzero baryon number density [16] have been improved considerably over the past few years. These

The above value of T_c is close to experimentally determined freeze-out temperatures in high energy runs at RHIC [10]. First studies of the dependence of the transition temperature on a nonzero baryon chemical potential (μ_B) [11] are also in good agreement with experimental findings, that are based on a comparison of experimental data for particle yields with hadron resonance gas model calculations. This seems to suggest that in a wide range of baryon chemical potentials, $\mu_B \lesssim 200$ MeV, which corresponds to beam energies at RHIC of $\sqrt{s_{NN}} \gtrsim 20$ GeV, calculations of the transition temperature based on next-to-leading order Taylor expansions of thermodynamic observables can be controlled and can be related to experimental findings at the time of freeze-out. There is, however, the tendency that the crossover temperature for the QCD transition and the exper-




FIG. 3. Left: The ratio of the entropy density to its ideal gas value, calculated with asqtad, p4 and stout improved staggered fermion action and compared with the resummed perturbative calculations (dashed lines). We also show this ratio for the strongly coupled $N = 4$ supersymmetric plasma obtained by AdS/CFT correspondence as the black line (see text). Right: The trace anomaly calculated for the HISQ action [17], compared with the stout continuum estimate.

calculations have been performed using improved staggered fermion formulations and now include a wide range of lattice spacings corresponding to temporal extent $N_\tau = 4$ to 12 (the lattice spacing a is related to N_τ and the temperature T as $a = 1/(N_\tau T)$). The HotQCD collaboration calculated the EoS using two different discretization schemes, the so-called asqtad and p4 improved staggered formulations, on lattices with temporal extent $N_\tau = 6$ and 8. A comparison of these results with a continuum estimate of the entropy density based on calculations with yet another staggered fermion discretization scheme, the stout action [14], and lattices with temporal extent $N_\tau = 6, 8$ and 10 is shown in Fig. 3. All these calculations have in common that the entropy density rises rapidly in a narrow temperature interval, reflecting the liberation of new degrees of freedom at the QCD transition temperature. At higher temperatures the entropy density quickly reaches a value that is close to the entropy density of a noninteracting quark-gluon gas. Current calculations performed with different discretization schemes, however, also show distinct differences. For the entropy density, the differences between the continuum extrapolated results obtained with the stout action and current HotQCD results are understood to be due to so-called taste symmetry breaking effects, which are known to be large for asqtad and p4 actions at low temperature. The discrepancies at high temperature, on the other hand, are at present not well understood. In this temperature range effects of taste symmetry breaking should not be important. On the other hand, at sufficiently high temperatures the stout action has the same large cutoff effects as unimproved staggered fermions, while the leading order cutoff effects are eliminated in the asqtad and p4 discretization schemes.

At present the highly improved staggered quark (HISQ) discretization scheme, used by the HotQCD collaboration in their calculations of the QCD equation of state, seems to be best suited for reducing lattice discretization errors at all temperatures. It is designed to reduce drastically taste violation effects at low temperature (more efficiently than in the stout scheme) and has the same small discretization errors at high temperature as the asqtad action. The HISQ action currently is used by the HotQCD collaboration to refine existing calculations of the equation of state. The basic observable here is the trace anomaly, *i.e.*, the

difference between energy density and three times the pressure, $\epsilon - 3p$, which is calculated on lattices with temporal extent varying from $N_\tau = 4$ to $N_\tau = 12$. The lattice cutoff thus varies by a factor three, which will allow a systematic analysis of the extrapolation to the continuum limit. Some results from this ongoing study are shown in the right hand part of Fig. 3 and are compared with a continuum estimate based on calculations with the stout action. While results obtained with the HISQ and stout discretization schemes agree well at low temperature, confirming that both schemes control taste violation effects well, differences similar to those seen in calculations with the asqtad and p4 actions persist at higher temperatures. At temperatures $T > 300\text{MeV}$ the HISQ results agree well with previous HotQCD results obtained with p4 and asqtad action. Furthermore, the HISQ and asqtad $N_\tau = 12$ data agree within errors. In the temperature region $190\text{ MeV} < T < 350\text{ MeV}$ the stout data for $\epsilon - 3p$ are, however, significantly below the HISQ data. Differences between the stout and HISQ calculations show up most prominently in the vicinity of the maximum of $(\epsilon - 3p)/T^4$, which occurs well into the high temperature phase at about $1.2T_c$. In order to perform a controlled continuum extrapolation in this temperature range, obviously one needs to reduce statistical errors in order to isolate clearly those effects coming from a systematic cutoff dependence of the data. It is expected that this can soon be accomplished also within the HISQ discretization scheme by making use of the new generation of Petaflop/s computing resources at Leadership Class computing facilities. Clearly, achieving agreement on continuum extrapolated results for the QCD equation of state must have high priority.

In Fig. 3 we also compare lattice QCD results for the entropy density with resummed perturbative calculations. The results obtained with the p4 and asqtad action agree well with the resummed perturbative result while the stout continuum estimate is below it. The fact that the entropy density is close to the ideal gas value and is well described by resummed perturbative calculations is often interpreted as an indication of the weakly interacting nature of the quark-gluon plasma (QGP). On the other hand in a strongly coupled $N = 4$ supersymmetric gluon plasma the entropy density calculated using AdS/CFT correspondence also deviates from the ideal gas limit by only 25%. It is interesting that in the temperature range relevant for the discussion of heavy ion experiments at RHIC and LHC, the stout results are close to strong coupling results predicted by AdS/CFT. This, of course, will change at higher temperatures. Nonetheless, this makes it clear that precise lattice results for the equation of state are needed to distinguish between the strongly and weakly interacting nature of a QGP. As the entropy density, as well as pressure and energy density, are deduced from the trace anomaly, making use of standard thermodynamic relations, it is mandatory to arrive at firm results for the latter. This will also provide other bulk thermodynamic observables, such as the velocity of sound, that are of relevance for the hydrodynamic modeling of the expansion of dense matter created in heavy ion collisions.

The advances made in studies of bulk QCD thermodynamics, however, do not complete the analysis of the QCD transition temperature and the equation of state. At present, essentially all continuum extrapolated results for the QCD transition and bulk thermodynamics have been obtained using a particular discretization scheme – the staggered fermion formalism. In general it is desirable to cross-check these results within another fermion discretization scheme. Domain Wall fermions (DWF) are particular attractive in this respect as the DWF scheme maintains exact chiral symmetry even at finite values of the cutoff, which may be of importance for the analysis of the QCD transition (see also section II.D). Moreover, even in the context of staggered fermions, extending the current results for the QCD equation of state to nonvanishing chemical potentials, (*i.e.*, nonvanishing strangeness chemical potential

as well as nonvanishing baryon chemical potential), is of great importance for a better understanding of the QCD phase diagram and the comparison with experimental findings at lower beam energies. Future calculations with nonzero chemical potentials, described in the next section, will also provide the necessary input for a calculation of the equation of state up to $\mathcal{O}(\mu_B^6)$ in the baryon chemical potential on lattices with temporal extent up to $N_\tau = 12$. This will greatly advance our current knowledge, which at present is limited to coarse lattices [16, 18] or leading order corrections only [19]. Furthermore, the inclusion of a physical charm quark mass [20] will be of importance for the physics of the early universe. We will also need to further analyze the influence of large electric and magnetic fields on bulk thermodynamics and the transition temperature, [21] which has implications for phase transitions in the early universe as well as heavy ion collisions [22].

B. The critical endpoint

Whether or not a true second order phase transition point, the critical endpoint [23], exists in the QCD phase diagram at nonzero values of the baryon chemical potential is one of the most exciting questions in current studies of QCD thermodynamics. At RHIC the beam energy scan (BES) program is devoted to this question. Much of this experimental program is motivated by lattice QCD studies that showed the sensitivity of fluctuations of conserved charges and their higher order cumulants to changes in temperature and baryon chemical potential [24, 25]. In particular, ratios of higher order cumulants of conserved charge fluctuations are interesting as they directly reflect the most relevant degrees of freedom that are carriers of these charges in different temperature regimes [26] and, moreover, are also accessible in heavy ion experiments [27]. The BES at RHIC produced first results on fluctuations of net proton numbers and net electric charge [28–30] that now can be compared with lattice QCD calculations [31, 32]. In Fig. 1 we show a schematic plot of the QCD phase diagram including a critical point and the parameter range of the experimental program at RHIC that is currently pursued in search for hints for the existence and location of the critical point.

Lattice QCD calculations at nonzero values of the chemical potentials are difficult, since at nonzero net baryon number density the fermion determinant becomes complex, resulting in the so-called fermion sign problem. Our standard Monte-Carlo simulation techniques rely on the existence of a positive integration kernel that can be interpreted as a probability density. But those techniques can not be applied. Various approaches have been exploited to avoid the limitations in finite density lattice QCD calculations that are caused by the sign problem, at least partially. For small values of the chemical potential, the regime of small chemical potential can be reached with a systematic Taylor series expansion. The expansion coefficients are generalized susceptibilities constructed from moments of net baryon number, electric charge and strangeness fluctuations. The expansion coefficients in the Taylor series are evaluated at vanishing values of the chemical potential, but can eventually be used to estimate the radius of convergence of the Taylor series. This method provides an estimate for the location of a possible critical point, which can be systematically improved when higher order expansion coefficients become accessible. At present such estimates [33] are at best indicative; they are based on calculations performed on coarse lattices with unimproved staggered fermion actions and unphysical values of the light quark masses. Continuum

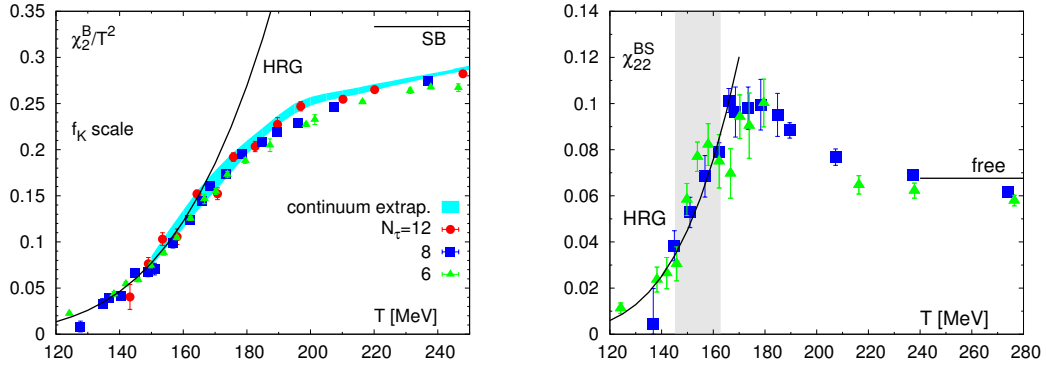


FIG. 4. Quadratic fluctuations of net baryon number (left) and the correlation between quadratic fluctuations of net baryon number and strangeness (right). The latter also is part of the next-to-leading order μ_B^2 -correction to net strangeness fluctuations at nonzero baryon chemical potential and is similar to the B-S correlation discussed by V. Koch *et al.* as a signature for QGP formation [27]. The gray band gives the location of the chiral transition temperature, $T_c = 154(9)$ MeV.

extrapolated results at (almost) physical values of the light quark masses and a physical value of the strange quark mass at present exist only for quadratic fluctuations of conserved charges [34, 35]. Some preliminary results for fourth order cumulants of charge fluctuations have been presented only recently [32, 36]. We show some results for quadratic and quartic fluctuations of net baryon number and electric charge in Fig. 4.

A second generation of lattice QCD calculations of higher order cumulants of baryon number and electric charge fluctuations is currently underway. Although these new calculations use the staggered fermion discretization schemes with strongly reduced cutoff effects, *i.e.*, the HISQ action, these calculations are far from final. Calculations closer to the continuum limit and with physical values of the quark masses are still needed. Furthermore, in order to make direct comparisons between lattice QCD calculations and experimental measurements we also need to take into account nonzero values of the strangeness and electric charge chemical potentials [31]. This will allow us to adjust the theoretical calculations to conditions met in heavy ion collision experiments, namely strangeness neutrality and the isospin asymmetry of the colliding nuclei. In particular, the analysis of electric charge fluctuations, which at low temperature as well as in the transition region receive large contributions from light pions, requires calculations on large lattices, close to the continuum limit, in order to suppress the influence of cutoff effects that result from taste symmetry violations in the staggered discretization scheme.

Taylor expansions can also be used to calculate charge fluctuations at nonzero values of the chemical potential. Complete next-to-leading order calculations of up to fourth order cumulants of conserved charge fluctuations (*i.e.*, mean, variance, skewness and kurtosis), that are currently being measured in the BES at RHIC, require knowledge of up to sixth order susceptibilities. This is sufficient to control properties of cumulants in the range of baryon chemical potentials $\mu_B \lesssim 200$ MeV [31]. It allows us to compare with the experimental results for beam energies $\sqrt{s_{NN}} \gtrsim 20$ GeV. A calculation of eighth order susceptibilities will be necessary to confirm the robustness of next-to-leading order results and to extend the range of validity of lattice QCD calculations to larger values of μ_B/T ,

corresponding to the smaller beam energies that are already accessible to the BES at RHIC and that will be studied in even more detail at future heavy ion facilities. Gaining control over eighth order susceptibilities will also help in estimating the radius of convergence related to the location of the critical point.

Calculations of higher order cumulants of net charge fluctuations typically require generation of $O(10^5)$ gauge-field configurations. On a sub-sample of $\mathcal{O}(10^4)$ configurations, higher order susceptibilities are calculated through multiple inversions of the fermion Dirac matrix using $\mathcal{O}(10^3)$ random source vectors. The calculation of higher order cumulants at a single value of the temperature thus typically involves several million matrix inversions. The number of random vectors plus gauge-field configurations needed to reach comparable statistical errors in different n^{th} order cumulants grows exponentially. Current calculations with the HISQ action suggest that it grows like $10^{n/2}$. Moreover, the computational effort per set of random vectors increases approximately as 1.5^n . The overall computational effort for going from a sixth order calculation to eighth order, keeping relative errors constant, thus increases roughly by two orders of magnitude. This shows that in addition to the likely hardware improvements, resulting in an order of magnitude increase in computational capabilities over the next five years, substantial algorithmic improvement is also required.

While it is of interest to explore the generic structure of higher order cumulants in a large temperature range, making contact with perturbation theory at high temperature [37–39] and the HRG model calculations at low temperatures [26], it clearly is most relevant to obtain high quality results in a rather narrow temperature range in which freeze-out is expected to occur in the ongoing heavy ion collision experiments, $150 \text{ MeV} \lesssim T \lesssim 170 \text{ MeV}$. In this temperature range calculations of cumulants up to eighth order should be performed in the future. This will require algorithmic improvements as well as increases in hardware resources. Before going to eighth order the current program on the calculation of sixth order coefficients needs to be completed. Continuum extrapolated results, using three lattice spacings corresponding to temporal extents $N_\tau = 8, 12$ and 16 , for susceptibilities should be obtained for at least five temperature values within this relevant temperature range using the HISQ action at physical values of the light and the strange quark masses. Such calculations will need substantial computing resources for the generation of gauge-field configurations as well as the calculation of Taylor expansion coefficients. In particular the latter is nowadays done very efficiently on general purpose graphics processing units (GPUs). State-of-the-art calculations of higher order cumulants continuously make use of $\mathcal{O}(500)$ GPUs. In order to advance to the calculation of eighth order cumulants one will need to exploit computing platforms with several thousand GPUs.

C. The properties of the quark-gluon plasma

Experiments at the Relativistic Heavy Ion Collider, Brookhaven National Laboratory, have revealed that even at temperatures moderately above the QCD transition temperature, the quark-gluon plasma exhibits far more complex properties than that expected for a weakly interacting thermal medium described by perturbative QCD [2, 3]. In particular, the QGP possesses remnant confining features, *i.e.*, some hadronic bound states involving heavy quarks continue to exist in this regime. Furthermore, these highly non-perturbative

properties of the QGP also results in intriguing transport properties leading to many unexpected experimental observations: the smallness of the ratio of shear viscosity and entropy leads to (almost) perfect fluidity of the QGP. Understanding the properties of this genuinely nonperturbative regime in the high temperature phase of QCD is a theoretical challenge that requires a genuine nonperturbative approach, namely realistic lattice QCD calculations performed close to the continuum limit using physical values of the dynamical quark masses.

Due to the smallness of the electromagnetic coupling and the limited volume of the hot medium generated in a heavy ion collision, dileptons and photons escape the medium without any subsequent interaction. They thus carry information not only about the temperature close to the hadronization transition but also about the early stages of the evolution. Studies of electromagnetic probes, such as the dilepton and photon emission rates, have produced many fascinating experimental results concerning the properties of the quark-gluon plasma close to the QCD transition. Particularly, the PHENIX collaboration has recently observed a large enhancement of low mass dilepton emission [40] over conventional hadronic sources, which, however, does not seem to be present in the data reported by the STAR collaboration [41]. Furthermore, experimental results also indicate that low energy photons may be emitted from a thermal medium having temperatures moderately above the QCD transition temperature [40, 42], and these photons also exhibit large collective flow [42, 43]. Theoretical studies indicate [44] that an explanation for these experimental results may lie in the enhancement of photon emission rates for temperatures just above the QCD transition temperature, calling for first-principle-based nonperturbative calculations of the photon/dileptons emissivity in the vicinity of the transition region.

Heavy quarks, such as charm and bottom, are another unique set of probes for properties of the quark-gluon plasma. Owing to their large masses, they cannot be produced thermally inside the QGP: they are produced at the primordial stages having an initial energy spectrum far from the thermal spectrum. Experimental observations show surprisingly large collective flow and energy loss of the charm quarks [45, 46], indicating rapid thermalization of charm quarks inside the QGP. Theoretical studies [47, 48] suggest that such a large collective flow and energy loss of charm quarks can be understood if the charm quark diffusion constant is almost an order of magnitude smaller than the perturbative estimates [49]. Thus, these puzzling experimental facts can only be explained by performing nonperturbative, first-principle QCD calculations of the charm diffusion constant in the QGP.

Lattice calculations of heavy quark bound states have advanced considerably in recent years. Finite temperature calculations in quenched QCD are now possible on very large lattices. The largest lattices exploited reach sizes of $128^3 \times 96$ [50], which is comparable to lattice sizes used today in zero temperature hadron spectroscopy, and have discretization constants (lattice spacings) as small as 0.01 fm. Such large lattices with a high resolution allow us to study the thermal heavy quark correlation function and provide enough information on the Euclidean time dependence of these correlators to perform a statistical analysis based on the maximum entropy method (MEM) [51]. The MEM analysis provides information on spectral functions as well as transport coefficients. Some recent results for the J/ψ spectral function are shown in Fig. 5 (left). They lead to the conclusion that all charmonium states have melted at $1.5T_c$. First attempts to generalize these calculations to properties of bound states at nonzero momenta [52] are underway, and also first attempts to study heavy quark bound states, including the bottomonium system, in QCD with a dynamical light quark sector have been started [53]. The latter, however, are at present limited to studies on rather




FIG. 5. The spectral function of the charmonium ground state at different values of the temperature (left) [50] and the thermal production rate for dileptons with total vanishing momentum that arises from quark anti-quark annihilation (right) in a thermal heat bath [54]. Both figures are based on calculations performed in quenched QCD.

small lattices. In order to have sufficient information on the Euclidean time dependence of the correlation functions such calculations are performed on highly anisotropic lattices.

Thermal gauge-field configurations on lattices with large temporal extent also permit an analysis of spectral functions in the light quark sector. Aside from showing explicitly that light quark bound states in different quantum number channels melt in the high temperature phase, the vector spectral function is of particular interest. It provides direct information about the production rate for dilepton pairs arising from quark anti-quark annihilation in a hot thermal medium. We show results from a recent calculation on quenched QCD lattices of sizes up to $128^3 \times 48$ [54] in Fig. 5(right). At a temperature $T \simeq 1.5T_c$ these continuum extrapolated results are in good agreement with resummed (hard thermal loop) perturbation theory. An extension of these calculations to nonzero momentum is underway and will also give access to thermal photon rates. Of course, in the light quark sector it will be even more important than for heavy quarks to include contributions from dynamical light quarks in these calculations.

A problem closely related to the analysis of light and heavy quark bound state properties is the calculation of light and heavy quark transport properties, *e.g.*, the electrical conductivity and the heavy quark diffusion constant. This requires gaining control over the low frequency region of spectral functions, which is even more sensitive to the large distance behavior of spectral functions and thus requires large lattices and good control over statistical errors. Usually, one needs an ansatz for the functional form of the spectral function in this regime ($\rho(\omega) \sim \omega$) to extract transport coefficients. First results on the electrical conductivity [54] and the heavy quark diffusion constant [50] heavy been obtained in connection with the recent studies of light and heavy quark correlation functions on large quenched lattices.

Lattice QCD studies of thermal hadron spectra, transport coefficients, such as the electrical conductivity and charm diffusion constants as well as the dilepton/photon emissivity all require knowledge of the spectral functions obtained from the analysis of various imaginary time correlation functions [55]. Although this is an ill-posed problem, statistical techniques based on the maximum entropy method have been applied quite successfully to such prob-

lems in the context of finite temperature QCD [50, 51]. In order to do so one, however, requires numerical results on rather large lattices with little statistical noise. For this reason most results on thermal masses and transport coefficients have so far been obtained in quenched QCD. Extending these calculations to QCD with dynamical light quark degrees of freedom on similar size lattices will be an important task during the next years.

The largest influence of dynamical quark degrees of freedom will become visible in the transition region. A calculation of spectral functions thus can focus on the temperature range $T/T_c \simeq (1 - 1.5)$. For three to four temperature values one should analyze correlation functions at two values of the cutoff that correspond to lattices with temporal extent $N_\tau = 32$ and 48, respectively. A third data set for $N_\tau = 16$ will be available from the finite density QCD calculations discussed in section II.B. Eventually one wants to extend such an analysis to even larger temporal lattices, $N_\tau \simeq 64$.

The transport properties related to light and heavy quark diffusion that we discussed so far are quite distinct from the shear and bulk viscosities that enter viscous hydrodynamics and play a central role in the interpretation of flow properties of hadrons observed in heavy ion collisions. The calculation of shear and bulk viscosities on the lattice [55–57] requires the analysis of observables involving the gluonic energy-momentum tensor. Unlike the hadronic correlation functions discussed so far, the gluonic correlators are much more noisy and require much higher statistics. While in the former case large computing resources are required to invert fermion matrices on a large set of gauge-field configurations, in the latter case these resources are needed to generate an even larger set of gauge-field configurations. The calculation of gluonic correlation functions itself, however, is simple and their analysis makes use of the same statistical tools (MEM) also applied to the analysis of hadronic correlation functions. To make progress with calculations of shear and bulk viscosities one thus will perform calculations at one or two temperature values at which the number of gauge-field configurations needs to be increased by an order of magnitude.

D. The chiral limit

Many fundamental features of the QCD phase diagram can be understood in terms of global symmetries of the QCD Lagrangian that are either anomalously (axial symmetry) or spontaneously (chiral symmetry) broken. First principle nonperturbative studies of the fundamental symmetries of QCD not only enhance our qualitative and quantitative knowledge of the phase structure of QCD, but also improve our understanding of the mechanisms that lead to phase transitions in quantum field theories in general.

The spontaneous breaking of the global chiral symmetry induces a singular term in the QCD partition function that increasingly dominates the thermodynamic observables at smaller values of the quark masses. In order to quantify the influence of the chiral symmetry at the physical values of the quark masses, one needs to extend calculations to lower quark masses closer to the singular point. During recent years calculations at and below the physical quark mass values have become possible with improved staggered fermion actions. For the first time this has allowed a convincing demonstration of the existence of an underlying universal structure in thermodynamic functions that is consistent with the scaling properties expected for a phase transition in the three-dimensional $O(4)$ universality class [7].

The exploitation of these universal features permits a well defined characterization of the crossover transition at nonvanishing light quark masses. It leads to a determination of the dependence of the QCD transition temperature at vanishing values of the chemical potential [8].




FIG. 6. Scaling behavior of the normalized chiral condensate, $M = m_s \langle \bar{\psi}\psi \rangle / T^4$, for different values of the light-to-strange quark mass ratio m_l/m_s on coarse lattices with temporal extent $N_\tau = 4$ [7].

of the QCD transition, leading to possible experimentally observable consequences for net charge fluctuations measured in the BES at RHIC. Although the results obtained with improved staggered fermion actions are quite promising, so far studies of universal properties that require calculations with less-than-physical quark masses have been restricted to coarse lattices due to large computational costs. These calculations need to be validated through the use of the highly improved staggered quark action and supplemented by studies closer to the continuum limit. Similar to the calculations discussed in section II.B. these calculations will rely on the generation of gauge-field configurations on Blue Gene/Q type machines and extensive use of GPUs for data analysis. For the scaling analysis one will consider chiral susceptibilities and baryon number susceptibilities up to sixth order, as these are the first to diverge in the chiral limit. These data will also allow us to calculate the $\mathcal{O}(\mu_B^4)$ corrections to the chiral phase transition line as a function of the baryon chemical potential as well as the strangeness chemical potential. They will provide solid estimates of the systematic errors induced by truncating a Taylor series expansion and will improve the reliability of estimates for the range of validity of the truncated series for the phase transition temperature. As such we will be able to strengthen the comparison between freeze-out temperatures measured in heavy ion experiments and the transition temperature calculated in QCD.

So far we have discussed universality in the restoration of $SU(2) \times SU(2)$ chiral symmetry. The $U(1)$ axial symmetry of the QCD Lagrangian, on the other hand, is anomalously broken. Therefore, it does not generate a singular term in the QCD partition function even in the limit of vanishing quark masses. However, as noted already in the seminal work on the phase structure of QCD in the chiral limit [59], qualitative changes in the QCD phase structure can occur if the anomalous breaking of the axial symmetry becomes sufficiently weak. For instance, it has been argued that the QCD transition may become first order if the amount of axial symmetry breaking becomes small enough through the suppression of topological charges due to the color screening induced at high temperatures [59]. The axial anomaly also plays a crucial role in defining the qualitative aspects of the phase structure of dense QCD, specially on the existence/location of the QCD chiral critical point




FIG. 7. The transformation properties of various scalar and pseudo-scalar hadrons under chiral rotations and the manifestation of these symmetry transformations in corresponding susceptibilities.

[60]. The breaking at low temperatures and the restoration of some of the chiral symmetries at high temperatures has an impact on the hadron spectrum in QCD in the scalar as well as vector channels and can be probed in lattice QCD calculations either by calculating directly thermal hadron masses or the corresponding susceptibilities that are obtained by integrating correlation functions with specific hadronic quantum numbers over space-time [61]. This is illustrated schematically in Fig. 7.

Beside the phenomenological consequences that result from the temperature dependence of the axial symmetry breaking there remain many unanswered deep theoretical questions regarding the axial anomaly in QCD under extreme conditions; particularly the mechanism of axial anomaly breaking after the chiral symmetry restoration and its manifestation in the Dirac eigenvalue spectrum. To address these intriguing phenomenological as well as theoretical issues via nonperturbative lattice QCD, it is obviously preferable to utilize a discretization scheme that preserves exact chiral symmetry and reproduces the correct axial anomaly even at nonvanishing values of the lattice spacing. The domain wall fermion (DWF) formulation is such a discretization scheme ideally suited for lattice QCD studies related to the chiral and anomalous properties of QCD. Although in the past one had to refrain from using DWF due to its prohibitive computational costs, over the past few years owing to tremendous algorithmic advancements coupled with the advent of increasing computing power, it is becoming increasingly possible to perform realistic lattice QCD simulations utilizing the DWF formalism. First results on the QCD transition [12] and studies related to the mechanism of the axial symmetry breaking at finite temperature [62] are indeed encouraging. These exploratory studies indicate that the relation between low-lying eigenvalues of the Dirac operator, enforced by the exact chiral symmetry of DWF, and the topological structure of gauge-field configurations becomes evident. In the future these DWF studies need to be extended toward the chiral limit by utilizing smaller quark masses for detailed explorations of the chiral symmetry and the axial anomaly in QCD. In such studies a detailed comparison between the distribution of low-lying eigenvalues of the Dirac operator and the topological structure of gauge-field configurations will be performed. Results can also be compared with the temperature dependence of various susceptibilities that are sensitive to the restoration or breaking of chiral symmetries.

The focus of calculations with DWF will be on exploring the consequences of exact chi-

ral symmetries at finite values of the lattice cutoff rather than trying to perform also the continuum limit. Most calculations thus can be performed on five dimensional lattices with temporal extent $N_\tau = 8$. Also, in these calculations it will be important to use smaller-than-physical quark mass values in order to control the approach to the chiral limit. In order to calculate the QCD transition temperature within the DWF formalism, additional calculations on lattices with large temporal extent will be necessary. Such calculations obviously are in need of large computational resources on leadership class computers.

III. OUTLOOK

A. Scientific challenges

Calculations up to the present have yielded reliable results for the QCD transition temperature at vanishing baryon chemical potential. We have some understanding of the dependence of the transition temperature on small values of the baryon chemical potential and made significant progress in calculating the QCD equation of state and related bulk thermodynamic properties.

Completing these calculations at physical values of the quark masses, and establishing the properties of strong interaction matter at vanishing net baryon number density in the chiral limit, making use also of calculations performed with chiral fermions (Domain Wall Fermions) will *define the anchor point for all studies of the QCD phase diagram as a function of temperature and net baryon number density. It will put an indisputable lower bound on the temperature at which hadron matter transforms into a QGP and will establish a reliable starting point for extensions of these calculations into the regime of nonvanishing baryon number density. In combination with calculations using values of light and heavy quark masses as realized in nature, this will quantify the role of chiral symmetry breaking and confinement in the thermodynamics of strong interaction matter. The equation of state will be the basic equilibrium input to a microscopic description of the rapidly expanding and cooling dense matter formed in a heavy ion collision* [1].

Calculations at nonzero baryon chemical potential have already had a strong influence on the experimental beam energy scan program pursued at RHIC to search for evidence of the postulated critical endpoint in the QCD phase diagram. Refining the existing calculations of quadratic and quartic cumulants of net baryon number, electric charge and strangeness cumulants and extending these calculations to even higher orders will help provide estimates for the location of a critical point or extend to larger values the excluded range of baryon chemical potentials, which today reaches up to $\mu_B \simeq 200$ MeV.

Calculations of spectral functions and the application of the maximum entropy method used to extract information on the fate of light and heavy quark bound states in the quark-gluon plasma as well as transport coefficients are highly advanced in pure gauge theory calculations. Including the influence of dynamical light quark degrees of freedom will be a major challenge during the next years.

Calculations of transport coefficients will provide fundamental insight into the structure of hot and dense matter. It will allow us to quantify aspects of the extent to which the phenomenologically successful modeling of heavy ion collisions has a solid foundation in QCD; i.e., whether a near-equilibrium QGP described by QCD indeed equilibrates rapidly and can be characterized as an almost-perfect fluid. Detailed information on the spectral function would confirm whether or not the QGP is strongly coupled at RHIC, and by varying the temperature in the lattice QCD calculations, scientists will learn how much the temperature has to be increased before the plasma becomes weakly coupled. This question will be of central importance in comparing the heavy-ion data obtained at the RHIC and the LHC experiments because the temperature in the latter will be about a factor 1.5 to 2 higher [1].

B. Computational challenges

All of the ongoing and future research projects outlined in the previous sections are in need of large computational resources. At present, members of USQCD are making use of dedicated hardware funded by the DOE through the LQCD-ext Computing Project, as well as a Cray XE/XK computer, and IBM Blue Gene/Q and Blue Gene/P computers, made available by the DOE's INCITE Program. During 2013, USQCD, as a whole, expects to sustain approximately 300 Tflop/s on these machines. Subgroups within USQCD also make use of computing facilities at the DOE's National Energy Research Scientific Computing Center (NERSC), the Lawrence Livermore National Laboratory (LLNL), and centers supported by the NSF's XSEDE Program. The BNL group furthermore has access to substantial additional computing capability on the Blue Gene/Q at BNL, the Blue Gene/Q prototypes of the RIKEN-BNL research center as well as through international collaborations. The Bielefeld-BNL Collaboration has grants at the European PRACE centers and utilizes a 400 GPU-cluster in Bielefeld.

For some time, the resources we have obtained have grown with a doubling time of approximately 1.5 years, consistent with Moore's law, and this growth rate will need to continue if we are to meet our scientific objectives. The software developed by USQCD under our Nuclear Physics SciDAC grant enables us to use a wide variety of architectures with very high efficiency, and it is critical that our software efforts continue at their current pace. Over time, the development of new algorithms has had at least as important an impact on our field as advances in hardware, and we expect this trend to continue, although the rate of algorithmic advances is not as smooth or easy to predict as that of hardware.

Almost all of the calculations outlined in the previous section require still more numerous or larger computational resources and further code development to exploit them. As the chiral critical point is approached, the computational expense grows rapidly both because the fermion matrices in the calculation become more ill-conditioned and because the physical size of the lattice must be increased to avoid distortions from finite volume effects. Multi-GPU architectures and resources with many thousands of GPUs or new multi-core CPU architectures are needed to carry out such calculations. Typically the time requirement for such calculations is dominated by the analysis part as many thousand matrix inversions with different source vectors on the same gauge-field configurations are needed. The number

of floating point operations, expressed in terms of Teraflop/s-years, needed to analyze the approach to the chiral limit with pion masses about half of their physical value are estimated in Table I.

T -values	m_s/m_l	lattice size	configuration generation [TFlop/s-years]	configuration analysis [TFlop/s-years]
5	40	$32^3 \times 6$	20	70
5	60	$40^3 \times 6$	49	205
5	80	$48^3 \times 6$	112	480

TABLE I. Requirements for the generation of 10^5 gauge-field configurations used for a scaling analysis at less-than-physical values of the quark masses at five values of the temperature close to the transition temperature in the chiral limit, *i.e.*, temperatures between 140-160 MeV. Also given is the number of floating point operations needed to calculate up to sixth order cumulants on a sample of 10^4 configurations. Repeating these calculations at smaller lattice spacings ($N_\tau = 8$) to control cutoff effects in the continuum limit would require a factor five larger computing resources. Based on today’s state-of-the-art software and hardware at leadership class facilities, this project would require access to 1.5 Blue Gene/Q racks and 5000 GPUs for one year to be completed.

The calculation of higher order cumulants is a huge “capacity” problem. Many thousands of GPUs are needed to accomplish the work, although each calculation may require only a small number of them. Algorithmic development is also needed. Since such calculations require repeatedly solving the Dirac equation on the same gauge-field configuration with many thousands of different sources, deflation algorithms based on removing low-lying eigenmodes should help reduce the computational cost considerably [63]. In Table II we summarize the number of floating point operations needed to analyze up to eighth order cumulants of conserved charges and to perform the continuum limit in a narrow temperature interval relevant for the comparison with heavy ion experiments.

T -values	lattice size	configuration generation [TFlop/s-years]	configuration analysis [TFlop/s-years]
5	$32^3 \times 8$	21	630
5	$48^3 \times 12$	49	1990
5	$64^3 \times 16$	144	1010 (sixth order)

TABLE II. Requirements for the generation of 10^5 gauge-field configurations used for the calculation of conserved charge fluctuations up to eighth order with physical values of the quark masses at five values of the temperature between 150-170 MeV and for three values of the lattice spacings. Also given are the resources needed to calculate up to eighth order cumulants on a sample of 10^4 configurations using $2 \cdot 10^3$ random source vectors for the inversions of fermion matrices. Limiting the analysis on the $N_\tau = 12$ lattices to sixth order only would reduce the resource requirement by a factor 2.

The calculation of thermal masses and transport properties requires lattice with large temporal extent, which, in turn, requires large multi-GPU architectures, such as Titan at Oak

Ridge National Laboratory or an equivalent cluster based on Intel MIC processors. Of course, appropriate software for exploiting such heterogeneous computing resources must be developed. Multigrid [64] and GPU-oriented domain-decomposition algorithms [65] should help reduce the computational cost. In Table III we summarize the resources needed to perform calculations of transport coefficients obtained from hadronic correlation functions on an isotropic lattice of size $N_\sigma^3 \times N_\tau$. Requirements for the calculation of shear and bulk viscosities are not tabulated but will require about an order of magnitude more resources.

T -values	lattice size	[TFlop/s-years]
3	$128^3 \times 32$	60
3	$192^3 \times 48$	315
3	$256^3 \times 64$	1150

TABLE III. Requirements for the generation and analysis of 10^4 gauge-field configurations and the determination of spectral functions of light and heavy thermal correlation functions. The fraction required for analysis is of the order of 1% only. Using anisotropic lattices with spatial extent of $N_\sigma = 128$ and temporal extents $N_\tau = 32 - 96$ for this project can reduce the required resources by a factor 3.

All calculations described above and the computational needs summarized in the tables can be performed with the highly improved staggered fermion action. For the study of thermal masses and transport coefficients one will also perform some of the calculations using Wilson fermions. As outlined in the USQCD white paper on *Lattice QCD for Cold Nuclear Physics* such calculations on lattices sizes listed in Table III are feasible and require about a factor four more computing resources.

For the analysis of chiral properties of QCD, in particular physics related to the breaking of the axial anomaly, one also wants to perform calculations with domain wall fermions. Such calculations also require large multiprocessor leadership-class architectures, such as the Argonne Blue Gene/Q. Due to the large computational needs for calculations with domain wall fermions, one would, in the next 5 years, not attempt to analyze charge fluctuations but would focus on an analysis of chiral susceptibilities and topology. The requirements for such calculations are summarized in Table IV.

T -values	m_s/m_l	lattice size	configuration generation [TFlop/s-years]
5	40	$48^3 \times 8$	145
5	60	$64^3 \times 8$	290
5	80	$96^3 \times 8$	1150

TABLE IV. Requirements for the generation of 10^4 gauge-field configurations used for a scaling analysis with domain wall fermions at less-than-physical values of the quark masses at five values of the temperature close to the transition temperature in the chiral limit, *i.e.*, temperatures between 140-160 MeV. Calculations on lattices of size $32^3 \times 8$ at the physical quark mass ratio $m_s/m_l \simeq 27$ at a larger set of temperature values are currently being carried out by HotQCD.

C. Computational priorities

The projects outlined in the previous subsection are in need of substantial computing resources. Not all of them can be started right away. We have indicated how they can be adjusted to actually available resources by performing them in smaller steps. We prioritize the computational projects according to the potential payoff for heavy ion and early universe phenomenology and for our theoretical understanding of field theory in general:

- **Equation of state including charm quarks.** Calculations already underway are aimed at resolving disagreements over the equation of state in the absence of charm quarks. Also calculations are underway treating charm quarks in thermal equilibrium. Methods for dealing with out-of-equilibrium charm will soon be developed.
- **Higher order cumulants of conserved charges.** These calculations have direct bearing on experimental measurements and the search for the critical endpoint. Calculation of lower-order cumulants is currently underway. With the expected increase in computing power over the next several years, the next order or two can be done.
- **Study of the chiral critical point with lighter-than-physical up and down quarks.** This study should use both the highly improved staggered quark formulation and the domain wall fermion formulation. These calculations simply await more computing power.
- **Calculation of transport coefficients and of the viability of charmonium and bottomonium with quarks included in the statistical ensemble.** The first goal with more computing power will be to confirm the reliability of the methods, particularly with light quarks included in the statistical ensemble.

IV. SUMMARY OF ACCOMPLISHMENTS AND CHALLENGES

In this section we provide an executive summary of the recent scientific accomplishments and challenges described in more detail above.

A. Accomplishments

- The QCD transition and equation of state
 - **Transition temperature at vanishing baryon number density.** There is a reasonably good consensus on the transition temperature (to the extent a crossover temperature is meaningful). Results compare roughly to the phenomenologically determined freeze-out temperature from experiments with heavy ion collisions.

- **Transition temperature at nonvanishing baryon number density.** The curvature of the transition temperature as a function of small chemical potential is known to leading order in the square of the baryon number chemical potential. Significant differences with freeze-out temperature start to show up at baryon number chemical potentials larger than $\mu_B \simeq 200$ MeV, but more work is needed to make a reliable comparison.
 - **The equation of state with up, down, and strange quarks.** Our knowledge of the temperature dependence of the energy density, pressure, and entropy has vastly improved, but a consensus on the extrapolation to zero lattice spacing is lacking and needs further work to settle the issues.
 - **Equation of state including charm quarks.** We have early results for the contribution of charm quarks to the equation of state, but more work is needed from experiment, phenomenology, and numerical simulation to understand the role of possibly out-of-equilibrium charm quarks in the thermodynamics of heavy ion collisions and the importance of the charm contribution in cosmological applications.
- The critical endpoint
 - **Computational evidence for the critical endpoint.** The evidence is not solid. We still lack a good computational method for simulating directly at nonvanishing chemical potential where a critical chiral endpoint is supposed to occur. Reliable, but indirect methods use a Taylor series expansion at vanishing chemical potential. The calculational effort grows exponentially with the order of the expansion. Further terms can help, but require significantly more computation.
 - **Fluctuations in conserved charges.** The coefficients of the Taylor series also determine fluctuation cumulants in conserved charges. They are measured in experiment and may give indications of a nearby critical endpoint. Second order cumulants are reasonably well computed by now. More refined calculation of higher order cumulants as well as cumulants at nonzero baryon number chemical potentials is needed for comparison with experimental results.
 - The properties of the quark-gluon plasma
 - **Survival of heavy quarkonium.** It has been reasonably well established that charmonium and bottomonium survive at temperatures at least slightly above the transition. There is still controversy about how far above the transition temperature one can go before those states melt. Results have direct bearing on dilepton emission. More solid results that also take into account dynamical light quark contributions will come with increased computing power.
 - **Heavy quark diffusion.** There are first exploratory results for heavy quark diffusion constants that are substantially smaller than perturbative results, but are in reasonable agreement with experimental findings that suggest an efficient thermalization of heavy quarks. So far all lattice calculations neglect contributions from dynamical light quarks. This will come with increased computing power.

- **Shear and bulk viscosity.** The electrical conductivity and shear and bulk viscosities are very important for phenomenology. Obtaining reliable results is computationally demanding as one has to deal with the gluonic part of the energy-momentum tensor, which is strongly fluctuating. Results from calculations in a statistical ensemble with only gluons suggest that the ratio of shear viscosity to entropy density is very low, but these results need refinement and our knowledge of transport with both gluons and quarks in the statistical ensemble is even more meager.
- **Electrical conductivity and thermal dilepton/photon rates.** Results on the thermal dilepton rate at vanishing momentum of the dilepton pair have been obtained in the absence of dynamical quark degrees of freedom. These calculations agree fairly well with hard thermal loop perturbation theory. Unlike the latter they, however, also lead to a finite electrical conductivity. An extension of these calculations to nonzero momenta is needed to get access to thermal photon rates. In the absence of quark degrees of freedom such calculations are possible now and will soon be available. In all cases, however, it will be necessary to include dynamical quark degrees of freedom in the calculations.
- The chiral limit
 - **Critical universality.** We have first indications that, close to the transition, the light quark chiral condensate and its susceptibility exhibit scaling properties consistent with the expected $O(4)$ universality class. Further study is needed to confirm this, particularly with fermion formulations with good chiral properties.
 - **Restoration of the axial $U(1)$ symmetry.** Indications to date suggest that this symmetry is not restored at the transition. These results need solid confirmation, since they could have direct bearing on the order of the transition.

B. Challenges

Clearly, we made substantial progress in understanding the thermodynamics of strong interaction matter, its phase structure and the role of exact symmetries of the underlying quantum field theory, quantum chromodynamics. Nonetheless there are still a number of challenging open questions that need to be answered. Aside from the analysis of the equation of state, which is on its way and soon will provide a solid basic input into the hydrodynamic modelling of the equation of state, reaching better understanding of the phase structure at nonzero chemical potential is most urgent. The calculations envisaged for the coming years will at least settle the question whether the existing estimates for a location of a critical point, which are based on lattice calculations performed on rather coarse lattices and thus suffer from large discretization errors, can be confirmed and survive the continuum limit. Furthermore, calculations with less-than-physical quark masses will allow us to establish firm results on the QCD phase diagram in the chiral limit. This may well be the only limit, in which strong interaction matter has a true phase transition.

The new generation of computers installed at leadership class computing facilities in Argonne and Oak Ridge National Laboratory as well as Lawrence Livermore National Laboratory in

combination with the resources operated by USQCD at National Laboratories offer many opportunities to bring us closer to answers for these challenging questions. Nonetheless, we are also in need of substantial resources on future even more advanced computing platforms.

Acknowledgements: We gratefully acknowledge contributions to this report from Alexei Bazavov, Peter Petreczky and Swagato Mukherjee, as well as valuable comments and suggestions made by Volker Koch, Krishna Rajagopal, Ralf Rapp and Nu Xu.

-
- [1] Scientific Grand Challenges, *Forefront Questions in Nuclear Science and the role of high performance computing*, U.S. Department of Energy 2009.
 - [2] RHIC white papers, Nucl. Phys. A **757**, 1 (2005).
 - [3] B.V. Jacak and B. Müller, Science **337**, 310 (2012).
 - [4] C. DeTar and U. M. Heller, Eur. Phys. J. A **41**, 405 (2009).
 - [5] M. M. Aggarwal *et al.* [STAR Collaboration], arXiv:1007.2613 [nucl-ex].
 - [6] G. Odyniec, Acta Phys. Polon. B **43**, 627 (2012).
 - [7] S. Ejiri, *et al.*, Phys. Rev. D **80**, 094505 (2009) [arXiv:0909.5122 [hep-lat]].
 - [8] A. Bazavov, T. Bhattacharya, M. Cheng, C. DeTar, H. T. Ding, S. Gottlieb, R. Gupta and P. Hegde *et al.*, Phys. Rev. D **85**, 054503 (2012) [arXiv:1111.1710 [hep-lat]].
 - [9] Y. Aoki, S. Borsanyi, S. Durr, Z. Fodor, S. D. Katz, S. Krieg and K. K. Szabo, JHEP **0906**, 088 (2009)
 - [10] M. M. Aggarwal *et al.* [STAR Collaboration], Phys. Rev. C **83**, 034910 (2011).
 - [11] O. Kaczmarek, *et al.*, Phys. Rev. D **83**, 014504 (2011) [arXiv:1011.3130 [hep-lat]].
 - [12] M. Cheng, *et al.*, Phys. Rev. D **81**, 054510 (2010) [arXiv:0911.3450 [hep-lat]].
 - [13] A. Bazavov *et al.* [HotQCD Collaboration], Phys. Rev. D **86**, 094503 (2012) [arXiv:1205.3535 [hep-lat]].
 - [14] Y. Aoki, Z. Fodor, S. D. Katz and K. K. Szabo, JHEP **0601**, 089 (2006).
 - [15] A. Bazavov [HotQCD Collaboration], Phys. Rev. D **80**, 014504 (2009) [arXiv:0903.4379 [hep-lat]].
 - [16] S. Ejiri, F. Karsch, E. Laermann and C. Schmidt, Phys. Rev. D **73**, 054506 (2006) [hep-lat/0512040].
 - [17] A. Bazavov [HotQCD Collaboration], arXiv:1210.6312 [hep-lat].
 - [18] C. Bernard, C. E. DeTar, L. Levkova, S. Gottlieb, U. M. Heller, J. E. Hetrick, R. Sugar and D. Toussaint, Phys. Rev. D **77**, 014503 (2008) [arXiv:0710.1330 [hep-lat]].
 - [19] S. Borsanyi, G. Endrodi, Z. Fodor, S. D. Katz, S. Krieg, C. Ratti and K. K. Szabo, JHEP **1208**, 053 (2012) [arXiv:1204.6710 [hep-lat]].
 - [20] M. Cheng [RBC-Bielefeld Collaboration], PoS LAT **2007**, 173 (2007) [arXiv:0710.4357 [hep-lat]];
S. Borsanyi, G. Endrodi, Z. Fodor, S. D. Katz, S. Krieg, C. Ratti, C. Schroeder and K. K. Szabo, PoS LATTICE **2011**, 201 (2011) [arXiv:1204.0995 [hep-lat]].
 - [21] M. D'Elia, S. Mukherjee and F. Sanfilippo, Phys. Rev. D **82**, 051501 (2010) [arXiv:1005.5365 [hep-lat]].
 - [22] D. E. Kharzeev, L. D. McLerran and H. J. Warringa, Nucl. Phys. A **803**, 227 (2008) [arXiv:0711.0950 [hep-ph]].

- [23] M. A. Stephanov, K. Rajagopal and E. V. Shuryak, Phys. Rev. Lett. **81**, 4816 (1998).
- [24] C. R. Allton, M. Doring, S. Ejiri, S. J. Hands, O. Kaczmarek, F. Karsch, E. Laermann, and K. Redlich, Phys. Rev. D **71**, 054508 (2005), [hep-lat/0501030].
- [25] R.V. Gavai and S. Gupta, Phys. Rev. D **71**, 114014 (2005).
- [26] S. Ejiri, F. Karsch and K. Redlich, Phys. Lett. B **633**, 275 (2006) [hep-ph/0509051].
- [27] V. Koch, arXiv:0810.2520 [nucl-th].
- [28] M. M. Aggarwal *et al.* [STAR Collaboration], Phys. Rev. Lett. **105**, 022302 (2010).
- [29] X. Luo [STAR Collaboration], arXiv:1210.5573 [nucl-ex].
- [30] J. T. Mitchell [PHENIX Collaboration], arXiv:1211.6139 [nucl-ex].
- [31] A. Bazavov, H. -T. Ding, P. Hegde, O. Kaczmarek, F. Karsch, E. Laermann, S. Mukherjee and P. Petreczky *et al.*, Phys. Rev. Lett. **109**, 192302 (2012) [arXiv:1208.1220 [hep-lat]].
- [32] S. Borsanyi, arXiv:1210.6901 [hep-lat].
- [33] S. Datta, R. V. Gavai and S. Gupta, arXiv:1210.6784 [hep-lat].
- [34] S. Borsanyi *et al.*, JHEP **1201**, 138 (2012).
- [35] A. Bazavov *et al.* [HotQCD Collaboration], Phys. Rev. D **86**, 034509 (2012) [arXiv:1203.0784 [hep-lat]].
- [36] C. Schmidt, arXiv:1212.4283 [hep-lat].
- [37] J. P. Blaizot, E. Iancu and A. Rebhan, Phys. Lett. B **523**, 143 (2001) [hep-ph/0110369].
- [38] J. O. Andersen, S. Mogliacci, N. Su and A. Vuorinen, arXiv:1210.0912 [hep-ph].
- [39] N. Haque, M. G. Mustafa and M. Strickland, arXiv:1212.1797 [hep-ph].
- [40] A. Adare *et al.*, [PHENIX Collaboration], Phys. Rev. Lett. **104**, 132301 (2010); Phys. Rev. C **81**, 034911 (2010).
- [41] F. Geurts [STAR Collaboration], arXiv:1210.5549 [nucl-ex].
- [42] D. Lohner [ALICE Collaboration], arXiv:1212.3995 [hep-ex].
- [43] A. Adare *et al.* [PHENIX Collaboration], Phys. Rev. Lett. **109**, 122302 (2012);
- [44] H. van Hees, C. Gale and R. Rapp, Phys. Rev. C **84**, 054906 (2011).
- [45] A. Adare *et al.*, [PHENIX Collaboration], Phys. Rev. C **84**, 044905 (2011).
- [46] S. LaPointe [ALICE Collaborations], arXiv:1209.6198 [nucl-ex];
L. Massacrier [ALICE Collaboration], arXiv:1208.5401 [nucl-ex].
- [47] G. D. Moore and D. Teaney, Phys. Rev. C **71**, 064904 (2005);
H. van Hees, M. Mannarelli, V. Greco and R. Rapp, Phys. Rev. Lett. **100**, 192301 (2008);
S. Cao, G.-Y. Qin and S. A. Bass, arXiv:1209.5405 [nucl-th].
- [48] For a recent review see: R. Rapp and H. van Hees, R. C. Hwa, X.-N. Wang (Ed.) Quark Gluon Plasma 4, World Scientific, 111 (2010) [arXiv:0903.1096 [hep-ph]].
- [49] S. Caron-Huot and G. D. Moore, Phys. Rev. Lett. **100**, 052301 (2008).
- [50] H. T. Ding, A. Francis, O. Kaczmarek, F. Karsch, H. Satz and W. Soeldner, Phys. Rev. D **86**, 014509 (2012) [arXiv:1204.4945 [hep-lat]].
- [51] M. Asakawa, T. Hatsuda and Y. Nakahara, Prog. Part. Nucl. Phys. **46**, 459 (2001).
- [52] H. -T. Ding, arXiv:1210.5442 [nucl-th].
- [53] G. Aarts *et al.*, JHEP **1111**, 103 (2011) [arXiv:1109.4496 [hep-lat]].
- [54] H. -T. Ding, A. Francis, O. Kaczmarek, F. Karsch, E. Laermann and W. Soeldner, Phys. Rev. D **83**, 034504 (2011) [arXiv:1012.4963 [hep-lat]].
- [55] H. B. Meyer, Eur. Phys. J. A **47**, 86 (2011) [arXiv:1104.3708 [hep-lat]].
- [56] S. Sakai and A. Nakamura, PoS LAT **2007**, 221 (2007) [arXiv:0710.3625 [hep-lat]].
- [57] H. B. Meyer, Prog. Theor. Phys. Suppl. **174**, 220 (2008) [arXiv:0805.4567 [hep-lat]].
- [58] M. Cheng, *et al.*, Phys. Rev. D **79**, 074505 (2009) [arXiv:0811.1006 [hep-lat]].

- [59] R. D. Pisarski and F. Wilczek, Phys. Rev. D **29**, 338 (1984).
- [60] T. Hatsuda, M. Tachibana, N. Yamamoto and G. Baym, Phys. Rev. Lett. **97**, 122001 (2006);
J. -W. Chen, K. Fukushima, H. Kohyama, K. Ohnishi and U. Raha, Phys. Rev. D **80**, 054012 (2009).
- [61] R. Rapp, Acta Phys. Polon. B **42**, 2823 (2011) [arXiv:1110.4345 [nucl-th]].
- [62] A. Bazavov *et al.* [HotQCD Collaboration], Phys. Rev. D **86**, 094503 (2012) [arXiv:1205.3535 [hep-lat]].
- [63] A. Stathopoulos, A. M. Abdel-Rehim and K. Orginos, J. Phys. Conf. Ser. **180**, 012073 (2009).
- [64] R. Babich, J. Brannick, R. C. Brower, M. A. Clark, T. A. Manteuffel, S. F. McCormick, J. C. Osborn and C. Rebbi, Phys. Rev. Lett. **105**, 201602 (2010) [arXiv:1005.3043 [hep-lat]].
- [65] M. Luscher, Comput. Phys. Commun. **156**, 209 (2004) [hep-lat/0310048].