Opportunities for Lattice QCD Thermodynamics

Lattice QCD Executive Committee

R. Brower, (Boston U.) N. Christ (Columbia U.), M. Creutz (BNL),
P. Mackenzie (Fermilab), J. Negele (MIT), C. Rebbi (Boston U.),
D. Richards (JLab), S. Sharpe (U. Washington), and R. Sugar (UCSB)

Abstract

Over the next decade high performance computing resources will reach the petaflops scale. Coupled with current and planned experiments at RHIC, FAIR, and the LHC, these computational resources will offer significant opportunities for the advancement of our understanding of the properties of strongly interacting matter at high temperatures and densities. We describe expected quantitative and qualitative gains in (1) our knowledge of the equation of state at zero and nonzero density, (2) the phase diagram of QCD at zero and nonzero density, and (3) the structure of the plasma phase, including excited states and transport properties. We describe several key computational projects for achieving these gains and estimate the computational cost in units of teraflops-years (TF-y).

1 QCD at nonzero temperature and density

The properties of strongly interacting matter at nonzero temperature and baryon number density are being studied in heavy ion experiments at the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory (BNL). In the near future these experiments will be extended to even higher energies and temperatures at the Large Hadron Collider (LHC). By contrast, at BNL and at the future European heavy ion facility FAIR, a series of new low energy experiments is planned that will allow us to study such matter at moderate temperatures, but high baryon number density. The former physical conditions occurred in the early universe; the latter may approximate the environment in the interior of dense stellar objects such as neutron stars.

Under extreme conditions of high temperature or high baryon number density strongly interacting matter is expected to have a rich phase structure as indicated in Figure 1. Quantifying the drastic changes in the interaction among elementary particles that go along with such phase changes requires large scale numerical calculations.

Numerical studies of lattice QCD can provide a wealth of new information about properties of strongly interacting matter. Lattice QCD is likely to have a particularly strong impact on current and future experimental studies as well as the phenomenological modeling of hot and dense matter in the following three areas:

- Lattice calculations can provide detailed information about basic bulk thermodynamic properties: the equation of state, energy and entropy density, the pressure, and the velocity of

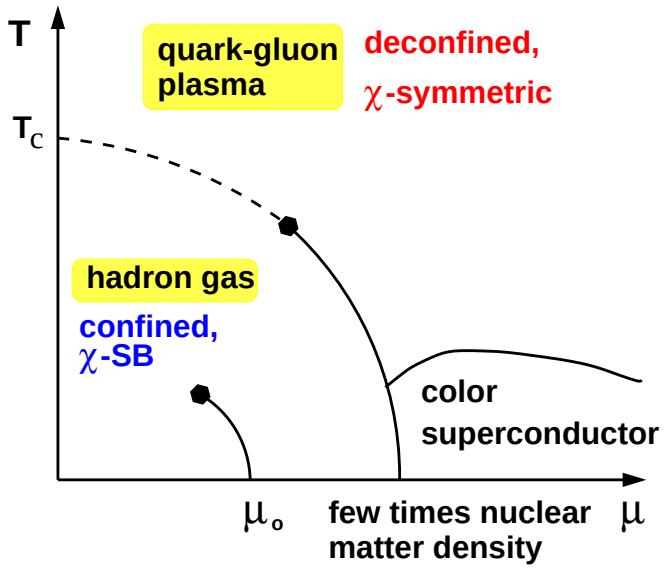

Figure 1: Phase diagram of strongly interacting matter

sound and basic structural properties: plasma modes and transport coefficients. These quantities are crucial input to the analysis of many experimental observables that characterize the formation of hot and dense matter in heavy ion collisions, and they are crucial for the hydrodynamic modeling of its time evolution [1, 2]. For example a precise knowledge of the equation of state is needed for a quantitative description of the expansion process and in the theoretical modeling of almost all experimental observables, including the hydrodynamic modeling of recent findings at RHIC, such as elliptic flow [3], viscosity [4] and quarkonium suppression [5] .

- Lattice calculations currently provide the only ab initio, quantitative method for determining the phase diagram of strongly interacting matter (Fig. 1), which, aside from the case of vanishing baryon number density (vanishing quark chemical potential), is largely unexplored. In particular, confirming the existence of a second order phase transition point in the phase diagram and subsequently determining its location accurately can only be achieved through demanding numerical calculations. Experiments at RHIC and FAIR are under consideration that would search for this critical point. Quantitative predictions from lattice calculations are needed.

- Lattice simulations of strongly interacting matter are limited to thermodynamic equilibrium and small deviations from it. Effective models help us develop insight and extend our understanding of the dynamical processes occurring in heavy ion collisions. Lattice calculations are essential for validating and constraining a variety of models ranging from hadronic resonance gas models at low temperature and quasi-particle models at high temperature to perturbative approaches at very high temperature [6].

2 Quark-gluon plasma equation of state

Lattice methods for determining the equation of state (EoS) are well developed, but numerically intensive. Our present knowledge of the continuum EoS comes with statistical errors of order 15% and probably comparable systematic errors. A combined error of order 5% at a physical light quark mass would provide a solid foundation for hydrodynamical modeling. This goal is easily feasible with petaflop resources.

The extrapolation of lattice results to the continuum requires carrying out calculations at a series of small enough lattice spacings a that the extrapolation is well controlled. For the EoS, pushing to smaller lattice spacing is expensive but feasible. To compute the energy density, pressure, and entropy requires a vacuum subtraction. That is, the simulation at $T > 0$ produces an unrenormalized value of the thermodynamic quantity, and the physically useful value is obtained by subtracting the corresponding zero temperature value. The procedure is straightforward, but the subtraction entails a loss of significance that worsens rapidly as the lattice spacing a is reduced. In fact, the fractional difference decreases as a^4. Since the numerical simulation estimates the quantities in the subtraction statistically with an error proportional to $1/\sqrt{N}$, the statistical sample size N must grow as a^{-8} to achieve the same accuracy in the result. The cost of obtaining a single statistically independent sample also grows with a high negative power of a. With these considerations, we estimate, conservatively, a net cost that scales as a^{-11}. Clearly, we require a careful analysis of discretization errors to make the most of simulation results at small a.

We divide the temperature scale (in units of the crossover temperature T_c) into three qualitatively different regions, each with its distinct numerical demands and impact on modeling.

- Low temperature region ($T < 0.95T_c$). Resonance gas models are often used to model strongly interacting matter in this range [7]. A reliable lattice calculation is needed to validate these models [8, 9].

- Transition region ($0.95T_c < T < 1.5T_c$). The crossover is strongly influenced by a nearby phase transition that restores chiral symmetry. The plasma is strongly interacting and only nonperturbative methods are applicable.

- High temperature region ($1.5T_c < T$). In this region one may hope to use resummed perturbation theory to characterize the plasma, but its reliability is unknown.

Figure 2 shows results for the EoS based on recent calculations with improved staggered fermions. On the left we show the difference of energy density and three times the pressure [10, 11], which is sometimes called the interaction measure. It summarizes our current knowledge of the EoS at low and high temperature at vanishing chemical potential. On the right we compare the ratio of pressure (p) to energy density (ε) at zero baryon number and at nonzero baryon number along curves of fixed entropy (S) per quark (or baryon) ($N_B = N_q/3$) [12]. Here the temperature is given in units of the crossover temperature T_0. In the hydrodynamic modeling of the expansion of dense matter created in a heavy ion collision, an accurate equation of state is of particular importance.

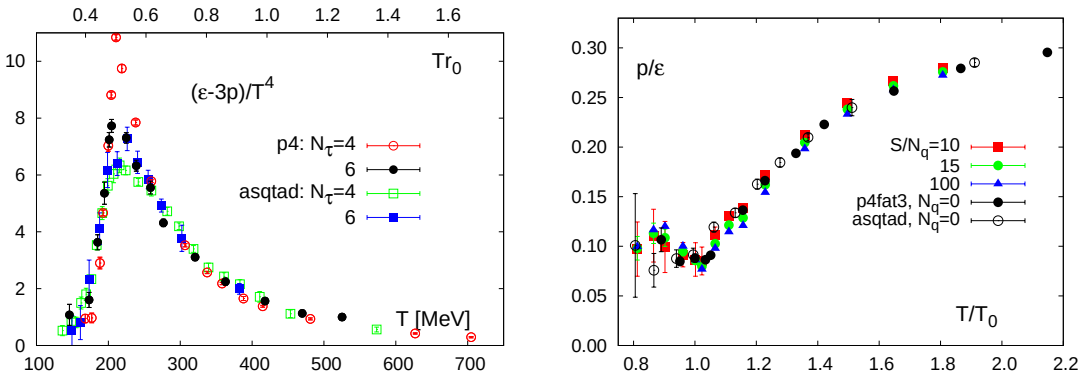

Figure 2: The equation of state of strongly interacting matter.

It is apparent from these figures that the qualitative features of the transition from the low to high temperature regime of QCD are quite well characterized, but that much more detailed studies are needed in all three temperature regimes to achieve control at the 5% level. Neither the approach to the perturbative regime at high temperature nor the consistency with resonance gas physics at low temperature is established on a quantitative level. Simulations with resources on the petaflop scale could have a tremendous impact here.

In the following subsections we discuss a strategy to reach the goal of 5% accuracy in the EoS at the physical light quark mass in each of the three temperature regimes listed above. In each regime we set lattice parameters to cover approximately the same range of lattice sizes, so a uniform extrapolation to the continuum will be possible.

Achieving this goal requires assessing and gaining control of cutoff effects (artifacts of the lattice approximation) as the lattice spacing is reduced. We discuss two popular ways to put quarks (fermions) on the lattice, namely "staggered" fermions and "domain wall" fermions with different sources of artifacts.

Staggered fermions suffer from a phenomenon called species doubling. Uncorrected, the resulting theory has four times as many quark species as desired. The standard expedient takes the fourth root of the "fermion determinant", which approximately corrects the multiplicity of quark species in the statistical ensemble. This formulation also distorts the desired chiral symmetry of the theory. The error in the approximation is widely believed to vanish in the continuum limit, but there is no proof [13].

The domain wall formulation does not suffer from species doubling. In addition to the usual lattice spacing, its lattice artifacts are controlled by a parameter L_s, which determines the degree to which the desired chiral symmetry is well approximated. The quality of that approximation is measured by the "residual mass" of the lightest quark. The residual mass should be small, which requires making L_s large and increases the computational cost.

For a given lattice spacing and set of physical parameters, a high quality domain wall simulation is far more expensive than the corresponding staggered fermion simulation. Thus we currently

have far more detailed results from staggered fermions, and the largest part of the computations we envision in this report are for staggered fermions. Nonetheless, well-chosen simulations with domain-wall fermions are essential for checking results from staggered fermions, particularly for quantities and phenomena sensitive to chiral symmetry.

In the subsections below we give estimates of the computational cost to achieve the desired accuracy. The estimates are summarized in Table 1. A brief explanation of some of the simulation parameters is in order. The simulation temperature is determined from a, the lattice spacing, and N_τ, the extent of the lattice in the Euclidean time direction, according to $T = 1/(N_\tau a)$. Thus at any given temperature, the approach to the continuum requires a larger N_τ.

2.1 Resonance gas regime: $T < 0.95T_c$

Contemporary simulations in the important hadronization temperature range of 150 - 200 MeV have been done at best at a fairly coarse lattice spacing of about 0.15 fm [10, 14, 15]. Here significant artifacts of the lattice formulation are likely to appear; their nature depends on the formulation. With the most extensively studied staggered fermion formulation, the meson spectrum is plagued by the increasing distortion of "taste multiplets". When the hadron spectrum is inaccurate, one may question its description of the hadronic medium for temperatures nearing the crossover T_c.

A simulation at $N_\tau = 12$ corresponds to a lattice spacing a of $0.09 - 0.12$ fm in the temperature range $140 - 180$ MeV. Measurements of the hadronic spectrum in this range show the expected scaling with decreasing a. Based on ongoing simulations at $N_\tau = 8$, we estimate a cost of 85 TF-y to carry out a comparable simulation at $N_\tau = 12$ for $T < T_c$. Such a simulation, combined with existing and ongoing simulations, should make it possible to reduce the uncertainty in the continuum extrapolation in this important region to 5%.

2.2 Transition temperature and thermodynamics in its vicinity: $0.95T_c < T < 1.5T_c$

This is the dramatic crossover region. The crossover is an echo of a nearby phase transition that restores chiral symmetry. Since the equation of state p/ε has a minimum at or close to the transition temperature T_c, the velocity of sound is small in this regime. The expanding and cooling dense matter created in a heavy ion collision thus spends a long time in this temperature regime. A determination of the transition temperature and the corresponding energy and entropy densities is of great importance.

In Fig. 3 we show a collection of recent results for the transition temperature as function of the light pseudo-scalar (pion) mass. Although there is good overall agreement the transition temperature still is not not known to better than 10%. Reducing the error below 5% requires additional calculations and the cross-check of results obtained with different fermion actions.

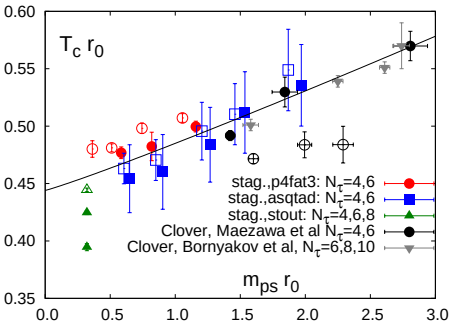

Figure 3: The transition temperature in (2+1)-flavor QCD as compiled in [11] obtained with different staggered fermion actions (asqtad [16], p4fat3 [17], stout [18]) and results obtained in two-flavor QCD with improved Wilson fermions (clover fermions) [19, 20]. Shown is the transition temperature in terms of a scale parameter (r_0) extracted from the slope of the static quark potential, versus the lightest pseudo-scalar meson mass, also expressed in units of r_0. Open symbols correspond to results on lattices with temporal extent $N_\tau = 4$. Filled symbols correspond to $N_\tau = 6$ or larger.

A key question is whether the distortion of chiral symmetry at nonzero lattice spacing in the staggered fermion formulation modifies the crossover temperature and the peak in chiral susceptibility. To answer this question requires good control of the continuum limit in the staggered fermion simulation and a companion calculation with chiral domain wall fermions.

In this region it should suffice to carry out the staggered fermion simulation at $N_\tau = 10$. Together with results from existing and ongoing calculations at $N_\tau = 6$ and 8, we should be able to do a controlled extrapolation to the continuum. A calculation at five temperatures in this critical region would cost 75 TF-y.

The domain wall formulation does not have the spectral lattice artifacts of staggered fermions. However, it is far more computationally expensive. The quality of its chiral behavior is measured by the "residual quark mass" parameter. At $N_\tau = 10$ in the transition region, we know from other tests that this quantity is small enough to assure good control of the chiral behavior. Our goal is to check the determination of the crossover temperature T_c and the strength of the peak in the chiral susceptibility. A calculation at seven temperatures in this critical region on a $48^3 \times 10$ lattice with $L_s = 32$ at twice the physical light quark mass would be sufficient and would cost 100 TF-y.

2.3 Perturbation theory and the high temperature limit of QCD: $1.5T_c < T < 5T_c$

Heavy ion experiments at the LHC will generate matter with initial energy densities as high as 1 TeV/fm^3, which may correspond to temperatures as high as $(800-1000)$ MeV. Even at these high temperatures, straightforward high-temperature perturbation theory is not sufficiently convergent to give a quantitative prediction. Refined resummation techniques have been developed to deal

with this problem. These techniques have been quite successful in the purely gluonic sector of QCD but are poorly developed for QCD with dynamical quarks. To make progress here requires accurate numerical results from lattice calculations.

In this temperature regime quark mass thresholds for the charm quarks become relevant. Perturbative calculations suggest that the charm quark contribution to the EoS can be significant already at temperatures of a few times the transition temperature [21]. In order to take into account the contribution of heavy quarks some progress can be made in a quenched approximation to the heavy quark sector. If one aims, however, at accuracies on the few percent level, a dynamical simulation will be needed.

Although including heavy quarks in a dynamical simulation is relatively cheap, a fine lattice is needed to resolve their contribution accurately. In order to incorporate the contribution of a charm quark in a dynamical simulation, the inverse lattice spacing should be much larger than the heavy quark mass m_c. This requires simulations on lattices with temporal extent $N_\tau = 10$ or even $N_\tau = 12$. For example, on a lattice with temporal extent $N_\tau = 10$, the inverse lattice spacing at $T \simeq 2T_c \simeq 0.4$ GeV is $a^{-1} = 4$ GeV. Simulations on $N_\tau = 8$ and 10 lattices should then allow an analysis of cut-off effects systematically in dynamical simulations that also include the charm quark sector.

To control the EoS at these high temperatures is computationally highly demanding. As is obvious from the left panel of Fig. 2, the signal is an order of magnitude smaller here than close to the transition region. Nonetheless, it is feasible to calculate the EoS on a lattice with temporal extent $N_\tau = 10$ at four temperatures in the range $2T_c \leq T \leq 4T_c$ to establish the temperature dependence of $(\varepsilon - 3p)$. Combining such a calculation with results obtained on lattices with temporal extent $N_\tau = 6$ and 8 will permit a controlled extrapolation to the continuum limit, where systematic errors will be below the 1% level. Overall errors will then be entirely controlled by the statistical error that can be achieved. Based on the experience with current simulation parameters it is conceivable that errors below 5% can be reached for $(\varepsilon - 3p)$. We estimate a cost of about 150 TF-y.

2.4 Equation of state at nonzero density

Heavy ion collisions occur in a baryon-rich environment, whereas lattice simulations are naturally suited for zero baryon density, *i.e.* zero baryon chemical potential. For technical reasons direct simulation at nonzero density and appropriately large lattice volume is extremely difficult. To reach a small, nonzero baryon number density, one constructs the Taylor series expansion in the chemical potential [22]. The coefficients of the series are evaluated in a standard simulation at zero chemical potential. This method is effective for the relatively low baryon number densities of heavy ion collisions. As more terms in the Taylor are calculated, it becomes possible to push to higher chemical potential.

Upcoming low energy runs at RHIC will achieve higher baryon density. Our goal is to determine the equation of state at the relevant densities to 5% accuracy in $2+1$ flavor QCD. The procedure for extending the equation of state to nonzero chemical potential by means of a Taylor expansion is well developed [22, 23, 24]. Typically, one works directly with the pressure. How-

ever, to determine the EoS at fixed entropy per baryon, as it is needed to describe the expanding dense matter created in heavy ion collisions, one also has to calculate Taylor expansion coefficients for the energy and entropy densities. This requires good control over the temperature dependence of the expansion coefficients for the pressure.

The state of the art of such calculations is given in Fig. 2(right). Coefficients of the Taylor expansion of the pressure are constructed from expectation values of products of quark loop operators on the equilibrium, zero density ensemble. Current calculations have been carried to sixth order for two light quark flavors. Algebraic expressions for the coefficients become rapidly complicated at higher order. Fortunately, the coding at any high order can be automated.

The calculation reuses the lattices generated in the equation of state study. It involves a large set of inversions of the fermion matrix with different random starting vectors. The number of random vectors needed to reach comparable statistical error in different orders N of the Taylor expansion grows exponentially, *i.e.* roughly like 4^N. Moreover, the computational effort per set of random vectors increases approximately as 1.5^N. The overall computational effort thus rises like 6^N.

The main domain of interest here is the temperature region $T < T_c(\mu_q = 0)$. This also is the computationally most difficult region. Based on current studies on lattices with temporal extent $N_\tau = 4$ we estimate the resources needed to calculate the expansion up to 8th order with an accuracy of 10% and add the effort required to reach 20% accuracy for the 10th order coefficient. The latter would automatically imply that the 8th order coefficient is obtained with an accuracy of about 3%. As errors of these coefficients are strongly correlated and the square root of ratios of subsequent coefficients is needed to estimate the radius of convergence of the Taylor series this suffices to get estimates for the convergence radius with a statistical error of 10%.

A reasonable strategy would be to perform an analysis up to 8th order for four temperature values below $T_c(\mu_q = 0)$, and add one 10th order calculation at the estimated chiral critical temperature to this. The 8th order calculations would require about 150 TF-y and the single 10th order calculation requires 340 TF-y. The latter would be four times more expensive, if one also aims at a 10% error on the expansion coefficient, *i.e.* a 5% error on the estimate for the radius of convergence and hence the location of the chiral critical point in the QCD phase diagram.

3 The phase diagram of strongly interacting matter

3.1 Phase boundary at zero baryon number density

It is now widely accepted that at physical quark masses and zero baryon density the transition in QCD from hadronic matter to a quark plasma is a rapid, but analytic crossover, rather than a phase transition. At other values of the quark masses and at nonzero baryon number density there are phase boundaries as indicated in the left panel of Fig. 4. For example, in the scenario of this figure, as the light quark mass is lowered toward the left from the physical point, a critical line is encountered. On this line we get a genuine high temperature phase transition with critical behavior.

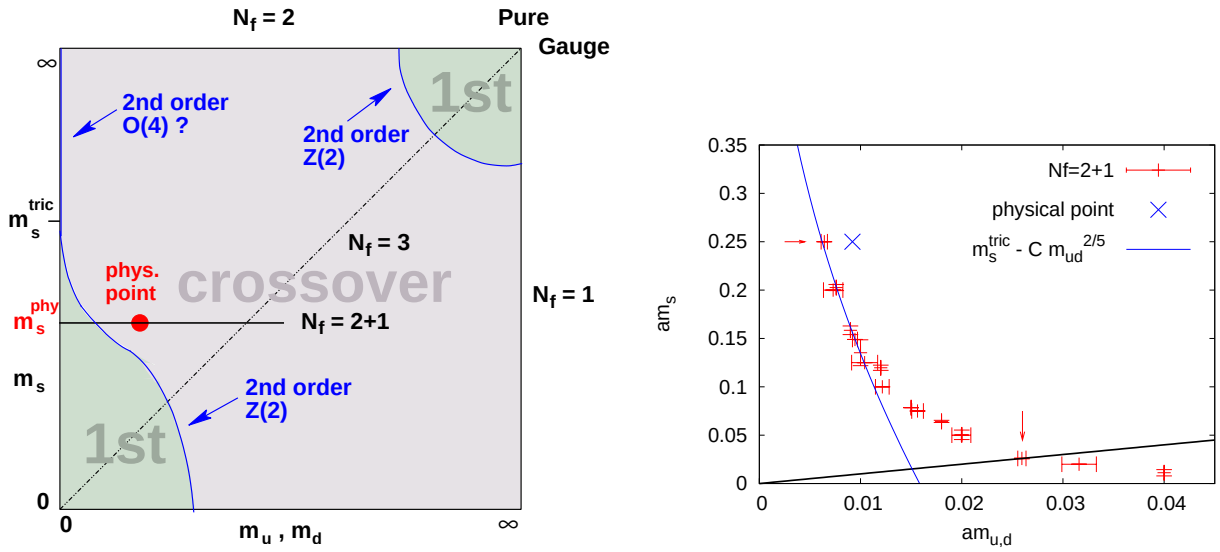

Figure 4: Left panel: sketch of the expected phase boundaries at zero chemical potential as a function of degenerate up and down quark masses. The physical point is plotted as a dot in the crossover region. To the left and below the 2nd order $Z(2)$ boundary a high temperature first order phase transition occurs. Right panel: result of an actual measurement of a portion of the 2nd order $Z(2)$ phase boundary from Ref. [25]. The axes give bare quark masses in lattice units and the blue cross marks the physical point.

When a strange quark is present, that critical line may occur at a nonzero light quark mass. For light quark masses below that line, the high temperature phase transition is first order.

This phase boundary has been mapped out by de Forcrand and Philipsen on rather coarse lattices ($N_\tau = 4$) with unimproved staggered fermions [25] as shown in the right panel of Fig. 4. It is well known that lattice artifacts have a strong influence on the location of the phase boundary [26, 27]. Thus it is important to push to smaller lattice spacing and explore more systematically the influence of explicit chiral symmetry on the the transition parameters.

Clearly, what is needed is a systematic characterization of the phase boundary with improved staggered fermions, proceeding from $N_\tau = 4$ to $N_\tau = 6$. This should be done at about ten different sets of quark mass values (m_{ud}, m_s), as in the unimproved example of Fig. 4. We estimate the cost to be about 200 TF-y.

The sensitivity of these results to the chiral sector of QCD should be studied by repeating the analysis at a few selected parameter sets (m_{ud}, m_s) with a chiral fermion formulation of QCD, namely the domain wall fermion method. Its high computational cost has prevented us from using it to characterize the phase boundary. However, with petascale resources it will be possible to do so. It will be feasible to carry out such simulations at least for three-degenerate-flavor QCD, *i.e.* with $m_s = m_{ud}$. To carry out such a simulation with all three quark masses equal to 1/20 of the physical strange quark mass on a $48^3 \times 10$ lattice at $L_s = 32$ would require 40 TF-y.

3.2 Soft plasma modes at the transition

Soft (massless) modes in the theory play an important role at the second order phase transitions. They may even dominate the evolution of the plasma at these temperatures. These modes can be studied at the same time we explore the phase boundary.

3.3 Phase boundary at nonzero baryon number density

At small, possibly even physical values of the strange quark mass, the phase boundary occurs at nonzero light quark mass as indicated by the $Z(2)$ line in Fig. 4. What happens to this line as the baryon number density is increased? Does it move toward the physical point or away from it? These two scenarios are sketched in Fig. 5. In the scenario on the left, as the chemical potential is increased at fixed physical quark mass, a critical phase boundary is encountered, beyond which the high temperature phase transition is first order.

Critical behavior would give rise to observable effects. A strongly first order phase transition would have dramatic observable effects, including a phase separation and metastable states. Which of the two scenarios in Fig. 5 is correct? A large experimental program at RHIC and a new heavy ion facility in Europe (FAIR) will be devoted to this question. Giving firm answers to this question through lattice calculations can have a tremendous impact on the final layout of these experiments.

A set of statistical quantities called "Binder cumulants" measure fluctuations at the transition temperature and help to distinguish a crossover from a genuine phase transition. De Forcrand and Philipsen have shown that simulations at imaginary chemical potential are helpful in determining the curvature of the critical surface [25]. Introducing a nonzero imaginary chemical potential is expensive. It requires generating a new set of gauge configurations for every simulation value of the imaginary chemical potential. A Taylor expansion of the Binder cumulants is also effective. The Taylor coefficients cost more to calculate. For each parameter set (m_{ud}, m_s) at which this analysis is to be performed, the computational effort discussed in the previous section would be doubled. Even so, the net expense is far less, since it is not necessary to generate a new ensemble of gauge configurations. We simply reuse the ones discussed in the previous section. Doing this analysis for only every second set (m_{ud}, m_s) of quark masses should give a fair picture of the curvature of the critical surface. This calculation will thus require about 100 TF-y.

4 Structure of the QGP

4.1 In-medium properties of hadrons

Deconfinement implies the dissolution of hadrons into their constituents. Thus one would expect that an experimental signal for deconfinement in heavy ion collisions is the disappearance of the

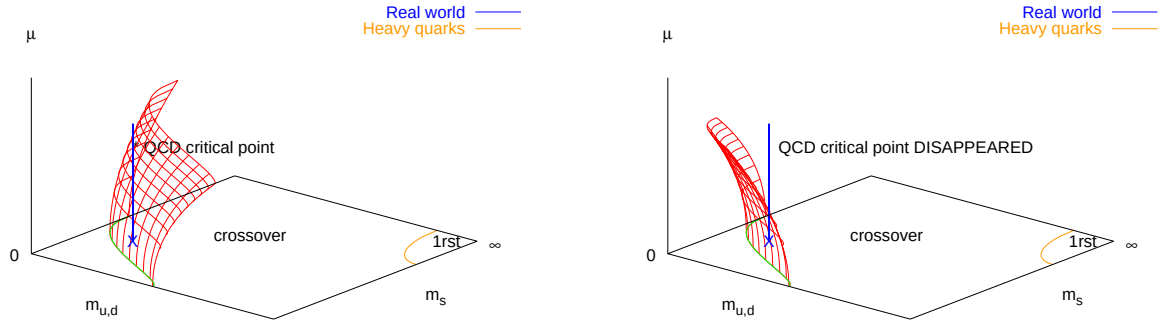

Figure 5: Two possible alignments of the chiral critical surface at low chemical potential from [25]. Left: the scenario permitting a first order phase transition at high densities and temperatures. Right: the scenario allowing only a crossover.

charmonium and bottomonium peaks in dilepton production [28]. Lattice simulations and the analysis of experimental measurements suggest, however, that hadronic matter is strongly interacting at temperatures well above T_c. Consequently quarkonium production is suppressed to a degree that depends on temperature [29, 30, 31].

A major puzzle from RHIC experiments is the large degree of collectivity, *i.e.* large flow. If this happens because the system is thermalized to a good approximation, the degree of quarkonium suppression should provide an estimate for the temperature. So far, relatively little is known about quarkonium properties at nonzero temperature; therefore, lattice information is crucial. If the system is locally thermalized, it should have a very low viscosity, *i.e* very small mean free path to produce the observed flow. Furthermore, recent experimental results from RHIC on open charm production indicate that even heavy quarks have a small mean free path. Thus estimating the heavy quark diffusion is also important.

Lattice simulations provide an indirect means for determining the thermal suppression of quarkonium production and for determining transport coefficients. The calculation involves measuring the imaginary time correlation function of the appropriate heavy-quark current $J(t,x,T)$ and deducing from it the real-time spectral function $\rho(\omega,p,T)$. The choice of current (scalar, pseudoscalar, vector, etc.) determines the channel of interest. Dilepton rates are obtained from the electromagnetic current correlator. The measured correlator has the general form

$$G(t,p,T) = \int d^3x \exp(ip\cdot x)\langle J(t,x,T)J(0,0,T)\rangle. \tag{1}$$

Deriving the spectral functions from measurements of the lattice correlators is difficult. The lattice correlator $G(t,p,T)$ is known on a discrete set of imaginary time values t, spatial momenta p, and temperature T, whereas the unknown spectral function $\rho(\omega,p,T)$ has support, in principle, for all real frequencies ω. The Euclidean correlator is related to the spectral function through

$$G(t,p,T) = \int_0^\infty \frac{d\omega}{2\pi}\frac{\cosh[\omega(t-1/2T)]}{\sinh(\omega/2T)}\rho(\omega,p,T). \tag{2}$$

Peaks at nonzero frequency in ρ correspond to plasma excitations. In principle transport coefficients are obtained from widths of peaks in ρ near zero frequency ω. For example, the light quark vector current correlator measures the electric conductivity [32, 33], and the heavy quark vector current correlator measures the heavy quark diffusion coefficient. For light quarks this method is difficult in practice, since the transport contributions are not easily separated from the resonance contributions. For heavy quarks, however, the spectral separation is much easier and recent promising attempts exploit this [34]. The heavy quark diffusion coefficient is indirectly related to the shear viscosity.

To extract information about $\rho(\omega, p, T)$ from lattice measurements, it is common to impose additional constraints, either from model assumptions about the spectral function or from additional conditions, as in, for example, the popular maximum entropy method, which minimizes the deviation of the predicted spectral function from a featureless reference model. Clearly, the more imaginary time values t at which the correlator is known and the more precisely the correlator is known, the less one must depend on additional constraints. Experience in similar condensed matter physics applications suggests that to obtain physically useful results requires several dozen imaginary time values.

Recent lattice calculations for a pure gluon plasma suggest that the charmonium signal persists to $1.5T_c$ or higher [30, 31]. Figure 6 shows results of a recent calculation [31] including error bars. These results were obtained on anisotropic lattices with up to 40 imaginary time values.

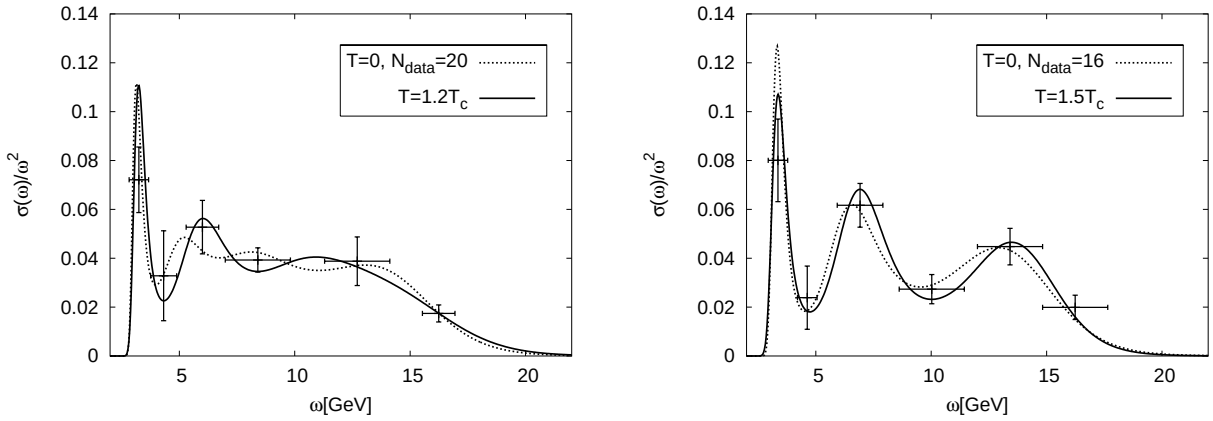

Figure 6: Spectral function for the η_c correlator from [31]. The error bars indicate the uncertainty in the determination of the features of the specral function.

There is clearly room for improvement. We need calculations with of the order of 100 or more points. For example, to carry out a quenched simulation on anisotropic $128^3 \times N_\tau$ lattices at seven temperatures would cost 15 TF-y. A companion simulation on $48^3 \times N_\tau$ lattices with both light and strange quarks included in the ensemble would cost 100 TF-y.

5 Conclusion

The advent of petaflops-scale computing promises dramatic gains in our understanding of the properties of strongly interacting matter at high temperatures and densities. We have described a program of lattice calculations that will (1) allow us to determine the equation of state of strongly interacting matter to an accuracy of 5%, (2) locate the critical surface of the QCD phase diagram at zero baryon density and predict its curvature as the baryon density is increased, (3) advance our understanding of the structure of the quark-gluon plasma and (4) determine some key transport coefficients. The first goal will provide essential, solid input for hydrodynamical modeling of heavy ion collisions, the second and third could represent a potential breakthrough by moving us from a qualitative to a quantitative understanding of the phase diagram and of the survivability of hadrons at high temperature, and the fourth and most ambitious goal could very well give us the first reliable lattice result for a transport coefficient of the quark-gluon plasma.

Acknowledgements: We thank Carleton DeTar, Frithjof Karsch and Robert Mawhinney for their contributions to this paper.

Project	lattice	temps	quark masses	trajecs	cost (TF-y)
EoS: $\mu = 0$, $T < 0.95T_c$	$48^3 \times 12$	7	2	100,000	35
EoS: $\mu = 0$, $T < 0.95T_c$	48^4	7	2	25,000	50
EoS: $\mu = 0$, $0.95T_c < T < 1.05T_c$	$48^3 \times 10$	5	2	100,000	25
EoS: $\mu = 0$, $0.95T_c < T < 1.05T_c$	48^4	5	2	25,000	50
EoS: $\mu = 0$, $2T_c < T < 4T_c$	$64^3 \times 10$	4	2	100,000	50
EoS: $\mu = 0$, $2T_c < T < 4T_c$	64^4	4	2	25,000	100
EoS DWF: $\mu = 0$, $0.95T_c < T < 1.05T_c$	$48^3 \times 10 \times 32$	4	1	50,000	100
EoS: $\mu > 0$ $T < 0.95T_c$ 8th order	$32^3 \times 8$	4	1	50,000	150
EoS: $\mu > 0$ $T < 0.95T_c$ 10th order	$32^3 \times 8$	1	1	50,000	340
phase boundary, $\mu = 0$	$32^3 \times 6$	4	10	10,000	200
phase boundary, $\mu = 0$, DWF	$48^3 \times 10 \times 32$	4	4	10,000	40
phase boundary $\mu > 0$	$32^3 \times 6$	4	4	10,000	100
spectral function, quenched	$128^3 \times N_\tau$	7	1	10,000	15
spectral function, dynamical	$48^3 \times N_\tau$	7	1	10,000	100

Table 1: Summary of simulation parameters and cost estimates. Cost estimates are based on current experience at $N_\tau = 6$ and 8. The computational effort is assumed to scale with decreasing lattice spacing as a^{-11} with quark masses fixed in physical units. Simulations labeled $\mu = 0$ are at zero quark number density. Simulations labeled $\mu > 0$ imply a Taylor expansion in chemical potential to reach small nonzero densities. Temperature ranges are expressed in terms of T_c, the relevant crossover temperature. The parameter "trajecs" measures the size of the statistical sample needed. The lattice dimension for the domain wall fermion simulations (DWF) includes the "fifth dimension" L_s parameter.

References

[1] P. F. Kolb, J. Sollfrank, and U. W. Heinz, Phys. Rev. **C62**, 054909 (2000), hep-ph/0006129.

[2] P. Huovinen and P. V. Ruuskanen (2006), nucl-th/0605008.

[3] C. Nonaka, B. Muller, M. Asakawa, S. A. Bass, and R. J. Fries, Phys. Rev. **C69**, 031902 (2004), nucl-th/0312081.

[4] M. Asakawa, S. A. Bass, and B. Muller, Phys. Rev. Lett. **96**, 252301 (2006), hep-ph/0603092.

[5] T. Gunji, H. Hamagaki, T. Hatsuda, and T. Hirano (2007), hep-ph/0703061.

[6] J. P. Blaizot, E. Iancu, and A. Rebhan, Phs. Lett. **B470**, 181 (1999), hep-ph/9910309.

[7] P. Braun-Munzinger, K. Redlich, and J. Stachel, in *Quark Gluon Plasma 3*, edited by R. C. Hwa and X.-N. Wang (World Scientific, 2004).

[8] F. Karsch, K. Redlich, and A. Tawfik, Eur. Phys. J **C29**, 549 (2003), hep-ph/0303108.

[9] F. Karsch, K. Redlich, and A. Tawfik, Phys. Lett. **B571**, 67 (2003), hep-ph/0306208.

[10] C. Bernard *et al.* (MILC) (2006), hep-lat/0611031.

[11] F. Karsch (2007), hep-ph/0701210.

[12] S. Ejiri, F. Karsch, E. Laermann, and C. Schmidt, Phys. Rev. **D73**, 054506 (2006), hep-lat/0512040.

[13] S. R. Sharpe, PoS **LAT2006**, 022 (2006), hep-lat/0610094.

[14] F. Karsch, E. Laermann, and A. Peikert, Phys. Lett. **478**, 447 (2000), hep-lat/0002003.

[15] Y. Aoki, Z. Fodor, S. D. Katz, and K. K. Szabo, JHEP **01**, 089 (2006), hep-lat/0510084.

[16] C. Bernard *et al.* (MILC), Phys. Rev. **D71**, 034504 (2005), hep-lat/0405029.

[17] M. Cheng *et al.*, Phys. Rev. **D74**, 054507 (2006), hep-lat/0608013.

[18] Y. Aoki, Z. Fodor, S. D. Katz, and K. K. Szabo, Phys. Lett. **B643**, 46 (2006), hep-lat/0609068.

[19] V. G. Bornyakov *et al.*, PoS **LAT2005**, 157 (2006), hep-lat/0509122.

[20] Y. Maezawa *et al.* (2007), hep-lat/0702005.

[21] M. Laine and Y. Schroder, Phys. Rev. **D73**, 085009 (2006), hep-ph/0603048.

[22] S. Ejiri *et al.*, Prog. Theor. Phys. Suppl. **153**, 118 (2004), hep-lat/0312006.

[23] C. R. Allton *et al.*, Phys. Rev. **D71**, 054508 (2005), hep-lat/0501030.

[24] R. V. Gavai and S. Gupta, Phys. Rev. **D68**, 034506 (2003), hep-lat/0303013.

[25] P. de Forcrand and O. Philipsen, JHEP **01**, 077 (2007), hep-lat/0607017.

[26] F. Karsch, E. Laermann, and C. Schmidt, Phys. Lett. **B520**, 41 (2001), hep-lat/0107020.

[27] F. Karsch *et al.*, Nucl. Phys. Proc. Suppl. **129**, 614 (2004), hep-lat/0309116.

[28] T. Matsui and H. Satz, Phys. Lett. **B178**, 416 (1986).

[29] S. Datta, F. Karsch, P. Petreczky, and I. Wetzorke, Phys. Rev. **D69**, 094507 (2004), hep-lat/0312037.

[30] M. Asakawa, and T. Hatsuda, Phys. Rev. Lett, **92**, 012001 (2004), hep-lat/0308034.

[31] A. Jakovac, P. Petreczky, K. Petrov, and A. Velytsky, Phys. Rev. **D75**, 014506 (2007), hep-lat/0611017.

[32] S. Gupta, Phys. Lett. **B597**, 57 (2004), hep-lat/0301006.

[33] G. Aarts, C. Allton, J. Foley, S. Hands, and S. Kim (2007), hep-lat/0703008.

[34] P. Petreczky, and D. Teaney, Phys. Rev. **D73**, 014508 (2006), hep-ph/0507318.