PROJECT NARRATIVE

1 SIGNIFICANCE OF RESEARCH

Background: A central goal of the experimental program at the Relativistic Heavy-Ion Collider (RHIC) of the Brookhaven National Laboratory (BNL) and at the Large Hadron Collider (LHC) at CERN, Switzerland is the exploration of the phase-diagram of strong-interaction QCD matter. The baryon-rich quark-gluon plasma (QGP) created at lower RHIC energies may experience a sharp first-order phase transition as it cools, with bubbles of QGP and bubbles of hadrons coexisting at a well-defined temperature. This region of co-existence ends in a critical point [1], where QGP and ordinary hadron-matter become indistinguishable. To search for the critical point in the QCD phase-diagram (see Figure 1) major experimental effort— the high-luminosity phase II of the RHIC Beam Energy Scan (BES-II) [3]— will be undertaken during the period 2019-2021.

On the other hand, QGP created at LHC and top RHIC energies consists of almost as much anti-matter as matter characterized by the nearly vanishing baryon-number chemical potential. Hot-dense lattice QCD calculations predicts that under these conditions the transition from the QGP to a hadron gas occurs through a smooth crossover, with many thermodynamic properties changing dramatically but continuously within a narrow range of temperature. It is conjectured that the pseudo-critical crossover in the physical world arises out of an underlying true chiral phase transition in the limit of vanishing quark masses. To what extent this conjecture is true can be studied experimentally by measuring non-Gaussian cumulants of conserved charge fluctuations [4]. This experimental experimental challenge will be undertaken at LHC with the new high-luminosity upgrade scheduled for 2021.

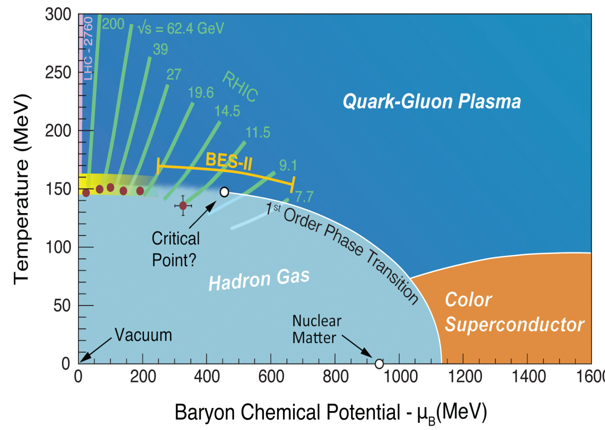

Figure 1. *A schematic phase diagram of QCD including and the range covered by the BES-II at RHIC, scheduled for 2019-21 [2].*

Motivation: With extensive upgrades of detectors and accelerator technologies the heavy-ion experiments at RHIC and LHC is about to enter, starting 2019, a new regime of high-luminosity studies. These will lead to unprecedented statistical accuracy that characterize properties of strong-interacting matter in the transition region between hadronic matter and QGP. The scheduled new heavy-ion runs at RHIC and LHC will lead to new experimental results on the transition at vanishing as well as non-vanishing net baryon-number densities, which in both experiments gets characterized by properties of higher order non-Gaussian cumulants of net charge fluctuations [5, 6, 7, 8], *i.e.* fluctuations of net baryon-number, net electric charge and net strangeness.

Goals and anticipated results: The non-Gaussian cumulants of net charge fluctuations also can be computed starting from the fundamental theory of quarks and gluons— quantum chromodynamics (QCD)— using large-scale numerical calculations of the lattice-regularized version of the theory. The goal of the proposed research is to provide quantitative QCD-based results on properties of strong-interaction matter in the vicinity of the pseudo-critical temperature that characterizes the transition from a phase dominated by hadrons as fundamental degrees of freedom to the QGP phase. With this project we focus on the QCD-based calculation non-Gaussian cumulants of net charge fluctuations that can directly be confronted with experimental findings in ultra-relativistic heavy-ion experiments at RHIC and LHC. Furthermore, QCD-based calculations of these cumulants will enable determination of phase-structure and bulk thermodynamic properties hot-dense strong-interaction matter from the fundamental theory of quarks and gluons. The important anticipated physics outcomes of the proposed research are:

– Continuum-extrapolated non-Gaussian fluctuations of various conserved charges.
– A more constrained location of the QCD critical point and continuum-extrapolated chiral crossover temperature at non-vanishing baryon densities.
– Continuum-extrapolated QCD equation-of-state for non-vanishing baryon densities.
– Identify degrees of freedom near the QCD crossover and freeze-out

Relevance and timeliness: The mission of the DOE's Nuclear Physics program is to discover, explore, and understand all forms of nuclear matter. As outlined in the 2015 NSAC Long Range Plan [2], a key component of the mission of this program is mapping the phase diagram of nuclear matter. The RHIC machine is uniquely suited for this doping scan. BES-II's high-statistics exploration of the phase diagram is a major priority RHIC and entire three year's (2019-21) of RHIC operation will be devoted for this purpose. The goal of the BES-II program is to discover a critical point in the QCD phase diagram if this landmark lies in the experimentally accessible region. Whether or not it is, BES-II program is uniquely positioned to map out how properties of hot QCD matter, including its equation-of-state, transport properties, and its baryonic and topological fluctuations, change with baryon density. On the other hand, the complementary program of at LHC, especially after its 2021 high-luminosity upgrade, will probe the transition region more deeply, by going beyond the traditional measurements, with higher order non-Gaussian fluctuations.

The objective of the proposed research is to understand the phase-structure of hot-dense QCD matter, as well as its equilibrium properties arise from the fundamental theory of quarks and gluons. In addition to making fundamental advancement in our understanding of strong-interaction matter, the proposed calculations provide QCD-results for quantities that can be directly confronted with experiments both at RHIC and LHC. Furthermore, the proposed computations will serve as key inputs for the dynamical modeling that will afford us a quantitative understanding from BES-II measurements of how the properties of hot QCD matter change with baryon density. Thus, the proposed research not only at the forefront of fundamental scientific advancement but also is very relevant to the mission of the US DOE Office of Science's Nuclear Physics program and perfectly in-line with the 2015 Long Range Plan for Nuclear Science [2].

This research program is also especially timely. The BES-II program is expected to commence in a year from now and is planned to continue at least till 2021. The LHC high-luminosity upgrades are expected to be completed by 2021. Furthermore, calculations undertaken as part of this project will provide the essential QCD inputs that are required urgently for dynamical modeling of heavy-ion collisions at BES energies, already underway within the framework of the DOE-funded BEST Collaboration [9]. Without having access to adequate computing time on Summit all the urgently needed lattice QCD calculations for RHIC BES-II will never materialize within the lifetime of the experiment, and before the start of the high-luminosity LHC runs.

Previous INCITE awards: In the past the project team has continuing access to computing time on Titan (and also a bit on Mira) supercomputer since 2014 through 2 USQCD-led INCITE awards. The continuing USQCD-led INCITE award on Titan, in principle, could have been renewed for another year. However, following the new USQCD policy that grant will not be renewed on Titan for the third year; instead we are applying for a new 3-year INCITE award for Summit focusing on solely on the hot-dense lattice QCD activities of USQCD consortium. The previous INCITE award led to the calculations of QCD equation-of-state (EoS) at vanishing baryon density. The results were published in Ref. [10], which has over 500+ citations to-date (according to Google Scholar), but has no direct bearing on the proposed research. The continuing INCITE award, as well as part of the previous one resulted in computations of the QCD EoS [11] and non-Gaussian fluctuations of net baryon-number [12], both at moderate values of baryon densities. Both these works were published last year, and Ref. [11] has already accumulated 75+ citations (according to Google Scholar) within a year. Refs. [11, 12], as well as the recently presented [13] (manuscript under preparation) work on the QCD phase-structure have prepared the ground for the proposed research, as will be described in Section 2.2.

2 RESEARCH OBJECTIVES AND MILESTONES

2.1 Lattice QCD at non-zero baryon density: Taylor expansion method

The proposed research aims to determine the thermodynamic properties of baryon-rich hot QCD matter in a completely model-independent manner, directly from the underlying fundamental theory of QCD. We will use large-scale numerical computations, based on state-of-the-art lattice QCD techniques. Over the last decade lattice QCD has proven to be the most successful technique for model-independent, first-principle calculations of the phase structure and properties of hot QCD matter; for recent reviews see Refs. [14, 15].

The standard Monte-Carlo sampling employed during lattice QCD calculations is inapplicable for a non-zero value of the baryon chemical potential, as in this case, the sampling weight, *i.e* the QCD action, becomes a complex number, and, thus, cannot be interpreted as a probability density. Although there are some recent developments towards performing lattice QCD simulations at non-zero baryon chemical potential using the complex Langevin and the Lefschetz thimble methods, at present none of them are mature enough to be applied for QCD calculations with physical quark masses and close to the thermodynamic and continuum limits [16, 17, 18, 19, 20]. At present, the two most viable and successful method for undertaking lattice QCD calculations at non-zero baryon, charge or strangeness chemical potentials for QCD with physical quark masses and close to the continuum and thermodynamic limit are— (i) simulations of QCD at imaginary values of the chemical potentials [21, 22] combined with ansätze for the analytic continuation of observables of interest, (ii) systematic expansions of the QCD partition function and observables in Taylor series set up around vanishing values of the chemical potentials [23, 24]. While the former allows for direct numerical simulations, it needs to deal with ansätze for analytic continuation for which the relevance of systematic effects are difficult to control. In practice these approaches currently use polynomial ansätze for the analytic continuation, which are equivalent to Taylor series. On the other hand, a direct, high precision, calculation of the expansion coefficients appearing in a Taylor series has the advantage that it provides direct control over statistical and systematic effects contributing in different ways to different expansion coefficients.

We aim at precision calculations of higher order cumulants of net baryon-number (B), electric charge (Q), and strangeness (S) fluctuations and their mutual cross-correlations. These cumulants are obtained as derivatives of the logarithm of the QCD partition function, $Z(T, V, \mu_B, \mu_Q, \mu_S)$, with respect to the chemical potentials $\mu_{(B,Q,S)}$, evaluated at vanishing $\mu_{(B,Q,S)}$ [25],

$$\chi_{ijk}^{BQS}(T) = \frac{1}{VT^3} \frac{\partial^{(i+j+k)} \ln Z(T, V, \mu_B, \mu_Q, \mu_S)}{\partial(\mu_B/T)^i \partial(\mu_Q/T)^j \partial(\mu_S/T)^k}\bigg|_{\mu_{(B,Q,S)}=0} . \tag{1}$$

These cumulants can be calculated with standard Monte Carlo techniques.

Higher order cumulants, evaluated at vanishing values of the charge chemical potentials, enter as expansion coefficients in the calculation of the Taylor series for the QCD equation of state, i.e. the pressure (P) and its derivatives with respect to temperature and chemical potentials, which then gives the energy density and conserved charge densities, respectively,

$$\frac{P(T,\mu_B)-P(T,0)}{T^4}=\frac{\chi_2^B(T)}{2}\left(\frac{\mu_B}{T}\right)^2\left[1+\frac{1}{12}\frac{\chi_4^B(T)}{\chi_2^B(T)}\left(\frac{\mu_B}{T}\right)^2+\frac{1}{360}\frac{\chi_6^B(T)}{\chi_2^B(T)}\left(\frac{\mu_B}{T}\right)^4+\cdots\right]. \tag{2}$$

They also allow to calculate cumulants $\tilde{\chi}_{ijk}^{BQS}(T,\mu_B,\mu_Q,\mu_S)$ at non-zero values of the chemical potentials by setting up appropriate Taylor expansions, *e.g.* for $\mu_Q=\mu_S=0$,

$$\tilde{\chi}_{2n}^B(T,\mu_B) = \chi_{2n}^B(T)+\frac{1}{2}\chi_{2n+2}^B(T)\left(\frac{\mu_B}{T}\right)^2+\cdots\ , \tag{3}$$

$$\tilde{\chi}_{2n}^X(\mu_B,T) = \chi_{2n}^X(T)+\frac{1}{2}\chi_{2,2n}^{BX}(T)\left(\frac{\mu_B}{T}\right)^2+\cdots\ ,\quad X=Q,\ S\ . \tag{4}$$

Similar expansions, of course, are obtained also for non-zero values of $\mu_Q,\ \mu_S$, which allows to adjust these chemical potentials to take care of conditions met in heavy ion experiments, e.g. a strangeness neutral environment with a particular electric charge to baryon number ratio. Of particular interest for a comparison with experimental measurements of higher order cumulants at RHIC are Taylor expansions for ratios of cumulants, e.g. the ratios of 4^{th} to 2^{nd} order cumulants that give the kurtosis, $\kappa_{(B,Q,S)}$, times variance $\sigma^2_{(B,Q,S)}$ of distributions of net proton-number, strangeness or electric charge,

$$\kappa_B\sigma_B^2=\frac{\tilde{\chi}_4^B}{\tilde{\chi}_2^B}=\frac{\chi_4^B}{\chi_2^B}+\frac{1}{2}\left(\frac{\chi_6^B}{\chi_2^B}-\left(\frac{\chi_4^B}{\chi_2^B}\right)^2\right)\left(\frac{\mu_B}{T}\right)^2+O(\mu_B^4) \tag{5}$$

and similarly for the strangeness and electric charge kurtosis,

$$\kappa_X\sigma_X^2=\frac{\tilde{\chi}_4^X}{\tilde{\chi}_2^X}=\frac{\chi_4^X}{\chi_2^X}+\frac{1}{2}\left(\frac{\chi_{24}^{BX}}{\chi_2^X}-\frac{\chi_4^X\chi_{22}^{BX}}{(\chi_2^X)^2}\right)\left(\frac{\mu_B}{T}\right)^2+O(\mu_B^4)\ ,\quad X=Q,\ S\ . \tag{6}$$

Similar equations hold for the ratio of 3^{rd} and 2^{nd} order moments, which relate to the skewness of net proton-number, strangeness or electric charge distributions, $S_X\sigma_X=\tilde{\chi}_3^X/\chi_2^X$.

While lattice QCD results for the skewness and kurtosis of net baryon-number distributions have been obtained in next-to-leading order in μ_B^2, corresponding results for the strangeness and electric charge cumulant ratios, which also get measured at RHIC [26, 27], have not yet been obtained. Although results for the lower order cumulants have been presented by us [28, 38, 30], a reliable analysis of *e.g.* $\kappa_Q\sigma_Q^2$ requires calculations on larger lattices, as we will explain in the following.

2.2 Achievements, progress and status

Utilizing our past and continuing INCITE awards for the Titan supercomputer we have made significant progress [11, 12, 13] toward achieving the physics goals outlined before.

Non-Gaussian fluctuations of conserved charges: Higher order non-Gaussian cumulants of conserved charge fluctuations the become increasingly sensitive to the proximity of criticality, therefore, constitute good probes of the QCD critical point [31, 32, 33], as well as the pseudo-criticality at $\mu_B=0$ [4]. However, as described before, to obtain these susceptibilities for $\mu_B\neq 0$ one must further Taylor expand these susceptibilities in powers of μ_B,and compute even higher order susceptibilities. Our recent lattice QCD calculations [28, 12] of the non-Gaussian fluctuations of the net baryon-number seem to describe the experimental results of RHIC BES for beam energies $\sqrt{s}<19.6$ GeV. An example of these calculations is shown in Figure 2.

A direct comparison between the net proton-number fluctuations measured in a heavy ion experiment and the net baryon-number fluctuations calculated in QCD thermodynamics is difficult for many reasons. In order to establish these observables as a viable tool for detecting a critical point in the QCD phase diagram, it thus is important to firmly establish the well defined QCD baseline for these observables that can be obtained in lattice QCD calculations. This requires on the one hand to improve now existing results for the leading and next-to-leading order corrections to net baryon-number fluctuations and, on the other hand, obtain corresponding results for net-strangeness and electric charge fluctuations. The latter faces new challenges in numerical calculations that arise from the fact, that electric charge fluctuations are dominated by fluctuations of light pions, which in numerical calculations give rise to a strong volume dependence as well as severe discretization errors, which require a careful monitoring of the approach to continuum limit results.

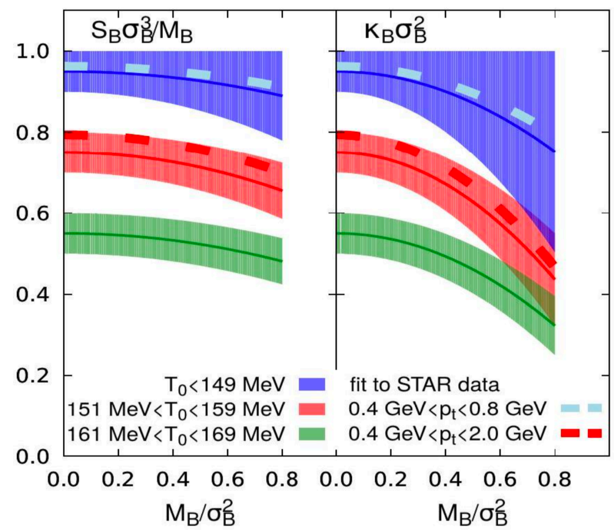

Figure 2. *Couple of ratios consisting of mean (M_B), variance (σ_B), skewness (S_B) and kurtosis (κ_B) of net baryon-number from recent lattice QCD calculations [12]. Dashed lines show the fits to experimental results from the STAR Collaboration for the corresponding ratios of net proton-number fluctuations.*

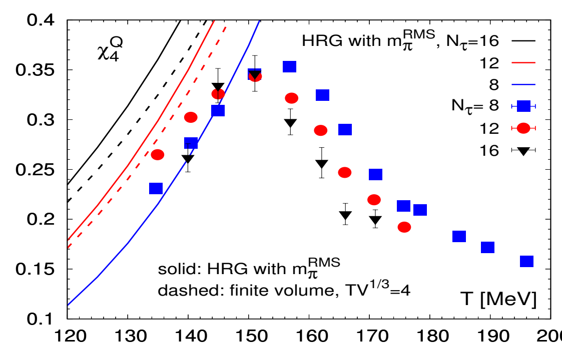

Figure 3. *Fourth order cumulant of net electric-charge fluctuations (HotQCD preliminary) calculated on lattices of size $N_\sigma^3 \times N_\tau$ with $N_\sigma = 4N_\tau$. Solid lines show results from hadron resonance gas (HRG) model calculations where the pion mass has been replaced by the RMS mass on a given size lattice. Dashed lines show the influence of finite volume effects.*

In corresponding calculations of NLO corrections of the expansion of skewness and kurtosis ratios for strangeness and electric charge distributions the complexity is similar to case of net baryon-number fluctuations. Additional problems, however, arise in the calculation of the leading order contributions. In the case of electric charge fluctuations the dominant contributions arise from fluctuations of light pion the dominant contributions arise from fluctuations of light pions. Statistical errors in a lattice QCD calculation thus are generally much easier to control. However, results are quite sensitive to finite volume effects [34] because of the large correlation length associate with the small pion mass, $\xi \sim 1/m_\pi$. Additional difficulties arise from the specific set-up used for most high statistics finite temperature calculations performed today with a physical spectrum of up, down and strange quarks. These calculations are performed in a so-called staggered fermion discretization scheme. Lattice QCD calculations with staggered fermions are fast, but have the disadvantage that the light pion spectrum is distorted by so-called taste violations. Only one pseudo-scalar is a Goldstone particle and thus is a light pion, while all other pseudo-scalar modes receive a large mass due to the explicit breaking of the $SU(2)_R \times SU(2)_L$ flavor symmetry in the staggered fermion formulation. This give rise to an average pion mass (root mean square (RMS) mass with approaches the physical pion mass value only in the limit of vanishing lattice-spacing, $m_\pi^{RMS} = \frac{1}{4}\sqrt{M_{\gamma_5}^2 + M_{\gamma_0\gamma_5}^2 + 3M_{\gamma_i\gamma_5}^2 + 3M_{\gamma_i\gamma_j}^2 + 3M_{\gamma_i\gamma_0}^2 + 3M_{\gamma_i}^2 + M_{\gamma_0}^2 + M_1^2} = m_\pi + O(a^2)$, where M_y denote various taste partner of the physical pion with mass m_π. For this reason the calculation of cumulants of electric charge fluctuations and their continuum extrapolation require lattices with large temporal extent and eventually also a large spatial volume. As can be seen in Figure 3, taste violations are significant in the low temperature hadronic phase of QCD and strongly reduce quartic cumulants of net electric-charge fluctuations. Large lattices are needed to control the continuum limit in this temperature range and firmly establish the deviations from hadron resonance gas models already at temperatures as low as 140 MeV, which currently available lattice QCD calculations seem to suggest. A similar, although less pronounced behavior also is found for strangeness fluctuations in the $|S| = 1$ sector, which is largely dominated by contributions from kaons. It is, thus, clear that controlled continuum extrapolations of non-Gaussian cumulants demands high-statistics computations with finer lattices with temporal extents $N_\tau \geq 16$.

Phase-boundary in the $T - \mu_B$ plane: The chiral crossover temperature, $T_c(\mu_B)$, of QCD is one of the most fundamental scale in the Standard Model of physics, and tells us the thermal condition when the Universe transitioned from a soup of quarks and gluons to the nuclear matter present today. Computations of the continuum-extrapolated value of $T_c(\mu_B = 0)$ for physical QCD is a highly noteworthy achievement of hot-dense lattice QCD calculations, and provides the hadronization temperature for all dynamical modeling of heavy-ion collisions. Recently, with the help of our INCITE awards we have made a more precise determination of continuum extrapolated $T_c(\mu_B = 0) = 156.5(2.0)$ MeV [13]. This needed precise computations of the chiral observables at $\mu_B = 0$ on finer lattices $N_\tau \geq 16$. The chiral crossover temperature for $\mu_B > 0$ also can be determined by Taylor-expanding the chiral observables [13]. It is important to note here that computation of the coefficients of the μ_B-series of the chiral-observables are significantly more challenging than that for pressure or energy density. Formally, computations of $O[(\mu_B/T)^n]$ corrections to chiral-observables are equivalent to computations of $O[(\mu_B/T)^{n+2}]$ corrections to the equation-of-state and, thus, demand even higher statistics.

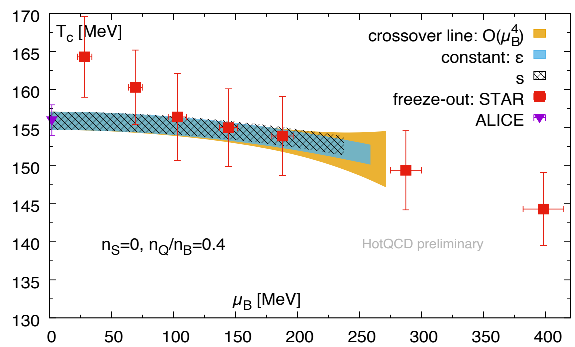

Figure 4. *QCD phase-boundary in the $T - \mu_B$ plane [13]. Also shown are lattice QCD for the the lines of constant energy and entropy densities [11] and the experimental results of the ALICE Collaboration at the LHC and the STAR Collaboration at RHIC for the freeze-out line extracted from particle yields.*

Figure 4 shows our recent results for $T_c(\mu_B)$ using lattices up to $N_\tau = 12$. These results need at least two fold improvements: (i) a more controlled continuum extrapolation, at least, for lower baryon densities, and (ii) extension to larger values of μ_B. The first can be achieved by high-statistics computations of, at least, up to 4$^{\text{th}}$ order Taylor-coefficients with finer $N_\tau \geq 16$ lattices. Presently, the second target only can be pursued with a extremely high-statistics computations of larger than 6$^{\text{th}}$ order Taylor-coefficients on a moderately fine, $N_\tau = 8$, lattice.

Constraining the location of the QCD critical point: By analyzing the radius-of-convergence of the Taylor-series in μ_B of the QCD partition function it is possible to constrain location of the singularity nearest to $\mu_B = 0$ axis, *i.e.* the location of the QCD critical point in the $\mu_B - T$ plane. For example, constraints on the location of the critical point can be obtained by analyzing the radius of convergence, $r_{2n}^\chi = |2n(2n-1)\chi_{2n}^B/\chi_{2n+2}^B|^{1/2}$, of the baryon-number susceptibility (χ_2^B), which diverges at the critical point. This shows that at least for high orders subsequent expansion coefficients need to rise like n^2 in order to obtain estimators for the radius of convergence that stay finite in the limit $n \to \infty$. Moreover, in order for the singularity that gives rise to a finite radius of convergence to also correspond to a phase transition point, one needs to find that all expansion coefficient stay positive beyond some order of the expansion.

In Figure 5 (Top) we the present estimates for the radius of convergence from our recent computations [12, 11] of up to 6th order Taylor-coefficients for lattice size 8×32^3 (lattice spacing ≈ 0.1 fm). The lower bounds of r_4^χ shown in Figure 5 are determined by the upper bounds of χ_6^B/χ_2^B, shown in Figure 5 (bottom). Firstly, improving the errors at higher temperatures and establishing beyond doubt that the ratio χ_6^B/χ_2^B is indeed negative for $T \gtrsim 150$ MeV would strengthen the assertion that a finite radius of convergence in this high temperature region does not correspond to a singularity at real values of the chemical potential and thus does not correspond to the searched for location of the critical point in the QCD phase diagram. Furthermore, it's obvious, more precise calculations of the 6th order and quantitative knowledge of the 8th Taylor-coefficients will provide more stringent constraints on the location of the critical point in the QCD phase diagram, vastly improving the present status. For example, assuming no change in the present central-values, an increase in the total statistics by $4\times$ will reduce the errors on the 6th by $2\times$ and also provide a quantitative estimate for upper bounds of 8th; leading to an increase in the lower bounds of convergence-radius estimator by $1.5\times$ to $\mu_B/T \approx 3$ and disfavoring the entire range of baryon chemical potential accessible to BES-II experiments.

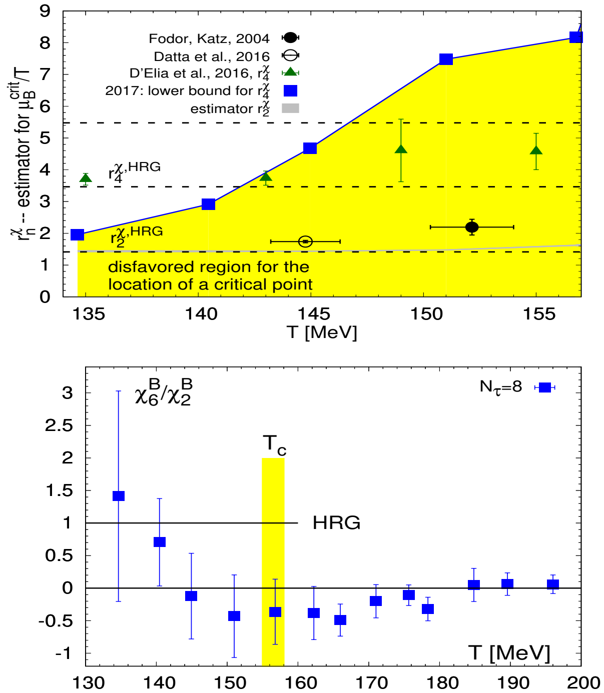

Figure 5. *(Top) Estimators for the radius-of-convergence of the Taylor-series for net baryon-number fluctuations (HotQCD preliminary). The blue squares depict the lower bounds for the estimator r_4^χ (see text). Dotted lines labeled with superscripts HRG indicate values of the corresponding estimators for a gas of hadrons without a critical point. (Bottom) The ratio of 6th to 2nd order Taylor-coefficients. The upper bounds of this ratio determines the lower bounds of r_4^χ.*

Equation-of-state for $\mu_B > 0$: The hydrodynamic modeling of matter formed in RHIC at varying beam energies demands the knowledge of the QCD equation-of-state (EoS) over the entire range of the QCD calculations at $\mu_B = 0$ is now well established [10] and is a standard component of state-of-the art hydrodynamic modeling of heavy ion collisions. A large range of μ_B/T explored by the BES is accessible to lattice QCD calculations using the Taylor expansions technique for pressure, energy density *etc.*. Recently, with the help of our INCITE awards, in Ref. [11] we have presented lattice QCD EoS from Taylor expansions up to $O[(\mu_B/T)^6]$. These lattice QCD EoS for $\mu_B > 0$ remains well-controlled for $\mu_B/T \lesssim 2$, suitable for dynamical modeling of RHIC's BES down to energies $\sqrt{s} \lesssim 14.5$ GeV. An independent lattice QCD study [35], relying on analytic continuation from purely imaginary chemical potentials, also arrived at very similar results and conclusions. However, for our present lattice QCD EoS results proper con-

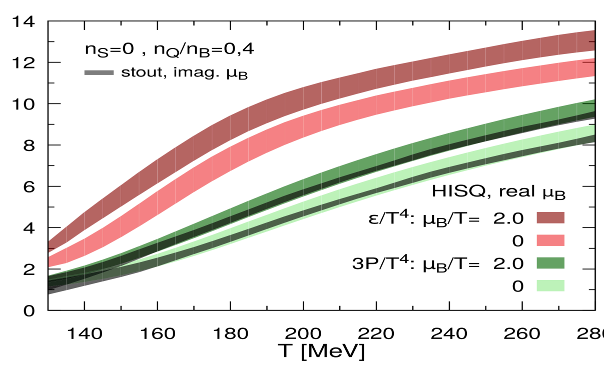

Figure 6. *Comparisons between pressure and energy density of QCD at $\mu_B = 0$ and $\mu_B/T = 2$. These results for QCD equation-of-state were obtained from lattice QCD calculations up to 6th order in μ_B [11].*

tinuum extrapolations were carried out only up to $O[(\mu_B/T)^2]$ using 3 or more successively finer lattice-spacing. For the higher $O[(\mu_B/T)^6]$ and $O[(\mu_B/T)^4]$ corrections we only have 'continuum estimates', which are based only on our 2 coarsest lattice-spacing corresponding to temporal extents $N_\tau = 6, 8$. To obtain the continuum extrapolated results for the QCD EoS at $\mu_B > 0$ precise, high-statistics calculations of the $O[(\mu_B/T)^4]$ and $O[(\mu_B/T)^6]$ corrections are required for finer lattices with temporal extents up to $N_\tau \geq 16$. However, to extend the computations of QCD EoS down up to $\mu_B/T \lesssim 3$, *i.e.* the full range of μ_B covered by BES=II, we have to reply on extremely high-statistics computations of larger than 6th order Taylor-coefficients on a moderately fine, $N_\tau = 8$, lattice.

Degrees of freedom near QCD crossover and freeze-out: In heavy-ion collision experiments the only tunable parameter is the collision energies of the heavy-ion beams, $\sqrt{s}$. However, to gain access to the information on the QCD phase diagram, this tunable parameter needs to be related to the thermodynamic variables, T and μ_B. Experimental explorations of pseudo-critical (crossover) behavior at LHC and search for the QCD critical point at RHIC BES-II will rely on proximity of the freeze-out line to the phase-boundary in the $T-\mu_B$ plane, see for example Figure 4. Thus, in addition to the phase-boundary, also reliable knowledge of the freeze-out line is crucial. Traditionally, the $\sqrt{s} \leftrightarrow (T, \mu_B)$ mapping has been done by comparing experimentally measured particle-yields with hadron resonance gas (HRG) models [36], and, therefore, need knowledge about the degrees of freedom in a hot-dense hadron gas in the vicinity of the QCD crossover.

Novel combinations of non-Gaussian cumulants also project onto partial pressures in different quantum number channels [30, 37, 38, 39], and, therefore, can be used to probe the degrees of freedom in the vicinity of the QCD crossover. Figure 7 depicts an example, clearly showing that above $T \gtrsim 145$ MeV (ideal) HRG fails to properly describe the degrees of freedom. However, also it is clear from the figure that large discretization artefacts are in play and to arrive at reliable conclusions continuum-extrapolated results using finer $N_\tau \geq 16$ lattices are needed. Furthermore, also it is possible to determine the freeze-out line by directly comparing the experimentally measured ratios non-Gaussian fluctuations of net charge with the corresponding lattice QCD results [40, 41]. However, as discussed before, reliable continuum extrapolation of higher order cumulants of net charge fluctuations also requires computations with finer $N_\tau \geq 16$ lattices.

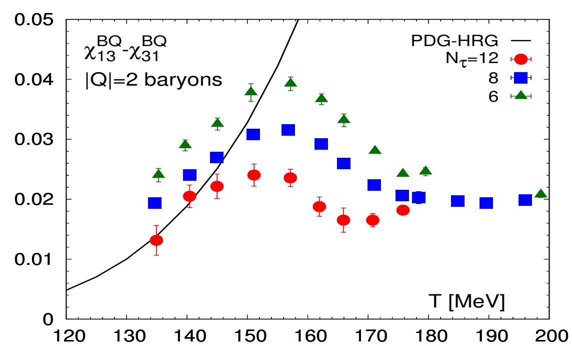

Figure 7. *Partial pressure (HotQCD preliminary) of the charge 2 baryons in the vicinity of the QCD crossover at $\mu_B = 0$, compared with that from HRG model.*

2.3 Planned computational campaigns and milestones

The calculation of high-order cumulants of net charge fluctuations probes the tails of the probability distributions for these fluctuation observables. They are thus extremely statistics-demanding. Experience gained with calculations of these observables on finite temperature lattices of size $N_s^3 \times N_\tau$, with $N_s/N_\tau = 4$ and $N_\tau = 6,\ 8,\ 12$ show that of the order of 100000 statistically-independent gauge-field configurations are needed to obtain Taylor expansion coefficients of 4th order with sufficiently small errors. An order of magnitude larger statistics is needed to add another even order expansion coefficient in these expansions or put stringent bounds on the location of a possible critical point in the QCD phase diagram. So far, about half a million independent gauge field configurations have been generated by us at several temperature values on lattices of size $32^3 \times 8$ that allowed us to constrain the region in the $T - \mu_B$ plane in which a critical point may be located. Clearly, present computing capabilities will not allow for a full-fledged continuum extrapolations for greater than 6th order cumulants; these calculations

target machine: Summit					
year	lattice size	task	# T-values	total # conf.	Summit node-hr
2019	$64^3 \times 16$	gauge-conf. generation	3	300K	320K
		calculation of cumulants	3	300K	290K
	$32^3 \times 8$	gauge conf. generation	2	2M	140K
		calculation of cumulants	2	2M	250K
2020	$64^3 \times 16$	gauge-conf. generation	3	300K	320K
		calculation of cumulants	3	300K	290K
	$32^3 \times 8$	gauge conf. generation	2	2M	140K
		calculation of cumulants	2	2M	250K
2021	$96^3 \times 24$	gauge-conf. generation	2	120K	670K
		calculation of cumulants	2	120K	330K

Table 1. *An overview of the computing campaigns planed during the next three years. For further details we refer to the Milestone Table.*

will be only limited moderately fine lattice-spacing, such as $N_\tau = 8$.

Based on the status of the research presented in Section 2.2 and the goals listed in Section 1 it is clear that we need to pursue a two-pronged approach:

- Controlled continuum extrapolations of up to 4$^{\text{th}}$ order cumulants using fine lattice-spacing corresponding to $N_\tau \geq 16$.
- Extremely high-statistics computations $\geq$ 6$^{\text{th}}$ order cumulants using a single moderately fine lattice-spacing corresponding to $N_\tau = 8$.

Overview of the computational campaigns and milestones are provided in Table 1.

3 COMPUTATIONAL READINESS

3.1 Computational approach

In this project we will perform calculations for QCD with two degenerate light quarks and a heavier strange quark. All quark masses have been tuned to their physical values using zero temperature calculations performed by us in the past. The numerical investigations utilize the Highly Improved Staggered Quark (HISQ/tree) discretization scheme in the fermion sector. This discretization scheme is known to have the smallest taste breaking effects at non-vanishing lattice spacing among the staggered type quark actions. In the gluon sector a Symanzik improved gauge action is used that eliminates discretization errors of $O(a^2)$ at the tree level. This set-up has been used by us for several QCD thermodynamics calculations. The gauge field configurations are generated by using a RHMC algorithm with a three-scale integrator which naturally profits from the Hasenbusch trick. The integrator is of Omelyan type on the largest scale and uses leapfrog on the two smaller ones. The inverter is a multi-shift Conjugate Gradient (CG). The RHMC program has been optimized by us for the use on NVIDIA GPU cards.

Our analysis programs developed for the calculation of higher order Taylor series expansion coefficients have been tuned for GPUs as well as CPUs. We currently use this software package on machines equipped with NVIDIA GPUs, including Titan at ORNL, the V100 cluster in Wuhan and an older GPU-cluster in Bielefeld, equipped with K20 cards. The analysis software greatly profits from a CG inverter that is optimized for the inversion of sparse (fermion) matrices using several right hand sides [42, 43], which has been developed by M. Wagner (NVIDIA) while he was a post-doc at Bielefeld University and in a similar version now also is part of the QUDA library of NVIDIA. This inverter gives a 4-fold increase in speed over the single right hand side inverter. Furthermore, the inverter uses deflation as a preconditioner which reduces the number of conjugate gradient steps in the matrix inversion by a factor five to eight, depending on the number of eigenvectors for small eigenmodes that one can afford to store. For this purpose we calculate $O(200)$ of the lowest eigenvalues of the fermion matrix and deflate the Krylov space to be explored by the Conjugate Gradient inverter from the eigenspace of the corresponding eigenvectors by means of a Galerkin projection. The low-lying eigenvalues and the corresponding eigenvectors are calculated using the Kalkreuter-Simma algorithm. To estimate the traces of the inverse of fermion matrices that enter the set of relevant observables we use stochastic estimators. This requires $O(1000)$ matrix inversions with different sources on identical gauge field configurations, which introduces a high degree of parallelism in our analysis.

The computational campaigns performed in this project consist of two parts:

(i) We will generate data sets for several values of the temperature on lattices of sizes $N_\sigma^3 \times N_\tau$ with $N_\tau = 8,\ 16,\ 24$ and spatial lattice sizes $N_\sigma/N_\tau = 4$ using our RHMC program for (2+1)-flavors with the strange quark mass tuned to its physical value and the two light quark masses taken to be degenerate and tuned to a physical value for the light Goldstone pion mass. After thermalization, gauge field configurations generated with the RHMC program will be stored after every 10^{th} RHMC trajectory of unit length for future analysis of other observables.

(ii) We will analyze these data by calculating the required expectation values which involves the inverse of Dirac matrices D and its derivatives up to n^{th} order with respect to the quark chemical potentials. This requires the calculation of operators like $\mathrm{Tr}[D^{-1}D_iD^{-1}D_j..]$, where D_i denotes the i^{th} derivative of D with respect to a quark chemical potential. These $O(20)$ basic operators then get combined to construct all relevant n^{th} order cumulants of $(B,\ Q,\ S)$ fluctuations. This requires multiple inversions of the large sparse matrix, D^{-1}, which we do using a deflated CG algorithm, that utilizes $O(200)$ precalculated eigenvectors of the Dirac matrix D.

Both programs required for both the parts are ready to be used on GPUs as well as CPUs, they have been optimized and tested for the V100 GPUs used in Summit. For lattices up to size $64^3 \times 16$ the generation as well as the analysis of gauge field configurations can be done on single GPUs. Typical production code used in (i), however, will run on two GPUs per stream. We will generate gauge field configurations in multiple streams simultaneously for different values of the temperature. Calculations in part (ii) are independently performed on each of the gauge field configurations. The calculation of various observables can be started once a sufficient number of configurations are generated. This typically will result in focused analysis campaigns performed on large numbers of GPUs in parallel.

In parallel to the production of $64^3 \times 16$ gauge field configurations on the GPUs we will run most of our production of $32^3 \times 8$ configurations on the Summuit CPU. The time needed to generate a new configuration is in this case about a factor 16 faster which makes synchronization of the runs for small and large lattices easy.

3.2 Code performance

We have ported our highly optimized software-suite for Titan to Summit. This software-suite is based on C++, CUDA, OpenMP and MPI. The most important ingredient in a calculation of high-order cumulants is an efficient, highly optimized CG routine. On one Volta GPU of the Summit machine our CUDA-based GPU code can carry out inversions for up to 11 right hand sides, and our OpenMP threaded C++ CPU code can accommodate up to 64 right hand sides for the IBM POWER 9 processor of a Summit node. Our HISQ CG inverter achieves a performance of 2 TFLOP/s on a single Summit-GPU, which is about 10 times higher than that for Titan. The performance of our HISQ CG inverter for a single Summit-CPU is about 500 GFlop/s, which about 8 times higher than that for Titan. Our CPU- and GPU-codes are integrated together through intra-node OpenMP threading and can be executed simultaneously, taking full advantage of the CPU-GPU heterogeneous architecture.

The generation of gauge field configurations, using a Rational Hybrid Monte Carlo (RHMC) algorithm is dominated to 70% by a memory-bandwidth-bound multi-shift CG using a single RHS. The remaining 30% are spend inside cache-bandwidth-bound force calculations. For a single Summit-GPU our RHMC runtime is about (6-11) times lesser than for Titan, with the lower number referring to our smaller lattices and the larger number referring to our large lattices, which take the largest fraction of the computing time requested for this project.

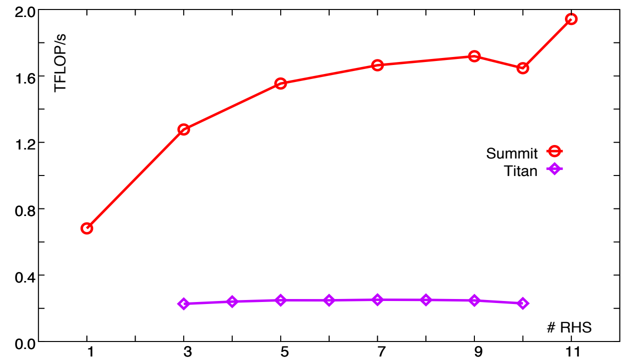

Figure 8. *Performance of our HISQ CG inverter for multiple right hand sides (RHS) on one Volta GPU.*

On 1000 Summit nodes performance of our code reaches reaches about 12 PFlop/s, which is about 50 times greater than for Titan and amounts to about 28% of the theoretical peek-speed of Summit.

3.3 Parallel performance

When running production code on a large number of nodes our gauge field generation campaigns generate many independent data streams. Deviations from perfect scaling arise from the fact that a larger number of initial thermalization/decorrelation runs are needed to make the production in different streams truly independent. The resulting scaling of our code for the production campaigns planned for this project is shown in Figure 9.

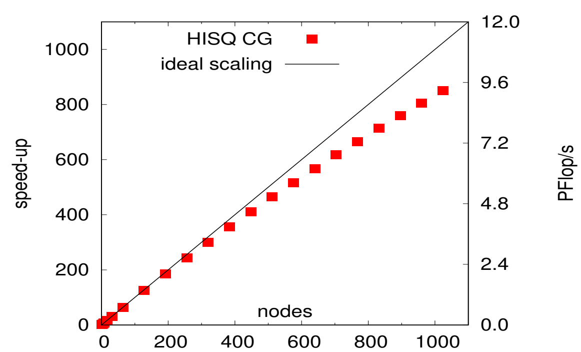

Figure 9. *Speed-up of our production code running on up to 1000 Summit nodes using 2 GPUs per stream. In this figure we assume that we will produce 100K independent gauge field configurations per temperature value. It takes into account 10% overhead per stream that arises from the larger number of streams, that is used when running production code on larger numbers of nodes, as well as a 10% reduction in performance over the single GPU performance of our code.*

3.4 Developmental work

During the next years we will continue to improve our software. The PI of this proposal is also the Co-PI of the Scientific Discovery through Advanced Computing (SciDAC) award “Computing the Properties of Matter with Leadership Computing Resources”, funded by Office of Nuclear Physics and Office of Advanced Scientific Computing, US Department of Energy for the period 2017-2021; and leads the software development efforts for hot-dense lattice QCD. With the help of SciDAC award, combining BNL and Michigan State University, efforts of 2 postdoctoral FTEs and 1 graduate student and 1 FTE will be fully devoted to hot-dense lattice QCD software developments for Summit-like pre-exascale systems, as well as for yet-to-come exascale systems. These efforts helped us to quickly port our Titan code for Summit, which already shows quite good performance— on 1000 Summit-nodes performance of our code is about 50 times greater than that for Titan. We will continue to improve and optimize our code. In particular, we will focus on optimizations of the multi-GPU performance of our code within a single Summit-node. We also will develop a multi-GPU code with intra-node communication, which we plan to use for the generation of the large lattices we are going to produce in year

3. So far no optimization specific to our code suite has been attempted for the 6-GPU configuration of the Summit-nodes. Furthermore, our software is being developed in collaboration with the lattice QCD group at Bielefeld University, Germany, and Central China Normal University in Wuhan, China. Both these groups have their own GPU-clusters and posses significant expertise in GPU-software developments.

3.5 Computing time requirements

Computing cost: Based on the realistic benchmarks of our code for Summit presented in the previous section:

- $32^3 \times 8$: Generation of one gauge-configuration, separated by 10 RHMC trajectories of unit length, takes ~ 0.07 Summit-node-hr. Measurements of all the operators needed for up to 10$^{\text{th}}$ order cumulants on each gauge-configuration takes ~ 0.125 Summit-node-hr.
- $64^3 \times 16$: Generation of one gauge-configuration, separated by 10 RHMC trajectories of unit length, takes ~ 1.06 Summit-node-hr. Measurements of all the operators needed for up to 6$^{\text{th}}$ order cumulants on each gauge-configuration takes ~ 0.96 Summit-node-hr.
- $96^3 \times 24$: Generation of one gauge-configuration, separated by 10 RHMC trajectories of unit length, takes ~ 5.6 Summit-node-hr. Measurements of all the operators needed for up to 4$^{\text{th}}$ order cumulants on each gauge-configuration takes ~ 2.75 Summit-node-hr.

Based on the timings presented and our run-plan outlined before, Table 1 show the breakdown of computational cost for the proposed project.

Data requirements: In this project we generate data sets, so-called gauge field configurations, that need to be stored for further processing, i.e. the calculation of various observables that are part of this project and future calculations that can utilize these data set. We plan to generate gauge field configurations on three different size lattices, $32^3 \times 8$, $64^3 \times 16$ and $96^3 \times 24$. The three different types of data sets require 0.05 GB, 0.8 GB and 4 GB per configuration, respectively. We plan to generate and store 2M, 600K and 120K configurations of the different sizes. In total we thus will produce $(100+500+480)$ TB $\simeq 1.1$ PB of data, which need to be archived. During the computing and analysis campaigns planned in each year 500 TB need to remain available throughout the year. The amount of data will rise linearly over the first 3 quarters of each year. We will start archiving analyzed data sets in the second half of each year. We imagine that the data will need to stay archived for about 5 years, before a new generation of data gathered on new compute hardware will make their re-use uninteresting.

At some point the data sets generated as part of this project will also be used by other members of the HotQCD collaboration and the international lattice QCD community. We thus expect that the data transfer out of ORNL will amount to be of the order of 30TB per month. Incoming data flux will be significantly smaller.

Data sharing:

- The data sets generated on the large $64^3 \times 16$ and $96^3 \times 24$ lattices will also be used in other HotQCD projects. In particular, they will be used to update our results on the QCD equation of state, to calculate screening masses of light and heavy hadrons as well as the charm contribution to fluctuation observables, and to analyze the temperature dependence of topological charge fluctuations.
- Data transfer to other LCFs and USQCD cluster installations at BNL, JLAB, and Fermilab will use GLOBUS.
- Upon request from other user groups we will transfer these data from ORNL.

3.6 Use of resources requested

During the first two years we plan to perform in parallel two production campaigns on two different lattice sizes, $32^3 \times 8$ and $64^3 \times 16$. As shown in the Milestone Table these production campaigns will be spread linearly of the first 9 months of year 1 and 2. The analysis of gauge field configurations generated in these runs will presumably start in the 2nd quarter of each year and continue until the end of the year. Depending on the load of the machine we intend to finish production of our data sets as quickly as possible.

Production jobs will utilize 1000 nodes per run. These jobs will utilize all 6 GPUs of a node simultaneously as well as the Summit CPU. Each run will generate about 1 TB output after 6 hours. Typical run time requests for the production jobs will be between 8 and 20 hours, *i.e.* they will require 8,000-20,000 node-h.

The analysis jobs can run be run in the same mode, i.e. utilizing 1000 nodes for about 10 hours. In case it is of advantage for machine utilization we can run the analysis also in dedicated campaigns on the entire machine, once about 30,000 data sets have been generated.

In year 3 we will generate the first data sets on a large $96^3 \times 24$ lattice. The number of configurations that can be generated in this campaign will be sufficient to calculate up to 4^{th} order cumulants. Also these calculations will utilize 1000 nodes per job.

Bibliography

[1] M. Asakawa, K. Yazaki, Nucl. Phys. A **504**, 668 (1989);
A. M. Halasz, A. D. Jackson, R. E. Shrock, M. A. Stephanov, J. J. M. Verbaarschot, On the phase diagram of QCD, Phys. Rev. D **58**, 096007 (1998). Nucl. Phys. A **967**, 453 (2017)

[2] D. Geesaman et al, The 2015 Long Range Plan for Nuclear Science

[3] D. Keane, The Beam Energy Scan at the Relativistic Heavy Ion Collider, J. Phys. Conf. Ser. **878**, 012015 (2017).

[4] B. Friman, F. Karsch, K. Redlich and V. Skokov, Fluctuations as probe of the QCD phase transition and freeze-out in heavy ion collisions at LHC and RHIC, Eur. Phys. J. C **71** (2011) 1694 [arXiv:1103.3511 [hep-ph]].

[5] A. Rustamov [ALICE Collaboration], Net-baryon fluctuations measured with ALICE at the CERN LHC,

[6] M. M. Aggarwal *et al.* [STAR Collaboration], Higher Moments of Net-proton Multiplicity Distributions at RHIC, Phys. Rev. Lett. **105**, 022302 (2010).

[7] X. Luo [STAR Collaboration], Energy Dependence of Moments of Net-Proton and Net-Charge Multiplicity Distributions at STAR, PoS CPOD **2014**, 019 (2015).

[8] P. Braun-Munzinger, A. Kalweit, K. Redlich and J. Stachel, Confronting fluctuations of conserved charges in central nuclear collisions at the LHC with predictions from Lattice QCD, Nucl. Phys. A **956**, 805 (2016)

[9] Beam Energy Scan Theory Collaboration (BEST)

[10] A. Bazavov *et al.* [HotQCD Collaboration], Equation of state in (2+1)-flavor QCD, Phys. Rev. D **90**, 094503 (2014) [arXiv:1407.6387 [hep-lat]].

[11] A. Bazavov *et al.*, The QCD Equation of State to $\mathcal{O}(\mu_B^6)$ from Lattice QCD, Phys. Rev. D **95**, 054504 (2017) [arXiv:1701.04325 [hep-lat]].

[12] A. Bazavov *et al.* [HotQCD Collaboration], Skewness and kurtosis of net baryon-number distributions at small values of the baryon chemical potential, Phys. Rev. D **96**, 074510 (2017) [arXiv:1708.04897 [hep-lat]].

[13] P. Steinbrecher, talk at Quark Matter 2018, Venice, Italy. (Manuscript under preparation.)

[14] H. T. Ding, F. Karsch and S. Mukherjee, Thermodynamics of strong-interaction matter from Lattice QCD, Int. J. Mod. Phys. E **24**, no. 10, 1530007 (2015) [arXiv:1504.05274 [hep-lat]].

[15] R. A. Soltz, C. DeTar, F. Karsch, S. Mukherjee and P. Vranas, Lattice QCD Thermodynamics with Physical Quark Masses, Ann. Rev. Nucl. Part. Sci. **65**, 379 (2015) [arXiv:1502.02296 [hep-lat]].

[16] G. Aarts, Developments in lattice quantum chromodynamics for matter at high temperature and density, Pramana **84**, no. 5, 787 (2015) [arXiv:1312.0968 [hep-lat]].

[17] G. Aarts, Complex Langevin dynamics and other approaches at finite chemical potential, PoS LATTICE **2012**, 017 (2012) [arXiv:1302.3028 [hep-lat]].

[18] M. Cristoforetti *et al.* [AuroraScience Collaboration], New approach to the sign problem in quantum field theories: High density QCD on a Lefschetz thimble, Phys. Rev. D **86**, 074506 (2012) [arXiv:1205.3996 [hep-lat]].

[19] M. Cristoforetti, F. Di Renzo, A. Mukherjee and L. Scorzato, Monte Carlo simulations on the Lefschetz thimble: Taming the sign problem, Phys. Rev. D **88**, no. 5, 051501 (2013) [arXiv:1303.7204 [hep-lat]].

[20] Z. Fodor, S. D. Katz, D. Sexty and C. TÃűrÃűk, Complex Langevin dynamics for dynamical QCD at nonzero chemical potential: A comparison with multiparameter reweighting, Phys. Rev. D **92**, no. 9, 094516 (2015) [arXiv:1508.05260 [hep-lat]].

[21] P. de Forcrand and O. Philipsen, The QCD phase diagram for small densities from imaginary chemical potential, Nucl. Phys. B **642**, 290 (2002) [hep-lat/0205016].

[22] M. D'Elia and M. P. Lombardo, Finite density QCD via imaginary chemical potential," Phys. Rev. D **67**, 014505 (2003) [hep-lat/0209146].

[23] C. R. Allton, S. Ejiri, S. J. Hands, O. Kaczmarek, F. Karsch, E. Laermann and C. Schmidt, The Equation of state for two flavor QCD at nonzero chemical potential, Phys. Rev. D **68**, 014507 (2003) [hep-lat/0305007].

[24] R. V. Gavai and S. Gupta, Pressure and nonlinear susceptibilities in QCD at finite chemical potentials, Phys. Rev. D **68**, 034506 (2003) [hep-lat/0303013].

[25] C. R. Allton, M. Doring, S. Ejiri, S. J. Hands, O. Kaczmarek, F. Karsch, E. Laermann and K. Redlich, Thermodynamics of two flavor QCD to sixth order in quark chemical potential, Phys. Rev. D **71**, 054508 (2005) [hep-lat/0501030].

[26] L. Adamczyk *et al.* [STAR Collaboration], Beam energy dependence of moments of the net-charge multiplicity distributions in Au+Au collisions at RHIC, Phys. Rev. Lett. **113**, 092301 (2014) [arXiv:1402.1558 [nucl-ex]].

[27] L. Adamczyk *et al.* [STAR Collaboration], Collision Energy Dependence of Moments of Net-Kaon Multiplicity Distributions at RHIC, arXiv:1709.00773 [nucl-ex].

[28] A. Bazavov *et al.*, Curvature of the freeze-out line in heavy ion collisions, Phys. Rev. D **93**, no. 1, 014512 (2016) [arXiv:1509.05786 [hep-lat]].

[29] A. Bazavov *et al.*, Additional Strange Hadrons from QCD Thermodynamics and Strangeness Freezeout in Heavy Ion Collisions, Phys. Rev. Lett. **113**, no. 7, 072001 (2014) [arXiv:1404.6511 [hep-lat]].

[30] A. Bazavov *et al.*, Strangeness at high temperatures: from hadrons to quarks, Phys. Rev. Lett. **111**, 082301 (2013) [arXiv:1304.7220 [hep-lat]].

[31] M. A. Stephanov, Non-Gaussian fluctuations near the QCD critical point, Phys. Rev. Lett. **102**, 032301 (2009) [arXiv:0809.3450 [hep-ph]].

[32] C. Athanasiou, K. Rajagopal and M. Stephanov, Using Higher Moments of Fluctuations and their Ratios in the Search for the QCD Critical Point, Phys. Rev. D **82**, 074008 (2010) [arXiv:1006.4636 [hep-ph]].

[33] M. A. Stephanov, On the sign of kurtosis near the QCD critical point, Phys. Rev. Lett. **107**, 052301 (2011) [arXiv:1104.1627 [hep-ph]].

[34] F. Karsch, K. Morita and K. Redlich, Effects of kinematic cuts on net-electric charge fluctuations, Phys. Rev. C **93**, no. 3, 034907 (2016) [arXiv:1508.02614 [hep-ph]].

[35] J. N. Guenther, R. Bellwied, S. Borsanyi, Z. Fodor, S. D. Katz, A. Pasztor, C. Ratti and K. K. SzabÃş, The QCD equation of state at finite density from analytical continuation, Nucl. Phys. A **967**, 720 (2017) [arXiv:1607.02493 [hep-lat]].

[36] A. Andronic, P. Braun-Munzinger, K. Redlich and J. Stachel, The thermal model on the verge of the ultimate test: particle production in Pb-Pb collisions at the LHC, J. Phys. G **38**, 124081 (2011) [arXiv:1106.6321 [nucl-th]].

[37] A. Bazavov *et al.*, The melting and abundance of open charm hadrons, Phys. Lett. B **737**, 210 (2014) [arXiv:1404.4043 [hep-lat]].

[38] A. Bazavov *et al.*, Additional Strange Hadrons from QCD Thermodynamics and Strangeness Freezeout in Heavy Ion Collisions, Phys. Rev. Lett. **113**, no. 7, 072001 (2014) [arXiv:1404.6511 [hep-lat]].

[39] S. Mukherjee, P. Petreczky and S. Sharma, Charm degrees of freedom in the quark gluon plasma," Phys. Rev. D **93**, no. 1, 014502 (2016) [arXiv:1509.08887 [hep-lat]].

[40] A. Bazavov *et al.*, Freeze-out Conditions in Heavy Ion Collisions from QCD Thermodynamics, Phys. Rev. Lett. **109**, 192302 (2012) [arXiv:1208.1220 [hep-lat]].

[41] S. Mukherjee and M. Wagner, Deconfinement of strangeness and freeze-out from charge fluctuations, PoS CPOD **2013**, 039 (2013) [arXiv:1307.6255 [nucl-th]].

[42] S. Mukherjee, O. Kaczmarek, C. Schmidt, P. Steinbrecher and M. Wagner, HISQ inverter on Intel$^{\circledR}$ Xeon Phi^{TM} and NVIDIA$^{\circledR}$ GPUs, PoS LATTICE **2014**, 044 (2015) [arXiv:1409.1510 [cs.DC]].

[43] O. Kaczmarek, C. Schmidt, P. Steinbrecher and M. Wagner, Conjugate gradient solvers on Intel Xeon Phi and NVIDIA GPUs, arXiv:1411.4439 [physics.comp-ph].