PROJECT NARRATIVE

We request a three-year allocation on the Oak Ridge Leadership Computing Facility's IBM/NVIDIA supercomputer, Summit, and on the Argonne Leadership Computing Facility's Cray/Intel supercomputer, Theta, in order to address fundamental questions in high-energy and nuclear physics. The proposed calculations directly support many large, world-wide experimental efforts in these fields. This proposal is submitted on behalf of the USQCD Collaboration, which consists of nearly all of the high-energy and nuclear physicists in the United States working on the numerical study of lattice field theories, compromising approximately 150 Ph.D. level scientists. It is naturally organized into four separate thrusts: *quark and lepton flavor physics*, *lattice gauge theories beyond the Standard Model*, *hadrons and nuclei*, and *hot-dense qcd*. They share common underlying methodology as well as software that the USQCD Collaboration has steadily developed over two decades. Three of the thrusts focus on quantum chromodynamics (QCD) while one explores other lattice gauge theories that might play a role in the electroweak sector of the Standard Model of particle physics.

This proposal is a successor to a previous three-year USQCD proposal (PI: Paul Mackenzie) that ran on Titan, Mira, and Theta. With this support, the USQCD team carried out numerous calculations of unprecedented precision and scope. Some highlights include sub-percent determination of the quark masses (fundamental parameters of nature), transformational development in calculations of hadron and nuclear structure, the inspiration for a beam-energy scan at heavy-ion colliders, and the emergence of a light scalar bound state with properties like that of the Higgs boson in certain non-QCD gauge theories. All four thrusts have experience with Summit via 2019 INCITE allocations (flavor physics, PI Norman Christ; hot-dense QCD, PI Swagato Mukherjee), early science running, and/or the exascale computing project.

The PI and co-PIs are the members of the USQCD Executive Committee (EC), with the PI being the current Chair. The calculations proposed and the balance among them represent the EC's careful consideration of the interests of the lattice-gauge-theory community and the needs of the high-energy and nuclear experiments that rely on lattice calculations. A summary of the request, split among the four thrusts, is given in Table 1. Part of the thrust *hadrons and nuclei* is split off and labeled "clover." Here, we propose to generate several ensembles of lattice gauge-field configurations with 2 + 1 flavors of sea quark. These ensembles will enable many calculations not proposed here over the coming decade. The other thrusts also generate ensembles, but they are either more specialized or a comparatively modest extension of already existing data.

Table 1. Summary of requests for Summit (in node-h) and Theta (in node-h). Here, "BSM" stands for "lattice gauge theories beyond the Standard Model"

Year	Machine	Flavor	BSM	Clover	Nuclear	Hot dense	Total
2020	OLCF Summit	1,195,000	300,000	542,000	469,000	563,000	3,069,000
	ALCF Theta	3,125,000	1,562,500	0	0	0	4,687,500
2021	OLCF Summit	1,206,000	300,000	333,000	524,000	563,000	2,926,000
	ALCF Theta	3,125,000	1,562,500	0	0	0	4,687,500
2022	OLCF Summit	1,199,000	300,000	0	943,000	563,000	3,005,000
	ALCF Theta	3,125,000	1,562,500	0	0	0	4,687,500
Totals	OLCF Summit	3,600,000	900,000	875,000	1,936,000	1,689,000	9,000,000
	ALCF Theta	9,375,000	4,687,500	0	0	0	14,062,500

The goals of this proposal are well-aligned with the strategic plans of the high-energy and nuclear communities in the U.S. These plans have been set out in the report of the Particle Physics Project Prioritization Panel (P5, in the case of high-energy physics) and the long-range plan of the Nuclear Science Advisory Committee. In this way, the calculations proposed here leverage enormous investments in experimental physics (≈ $700M for high-energy physics and ≈ $600M for nuclear physics). Several experimental measurements cannot be interpreted without the calculations discussed in this proposal. Mnay of the calculations look to the future: for example, calculations of parton distribution functions, which will (in the exascale era) inform future running at the Large Hadron Collider and will sharpen the case for the Electron-Ion Collider.

For convenience, the next two sections of this proposal is organized around the four thrusts.

1 SIGNIFICANCE OF RESEARCH

1.1 Quark- and lepton-flavor physics

1.1.1 Heavy-quark flavor: B-meson decays

For the past decade, lattice QCD flavor-physics calculations have made critical contributions to determinations of CKM matrix elements (fundamental parameters of the Standard Model (SM) that control transitions between quark flavors) and to searches for beyond-the-SM effects in flavor physics. In particular, our recent results for the leptonic B, B_s, D, and D_s decay constants [1, 2] with crucial support from USQCD INCITE awards have sub-percent errors and are 3 to 10 times more precise than any previous lattice-QCD calculation. Semileptonic decays of these mesons permit a still more precise determination of the CKM matrix elements. From the Belle II and LHC experiments, we will soon have much more accurate experimental measurements of these processes, which demands timely, commensurate precision improvements in theory. We focus here on the form factors for semileptonic decays, which are needed for $|V_{ub}|$, $|V_{cb}|$, $|V_{cd}|$, and $|V_{cs}|$ determinations and for rare decay phenomenology. In order to meet the experimental precision goals, we will use the same lattice actions and methodology as in our recent high-precision calculation of the leptonic decay constants.

1.1.2 Lepton flavor: muon $g_\mu - 2$

The anomalous magnetic moment of the muon $a_\mu = (g_\mu - 2)/2$ is measured in experiment to 0.54 ppm. A 2006 measurement [3] disagrees with SM predictions at the level of 3.5 to 4σ, which has generated considerable excitement over the prospects of discovering new physics. A major experimental effort is underway at Fermilab to reduce the measurement uncertainties by a factor of four. Likewise, an urgent theoretical effort is underway to achieve the essential commensurate improvement in the Standard-Model predictions where the largest uncertainty comes from QCD effects [4, 5, 6, 7, 8, 9, 10, 11, 12, 13]. Accordingly, we have embarked on a multi-year effort to calculate the hadronic vacuum polarization (HVP) and hadronic light-by-light (HLbL) scatttering contributions to a_μ to high precision. At present our HVP results are good to ≈ 2% accuracy [10, 11]. Our goal with help from the proposed calculations is to reduce the uncertainty in the HVP contribution to below 0.5% (which would be competitive with the best phenomenological methods) and the smaller HLbL part to 10%.

1.1.3 Second-order weak processes

Rare phenomena that occur only at second order in the weak interactions of the SM provide important opportunities to discover new phenomena that lie outside the SM. However, the SM predictions for these processes must be known. Here we propose to study two such phenomena: the mass difference, ΔM_K between the long- and short-lived neutral K mesons and the long distance (LD) contribution to the measure

of indirect CP violation, ϵ_K. Both ΔM_K and ϵ_K involve a strangeness change by two units and hence in the SM are very small second-order effects. At 10% precision, ΔM_K is sensitive to physics at an energy scale as high as 10^4 TeV. The major SM contribution to ΔM_K comes from the charm energy scale where lattice QCD is needed to improve the 36% errors found in a NNLO perturbative calculation [14]. Likewise the LD part of ϵ_K is required to test the SM prediction at a precision of $\lesssim 5\%$ [15]. Lattice QCD is the only approach able to determine ΔM_K and ϵ_K with controlled systematic errors [16, 17, 18].

1.2 Lattice gauge theories beyond the Standard Model

In spite of the dramatic discovery of the Higgs boson and the increasingly accurate tests of the SM at the Large Hadron Collider (LHC), it is generally acknowledged that the SM must be a low energy effective limit of a yet undiscovered more fundamental theory. Missing ingredients of the SM include a "natural" mechanism to stabilize the mass of the "elementary" Higgs scalar, the absence of a candidate for dark matter (DM), an explanation of the matter asymmetry of the Universe, and of course a connection to quantum gravity theory. For each of these problems, theorists have posited new strong gauge dynamics, for which lattice simulations are an appropriate method to study their non-perturbative dynamics. Our proposed beyond the SM (BSM) efforts seek to advance our understanding of possible strongly-coupled extensions to the SM for composite Higgs and composite DM sectors.

We should note that the BSM lattice community has broader range of projects including for conformal field theories (CFT), symmetry breaking in supersymmetric (SUSY) theories, a range of AdS/CFT studies to probe quantum gravity as the dual of strong coupling gauge theory. The choice of composite Higgs and DM theories in this proposal is made because of the potential of current INCITE resources to clarify the underlying dynamics and make an impact on phenomenology.

1.2.1 Strongly-Coupled Composite Higgs Dynamics

Over the last few years there is now strong evidence [19] that gauge theories with a large number of fermions do exhibit a dynamical mechanism near to the conformal window (CW), which stabilizes a parametrically low mass composite scalar that might be able to play the role of the SM Higgs. The precise nature of the mechanism is poorly understood. The first project seeks to study two candidate SU(3) gauge theories with sextet vs fundamental fermions respectively and to parametrized the low energy effective theory of the Higgs-like scalar, the Nambu-Goldstone modes required give mass to the W and Z bosons, and further low-lying resonances. Characterizing the effective theory with substantial lattice simulation can either eliminate the model or guide experimental investigations.

For these theories, our focus is on the CW and its extensions to the near-conformal paradigm. There are three very interesting scenarios we propose to investigate: i.) Spontaneous conformal breaking combined with chiral symmetry breaking χSB driven by near-conformal behavior giving rise to an appropriate *dilaton EFT* as the CW is approached [20, 21]. The near-conformal dilaton EFT is built from the expansion around the conformal infra-red fixed point when the number of flavors exceeds its critical value. ii.) In the absence of spontaneous scale symmetry breaking, an appropriate *chiral EFT* requires the inclusion of a light σ-particle which may reach the linear σ-model limit for very small pion masses [22]. iii.) A new paradigm, based on the interpretation of the walking β-function as pinched between a conjugate pair of complex conformal fixed points [23, 24], will require the development of its infrared EFT derived from nearby complex CFT fixed points. The choice of two models balances the need for large resources on both while making it possible to distinguish generic from model-dependent features.

1.2.2 Composite Dark Matter Early Universe Phase Transitions

Viable composite DM models have been established by lattice calculations [25, 26, 27, 28]. These models generically deconfine at high temperature and the phase transitions can be strongly first order [29, 30], with consequences for the evolution of early universe. If the thermal phase transition is first order, then the transition is violent, with bubbles of the low-temperature phase nucleating, expanding, and colliding with one another. The rapid changes in matter-energy density during such a phase transition can act as a source for gravitational waves [31, 32]. Also, strong first-order transitions are one of the Sakharov conditions [33] for generating matter asymmetry. Lattice BSM studies leverage decades of experience in QCD thermodynamics to compute the relevant quantities in DM models, such as the latent heat of the transition.

1.3 Hadrons and nuclei

1.3.1 Hadron Tomography

A key goal of the nuclear physics program is a first-principles description of nucleon and hadron structure from QCD. Our knowledge of hadron structure is encapsulated in a variety of quantities. From the earliest observations of Bjorken scaling in deep inelastic scattering, a one-dimensional longitudinal description of the nucleon has been provided through the unpolarized and polarized *parton distribution functions* (PDFs). In contrast, the transverse distribution of charge and currents was probed experimentally in elastic scattering, encapsulated in the electric and magnetic form factors. More recently, new quantities have been discovered correlating both the longitudinal and transverse structures: the *generalized parton distributions* (GPDs) describing hadrons in longitudinal fractional momentum x and transverse impact-parameter space, and the *transverse-momentum-dependent distributions* (TMDs) providing a description in x and transverse momentum space. These new descriptions have opened new vistas on the nucleon, enabling us to understand its spin decomposition and allowing the orbital angular momentum to be discerned.

Recently, new ideas have been proposed that allow for the first time lattice QCD calculations of the full x-dependence of PDFs, GPDs, and TMDs [34, 35, 36, 37, 38]. The work proposed here will explore these new methods in the lightest hadron systems, including the pion and nucleon. In particular, we will calculate the one- and three-dimensional images of the pion and the nucleon and reveal vital information about the quark and gluon degrees of freedom in hadrons. This program is well aligned with the 12 GeV experimental program at Jefferson Lab. The large-x region is precisely where our computational methods are most effective, and thus will have direct impact on the approved experiments measuring the large-x behavior of the pion, kaon and nucleon structure functions. Our project also works in close cooperation with the JAM collaboration with the aim to further enhance the fidelity of the determination of pion and nucleon PDFs from experimental data. More generally, this proposal will establish the theoretical and computational underpinnings for future calculations of the structure of the hadrons. Extending these calculations into matrix element computations [39] of hybrid states which manifest gluonic degrees of freedom explicitly will provide an entirely new view of the role of glue in hadrons, allowing theory to both predict and confront measurements from a future electron-ion collider (EIC) which will have the mission of uncovering the role of quarks and gluons in hadrons.

1.3.2 The physics of nuclei

Nuclei make up the majority of the visible matter in the universe; knowledge of nuclear structure and interactions thus underlies our understanding of the world we live in. Our knowledge of nuclei is derived from decades of innovative experiments, yet our theoretical understanding of the emergence of nuclei from quarks and gluons lags this impressive phenomenology. A central part of the mission of the DOE Office of Nuclear Physics is to explore the frontiers of nuclear physics both through new and more precise

experiments and by deepening our theoretical understanding of nuclear structure and interactions. At the same time, the DOE Office of High-Energy Physics is invested in many current and future experiments seeking to probe the limits of known physics using nuclei as targets. These experiments require detailed knowledge of nuclear structure to maximize their impact. The proposed work addresses these missions by providing first-principles calculations of the properties and interactions of light nuclei from the SM.

Building on our recent work pioneering the *ab initio* studies of the spectrum [40, 41, 42, 43, 44, 45, 46], properties [47, 48, 49, 50], and interactions [51, 52, 53, 54, 55, 56] of light nuclei, we will study the quark and gluon structure of light nuclei, and calculate matrix elements needed to control systematic uncertainties in upcoming experiments searching for dark matter and testing the fundamental symmetries of nature (for example searches for time-reversal symmetry violation through the observation of a nuclear electric dipole moment, and lepton number violation through double β decay searches). The proposed calculations build upon our successful previous study of light nuclei at heavier-than-physical values of the quark masses, and are made possible by recent advances in lattice-QCD algorithms and software enabled by the SciDAC program. In particular, we will calculate the ground-state spectrum and matrix elements of nuclei up to $A = 4$ in multiple lattice volumes, enabling the infinite volume binding energies and structure of these nuclei to be determined from first principles using effective field theory (EFT) methods [57]. These calculations are imperative to demonstrate a concrete understanding of nuclei from the SM.

1.4 Hot-dense lattice QCD

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 [58]. The quark-gluon plasma (QGP) created at LHC and top RHIC energies consists of almost as much antimatter as matter, characterized by the nearly vanishing baryon-number chemical potential. Hot-dense lattice QCD calculations predict that under these conditions the transition from the QGP to a hadron gas occurs through a smooth crossover [59]. 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 [60]. To what extent this conjecture is true can be studied experimentally by measuring non-Gaussian cumulants of conserved charge fluctuations [61]. These experimental measurements have already been undertaken at LHC [62], and much improved results will follow after the 2021 high-luminosity upgrade. On the other hand, the baryon-rich QGP created at lower RHIC energies may experience a sharp first-order phase transition as it cools, with bubbles of QGP and of hadrons coexisting at a well-defined temperature. This region of co-existence ends in a critical point [63], where QGP and ordinary hadron-matter become indistinguishable. The search for the critical point is a major experimental effort, the Beam Energy Scan (BES), is currently underway at RHIC and will continue until 2021. The non-Gaussian cumulants mentioned above are currently being measured in BES [64].

Non-Gaussian cumulants of conserved charge fluctuations also can be computed directly from QCD, using large-scale numerical calculations. The goal of the proposed research is to provide quantitative QCD-based results for these cumulants in the vicinity of the pseudo-critical temperature that characterizes the phase transition from hadrons to the QGP. These results will enable determination of phase-structure and bulk thermodynamic properties hot-dense strong-interaction matter [65] and, thereby, crucially facilitate experimental programs at RHIC and LHC. More specifically, the proposed QCD calculations will provide i) Non-Gaussian fluctuations of various conserved charges, the equilibrium QCD baseline for the corresponding experimentally measured ones; ii) knowledge of interactions and degrees of freedom relevant to determine the chemical freeze-out parameters in the vicinity of the QCD phase boundary;

iii) QCD equation-of-state at non-vanishing baryon density; iv) chiral crossover temperatures at non-vanishing baryon densities and constraints on the location of the QCD critical point.

2 RESEARCH OBJECTIVES AND MILESTONES

2.1 Quark- and lepton-flavor physics

2.1.1 Heavy-quark flavor: B-meson decays

Last year we began a multi-year effort to calculate in lattice QCD the complete set of decay form factors for tree-level and rare $B \to \pi$, $B_s \to K$, $B \to K$, $B \to D$, $B_s \to D_s$, and $B \to D^*$ transitions, as well as for the tree-level $D \to K$, $D \to \pi$, and $D_s \to K$ transitions. Our uncertainties in the form factors for B-meson decays will be at the 1% level and the sub-percent level for D-meson decays. These calculations will enable determinations to unprecedented precision of the CKM elements $|V_{ub}|$, $|V_{cb}|$, $|V_{cd}|$, and $|V_{cs}|$ from experimental measurements of the corresponding differential decay rates. They will also provide crucial support for ongoing searches for new physics in rare flavor-changing neutral-current decays. Our precision goals for this project are guided by our recent high-precision work on the kaon, D, and B-meson decay constants [1, 2] and the methods employed in earlier studies of semileptonic decays [66, 67, 68, 69, 70, 71]. The key new elements are i) a large set of gauge-field configurations at lattice spacings ranging from 0.03–0.15 fm with physical-mass up, down, strange, and charm sea quarks; ii) a highly improved fermion action for all quarks. The multiyear effort is progressing from coarser to finer lattices. We now propose calculations at 0.06 fm and to start those at 0.042 fm, which will be finished on an exascale machine.

2.1.2 Lepton flavor: muon $g_\mu - 2$

Because of the importance of $g_\mu - 2$, calculations using both staggered and chiral quarks are crucial. Using staggered fermions we will calculate the contribution of leading-order hadronic vacuum polarization (HVP) at a lattice spacing of 0.06 fm. We have found that a large statistical sample is essential [11], so we propose to enlarge our ensemble at 0.06 fm and compute the "connected" part of the HVP on them. To meet our precision objectives, we must also reduce uncertainties in the smaller "disconnected" HVP contribution. We propose to do this on Theta with measurements on existing gauge configurations at 0.09 fm.

Using chiral fermions we will extend the HVP calculation to a new $96^3 \times 192$ ensemble, reducing the overall error to 0.3–0.4%. This will be accomplished by computing the two-point correlation function with all-to-all methods and by using distillation methods to determine the long-distance contribution of multi-pion states, both on 40 configurations. For the chiral HLbL calculation we propose to use Theta to measure the disconnected contribution on a mixture of $32^3 \times 64$, $48^3 \times 96$ and $64^3 \times 128$ configurations, reducing the disconnected error by a factor of two.

2.1.3 Second order weak processes

The proposed ΔM_K calculation relies on the new ensemble of $96^3 \times 192$ with the reduced lattice spacing of $a \approx 0.07$ fm now being created with a 2019 Summit Incite award. The ≈ 400 molecular dynamics (MD) time units expected from this earlier award are insufficient for the proposed study. Therefore, we propose to extend this ensemble, adding 200 time units each year. (This extended ensemble will benefit many projects over many years.) Next we must determine the masses of the lowest energy two- and three-quark states as well as physical quantities such as f_π, f_D, etc. to establish the scale and required physical quark masses for this ensemble. This is listed as "standard observables" in the table of milestones. The calculation of ΔM_K will be carried out over three years yielding results from 38 configurations. In the second year we would begin the calculation of the LD part of ϵ_K using physical quark masses on fifty $a \approx 0.070.86$ fm, $64^3 \times 128$ configurations. Our goal for the calculation of ΔM_K is a 20% statistical error at this smaller lattice spacing

and the resulting control of the finite lattice spacing error that will come from a comparison between this new calculation and the Mira INCITE calculation completed last year [72].

Table 2. Computational cost for the staggered-quark parts of this project.

$\approx a$ (fm)	$N_s^3 \times N_t$	calculation	Summit (node-h)	Theta (node-h)
0.06	$96^3 \times 192$	generate 1500 cfg	225,000	0
0.06	$96^3 \times 192$	$g_\mu - 2$ conn 1500 cfg	375,000	0
0.06	$96^3 \times 192$	form factor 1000 cfg	850,000	0
0.042	$144^3 \times 288$	form factor 100 cfg	350,000	4,687,500
Summit total			1,800,000	4,687,500

Table 3. Computational cost of the chiral-quark part of this project.

quantity	number of gauge configurations	cost	
		Summit (k node-h)	Theta (M core-h)
Ensemble generation	600	560,000	0
Standard observables	32	420,000	0
ΔM_K	28	330,000	0
ϵ_K(LD)	46	80,000	0
$g_\mu - 2$ HVP	38	410,000	0
$g_\mu - 2$ HLbL	600	0	4,687,500
Total		1,800,000	4,687,500

2.2 Lattice gauge theories beyond the Standard Model

2.2.1 Strongly-coupled composite higgs dynamics

The research objective is to study two specific composite Higgs models based on the SU(3) gauge group with Dirac fermions in different color and flavor representations. Both models are near-conformal and have a light composite scalar boson. The goal is to constrain the low-energy EFT of the light scalar and determine how its structure is influenced by the near-conformal nature of the underlying gauge theory. In the model with two flavors in the sextet representation of SU(3), the goal is to produce new configurations where the light scalar and pion masses are set by the box size $M \times L \sim 1$, the so-called ϵ regime of the EFT, which is an approach to extracting EFT data from computing the Dirac eigenvalue spectrum.

In the models with fermions in the fundamental representation of SU(3), we take two approaches. First, we keep $N_f = 8$ with equal masses and try to make the composite pion and scalar masses as light as possible by tuning the fermion masses as light as possible holding $M \times L \approx 5.3$, outside the ϵ regime, on $96^3 \times 192$ volumes using staggered fermions. We will use these configurations to perform calculations that constrain the EFT. The model is somewhat unphysical, because it produces too many light composite pions, but understanding the structure of the EFT will shed light on the composite-Higgs mechanism. Second, we split $N_f = 4 + 6$ flavors into light and heavy multiplets and explore how to keep the composite scalar mass light while increasing the splitting between the heavy and light multiplets. This technique has the added benefit of reducing the number of unwanted light pions. We will generate $48^3 \times 96$ configurations with $N_f = 4 + 6$ of Möbius domain wall fermions, because the better chiral symmetry properties are essential in the mass-splitting procedure, justifying the extra expense.

Table 4. Computational cost of the BSM part of this project.

Quantity	Number of gauge configurations	Cost (node-h) Summit	Theta
Composite Higgs - Sextet	1,600	450,000	0
Composite Higgs - Fundamental (light)	450	225,000	0
Composite Higgs - Fundamental (split)	900	0	4,687,500
Composite DM - Critical T	2,100	75,000	0
Composite DM - Latent heat	1,000	150,000	0
Total		900,000	4,687,500

2.2.2 Composite DM early universe phase transitions

We work in the context of a specific composite dark matter model, "stealth dark matter," based on an SU(4) gauge theory with $N_f = 4$ mass degenerate fermion species [26, 27, 28], where a first-order thermal phase transition is expected [29, 30]. Such a transition could produce significant observable effects as the early universe cools, including a new source of primordial gravitational waves [31, 32] and an opportunity for matter asymmetry generation due to out-of-equilibrium interactions [33] during the transition. We begin with several parallel simulations at different values of the physical parameters carried out in order to precisely locate the phase transition temperature and verify the transition is first order. We plan to study three different values of the bare gauge coupling and seven different fermion masses, for a total of 21 ensembles on a volume of $48^3\times24$, refining an initial study on smaller volumes already underway on other resources. Once the transition is located, a follow-up lattice calculation with large statistics at the transition point can be carried out in order to determine its characteristics, such as the latent heat and surface tension which enter into the hydrodynamics of gravity wave generation. For this study we plan a lattice volume of $64^3\times24$, with an increased aspect ratio to study and remove finite-volume corrections.

2.3 Hadrons and nuclei

2.3.1 Hadron Tomography

We will calculate the non-local quark matrix elements required to generate the correlation functions for obtaining pion, nucleon and excited state PDFs, GPDs and TMDs. We plan to employ a novel methodology based on *distillation* [73, 74]. This methodology, has been shown to offer excellent control of excited state contamination as well as offering a significant reduction to statistical fluctuations resulting in accurate determinations of the matrix elements that are required in order to extract the physics observables. Furthermore, the basic building blocks of this computation can be recycled for future use for related computations. In Fig. 1 we present the valence pion PDF from our recent calculation at one lattice spacing. We plan to perform these computations on three lattice spacings in order to achieve a controlled continuum limit extrapolation as well as reach smaller values of momentum fraction x. A fourth lattice spacing will be used for first calculations of excited state structure. The estimated computational cost arise from both the computation of quark propagators as well as contractions of these propagators to produce the relevant correlation functions. In addition we estimate computer time for computing the so-called disconnected diagrams that arise from flavor singlet matrix elements. These diagrams will be shared with the "physics of nuclei" which also focuses on relevant flavor singlet matrix elements.

A major part of the estimated resources will also be devoted to generating the gauge configuration ensembles required for our calculations. In particular, we will produce configurations at three temporal lattice spacings of 0.091 fm, 0.072 fm, 0.06 fm and 0.035 fm, at physical pion masses, that will enable us

to perform continuum extrapolations. These ensembles will be stored and shared with the rest of the collaboration for future projects on hadron structure and spectroscopy.

2.3.2 The physics of nuclei

We plan to calculate the lowest few finite-volume energy levels of atomic number $A = 1, 2, 3, 4$ nuclear and hyper-nuclear systems in two large volumes (complementing currently running calculations at a smaller volume) that will constrain effective field theory extrapolations to infinite volume [75] and allow determinations of the nucleon-nucleon and hyperon-nucleon scattering phase-shifts and three-body interactions[76, 77, 78, 79, 80] at low energy. Demonstrating that the ground states of nuclei can be calculated at close-to-physical quark masses is a vital step forward and is a primary focus.

We also plan to calculate the matrix elements of the scalar, axial and tensor currents and the corresponding transition matrix elements in the same light nuclei in two volumes. Knowledge of these matrix elements is important for precise determinations of the low-energy pp fusion cross-section needed for stellar modelling, interpretations of dark matter direct-detection experiments, and of tests of fundamental symmetries through searches for nuclear electric dipole moments and double-β decay processes.

The estimated resources are summarized in Table 5.

Table 5. Computational cost for the clover-quark part of the project in Summit node-h.

$\approx a$ (fm)	$N_s^3 \times N_t$	Gauge generation	Disconnected	Nucleon Tomography	Light nuclei
0.091	$64^3 \times 128$	–	27,000	207,000	166,000
0.091	$96^3 \times 192$	250,000	135,000	–	365,000
0.072	$72^3 \times 192$	167,000	57,000	332,000	–
0.060	$96^3 \times 256$	333,000	179,000	399,000	–
0.035	$48^3 \times 512$	125,000	13,000	56,000	–
Total	**2,812,000 Summit node-h**				

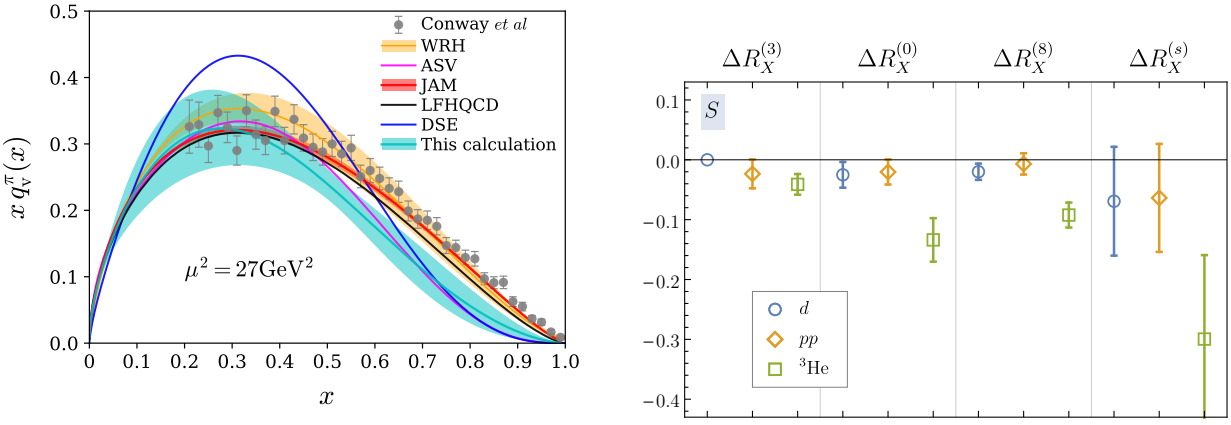

Figure 1. (L) The pion momentum fraction distribution $xq_v^\pi(x)$ from our recent LQCD calculation [81] (cyan band) in comparison with extractions from experiment [82] (gray data points with uncertainties), next-to-leading order (NLO) fits [83, 84, 85] (orange band, magenta curve, and red band), and model calculations [86, 87] (black and blue lines). (R) The shift of scalar matrix elements in various nuclei, deuteron (d), dinucleon (pp) and ^{3}He, separated according to flavor decomposition [54].

2.4 Hot-dense lattice QCD

The hot-dense lattice QCD project team had continuing access to computing time on Titan since 2014 through USQCD-led INCITE awards. These awards led to: i) Calculations of QCD equation-of-state (EoS) at high temperature T, but for vanishing baryon chemical potential μ_B [88]; these calculations provide the baseline for calculations of EoS for $\mu_B > 0$ but have no direct bearing on the presently proposed research. ii) Calculations of up to 4^{th} order cumulants of net-baryon number fluctuations for $\mu_B/T \lesssim 1.5$ [89], reproducing the qualitative trend found in the experimental results. iii) Calculations of the QCD EoS for $\mu_B/T \lesssim 2$ using up to 6^{th} order cumulants of conserved charge fluctuations [90].

The hot-dense lattice QCD project team also won a large INCITE award for 2019 on Summit. Within the first few months we have completely used up this allocation and made significant progress: i) We have determined the QCD crossover temperature and, thereby, the QCD phase boundary in the T-μ_B plane for μ_B/T. These results have been published [59] and are summarized in Fig. 2 (left). We found that for $\mu_B \lesssim 300$ MeV, the chemical freeze-out in heavy-ion collisions takes place in the vicinity of the QCD phase boundary, which coincides with the lines of constant energy and entropy densities. ii) We have extended our calculations of the QCD EoS of Ref. [90] up to $\mu_B/T = 2.5$ by a) improving the precision of $\mathcal{O}(\mu_B^6)$ corrections for the $N_\tau = 8$ lattices; b) adding $\mathcal{O}(\mu_B^8)$ corrections for the $N_\tau = 8$ lattices; c) adding $\mathcal{O}(\mu_B^6)$ corrections for $N_\tau = 12$ lattices. As an example, in Fig. 2 (middle) we show the difference in QCD pressure at $\mu_B/T = 2.5$ and that at $\mu_B = 0$ up to various orders in μ_B and for two different lattice spacings. Thus, we now have controlled QCD calculations for the transition temperature and EoS for almost the entire range of μ_B to be scanned in the BES program as shown by the STAR points in Fig. 2 (left).

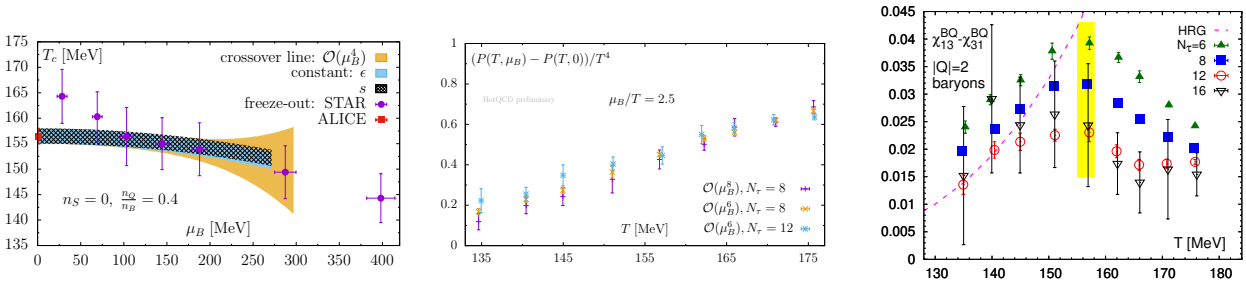

Figure 2. Left: Phase boundary of QCD in the T-μ_B plane compared with the chemical freeze-out parameters extracted from relativistic heavy-ion collision experiments [59]. Middle: QCD pressure at $\mu_B/T = 2.5$ (HotQCD preliminary); see text for details. Right: Contribution of charged Δ baryons to QCD thermodynamics for various lattice spacings, compared with the expectation from hadron resonance gas (Model) model (HotQCD preliminary); the yellow band indicates the QCD chiral crossover temperature.

One key critical physics issue that still needs urgent attention pertains to the question: What is a reliable translation between the colliding energy, $\sqrt{s}$, of the heavy ions (the only adjustable quantity in the lab) and the thermodynamic variables $\{T, \mu_B\}$ at freeze-out? (Freeze-out is the stage of the heavy-ion collision when the QGP re-hadronizes and the hadrons are observed in the detectors.) Without a trustworthy version of the $\sqrt{s} \leftrightarrow \{T, \mu_B\}$ mapping, it will not be possible to transcribe the experimentally obtained knowledge onto the QCD phase diagram. Traditionally this has been done by fitting the experimentally measured hadron yields with a simple hadron resonance gas (HRG) model [91], *i.e.*, the thermodynamics of an uncorrelated gas of hadrons having vacuum masses. In a previous study, we have proposed and demonstrated that this model dependence can be eliminated by comparing directly the experimentally measured non-Gaussian cumulants of net charge fluctuations with their lattice-QCD counterparts [92]. Alternatively, one can use refined

versions of the HRG models, in which the relevant degrees of freedom and effective hadronic interactions in the vicinity of the QCD transition are more accurately incorporated by validating the models against lattice-QCD results [93, 94, 95, 96]. Previously, we demonstrated how such the validation can be implemented with non-Gaussian fluctuations of conserved charges [97, 98]. Accurate computations of the non-Gaussian fluctuations of conserved charges in the vicinity of the QCD transition demand very fine lattice spacings in order to control the to continuum-limit extrapolation. Figure 2 (right) shows an example of the lattice-spacing dependence of the contributions of the charged Δ baryons to the QCD thermodynamics, constructed out of the non-Gaussian cumulants of net charge fluctuations. The figure also exemplifies how far off the simple HRG model is from the QCD calculations in the vicinity of the QCD transition. Thus, the objective of this proposal is to complete this program by finishing the computations on the finest $N_\tau = 16$ lattices by increasing the statistics from the current $O(10,000)$ gauge configurations per temperature to $O(100,000)$, so that the statistical errors become comparable to that on the coarser lattices. The computing resources required to accomplish this objective are given in Table 6.

Table 6. Hot-dense lattice QCD resource requirement in Summit-node-hr for computations of up to 6$^{\text{th}}$ order cumulants of conserved charge fluctuations on $64^3 \times 16$ lattices using 100K gauge configurations per temperature, generated using 2+1 flavors of HISQ with physical quark masses.

Year	# of temperatures	Configuration generation	Cumulant computations	Total
1	3	313,000	250,000	563,000
2	3	313,000	250,000	563,000
3	3	313,000	250,000	563,000
All	9	939,000	750,000	1,689,000

3 COMPUTATIONAL READINESS

3.1 Use of Resources Requested

3.1.1 Quark- and lepton-flavor physics

The resource estimates are based on timings of the same or somewhat smaller calculations on Mira. The cost was converted to Summit (Theta) node-hours, based on benchmarks of the same modules. The ALCF Theta $g-2$ project will require four hours on 128 nodes to complete one gauge configuration. On OLCF Summit, the HISQ gauge-configuration-generation component will require 48 nodes and take approximately 3 hours to produce a new gauge configuration. The HISQ $g-2$ analysis component of the project will require 6 GPUs on 24 Summit nodes to analyze one gauge configuration for up to eight hours. The heavy-quark form-factor analysis component will need 48 GPU nodes for up to fifteen hours. There will be several thousand such jobs with a steady burn rate. Since run times are uniform, the analysis jobs can be bundled efficiently to suit queue requirements. We have already run bundled jobs on lattices of this size on 96 nodes. We will use the burst buffer for intermediate storage of 20 TB worth of propagators.

Requirements for I/O external to a job are very modest. On Summit we will be reading and possibly rereading thousands of 49-GByte gauge-configuration files over the course of the project. On Theta, the gauge-configuration sizes are more modest: 7 GB each. These files will be transferred with Globus Online from archives at Fermilab and staged for analysis on Summit and Theta. Correlator output will be transferred for post-processing elsewhere.

The domain-wall fermion ensemble generation and ΔM_K calculations are capability jobs requiring 20% of

Summit. For example, the former is run as multi-trajectory jobs with each trajectory taking one hour on 1024 nodes, using all six GPUs. The DWF $g_\mu - 2$ and ΔM_K jobs have substantial IO requirements and must load 40 TByte eigenvector files.

3.1.2 Lattice gauge theories beyond the Standard Model

The BSM ensemble generation codes run most efficiently within the expected wall time limit of 12 hours on either 128 or 256 Summit nodes. The largest ensemble ($96^3 \times 192$) can be run on as many as 1024 Summit nodes with only a slight loss in performance due to the good strong scaling of the staggered inverter. BSM lattice generation on Theta uses 384 KNL nodes.

3.1.3 Hadrons and nuclei

The resource estimates in Table 5 are for gauge generation of 10k trajectories, providing 500 configurations. The estimates for $72^3 \times 192$ at 0.072 fm are based on trajectory timings using 576 Summit nodes. The timings for $96^3 \times 192$ used node counts up to 1024 nodes, while 512 are optimal. The timings for $a = 0.035$ fm are based on production runs at 576 nodes. Disconnected diagrams occurring in three-point functions are common to all the hadron and nuclear calculations, with timings based on 2048 inversions per configuration using 64 nodes. The resource estimates for nucleon matrix elements are based on 16k inversions and contraction counts using 64 nodes, using multiple source and sink separations for inversions. The nuclear matrix elements will use half that number but with larger contraction requirements consuming the propagators. Analysis jobs will be aggregated to partition sizes up to about 1000 nodes with only a small degradation in performance.

I/O requirements are modest. Gauge generation produces files about 30 GByte per configuration. For the analysis, propagators are held in memory and consumed within the contractions. The output are small files, less than 1 GB per configuration.

3.1.4 Hot-dense lattice QCD

We will use our highly optimized, well-tested HotQCD software-suite based on C++, CUDA, OpenMP and MPI. This software runs successfully on Summit under our 2019 INCITE allocations. The resource requirements quoted in Table 6 are based on the benchmarks of HotQCD code on Summit for $64^3 \times 16$ lattices with physical quark masses: (i) Generation of one gauge-configuration, separated by 10 RHMC trajectories of unit length, takes ~ 1.05 Summit-node-hr. (ii) Measurements of all the operators needed for up to 6^{th} order cumulants on each gauge-configuration takes ~ 0.83 Summit-node-hr. Gauge configuration generation jobs use about 300 nodes at a time, and typical runtime of a single job will be around 12 hours. Typical cumulants measurement jobs will run on about 1000 Summit nodes for about 10 hours, simultaneously using both GPUs and CPUs. Each year we will generate about 250 TB of gauge configurations, which must be archived for long term-use. Additionally, during the cumulants measurements, 500 TB need to remain available throughout the year for the gauge configurations and the eigenvectors needed for the deflated conjugate gradient inverter.

3.2 Computational Approach

3.2.1 Quark- and lepton-flavor physics

This project uses two optimized lattice QCD code bases, the CPS code for domain-wall fermions and the MILC code for Highly-Improved Staggered Fermions. Both use MPI for inter-GPU communication (inter-node on Theta), OpenMP for threaded operation on the Power9 CPUs (and on KNL processors in the case of Theta). For optimum performance on Summit, they use QUDA, a CUDA-based library created for

lattice QCD by NVIDIA collaborators. On Theta our codes use Grid modules, which are specifically optimized for that architecture. Job submission and management is controlled by scripts. Temporary files are kept on the SSD for rapid I/O. Postprocessing is handled elsewhere.

3.2.2 Lattice gauge theories beyond the Standard Model

The BSM ensemble generation codes use the same USQCD software libraries and QUDA GPU libraries on Summit (or Grid libraries on Theta) to achieve comparable performance as QCD codes used elsewhere. At the highest level, specific application packages implement the composite Higgs/Dark Matter models: LatHC for the sextet composite Higgs, NIM/QEX for fundamental composite Higgs and composite Dark Matter, and Grid for mass-split studies on Theta. Job submission is handled by scripts and postprocessing of output files is handled on other resources.

3.2.3 Hadrons and nuclei

The Chroma [99] code base along with the QUDA library [100, 101, 102] is used for the ensemble generation with the clover action. Chroma and QUDA have been well optimized for GPUs for the propagator and contraction operations parts of the project. We will use the accelerated version of QDP-JIT [103] where the lattice-wide C++ expression template implementation generates code Just-in-Time (JIT) for the GPUs. This version can take advantage of the GPU-direct feature of GPU aware MPI implementations such as Spectrum MPI. Chroma's gauge generation integrator code has been significantly accelerated with the inclusion of the Adaptive Aggregation Multi-grid (AAMG) solver from QUDA [102, 104] among other optimizations. The propagator codes use the AMG code as well as JIT and other low level optimizations for contractions of solution vectors on GPUs.

3.2.4 Hot-dense lattice QCD

The gauge field configurations with HISQ sea quarks are generated with an RHMC algorithm using a three-scale integrator and 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 to run on 2–6 GPUs, depending on the lattice sizes, communicating through NVLINK. Our analysis programs for the higher-order Taylor-series expansion coefficients have been tuned for both GPUs and CPUs. The analysis software exploits a CG inverter optimized for inverting sparse (fermion) matrices using several right hand sides [105, 106]. On one Summit Volta GPU, our CUDA-based GPU code can carry out inversions for up to 12 right hand sides, and our OpenMP threaded C++ IBM POWER 9 CPU code can accommodate up to 64 right hand sides. The mult-right hand side GPU inverter also has been optimized to run on 2–6 GPUs, depending on the lattice sizes, communicating through NVLINK. Furthermore, the inverter uses deflation as a preconditioner which reduces the number of conjugate gradient steps in the matrix inversion by a factor 5–8, depending on the number of eigenvectors for small eigenmodes one can store. For this purpose we calculate $O(200)$of the lowest eigenvalues of the fermion matrix and deflate the Krylov space from the eigenspace of the corresponding eigenvectors by a Galerkin projection. To estimate the traces of the inverse of fermion matrices entering the relevant observables, we use stochastic estimators, which require $O(1000)$ matrix inversions with different sources on identical gauge field configurations, which introduces a high degree of parallelism.

3.3 Parallel Performance

3.3.1 Quark- and lepton-flavor physics

The multimass solver accounts for 60–80% of the cost of our calculations, so it is a key performance indicator. Weak scaling results in Table 7 are from Summit in early 2018. Our fermion formulation (HISQ)

has low arithmetic intensity, so performance is network limited. Since then, there have been considerable improvements in system software, our codes, and run parameters. In a recent Summit calculation, the code sustained 1.3 TF/s per node in double precision on 96 nodes with a bundling factor of four.

Table 7. The second row shows weak-scaling performance for the MILC/QUDA/HISQ double-precision multimass solver with a local volume of 32^4 per GPU. The third row shows CPS/QUDA/DWF weak scaling with a local volume of $12^3 \times 16$ per GPU, done in half-precision with correction steps to give double-precision accuracy.

Code	Nodes→	16	32	64	128	256	512	1024
MILC/QUDA/HISQ	(GF/s-node)	432	426	306	300	282	258	228
CPS/QUDA/DWF	(GF/s-node)	—	—	1990	—	—	1590	1550

Over the past year the performance listed in Table 7 for the DWF code has improved by a factor of 2.1 from a combination of application and system software upgrades. The most interesting application software improvement arises from the MSPCG algorithm which uses the tensor cores in a local, preconditioning step [107].

3.3.2 Lattice gauge theories beyond the Standard Model

For the BSM lattice generation codes on Summit, the QUDA staggered solver accounts for about 80% of the total cost of our calculations. In recent tests on Lassen at LLNL (similar to Summit except 4 GPU/node instead of 6) we are sustaining between 1.3 and 2.2 TFlops/node in strong scaling tests on $96^3 \times 192$ lattices, comparable to what has been seen in the MILC/QUDA/HISQ tests on Summit. We anticipate the additional two GPUs/node on Summit will boost the performance per node by nearly 50% given the strong on-node scaling performance of QUDA due the fast on-node interconnections between GPUs via NVLink.

3.3.3 Hadrons and nuclei

In Table 8 we list the strong scaling results for Chroma gauge generation on a $72^3 \times 192$ lattice with a slightly heavier mass, $m_\pi \approx 270$ MeV, than would be used in production. The table indicates that strong scaling effects become significant at and beyond 192 nodes (1152 GPUs). Strong scaling is limited by the coarsest grid within the adaptive multigrid solver used in the gauge integrator. Previous timings, using an inferior inverter, which combined a GCR solver with a domain decomposed preconditioner, would allow for strong scaling up to 576 nodes. Overall, the use of multigrid reduces resource costs. Given the same features of AMG, the propagator calculations are run on as few nodes as possible to improve efficiency. Analysis on $64^3 \times 128$ with 16 nodes obtains 36 Tflops/s, and propagators on $72^3 \times 192$ with 32 nodes realizes 60 Tflops/s.

Table 8. Strong scaling timings for Chroma gauge generation with lattice size $72^3 \times 192$ sites.

Nodes	48	64	192	576
sec/trajectory	825	675	535	400

3.3.4 Hot-dense lattice QCD

Figure 3 shows multi-GPU performance of the HotQCD codes on Summit for the CG inverters used for the gauge generation (left) and cumulants measurements (right). For the proposed $64^3 \times 16$ lattice generation performance of the HotQCD code for 4 GPUs is about 70% of the ideal scaling, and that for measurements is about 90% for 2 GPUs. Since we need to generate 100k gauge configurations, *i.e.*, about 1M RHMC

trajectories per temperature, this can only be achieved through many independent streams of gauge generations for each temperature. For each temperature we will use 200 independent gauge generation streams, with each stream generating about 5000 RHMC trajectories. On a single Summit node we will generate 2 streams, one running using 4 GPUs and the other 2 GPUs. All 600 generation streams combining 3 temperatures will be tied together through intra-node OpenMP and inter-node MPI into one single job of 300 Summit nodes. Each cumulants measurement job for a single configuration will run on 2 GPUs, measuring 3 configurations on a node. By doing a simultaneous measurements on 3000 configurations, once again tied together into one single job as mentioned before, we will run 1000 Summit node jobs. This strategy has been implemented very successfully within our on-going INCITE runs.

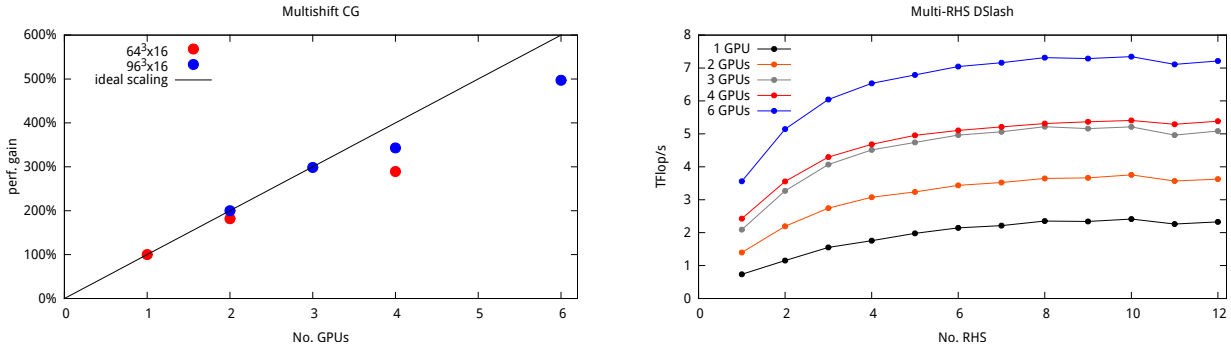

Figure 3. Left: Multi-GPU performance of HotQCD multi-shift HISQ CG inverter on Summit for two different lattice sizes. Right: Multi-GPU performance of HotQCD multi right hand side HISQ CG inverter on Summit for $96^3 \times 16$ lattice.

3.4 Developmental Work

Some of our co-PIs and team members are participating in a joint Nuclear Physics/ASCR SciDAC-4 project. This project has been supporting the development of the gauge generation and analysis suites for near term Leadership computing facilities, including Summit. SciDAC-4 provides the support for adapting workflow software for managing large scale analysis campaigns, some of the nuclear physics matrix element capability as well the Hot QCD portion, including the calculation of transport coefficients.

The PI, several co-PIs, and many team members of this proposal are active participants in the LatticeQCD Application Development component of the DOE Exascale Computing Project (ECP). Although the long-term goal of the ECP is to ready codes for the Aurora and Frontier exascale computers, the OLCF Summit computer is being used as a gauge of progress. Some of us also participate in the NERSC Perlmutter NESAP. Thus all major code bases in this proposal, namely, CPS, Chroma, MILC, and HotQCD—and the packages that these codes depend on, namely, QUDA and Grid—are under intensive and continual improvement, both in portability and performance. The ECP project also supports research and development of new, efficient algorithms for our projects, including multigrid solvers, graphical complexity reduction for multiquark calculations, and Fourier acceleration of gauge-configuration generation.

4 REFERENCES

[1] **Fermilab Lattice, MILC** Collaboration, A. Bazavov *et. al.*, *Charmed and light pseudoscalar meson decay constants from four-flavor lattice QCD with physical light quarks*, *Phys. Rev.* **D90** (2014) 074509 [1407.3772].

[2] **Fermilab Lattice, MILC** Collaboration, A. Bazavov *et. al.*, *B- and D-meson leptonic decay constants from four-flavor lattice QCD*, *Phys. Rev.* **D98** (2018) 074512 [1712.09262].

[3] **Muon** $g-2$ Collaboration, G. W. Bennett *et. al.*, *Final report of the Muon E821 anomalous magnetic moment measurement at BNL*, *Phys. Rev.* **D73** (2006) 072003 [hep-ex/0602035].

[4] **HPQCD** Collaboration, B. Chakraborty, C. T. H. Davies, G. C. Donald, R. J. Dowdall, J. Koponen, G. P. Lepage and T. Teubner, *Strange and charm quark contributions to the anomalous magnetic moment of the muon*, *Phys. Rev.* **D89** (2014) 114501 [1403.1778].

[5] T. Blum, N. Christ, M. Hayakawa, T. Izubuchi, L. Jin, C. Jung and C. Lehner, *Connected and leading disconnected hadronic light-by-light contribution to the muon anomalous magnetic moment with a physical pion mass*, *Phys. Rev. Lett.* **118** (2017) 022005 [1610.04603].

[6] **HPQCD** Collaboration, B. Chakraborty, C. T. H. Davies, P. G. de Oliviera, J. Koponen, G. P. Lepage and R. S. Van de Water, *The hadronic vacuum polarization contribution to a_μ from full lattice QCD*, *Phys. Rev.* **D96** (2017) 034516 [1601.03071].

[7] **Budapest-Marseille-Wuppertal** Collaboration, S. Borsanyi *et. al.*, *Hadronic vacuum polarization contribution to the anomalous magnetic moments of leptons from first principles*, *Phys. Rev. Lett.* **121** (2018) 022002 [1711.04980].

[8] T. Blum, N. Christ, M. Hayakawa, T. Izubuchi, L. Jin, C. Jung and C. Lehner, *Using infinite volume, continuum QED and lattice QCD for the hadronic light-by-light contribution to the muon anomalous magnetic moment*, *Phys. Rev.* **D96** (2017) 034515 [1705.01067].

[9] **Fermilab Lattice, HPQCD, MILC** Collaboration, B. Chakraborty *et. al.*, *Strong-isospin-breaking correction to the muon anomalous magnetic moment from lattice qcd at the physical point*, *Phys. Rev. Lett.* **120** (2018) 152001 [1710.11212].

[10] **RBC, UKQCD** Collaboration, T. Blum, P. A. Boyle, V. Gülpers, T. Izubuchi, L. Jin, C. Jung, A. Jüttner, C. Lehner, A. Portelli and J. T. Tsang, *Calculation of the hadronic vacuum polarization contribution to the muon anomalous magnetic moment*, *Phys. Rev. Lett.* **121** (2018) 022003 [1801.07224].

[11] **Fermilab Lattice, HPQCD, MILC** Collaboration, C. T. H. Davies *et. al.*, *Hadronic-vacuum-polarization contribution to the muon's anomalous magnetic moment from four-flavor lattice QCD*, 1902.04223.

[12] A. Gérardin, M. Cè, G. von Hippel, B. Hörz, H. B. Meyer, D. Mohler, K. Ottnad, J. Wilhelm and H. Wittig, *The leading hadronic contribution to $(g-2)_\mu$ from lattice QCD with $N_f = 2+1$ flavours of O(a) improved Wilson quarks*, 1904.03120.

[13] C. Aubin, T. Blum, C. Tu, M. Golterman, C. Jung and S. Peris, *Light quark vacuum polarization at the physical point and contribution to the muon $g-2$*, 1905.09307.

[14] J. Brod and M. Gorbahn, *Next-to-next-to-leading-order charm-quark contribution to the CP violation parameter ϵ_K and ΔM_K*, *Phys. Rev. Lett.* **108** (2012) 121801 [1108.2036].

[15] A. J. Buras, D. Guadagnoli and G. Isidori, *On ϵ_K beyond lowest order in the operator product expansion*, *Phys. Lett.* **B688** (2010) 309–313 [1002.3612].

[16] **RBC, UKQCD** Collaboration, N. H. Christ, T. Izubuchi, C. T. Sachrajda, A. Soni and J. Yu, *Long distance contribution to the K_L-K_S mass difference*, *Phys. Rev.* **D88** (2013) 014508 [1212.5931].

[17] Z. Bai, N. H. Christ, T. Izubuchi, C. T. Sachrajda, A. Soni and J. Yu, *K_L-K_S mass difference from lattice QCD*, *Phys. Rev. Lett.* **113** (2014) 112003 [1406.0916].

[18] N. H. Christ and Z. Bai, *Computing the long-distance contributions to ε_K*, *PoS* **LATTICE2015** (2016) 342.

[19] O. Witzel, *Review on Composite Higgs Models*, *PoS* **LATTICE2018** (2019) 006 [1901.08216].

[20] W. D. Goldberger, B. Grinstein and W. Skiba, *Distinguishing the Higgs boson from the dilaton at the Large Hadron Collider*, *Phys. Rev. Lett.* **100** (2008) 111802 [0708.1463].

[21] T. Appelquist, J. Ingoldby and M. Piai, *Dilaton EFT Framework For Lattice Data*, *JHEP* **07** (2017) 035 [1702.04410].

[22] **LSD** Collaboration, T. Appelquist *et. al.*, *Linear Sigma EFT for Nearly Conformal Gauge Theories*, *Phys. Rev.* **D98** (2018), no. 11 114510 [1809.02624].

[23] V. Gorbenko, S. Rychkov and B. Zan, *Walking, Weak first-order transitions, and Complex CFTs*, *JHEP* **10** (2018) 108 [1807.11512].

[24] V. Gorbenko, S. Rychkov and B. Zan, *Walking, Weak first-order transitions, and Complex CFTs II. Two-dimensional Potts model at $Q > 4$*, *SciPost Phys.* **5** (2018), no. 5 050 [1808.04380].

[25] **Lattice Strong Dynamics (LSD)** Collaboration, T. Appelquist *et. al.*, *Lattice Calculation of Composite Dark Matter Form Factors*, *Phys. Rev.* **D88** (2013), no. 1 014502 [1301.1693].

[26] **Lattice Strong Dynamics (LSD)** Collaboration, T. Appelquist *et. al.*, *Composite bosonic baryon dark matter on the lattice: SU(4) baryon spectrum and the effective Higgs interaction*, *Phys. Rev.* **D89** (2014), no. 9 094508 [1402.6656].

[27] T. Appelquist *et. al.*, *Stealth Dark Matter: Dark scalar baryons through the Higgs portal*, *Phys. Rev.* **D92** (2015), no. 7 075030 [1503.04203].

[28] T. Appelquist *et. al.*, *Detecting Stealth Dark Matter Directly through Electromagnetic Polarizability*, *Phys. Rev. Lett.* **115** (2015), no. 17 171803 [1503.04205].

[29] R. D. Pisarski and F. Wilczek, *Remarks on the Chiral Phase Transition in Chromodynamics*, *Phys. Rev.* **D29** (1984) 338–341.

[30] A. Butti, A. Pelissetto and E. Vicari, *On the nature of the finite temperature transition in QCD*, *JHEP* **08** (2003) 029 [hep-ph/0307036].

[31] P. Schwaller, *Gravitational Waves from a Dark Phase Transition*, *Phys. Rev. Lett.* **115** (2015), no. 18 181101 [1504.07263].

[32] C. Caprini *et. al.*, *Science with the space-based interferometer eLISA. II: Gravitational waves from cosmological phase transitions*, *JCAP* **1604** (2016), no. 04 001 [1512.06239].

[33] A. D. Sakharov, *Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe*, *Pisma Zh. Eksp. Teor. Fiz.* **5** (1967) 32–35. [Usp. Fiz. Nauk161,no.5,61(1991)].

[34] X. Ji, *Parton physics on a euclidean lattice*, *Phys. Rev. Lett.* **110** (2013) 262002 [1305.1539].

[35] X. Ji, *Parton Physics from Large-Momentum Effective Field Theory*, *Sci. China Phys. Mech. Astron.* **57** (2014) 1407–1412 [1404.6680].

[36] Y.-Q. Ma and J.-W. Qiu, *Extracting parton distribution functions from lattice QCD calculations*, *Phys. Rev.* **D98** (2018), no. 7 074021 [1404.6860].

[37] A. V. Radyushkin, *Quasi-parton distribution functions, momentum distributions, and pseudo-parton distribution functions*, *Phys. Rev.* **D96** (2017), no. 3 034025 [1705.01488].

[38] K. Orginos, A. Radyushkin, J. Karpie and S. Zafeiropoulos, *Lattice QCD exploration of parton pseudo-distribution functions*, *Phys. Rev.* **D96** (2017), no. 9 094503 [1706.05373].

[39] A. Baroni, R. A. Briceño, M. T. Hansen and F. G. Ortega-Gama, *Form factors of two-hadron states from a covariant finite-volume formalism*, 1812.10504.

[40] S. R. Beane, P. F. Bedaque, K. Orginos and M. J. Savage, *Nucleon-nucleon scattering from fully-dynamical lattice QCD*, *Phys. Rev. Lett.* **97** (2006) 012001 [hep-lat/0602010].

[41] **NPLQCD** Collaboration, S. R. Beane, E. Chang, W. Detmold, H. W. Lin, T. C. Luu, K. Orginos, A. Parreno, M. J. Savage, A. Torok and A. Walker-Loud, *The deuteron and exotic two-body bound states from lattice QCD*, *Phys. Rev.* **D85** (2012) 054511 [1109.2889].

[42] S. R. Beane, E. Chang, S. D. Cohen, W. Detmold, H. W. Lin, T. C. Luu, K. Orginos, A. Parreno, M. J. Savage and A. Walker-Loud, *Hyperon-nucleon interactions and the composition of dense nuclear matter from quantum chromodynamics*, *Phys. Rev. Lett.* **109** (2012) 172001 [1204.3606].

[43] **NPLQCD** Collaboration, S. R. Beane, E. Chang, S. D. Cohen, W. Detmold, H. W. Lin, T. C. Luu, K. Orginos, A. Parreno, M. J. Savage and A. Walker-Loud, *Light nuclei and hypernuclei from quantum chromodynamics in the limit of SU(3) flavor symmetry*, *Phys. Rev.* **D87** (2013), no. 3 034506 [1206.5219].

[44] **NPLQCD** Collaboration, S. R. Beane *et. al.*, *Nucleon-nucleon scattering parameters in the limit of SU(3) flavor symmetry*, *Phys. Rev.* **C88** (2013), no. 2 024003 [1301.5790].

[45] K. Orginos, A. Parreno, M. J. Savage, S. R. Beane, E. Chang and W. Detmold, *Two nucleon systems at $m_\pi \sim 450$ MeV from lattice QCD*, *Phys. Rev.* **D92** (2015), no. 11 114512 [1508.07583].

[46] M. L. Wagman, F. Winter, E. Chang, Z. Davoudi, W. Detmold, K. Orginos, M. J. Savage and P. E. Shanahan, *Baryon-baryon interactions and spin-flavor symmetry from lattice quantum chromodynamics*, *Phys. Rev.* **D96** (2017), no. 11 114510 [1706.06550].

[47] S. R. Beane, E. Chang, S. Cohen, W. Detmold, H. W. Lin, K. Orginos, A. Parreno, M. J. Savage and B. C. Tiburzi, *Magnetic moments of light nuclei from lattice quantum chromodynamics*, *Phys. Rev. Lett.* **113** (2014), no. 25 252001 [1409.3556].

[48] **NPLQCD** Collaboration, E. Chang, W. Detmold, K. Orginos, A. Parreno, M. J. Savage, B. C. Tiburzi and S. R. Beane, *Magnetic structure of light nuclei from lattice QCD*, *Phys. Rev.* **D92** (2015), no. 11 114502 [1506.05518].

[49] F. Winter, W. Detmold, A. S. Gambhir, K. Orginos, M. J. Savage, P. E. Shanahan and M. L. Wagman, *First lattice QCD study of the gluonic structure of light nuclei*, *Phys. Rev.* **D96** (2017), no. 9 094512 [1709.00395].

[50] W. Detmold, K. Orginos, A. Parreno, M. J. Savage, B. C. Tiburzi, S. R. Beane and E. Chang, *Unitary Limit of Two-Nucleon Interactions in Strong Magnetic Fields*, *Phys. Rev. Lett.* **116** (2016), no. 11 112301 [1508.05884].

[51] **NPLQCD** Collaboration, S. R. Beane, E. Chang, W. Detmold, K. Orginos, A. Parreno, M. J. Savage and B. C. Tiburzi, Ab initio *calculation of the $np \to d\gamma$ radiative capture process*, *Phys. Rev. Lett.* **115** (2015), no. 13 132001 [1505.02422].

[52] M. J. Savage, P. E. Shanahan, B. C. Tiburzi, M. L. Wagman, F. Winter, S. R. Beane, E. Chang, Z. Davoudi, W. Detmold and K. Orginos, *Proton-proton fusion and tritium β-decay from lattice quantum chromodynamics*, 1610.04545.

[53] P. E. Shanahan, B. C. Tiburzi, M. L. Wagman, F. Winter, E. Chang, Z. Davoudi, W. Detmold and K. J. Orginos, *Isotensor Axial Polarizability and Lattice QCD Input for Nuclear Double-β Decay Phenomenology*, *Phys. Rev. Lett.* **119** (2017), no. 6 062003 [1701.03456].

[54] **NPLQCD** Collaboration, E. Chang, Z. Davoudi, W. Detmold, A. S. Gambhir, K. Orginos, M. J. Savage, P. E. Shanahan, M. L. Wagman and F. Winter, *Scalar, axial, and tensor interactions of light nuclei from lattice QCD*, *Phys. Rev. Lett.* **120** (2018), no. 15 152002 [1712.03221].

[55] B. C. Tiburzi, M. L. Wagman, F. Winter, E. Chang, Z. Davoudi, W. Detmold, K. Orginos, M. J. Savage and P. E. Shanahan, *Double-β decay matrix elements from lattice quantum chromodynamics*, *Phys. Rev.* **D96** (2017), no. 5 054505 [1702.02929].

[56] W. Detmold and D. J. Murphy, *Nuclear matrix elements for neutrinoless double beta decay from lattice QCD*, *PoS* **LATTICE2018** (2019) 262 [1811.05554].

[57] N. Barnea, L. Contessi, D. Gazit, F. Pederiva and U. van Kolck, *Effective field theory for lattice nuclei*, *Phys. Rev. Lett.* **114** (2015), no. 5 052501 [1311.4966].

[58] D. Geesaman *et. al.*, *The 2015 long range plan for nuclear science*, 2015. https://science.osti.gov/-/media/np/nsac/pdf/2015LRP/2015_LRPNS_091815.pdf.

[59] A. Bazavov *et. al.*, *Chiral crossover in QCD at zero and non-zero chemical potentials*, *Phys. Lett.* **B795** (2019) 15–21 [1812.08235].

[60] H. T. Ding *et. al.*, *The chiral phase transition temperature in (2+1)-flavor QCD*, 1903.04801.

[61] 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.* **C71** (2011) 1694 [1103.3511].

[62] **ALICE** Collaboration, A. Rustamov, *Net-baryon fluctuations measured with ALICE at the CERN LHC*, *Nucl. Phys.* **A967** (2017) 453–456 [1704.05329].

[63] A. M. Halasz, A. D. Jackson, R. E. Shrock, M. A. Stephanov and J. J. M. Verbaarschot, *On the phase diagram of QCD*, *Phys. Rev.* **D58** (1998) 096007 [hep-ph/9804290].

[64] A. Bzdak, S. Esumi, V. Koch, J. Liao, M. Stephanov and N. Xu, *Mapping the phases of quantum chromodynamics with beam energy scan*, 1906.00936.

[65] H.-T. Ding, F. Karsch and S. Mukherjee, *Thermodynamics of strong-interaction matter from lattice QCD*, *Int. J. Mod. Phys.* **E24** (2015), no. 10 1530007 [1504.05274].

[66] **Fermilab Lattice, MILC** Collaboration, J. A. Bailey *et. al.*, *$|V_{ub}|$ from $B \to \pi\ell\nu$ decays and (2+1)-flavor lattice QCD*, *Phys. Rev.* **D92** (2015) 014024 [1503.07839].

[67] **Fermilab Lattice, MILC** Collaboration, J. A. Bailey *et. al.*, *$B \to \pi\ell\ell$ form factors for new-physics searches from lattice QCD*, *Phys. Rev. Lett.* **115** (2015) 152002 [1507.01618].

[68] **Fermilab Lattice, MILC** Collaboration, J. A. Bailey *et. al.*, *$B \to Kl^+l^-$ decay form factors from three-flavor lattice QCD*, *Phys. Rev.* **D93** (2016) 025026 [1509.06235].

[69] **Fermilab Lattice, MILC** Collaboration, E. Gámiz *et. al.*, *Kaon semileptonic decays with $N_f = 2 + 1 + 1$ HISQ fermions and physical light-quark masses*, *PoS* **LATTICE2016** (2016) 286 [1611.04118].

[70] **Fermilab Lattice, MILC** Collaboration, T. Primer *et. al.*, *D meson semileptonic form factors with HISQ valence and sea quarks*, *PoS* **LATTICE2016** (2017) 305.

[71] **Fermilab Lattice, MILC** Collaboration, A. Bazavov *et. al.*, *Determination of $|V_{us}|$ from a lattice-QCD calculation of the $K \to \pi\ell\nu$ semileptonic form factor with physical quark masses*, *Phys. Rev. Lett.* **112** (2014) 112001 [1312.1228].

[72] B. Wang, *Results for the mass difference between the long- and short-lived K mesons for physical quark masses*, *PoS* **LATTICE2018** (2019) 286 [1812.05302].

[73] C. Egerer, D. Richards and F. Winter, *Controlling excited-state contributions with distillation in lattice QCD calculations of nucleon isovector charges g_S^{u-d}, g_A^{u-d}, g_T^{u-d}*, *Phys. Rev.* **D99** (2019), no. 3 034506 [1810.09991].

[74] C. J. Shultz, J. J. Dudek and R. G. Edwards, *Excited meson radiative transitions from lattice QCD using variationally optimized operators*, *Phys. Rev.* **D91** (2015), no. 11 114501 [1501.07457].

[75] M. Lüscher, *Volume dependence of the energy spectrum in massive quantum field theories 2: Scattering states*, *Commun. Math. Phys.* **105** (1986) 153–188.

[76] S. R. Beane, W. Detmold and M. J. Savage, *n-boson energies at finite volume and three-boson interactions*, *Phys. Rev.* **D76** (2007) 074507 [0707.1670].

[77] R. A. Briceño and Z. Davoudi, *Three-particle scattering amplitudes from a finite volume formalism*, *Phys. Rev.* **D87** (2013), no. 9 094507 [1212.3398].

[78] M. T. Hansen and S. R. Sharpe, *Relativistic, model-independent, three-particle quantization condition*, *Phys. Rev.* **D90** (2014), no. 11 116003 [1408.5933].

[79] U.-G. Meißner, G. Rios and A. Rusetsky, *Spectrum of three-body bound states in a finite volume*, *Phys. Rev. Lett.* **114** (2015), no. 9 091602 [1412.4969]. [Erratum: Phys. Rev. Lett.117,no.6,069902(2016)].

[80] R. A. Briceño, J. J. Dudek and R. D. Young, *Scattering processes and resonances from lattice QCD*, *Rev. Mod. Phys.* **90** (2018), no. 2 025001 [1706.06223].

[81] R. S. Sufian, J. Karpie, C. Egerer, K. Orginos, J.-W. Qiu and D. G. Richards, *Pion Valence Quark Distribution from Matrix Element Calculated in Lattice QCD*, *Phys. Rev.* **D99** (2019), no. 7 074507 [1901.03921].

[82] J. S. Conway *et. al.*, *Experimental study of muon pairs produced by 252-GeV pions on tungsten*, *Phys. Rev.* **D39** (1989) 92–122.

[83] K. Wijesooriya, P. E. Reimer and R. J. Holt, *The pion parton distribution function in the valence region*, *Phys. Rev.* **C72** (2005) 065203 [nucl-ex/0509012].

[84] M. Aicher, A. Schafer and W. Vogelsang, *Soft-gluon resummation and the valence parton distribution function of the pion*, *Phys. Rev. Lett.* **105** (2010) 252003 [1009.2481].

[85] P. C. Barry, N. Sato, W. Melnitchouk and C.-R. Ji, *First Monte Carlo global QCD analysis of pion parton distributions*, *Phys. Rev. Lett.* **121** (2018), no. 15 152001 [1804.01965].

[86] **HLFHS** Collaboration, G. F. de Teramond, T. Liu, R. S. Sufian, H. G. Dosch, S. J. Brodsky and A. Deur, *Universality of generalized parton distributions in light-front holographic QCD*, *Phys. Rev. Lett.* **120** (2018), no. 18 182001 [1801.09154].

[87] C. Chen, L. Chang, C. D. Roberts, S. Wan and H.-S. Zong, *Valence-quark distribution functions in the kaon and pion*, *Phys. Rev.* **D93** (2016), no. 7 074021 [1602.01502].

[88] **HotQCD** Collaboration, A. Bazavov *et. al.*, *Equation of state in (2+1)-flavor QCD*, *Phys. Rev.* **D90** (2014) 094503 [1407.6387].

[89] **HotQCD** Collaboration, A. Bazavov *et. al.*, *Skewness and kurtosis of net baryon-number distributions at small values of the baryon chemical potential*, *Phys. Rev.* **D96** (2017), no. 7 074510 [1708.04897].

[90] A. Bazavov *et. al.*, *The QCD equation of state to* $\mathrm{O}(\mu_B^6)$ *from lattice QCD*, *Phys. Rev.* **D95** (2017), no. 5 054504 [1701.04325].

[91] A. Andronic, P. Braun-Munzinger, K. Redlich and J. Stachel, *Decoding the phase structure of QCD via particle production at high energy*, *Nature* **561** (2018), no. 7723 321–330 [1710.09425].

[92] A. Bazavov *et. al.*, *Freeze-out conditions in heavy ion collisions from QCD thermodynamics*, *Phys. Rev. Lett.* **109** (2012) 192302 [1208.1220].

[93] P. Huovinen and P. Petreczky, *Hadron resonance gas with repulsive interactions and fluctuations of conserved charges*, *Phys. Lett.* **B777** (2018) 125–130 [1708.00879].

[94] V. Vovchenko, M. I. Gorenstein and H. Stoecker, *van der Waals interactions in hadron resonance gas: From nuclear matter to lattice QCD*, *Phys. Rev. Lett.* **118** (2017), no. 18 182301 [1609.03975].

[95] V. Vovchenko, M. I. Gorenstein and H. Stoecker, *Finite resonance widths influence the thermal-model description of hadron yields*, *Phys. Rev.* **C98** (2018), no. 3 034906 [1807.02079].

[96] A. Andronic, P. Braun-Munzinger, B. Friman, P. M. Lo, K. Redlich and J. Stachel, *The thermal proton yield anomaly in Pb-Pb collisions at the LHC and its resolution*, *Phys. Lett.* **B792** (2019) 304–309 [1808.03102].

[97] A. Bazavov *et. al.*, *Additional strange hadrons from qcd thermodynamics and strangeness freezeout in heavy ion collisions*, *Phys. Rev. Lett.* **113** (2014), no. 7 072001 [1404.6511].

[98] A. Bazavov *et. al.*, *The melting and abundance of open charm hadrons*, *Phys. Lett.* **B737** (2014) 210–215 [1404.4043].

[99] **SciDAC, LHPC, UKQCD** Collaboration, R. G. Edwards and B. Joo, *The Chroma software system for lattice QCD*, *Nucl.Phys.Proc.Suppl.* **140** (2005) 832 [hep-lat/0409003].

[100] R. Babich, J. Brannick, R. C. Brower, M. A. Clark, T. A. Manteuffel, S. F. McCormick, J. C. Osborn and C. Rebbi, *Adaptive multigrid algorithm for the lattice Wilson-Dirac operator*, *Phys. Rev. Lett.* **105** (2010) 201602 [1005.3043].

[101] M. A. Clark, R. Babich, K. Barros, R. Brower and C. Rebbi, *Solving lattice QCD systems of equations using mixed precision solvers on GPUs*, *Comput. Phys. Commun.* **181** (2010) 1517 [0911.3191].

[102] M. A. Clark, B. Joo, A. Strelchenko, M. Cheng, A. Gambhir and R. Brower, *Accelerating lattice QCD multigrid on GPUs using fine-grained parallelization*, 1612.07873.

[103] F. T. Winter, M. A. Clark, R. G. Edwards and B. Joó, *A framework for lattice QCD calculations on GPUs*, in *Proceedings of the 28th IEEE International Parallel and Distributed Processing Symposium*, pp. 1073–1082, 2014. hep-lat/1408.5925.

[104] M. A. Clark, A. Strelchenko, A. Vaquero, M. Wagner and E. Weinberg, *Pushing memory bandwidth limitations through efficient implementations of block-Krylov space solvers on GPUs*, *Comput. Phys. Commun.* **233** (2018) 29–40 [1710.09745].

[105] S. Mukherjee, O. Kaczmarek, C. Schmidt, P. Steinbrecher and M. Wagner, *HISQ inverter on Intel® Xeon Phi™ and NVIDIA® GPUs*, *PoS* **LATTICE2014** (2015) 044 [1409.1510].

[106] O. Kaczmarek, C. Schmidt, P. Steinbrecher and M. Wagner, *Conjugate gradient solvers on Intel® Xeon Phi™ and NVIDIA® GPUs*, in *Proceedings, GPU Computing in High-Energy Physics*, pp. 157–162, 2015. 1411.4439.

[107] D. Guo, R. Mawhinney and J. Tu, *Solving DWF Dirac equation using multisplitting preconditioned conjugate gradient*, 1804.08593.