PROJECT NARRATIVE

1. SIGNIFICANCE OF RESEARCH

We request a three-year allocation on the Argonne Leadership Computing Facility's Blue Gene/Q, Mira, and on the Oak Ridge Leadership Computing Facility's Cray XK7, Titan, in order to address fundamental questions in high-energy and nuclear physics. The calculations we propose are directly supportive of the very 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 efforts or frontiers (intensity, energy, cold nuclear physics, and hot nuclear physics), as described below.

The advent of petascale computers, such as Mira and Titan, has made a transformational impact on our fields. Indeed, as described in more detail below, these resources have been used by the USQCD collaboration to transform the theoretical landscape in both high-energy and nuclear physics by performing challenging numerical calculations that were previously not possible and by obtaining a number of important results with unprecedented precision. Members of USQCD also continue to play leading roles in the development of algorithms and community codes. In this proposal we seek to build on these achievements to make possible further major advances in high-energy and nuclear physics.

A central goal of the experimental high-energy physics program is to search for direct and indirect signs of new particles and forces. High-precision experiments at the Intensity Frontier aim to search for quantum-mechanical effects of new particles that give rise to tiny deviations from standard-model expectations. The interpretation of these experiments, however, requires reliable and equally precise theoretical predictions. In almost all cases, the precision of these tests are presently limited by our theoretical knowledge of the relevant strong-interaction effects. Lattice Quantum Chromodynamics (QCD) provides the only systematic method for calculating the needed hadronic matrix elements and other QCD parameters with controlled uncertainties. The objective of USQCD is to bring the lattice-QCD errors down to, or below, the experimental ones, in order to maximize the new-physics discovery potential of current and future experiments at the Intensity Frontier.

Experiments at the Energy Frontier aim to directly produce new particles by colliding protons (or other particles) at the highest energies accessible by accelerators. Many proposed beyond-the-standard model theories introduce new strong dynamics to provide a natural explanation for the lightness of the Higgs boson. Lattice investigations are needed to determine the signatures of these new-physics models at the Large Hadron Collider (LHC) so that experimentalists and model builders can identify which, if any, are compatible with data from the ATLAS and CMS experiments. The beyond-the-standard model effort of USQCD is developing quantitative tools for studying new strongly-coupled gauge theories, which may behave quite differently than QCD. The current emphasis is on two well-motivated classes of models - composite-Higgs theories and supersymmetric theories with dynamical SUSY breaking. With the numerical tools in place, lattice-gauge-theory can provide essential quantitative input to aid experiments at the LHC or future colliders in discovering TeV-scale strong dynamics.

USQCD's program in high-energy physics is closely tied to the $700M per year U.S. experimental program. Members of USQCD participated in the APS DPF Snowmass 2013 Community Planning Study, including as the conveners of several working groups, and USQCD's goals are well-aligned with the five "science drivers" enumerated in the DOE HEPAP Particle Physics Prioritization Panel (P5). For example, the proposed calculations of the hadronic contributions to the muon anomalous magnetic moment are needed to identify new-physics effects at the new Muon g-2 Experiment at Fermilab. The quark-flavor-physics calculations support the rare-kaon-decay experiments at CERN and the J-PARC facility in Japan, B-physics experiments at CERN and the KEKB accelerator in Japan, and charm-physics experiments at China's BEPC II collider. Improved calculations of the charm- and bottom-quark masses and strong-coupling constant will be needed for experiments at future colliders such as an ILC or a high-luminosity LHC to identify new-physics effects in Higgs-boson decays. USQCD's work exploring the properties of theories beyond the standard model which involve strong dynamics is laying the essential foundation to investigate any strongly-interacting new-physics sectors discovered at the LHC or other future high-energy colliders.

USQCD's program in nuclear physics is closely tied to the $600M per year U.S. experimental program and to reaching the scientific milestones established by NSAC. It also provides crucial support to other areas of nuclear and particle theory. The component of our program focused on the spectrum and structure of the mesons and baryons, including the study of exotic states, aligns well with, and motivated aspects of, the 12 GeV program at Jefferson Lab, and is evolving to address the anticipated needs of a future electron-ion collider (EIC). Our work on light nuclei, hypernuclei and nuclear forces complements and supports the FRIB facility under construction at Michigan State University, and provides input for studies of dense astrophysical environments. The fundamental symmetries component, which is intimately tied to the hadron structure and nuclear forces components, is essential to efforts to observe and subsequently understand the violation of time-reversal invariance and searches for non-standard model physics at neutron facilities at ORNL and LANL. It is also important to improving estimates of lepton- number violating neutrinoless double-beta decay rates of nuclei in support of the planned ton-scale $0\nu\beta\beta$ experiment.

Calculations performed as part of the USQCD nuclear-physics effort also support the Intensity-Frontier and Cosmic-Frontier experimental programs. Improved neutrino-nucleus cross sections, which depend upon nonperturbative nucleon- and nuclear-decay matrix elements, are needed by DUNE to observe CP violation in neutrino oscillations. If charged-lepton flavor violation is observed by Mu2e, nuclear matrix elements will be needed to interpret the measurements as constraints on underlying new-physics models. These same matrix elements are also needed to interpret dark-matter detection experiments in which the dark-matter particle scatters off a nucleus.

The lattice-QCD program in high-temperature QCD will be crucial to the coming heavy-ion programs at RHIC and the LHC. The higher order moments of baryon number, electric charge and strangeness fluctuations that we will calculate will be measured at the LHC and in the RHIC beam-energy scan to find direct experimental evidence for the QCD phase transition and the possible existence of a critical point at non-zero baryon number density. Our calculations will provide results on the electric charge kurtosis in a large range of baryon chemical potentials and will further constrain the location of a possible critical point in the QCD phase diagram.

Lattice QCD at the Intensity Frontier: A major component of our work under recent INCITE grants has been the study of decays and mixings of strongly interacting particles containing strange, charm, and bottom quarks, in which the quarks change type, or "flavor." This program has been very successful and produced numerous significant results. A central goal has been to determine elements of the Cabibbo-Kobayashi-Maskawa (CKM) matrix, which are the transition couplings between quark flavors. The standard model predicts relations between the elements of the CKM matrix. Hence, precise determinations of the CKM elements from different processes enable tests of these relations and provide a window to search for new physics. Using INCITE resources, members of USQCD have recently calculated the hadronic parameters for semileptonic and leptonic decays of pions and kaons with total theoretical uncertainties at an unprecedented $0.23-0.33\%$, yielding the most precise determinations to-date of the CKM element $|V_{us}|$, and enabling a stringent tests of the unitarity of the first row of the CKM matrix. They have also calculated the hadronic parameters for leptonic D and D_s meson decays, and for semileptonic B-meson decays yielding new determinations of $|V_{cd}|$, $|V_{cs}|$, and the most precise determinations of $|V_{ub}|$ and $|V_{cb}|$ to-date. Some quark flavor-changing processes are mediated in the standard model by the exchange of more than one particle.They are therefore suppressed in the standard model and may receive significant contributions from the exchange of new heavy particles that are large enough to be observed by current or planned experiments. Using INCITE resources, members of USQCD have recently calculated the hadronic parameters of rare semileptonic B meson decays and of neutral B-meson mixing matrix elements, yielding the most precise-to-date determinations of $|V_{ts}|$ and $|V_{td}|$. Similar second-order processes are also of great interest in the kaon system. However, the calculation of hadronic amplitudes for ΔM_K, the long-distance contribution to ε_K, the direct CP violation parameter ε', and other related quantities is much more challenging due to the presence of two pions in the intermediate or final states of the amplitudes. Nevertheless, USQCD members have performed the first complete lattice calculations of several of these quantities, yielding, for example, the first ever standard-model calculation of ε' with control over the theoretical uncertainties, and first-principles explanation of the $\Delta I = 1/2$ rule.

The new lattice-QCD results listed above have, when combined with experimental measurements, revealed many 2–3σ discrepancies that may indicate the presence of new physics. During the coming decade, quark-flavor experiments focused on hadrons with bottom and charm quarks will continue both at new e^+e^- machines in China and Japan, and at the LHC, enabling stronger standard-model tests and more sensitive new-physics searches. Further, a new set of experiments is being mounted at CERN and in Japan to search for new physics in rare kaon decays. These anticipated quark-flavor experiments will require even more precise lattice-QCD calculations than we have performed to date, as well as calculations of new processes.

Other upcoming experiments at the Intensity Frontier, including many based in the U.S., will also depend on lattice QCD. The Muon $g-2$ Experiment at Fermilab expects a four-fold reduction in the experimental uncertainty of the muon magnetic moment. The leading theoretical uncertainties, stemming from hadronic contributions to the muon magnetic moment, will dominate the total uncertainty unless they are improved, and lattice QCD offers the only feasible path for doing so. Matrix elements of protons and neutrons are needed to interpret constraints on CP violation from limits on electric dipole moments, to aid the search for baryon-number violation in proton decay, and even to guide searches for dark matter and axions at the Cosmic Frontier. If charged-lepton flavor violation is observed at the Mu2e experiment at Fermilab or elsewhere, nuclear matrix elements will be needed to interpret the measurements as constraints on underlying new-physics models. The observation of CP violation in neutrino oscillations at DUNE will require nucleon and nuclear matrix elements with greater precision than presently available. Thus, lattice-QCD calculations will continue to be an essential theoretical adjunct to the experimental high-energy physics program.

Lattice Gauge Theory at the Energy Frontier: The first run of the Large Hadron Collider (LHC) led to the discovery of the Higgs boson without any additional hints of beyond-the-standard-model (BSM) physics and without evidence for new dark matter particles on the electroweak scale (WIMPs). Now Run 2 of the LHC has begun an exciting new program searching for BSM physics, with significantly improved discovery potential after the energy upgrade of the accelerator. Hints of Higgs compositeness or supersymmetry may appear early. Future confirmation of the observed diphoton excess at 750 GeV, or of new diboson resonances, could provide the first signatures of a composite Higgs sector. Alternatively, observation of final states with many jets and a large amount of missing energy would favor a supersymmetric interpretation, since composite Higgs models do not typically produce this signature.

All of these scenarios involve new strongly-coupled physics, either directly in the case of a composite Higgs scenario, or indirectly for supersymmetric models with dynamical supersymmetry breaking. Our proposal is a coordinated project to provide lattice input for these BSM scenarios, helping phenomenological studies for the LHC and interpreting new experimental results.

Lattice gauge theorists in the last three years have investigated a number of strongly-interacting gauge theories, finding evidence in certain theories for emergence of light scalar resonances that may serve as composite Higgs candidates. The BSM group of USQCD has led the world-wide effort in this development. Deeper theoretical understanding of these results, including studying the interplay of chiral symmetry breaking with scale symmetry breaking and its dilaton footprint, requires new effort with significant resources.

Equally exciting is the alternate pseudo Nambu-Goldstone boson (PNGB) scenario of the composite Higgs where it is identified as the massless Goldstone scalar from the vacuum alignment of spontaneous chiral symmetry breaking. After mass generation the scalar would remain parametrically light and partial compositeness of the fermions could offer insight into a natural mechanism of fermion mass generation. The two complementary composite-Higgs scenarios are addressed in our program in a coordinated way including common phenomenological studies, and important investigations of composite-dark matter implications of such models.

We are also exploring supersymmetric gauge theories using lattice methods. Supersymmetry provides a natural explanation for light scalar particles because in the absence of supersymmetry breaking the scalars are degenerate with fermions. Moreover, the only consistent theory of quantum gravity that we have available to us, string/M theory, requires supersymmetry, so a unification of all four forces of nature predicts supersymmetric gauge theories at some energy scale. In terms of making predictions for what can be observed in particle physics experiments (either high precision or accelerators), it is crucial to better understand the mechanisms of dynamical supersymmetry breaking, which invariably involve strongly coupled supersymmetric gauge theories. The formation of condensates, and other nonpertubative dynamics, play an important role in this regard. N=1 deformations and the Coulomb branch of the N=4 super Yang-Mills theory studied here may have some application in this arena, in addition to the ways in which our results might impact understanding string theory through the gauge-gravity duality (AdS/CFT).

Cold Nuclear Physics: The USQCD program in cold nuclear physics supports and complements much of the low- and medium-energy nuclear physics experimental program, and is crucial to the success of major elements of the ongoing and planned high-energy physics program. Over the next few years, this program will evolve to provide support for the planned EIC and ton-scale $0\nu\beta\beta$ experiments that are an emerging priority in the field, as identified in the 2015 NSAC Long Range Plan. Hadron spectroscopy, the study of the spectrum of mesons and baryons, plays a pivotal role in developing our understanding of QCD. Results from recent experiments have upended the conventional picture that hadrons have a simple classification scheme, which has ignited a firestorm of interest. Lattice-QCD calculations are critical to this effort, suggesting the existence of new types of particles, and address a unique challenge for our understanding of QCD: how does the complexity of the hadron spectrum emerge from the interactions of the quarks and gluons. In particular, a key goal is to identify the presence of “hybrid” mesons and baryons, exotic states of QCD in which the gluonic degrees of freedom are manifest, and whose discovery is the primary motivation of the GlueX experiment at Jefferson Lab.

A precise knowledge of the structure of the light hadrons, such as the pion, proton and neutron, is essential for many aspects of subatomic physics, from the electronic structure of atoms through to the response of detectors to neutrinos, through to the design of the next-generation hadron colliders. Such an understanding includes the complete and precise decomposition of the mass of the proton, the decomposition of its spin, and a complete three dimensional tomography of the nucleon to be realized through the determination of generalized parton distribution functions (GPDs) and transverse momentum dependent distribution functions (TMDs). These studies support and complement the extensive experimental programs at JLab, at RHIC and at the LHC, and will be essential in mapping the gluonic structure of nucleons and nuclei in the EIC era. They will be an important part of the newly funded DOE Nuclear Theory Topical Collaboration on TMD related hadron structure. The nature and strength of the interactions of the light hadrons with the carriers of the weak force, dictated by their structure, is critical to the success of experiments focused on determining the properties of neutrinos. Related matrix elements are also critical in searches for dark matter and other BSM physics.

Lattice-QCD calculations play a critical role in refining the nuclear forces and interactions that determine the nature of visible matter. Providing the fundamental input into nuclear many-body calculations, these forces are essential in predicting the structure and decays of nuclei, nuclear reaction rates, the response of nuclei in electroweak processes, and the behavior of matter in extreme astrophysical environments. With lattice-QCD calculations, quantities that are difficult (or impossible) to access experimentally can be determined, thereby providing theoretical support that is critical in capitalizing on the investments made in the US experimental program. The forces that describe the stable and long-lived nuclei have been tightly constrained through decades of experiment. However, the uncertainties in the nuclear forces that determine the structure and behavior of neutron-rich nuclei, and of the matter that forms in core-collapse supernovae and other dense environments, are limiting the precision of predictions for the structure and dynamics of these systems. The FRIB experimental program will reduce some of these uncertainties through direct measurements of processes involving short-lived, neutron rich nuclei. To complement this effort, lattice-QCD calculations aim to refine other components of the forces through studies of neutron-rich systems and hypernuclei that are not accessible in experiment.

Generating the observed asymmetry between matter and antimatter in the universe requires the non conservation of both baryon number (B) and charge conjugation and parity combined (CP) during non-equilibrium dynamics in the earliest moments of the universe. The CP-violation and B+L violation present in the standard model is insufficient to produce the observed asymmetry, and a number of current and planned experiments are designed to explore T, CP, B, and L violating processes. In particular, the new U.S.-based DUNE experiment that is expected to start operation in 2024 will search for B-violating proton decay, and neutron-antineutron oscillation experiments are also envisaged. An experiment to measure the neutron Electric Dipole Moment (nEDM), a probe of time-reversal violation, is under development at the SNS at Oak Ridge. A new flagship initiative in nuclear physics is the development and deployment of a ton-scale experiment to search for $0\nu\beta\beta$ of nuclei that, if observed, would provide explicit evidence of the violation of lepton number. In order to constrain the nature of the associated new physics, the magnitude of the neutron EDM and the rates of proton and $0\nu\beta\beta$ decays from the possible new operator structures must be calculated.

QCD at High Temperature and Densities: The thermodynamic properties of matter interacting through the strong force are described by Quantum Chromodynamics (QCD). The only approach capable of providing quantitative results on the collective (thermal) behavior of matter at present is lattice QCD. Experimentally, the phase structure of strongly interacting matter is studied in heavy ion collisions performed at the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory (BNL) and the LHC at CERN, Switzerland. While the latter allows to explore properties of matter at the highest energies and almost vanishing net baryon number, RHIC is capable of colliding heavy ions over a wide range of beam energies that allows to study matter with varying net baryon densities or, equivalently, net baryon chemical potential.

A central goal in theoretical and experimental studies of strongly interacting matter at nonzero temperature and baryon number density is to explore its phase diagram. A first beam-energy scan (BES-I), varying the center-of-mass energy of the colliding beams in the range 7.7 GeV $\leq \sqrt{s_{NN}} \leq$ 200 GeV, has been performed at RHIC in 2010/11 and 2014. This allowed for a first exploration of the properties of matter in a range of baryon chemical potentials $\mu_B \leq 450$ MeV, or equivalently up to about $\mu_B = 3T_c$, with T_c denoting the transition temperature from hadronic matter to a quark-gluon plasma at $\mu_B = 0$. BES-I led to a number of intriguing signals, in particular a pronounced non-monotonic behavior in the $\sqrt{s_{NN}}$-dependence of higher order cumulants of net-proton fluctuations, that has been expected to occur in the vicinity of a critical point. Strengthening evidence for the existence of such a critical point is a central goal of the next beam energy scan (BES-II) scheduled at RHIC for 2019/20 as well as theoretical lattice-QCD calculations. This constitutes a central goal of the physics program at RHIC, as formulated in the 2015 NSAC Long Range Plan.

Higher order cumulants of conserved charges, i.e. net-baryon number, net-electric charge and net-strangeness, have successfully been calculated in lattice QCD. These calculations provided first results on the equation of state at non-zero baryon chemical potential and allowed for a characterization of freeze-out conditions at $\sqrt{s_{NN}} \geq$ 27 GeV. Of particular interest for the comparison with experimental data are up to fourth-order cumulants. However, in order to strengthen theoretical evidence for the existence of a critical point, one needs to go further and get control also over higher order cumulants. First results on sixth-order cumulants exist and the calculation of eighth-order cumulants is desirable to reach control over the convergence properties of Taylor expansions of bulk thermodynamic observables (equation of state) in the entire parameter range accessible in experiments at RHIC. This will also provide Taylor series expansion of cumulants of conserved charge fluctuations. The fourth-order cumulants (kurtosis) measured in experiments can then be calculated in next-to-next-to-leading order Taylor expansions, which is necessary if one wants to extract a possible non-monotonic behavior in the density dependence of these fluctuation observables directly from QCD. Given the huge decrease of statistical errors expected to result from the experimental upgrades performed at RHIC prior to BES-II, solid results on the temperature and density dependence of cumulants of charge fluctuations are needed for comparison (and hydrodynamic modeling) of the experimental results.

2. RESEARCH OBJECTIVES AND MILESTONES

Lattice QCD at the Intensity Frontier: We will continue the calculation of those standard-model parameters and hadronic matrix elements needed to interpret high-precision measurements as standard-model tests and new-physics searches, focusing on delivering results needed by current and upcoming experiments. With the requested resources, we plan to reduce the errors on our existing results in quark-flavor physics. Petascale computers, such as Mira and Titan, are enabling us for the first time perform calculations with the physical masses for all of the quarks entering our simulations. This significant lattice-QCD milestone is substantially improving the precision of our calculations. For some quantities, we will include for the first time effects from electromagnetism and strong isospin-breaking that have until now been small enough to omit. With the requested resources, we will also extend our Intensity-Frontier program to new processes such as the muon anomalous magnetic and rare semileptonic kaon decays, both of which are going to be measured with greater precision within this decade. These new initiatives are only possible because of the availability of petascale computers.

For our Intensity-Frontier work, we will generate gauge-configuration ensembles with two improved gauge actions, domain wall (DWF) and highly improved staggered quark (HISQ). Each has advantages for different aspects of our work. Further, when performing calculations that cannot be checked by direct comparison with experiments, it is essential to validate our results by obtaining them with more than one action. Below we briefly describe each action and the ensembles we propose to generate with them, and summarize the physics projects for which we will use them. The parameters of these ensembles and the resources needed to generate them are given in Tables 1 and 2.

Domain wall quarks: Over the past four years, the substantial capability of Mira has allowed us to create gauge-field ensembles and to carry out extensive calculations using lattice chiral quarks with their physical masses. As a result, we have computed many important quantities in kaon physics with errors on the 1% level or better such as the leptonic decay constants, semileptonic form factors, and mixing matrix elements. We also have computed for the first time in lattice QCD the matrix elements for $K \to \pi\pi$ decays with all uncertainties controlled, obtaining a $\sim 20\%$ total error. These accomplishments open two very promising research directions that will be enabled by the present INCITE proposal.

First, we propose to augment the methods of lattice QCD to include the effects of electromagnetism in kaon and pion decays, allowing important quantities such as the CKM matrix elements $|V_{ud}|$ and $|V_{us}|$ to be determined from experimental measurements to accuracies well below 0.5%. Given the long-range character of the electromagnetic interactions, it is important to work on large physical volumes. Fortunately, for the combination of a gauge action which includes the dislocation suppressing determinant ratio (DSDR) and the Möbius DWF action, we observe discretization effects at the percent level for lattice spacings as coarse as $a \approx 0.2$ fm. Consequently, we can create ensembles with physical-mass, chiral fermions and volumes of $(6\text{ fm})^3$ or larger at reasonable computational cost to be used in calculations which include electromagnetism. We plan to create three ensembles to support this effort. Two will have the same lattice spacing, $a \approx 0.2$ fm, but different spatial sizes of $(4.7\text{ fm})^3$ and $(6.3\text{ fm})^3$, to estimate finite-volume errors. The third, with a lattice spacing of $a \approx 0.13$ fm and a $(6.3\text{ fm})^3$ box, will allow the estimation of nonzero-lattice-spacing errors. The two $(6.3\text{ fm})^3$ ensembles require Mira's capability and are listed in Table 1. The smallest ensemble will be generated with local resources. These ensembles will be an essential part of planned future calculations that include electromagnetic effects.

Project	Existing	2017	2018	2019	Total
Configuration generation					
$32^3 \times 64$, $a \approx 0.2$ fm	10	150	50	0	210
$48^3 \times 96$, $a \approx 0.13$ fm	0	50	50	0	100
$64^3 \times 128$, $a \approx 0.082$ fm	40	14	17	24	95
Analysis					
ΔM_K	26	30	30	20	106
$K^+ \to \pi^+\nu\bar{\nu}$	0	7	7	6	20
Total Time (10^6 Mira core-hours)		178	183	181	

Table 1: Resources requested for projects using chiral fermions with physical masses. For each project, the second column shows the number of configurations that have already been generated or analyzed, the next three columns show how many will be generated or analyzed in each year and the last column shows the total number that will be completed by the end of the proposed work. The last row gives the number of Mira core-hours requested for each year. For the first two ensemble generation projects the independent configurations are separated by four time units while for the third, this separation is increased to 40 time units.

Second, we propose to pursue another research direction closely related in physics objective and technique to the first, which also targets second-order electroweak phenomena but now involving the exchange of two heavy weak bosons instead of a heavy weak boson and a photon. Specifically, we propose the first realistic calculation of the mass difference between the long- and short- lived K mesons, ΔM_K, and the amplitude for the rare kaon decay $K^+ \to \pi^+\nu\bar{\nu}$. The former is the smallest relative energy difference ever measured and offers unique sensitivity to energy scales as high as 1000 TeV, one hundred times larger than those accessible to the LHC at CERN. However, because this quantity has not yet been reliably computed in the standard model, only a very approximate, $\sim 50\%$ comparison between theory and experiment has been possible, even though the experimental value $\Delta M_K = 3.483(6) \times 10^{-12}$ MeV is quite accurate.

The rare kaon decay is the principal objective of the new NA62 experiment now underway at CERN. This process can occur in the standard model only at second order in the weak interactions. The resulting very small size of the expected decay rate, makes this an important place to search for new phenomena that are not part of the standard model. Although the dominant, standard-model contribution to this decay comes from short distances and can be computed using QCD perturbation theory, we estimate that the contribution from scales at or below that of the charm quark is on the order of 25%, making a lattice calculation of these effects important given the NA62 target accuracy of 10%.

The charm quark plays an essential role in both of these calculations and the $a \approx 0.082$ fm, $64^3 \times 128$ ensemble previously generated on Mira is required if the charm-quark discretization errors are to be kept at the 20% level. Thus, both of these calculations can only be performed using 8-rack Mira partitions or larger. We have begun the ΔM_K calculation as part of our current INCITE project and carried out extensive exploratory $K^+ \to \pi^+ \nu\bar{\nu}$ studies on the 512-node USQCD BG/Q machine at BNL. We have efficient BG/Q code that aggressively exploits low-mode deflation using 2000 eigenvectors, all-to-all propagators to evaluate the disconnected diagrams and the all-mode-averaging variance reduction method. The bottom panel of Table 1 lists the resources required for these three analysis projects, including an extension of the existing $64^3 \times 128$ ensemble needed for the ΔM_K and rare kaon decay measurements.

Project	Existing	2017	2018	2019	Total
Configuration Generation	700	150	150	0	1000
B and *D* Decays	0	250	250	500	1000
Muon g-2	400	200	200	200	1000
B-Mixing	0	0	0	1000	1000
Total Time (10^6 Mira core-hours)		140	140	140	

Table 2: Resources requested for the generation and analysis of $a \approx 0.06$ fm gauge-field configurations with physical-mass HISQ quarks. For each project, the second column shows the number of configurations that have already been generated or analyzed, the next three columns show how many will be generated or analyzed in each INCITE year, and the last column shows the total number that will be completed by the end of the proposed work. The last row gives the number of Mira core-hours requested for each year. We measure the pseudoscalar-meson masses and leptonic decay constants as the configurations are created, and the resources for doing so are included under "configuration generation."

HISQ quarks. Staggered-quark actions require significantly fewer floating-point operations to generate gauge-field configurations with a given lattice spacing and set of light-quark masses than other improved actions. Consequently, they have enabled precise determinations of numerous parameters of the standard model. Indeed, the most precise results for the strong coupling constant, the charm- and bottom-quark masses, and several CKM matrix elements have been obtained using staggered quarks. For these quantities, chiral symmetry plays a less important role, but large lattice volumes, made accessible by the lower computational cost of staggered quarks, are essential for accurate control of finite lattice-spacing errors.

Over the last several years, we have generated a library of gauge-field ensembles using the highly improved staggered quark (HISQ) formulation of lattice quarks. These HISQ ensembles include four flavors of quarks: up, down, strange and charm. We tune the strange and charm quark masses m_s and m_c to their physical values, and for each lattice spacing work with three values of the light quark mass, $m_l = (m_u + m_d)/2$, including its physical value. In the first phase of this work, we are generating ensembles with four lattice spacings: $a \approx 0.15$, 0.12, 0.09 and 0.06 fm. Our goal is to obtain 1,000 equilibrated configurations for each ensemble, and we have reached this goal for all of the ensembles with $a \geq 0.06$ fm, except for the $a \approx 0.06$ fm ensemble with physical value of m_l. We have the resources in hand to analyze all of the ensembles smaller than the $a \approx 0.06$ fm, physical-quark-mass one. Thus, in this proposal, we request time to generate the 300 configurations needed to complete $a \approx 0.06$ fm, physical-quark-mass ensemble, as well as to perform several analysis calculations on them that require the capability and capacity of Mira.

Table 2 shows the calculations we propose on the $a \approx 0.06$ fm, physical-quark-mass ensemble, and the resources required in each year to carry them out. These include the contributions of hadronic vacuum polarization to the muon anomalous magnetic moment $(g-2)$, the leptonic decay constants and semileptonic form factors for decays of D and B mesons, and the mixing of neutral B mesons with their anti-particles. With the inclusion of precise data from the $a \approx 0.06$ fm, physical quark mass ensemble to anchor extrapolations to zero lattice spacing, we expect to obtain state-of-the-art results for all of these quantities.

The determination of the muon $g-2$ is an especially high priority for the Intensity-Frontier program given the imminent start of the Muon $g-2$ Experiment at Fermilab in around 2017–2018. It has been measured to better than a part in a million, with a similar uncertainty quoted for the theoretical value. The measured $(g-2)_\mu$ differs from the standard-model expectation by about 3σ, and the new experiment is expected to reduce the experimental error by a factor of four. The leading theoretical uncertainties stem from hadronic contributions, and lattice QCD offers the only feasible path for reducing the errors on these contributions to the precision needed to interpret the upcoming experiment. Nevertheless, such calculations pose a number of challenges, which has spurred the development of several new lattice-QCD methods. Here, we focus on the hadronic-vacuum-polarization contribution to $(g-2)_\mu$ using a promising method that has already yielded this contribution to $(g-2)_\mu$ with about a 2% error on ensembles with coarser lattice spacings. Our goal is to reduce the theory error on the hadronic-vacuum-polarization contribution to $(g-2)_\mu$ to the sub-percent level needed to match the anticipated experimental precision, and the analysis of the 0.06 fm ensemble with this proposal will bring us significantly closer to this target.

Project	Lattice	Lattice Scale	2017	2018	2019
Sextet Composite Higgs	$64^3 \times 128$	0.020/F–0.030/F	5000	4000	3000
Fundamental 8/10 Flavors	$48^3 \times 96$	0.019/F–0.021/F	4000	4000	2000
$\mathcal{N} = 4$ SUSY	24^4	Coupling $g = 1$	0	0	$2 \times 10 \times 4000$
Total Time (10^6 Mira core-hours)			80	80	80

Table 3: Resources requested to generate SU(3) gauge configurations on Mira for the composite Higgs in the two flavor Sextet project, the 8 and 10 flavor fundamental project, and the $\mathcal{N} = 4$ super Yang-Mills (SYM) project. In the second column, we give the lattice sizes and and in the third column the lattice spacing in units of the Goldstone decay constant F in the chiral limit, which sets the Electroweak scale. The last three columns give the number of configurations we plan to generate in a given year for each project. For the $\mathcal{N} = 4$ SUSY project a total of 2×10 lattices will be generated in order to numerically integrate the free energy with two types of boundary conditions to determine the masses of the BPS states. The last row gives the number of Mira core-hours requested each year.

Lattice Gauge Theory at the Energy Frontier: In Year One and Year Two we propose to investigate the composite Higgs mechanism in two models which provide complementary viewpoints of strongly coupled gauge theories. The near-conformal model with a fermion doublet in the two-index symmetric (sextet) representation of the SU(3) color gauge group provides nonperturbative insight into the strong dynamics of a light scalar closest to the conformal window. Recent results show that the model has the very small non-zero β-function relative to other theories that have been investigated numerically. This theory provides a natural setting for a light composite Higgs particle without fine tuning.

Compelling realizations of the pseudo Nambu-Goldstone boson (PNBG) scenario have been recently proposed based on an SU(3) gauge theory with 4 massless and N massive fermions, all in the fundamental representation. In these models, the 4 light flavors are responsible for producing the composite-Higgs doublet, whereas the influence of the N heavy flavors is necessary to produce the observed SM fermion masses, particularly the top-quark mass. The unknown dependence of the low-energy Higgs sector physics on N is of particular interest. We have significant expertise in the $4+8$ version of this model, while the relatively QCD-like $4+0$ system has been studied previously on the lattice. We propose in Year One to investigate the 4+4 flavor system, which draws on our experience both with the 4+8 theory and with the degenerate 8-flavor theory. Domain-wall fermions, which provide enhanced chiral symmetry at finite lattice spacing, are crucial for the study of certain quantities specific to the PNGB scenario, such as the electroweak contribution to the Higgs potential. Extension to the 4+6 flavor theory in Years Two and Three will allow systematic study of the N dependence.

Both composite-Higgs scenarios share a number of common physics goals, despite the difference in dynamical mechanism. Important milestones include: calculation of the low-lying resonance spectrum, which can be used to predict new particles at the LHC and future colliders; calculation of the S-parameter, to compare with electroweak precision constraints; and investigation of Goldstone scattering, which can be used to predict deviations from WW scattering due to the composite sector. Our program in Year Three will also include study of the composite dark matter candidates in each of the models under investigation.

In Year Three, we will shift our focus to supersymmetry, using 36 million hours to study S-duality in N=4 super Yang-Mills using the twisted lattice formulation, which has a minimum of fine-tuning, and an exact supersymmetry. For this study we will use twisted and C-periodic boundary conditions to measure the masses of W bosons and magnetic monopoles on the Coulomb branch of the theory. These are symmetry-protected quantities so we have exact predictions for the masses coming from the continuum. They are related through S-duality, in which the two states are exchanged, so by measuring these masses we will have a nontrivial check on this remarkable relation.

Cold Nuclear Physics: The main objective of this proposal in cold nuclear physics is to perform calculations of important quantities in nuclear physics at the physical pion mass. These calculations will provide a direct connection between hadronic spectroscopy, structure, the physics of nuclei, and fundamental symmetries with the underlying theory of the strong interactions, directly supporting the U.S. nuclear and particle physics experimental programs and providing critical inputs into multiple areas of theoretical nuclear and particle physics. These calculations will be enabled by the production of ensembles of gluon-field configurations at the physical pion mass using capability resources requested in this proposal, which will be used to calculate quark probability amplitudes, followed by hadronic and nuclear correlation functions using USQCD capacity hardware. Sophisticated large-scale analysis techniques will then be applied to these correlation functions to extract the physics that drives our research.

Our previous work has provided strong indications of the presence of exotic mesons in the energy reach of GlueX and the presence of non-exotic "hybrid" mesons, those in which the gluons play an essential role. An excited baryon spectrum at least as rich as that of the non-relativistic quark model is now known to exist, and a common mechanism giving rise to hybrid mesons and baryons has been identified. Previous calculations with heavier-than-physical quark masses (corresponding to pion masses of approximately 240 MeV and 390 MeV) have demonstrated that the finite-volume of the lattice can be used to map the energy dependence of scattering amplitudes, enabling resonance properties and decay rates, to be directly computed from QCD. The proposed program at the physical pion mass will be able to directly confront experiment and, in the case of the excited meson spectrum, have the potential to *predict* the spectrum before it is measured in experiment. Furthermore, calculations of the decays and electromagnetic properties of excited states will play an essential role in experimental analyses.

The ensembles of physical-mass configurations will make it possible to calculate the structure of the nucleon from first principles, explicitly eliminating the systematic uncertainties associated with chiral extrapolations. Pioneering calculations near the physical quark masses have shown that key properties of the nucleon including the charge and magnetic radii, magnetic moment, and quark momentum fraction approach their observed values, albeit with errors $\sim 5\%$ or larger. The definitive and high-precision calculations of nucleon structure that we are proposing will have decisive impact – resolving the $\sim 7\sigma$ discrepancy between experimental measurements of the proton charge radius $\sqrt{\langle r^2 \rangle_p}$ using the muonic Lamb shift and electron-proton scattering. These calculations will reveal the quark and gluon contributions to the nucleon spin, which are a central focus of present experiments at RHIC, JLab, and essential to the experimental program(s) anticipated at a future EIC. Further, they will refine the three-dimensional picture of the nucleon by providing the moments of generalized parton distributions (GPDs), complementing experimental measurements of GPDs at Jefferson Lab. The axial structure of the nucleon, encapsulated in its form factors that will be calculated, defines its response to neutrino fluxes, which is an important component of the response of nuclei. Our plans also include determining the, as yet unmeasured, nucleon scalar and tensor charges, which are essential in interpreting results from the ultra-cold neutron experiments at LANL that are searching for physics beyond the standard model. Finally, matrix elements relevant to the interpretation of neutron EDM experiments will also be determined.

Significant progress has been made in calculating the two-nucleon and hyperon-nucleon interactions and the spectra of the light nuclei at unphysical quark masses. Calculations have been completed for pion masses of approximately 805 MeV and 450 MeV, and are currently being performed at 300 MeV. Important aspects of the strong interactions have been determined with these calculations, and with those from a collaboration in Japan. In particular, it is now known that nuclei become more deeply bound with increasing quark masses. During the last two years, calculations of the magnetic interactions of the light nuclei have been performed, to determine nuclear magnetic moments and magnetic polarizabilities, and the first calculation of the cross

Year	$\approx a$ (fm)	m_π (MeV)	Dims.	MC traj.	Titan Core-hrs
2017	0.076	140	$64^3 \times 128$	4500	145M
2017	0.076	140	$72^3 \times 192$	750	55M
2018	0.076	140	$72^3 \times 192$	2660	200M
2019	0.076	140	$72^3 \times 192$	2660	200M

Table 4: Resources requested to generate Clover-Wilson gauge-field configurations for calculations in cold nuclear physics. The second column column gives the approximate lattice spacing in fm, the third the pion mass, which is close to its physical value, and the third the dimensions of the lattice. The fourth column gives the number of Monte Carlo steps we propose to generate, and the last column the number of Titan core-hours required for the calculation.

section of an inelastic nuclear reaction, $np \to d\gamma$. From the computed magnetic moments, it can be concluded that nuclei, over a wide range of quark masses, resemble collections of "weakly" interacting nucleons, as in nature, and that the deviations from the naive shell-model are small and consistent with experiment. Related calculations of the axial-current matrix elements in light nuclei are now being performed in order to aid nuclear many-body theorists reproduce the axial charges of nuclei and to refine calculations of $0\nu\beta\beta$ rates in nuclei. With the configurations generated in this proposal, we plan to address each of these aspects of light nuclei at the physical pion mass.

During the first year, we will produce *isotropic* gauge-field configurations generated with $n_f = 2+1$ flavor Wilson-clover (W) fermion action of size $64^3 \times 128$, with $m_\pi \sim 140$ MeV, and a lattice spacing $a \sim 0.076$ fm, which is sufficiently fine to render previously employed anisotropy in hadronic spectroscopy calculations redundant. We will also begin production of an ensemble of Wilson-clover gauge-field configurations at a larger volume of $72^3 \times 192$. In the second year, we will continue the production of the ensemble of Wilson-clover gauge-field configurations on the larger volume in order to reduce the statistical uncertainties associated with each of the physics observables of interest. In the third year, we intend to complete the generation of the ensemble of $72^3 \times 192$ lattices. By combining these calculations with those performed in the smaller spatial volume, but at the same pion mass and lattice spacing, the systematic uncertainty associated with working in a finite volume can be parametrically reduced, leaving only calculations at smaller lattice spacings needed to provide a complete quantification of associated uncertainties. The resources required for this program of calculations are shown in Table 4.

During each of the three years, concurrent generation of correlation functions will be performed on USQCD capacity resources as well as with ALCC awards to address the physics we have presented in the areas of hadron spectroscopy, spectrum, nuclear forces and fundamental symmetries programs.

In summary, enabled by the resources we are requesting in this proposal, we will undertake a comprehensive program in lattice QCD that is critical to several thrusts of the U.S. nuclear physics program. In the hadron spectrum component of the program, we propose to calculate the role of tetra-quark contributions to the meson spectrum; the momentum-dependent phase shifts for the lowest-lying resonances for both the iso-scalar and iso-vector states; radiative transitions between excited mesons in order to estimate the photo-production rates at the GlueX experiment; and phase shifts in meson-baryon scattering. In the nucleon-structure component of the program, we propose to calculate vector and axial form factors, including scalar, axial, and tensor charges; moments of the x-dependence parton distributions; and TMDs and moments of GPDs facilitating study of the three-dimensional structure of the proton. In the nuclei, hyper-nuclei and nuclear forces component of the program, we propose to calculate the lowest energy levels of the light nuclei, including the deuteron, ^{3}He, ^{4}He and their hyper-nuclear partners; the nucleon-nucleon scattering parameters and the multi-nucleon forces. We will also calculate the magnetic moments and axial charges of the light nuclei, determining the the short-distance two-nucleon-axial current interaction, L_{1A}, contributing to the proton-proton fusion solar reaction, as well as the magnetic and axial polarizabilities. In the area of fundamental symmetries, the strengths of interactions that can be generated by physics beyond the standard model will be refined, including those related to the search for the violation of time-reversal through electric dipole moments, and initiating calculations required for improved estimates of $0\nu\beta\beta$ decays rates of nuclei.

QCD at High Temperature and Densities:

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 fluctuations observables on small lattices of size $N_\sigma^3 \times N_\tau$, with $N_\sigma/N_\tau = 4$ and $N_\tau = 6,\ 8$, shows that $O(1\ million)$ Monte-Carlo trajectories are needed to obtain Taylor expansion coefficients of 6^{th} Taylor expansion coefficients with sufficiently small errors to be useful for an evaluation of the equation of state. Even higher statistics is needed to constrain the location of a possible critical point in the phase diagram of strong interaction matter. These calculations need to be continued on larger lattices, closer to the continuum limit. Some work on lattices with temporal extent $N_\tau = 12$ are currently going. We intend to contribute to these calculations by adding new data at two low temperature values, and prepare ground for calculations on larger lattices with temporal extent $N_\tau = 16$. Taking into account prospects of using new generations of leadership class computers in later years of this project we optimized our workflow such that the gauge field generation and data analysis can run in multiple streams and the central matrix inversion routines are optimized for single node performance.

During the first year we plan to focus on the generation of gauge field configurations and the subsequent calculation of up to 8^{th} order cumulants on lattices of size $48^3 \times 12$. For this purpose we will perform calculations with the Highly Improved Staggered Quark (HISQ) action. For each temperature value we will generate $O(100.000)$ gauge field configurations separated by 10 Monte Carlo steps. We have developed highly optimized programs that ran very efficiently on Titan at ORNL. The need for $O(1000)$ matrix inversions performed on identical gauge field configurations makes these calculations particularly well suited for simultaneous calculations on thousands of GPUs.

During years 2 and 3 we intend to generate also gauge field configurations on larger lattices of size $N_\sigma^3 \times N_\tau$ with varying aspect ratio $3 \le N_\sigma/N_\tau \le 6$ and $N_\tau \le 16$. This will allow to get control over finite volume effects in the calculation of cumulants of electric charge fluctuations. While the finite volume studies will be restricted to the calculations of up to 4^{th} order cumulants we intend to improve statistics on higher order cumulants on lattices of size $48^3 \times 12$ in order to further constraint the possible location of a critical point in the QCD phase diagram. Details of the run parameters, in particular the temperature values at which these calculations will be performed, will depend on the experience gained with the runs scheduled for year 1. We plan to use BG/Q resources for the generation of new gauge field configurations on our largest lattices, $64^3 \times 16$.

Year	Lattice	Project	Total configurations generated/analyzed	Titan Core-hrs	Mira Core-hrs
2017	$48^3 \times 12$	gauge conf. generation	160,000	53	-
		calculation of cumulants	160,000	47	-
2018	$48^3 \times \{12,\ 16\}$	gauge conf. generation	100,000	40	-
	$48^3 \times \{12,\ 16\}$	calculation of cumulants	110,000	35	-
	64×16	gauge conf. generation	10,000	-	25
2019	$48^3 \times \{12,\ 16\}$	gauge conf. generation	100,000	40	-
	$48^3 \times \{12,\ 16\}$	calculation of cumulants	110,000	35	-
	64×16	gauge conf. generation	10,000	-	25

Table 5: Resources requested for the study of QCD at high temperatures. The second column gives the lattice dimensions, and the fourth column the total number of configurations to be generated or analyzed. The fifth and six columns give the number of Titan and Mira core-hours, required for the task in units of millions of core-hours.

3. COMPUTATIONAL READINESS

3.1 Leadership Classification:

The work we propose requires the capability and capacity of leadership class computers. The generation of gauge configurations, especially ones with physical pion masses, presents a capability challenge. This component of our work is the major bottleneck in our calculations, because each configuration in an ensemble is generated from the previous one. Thus, configuration generation requires long runs on large partitions of the most capable available computers. Configuration generation also presents a capacity problem, as it must be carried out for a range of lattice spacings in order to perform extrapolations to the continuum limit, and for several different formulations of quarks on the lattice. Analysis routines also present a capacity challenge because of the large number of configurations on which they must be performed. However, the analysis of different configurations can be carried out in parallel, so leadership class computers are needed for this phase of our work only for the most challenging gauge ensembles

3.2 Computational Approach:

3.2.1 Underlying mathematical formulation: In lattice QCD and lattice field theories that have been proposed for the study of physics beyond the standard model, physical observables can be expressed in terms of Feynman path integrals on regular four-dimensional, hyper-cubic lattices,

$$\langle O \rangle = \frac{\int \mathcal{D}U \, O(U) \exp[-S(U)]}{\int \mathcal{D}U \exp[-S(U)]} \, .$$

Here, U represents the gauge field, which consists of one matrix assigned to each link of the lattice. Typical state of the art lattices now have $O(10^7 - 10^9)$ matrices. For QCD, each of the components of U is a 3×3 complex matrices that is a member of the group SU(3), whereas for beyond the standard model theories elements of U may belong to other groups. $\mathcal{D}U$ represents an integral over all of the elements of U with the Haar measure, and $S(U)$ is the Euclidean action of the lattice theory after the quark fields have been integrated out. O is the physical observable being studied, and $O(U)$ is the value of O in the gauge configuration U. The first step is to use importance sampling techniques to generate an ensemble of gauge configurations U_i, $i = 1, \ldots N$, with the probability distribution

$$P(U_i) \sim \exp[-S(U_i)] \, .$$

Once an ensemble of representative gauge configurations is available, an unbiased estimator for the observable O is given by

$$\langle O \rangle = \frac{1}{N} \sum_i^N O(U_i) \, .$$

3.2.2 Algorithms and numerical techniques: The act of integrating out the quark fields makes $S(U)$ highly non-local since it involves the determinants of the Dirac operators for the quarks. Local algorithms in which one updates one element of U at a time are prohibitively expensive. We make use of the Rational Hybrid Monte Carlo (RHMC) algorithm in our work. It dramatically reduces the time needed to generate configurations relative to earlier algorithms, especially for the configurations with light quark masses, which are our main interest. The RHMC algorithm requires that one numerically integrate a set of molecular dynamics equations, which are non-linear, coupled first order differential equations. We do so using either the a multi-time step Omelyan integrator or the order $(\Delta t)^4$ force gradient integrator. At each step of the integration one must solve a set of linear equations of the form $(M + \sigma_j I)x_j = y_j$, where I is the unit matrix, the σ_j are real positive numbers, and M is either the Dirac operator for a quark, a large sparse matrix of dimension of order the volume of the lattice, or the product of the Dirac operator with its Hermitian conjugate. In the former case, this shifted linear system of equations is solved by a multi-shift or multi-mass version of the biconjugate gradient stabilized (BiCGSstab) algorithm, and for the second by a multi-mass version of the conjugate gradient (CG) algorithm. These solvers, along with the calculation of the force that enters into the integration of the molecular dynamics equations, consumes the bulk of the floating point operations in configuration generation. The majority of the floating point operations in the analysis calculations go into solving linear equations of the form $Mx = y$, for which we employ the usual single mass BiCGstab or CG algorithm, combined with various acceleration techniques described below.

Petascale computers, such as Mira and Titan, are enabling us to carry out simulations with smaller lattice spacings than previously possible using, for the first time, physical values of all quark masses. This development is greatly improving the precision of our calculations, but it requires major improvements in our solvers. To this end we have developed adaptive multigrid, all-mode-averaging, and deflation algorithms for our work. These algorithms have given speedups of as much as a factor of twenty for some of our most challenging cases.

3.2.3 Programming languages, libraries and other software: Under grants from the DOE's SciDAC Programs, the USQCD Collaboration developed the QCD Applications Program Interface (QCD API), which will be used in this project. The QCD API has three layers. The first consists of two libraries, QLA and QMP. QLA contains routines for linear algebra operations involving the complex, 3×3 matrices that are the basic elements in lattice QCD calculations, and other matrix operations for gauge theories in the BSM investigations. These operations are local to lattice sites or links, and do not involve inter-processor communication. There are both C and C++ versions of QLA, and some key routines have been written in assembly language for various processors. QMP is a communication library containing routines with the functionality of that portion of MPI relevant to lattice QCD calculations. It also has extensions that 1) partition the QCD space–time lattice, and map it onto the geometry of the hardware network; and 2) contain specialized routines designed to aid in the use of low level protocols on parallel computer networks, such as SPI on the Blue Gene/Q. A version of QMP that runs over MPI has been developed to insure portability of codes.

The second layer of the QCD API is QDP, a library of data parallel operations that are built on QMP and QLA. QDP makes it straightforward to develop highly efficient, portable codes. Versions exist in both C and C++. The third layer of the QCD API, QOP, is a set of highly optimized routines that perform the bulk of the floating point operations in our work. QOP routines can be called from any code that conforms to the QCD API. Finally, an I/O library, QIO, enables users to read and write the different types of files that arise in our work in standard formats, which enables sharing of the large data sets that are created in the generation of gauge configurations and in their analysis. QIO supports a logical partitioning of the computer into I/O partitions with one core per partition handling I/O for the data in just that partition. Thus, our codes can read and write data in multiple files. The files thus created can be flattened into one large file offline on a single processor machine. There are no unusual memory requirements for this process. Likewise, for input one of our large files can be fragmented on a single-processor machine, copied to local disks, and read by the individual partitions. By tuning the size of the I/O partitions, we can maximize the I/O bandwidth and avoid contention.

We have made a major effort to develop lattice gauge theory code for GPUs. The code has been incorporated into a new library, QUDA (QCD in CUDA). The first step was to develop code for all Dirac solvers commonly used in QCD (Clover-Wilson, HISQ/asqtad, DWF, and twisted mass). The GPU solvers have already had a major impact on our physics analysis work. We have developed GPU codes for the full evolution of Clover-Wilson and HISQ gauge configurations. The major QUDA codes run on multiple GPUs with support for partitioning of the lattice in all four space-time dimensions. The QUDA library incorporates a number of sophisticated optimizations. For example, core GPU kernels are subject to an automated tuning procedure which optimizes the assignment of thread blocks to the streaming multi-processors in a GPU. GPUDirect is also utilized, enabling peer-to-peer communication between GPUs sharing a PCIe bus.

3.2.3 Parallel programming model: Our lattices are divided into identical four-dimensional sub-lattices, each of which is assigned to a node. Communications between nodes are handled by message passing, MPI in the case of Crays, and SPI in the case of the Blue Gene/Q. Threaded codes, OpenMP, CUDA threads, are used on nodes and on GPUs.

3.2.4 Project workflow: As indicated in section 3.1 above, we plan to perform configuration generation on the leadership class computers. This component of our work requires long runs on large partitions of the most capable available computers. Although physics analysis on a configuration generally requires more floating point operations than were needed to generate it, configurations can be analyzed in parallel within a single job or multiple ones. In addition, the same configurations can be used to study a wide range of phenomena, again in parallel. Thus, physics analysis usually presents a capacity problem, rather than a capability one. Therefore, to the extent possible physics analysis on configurations will be performed on dedicated clusters funded for our collaboration by the DOE, and with time at other supercomputer centers. However, for the most challenging configurations, analysis work must also be carried out on leadership class machines, if results are to be obtained in a timely fashion. The analysis jobs are dominated by sparse matrix inversions using the BiCGStab or CG algorithms, so they are also suitable for running on large partitions.

3.2.5 Workflow software: We have written a number of scripts which help to manage workflow. These scripts provide input parameters for jobs and in many cases submit new jobs when previous ones finish. In our analysis work, we use scripts to combine calculations of different observables in order to reuse expensive quark propagators whenever possible.

3.2.6 I/O requirements: In configuration generation jobs, one reads a single binary file consisting of the gauge configuration at restart, and writes this file for check pointing and to save configurations for physics analysis. The largest files we will deal with in the proposed work will contain the $96^3 \times 192$ HISQ configurations, which are approximately 48.9 GB in size. As indicated in Section 3.2.3, we will use the QIO library to read and write files in parallel. This can be done in one to two minutes so I/O consumes a very small fraction of the time used for configuration generation. Furthermore, our analysis work will also not strain the I/O systems of Mira or Titan. For an analysis job on the $96^3 \times 192$ HISQ gauge configurations, one would read in a single 48.9 GB configuration file, and write a number of quark propagator files which are a factor of three larger. As is the case for gauge configuration files, the propagator files are read and written in parallel in binary.

3.3 Development Work and Code Performance: We have an ongoing program to continue optimization of our codes for the Blue Gene and Cray architectures, which is funded by our DOE SciDAC grants. Members of USQCD are also collaborating with Intel engineers under three different NESAP grants to prepare our codes for the coming KNL architecture.

Members of USQCD at Brookhaven National Laboratory and Columbia University, along with colleagues at the University of Edinburgh, participated in the design of the Blue Gene/Q, and at a very early stage developed high performance code for QCD with DWF quarks. The DWF solver runs at 30.5% of peak, while code for the full evolution of gauge configurations and analysis runs at up to 25% of peak with exact performance depending on the specific problem. The scaling behavior of the solver is illustrated in the left hand panel of Fig. 1.

Members of USQCD at Jefferson Laboratory have a long term project to develop high performance code for QCD with Clover-Wilson quarks. This work has included the development of new algorithms to extend the scaling of the Clover-Wilson solver, and to run full gauge configuration generation on Titan's XK nodes. Their success is illustrated in the right hand panel of Fig. 1. The work on lattice generation code is based on Just in Time (JIT) compilation of the Chroma code.

Hand-written code for the HISQ solver achieves up to 23% of peak on Mira, and exhibits excellent weak scaling. We are working on a new code framework that will be better optimized for future architectures, such as KNL, and may also bring performance improvements for the BG/Q.

Our other codes also exhibit near perfect scaling with performances that vary with the lattice action and the specific problem being studied. For example, the thermodynamics code, which makes use of the GPUs and CPUs in parallel, achieves 23% of theoretical peak on Titan with excellent weak scaling through 4096 nodes. The beyond the standard model work with the sextet action uses HISQ fermions, and achieves a performance of approximately 30% of peak on $64^3 \times 128$ lattices on Mira, and demonstrates excellent strong scaling through 16 racks. The work on the composite Higgs model in the fundamental representation makes use of the DWF code discussed above, which achieves approximately 25% of peak on Mira.

3.4 Use of Resources Requested:

3.4.1 Job Description: In Section 2 we provide tables showing the specific calculations we plan in each of the four areas of research we propose. The text in that section describes each of the projects, and indicates how these calculations are related to the goals set out in Section 1.

3.4.2 Processor/core Utilization: The bulk of the jobs we plan for Mira will run efficiently on four to eight racks, and the four rack ones can be bundled to run on eight racks. We have a limited number of jobs that will run most efficiently on two racks. Our jobs on Titan run on 4096 nodes or more. Our configuration generation jobs consist of long streams in which each configuration evolves from the one before it. Minimum running times range from one to two hours, but we on both machines we would prefer to run jobs as long as allowed by the queue limits. We also prefer to run long jobs for the analysis projects.

3.4.3 Calculation of Resource Request: Our resource requests for each project in each year are shown in the tables of Section 2. The resources needed to generate a single molecular dynamics time unit (Monte Carlo step) and to analyze a single configuration are known from a production runs already in progress with the desired parameters, or have been determined by timing runs with the production code and planned parameters. Our estimates of the number of time units needed to reach our physics objectives are based on past experience. They vary significantly from calculation to calculation.

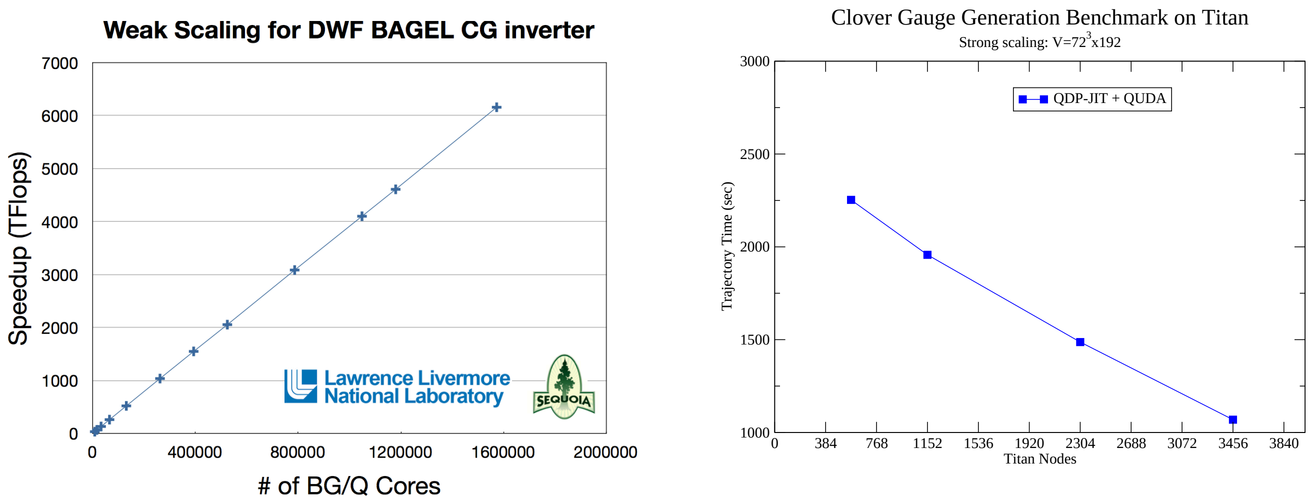

Figure 1: The left panel shows the performance in teraflop/s of the DWF solver on the Lawrence Livermore National Laboratory Blue Gene/Q, Sequoia, as a function of the number of compute cores on up to 98 Blue Gene/Q racks. This is a weak scaling plot with the number of lattice points per core fixed at 8^4. The right panel shows the performance of the Clover-Wilson solver on Titan for the $72^3 \times 192$ lattice we plan to use in production runs. This is a strong scaling test in which the lattice size remains fixed as the number of XK nodes (GPUs) increases.

3.4.4 Anticipated Annual Burn Rate: Production codes exist and have been thoroughly tested for all of the calculations proposed, except the SUSY model to be studied in 2019. Work on the SUSY code is underway. Further algorithm development work will be performed for on workstations and clusters. So, all of the time requested at ALCF and OLCF will go into production jobs. We would prefer to use all of our requested time for configuration generation as early as possible in the year so that we can exploit the configurations for physics analysis at the earliest possible time. During the past three years we have found it possible to front load our work at both centers, and we will attempt to do so again under this grant. If that is not possible, then we will run at steady state throughout the year.

3.5 Parallel Performance: In performing lattice QCD calculations one divides the lattice into a set of identical sub-lattices, and assigns one sub-lattice to each core or node. The number of lattice points on a core or node is often referred to as the local volume. In most cases, the calculations performed on the various cores are identical, and only require data from a few neighboring cores. The major exceptions are the global sums that enter the conjugate gradient algorithm used for sparse matrix inversion. Thus, solvers for the Dirac operator provides the most stringent test of the scaling of our codes. It also gives a good indication of their overall performance. We illustrate the scaling behavior of our codes in Fig. 1, and indicate the performance of each of our codes in Section 3.3.