Office of High Energy Physics
Report on the

LQCD-ext III/NP LQCD Initiative
2023 Progress Review

May 22-24, 2023

Executive Summary

The Progress Review of the Lattice Quantum Chromodynamics extension III (LQCD-ext III) research program of the Office of High Energy Physics (HEP) and the Office of Nuclear Physics (NP) LQCD Initiative of the NP was held on May 22-24, 2023, via ZOOM. The purpose of the review was to assess the program's progress towards its overall scientific and technical goals, and to assess the role of the National Lattice Chromodynamics Collaboration (USQCD) collaboration in governing the usage of the program's hardware. In particular, the LQCD-ext III/NP LQCD Initiative team was instructed to address seven charges:

1. The continued significance and relevance of the LQCD-ext III program and the NP LQCD Initiative, with an emphasis on its impact on the experimental programs of the Department of Energy (DOE), NP and HEP.

2. Progress toward scientific and technical milestones.

3. The status of the technical design and proposed technical scope for fiscal year (FY) 2023-2024.

4. The merits of updating the HEP LQCD-ext III research program and the NP LQCD Initiative research program to include the construction and operation of dedicated hardware.

5. The feasibility and completeness of the proposed budget and schedule, including the preliminary FY2025-2029 5-year plan for the NP LQCD Initiative.

6. The effectiveness of the proposed management structure, and responsiveness to any recommendations from last year's review.

The USQCD collaboration also addressed the charge:

1. The effectiveness of USQCD in allocating the LQCD resources to its community of lattice theorists, the scientific impact of this research on the entire HEP and NP communities and the status, operational procedures, and related activities of the USQCD collaboration itself.

In general, the review panel was very impressed with the technical and scientific achievements of LQCD-ext III/NP LQCD Initiative and USQCD, but they presented several suggestions and recommendations. They noted the impacts of the program's simulations on experimental programs in precision measurements of the Standard Model (SM), Beyond the Standard Model (BSM) model building, hadronic matrix elements, nuclear structure, and neutrino interactions. The review singled out the USQCD's organization of the HEP theory community to present a consensus for the Fermi National Accelerator Laboratory (FNAL) Muon g-2 (anomalous magnetic moment) experiment (muon g-2) measurement now being made by E896 at FNAL. The timeline for publishing lattice predictions for the muon g-2 was examined in detail and its coordination with the publication of experimental results was noted and commended. The review panel had two specific recommendations:

1. LQCD-ext III, with the guidance of HEP, should seriously consider taking the lead in the creation of one or several topical collaborations, following the successful program of topical collaborations by the NP, focused on neutrino nucleus interactions with the goal to provide experimentally relevant predictions.

2. The review team endorsed the program's request to allow the purchase of dedicated hardware for LQCD capacity computing in appropriate circumstances. The review team argued that this action would mitigate the risk associated with Institutional Cluster (IC) updates.

The HEP and NP requested that the program, LQCD-ext III/NP LQCD Initiative, present plans to implement these two recommendations within the next 6 months at one of their regular monthly meetings with HEP and NP.

Introduction and Background

The DOE Offices of Advanced Scientific Computing Research (ASCR), HEP and NP have been involved with the USQCD in hardware acquisition and software development since 2001. The LQCD IT hardware acquisition and operations project, which started in 2006 and ran through 2009, operated a "Quantum Chromodynamics-on-a-chip" machine at Brookhaven National Laboratory (BNL), and built and operated special-purpose commodity clusters at the FNAL and the Thomas Jefferson National Accelerator Facility (TJNAF/JLAB). The project's

4-year budget was $9.2M. LQCD met its 2009 project goal of providing 17.2 Teraflops of sustained computer power for lattice calculations. The clusters designed, constructed, and operated by LQCD complement the lattice community's access to supercomputers. Supercomputers produce the gauge configurations and quark propagators of Quantum Chromodynamics and the clusters and other hardware platforms of LQCD run the programs that analyze those configurations and compute matrix elements and predict cross-section and rates of decay processes.

The hardware project, organized by the USQCD collaboration of ~120 lattice computational physics theorists, successfully completed its original 4-year allocation. The collaboration then proposed and was granted an extension project, LQCD-ext, which ran from FY2010-2014. LQCD-ext worked with a robust budget of $22.9M. The project pioneered the use of graphics processing units (GPU) and this new "disruptive" technology helped the project exceed its original milestones by a wide margin.

The second extension of the project, LQCD-ext II, was described by the USQCD collaboration in a proposal entitled, "LQCD-ext II Computational Resources for Lattice QCD: 2015-2019" dated October 23, 2013. This document presented the scientific objectives, the computational strategy, and the hardware requirements of the LQCD-ext II project. The scientific content of the proposal was reviewed successfully on November 8, 2013, and the scientific vision and specific goals of the project were enthusiastically endorsed by a panel of scientific experts. The reviewers recommended full funding, $23.4M for the 5-year period. However, due to budget constraints, the HEP and NP provided budget guidance to the project team of between $14M and $18M for the 5-year period, well below the project's request. These plans became the basis for the project team's planning for LQCD-ext II. That project passed its critical decision (CD)-1 review on February 25, 2014, and was granted CD-1 approval on May 1. It held its CD-2/3 review on July 10 and was approved on October 1, 2014.

The budget planning for the LQCD-ext II project was of some concern to the review panels of 2014 and 2015. The original 5-year budget of $23.4M ($4.68M per year) proposed by the collaboration and endorsed by the November 8, 2013, Science Review resulted in the following anticipated Teraflops profile from FY2015 to FY2019:

Full Funding Scenario ($23.4M)	FY2015	FY2016	FY2017	FY2018	FY2019
Planned computing capacity of new deployments, TeraFlops	165	233	330	467	660

However, funding at the $14M level followed funding profile:

	FY2015	**FY2016**	**FY2017**	**FY2018**	**FY2019**	**Total**
HEP	1.0	2.0	2.0	2.0	2.0	9.0
NP	1.0	1.0	1.0	1.0	1.0	5.0
Total	**2.0**	**3.0**	**3.0**	**3.0**	**3.0**	**14.0**

The estimated Teraflops profile was reduced to:

Reduced Funding Scenario ($14.0M)	FY2015	FY2016	FY2017	FY2018	FY2019
Planned computing capacity of new deployments, TeraFlops	0	107	160	244	358

which was a 53 percent reduction in compute power compared to the full funding scenario. This reduction in computing capacity challenged USQCD to maintain its productivity, its balance with its Leadership Class computing allocations and its international standing. The 2014 review panel commented on these developments since they influenced the use and productivity of the FY2014 hardware acquisitions they endorsed. The 2015, 2016 and 2017 review panels also commented on the extra challenges that constrained funding placed on the project and they noted that any additional funding would directly increase the project's hardware acquisition plans. LQCD-ext II managed its computing resources wisely over this time period and the productivity of the project increased accordingly even with a less-than-optimal budget.

Over the course of the project and its extension, 2006-present, the hardware acquisition strategy of LQCD had been essential to its success. Each year the project's technical personnel benchmark the kernels of the QCD code on the newest cluster, GPU and supercomputer hardware, and the winner of the price-to-performance competition becomes next year's provider.

The usage of the hardware procured by LQCD has been governed by the USQCD collaboration through its Executive Committee (EC) and Scientific Program Committee (SPC). In addition, the collaboration organizes the community's access to the DOE Leadership Class Supercomputers available through the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. Members of the USQCD collaboration submit proposals through the EC for computer time, some on the Leadership Class machines for large-scale capability computing, and some on the dedicated clusters of LQCD for large scale capacity computing. Allocations on the dedicated clusters of LQCD are awarded by the SPC based on a merit system. Three classes of applications for computer time allocations on the dedicated LQCD hardware are distinguished, these being large-scale mature projects (allocation class A), mid-sized projects (allocation class B), and exploratory projects (allocation class C). Suitable computer platforms are assigned to the various projects upon approval. The clusters of the hardware project analyze and compute matrix elements from the gauge field configurations generated on Leadership Class machines. This strategy requires a balance between the compute power of the clusters and the Leadership Class machines.

Following recommendations from past reviews, a Science Advisory Board (SAB) was formed in 2013 and has participated in the USQCD allocation process. The SAB brings the perspective of the broader HEP and NP community into the high-level decision-making processes of USQCD and is meant to guarantee that the goals of the lattice effort reflect the diverse needs, challenges and interests of high energy and nuclear researchers. The SAB consists of 7 members, 4 experimentalists and 3 theorists. They comment on the science goals of USQCD, the effectiveness and fairness of the allocation process and participate in the annual all-hands meeting.

In addition to the original hardware project LQCD, USQCD has also played a role in software development through the Scientific Discovery through Advanced Computing (SciDAC) program. USQCD was awarded a SciDAC-I grant (2001-2006) which was used to develop efficient portable codes for QCD simulations. USQCD was subsequently awarded a second "SciDAC-II" grant (2006-2011) to optimize its codes for multi-core processors and create a physics toolbox. These SciDAC grants supported efforts to provide a user interface to lattice QCD which permits the user to carry out lattice QCD simulations and measurements without the need to understand the underlying technicalities of the lattice formulation of relativistic quantum field theories and its implementation on massively parallel computers. In 2012 USQCD submitted two proposals to the SciDAC-III program, and both were funded, one through NP and ASCR, and the other through HEP and ASCR. In 2017 USQCD submitted two proposals to the SciDAC IV competition, and the NP proposal was funded.

USQCD organized and submitted a proposal, the Exascale Computing Project, to ASCR's Exascale program. That effort was funded and USQCD is actively involved in preparing for the next era in computer power.

The precision and relevance of the lattice community's calculations have improved steadily over the years. Lattice calculations now come with detailed error analyses. The experimental community has taken note of this important development and looks to lattice calculations in their planning. In order to impact the experimental and theoretical programs of NP and HEP, the collaboration has been encouraged to organize workshops where it can interact with the other communities and actively disseminate its program. These workshops have been successful in engaging a wider audience for the lattice calculational program and now USQCD's contributions to many annual workshops are widespread and generally significant and appreciated. In 2018 HEP converted the lattice hardware project to a research program. The project was told to move away from dedicated hardware and begin supporting and using laboratory-based Institutional Clusters. The motivation for the policy change was HEP's intention of providing a more level playing field in computing across its entire program. NP had different priorities and decided to continue supporting dedicated hardware at TJNAF, outside of the HEP project. However, to maintain the scientific productivity of the effort, the engagement of USQCD in the lattice effort was not changed and the allocation process for its members to use both HEP and NP facilities remained unchanged.

BNL indicated strong interest in supporting lattice gauge theory in 2018. David Lissauer, the Deputy Associate Laboratory Director for Nuclear and Particle Physics, and Kerstin Kleese van Dam, the Director of the lab's Computation Science Initiative, participated actively in the 2018 annual progress review. BNL staff indicated their commitment to the Institutional Cluster funding model. The LQCD-ext II project team indicated that interactions with BNL had been productive and successful. The BNL staff endorsed the project team's commitment to discover and procure the most cost-effective hardware chips for lattice gauge theory simulations each year. Several meetings between the project team and the FNAL computing division were scheduled over the next few months to formulate a plan to work together to design, purchase hardware and use institutional clusters.

Toward the end of the LQCD-ext II project in 2018-2019, the project and USQCD put together a plan for the next 5-year period, 2020-2024. USQCD prepared seven whitepapers describing its 5-year plans in all its research subfields and the whitepapers were published in a special edition of a European research journal. That plan was reviewed in a special comprehensive HEP review which pitted the lattice effort against all the theory research efforts supported by HEP. This review occurred in Rockville, MD, on July 9-10, 2019. A review team of 7 outstanding theorists, under the direction of Bill Kilgore, the program manager of HEP theory, evaluated the lattice program and recommended that it continue for another 5 years. They supported increasing its funding from the requested amount of $2.0M per year to $2.5M per year to address storage needs and new growth opportunities. The review team also endorsed the institutional cluster model for the program.

The second review of the new 5-year cycle occurred on May 18-19, 2021. This review took place via ZOOM. The review consisted of 2 days of presentations and a third half-day of questions and answers, report writing, and a closeout session.

Six expert reviewers from high energy theory, phenomenology, and computer science participated in the review. The review began with a presentation by Andreas Kronfeld, spokesperson for USQCD, which gave an overview of the USQCD collaboration and the LQCD-ext III project. His deputy, Robert Edwards, followed with more details on the collaboration, its structure, governance and accomplishments in science and personnel. Then there were 4 presentations on the scientific topics which comprise lattice gauge theory. Finally, Andreas Kronfeld returned with a discussion of USQCD's plans for the future.

This review report presents the results of the most recent May 22-24, 2023, virtual review of the program, the third review of the new 5-year cycle. The review was expanded to include the NP LQCD Initiative which supports the hardware at JLAB and provides $1M/year. The review consisted of 2 days of presentations and a third half-day of questions and answers, report writing, and a closeout session. The Appendices to this report provide additional detailed material relating to the review: Appendix A contains the charge letters to the LQCD program management team, Appendix B lists the reviewers and DOE participants, and Appendix C contains the agenda.

Seven expert reviewers from high energy and nuclear physics theory, phenomenology, and computer science participated in the review. The review began with a presentation by Robert Edwards, spokesperson for USQCD, and followed the format of the 2021 review in close detail, as shown in Appendix C.

The remaining sections of this report present the findings, comments, and recommendations of the review committee for each of the seven charge elements that the LQCD-ext III/NP LQCD Initiative team was asked to address in their charge letters.

Analysis of the seven charge points

1. The Program's Continued Significance and Relevance

Findings

The USQCD Collaboration has been and continues to be extremely influential in the paths taken by the nuclear and particle physics communities in the United States. Members of the collaboration are involved actively and significantly in the Nuclear Science Advisory Committee

(NSAC) and the Particle Physics Project Prioritization Panel (P5) process in many ways:

a) through participation

b) occupying numerous leadership positions and

c) producing results impacting with high citations, which is the most direct evidence of this group guiding the planning process in the country.

The Collaboration has also been actively organizing workshops around the recent and ongoing P5 and NSAC long range planning process, respectively. Specific subfields of this group's influence include hadron structure, physics of light and heavy quarks which directly impact our knowledge of precision electroweak and exploration of beyond-the-Standard Model physics.

Comments

One hundred and seventy people are involved in the entire USQCD collaboration. Communication between them, their elected management, the DOE site offices, and the DOE is working very well and has resulted in a very successful enterprise sustained over 15+ years. Everyone involved is to be commended.

Since lattice QCD simulations have been possible at physical quark masses, LQCD is in an era where its results are directly impacting particle and hadron phenomenology; see for instance the Flavor Lattice Averaging Group (FLAG) review. USQCD is strongly represented overall in that review, and sometimes dominant. Machine-learning based methods already have interesting applications in computing observables or in the data analysis stage for QCD. There is potential for interdisciplinary exchanges of methods and ideas with other areas of science.

Recommendations

None

2. The Progress toward Scientific and Technical Milestones

Findings

Overall, significant progress is being made towards reaching the milestones laid out in the program's proposal. A "living" website showing timelines and the relation to experiments has been created.

The allotment of computing resources is dictated by well-defined and well-understood rules developed within the collaboration. Computing allocations get a monthly review. Reductions from the initial computing allocations are rare but do happen based on monitoring the burn-rate. Most common causes of reductions are due to fluctuations in personnel working on a topic or code issues that sometimes remain unresolved.

Specific comments on the various areas are listed below:

Light quark and lepton physics:

For muon g-2, a notable advance is the agreement between several lattice calculations of the Hadron Vacuum Polarization (HVP) at various length scales, showing tension with the data-based methods. This solidifies the previously reported tension. Progress towards the ultimate per-mille accuracy goal for the HVP contribution is steady. Substantial progress on the Light-by-Light contribution, and the target precision goal is nearly achieved.

The first calculation of $\pi^0 \rightarrow e^+ e^-$, which serves as a benchmark for the more difficult calculation of $K_L \rightarrow$ $mu^+ mu^-$, the latter having additional contributions from weak and two photon exchange. The experimental measurement is extremely accurate. LQCD will push for a reduction in uncertainty to match the experimental accuracy.

Steady progress on the calculation of K→pi pi weak decay amplitudes was reported. The introduction of a new approach using periodic boundary conditions allows very important cross checks. Results for the pi-pi scattering amplitude in the isospin I= 0,2 channels at physical quark masses agree with experiment.

New methods to calculate a range of new quantities are being pioneered: long-distance contribution to ϵ_K, mass splitting of K_L-K_S, and many others.

Beyond the Standard Model (BSM):

A broad program studying "cousins" of QCD was presented, with different numbers of colors and flavors and representations, as well as with supersymmetry, with the aim of finding new classes of phenomena that arise nonperturbatively, as well as determining nonperturbative matrix elements relevant for dark-matter searches in particular classes of models.

Novel methods introduced in the study of BSM theories are feeding back into LQCD calculations.

Heavy Quark Flavor Physics:

LQCD is a vital contributor to knowledge of B and D heavy quark physics at precisions that could expose physics beyond the SM. The groundwork has been laid out for precision predictions which will be comparable to upcoming precise experiments.

Hadron Physics:

The impressive power of recent lattice calculations was displayed in the talks, especially with regards to the Parton Distribution Functions (PDF) of nucleons and pions. The uncertainty estimates, if reliable, are also impressive. Lattice inputs to global PDF fits used for Large Hadron Collider (LHC) predictions are starting to have experimental impact (e.g., for strangeness asymmetry). This influence should grow as the reach of lattice predictions improves. However, in the short- and intermediate-term, it may be difficult to improve on the uncertainties for Higgs physics at the LHC as PDF uncertainties are already at the level of 1-2 percent in the relevant Bjorken x ranges.

Neutrino Physics:

First results on axial form factors of nucleons and the axial charge of He_3 have been obtained. Considerable work on g_A quenching has been undertaken. USQCD is aware that nucleon-level quantities are only one ingredient of many to make predictions relevant for neutrino experiments with respect to either neutrino nucleus scattering or neutrino-less double beta decay. Considerable effort has been demonstrated to connect to all the stakeholders and contributors towards that goal.

Strong coupling constant:

This is a major success of previous rounds of lattice calculations, with the lattice providing one of the most accurate determinations of α_S, as reflected in the FLAG report and the Particle Data Group (PDG). The PDG strong coupling determination averages the FLAG lattice result with the average of all other strong coupling determinations. It is notable that the two PDG components agree both for the central value and the level of uncertainty. However, a new method based on Wilson-flow is now being introduced (imported from lattice BSM studies by USQCD), and the first results find some tension with earlier LQCD results. Understanding this will be a focus of calculations over the next few years. With precision comparisons to LHC data increasingly reliant on higher order QCD calculations, the current uncertainty on the strong coupling constant is already one of the dominant contributions to cross section uncertainties at the LHC.

Hadron spectroscopy and resonance properties:

Calculational methods have advanced—based on both theoretical and algorithmic progress, as well as increases in computational resources—such that the first determinations of scattering phase shifts at the physical quark mass with controlled errors are now available (for pi-pi scattering). For other mesonic resonances (e.g. the f^0/sigma) there is considerable progress towards a first-principles calculation, although so far at heavier-than-physical quark masses. Methodology now exists for determining the couplings of resonances to electroweak probes (e.g., gamma+pi -> rho). Exotic states are accessible at large quark masses and give qualitative guidance to experimental searches. Formalism for studying resonances with three particle decays has been developed and pioneering applications have been undertaken. A comprehensive 5–10-year plan of development was presented, in which the contact with experimental programs will significantly increase.

Comments

The collaboration is to be commended for the clarity and organization of the presentations, and the comprehensive answers to the reviewer's questions. A number of top highlights were shown and an impressive set of activities at the forefront of research in lattice field theory were described. The impact in the heavy-quark sector, in the muon g-2, in the spectroscopy of exotics and in some aspects of nucleon structure is particularly high. The breadth of the program is very impressive.

BSM represents about 10 percent of the USQCD program but trains a disproportionate number of young scientists. If/when BSM physics is discovered, this workforce will be extremely useful. There is a positive synergy between BSM studies and the core work on LQCD, with studies of the former introducing new methods to apply to the latter.

In the nucleon sector, there will likely be a "consolidation period", where the same PDFs are recomputed with higher statistics, smaller lattice spacings, better control of excited-state effects, etc. In that period, the quoted errors might not shrink much, but instead become more robust.

Neutrino physics is a relatively new application of LQCD, and the first encouraging results have been obtained spearheaded by early career researchers. This is commendable. It is important going forward to ensure that this effort is broadly supported within the collaboration. Nuclear response and matrix elements of large nuclei is a very difficult problem and thus, the goal must be to create a community of researchers both within LQCD and the community at large which can bring together the necessary expertise to eventually provide experimentally relevant predictions. Within NP, topical collaborations have proven to be an invaluable framework allowing for the collaboration of experts across disciplines, e.g., in the case of matrix elements for neutrino-less double beta decay. It, therefore, seems plausible that topical collaborations also can play a useful and essential role for neutrino nucleus scattering. Without a supporting community the efforts by LQCD will remain insular and will not have the desired impact on the experimental program.

There is always a balance to strike between having a broad program and focusing on a smaller number of physics projects. Given the bottom-up approach of USQCD in allocating computing time, its tendency is naturally to err on the side of having lots of projects. In some cases, it might be worth enforcing stronger focus "from above", e.g., allocating resources only to the part anticipated to be the most impactful within an application. Given the oversubscription of USQCD-managed resources, this must be happening already to some extent.

Recommendations

DOE HEP and LQCD-ext III should seriously consider taking the lead in the creation of one or several topical collaborations focused on neutrino nucleus interactions with the goal to provide experimentally relevant predictions. The NP Topical Collaborations could serve as a model here.

3. The Status of the Technical Design and Proposed Technical Scope for FY2023-2024

Findings

FNAL has completed a technical design for the LQ2 system and has awarded a contract to a vendor (KOI). The LQ2 resource is scheduled to be online and available to USQCD in October 2023.

JLAB has collected responses to a Request for Information for a FY2023 purchase. Hardware benchmarking is planned for the summer of 2023 with equipment delivery anticipated in January 2024.

BNL is scheduled for the "pong" part of the schedule which will procure its next system in FY2024. Specific hardware details are not yet determined. USQCD is using a "ping-pong" method of procurements at FNAL and BNL, meaning that procurements alternate on an annual basis between the two labs.)

Comments

The technical designs are being done in close coordination with other common resources at the labs (networks, storage, cooling) even though the model is shifting towards dedicated compute resources. This is a good and necessary method of planning, and we encourage this to continue to ensure optimal designs for LQCD-specific resources going forward.

Releasing a fair and open Request for Proposal for computer procurements is a desirable aim to ensure the advantages obtained from competition. LQCD should carefully examine the computing landscape at any given time in the "ping-pong" model of procurement to evaluate the feasibility of fielding an appropriate amount of resources of any given type given price constraints. Specifically, this information should inform the effort expended on benchmarking against various designs where the likelihood exists that those platforms will prove to be prohibitively expensive.

Recommendations

None

4. The Merits of Including the Construction and Operation of Dedicated Hardware

Findings

The reviewers heard several examples where USQCD collaboration members were directly involved in development and construction of dedicated hardware with companies and the overall impact this had on QCD research.

For USQCD to have new computational resources at BNL in the near future, it will be necessary for the project to purchase dedicated cluster hardware, rather than use the previous IC model. This requires a change of policy from DOE-HEP, bringing it into agreement with the policy applied to NP funding.

Comments

Collaboration with industry partners for the development of future improved hardware should be strongly encouraged.

The required change of policy to buy dedicated hardware at BNL seems sensible and does not, in practice, lead to substantial changes. We encourage this change.

Recommendations

We recommend allowing the flexibility to buy dedicated hardware in appropriate circumstances. This mitigates the risk associated with IC updates.

5. The Feasibility and Completeness of the Proposed Budget and Schedule

Findings

The "ping pong" funding schedule has been implemented for FNAL and BNL to improve the effectiveness of funds by increasing the size of purchases resulting in an economy of scale. JLAB has adopted a similar technique by alternating compute and storage purchasing years.

The proposed budgets for the FY2025-2029 period are \$3M/year for HEP and \$1.5M/year for NP. These are increases of \$0.5M/year in both cases from the FY2023-2024 budgets.

The bi-annual computer purchases at each lab (FNAL, BNL, JLAB) will be (roughly) guided by Moore's Law.

Comments

The purchase model has worked effectively. The choice of whether to purchase computer processing unit-based or GPU-based machines is based on demand as gauged by the proposals.

The proposed use of funds for the remainder of the present proposals is reasonable and efficient.

We support the aims of the preliminary 5-year plan for NP-LQCD, since LQCD will play an ever-increasing role in NP, particularly with the upcoming Electron-Ion Collider.

Recommendations

None

6. The Effectiveness of the Proposed Management Structure, and Responsiveness to Recommendations

Findings

In response to a 2021 recommendation, a detailed timeline for results in the context of the HEP and NP experiments was created and is maintained on GitHub here: https://grokqcd.github.io/USQCD-theory-and-experimental-time-lines/

In response to a 2021 recommendation, a Diversity Equality and Inclusion (DE&I) survey was conducted, and results were summarized and made available to the USQCD collaboration.

In response to a 2021 recommendation, 5 new questions were added to the annual User Survey to probe the user experience with the SPC allocation process, fairness, and scientific impact.

The CDEI (Diversity Equality and Inclusion Committee) has developed an outline of how to deal with violations of the code of conduct, although, to date, no violations have been reported to the committee.

The site managers meet regularly with DOE program managers to provide frequent feedback and updates. The DOE program managers also maintain frequent communication with the EC and SPC of USQCD to ensure user's needs are being met.

A DE&I-related HEP training program for graduate students has been funded and is being set up.

Comments

The detailed timeline for results is an excellent addition to the documentation.

While, to date, the system set up by the CDEI to deal with complaints or issues has not been tested, several issues will need to be addressed when it is, and it would be good to prepare for this by further consultations with DE&I officers at the labs. One tricky issue in this regard is whether and how to notify home institutions in the case of sanctions. Another issue is the extent to which any sanctions are made known to the entire collaboration, i.e., the transparency of procedures.

The successful proposal for a training program is to be commended. We look forward to hearing more about the details of this program next year.

Recommendations

None

7. The Effectiveness of USQCD

Findings

USQCD manages and allocates time on resources at FNAL, BNL, and JLAB for the collaboration. These resources are augmented at the capability scale by additional resources at leadership computing facilities. The capacity resources managed by LQCD have represented between roughly 50 percent and 25 percent of the total time used by the collaboration over the past decade.

Allocations are made primarily by an SPC, which provides a recommendation for allocations to the EC. The SPC handles allocations of Type A and B (large and small), while development and testing allocations (Type C) are delegated to site managers.

Type A allocations are made each July based on available resources in the collaboration, with a 10 percent reserve held for Type B and C allocations to be considered on a rolling basis. For

FY2023-2024, this amounted to 384 Sky core-hours. Yearly requests are typically 2x what can be accommodated. However, proposals are rarely declined, the collaboration preferring to reduce allocation amounts for proposals, according to size and collaboration priorities.

Comments

The procedures for the allocation process are clear and communicated to all proposers.

The yearly allocation cycle seems appropriate to ensure that the scientific aims of the collaboration are pursued, and the resulting scientific impact is maximized.

The reliance on leadership computing facilities as a primary source of gauge generation is recognized as a source of possible risk. Estimates for available time based on historical allocations on Leadership Class Facilities (LCF) resources are used for planning but are not assumed. Planning includes LCF allocation fractions for LQCD as low as 4 percent of the total available hours per year. This is likely a reasonable lower limit.

Despite this reasonableness, LQCD might consider adding the possibility of a marked reduction in LCF allocations in any given allocation year to the formal risk register.

The smooth running of operations of USQCD is a very nice example of NP and HEP physicists sharing resources. It is very appropriate for things to continue to be run in this fashion given the many overlapping physics goals and common computational techniques.

The early career allocation program track evaluates proposals from early career researchers along with all other proposals, providing context and allowing scientific priorities to be considered straightforwardly. This is likely a best practice for these kinds of programs, and matches what is done, e.g., in the INCITE program.

As for the response rate to surveys, perhaps it would help to point out to USQCD members (raise awareness) that responding strengthens the case made to DOE for continued funding as part of this and future reviews.

Recommendations

None

APPENDIX A. Charge Letters

<u>Charge Letter to the LQCD-ext III Project Team</u>

Department of Energy
Office of Science
Washington, DC 20585

Dr. Josephine Fazio
LQCD Contractor Project Manager
Fermi National Accelerator Laboratory
Mail Station: 228 (WH 5W)
P.O. Box 500
Batavia, IL 60510

Dear Dr. Fazio:

The Department of Energy (DOE) Office of High Energy Physics (HEP) plans to conduct an Annual Progress Review of the Lattice Quantum Chromodynamics (LQCD-ext III) Computing Program on May 22-24, 2023, virtually, using ZOOM. This year's review will also include the LQCD-ext III cluster program at Thomas Jefferson National Accelerator Facility (TJNAF or JLAB) which is funded by the Office of Nuclear Physics (NP) since this program uses the same allocation process within a common collaboration and governance structure. A review panel of experts in high energy physics, nuclear physics, project management and computer science is being convened for this task.

John Kogut of HEP is responsible for this review; he will be assisted by Bill Kilgore, the Theory Program Manager of HEP.

Each panel member will evaluate background material on the LQCD-ext III research program and attend all the presentations at the May 22-24, 2023 review. The focus of the 2023 LQCD-ext III Annual Progress Review will be on understanding:

- The continued significance and relevance of the LQCD-ext project, with an emphasis on its impact on the experimental program of the DOE Offices of High Energy and Nuclear Physics.

- The progress toward scientific and technical milestones.

- The status of the technical design and proposed technical scope for FY2023-2024.
- The merits of updating the HEP LQCD-ext III research program to include the construction and operation of dedicated hardware in addition to the existing program of funding Institutional Clusters at Fermilab and Brookhaven National Laboratory.
- The feasibility and completeness of the proposed budget and schedule.
- The effectiveness of the proposed management structure, and the responsiveness to any recommendations from the last progress review.

Since LQCD-ext III provides computer cycles that are distributed by the US Lattice Quantum Chromodynamics (USQCD) collaboration, the panel members will also consider:

- The effectiveness of USQCD in allocating the LQCD-ext III resources to its community of lattice theorists, the scientific impact of this research on the entire HEP and NP communities and the status, operational procedures and related activities of the USQCD collaboration itself.

The 3 days of the review will consist of presentations and executive sessions. The third day will include an executive session and preliminary report writing; a brief close-out will conclude the review. Preliminary findings, comments, and recommendations will be presented at the close-out. You should work with John Kogut to generate an agenda which addresses the goals of the review.

Each panel member will be asked to review those aspects of the LQCD project listed above which are within their scope of expertise and contribute his/her findings to the closeout and final review reports. John Kogut, the Federal Project Manager, will be responsible for the final review report. That report will have recommendations for your consideration that you and USQCD should respond to in a timely fashion.

Please set up a web site for the review with relevant background information on LQCD-ext III, links to the various LQCD-ext III sites the collaboration has developed and distribute relevant background and research materials to the panel at least 2 weeks prior to the review. Please coordinate these efforts with John Kogut so that the needs of the review panel are met.

We greatly appreciate your willingness to assist us in this review. We look forward to a very informative and stimulating review.

Sincerely,

Regina Rameika
Associate Director of Science
for High Energy Physics

Charge Letter to the NP LQCD Initiative Project Team

Department of Energy
Office of Science
Washington, DC 20585

March 9, 2023

Dr. Josephine Fazio
LQCD Contractor Project Manager
Fermi National Accelerator Laboratory
Mail Station: 228 (WH 5W)
P.O. Box 500
Batavia, IL 60510

Dr. Robert Edwards,
Principle Investigator
NP LQCD Initiative
Thomas Jefferson National Accelerator Facility
12000 Jefferson Ave, Newport News, VA 23606

Dear Dr. Fazio and Dr. Edwards:

The Department of Energy (DOE) Office of Nuclear Physics (NP) plans to conduct a Progress

Review of Nuclear Physics Lattice QCD Computing Initiative (NP LQCD Initiative), in conjunction with the Office of High Energy Physics (HEP) Annual Progress Review of the Lattice Quantum Chromodynamics (LQCD-ext III) Computing Program, on May 22-24, 2023, using ZOOM. The NP LQCD Initiative funds the LQCD cluster program at Thomas Jefferson National Accelerator Facility (TJNAF or JLAB) as part of the LQCD cluster program. It follows the same allocation process within a common collaboration and governance structure. A review panel of experts in high energy physics, nuclear physics, project management and computer science is being convened for this task.

Paul Sorensen is the NP representative for this review, and will coordinate with John Kogut of HEP to have a joint LQCD-ext III and NP LQCD Initiative review.

Each panel member will evaluate background material on the LQCD research program, including those for NP LQCD Initiative, and attend the presentations at the May 22-24, 2023 review. The focus of the 2023 NP LQCD Initiative Progress Review will be on understanding:

- The continued significance and relevance of the NP LQCD Initiative, with an emphasis on its impact on the experimental program of the DOE Offices of High Energy and Nuclear Physics.

- The progress toward scientific and technical milestones.

- The status of the technical design and proposed technical scope for FY2023-2024.

- The merits of updating the NP LQCD Initiative research program, including the construction and operation of dedicated hardware.

- The feasibility and reasonableness of the preliminary FY25-29 5-year plan for NP LQCD Initiative.

- The effectiveness of the proposed management structure, and the responsiveness to any recommendations from the last progress review.

Since NP LQCD Initiative provides computer cycles that are distributed by the US Lattice Quantum Chromodynamics (USQCD) collaboration, the panel members will also consider:

- The effectiveness of USQCD in allocating the NP LQCD Initiative resources to its community of lattice theorists, the scientific impact of this research on the entire HEP and NP communities and the status, operational procedures and related activities of the USQCD collaboration itself.

The 3 days of the review will consist of presentations and executive sessions. The third day will include an executive session and preliminary report writing; a brief close-out will conclude the review. Preliminary findings, comments, and recommendations will be presented at the close-out.

Each panel member will be asked to review those aspects of the LQCD project listed above which are within their scope of expertise and contribute his/her findings to the closeout and final review reports. That report will have recommendations for your consideration that you and USQCD should respond to in a timely fashion.

Please set up a web site for the review with relevant background information on NP LQCD Initiative, and distribute relevant background and research materials to the panel at least 2 weeks prior to the review.

We greatly appreciate your willingness to assist us in this review. We look forward to a very informative and stimulating review.

Sincerely,

Timothy J. Hallman
Associate Director of Science for Nuclear Physics
Office of Science

APPENDIX B. Reviewers and Participants

Reviewers for LQCD-ext III/NP LQCD Initiative Annual Progress Reviewers 2023

Professor Abhay Deshpande
Department of Physics & Astronomy
Stony Brook University
C102 Physics Building
Stony Brook, NY 11794-3800
Tel. (631) 632-8109
abhay.deshpande@stonybrook.edu

Professor Patrick Huber
Director of the Center for Neutrino Physics
Virginia Polytechnic Institute and State
University
117 Robeson Hall, MC 0435
850 West Campus Drive
Blacksburg, VA 24061
Tel. (540) 231-8727
pahuber@vt.edu

Professor Joey Huston
Department of Physics & Astronomy
Michigan State University
3230 Biomedical & Physical Sciences
Building
East Lansing, MI 48824-1116
Tel. (517) 884-5564
huston@msu.edu

Dr. David Lawrence
Computational Science & Technology Division
Thomas Jefferson National Accelerator Facility
12000 Jefferson Avenue, Suite 4
Newport News, VA 23606-4468
Tel. (757) 269-5567
davidl@jlab.org

Dr. Bronson Messer
Director of Science and Acting Section Head
Oak Ridge National Laboratory
1 Bethel Valley Road
Oak Ridge, TN 37830
Tel. (865) 241-7340
bronson@ornl.gov

Professor Stephen R. Sharpe
Department of Physics
University of Washington
Physics-Astronomy Building
Box 351560
Seattle, WA 98195-1560
Tel. (206) 685-2395
srsharpe@uw.edu

University-Professor Harvey Meyer
Johannes Gutenberg-Universität Mainz
Johann-Joachim-Becher Weg 45
D-55099 Mainz
Office: Staudinger Weg 18, 3rd floor, room 03-105
Tel.: + 49 (0)6131 39-20350
Fax: + 49 (0)6131 39-25474
meyerh@uni-mainz.de

<u>Attending DOE program managers</u>

J. Kogut, LQCD-ext III HEP Federal Program Director
Glen Crawford, Director of Research and Technology Division of HEP
Bill Kilgore, Program Manager of Theory, HEP
Paul Sorensen, Program Manager for Fundamental Symmetries, NP

APPENDIX C. Agenda

Review Agenda

Day One – May 22nd - Monday

PDT	CDT	EDT	
08:00	10:00	11:00	Executive Session (45 min) – *John Kogut, Paul Sorensen & Committee*
08:45	10:45	11:45	Welcome and Logistics (5 min) – *Jo Fazio*
08:50	10:50	11:50	USQCD, NP goals, FY21 recommendations & responses- (15 min)– *Robert Edwards*
08:50	11:05	11:50	Long range computational plan for USQCD (15 min)– *Robert Edwards*
09:20	11:20	12:20	USQCD, HEP goals & governance (30 min) – *Tom Blum*
09:50	11:50	12:50	Break (10 min)
10:00	12:00	13:00	LQCD-ext III Management and Performance (30 min) – *Jo Fazio*
10:30	12:30	13:30	NPPLCI/Nuclear and Particle Physics (30 min) - *Amitoj Singh*
11 :00	13 :00	14 :00	Science 1: Light Quark and Lepton Physics (20 min) – *Luchang Jin*
11:20	13:20	14:20	Science 2: Beyond the Standard Model (20 min) – *George Fleming*
11:40	13:40	14:40	Science 3: Heavy Quark Flavor Physics (20 min)- *Will Jay*
12:00	14:00	15:00	Executive Session (30 min) – *John Kogut, Paul Sorensen & Committee*
12:30	14:30	15:30	Adjourn

Day Two – May 23rd - Tuesday

PDT	CDT	EDT	
08:00	10:00	11:00	Science 4: Hadron structure (20 min) - *Chris Monahan*
08 :20	10:20	11:20	Science 5: Nuclear structure (20 min) - *Phiala Shanahan*
08 :40	10:40	11 :40	Science 6: Spectroscopy (20 min) – *Raul Briceno*
09:00	11:00	12:00	Break (10 min)
09:10	11:10	12:10	USQCD User Survey & DEI Results (40 min)– *Jo Fazio & Will Detmold*
09:50	11:50	12:50	FNAL Plans & Deployment (15 min) *Jim Simone*
10:05	12:05	13:05	BNL Performance & Plans (15 min) – *Tony Wong*
10:20	12:20	13:20	Executive Session (60 min) – *John Kogut & Committee*
11:20	13:20	14:20	Committee Adjourns, *but …*

Committee Request for Information (60 min) – *John Kogut & Project Leadership*

12:20	14:20	15:20	Project Team Meeting (30 min) – *Entire Project Team*

Day Three – May 24th - Wednesday

PDT	CDT	EDT	
08:30	10:30	11:30	Response to Committee Questions & Discussion (90 min)
10:00	12:00	13:00	Break (10 min)
10:10	12:10	13:10	Executive Session & Preliminary Report Writing (120 min)
12:10	14:10	15:10	Closeout (60 min)
13:10	15:10	16:10	Adjourn