Offices of High Energy Physics
Report on the

LQCD-ext III

2020 Annual Progress Review

September 9-10, 2020

Executive Summary

The Annual Progress Review of the LQCD-ext III (Lattice Quantum Chromodynamics extension III) research program was held on September 9-10, 2020, via ZOOM. The purpose of the review was to assess LQCD-ext III's progress towards its overall scientific and technical goals, and to assess the role of the USQCD collaboration in governing the usage of the program's hardware. In particular, the LQCD-ext III team was instructed to address five charges:

1. the continued significance and relevance of the LQCD-ext III program, with an emphasis on its impact on the experimental program supported by the DOE Office of High Energy (HEP);
2. the progress toward scientific and technical milestones as presented in the LQCD-ext III's Execution Plan;
3. the status of the technical design and proposed technical scope for FY 2020-2021 for the program;
4. the feasibility and completeness of the proposed budget and schedule for the program;
5. the effectiveness with which the LQCD-ext III program has addressed the recommendations from last year's review.

The USQCD collaboration addressed the charge:

6. 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 community, 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 and USQCD. The impacts of LQCD-ext III simulations on experimental programs in precision measurements of the Standard Model (SM), Beyond the Standard Model (BSM) model building, hadronic matrix elements, and neutrino interactions were all evaluated. The impact of the LQCD research program has grown dramatically over the last several years. The USQCD's organization of the theory community to present a consensus for the muon g-2 measurement now being made by E896 at Fermilab was singled out and heavily praised. This development has been driven by algorithmic improvements and the use of new hardware platforms, including LQCD-ext III's mastery of Graphical Processing Units (GPUs). The governance of the program by the USQCD collaboration was judged to be effective and fair. The organization of the USQCD into an Executive Committee (EC) and a Science Policy Committee (SPC) was also praised. The leadership of Andreas Kronfeld as the USQCD spokesperson for the last three years with Robert Edwards as his deputy was commended. Previous review panels had suggested that USQCD consider electing more junior lattice gauge theorists into higher positions of the collaboration, and those requests have been largely met. The review team produced some suggestions for improving the LQCD-ext III program even further, but there were no formal recommendations.

Introduction and Background

The DOE Offices of Advanced Scientific Computing Research (ASCR), High Energy Physics (HEP) and Nuclear Physics (NP) have been involved with the National Lattice Quantum Chromodynamics Collaboration (USQCD) in hardware acquisition and software development since 2001. The Lattice Quantum Chromodynamics IT hardware acquisition and operations project ("LQCD"), which started in 2006 and ran through 2009, operated a "Quantum Chromodynamics-on-a-chip" (QCDOC) machine at BNL, and built and operated special-purpose commodity clusters at the Fermi National Accelerator Laboratory (FNAL) and the Thomas Jefferson National Accelerator Facility (TJNAF). The project's four-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 four-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 GPUs 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 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 five 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 five 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 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 Oct. 1, 2014.

The budget planning for the LQCD-ext II project was of some concern to the review panels of the 2014 and the 2015. The original five 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**	**FY 2018**	**FY 2019**	**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% 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 INCITE (Innovative and Novel Computational Impact on Theory and Experiment) 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 seven members, four experimentalists and three 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 (ECP), 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. There are typically 2-3 such workshops each year and they have been successful in engaging a wider audience for the lattice calculational program.

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, in order 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.

Brookhaven 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 (CSI), 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 Fermilab 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 LQCD-ext II project in 2018-2019, the project and USQCD put together a plan for the next five year period, 2020-2024. USQCD prepared seven whitepapers describing its five year plans in all of 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 seven 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 five 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 September 9-10, 2020 review was the first review of the new five year cycle. This review took place via ZOOM. The review consisted of one day of presentations and a second 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: App. A contains the charge letter to the LQCD-ext III management team, App. B lists the reviewers and DOE participants, and App. C contains the agenda and links to the talks.

Five 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 four presentations on the scientific topics which comprise lattice gauge theory. These are discussed in more detail below. Management talks followed and finally Andreas Kronfeld returned with a discussion of USQCD's plans for the future.

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

Continued Significance and Relevance

The LQCD-ext III program supports activities in four research areas:

1) QCD for Precision Particle Physics. Precision calculations which are relevant to the determination of standard model parameters extracted from heavy quark processes have been a major element in lattice calculations for many years. Calculations of decay constants and form factors which are essential for the extraction of CKM elements from experimental data and for looking for hints of new physics are continuing with ever increasing precision. Strong interaction matrix elements and scattering processes that are relevant to experiments at the Intensity Frontier, including the muon g-2 and the muon to electron conversion experiments at Fermilab, numerous kaon physics processes which are used to extract fundamental Standard Model parameters from various decay rates and scattering amplitudes, and low energy neutrino-nucleon cross-sections which are crucial to extracting results from neutrino oscillation experiments in progress at Fermilab, are new focus areas of lattice calculations. Thomas Blum of the University of Connecticut summarized this subfield of lattice gauge theory at the review. He emphasized the alignment of the lattice calculations with the growing set of experiments and projects in the near term Intensity Frontier program. He explained that recent algorithmic improvements in the muon g-2 calculational program have produced a sufficiently accurate lattice calculation that has allowed the theory community to publish a consensus Standard Model prediction (2006.04822[hep-ph]) of the muon g-2 before the experiment's publication scheduled for 2020-22. Recent improvements in the most difficult part of the calculation, light-by-light scattering, have been successful in reducing its error bars to a level comparable to that expected of the Fermilab experiment. In addition, several groups have improved the lattice calculation of the required vacuum polarization matrix elements to the required accuracy. Recent improvements in bottom and charmed meson decays, CP violation parameters in Kaon decays and rare Kaon decays were also reviewed. Significant improvements in all these calculations are expected over the next few years.

2) Beyond the Standard Model. Exploratory calculations based on "beyond the standard model" (BSM) theories, which in many cases are strongly coupled field theories, for which lattice gauge theory is at present the only effective technique for extracting quantitative predictions, constitute the main area of lattice calculations in this subfield. The emphasis has been on composite Higgs models, composite models of Dark Matter and lattice versions of Supersymmetry. Several of the most interesting models are “almost” conformal although they employ familiar gauge groups (SU2, SU3, SU4,…) but have many species of massless “quarks” in various representations of the gauge group. Calculations which accommodate the Higgs at 125 GeV/c^2 as a pseudo-Goldstone boson also predict additional states accessible to the LHC 14 TeV run. Investigative studies of supersymmetry are also underway. GPU clusters are proving particularly useful in these studies. George Fleming of Yale University summarized this subfield of lattice gauge theory at the review. He emphasized that this work is exploratory and only accounts for ~10 % of the total USQCD efforts. Over the last year there have been several publications in referred journals in this subfield. In addition, there are regular workshops annually where the lattice community interacts with theorists, phenomenologists and experimentalists in the field.

3) Nucleon Matrix Elements. Hadronic physics quantities such as the spectrum of hadrons, form factors, moments of structure functions, hadron-hadron interactions and scattering make up this subfield. Several of these calculational programs are well aligned with experiments planned for the 12 GeV upgrade of the Continuous Electron Beam Accelerator Facility (CEBAF) at TJNAF, including the spectroscopy of exotic mesons relevant to the GlueX project. The advent of peta-scale computing has led to calculations with physical pion masses so chiral extrapolations and the attendant uncertainties are no longer issues. Tanmoy Bhattacharya of Los Alamos National Laboratory summarized this subfield of lattice gauge theory at the review. Recent developments include: coupled channel resonance calculations, parton distribution functions, nucleon form factors, electric dipole moment calculations and fundamental symmetry breaking, nuclear double beta decay, and nuclear and hyper-nuclear forces. The accuracy and prospects for improvements in these calculations were also reviewed, Many of the simpler calculations are done at the physical pion mass, but others are still restricted to heavier (~300+ Mev/c^2) unphysical pions. The productivity of this lattice subfield has been strong in the last several years. The job market in the field has also shown considerable growth over the same time period.

4) Neutrino Physics. Neutrino physics is a central part of the HEP domestic experimental physics program. More accurate neutrino nucleon cross sections are needed to efficiently extract information from neutrino oscillation experiments. The lattice is engaged in elastic, transition and inclusive cross section calculations. There is also work on light nucleus calculations. Considerable progress was reported in nucleon form factor and transition form factor calculations over a wide range of momentum transfers and energies. Huey-Wen Lin of Michigan State University summarized this subfield of lattice gauge theory at the review. There are regular topical workshops on neutrino-nucleus calculations involving lattice physicists, experimentalists, phenomenologists and nuclear physicists.

The reviewers reported findings on these four scientific areas:

Findings

Flavor: The hadronic contributions to (g-2) of the muon were obtained using the USQCD resources. FNAL-HPQCD-MILC (2019) and RBC/UKQCD (2018) evaluated the hadronic vacuum polarization (HVP), while RBC performed the first computation of the hadronic light-by-light (HLbL) contribution. The first HLbL calculation was done on USQCD resources from 2015, with a more precise recent result using ALCF Mira in 2020 which was included in the recent Muon g-2 Theory Initiative average.

FNAL-MILC performed subpercent determinations of heavy quark decay constants (2018) and a percent level determination of semileptonic form factors (2019). Inclusion of QED corrections is being worked on.

There is a long-term effort in hadronic and semileptonic kaon decays. A milestone determination of epsilon' was achieved by RBC/UKQCD in 2020.

BSM: Several groups are exploring Beyond the Standard Model topics using lattice field theory calculations. These include composite Higgs, strongly-coupled (composite) dark matter, applications to gauge-gravity duality and the conformal bootstrap. These efforts are a small portion of USQCD resources (10%). Composite Higgs theories use SU(3) with many flavors (10 fundamentals or 2 sextets) that may approach the quasi-conformal regime. The recent result from the 10 flavor theory suggests the theory is near conformal. Dark matter explorations have focused on a particular model, SU(4) with four flavors. Recent results have shown that the four flavors need to be quite heavy in order to have a first order phase transition. When combined with other constraints, if dark matter were realized by this theory, with such heavy dark fermions, the first order phase transition would lead to gravitational wave signals at LISA. In the exploration of new lattice formulations there is an active effort to directly compare lattice calculations with the conformal bootstrap and gauge-gravity duality, where a recent result shows the correct spectrum of CFT scaling dimensions.

Neutrino physics: The most progress in neutrino scattering on the lattice is on the quasi-elastic form factors, certainly relevant to the US experimental program. Progress has been made on the axial vector coupling and axial-charge radius squared calculations and their associated errors. A projected milestone is to reduce the uncertainty in lattice calculation of the axial-charge radius-squared to 20%. Less progress has been made on hadronic transition form factors and structure function evaluations.

Global comparison: The lattice communities in Europe (including UK) and Japan are comparable in size to the USA lattice community, or even bigger. They are not organized in structures similar to USQCD, although there is a beginning of such an effort in Europe ("EuroLat"). Unlike USQCD these communities are not necessarily aligned with the experimental program. USQCD's existing structure and the alignment with the experimental program are important advantages. It allows, for example, for USQCD to address better and more quickly the needs of the experimental program.

Nucleon matrix elements: Calculations of the relevant matrix elements were presented that contribute to nEDM have made significant advances utilizing USQCD resources. Progress is continuing with multiple teams however the chiral continuum extrapolation is still in progress.

Comments

Flavor: The HVP and hadronic light by light contributions to (g-2)_mu is a crucial program, needed for the interpretation of the upcoming (g-2)_mu experimental result. The USQCD efforts are world leading. In view of the BMW estimate of HVP determination, the upcoming USQCD results will be crucial in order to understand the systematic errors in different determinations.

The outlook for g-2 is bright: sub-percent for HVP contribution and sub-10 percent for HLbL are now within reach. These efforts should be continued with high priority.

The heavy quark decay constants and form factors are a cornerstone input to the CKM standard unitarity triangle determination. In this decade two dedicated flavor experiments, Belle II and LHCb, are expected to improve the precision of flavor observables. They might also resolve some long-standing flavor physics tensions that might signal new physics. To match this experimental effort, it is both important and timely that similar progress is done on the theoretical side. A substantial part of the experimental program at LHCb and Belle II relies on the continued progress in lattice QCD determinations of these hadronic parameters. Lattice calculations are likely to play a big role in the coming decade in this effort.

The relatively precise determination of the epsilon’ parameter in kaon mixing is a seminal achievement, especially after the improved results from this year, where the de-contamination from higher excitations was achieved.

The continued progress on epsilon’ as well as the projected progress on epsilon_K and Delta mK, as well as the projected heavy quark improvements are highly desirable due to their impact on the experimental physics program.

Neutrino physics: Improving our understating of the neutrino-nucleus interaction is crucial for the USA neutrino program. Lattice QCD is expected to play a very important role in improving the theoretical handle on this interaction. For example, lattice QCD is arguably the only way we can access in the near future the axial form-factor which is important for the quasi-elastic region of interaction. Similarly, the resonance region of interaction, very relevant for DUNE, requires the knowledge of transition form factors. These are hard to access using other methods apart from lattice QCD, but it will require innovative approaches that may involve time scales of 5-10 years.

PDF calculations remain challenging, and it is not clear how the PDF calculations will impact the experimental neutrino program. The structure function approach may be a more fruitful direction since the low Q region is not directly probing quark substructure.

A discussion of the potential impact of the lattice PDF work on global-analysis fits to data would have been useful at this review.

The neutrino program depends on theoretical contributions that involve both nucleons and nuclei. Lattice contributions will play an important role, with the quasi-elastic scattering process with nucleons most clearly visible at present. Improvements of the nucleon component will impact assessments of nuclear effects.

The impact of lattice results would be stronger by maintaining and enhancing connections with the neutrino cross section community through Neutrino Scattering Theory Experiment Collaboration, the International Workshop on Neutrino Interactions workshops and other venues that bring experiment/theory and HEP/NP groups together. The Letters of Intent at Snowmass are a start. The project should develop a plan to increase its impact on this research area which is becoming a larger portion of the domestic high energy physics program.

BSM: USQCD presented a very interesting BSM program, especially given that only 10% of the LQCD-ext III program's resources are going towards it. The type A, B, C classification of projects seems to be effective for BSM. The recent result from the 10 flavor theory suggests that the theory is near conformal, and continued exploration of this seems warranted purely for the theory exploration of near conformal theories (even if this may not be an ideal composite Higgs theory).

The outlook for progress on BSM research is encouraging. Exploration of near-conformal theories is of continued interest, potentially to composite Higgs theories. Exploration of composite dark matter is novel, but rather specific to a particular model. The project should present these results with the widest applicability [e.g. applicable at large Nc]. Ongoing studies of theories beyond QCD maintain important connections between the theory program (formal theory) and lattice researchers.

<u>Nucleon matrix elements:</u> Calculations of the matrix elements for the topological and Weinberg terms are well in progress for nEDM, and the chiral continuum extrapolation needs continued effort. Preliminary calculations of quark CEDM were presented, though there is some discrepancy between the RBC/UKQCD and PNDME, that could be due to different pion masses. It would be helpful to understand the origin of the different results (scaling RBC/UKQCD by a factor of 0.25).

Recommendations

None.

Progress towards Scientific and Technical Milestones

Findings

USQCD determination of heavy quark decay constants are significantly better than projections, while the determinations of form factor are somewhat less precise than was anticipated in projections from 2013.

Comments

The latest results on heavy quark form factors are from early 2010s, a situation that is at least partially due to fluctuations in the availability of postdocs to carry out the calculations.

The searches for new physics via B meson mixing measurements is dominated by lattice errors for mass differences. Any improvements in the bag parameters by lattice would translate directly to better bounds on new physics (or a discovery of new physics effects).

Given the importance of heavy quark form factors for V_{ub} and V_{cb} determinations and the persistent puzzle of the difference between exclusive and inclusive measurements, it would be desirable for the project to improve the precision of lattice calculations of these quantities.

Devoting adequate resources to the determination of B meson bag parameters has considerable potential to improve the reach of new physics searches.

Recommendations

None.

Technical Design and Scope for FY2020/21

Findings

USQCD utilizes resources at three labs BNL, FNAL, and JLAB. The JLAB resources are funded from the NP program.

Both FNAL and BNL presented their acquisition status and expansion plans. There is a strong demand for both CPU and GPU based clusters with requests exceeding available resources.

Both sites follow a “best value” institutional cluster approach to acquisitions and have well defined procurement processes.

At FNAL the LQ1 cluster was installed and became available to users a short time afterwards. Soon after becoming available, the cluster was fully utilized.

The project showed that they have well defined metrics that allow comparison of allocation and usage. These metrics show that the project is meeting its goals for delivering high performance computing resources that are highly utilized.

Software is outside the scope of this project but there is a strong coupling with other projects that fund software development. Considerations of the software developers do seem to be part of hardware technology choices.

When compared with high volume experiments the long term storage needs for LQCD are modest. The model presented for long term data preservation is to rely on the labs.
BNL is in the process of constructing a new datacenter.

Comments

The two labs appear to have a well defined process that includes common benchmarks and sharing of technology decisions in a way that best serves the needs of the community. The existence of quality benchmarks and a technology evaluation process, as well as a set of metrics for measuring delivered performance, is an indicator of a mature and well run project.

The level of LQ1 cluster utilization so soon after installation demonstrates that the technology choice was appropriate for the community

In the LQCD community software stacks are highly optimized to produce the highest performance for a particular hardware architecture. With several different CPU and GPU chip manufacturers there is a risk of locking users to clusters where their code has the highest performance. Although a novel architecture may have higher performance it may not necessarily be a good fit for the broad community.

Long term data preservation technology choices are driven by future data use patterns. Infrequently, or sparsely, accessed data is much easier to deal with than data which is frequently accessed. Some LQCD data falls into the last category and can remain useful for decades. The plan of relying on the labs for data storage is sound but should be carefully monitored since, as LQCD data is a small percentage of a lab’s archive, LQCD is not driving the storage technology choices at the labs.

The project is to be commended on its performance on its technical scope. It should continue with its present methods and plans.

Recommendations

None.

Feasibility and Completeness of Budget and Schedule

Findings

The presentations demonstrated a record of providing resources on schedule and within budget.

The allocation and utilization charts showed a usage that was a close to, or exceeded, allocation.

Within a short time of the LQ1 becoming available it was fully utilized.

Comments

It is highly commendable that the project enables efficient use of resources so that they are, in most cases, fully (or even over) utilized. This is also concerning because it implies that there is a demand that exceeds what is delivered. This would seem to indicate that the project is constrained by budget.

With computational resources it is frequently the case that they are in high demand. It is often difficult to determine the scale of resource that is required. USQCD does have very good metrics on cluster use. The project should carefully determine if the resource levels that the project provides are adequate for the science scope. If this is not the case, then the project should consider a request for increased funding.

Recommendations

None.

Effectiveness of Management Structure and Responsiveness to past Recommendations

Findings

USQCD is a federation of lattice collaborations that include Fermilab Lattice, HotQCD, HPQCD, HadSpec, HLPC, LSD, MILC, NPLQCD, RBC, etc. that coordinates computing resources to optimize the science program for researchers (both HEP and NP at DOE). Approximately 160 people with 100 participating in physics proposals in a typical year.

Provides a platform for junior participants (postdocs and junior faculty) to try their ideas on USQCD clusters following two “Types” of proposals (A and B).

USQCD has an Executive Committee that balances interests of HEP / NP, labs, science collaborations, and general physics areas. The Scientific Program Committee balances scientific interests, supporting small collaborations and junior PIs. The Executive Committee consists of about 10 members, including a few junior members; the Scientific Program Committee consists of about 7 members. The SPC collects and reviews proposals.

Communications and reporting rely on regular site manager and program office meetings, and annual review and a all-hands meeting.

The project demonstrated a mature work planning and organization structure.

Risk management is performed and documented.

Performance indicators and milestones were shown that indicate that the project is functioning effectively.

Despite the current health crisis that severely impacted on site work FNAL staff were able to meet the installation timeline for the LQ1 expansion in March 2020.

Comments

The management structure of the project was appropriate for a multi lab project with local site managers at each lab.

The project management presentation was very thorough and demonstrated that this is a mature and well managed project. The management practices that have been carried over in the transition from dedicated machines to institutional computing resources look to be very effective to document the budgeting, monitoring and evaluation of the program and to keep the risks low.

The list of metrics and the way that they are utilized was very impressive.

FNAL, and BNL, are commended for their commitment to providing computational resources under very difficult circumstances in 2020.

The project management has proven to be responsive and effective. It should maintain its present methods and activities in the future.

Recommendations

None.

Effectiveness of USQCD, Scientific Impact, Procedures and Related Activities

Findings

For 2020-21, USQCD has allocated resources to approximately 30 projects in NP, several with HEP / NP overlap, Flavor physics / muon g-2, thermodynamics and BSM. The top 3 uses of the LQCD resource are NP (Meson Resonances and Couplings); HEP/NP (Parton Distribution Functions); and thermo (Chiral limit of 2+1-flavor QCD and the axial anomaly at high-T).

Members of USQCD have written three Letters of Interest in neutrino physics to the Snowmass 2021 process, in collaboration with nuclear theorists and neutrino physics phenomenologists.

Comments

There appears to be an excellent diversity of projects for 2020-21 with resources allocated broadly across the spectrum of topics. This suggests an excellent use of LQCD hardware to enable a large number of people to carry out early calculations.

It is very encouraging that the members of USQCD have started a concerted effort to collaborate with nuclear theorists in order to determine in which way lattice calculations can have the largest impact in reducing errors for neutrino cross section predictions or on bounding the nonstandard properties of neutrinos.

USQCD should continue its community-building efforts and strengthen those between lattice, nuclear physics, and neutrino experiment.

We encourage USQCD to work with the laboratory offices on Diversity, Equity and Inclusion & Science, Technology, Engineering and Math outreach activities.

Recommendations

None.

APPENDIX A

Charge Letter to the LQCD-ext III Project Team

Dr. W. Boroski

LQCD Contractor Project Manager

Fermi National Laboratory

Mail Station: 127 (WH 7W)

P.O. Box 500
Batavia, IL 60510-0500

Dear Dr. Boroski:

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 19-20, 2020, at the Brookhaven National Laboratory (BNL). A review panel of experts in high energy 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 19-20 review. The focus of the 2020 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 Office of High Energy Physics;

- The progress toward scientific and technical milestones;

- The status of the technical design and proposed technical scope for FY 2020-21;
- The feasibility and completeness of the proposed budget and schedule;
- The effectiveness of the proposed management structure, and responsiveness to any recommendations from last year's review.

Since LQCD-ext III provides computer cycles that are distributed by the 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 community and the status, operational procedures and related activities of the USQCD collaboration itself.

The two days of the review will consist of presentations and executive sessions. The second 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 write an individual report on his/her findings. These reports will be due at the DOE two weeks after completion of the review. John Kogut, the Federal Project Manager, will accumulate the reports and compose a final summary report based on the information in the letters. That report will have recommendations for your consideration that you and USQCD should respond to in a timely fashion.

Please designate a contact person at BNL for the review panel members to contact regarding any logistics questions. Word processing, internet connection and secretarial assistance should be made available during the review. You should 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 two 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 at BNL.

Sincerely,

James Siegrist

Associate Director

Office of High Energy Physics

APPENDIX B

Reviewers for LQCD-ext III Annual Progress Reviewers 2020

1. Jure Zupan (University of Cincinnati, Flavor Physics) jure.zupan@uc.edu
2. Graham Heyes, (associate of Amber, running SciComp at JLAB) heyes@jlab.org
3. Graham Kribs (University of Oregon, BSM) kribs@uoregon.edu
4. Gil Paz (Wayne State, Low Energy Hadronic and Flavor Physics) gilpaz@wayne.edu
5. Mary Hall Reno (University of Iowa, Neutrinos) mary-hall-reno@uiowa.edu

Attending DOE program managers

J. Kogut LQCD-ext III HEP Federal Program Director
Bill Kilgore, Theory Program Manager

APPENDIX C

Review Agenda

September 9th (Central Time)

08:30 Executive session (60 min) – *John Kogut*

09:30 Welcome and Logistics (5 min) – *Bill Boroski*

09:35 Scientific Goals of the USQCD Collaboration (25 min) – *Andreas Kronfeld*

10:00 USQCD Governance (20 min) – *Robert Edwards*

10:20 Break (10 min)

10:30 Science Talk 1: Quark and Lepton Flavor Physics (30 min) – *Thomas Blum*

11:00 Science Talk 2: Beyond the Standard Model (30 min) – *George Fleming*

11:30 Science Talk 3: Nucleon Matrix Elements for HEP and NP (25 min) – *Tanmoy Bhattacharya*

11:55 Lunch / Executive Session

1:00 Science Talk 4: Lattice QCD for Neutrino Physics (25 min) – *Huey-Wen Lin*

1:25 Response to Recommendations from 2019 LQCD Review (I) (15 min) – *Andreas Kronfeld*

1:40 Response to Recommendations from 2019 LQCD Review (II) (15 min) – *Jo Fazio*

1:55 LQCD-ext III Management and Performance (50 min) – *Bill Boroski*

2:45 Break (10 min)

2:55 FY19/20 Acquisition Status and FNAL Performance (30 min) – *Amitoj Singh*

3:25 FY20/21 FNAL Institutional Cluster Expansion Planning (20 min) – *Amitoj Singh*

3:45 BNL Institutional Cluster Facility Performance and Plans (20 min) – *Tony Wong*

4:05 Technical Design and Proposed Technical Scope for FY20-21 (20 min) – *Bill Boroski*

4:25 Executive Session (75 min)

5:40 Committee request for additional information – *John Kogut / Project Leadership*

6:00 Project Team Meeting (30 min) – *Andreas/Bill/Project Team*

September 10th (Central Time)

8:30 Response to committee questions and discussion (90 min)

10:00	Break (10 min)
10:10	Executive Session / Preliminary Report Writing
12:00	Lunch
1:00	Closeout
2:00	Adjourn