Offices of High Energy Physics and Nuclear Physics
Report on the

LQCD-ext II
2018 Annual Progress Review

May 21-22, 2018

Executive Summary

The Annual Progress Review of the LQCD-ext II (Lattice Quantum Chromodynamics extension II) project was held on May 21-22, 2018, at the Brookhaven National Laboratory (BNL). The purpose of the review was to assess LQCD-ext II's progress towards their overall scientific and technical goals, and to assess the role of the USQCD collaboration in governing the usage of the project's hardware. In particular, the LQCD-ext II team was instructed to address five charges:

1. the continued significance and relevance of the LQCD-ext II project, with an emphasis on its impact on the experimental programs supported by the DOE Offices of High Energy (HEP) and Nuclear Physics (NP);
2. the progress toward scientific and technical milestones as presented in the LQCD-ext II's Project Execution Plan;
3. the status of the technical design and proposed technical scope for FY 2018-2019 for the project;
4. the feasibility and completeness of the proposed budget and schedule for the project;
5. the effectiveness with which the LQCD-ext II project 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 II 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.

There were several special elements in this year's review. The scope of the project has changed to institutional clusters only, and will continue to be supported with HEP funding. HEP will be treating it as a research program rather than a project in the future. NP will support LQCD hardware at Thomas Jefferson National Accelerator Facility (TJNAF) off-project. The allocation process executed by USQCD for its members to use these facilities remains unchanged.

As a consequence, NP's interest in the review was limited to the hardware acquisitions and operations at TJNAF while HEP focused only on the operations of the institutional clusters at BNL and the plans for Institutional clusters at Fermilab. NP asked Chip Watson, the TJNAF contract program manager for the project, to present on the TJNAF operations over the last 12 months and HEP asked Tony Wong, a computer specialist at BNL, to present on the operations of the BNL Institutional Cluster. Another special element at the review was a request that USQCD update the status of their whitepapers presenting their plans for capacity computing over the next five year period, 2020-2024. HEP plans a full science review of the lattice program in summer 2019 where these plans will be evaluated by a team of theorists.

Four expert reviewers from the nuclear physics, high energy physics and computer science communities heard presentations on project management, computing hardware acquisitions and operations, organization of the USQCD collaboration, scientific progress, allocation of resources, and the dissemination of scientific results.

In general, the review panel was very impressed with the technical and scientific achievements of LQCD-ext II and USQCD. The impact of LQCD-ext II simulations on experimental programs in precision measurements of the Standard Model (SM), Heavy Ion collisions and hadron spectroscopy has grown dramatically over the last few years. These developments have been driven by algorithmic improvements and the use of new hardware platforms, including LQCD-ext II's early mastery of Graphical Processing Units (GPUs). The governance of the projects 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 election of Andreas Kronfeld to be the USQCD spokesperson for three years and have Robert Edwards as his deputy was praised. Previous review panels had suggested that USQCD consider electing more junior lattice gauge theorists into higher positions of the collaboration, and the election of Kronfeld and Edwards addressed these suggestions. Similar to past review panels, this year's reviewers did have several suggestions and recommendations that the project team and USCD should address. They are:

1. Each of the reviewers recommended that the remaining HEP FY2018 project funds ($0.85M) be released as soon as possible in light of the very positive impressions made at the review.

2. The project should work with Fermilab to initiate the development of a program of Institutional Clusters.

3. The project should develop procedures to document scientific milestones uniformly over all the LQCD areas so that the project can track their annual progress quantitatively and present it more thoroughly at each review.

4. Given the direct relevance of lattice gauge calculations to the experimental community, it would be valuable to enlist experimental physicists to advocate for the project during future reviews and in the next multi-year extension proposal due in 2019.

5. The project team should formulate a written plan to address the decreasing satisfaction articulated in the Compute Facility Satisfaction user survey results and present it to the DOE within two months.

Further guidance and additional details of these suggestions can be found in the body of the review report.

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 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 hardware project is now in its second extension, LQCD-ext II, which runs from FY2015-2019.

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 OHEP and ONP provided budget guidance to the project team of either $14M or $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 October 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 Annual Review. The original five year budget of $23.4M ($4.68M per year) proposed by the collaboration and endorsed by the November 8, 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 challenges 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 influence 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 places on the project and they noted that any additional funding would directly increase the project's hardware acquisition plans.

Over the course of the project and its extension, 2006-present, the hardware acquisition strategy of LQCD has been essential to its success. Each year the project's technical personnel benchmarks 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 is successful only if there is 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 has also organized and submitted a proposal to ASCR's Exascale program. That effort has been 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 error analyses which are fully defensible. 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.

The 2018 Annual Progress Review of LQCD-ext II took place at the BNL on May 21-22. 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 II management team, App. B lists the reviewers and DOE participants, and App. C contains the agenda and links to the talks.

There were several special elements in this year's review. As a consequence of the scope of the project focusing on Institutional Clusters, the sponsorship of the project has changed. HEP will continue to support the project focused on institutional clusters at BNL and FNAL. NP will support dedicated hardware at TJNAF, outside of the HEP project. The allocation process executed by USQCD for its members to use these facilities remains unchanged. NP's interest in the 2018 review was limited to the hardware acquisitions and operations at TJNAF while HEP focused only on the operations of the institutional clusters at BNL and the plans for Institutional clusters at Fermilab. NP asked Chip Watson, the TJNAF contract program manager for the project, to present on the TJNAF operations over the last 12 months and HEP asked Tony Wong, a computer specialist at BNL, to present on the operations of the BNL Institutional Cluster. Another special element at the review was a request that USQCD update the status of their whitepapers presenting their plans for capacity computing over the next five year period, 2020-2024. HEP plans a full science review of the lattice program in the summer of 2019 where these plans will be evaluated by a team of theorists.

Brookhaven indicated strong interest in supporting lattice gauge theory at the 2018 review. 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 review. BNL staff indicated their commitment to the Institutional Cluster funding model. The LQCD-ext II project team indicated that interactions with BNL have been both effective and flexible: the BNL staff fully endorses the project team's commitment to discover and procure the most cost effective hardware chips for lattice gauge theory simulations each year. It was unclear to the participants of the review if such a constructive attitude is present at Fermilab. Several meetings between the project team and the Fermilab computing division are scheduled over the next few months to resolve these issues and to formulate a plan. The 2018 review panel endorsed this action.

The review began with a presentation by Andreas Kronfeld, spokesperson for USQCD, which gave an overview of the USQCD collaboration and the LQCD-ext II 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 charge elements that the LQCD-ext II project team was asked to address in their charge letter.

Continued Significance and Relevance

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

1) QCD for 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 several 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. Christoph Lehner of BNL and lattice gauge theory's first Early Career awardee, 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 should produce a sufficiently accurate lattice calculation to improve the Standard Model theory prediction before the experiment's data analysis scheduled for 2018-19. Recent improvements in the most difficult part of the calculation, light-by-light scattering, have been very promising. In addition, several groups are improving the lattice calculation of the required vacuum polarization matrix elements. The lattice group hosted a theory workshop on the muon g-2 calculations on June 3-6, 2017 and has a organized a taskforce to optimize their activities on the muon g-2 calculations.

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. Ethan Neil summarized this subfield of lattice gauge theory at the review for the third consecutive year. He emphasized that this work is exploratory and only accounts for 5-10 % of the total USQCD efforts. Over the last year there have been seven publications in referred journals in this subfield. In addition, there was a workshop at Boston University in April, 2017, where the lattice community interacted with theorists, phenomenologists and experimentalists in the field.

3) Cold Nuclear Physics. Hadronic physics quantities such as the spectrum of hadrons, form factors, moments of structure functions, hadron-hadron interactions and scattering make up this subfield. In addition, 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. Other calculations focus on the program planned for the Facility for Rare Isotope Beams (FRIB). The advent of peta-scale computing will lead to calculations with physical pion masses so chiral extrapolations and the attendant uncertainties will no longer be relevant. This will lead to a new era in hadronic structure and spectroscopy calculations and make lattice simulations even more relevant to NP's experimental program. Will Detmold 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, Some of the simpler calculations are done at the physical pion mass, but others are restricted to heavier (~300+ Mev/c^2) unphysical pions. The productivity of this lattice subfield has been strong in the last 12 months with 6 Physical Review Letters. The job market in the field has also shown considerable recent improvements with five junior faculty appointments in 2017/2018..

4) QCD Thermodynamics. Calculations of the properties of QCD at finite temperature and baryon density, which is explored experimentally in relativistic heavy ion collisions, are critical to this subfield. These simulations are having an impact on the run plans of RHIC at BNL. The equation of state of the quark-gluon plasma is an essential input into the analysis of experimental data and the development of phenomenological models of final states. Recent calculations have focused on the critical temperature for the formation of the quark-gluon plasma, the critical point, the freeze-out lines, the velocity of sound and its temperature dependence, susceptibilities, and thermal di-leptons. Calculations of the Equation of State and the Transition Temperature are now considered "mature" and definitive. Several lattice calculations address questions posed in the NSAC Long Range Plan 2007. As lattice calculations become more accurate and ambitious, they are having an ever larger impact on the experimental NP program at RHIC and other worldwide facilities. Considerable progress in exploring the QCD phase diagram using charge fluctuations has been achieved since 2014, with an emphasis on the computation of freeze-out lines that should impact the next set of runs (Beam Energy Scan II) scheduled for RHIC. It is hoped that the next RHIC run will explore that portion of QCD parameter space in chemical potential and temperature where the critical endpoint of the first order transition line between the quark-gluon plasma and the cold hadronic phase exists. P. Petreczky summarized this subfield of lattice gauge theory at the review for the second consecutive year. Production calculations use Taylor expansions to address non-zero baryon chemical potentials. Progress on first principle algorithms remain exploratory. There was a topical workshop on lattice calculations and heavy-ion collisions at BNL in February, 2017.

The reviewers reported findings on these four scientific areas:

Findings

General Inclusive Findings

The USQCD scientific program covers a range of theory topics that are of specific interests to the experimental programs supported by the DOE Office of High Energy Physics.

The scientific community presented many examples of such research activities organized under four broad fields: Cold Nuclear Physics, QCD Thermodynamics, Particle Physics and Beyond the Standard Model investigations, the latter two of which specifically address experimental programs supported by the HEP Office.

LQCD calculations were illustrated to be important for advancing a broad range of investigations in HEP experimental programs such as the muon g-2 experiment, at the LHC, in neutrino experiments, and in searches for Dark Matter.

QCD for Particle Physics

This is a broad program which naturally overlaps and complements the Cold QCD and BSM programs. The lattice calculations are relevant for precision measurement experiments testing the Standard Model, in particular, muon g-2, B-physics etc. The calculations go hand in hand with experiments, enabling better understanding of QCD and searches for BSM signals.

The impact of LQCD in understanding QCD for HEP is strong. Of special interest is the rapid ramp up of calculations to improve the Standard Model prediction of the muon anomalous magnetic moment. The recent increased activity in this subfield was clearly a response to previous years' report recommendations. Recent advances in that effort suggest that LQCD will make a significant improvement to the theoretical uncertainty within the next year.

Beyond the Standard Model

Beyond the Standard Model research is a very interesting and versatile program. It is relevant for dark matter searches, or searches for hidden sectors as possible explanations of Higgs physics. The research in this direction is inherently speculative, but its importance goes beyond its direct applications as it improves our general understanding of strong coupling dynamics – the area where the lattice approach to quantum field theory has an indisputable edge.

Cold Nuclear Physics

Accomplishments of the LQCD-ext project in recent years have established not only new tools and methodology but are providing lattice calculations of hadron structure and spectroscopy resulting in new insights into the dynamics of the systems. These calculations provide a timely connection with flagship experimental programs in hadron spectroscopy such as GlueX at

Jefferson Lab. For example, new calculations of the meson spectrum include isospin content and flavor mixing, and states of high spin and exotic quantum numbers which suggest the presence of exotics in a regime accessible to GlueX. Beyond the calculations of hadronic spectra, applications of Luscher's finite volume method are producing new lattice calculations of the scattering dynamics of hadronic states.

Additional calculations include: first principle calculations of nucleon form factors and matrix elements, parton distributions and resonance spectroscopy which are relevant to a number of experimental programs at Jefferson Lab. There are also first principles nuclear force calculations which are relevant to structure calculations and will eventually be relevant to FRIB physics.

This program successfully trains new generations of scientists, as demonstrated by an impressive number (5) of junior faculty appointments in this area in 2017/2018.

<u>QCD Thermodynamics</u>
LQCD calculations of thermodynamic properties of Quantum Chromodynamics are providing essential inputs for theoretical models of the physics in the high temperature phase of QCD.

The Hot-Dense Lattice QCD Thermodynamics program is focused on the first-principles calculation of the lattice QCD equation of state at finite temperature and density which is a building block for all hydrodynamic evolution calculations needed to compare with and analyze experimental data from heavy-ion collision experiments. As such, QCD Thermodynamics is directly supporting the ongoing experimental program at RHIC. The close relationship with experiment is underscored by the fact that a Lattice Thermodynamics plenary talk is a necessary component of every Quark Matter meeting.

Comments

<u>General Inclusive Remarks</u>
The LQCD-ext project was presented as both significant and relevant to our understanding of and guidance for experimental research in those fields.

The speakers presented many studies of LQCD, ranging from clearly useful calculations of bound state phenomena to theoretical guidance on more speculative areas of research at the science frontiers.

There are good connections to the experimental community and the project for recognizes the value of enlisting input and project endorsement from their experimental colleagues.

The flagship calculation is the muon's anomalous magnetic moment. There is a 3 sigma discrepancy between the theoretical prediction and the value measured at Brookhaven, which has prompted a two pronged approach to confirming or refuting it. One is the Fermilab g-2 experiment (E-989), and the other is improving the Lattice QCD calculation of g-2 (particularly the hadronic vacuum polarization and hadronic light by light terms). This is an important and necessary step. The lattice community's Muon g-2 Theory Initiative is a good way to address this.

The impact of LQCD in general and USQCD in particular, extends beyond the leptonic flavor sector into the quark flavor sector. The speakers presented recent results in bottom and kaon physics, improvements in our understanding of direct CP violation, CKM elements and calculations of B decay anomalies.

The work being done in cold QCD provides input to the Jefferson Lab program, and in particular to GlueX. The lattice advances in scattering, meson properties, and decays are very important. The work in Hot QCD and the impact it is having on the Heavy Ion program is very impressive. In HEP, the breakthroughs that will lead to successful g-2 calculations are also remarkable. The work that has been done in heavy flavor physics and BSM is quite interesting.

Recommendations

None.

Progress towards Scientific and Technical Milestones

Findings

The project has been run within a multi-year plan to deploy and supply resources to the scientists of the USQCD collaboration, in which specific technical milestones (system uptime and availability) and to user satisfaction are listed in the appendices of the Project Execution Plan. The projected amount of CPU resources delivered was lowered to reflect the inability to execute the project procurement plan due to a budget hold in 2017.

Explicit milestones on resource availability (>=95%) were demonstrated to be met as CPU goals were consistently surpassed by actual resource delivery. Other goals, such as ticket resolution within 2 business days (>=95%), were showed to be largely met. An overall customer satisfaction (>=92%) was not met, with several measures, such as user documentation, user support, and staff responsiveness, having gone down significantly over the past year. Achievement of scientific milestones, other than resource delivery to the science community, was illustrated by the descriptions of progress made within the four broad fields of research which are being carried out by LQCD calculations.

Comments

The technical milestones of systems deployed, aggregate uptime, and cluster utilization were generally well presented and the project was shown to have met its goals.

The clusters appeared to be well managed and operated, with good overall utilization rates. It will be interesting to see whether the uptime and utilization remain high in the future Institutional Cluster (IC) era of the project.

Part of the review's charge was to determine progress towards scientific milestones. While each scientific area noted several impressive scientific achievements independently over the LQCD-ext II phase of the project, it would be useful to have a summary talk of scientific achievements that target scientific goals as presented in the second day of talks.

In formulating milestones, particularly when related to high-profile projects, it would be good to identify the resources needed to reach the goal. It would also be useful to know at what level of resources the results would no longer be valuable. This could be especially useful in times of tight or unknown resources as it provides solid arguments for specific levels of support.

It was clear that the "project culture" that the experimental community has learned to embrace is still somewhat new to the lattice theory community. The community feels one cannot "do science on a schedule" and is quite resistant to creating such a schedule. While it is difficult to manage scientific progress on a timeline, leaving the relationship between compute capacity and scientific progress and timelines vague makes it difficult for reviewers to fully appreciate their interconnectedness. That in turn makes it difficult to understand the effects of budget cuts in detail, and that then makes it difficult to defend these budgets. The USQCD collaboration would be improved if they firmed up their schedule, planning and timelines.

The technical milestones were crisp, and easy to understand. Progress is excellent, with only two substantial variances. One is interconnect delivery for the BNL Skylake cluster, which is unfortunate but with a single vendor (Mellanox) likely unavoidable, and the other is simply the delay caused by the release of only half the FY18 funding.

Recommendations

Each reviewer recommends that the second half of the FY18 funding be released to LQCD-ext II as soon as possible. This will allow the technical (and probably the scientific) progress to get back on track. Right now, the scientific progress appears to be on track because the JLab 16P system was more performant than anticipated, but this is not a reason to hold up funding.

The project should develop procedures to document scientific milestones uniformly over all the LQCD areas so that the project can track their annual progress quantitatively and present it more thoroughly at each review.

Technical Design and Scope for FY2018/19

Findings

The project chose to continue operations at BNL after its half-rack BlueGene-Q system was decommissioned at the end of FY2017. It evaluated the procurement of a BNL IC.
With funding frozen in 2017, those procurements were put on hold.

The evaluation of hardware needs of the project was revisited in early 2018, resulting in a 45%-30%-25% mix of SkyLake, KNL, and IC systems to be deployed at BNL. That evaluation was assisted with access to both BNL and JLAB resources. That plan was executed in February and continues with the ongoing SkyLake procurement.

Continued operation of project-owned resources at FNAL in FY2018-19 was assured but the addition of new resources within an IC model at FNAL is under discussion with FNAL management.

Comments

While there is a change in how the project is run at the three sites, past experience indicates that the project will make effective use of DOE funds, and new hardware procurement planning should be based on obtaining the remaining FY2018 funding from the DOE HEP Office.

In addition, the project's hardware evaluation process has been significantly helped by its multi-site operations.

The project would be best served if it continues to work on technical issues with the computing team at JLAB in support of the scientific mission of USQCD.

Two models were shown: JLab uses a dedicated compute cluster model, where the lattice community gets their own systems, and BNL has an institutional cluster model, where the lattice community gets a guaranteed amount of computing on shared systems. Both are working well. The institutional cluster model is intriguing, because it allows one to provision the systems for average use rather than the peak, reducing the cost of the system, or equivalently, allowing more resources to be obtained for the same amount of money. In this particular case, USQCD has been able to regularly take advantage of unused cycles, a benefit they do not have on dedicated hardware.

Andreas Kronfeld told the reviewers "The best institutional cluster model is one which is always on the cutting edge of hardware." Looking beyond 2019 and this review's charge, the reviewers were concerned with 2021 and beyond. Today the US LQCD community is using roughly equal amounts of computing on supercomputers (competitively awarded by ASCR under their INCITE and ALCC programs) and on resources they control through LQCD-ext II. Some of the latter resources are used for pre-and post- processing the work done via supercomputer. In 2021 the first exascale computer, Aurora (or A21) will come on line, providing 50x the computing power as Mira. This will induce a huge imbalance between what is available on supercomputers and what is available to pre-and post- processing the work done by these supercomputers. The community should prepare for this and be able to demonstrate that they can still do the science without needing fifty times as much dedicated computing.

Recommendations

Fermilab's plans are in flux as they transition to a new CIO. The reviewers recommend that Fermilab carefully examine the BNL institutional cluster model. The reviewers believe that Fermilab would discover that the advantages outweigh the disadvantages. It may even be beneficial for the Laboratories to coordinate: do both Labs need to have the exact same mix of single-CPU, multicore and GPU based computing?

Feasibility and Completeness of Budget and Schedule

Findings

The LQCD ext II project has been operated on a five-year funding schedule of $2M in FY15 and $3M in each of FY16 through FY19 including funds from both NP and HEP Offices and has currently received 65% of that 5-year budget.

The project has had to adjust its budget and schedule based on several factors including uncertainty in overall project funding, re-scope of deliverables from the NP and HEP offices, and an unresolved resource deployment model to be used at FNAL in the future.

The project noted the uncertainty regarding the development of an Institutional Cluster model at FNAL, but was optimistic with respect to the current FNAL management.

Comments

The project seems to have made reasonable adjustments to budget uncertainties and operational changes.

Current activities for resource deployment and operation at BNL seem stable and the project is actively working on resolving the model to be used at FNAL.

Recommendations

None.

Effectiveness of Management Structure and Responsiveness to past Recommendations

Findings

The management structure of LQCD-ext-II is composed of a project manager located at FNAL tasked with overseeing all project operations, under advisement from the USQCD Executive and Scientific Program Committees, and reporting to DOE program management.

The project manager is responsible for coordinating activities of the site-level manager and architects.

The presentation did not explicitly reference its responsiveness to last year's review.

Comments

The project management team is experienced and capable, and demonstrated a management structure that is both very detailed in its roles and responsibilities, and effective in its implementation.

Overall the reviewers were impressed by the management team, and commend them on the smooth transition of spokesperson, and the creation of a "chair line" for the spokesperson position, where the deputy becomes the new spokesperson after three years. The alternation of spokespeople from the HEP and NP communities ensures that both communities are represented in the management team, something especially important with the new funding model.

However, there is considerable risk in the management structure due to the lack of redundancy at the project management level after the recent loss of the project's assistant project manager. It is understandable that the project did not replace the assistant project manager given the uncertainty of project scope in terms of both long-term funding and a greater reliance on lab-based management within the IC model. However, this vulnerability should be addressed as soon as possible.

As mentioned above, it is clear that the "project culture" that the experimental community has learned to embrace is still somewhat new to the lattice theory community. For example, the project did not present the status of responses to last year's recommendations and suggestions until the review committee requested that information.

Recommendations

None.

Effectiveness of USQCD, Scientific Impact, Procedures and Related Activities

Findings

The review presentations from LQCD-ext II project described its strategy for resource procurement and deployment as being done to optimize the science impact of its computing resources for the USQCD Collaboration.

The collaboration, through its advisory activities to the project, helps the project benchmark potential hardware resources in support of the procurement process and determines the amount of available resources for allocation each year.

The mix of conventional and accelerated compute systems is one mechanism the project uses to optimize its resources relative to the computational work deemed critical to USQCD scientists.

The USQCD Scientific Program Committee (SPC) evaluates annual allocation requests based on their scientific merits and the allocations are awarded after approval of the Executive Committee and in consultation with the project site managers.

The allocation process includes a distribution of resources that is split roughly evenly between NP and HEP communities and more finely split based on relative impact on the current experimental activity in the US. The process is flexible to allow modest allocations to be awarded outside the normal annual process.

The project solicits an annual User Survey to assess user satisfaction with both the compute facilities and the allocation process. The feedback from the User Survey indicates a high user satisfaction with the allocation process. The Compute Facility Satisfaction has fallen from above 95% in 2014 to 88% in the most recent survey. This is below the KPI goal of 92%.

Comments

As stated above, the user survey indicates that the users' Compute Facility Satisfaction has fallen below the KPI goal of 92%. The proponents argued that this was because documentation needed to be improved and a lack of a common authentication procedure across the Labs. This may not be true: the details shown indicate that this decrease in satisfaction was in all six measured categories. The reviewers believe that the team should formulate a written plan to address the falling Compute Facility Satisfaction.

The Allocation Process Satisfaction is harder to interpret. It too is down: it's at 85%, close to its historical average but down from its peak of 91%. However, all four measured subcategories are up. There was an attempt to explain this as unhappiness with the total amount of computing available and users are not getting their full requests, but that doesn't quite mesh with the scores. Of all of the categories, the one that is most concerning is transparency. Ninety percent said the process was fair and transparent, which means 10% of the collaboration feels it is unfair and opaque. This is a large number, and is independent of the total number of cycles to be allocated. It is acceptable if a collaborator feels the outcome of the allocation process is *disappointing*, but it's another if the collaborator feels it is unfair. Perhaps the project team and USQCD should reach out to these people and discuss how the process could be improved.

One strength in the allocation process is that it is driven by the best possible science and not by preserving a fixed HEP/NP fraction. Another is that there are continuous opportunities for relatively small allocations. If someone gets a great idea, they don't have to wait up to a year to begin pursuing it.

The project is encouraged to continue taking annual user surveys.

Recommendations

The project team should formulate a written plan to address the decreasing satisfaction articulated in the Compute Facility Satisfaction user survey results and present it to the DOE within two months.

USQCD Plans Beyond FY2019

Findings

The project listed seven Whitepapers under preparation to elucidate its scientific goals over the coming years.

The Scientific Advisory Board (SAB), consisting of members of both the theory and experimental communities, are advising on the content of those Whitepapers.

It was noted that the writing of several of the Whitepapers were already underway and the time frame for completion was late Summer or early Fall of 2018.

Comments

As noted above, the project is commended for reaching out to the experimental community for endorsement of and guidance to the scientific goals of LQCD as illustrated by the make-up of the SAB.

The project should also be commended for including a Whitepaper that explicitly covers the challenges and achievements in computing from the LQCD perspective. Such a paper should address some of the pros and cons of running project-owned resources verses operating in a model, such as the IC one, where resources are shared with other scientific communities. Such a Whitepaper can be useful for the broader community of DOE supported science.

Recommendations

Given the direct relevance of lattice gauge calculations to the experimental community, it would be valuable to enlist experimental physicists to advocate for the project during future reviews and in the next multi-year extension proposal due in 2019.

APPENDIX A
Charge Letter to the LQCD-ext II Project Team

Dr. William 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 and the Office of Nuclear Physics plan to conduct an Annual Progress Review of the Lattice Quantum Chromodynamics (LQCD-ext II) Computing Program on May 21-22, 2018, at the Brookhaven National Laboratory (BNL). 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 the Office of High Energy Physics is responsible for this review; he will be assisted by Elizabeth Bartosz and Ted Barnes of the Office of Nuclear Physics.

Each panel member will evaluate background material on the LQCD-ext II project and attend all the presentations at the May 21-22 review. The focus of the 2018 LQCD-ext II 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 programs' support by the DOE Offices of High Energy Physics and Nuclear Physics;

- The progress toward scientific and technical milestones;

- The status of the technical design and proposed technical scope for FY 2018;

- 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 II 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 II 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.

We are also requesting USQCD present its plans for further capacity computing and USQCD should be prepared to answer the following question:

- What is the status of a Whitepaper presenting the research plan for FY2020-2025?

In addition, the Office of Nuclear Physics would like Chip Watson of the Thomas Jefferson National Accelerator Facility (TJNAF) to present the accomplishments and status of its LQCD-ext II facilities. Detailed instructions for this portion of the review will be communicated directly from the Office of Nuclear Physics to Chip Watson and his staff at TJNAF.

The two days of the review will consist of presentations and executive sessions. The later half of 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 II, links to the various LQCD-ext II sites the collaboration has developed, and distribute relevant background and project 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 of Science
for High Energy Physics

Timothy Hallman
Associate Director of the Office of Science
for Nuclear Physics

APPENDIX B

Reviewers for LQCD-ext II Annual Progress

Reviewers 2017

NP Participants:

1. R. Jefferson Porter

   NERSC & Nuclear Science Division

   Lawrence Berkeley National Laboratory

   rjporter@lbl.gov

2. Curtis A. Meyer MCS Associate Dean for Research

   Carnegie Mellon University

   Pittsburgh, PA 15213

   cmeyer@cmu.edu

HEP Participants:

3. Bill Kilgore, Theory Program Manager

   William.Kilgore@science.doe.gov

4. Tom Lecompte, ANL detailee Computational HEP

   thomas.lecompte@science.doe.gov

List of attending DOE program managers

J. Kogut (HEP, LQCD-ext II HEPProject Manager)

E. Bartosz (NP, LQCD-ext II NP Project Manager)

APPENDIX C

DOE Annual Progress Review of the Lattice Quantum Chromodynamics (LQCD) Computing Project May 21-22, 2018

Brookhaven National Laboratory
(Computational Science Initiative) Building 725, Meeting Room 2-124

May 21

08:30 Executive Session	(60 min)	
09:30 Welcome	(10 min)	**David Lissauer**
09:40 Logistics and Introductions	(5 min)	**Bill Boroski**
09:45 LQCD-ext II Overview & USQCD Governance	(25 min)	**Andreas Kronfeld**
10:10 LQCD-ext II Overview & USQCD Governance	(20 min)	**Robert Edwards**
10:30 Break	(20 min)	
10:50 Science Talk 1: Cold Nuclear Physics	(30 min)	**Will Detmold**

11:20	Science Talk 2: Heavy Ion Physics	(30 min)	**Peter Petreczky**
11:50	Science Talk 3: Beyond the Standard Model	(20 min)	**Ethan Neil**
12:10	Lunch / Executive Session	(50 min)	
13:00	Science Talk 4: QCD for HEP	(40 min)	**Christoph Lehner**
13:40	LQCD-ext II: Management and Performance	(60 min)	**Bill Boroski**
14:40	LQCD-ext II: FY17/18 Acquisition Plan & Status	(20 min)	**Bob Mawhinney**
15:00	Break	(20 min)	
15:20	LQCD-ext II: BNL Institutional Cluster Adoption	(20 min)	**Tony Wong**
15:40	LQCD-ext II: Accomplishments and Status of JLab Facilities	(20 min)	**Chip Watson**
16:00	Lattice QCD in FY19 and Beyond	(40 min)	**Andreas Kronfeld**
16:40	Executive Session	(60 min)	
17:40	Committee request for additional information		**John Kogut / Project Leadership**
18:00	Adjourn		
19:00	Dinner at Pure North Fork Craft Kitchen & Bar		

May 22

08:30	Response to committee questions and discussion	(90 min)
10:00	Break	(10 min)
10:10	Executive Session / Preliminary Report Writing	(110 min)
12:00	Lunch	(60 min)
13:00	Closeout	(60 min)
14:00	Adjourn	