Offices of High Energy Physics and Nuclear Physics
Report on the

LQCD-ext II
2017 Annual Progress Review

May 16-17, 2017

Executive Summary

The Annual Progress Review of the LQCD-ext II (Lattice Quantum Chromodynamics extension II) project was held on May 16-17, 2017 at the Fermi National Accelerator Laboratory (Fermilab). 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 and Nuclear Physics;
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 2017-2018 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 two special element in this year's review. The first was driven by the fact that Brookhaven National Laboratory (BNL) became a participant in this project in 2016, on par with the Thomas Jefferson National Accelerator Facility (TJNAF) and Fermilab. So, USQCD was asked to present the contributions that BNL has made to the project over the last twelve months. In particular, BNL is the first lab to contribute cycles on its Institutional Cluster to this project and there is considerable interest within OHEP and ONP in determining whether this computing model can address the project's needs more generally. P. Mackenzie, the chairperson of the executive committee of USQCD, was asked to address these questions in his presentation. The second special topic was USQCD's plans for the hardware project beyond the end of LQCD-ext II in FY2019. P. Mackenzie made a 45 minute presentation on this topic.

Six 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. Previous review panels have suggested that USQCD consider electing more youthful lattice gauge theorists into the higher positions of the collaboration in light of its aging demographics, and these recommendations are being enacted. Similar to past review panels, this year's reviewers did have several suggestions, some new and some advancing past recommendations, that the project team should address:

1. The project should consider shifting focus from specialized LQCD purchases to institutional based purchases, given the success of the BNL Institutional Cluster (IC) and budget pressures on the national NP and HEP programs. USQCD and the project should develop a plan to merge the LQCD purchase process, including requirements gathering, benchmarking, and acceptance, into the lab IC purchase process.

2. Given the growth of young researchers in the field, the collaboration should consider adding additional junior members to its executive and scientific program committees. The new directions for the project proposed in item 1 above suggest USQCD should consider the election of a new spokesperson and new personnel in its executive and science policy committees.

3. Data sharing (configurations) is part of the collaboration's charter. However, a data management plan was not presented at the review. The USQCD collaboration should develop such a plan and disseminate it at its All Hands meeting.

4. Since physics deliverables are the ultimate objective of the project, the definition and documentation of science milestones should be paramount. 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.

5. 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/or the next multi-year extension proposal past 2019.

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 Brookhaven National Laboratory (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 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 Annual Review. 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 produced the following 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 is 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 and 2016 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 which are pending.

The precision and relevance of the lattice community’s calculations have improved steadily over the years. 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 2017 Annual Progress Review of LQCD-ext II took place at Fermilab on May 16-17. 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.

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 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. Andreas Kronfeld summarized this subfield of lattice gauge theory at the review, for the third year running. 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 recent publication list of the Intensity Frontier lattice effort was not presented. The lattice group hosted a theory workshop on the muon g-2 calculations directly after this review, on June 3-6.

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 second 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 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. Many of these calculations are relevant to several NSAC Milestones. 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. Martin Savage 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 10 Physical Review Letters and 59 publications in referred journals in all. The job market in the field has also shown considerable recent improvements with five junior faculty appointments in 2017 and four in 2016.

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 dileptons. 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.

S. Mukherjee 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 topics that are of specific interests to the experimental programs supported by the DOE Offices of High Energy Physics and Nuclear Physics.

The LQCD calculations were shown to be very important for advancing a broad range of scientific investigations within the experimental programs supported by the NP office (at TJNAF, FRIB, RHIC/LHC and future EIC) and the HEP office (Muon g-2 experiment, experiments at the LHC, and in neutrino experiments.)

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 last year's 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.

The presentation at the review showed new developments on "conformal Higgs" theories.

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, the application of Luscher's finite volume method are producing new lattice calculations into the scattering dynamics of 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.

QCD Thermodynamics

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

General Inclusive Remarks

The LQCD-ext II program's continued significance, relevance, and impact on national and international experimental programs in nuclear and particle physics is beyond question.

The computational activities supported by the LQCD ext II project are both significant and relevant to our understanding of the experimental research in those fields supported by both the NP and HEP offices.

The mix of research that can be addressed with LQCD, ranging from the clearly useful calculations of bound state phenomena to system states in heavy-ion physics to guidance on more speculative areas of research at the science frontiers, is very impressive.

The relevance and impact of QCD Thermodynamics is particularly significant now since RHIC is gearing up for Beam Energy Scan Phase II. This warrants an increase of resource allocation and effort in this direction, at least in the short run, to ensure the success of RHIC BES-II.

A recent resurgence of interest and new results in the development of new methods dealing with the complex action/sign problem was noted. These methods are still not at a stage directly useful for lattice QCD. However, this high-risk/high-reward effort should receive a high priority.

The presentations at the review suggest that LQCD will make an important contribution to understanding and evaluating Fermilab's high precision muon g-2 measurements expected in 2018-19. The recent advances in the hadronic light-by-light calculation are very encouraging.

Given the direct relevance of these calculations to the experimental community, it might be valuable to enlist experimental physicists in advocating for the project during future reviews and/or the next multi-year extension proposal past 2019.

Recommendations

None.

Progress towards Scientific and Technical Milestones

Findings

Bill Boroski, the LQCD-ext contractor project manager, presented the management and performance information for the project. He presented 1. The project scope, organization, and budget, 2. Performance measures and metrics, 3. FY17 year-to-date performance results, 4. FY17 year-to-date financial results, and 5. The project summary.

The User Survey results were presented by the associate contract manager at Fermilab, Rob Kennedy.

The hardware selection strategy and acquisition plan for the FY2016-17 deployment at TJNAF was presented by Chip Watson, the technical contractor project manager at TJNAF.

The hardware selection strategy and acquisition plan for the FY2017 deployment at BNL was presented by R. Mawhinney, the co-site architect at BNL.

The project has received a total of $8M in funding in FY15, FY16 and FY17. The project's five year base budget is $14M.

A significant portion of the $8M (~55%) has gone into personnel. ~35% will be spent on hardware procurement in FY17.

The spending trend in FY17 is on-track.

The hardware performance achieved during the first three years of the project is exceeding the baseline goals, both on the conventional clusters and GP accelerated clusters.

The FY17 survey results indicated that the users were satisfied with the project's hardware choices and their operations. The scoring on the user survey exceeded expectations.

Operational milestones for facility up-time for conventional and accelerated-based systems were presented as 99.2% and 99.7% respectively, well above the target of 91.3%.

During the past year, the project procured and deployed new CPU hardware at TJNAF, done in a pair of purchases in late FY16 (procurement) and early FY17 (expansion of PO) to meet the overall project plan deployment schedule. Preparation for the procurement included an extensive evaluation of the cost vs performance of a number of architectures to meet a significant portion of LQCD processing tasks. From that evaluation, a homogenous Intel Phi Knights Landing cluster was selected for the TJNAF procurements.

Comments

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

Progress toward technical milestones of the LQCD ext II project was supported by the extensive hardware evaluation done for the TJNAF procurements, followed by a rapid deployment of the two-year plan during late FY16 and early FY17.

Overall, the clusters are well managed and operated, with good overall utilization rates.

Among the scientific milestones of the project are precise calculation of the masses of the lowest-lying meson and baryon states. Remarkable progress has been made with many calculations at or rapidly approaching the physical pion mass. These include LQCD calculations of physics beyond the Standard Model and the thermodynamics of strongly interacting matter under extreme conditions.

The percentage of active users responding to the survey has improved in FY17 compared to past years. This result satisfied a critical comment in the FY16 review report.

Physics deliverables are the ultimate objective of the project. The definition and documentation of science milestones seem to be quite good in some LQCD areas, and lacking in others. The project should develop procedures to document scientific milestones uniformly over all the LQCD areas so that the project can track their annual progress more quantitatively.

Recommendations

None.

Technical Design and Scope for FY2017

Findings

The project makes use of a procurement schedule that includes a rotation through the three sites while still allowing for lower-overhead purchases done in back-to-back years at one facility. The two year procurement at TJNAF finished in early 2017 with the next project procurement scheduled for late 2017 at BNL.

The project has chosen to continue operations at BNL after its half-rack BlueGeneQ system is decommissioned at the end of FY2017, in part because the project was given access to a BNL institutional cluster (IC), which has proved convenient and productive. The 2017 procurement will be for a new cluster operated at BNL.

Hardware selection is underway, drawing on the experience from the TJNAF purchases last year.

Chip Watson presented the technical performance of the FY17 cluster deployment at TJNAF and Bob Mawhinney did the same for BNL.

BNL and TJNAF are focusing on purchasing KNL machines for FY17/18.

KNL offers advantages over GPUs, especially in the ease of porting code to the architecture. There are issues, however, including performance degradation over time that requires periodic CPU resets to mitigate.

Comments

The project's decision to continue operations of at site at BNL seems well justified. In addition, there appears to be growing interest in Institutional Clusters.

Tracking the cost effectiveness of the BNL Institutional Cluster will be important in future planning.

It was noted during the presentations that sharing data (configurations) is part of the collaboration's charter. The strategy to achieve this was not presented. In fact, there was no discussion about a data management plan as part of the LQCD ext II project or the USQCD Collaboration. The USQCD collaboration should evaluate whether having a written data management plan within the project would serve to ensure that the integrity of the community's data is maintained.

The focus on KNL machines for new purchases at BNL and TJNAF seems appropriate given the installed base of GPU nodes. It appears that the project should consider shifting its focus from specialized LQCD purchases to institutional based purchases, especially given the success of the

BNL Institutional Cluster (IC) and constrained funding projections for both NP and HEP. In the past, LQCD was pushing the envelope on taking advantage of advanced architectures and so specialized machines were appropriate. Others, such as experiment HEP codes, were not prepared to use them. The rest of the world has now caught up, both in experimental HEP and NP codes (the Exascale era is pushing everyone in this direction) and deep learning training codes. Machines that were special for LQCD are now more generally useful. The challenge will be merging the LQCD purchase process including requirements gathering, benchmarking, and acceptance, into a lab IC purchase process. Managing risks will also be important. The labs, of course, must realize that LQCD will be a major user of such a facility and will need to give USQCD an influential seat at the acquisition table.

The FY17 acquisition plan presented at the review was very thorough, detailed and impressive.

Recommendations

None.

Feasibility and Completeness of Budget and Schedule

Findings

The LQCD ext II project is operated on a five-year funding schedule of $2M in FY15 and $3M in each of FY16 through FY19. This schedule allows the project to make multi-year plans for both operations and procurements needed to support its technical mission.

Although there are year-to-year variations, approximately 33% of the funds are expected to be used to purchase Compute and Storage systems, while about 55% are allocated for personnel efforts.

The project is three years into the five-year funding schedule and has delivered over 100% of their goals in terms of teraflop-years (TF-yrs).

The project's plan for new hardware at BNL (FY2017 and FY2018) and FNAL (FY2018 and FY2019) is designed to meets it overall goal for delivery of compute cycles to USQCD scientists.

Comments

The operations costs to the project have grown relative to its equipment costs over the life of the project. The project, however, has demonstrated that they have made a significant effort in itemizing and then analyzing those costs to confirm that they are justified. The project should be commended.

The recent new hardware installation and commissioning done at TJNAF illustrated the project's capability to meet its deployment schedule and the preparation for the new procurement at BNL appears to be following suit. In addition, the project's utilization of lower-overhead multi-year procurements at each site illustrates the attention to detail the project employs for managing cost savings to the benefit of its scientific productivity.

With regards to the feasibility and completeness of the proposed budget and schedule, the recent and planned acquisitions of KNL clusters is a move from specialized to more accessible hardware. Continuing along this track in the near term may provide additional opportunities to better utilize institutional cluster resources. This puts into question the continued need for additional GPUs. However, it seems reasonable to maintain a fraction of the total resources based on GPUs since no new code development is required to make substantial use of existing GPU based resources and since the next large increase in US computing capacity will be the GPU-accelerated Summit machine at ORNL.

Recommendations

None.

Effectiveness of Management Structure and Responsiveness to past Recommendations

Findings

The management structure of LQCD-ext-II is composed of a project (CPM) and assistant project (ACPM) 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 managers from the High Energy Physics and Nuclear Physics Offices.

The CPM and ACPM are responsible for monitoring the budgets, tracking milestones and coordinating activities at the independent sites. At least one Site Manager and one Site Architect are identified at each of the three sites where project compute and storage systems are deployed.

The Site Manager is primarily responsible for site operations and administration, while the Site Architect is responsible for evaluation and deployment of new hardware.

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. In addition, there is good overlap in responsibilities at all levels of management (CPM and ACPM at the project level and Site Manager and Site Architect at the site level) to ensure that the project can efficiently manage personnel transitions.

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 maximize its resources relative to the computational work deemed critical to USQCD scientists.

The USQCD Scientific Program Committee 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 the NP and HEP communities and more finely split based on relative impact on the current experimental activity in the US.

The project solicits an annual User Survey to assess user satisfaction with both the compute facilities and the allocation process. The FY2017 results were based on 73 responses and showed a 93% satisfaction with the compute facility and 85% satisfaction rate with the allocation process.

Comments

The procedures for making allocations appear clear to the scientific community and the project has instituted a reasonable non-use allocation reduction algorithm to help insure that a backlog of allocations which cannot be met does not build up.

The overall computing needs of the community were shown to be extensive and the project seems to be active in searching out ways to increase those resources, such as a more aggressive strategy for INCITE application and backfilling use of LCFs. These activities are encouraged.

The feedback from the User Survey indicates a high user satisfaction with the project and its allocation process. The project is encouraged to continue taking such surveys. One suggestion to improve feedback to the project is to hold a user-organized session during the annual All Hands Meeting to discuss user perspectives of the allocation process and the facility operations. Such a session, if actually user motivated and well attended, may be a good way to more clearly capture any common user pain points for using the facilities.

The management structure of the program is effective. The executive committee organizes the physic goals as outlined in the proposal and white papers and a separate scientific program committee is tasked to implement the program. The recent addition termed junior member to the executive committee provides a mechanism for greater representation of young researchers. Given the growth of young researchers in the field, the collaboration should consider adding additional junior members to its executive and scientific program committees.

If the project moves to Institutional Clusters as the main provider of cycles in its capacity computing model, then USQCD should consider the election of a new spokesperson and new personnel in its executive and science policy committees to reflect this new approach. The present management of USQCD appears to be wedded to the model of specialized, dedicated computers. The new influx of active junior physicists, especially through its growing and most productive element, Cold Nuclear Physics, should be given the chance to reinvigorate and refocus this project.

Recommendations

None.

USQCD Plans Beyond FY2019

Findings

A process for preparing the next proposal was presented - mainly producing science whitepapers and then a vision for the entire US community. The existing USQCD hardware portfolio will be taken into account when planning beyond FY19 purchases. Software support is planned to come from SciDAC and Exascale funding. The role of institutional clusters will be discussed. Running on the Leadership Class Facilities will also be discussed. It is noted that USQCD is very effective when using opportunistic computing at the LCF machines by bundling and backfilling.

Comments

USQCD will prepare a proposal for hardware purchases beyond FY19. USQCD should seriously consider institutional clusters. The project's rationale for purchasing their own hardware made more sense when they were first adopters of new architectures. This position is no longer true. An option for an FY19+ proposal could be to request funds to equip institutional clusters with features that may not be purchased otherwise, such as fast highly interconnected network systems. Such features would likely not harm non-LQCD users, but, as stated in the Future Plans presentation by P. Mackenzie, may be crucial for LQCD codes.

APPENDIX A

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

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 and the Office of Nuclear Physics plan to conduct an Annual Progress Review of the Lattice Quantum Chromodynamics (LQCD-ext II) Computing Project on May 16-17, 2017, at the Fermi National Accelerator Laboratory (FNAL). 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 16-17 review. The focus of the 2017 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 2017;

- 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 questions:

- Will USQCD be requesting a further extension of the IT Hardware project LQCD-ext II beyond FY2019?

- If so, what is the status of a Whitepaper presenting the research plan for FY2020-2025?

- If not, what are USQCD's plans for the ramp-down of LQCD-ext II?

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 FNAL 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 FNAL.

Sincerely,

James Siegrist
Associate Director
Office of High Energy Physics

Timothy Hallman
Associate Director
Office of Nuclear Physics

APPENDIX B

Reviewers for LQCD-ext II Annual Progress Reviewers 2017

1. HEP Theory

   Bill Marciano (BNL) marciano@bnl.gov marciano@quark.phy.bnl.gov

2. HEP Experiment/Computing

   David Asner (PNNL) David.Asner@pnnl.gov

3. HEP Experiment/Computing

   Adam Lyon (FNAL) lyon@fnal.gov

4. NP Theory

   Misha Stephanov (UIC) misha@uic.edu

5. NP Experiment

   Paul Eugenio (FSU) eugenio@fsu.edu

6. NP Computing

   Jeff Porter (LBNL, NERSC) RJPorter@lbl.gov

List of attending DOE program managers

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

T. Barnes (NP, Theory and Computation)

List of remote DOE program managers (ZOOM)

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

APPENDIX C

DOE Annual Progress Review of the
Lattice Quantum Chromodynamics (LQCD) Computing Project

LQCD-EXT II

May 16-17, 2017

Fermi National Accelerator Laboratory

Wilson Hall, Room: WH 1-East

Agenda

May 16

8:30 Executive Session (45 min)
9:15 Logistics and Introductions (5 min) – *Bill Boroski*
9:20 Welcome (15 min) – *Rob Roser, CIO, Fermi National Accelerator Laboratory*
9:35 Overview (55 min) – *Paul Mackenzie*
10:30 Break (15 min)
10:45 Science Talk 1: Cold Nuclear Physics (30 min) – *Martin Savage*
11:15 Science Talk 2: Heavy Ion Physics (30 min) – *Swagato Mukherjee*
11:45 Science Talk 3: Beyond the Standard Model (20 min) – *Ethan Neil*
12:05 Lunch / Executive Session

1:05 Science Talk 4: QCD for HEP (40 min) – *Andreas Kronfeld*
1:45 LQCD-Ext II: Project Management and Performance (45 min) – *Bill Boroski / Rob Kennedy*
2:30 LQCD-Ext II: TJNAF Acquisition Review: 16p (15 min) – *Chip Watson*
2:45 Coffee Break (15 min)
3:00 LQCD-Ext II: FY17 BNL Hardware Acquisition - Part 1 and Part 2 (30 min) – *Bob Mawhinney / Alex Zaytsev*

3:30 USQCD Beyond FY19 (45 min) – *Paul Mackenzie*
4:15 Executive Session (60 min)
5:15 Committee request for additional information – *Review Committee & Project Leadership*
6:00 Adjourn
6:30 Dinner: *20-30 min drive from Fermilab*

May 17

8:30 Executive Session (30 min)
9:00 Committee questions and discussion (60 min)
10:00 Break (15 min)
10:15 Executive Session / Preliminary Report Writing (105 min)
12:00 Lunch
1:00 Closeout (60 min)
2:00 Adjourn