Department of Energy
Office of Science
Washington, DC 20585

JUN 18 2014

Dr. W. Boroski
LQCD-ext Contract Project Manager
Fermi National Laboratory
P.O. Box 500 MS-369 (FCC 2W)
Batavia, IL 60510-0500

Dear Mr. Boroski:

We have enclosed a copy of the report resulting from the DOE review of the LQCD-ext 2014 Annual Progress Review that was held at the Fermi National Accelerator Laboratory on May 15-16, 2014. We very much appreciate the work that the LQCD-ext project team and USQCD collaboration invested in preparation for this review and in the presentations to the review committee.

The review committee was very favorably impressed by the review and its associated materials. They did, however, have several comments, and suggestions that the LQCD-ext project team, and the USQCD collaboration should consider and respond to. The details of their findings, comments and recommendations can be found in the attached report. Please address the review committee's comments and suggestions in a response to this office within the next three weeks.

We hope that the review report is helpful to you in completing the LQCD-ext project and preparing for its continuation, LQCD-ext II. Congratulations for continuing to achieve excellent results.

Sincerely,

Jim Siegrist
Associate Director of Science
for High Energy Physics

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

Offices of High Energy Physics and Nuclear Physics
Report on the

LQCD-ext
2014 Annual Progress Review

May 15-16, 2014

Executive Summary

The Annual Progress Review of the LQCD-ext (Lattice Quantum Chromodynamics extension) project was held on May 15-16, 2014 at the Fermi National Accelerator Laboratory (FNAL). The purpose of the review was to assess LQCD-ext'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 team was instructed to address five charges:

1. the continued significance and relevance of the LQCD-ext 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's Project Execution Plan;
3. the status of the technical design and proposed technical scope for FY 2012-2013 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 project has addressed the recommendations from last year's review.

The USQCD collaboration addressed the charge:

> The effectiveness of USQCD in allocating the LQCD-ext 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.

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 dissemination of scientific results. In general, the review panel was very impressed with the technical and scientific achievements of LQCD-ext and USQCD. The impact of LQCD-ext simulations on experimental programs in precision measurements of the Standard Model (SM), Physics Beyond the Standard Model (BSM), 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'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 review panel suggested that USQCD consider involving more mid-career lattice gauge theorists in the higher positions of the collaboration in light of its aging demographics. Finally, the review panel was impressed by the results and trends evident in the annual user survey.

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). LQCD met its 2009 project goal of providing 17.2 Teraflops of sustained computer power for lattice calculations.

The hardware acquisition strategy of LQCD was essential to its success. Each year the project's technical personnel benchmarked the kernels of the QCD code on the newest cluster and supercomputer hardware, and the winner of the price-to-performance competition became 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.

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.

USQCD proposed to extend the work of LQCD beyond 2009, and submitted the proposal "LQCD-ext Computational Resources for Lattice QCD: 2010-2014" to the DOE in the spring of 2008. The scientific content of the proposal was reviewed successfully on January 30, 2008 at the Germantown facility, and the scientific vision and specific goals of the project were enthusiastically endorsed in full by a panel of scientific experts. The proposal requested funding of $22.9M over a five year period to achieve the specified scientific goals.

In the January 30, 2008, review, USQCD argued that the purchase, construction and operation of mid-scale computer hardware was a critical component of the overall strategy to extract physics predictions from lattice Quantum Chromodynamics. That strategy depends on access to the largest Leadership Class machines for the generation of large lattice gauge configurations. These configurations are then analyzed for accurate predictions of matrix elements and spectroscopy on the mid-scale computers operated by LQCD, and results of interest to the experimental and theoretical communities in high energy and nuclear physics are obtained. These mid-scale LQCD computers are also used to generate smaller gauge configurations which are critical to studies of Quantum Chromodynamics in extreme environments (e.g. high temperature and density); these are used to ultimately interpret data from the heavy ion physics program at the Relativistic Heavy Ion Collider (RHIC) at BNL, which is operated by the Office of Nuclear Physics. Many of these calculations are not suited for Leadership Class machines, but run efficiently on mid-scale platforms. Several computer scientists at the January review carefully evaluated and then endorsed the mix of computers advocated by USQCD. The review panel also assessed USQCD's claim that the accuracy of some of its predictions rival the accuracy of the present generation of experiments now running at DOE HEP and NP facilities. The review panel also analyzed USQCD's claim that the proposed project, LQCD-ext, was needed to maintain this parity in the future.

The LQCD-ext project then entered the DOE Critical Decision review process. The CD-0 Mission Need Statement for LQCD-ext was approved on April 14, 2009.

The CD-1, Approve Alternative Selection and Cost Range, readiness review occurred at Germantown on April 20, 2009. The review evaluated the LQCD-ext project's documents on conceptual design, acquisition strategy, project execution plan, integrated project team, preliminary system document, cyber security plan, and quality assurance program.

The LQCD-ext team updated its documents following recommendations from the CD-1 review panel and received formal CD-1 approval on August 27, 2009, through a paper Energy Systems Acquisition Advisory Board (ESAAB) review.

The CD-2/3, Approve Performance Baseline/Start of Construction, readiness review occurred at Germantown on August 13-14, 2009. Final approval for the project was granted on October 28, 2009.

The Offices of High Energy Physics and Nuclear Physics developed the following planning budget for the LQCD-ext CD-2/3 review:

Table 1. Planning Budget for LQCD-ext (in millions of dollars)

	FY2010	FY2011	FY2012	FY 2013	FY 2014	Total
HEP	2.50	2.50	2.60	3.10	3.20	13.90
NP	0.50	0.75	1.00	1.00	1.00	4.25
Total	**3.00**	**3.25**	**3.60**	**4.10**	**4.20**	**18.15**

The TPC of $18.15 left the LQCD-ext project $4.75M short of the figure of $22.9M which had been supported by the scientific review of January 30, 2008, and which USQCD had also included in their original proposal. This shortfall was subsequently effectively addressed by the successful application by the Office of Nuclear Physics for $4.96M of funding through the American Recovery and Reinvestment Act of 2009 (ARRA) to build a 16 teraflop commodity cluster at TJNAF and operate it for four years. Although this effort was not a formal part of the LQCD-ext project, the resulting hardware at TJNAF was governed by USQCD using exactly the same procedures that applied to LQCD-ext, and the acquisition, deployment and operation of this hardware was tracked on a monthly basis by the same team that was running LQCD-ext. In this manner the Offices of High Energy Physics and Nuclear Physics monitored the full scope of the science effort put forward in the USQCD proposal "LQCD-ext Computational Resources for Lattice QCD: 2010-2014". It was agreed that the two efforts (LQCD-ext and LQCD/ARRA) would share Annual Progress Reviews.

The LQCD-ext project team argued at the CD-2/3 review that the budget in Table 1 would support the new deployments and operations described in Table 2 below:

Table 2: Performance of New System Deployments, and Integrated Performance

	FY 2010	FY 2011	FY 2012	FY 2013	FY 2014
Planned computing capacity of new Deployments, teraflops	11	12	24	44	57
Planned delivered Performance (TJNAF + FNAL + QCDOC), teraflop-years	18	22	34	52	90

The original computing goal for the LQCD/ARRA project was 16 teraflops (sustained), from a single cluster at TJNAF. The project team initially estimated that $3.2M would be used for hardware (to be operated for four years), and that labor costs for deployment, operations and management would be $1.2M, with incidental costs for disc space, spares, travel and misc. The project would require the addition of one position at TJNAF. Subsequently, a more quantitative and detailed cost breakdown was developed, which follows in Table 3:

Table 3: LQCD/ARRA Project Funding (in dollars)

Budget	FY09	FY10	FY11	FY12	FY13	Total
Steady State Operations	-	237,406	283,279	294,370	305,905	**1,120,960**
Hardware Deployment	1,929,280	1,817,423	-	-	-	**3,746,703**
Project Management	26,000	27,040	14,061	14,623	15,208	**96,932**
Total	**1,955,280**	**2,081,870**	**297,340**	**308,993**	**321,113**	**4,964,596**

Planning for hardware acquisition for these projects was, however, strongly affected in FY 2010-11 by a "disruptive technology" development in the field of PC chips. Although the first year of acquisitions had been planned assuming commodity cluster technologies, the development of Graphical Processor Units (GPUs) for the commercial gaming industry opened new opportunities for these projects. GPUs consist of several hundred cores per chip, and are the source of the high resolution interactive graphics capabilities needed for video game entertainment. Typically GPUs are capable of an order of magnitude more processing per second than general duty desktop CPUs. GPUs however are difficult to program at this time, and are unbalanced (too little memory per core) for general purpose applications. However, certain low-memory but computationally intensive, highly parallel algorithms can take advantage of a GPU's floating point capabilities, and can thus run 10-100 times faster on GPUs than on a CPU of comparable clock period. Lattice QCD calculations are dominated by one such algorithm; typically ~ 90%+ of the CPU time in lattice QCD is expended in inverting a sparse matrix, which is the Dirac operator that describes the dynamics of virtual quarks in QCD. Anticipating these developments, LQCD/SciDAC had been developing software for several years to run lattice algorithms on GPUs, and the fruits of that effort were then apparent in the GPU hardware acquired for LQCD/ARRA. Two complete physics projects ran on a GPU cluster at TJNAF during the GPU cluster's first year of availability. That number grew to ~ 9 projects in the second year, and continued to increase annually. The price performance on GPUs is ~$0.0l/megaflop which compares to $0.15-0.22/megaflop for the best CPU hardware. This development constituted an important new alternative in the hardware acquisition strategy of these projects, and was assessed in detail by previous review teams. These reviews have contributed several observations about this development:

1. The success of the hardware projects is very sensitive to the continuance of the LQCD/SciDAC software grant, because this is where the software will be developed that will eventually make GPUs more generally useful to the science community;

2. A mix of CPU and GPU clusters will be needed in the short term, because many lattice scientific applications are not ready to be ported to GPUs, but will be much more productive if and when that happens;

3. The initial estimates of the teraflop rating of clusters that can be built for $22.15M proved to be considerably higher than the original planning figures shown in Table 2, but it was hard to estimate new milestones at that time;

4. The scientific output and impact of these projects might be considerably higher than was initially assumed; and

5. The risks associated with using the new GPU hardware would exceed that of the more familiar CPUs.

All of these considerations became part of the discussions regarding plans for hardware acquisitions in FY 2010-12. Several of the observations and predictions quoted above have been confirmed: The ARRA GPU cluster is sustaining ~76 teraflops on a fairly diverse set of physics projects, which exceeds the project's original milestone by a factor of 76/16 ~ 4.75. The LQCD-ext project installed a GPU cluster in 2013 at FNAL to meet the extra demand coming from proposals submitted to USQCD over the previous 12 months. At the termination of the ARRA project in 2012 the ARRA hardware was incorporated into the LQCD-ext project, bringing its budget up to $23.115M, in line with the initial proposal from USQCD.

LQCD-ext is now operating a ½ rack BlueGene/Q supercomputer at Brookhaven National Laboratory (BNL) and has built and operated several commodity CPU clusters and GPU clusters at Fermi National Accelerator Laboratory (FNAL) and the Thomas Jefferson National Accelerator Facility (TJNAF). This aggregate of computers sustains ~165 TeraFlops for realistic physics codes. The project is approaching its completion and has exceeded its leading scientific and technical milestones.

The LQCD-ext project terminates at the end of fiscal 2014. An extension of the LQCD-ext project was proposed by USQCD and it is described in the proposal entitled "LQCD-ext II Computational Resources for Lattice QCD: 2015-2019" dated October 23. 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 of $23.4M. These plans have become 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 is now preparing for its CD-2/3 review on July 10.

The budget planning for the LQCD_ext II project was of some concern to the review panel of this 2014 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, if the project is funded at the $14M level, with the following 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**

as proposed by the Offices of High Energy Physics and Nuclear Physics in light of constrained, anticipated federal funding, then the estimated Teraflops profile is 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 will challenge USQCD to maintain its productivity, its balance with its Leadership Class computing allocations and its international standing. The review panel commented on these developments since they will influence the use and productivity of the FY2014 hardware acquisitions they were reviewing.

The Annual Progress Review of LQCD-ext took place at FNAL on May 15-16, 2014. 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 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 detail the findings, comments, and recommendations of the review committee for each of the six charge elements that the LQCD-ext project team was asked to address.

LQCD-ext Review

Continued Significance and Relevance

Findings

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

1) Intensity Frontier. 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. Ruth Van de Water summarized this subfield of lattice gauge theory at the review. She emphasized the alignment of the lattice calculations with the growing set of experiments and processes in the near term Intensity Frontier program.

2) Energy Frontier. 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 “walking” Technicolor models in which strong dynamics of new generations of quarks and gauge fields generate a composite Higgs which breaks electroweak symmetry. Unlike QCD, these theories are

"almost" conformal. Calculations which accommodate the Higgs at 125 GeV/c^2 as a pseudo-Goldstone boson and predict additional states accessible to the LHC 14 TeV run were presented. Investigative studies of supersymmetry are also underway. GPU clusters are proving particularly useful in these studies. Richard Brower summarized this subfield of lattice gauge theory at the review.

3) Hadonic Spectroscopy and Form Factors. 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.

4) Extreme Environments. 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 line, 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. Frithjof Karsch summarized this subfield of lattice gauge theory at the review.

USQCD's scientific goals are focused on carrying out world-leading computations of quantities that are of importance to the HEP and NP programs.

USQCD has presented a plan of calculations for the next five years which focuses on particular matrix elements and processes. The plan predicts the improvement of error bars in the lattice calculations over that period. In many cases the theoretical error bars are competitive with the experimental error bars.

USQCD produces White Papers on their various focus areas. These papers present short term and long term goals of the program. Tables of calculations with the expected errors of various physical quantities are regularly reported in these reports.

Lattice simulation is the only known way to accurately calculate equilibrium properties of hot QCD matter that is produced in the collisions at RHIC.

USQCD continues to have workshops with the experimental and theory communities to widen the impact of lattice simulations by actively engaging with complementary communities of researchers. The full slate of these workshops was included in the overview talk presented by Paul Mackenzie.

Comments

The review panelists remarked that the LQCD-ext project continues to be of the highest relevance to the nuclear and high energy physics communities. It tackles important calculations that require high performance scientific computing and its results are indispensable for the interpretation of measurements made by DOE-supported experimental programs and for determining these programs' future directions.

On the high intensity and high temperature frontiers, LQCD-ext calculations have enabled analyses and have led to understandings that would have been impossible otherwise. The ongoing extraction of the quark-gluon plasma viscosity (a transport coefficient that has received wide-spread interest also from other physics sub-communities) from collective flow data collected in RHIC and LHC experiments would be impossible without the (by now) very precise knowledge of the QCD equation of state gained by lattice simulations.

Our present knowledge of the degree of precision with which the Standard Model describes particle physics phenomenology could not have been achieved without the precise calculations of various transition matrix elements and form factors that were possible only with LQCD-ext.

The major goal of experimental studies in flavor physics is the measurement of new physics effects through quantum loops, which entails measurements of deviations from the Standard Model in a broad set of rare processes. On the experimental front, major leaps are expected in the precision of observables measured in current and future experiments. For these measurements to be fruitful in the new physics search endeavor, a similar leap is also needed on the theory front. For a large set observables, lattice simulations hold the key: precise and validated lattice QCD calculations of the SM predictions are absolutely critical to progress in the field. In view of the importance and the potential impact of these calculations, it is extremely important that the USQCD community continue, and further enhance, the priority level of their effort to validate the lattice QCD calculations against experimental measurements, where possible, and explain and communicate the basis of their estimated uncertainties to experimentalists. USQCD is encouraged to do more to help the outside physics community of experimentalists and phenomenologists better understand the level of robustness of some of its simulation results.

The impact of lattice QCD simulations on the experimental programs measuring the structure of atomic nuclei in or close to their ground states, and on high energy experiments at the Energy

Frontier, is still limited, since these developments are quite recent. However, these newer areas of lattice QCD applications are growing vigorously, with many ideas now having matured to a level where meaningful large scale simulations are being performed and have produced valuable semi-quantitative insights. Quantitative results, at physical values of the input parameters, are on the horizon. The US lattice community has propelled itself into a position of worldwide leadership in these novel directions while maintaining a strong pole position in the more mature research area at the Intensity Frontier and sharing the leadership with European and Japanese groups in high temperature QCD.

Without the LQCD-ext project, it is difficult to imagine how the US lattice community would have been able to move into such leadership positions, or how it would be able to defend them in the future. USQCD's access to hardware, both capability and capacity, will be essential for this favorable situation to continue.

Work on calculations of the hadronic contributions to g-2 should receive a high priority, as this will be crucial in trying to understand if the current discrepancy between theory and experiment is a signal for new physics. Continued work on nucleon properties for the muon→electron conversion experiment is also very important.

Work on cold nuclear structure, nucleon-nucleon interactions and spectroscopy is largely on track to meet its milestones and to impact the GlueX experiment. The areas of ground state properties (form factors, momentum distributions, spin structure, etc.) is experiencing the convergence of technological capabilities and physics interests. Lattice calculations will be particularly relevant to generalized parton distribution and transition form factors which will be measured in the 12 GeV program at TJNAF. Resurgent efforts to "solve" nuclear physics within the scope of lattice calculations are extremely exciting and important. Using lattice methods with effective field theories to compute the structure of medium and heavy weight nuclei promises to revolutionize the field of many-body nuclear physics.

Several members of the panel were impressed by the choice of the high priority calculations that USQCD is engaged in. They remarked that USQCD appears to be addressing the highest priorities of the HEP and NP experimental programs.

One reviewer remarked that "lattice QCD is in a Golden Age of productivity….Future progress should be strongly supported, and all efforts to gain full funding for the follow-up program should be encouraged…".

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. The presentation covered the project organization, scope, budget and the performance milestones and metrics, including the FY2013 and FY2014 performance and financial results, a summary of the user survey and the FY2014 hardware acquisition strategy.

The project met all of its major milestones and its user satisfaction scores improved markedly compared to the 2013 Annual Progress Review.

The hardware selection strategy and acquisition plan for the FY2014 deployment was presented by Don Holmgren, the technical contractor project manager at FNAL. The FY2014 equipment budget is $1.80M, with an additional $150,000 in carryover, and will be used to purchase a mix of conventional CPU hardware and accelerated GPU hardware. Vendor bids have been received and are under evaluation. The Release to Production date is expected to be September 15.

The FY2013 Continuing Resolution (CR) has caused some scheduling delays compared to the initial planned deployment schedule at FNAL. These issues should not have lasting impacts on the project.

The Infiniband CPU clusters and the GPU clusters are running with excellent availability at both FNAL and TJNAF. The CPU Infiniband clusters are typically available ~99% of the time and utilization is high. The GPU cluster performance has also been excellent and their utilization has been generally good with the cycles devoted to several large, demanding projects.

Comments

All members of the panel found the technical leadership of LQCD-ext to be very experienced and strong.

Progress toward the scientific and technical milestones is good. There were no significant slips relative to projections, beyond delays caused by the FY2013 CR.

The clusters, both CPU and GPU varieties, appear to operate effectively and efficiently. All the hardware platforms are well utilized, and idle time and down time are minimal.

A summary table of the physics milestones including projected results, actual results and future goals would be useful at each Annual Review. There are such tables buried in the various USQCD white papers, but it would be useful if the project could make them more accessible.

Recommendations

None.

Technical Design and Scope for FY2014

Findings

The project is currently considering eight vendor bids for conventional CPU and accelerated GPU hardware for deployment at FNAL.

Presentations were made by Robert Mawhinney on the FY2013 deployment and operations of a ½ rack BlueGene/Q at BNL and by Amitoj Singh on FY2013cluster deployment and operations at FNAL.

Don Holmgren presented the acquisition strategy and history of the project's procurements. The annual alternatives analysis for FY2014 was also presented. The deployments of all the systems and their subsequent operational statistics were reviewed.

Comments

The panelists commented that the project's just-in-time acquisition strategy and benchmarking protocols have been very successful in the past and should continue that track record.

Overall the technical planning is excellent.

Program management has done a good job of staying on top of technical developments in hardware and in modifying its hardware purchase plans accordingly, while taking into account the heterogeneous needs of the user community.

Recommendations

None.

Feasibility and Completeness of Budget and Schedule

Findings

The project has continued to meet its major budget milestones and schedule.

Comments

Schedule challenges have been encountered due to the federal budget cycle. The project has done a commendable job in working through these issues. This will remain an issue in the future and is being incorporated into the acquisition plans.

The reduced guidance provided by NP and HEP for the next five year cycle of the hardware project is a major concern. Even small improvements compared to the currently envisioned budget of LQCD-ext II would make a disproportionately large impact on the science results. The project should think creatively about how to increase the hardware fraction in its next five year cycle.

Recommendations

None.

Effectiveness of Management Structure and Responsiveness to past Recommendations

Findings

The recommendations from the review last year were considered and acted upon in several cases.

Results of the User Survey were presented by Robert Kennedy, the Associate Contractor Project Manager. The survey's overall satisfaction ratings exceeded the project's target goal of 92%. The year-to-year trends in the survey responses were favorable. The Allocation Process Satisfaction portion of the survey was considerably improved over the results of the previous year.

Following recommendations from past reviews, a Science Advisory Board (SAB) has been formed and has participated in the USQCD allocation process. 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.

Comments

The panelists commended the project team and USQCD on the management of the project's hardware.

The SPC is engaging younger (mid-career) members of the USQCD community in a positive, productive fashion.

The collaboration should continue to think about the best balance between sparing younger (but already permanent rank) scientists from bureaucratic responsibilities and ensuring the best long-term future of the US lattice QCD community.

The SAB is off to a good start and should insure that USQCD remains in touch with the interests of the broader physics community.

The results of the User Survey showed a high level of satisfaction of the lattice community with the current practices of USQCD and its allocation process in particular.

Recommendations

None.

Effectiveness of USQCD, Scientific Impact, Procedures and Related Activities

Findings

The USQCD collaboration consists of ~163 physicists with ~100 participating in physics proposals in 2014.

USQCD operates as a "collaboration of collaborations". It organizational structure consists of an Executive Committee (EC), a Science Policy Committee (SPC) and a more recently formed Science Advisory Board (SAB).

The Executive Committee (EC) of USQCD consists of 10 members, with one member rotating off per year. The main job of the EC is the preparation of proposals to the Leadership Class Supercomputer Centers and to the SciDAC program.

The Science Policy Committee (SPC) consists of seven members, with one or two rotating off per year. The main job of the SPC is the allocation of time to individual proposals on the Leadership Class machines as well as the capacity machines of the LQCD-ext project.

The Science Advisory Board (SAB) consists of seven members, four experimentalists and three theorists, who comment on and suggest revisions to USQCD's short term and long tern science goals. They also read and comment on specific proposals and allocations, monitor the effectiveness and fairness of the allocation process and participate in the annual all-hands meeting. An active SAB guarantees that the perspective of the wider physics community provides guidance to the project.

The physics goals of USQCD are set out in proposals and white papers which are organized by the EC in consultation with the SPC. In 2013 several white papers were completed with 23 authors.

The computing resources governed by USQCD are typically ~30% to ~50% oversubscribed by requests from the collaboration members.

USQCD was awarded 1183M hours on Leadership Class machines in 2014. The CPU clusters of LQCD-ext provided 411M hours and the GPU clusters provided 646M hours.

In 2012, USQCD won two SciDAC grants, each of ~$1M in size from NP and HEP separately. These SciDAC grants continue to provide support for the hardware project.

USQCD presented demographic information on the entire collaboration. Of particular interest was the fact that the job market for junior faculty lattice theorists at US universities has improved over the last several years, with four positions attained in 2013 and two more in 2014. This success addresses a concern of previous review panels.

USQCD continues to organize and run workshops with the experimental, phenomenological and theoretical physics communities in both NP and HEP. The latest workshops were in Beyond the Standard Model physics and QCD in Extreme Environments.

Comments

Several review panelists remarked that USQCD leadership has responded extensively and convincingly to questions raised by this and previous review committees about the effectiveness of the USQCD management structure and resource allocation procedures.

The review panel was impressed by the quality of the scientific program and its output and credited the SPC with doing an excellent job in deciding how to allocate resources according to a balanced mixture of proposal merit, demonstrated track record, and importance to the mission of the collaboration in supporting the experimental programs while at the same time encouraging the testing of new ideas through the class B (mid-sized) and class C (small, exploratory) proposals.

The SPC is balanced in age and across subfields; term limits of three to four years ensure that its membership rotates regularly. The EC rotates more slowly and has fewer younger members, as is natural for the task of steering the collaboration over the long term and in developing strategies for the future.

The engagement of the younger, mid-career persons on the SPC was recognized by the review panel as an important development which should be emphasized and continued.

It is important that the EC remains responsive to long-term changes in the field and its mission. The founding of the SAB with a membership of knowledgeable representatives of the experimental and theoretical communities was praised by the review panel. This action occurred following a recommendation of the 2013 annual progress review.

On the whole the system seems to be functioning well. The data presented at the review indicates that the computing resources are oversubscribed. The current procedure for allocating and monitoring the use of resources was judged to be sound. The user surveys show that the collaboration is generally satisfied with the outcome of the resource allocation and that the resource allocation has become more transparent and responsive, as recommended by previous review panels.

The SciDAC-III program is critical to the success of the project. The BG/Q and the GPU and MIC clusters must be useable by a large fraction of the LQCD-ext community for the project to continue increasing its impact on the HEP and NP programs. The SciDAC-III software packages will be essential for the project to meet its science milestones.

Some statistics were presented showing the investments in computing resources for lattice calculations made in the US, Europe and Japan. US computing resources are approximately 1/2 - 1/3 of the entire combined world resources but productivity of the US effort appears equal or stronger under various metrics (precision, number of publications, citations, number of talks). This result suggests that the US community is utilizing its resources in hardware, software and personnel, more effectively than its international competitors.

In answer to questions from previous review panels, it was noted that USQCD has produced ~70 Ph.D.'s over the last 10 years. The review panel considered this productivity impressive and suggested that the collaboration compile these statistics annually.

USQCD has a Charter which was negotiated with the DOE. However, it does not appear to have a constitution with formal bylaws. This structure could be useful for institutional memory, succession planning, and allocation procedures, especially if USQCD is challenged with diminishing resources in the future. USQCD and the DOE should address this short-coming.

The USQCD sponsored workshops have added to the impact and visibility of lattice calculations. USQCD's participation in the Snowmass and P5 process of the Office of High Energy Physics was thorough and successful.

The reduced budget proposed for the hardware project over the next five years will challenge USQCD's productivity and organization. Competition for resources by members of the collaboration will become more intense. The organization of USQCD may be severely tested. The review panel noted that the proposed $14M budget for the FY2015-19 period will result in no new hardware procurements in FY2015. This may lead to an imbalance in capability and capacity resources and a falloff in the collaboration's productivity.

Recommendations

None.

APPENDIX A

Charge Letter to the LQCD-ext Project Team

Dr. W. Boroski
LQCD Contractor Project Manager
Fermi National Laboratory
Mail Station: 127 (WH 7W)
P.O. Box 500
Batavia, IL 60510-0500

Dear Dr. Boroski:

The Department of Energy (DOE) Office of High Energy Physics and the Office of Nuclear Physics plan to conduct an Annual Progress Review of the Lattice Quantum Chromodynamics (LQCD-ext) Computing Project on May 15-16, 2014, 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 Kawtar Hafidi and Ted Barnes of the Office of Nuclear Physics.

Each panel member will evaluate background material on the LQCD-ext project and attend all the presentations at the May 15-16 review. The focus of the 2014 LQCD-ext 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 as presented in the project's IT Exhibit 300;

- The status of the technical design and proposed technical scope for FY 2014;

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

The 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, links to the various LQCD-ext 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 2014 LQCD-ext Annual Review (FNAL May 15-16)

Computer specialist:
Amber Boehnlein (SLAC)
amber@slac.stanford.edu

HEP theorist:
Zoltan Ligeti (LBNL)
zligeti@lbl.gov

NP theorist:
Steve Godfrey (Carleton U., Ottawa)
Stephen.Godfrey@carleton.ca

NP experimentalist:
Reinhard Schumacher (CMU)
schumacher@ernest.phys.cmu.edu

List of attending DOE program managers

J. Kogut (HEP, LQCD-ext Federal Project Director)

T. Barnes (NP)

K. Hafidi (NP, LQCD-ext NP Project Manager)

APPENDIX C

Review Agenda

DOE Annual Progress Review of the
Lattice Quantum Chromodynamics (LQCD) Computing Project
LQCD-Ext
May 15-16th, 2014

Fermi National Accelerator Laboratory
Location: Wilson Hall (Room: Comitium, WH2SE)

Agenda

May 15
08:30 Executive session (45 min)
09:15 Welcome (15 min) – *Nigel Lockyer*
09:30 Logistics and Introductions (5 min) – *Bill Boroski*
09:35 LQCD-ext Overview & USQCD Governance (50 min) - *Paul Mackenzie*
10:25 Break (15 min)
10:40 USQCD Allocation Process (20 min) – *Robert Edwards*
11:00 Science Talk 1: Lattice QCD for HEP: Standard-Model Parameters and Matrix Elements (40 min) – *Ruth Van de Water*
11:40 Science Talk 2: Lattice Field Theory Beyond the Standard Model (20 min) – *Rich Brower*
12:00 Lunch / Executive Session
1:00 Science Talk 3: High Temperature/Density QCD (30 min) – *Frithjof Karsch*
1:30 Science Talk 4: Hadron Spectroscopy, Structure and Interactions (30 min) – *Martin Savage*
2:00 LQCD-Ext Project: Management and Performance (60 min) – *Bill Boroski*
3:00 Coffee Break (20 min)
3:20 LQCD-Ext: Technical Performance of FY2013 BG/Q Deployment (20 min) - *Bob*

Mawhinney
3:40 LQCD-Ext: Technical Performance of FY2013 Cluster Deployment (20 min) – *Amitoj Singh*
4:00 LQCD-Ext: FY2014 Hardware Plan & Cluster Deployment Status (30 min) – *Don Holmgren*
4:30 LQCD-Ext: Future Planning - FY2015 and Beyond (15 min) – *Bill Boroski*
4:45 Executive Session (60 min)
5:45 Committee request for additional information - *Committee/Project Leadership*
6:00 Adjourn
7:00 Dinner

May 16
8:30 Executive Session (30 min)
9:00 Committee questions and discussion (60 min)
10:00 Break (10 min)
10:10 Executive Session / Preliminary Report Writing
12:00 Lunch
1:00 Closeout
2:00 Adjourn