Department of Energy
Washington, DC 20585

JUL 25 2013

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

Dear Dr. Boroski:

We have enclosed a copy of the report resulting from the Department of Energy review of the Lattice Quantum Chromodynamics (LQCD-ext)/American Recovery and Reinvestment Act (ARRA) 2013 Annual Progress Review that was held at the Thomas Jefferson National Accelerator Facility on May 9-10, 2013. We very much appreciate the work that the LQCD-ext/ARRA project team and National Lattice Quantum Chromodynamics Collaboration (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/ARRA project team, the ARRA project team, and the USQCD collaboration should consider and respond to. The details of their findings, comments, and recommendations can be found in the enclosed 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 continuing the LQCD-ext/ARRA project. Congratulations for continuing to achieve excellent results.

Sincerely,

James Siegrist
Associate Director of Science
for High Energy Physics

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

Enclosure

Offices of High Energy Physics and Nuclear Physics
Report on the

LQCD-ext/ARRA
2013 Annual Progress Review

May 9-10, 2013

Executive Summary

The Annual Progress Review of the LQCD-ext (LQCD extension) and the LQCD ARRA (American Recovery and Reinvestment Act) projects was held on May 9-10, 2013 at the Thomas Jefferson National Accelerator Facility. The purpose of the review was to assess LQCD-ext/ARRA's progress towards their overall scientific and technical goals, to close out the LQCD ARRA project, which was completed in 2012 and has been merged with LQCD-ext into LQCD-ext/ARRA, and to assess the role of the USQCD collaboration in governing the usage of the hardware. In particular, the LQCD-ext/ARRA team was instructed to address five charges:

1. the continued significance and relevance of the LQCD-ext/ARRA projects, 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/ARRA's Project Execution Plans;
3. the status of the technical design and proposed technical scope for FY 2012-2013 for both projects;
4. the feasibility and completeness of the proposed budget and schedule for each project;
5. the effectiveness with which the LQCD-ext/ARRA projects have addressed the recommendations from last year's review.

The LQCD ARRA project addressed the charge:

> The status of the LQCD-ARRA project in achieving project goals and performance specifications.

The USQCD collaboration addressed the charge:

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

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 dissemination of scientific results. In general, the review panel was very impressed with the technical and scientific achievements of LQCD-ext/ARRA and USQCD. The impact of LQCD-ext/ARRA 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/ARRA's early mastery of Graphical Processing Units (GPUs). The performance of the LQCD-ext/ARRA project was given high grades and the project's procedure for acquiring the most cost effective hardware annually was endorsed. The review panel commented that the LQCD ARRA project had met or exceeded its milestones in all cases and should be commended for its early adoption of GPUs, which allowed the project to exceed its teraflop performance goal by over a factor of five. 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. However, the review panel suggested that USQCD consider implementing a mechanism to promote a regular turnover of its EC members through a democratic process which would involve the entire collaboration. Several reviewers encouraged USQCD to elect one or two members of the EC by popular vote at the annual all hands meeting. Finally, the review panel was impressed by the results and trends evident in the annual user survey, although some of the scores were based on limited statistics and were inconclusive. The review panel encouraged USQCD to make its allocation policies as transparent as possible, and to share negative reviews and comments and discuss the issues involved with the relevant PIs. In general the reviewers were positively impressed by all aspects of the program discussed in the review, and did not want to "micromanage a program that they believe is working well."

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 Board and Scientific Program Committee (allocations board). 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 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. The resources are awarded based on a merit system. Three classes of applications for computer time allocations 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.

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 this 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; this document is the third of these reports.

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 has been developing software for several years to run lattice algorithms on GPUs, and the fruits of that effort are now 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 is expected to continue to increase. 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 will probably 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 is now installing a GPU cluster at FNAL to meet the extra demand coming from proposals submitted to USQCD over the past 12 months.

The Annual Progress Review of LQCD-ext and LQCD/ARRA took place at TJNAF on May 9-10, 2013. 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/ARRA 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 five charge elements that the LQCD-ext/ARRA project team was asked to address, the charge detailing the closeout of the ARRA project, and the charge given to USQCD.

LQCD-ext/ARRA Review

Continued Significance and Relevance

Findings

The LQCD-ext/ARRA 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, 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 "simple" 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. Julius Kuti 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 aimed at quantities which will be studied at 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). This program is also attuned to the NSAC milestones. 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. 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.

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.

Comments

The review panelists remarked that the LQCD-ext/ARRA 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/ARRA calculations have enabled analyses and led to understanding 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/ARRA.

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.

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/ARRA 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.

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 recent resolution of the long-standing discrepancy between lattice and experiment for g_A is a significant advance, as is the identification of the rho resonance in phase shifts (although at an unphysical pion mass).

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.

Recommendations

None.

Progress towards Scientific and Technical Milestones

Findings

Bill Boroski, the LQCD-ext/ARRA 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 FY2012 and FY2013 performance and financial results, a summary of the user survey and the FY2014 hardware acquisition strategy.

Two changes in the Project's organization include: 1. Kawtar Hafidi is now the NP Project Manager, and 2. Robert Kennedy is now the Associate Contract Project Manager and is directly responsible to Bill Boroski.

The project met all of its milestones with the minor exceptions of its user satisfaction score (which it narrowly missed) and its deployment milestone due to flooding in Thailand which affected the entire chip industry in 2012.

The selection strategy for the FY2014 deployment was presented by Don Holmgren, the technical contractor project manager at FNAL. The hardware options are 1. A partial rack of BlueGene/Q, 2. Infiniband clusters based on Intel or AMD processors, and 3. Accelerated clusters based on either Nvidia "Kepler" GPU chips or the Intel "MIC" architecture. The needs of the USQCD collaboration will be meshed with the results of benchmarks on the three platforms to arrive at a preferred system by mid-August. This information will be presented to the Federal Program Manager for his approval.

Comments

All members of the panel found the technical leadership of LQCD-ext/ARRA to be very impressive.

Progress toward the scientific and technical milestones is good. There were no significant slips relative to projections, beyond delays caused by lack of timeliness in the short-term availability of resources due to extraneous circumstances (funding and chip development). Several milestones were achieved sooner than anticipated. Martin Savage mentioned an anticipated modest delay for reaching the HP10 (2014) DOE Milestone for Hadronic Physics, due to resource limitations.

The clusters, both CPU and GPU varieties, appear to operate effectively and efficiently.

Recommendations

None.

Technical design and scope for FY2013-14

Findings

The project is currently considering three basic hardware options for FY13-14: BG/Q, conventional infiniband clusters based on either Intel Sandy Bridge or AMD Bulldozer chips, and accelerated clusters using either Nvidia Kepler GPUs or Intel MIC.

Presentations were made by Chip Watson on FY12 deployments, by Don Holmgren on the FY13 and FY14 acquisition strategies, alternative analyses and scheduling plans, and by Robert Mawhinney on the BlueGene/Q acquisition, and operations at BNL.

The project must submit a funding request by mid-August that specifies the budget allocation for each lab. Don Holmgren reviewed the procedures and timetable that the project is planning in order to meet this deadline.

The project has access to early and prototype hardware and is benchmarking the various alternatives. Don Holmgren and Chip Watson reviewed the details of these processes.

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.

Intel MIC could offer a potentially attractive option since codes should port to it more easily than to GPU clusters and it should provide a higher price-performance than conventional CPU clusters.

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/ARRA community for the project to continue increasing its impact on the HEP and NP programs. The SciDAC-III software packages will be crucial for the project to meet its science milestones.

Recommendations

None.

Feasibility and Completeness of Budget and Schedule

Findings

The project has continued to meet its 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 those issues. This will remain an issue in the future and is being incorporated into the acquisition plans.

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.

The broadened membership of the SPC, which now includes an experimentalist and a theorist, was duly noted.

Following recommendations from past reviews, a Science Advisory Board (SAB) has been formed. It 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.

The development of a "Speaker's Bureau", headed by Andreas Kronfeld, to publicize the successes and breakthroughs achieved by LQCD-ext/ARRA and to engage and promote the junior members of the collaboration was commended by the review panel.

The Infiniband cluster and the GPU clusters are running with excellent availability. The Infiniband cluster is available ~99% of the time and utilization is high, and the GPU cluster performance has been excellent and its utilization has been generally good with the cycles devoted to several large demanding projects.

Comments

The panelists commended the project team and USQCD on the management of the clusters.

Responding to comments from last year's review, the SPC is engaging younger members of the USQCD community in a positive, productive fashion.

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

Recommendations

None.

LQCD ARRA Closeout Review.

The status of the project in achieving project goals and performance specifications.

Findings

The contractor project manager, Chip Watson, presented an overview, management and performance summary of the LQCD/ARRA project.

Acquisition and deployment is complete. The project completed its hardware component generally on schedule. The LQCD/ARRA project merged with the LQCD-ext project in FY2012 to optimize resources.

The project exceeded its milestones within its budget and schedule. It reached its project goals significantly faster than anticipated and in the end exceeded them by a significant margin.

The LQCD/ARRA project, by adopting GPU systems (a "disruptive technology"), achieved considerable gains over the originally projected goals. In particular, instead of a projected 16 Teraflops, the project achieved an effective 95 Teraflops with a fixed hardware investment.

The resources which were deployed as part of the ARRA project are: 1. Conventional systems (CPU clusters) consisting of 544 nodes, 2. GPU accelerated systems consisting of 123 quad GPU nodes, 3. Xeon Phi Accelerated systems consisting of 12 quad nodes, and 4. A file system consisting of 18 servers.

The ARRA project took advantage of lessons learned from the LQCD-ext hardware project. It used the build-to-cost procurement strategy which allowed it to keep contingency funding very low (~5% for labor). Allocation of its resources was determined by USQCD processes, the same ones applied to hardware purchased by the LQCD-ext project.

Use of the GPU capability ramped up quickly and the GPUs are well utilized, as monitored in real time at TJNAF. GPU clusters are now an established part of the LQCD-ext/ARRA acquisition plan. Efficient use of the GPU clusters depends on software developed under the SciDAC program. The third acquisition of the ARRA project, the Xeon Phi system, consists of an accelerated architecture, "MIC", developed by Intel which promises to provide accelerated cycles competitive with Nvidia GPU's but with user friendly software and a more conventional computing environment. This hardware option is now a contender in the annual acquisition competition of the LQCD-ext/ARRA project.

Comments

The project was a model of good planning and sensible risk taking (it pioneered the development of GPU clusters for lattice gauge theory).

The implementation of GPUs to serve the USQCD community has been very impressive. This work is pioneering in many ways, and to implement it under the time constraints of a Recovery Act project was admirable.

One reviewer remarked that the idea of devoting "disruptive" technologies (GPUs) to the project and implementing it successfully was "a stroke of genius".

The project team presented several examples where the USQCD community's access to GPUs enabled exploration of science areas previously constrained by access to adequate computational resources.

Recommendations

None.

USQCD Review

Effectiveness, scientific impact, operational procedures and related activities

Findings

The USQCD collaboration consists of ~163 physicists with ~100 participating in physics proposals in 2013.

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/ARRA project.

Following recommendations from past reviews, a Science Advisory Board (SAB) has been formed. It 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 they participate in the annual all-hands meeting.

The computing resources governed by USQCD are typically ~30% to ~70% oversubscribed by requests from the collaboration members.

USQCD was awarded 430M hours on Leadership Class machines in 2013. The CPU clusters of LQCD-ext/ARRA provided 469M hours and the GPU clusters provided 385M hours.

In 2012, USQCD won two SciDAC grants, each of ~$1M in size from NP and HEP separately.

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. This success addresses a concern of previous review panels.

USQCD has formed a Speakers Committee to find prime opportunities for junior members of the collaboration.

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 previous review committees and the present one 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 Scientific Program Committee 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 Scientific Program Committee is balanced in age and across subfields; term limits of three to four years ensure that its membership rotates regularly. The Executive Committee 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, untenured persons on the SPC was recognized by the review panel as an important development which should be continued.

It is important that the Executive Committee remains responsive to long-term changes in the field and its mission. To feel the pulse of the LQCD community, it may consider adding one or two term-limited members that are elected by the entire collaboration.

The founding of the Scientific Advisory Board (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 last year's 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 seems sound. The user surveys show that the collaboration is generally (~80%) satisfied with the outcome of the resource allocation.

The fact that there is a continuing stream of quality publications appearing in the four target physics areas is proof that the scientific goals of the allocation process are being attained.

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 a question from the review panel, it was noted that USQCD has produced ~60 Ph.D.'s over the last 10 years. The review panel considered this productivity impressive and suggested that the collaboration compile these statistics annually.

The previous review in 2012 raised issues regarding the governance of the hardware constructed and operated by LQCD-ext/ARRA and, in particular, the membership and tenure periods of the Executive Board (EB) and the Scientific Program Committee (SPC). After adopting several recommendations from that review, the reviewers are now more optimistic. The current governance method is well suited to achieving several goals: finding people for the EB and SPC who are well qualified to lead, who are willing to expend the time and the energy necessary to do it, who have a vision of the field, and who are compatible in temperament and goals with the other members of the leadership team. However, the review panel noted that following present procedures the EC could turn over merely by replicating itself, thereby excluding the possibility that people with radically different, but useful ideas, could join the leadership team. The leadership might not be adequately sensitive to the opinions of the younger members of USQCD. However, in light of the success of USQCD in governing itself and the hardware project, the review panel does not think major changes are required. It does, however, urge USQCD leadership to continue to think about these issues and fine tune its governance processes accordingly.

The USQCD sponsored workshops have added to the impact and visibility of lattice calculations. It is important to continue and even expand these efforts, if possible. The US experimental program will be evolving rapidly over the next few years and the lattice community must continue to stay abreast of those developments. Participation in the Snowmass process, the upcoming new P5 process and related activities within HEPAP are all important here. Perhaps the lattice community should lobby for an increased role in DOE advisory committees such as HEPAP and NSAC.

Recommendations

None.

APPENDIX A

Charge Letter to the LQCD-ext/ARRA 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 9-10, 2013, at the Thomas Jefferson National Accelerator Facility (TJNAF). 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 9-10 review. The focus of the 2013 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 2013;

- 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.

In addition, we will also be using this review to assess the status of the project at TJNAF to construct and operate LQCD clusters, which was funded under the American Recovery and Reinvestment Act (ARRA) of 2009. This project was completed in 2012. The panel is asked to assess:

- The status of the LQCD-ARRA project in achieving project goals and performance specifications.

Because ARRA funding is subject to special scrutiny, the LQCD-ARRA project will receive a separate report. Chip Watson, the Contractor Project Manager should present the relevant documentation in order to allow the panel to perform the above assessment.

The LQCD-ARRA clusters are subject to the same governance as those of LQCD-ext. Consequently, those two efforts are now merged into one operation under your direction. Since LQCD-ext computer cycles are allocated 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-ext and LQCD-ARRA projects 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 TJNAF 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 TJNAF.

Sincerely,

James Siegrist
Associate Director
Office of High Energy Physics

Timothy Hallman
Associate Director
Office of Nuclear Physics

APPENDIX B

Reviewers for 2013 LQCD-ext Annual Review (TJNAF May 9-10)

NP Reviewers

Theory: Ulrich Heinz (Ohio State) Heinz@mps.ohio-state.edu

Experiment: Matt Shepherd (Indiana University) mashephe@indiana.edu

Computing: Charles Maguire (Vanderbilt) charles.f.maguire@vanderbilt.edu

HEP Reviewers

Theory: Geoff Bodwin gtb@hep.anl.gov

Experiment: Hassan Jawahery jawahery@physics.umd.edu jawahery@slac.stanford.edu

Computing: Jay Srinivan jay@nersc.gov

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

May 9

08:30	Executive session (45 min)
09:15	Welcome (10 min) – *Hugh Montgomery*
09:25	Logistics and Introductions (5 min) – *Bill Boroski*
09:30	LQCD Overview & USQCD Governance (50 min) - *Paul Mackenzie*
10:20	Break (15 min)
10:35	USQCD Allocation Process (15 min) – *Robert Edwards*
10:50	Science Talk 1: Intensity Frontier Lattice Gauge Theory (30 min) – *Ruth Van de Water*
11:20	Science Talk 2: Lattice Gauge Theory for the Energy Frontier (30 min) – *Julius Kuti*
11:50	Science Talk 3: Hadron Spectroscopy, Structure and Interactions (30 min) – *Martin Savage*
12:20	Lunch / Executive Session
1:30	Science Talk 4: High Temperature/Density QCD (30 min) – *Frithjof Karsch*
2:00	LQCD-ARRA Project: Project Summary and Closeout Report (40 min) – *Chip Watson*
2:40	LQCD-Ext Project: Management and Performance (30 min) - *Bill Boroski*
3:10	Coffee Break (20)
3:30	LQCD-Ext: Technical Performance of FY2012 Deployments (30 min) – *Chip Watson*
4:00	LQCD-Ext: FY2013 Hardware Plan & Cluster Deployment Status (30 min) – *Don Holmgren*
4:30	LQCD-Ext: FY2013 BlueGene/Q Deployment Status (15 min) – *Bob Mawhinney*
4:45	LQCD-Ext: Proposed Selection Strategy for FY2014 Deployment (30 min) – *Don Holmgren*
5:15	Executive Session (60 min)
6:15	Committee request for additional information - *Committee/Project Leadership*
6:30	Adjourn
7:00	Dinner

May 10

8:30	Executive Session (30 min)
9:00	Committee questions and discussion (60 min)
10:00	Break
10:10	Executive Session / Preliminary Report Writing
12:00	Lunch
1:00	Executive Session / Closeout Preparation
2:00	Closeout
3:00	Adjourn