Offices of High Energy Physics and Nuclear Physics
Report on the

LQCD-ext/ARRA
2012 Annual Progress Review

May 16-17, 2012

Executive Summary 4
Introduction and Background 5
Continued Significance and Relevance 9
Findings 9
Comments 10
Recommendations 12
Progress towards Scientific and Technical Milestones 12
LQCD-ext 12
Findings 12
Comments 13
Recommendations 13
LQCD/ARRA 13
Findings 13
Comments 14
Technical design and scope for FY2012-13 14
LQCD-ext 14
Findings 14
Comments 15
Recommendations 16
LQCD/ARRA 16
Findings 16
Comments 16
Feasibility and Completeness of Budget and Schedule 16
LQCD-ext 16
Findings 16
Comments 16
LQCD/ARRA 17

Findings.. 17
Comments 17
Effectiveness of Management Structure and Responsiveness to past Recommendations............ 17
LQCD-ext 17
Findings.. 17
Comments 17
LQCD/ARRA 18
Findings.. 18
Comments 19
APPENDIX A.. 20
APPENDIX B 22
APPENDIX C 23

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 16-17, 2011 at the Brookhaven National Laboratory (BNL). The purpose of the review was to assess the projects' progress towards their overall scientific and technical goals. Six expert reviewers from the nuclear physics, high energy physics and computer science communities heard presentations on scientific progress, computing hardware acquisitions and operations, allocation of resources, and dissemination of scientific results. In particular, the LQCD-ext/ARRA teams were 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 projects' 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.

In general, the review panel was very impressed with the technical and scientific achievements of the LQCD-ext/ARRA efforts. The impact of LQCD simulations on experimental programs in precision measurements of the Standard Model (SM), Heavy Ion collisions and 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). However, the review panel felt that the project is not using its hardware optimally, so they recommended that within each of its subfields, LQCD-ext/ARRA should focus on achieving a small number of high impact results rather than executing many calculations of lesser impact. In addition, the review panel had four comments and associated suggestions concerning the governance of the project that they felt are important to the long term viability of LQCD-ext/ARRA: 1. USQCD, the governing collaboration of the LQCD-ext/ARRA project, presented demographic statistics that showed that the field attracts students and post-docs, but universities rarely hire them into junior faculty positions. It would benefit LQCD-ext/ARRA if USQCD could develop strategies to address this shortcoming. 2. The review panel felt that the Scientific Program Committee, which oversees resource allocations for USQCD, would be more effective if it includes representatives of the relevant experimental communities. 3. USQCD could publicize the achievements of LQCD-ext/ARRA more effectively if it formed a speakers' bureau modeled after successful HEP/NP experimental collaborations' bureaus, and 4. USQCD would be a more potent governing body is it drafted a constitution that codifies a periodic rotation of all its leadership and executive positions based on a democratic process.

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 (LQCD) IT hardware acquisition and operations activity, 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 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 collaboration 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 Leadership Class supercomputers available through the INCITE program. Members of the USQCD collaboration submitted proposals for computer time, some on the Leadership Class machines for large scale capacity computing, and some on the dedicated clusters of LQCD for large scale capability computing. The resources were awarded on a merit system. Three classes of computer projects have been considered, ranging from large scale mature projects (allocation class A) to mid-sized projects (allocation class B) to exploratory projects (allocation class C). Suitable computer platforms were assigned to the various projects.

In addition to the hardware project LQCD, USQCD has 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 developed efficient portable codes for QCD simulations. USQCD had a 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. USQCD has submitted proposals to the SciDAC-III program.

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

In the January 30, 2008, review, USQCD argued that the mid-scale computer hardware purchased, constructed and operated by LQCD was a critical portion of its overall strategy to produce physics predictions of 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 of LQCD, and results of interest to the experimental and theoretical communities in high energy physics and nuclear physics are obtained. The mid-scale hardware of LQCD also produces smaller gauge configurations which are critical to studies of Quantum Chromodynamics in extreme environments that are relevant to 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 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 it received formal CD-1 approval on August 27, 2009, through a paper Energy Systems Acquisition Advisory Board (ESAAB) presentation and 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 produced a planning budget for the LQCD-ext CD-2/3 review which read:

Table 1. Planning Budgets 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 was supported by the scientific review of January 30, 2008, and which USQCD had estimated in their original whitepaper. This shortfall was subsequently addressed, however, by the request of the Office of Nuclear Physics for $4.96M of funding through the American Recovery and Reinvestment Act of 2009 (ARRA) to build a 16 Tflop/s commodity cluster at TJNAF and operate it for four years. Although this effort is not a formal part of this LQCD-ext project, the resulting hardware at TJNAF is being governed by USQCD using exactly the same procedures that apply to LQCD-ext and the acquisition, construction and operation of this hardware is being tracked on a monthly basis by the same team that is running LQCD-ext. In this way, the Offices of High Energy Physics and Nuclear Physics are monitoring 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 and this report is the third in a series.

LQCD-ext argued at the CD-2/3 review that the budget of Table 1 would support the new deployments and operations of equipment contained in Table 2:

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, Tflop/s	11	12	24	44	57
Planned delivered Performance (TJNAF + FNAL + QCDOC), Tflop/s-yr	18	22	34	52	90

The original computing goal for the LQCD/ARRA project was 16 Tflops (sustained) from a single cluster at TJNAF. The project team initially estimated that $3.2M would be used for hardware that would 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 and it reads:

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**

However, the planning for hardware acquisition for LQCD-ext/ARRA has been strongly affected in FY2010-11 by a “disruptive technology” development in the field of PC chips. Although the first year of acquisitions were expected to be based on commodity cluster technologies, the development of Graphical Processor Units (GPU) for the commercial gaming industry has given new opportunities to these projects. GPUs consist of several hundred cores per chip and are the heart of high resolution interactive graphics capabilities needed for video game entertainment. Typically they are capable of an order of magnitude more processing per second than general duty desktop CPUs. However, they are difficult to program at this time and are unbalanced (too little memory per core) for general purpose applications. However, low memory but compute intensive and highly parallel algorithms, such as the heart of lattice QCD where 90%+ of the CPU time is spent in inverting a sparse matrix, the Dirac operator describing the dynamics of virtual quarks of QCD, can take advantage of a GPU’s floating point capabilities and can run 10-100 times faster than on a CPU of comparable clock period. 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 GPU hardware ordered for LQCD/ARRA. Two complete physics projects ran on a GPU cluster at TJNAF during the GPU cluster’s first year of availability, but that number has grown to ~9 in the second year and is expected to continue to increase in the near term. Their price performance is ~$0.0l/Mflops which compares to $0.15-0.22/Mflops for the best CPU hardware. This development constituted an important new alternative in the hardware acquisition strategy of LQCD-ext/ARRA and was considered in detail by previous review teams. These reviews have had several observations about this development: 1. The success of the hardware project LQCD-ext/ARRA 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 for LQCD-ext/ARRA because many lattice scientific applications are not ready to be ported to GPUs but would be greatly more productive if and when that happens; 3. The initial estimates of TFlops of clusters that can be built for $22.15M will probably be considerably higher than the planning figures shown above, but it is hard to estimate new milestones at this time; 4. The scientific output and impact of LQCD-ext/ARRA may be considerably higher than originally planned for; and 5. The risk associated with the new GPU hardware will exceed that of the more familiar CPUs. All these considerations became part of the discussions of the planning for LQCD-ext/ARRA in FY2010-12 . Several of these observations have met with fruition: The ARRA GPU cluster is sustaining ~76 Tflops on a fairly diverse set of physics projects, beating 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 BNL on May 16-17, 2012. 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 five sections of this report detail the findings, comments, and recommendations of the review committee for each of the charge elements that the LQCD-ext/ARRA project teams were asked to address.

Continued Significance and Relevance

Findings

Paul Mackenzie, the chairperson of the executive committee of the USQCD collaboration, presented the overview of LQCD-ext/ARRA.

The LQCD-ext/ARRA program supports activities in several research areas:

1) Precision calculations relevant to the determination of standard model parameters from heavy quark processes. 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 have been made. Ruth Van de Water summarized this subfield of LQCD at the review.

2) Exploratory calculations based on "beyond the standard model" (BSM) theories, for which LQCD is at present the only effective technique for extracting quantitative predictions. 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. GPU clusters are proving useful in these studies. Julius Kuti summarized this subfield of LQCD at the review

3) Hadronic physics quantities such as the spectrum of hadrons, form factors, moments of structure functions, hadron-hadron interactions and scattering. 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. Kostas Orginos summarized this subfield of LQCD at the review

4) Calculations of the properties of QCD at finite temperature and baryon density; this regime is explored experimentally in relativistic heavy ion collisions. These simulations are having an impact on the run plans of the Relativistic Heavy Ion Collider (RHIC) at BNL. Peter Petreczky summarized this subfield of LQCD at the review

USQCD’s scientific goals are focused on carrying out world-leading computations of quantities that are of importance to the experimental high energy physics (HEP) and nuclear physics (NP) programs.

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

LQCD continues to have workshops with the experimental and theory communities to widen its impact and engage in communications with complementary communities of researchers to enhance its influence and impact. The most recent workshops on QCD were mainly focused on Nuclear Physics spectroscopy which is relevant to the experimental program at TJNAF and High Energy Physics “Beyond the Standard Model” topics.

The demographics of the lattice gauge community is presented at each annual review. This year’s statistics showed that there is a net in-flow into the field at the graduate student and post-doctoral level, but there is a net out-flow at the junior faculty level. 5 out of 16 HEP researchers found jobs in research institutions over the last 10 years. The corresponding statistic for NP is 7 out of 11. Over the last four years no lattice gauge researcher has been hired into a tenure track position at a US institution. No top ranked US institution has hired a lattice gauge theorist over the last 10 years.

Comments

Lattice gauge QCD has been transformed in the past several years both by algorithmic developments and the harnessing of high performance computing technology. In particular, moving from the quenched approximation to unquenched opened up the ability to produce reliable precision calculations of quantities vital in moving forward the field's understanding of fundamental physics (High Energy Physics (HEP), Nuclear Physics, Astrophysics and Cosmology), that are otherwise not amenable to calculation by analytical methods.

The renaissance in lattice gauge QCD is also bringing an expanding scope to the possible calculations. As an example it is finding new application in the study of possible new strongly coupled systems at the LHC. There are demonstrably broader impacts in the technology transfer of the algorithmic, software and hardware developments to the high performance computer industry and to other fields of science. For these reasons it is important that the US continue to maintain a vibrant, competitive, well funded lattice QCD community.

In the HEP component of the LQCD, there is good incremental progress in a number of important calculations such as the exclusive semi-leptonic decays relevant to the extraction of V_{ub}. In the other fields there is similarly good progress in a broad set of calculations. Overall the panel judged the progress to be good but not outstanding.

Several members of the panel expressed concern that the breadth of the program is leading to a large number of calculations but there is a not a focus on a small set of particularly high impact results. As a consequence the panel asked each of the subfields to provide the top three possible high impact calculations which they duly did. It was notable that there was not complete agreement with the relevant experts on the panel on these choices and their ranking, particularly in respect to what was most significant to experimentalists.

The LQCD program is essentially controlled by the Scientific Policy Committee (SPC) of the USQCD collaboration. The SPC consists exclusively of LQCD theoreticians who review proposals for allocation of time on the clusters. The review panel strongly suggests that the SPC be expanded to include knowledgeable representatives of the relevant experimental communities. The expanded SPC should define the set of high impact results and make those known to the entire USQCD and related experimental communities.

As noted above, the panel was not entirely convinced that the prioritization of high impact results provided by the USQCD representatives would find general agreement among the experimental community. However it is important to ensure that a significant fraction of the program is focused on the high impact results relevant to the experimental community. The SPC is proposal driven rather than setting its priorities and allocating according to those priorities. Several panelists were concerned that with this model there is a tendency to try to be collegial and equitable in the allocation process.

Several panelists felt that the program is very vulnerable to a budget cut if it continues in its present fashion.

Several panelists were alarmed at the demographics which showed that US institutions are not hiring lattice gauge theorists into junior faculty positions at a rate that could sustain the field. They strongly suggested that USQCD develop a strategy to encourage US universities to hire lattice gauge theorists into junior faculty positions. The Bridge Program developed by TJNAF and its member universities could serve as a model for an expanded effort. It was noted that the TJNAF program in this area has been quite successful and has helped NP grow in the lattice field.

Recommendations

Each subfield within the LQCD program (Precision calculations of parameters for the Standard Model, Relativistic Heavy Ion Collisions, etc) should be tasked with providing a small number of high impact results deliverable prior to the 2014 funding renewal. The SPC should give priority in its allocation decisions to those proposals which will most likely further the program of high impact results.

Progress towards Scientific and Technical Milestones

LQCD-ext

Findings

Bill Boroski, the LQCD-ext contractor project manager, presented the management and performance information for the project. The presentation covered the project scope, budget and organization, including the FY11 and FY12 performance and financial results, a summary of the user survey and the FY13 hardware acquisition strategy.

The project is meeting all of its 19 IT E-300 milestones with the minor exceptions of its user satisfaction score (which it narrowly missed) and its deployment milestone due to the continuing resolution.

Of note, the project exceeded its FY11 milestone of delivering 22.0 TFlop-years by 43% by delivering 31.48 TFlop-yrs. It is on schedule to exceed its FY12 performance milestone as well.

The project is planning to formulate a new project goal that better measures the effectiveness of both conventional CPU and GPU clusters.

The selection strategy for the FY13 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 is information will be presented to the Federal Program Manager for his approval.

Comments

With some additional effort on improving user documentation or simplifying access (especially at FNAL), the user satisfaction score could be improved.

All members of the panel found the technical leadership of LQCD to be very impressive.

LQCD-ext/ARRA has exceeded its milestones through harnessing new technologies such as GPUs.

The clusters appear to operate effectively and efficiently.

On panelist had a minor criticism which was that only one of the clusters (TJNAF) used prioritized batch queues. These are quite standard on most processor farms. They provide an effective way to balance loads and can be used to implement recommendation 3 above.

Recommendations

The clusters should implement prioritized batch queues to place lower priority jobs on slower processors and to more rapidly execute the high priority calculations.

In the past, the project has relied on a few simple kernels to measure system performance. The project has now modified that approach to account for overheads when using accelerators like GPUs. The project has established a GPU effective performance (GPUEP) metric for its procurements. The project should adopt this approach for tracking progress towards it technical milestones. This will allow the project to more accurately measure the improvements from new architectures such as GPU or Intel MIC. However, this requires developing a conversion approach that is defensible.

LQCD/ARRA

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 will merge with the LQCD-ext project in FY13 to optimize resources.

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 Tflops, the project is achieving an effective 84 Tflops with a fixed hardware investment. .

Use of the GPU capability ramped up quickly and the GPUs are well utilized.

Comments

The new GPU technology necessitated larger than expected commitments of manpower resources especially for software development (reprogramming in CUDA). These manpower costs for software development have been borne successfully by other funding sources.

The implementation of GPUs to serve the LQCD 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 is admirable.

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

Technical design and scope for FY2012-13

LQCD-ext

Findings

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

The slow-down in the rate of improvement in conventional clusters would likely consume most of the project's scope contingency. BG/Q and GPU clusters offer the best chance for meeting the targets, but not all applications have been ported to these platforms and some ports will likely not see the same benefits as previous efforts (i.e., early GPU ports).

If BG/Q is selected then it will be sited at BNL, while clusters will be sited at FNAL. Despite the fact that the best option is not yet known, the project must submit a funding request by mid-August that specifies the budget allocation for each lab.

The project has access to early and prototype hardware and is benchmarking the various alternatives.

The project has elected to distribute resources across three sites (BNL, FNAL, and TJNAF). This approach allows it to better leverage resources at each site, insure that one site is not overburdened with support, and helps to insure that some resources are always available to the project. However, this approach does mean that data may need to be moved between the sites in order for scientists to accomplish their goals. All of the sites are connected through the DOE ESnet at 10Gb. However, many of the storage resources associated with the project are connected at 1Gb.

Comments

The panelists commented that the project's just-in-time acquisition strategy and benchmarking protocols make good sense.

The development of compilers to allow more calculations to run on the GPU cards will allow more efficient use of resources. The addition of an IBM Blue-Gene/Q rack at BNL will enhance the program and is cost effective since the servers and disks come for free.

Overall the technical planning is excellent.

The project essentially has a trigger date of August 15th. The project should establish some threshold for deciding whether or not to pursue the BG/Q route or cluster route, since this must be known in order to make the appropriate budget request.

There are some potential platforms that were excluded (Cray Cascade) but otherwise the list looks complete. The project should continue to focus on benchmarking on early hardware. Intel MIC could offer a potentially attractive option since codes should port to it more easily and it should provide a higher price-performance than conventional clusters.

The likely need to leverage GPUs or other accelerator-based approaches in order to meet project milestones does create a critical dependency on external development activities. While there is significant engagement across the LQCD community, some risk remains that the project will fail to meet its targets unless these other activities are successful. Investments by DOE-SC (through SciDAC-3) are critical to the success of the entire LQCD community and, especially, this project.

From a non-technical point of view, the dependency on the SciDAC-3 program for software development is a potential problem, especially since the global level of funding is clearly decreasing. This could lead to a situation where the community does not have the resources to perform software developments, while software is clearly identified as one of the current major challenges. This risk is exacerbated by the difficulties the community is facing to retain expertise. In fact, the dependency on key project members and the technological choice of relying on a proprietary solution like CUDA (CUDA is the programming language developed and maintained by Nvidia), could lead to challenges that will be very difficult to address (for instance if it is necessary to move away from CUDA and key members having the expertise are leaving at the same time). Therefore, the LQCD community should consider replacement solutions for CUDA.

Some sites (TJNAF) are charging for storage usage using a conversion factor to translate storage into core hours. This is useful since it creates an incentive for users to monitor and control storage usage. This should be adopted at the other sites, if feasible.

It may become necessary to upgrade servers at each site to provide full 10Gb access to storage resources associated with the project. This includes both archival storage and file system storage. Some of the storage resources are already accessible at 10Gb, but other end--points need to be upgraded so that scientists can take full advantage of the available bandwidth.

Recommendations

The project should evaluate establishing a method for appropriately scaling the IO subsystem (both capacity and bandwidth). This should be coupled to application needs and should reflect any expected changes in behavior. This could be based on a system characteristic like total system memory or specific applications that are particularly IO demanding.

LQCD/ARRA

Findings

The LQCD/ARRA resources were placed at TJNAF. The acquisition phase of the project is complete. It is in the operations phase. It will merge with LQCD-ext on Oct 1, 2013.

Comments

The merge of the LQCD/ARRA project with the LQCD-ext project is clear and well planned, and should not experience any problem.

Feasibility and Completeness of Budget and Schedule

LQCD-ext

Findings

The project has continued to meet its technical milestones and schedule.

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.

The FY12 acquisition is taking advantage of the experience from the ARRA procurement. A mix of conventional clusters and accelerated clusters are being deployed.

Comments

The FY12 acquisition is blocking on more accurate dates for the availability of Nvidia Kepler GPUs. The project should establish some trigger date for electing to use existing Fermi cards versus waiting for the next generation.

The FY13 acquisition does present some potential risk to the project on meetings its scope (even with scope contingency).

LQCD/ARRA

Findings

The project met or exceeded its milestones within its budget and schedule. It will merge with LQCD-ext on Oct 1, 2013.

Comments

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

Effectiveness of Management Structure and Responsiveness to past Recommendations

LQCD-ext

Findings

There was a suitable response to the one recommendation from the review last year (the user survey was simplified).

Comments

The panelists saw no problems in the development and management of the clusters.

In terms of the broader USQCD program, the panel was interested to judge its effectiveness by looking at the program's final output and how it compares with similar lattice QCD collaborations in Europe and Asia. The lattice community in the US consists of about 125 people from both universities and national labs. In terms of resources, publications and citations some partial data was presented that indicate that the US is competitive.

However it was shown that the US is severely deficient in the career development of young scientists. There is a net inflow of students and post-docs and a net outflow of faculty and staff. In other words the DOE funding is being used to train a generation of scientists who then take their physics and computing expertise elsewhere in the world.

Several panelists suggested that the USQCD is not taking enough responsibility for the career development of its younger physicists. It is recognized that within the broader HEP theory community that there is a stigma attached to lattice QCD. One speaker stated that it was a constant struggle to convince this community that they were not computer jockeys rather than legitimate HEP theorists.

The panelists commented that the possibility of younger physicists obtaining permanent jobs is enhanced if they show leadership potential. The best way to show ability in leadership is to give physicists the opportunities to lead and to give them career development paths within the collaboration that increase their opportunities. The hierarchical nature of academia tends to lead to entrenched leadership and static management unless a periodic overhaul is forced on it. This leads to limited opportunities for physicists of all levels to lead and can be particularly destructive to young people. The tool for doing this is a set of bylaws or a constitution. The review panel, therefore, suggested that the USQCD collaboration draft a constitution that codifies a periodic rotation (every 2 or 3 years) of all its leadership and executive positions based on a democratic process.

The panelists agreed that there has been a seminal change in lattice QCD recently. However, the panelists do not believe these developments have been effectively communicated to the theory community nor, indeed, to the broader physics community. There was general agreement that the USQCD collaboration must engage in a campaign to educate the community

Several panelists suggested that University Colloquia offer an especially important venue for the promotion of the entire field of lattice QCD. However this is a particular type of talk that requires a speaker who can communicate effectively to a wider audience

The review panel strongly suggested that the USQCD collaboration form a speakers bureau modeled along the lines of a typical HEP/NP experimental collaboration's bureau. Its role would be multiple. First, to acquire talking slots at major conferences, national labs and university colloquia. Second, to exercise internal quality control of those talks. This takes the form of practice talks attended by your peers in which criticism is invited and the talk is refined until approved by the speaker's bureau. Third, to promote physicists appropriately at different points in their career by allocating them talks to enhance their visibility.

The speakers' bureau of USQCD should identify a set of speakers who possess a broad knowledge of the field and a demonstrated ability to engagingly communicate to a wider audience. They should communicate the important recent developments and accomplishments of the field, emphasize its broadening scope and synthesis of technology and theory, and especially the increased investment of funding being made. A suggested goal is that 15 colloquia are given in each of the next three years.

LQCD/ARRA

Findings

The LQCD/ARRA management team is modeled after the LQCD-ext team.

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

Comments

The management is performing well. They will merge efforts with LQCD-ext at the beginning of FY13.

They presented a coherent, consistent and complete set of summaries for the review.

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 16-17, 2012, at the Brookhaven National Laboratory (BNL). A review panel of experts in high energy physics, nuclear physics, project management and computer science is being convened for this task.

John Kogut of the Office of High Energy Physics is responsible for this review; he will be assisted by Helmut Marsiske 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 16-17 review. The focus of the 2012 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 plans for, and progress on, the construction and operation of the Thomas Jefferson National Accelerator Facility (TJNAF) LQCD cluster which is funded by the American Recovery and Reinvestment Act (ARRA) of 2009. We are consolidating these reviews because the LQCD ARRA cluster will be operated by the USQCD collaboration like any other hardware platform of the LQCD-ext project. However, since ARRA funding is subject to special scrutiny, it will receive a separate progress report. Chip Watson, the Contractor Project Manager for the LQCD ARRA cluster, should present the relevant information in the LQCD ARRA project documentation in order to allow the panel to evaluate the project according to the above charge elements.

Each panel member will be asked to review these aspects of the LQCD-ext and LQCD ARRA projects 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.

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 Chip Watson and John Kogut to generate an agenda which addresses the goals of the review.

Please designate a contact person at BNL for the review panel members to contact regarding any logistics questions. Word processing, internet connection and secretarial assistance should be made available during the review. You should set up a web site for the review with relevant background information on LQCD-ext, 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 BNL.

Sincerely,

James Siegrist
Associate Director
Office of High Energy Physics

Timothy Hallman
Associate Director
Office of Nuclear Physics

APPENDIX B

Reviewers for 2012 LQCD-ext Annual Review (BNL May16-17)

NP Reviewers

Berndt Mueller mueller@phy.duke.edu

Paul Eugenio eugenio@fsu.edu

Shane Canon scanon@lbl.gov

HEP Reviewers

Tom Appelquist thomas.appelquist@yale.edu

Geoffroy R. Vallee valleegr@ornl.gov

Colin Jessop cjessop@nd.edu

List of attending DOE program managers

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

T. Barnes (NP)

APPENDIX C

https://indico.fnal.gov/internalPage.py?pageId=2&confId=5414

DOE Annual Progress Review of

Lattice Quantum Chromodynamics (LQCD) Computing Projects

LQCD-Ext and LQCD-ARRA

May 16-17, 2012

Brookhaven National Laboratory

Berkner Hall B

AGENDA

May 16

08:30 Executive session (45 min)

09:15 Welcome (10 min) – *Tom Schlagel*

09:25 Logistics and Introductions (5 min) – *Bill Boroski*

09:30 LQCD Overview (45 min) - *Paul Mackenzie*

10:15 Break (15 min)

10:30 Science Talk 1: Fundamental Parameters of the Standard Model (30 min) – *Ruth Van der Water*

11:00 Science Talk 2: Beyond the Standard Model Physics (30 min) – *Julius Kuti*

11:30 Science Talk 3: Hadron Spectroscopy, Structure and Interactions (30 min) – *Kostas Orginos*

12:00 Lunch / Executive Session

1:00 Science Talk 4: High Temperature/Density QCD (30 min) – *Peter Petreczky*

1:30 LQCD-Ext Project: Management and Performance (30 min) - *Bill Boroski*

2:00 LQCD-ARRA Project: Management and Performance (20 min) – *Chip Watson*

2:20	LQCD-ARRA Technical Performance (30 min) – *Chip Watson*
2:50	Combined Break and Compute Facility Tour (40 min)
3:30	LQCD-Ext: Technical Performance of FY2011 Deployments (15 min) - *Don Holmgren*
3:45	LQCD-Ext: Hardware Acquisition Plan & Status for FY2012 (30 min) – *Chip Watson*
4:15	LQCD-Ext: Prototype IBM BG/Q Overview (15 min) – *Frank Quarant*
4:30	LQCD-Ext: Proposed Selection Strategy for FY2013 Deployment (30 min) – *Don Holmgren*
5:00	Executive Session (60 min)
6:00	Committee request for additional information - *Committee/Project Leadership*
6:30	Adjourn
7:00	Dinner

May 17

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
