Office of High Energy Physics
Report on the

LQCD 2008 Annual Review

May 21, 2008

Executive Summary

The Lattice Quantum Chromodynamics (LQCD) project supports the development and operation of a large scale dedicated computing facility for simulations of lattice gauge theories describing nuclear and high energy physics. It consists of a special purpose "Quantum Chromodynamics-on-a-chip" (QCDOC) computer at Brookhaven National Laboratory (BNL) and commodity clusters at Fermi National Accelerator Laboratory (FNAL) and Thomas Jefferson National Accelerator Facility (JLAB). This facility is available to lattice gauge theorists at national laboratories and universities throughout the United States and is managed by the US Lattice Quantum Chromodynamics (USQCD) collaboration. The LQCD project started in FY 2006 and is to be completed in FY 2009. The funding for this project supports the acquisition and operation of ~13 Tflops that, when combined with existing hardware, a 4.2 Tflops QCDOC at BNL, will yield a system capable of over 17 Tflops. The total cost for the four year project is $9.2M. The project's budget provides $2.5 M in FY 2008 to continue the project, and $1.7M in FY 2009 to complete it.

LQCD is charged with the task of finding the most cost effective hardware platform each year for dedicated simulations of Lattice Quantum Chromodynamics codes. In FY 2008 the project will begin installing its largest cluster at FNAL. Early in FY 2009, that cluster will be completed using the hardware funds from that fiscal year with the goal of creating a cluster that will sustain over 6.2 Tflops.

The LQCD collaboration was instructed to address five charges in its Annual Project Review, which occurred at BNL on May 13-14, 2008:

- The continued significance and relevance of the LQCD project, with an emphasis on its impact on the experimental programs supported by the Offices of High Energy and Nuclear Physics of the DOE;
- 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 2008-2009;
- The feasibility and completeness of the proposed budget and schedule;
- The effectiveness with which LQCD has addressed the recommendations from last year's review.

The review panel reported that the LQCD collaboration had addressed the five charges in their written as well as their oral presentations and that they had made positive impressions in most cases. The review panel had suggestions and recommendations

on the LQCD collaboration's responses to the five charges. The most important ones which require a response from USQCD are listed here:

1. USQCD should consider including experimenters and theorists from outside the lattice community on the USQCD Executive Committee and Allocations Board.

2. USQCD should encourage more simulation work on the charmonium and open-charm spectra, in light of the states recently discovered at the B factories. USQCD should similarly encourage spectroscopy calculations (light-JPC exotics, etc.) that are relevant to the 12 GeV upgrade at JLAB, since this is currently the highest DOE NP experimental priority.

3. In allocating time to small exploratory projects, USQCD should give special emphasis to physics beyond the Standard Model.

4. USQCD uses a "bottom up", proposal-driven allocation process. There is, therefore, no process to guarantee that the LQCD facilities will be used to meet the priorities of the broader High Energy and Nuclear Physics communities. USQCD might consider allocations of some assigned activities to assure the physics community that some high priority opportunities are not missed.

5. LQCD stated that an additional 0.65 FTE is needed at FNAL and JLAB to support the running clusters. The dollar amounts of this budget change should be determined and this plan should be presented to the executive board for approval. The impact on the hardware acquisitions scheduled for FY 2008 and 2009 should be quantified.

6. The schedule contingency tied to the uncertainty in the availability of this year's most cost effective chip technology, Intel's Nehalem, in late FY 2008 should be developed and the risks to the project's milestones should be clarified.

7. The informal contributions of space and power that the labs make to LQCD should be tracked quantitatively and, if necessary to insure stability of the project, should be formalized through amendment of the present Memorandum of Understanding (MOU).

8. The FY 2007 user surveys indicated low user satisfaction with the transparency of the allocation process. Additional more specific surveys should be pursued by LQCD to identify and resolve the problem.

Additional suggestions, recommendations and discussion points are in the text of this report.

Table of Contents

1. Introduction

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 LQCD project, hardware acquisition and operations activity, which started in 2006 and runs through 2009, has operated a "Quantum Chromodynamics-on-a-chip" (QCDOC) machine at BNL, and has built and operates special purpose clusters at FNAL and JLAB with the goal of providing over 17 Teraflop/s of sustained computer power to lattice calculations.

The hardware acquisition strategy of LQCD is essential to its success. Each year the collaboration benchmarks the kernels of the QCD code on the newest cluster hardware and the winner of the price-to-performance competition becomes that year's provider. The FY 2008 budget for LQCD is $2.5M in support of hardware and operations. The competition among the hardware providers is well underway, and a purchase is scheduled for Fall of CY 2008. This acquisition will be combined with a similar one early in FY 2009 to assemble one large cluster at Fermilab, with a performance goal of 6.2 Tflops. The FY 2009 budget is $1.7M.

The usage of hardware produced by LQCD is governed by the USQCD collaboration through its executive board and allocations committee. Members of the USQCD collaboration submit proposals for computer time, some on general purpose supercomputers run by National Energy Research Scientific Computing Center (NERSC), National Nuclear Security Administration (NNSA) and the National Science Foundation (NSF) and some on the dedicated clusters. The resources are awarded on a merit system. Three classes of computer projects are 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 are 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 now has a SciDAC II grant (2006-2011) which will optimize its codes for multi-core processors and create a physics toolbox. These SciDAC grants provide a user interface to lattice QCD which permits the user to make lattice QCD simulations and measurements without the need to understand the underlying technicalities of the lattice formulation of relativistic field theories and its implementation on massively parallel computers.

LQCD proposes to extend its work beyond 2009, and has submitted a proposal for the 2010-2014 time period, named "LQCD-II." That project is under consideration by the Offices of Nuclear Physics and High Energy Physics.

The Annual Physics Project Review of LQCD occurred at BNL on May 13-14, 2008. 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 present the details of the review. The remaining five sections of this report detail the findings, comments and recommendations of the review committee for each of the charges that the LQCD collaboration was asked to address.

2. The continued significance and relevance of the LQCD Project

Findings and Comments:

1. The LQCD project has an impressive record of scientific accomplishments, particularly in extracting the fundamental parameters of the Standard Model (CKM matrix), the light quark masses and the strong coupling constant.
2. LQCD has made less impressive contributions to hadron spectroscopy, structure and interactions, although work on determining the quark contributions to the nucleon spin has progressed and may lead to a resolution of the "proton spin" crisis.
3. LQCD contributions to QCD in extreme environments have been limited to small volume simulations, although they have been steadily improving. The work on environments rich in baryons has been disappointing because of the lack of an adequate algorithm to overcome the absence of positivity of the fermion determinant here.
4. LQCD is beginning simulations of models which go beyond the Standard Model, but none of the models under consideration are realistic.
5. The number of users of the LQCD clusters has outstripped early growth estimates by a large factor, indicating the attractiveness of the program.

Recommendations:

1. USQCD should consider appointing experimenters and theorists from outside the lattice community to the USQCD Executive Committee and Allocation Board. This would broaden the scientific program and allow a wider community to influence the prioritization of physics topics and the allocation process.
2. USQCD should continue its workshops with other segments of the high energy and nuclear physics communities. It should also continue its series of summer schools to encourage the growth of the field by attracting talented young physicists.
3. As the accuracy of LQCD simulations have improved, small discrepancies between alternative methods and discrepancies with experimental results are becoming apparent. The sources of these problems should be identified. The independence of the members of gauge ensembles should be monitored

closely, and the results of such studies should be included in the stated errors of the resulting matrix elements.

4. USQCD should encourage planning within the community to ensure that analytic calculations in chiral perturbation theory are completed in a timely fashion.
5. USQCD should encourage more work on the charmonium and open-charm spectra, in light of the states recently discovered at the B factories. USQCD should similarly encourage spectroscopy calculations (light JPC exotics etc.) that are relevant to the 12 GeV upgrade at JLAB, since this is currently the highest DOE NP experimental priority.
6. USQCD should encourage the calculation of transport coefficients in finite temperature simulations, since these quantities are crucial to different theoretical approaches to the subject, and are central to the experimental programs at the Relativistic Heavy Ion Collider (RHIC) and elsewhere.
7. In allocating time to Type C projects, USQCD should give special emphasis to exploratory work on physics beyond the Standard Model.
8. USQCD uses a "bottom up", proposal-driven allocation process. There is, therefore, no process to guarantee that the LQCD facilities will be used to meet the priorities of the broader High Energy Physics and Nuclear Physics communities. Several of the recommendations above address this concern in part, but USQCD might consider developing a more definite roadmap outlining actual commitments of groups to particular calculations, with projected estimates of precisions. In particular, USQCD might consider a process that has been applied to large experimental collaborations, specifically providing allocations for some assigned activities to insure the physics community that specific high priority opportunities are not missed.
9. USQCD should become more systematic in making physical quantities (and their associated error matrices) publicly available before chiral and/or continuum extrapolations, to allow future improvements in these areas to be propagated back to earlier results.
10. The number of post-docs, graduates and undergraduates involved in LQCD research should be better documented, in order to understand the impact the project is having on the demographics of NP and HEP.

3. The progress toward scientific and technical milestones.

Findings and Comments:

1. LQCD has made good progress toward its scientific goals.
2. LQCD has exceeded its Tflops goals to this date and is in good position to exceed them by the end of the project in FY 2009.

4. The status of the FY 2008-2009 technical design and technical scope.

Findings and Comments:

1. LQCD has successfully executed its hardware acquisition process, based on an optional price-performance metric.
2. The strategy of evaluating and possibly acquiring Nehalem hardware in FY 2008-2009 is sound. However, the uncertainty and risk in meeting project milestones are considerable.

Recommendations:

1. The schedule contingency and risk associated with the uncertainty in the availability of the Nehalem technology should be clarified.

5. The feasibility and completeness of the proposed budget and schedule.

Findings and Comments:

1. LQCD generally has been within budget and schedule constraints in the past.
2. The labs make considerable informal but crucial contributions to LQCD in power and space, which are not documented. There is no guarantee that this contribution will continue.
3. LQCD stated that an additional 0.65 FTE is needed at FNAL and JLAB.

Recommendations:

1. LQCD should determine the dollar amounts of this budget change in Findings 3 and it should be presented to the USQCD executive board for approval. The funds would have to be taken from the project's hardware acquisition budget, and LQCD should verify their claim that the findings change will not seriously compromise the hardware performance of the planned cluster.
2. The informal contributions of power and space that the labs make to LQCD should be tracked quantitatively and, if necessary to insure stability of the project, should be formalized through amendment of the present MOUs.

6. The effectiveness with which the FY 2007 recommendations have been addressed.

Findings and Comments:

1. The presentations at the review indicated clearly that LQCD has addressed all the recommendations satisfactorily.
2. The user survey response, 54 out of a possible 60, was outstanding, and most of the responses were positive. However, some (less than 10) users expressed reservations about the transparency of the allocation process.

Recommendations:

1. The user surveys indicated that the transparency of the allocation process could be improved. Additional more specific surveying should be pursued by LQCD to pinpoint the source of the problem and to remedy it.

Appendix 1. Links to LQCD Presentations at the Review

The LQCD talks presented at the review and associated background material can be found on the web at http://projects.fnal.gov/LQCD/reviews/May2008Review.

Appendix 2. Charge to the LQCD Collaboration

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

Dear Dr. Boroski:

The Office of Nuclear Physics and the Office of High Energy Physics plan to conduct an Annual Progress Review of the Lattice Quantum Chromodynamics (LQCD) Computing Project on May 13-14, 2008, at the Brookhaven National Laboratory (BNL). A review panel of experts in high energy physics, nuclear physics and computer science is being convened for this task.

Each panel member will evaluate background material on the LQCD project and attend all the presentations at the May 13-14 review. The focus of the 2008 LQCD Annual Progress Review will be on understanding:

- The continued significance and relevance of the LQCD project, with an emphasis on its impact on the experimental programs support by the Offices of High Energy and Nuclear Physics of the DOE;

- 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 2008-2009;

- The feasibility and completeness of the proposed budget and schedule;

- The effectiveness with which LQCD has addressed the recommendations from last year's review.

Each panel member will be asked to review these aspects of the LQCD project and write an individual report on his/her findings. These reports will be due at the Department of Energy 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 first day of the review will consist of presentations and executive sessions. The second day will include an executive session and preliminary report writing; a brief close-out will occur in the early afternoon. Preliminary findings, comments, and recommendations will be presented at the close-out. You should work with John Kogut to make an agenda which can accommodate these goals.

The panel members will be instructed to contact Eric Blum at BNL at (631) 344-2438 or E-mail: blum@bnl.gov 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, links to the various LQCD 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,

Dennis Kovar
Acting Associate Director
High Energy Physics

Jehanne Simon-Gillo
Acting Associate Director Office of
Office of Nuclear Physics

Appendix 3. Agenda of the Review

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

Brookhaven National Laboratory
May 13-14, 2008

May 13

Start	Dur.	Subject
8:30	*0:30*	Executive Session
9:00	*0:10*	Welcome
9:10	*1:20*	Overview (Bob Sugar)
10:30	*0:15*	Break
10:45	*0:45*	Fundamental Parameters of the Standard Model (PaulMacKenzie)
11:30	*0:45*	Hadron Spectroscopy, Structure and Interactions (David Richards)
12:15	*1:00*	Lunch
1:15	*0:45*	High Temperature/Density QCD (Frithjof Karsch)
2:00	*1:00*	Project Management (Bill Boroski)
3:00	*0:15*	Break
3:15	*0:45*	Technical Design and Scope for FY08-09 Procurement (Don Holmgren)
4:00	*0:30*	Responses to Scientific Recommendations from the 2007 Review (Bob Sugar)
4:30	*0:15*	Responses to Technical Recommendations from the 2007 Review (Bill Boroski)
4:45	*1:00*	Executive Session
5:45	Adjourn	
6:30	Dinner	

May 14

Start	Dur.	Subject
8:30	*0:30*	Executive Session
9:00	*1:00*	Committee questions and discussion
10:00	*2:00*	Executive Session and Preliminary Report Writing
12:00	*1:00*	Lunch
1:00	*0:30*	Closeout
1:30	Adjourn	

Appendix 4. Reviewers and Observers

High Energy Experimentalist

Lawrence K. Gibbons, Cornell
lkg@mail.lepp.cornell.edu

High Energy Theorist

Geoff Bodwin (ANL)
gtb@hep.anl.gov

Computer Scientist (Project Management)

Stephen L. Scott (ORNL)
scottsl@ornl.gov

Nuclear Theorist

Simon Capstick (FSU)
capstick@mail.physics.fsu.edu

Nuclear Experimentalist

Alex Dzierba (IU and JLAB)
dzierba@indiana.edu

DOE Participants

LQCD Federal Project Manager: John Kogut John.Kogut@science.doe.gov

LQCD Project Manager (NP) Ted Barnes: Ted.Barnes@science.doe.gov

ASCR Project Manager Vince Dattoria: Vince.Dattoria@science.doe.gov