Department of Energy
Washington, DC 20585

July 5, 2007

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

Dear Dr. Boroski:

The Office of High Energy Physics and the Office of Nuclear Physics conducted an Annual Progress Review of the ongoing Lattice Quantum Chromodynamics (LQCD) Computing Project. As you are aware, this review was held on May 14-15, 2007, at the Thomas Jefferson National Accelerator Facility (JLAB).

Each panel member was asked to evaluate and comment on any relevant aspect of LQCD. However, the particular focus of the LQCD Annual Progress Review was on understanding:

- The continued significance and relevance of the LQCD project, with an emphasis on its scientific goals and its service to a growing user community;
- 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 effectiveness of the management structure.

The co-chairpersons of the review, Maarten Golterman and John Kogut, (the LQCD Federal Project Manager), have written a report based on individual comments and letters from the reviewers and it is attached. We would like you to use this report as guidance for continuing and improving the LQCD project.

The report is generally quite positive. You and the rest of the project management team were singled out for good planning, reporting and hardware acquisition strategies and technical support. The scientific productivity and impact of LQCD were rated highly. Strong interactions with the experimental communities, involving flavor physics, nucleon structure and QCD thermodynamics, were noted and encouraged to grow. The workshop which brought together members of USQCD using the LQCD clusters and the phenomenologists working on BaBar at SLAC was considered to be a model of possible future activities.

There are, however, several areas where LQCD could improve. These are discussed in detail in the report and include the need for user surveys of hardware and system support, hardware monitoring to guarantee the full utilization of the clusters and more outreach by the allocation board to bring new junior users into the program.

Congratulations on your first two years of the project. We look forward to your response to this report and your continued success.

Sincerely,

Robin Staffin
Associate Director
Office of High Energy Physics

Dennis Kovar
Associate Director
Office of Nuclear Physics

Report of the DOE Office of High Energy Physics and the Office of Nuclear Physics Review of LQCD May 14-15, 2007

Introduction

The DOE Offices of High Energy Physics and Nuclear Physics convened an external committee to carry out an on-site review of the Lattice Quantum Chromodynamics (LQCD) project at the Jefferson Laboratory on May 14 and 15, 2007. This document constitutes the summary report of that review. The committee members were David Kaplan, Director of the Institute of Nuclear Theory, Rob Pisarski of BNL, Mark Alford of Washington University, Laurent Lellouch of Marseilles University, Doug Olson of LBNL and NERSC, David Kirkby of UC Irvine and Maarten Golterman of San Francisco State, who co-chaired the review with John Kogut, Federal Program Manager for LQCD. DOE observers were Sidney A. Coon, the Project Monitor for LQCD from the ONP and Vince Dattoria of ASCR. Bob Sugar and other members of USQCD and the executive board of LQCD were in attendance. V. White, who is in charge of the FNAL computer facilities, attended the entire review.

Over the course of 1 1/2 days, the committee heard a series of presentations on LQCD which covered the scientific productivity and impact of the project, the user allocation mechanism and the user experience, the project's management and acquisition strategy and its hardware milestones. These presentations, related documentation, and conversations at the review form the basis of the review committee's evaluation.

There was general agreement between the members of the committee and the consensus opinion of LQCD was very positive.

This report will be organized according to the topics in the charge letter to the LQCD project

1. *The continued emphasis and relevance of the LQCD project, with an emphasis on its scientific goals and its service to a growing user community*

The USQCD collaboration has been successful in providing an excellent large-scale computing infrastructure, through which a number of collaborations within the USQCD community has produced, and will produce first-rate scientific results. The infrastructure consists of the hardware installed at Brookhaven, Fermilab, and Jefferson Lab, as well as user-friendly software packages provided under the auspices of grants funded by the DOE SciDAC program. The organizational structure appears to be optimal in providing a platform for large-scale computing to the US lattice QCD community. In particular, the communication between scientists, technical staff and administrators associated with the project appears to be excellent, and highly efficient.

The SPC (scientific program committee) has been set up to allocate time on the various LQCD facilities on an annual basis. While most time is allocated on a competitive basis to the various large collaborations involved in the major scientific goals of the US lattice effort (class A and class B), it is notable that a small amount of time (5%) is made available to innovative, non-mainstream projects (class C), with an absolute minimum in bureaucratic overhead (class-C projects are informally allocated, and do not have to go through the annual evaluation by the SPC). The committee strongly supports this approach, as it stimulates innovation, and it provides junior colleagues, or others new to the field of lattice gauge theory with an excellent opportunity to try out new ideas. We also noted that all projects submitted to the SPC (large or small scale) get some allocation, and almost all get a substantial fraction of the requested allocation. This suggests that the SPC is making a real effort to be inclusive, while maintaining high scientific standards. In general, USQCD has worked hard to provide an infrastructure with relatively easy access to newcomers. This has already led to the pursuit of important new research. A very good example is the multi-baryon nuclear physics project.

Large-scale projects carried out on LQCD facilities under the auspices of USQCD fall into three categories: a) Fundamental Parameters of the Standard Model; b) Nucleon structure, spectroscopy and hadronic interactions; and c) QCD at non-zero temperature and non-zero baryon density.

a) Fundamental Parameters of the Standard Model

In the area of weak-matrix elements and fundamental parameters, the committee found that the USQCD is leading the world competition, making a very efficient use of computational resources which are not superior to, but more or less competitive with, those in the rest of the world (without LQCD this would not be the case). This work is of great relevance for enabling precision experimental tests of quark flavor mixing and CP violation as predicted by the standard model, which may reveal the presence of new physics. There are strong and direct interactions with the relevant experiments, such as BaBar, CDF, D0, Cleo-c, etc. Several predictions have been made from "first-principle" computations, and they have been verified by experiment. A number of other quantities for which the experimental values have already been known for a long time, such as α_s, quark masses, f_π and f_K have been computed from first principles with unprecedented precision. Lattice computations of a number of weak matrix elements are being carried out with a precision that makes it possible to over-constrain the Standard Model, thus providing a window on possible new physics complementary to that provided by direct searches. The presenters for USQCD made a convincing case that already within a few years the necessary precision will be available, and that beyond that errors on many quantities can be reduced even more, making the comparison with experiment even more stringent.

All quantities for which results with controlled errors exist have been computed using

Lattice QCD with staggered fermions. This is part of the reason that USQCD has been able to take the lead in this direction. At present, there is also a large-scale effort to use domain-wall fermions, which are more expensive computationally (by a factor of about 10 to 20). This is important for two reasons. First, computations with domain-wall fermions will serve to validate the results obtained with staggered fermions, much as one usually would like to validate experimental results independently. A particular reason (but not the only) for doing so is that the use of the staggered-fermion method involves a theoretical issue (the "rooting trick"). Second, domain-wall fermions will be able to provide us with high-precision computations of a number of important quantities (such as B_K and quantities related to the physics of nucleons) for which it is much harder to reach the desired precision with staggered fermions.

With respect to this effort, the committee has the following suggestions:

- We encourage USQCD to work hard on adding new quantities to the list, in particular those relevant for the beyond the Standard Model physics program at the Tevatron and LHC. Examples include matrix elements relevant for supersymmetric extensions and proton decay, as well as three-particle matrix elements. Many of these extensions are already being pursued.
- It is important to take into account correlations in the Unitarity-Triangle analysis resulting from the fact that various lattice inputs move in a correlated way under the extrapolations in lattice spacing and quark masses.
- While the importance of validation between different fermion methods was emphasized by the presenters, we would like to see a clearer outline on how the comparison will be carried out. What will be compared, to what precision can this be done, and what will be learned from the comparison?
- Contact with experimentalists is crucial, and we encourage the organization of workshops bringing together (lattice) theorists and experimentalists. The workshop held at SLAC in September 2006 was an excellent start.

b) Nucleon structure, spectroscopy and hadronic interactions

The committee was impressed by the enthusiasm of members of the USQCD collaboration to attack difficult problems in nuclear physics, such as NN and nucleon-hyperon interactions, radiative transitions and baryon resonances. It appears that for a number of important, basic quantities (such as g_A and moments of structure functions) it will finally be possible to reach light enough pion masses to see the non-analytic terms predicted by chiral perturbation theory. This is crucial in order to have confidence in one's control over the systematic error associated with the chiral extrapolation – a long-standing problem for the numerical computation of nucleonic quantities. The USQCD collaboration made a good case that this will be in reach with the resources provided through the LQCD computing project, at least for a number of interesting quantities.

With respect to this effort, the committee has the following suggestions:

- With respect to baryon spectroscopy, interesting results have been obtained on the resonance spectrum in the quenched approximation at a rather large pion mass. However, many of the resonances in this approximation may be stable, and it is not clear why the analysis applied there will carry over to the unquenched case with light quark masses. The analysis tools developed so far are certainly essential, but may have to be augmented with other tools, such as using volume dependence.

- It is very important to develop tools for dealing with "disconnected" diagrams. The USQCD collaboration is of course well aware of this, and we encourage work in this direction.
- Again, direct contact between lattice theorists and experimentalists is very important. This type of contact seems well established at Jefferson Lab; and theorists are actively involved in the experimental planning at Jefferson Lab. It is worth thinking about workshops along the lines of the SLAC workshop mentioned above.

c) QCD at finite temperature and non-zero baryon density

Again, the committee was very impressed by the quality of research done under USQCD auspices in the field of high-temperature and high-density physics. There is a very strong focus on what is relevant for experiment, notably RHIC, but also future experiments in this direction, such as LHC (CERN, Switzerland) and FAIR (GSI, Darmstadt, Germany). Notable examples are the computation from first principles of the cross-over temperature from hadronic matter to the quark-gluon plasma, which is being done at more or less physical values of the quark masses, and the equation of state, in which discretization errors start to be controlled (by use of finer lattice spacings, as well as by comparison between different discretizations). Very important work is being done on the curvature of the critical surface at small chemical potential. This has direct impact on the determination of a possible critical end-point at non-zero baryon density, where the situation is very confused at present. USQCD made a convincing case that new and solid results on the transition temperature and the equation of state, including that at non-zero baryon density, can be obtained within the time span of the LQCD project.

A recommendation:

- While contact with experiment is crucial, it should also be interesting to keep an eye on finite-temperature results coming from string-theoretical approaches to QCD. A comparison would possibly provide tests of these approaches, which is not unimportant, considering the effort directed toward these approaches by the theoretical community.

A few general recommendations:

- The committee very much supports the policy that all gauge configurations generated within the LQCD project are made available publicly. However, there does not seem to be a uniform policy about when to make these configurations available. The committee recommends that a uniform policy be put in place, with a minimum delay between generation and availability, 6 months for instance.
- Class C allocations are at present undersubscribed. The committee recommends that USQCD find ways to more widely disseminate information about the availability of this type of allocation.

- It would be instructive and useful for any review committee to have a list of scientific publications that have been made possible through the LQCD effort.

In general, USQCD presented quite detailed projections for work in lattice QCD beyond the LQCD computing project. One aspect of this is relevant for this review, and that is the leverage that work carried out under LQCD will provide for large-scale lattice QCD projects beyond LQCD. Members of USQCD have already been awarded time on INCITE with a proposal based on work done with LQCD resources.

The LQCD project appears to be run very well, and with great enthusiasm by all involved, including scientists, technical staff and administrators. Contact with experiment is strong, partly because LQCD has computational and human resources spread over three national laboratories. The resources provided through the LQCD project are crucial for the US lattice QCD community to stay internationally competitive. This will certainly remain true beyond the final year of the LQCD project, 2009, and the committee believes that an increase in computational resources beyond 2009 should be strongly encouraged, building on the success of the 2006-09 LQCD project. There are several reasons for the need to continue beyond 2009, among which:

- While LQCD will provide a number of controlled computations of quantities of interest (as described above), many more quantities will be accessible with computational resources beyond those of LQCD (in all three subfields), and it will be possible to reduce errors on quantities already computed through LQCD.
- For many quantities, domain-wall fermions will probably turn out to be the method of choice, because of the fact that this formulation of lattice QCD retains important continuum symmetries (flavor and chiral symmetries). However, domain-wall fermions are an order of magnitude more expensive, and resources beyond LQCD will be needed in order to exploit their potential.

2. ***The progress toward scientific and technical milestones as presented in the projects IT Exhibit 300***

Progress toward these goals appeared to be satisfactory to the committee (see also below). In particular, the Tflops-yrs projected thus far have been delivered, and there is no reason that this would not continue to be true during the remaining part of the LQCD project. The technical staff appears to be on top of the game when it comes to purchasing decisions, from benchmarking and cost analysis to contact with vendors.

The presenters for USQCD provided a detailed explanation of the delay with the installation of the next (7n) cluster at Jefferson Lab. The delay appears to be caused by two problems: one being the fact that USQCD has had to operate under a Congressional Continuing Resolution, and the other a small but significant deviation from "Moore's law" estimating the cost/performance ratio of future hardware. LQCD has turned these problems into an advantage, by planning to upgrade part of this cluster to newer quad-core processors, which are projected to help it make possible to meet the projected milestones for the year 2007 even with the installation delay. There are some uncertainties in the modeling of scaling of this newer architecture, and the committee urges LQCD to be as timely as possible with exploring the quad-core benefits, in order to maximize their effectiveness. But the committee is satisfied that this part of the project will be under control so that (in fact probably more than) the projected Tflops-yrs will be delivered.

Together with the planned acquisitions beyond 2007 (see point 4. below) the LQCD project is on track to deliver at least the number of promised Tflops-yrs by the end of each calendar year during the project. In total, the current conservative projection is to deliver 15.9 Tflops-yrs by September 2009, to be compared with the 15 Tflops-yrs performance goal of OMB exhibit 300.

While progress appears to be on track, the committee would like to make the following suggestions:

- While the SPC has explained convincingly how they aim to be fair and inclusive, while maintaining high scientific standards, it would be useful to obtain feedback from the users community on the perceived fairness of the allocation process. This could for instance be done through a questionnaire. Maybe the NERSC user survey can be used as an example.
- If possible, it would also be useful to track job failure rates, and in particular find out which part is due to hardware errors, and which part due to user error. This would help in determining whether instruction and assistance to users are optimal, and whether the SciDAC software is of optimal service to users. It may also help identifying users who do not use their allocation efficiently.

There is anecdotal evidence that all these aspects are working reasonably well, but a more systematic effort in collecting feedback and performance data would give more information, and help remedying any problems.

3. ***The progress in meeting the goals of the project's proposed budget schedule, including a discussion of workforce costs***

The project has on average been on budget and schedule so far. The partial delay of the deployment of the planned cluster this year at Jefferson Lab was explained in detail, see point 2 above.

4. ***The status of the technical design and proposed technical scope for FY2008-2009***

USQCD is planning to combine the 2008 and 2009 acquisitions into one single large cluster at Fermilab, with a planned release for production in January 2009, and a projected performance of 6.2 Tflops (by conservative estimate). While Moore's law would make it likely for faster hardware to be available in 2009, integrated physics production would be greater over the first three years of operation with the combined acquisition. The May 2006 Review committee has endorsed this combined acquisition approach. A detailed presentation at this review by D. Holmgren convinced our committee that this change in plans indeed makes sense, and has been very well thought through.

Maarten Golterman and John Kogut co-chairs May 23, 2007