Lattice QCD Computing Project (LQCD)

Response to Recommendations from the 2007 Annual Progress Review of the LQCD Computing Project

Compiled by

Robert Sugar
Chair, LQCD Executive Committee

William Boroski
Contractor Project Manager, LQCD Computing Project

May 7, 2008

INTRODUCTION

On May 14-15, 2007, the U.S. Department of Energy 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. The review was held at the Thomas Jefferson National Accelerator Facility and co-chaired by Maarten Golterman and John Kogut (the LQCD Federal Project Manager).

The review resulted in a written report that identified several areas for improvement. In particular, the report contained 14 recommendations to help improve the project's effectiveness and impact.

This document summarizes the LQCD Computing Project's response to these recommendations and where appropriate, the actions taken to implement specific recommendations.

RESPONSE TO RECOMMENDATIONS

Recommendation #1:

"We encourage the USQCD to work hard on adding new quantities to the list, in particular those relevant for the beyond the Standard Model physics programs at the Tevatron and LHC. Examples include matrix elements relevant for super symmetric extensions and proton decay, as well as three-particle matrix elements. Many of these extensions are already being pursued."

Response: A growing number of USQCD projects address theories for physics beyond the Standard Model. This trend was very much in evidence in the last round of USQCD proposals. Projects include studies of matrix elements relevant to proton decay; the Peskin-Takeuchi S parameter; theories describing new strong dynamics, such as walking technicolor; supersymmetric theories; non-perturbative Higgs physics; and a study of the density of strange quarks in the proton relevant to the detection of the MSSM neutralino.

Recommendation #2:

"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 spacings and quark masses."

Response: This will become an important consideration in the future when there are several results from the same set of gauge configurations, but it is not an issue at the moment.

Recommendation #3:

"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?"

Response: There are two important, but closely related, reasons for comparing results from different fermion methods:

1) To demonstrate that one does have full control over systematic errors, some of which are expected to differ among the fermion methods.
2) To add confidence to the determination of quantities that are particularly important for the experimental programs in high energy and nuclear physics. This would be especially important if one obtained results that were not in agreement with experiment.

Since different fermion formulations have different lattice artifacts, it is necessary to make comparisons in the continuum limit.

Among the quantities currently being calculated with more than one fermion method are the leptonic decay constants of the π, K, D, D_s, B and B_s mesons, the $K^0 - \text{Kbar}^0$ mixing parameter B_K, the quark masses $m_l = (m_u + m_d)/2$, and m_s, the low energy constants of the chiral lagrangian, the heavy quark potential parameters r_0 and r_1, hadron structure observables, the equation of state for high temperature QCD, and the cross-over temperature between ordinary hadronic matter and the quark-gluon plasma. The accuracy with which the comparisons will be carried out in the next few years depends on the quantity being studied, ranging from the percent level for the leptonic decay constants of the π and K mesons, to the five percent level for equation of state and cross-over temperature.

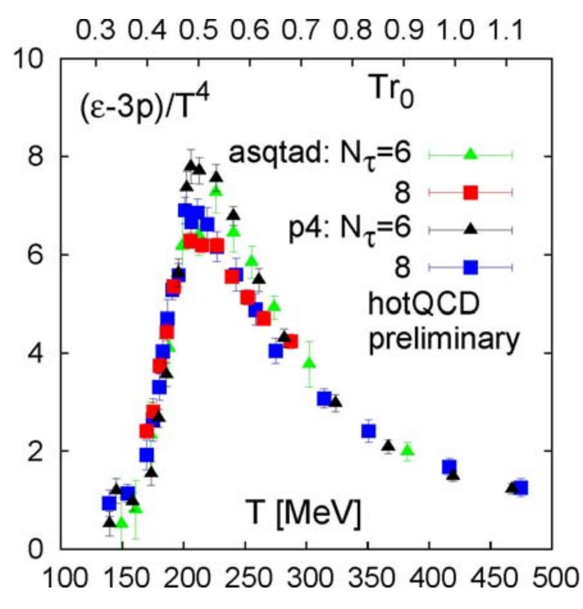

Figure 1. Comparison of the interaction measure $I = \varepsilon - 3p$ using the Asqtad and P4 formulations of lattice quarks by the HotQCD Collaboration.

Recommendation #4:

"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."

Response: The workshop referred to by the 2007 Review Panel involved lattice gauge theorists working on weak interaction matrix elements and experimentalists from the BaBar Collaboration. Another workshop titled "Lattice QCD Meets Experiment" was held at Fermilab on December 10-11, 2007. Talks were given by lattice gauge theorists and representatives of all of the major high energy physics experiments in the US for which calculations of weak matrix elements are important: CDF, D0, BaBar and CLEO-c. Participants described the current status of their calculations or experiments, the expectations for the future, and indicated what each will need from the other in the future.

A series of workshops has been held at BNL to bring together experimentalists working at RHIC and theorists studying QCD thermodynamics on the lattice. The latest of these were "Can We Discover the QCD Critical Point at RHIC?", which took place on March 9-10, 2006, and "Understanding QGP through Spectral Functions and Euclidean Correlators", which was held at BNL on April 23-25, 2008. Another workshop titled "Critical Point and Onset of Deconfinement" is scheduled for March 16–20, 2009 in conjunction with a planned low-energy scan at RHIC in FY 2010 to search for this critical point.

A workshop titled "Lattice Gauge Theory for LHC Physics" was held at Livermore on May 2-3. It brought together lattice gauge theorists with other physicists interested in Beyond the Standard Model Physics to discuss the potential for lattice gauge theory calculations in this area.

Recommendation #5:

"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."

Response: Members of USQCD organized a workshop titled Synergy between Experiment and Lattice QCD in Exploring Hadron Structure, which was held at the Institute for Nuclear Theory, on April 24-25, 2006. A workshop titled Lattice QCD and Experiment: revealing the structure and interactions of hadrons, which aims to bring together experimentalists and lattice theorists interested in these issues, will be held at JLab on November 21-22, 2008.

Recommendation #6:

"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."

Response: A crucial element of the program of anisotropic-clover lattice generation is the use of two (or more) spatial volumes as the quark masses approach the regime where hadrons become unstable under the strong interactions. Another important element of the anisotropic-clover program is to develop stochastic methods for the evaluation of "all-to-all" propagators, with the aim of constructing multi-quark operators that might be expected to couple strongly to two particle states.

Recommendation #7:

"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."

Response: The development of tools for calculating disconnected diagrams is important for many aspects of our research. These include the work on the baryon spectrum mentioned above, the calculation of masses of flavor singlet particles, and the strange quark contribution to the nucleon form factors. A number of techniques are actively being investigated for variance reduction, including dilution, explicit treatment of low eigenmodes and low orders in the hopping parameter expansion, and multigrid preconditioning. Resources were provided to several groups working in this area in the last round of USQCD allocations.

Recommendation #8:

"While contact with experiment is crucial, it would 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 towards these approaches by the theoretical community."

Response: A critical issue is the determination of the bulk and shear viscosities of the quark gluon plasma. Pioneering work in this area has been carried out in two independent USQCD efforts, those of Frithjof Karsch and his collaborators, and of Harvey Meyer. In the recent USQCD allocation process, Meyer received an allocation specifically to advance his calculations of the bulk and shear viscosities. Work on these quantities is also planned by Karsch and his colleagues in the RBC collaboration.

Recommendation #9:

"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."

Response: Our policy is that gauge configurations generated with USQCD resources will be made available to all members of USQCD as they are generated, and that they will be made available to the international lattice gauge theory community via the International Lattice Data Grid as soon as the first paper describing or making use of the configurations has been published.

Recommendation #10:

"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."

Response: The availability of Type C allocations is announced each year by the Scientific Program Committee (SPC) in its Call for Proposals, which is distributed via the USQCD email list. To respond to the undersubscription, the Chair of the SPC sent an email message in January (half-way through the allocation year), to remind the Collaboration of this mechanism for launching projects. The email explicitly asked faculty to forward the information to students and postdocs who might not yet be on the list. Meanwhile, the Executive Committee asked senior collaborators to encourage their postdocs and graduate students to join USQCD, and thereby be included in the email list. This has happened, and the list now has over 150 entries.

We have emphasized those just starting in the field, because they are particularly likely to benefit from Type C allocations. At this year's All Hands' Meeting, the Chairs of both the SPC and of the Executive Committee stressed that Type C proposals are an under-utilized way of obtaining computer resources whenever new ideas arise. Finally, this year's Call for Proposals increased the maximum allocation for a Type C proposal from 800 6n-equivalent node-hours to 4000, although the old limit was never strictly enforced.

Recommendation #11:

"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."

Response: A list of papers describing research that made use of LQCD Computing Project hardware, or of gauge configurations generated with this hardware, can be found on the Review Web Site and on the memory sticks that were distributed to Panel Members. Note that papers reporting on work done with the SciDAC-1 clusters or the QCDOC prior to the start of the LQCD Computing Project are not included.

Recommendation #12:

"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 user 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."

Response: An online user survey was conducted by the LQCD Computing Project in August/September 2007. The survey solicited user feedback in 23 subject areas, including the scientific proposal and allocation process. Responses were received from 54 users out of an estimated user base of 60 individuals.

- 70% of the respondents felt that the allocation process helps maximize scientific output.
- 98% of the respondents felt that the time allowed for proposal preparation was adequate.
- 83% found the Call for Proposals to be adequately clear (i.e., no additional clarification needed).

However, on a scale of 1 to 5, with 1 being very dissatisfied and 5 being very satisfied, the transparency of the allocation process was rated 3.05 (61% satisfaction rating) and the fairness of the allocation process was rated 3.15 (63% satisfaction rating). The survey results clearly indicate that some improvement is needed to improve the transparency and perceived fairness of the allocation process by the user community.

To address this, Andreas Kronfeld, Chair of the Scientific Program Committee, discussed the results of the survey at the 2008 USQCD All-hands Meeting. He shared with the community the free-form comments that were provided by survey respondents. During his presentation, he described the proposal and allocation process in an attempt to address transparency concerns. Following his presentation, he chaired a round-table discussion and encouraged questions and comments regarding the allocation process.

We plan to conduct a follow-up survey in late-summer 2008 to assess the perception of the community to the 2008 allocation process. In the ’08 survey, we will structure questions related to the proposal/allocation process differently, in an attempt to gain more insight into how to better improve transparency and perceived fairness going forward.

Recommendation #13:

"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 identify users who do not use their allocation efficiently."

Response: In order to automatically trace job failures in sufficient detail for categorization, software development is necessary. Such software development is out of scope of this project. However, this task has been undertaken by the cluster reliability subproject of the SciDAC II Lattice QCD Computing project. As of late February 2008, detailed automated tracking of most aspects of operation, including scheduler activity, hardware failure information, and process information, was put in place on the Fermilab Pion and Kaon cluster. The subproject continues to work on the software needed to interpret this recorded data and will begin producing reports by the end of FY08. This software will be ported to the JLab clusters in FY09.

QCDOC does not permit the use of exit codes for this purpose, so quantitatively tracking the rate and nature of job failure rates on QCDOC is complicated. To obtain this information would require modifying the OS, which would be time-consuming and costly. Given the cost of this effort and the nature of jobs run on QCDOC, we do not feel that the benefits of an automated job failure tracking system for QCDOC justify the cost, given the small number of users affected.

Recommendation #14:

a) *"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."*

b) *"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."*

c) *"The resources provided through the LQCD project are crucial for the US lattice QCD community to stay internationally competitive. This will 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 2006-2009 LQCD project."*

Response: The three statements by the 2007 Review Panel quoted above point to the need for greater computational resources than can be provided by the LQCD Computing Project. We agree with all of them, and we have taken steps to meet this need during the past year. As indicated previously, we have obtained a three year grant from the DOE Incite Program for time on its leadership class supercomputers at ANL and ORNL. This grant is already having a major impact on our work.

We have submitted a proposal to the High Energy Physics and Nuclear Physics Programs of the DOE for a follow on to the LQCD Computing Project for the five year period FY 2010 through FY 2014. This proposal requests a hardware budget of $3M per year and an operating budget starting at $1.45M in FY 2010, and increasing by 4% per year. This budget would enable the acquisition of hardware sustaining 35 teraflop/s in FY 2010, growing to 255 teraflop/s in 2014. The proposal has been reviewed by a panel of high energy and nuclear physicists, and computer scientists, which strongly recommended funding it.

We have also submitted a proposal to the NSF for early access to its petascale computer, Blue Waters, which is expected to become available in 2011.

Finally, we note that sub-groups within USQCD have computing resources that, while not available to the entire collaboration, support its overall scientific program. These include the QCDOC at the Riken Brookhaven Research Center, allocations on NSF and NERSC supercomputers, and access the the large BlueGene/L at LLNL.