Lattice QCD Computing Project (LQCD)

Response to Recommendations
from the
2009 Annual Progress Review of the LQCD Computing Project

Compiled by

Paul Mackenzie
Chair, LQCD Executive Committee

William Boroski
Contractor Project Manager, LQCD Computing Project

April 26, 2010

INTRODUCTION

On June 4-5, 2009, 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 Fermi National Accelerator Laboratory and resulted in a written report that contained twelve recommendations to help improve the project's effectiveness and impact. Eleven recommendations were associated with the scientific program and one recommendation was associated with technical aspects of the computing project.

This document summarizes the LQCD Computing Project's response to these recommendations and where appropriate, the actions taken to implement specific recommendations. Neither the report sections nor the recommendations were uniquely numbered in the review report, so we have numbered the sections and recommendations sequentially in the order they appear in the review report, and have adopted the following numbering scheme when tracking resolution:

ReviewReportSectionNumber.RecommendationNumber

RESPONSE TO RECOMMENDATIONS

1.0 Continued Significance and Relevance

Recommendation 1.1

The expanded workshop program with the wider NP and HEP communities has been a valuable development. The impact of these workshops on the USQCD priorities and those of the other communities should be more explicitly identified and reported.

Response: The needs of the HEP and NP experimental programs are the most important drivers of USQCD priorities in our proposals and in year to year allocations of USQCD resources. Meetings involving members of USQCD and of the experimental communities have proven to be a valuable opportunities for improving our understanding of these needs. We have held a series of these in flavor physics, nuclear physics, beyond-the-Standard-Model physics, and the thermodynamics of QCD plasmas. We plan hold more in all these areas. These workshops affect USQCD's priorities by providing us with insight into which of the calculations we do can have the biggest impact on the experimental programs. To help us understand how the workshops affect the priorities of the experimental communities, at each USQCD All-Hands Meeting, we ask an experimentalist from a recent workshop to report on it to the USQCD Collaboration. The most recent of these was the report at the April 16-17, 2010, All-Hands Meeting by Paul Sorenson of the Star Collaboration on the workshop On the interplay between "Lattice QCD calculations and Heavy Ion Experiments: Critical Point and Onset of Deconfinement".

Jochen Dingfelder of BaBar writes to us on the impact of the workshops on BaBar, "I do believe that the combined lattice+experiment workshops are extremely useful... If I should think of some specific examples for the exclusive B->pi l nu analyses,
this is what would come to my mind:

- The use of the z-expansion was recommended to us during a LQCD+BaBar workshop, and we are now using it as default parametrization in BaBar B->pi l nu publications.
- We have received a lot of good feedback on how to use the LQCD predictions correctly in our semileptonic analyses.
- Discussions about correlations of theory errors or between different LQCD calculations are very helpful for averaging of results, e.g. by HFAG.
- IMO these workshops help to gain confidence in the way the LQCD predictions are used by experimentalists and the exp. results by theorists."

Recommendation 1.2:

USQCD, in particular through the multiple instruction multiple data (MIMD) lattice collaboration (MILC), has led the open dissemination of lattice data, with great benefits to the field. This should continue, and USQCD should maintain a public data release policy that builds on and sustains the international lattice data grid (ILDG) agenda.

Response: It has been USQCD practice for some time that groups generating large gauge ensembles share them within the collaboration. The USQCD Executive Committee has proposed that USQCD adopt the policy that all significant ensembles of gauge configurations generated with USQCD resources be made publicly available through the International Lattice Data Grid no later than six months after the first publication in a refereed journal of a paper that makes use of them by the group that generated them. This policy was discussed and agreed to by the collaboration at the 2010 USQCD All-Hands Meeting earlier this month, April 16-17, 2010.

Recommendation 1.3:

USQCD should have a data curation strategy including disaster recovery.

Response: It is USQCD's policy that all significant ensembles of gauge configurations generated with USQCD resources be stored in at least two physically separated locations.

Recommendation 1.4:

Wider exploitation of lattice results beyond the lattice community might be facilitated by making appropriate intermediate results and tools available to phenomenonologists and experimenters. Understanding how to do this effectively is an important challenge for the field, which USQCD should address.

Response: We are delighted to work with groups interested in examining intermediate lattice data. However, as the recommendation implies, it is not straightforward to know how to do this in a way that is useful for the customer. For such efforts to be productive, it is important to closely involve those interested in consuming these intermediate results and using the possible tools. To date, few possible users of such intermediate results have come forward.

One example was a request from Lawrence Gibbons of the CLEO Collaboration to the Fermilab and MILC collaborations. He wished to combine raw lattice theory data for the form factors in the decay B->pi l nu in a global fit with experimental data from CLEO, BaBar and Belle. The intermediate results he hoped to obtain did not exist and would have had to be generated. The raw data that did exist contained, for example, discretization artifacts in chiral extrapolations that needed to be analyzed and removed by lattice experts. Including them in a global fit would have required substantial collaboration with a lattice expert such as the postdoc who did the original lattice analysis. We discussed with him how this might work but did not arrive at a plan before people got busy with other things.

Another example is a request from a group of European phenomenologists including Veronique Bernard, Emilie Passemar and Christoph Haefeli. Their request to the RBC-UKQCD collaboration sought detailed information about the correlations between meson masses and decay constants computed at a variety of valence and sea quark masses. The RBC-UKQCD response was to compute and provide values with errors for the specific products and ratios which Bernard and collaborators requested.

Since our intermediate data are in general in a form different than what experimenters and phenomenologists are used to, we do not know how to find a general way of doing this without working with specific interested customers. When we do have such potential customers, we are delighted to look for ways to be more useful to them.

Recommendation 1.5:

Focusing the major investment of resources on two of several possible lattice actions is necessary, but carries risk. International competitors have made different choices. Long term plans should explicitly manage this risk, e.g. by fostering small scale exploratory work on other actions, recognizing natural break points in configuration production as opportunities for strategic review, and supporting algorithm development with both effort and computing resources.

Response: USQCD makes major investments in the generation of gauge configurations with three formulations of lattice quarks: anisotropic-clover, domain wall and improved staggered. We agree that it is important to foster small scale exploratory work on other actions and on algorithm development and we are doing just that. During the past year, USQCD provided modest resources for an exploratory study of the highly improved staggered quark (HISQ) action. Because this study was very successful, and because our long term project to generation gauge configurations with improved staggered (asqtad) quarks reached a natural break point, we ended the generation of asqtad gauge configurations and started a new project to generate configurations with the HISQ action. Similarly, during the past year, modest USQCD resources where invested in the development of an improved algorithm for generating gauge configurations with domain wall quarks. This algorithm has now been incorporated into our production runs.

Recommendation 1.6:

In cases in which collaboration with experiment and phenomenonology are needed to achieve the final science goals, it would be useful to have plans that outline how the ultimate goals will be reached. The experimental program at the 12-GeV upgrade at TJNAF is a case in point. A plan should address the experimental goals quantitatively, and decide which are achievable in the near term with foreseeable resources, and which will have to wait for another generation of computer resources.

Response: We focus on two components of the physics program of the 12 GeV upgrade where precise lattice calculations will be essential to achieve and to capitalize on its goals, namely to investigate the excited state spectrum of QCD,, and to study hadron structure and the origin of spin in the nucleon.

The primary task for lattice QCD in spectroscopy is to present precise calculations of the low-lying states of the theory that can confront, and ideally predict, anticipated results from experiment. Under the auspices of USQCD, the Hadron Spectrum collaboration has a program of calculations that address that aim. Its work is already yielding results of direct relevance to the experimental program: *a calculation of the low-lying isovector meson spectrum for pion masses down to 400 MeV which demonstrates the existence of states with exotic quantum numbers in a range likely to be accessible to GlueX, and indeed even suggests the existence of non-exotic states with manifest gluonic degrees of freedom.*

A major challenge is to extend calculations to the physical light-quark masses. This will require that unstable states and multi-hadron states be treated. Work is in progress on these problems, exploiting the disruptive computational technology provided by General-Purpose GPUs, and preliminary results obtained for the (unbound) I=2 pi pi phase shift. demonstrating the feasibility of the calculation. The USQCD collaboration is proposing a program of gauge generation that we anticipate will see calculations approaching the physical light-quark masses by the time of the first physics results from JLab@12GeV in FY15.

The USQCD collaboration has performed a series of calculations that address the hadron structure program at JLab and at RHIC spin, including calculations of the electromagnetic form factors, the contribution of moments of parton distributions, and calculations of moments of Generalized Parton Distributions (GPDs). As in the case of the spectroscopy program, USQCD calculations had already impacted the experimental program: a calculation of the orbital angular momentum carried by the u and d quarks in a proton shows that the total orbital angular momentum is small, but that carried by the individual quark flavors is substantial, a result in accord with experimental bounds obtained by DVCS off the neutron at Jefferson Laboratory, and DVCS off the proton at HERMES.

Here lattice QCD faces three challenges if it is to capitalize on the future experimental program: a) the inclusion of disconnected diagrams, enabling access to flavor-singlet hadron structure; b) performing calculations at the physics light-quark masses and c) a program of calculations of the new quantities that will be explored at JLab@12GeV. Work is in progress to address all three of these challenges.

Recommendation 1.7:

The BSM work is a strategically important initiative. As it exploits a growing fraction of the computing resources available, its direction and impact should be reviewed.

Response: We continue to review the progress and goals of Beyond the Standard Model strongly interacting lattice field theory. To guide in choosing initial projects, we initiated regular workshops with representatives from the experimental and model building communities. The first workshop, ``Lattice Gauge Theory for LHC Physics'', was held at LLNL, May 2-3, 2008 (see http://www.yale.edu/LSD/workshop08). A second was held at Boston University, Nov. 6-7, 2009 (see http://www.yale.edu/LSD/workshop/). In addition a three week Aspen 2010 summer workshop on ``Strong Dynamics Beyond the Standard Model" will take place May 23 to June 13. In view of the very wide range of theoretical models proposed for electroweak symmetry breaking and new TeV physics, these exploratory lattice projects are also diverse. However there is a consensus to initially focus in depth on a few of the most promising scenarios, such as the conjectured near conformal (or walking) "technicolor" theories, and on N = 1 super Yang Mills sector of the MSSM. The objective is to begin to make quantitative prediction for spectra and departures from standard model predictions. This gradual expansion in BSM lattice investigations is, we believe, a prudent approach, which in the event of the discovery of new strong dynamics at the LHC, will provide a robust lattice capability to assist in the difficult task of constraining the theoretical interpretation of experimental data.

2.0 Progress Towards Scientific and Technical Milestones

Recommendation 2.1:

LQCD is dependent on expertise and facilities at the national labs. This joint effort has been very successful in enabling efficient exploitation of rapidly evolving technologies, and it should be sustained. However, this association obscures the total cost of facilities. The total cost of procuring and operating LQCD systems, together with their environmental impact, may become more important in the future. Consequently, the project should track these costs in future procurements and operating models.

Response: We agree that the LQCD computing project, and subsequently the USQCD collaboration, benefit from the strong relationship with the national labs hosting LQCD computing facilities. The relationship between the project and the laboratories is defined in formal signed MOUs that describe the procurement and operating costs that will be covered using project funds, and the facility infrastructure costs that will be provided by each respective laboratory. For example, the following excerpt from the Fermilab MOU defines facility and utility costs that will be provided to the project. Similar language is contained in the JLab and BNL MOUs.

> *7.3. Facilities and Equipment*
>
> *Adequate facility infrastructure will be made available to the LQCD project to carry out the implementation and operation of the LQCD computing system at the Fermilab site. Fermilab agrees to pay for all facility and utility costs, such as the power needed to support the computing and HVAC systems.*

Given this relationship, the need for the project to explicitly track procurement support and facility infrastructure costs in terms of dollars is not obvious. At some level, procurement and infrastructure costs are already factored into the overhead rates charged to the project for the use of laboratory services and facilities. Therefore, we feel that it is more effective to forecast and track facility infrastructure costs in units pertinent to the compute facility managers at each site (e.g., ft^2 for space and KW for power and cooling), which helps ensure that LQCD power and space needs are factored in to future facility planning.

To support proper planning and to document facility resource usage, each site manager is responsible for maintaining an inventory of the computing hardware in operation, which includes the power, space and cooling needs for that facility. Each site manager also maintains a historical record of computing hardware that has reached end of life and has been decommissioned. If it does becomes necessary at some point in the future to explicitly denote physical infrastructure costs, one could apply appropriate conversion factors for each site to convert physical units to dollars (e.g., $\$/ft^2$ or $/KW).
The following sections provide more detail regarding the current handling of costs associated with procurement activities, and the type of information maintained by the project to forecast and track computing facility power and space needs.

Procurement Costs

LQCD procurements at each site follow the respective laboratory's procurement policies and procedures, and utilize services provided by the laboratory's purchasing, receiving, and inventory control departments. Following standard practice, the costs related to purchasing, receiving, inventory control, and so forth are recovered via the overhead structure. Since the LQCD project is charged overhead at standard laboratory rates for these services, the total cost of procurement activities is already included in the total project cost. Accounting reports for past procurements explicitly show overhead charges related to procurement activities, and planning activities for future procurements take into account the overhead costs associated with procurement activities.

Quantified Power Needs

Each LQCD site manager maintains records of the power requirements for LQCD compute facilities at their respective sites. For example, the following table summarizes the laboratory power contribution for clusters deployed at FNAL and JLab. Note that the JLab 9q, 10q, 9g, and 10g clusters are not part of the LQCD-ext project, but are a part of the LQCD ARRA project. Note also that the FNAL QCD and Pion machines were decommissioned in March 2010, and that the JLab 6n machine will be decommissioned before the end of 2010.

Table 1. Compute Facility Power Requirements for Clusters at FNAL and JLab

Cluster Name	Date	Node Cnt	Power/Node (W)	Total Compute Nodes Power (KW)	Cooling Power Factor	Total Compute Facility Power (KW)
FNAL Deployments						
QCD	Jun-04	128	147	18.8	1.5	28.2
Pion	Dec-05	520	176	91.5	1.5	137.3
Kaon	Oct-05	600	275	165.0	1.5	247.5
Jpsi	Apr-09	864	300	259.2	1.7	440.6
FNAL Sub-total		2,112		534.5		853.6
JLab Deployments						
6n	Jan-06	260	180	46.8	1.5	70.2
7n	Jun-07	396	300	118.8	1.5	178.2
9q	Jan-10	320	280	89.6	1.5	134.4
9g	Feb-10	65	800	52.0	1.5	78.0
10q (est)	Jun-10	224	280	62.7	1.5	94.1
10g (est)	Sep-10	52	1000	52.0	1.5	78.0
JLab Sub-total		1,317		421.9		632.9

Note: The power totals shown in the column "Total Compute Nodes Power" are for the compute nodes directly, and do not include ancillary items such UPS power loss, nor power used by the A/C system. Multiplying these values by the Cooling Power Factor (CPF) provides a conservative estimate of total power required. The CPF for J-Psi is slightly higher because that computer room is only partially occupied and both the cooling and UPS systems are not as efficient as they will be once additional systems are installed.

In addition to the cluster deployments at FNAL and JLab, power requirements for the 12,288-node QCDOC machine deployed at BNL are as follows:

- The water-cooled QCDOC crates use 11 KW each and there are 12 crates. Estimated power requirement for all crates = 132 KW.
- Additional power is required for the front-end hosts, file servers, air-cooled crates, and other supporting hardware.
- Total laboratory power contribution for QCDOC is of order 200 KW.

Quantified Space Needs

As shown in Table 2, the cluster deployments at FNAL and JLab require approximately 1,620 ft^2 and 1040 ft^2, respectively. Note that the estimated floor space taken by a rack position is ~5 tiles, each measuring 4 ft^2.

At each host site, clusters are sited in available, suitable compute facility space. Clusters at FNAL are housed in three computer rooms; clusters at JLab are housed in a single facility.

Table 2. Floor Space Requirements for Clusters at FNAL and JLab

Cluster Name	Date	Node Cnt	# of Rack Positions	Floor Area (ft^2)
FNAL Deployments				
QCD	Jun-04	128	6	120
Pion	Dec-05	520	22	440
Kaon	Oct-05	600	31	620
Jpsi	Apr-09	864	22	440
FNAL Sub-total		2,112		1,620
JLab Deployments				
6n	Jan-06	260	7	140
7n	Jun-07	396	15	300
9q	Jan-10	320	10	200
9g	Feb-10	65	7	140
10q (est)	Jun-10	224	7	140
10g (est)	Sep-10	52	6	120
JLab Sub-total		1,317		1,040

At BNL, the QCDOC machine requires approximately 100 ft^2 of floor space in the computer room directly. Additional floor space is required in adjacent mechanical areas for supporting equipment such as dedicated heat exchangers for the water-cooled machine.