Offices of High Energy Physics and Nuclear Physics
Report on the

LQCD 2009 Annual Review

June 4-5, 2009

Executive Summary

The Annual Review of the LQCD IT initiative was held on June 4-5, 2009 at Fermilab. The purpose of the review was to assess the collaboration's progress towards its overall scientific and technical goals. Five expert reviewers from the nuclear, high energy physics and computer science communities heard presentations on scientific progress, computer acquisitions and operations, allocation of resources, and dissemination of scientific results. In particular, the LQCD collaboration was instructed to address five charges:

1. 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;
2. The progress toward scientific and technical milestones as presented in the project's IT Exhibit 300;
3. The status of the technical design and proposed technical scope for FY 2008-2009;
4. The feasibility and completeness of the proposed budget and schedule;
5. 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 met or exceeded expectations in almost all cases. The significance and relevance of the LQCD calculations to both the high energy physics and nuclear physics programs have grown dramatically since the LQCD initiative began in 2006. LQCD has exceeded its quantitative technical milestones by several percent over the past twelve months. The review panel endorsed LQCD's benchmarking and procurement policies and remarked that they met the price/performance goals set out by the initiative. These efforts are running slightly ahead of schedule and under budget, which has allowed a preliminary investigation of the use of graphical processing units (GPUs), which may lead to considerable benefits in LQCD applications in the near future. The coordination of the LQCD project with USQCD SciDAC grants was considered a very productive effort. The allocations procedures were seen to be fair and well executed, and the user survey was judged to be very effective. The workshops and other outreach efforts to both the NP and HEP communities were praised, although more effort of this sort was recommended, since lattice calculational methods are still undervalued by broader audiences.

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 Lattice Quantum Chromodynamics (LQCD) IT hardware acquisition and operations activity, which started in 2006 and runs through 2009, has operated a "Quantum Chromodynamics-on-a-chip" (QCDOC) machine at Brookhaven National Laboratory (BNL), and has built and operates special purpose clusters at the Fermi National Accelerator Laboratory (FNAL) and the Thomas Jefferson National Accelerator Facility (TJNAF) with the goal of providing 17.2 Teraflops of sustained computer power for 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 and supercomputer hardware, and the winner of the price-to-performance competition becomes that year's provider. The FY 2009 budget for LQCD is $1.7M in support of hardware and operations. The FY 2009 acquisition has been combined with that of FY 2008 to assemble one large cluster at Fermilab, with a performance goal of 6.2 Teraflops.

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 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.

LQCD proposes to extend its work beyond 2009, and has submitted a proposal for an extended project known as "LQCD-ext" for the period 2010-2014. That follow-on project is under consideration by the Offices of Nuclear Physics and High Energy Physics.

The Annual Physics Project Review of LQCD took place at Fermilab on June 4-5, 2009. 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 LQCD 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 questions that the LQCD collaboration was asked to address.

Continued significance and relevance

Findings

The program supports activities in several research areas:

1) Precision calculations relevant to the determination of standard model parameters from heavy quark processes.

2) Exploratory calculations based on "beyond the standard model" theories, for which LQCD may be the only effective technique for extracting quantitative predictions.

3) Hadronic Physics quantities such as the spectrum of hadrons, form factors, moments of structure functions, hadron-hadron interactions and scattering.

4) Calculations of the properties of QCD at finite temperature and baryon density; this regime is explored experimentally in relativistic heavy ion collisions.

USQCD's scientific goals are focused on carrying out world-leading computations of quantities that are of critical 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 the Relativistic Heavy Ion Collider (RHIC). USQCD activity in QCD thermodynamics has grown through the LQCD initiative, and these results are now among the most highly cited in this field. For example, this work has lead to the world's best result for the equation of state over a large temperature range, with almost physical quark masses. However, there are challenges in connecting the experimental measurements to fundamental QCD physics. The path forward here appears to involve the development of hydrodynamic models; it is clear that much work remains before LQCD simulations can be related to RHIC phenomenology and experiments in a quantitative manner.

USQCD work on hadron spectrum, structure and scattering is also world-leading, and is very well aligned with the NP long range plan. There is a growing recognition that lattice simulations are crucial for meeting NP goals: With regard to the Jefferson Lab Scientific Program and the 12-GeV upgrade, much of the experimental program relies on lattice calculations (for example of exotic meson spectroscopy, photocouplings, and nucleon structure) to interpret the planned measurements. The LQCD effort is now a crucial part of this JLAB upgrade project, and LQCD predictions of hadron properties will likely be of increasing importance as the project develops. However, there are daunting technical challenges in these LQCD simulations, and because of the sensitivity of the results to quark mass and lattice size and resolution, physically relevant simulations are needed with lighter quarks, larger volumes, and higher statistics.

In HEP, USQCD focuses on the determination of quark masses and Cabbibo-Kobayashi-Maskawa (CKM) parameters for precision tests of the Standard Model (SM). This effort has produced many of the best results available today, which largely meet or exceed the original aspirations of the LQCD initiative. Unitarity triangle analyses for example now take these results very seriously. A two-sigma discrepancy in one unitarity triangle constraint has now been resolved through better control of the lattice computations, and an intriguing 3-sigma deviation from experiment for the Ds leptonic decay constant has recently emerged.

USQCD has focused most of its resources on two particular choices of lattice QCD actions, which are 1) staggered, and 2) domain-wall fermions (DWF). (There are also some specialized efforts using anisotropic clover actions.) The resulting ensembles of configurations and quark propagators support many physics projects. Their parameters favor different physics applications, but they are sufficiently similar to allow essential checks that these different formulations ultimately give the same answers for physical quantities.

The USQCD collaboration also plays a leading role in their exploratory work on some of the more popular "beyond the standard model" (BSM) candidate theories; these include technicolor theories of electro-weak symmetry breaking (EWSB), which require non-perturbative dynamics, and investigations of lattice supersymmetry (SUSY). This rapidly growing activity (which currently uses ca. 7% of LQCD resources) is broadening the relevance of lattice simulations to the wider HEP program.

Comments

USQCD is very successfully managing a large integrated program of hardware acquisition and operation, algorithm and software development, data management, and scientific research. This effort is an example of best practice across all computational science disciplines, and this collaboration is doing innovative work on all aspects of this program. Many of their physics results are world-leading.

The LQCD team effort is very impressive. They have managed a large project that spans the acquisition of computer hardware, maintenance and operation of hardware and making it available to users with a very high level of reliability, development of software tools both to take advantage of diverse hardware and to lower the barrier to entry by new people, and finally developing a large and extensive scientific program with the tools needed to allocate resources across the community. All of this appears to have been done in a fashion that exceeded the project goals for computational resources as well as coming in slightly under budget. The LQCD effort is to be commended for its successes.

In the case of the N* program, LQCD chose to calculate the "easy" quantities first; for future N* applications phenomenological input will likely be required. Is there a plan for this? In the case of GlueX, there are a number of results that are needed for the experimental program. The exotic meson photocouplings now being obtained by LQCD are a nice start, but it is unclear if these are sufficiently accurate, and whether the available computing resources are adequate for completion of these calculations.

Recommendations

The expanded workshop program with the wider NP and HEP communities has been a valuable development. The impact of these workshops on USQCD priorities and those of the other communities should be more explicitly identified and reported. It might be helpful to develop and maintain a public physics roadmap. This could list where there is scope for focusing effort, both inside and outside the lattice community, to maximize the impact of lattice simulations on progress in the field, particularly where there is a possibility of discoveries of new physics.

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. Generally, USQCD should have a data curation strategy including disaster recovery. Wider exploitation of lattice results beyond the lattice community might be facilitated by making appropriate intermediate results and tools available to phenomenologists and experimenters. Understanding how to do this effectively is an important challenge for the field, which USQCD should address.

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.

In the cases in which collaboration with experiment and phenomenology 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.

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.

Progress towards scientific and technical milestones

Findings

The LQCD Project has made excellent progress on its scientific goals, and is on track to exceed their original goals in all cases.

The LQCD Project exceeded all performance goals for systems and support in FY 2008, with high customer satisfaction levels. It is on track to achieve elevated goals for FY 2009.

The LQCD hardware procurement strategy is responsive to rapidly evolving technology, and has delivered excellent value for the money. The total Teraflops deployed has exceeded the baseline targets every year, and is on track to do so this year.

Utilization has typically been above 90% for all machines, which is very good.

USQCD is making very efficient use of a wide range of high performance computer (HPC) architectures, and this has stimulated community-based open QCD software development, enabling USQCD to rapidly exploit new systems when they give competitive advantages (e.g. through INCITE). The result has been that USQCD has had access to more than double the machine cycles provided by LQCD.

Comments

The LQCD initiative monitors the usage of its clusters effectively. It sometimes finds large usage drops over short periods of time. The project should investigate the causes of these rare irregularities.

USQCD requires a hardware strategy that reflects the growing architectural complexity of computer systems, particularly the growth of parallelism at chip and system levels. This affects physics analyses as much as configuration generation. It will be increasingly important to integrate hardware with algorithm and software planning and development. Continued SciDAC support for software development is therefore vital to the success of LQCD and subsequent similar projects.

Recommendations

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 the 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.

Technical design and scope for FY08-09

Findings

The procurement process was thorough, and is an example of excellent practice. This process led to a careful evaluation of the options for maximizing delivered computing capability over the 3.5 year lifetime of the LQCD initiative, which was successfully implemented. The resulting system exceeded its target capability.

The FY 2008 and 2009 procurements were combined. A 5.75 Teraflop JPsi cluster (quad-core Opteron) at FNAL started physics production in Jan 2009, which produced additional resources ahead of schedule. The delay of the availability of the Intel Nehalem chipset and the results of benchmarking led to the decision to expand the Opteron system in Apr 2009 by 2.65 Teraflops, using FY 2009 funds. The upgraded system began operation in mid-April. Price performance continued to follow Moore's law, which led to the achievement of 8.4 Teraflops for QCD simulations with light quarks, exceeding the milestone of 6.2 Teraflops.

Four Nvidia quad-GPU systems were purchased using the remaining funds, and were attached to 6 JPsi nodes (to be used for propagator calculations).

Comments

It will be increasingly important to explore new architectures; the procurement of a small GPU system is welcomed.

Budget and schedule

Findings

The LQCD Project is within budget and on schedule. Costs are projected to be 2% under budget at the end of FY 2009.

Staff support levels are appropriate. They were increased slightly following a review of the staffing model; this was accomplished by moving a small amount of funding from hardware.

Recommendations from previous review

Findings

USQCD has carefully considered and responded to the recommendations of the last review. There is some confusion about what actions by USQCD should follow, particularly in cases in which they have inadequate effort levels to pursue particular physics suggestions.

Particular importance is attached to interactions with the wider NP and HEP communities and to the alignment of priorities. USQCD has responded to this challenge, although there is scope for more activity.

USQCD has increased the number of joint workshops with experimentalists and theorists outside the lattice community across all areas of activity. Directions for USQCD are discussed at the annual All Hands meeting, at which the outside workshop coordinators are invited to present. There is already an example (the discrepancy in the Ds leptonic decay constant) in which a physics result was responded to by a reprioritization of LQCD resources.

USQCD's allocation procedure is thorough, open, and fair. It is responsive to new developments, such as BSM physics, and supportive of a wide range of project sizes and durations.

Comments

One of the recommendations from the 2008 review was that Charmonium/Open Charm & GlueX/JLab physics be given more priority. USQCD considered this point, but it was felt that it was difficult to "tell people what to do". Clearly one cannot force people to work on topics which do not interest them, but in cases like this, allocation requests for calculations central to the broader missions of NP and HEP could be given a high priority if resources permit. USQCD should consider a positive feedback system to drive priority calculations.

The transparency of the allocation process was a concern of users in the past, and seems to still be a concern, *albeit* at a lower level. USQCD should consider providing feedback to the proponents of a proposal to explain the reasons for their allocation decision. A "scientific priority" rating based on the overall needs of the field could be implemented within this process.

Lattice QCD has reached a point at which it is delivering results of direct relevance to experiment and to other areas of theory. USQCD has a responsibility for raising awareness of this by presenting its results appropriately, particularly so that error bars and systematics can be understood by non-experts. Broader engagement between the communities should be encouraged, from graduate student training upwards.

Broadening the membership of the USQCD Executive Committee to include influential experimental physicists in order to increase awareness of the lattice calculations and increase their impact in a wider NP and HEP community has been recommended in the past. In itself such an action may not be successful unless particularly effective new members are chosen. USQCD should address the problem of wider dissemination of its results to increase their impact and to involve a larger fraction of the NP and HEP communities in their activities.

APPENDIX A

Department of Energy
Washington, DC 20585

MAY 1 2009

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 Department of Energy (DOE) 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 initiative on June 4-5, 2009, at the Fermi National Accelerator Laboratory (FNAL). A review panel of experts in high energy physics, nuclear physics, project management and computer science is being convened for this task.

Each panel member will evaluate background material on the LQCD initiative and attend all the presentations at the June 4-5th review. The focus of the 2009 LQCD Annual Progress Review will be on understanding:

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

- The progress toward scientific and technical milestones as presented in the initiative's Information Technology Exhibit 300;

- The status of the technical design and proposed technical scope for FY 2009;

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

- 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 initiative and write an individual report on his/her findings. These reports will be due to 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 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.

Please designate a contact person at FNAL 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, 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 FNAL.

Sincerely,

for Dennis Kovar
Associate Director of Science
for High Energy Physics

Eugene A. Henry
Acting Associate Director of Science
for Nuclear Physics

APPENDIX B

List of Reviewers

Richard Kenway (University of Edinburgh) r.d.kenway@ed.ac.uk

Stephen L. Scott (Oak Ridge National Lab) scottsl@ornl.gov

Eric Swanson (Pittsburgh) swansone@pitt.edu

Hassan Jawahery (University of Maryland) jawahery@physics.umd.edu

Curtis Meyer (Carnegie Mellon University) cmeyer@cmu.edu

List of DOE program managers

T. Barnes (NP)

List of observers

Robert Lindsay (ASCR)

Jay Stroup and Eric Linderman (EA/CPIC manager)

APPENDIX C

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

June 4-5, 2009

Fermi National Accelerator Laboratory
One East Conference Room, Wilson Hall

Agenda

June 4, 2009

08:30 - Executive session (30 min)

09:00 - Logistics and Introductions (10 min) - *Bill Boroski*

09:10 - Welcome (10 min) - *Vicky White*

09:20 - Project Overview (45 min) - *Paul Mackenzie*

10:05 – Fundamental Parameters of the Standard Model (40 min) - *Andreas Kronfeld*

10:45 - Break

11:00 - Hadrôn Spectroscopy, Structure and Interactions (40 min) - *John Negele*

11:40 - High Temperature/Density QCD (40 min) - *Frithjof Karsch*

12:20 - Lunch

1:00 - Beyond the Standard Model Physics (40 min) – *Simon Catterall*

1:40 - Project Management and Performance (60 min) - *Bill Boroski*

2:40 - Break

2:55 - Technical Design and Proposed Technical Scope for FY2009 (40 min) - *Don Holmgren*

3:35 - Responses to Scientific Recommendations from the 2008 Review (30 min) - *Paul Mackenzie*

4:05 - Responses to Technical Recommendations from the 2008 Review (20 min) - *Bill Boroski*

4:25 - Executive Session

5:30 - Committee request for additional information - *Committee/Project Leadership*

6:00 - Adjourn

6:30 - Dinner

June 5, 2009

08:30 - Executive session (30 min)

09:00 - Committee questions and discussion (60 min)

10:00 - Break

10:10 - Executive Session and Preliminary Report Writing (2 hrs)

12:00 - Lunch

1:00 - Closeout

2:00 - Adjourn