# Computational Resources for Lattice QCD

# 2010–2014

## Lattice QCD Executive Committee

R. Brower, (Boston U.) N. Christ (Columbia U.), M. Creutz (BNL),
P. Mackenzie (Fermilab), J. Negele (MIT), C. Rebbi (Boston U.),
D. Richards (JLab), S. Sharpe (U. Washington), and R. Sugar (UCSB)

## I. Project Motivation

In this paper we propose a Lattice QCD Computing Project covering the period FY 2010–2014. This project is a follow on to the current project which ends in FY 2009. Over the past five years the Office of Science has funded dedicated computers for the study of quantum chromodynamics (QCD): clusters optimized for QCD through the SciDAC-1 Program and the current Lattice QCD Computing Project, and the specially designed QCDOC as a stand alone project. These dedicated computers are available for use by members of the USQCD Collaboration, which consists of nearly all the high energy and nuclear physicists in the United States engaged in the numerical study of QCD. They have enabled major progress in the study of QCD, and brought the field to a point where it is now providing accurate determinations of a wide range of quantities of importance to the Office of Science's experimental programs in high energy and nuclear physics. We propose a mix of access to the DOE's leadership class computers and acquisition of dedicated machines in order to maximize the science output and minimize the cost.

Enormous progress has been made in Lattice QCD over the last five years through the use of improved formulations of QCD on the lattice, which reduce systematic errors due to finite lattice spacing artifacts and chiral symmetry breaking; the development of new algorithms, which reduce the number of floating point operations needed for some studies by factors as large as six; the development of software under the DOE SciDAC Program, which enables effective use of a wide range of existing computers, and rapid porting of efficient code to new ones; and the greatly increased power of computers available for the study of QCD. Some key quantities have been calculated with uncertainties as small as a few percent, and many others of great importance to experimental programs in high energy and nuclear physics can be determined to similar accuracies given sufficient computing resources. Below we briefly describe areas that we believe will be strongly impacted by the proposed Project. White papers providing more detailed discussions of each of these areas are in preparation. We will not be able to complete all of the calculations set out below within the period of the proposed Project. Scientific priorities will be set each year by the USQCD Scientific Program Committee through its responses to proposals for the use of USQCD computational resources. This is the process which is being used in the current Project.

**Fundamental parameters of the Standard Model:** One of the central aims of calculations using lattice QCD is to determine the underlying parameters of the Standard Model by stripping away the effects of the strong interactions. In particular, one would like accurate determinations of the quark masses, the strong coupling constant $\alpha_S$, and the values of the weak transition couplings between quarks—i.e. the elements of the Cabibbo-Kobayashi-Maskawa (CKM) matrix. These quantities, along with the unknown Higgs mass and coupling, and the well known electroweak coupling and mixing angle, are the parameters of the $SU(3) \times SU(2) \times U(1)$ Lagrangian that defines the Standard Model. Particularly exciting is the possibility of determining different, inconsistent values of the CKM matrix elements from different decay processes. This would indicate a breakdown in the Standard Model and thus the need for new physics. This approach is complementary to the direct discovery searches to be undertaken at the Large Hadron Collider (LHC) at CERN. To be successful it requires reliable and precise lattice QCD calculations.

The last five years have seen lattice QCD calculations mature to the point that accurate determinations of some of the fundamental parameters are possible at the percent level with all errors controlled. This has allowed stringent tests of the methodology by comparing precision results to experimental values, e.g. those for $f_\pi$, $3M_\Xi - M_N$, and the masses of heavy-light and heavy-heavy mesons. Another important test is that the lattice determination of $\alpha_S$, which has an error smaller than 1%, and agrees with that determined from high-energy experiments. The up, down and strange quark masses have been determined to an accuracy of better than 10%, providing input for models of physics beyond the Standard Model. There have also been successful *pre*-dictions of the mass of the $B_c$ meson, the leptonic decay constant of the $D^+$ and $D_s$ mesons, and the shape and normalization of the $D \to K$ semileptonic form factor.

Continued progress in lattice calculations will allow precision calculations of the more complicated matrix elements needed to thoroughly probe the CKM sector of the Standard Model. Core examples are the matrix elements of the four-fermion operators that determine the rates of $K - \overline{K}$ mixing (with CP-violation), $B - \overline{B}$ and $B_s - \overline{B}_s$ mixing. Experimental measurements of these rates, coupled with knowledge of the CKM matrix from recent advances at B-factories, allow one to predict what the matrix elements should be if the Standard Model is correct, with a precision varying from 5-12%. Present lattice estimates are somewhat less precise, and do not have all errors controlled. Computational resources at the 100 Tflop-year level will allow control of all errors for these matrix elements, at an expected accuracy of less than 5%. (One Tflop-year is the number of floating point operations a computer sustaining one teraflop/s on production code will produce in one year). Further increases in resources will allow the errors to be systematically reduced, and the calculations extended to a wide variety of other matrix elements.

A significant portion of the error in these envisaged calculations comes from the need to do an extrapolation in the light quark mass. When resources approach the 500 Tflop-year level, calculations using physical light quark masses will be possible, obviating the need for this extrapolation. This will mark a major milestone in lattice simulations. Another important milestone is the ability to calculate quark-disconnected diagrams with sufficient precision to determine the rate of $K \to \pi\pi$ and related decays from first principles. It is harder to estimate the resources needed to attain this goal, but it is likely at the 500 Tflop-year level as well.

**QCD Thermodynamics:** The properties of strongly interacting matter at nonzero temperature and baryon number density are being studied in heavy ion experiments at the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory (BNL). In the near future these experiments will be extended to even higher energies and temperatures at the Large Hadron Collider (LHC). By contrast, at BNL and at the future European heavy ion facility FAIR, a series of new low energy experiments is planned that will allow us to study such matter at moderate temperatures, but high baryon number density. The former physical conditions occurred in the early universe; the latter may approximate the environment in the interior of dense stellar objects such as neutron stars.

Under extreme conditions of high temperature or high baryon number density strongly interacting matter is expected to have a rich phase structure. Quantifying the drastic changes in the interaction among elementary particles that go along with such phase changes requires large scale numerical calculations.

Numerical studies of lattice QCD can provide a wealth of new information about properties of strongly interacting matter. Lattice QCD is likely to have a particularly strong impact on current and future experimental studies, as well as the phenomenological modeling of hot and dense matter, in the following three areas:

- Lattice calculations can provide detailed information about basic bulk thermodynamic properties: the equation of state, energy and entropy density, the pressure, and the velocity of sound and basic structural properties: plasma modes and transport coefficients. These quantities are crucial input to the analysis of many experimental observables that characterize the formation of hot and dense matter in heavy ion collisions, and they are crucial for the hydrodynamic modeling of its time evolution. For example a precise knowledge of the equation of state is needed for a quantitative description of the expansion process and in the theoretical modeling of almost all experimental observables. Lattice methods for determining the equation of state (EoS) are well developed, but numerically intensive. Our present knowledge of the continuum EoS comes with statistical errors of order 15% and probably comparable systematic errors. A combined error of order 5% would provide a solid foundation for hydrodynamical modeling, and is easily feasible with petascale computing resources.

- Lattice calculations currently provide the only *ab initio*, quantitative method for determining the phase diagram of strongly interacting matter, which, aside from the case of vanishing baryon number density (vanishing quark chemical potential), is largely unexplored. In particular, confirming the existence of a second order phase transition point in the phase diagram and subsequently determining its location accurately can only be achieved through demanding numerical calculations. Experiments at RHIC and FAIR are under consideration that would search for this critical point. Quantitative predictions from lattice calculations are needed, and will require computing resources totaling hundreds of Tflop-years.

- Lattice simulations of strongly interacting matter are limited to thermodynamic equilibrium and small deviations from it. Effective models help us develop insight and extend our understanding of the dynamical processes occurring in heavy ion collisions. Lattice calculations are essential for validating and constraining a variety of models ranging from hadronic resonance gas models at low temperature and quasi-particle models at high temperature to perturbative approaches at very high temperature. These studies too will require hundreds of Tflop-years.

**Hadron Structure, Spectroscopy, and Interactions:** Understanding how the structure, spectroscopy, and interactions of hadrons emerge from QCD is one of the central challenges of contemporary nuclear physics. The advances in lattice QCD over the last five years have led to its emergence as a powerful quantitative tool for understanding problems in nuclear physics, and making the precise calculations that the experimental program demands.

The internal quark and gluon structure of the nucleon is a defining component of hadronic physics just as the structure of the hydrogen atom is of atomic physics. Hence, one goal of lattice QCD is precision calculation of fundamental experimental quantities characterizing the nucleon, such as the electromagnetic form factors, moments of parton densities and moments of Generalized Parton Distributions (GPDs). A detailed knowledge of the meson and baryon spectra from first principles will distill the key degrees of freedom needed to describe the bound states of the theory. There is currently an intense experimental effort to determine baryon resonances; the GlueX Collaboration's quest to produce exotic mesons is a flagship component of the 12 GeV upgrade at Jefferson Laboratory. These programs demands a commensurate effort to predict and understand the hadronic spectrum from first-principle calculations using lattice QCD. Rigorously computing the properties and interactions of nuclei remains a major challenge. We have a precise phenomenology of the strong interaction in the non-perturbative regime, but little understanding of how this arises from QCD. Lattice QCD will be vital to gaining this understanding. Finally, lattice QCD will be essential in realizing many of the DOE Milestones in Hadron Physics.

Accomplishments over the last five years have laid the ground work for the ambitious program below. Salient achievements in hadron structure include delineating the contributions of quark spin and orbital angular momentum to the spin of the nucleon, and the calculation of the nucleon axial vector charge to within 7%. Spectroscopy has witnessed the use of correlation-matrix techniques to determine the spectrum of pure Yang-Mills glueballs, and the efficacy of applying the method to states containing quarks has been established. Meson-meson scattering lengths with pion masses below 300 MeV have been calculated, the first prediction of $K\pi$ scattering lengths made, and the first full-QCD studies of nucleon-nucleon scattering lengths performed.

As in the case of determination of the fundamental parameters of the Standard Model, continued progress will exploit our ability to perform calculations closer to, and eventually at, the physical quark masses. Fully chiral studies of hadron structure, at a lattice spacing around 0.1 fm and pion masses down to 180 MeV, will yield errors on key observables down to about 5%. Such calculations would require 100 Tflop-years. Sufficiently high statistics, requiring perhaps several 100 Tflop-years, will enable the demanding disconnected diagrams to be computed, allowing the contributions of the different quark flavors to be delineated, and exposing the role of gluons. Finally, resources of around one Petaflops-year will enable milestone calculations at the physical light-quark masses to be performed.

Calculations of the baryon and meson resonance spectrum using anisotropic clover lattices at pion masses as low as 180 MeV, will require around 100 Tflop-years. Calculations of the properties of resonances will both provide insight into their structure, and inform the expected production rates at GlueX and other experiments. The achievement of calculations at the physical light-quark masses using anisotropic lattices will require around 300 Tflop-years.

In the study of hadron interactions, important goals will be high-precision calculations of the $\pi\pi$, $K\pi$, $KK$, $NN$ and $YN$ scattering lengths. These calculations will employ lattices generated for other projects, and require around 100 Tflop-years of additional computing resources. A precise computation of the deuteron binding energy and other properties of the deuteron at the physical pion mass may well be possible with 1 Petaflops-year, a remarkable achievement critical for satisfying the Department of Energy's milestone for 2014.

**New directions in the LHC era:** The LHC era is likely to reveal new non-perturbative physics beyond the QCD sector of the Standard Model. Theorists have proposed a wide variety of possible scenarios for this new physics. To understand the options and the experimental signatures that will discriminate among them is likely to require non-perturbative investigations of lattice field theories. A similar exploratory approach has proven very useful in lower dimensional condensed matter systems, for which the computational requirements are much smaller. Fortunately, the LHC era coincides with increasing access to petascale computers, so this approach will be possible for four-dimensional quantum field theories. There are three main scenarios envisioned for physics in the TeV energy region:

- **Standard Higgs formulations of the Standard Model:** The first scenario is the discovery of the Standard Model Higgs with little hint of its origin. This is, perhaps, the least exciting scenario. If it turns out to be the case, lattice QCD will continue to play a central role in high precision tests of the Standard Model, as discussed above. Related topics that have already and will continue to receive attention include bounds on the mass of the Higgs boson, hadronic corrections to proton decay, and the connection between electric dipole moments and strong CP violation. Additional topics include a quantitative treatment of symmetry restoration in the early universe, and whether the coupling of Higgs to the top quark will involve appreciable non-perturbative physics.
- **Supersymmetric quantum field theories:** The second scenario involves the discovery of supersymmetry with its attendant zoo of new particles. In this case the need for lattice field theory to incorporate supersymmetry, and to investigate its corresponding breaking pattern and vacuum structure will become paramount. Placing supersymmetric field theories on the lattice is not trivial because the full symmetry is an extension of the Poincare group, which is broken by the lattice itself. However, there are several lattice methods that recover supersymmetry in the continuum limit. Supersymmetric field theories also play a pivotal theoretical role in understanding the AdS/CFT duality which has stimulated a broad range of new approaches to the dual string formulation of QCD, as well as extra dimensional model building.
- **New strong dynamics:** The third scenario is the discovery of a new strong dynamics. This would be an ideal result for lattice field theory. The Higgs may well be most cleanly described as a composite arising in a new strongly coupled gauge field theory. Examples of such a theory are technicolor, Higgsless models, extra-dimensional (Randall-Sundrum) models. In these models the lattice would also be useful to precisely compute electroweak variables, such as the S and T parameters. At present the only other non-perturbative tools are qualitative in nature based on the AdS/CFT paradigm.

Regardless of which of these scenarios plays out, it is important to emphasize that the investigation of quantum field theory using lattice methods can enable one to explore fundamental issues well beyond the TeV range accessible at LHC energies. Indeed the lattice approach has already established a distinguished record in this regard. For example in understanding confinement and spontaneous chiral symmetry breaking in QCD. Future topics should include the strong/weak coupling duality of the Maldacena AdS/CFT conjecture, model building methods of deconstruction from high dimensions, triviality and ultraviolet completion, the large $N_c$ limit of Yang Mills theory (including QCD), and matrix model reductions. Particle physicists are only beginning to gain a deeper appreciation of the non-perturbative complexities of relativistic quantum field theory. Lattice simulations will inevitably continue to play a major role in this broad enterprise.

## II. Project Plan

We are convinced that during the period of the proposed Project, FY 2010–2014, a mixture of access to the DOE's leadership class computers and acquisition of computers dedicated to the study of QCD will maximize scientific output and minimize cost.

Lattice QCD calculations proceed in two steps. In the first one performs Monte Carlo calculations to generate gauge configurations, which are representative samples of the QCD ground state. These configurations are stored, and in the second step they are used to calculate a wide variety of physical quantities. Configuration generation is the most computationally intensive part of our work, and in most cases limits the rate of progress. Since it involves a Markov chain, it must be carried out in a small number of streams. It is thus highly desirable to bring to bear the most powerful available computers on individual calculations. For this reason, lattice generation is best done on leadership class computers. On the other hand, a single analysis project typically requires significantly fewer floating point operations than were needed to generate the gauge configurations with which it is run. The analysis of different configurations are independent of each other, and can be run in parallel. For this reason, many, but not all, of our analysis campaigns can be run on machines that are less powerful than the leadership class ones. We propose to acquire dedicated computers, such as commodity clusters, for this phase of our work because they are more flexible and better lend themselves to the size of many analysis calculations than the more powerful leadership class machines. We believe that we will do the most science by using the computers that are best suited for each phase of our work.

All gauge configurations generated with USQCD resources are made available to the entire collaboration for analysis work, and we are working with physicists in Europe, Japan and the United Kingdom to share configurations on an international scale. By sharing these very valuable resources we are advancing the field at a much more rapid rate than would otherwise occur. As more realistic configurations have become publicly available, the fraction of our resources that go into analysis has grown. A few years ago, the bulk of the resources in any lattice QCD calculation went into configuration generation, whereas in FY 2006, approximately half the USQCD allocations went to analysis projects. The trend towards using a greater fraction of computing resources on analysis projects is again visible in the FY 2007 proposals for use of USQCD computing resources. There are several reasons for this shift. The more realistic gauge configurations now being produced enable calculations that could not previously be undertaken, and have in some cases yielded results of of unprecedented accuracy. The potential for further exciting progress has encouraged ever more ambitious projects. Thus, although no single analysis project may exceed the cost of generating the configurations it makes use of, the large number of projects undertaken with each configuration ensemble will certainly keep the cost of analysis at or above that of configuration generation. In drawing up our plans for the proposed Project we have projected that analysis will require more computing resources than lattice generation.

The history of our field indicates that at least as much is gained by advances in algorithms as by advances in hardware technology, and we expect this trend to continue. Indeed, in the last few years the introduction of the Rational Hybrid Monte Carlo (RHMC) algorithm has reduced the number of floating point operations needed to generate gauge configurations with light quarks by factors of four to six compared to algorithms in use a few years ago. The gauge configurations we expect to generate during the course of the Project could not be produced by the proposed resources in any reasonable amount of time without the RHMC algorithm. Similarly, under the Office of Science's SciDAC-1 Program we developed a QCD Applications Programing Interface (QCD API) which enables members of our field to quickly develop highly efficient and portable code for the study of QCD. Under a SciDAC-2 grant we are extending the QCD API, and optimizing it for a wider class of architectures and processors. We will therefore be ready to use new leadership class and dedicated computers effectively as they come on line. We are confident that our continued work on algorithms and software will significantly enhance our productivity during the proposed Project, extending the range of science we will be able to do.

Over the last few years members of USQCD have been allocated approximately 10% of the cycles on the NERSC IBM SP, Seaborg. In the current year our collaboration has an allocation through the INCITE Program that amounts to approximately 10% of the cycles on the Oak Ridge National Laboratory (ORNL) Cray XT4, Jaguar. We request that our collaboration receive at least 10% of the cycles on the leadership class computers at Argonne National Laboratory (ANL) and ORNL. To see what this might mean in terms of scientific output, we note that the President's FY 2008 budget calls for the installation of a leadership machine at ORNL with a peak performance of 1,000 teraflop/s, and another at ANL with a peak performance of 250 to 500 teraflop/s by the end of FY 2008, providing approximately 1,375 teraflop/s (peak) for use in FY 2009. (We take the mid-point in the range of the ANL machine in making the estimates below). Although we are aware that the Office of Science has more ambitious plans for expanding its leadership class facilities, we make the conservative assumption that they will grow in accordance with Moore's law between in FY 2010–2014. Our full configuration generation codes currently sustains 15% of peak on these platforms. We believe that this performance will improve in the future, but use it in the following estimates. Under these assumption we would obtain a throughput from leadership class machines of 33 Tflop-Years in FY-2010 growing to 208 Tflop-Years in FY 2014. We believe that our assumptions regarding the leadership class machines are quite conservative. If they increase in performance more rapidly than we have projected, or if we obtain more than 10% of the time allocated on them, then our research will accelerate proportionately.

Our proposed process for acquiring dedicated computers follows that of the current Lattice QCD Computing Project. That is, each year we will acquire the hardware that best advances our science. As in the current Project, we propose to locate the hardware at Brookhaven National Laboratory (BNL), Fermi National Accelerator Laboratory (FNAL) and Thomas Jefferson National Accelerator Facility (JLab). In the first two years of the current project, commodity clusters optimized for the study of QCD have been the hardware of choice, and we base our price-performance estimates on them. However, if in any year a commercial machine or a specially designed one should prove more cost-effective than clusters, we would obtain it and accelerate our scientific program with the resulting gain in performance. Since we began building clusters under the SciDAC-1 Program, their price-performance has improved in accordance with Moore's law with a 1.5 year doubling time. This is illustrated in Fig. 1, where we show the price-performance of the clusters we have built as a function of time on a log plot. Performance is measured as the average of that sustained by the sparse matrix inversion routines for the Domain Wall and Improved Staggered quark actions. The latest point on this curve is for the cluster being installed at JLab this spring, and corresponds to a price-performance of $0.60 per sustained Mflop/s.

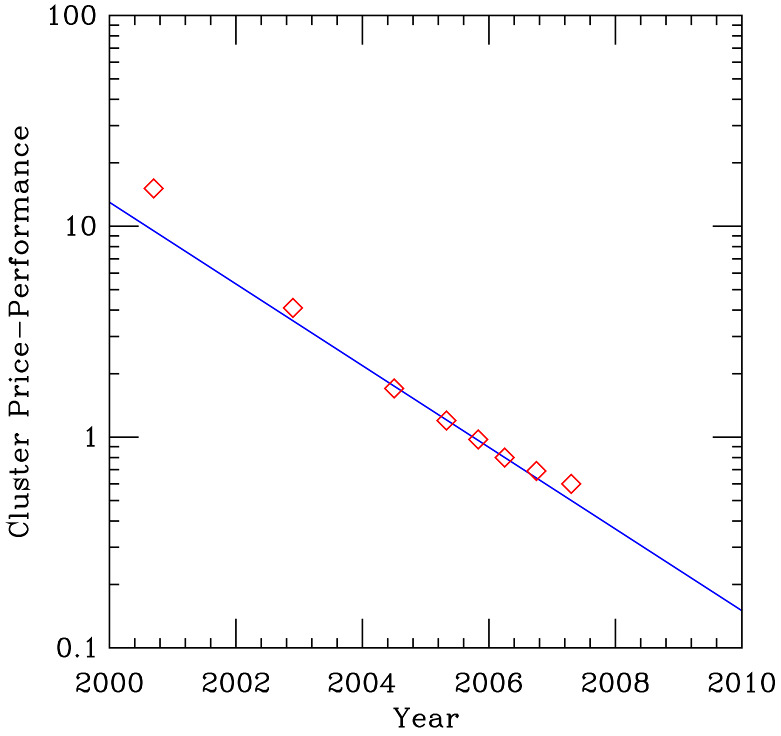

Figure 1: The price-performance in dollars per sustained Mflop/s of clusters built under the SciDAC-1 Program and the Lattice QCD Computing Project as a function of time. The red diamonds are the price-performance of individual clusters, and the blue line a fit to the data.

Assuming that Moore's law continues to hold, the price-performance for clusters will reach $0.15 per sustained Mflop/s in 2010 and $0.024 in 2014. Of course the extrapolation to 2014 has considerable uncertainty attached to it. We have found that the useful lifetime of a cluster is three years, so in 2010 the FY 2007-2009 clusters (or commercial machines) acquired in the current project will still be in service. They are planned to have a total throughput of 9.4 teraflop/s. We also expect the QCDOC, which sustains 4.2 teraflop/s under the same measure as used for clusters, to remain in service through FY 2011. With these assumptions, we show in Table 1 the total throughput for the dedicated hardware in each fiscal year, once that year's installation is completed. Here we assume a fixed hardware budget of $3.0 million per year, the growth in computing power results entirely from Moore's law. We also show in this Table the throughput from the leadership class machines expected throughout that year. Under this plan the installation of dedicated hardware will lag the increase in computing power of the leadership class machines by of order six months, which is reasonable since time is needed to generate new gauge configurations before analysis on them begins.

| Fiscal Year | Dedicated Hardware (Teraflop/s) | Leadership Class Machines (Teraflop/s) |
|---|---|---|
| 2010 | 34 | 33 |
| 2011 | 61 | 52 |
| 2012 | 100 | 82 |
| 2013 | 161 | 131 |
| 2014 | 256 | 208 |

Table 1: Expected throughput for dedicated hardware (column 2) and leadership class computers by fiscal year with a hardware budget of $3.0 million per year.

In FY 2007 we will spend approximately $1.5 million for hardware and $1.0 million for operations. An increase in the hardware budget of the amount we propose would require an increase in the operating budget to $1.3 million in FY 2007 dollars. Assuming a 4% per year cost of living increase for operations would result in the yearly budgets shown in Table 2.

| Fiscal Year | Hardware Budget | Operations Budget | Total Budget |
|---|---|---|---|
| 2010 | 3.00 | 1.46 | 4.46 |
| 2011 | 3.00 | 1.52 | 4.52 |
| 2012 | 3.00 | 1.58 | 4.58 |
| 2013 | 3.00 | 1.64 | 4.64 |
| 2014 | 3.00 | 1.71 | 4.71 |
| Totals | 15.00 | 7.91 | 22.91 |

Table 2: Budgets for the proposed Project in millions of dollars.

It will be noted that even the very significant computing power set out in Table 1 will not be sufficient to achieve all of the science goals discussed in Section I during the course of the proposed Project. Indeed, to do so will require a total of several thousand Tflop-Years. It is therefore important to have a process for setting scientific priorities that will enable us to make optimal use of the available resources. We propose to use the process in place for the current Lattice QCD Computing Project. Each year the USQCD Scientific Program Committee issues a call for proposals for use of the collaboration's dedicated computers. It also invites proposals for projects that would be appropriate for use of the DOE's leadership class computers. On the basis of these proposals, the Scientific Program Committee draws up a draft plan for the coming year. This plan is discussed at the yearly All Hands' meeting of the USQCD Collaboration. Following that discussion the Scientific Program Committee allocates time on USQCD resources, and transmits to the Lattice QCD Executive Committee the priorities of the Collaboration for use of the Leadership Class machines. It is the responsibility of the Executive Committee to submit proposals for the use of these computers on behalf of the Collaboration.

We have not tried to project scientific opportunities or computational requirements beyond 2014, which is, after all, seven years away. However, the numerical study of QCD is an ongoing subject which will not come to an end in FY 2014. In addition to completing the calculations discussed in Section I, there will undoubtedly be exciting opportunities in FY 2015 and beyond that will require both access to the most powerful leadership class computers and the acquisition of dedicated machines. For example, new experimental data from the Large Hadron Collider (LHC) at CERN is likely to bring to the fore the study of strongly coupled field theories that go beyond the Standard Model, but will be susceptible to study by the techniques that have been developed for QCD. So, we propose that dedicated computers for the study of lattice field theory be considered an ongoing component of the Office of Science budgets.

## III. International Efforts

Like many fields of science, Lattice QCD is an international activity with very strong programs in Europe, Japan and the United Kingdom. The Lattice QCD Computing Project, the funding of the QCDOC and the two SciDAC grants have enabled physicists in the United States to maintain a leadership position in the field. However, physicists in other countries also recognize the scientific opportunities that have arisen in Lattice QCD, and are moving aggressively to obtain the computing resources necessary to capitalize on them. Table 3 shows estimates of the computing resources currently available for the study of Lattice QCD in the countries that are major participants in the field. The estimates for other countries were obtained by making inquiries of senior physicists in each of them. Approximately half of the United States resources labeled National Centers come from an allocation to the USQCD Collaboration by the DOE's INCITE Program, while the other half come from allocations to individuals or groups by NSF Centers and NERSC. Two computers located in the United States, but not available to the full USQCD Collaboration, are not shown in the table. One is a QCDOC located at the Riken Brookhaven Research Center, which was funded by the Riken Institue of Japan, and is used jointly by some Japanese and United States physicists. The other is a one rack IBM BlueGene/L located at MIT, which is used partly for software development and partly for studies of Lattice QCD by MIT physicists. It is clear that without the Lattice QCD Computing Project the United States would have fallen significantly behind Japan and Germany.

| Country | Sustained Teraflop/s |
|---|---|
| Germany | 10–15 |
| Italy | 5 |
| Japan | 14–18 |
| United Kingdom | 4–5 |
| Unites States | |
| LQCD Project | 9 |
| National Centers | 2 |
| US Total | 11 |

Table 3: Computing resources in sustained teraflop/s estimated to be available for the study of Lattice QCD in various countries, as of February, 2007.