

Overview of the Lattice QCD Computing Project

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)

1 Introduction

The mission of the Lattice QCD (LQCD) Computing Project is to acquire and operate dedicated computing hardware for the study of quantum chromodynamics (QCD). Hardware is located at Brookhaven National Laboratory (BNL), Fermi National Accelerator Laboratory (FNAL) and the Thomas Jefferson National Accelerator Facility (JLab). The Project began in fiscal year (FY) 2006, and is scheduled to run through FY 2009. It is funded jointly by the Department of Energy's Offices of High Energy Physics and Nuclear Physics with a total four year budget estimated to be \$9,200,000. The budgets for the first two years of the project, and the estimated budgets for the last two years are shown in Table 1 broken down into funds for hardware acquisition and operations.

Year	FY2006	FY2007	FY2008	FY2009	Total
Hardware	\$1,850	\$1,592	\$1,630	\$798	\$5,870
Operations	\$650	\$908	\$870	\$902	\$3,330
Total	\$2,500	\$2,500	\$2,500	\$1,700	\$9,200

Table 1: Budget for the LQCD Computing Project by fiscal year in thousands of dollars.

This project is the key element in the effort of the USQCD Collaboration to build the computational infrastructure needed for the study of QCD. USQCD consists of nearly all the high energy and nuclear physicists in the United States involved in numerical studies of QCD and other strongly coupled field theories relevant to the understanding of fundamental physics. The Collaboration was formed in 1999 with the goal of developing the computational resources needed for these studies. It is open to all physicists in the United States, and the computational infrastructure it develops is available to all of its members. In addition to sharing the infrastructure, members of USQCD also share the computationally expensive gauge configurations generated with it. However, the physics calculations performed with the gauge configurations are carried out by groups or individuals within USQCD. In this way we attempt to maximize scientific output by giving free rein to the different approaches and expertise of individuals within the Collaboration, while minimizing duplication of effort.

To date, the USQCD Collaboration has undertaken three major infrastructure projects in addition to the LQCD Computing Project: 1) the construction of the QCDOC, a computer specially designed

for the study of QCD by a group of physicists centered at Columbia University; 2) the design and construction of commodity clusters optimized for the study of QCD under a grant from the DOE's SciDAC-1 Program; and 3) the development of community software, under grants from the DOE's SciDAC-1 and SciDAC-2 Programs, that enables members of USQCD to make effective use of a wide range of high performance computers including the QCDOC, commodity clusters and the DOE's Leadership Class computers. The Lattice QCD Computing Project is the subject of the current review, but the construction of the QCDOC and the SciDAC-1 clusters, and the software effort are not. However, the QCDOC and SciDAC-1 clusters are operated under the Project, and the SciDAC software enhances the productivity of the Project's hardware.

In Section 2 of this document we set out the scientific goals that are supported by the LQCD Computing Project, and describe progress towards reaching them. Then, in Section 3 we provide an overview of the existing hardware and that which is planned for the immediate future, and in Section 4 we describe the management structure of the project. Section 5 contains a description of the SciDAC software effort, and Section 6 a brief description of international efforts to obtain computing resources for the study of Lattice QCD. More detailed information regarding the material in Sections II-IV is contained in other documents on the review website.

2 Scientific Goals and Accomplishments

Enormous progress has been made in Lattice QCD over the last five years, enabled in large part by the dedicated computing facilities provided to USQCD by the DOE. Other important factors in this progress have included 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; and 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. 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. There have been several successful predictions which were later confirmed by experiment.

Up to now, work within the USQCD Collaboration has focused on three areas which we believe have the greatest potential for large and immediate impact on the experimental programs in high energy and nuclear physics: the determination of fundamental parameters of the Standard Model; QCD thermodynamics; and hadron structure, spectroscopy and interactions. Each of these areas contains important problems that can be addressed with the computers provided by the LQCD Computing Project, and each also contains important problems that will require computers with throughputs of hundreds or even thousands of teraflop/s. Recently, some exploratory projects have been initiated to study lattice field theories that may be relevant to physics beyond the Standard Model. This area is much less advanced than the more mature studies of QCD, but with the Large Hadron Collider (LHC) at CERN about to come on line, we believe it is important to have such exploratory studies underway. Below we provide an overview of recent progress and opportunities in each of these areas. White papers providing more much detailed discussions are available on the review website. The white papers consider the period between the present and 2014. They there-

fore discuss some projects that will require significantly greater computer power than is provided by the current LQCD Computing 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 α_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 LHC. 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_π , $3M_\Xi - M_N$, and the masses of heavy-light and heavy-heavy mesons. Another important test is that the lattice determination of α_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 predictions of the mass of the B_c meson, the leptonic decay constants of the D^+ and D_s mesons, and the shape and normalization of the $D \rightarrow K$ 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 $K - \bar{K}$ mixing (with CP-violation), $B - \bar{B}$ and $B_s - \bar{B}_s$ 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. An important goal is to obtain control of all errors for these matrix elements at an accuracy of 5%. Increases in computing resources will allow the errors to be systematically reduced, and the calculations extended to a wide variety of other matrix elements.

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 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 an important goal.
- 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.
- 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.

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 calculations 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 demand 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 an accuracy of 7%. In spectroscopy correlation-matrix techniques have been used to determine the spectrum of pure Yang-Mills glueballs, and the efficacy of applying this approach 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%. Sufficiently high statistics 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. Calculations of the baryon and meson resonance spectrum using anisotropic clover lattices with light pion masses will both provide insight into their structure, and inform the expected production rates at GlueX and other experiments. 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.

Physics Beyond the Standard Model: 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. Current work involves exploratory studies, such as the development of techniques for studying supersymmetric theories on the lattice, in preparation for possible large scale projects.

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 had a number of successes in exploring such issues, 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.

3 Lattice QCD Computing Project Hardware

The mission of the LQCD Computing Project is to acquire and operate dedicated computers for the study of Lattice QCD. Hardware is located at Brookhaven National Laboratory (BNL), Fermi National Accelerator Laboratory (FNAL) and the Thomas Jefferson National Accelerator Facility (JLab). A list of the computers now in operation, their location, their number of nodes and their performance in sustained teraflop/s is given in Table 2. The average of the performances of the

inverters of the Dirac operator for domain wall (DWF) and improved staggered (Asqtad) quarks is used as the measure of performance of the computers. These two codes consume a significant fraction of our computing resources, and are representative of our full codes. Benchmarks for them provide a straightforward way to compare computers. The first four computers listed in Table 2, QCD, 3g, 4g, and Pion, are commodity clusters constructed under the SciDAC-1 grant with supplementary funding from FNAL, JLab and the DOE's Office of High Energy Physics. The QCDOC was built with funding from the DOE's Offices of Advanced Scientific Computing Research, High Energy Physics and Nuclear Physics. The 6n and Kaon clusters were acquired with FY 2006 funds from the LQCD Computing Project. Both of these clusters also contain hardware purchased with SciDAC-1 funds. The 6n cluster has 3.0 GHz Pentium-D dual-core nodes, and the Kaon cluster 2.0 GHz dual-socket, dual-core Opteron nodes. Both employ Infiniband networks. The procurement of the FY 2007 hardware is in progress, and the detailed acquisition plan is located on the review website. It will be an Infiniband cluster located at JLab with either dual-socket, dual-core or dual-socket, quad-core Opteron nodes. The decision regarding the choice of nodes will be made after the quad-core processors become available for testing, which we expect to happen prior to the review.

Computer	Site	Nodes	Performance (teraflop/s)
QCD	FNAL	127	0.15
3g	JLab	128	0.10
4g	JLab	384	0.54
Pion	FNAL	518	0.86
QCDOC	BNL	12,288	4.20
6n	JLab	256	0.55
Kaon	FNAL	600	2.56

Table 2: The computers currently operated by the Lattice QCD Computing Project.

The OMB-300 milestones for this project and the progress made towards meeting them are set out in a separate document on this website. One of the major milestones is the performance of the new hardware deployed each fiscal year. The milestones for FY 2006, 2007, 2008 and 2009 are 2.0, 2.9, 4.2 and 3.0 teraflop/s, respectively. Again, performance is measured as the average of that for the DWF and Asqtad inverters. The components of 6n and Kaon purchased with LQCD Computing Project funds have a total throughput of 2.6 teraflop/s, easily exceeding the FY 2006 deployment milestone. If the quad-core Opterons meet our expectations, then the FY 2007 acquisition will also meet its milestone.

A second major milestone is the delivered TF-years. (One TF-year is the number of floating point operations produced by a computer sustaining one teraflop/s in an 8,000 hour "year"). The FY 2007 milestone is to deliver 9.0 TF-Years. Through February 28, 2007 a total of 3.83 TF-Years had been delivered, which puts the project on pace to deliver 9.19 TF-Year for the full fiscal year without any contribution from the FY 2007 acquisition. With full deployment of quad-core processors, the total cycles delivered could be as high as 9.93 TF-years.

At the review we will present possible strategies for hardware acquisitions in the last two years of the Project. At the FY 2006 review, we proposed that FY 2008 and 2009 hardware funds be combined to acquire a single computer in FY 2008 to be located at FNAL. This approach would enable us to obtain a single large machine capable of addressing some of our most challenging problems at an early date. It would also reduce the overhead associated with the acquisition process. The FY 2006 Review Panel recommended that this proposal be followed, but it may require that FY 2009 hardware funds be moved forward to FY 2008.

4 Project Management

The Contract Project Manager, William Boroski (FNAL), has responsibility for the overall management of the Project. He insures that the Project is well defined via a work breakdown structure, and tracked via milestones. He is the primary interface to the DOE for financial matters, reporting and reviews of the Project. Mr. Boroski is assisted by the Associate Contract Project Manager, Bakul Banerjee (FNAL). She maintains the Project's work breakdown structure and all other controlled documents related to its management, and tracks expenditures and progress in achieving milestones. Each of the host laboratories has a site manager or managers: Eric Blum (BNL), Don Holmgren and Amataj Singh (FNAL) and Chip Watson (JLab). The site managers are responsible for the hardware deployment and operations at their laboratories. They are also responsible for developing and executing the components of the work breakdown structure relevant to their sites, and for insuring that appropriate commitments by their laboratory are obtained and carried out.

The Lattice QCD Executive Committee provides leadership for the USQCD's effort to develop computational infrastructure for the study of QCD. Its role in the LQCD Computing Project is to set scientific goals, and determine the computational resources needed to achieve these goals. The members of the Executive Committee are Richard Brower (Boston U.), Norman Christ (Columbia U.), Michael Creutz (BNL), Paul Mackenzie (FNAL), John Negele (MIT), Claudio Rebbi (Boston U.), David Richards (JLab), Stephen Sharpe (U. of Washington), and Robert Sugar (UC Santa Barbara, Chair). Robert Sugar, the Chair of the Executive Committee serves as scientific spokesperson for the Project. He is the principal point of contact for the DOE on scientific matters related to the project, and the liaison between the Executive Committee and the Contract Project Manager.

The Scientific Program Committee assists the LQCD Executive Committee in providing scientific leadership for the USQCD infrastructure effort. It monitors the scientific progress of the effort, and provides leadership in setting new directions. The major activity of the Committee is to allocate computational resources operating by the LQCD Computing Project. It begins this process by issuing a call for written proposals once per year. Oral presentations of large proposals are made and discussed at the annual USQCD all hands meeting, which is organized by the Scientific Program Committee. Following this meeting the Committee makes final allocations. This year's call for proposals, which sets out the process in detail can be found on the review website. Starting this year, the Committee has also used the proposal process to make recommendations to the Executive Committee regarding priorities for proposals for use of the DOE's Leadership Class computers. Members of the Scientific Program Committee are appointed by the LQCD Executive Committee for terms of three years. The current members are Tom Blum (U. Connecticut), Andreas Kronfeld (FNAL, Chair), Colin Morningstar (Carnegie Mellon U.), John Negele (MIT), Stephen Sharpe (U.

Washington), Junko Shigemitsu (Ohio State U.), Doug Toussaint (U. Arizona).

The Change Control Board assures that changes to the project are managed with the primary focus on the advancement of the scientific goals. It must approve all changes resulting in cumulative increases of more than \$125,000, or a delay of more than one month of any level 1 milestone or three months of any level 2 milestone. The members of the Change Control Board are Eric Blum (BNL), William Boroski (FNAL), Steven Gottlieb (Indiana U.), Robert Sugar (UC Santa Barbara, Chair), Victoria White (FNAL) and Roy Whitney (JLab).

5 SciDAC Software

The SciDAC software effort is not part of the LQCD Computing Project; however, as previously noted, it enhances the productivity of the Project's hardware. Under its SciDAC grants, the USQCD Collaboration created a unified programming environment that enables its members to achieve high efficiency on terascale computers, including commodity clusters, the QCDOC and the DOE's leadership class computers. Principal design goals were to enable users to quickly adapt codes to new architectures, easily develop new applications and incorporate new algorithms, and preserve their large investment in existing codes. These goals were achieved through the development of the QCD Applications Programming Interface (QCD API).

The QCD API has a layered structure which is implemented in a set of independent libraries. Level 1 provides the code that controls communications (QMP) and the core single processor computations (QLA). To obtain high efficiency on terascale computers this layer may have to be specifically optimized for a given platform. However, portable versions exist in C and C++ using MPI for communications. Level 2 (QDP) is a data parallel language that enables the user to quickly develop efficient, portable code. By building on QMP and QLA, the details of communications buffers, synchronization barriers, vectorization over multiple sites on each node, etc. are hidden from the user. A very large fraction of the resources in any lattice QCD simulation go into a few computationally intensive subroutines, most notably the repeated inversion of the Dirac operator, a large sparse matrix. To obtain the level of efficiency at which we aim, it is necessary to optimize these subroutines for each architecture. Level 3 (QOP) consists of these routines. They can be called from QDP, as well as from general C or C++ code that conforms to the QCD API.

A very large fraction of the computing resources used in lattice QCD calculations go into Monte Carlo simulations that generate representative configurations of the QCD ground state. The same configurations can be used to calculate a wide variety of physical quantities of interest in high energy and nuclear physics. Because of the large resources needed to generate configurations, the USQCD Collaboration has agreed to share all of those that are generated with DOE resources. To enable this sharing we have created standards for file formats, and written an I/O library (QIO) that adheres to them. We are charter members of the International Lattice Data Grid (ILDG), which has developed a basic set of meta-data and middleware standards to enable international sharing of data. QIO conforms to these standards, and therefor enables members of USQCD to archive and retrieve data on the ILDG.

An important activity under the USQCD’s SciDAC-2 grant will be to optimize the QCD API for computers with multi-core processors. For dual-core processors we can treat each core as an individual processor with little, if any, loss in efficiency. However, for processors with four or more cores, multi-threaded code has the potential to provide significant improvement in performance. Current trends indicate that such processors will certainly be used in the FY 2008/2009 hardware acquisition, if not in the FY 2007 one. So, a multi-threaded version of the QCD API has the potential to markedly enhance the performance of the Project’s hardware.

6 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 Institute 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.