Computational Infrastructure for Lattice Gauge Theory

A Strategic Plan

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.),
S. Sharpe (U. Washington), R. Sugar (UCSB) and W. Watson, III (JLab)

April 4, 2002

Executive Summary

The goals of our research are to understand the physical phenomena encompassed by quantum chromodynamics (QCD) and to make precision calculations of the theory's predictions. To do so requires terascale simulations within the framework of lattice gauge theory. Such simulations are necessary to solve the fundamental problems in high energy and nuclear physics that are at the heart of the Department of Energy's large experimental efforts in these fields. A computational capability of tens of teraflops (sustained) is needed to achieve our near term scientific goals. In this document we set out our plans for developing this capability in partnership with the DOE. By taking advantage of special features of lattice QCD calculations, we will achieve the required computational power at a fraction of the cost that would be required with conventional high performance computing platforms.

Major goals of the DOE's experimental program in high energy and nuclear physics are i) verifying the the Standard Model or discoverying its limits, ii) determining the properties of hadronic matter under extreme conditions, and iii) understanding the structure of nucleons and other hadrons. Lattice QCD calculations are essential to research in all of these areas.

A central focus of experiments at U.S. high energy physics facilities is precision testing of the Standard Model. The ultimate aim of this work is to find deviations from this model—a discovery which would require the introduction of new physical principles to describe matter at the shortest distances. Many of these tests require, in addition to precise experimental measurements, accurate evaluation of the effects of the strong interactions on processes induced by the electroweak interactions. Such evaluations require multi–teraflops lattice QCD simulations. We estimate that a computer sustaining 0.5 teraflops for a year would lead to a reduction of error in the evaluation of crucial weak matrix elements by about a factor of two, and that a machine sustaining 10 teraflops for a year would halve the uncertainties again. Such a reduction in errors will be crucial for the interpretation of experimental results from the SLAC B-factory, the Tevatron B-meson program at FNAL and the proposed CLEO-c program at Cornell.

At low temperatures and densities quarks and gluons are confined in elementary particles, such as neutrons and protons. At very high temperatures and densities one expects a phase transition or crossover from this ordinary state of strongly interacting matter to a plasma of quarks and gluons, which is believed to have been a dominant state of matter in the early development of the universe and a possible central component of neutron stars today. A primary physics goal of the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory is the discovery and characterization of the quark-gluon plasma. The computational facilities we propose would allow theorists to determine the nature of the transition, the properties of the plasma, including its strange quark content and equation of state, subjects which are critical to the RHIC physics program, and which can only be addressed from first principles through lattice QCD calculations.

The third major scientific goal is to achieve a quantitative, predictive understanding of the structure and interactions of hadrons. The internal structure of the nucleon is a defining problem for hadron physics just as the hydrogen atom was for atomic physics. Indeed, the DOE Strategic Plan specifically highlights the goal of developing a quantitative understanding of how quarks and gluons provide the binding and spin of the nucleon based on QCD. Major experimental efforts in recent years at Bates, JLab, SLAC, FNAL, DESY, and CERN have provided rich and precise measurements of the quark and gluon structure of the nucleon, and proposed experiments, such as the RHIC spin program, promise to reveal even greater detail. A terascale computational facility, together with recent advances in lattice field theory, will make it possible to calculate this nucleon structure directly from QCD, and thus realize the full physics potential of major accelerators and detectors.

We propose to construct a distributed computing facility for lattice QCD with major hardware located at Brookhaven National Laboratory, Fermi National Accelerator Laboratory, and Thomas Jefferson National Accelerator Facility. Two types of hardware platforms will be utilized: (1) the QCDOC, a specially designed computer that combines computation and communication capabilities on a single chip, and (2) large-scale commodity clusters, which are configured to maximize their cost–effectiveness for lattice QCD calculations. The software infrastructure needed to achieve very high efficiency on these platforms is under development with support from the DOE SciDAC program. The construction and maintenance of computational infrastructure for lattice QCD will require a sustained, long term effort. We propose to ramp up spending for production hardware during the 2003–2005 fiscal years, reaching a plateau of approximately \$9.4M per year in FY 2006 and beyond. By the end of 2006 we expect our distributed facility to sustain more than 20 teraflops on lattice QCD calculations. We emphasize that this number represents sustained, not peak, performance. Beyond 2006 continued funding and Moore's law will enable us to steadily increase the capabilities of the facility.

The project to construct, staff and operate a distributed computing facility for lattice QCD will be managed under the administrative structure created for our current SciDAC grant. It consists of an Executive Committee, which has overall responsibility for the project; a Scientific Program Committee, which monitors the scientific progress of the project, provides leadership in setting new directions, and will allocate resources on all hardware built in this project; and an Oversight Committee, which reviews progress in implementing the plans of the collaboration, reviews plans for the acquisition of hardware, and makes recommendations regarding alternative approaches or new directions for the project. Our collaboration comprises almost all of the senior lattice gauge theorists working in the United States, as well as several computer scientists. The hardware and software infrastructure developed under this project will be open to the entire U.S. lattice QCD community.

Time is critical. Experimental results will be produced in the next few years at a cost of hundreds of millions of dollars. Their interpretation will require large scale lattice QCD calculations. Theorists in Europe and Japan presently have computational resources exceeding those now available in the U.S., and are moving rapidly to secure new resources comparable to those we propose. If U.S. lattice gauge theorists are to play a significant role in the major advances expected in this area over the next decade, we must act now.

1 Introduction

The over–arching goals of our research are to understand the physical phenomena encompassed by quantum chromodynamics (QCD) and to make precision calculations of the theory's predictions. To do so requires terascale simulations within the framework of lattice gauge theory. Such simulations are necessary to solve the fundamental problems in high energy and nuclear physics that are at the heart of the Department of Energy's large experimental efforts in these fields. Virtually all of the members of the U.S. lattice QCD community are working together to construct the hardware and software infrastructure needed for this research. We have recently been awarded a three year grant under the Department of Energy's SciDAC program for software development and hardware prototyping. In this document we set out our plans for the construction and operation of terascale hardware facilities.

The twentieth century was an era of striking progress towards comprehending the fundamental structure of matter, beginning with the discovery of quantum mechanics and atomic physics, progressing to nuclear physics, and culminating with the Standard Model of high energy physics. However, traditional analytical tools have proven inadequate to extract many of the predictions of QCD. Our understanding of nature will remain fundamentally deficient until we know how the rich and complex structure of strongly interacting matter, which comprises most of the known mass of the universe, arises from the interactions among quarks and gluons.

Perturbation theory has proven to be a powerful tool for studying the electro-weak interactions and those aspects of QCD that are controlled by short distance effects, for which interactions among quarks and gluons are weak. However, most QCD phenomena involved large distance effects for which interactions among quark and gluons are strong, and perturbation theory is therefore not applicable. At present, the only method to extract predictions of QCD in the non-perturbative regime from first principles and with controlled systematic errors is through large scale numerical simulations of lattice gauge theory. Lattice gauge theory calculations have demonstrated important qualitative features of QCD, such as quark confinement and chiral symmetry breaking. They have also yielded quantitative results of increasing accuracy. Recent refinements of numerical algorithms coupled with major increases in the capabilities of massively parallel computers have brought these simulations to a new level. It is now possible to calculate a few crucial quantities to an accuracy of a few percent. The strong coupling constant and the masses of the c and b quarks are notable examples. Furthermore, the experience we have gained allows confident predictions of the computing resources required for accurate determinations of a broad range of fundamental quantities. A very substantial fraction of the DOE's experimental programs in high energy and nuclear physics is aimed at carrying out precision tests of the Standard Model, understanding the structure of nucleons and other hadrons, and determining the properties of hadronic matter under extreme conditions. Lattice gauge theory calculations are essential to this research. The terascale computing facilities we propose to construct are required if these calculations are to reach the needed accuracy in step with the experiments they support. In Section 2 of this document we set out our scientific goals, and identify projects that will receive high priority for early use of the proposed facilities. We estimate the resources that will be required to make substantial progress on these projects. A list of major achievements, as well as near term prospects for future ones is given in Appendix A.1, along with the computer resources needed to achieve them.

The size of the computing resources we seek to construct is driven by our scientific objectives. It is clear from the discussion in Section 2 that in order to provide support to the experimental programs in high energy and nuclear physics in a timely fashion, and to keep pace with the ambitious plans of our colleagues in Europe and Japan, the U.S. lattice QCD community requires computing resources capable of sustaining tens of teraflops within the next few years. By taking advantage of simplifying features of lattice QCD, such as regular grids and the well understood influence of each site on its neighbors which leads to uniform, predictable communications, it is possible to construct computers for lattice QCD that are far more cost effective than general purpose supercomputers, which must perform well for a wide variety of problems including those requiring irregular or adaptive grids, non-uniform communication patterns, and massive input/output capabilities. We are therefore convinced that it is possible to address our scientific goals in a significant manner at a cost of less than \$10M per year.

We have identified two computer architectures that promise to meet the needs of lattice QCD. One is the QCDOC, the latest generation of highly successful Columbia/ Riken/ Brookhaven National Laboratory (BNL) special purpose computers, which is being developed at Columbia University in partnership with IBM. The other is commodity clusters, which are being specially optimized for lattice QCD at Fermi National Accelerator Laboratory (FNAL) and Thomas Jefferson National Accelerator Facility (JLab). This two track approach will position us to to exploit future technological advances, and enable us to retain flexibility. Each architecture has its own strengths, and either may prove optimal for different aspects of our work. The QCDOC project is expected to provide very powerful computing platforms within the coming year at a cost of approximately \$1 per sustained megaflops (Mflops). These platforms are likely to be particularly effective at generating the computationally expensive dynamical quark lattices that are crucial for our research. The clusters will allow us to take advantage of the rapid advances in the commodity computing market to build increasingly cost-effective machines over time. The well developed software packages and flexible communications systems of clusters make them particularly advantageous for the more complex physics applications and for the development of new algorithms and computational techniques crucial for the advancement of our field. We discuss both architectures in Section 3. The development of the software infrastructure needed to obtain very high efficiency on these platforms is in progress under our SciDAC grant. It is briefly described in Section 4.

The concept of a topical computing facility, as set out in the Office of Science's computing plan, is particularly well suited for lattice QCD. We propose to construct a distributed facility with major hardware located at BNL, FNAL and JLab. Initially BNL will focus on the QCDOC, while FNAL and JLab will concentrate on clusters. This approach will allow us to take advantage of the very considerable expertise at each of the participating laboratories, while building platforms of the appropriate size to meet our research objectives. The facilities planned for each of the laboratories are described in Section 5.

The construction and maintenance of computational infrastructure for lattice QCD will require a sustained, long term effort. At present, development work for the QCDOC is being funded through the HEP base program, and that for the clusters through our SciDAC grant. The SciDAC grant provides funds to develop software that will enable the U.S. lattice QCD community to exploit both types of computers productively. We propose to ramp up spending for production hardware during the 2003–2005 fiscal years, reaching a plateau of approximately \$9.4M per year in FY 2006. During 2003 and 2004 we will focus on the QCDOC, because it is expected to provide the most capable and cost-effective hardware during this period. Extrapolations indicate that by 2005 clusters can surpass the QCDOC in cost effectiveness, so we anticipate switching our focus to them for 2005 and 2006. Given the computer industry's rapid rate of development, it is essential that we re-evaluate our plans each year, and when necessary re-allocate resources to maximize the scientific return on the investments. Before building any new, large platform, we will construct a smaller prototype to evaluate the performance of production codes. We will select machines (and their configurations) to maximize physics production, and will target software developments to areas with the greatest potential for further performance gains. The funding profile we propose is set out in Section 6.

The project to construct, staff and operate a distributed computing facility for lattice QCD will be managed under the administrative structure created for our current SciDAC grant. The Executive Committee will have overall responsibility for the project. The Scientific Program Committee will monitor the scientific progress of the project, and provide leadership in setting new directions. It will also allocate the resources on all hardware built in this project. The Oversight Committee will review progress in implementing the plans of the collaboration, review plans for the acquisition of hardware, and make recommendations regarding alternative approaches or new directions for the project. The members of these committees are given in Appendix A.2, and the senior members of the collaboration in Appendix A.3.

2 Physics Goals and Required Computational Resources

Our goal is to understand the physical phenomena encompassed by QCD, and to make precision calculations of its predictions. We will follow a multifaceted approach in which we calculate both quantities that can be compared with existing experimental results, so as to calibrate the accuracy of our methods, and use the same methods to predict quantities yet to be measured. In particular, our calculations will allow increasingly precise tests of the electroweak sector of the Standard Model, allow us to determine the properties of hadronic matter under extreme conditions, and understand nucleon structure and interactions. We describe each of these areas below. Furthermore, we expect the calculational techniques and computational infrastructure we develop for the study of QCD to be applicable to strongly coupled theories which go beyond the Standard Model, such as supersymmetric gauge theories, chiral gauge theories, and string theory. Indeed, work in these directions is already in progress by some members of our community.

2.1 Precision Testing of the Standard Model

A central focus of experiments at U.S. high energy physics facilities is precision testing of the Standard Model. The ultimate aim of this work is to find deviations from this model—a discovery which would require the introduction of new physical principles to describe matter at the shortest distances. Many of these tests require, in addition to precise experimental measurements, accurate evaluation of the effects of the strong interactions on processes induced by the electroweak interactions. Such an evaluation requires lattice QCD, the only known method which can systematically reduce all sources of error. These computations are one of the major physics focuses of our research plan, and a crucial companion to the U.S. experimental program in high energy physics.

The technical challenge is to calculate quantities known as "weak matrix elements" (matrix elements of electroweak operators between hadronic states). Each such matrix element, when combined with a particular experimental quantity, gives a direct measurement of an underlying parameter of the Standard Model. If multiple measurements of these parameters disagree, new physical principles are required. Table 1 summarizes the present situation for four key matrix elements. For three of the four, lattice calculations lag well behind experiment. The impact of the larger lattice errors is shown in Figure 1. For the Standard Model to be correct, the parameters ρ and η are constrained to lie in the region of overlap of the solidly colored bands. The figure on the left shows the constraints as they exist today. The figure on the right shows the constraints as they would exist with no improvement in the experimental errors, but with lattice gauge theory errors reduced to 3%.

Terascale computers will lead to an enormous advance in lattice calculations of such matrix elements, thus bringing the theoretical precision much closer to that of the experiments. We estimate that a computer sustaining 0.5 teraflops (Tflops) for a year would lead to a reduction in the uncertainties in the quantities listed in the first three rows of Table 1 by about a factor of two, and that a machine sustaining 10 Tflops for a year would halve the uncertainties again. These are conservative estimates, relying on the measured scaling properties of existing algorithms. Such a reduction in errors—coupled with improvements in experimental results from the SLAC B-factory and the Tevatron B-meson program—will narrow the corresponding bands in Figure 1 by four or more, allowing the possibility of inconsistent determinations of ρ and η.

Measurement	CKM Matrix Element	Hadronic Matrix Element	Expt. Error	Current Lattice Error	Lattice Error 0.5 TF-Yr	Lattice Error 10 TF-Yr
ΔM_{B_d} ($\bar{B}B$ mixing)	$\lvert V_{td}\rvert^2$	$f_{B_d}^2 B_{B_d}$	4%	35%	18%	9%
$\Delta M_{B_s}/\Delta M_{B_d}$	$\lvert V_{ts}/V_{td}\rvert^2$	$f_{B_s}^2 B_{B_s}/f_{B_d}^2 B_{B_d}$	Not yet measured	10%	5%	3%
ε ($\bar{K}K$ mixing)	$\mathrm{Im}\,V_{td}^2$	B_K	2%	20%	10%	5%
$B \to \binom{\rho}{\pi} l\nu$	$\lvert V_{ub}\rvert^2$	$\left\langle \begin{smallmatrix}\rho\\ \pi\end{smallmatrix} \right\rvert (V-A)_\mu \lvert B\rangle$	25%	Calc. in progress	15%	5–10%

Table 1: *Impact of lattice QCD on the determination of CKM matrix elements. In the table above, f_X is the leptonic decay amplitude for the indicated meson, and B_X is proportional to the matrix element of $\Delta S = 2$ or $\Delta B = 2$ four–quark operators. The last two columns show the improvements in lattice errors that we estimate would be obtained with computers sustaining 0.5 and 10 Tflops for one year.*

There are many other matrix elements that can be used in a similar way to test the Standard Model, but for which the lattice calculations are less advanced. Examples include semileptonic form factors of D mesons (e.g. $D \to \pi\ell\nu$ which can be used to study $c \to d$ transitions), radiative form factors of B mesons (e.g. $B \to K^*\gamma$), CP-violating parts of $K \to \pi\pi$ decay amplitudes (measured by ε'/ε), and the electric dipole moment of the neutron (of measurable size in some extensions of the standard model). In these cases, a terascale facility will enable substantial improvements, and provide an essential step towards reaching the desired accuracy.

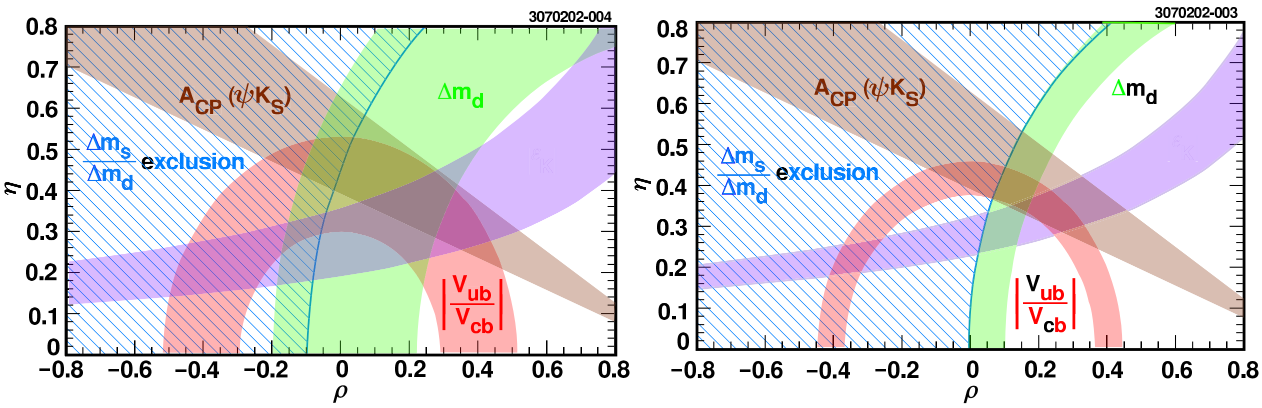

Figure 1: Constraints on the Standard Model parameters ρ and η (one sigma confidence level). For the Standard Model to be correct, they must be restricted to the region of overlap of the solidly colored bands. The figure on the left shows the constraints as they exist today. The figure on the right shows the constraints as they would exist with no improvement in the experimental errors, but with lattice gauge theory uncertainties reduced to 3%. R. Patterson, Cornell University.

There are also a large number of measured hadronic properties that will be calculable to high precision using a terascale facility, and which can be used to calibrate our methods. Such tests will be important demonstrations of the reliability of lattice calculations. Examples of properties that can be used for calibration include the masses and decay constants of B and D mesons and corresponding baryons, of charmonium and bottomonium states (bound states, respectively, of charm or bottom quarks and their antiparticles), and of hadrons composed of light quarks. An accuracy of a few percent is expected for some of these quantities with a terascale facility. We note that this calibration will be significantly sharpened if the proposed CLEO-c program goes forward at Cornell.

An important feature of the lattice methodology is its flexibility. Particle physicists will no doubt discover other interesting matrix elements to calculate, and it is often the case that these can be "piggybacked" on previous calculations with little overhead. Thus the terascale facility we propose will provide a flexible database which can be reused repeatedly as new ideas appear.

Finally, we note that our calculations will also provide precise values for other fundamental parameters of the Standard Model—the quark masses and the strong coupling constant, α_S. Precise results for these are needed to differentiate between competing models of flavor physics and electroweak symmetry breaking, and lattice simulations provide the only method for doing such calculations. Indeed, the lattice results for the c and b quark masses are already very accurate (e.g. the error is 2% for m_b). The light quark masses—m_u, m_d and m_s—are more difficult to calculate, requiring extensive simulations with light dynamical quarks. Present errors, estimated to be 25%, will be substantially reduced by a terascale facility.

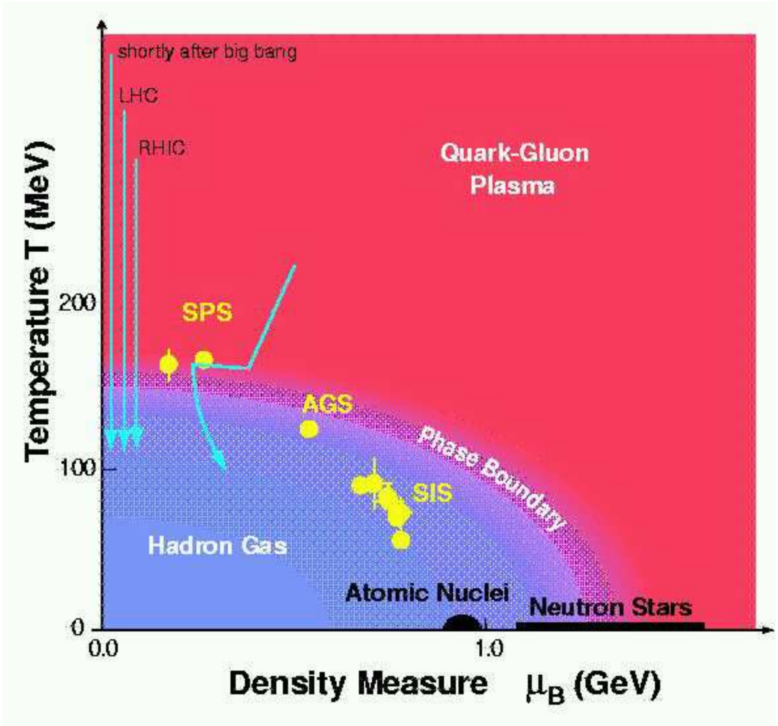

Figure 2: Last seen a few microseconds after the big bang, the quark gluon plasma is the quarry of the RHIC facility, and can be explored from first principles using lattice gauge theory. http://www-aix.gsi.de/˜alice/phase-diag.jpg

2.2 Simulating the Quark Gluon Plasma

At low temperatures and densities quarks and gluons are confined in elementary particles, such as neutrons and protons. At very high temperatures or densities one expects a phase transition or crossover from this ordinary state of strongly interacting matter to a plasma of quarks and gluons. Such a plasma is believed to have been a dominant state of matter in the early development of the universe, and it may exist today in the cores of neutron stars. A primary physics goal of the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory is the discovery and characterization of the quark-gluon plasma. There are a number of issues that are critical to the RHIC physics program which can only be addressed from first principles through lattice gauge theory calculations. These include the nature of the transition, the properties of the plasma, including strange quark content, and the equation of state. Their study will be a major focus of our work.

One of the major advances in lattice QCD in recent years has been the development of improved actions, that is improved methods of placing the continuum theory on the lattice. These improved actions vastly increase the accuracy of lattice calculations for a given amount of computing resources. Members of our collaboration have played important roles in this work. We propose to take advantage of our recent successes with improved actions and the enormous power of the proposed facilities to carry out a definitive study of the quark-gluon plasma at high temperatures and zero baryon density with a realistic quark ensemble and vastly reduced discretization artifacts. The results are expected to give valuable assistance to the RHIC experimental program. The initial objectives of this work will be to map the phase diagram in temperature and quark mass for up, down, and strange quarks, determine the order of the phase transition, and obtain an accurate determination of the temperature of the crossover. The next step will be the determination of the equation of state of the plasma, including strange quark content, and the prediction of real-time excitations of the plasma. These objectives have obvious relevance to the analysis of RHIC experiments. We will also seek to understand the role of instantons in the phase transition, and to measure the strength of the axial $U(1)$ anomaly. This information is needed to formulate phenomenological models that extrapolate to regimes inaccessible to lattice gauge theory.

The development of highly improved actions for lattice QCD has opened the possibility of carrying out realistic simulations of hadronic matter in the vicinity of the quark-gluon transition. However, significant enhancement in computing power is needed for this work. Based on our extensive experience in the study of QCD thermodynamics with simpler actions, this will be a multi-year project, which will require long-term use of the proposed multi-Tflops facilities.

The study of QCD at finite baryon density is a very important problem, which will require the development of new algorithms. Work in this direction will be an important part of our effort, and will require significant computing resources.

2.3 Structure and Interactions of Hadrons

The third major scientific goal of our collaboration is to achieve a quantitative, predictive understanding of the structure and interactions of hadrons.

The internal structure of the nucleon is a defining problem for hadron physics just as the hydrogen atom is for atomic physics. Indeed, the DOE Strategic Plan specifically highlights the goal of developing a quantitative understanding of how quarks and gluons provide the binding and spin of the nucleon based on QCD. Major experimental efforts in recent years at Bates, JLab, SLAC, FNAL, DESY, and CERN have provided rich and precise measurements of the quark and gluon structure of the nucleon, and proposed experiments such as the RHIC spin program promise to reveal even greater detail. With recent advances in lattice field theory, it is now possible to calculate this nucleon structure directly from QCD, so that multi-terascale lattice calculations have become an essential tool to obtain the full physics potential of major accelerators and detectors.

A wealth of experimental observables can be calculated on the lattice. Electromagnetic form factors measured in elastic electron scattering characterize the distribution of charge and magnetization arising from all the quarks in the nucleon, and parity violating electron scattering experiments further reveal the specific contributions to these quantities from strange quarks. Deep inelastic scattering of electrons, muons, and neutrinos measures structure functions characterizing the light cone quark density, quark spin density, and gluon density as a function of momentum fraction, and the moments of these distributions can be calculated on the lattice. A particularly important example is the lowest moment of the spin density, which measures the fraction of the nucleon spin carried by the spin of quarks. Indeed, the only way to fully resolve the so-called "spin crisis" which arose when experiments showed that only about 20% of the spin of the nucleon originates from quark spins is to calculate in lattice QCD how the total spin is divided between quark and gluon spin and orbital angular momentum.

Calculations with limited computational resources have already established the methodology to calculate the nucleon form factor, the contributions of strange quarks in the nucleon, and moments of the quark density, spin, and transversity distributions in the nucleon. These calculations have highlighted the fact that the pion cloud plays an essential role in all these observables, and that quantitative agreement with experiment will only be possible in the next generation of calculations which are performed on sufficiently large lattices at sufficiently low quark masses that the physics of the pion cloud is accurately included. Based on the known scaling properties of present algorithms and using the analytical tools of chiral perturbation theory, the relevant calculations in full QCD with dynamical quarks could be carried out in one year of production work on a computing facility sustaining 10 Tflops.

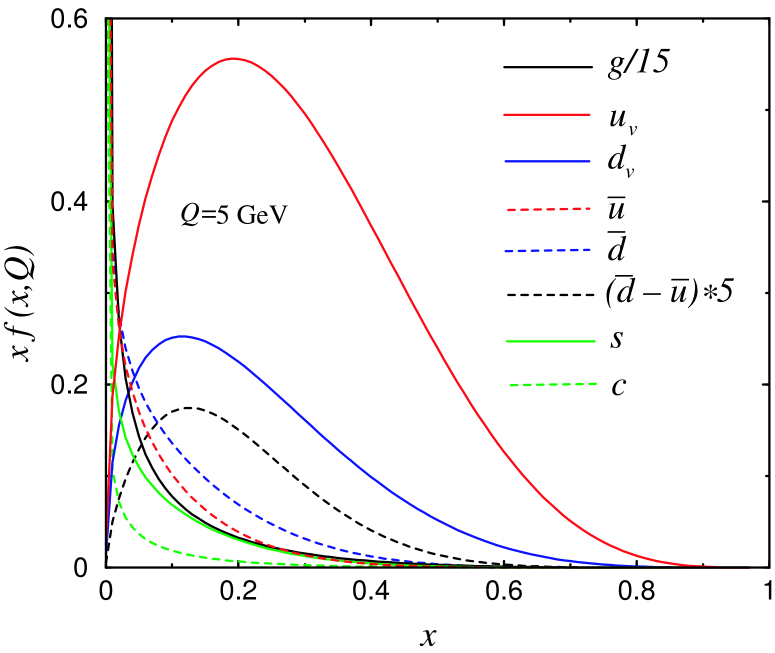

Figure 3: High energy scattering experiments have measured the distribution of quarks and gluons in the proton. The measured distributions $f(x,Q)$ for six types of quarks and antiquarks (denoted u, d, $\bar{u}$, $\bar{d}$, s, and c) and for gluons (denoted g) are shown in the figure. Multi-teraflops facilities will enable calculation of the moments of these distributions from first principles, providing a fundamental understanding of the structure of the proton.

Spectroscopy is the classic tool for discovering the relevant degrees of freedom of a physical system and the forces between them. One of the fascinating features of QCD is that it offers the possibility of a richer range of hadronic states than has yet been observed experimentally, so that precise lattice calculations can play a pivotal role in helping guide experimental searches. Lattice calculations will study the number and structure of hadronic excited states, as well their transition form factors. The presence or absence of hadrons with exotic quantum numbers, the nature of glueballs, and the overlap between model trial functions and exact hadron states will provide insight into the role of flux tubes, dibaryons, and the inner workings of QCD. Again, exploratory calculations of the lowest negative parity N^* state, a comprehensive calculation of the glueball spectrum, calculation of hybrid mesons, and a study of the existence of the H particle show that the requisite methodology is ready. Precision multi-terascale calculations will provide crucial insight into current and future hadron spectroscopy.

Currently, there is no fundamental understanding of the very foundation of nuclear physics, the nucleon-nucleon interaction. Significant insight into the role of gluon exchange, quark exchange, meson exchange, and the origin of short range repulsion will be obtained by lattice calculations of the adiabatic potential between heavy-light systems, that is, mesons or baryons containing a single heavy quark in addition to other light quarks or antiquarks.

In addition to calculating observables to compare with experiment, lattice calculations are invaluable to obtain insight into fundamental aspects of QCD. Current lattice techniques can study the role of instantons and their associated zero modes in chiral symmetry breaking, the role of center vortices and magnetic monopoles in confinement, and calculation of the parameters entering chiral perturbation theory. The lattice also allows theorists to answer interesting theoretical questions inaccessible to experiment, such as how the properties of QCD change with the number of colors, quark flavors, or quark masses. Recently developed techniques may also enable study of the phases of dense hadronic matter and the transitions between them.

A more detailed description of research opportunities in hadron structure and interactions may be found on the web at `ftp://www-ctp.mit.edu/pub/negele/LatProp/`

3 Hardware Plans

3.1 Deployment Strategy

As detailed in the previous section, a computational capability of tens of Tflops is needed to achieve our near-term scientific goals. To meet this need we plan an aggressive hardware deployment project in partnership with the Department of Energy. By taking advantage of the special features of lattice QCD calculations, we will achieve the required computational power at a fraction of the cost of typical high performance computing platforms. We plan to follow two independent approaches, which both promise to lead, in a complementary manner, to low-cost, massive computational capabilities. Specifically, we intend to exploit (1) the QCDOC, a specially designed machine that combines computation and communication capabilities on a single chip, and (2) large-scale commodity clusters, which are configured to maximize their cost-efficiency for lattice QCD calculations. By differentiating our hardware strategy in this way, we will be able to follow technical developments, and redirect our efforts, if necessary, so as to achieve the maximum return on the investment in the project. The diversification entails another major advantage. With the broad range of investigations we plan, different calculations will be better supported on different platforms, and we will be able to tailor our computational strategies accordingly.

3.2 QCDOC Overview

As has been discussed above, the regular character of lattice QCD calculations can be exploited to achieve very significant cost/performance advantages. While the most demanding calculations easily require thousands of closely-coupled processors, the underlying nearest-neighbor communication pattern allows the use of a mesh-style network architecture. Such mesh networks can provide both low latency and high bandwidth at very reasonable cost. By exploiting the present capabilities of commercial chip design, it is possible to fabricate an individual computing node which includes this network electronics and is contained in a single chip. It is this combination of processor, network and memory on a single chip which lies at the heart of the QCDOC (Quantum Chromodynamics on a Chip) architecture.

Targeting a cost per sustained performance of $1/megaflops, the QCDOC design plans a single processing node which occupies less than 6 in^3 (including up to 2 Gbytes of off-chip memory) and consumes roughly 3 watts. This combination of price, power and packaging permits the economical construction of the very large, 10K processor, 5 Teraflops machines required for significant advances on many of the critical physics topics addressed in this plan.

The group centered at Columbia has pioneered the design and construction of special purpose machines for the study of QCD. The present QCDSP machines at Columbia and Brookhaven have provided a sustained 300 Gflops for lattice QCD calculations for the past four years and even now represent one of the two largest facilities for lattice QCD simulations internationally. (The QCDSP machine at BNL won the Gordon Bell prize for price performance at the November 1998 Supercomputing Conference, SC98.) With substantial funding from the U.S. DOE, the RIKEN Laboratory of Japan and the UKQCD collaboration, the design of this next QCDOC machine, begun in the fall of 1999, is now nearing completion with first prototype chips expected this summer. This design work represents a very fruitful collaboration between Columbia, RIKEN, UKQCD and the IBM corporation which is exploiting this architecture for its own future research machines (`http://www.research.ibm.com/bluegene/index.html`).

The QCDOC architecture is a natural evolution of that used in the QCDSP machines. Individual processing nodes are PowerPC-based and interconnected in a 6-dimensional mesh with the topology of a torus. A second, Ethernet-based network provides booting and diagnostic capability, as well as more general I/O. The entire computer will be packaged in a style that provides good temperature control, a small footprint and easy accessibility. Central to this design is IBM's technology which makes possible the high-density, low-power combination of an industry-standard RISC processor with 64-bit floating point, embedded DRAM, 500 MHz communications and the wide array of pre-designed functions needed to assemble the complete, functioning unit.

This carefully balanced architecture permits demanding, full QCD calculations to be distributed over many processors, providing the network bandwidth required for the resulting, large surface to volume ratios. Since the design is now nearing completion, the performance of real code can be demonstrated with gate-level simulation. The 500 MHz floating point unit achieves 84% efficiency for the Wilson Dirac operator when operating from cache for double-precision arithmetic, *i.e.* 0.84 Gflops. Using the large 8 Gbyte/sec bandwidth between processor and on-chip memory, this same Wilson Dirac operator can be evaluated with 78% efficiency. Finally with latencies in the range of a few hundred nsecs and an aggregate off-node bandwidth of 12 Gbits/sec, current gate-level simulations achieve a 49% efficiency for the Wilson Dirac operator in the most demanding case of 2^4 sites/node.

Critical to the effectiveness of this QCDOC architecture is the SciDAC software development effort. The creation of standard, community-endorsed communication interfaces as well as both high- and low-level arithmetic routines is supporting an evolution of the very large and valuable U.S. QCD code base which will permit its efficient execution on QCDOC machines as well as specifically optimized cluster network architectures. The resulting efficient interoperability will provide very important advances in scientific productivity for the entire U.S. lattice QCD effort.

We expect these QCDOC machines to play a critical role in U.S. physics for the first 3-4 years covered by this strategic plan. Toward the end of that period, we expect technological advances to offer even greater opportunities for this system-on-a-chip approach to massively parallel computing. In particular, we expect to seriously consider a follow-on machine constructed using the IBM Bluegene/L technology which may achieve \$0.25/Mflops (double precision) for QCD, and permit the construction of a machine of scale 25 teraflops within the budget profile outlined in this plan.

3.3 Clusters Overview

Over the last several years, the intense competition for high performance desktop computers has lead to enormous gains in processor performance. At the same time, the market for specialized interconnects to support parallel computing has made it possible to assemble cost effective clusters from commodity components with the performance required for QCD calculations. An additional strong benefit to the lattice community flowing from the exploitation of clusters is the ease with which university clusters can be used in software development, production running, and data analysis, exploiting the identical software used at the national facilities. For these reasons, commodity clusters form an essential part of the long range plan for lattice computing.

Each year, market leading components will be used to build a leading-edge cluster optimized to the requirements of lattice QCD (lean memory, lean disk, low latency semi–commodity cluster interconnect). The ensemble of clusters will be presented to the users as a coherent computing resource, independent of system version and cluster location, through the combination of platform specific libraries and sophisticated grid software. As processor performance advances, the performance of the latest clusters will increase, extending the physics problems accessible to the lattice theorists each year.

Single chip processor performance increases as Moore's Law (60% per year), with current Intel Pentium 4 chips sustaining more than 2 Gflops single precision on small CPU–intensive lattice kernels, and 1.6 Gflops for the full Wilson inverter. Intel has already demonstrated chips running at 2.5 times today's clock speeds, leading us to expect single chip performance exceeding 5 Gflops within two to three years.

Chip performance is limited by memory bandwidth, which tends to increase with time less rapidly than CPU speed, but this limiting effect is softened by a growing use of, and size of, full speed on–chip level 2 (L2) cache. The impact of the memory bottleneck today is that in a dual processor box and for problems which mostly fit into the L2 cache, the second processor adds only 25% – 40% in performance. However, a second processor adds only 15% – 20% to the cost, so it can still be worthwhile. Quad–processor architectures have even higher memory bandwidths allowing the use of four processors, and are expected to be cost competitive with duals in 2004 or 2005 (the choice between duals and quads will be made based upon market conditions). Today the market is making the transition to double data rate SDRAM (interleaved to achieve 3.2 GB/s peak bandwidth), and Rambus memories will soon double in bandwidth to over 6 GB/s, so healthy memory bandwidth gains can be expected over the next several years. In this proposal, sustained multi–processor node performance is conservatively predicted to grow at approximately 40% (instead of 60%) per year for the next two years, and somewhat more slowly thereafter, to account for memory bandwidth limitations.

Cluster interconnects introduce additional bandwidth and latency related inefficiencies. The bandwidth limitations presently come from the external bus architectures, currently PCI–64/66 (512 Mbytes/s), with PCI–X (1 GB/sec) emerging this year and doubling in performance within two years. By 2005, Infiniband 12X is expected to be the market leader, providing 6 Gbytes/sec of bandwidth per node, thereby keeping pace with processor gains. These high bandwidth buses coupled with high performance interconnects should allow us to reduce the inefficiency introduced by clustering to roughly 10%, although the real optimization on design will be the overall price performance of the entire cluster (CPU, memory, interconnect). Network optimization will require high efficiency in overlapping communications and computation, which will be facilitated by the communications libraries being constructed under our SciDAC grant.

As already stated, the proposed national facility consists of three sites, leveraging the existing computing infrastructure at BNL, FNAL, and JLab. For the two cluster–based centers, the strategy will be a continual deployment of new capacity, adding one cluster at each site per year. With a constant funding profile, this will yield ever increasing capacity (following Moore's Law), with older systems being replaced by newer ones in a 3–4 year timescale. This strategy is the same as that being followed at a number of science computing centers across DOE and other agencies.

In a given year, if market conditions are stable, the systems at FNAL and JLab may be identical; if multiple viable technologies are available with different strengths and weaknesses, or if procurements differ in time by more than 3 months, they may be different. For at least the next several years, the sizes of the lattices that will be used, and the increased flexibility of multiple independent systems, support this multi–cluster (multi–site) strategy.

A cluster will consist of multiple racks of compute nodes interconnected via a high speed cluster interconnect such as Myricom's Myrinet (the current market leader). Each node will have a lean memory configuration with the FNAL and JLab clusters possibly choosing different parameters to optimally cover a range of application requirements, and will have only a small disk of minimal performance, again based upon application requirements. In addition, each node will have a standard network connection (today 100 Mbit ethernet, within 2–3 years 1 Gbit ethernet). The standard network will link the cluster to a set of conventional file servers, some of which will serve as a disk cache for input and output data files backed up by the host laboratory's tertiary storage system.

The table below gives the most significant cluster parameters, and the expected size and cost of the hardware at each of the two cluster sites. Actual procurements may differ, using a build-to-cost strategy and market optimizations. Additional nodes beyond what is shown in the table will be procured as hot spares.

The price–performance ratio in \$/Mflops and the overall performance assumes a full-cluster job with Wilson quarks.

Current and expected next generation interconnects will provide sufficient bandwidth to support clusters of approximately a thousand nodes, which will enable theorists to study many problems of high interest today. By 2007, it will be possible to construct clusters of several thousand nodes and multiple tens of teraflops, constrained in size only by the funding profile, permitting theorists to address even larger problems. As part of the SciDAC project, FNAL and JLab will be working together to select or develop various software tools to manage large clusters (bios upgrades, disk replication, booting, monitoring and diagnostics, etc.). This sharing will continue into the operational phase of the large facilities.

	FY 2004	FY 2005	FY 2006
CPUs/Node	2	4	4
Gflops/Node	3.2	8.0	10.0
Nodes	192	384	512
Link Bandwidth (MB/s)	300+300	$2\times$(400+400)	$2\times$(500+500)
Link Latency μsec	6	5	4
Performance (Tflops)	0.6	2.5	4.5
Hardware Cost ($M)	0.7	2.5	3.2
$/Mflops	1.2	0.9	0.7

Table 2: Cluster Parameters

3.4 Distributed Facility Integration

The multiple machines (QCDOCs at BNL and Columbia, and various generations of clusters at FNAL and JLab) will be presented to the user as a coherent computing resource. Users will be able to submit jobs to a particular resource based upon knowledge of application dependent advantages, or to a virtual batch queue which will deploy the job to the platform which can execute it the soonest, constrained by allocations and other policy.

Input and output files will be managed using data grid software, particularly leveraging the experiences and developments of the Particle Physics Data Grid Collaboratory SciDAC project in which JLab is an active participant, as well as developments coming from other groups within the Global Grid Forum. In particular, there will be a location independent global name space for files, so that files can be referenced and retrieved from any site.

Batch jobs requiring large input files will tend to execute where the source data is located, but the system will be capable of automatically migrating data as appropriate, tracking all replicas for future use.

Generated data will be cataloged in a global, searchable database, which can be queried both through a web portal or through programming interfaces. Tools will be available to automatically import the meta-data associated with a particular data set, such as a lattice configuration, ensuring that valuable data is accessible to a large community.

4 SciDAC Software Infrastructure Project

The goal of the SciDAC software infrastructure project is to create a unified programming environment that will enable the U.S. lattice community to achieve very high efficiency on the multi-terascale computer architectures targeted in our plan. To accomplish this a national effort is under way to develop uniform standards so that the time and energy of the lattice QCD community will not be consumed by porting and optimizing application code as the architectures evolve over time. Moreover, with our two architecture approach, it is essential to have in place the capability to quickly move applications between platforms to optimize utilization of resources and balance the load. The ability to easily share data and rapidly implement new algorithms will substantially reduce costs, and increase the speed with which we can carry out the ambitious physics program described above.

The Software Coordinating Committee provides overall leadership of the software effort. Its members are listed in Appendix A.2. The initial focus of the Committee has been to define a QCD Application Programming Interface (API) so that a uniform, highly optimized, communication and linear algebra layer is provide to the entire QCD community. As this API is adopted, the work to develop optimized implementations and to port the major application codes has accelerated. The entire purview of the Committee includes the following tasks: QCD API and code library; optimized network communications; optimized Lattice QCD kernels; application porting and optimization; data management and documentation; and executions environment. Detailed information on the project and access to published standards and code can be found at the URL www.lqcd.org. The goal is to have the infrastructure in place by the end of 2003, with essential parts of it deployed as early as this summer.

4.1 QCD Application Interface

The QCD API has been the central focus of the Software Co-ordinating Committee, as it provides the overall software matrix within which much, but not all, of the other tasks are found.

QCD-API Level Structure

Level 3
Dirac Operators, CG Routines, etc. (Critical sections of application codes)

QDP_xxx	Level 2
Data Parallel QCD Lattice Operations (overlapping Algebra and Messaging) e.g. A = SHIFT(B, mu) * C; Global sums, etc	
Lattice Wide Linear Algebra (No Communication) e.g. A = B * C	Lattice Wide Data Movement (Pure Communication) e.g. Atemp = SHIFT(A, mu-dir)

QLA_xxx	Level 1	QMP_xxx
Single Site & Vector Lin Alg API e.g. SU(3), Dirac algebra, etc.		Message Passing API (Lattice Geometry into Network)

As indicated in the table, the design for the QCD API is as follows. The application code sits above level 2 of the API. level 3 represents plug–ins to the application code of specially optimized subroutines for computationally intensive segments of the calculation, such as the conjugate gradient solvers, which may call level 1 code directly. The middle layer, level 2, is the fundamental data-parallel interface presented to the application programmer (function names **QDP_xxx**). To shield the application programmer from architectural dependencies, this level 2 API is built on low level, level 1, linear algebra and message passing API's (function names **QLA_xxx** and **QMP_xxx** respectively). In general, assembly level routines are restricted to a small kernel in level 1 so that the remaining C or C++ code can be quickly ported to new platforms. The code is open software so that when optimization or flexibility requires it, the application programmer is allowed to call level 1 routines directly.

Some highlights of the present projects under way in the SciDAC software infrastructure for QCD are the following. The message passing API design has been completed and both C and C++ interfaces have been implemented on top of MPI for testing. Implementations optimized for the QCDOC and for Myrinet (on top of GM) are nearing completion, so this critical layer will be ready for use by applications for both architectures in the summer of 2002. The level 1 linear algebra routines are being generated in pure C at first, and the basic data parallel design document for level 2 has been largely written and agreed on. Plans are being put into place to fully port the Lattice Hadron Physics Collaboration's code base (SZIN *et. al.*) to sit on top of level 2. The development of quark inverters highly optimized for the QCDOC is nearing completion, and the initial benchmarks on the simulator already are within the design range of the machine. Thus, in a short time good progress on the QCD-API has been made.

4.2 Broader Infrastructure Project

There are a host of other tasks pertinent to the full range of SciDAC software infrastructure components listed above, which are in various stages. For example JLab and FNAL, with collaborators in a number of universities, are optimizing code for the Pentium 4 using its SSE instructions, and software tools are being assembled and adapted for improving the execution environment for large clusters. Computer scientists are developing tools that can instrument scientific applications. For instance, a particular tool, SvPablo, has been used to instrument the MILC code to assess the impact of new architectures on performance. Recently, work has begun in the areas of data management and documentation. The Software Coordinating Committee is studying the use of XML, both as a tool to control data file I/0 and to locate relevant data using web based search methods employing SQL. However, the Committee is taking a pragmatic approach. It seeks to avoid "reinventing the wheel" and to always maintain a smooth transition from existing practices, so that users can focus on the physics, and the valuable legacy of 25 years of software development in the USA lattice QCD community is leveraged. Our test of successful software infrastructure for QCD is that it increases physics productivity and promotes the voluntary adoption of more uniform standards by the community.

5 Physical Facilities

5.1 QCDOC Facilities

An important part of our strategic plan is the staged deployment of the QCDOC computers. This will be accomplished in three steps:

Prototype: In this first stage a prototype machine of a few hundred nodes and performance of ≈ 100 Gflops will be constructed and used to run benchmark code to verify both the QCDOC design and its utility to support the physics goals of the collaboration. This prototype stage is also

Item	Cost
FY 2003	
1.5 Teraflops development machine	\$1500
BNL construction costs	
Site preparation	\$250
Technicians/Facility Management	\$150
Total	\$400
Components for 5.0 Teraflops BNL machine	\$600
Total	**\$2500**
FY 2004	
5 Teraflops QCDOC construction	\$4400
Operating costs	
Personnel (including fringe, support and burden):	
0.25 FTE Facilities Manager	\$52
1.5 FTE Technician	\$234
1 FTE System Administrator	\$186
1 FTE Software Development	\$186
1 FTE Software Support/Training/Librarian	\$163
Total personnel	\$821
Electricity	\$100
Disk and server upgrades, supplies	\$79
Total operating	\$1000
Total	**\$5400**
FY 2005	
Total	**\$1000**
FY 2006	
Total	**\$1000**

Table 3: Budget required for QCDOC construction and support (all numbers are given in thousands of dollars).

an important part of the hardware development effort and is already funded by the DOE. This machine should be available in the late summer or early fall of 2002.

Development: This 1.5 Teraflops machine will be constructed at Columbia as soon as the hardware design has been verified and the benchmarks above have been successfully passed. A combination of Columbia and BNL personnel as well as the SciDAC team at BNL will provide user support, allowing this very significant resource to be used by members of the U.S. community for physics projects selected by the collaboration's Scientific Program Committee. This machine will provide internationally competitive resources to the U.S. lattice community and a very effective test of the ability of the QCDOC hardware and software to advance the physics goals of the collaboration. We anticipate that construction of this machine will begin late in 2002 and that the computer will be available for general use in the Spring of 2003. The cost of the machine is \$1.5M. It can be built through Columbia without overhead charges.

Production: By constructing a machine of at least 5 Teraflops sustained capacity, this production stage will provide the U.S. community with a resource that is comparable to the 5 Teraflops machines that are now funded for the UKQCD collaboration, the RIKEN BNL Research Center (both also QCDOC machines), the 5 Teraflops machine being planned in Germany and the $\approx$ 5 Teraflops APE NEXT computers anticipated in Italy. We would go forward with this major step only after the development machine has shown substantial success. The construction of this $5M machine would be started late in 2003 and should be available for full-scale physics use early in 2004. Critical to realizing the full physics potential of this resource is the provision of adequate staffing to maintain and evolve flexible and efficient application software, appropriate run-time support and a convenient batch and interactive user environment.

The required funding profile, for the period beginning in FY 2003 is presented in Table 3. In addition to funds for the development machine, the FY 2003 budget includes BNL site preparation costs and a 1.0 FTE technician and 0.25 FTE facilities manager who will participate in the assembly of the development machine and oversee the site preparation and initial installation at BNL. Significant savings are realized in site preparation at BNL since the Lab has already budgeted $1.7M for renovation and plant improvements to support the installation of both the 5 Tflops machine described here and the 5 Tflops RIKEN machine that is planned during the same period. (Further savings result in the later operating costs since the required 0.5 FTE facility manager and two dedicated technicians have been divided between these two projects.)

The FY 2004 budget includes the remaining construction funds for the 5 Tflops machine. This procurement will be carried out through Columbia and will not incur overhead. The remaining operating costs of $1M in 2004 will provide the personnel needed to support large-scale community use of this substantial resource. This level of system management, user support and software maintenance and development are essential functions, required for this facility to realize its full physics potential. These operating costs continue for the remainder of the period covered by this plan.

5.2 Cluster Facilities

The clusters will be installed within computing environments at national labs with the necessary network connections, tertiary storage systems, and staff expertise. The required network bandwidth and storage capacity (silo) will be much smaller than the laboratories' experimental programs so there are no large related up-front costs associated with a new site. Further, both JLab and FNAL have considerable experience in running clusters of commodity nodes as part of their experimental physics programs.

The FNAL facility is located in the New Muon Laboratory, in the fixed target area. Computer ready space (with raised floors, electrical supply, and air conditioning) has been set aside for lattice QCD clusters in the electronics areas of completed experiments. An 80 node cluster has been in operation in this area since the end of 2000. It uses dual Pentium III nodes connected with a Myrinet 2000 high performance network. It is in use for physics by members of the FNAL, MILC, and Cornell lattice collaborations. Tertiary storage of 20 terabytes is currently available for lattice users on FNAL's Enstore tape robots. Long term tape storage of 2.4 petabytes for experiments will be installed in 2002, on which future lattice needs will be a small perturbation.

Under the SciDAC grant, this cluster will be upgraded to around 512 nodes, which will be administered under SciDAC. This year, it will be upgraded to around 256 nodes using a Myrinet 2000 network, and almost certainly dual Pentium 4 nodes. In 2003, the facility will be augmented with an additional 256 node cluster (as funding permits). A likely candidate for the network will be a

Item	Cost
FY 2004	
Cluster Hardware	\$700
Operating Costs	
Personnel (including fringe, support and burden):	
0.25 FTE Facilities Manager	\$50
1.0 FTE Systems Administrator	\$140
1.0 FTE Software Staff/User Support	160
Site preparation	\$40
Electricity	\$40
Disk cache, tapes, spares, supplies	\$20
Total	**\$1,150**
FY 2005	
Cluster Hardware	\$2,600
Operating Costs	
Personnel (including fringe, support and burden):	
0.5 FTE Facilities Manager	\$100
1.0 FTE Systems Administrator	\$140
1.0 FTE Hardware Technician/Sysadmin help	\$120
2.0 FTE Software Staff/User Support	320
Site preparation	\$400
Electricity	\$170
Disk cache, tapes, spares, supplies	\$100
Total	**\$3,950**
FY 2006	
Cluster Hardware	\$3,100
Operating Costs	
Personnel (including fringe, support and burden):	
0.5 FTE Facilities Manager	\$100
1.0 FTE Systems Administrator	\$140
1.0 FTE Hardware Technician/Sysadmin help	\$120
2.0 FTE Software Staff/User Support	320
Site preparation	\$50
Electricity	\$250
Disk cache, tapes, spares, supplies	\$100
Total	**\$4,180**

Table 4: Cost per year for each cluster site. (All numbers are given in thousands of dollars).

256 port Myrinet switch, if it appears on the market as announced.

Space for 2,800 narrow (2U) nodes is currently set aside for lattice QCD clusters. More space may be obtained in the New Muon building as required. The present cluster was constructed using around two FTEs provided by FNAL (pre-SciDAC). The staff is being increased to three FTEs under SciDAC and would rise to four or five in the long term, as outlined below.

The JLab facility will be housed within the lab's computer center. Like FNAL, JLab will deploy as part of the SciDAC lattice computing project 512 nodes during 2002 and 2003 in existing computer room space (raised floors, amply power, A/C). The first half of this system will be Pentium 4 nodes with Myrinet 2000, and be operational by summer 2002. The clusters proposed for 2004 will also be installed in this same room. The larger clusters in 2005 and 2006 are to be installed in a planned expansion of CEBAF Center, which will house a much enlarged computer room (again with raised floors, ample power and air conditioning).

The long–term expenses for future facilities beyond the computers themselves will consist of disk space (cache), storage media, additional tape drive(s), power, site preparation, and staff. These costs are detailed in Table 4.

By FY 2005, the FNAL and JLAB staffs will consist of 1 system administrator, 1 technician (who can assist with system administration), 2 system programmers (who will develop and maintain low level, high performance, application libraries and cluster operations software), and a one-half time facility manager. Overlapping administration and system programming tasks will be shared between FNAL and JLab. This lean staffing is possible only because network security, account management, and a number of other system administration functions will leverage the existing staff at each site.

6 Funding Profile Summary

Our immediate objective is to construct the hardware and software infrastructure that the U.S. lattice gauge theory community needs to carry out terascale simulations of QCD. Work on software development is in progress under our current SciDAC grant. The SciDAC grant also provides funds to build prototype clusters, and the HEP Division is supporting development work for the QCDOC. The purpose of this document is to describe our plans for the construction of the terascale computers required to reach our scientific goals, and to set out the funding profile needed to reach these goals in time to support the experimental programs in high energy and nuclear physics, and to maintain the U.S. leadership position in lattice QCD.

We know the computing resources needed to carry out the research program outlined in Section 2 from extensive experience. It is imperative that we act expeditiously to construct these resources because many of the experiments at which our work aims are in progress or planned for the near future. Furthermore, our colleagues in Europe are in the process of acquiring computing resources very similar in scope to those we propose. The APE project expects to construct terascale class machines in 2003 for laboratories and universities in a number of European countries, DESY has announced plans to obtain a 20 Tflops (peak) APEnext in 2004. In addition, the UKQCD collaboration expects to install a 5 Tflops QCDOC in 2003. Our Japanese colleagues already have computers with a throughput of close to 1 Tflops, and have ambitious plans for upgrading their facilities. Thus, if U.S. lattice gauge theorists are to play a significant role in the major advances expected in this area over the next decade, we must act now.

The construction and operation of a distributed computational computing facility for lattice QCD will be a long term project. We propose to ramp up funding during the 2003–2005 fiscal years, leveraging funds that have already been committed by the SciDAC Program, HEP and the participating laboratories. We propose that funding plateau in fiscal year 2006 and beyond. The table below shows the proposed funding profile, and the computing power it is expected to yield. These figures show the total cost of the project without distinguishing between new funds and continuation of current investments.

We expect to receive approximately $2.0M in 2003 through the SciDAC Program for our current software and cluster development projects. We request an additional $2.5M during this year of which $1.5M will be used to construct a QCDOC at Columbia University with a throughput in excess of 1 Tflops. This machine will allow U.S. lattice gauge theorists to make the earliest possible start on terascale computations. The remaining $1.0M we request for 2003 will be used for site preparation and the purchase of components for a 5 Tflops QCDOC which is to be built at BNL in 2004. In 2004 we request $7.7M, $5.4M for the construction and operation of the BNL QCDOC, and a total of $2.3M for the cluster projects at FNAL and JLab. The emphasis on the QCDOC during 2003 and 2004 is motivated by our expectation that it will be the most cost effective large scale platform during this time period. At the same time, it is necessary to invest enough in clusters to maintain the vitality of that part of the program, especially since they may be the preferred machines for the development of new algorithms and for some of our most complex calculations.

By 2005 we assume that the cost effectiveness of clusters will surpass that of the QCDOC, so our emphasis will switch to them. We propose a total budget of $9.0M in that year, of which $1.0M will pay for the operation of the BNL QCDOC, and the remainder for the construction and operation of large clusters at FNAL and JLab. The funding for 2006 is proposed to rise to the steady state level of $9.4M. In 2007 and beyond we plan to continue our two track approach, using advances in technology to steadily increase our computing capabilities with stable budgets.

In the table below we show the proposed funding profile for the fiscal years 2003 through 2006. The first three columns indicate the requested funding for each year in millions of dollars for the SciDAC project, the QCDOC and the clusters. We include both hardware and operating costs in these figures. The fourth column is the total funding for the year, again in millions of dollars, and the fifth column shows the expected total computing power in Tflops arising from this project and our current SciDAC grant. As above, the performance estimates are based on Wilson quarks.

Year	SciDAC	QCDOC	Clusters	Total	Tflops
2003	$2.0M	$2.5M		$4.5M	2
2004		$5.4M	$2.3M	$7.7M	8
2005		$1.0M	$8.0M	$9.0M	14
2006		$1.0M	$8.4M	$9.4M	22

Table 5: Funding Profile

A Appendices

A.1 Some Lattices QCD Successes

In this appendix we list some of the major achievements of lattice gauge theory, as well as near term prospects for future ones.

Key Results	**Date(s)**	**Computer Resources**
"Proof" of quark confinement	1974	None–Strong coupling expansion of the lattice QCD action
First simulation–demonstration that Monte Carlo simulation of field theory is feasible	1979	CDC-7600
Temperature of the chiral symmetry restoration phase transition to a quark–gluon plasma ($T_c = 150$ MeV)	Late 1980s	A few Gflops-years
Accurate light–hadron spectrum (quenched)	Early 1990s	10's of Gflops-years
Accurate light–hadron spectrum (unquenched)	Ongoing	100's of Gflops-years (so far)
Glueball spectrum–surprising answer was that spectrum starts near 2 GeV, much above the canonical 1 GeV	Mid 1990s	Many calculations, each of which used a few Gflops-years
Determination of $\alpha_s(M_Z)$ from the Υ spectrum to a precision of 4-5%; m_b and m_c to about the same precision. This is a good as the best results obtained from jet analyses	1995	10's of Gflops-years
Determination of $\alpha_s(M_Z)$, etc., to 1-2%, beyond the capability of any jet analysis	Near future	100's of Gflops-years (at least)

B and *D* meson decay constants and $\bar{B}B$ and $\bar{K}K$ mixing matrix elements, to a precision of about 10-20% the f_B and f_D leptonic decay constants came out much higher than expected (f_{D_s} is now confirmed). These results are regularly used in *B* physics analyses	Late 1990s	100's of Gflops-years
B and *D* meson matrix elements to a precision of about 5%; exclusive semi–leptonic *B* decay form factors. These are key inputs in the determination of CKM matrix elements from *B*–factory measurements	Near future	Teraflops-year(s)
ε'/ε hadronic matrix elements measuring CP violation	Near future	Teraflops-year(s)
Nucleon form factors and moments of quark distributions to a precision of 20-30%	Early 2000's	100's of Gflops-years
Nucleon form factors, moments of quark, gluon, and generalized parton distributions to a precision of 5%. These are key quantities measured at SLAC, Fermilab JLab, Bates, and RHIC-spin	Fairly near future	Teraflops-years
Quark–gluon plasma equation of state	Fairly near future	Teraflops-years

A.2 SciDAC Committees

In this appendix we list the members of the committees that are providing leadership for the project.

Executive Committee

Richard Brower	Boston University
Norman Christ	Columbia University
Michael Creutz	Brookhaven National Laboratory
Paul Mackenzie	Fermi National Acccelerator Laboratory
John Negele	Massachusetts Institute of Technology
Claudio Rebbi	Boston University
Stephen Sharpe	University of Washington
Robert Sugar (Chair)	University of California, Santa Barbara
William Watson, III	Thomas Jefferson National Accelerator Facility

Scientific Program Committee

Peter Lepage	Cornell University
Robert Mahwinney	Columbia University
Colin Morningstar	Carnegie Mellon University
John Negele	Massachusetts Institute of Technology
Claudio Rebbi (Chair)	Boston University
Stephen Sharpe	University of Washington
Doug Toussaint	University of Arizona
Frank Wilczek	Massachusetts Institute of Technology

Oversight Committee

Steven Gottlieb (Chair)	Indiana University
Anna Hasenfratz	University of Colorado
Gregory Kilcup	Ohio State University
Julius Kuti	University of California, San Diego
Robert Pennington	National Center for Supercomputer Applications
Ralph Roskies	Director, Pittsburgh Supercomputer Center
Terry Schalk	University of California, Santa Cruz

Software Co-ordinating Committee

Richard Brower (Chair)	Boston University
Carleton DeTar	University of Utah
Robert Edwards	Thomas Jefferson National Accelerator Facility
Donald Holmgren	Fermi National Accelerator Laboratory
Robert Mawhinney	Columbia University
Celso Mendes	University of Illinois
William Watson, III	Thomas Jefferson National Accelerator Facility

A.3 Senior Personnel

In this appendix we list the senior personnel who are participating in this project. They comprise nearly all of the senior lattice gauge theorists in the United States, as well as senior computer scientists and engineers who have agreed to join in the effort.

Claude Bernard	Washington University
Tanmoy Bhattacharya	Los Alamos National Laboratory
Richard Brower	Boston University
Thomas Blum	Brookhaven National Laboraoty
Matthias Burkardt	New Mexico State University
Shailesh Chandrasekharan	Duke University
Dong Chen	T.J. Watson Laboratories, IBM
Jie Chen	Thomas Jefferson National Accelerator Facility
Norman Christ	Columbia University
Michael Creutz	Brookhaven National Laboratory
Thomas DeGrand	University of Colorado
Carleton DeTar	University of Utah
Shao-Jing Dong	University of Kentucky
Zhihua Dong	Columbia University
Terrence Draper	University of Kentucky
Patrick Dreher	Massachusetts Institute of Technology
Anthony Duncan	University of Pittsburgh
Robert Edwards	Thomas Jefferson National Accelerator Facility
Estia Eichten	Fermi National Accelerator Laboratory
Aida El-Khadra	University of Illinois, Urbana
Rudolf Fiebig	Florida International University
Alan Gara	T.J. Watson Laboratories, IBM
Steven Gottlieb	Indiana University
Rajan Gupta	Los Alamos National Laboratory
Anna Hasenfratz	University of Colorado
Urs Heller	Florida State University
James Hetrick	University of Pacific
Donald Holmgren	Fermi National Accelerator Laboratory
Xiangdong Ji	University of Maryland
Gregory Kilcup	Ohio State University
Joseph Kiskis	University of California, Davis
John Kogut	University of Illinois, Urbana
Julius Kuti	University of Californa, San Diego
Andreas Kronfeld	Fermi National Accelerator Laboratory
Frank Lee	George Washington University
Peter Lepage	Cornell University
Keh-Fei Liu	University of Kentucky
Paul Mackenzie	Fermi National Accelerator Laboratory
Robert Mawhinney	Columbia University

Celso Mendes	University of Illinois, Urbana
Colin Morningstar	Carnegie Mellon University
John Negele	Massachusetts Institute of Technology
Shigemi Ohta	KEK and Riken BNL Research Center
Robert Pennington	National Center for Supercomputer Applications
Donald Petravick	Fermi National Accelerator Laboratory
Andrew Pochinsky	Massachusetts Institute of Technology
Claudio Rebbi	Boston University
Ronald Rechenmacher	Fermi National Accelerator Laboratory
Daniel Reed	University of Illinois, Urbana
David Richards	Thomas Jefferson National Accelerator Facility
Stephen Sharpe	University of Washington
Junko Shigemitsu	Ohio State University
James Simone	Fermi National Accelerator Laboratory
Donald Sinclair	Argonne National Laboratory
Amarjit Soni	Brookhaven National Laboratory
Robert Sugar	University of California, Santa Barbara
Eric Swanson	University of Pittsburgh
Harry Thacker	University of Virginia
Doug Toussaint	University of Arizona
Steven Wallace	University of Maryland
William Watson, III	Thomas Jefferson National Accelerator Facility
Tilo Wettig	Yale University
Uwe-Jens Wiese	Massachusetts Institute of Technology
Walter Wilcox	Baylor University