I. Research Objectives

We propose to continue the broad research program in Quantum Chromodynamics (QCD) that we have been engaged in for a number of years. This research addresses fundamental questions in high energy and nuclear physics, and is directly related to major experimental programs in these fields. It includes studies of the mass spectrum of strongly interacting particles, the weak interactions of these particles, and the behavior of strongly interacting matter under extreme conditions.

The Standard Model of High Energy Physics consists of two quantum field theories: the Weinberg-Salam Theory of the electromagnetic and weak interactions, and QCD, the theory of the strong interactions. It has been enormously successful in explaining a wealth of data produced in accelerator and cosmic ray experiments over the past thirty years. However, our knowledge of the Standard Model is incomplete because it has been difficult to extract many of the most interesting predictions of QCD, those that depend on the strong coupling regime of the theory, and therefore require non-perturbative calculations. Although there is little doubt that QCD is the correct theory of the strong interactions, non-perturbative QCD calculations are crucial for testing the weak interaction part of the Standard Model: In the absence of such calculations the strong effects completely obscure the weak physics one is trying to study. At present the only means of carrying out non-perturbative QCD calculations from first principles and with controlled errors is through large scale numerical simulations. These simulations are needed to obtain a quantitative understanding of the physical phenomena controlled by the strong interactions, to determine a number of the basic parameters of the Standard Model, and to make precise tests of the Standard Model's range of validity. Despite the many successes of the Standard Model, it is believed by high energy physicists that to understand physics at the shortest distances a more general theory, which unifies all four of the fundamental forces of nature, will be required. The Standard Model is expected to be a limiting case of this more general theory, just as classical mechanics is a limiting case of the more general quantum mechanics. A central objective of the experimental program in high energy physics, and of lattice QCD simulations, is to determine the range of validity of the Standard Model, and to search for new physics beyond it. Thus, QCD simulations play an important role in efforts to obtain a deeper understanding of the fundamental laws of physics.

Several years ago we used our NRAC allocations to develop an improved action (improved discretization) for lattice QCD [1], which significantly increases the accuracy of our simulations for a given amount of computing resources. This improved Asqtad action, coupled with increases in computing resources available to lattice gauge theorists, and major improvements in the performance of our code, have enabled us to make important progress in our research. The bulk of the computing resources in any lattice gauge theory calculation go into the generation of gauge configurations, which are snapshots of the ground state of QCD. These configurations are saved and then used to calculate a wide variety of physical quantities. Over the last several years, we have generated an extensive library of gauge configurations with the improved Asqtad action. It is necessary to work with a range of lattice spacings and light quark masses in order to perform extrapolations of physical quantities to the continuum and chiral limits. A table showing the current status of our configuration generation is contained in the attachment titled *Impact of the Asqtad Gauge Configurations*. During the coming year we propose to use NSF and DOE resources to generate additional configurations at a lattice spacing of 0.06 fm, the smallest with which we have worked to date. These configurations will significantly reduce errors in a wide variety of physical quantities of importance to the experimental programs in high energy and nuclear physics.

The level of accuracy we seek in our calculations is determined in large part by our studies of the weak interactions of strongly interacting particles. An international effort is in progress to determine the elements of the Cabibbo-Kobayashi-Maskawa (CKM) matrix, which determines how quarks, the fundamental strongly interacting particles, couple to the weak interactions. In most cases both an accurate experiment and lattice calculation are needed. By determining individual elements of the CKM matrix from different experiments and lattice calculations, one can place powerful constraints on the range of validity of the Standard Model. In almost all cases, the precision of the determination of the CKM matrix element is limited by the uncertainties in the lattice calculations, rather than in the experiments. Our objective is to bring the lattice errors down to, or below, the experimental ones. The importance of doing so can be seen from the fact that a significant fraction of the $750,000,000 per year that the United States spends on experimental high energy physics is devoted to the study of the weak decays and mixings of strongly interacting particles. It is clear that to fully capitalize on this investment, the lattice calculations must keep pace with the experimental measurements.

In order to maximize the physics output from the large investment of computing resources in the generation of Asqtad gauge configurations, we are making them available to other lattice gauge theorists for their research. In the attachment titled *Impact of the Asqtad Gauge Configurations* we list the many physical phenomena being studied with these configurations by other groups, and provide references to the papers that have resulted from this work. In order to demonstrate that we do have control over all sources of systematic errors, we and several other groups making use of our configurations have jointly calculated a set of quantities whose values are well known experimentally [2]. Some of our results are shown in Figure 1 of the Progress Report, which is contained in an attachment. In each case the calculations agree with experiment within statistical and systematic errors of 3% or less. Among the more recent validations of our approach have been our determination of the CKM matrix element V_{us} [3], and determination of the strong coupling constant [4] by the HPQCD and UKQCD Collaborations. In both cases the results were in agreement with, and had errors comparable to the world averages obtained by other methods. There have also been predictions that were subsequently confirmed by experiment. They include calculations of the leptonic decay constants of the D and D_s mesons [5] and the form factors for the semileptonic decays of these particles [6] in joint work by our group and the Fermilab Lattice Collaboration, as well as the determination of the mass of the B_c meson by the Fermilab Lattice and UKQCD Collaborations [7]. These calculations have have been featured in articles in *Nature* [8], *Science* [9], *Physics Today* [10], the CERN Courier [11] and FermiNews Today [12]. They indicate that we are in a position to determine a wide range of physical quantities of great importance to the phenomenology of high energy and nuclear physics.

In the attachment titled *Progress Report*, we review in greater detail the research we have carried out since submitting our last LRAC proposal, and show figures illustrating some of our recent work. Results from most of the projects discussed in the Progress Report will be presented at the international conference on lattice gauge theory, Lattice 2006, in July.

In Section II, we describe our code and recent work that has improved its performance. We also discuss the integration into our code of a new class of algorithms for generating gauge configurations [13] that will significantly accelerate our work. In Section III we give the justification for our resource request. We request a one year allocation with a total of 11,136,300 service units. A breakdown of our request by projects and computing platforms is given in the attachment titled *Resource Request by Projects*. Our request is to support the work of six principal investigators, their postdoctoral research associates and graduate students. Thus, it comes to 1,856,050 service units per principal investigator. Furthermore, as noted in the attachment *Impact of the Asqtad Gauge Configurations*, our configurations are being used by many other physicists, so the impact of our allocations goes well beyond our own work. A description of our local computing environments, current and pending allocations of supercomputing time, and the qualifications of the principal investigators is given in the attachment titled *Computer Resources*. A complete list of publications of our collaboration, the current members of our group, and the vita of the principal investigators are provided in the attachment *Publications and People*. In the remainder of this section we describe our research objectives in greater detail.

Generation of gauge configurations with improved gauge and quark actions: The ensembles of gauge configurations that we have generated are the basis for a wide variety of physics calculations by our group and by others; but, as is discussed below, in order to determine some of the most interesting quantities we are studying to an accuracy required by experiment, it is necessary to push our simulations to smaller lattice spacings and quark masses. We therefore propose to continue to generate gauge configurations using resources provided by the NSF and DOE.

Configurations are being generated with three flavors of dynamical sea quarks: up, down and strange. These three flavors of quarks are the only ones with masses comparable to the energy scale of QCD. They are therefore the only ones light enough to have a significant impact on realistic calculations. We are working with a one–loop Symanzik improved gauge action and the Asqtad staggered quark action [1]. Both the gauge and quark actions have all lattice artifacts removed through order a^2 (a is the lattice spacing) at the tree level, and are tadpole improved. So, the leading discretization errors are of order $a^2/log(a)$.

We take the masses of the up and down quarks to be equal, which has a negligible effect ($< 1\%$) on isospin-averaged quantities. The average value of the up and down quark masses, which we denote by m_l, is much smaller than other energy scales in QCD, and it is currently too expensive to perform simulations at its physical value. Instead one works with a range of values for m_l which are small enough so that one can perform extrapolations to its physical value with the aid of chiral perturbation theory [14]. On the other hand, the strange quark mass is heavy enough so that we can perform simulations at its physical value m_s, and we attempt to keep it fixed at this value as we vary the lattice spacing and m_l.

We currently have ensembles at four different lattice spacings, $a \approx 0.18$ fm, 0.15 fm, 0.12 fm and 0.09 fm with the light quark masses in the range $0.1 m_s \leq m_l \leq m_s$, and have begun runs with lattice spacing $a \approx 0.06$ fm and light quark masses $0.4 m_s$ and $0.2 m_s$ on $48^3 \times 144$ lattices. One of our major objectives for the coming year is to complete these two runs, as they will substantially reduce the errors in extrapolations of physical quantities to the continuum limit. Indeed, the leading discretization errors in the five ensembles with lattice spacings 0.18 fm to 0.06 fm are in the ratio 1.00/0.63/0.36/0.18/0.07, so in going from $a \approx 0.09$ fm to $a \approx 0.06$ fm, these errors will be reduced by a factor of nearly 2.6.

We are working with the Kogut-Susskind (staggered) formulation of lattice quarks because at present they enable dynamical simulations with significantly smaller quark masses than other formulations. A problem with Kogut–Susskind quarks is the breaking of taste symmetry. ("Taste" refers to the different ways one can construct the same physical particle in the Kogut-Susskind formalism). Although particles with different tastes become identical in the continuum limit, their masses can differ at finite lattice spacing. However, the Asqtad action has significantly better rotational and taste symmetry than the conventional Kogut–Susskind action [15]. (The rapid decrease of taste symmetry violation with decreasing lattice spacing is illustrated in Figure 10 of the Progress Report). Furthermore, one of us (CB) has developed techniques, based on the pioneering work of Lee and Sharpe [16] for incorporating the effects of taste symmetry breaking into chiral perturbation theory [17, 18, 19], thereby significantly improving our ability to make extrapolations to the chiral and continuum limits.

Another difficulty with staggered quarks is that each Kogut-Susskind field represents four tastes of quarks that become degenerate only in the continuum limit. So, in order to have one quark per flavor in the continuum, one uses the fourth root of the Kogut-Susskind determinant in performing simulations. It has been suggested that taking the fourth-root of the determinant at finite lattice spacing might give rise to unphysical non-localities that persist in the continuum limit [20], or to other unphysical effects [21]. However, the agreement of existing results with experiment [2, 3, 4, 5, 6, 7], as well as a growing body of direct studies of these issues [22, 23, 24, 25], give us confidence that no fundamental problem exists.

Physics of light pseudoscalars: Lattice computation of the properties of light pseudoscalar mesons (*i.e.*, π, K and η mesons) offers a unique opportunity to check our lattice methods to high ($\approx$ 2 to 3%) precision, and to calculate phenomenologically important physical quantities that are difficult or impossible to obtain with controlled errors by other methods. The advantages of this system stem firstly from the fact that we are able to compute quantities such as the π and K masses and decay constants at fixed lattice spacing and (larger than physical) quark mass with extremely high statistical accuracy: 0.1% to 0.7%, depending on the quark masses. Secondly, the dependence of these quantities on quark masses is governed by the formalism of chiral perturbation theory [14]. This formalism has been extended to the "partially quenched" case [26], where valence and sea quark masses are different, and, further, to the case of staggered quarks with discretization errors included ("staggered chiral perturbation theory" [16, 17, 18, 19] — SχPT). The latter allows us to fit the lattice data accurately, including the effects of $\mathcal{O}(a^2)$ lattice spacing errors, and then make a controlled chiral extrapolation (extrapolation to the physical value of m_l) followed by a controlled continuum extrapolation.

Using the above method, we have computed the leptonic decay constants of the π and K mesons, f_π and f_K, with total errors of under 3% (see the attached Progress Report). The results agree with experiment at this level, providing good evidence that we understand and can control our errors. A calculation of f_K/f_π enables one to determine the CKM matrix element V_{us}. Our current value is $|V_{us}| = 0.2223^{+26}_{-14}$, which is consistent with the 2004 world average value $|V_{us}| = 0.2200(26)$ determined by the Particle Data Group from experiment and non-lattice theory [27], as well as a recent update using new experimental results: $|V_{us}| = 0.2262(23)$ [28].

The same computational method, coupled with perturbative calculation of the mass renormalization constant, has allowed us to determine quark masses (the mass of the strange quark, m_s, the average mass of the u and d quarks, m_l, and the ratio m_u/m_d) and several of the Gasser-Leutwyler parameters, L_i [14]. Two of these quantities, in particular, urgently require precise lattice QCD evaluations: Uncertainty in m_s severely limits the theoretical precision of various phenomenological studies, including the important CP-violating quantity ε'/ε [29]. Determination of m_u/m_d addresses the long standing proposal [30] that $m_u = 0$. A mass-less up quark could have solved the “Strong CP Puzzle” [31]; however, our results rule that possibility out at the 10σ level. Our current results [2, 3, 32] for decay constants, quark masses, and L_i are presented in the Progress Report.

To go further, we need to have more configurations on finer lattices. We estimate that the configurations with lattice spacing $a \approx 0.06$ fm will allow us to reduce the systematic errors on f_π and f_K to 2% or better, which would be an important milestone for lattice QCD calculations. We expect corresponding improvements in other physical quantities of interest. In particular, our evaluation of $|V_{us}|$ should become more accurate than the current world average.

Weak decays of particles containing heavy quarks: At the SLAC and KEK B-factories, and at Fermilab, a concerted experimental effort is under way to determine elements of the CKM matrix through studies of the mixings and decays of B mesons. In addition, the properties of D mesons are being measured to high accuracy in the CLEO-c Program at Cornell. The CLEO-c measurements can provide powerful constraints on the CKM matrix, both directly (on matrix elements involving c quarks) and indirectly (through implications for b quark systems). Indeed, one hopes that by tightly over–constraining these CKM matrix elements, the range of validity of the the Standard Model will be determined, and new physics beyond it will be discovered. However, the experimental results do not in general determine the CKM parameters without lattice calculations of the effects of the strong interactions.

Our group and the Fermilab Lattice Collaboration are involved in an extensive joint study of the decays of pseudoscalar mesons with one light and one heavy quark. The main objects of our work are B (with a heavy b quark and light u or d antiquark), B_s (b quark and s antiquark), D (c quark and u or d antiquark) and D_s (c quark and s antiquark). We are studying both leptonic and semileptonic decays. Strong interaction effects in leptonic decays are characterized by the decay constants f_B, f_{B_s}, f_D and f_{D_s}. Semileptonic decays are characterized by various form factors $F(q^2)$, where q is the momentum transferred to the leptons. CLEO-c is providing precise measurements of the D and D_s leptonic and semileptonic decays. Comparisons of lattice and experimental results offer a unique opportunity to validate our approach, and insure that we do, in fact, have full control over systematic errors. Our successful predictions of the leptonic decay constants f_D and f_{D_s} [5] and the shape and normalization of the semileptonic form factors [6] for D mesons by us and our Fermilab collaborators is a first step in this process. We are now in a position to carry out similar calculations for the b-quark system, where the corresponding quantities have not been measured, and are unlikely to be in the near future. Lattice results for the leptonic decay constants f_B and f_{B_s} along with the $B^0-\bar{B}^0$ mixing parameters B_B and B_{B_s} with comparable precision would have a major impact on the determination of the poorly known CKM matrix element V_{td} from experimental measurements of B-$\bar{B}$ and B_s-$\bar{B}_s$ mixing. Similarly, accurate lattice determinations of semileptonic form factors for $B \to \pi l \nu$ and $B \to D^* l \nu$ would significantly reduce errors on the CKM matrix elements V_{ub} and V_{cb}, respectively. The calculation of these quantities are our major goals for the coming year. This work is in progress [33], and new results will be reported at the Lattice 2006 meeting in July. Our goal in the next few years is to determine the leptonic decay constants, the mixing parameters and the corresponding semileptonic form factors to an accuracy of 5% to 10%. We are using the same techniques in our studies of B decays and mixings that are being employed for our work on the D mesons [34].

Hadron mass spectrum with the improved action: Calculation of the spectrum of the light hadrons is an essential test of lattice simulations. The nucleon and Ω^- masses are precisely known, and can be computed accurately on the lattice, making them trenchant tests of our techniques. Moreover, lattice computations can shed light on some of the open questions regarding the nature of the light hadrons. For example, the nature of the $a_0(980)$ is still somewhat controversial — to what extent is it a quark-antiquark state, and to what extent a $K-\bar{K}$ molecule? Also, the quark model assignments of many of the excited states are not well established, and lattice computations should help nail them down. Lattice calculations are also important for understanding hadrons that are not explained by the naive quark model, namely hybrids and glueballs. These hadrons, especially those with exotic quantum numbers, are an important part of the experimental program at Jefferson Laboratory.

We have customarily computed the masses of the simplest hadrons, which can be produced from a point-like operator with staggered quarks, as we have generated the gauge configurations. This “point source” spectrum gives us the masses of the π, ρ, a_0, a_1, b_1 and nucleon, together with the corresponding states containing strange quarks. We will continue to calculate the masses of these particles in parallel with our configuration generation. In addition to being important in their own right, these calculations are useful in monitoring the progress of the runs.

An accurate determination of the mass of the nucleon has been a long term problem for lattice gauge theorists. The major difficulties are the extrapolation to the physical value of m_l and contamination by excited states, such as the delta. We expect to improve the accuracy of the extrapolation through the use of our new light quark mass, fine lattice spacing configurations, coupled with a staggered chiral perturbation theory calculation of the dependence of the nucleon mass on m_l being carried out by a graduate student, Jon Bailey. In addition to better fitting and extrapolation techniques, we expect to account for delta contamination by simultaneous fitting to nucleon and delta propagators, and we hope to eliminate the contamination due to other excited states by employing new nucleon operators of the type suggested by Bailey. We will experiment with these new operators and the new operators for the other light baryons [35]. We have calculated the Δ and Ω^- masses on all of the gauge configurations generated to date, and propose to do so on the ones we create during the coming year. In addition, we plan to apply the partial quenching technique that was so important for the study of light pseudoscalars to the study of baryons. As with the pseudoscalars, having a large number of data points with different valence quark masses should allow us to determine the many parameters that appear in chiral perturbation theory, be it staggered (as under investigation by Bailey) or in the continuum [36].

High temperature QCD with three flavors of quarks: At very high temperatures one expects to observe a phase transition or crossover from ordinary strongly interacting matter to a plasma of quarks and gluons. A primary motivation for the construction of the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory was to observe the quark–gluon plasma and determine its properties. The plasma was the state of matter in the early development of the Universe, and may be a central component of neutron stars today. The behavior of strongly interacting matter in the vicinity of the phase transition or crossover is inherently a strong coupling problem, which can only be studied from first principles through lattice gauge theory calculations. Among the issues that can uniquely be addressed by lattice calculations are the nature of the transition, the properties of the plasma, including strange quark content, and the equation of state. We are engaged in a multi–year effort to carry out detailed studies of these issues using the improved Asqtad action [37]. We have just completed a study of the equation of state of high temperature QCD at zero baryon density on lattices with four and six time slices [38]. Final results will be presented at the Lattice 2006 conference, and a journal article is in preparation [39]. This is the first study with an improved action and a realistic set of quarks with such small lattice spacings. Near and slightly below T_c there are still lingering uncertainties even on lattices with six time slices ($N_t = 6$), since lattice artifacts are large there. This temperature range is a phenomenologically important one, since it is relevant to the formation of confined hadrons as the quark-gluon plasma cools. We propose further calculations during the coming year to reduce lattice artifacts in this temperature range

All of our simulations to date have been at zero baryon density; however, in heavy ion collision experiments being conducted at RHIC, the baryon density is not zero. We are therefore extending our study of the equation of state to the small baryon densities relevant to these experiments using the Bielefeld-Swansea Taylor series method [40], which is expected to give a reliable extrapolation to the experimentally relevant regime. For this study we are reusing our existing archive of zero baryon density gauge configurations. The code was brought into production this Spring, and we will report preliminary results for $N_t = 4$ at the Lattice 2006 conference. We propose to reduce lattice artifacts by extending this study to the finer Euclidean time resolution $N_t = 6$ in the coming year.

Topological Charge and Susceptibility: In QCD gauge configurations may have nontrivial topology. The topological susceptibility characterizes the degree of non-triviality. Chiral perturbation theory gives a relationship between the topological susceptibility and the quark masses for small values of the masses. In particular, as the masses vanish, the susceptibility must also vanish with a slope determined by the pion decay constant and the number of quark species. As we mentioned above, our staggered fermion formalism achieves the correct number of quarks by taking the fourth root of the fermion determinant. Thus it is important to verify that the measured slope of the vanishing susceptibility is consistent with the correct number of quarks. As we generate our ensembles of gauge configurations, we have been measuring the topological susceptibility. We will continue to do so using local resources. So far our results have been consistent with the predictions of staggered chiral perturbation theory with the correct number of quarks [41]. An update is in preparation [42].

Non-perturbative evaluation of quark masses: Quark masses are fundamental parameters of the Standard

Model, but since free quarks do not exist because of confinement in QCD, they can not be directly measured in experiments. They can be inferred indirectly, through non-perturbative calculations, such as the precise pseudoscalar mass calculations done by us. However, such computations need a renormalization factor that depends on the details of the quark action used. Up to now, for Asqtad quarks the renormalization factor has only been computed in perturbation theory. Its value changed by about 12% when going from one loop to two loops in perturbation theory, more than the estimated uncertainty due to the unknown higher order contributions. Even using the two-loop result, the uncertainty in the renormalization factor constitutes the largest systematic error in our determination of the light quark masses. There exists a well established non-perturbative method to compute the renormalization factor [43]. We have implemented this method for the Asqtad quark action and intend to compute the renormalization factor non-perturbatively, removing the largest systematic error in our quark mass values, and making ours the most accurate determination of them.

II. Code Optimization and Algorithm Development

We have developed a family of highly portable codes for the study of QCD with conventional and improved Kogut-Susskind dynamical quarks, improved Wilson (Clover) dynamical quarks, and in the quenched approximation. Compilation options allow one to choose among different computer architectures, different formulations of lattice QCD, and different algorithms for carrying out simulations. The code and a manual describing it in detail are publicly available at the URL http://www.physics.utah.edu/ ˜detar/milc. We estimate about 15 downloads of the code per month, and we field several questions per month about it.

We continue to put significant effort into improving the performance and flexibility of our code. Several of us participate in the DOE-funded SciDAC project for lattice gauge theory, which, over the past years, has been developing a software framework for the US community of lattice gauge theorists. To take advantage of the considerable optimization effort of that project, Version 7 of the MILC code inter-operates with SciDAC modules and incorporates SciDAC optimization strategies. We have already seen considerable improvement in our production code. For example, the performance of the conjugate gradient subroutine for the inversion of large sparse matrices, which uses the bulk of the floating point operations in any lattice gauge theory calculation, increased from 408 to 580 Mflop/s per processor in dual-node, single-precision production runs on the NCSA Tungsten cluster. The sparse matrix inversion figures for the NCSA Teragrid machine and the PSC Cray XT3 are 680 MFlop/s and 970 Mflop/s, respectively.

We are currently optimizing our code for the BlueGene/L. The strategy is to convert time critical subroutines to use the SciDAC QCD libraries, and then take advantage of the SciDAC effort to optimize them for the BlueGene/L. This has already been done for the two subroutines that consume the largest number of floating point operations in our work, those that perform the sparse matrix inversion and the calculation of the fermion force. The performance of these routines has been improved by a factor of approximately 2.5 over standard C code. Future work will focus on improving the SciDAC linear algebra (QLA) and message passing (QMP) libraries on the BlueGeene/L, which should provide further performance gains ranging from 15% to 20%. Code developed up to now does not effectively use the dual floating point units of the BlueGene/L; however, work is in progress to generate code that will do so. Finally, the QMP library is currently implemented on the BlueGene/L by making calls to MPI. Much of the time this results in extra copies of the data being made, which reduces performance. To avoid this problem the QMP library is being rewritten to use the low level communications hardware directly.

Simulations of staggered quarks with fewer than four degenerate flavors (or Wilson quarks with non-degenerate flavors) require a weight function that contains fractional powers of the fermion determinant, which introduces complications in the molecular dynamics algorithms commonly used in lattice QCD. In the past this problem has been handled by introducing a random noise in the fermion force, the R-algorithm [44]. Recently a class of algorithms has been developed that approximates the fractional powers of determinants using rational functions [13]. These algorithms remove the need for a random noise in the fermion force, and allow elimination of errors from the integration step size in the molecular dynamics equations via a Metropolis accept/reject step. It turns out that although the errors in observables coming from the random noise in the R-algorithm are formally of the same order in the step size as the errors from the integration of the molecular dynamics equations, their magnitude is considerably larger. This means that much larger integration step sizes can be used in the rational function algorithms than in the R-algorithm.

In the past year, we have added the technology for rational function approximations to the MILC code suite, and have begun a program of testing different variants of this set of algorithms for the lattice sizes and parameters of our current and planned large simulations. This code is currently being used in our production run on $48^3 \times 144$ lattices with $a \approx 0.6$ fm and $m_l = 0.2\,m_s$. The variant that we are currently using runs two to three times faster than the R-algorithm code at the same lattice spacing and quark masses. Furthermore, there are a number of enhancements that we plan to work on in the coming year that are likely to lead to additional gains in performance.

III. Resource Justification

A breakdown of our request by projects and computing platforms is given in the attachment *Resource Request by Projects*. We request 3,300,000 service units on the PSC Cray XT3, 2,719,500 service units on the NCSA Tungsten Cluster, 2,295,000 service units on the PSC TCS1, 2,040,000 service units on the SDSC BlueGene/L, and 781,800 service units on the Teragrid Itanium Clusters.

Generation of lattices with improved gauge and quark actions: Our major lattice generation efforts in the next year will be the completion of the $m_l = 0.2m_s$ and $m_l = 0.4\,m_s$ ensembles at lattice spacing $a \approx 0.06$ fm run. We plan to finish the $0.4\,m_s$ run ensemble on the QCDOC at Brookhaven National Laboratory, and to generate 1,000 trajectories of the $0.2\,m_s$ one on the Kaon Cluster at Fermilab. We request time on the PSC Cray XT3 for 1,500 trajectories of the $0.2\,m_s$ ensemble. In production runs with this lattice spacing and quark mass on 1,536 processors of the XT3 our code achieves a performance of 1.5 Tflop/s. At this rate, one molecular dynamics trajectory takes 2,200 processor-hours. Thus, the 1,500 trajectories we propose to generate on the XT3 will require 3,300,000 service units.

Hadron mass spectrum and decays of light pseudoscalar mesons: We calculate the masses of the simplest light hadrons, the heavy quark potential and the partially quenched light pseudoscalar masses and decay constants as soon as we generate the gauge configurations. We request resources to analyze 250 new $a \approx 0.06$ fm, $m_l = 0.4\,m_s$ configurations during the coming year. On the NCSA Tungsten cluster each $m_l = 0.4\,m_s$ configuration requires 2,800 processor-hours for the light hadron spectrum, 2,000 processor-hours for the Ω^- and Delta masses, and 4,050 processor-hours for the pseudoscalar masses and decay constants. Thus, the total time for these calculations will be $250 \times 8,850 = 2,212,500$ processor-hours on the Tungsten cluster. We also request resources to run the light hadron spectrum and pseudoscalar codes on 100 $a \approx 0.06$, $m_s = 0.2\,m_s$ configurations. Both will have to be run in double precision. We estimate that the spectrum code will require 15,000 processor-hours per configuration on the SDSC BlueGene/L, and the light pseudoscalar code will require 5,400 processor-hours, based on the optimization effort discussed above. Thus, $100 \times 20,400 = 2,040,000$ processor-hours on the BlueGene/L.

Weak decays of particles containing heavy quarks: Our studies of the weak decays of D and B mesons on the $a \approx 0.12$ and 0.09 fm ensembles are being carried out on clusters at Fermilab. In this proposal we request time to perform this analysis on 150 of the $48^3 \times 144$, $a \approx 0.06$ fm, $m_l = 0.4\,m_s$ configurations. Each configuration requires 25,300 processor-hours on the PSC TCS1. We therefore request a total of $150 \times 15,300 = 2,295,000$ service units on TCS1 for this work.

High temperature QCD with three flavors of quarks: We expect to finish the calculation of the equation of state for non-zero baryon density on lattices with $N_t = 4$ by the end of our current allocation, and request time to extend this work to lattices with $N_t = 6$. We propose to analyze approximately 600 configurations at each of thirteen temperature points. Each configuration requires an average of 65 processor-hours on the Tungsten cluster, so our total request is 507,000 processor-hours on Tungsten. The second part of this project is to measure the equation of state for zero baryon density in the regime below the crossover temperature with the exact RHMC algorithm in order to eliminate the finite step size errors that dominate the uncertainty in our current results. We require seven high temperature runs of 6,000 trajectories each which will cost a total of 350,000 processor-hours on the Teragrid cluster, and four low temperature runs of 1,200 trajectories, which will cost 250,000 processor-hours on this platform. Thus, our total request for this work is 600,000 service units on the Teragrid cluster, and 507,000 service units on the Tungsten cluster.

Non-perturbative evaluation of quark masses: We propose to analyze 200 configurations from each of the ensembles with lattice spacings $a \approx 0.12$ fm and $a \approx 0.09$ fm and light quark masses $m_l = 0.4\,m_s$, $0.2\,m_s$ and $0.1\,m_s$. On the Teragrid cluster, the $a \approx 0.12$ ensembles will require an average of 28 processor-hours per configuration, while the $a \approx 0.09$ fm ensembles will require an average of 275 processor-hours per configuration. So, we request $600 \times (28 + 275) = 181,800$ service units on the Teragrid cluster for this work.

References

[1] The MILC Collaboration: C. Bernard *et al.*, Nucl. Phys. (Proc. Suppl.), **60A**, 297 (1998); Phys. Rev. **D 58**, 014503 (1998); G.P. Lepage, Nucl. Phys. (Proc. Suppl.), **60A**, 267 (1998); Phys. Rev. **D 59**, 074501 (1999); Kostas Orginos and Doug Toussaint (MILC), Nucl. Phys. (Proc. Suppl.), **73**, 909 (1999); Phys. Rev. **59**, 014501 (1999); Kostas Orginos, Doug Toussaint and R.L. Sugar (MILC), Phys. Rev. D **60**, 054503 (1999); The MILC Collaboration: C. Bernard *et al.*, Phys. Rev. **D 61**, 111502 (2000).

[2] The Fermilab, HPQCD, MILC, and UKQCD Collaborations: C.T.H. Davies *et al.*, Phys. Rev. Lett. **92**, 022001 (2004).

[3] The MILC Collaboration: C. Aubin *et al.*, Phys. Rev. D **70**, 114501 (2004); Nucl. Phys. B (Proc. Suppl.) **140**, 231 (2005); Proceedings of Science (Lattice 2005) 025 (2005).

[4] The HPQCD and UKQCD Collaborations: C. Davies *et al.*, Nucl. Phys. B (Proc. Suppl.) **119** (2003) 595; Phys. Rev. Lett. **95**, 052002 (2005).

[5] The Fermilab Lattice, MILC and HPQCD Collaborations: C. Aubin *et al.*, Phys. Rev. Lett.**95**, 122002 (2005). The prediction for the leptonic decay constant of the D meson was confirmed experimentally be the CLEO-c Collaboration, Phys. Rev. Lett. **95** 251801 (2005); and that for the D_s meson by the BaBar Collaboration, http://moriond.in2p3.fr/EW/2006/Transparencies/J.W.Berryhill.pdf

[6] The Fermilab Lattice and MILC Collaborations: C. Aubin *et al.*, Phys. Rev. Lett. **94**, 011601 (2005). Experimental confirmation was provided by The Focus Collaboration, Phys. Lett. **B607**, 233 (2005), and the Belle Collaboration, arXiv:hep-ex/0510003.

[7] The Fermilab Lattice and UKQCD Collaborations: I.F. Allison *et al.*, Phys. Rev. Lett. **94**, 172001 (2005). This prediction was confirmed by the CDF Collaboration, Phys. Rev. Lett. **96** 082002 (2006).

[8] Ian Shipsey, Nature **427** 591 (2004).

[9] A. Cho, Science, **300** 1076 (2003).

[10] C. DeTar and S. Gottlieb, Physics Today, p 45, February 2004.

[11] C. Davies, CERN Courier, **44 #5**, 23, June, 2004, http://www.cerncourier.com/main/article/44/5/13; and D. Cassel, CERN Courier **45 #6**, 14, July, 2005, http://www.cerncourier.com/main/article/45/6/14.

[12] http://www.fnal.gov/pub/today/archive_2004/today04-02-12.html, and http://www.fnal.gov/pub/today/archive_2005/today05-08-10.html.

[13] M.A. Clark, A.D. Kennedy and Z. Sroczynski, Nucl. Phys. (Proc. Suppl.) **140**, 835 (2005); M.A. Clark, Ph. de Frocrand and A.D. Kennedy, PoS LAT2005, 115 (2005).

[14] See, for example, J. Gasser and H. Leutwyler, Nucl. Phys. **B250**, 465 (1985).

[15] The MILC Collaboration: C. Bernard *et al.*, Phys. Rev. D **64**, 054506 (2001); MILC Collaboration: C. Aubin *et al.*, Phys. Rev. D **70**, 094505 (2004).

[16] W. Lee and S. Sharpe Phys. Rev. **D60**, 114503 (1999).

[17] C. Bernard, Phys. Rev. D **65**, 054031 (2002).

[18] C. Aubin, and C. Bernard, Phys. Rev. D 68, 034014 (2003); Phys. Rev. D **68**, 074011 (2003).

[19] C. Aubin, and C. Bernard, Nucl. Phys. (Proc. Suppl.) **140**, 491 (2005).

[20] K. Jansen, Nucl. Phys. B (Proc. Suppl.), **129**, 3 (2004); T. DeGrand, Int. J. Mod. Phys. A **19**, 1337 (2004).

[21] M. Creutz, arXiv:hep-lat/0603020.

[22] E. Follana, A. Hart and C.T.H. Davies (HPQCD Collaboration), Phys. Rev. Lett. **93** 241601 (2004); S. Dürr, C. Hoelbling and U. Wenger, Phys. Rev. D **70**, 094501 (2004); D.H. Adams, Phys. Rev. D **72**, 114512 (2005); F. Maresca and M. Peardon, arXiv:hep-lat/0411029.

[23] Y. Shamir, Phys. Rev. D **71**, 034509 (2005); and arXiv:hep-lat/0607007.

[24] C. Bernard, arXiv:hep-lat/0603011.

[25] C. Bernard, M. Golterman and Y. Shamir, arXiv:hep-lat/0604017; C. Bernard, M. Golterman, Y. Shamir and S. Sharpe, arXiv:hep-lat/0603027.

[26] C. Bernard and M. Golterman, Phys. Rev. D**49**, 486 (1994); S. Sharpe and N. Shoresh Phys. Rev. D**62**, 094503 (2000).

[27] S. Eidelman, *et al.* [Particle Data Group], Phys. Lett. B **592**, 1 (2004).

[28] E. Blucher, arXiv: hep-ph/05120. To appear in the Proceedings of the Third Workshop on the Unitarity Triangle: CKM2005, San Diego, CA March 15-18, 2005.

[29] A. J. Buras, M. Jamin and M. E. Lautenbacher, Phys. Lett. **B389**, 749 (1996).

[30] D. B. Kaplan and A. V. Manohar, Phys. Rev. Lett. **56**, 2004 (1986); A. G. Cohen, D. B. Kaplan and A. E. Nelson, JHEP **9911**, 027 (1999).

[31] G. t'Hooft, Phys. Rev. Lett., **37**, 8 (1976).

[32] HPQCD, MILC, and UKQCD collaborations: C. Aubin, *et al.*, Phys. Rev. D **70** 031504(R) (2004).

[33] The Fermilab Lattice, MILC and HPQCD Collaborations: C. Aubin *et al.*, manuscript in preparation.

[34] An complementary approach to the study of B decays and mixing is being taken by the HPQCD collaboration using our gauge configurations and the NRQCD formulation of heavy quarks on the lattice. The HPQCD Collaboration: M. Wingate *et al.*, Phys. Rev. Lett. **92**, 162001 (2004); and Phys. Rev. Lett. **95**, 212001 (2005).

[35] Jon A. Bailey, manuscript in preparation.

[36] J.-W. Chen and M. J. Savage, Phys. Rev. D **65**, 094001 (2002).

[37] The MILC Collaboration: C. W. Bernard *et al.*, Nucl. Phys. **A702**, 140 (2002); Nucl. Phys. **B** (Proc. Suppl.) **106**, 429 (2002); Nucl. Phys. B, (Proc. Suppl.), **119**, 523 (2003); Nucl. Phys. B (Proc. Suppl.), **129& 130**, 626 (2004); Nucl. Phys. B (Proc. Suppl.) **140**, 538 (2005); Phys. Rev. D **71**, 034504 (2005).

[38] The MILC Collaboration: C. Bernard *et al.*, Proceedings of Science (Lattice 2005) 156 (2005).

[39] The MILC Collaboration: C. Bernard *et al.*, manuscript in preparation.

[40] C.R. Allton *et al.*, Phys.Rev. D **71**, 054508 (2005), and references contained therein.

[41] B. Billeter, C. DeTar and J. Osborn, Phys. Rev. D **70**, 077502 (2004).

[42] The MILC Collaboration: C. Bernard *et al.*, manuscript in preparation.

[43] G. Martinelli, C. Pettori, C.T. Sachrajda, M. Testa and A. Vladikas, Nucl. Phys. **B445**, 81 (1995).

[44] S. Gottlieb, W. Liu, D. Toussaint, R.L. Renken, and R. Sugar, Phys. Rev. D **35**, 2531 (1987.)

[45] I.P. Omelyan, I.M. Mryglod and R. Folk, Comp. Phys. Comm., **151**, 272 (2003); T. Takaishi and P. de Forcrand arXiv/hep-lat0505020.

[46] J.C. Sexton and D.H. Weingarten, Nucl. Phys. B **380**, 665 (1992).