Searching for Physics Beyond the Standard Model: Strongly-Coupled Field Theories at the Intensity and Energy Frontiers

Lead Principal Investigator and DOE Contact: Paul Mackenzie
Address: Fermi National Accelerator Laboratory
106 (WH 3E)
Batavia, IL 60510-5011.
Email: mackenzie@fnal.gov
Phone: (630) 840-3347

Applicant/Institution: Fermi National Accelerator Laboratory
Street Address/City/State/Zip: Batavia, IL 60510-5011
Principal Investigator: Paul Mackenzie
Postal Address: 106 (WH 3E), Fermi National Accelerator Laboratory , Batavia, IL 60510-5011
Telephone Number: (630) 840-3347
Email: mackenzie@fnal.gov

Funding Opportunity Announcement Number: DE-FOA-0000580
DOE/Office of Science Program Office: Office of High Energy Physics
DOE/Office of Science Program Office Technical Contact: Dr. Lali Chatterjee

Participating Institutions and Principal Investigators:

Argonne National Laboratory*, James Osborn ‡
Boston University*, Richard Brower † ‡
Brookhaven National Laboratory*, Frithjof Karsch † ‡
Columbia University, Norman Christ †
Fermi National Accelerator Laboratory*, Paul Mackenzie † ‡
Indiana University*, Steven Gottlieb ‡
Lawrence Livermore National Laboratory*, Pavlos Vranas ‡
Syracuse University*, Simon Catterall ‡
University of Arizona*, Doug Toussaint ‡
University of Illinois*, Guochun Shi ‡
University of California, Santa Barbara, Robert Sugar †
University of California, San Diego, Julius Kuti †
University of Utah*, Carleton DeTar ‡
University of Washington, Stephen Sharpe †

* Institution submitting an application
† Project Principal Investigator, Member of the USQCD Executive Committee
‡ Institution Principal Investigator

Abstract

Project Title: Searching for Physics Beyond the Standard Model:
Strongly-Coupled Field Theories at the Intensity and Energy Frontiers

Project Director and DOE Contact: Paul Mackenzie
Address: Fermi National Accelerator Laboratory 106, Batavia, IL 60510-5011
Email: mackenzie@fnal.gov **Phone:** (630) 840-3347
Project Co-Director for Science: Stephen Sharpe, U. Washington (sharpe@phys.washington.edu)
Project Co-Director for Computation: Richard Brower, Boston University, (brower@bu.edu)

Participating Institutions and Principal Investigators:
Argonne National Laboratory*, James Osborn ‡
Boston University*, Richard Brower †‡
Brookhaven National Laboratory*, Frithjof Karsch † ‡
Columbia University, Norman Christ †
Fermi National Accelerator Laboratory*, Paul Mackenzie † ‡
Indiana University*, Steven Gottlieb ‡
Lawrence Livermore National Laboratory*, Pavlos Vranas ‡
Syracuse University*, Simon Catterall ‡
University of Arizona*, Doug Toussaint ‡
University of Illinois*, Guochun Shi ‡
University of California, Santa Barbara, Robert Sugar †
University of California, San Diego, Julius Kuti †
University of Utah*, Carleton DeTar ‡
University of Washington, Stephen Sharpe †

* Institution submitting an application; ‡ Institution Principal Investigator
† Project Principal Investigator, Member of the USQCD Executive Committee

In this proposal we request resources to develop the software and algorithmic infrastructure needed for the numerical study of quantum chromodynamics (QCD), and of theories that have been proposed to describe physics beyond the Standard Model (BSM) of high energy physics, on current and future computers. This infrastructure will enable users (1) to improve the accuracy of QCD calculations to the point where they no longer limit what can be learned from high precision experiments that seek to test the Standard Model, and (2) to determine the predictions of BSM theories in order to understand which of them are consistent with the data that will soon be available from the LHC. Work will include the extension and optimizations of community codes for the next generation of leadership class computers, the IBM Blue Gene/Q and the Cray XE/XK, and for the dedicated hardware funded for our field by the Department of Energy. Members of our collaboration at Brookhaven National Laboratory and Columbia University worked on the design of the Blue Gene/Q, and have begun to develop software for it. Under this grant we will build upon their experience to produce high efficiency production codes for this machine. Cray XE/XK computers with many thousands of GPU accelerators will soon be available, and the dedicated commodity clusters we obtain with DOE funding include growing numbers of GPUs. We will work with our partners in NVIDIA's Emerging Technology group to scale our existing software to thousands of GPUs, and to produce highly efficient production codes for these machines. We will collaborate with our nuclear physics colleagues in USQCD on software for Intel's Many Integrated Core (MIC) architecture. Work under this grant will also include the development of new algorithms for the effective use of heterogeneous computers, and their integration into our codes. It will include improvements of Krylov solvers and the development of new multigrid methods in collaboration with members of the FASTMath SciDAC Institute, using their HYPRE framework, as well as work on improved symplectic integrators.

Contents

1 Introduction

The long term goals of high energy physicists are to identify the fundamental building blocks of matter, and to determine the interactions among them that give rise to the physical world we observe. Major progress has been made towards these goals through the development of the Standard Model of high energy physics. It consists of two quantum field theories: the Weinberg-Salam theory of weak and electromagnetic interactions, and quantum chromodynamics (QCD), the theory of the strong interactions. The Standard Model has been enormously successful in explaining a wealth of data produced in accelerator experiments and astrophysical observations over the past thirty years. Despite these successes, it is believed by high energy physicists that to understand physics at the energy scales being probed by the Large Hadron Collider (LHC) will require a more general theory capable of resolving fundamental questions, such as the origin of mass, the nature of dark matter and energy, and the pattern of symmetries found in nature. 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.

Central objectives of the search for a more general theory of fundamental interactions are to determine the range of validity of the Standard Model, and to search for physical phenomena that will require new theoretical ideas for their understanding. Two complementary approaches are being employed. Accelerators with the highest available beam intensity are being used to make precision tests of the Standard Model in the hope of finding contradictions, while accelerators with the highest available energy are being used to search directly for physical phenomena not predicted by the Standard Model. Numerical simulations within the framework of lattice gauge theory play important roles in both of these approaches. In work at the *intensity frontier,* precise calculations of the effects of the strong interactions on weak and electromagnetic transition amplitudes are often needed for Standard Model tests. For most such cases, the precision of the lattice calculations lag behind those of the experiments. One of the major goals of this project is to improve the precision of these calculations to the point where they no longer limit what can be learned from experiments. The role of lattice gauge theory simulations in work at the *energy frontier* is somewhat different. A number of theories have been proposed to describe beyond the Standard Model (BSM) physics, most of which contain strong coupling regimes that can be studied from first principles only with lattice simulations. Our goal in this area is to determine the predictions of these theories in order to understand which of them are consistent with the data that will soon be available from the LHC. BSM calculations are at least as challenging as QCD ones because of the great range of scales that must be covered for these conformal or nearly conformal theories. Thus, in work at both the intensity and energy frontiers, lattice gauge theory simulations play important roles in efforts to obtain a deeper understanding of the fundamental laws of physics. Detailed descriptions of our scientific goals are presented in Section 2.

In this proposal we request resources to develop the software and algorithmic infrastructure needed to study QCD and BSM theories on current and future computers. The proposed work is a project of the USQCD Collaboration, which consists of nearly all of the high energy and nuclear physicists in the United States working on the numerical study of lattice gauge theories. It will be carried out by high energy physicists within USQCD, in collaboration with nuclear physicists within USQCD, members of the SciDAC FASTMath Institute and colleagues in the NVIDIA Emerging Applications Group. The roles of the FASTMath Institute and the NVIDIA Emerging Applications group in this proposal are spelled out in Secs. 5 and 6. Appendix A.2 contains a letter from Dr. David Luebke, the Director of Research of NVIDIA, expressing his company's goals for our joint work, and setting out its commitment of resources for this project, which include the majority of the time of one staff member, Mike Clark, to support our efforts to develop software to exploit graphics processing units (GPUs), and time of a second staff member, Ron Babich, to "leverage the lessons learned during this collaborative effort to guide the design of next generation GPU architectures from NVIDIA." Babich and Clark are lattice gauge theorists who recently moved to NVIDIA. They have played major roles in developing USQCD's GPU codes. Appendix A.3 contains a letter from Dr. George

Chiu, Senior Manager, Advanced Server Systems, IBM, expressing his company's intention to continue "its fruitful, 10+ year collaboration with USQCD scientists," focusing on the optimization of USQCD codes for the Blue Gene/Q, and the development of new algorithms to exploit its heterogeneous architecture. Appendix A.4 contains a letter from Mr. Joseph Curley, Director of Marketing, Technical Computing Group at Intel, indicating his company's intention to "proceed on collaborative interactions" with USQCD on the development of our codes for the Intel Many Integrated Core (Intel MIC) Architecture.

In Sec. 3 we briefly describe the computational approach used in our studies of QCD and BSM theories. Our calculations proceed in two steps. In the first, one uses Monte Carlo techniques to generate configurations of gauge fields in proportion to their weight in the Feynman path integrals that define the theories. These configurations are saved, and in the second step they are used to measure a wide range of physical quantities. The generation of gauge configurations presents a capability challenge which is best met on the most powerful available leadership class computers. The measurements on different configurations can be done in parallel, and present a capacity challenge. For the last six years the DOE Offices of High Energy and Nuclear Physics have jointly funded dedicated commodity clusters for USQCD under the Lattice QCD (LQCD) Computing Project, and its successors. These machines have primarily been used for measurement campaigns.

Our proposed work will build upon achievements of USQCD under grants from the DOE's SciDAC-1 and SciDAC-2 programs. With these grants we developed the QCD Applications Programming Interface (QCD API), a programming environment that enables users to quickly adapt existing codes to new architectures, easily develop new codes and incorporate new algorithms. It has greatly facilitated the efficient use of leadership class computers and commodity clusters. Indeed, the committees that have performed yearly reviews of the LQCD Computing Project have repeatedly emphasized the critical role our SciDAC software plays in that effort. The accomplishments of USQCD under our SciDAC grants are reviewed in Sec. 4.

In Sec. 5 we present an overview of our plans for software and algorithm development over the three year period of this proposal. The increasingly heterogeneous nature of the computers we expect to use during the coming years requires that we fully adopt a hybrid computing model in which threaded code is used on SMP nodes and GPUs, and message passing is used for communication between nodes. Work in this direction is in progress, but it needs to be completed and fully integrated into our production codes under this grant. In addition to optimizing our QCD codes, we propose to extend our software to cover BSM theories, which, although similar in structure to QCD, in most cases involve different gauge groups and/or fermion representations than QCD.

Particular emphasis will be placed on optimizing the QCD API and our community applications codes for the next generation of leadership class computers: the Blue Gene/Q, and the Cray XE/XK, and for commodity clusters with GPU accelerators. Members of USQCD at Brookhaven National Laboratory and Columbia University participated in the design of the Blue Gene/Q, and have already begun to develop code for it in collaboration with colleagues at the University of Edinburgh and at IBM. Under this grant we will build upon their experience to produce high-efficiency production codes for the Blue Gene/Q. We have been selected to take part in the Early Science Program of the Blue Gene/Q, Mira, which is scheduled to be installed at Argonne National Laboratory in 2012. Our goal is to have highly efficient production code ready to run when the Early Science Program begins.

Both the DOE and the NSF plan to make Cray XE/XKs with thousands of GPU accelerators available to the U.S. research community in the near future, Titan at Oak Ridge National Laboratory and Blue Waters at NCSA. We have a PRAC grant for early access to Blue Waters, and hope to obtain access to Titan through the DOE's INCITE Program. The clusters USQCD has recently obtained through the LQCD Computing Program have had growing numbers of GPUs, and we expect this trend to accelerate in the future. The Krylov solvers associated with the lattice Dirac operator consume the largest fraction of floating point operations in both configuration generation and measurement routines. There are a number of different formulations of the lattice Dirac operator, and we have GPU code for the solvers associated with all of those in common use. All but one of these codes runs on multiple GPUs, and exhibit excellent strong scaling on hundreds of GPUs. Under this grant we will work with our partners in NVIDIA's Emerging Applications Group to extend scaling to thousands of GPUs, and to produce highly efficient production codes for the Cray XE/XK and for clusters with GPU accelerators. We will also develop code for the Intel MIC architecture, which will be a strong candidate for use in future USQCD clusters A.4.

Work under this grant will also include the development of new algorithms for the effective use of heterogeneous computers and their integration into our codes. The advent of petascale computers should enable us to carry out simulations with smaller lattices spacings than previously possible using physical values of the quark masses. This development would greatly improve the precision of our calculations, but it requires major improvements in our solvers. We have recently succeeded for the first time in applying multigrid methods [1, 2, 3] to the solver for one formulation of lattice quarks, and propose to extend this work to other formulations, as well as to study other approaches. This work will be carried out in collaboration with members of the FASTMath SciDAC Institute, using their HYPRE framework. We also propose to work on the development of improved symplectic integrators which play a key role in our simulations.

In Sec. 6 we set out the tasks and milestones for each institution participating in this project. More detailed statements of work can be found in Appendix A.6, while individual budgets for each instituion are contained in Appendix A.5. Most of the work will involve collaborations between participating institutions. The management plan for this project is set out in Sec. 7, and our plans for outreach and dissemination of results in Sec. 8. Finally, the budget proposed for each participating institution in each of the three years of the grant, is contained in Sec. 9. Nearly all of the funds are for the support of the 8.7 FTE working on the software effort. Of this number 5.0 FTE will be for the support of postdoctoral research associates. The 5.0 FTE will be divided among seven individuals, with the remaining support coming from high energy physics research grants or their home institutions. The postdocs supported by this grant will therefore have training in high energy physics, and in state of the art software and algorithm development. Scientists trained in our field have gone on to very successful careers in the computer industry, as well as in academia, so this proposal has the potential to play a significant role in the training of the next generation of computational scientists.

2 Scientific Goals

2.1 Flavor Physics at the Intensity Frontier

A major component of the experimental program in high energy physics lies at the intensity frontier. It is devoted to making precise tests of the Standard Model (SM), in order to determine its range of validity, and to searching for indications of new physics beyond the SM. Many of these tests require both accurate experiments *and* accurate lattice QCD (LQCD) calculations of the effects of the strong interactions on electroweak processes. In most cases, the precision of the tests are limited by the uncertainties in the lattice calculations, rather than in the experiments. One objective of our calculations is to reduce lattice errors down to, or below, the level of the experimental errors.

The US LQCD community has developed a very successful program to calculate the needed electroweak matrix elements. Indeed, it is now appropriate to provide world averages for key quantities to the wider particle physics community [4, 5]. Results obtained by members of the USQCD collaboration play the leading role in the determination of the averages for many of these quantities. The software and algorithmic development funded by the SciDAC-2 award has been crucial to obtaining these results.

Under SciDAC-3, we propose to extend this work in two ways. First, to improve the errors on quantities for which results with fully-controlled errors exist, but for which the errors are still larger than or comparable to those from other sources. And, second, to expand our program of calculations to meet the needs of upcoming intensity-frontier experiments, for example the muon $g-2$ experiment at Fermilab, the Project X kaon program at Fermilab, LHC-b, Belle II and Super-B. We describe both these extensions below, focusing on opportunities where we expect LQCD calculations to play a key role in searches for (and possibly discovery of) new physics in the quark-flavor sector.

This effort is complementary to direct searches for new physics at the energy frontier, such as those underway at the LHC. Indeed, if indications for beyond-the-Standard Model (BSM) physics are found at the LHC, then one expects that there will also be significant new contributions to flavor-changing processes such as kaon and B-meson mixing. To determine the size of these contributions predicted by a given model of new physics one also requires LQCD calculations.

The period of the SciDAC-2 award has seen a rapid maturation in LQCD calculations relevant for flavor physics. At the time of our previous proposal, calculations by members of the USQCD collaboration had given the first results in which all sources of errors were controlled. Typical errors were at the few percent level, and the results provided validation of LQCD methods both by finding agreement with experimental quantities and by making predictions that were subsequently confirmed. Examples of successful "post-dictions" are the results for f_π, f_K, and splittings in charmonium and bottomonium systems presented in Ref. [6]. Examples of successful predictions are those of D meson semileptonic form factors, D- and D_s-meson leptonic-decay constants, and the mass of the B_c meson, as summarized in Ref. [7].

During the SciDAC-2 award, the number of lattice predictions has substantially increased, the accuracy of the results has significantly improved (with sub-percent accuracy achieved in some cases), and results have been obtained using several different methods for discretizing fermions (providing crucial cross-checks). It is thus now appropriate to provide world averages, as noted above. We mention several important examples in which USQCD has played the major role.

The kaon mixing parameter B_K is the hadronic matrix element needed to determine the constraint on the Cabibbo-Kobayashi-Maskawa (CKM) matrix arising from the observation of CP-violation in kaon mixing. We emphasize that this measurement was made in 1964 [8], and it is only now, after nearly 50 years, that we can use this result to study the properties of quarks. Specifically, combining the measurement with B_K (and some other known inputs) one can determine $\mathrm{Im}V_{td}^2$. (Here the quantity V_{ab} is an element of the 3×3 CKM matrix where the index a identifies one of the three "up"-type quarks, u, c, and t while b labels one of the three "down"-type quarks d, s and b.) At the time of our SciDAC-2 proposal, there was only a single calculation of B_K using the physical complement of light-quark loops (up, down and strange quarks), and several sources of error were not controlled. The total error was estimated to be $12-20\%$, It was forecast in the SciDAC-2 proposal that the error would be reduced to $\sim 4\%$ during the award. The present error, in fact, exceeds this goal. Using the average from Ref. [4] ("end of 2010"), which includes only published results using the physical complement of light-quark loops and which control all sources of error, the error is now 2.7%.[1] This is based on three results all obtained using USQCD resources, all using different methods and giving consistent values. The lattice result has improved enough that it is no longer the dominant source of uncertainty in the constraint on CKM matrix elements based on the measured CP-violation in kaon mixing.

The second quantity we highlight is ξ, the ratio of mixing matrix elements for strange and non-strange neutral B-mesons. This ratio, combined with the experimental result for $\Delta M_B/\Delta M_{B_s}$, allows one to determine $|V_{td}/V_{ts}|^2$. Based on a single calculation, the error in this ratio was estimated to be 8% at the time of our SciDAC-2 proposal, and we forecast that the error would be reduced to $\sim 4\%$ during the award. The present error of 2.6% [4] again exceeds this goal. This is based on 3 independent calculations, all using USQCD software and hardware resources, and all with a complete error budget.

[1] Not included in this average is a very recent result using Wilson fermions from the BMW collaboration which further reduces the overall error [9].

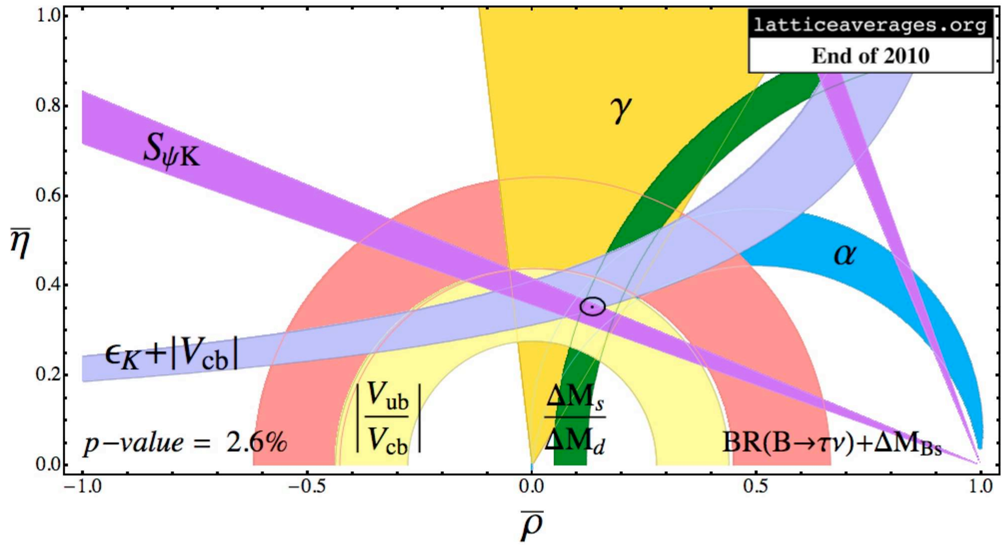

Figure 1: Global fit of the CKM unitarity triangle [4]. The constraints labeled $\varepsilon_K + |V_{cb}|$, $|V_{ub}/V_{cb}|$, $\Delta M_s/\Delta M_d$, and $\mathrm{BR}(B \to \tau\nu) + \Delta M_{B_s}$ all require LQCD input, while the others require minimal or non-lattice theoretical input. The solid ellipse encloses the 1σ region.

The impact of these and other LQCD results can be seen from Fig. 1, which shows a recent global fit to the parameters of the CKM unitarity triangle. New quark flavor-changing interactions or *CP*-violating phases would manifest themselves as apparent inconsistencies between measurements of the apex $(\overline{\rho}, \overline{\eta})$ that are predicted to be the same within the SM framework. Four of the constraints require LQCD input—for example, those labeled "$\varepsilon_K + |V_{cb}|$" and "$\Delta M_s/\Delta M_d$" require B_K and ξ, respectively. The widths of the bands requiring lattice inputs have been significantly reduced over the period of the SciDAC-2 award, primarily due to the reduction in LQCD errors. As can be seen from the figure, this combination of precise experiments and theoretical calculations has established the CKM paradigm of CP-violation at the few-percent level. At the same time, an $\sim 3\sigma$ tension in the fit has been revealed [10, 11, 12, 13]. This may indicate the presence of sources of CP-violation beyond the SM.

Two other highlights of recent lattice calculations are the determination of α_S and the quark masses. LQCD provides the most accurate result for α_S [14], and competitive results for the charm and bottom masses. LQCD provides the only *ab initio* determination of the up, down and strange masses [4, 5]. All these quantities are fundamental parameters of the SM, and are needed as inputs into models of new physics such as possible grand unified theories.

A major goal for the present proposal is to extend these successes to other hadronic matrix elements. The following matrix shows which hadronic processes can be used to obtain each CKM matrix element, in each case using LQCD calculations:

$$\begin{pmatrix} \mathbf{V_{ud}} & \mathbf{V_{us}} & \mathbf{V_{ub}} \\ \pi \to \ell\nu & K \to \ell\nu & B \to \ell\nu \\ & K \to \pi\ell\nu & B \to \pi\ell\nu \\ \mathbf{V_{cd}} & \mathbf{V_{cs}} & \mathbf{V_{cb}} \\ D \to \ell\nu & D_s \to \ell\nu & B \to D\ell\nu \\ D \to \pi\ell\nu & D \to K\ell\nu & B \to D^*\ell\nu \\ \mathbf{V_{td}} & \mathbf{V_{ts}} & \mathbf{V_{tb}} \\ \Delta M_d & \Delta M_s & \end{pmatrix}$$

Table 1: *Impact of improved LQCD calculations on the determination of CKM matrix elements.*

Quantity	CKM element	Present expt. error	Present lattice error	2014 lattice error	2020 lattice error
f_K/f_π	$\lvert V_{us}\rvert$	0.2%	0.6%	0.3%	0.1%
$f_+^{K\pi}(0)$	$\lvert V_{us}\rvert$	0.2%	0.5%	0.2%	0.1%
$D \to \pi\ell\nu$	$\lvert V_{cd}\rvert$	2.6%	10.5%	4%	1%
$D \to K\ell\nu$	$\lvert V_{cs}\rvert$	1.1%	2.5%	2%	$< 1\%$
$B \to D^{(*)}\ell\nu$	$\lvert V_{cb}\rvert$	1.8%	1.8%	0.8%	$< 0.5\%$
$B \to \pi\ell\nu$	$\lvert V_{ub}\rvert$	4.1%	8.7%	4%	2%
$B \to \tau\nu$	$\lvert V_{ub}\rvert$	21%	6.4%	2%	$< 1\%$
ξ	$\lvert V_{ts}/V_{td}\rvert$	1.0%	2.5%	1.5%	$< 1\%$
ΔM_s	$\lvert V_{ts}V_{tb}\rvert^2$	0.7%	10.5%	5%	3%

In Table 1 we compare, for many of these processes, the present experimental and lattice errors, and give an estimate of what LQCD can achieve in ~ 3 years as well as a forecast for the end of the decade. Note that in the few cases where the effects of low energy QCD can be estimated using non-lattice methods, the resulting uncertainties are at least as large as those in the present lattice determinations.

These forecasts assume computations running at the 10's of petaflop/s by 2014 and approaching the exascale by 2020. Achieving such computational throughput will require extensive software development, for example the extension of our high-performance code for the current, leadership class machines to the much more challenging multi-thread/multi-core architectures of next generation HPC machines such as the Blue Gene/Q, the adaptation of codes to an increasingly heterogeneous, multi-GPU environment such as Titan and Blue Waters, the development of improved multi-scale solvers to avoid critical slowing down as the quark masses are decreased, and the development of improved integrators such as the force-gradient method. This algorithm and software development will be a major focus of our SciDAC-3 program.

A major milestone achieved during the SciDAC-3 program will be the widespread use of simulations at the physical quark masses, first results from which have recently appeared [15, 16, 17]. We plan to undertake such simulations for both highly improved staggered quarks (HISQ) and domain wall fermions (DWF) and these play a major role in the reduction of errors forecast in Table 1.

Table 1 shows that LQCD calculations can by 2014 match the present experimental errors in many quantities and do substantially better than the alternative methods. Such improvements promise to push the SM beyond its limits—possibly raising the present $\sim 3\sigma$ tension to the 5σ discovery level. We give three examples of the essential role that LQCD calculations play in this program.

The first concerns the lattice results for B-$\bar{B}$ mixing. The lattice quantity ξ^2 determines the width of the $\Delta M_s/\Delta M_d$ band in Fig. 1. By 2014 we expect that this width should be roughly halved, which will significantly reduce the horizontal width of the allowed ellipse. It will also be possible to include a separate constraint from ΔM_s alone. This requires a LQCD calculation of the matrix element $f_{B_s}^2 B_{B_s}$. Only one calculation with the full complement of light sea quarks is available and, as noted in the bottom row of Table 1, it has an error of 10.5%. We expect several independent calculations to be completed by 2014, with errors roughly halved. Although the lattice error will still dominate over the sub-percent experimental error, the input will add an important further constraint on the unitarity triangle.

Our second example concerns the importance of improving the determination of $|V_{cb}|$ using lattice calculations of $B \to D^{(*)}\ell\nu$ form factors. $|V_{cb}|$ plays a key role in tightening unitarity constraints, since, now that

B_K is well determined, the width of the ε_K band in Fig. 1 arises principally from the uncertainty in $|V_{cb}|^4$. We expect a reduction in this width by about a factor of two by 2014. Reducing the uncertainty in $|V_{cb}|$ also has a major impact on searches for new physics in the rare decays $K^+ \to \pi^+\nu\bar{\nu}$ and $K_L \to \pi^0\nu\bar{\nu}$. The leading SM contribution to these decays is from penguin diagrams, and is thus very small, so that new physics can lead to substantial deviations from SM predictions. It is thus important to nail down the SM predictions; and these are also proportional to $|V_{cb}|^4$.

Our final example concerns the determination of $|V_{ub}|$, using lattice results for the $B \to \pi\ell\nu$ form factor and of f_B. At present there is a $\sim 2\sigma$ discrepancy between the value of $|V_{ub}|$ determined from lattice inputs and that determined from comparing inclusive B decays with perturbative QCD predictions based on heavy-quark effective theory. Reducing the lattice errors will allow us to determine whether this long-standing discrepancy is real—indicating new physics such as right-handed currents—or not. We forecast very significant improvements in the lattice results. This is in part because one will be able to work on very fine lattices, allowing the use of automatically normalized vector and axial currents.

The calculations discussed so far involve standard methods, and their improvement with increasing resources can be estimated with reasonable accuracy. We now turn to the second extension of our SciDAC-2 program, namely to calculations of new quantities. For most of these, the methodology is at an earlier stage of development, and it is hard to predict future errors. We expect these calculations to take up an increasing part of our computational effort over the next few years. We describe three main examples. (Other promising topics which are not discussed include $B \to K\ell^+\ell^-$, nucleon matrix elements required for dark matter searches, and hadronic contributions to the determination of the fine structure constant α.)

The first is the most straightforward. New physics is expected to enhance mixing in neutral kaon, D and B systems. This would show up, for example, as a failure of the unitarity triangle fit. If such a failure occurs, we need to be able to determine whether any given model of new physics, e.g. supersymmetry or extra dimensions, is consistent with the observed mixing. This involves both determining the SM contribution—which for K and B systems has been discussed above—and the contribution from new physics. The latter enters through four-fermion operators which have different Dirac structures than the left-handed operators of the SM. Thus, we must determine the generalizations of B_K, B_D and B_B for these new operators. The calculations are, in fact, being done in parallel with those for B_K etc., and we expect that results with comparable errors will be available during the SciDAC-3 award.

Our second example concerns kaon properties. It has been a long-standing aim of lattice calculations to calculate the $K \to \pi\pi$ decay amplitudes. Calculating the CP-conserving parts will allow us to test whether QCD explains the $\Delta I = 1/2$ rule (i.e. the dominance of the $I = 0$ final state pions over $I = 2$). Calculating the CP-violating parts will allow us to use the experimental result for direct CP-violation in kaon decays ($\varepsilon'_K/\varepsilon_K$) in order to further constrain the parameters of the SM. We have recently made very significant progress on the lattice calculations of these decay amplitudes using domain wall fermions. A complete methodology is now in place, and the path to precise calculations is clear. Indeed, results with $\sim 20\%$ precision for the $\Delta I = 3/2$ amplitudes have already been obtained [18], and a pilot calculation of all parts of the more challenging $\Delta I = 1/2$ amplitudes has been completed [19]. We expect that by 2014 the error in the latter amplitudes will have dropped to $\sim 30\%$.

Another quantity that can be used to constrain new physics is the CP-conserving part of the kaon mixing amplitude, ΔM_K. In the SM, this quantity receives a large contribution from the long-distance two pion intermediate states. This makes it much harder to calculate using lattice methods than the short-distance dominated CP-violating part (which is proportional to B_K). Nevertheless, a method for the calculation has recently been presented [20], and a pilot study undertaken [21]. We expect an error of $\sim 30\%$ by 2014.

An important area of overlap between the standard quantities where errors on the 1% level are targeted and new directions where the methods of lattice QCD can be applied but are not fully developed is the inclusion of electromagnetic effects and the isospin breaking difference between the masses of the up and down quarks. Here pioneering calculations have demonstrated that the methods of LQCD can indeed be extended to include the electromagnetic as well as the color gauge fields [22, 23, 24]. These methods will be further developed during SciDAC-3 and will be needed for some of the precision goals targeted in 2014.

Finally, we discuss how LQCD calculations will contribute to the calculation of the muon anomalous magnetic moment in the SM. At present, there is a $\sim 3.5\sigma$ discrepancy between the theoretical prediction and experiment (see, e.g., Ref. [25]). If this discrepancy holds up, it is a clear signal of new physics. The new FNAL experiment E989 plans to reduce the experimental error over the next decade from the present 0.5 parts per million (ppm) to 0.14 ppm. To make full use of this experimental effort, it is crucial to reduce the theoretical error to a similar level. This is a major challenge, because there are two hadronic contributions, involving strong interaction physics, that have uncertainties that are presently significantly larger than the planned final accuracy. These are the hadronic vacuum polarization contribution, which is presently obtained using data from $e^+e^- \to$ hadrons and has an error of about 0.4 ppm, and the light-by-light scattering contribution, which is estimated using various theoretical methods and has a comparable error. LQCD calculations can, in principle, provide improved results for both quantities, and methods for these calculations are under development [26, 27, 28, 29, 30]. It is very hard to forecast the future of such pioneering calculations, but we expect significant progress during the period of the SciDAC-3 award.

2.2 Physics Beyond the Standard Model at the Energy Frontier

The era of the Large Hadron Collider (LHC) is likely to expose new Beyond the Standard Model (BSM) physics with important non-perturbative aspects to explain the origin of electroweak symmetry breaking and the origin of mass in the universe. A complete understanding of the new physics will require non-perturbative lattice studies that can be effectively implemented with the USQCD hardware/software infrastructure. The BSM effort at the energy frontier of the LHC is complementary to indirect searches for new physics at the intensity frontier. Precise lattice calculations are capable of making important connections between the two frontiers since BSM models lead to direct predictions that can be observed at the intensity frontier. The BSM paradigm also offers potential new insight into the origin of the electroweak phase transition in the early universe relevant to the cosmic frontier.

In the Standard Model (SM) electroweak symmetry breaking is accomplished by coupling the theory to the elementary scalar Higgs field - and the related Higgs particle. From a theoretical viewpoint the problematic and unnatural fine tuning of the scalar Higgs field can be replaced by two known solutions: 1) replace the elementary Higgs boson by some new composite Higgs mechanism, or 2) assume that the Standard Model is embedded in an appropriate supersymmetric theory. In both cases new strong dynamics is required, either to break supersymmetry, thus ensuring that any new super particles are sufficiently heavy to have evaded detection thus far, or to provide the new strong force (with roots in past technicolor studies) necessary to break chiral symmetry and trigger a composite Higgs mechanism when the electroweak gauge interactions are turned on [31, 32, 33]. In the following subsections we will describe a lattice program designed to explore important nearly conformal BSM models and SUSY scenarios.

2.2.1 The composite Higgs mechanism close to the conformal window

For theories of dynamical weak symmetry breaking, such as technicolor, to be viable, they must have nearly conformal, or ”walking”, behavior over a large range of energy scales, as explained below. Walking technicolor theories generically have many techniquarks, and therefore many light techni-hadrons. The most copiously produced are likely to be singly-produced techni-rhos and techni-omegas, techni-hadron analogs of the spin-one meson states of QCD. These generically decay into pairs of Ws and Zs, or into a W or a Z and a pair of jets via a techni-pion. Singly-produced narrow resonances do not occur in SUSY theories with R parity, so discovery of such a signal would give high priority to non-perturbative lattice search for nearly conformal field theories representing a very interesting and promising class of BSM theories with the realization of the composite Higgs mechanism [34, 35, 36, 37, 38].

These theories are built on new fundamental particles, historically known as "techniquarks", that are different from the QCD quarks. They are massless and occur in an unknown number of N_c colors and N_f flavors. Candidate models also differ in the choice of fermion representation of the technicolor gauge group in the BSM Lagrangian. The theories exhibit a fundamental chiral symmetry on the Lagrangian level, which is dynamically broken in the vacuum by the new strong force, forming a condensate of fermions carrying the quantum numbers of the SM Higgs and providing the replacement for the old Higgs mechanism. The new strong force, similar to QCD, but operating at the TeV scale, plays a central role, replacing the SM Higgs particle as the source of electroweak symmetry breaking. The composite Higgs particle, if it can be observed in the new theory, will exhibit modified couplings to the electroweak gauge fields, like the dilaton of broken scale invariance close to conformality. The "techniquarks" are bound by the confining force of the new theory into heavy and colorless composite particles on the TeV scale providing interesting LHC signatures.

Early technicolor efforts employed models that were scaled up versions of QCD, but these models have since been ruled problematic by precision tests of the Standard Model. In contrast, nearly conformal Yang-Mills theories require good non-perturbative understanding of how precise properties of the theory depend on N_f, N_c and the fermion representation. At fixed N_c in a given fermion representation and low flavor number N_f, the models exhibit confinement and chiral symmetry breaking, which can in turn be used to describe electroweak symmetry breaking (EWSB). On the other hand, a perturbative study of the beta function indicates that these properties are lost for large N_f, giving way to conformal behavior in the infrared that can no longer be used for EWSB. The transition between these two phases occurs at some critical value $N_f^c(N_c)$ for a given fermion representation. Walking gauge coupling is expected just before $N_f^c(N_c)$ and the conformal phase is reached. Several such theories have been studied in the last few years and USQCD has played a major and successful role in these efforts [39, 40, 41, 42, 43, 44, 45, 46, 47]. The left side of Figure 2 is a visual summary of ongoing and planned USQCD BSM investigations of the theory space while the right side displays the first results on the variation of the S-parameter moving toward the conformal window.

Recent first-principles lattice calculations at $N_c = 3$ indicate that in the fundamental fermion representation N_f^c is close to $N_f^c = 12$ [40, 43, 44, 45, 49, 50]. A theory with N_f just below the critical value N_f^c may show approximately conformal infrared behavior, making it a good candidate model for walking technicolor. A very important added dimension to this exploration is the choice of the fermion representation. For example, N_f^c in the color sextet fermion representation has been found to be close to $N_f^c = 2$ [46, 47, 43]. There is reason to believe that phenomenologically viable models would prefer low N_f^c values, hence the significance of our proposed search in the theory space of higher fermion representations. Recent simulations with $SU(2)$ color gauge group have also made considerable progress exploring the BSM theory around the lower end of the conformal window [51, 52, 53, 54, 55, 56].

New BSM physics can affect low-energy precision measurements. Two such quantum effects are the contribution of the new interactions to the S parameter of electroweak precision tests and to flavor changing neutral currents (FCNCs). There are theoretical arguments that nearly conformal gauge theories might lead to an acceptable S parameter and naturally suppress contributions to FCNC's through a large anomalous dimension of the fermion condensate. All these constraints are deeply non-perturbative issues demanding large scale computations and the design of new BSM software infrastructure.

Composite Higgs model plans for the next three years include studies with two, three, and four colors. Studies with the $SU(2)$ color group will focus on the fundamental and adjoint fermion representations close to the conformal window. Studies with the $SU(3)$ color group will focus on the fundamental and two-index symmetric fermion representations. In addition, We plan to study the two-index symmetric fermion representation with four colors. Finally, we will extend our investigations of strongly coupled gauge theories

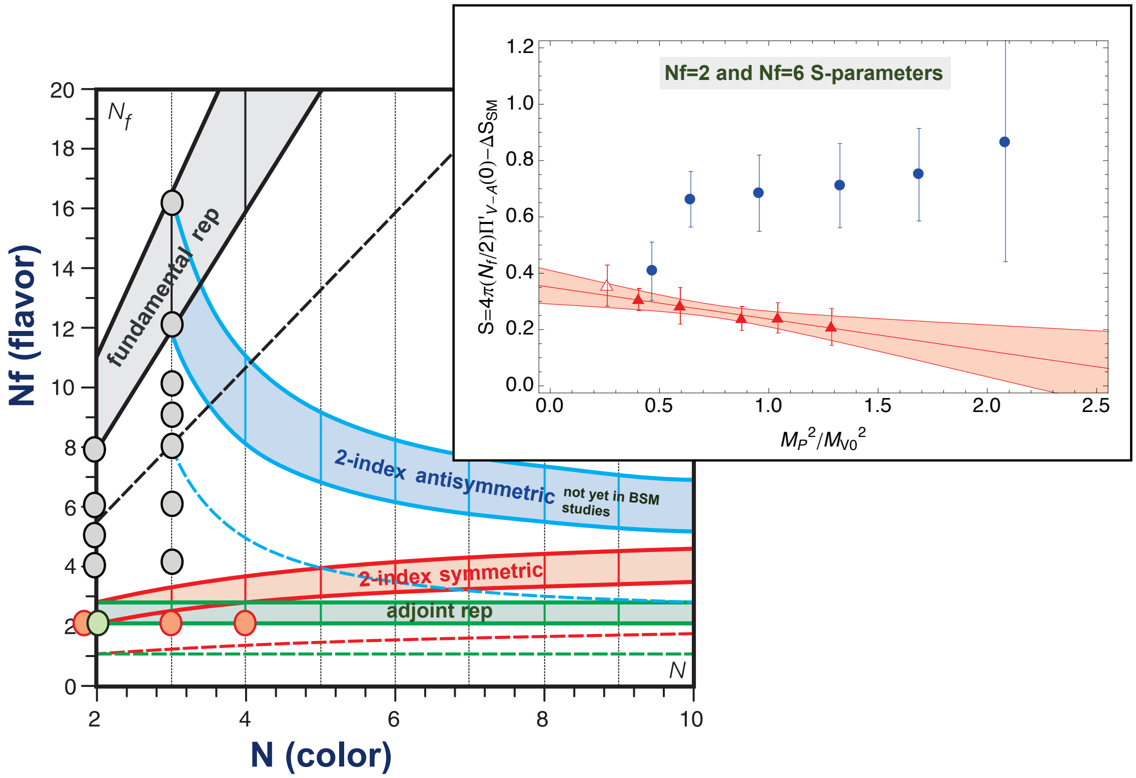

Figure 2: On the left side circles mark USQCD BSM activities in the large space of N_f and N_c with different $R(N_c)$ fermion representations color coded. Solid lines mark the upper and lower boundary of the conformal window from approximate analytic estimates [38], while the dashed line marks the lower boundary from the original Banks-Zaks prediction [34]. The right side plot of the S parameter from the LSD collaboration [48] is the first indication that the theory getting closer to the conformal window is not responding to the variation of N_f as it was expected from a scaled up version of QCD.

in several directions, including four-fermion interactions, like gauged NJL type theories, or models related to top quark condensates.

2.2.2 Supersymmetry and supersymmetry breaking

Supersymmetry (SUSY) was proposed as a possible symmetry of nature forty years ago. While it is an extension of the usual symmetries of space-time, such as translations and rotations, its consequences are dramatic; it predicts that every particle is accompanied by a *superpartner* of equal mass and charge but with spin differing by one-half. It thus unifies fermions with bosons. From a phenomenological point of view SUSY can explain why the Higgs particle responsible for breaking electroweak symmetry is light, while from a theoretical perspective it plays a crucial role in constructing consistent string theories, and may play an important role in understanding quantum gravity.

In supersymmetric models, SUSY is usually taken to be broken dynamically, with SUSY breaking communicated to the standard model sector by some mediation mechanism. Superpartners are expected to provide evidence of the mediation mechanism in their decays. For example, copious photon production in superpartner decay would be an evidence of gauge mediation. Light sleptons associated with superpartner decay would be a sign of anomaly mediation, and gravity mediation is associated with missing energy. If superpartners and evidence for a mediation mechanism is found, the investigation of the dynamical breaking of supersymmetry with lattice gauge theory will come to center stage.

Dynamical SUSY breaking and the soft parameters in the MSSM: It is straightforward to "supersymmetrize" the usual theories of particle physics; the Minimal Supersymmetric Standard Model (MSSM) is perhaps the most studied extension of the Standard Model. In this model the Higgs is naturally light since it is accompanied by a fermionic partner whose mass is protected by chiral symmetries. A great deal of effort has been expended on predicting signals for SUSY/MSSM at the LHC.

Of course the low energy world we inhabit is manifestly not supersymmetric; so a key component of any realistic theory of Beyond Standard Model physics must provide a mechanism for spontaneous supersymmetry breaking. In general a variety of no go theorems ensure that any such symmetry breaking must be non-perturbative in nature. In the MSSM the effects of this non-perturbative SUSY breaking is parametrized in terms of a series of "soft breaking terms" which are put in by hand and lead to a large parameter space and a lack of predictivity for the theory.

However, in general we might expect that these parameters are determined from dynamical breaking of SUSY at high energies in a "hidden sector". A supersymmetrized version of QCD - super QCD - with N_c colors and N_f massive flavors is a natural candidate for this hidden sector. For $N_c + 1 \leq N_f < \frac{3}{2}N_c$ it is thought that super QCD has long lived metastable SUSY breaking vacua [57]. The lifetimes of these metastable vacua can exceed the age of the Universe and ensure that the physical vacuum breaks supersymmetry. Within such a vacuum state non-perturbative phenomena, such as confinement and chiral symmetry breaking, precipitate a breaking of supersymmetry. Furthermore, if the quark masses are small compared to the confinement scale, these vacua have extremely long lifetimes. If the Standard Model fields are coupled to the hidden sector fields in an appropriate fashion, then such non-perturbative dynamics arising in the broken phase of this theory can feed down to yield soft supersymmetry breaking terms in the low energy effective theory - the MSSM.

Thus, a detailed understanding of the vacuum structure and strong coupling dynamics of super QCD can strongly constrain possible supersymmetric models of BSM physics, in some cases leading to detailed predictions of the soft parameters of the MSSM in terms of a handful of non-perturbative quantities obtained in the hidden sector super QCD theory. Lattice simulations of supersymmetric lattice QCD thus have the potential to play an important role in constraining the parameter space of the MSSM and in building realistic supersymmetric theories of BSM physics.

The technical problem that must be immediately faced in studying supersymmetric lattice theories is that supersymmetry is broken by discretization and it is non-trivial to regain SUSY as the lattice spacing is sent to zero with e.g. Wilson or staggered fermions. Luckily in the case of $\mathcal{N} = 1$ super Yang-Mills theory, which contains both gluons and their fermionic superpartners gluinos, this generic problem can be avoided by the use of domain wall fermions (DWF). In this case the exact lattice chiral symmetry of the fermion action ensures that SUSY is *automatically* recovered in the chiral limit.

Preliminary work by USQCD has already revealed a non-zero gluino condensate in the SU(2) theory in agreement with theoretical expectations [58, 59, 60]. However, to build a realistic theory capable of yielding the soft parameters of the MSSM, we will need to add N_f quarks and their scalar superpartners - squarks, and additionally extend the gauge group to a larger number of colors N_c.

To restore SUSY in the continuum limit now requires *both* use of DWF *and* tuning of parameters in the squark sector [61]. In principle this can be done by performing a series of runs over a grid in squark parameter space, and using (offline) reweighting techniques in the scalar sector to tune to the supersymmetric point.

We anticipate this program would proceed in a number of steps:

(1) Conduct studies of $\mathcal{N} = 1$ super Yang-Mills, which is arguably the simplest supersymmetric theory and the core of the MSSM. Studies of this theory for $SU(2)$ gauge group have already been started and the

relevant range of parameter space has been determined. A computation of physical observables, such as the gaugino condensate and low lying spectrum, are achievable within a couple of years. Code development to support an arbitrary number of colors of adjoint DWF will be necessary.

(2) Extend the previous calculations to super QCD. To tune to the supersymmetric limit a series of runs over a grid in the scalar (squark) parameter space will be necessary. DWF code with support for arbitrary numbers of flavors of fermions in the fundamental representation and additional Yukawa interactions will be needed. At each point in this parameter space a computational effort will be needed which is comparable to the $\mathcal{N}=1$ super Yang-Mills case.

Lattice supersymmetry and string theory: String theory evolved out of attempts to understand the strong interactions. However, these early string theories were quickly discarded with the development of QCD. In the last decade there has been a huge resurgence of interest in the connections between string theory and gauge theories such as QCD. This interest dates from a seminal paper by Juan Maldacena [62], which conjectured that a particular supersymmetric gauge theory – $\mathcal{N}=4$ super Yang-Mills – was dual or equivalent to type IIb string theory propagating on a five dimensional anti-de Sitter space.

The number of these so-called AdS/CFT dualities connecting QCD like theories with gravitational theories is now vast and examples exist in many dimensions, for many different space-times and for many types of conformal and non-conformal gauge theories. However, in most cases these dualities are conjectural in nature, and based on calculations in which the number of colors is taken large and the string theory is computed at low energy. Typically the dual Yang-Mills theories are strongly coupled and existing analytical techniques fail. This motivates the use of the lattice, which now offers a new non-perturbative tool to study theories of quantum gravity. This is the general focus of this second thread of work. Specifically we would like to develop codes to simulate $\mathcal{N}=4$ super Yang-Mills using two strategies.

(1) $\mathcal{N}=4$ super Yang-Mills is a special case of the super QCD theories discussed above, since it can be thought of as a special case of an $\mathcal{N}=1$ gauge multiplet coupled to three hyper-multiplets. We can thus hope to study it using domain wall fermions, fine tuning in this case exactly four quartic scalar operators. The computational cost is high, but the approach is somewhat conservative in the sense that the numerical algorithms are well understood and the calculations can piggyback on the LHC physics program described above.

(2) Over the last five years a number of exciting theoretical developments have taken place which have culminated in a lattice action for $\mathcal{N}=4$ super Yang-Mills in which the supersymmetry is exact even for non-zero lattice spacing [63]. This dramatically reduces the amount of fine tuning need to ensure the continuum limit is supersymmetric. To use this approach new codes would need to be developed. However, single core codes already exist and exploratory calculations have begun with some promising results [64].

These two approaches should allow us to make contact with the wealth of physics applications flowing from the AdS/CFT correspondence - for example the computation of black hole thermodynamics from gauge theory [65].

2.2.3 Physics driven BSM needs of SciDAC-3 software support

Our new SciDAC-3 software infrastructure will transform the early BSM lattice field theory program of USQCD into flexible rapid-response software solutions to answer challenges driven by new theoretical ideas and LHC discoveries. BSM lattice studies are extremely demanding computationally because N_c in the color gauge group $SU(N_c)$, the number of fermion "flavors" N_f, and the dimension of the group representation $R(N_c)$ for fermions vary in the theory space of the BSM paradigm. Some of the special needs of the new SUSY program were already outlined, and they also require the extensions of our application codes to general values of N_c, N_f and $R(N_c)$. There are some very specific physics driven goals of the new BSM software suite:

- With particular importance on the chiral properties of the BSM models, the domain wall, overlap, and staggered fermion solvers have to enable the use of several color groups in the fundamental, adjoint, two-index symmetric and two-index antisymmetric fermion representations.
- Using Random Matrix Theory, the determination of the low-lying eigenvalue spectrum and the related wave-functions of the fermion Dirac operators will help significantly in finite volume studies of BSM models. The eigenvalue spectrum is particularly important for the identification and detailed study of chiral symmetry breaking close to the conformal window. Thus, very large scale and robust eigen-solvers have to be designed and developed, addressing the issues of large grids, different lattice fermion choices, and different fermion representations.
- Studies of the running coupling will require the implementation of Schrödinger functional boundary conditions, twisted gauge and twisted fermion boundary conditions and renormalization group blocking methods in all the important fermion representations.

3 Computational Approach

In this section we briefly outline the computational approach used in our studies of QCD and BSM theories. They are quantum field theories defined in the four-dimensional space-time continuum. They have gauge symmetries analogous to that of quantum electrodynamics, except that the gauge groups are non-Abelian, SU(3) in the case of QCD, and more generally SU(N) for the BSM theories. In order to study these theories numerically one must reformulate them on four-dimensional lattices or grids. The gauge fields, which carry the forces, are represented by elements of the gauge group, with one element assigned to each link of the lattice. The matter fields, quarks in the case of QCD, are associated with the lattice points. As in all quantum field theories, physical observables can be expressed in terms of Feynman path integrals:

$$\langle O \rangle = \frac{\int \mathcal{D}U\, O(U)\, \exp[-S(U)]}{\int \mathcal{D}U\, \exp[-S(U)]}. \tag{1}$$

Here U represents the gauge field, and $\mathcal{D}U$ an integral over all of the components of U using the Haar measure. Because the matter fields are fermionic, they are represented in the path integrals by anti-commuting c-numbers, elements of a Grassman algebra. In the theories of interest, the action is quadratic in these variables, and they have been integrated out in the above expression. As a result, the effective action $S(U)$ is non-local. Indeed, it contains a term proportional to the logarithm of the determinant of the Dirac operator, which describes the propagation of a fermion in the gauge field U. On the lattice, the Dirac operator is a large sparse matrix with dimension proportional to the number of lattice sites. (The largest lattice we plan to work with in the next three years has just over 400 million sites). Finally, O is the physical observable being studied, and $O(U)$ is the value of O in the gauge configuration U.

The first step in the calculation is to use importance sampling techniques to generate an ensemble of gauge configurations, U_i, $i = 1, \ldots N$, with probability distribution

$$P(U_i) = \frac{\exp[-S(U_i)]}{\int \mathcal{D}U\, \exp[-S(U)]}. \tag{2}$$

Once an ensemble of representative gauge configurations is available, an unbiased estimator for any physical observable O is given by

$$\langle O \rangle = \frac{1}{N} \sum_{i}^{N} O(U_i). \tag{3}$$

One must generate ensembles of gauge configurations with several lattice spacings in order to perform extrapolations to the continuum (zero lattice spacing) limit, but each ensemble is used in a wide variety of calculations.

In generating an ensemble of gauge configurations, each configuration evolves from the one before it. It would be prohibitively expensive to use a local updating scheme in which one or a few links of the lattice are changed at a time because of the non-locality of $S(U)$. Instead, one uses algorithms, such as hybrid Monte Carlo (HMC) [66], or rational hybrid Monte Carlo [67], in which all components of U are updated simultaneously. The most time consuming step in this class of algorithms is the numerical integration of a set of molecular dynamics equations, which are non-linear, coupled first order differential equations of dimension proportional to the number of lattice sites. We are currently using a multi-time step symplectic Omelyan integrator [68] for this calculation; however, as described below, we propose to work on improved integrators for these equations.

At each step of the integration of the molecular dynamics equations one must solve a set of linear equations of the form

$$[D(U)^{\dagger}D(U)+\sigma_j I]x_j = b_j, \tag{4}$$

where $D(U)$ is the Dirac matrix, I is the unit matrix, and the σ_j are real positive numbers. These equations are currently solved using Krylov space methods. Their solution consumes the largest fraction of floating point operations in the generation of gauge configurations. The vast majority of the floating point operations in the measurement calculations, that is in the determination of $\langle O \rangle$ from Eq. 3, go into solving a linear equation of the form Eq. 4 with $\sigma_j = 0$, for which we again use Krylov space techniques. The systems of linear equations encountered in both configuration generation and measurements become increasingly ill-conditioned as the masses of the quarks decrease. For this reason, almost all QCD simulations up to now have used heavier than physical values of the masses of the two lightest quarks, the up and the down, and then relied on extrapolations in the light quark masses to obtain physical results. A major goal of this project is to perform simulations directly at the physical values of the light quark masses. The proposed work with the FASTMath SciDAC Institute to develop improved solvers for these equations will play an important role in reaching this goal.

There are a number of different formulations of lattice fermions in current use, all of which should give the same results in the continuum limit. In this project we will use domain wall fermions (DWF) and highly improved staggered quarks (HISQ) formulations for gauge configuration generation. In our measurement routines we will use these two formulations, as well as Wilson-Clover fermions for b-quarks, and overlap fermions for some of our studies of BSM theories. Each formulation has advantages for different components of our work. Furthermore, in order to demonstrate that we have reached the level of precision that we seek, it is important to obtain some results with different quark formulations. We have already done that in some instances, and will continue to do so.

The generation of gauge configurations is the chief bottleneck in our work. Because each configuration follows from its predecessor, this part of the calculation must be run in a single stream, or a small number of them. It is therefore highly desirable to generate configurations on the most capable available leadership class computers. By contrast, measurements can be performed on many configurations in parallel. This phase of our work poses a capacity challenge. It can be run on both leadership class computers, and on the dedicated clusters that are funded for our research through the LQCD Computing Project.

4 SciDAC-1 and SciDAC-2 Software and Algorithm Accomplishments

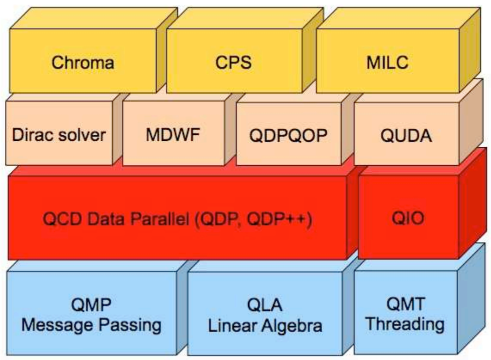

Figure 3: Levels of the QCD API

Under its SciDAC-1 and SciDAC-2 grants, the USQCD Collaboration developed software and algorithmic infrastructure for the numerical study of lattice gauge theories. This work was carried out jointly by high energy and nuclear physicists within USQCD, in collaboration with applied mathematicians and computer scientists. The software and its documentation is publicly available at the USQCD software web site *http://www.usqcd.org/usqcd-software*. The code has been widely adopted within the United States, and is used extensively abroad. It has been instrumental in our effective use of leadership class computers, and of the dedicated computers funded for USQCD through the LQCD Computing Project. The committees which review our hardware program on a yearly basis have consistently emphasized the importance of the work done under our SciDAC grants, and the need for their continuation. The work proposed for SciDAC-3 builds upon our accomplishments under our two SciDAC grants, so we briefly summarize them here.

4.1 The QCD Applications Programming Interface

Under our SciDAC-1 and SciDAC-2 grants, the USQCD Collaboration created the QCD Applications Programming Interface (QCD API), a unified programming environment that enables its users to quickly adapt existing codes to new architectures, easily develop new codes and incorporate new algorithms, and preserve their large investment in existing codes. It has greatly facilitated the efficient use of leadership class computers and commodity clusters. The QCD API was developed as a layered structure which is implemented in a set of independent libraries. It is illustrated in Fig. 3, which shows the three levels of the API and the application codes that sit on top of them. Extensions to these libraries and the maintenance of the API code is an ongoing activity as computer architectures and algorithms change. The API is a critical software underpinning for all of our community application codes, which requires maintenance, testing, version control, documentation and distribution.

Level 1 provides the code that controls communications and the core single processor computations. To obtain high efficiency, sometime much of this layer has to be written in hardware specific assembly language; however, versions exist in C and C++ using MPI for transparent portability of all application codes.

Message Passing: QMP defines a uniform subset of MPI-like functions with extensions that (1) partition the QCD spacetime lattice and map it onto the geometry of the hardware network, providing a convenient abstraction for the Level 2 data parallel API (QDP); (2) contain specialized communication routines designed to access the full hardware capabilities of computers, such as the Blue Gene line, and to aid optimization of low level protocols on cluster networks. New versions are developed as needed to accommodate changing architectures and algorithms. For example, as discussed below, hooks to combine message passing and threaded code are being added, as is the ability to work with multiple lattice geometries, which is needed for multigrid and domain decomposition algorithms.

Linear Algebra: All lattice QCD calculations make use of a set of linear algebra operations in which the basic elements are three-dimensional complex matrices, elements of the group SU(3). These operations are local to lattice sites or links, and do not involve interprocessor communications. The C implementation has about 19,000 functions generated in Perl, with a full suite of test scripts. The C++ implementation makes considerable use of templates, and so contains only a few dozen templated classes (the required specific

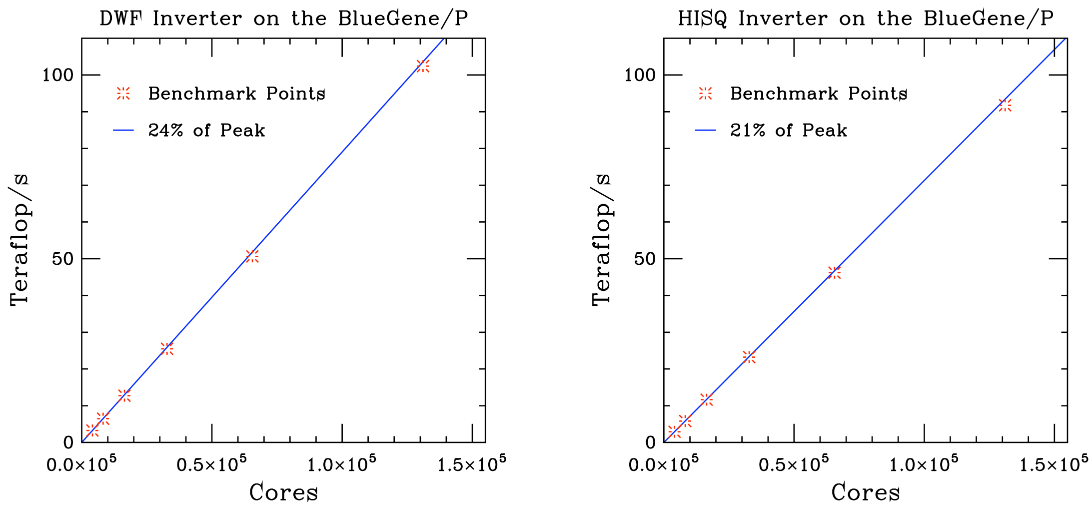

Figure 4: Performance of the Dirac solver on the Blue Gene/P in Tflops as a function of the number of cores for DWF quarks (left panel) and HISQ quarks (right panel). The red bursts are the benchmark points, and the solid blue lines indicate 24% and 21% of peak, respectively. These are weak scaling tests with the number of lattice points per core being held fixed at 6^4 for the DWF solver, and 8^4 for the HISQ solver.

classes are generated on demand by the compiler). For both C and C++ it is important to optimize the code for the most heavily used linear algebra modules.

Data Parallel Interface: Level 2 (QDP) contains data parallel operations that are built on QIO (parallel I/O) and QMP. There are both C and C++ versions of QDP. The C version is built on QLA. QDP allows extensive overlapping of communication and computation in a single line of code. By making use of the QMP and QLA layers, the details of communications buffers, synchronization barriers, vectorization over multiple sites on each node, etc. are hidden from the users, allowing them to focus on the physics, rather than the subtleties of parallel programming. QDP significantly accelerates the process of developing new codes and optimizing existing ones. It also lowers barriers for entry into the field by graduate students, postdoc and senior scientists from other fields.

Optimized Subroutines: Level 3 (QOP) consists of highly optimized code for a limited number of subroutines that consume a large fractions of the resources in any lattice gauge theory calculation. Most notable among these is the subroutine for the solution of the linear, sparse matrix equations involving the Dirac operator discussed in Sect. 3. To obtain the level of performance at which we aim, it is necessary to optimize these subroutines for each architecture. These routines are generally written with extensive assembly language coding, either employing hand coding or specialized tools, such as Bagel [69] and QA(0) [70], which were developed to generate optimized codes. The data mapping and cache efficiency is extensively tuned. In Fig. 4 we show the performance of the Dirac solver for DWF and HISQ quarks on the Blue Gene/P.

Data Management: QIO enables users to read and write the different types of files that arise in our work in standard formats. It supports a logical partitioning of the computer into I/O partitions with one core per partition handling I/O for the data in just that partition. Thus, in a suitable files system our codes can read and write data in parallel from/to a single file, or in any file system from/to multiple files, and these files can be flattened into one large one offline on a single processor machine. There are no unusual memory requirements for this process. By tuning the size of the I/O partitions, we can maximize the I/O bandwidth and avoid contention. In order to maximize the physics output from the very large computational resources that go into the generation of gauge configurations, we share all gauge configuration files that are created with USQCD resources. To enable this sharing we have created standards for file formats, which

QIO adheres to. In addition, we are charter members of the International Lattice Data Grid (ILDG), which established a basic set of meta-data and middleware standards to enable international sharing of data [71, 72], which are also adhered to by QIO.

Application Codes: There are three large, publicly available application code suites developed by members of USQCD that take advantage of the QCD API. Chroma was built directly on QDP++, while the Columbia Physics System (CPS) and the MILC code predate the API, but incorporate key features of it. As a result, all three applications suites benefit immediately from any extensions to or optimizations of the QCD API. Among them, these suites contain all of the codes required for the QCD configuration generation and measurement campaigns we intend to carry out over the next three years. The application code suites and their documentation can also be found at the USQCD software web site.

4.2 Recent Extensions of the QCD API

Although the QCD API and the application codes are highly portable, as we move to new computers, we typically have to upgrade the Level 1 and QOP routines. The advent of computer nodes with large numbers of cores and the use of GPU accelerators on the nodes have required that we develop threaded versions of our codes. Furthermore, we, and others in our field, regularly develop new algorithms which must be integrated into the API and the application codes. These developments require that we continually upgrade and extend the QCD API. Here we give a few highlights of this phase of our work.

Hybrid MPI/Threaded Code: It seems clear that in the near future computer nodes will contain large numbers of cores, and that for such machines we will need to employ a hybrid programming model in which communication between nodes is programmed in MPI or QMP, and work on nodes is performed with threaded code. At this early stage we do not believe that a "one size fits all" approach is possible, so we are experimenting with a variety of them. We have obtained early access to the Blue Gene/Q because members of our collaboration at Columbia University and Brookhaven National Laboratory, and our international collaborators at the University of Edinburgh, worked with colleagues at IBM on its design. They are well along in the development of code for domain wall fermions, and find that a hybrid MPI/OpenMP approach works well. They have also produced a highly optimized DWF solver using the Bagel tool [69], which produces assembly code for the Blue Gene/Q's PowerPC processors. Similar code for HISQ and Wilson-clover quarks will follow. For GPU accelerators, we are using CUDA threads on the GPUs combined with POSIX threads on the CPU, and MPI between nodes, while for computers with Intel and AMD multi-core processors, such as the Cray XE series, we have implemented a new threaded library, QMT. Our long range goal is to provide a single uniform data parallel interface so that the applications programmer does not need to be aware of the details of the hybrid code. In Fig. 5 we show strong scaling results for threaded code on NERSC's Cray XE6, Hopper, and on the Edge cluster at LLNL.

The QUDA GPU Library: Starting in 2008, we have explored high performance Dirac solvers in CUDA on NVIDIA GPUs [73]. This effort was initially supported by NSF funding, but has rapidly expanded into a major SciDAC project with the development of the QUDA (QCD in CUDA) library [74, 75, 76], and the rapid deployment of GPU accelerated clusters at Jefferson Laboratory and Fermilab. Our ability to respond rapidly to this new architecture demonstrates the advantage of our clear factorization of Level 3 solvers in the QCD API. At present the QUDA library has expanded to include all Dirac solvers used in QCD (Wilson-Clover, HISQ/asqtad, domain wall and twisted mass). The result has been a dramatic improvement in price/performance for a range of analysis work that is dominated by Dirac solvers. The most recent advance has been the extension of code from single to multiple GPUs. The multiple-GPU codes enables us to analyze the full set of lattices sizes generated by USQCD members with excellent weak scaling. In a paper presented to Super Computing 2011 we demonstrated that we have achieved good strong scaling on up to 256 GPUs for the HISQ/asqtad and Wilson-Clover solvers running on the Edge cluster at LLNL [77].

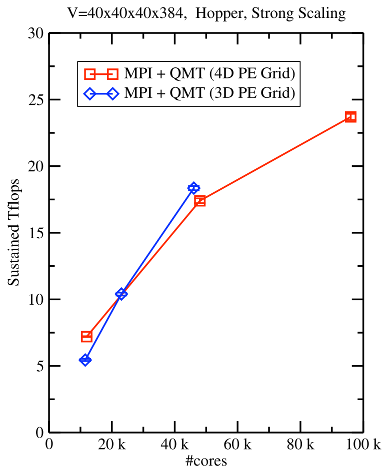

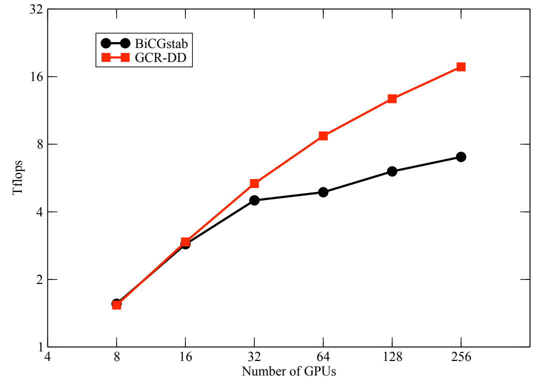

Figure 5: Strong scaling tests of the Wilson-clover inverter with threaded codes. On the left, results on the NERSC Cray XE6, Hopper, using hybrid code with MPI and the QMT library; and on the right, results on the LLNL Edge cluster for up to 256 GPUs with MPI and CUDA threads.

Solvers for the Dirac Operator: The solver for the lattice Dirac operator has traditionally been a dominant focus of algorithm and specialized software development because of its central role in all QCD codes. A large variety of Krylov solvers have been used with the conjugate gradient and BiCGStab being the current work horses in many production codes. Data layout to improve cache behavior and hand coded assembly kernels are commonplace. For example the Möbius [78, 79] domain wall fermion (MDWF) solver uses Morton ordering in its internal data representation [70], and the QUDA code employs specializes mappings and a novel mixed precision schemes from half (16 bits) to single (32 bits) to double (64 bits) in order to provide double precision accuracy with reduced data traffic between the processor and the memory. A new area of activity beginning to show great promise is the use of multigrid methods [80]. Lattice gauge theorists have attempted to apply multigrid methods to QCD for over twenty years [1]. In collaboration with applied mathematicians from TOPS, we have finally succeeded in formulating an adaptive multigrid solver for Wilson-clover [2]. In the left hand panel of Fig. 6, we show the speedup in the time for one additional solution provided by the multigrid solver compared with our best BiCGStab Krylov solver – nearly a 25x speed up as we move to the physical light quark mass limit. The multigrid algorithm has an overhead to construct its preconditioner, and in the right hand panel of Fig. 6 we show the number of solves of Eq. 4 with different right hand sides needed to amortize this overhead sufficiently so that the multigrid solver outperforms the BiCGStab one. In some measurement routines, such as those involving disconnected diagrams, hundreds of solves are required on each configuration, so the multigrid algorithm already offers a major improvement over BiCGStab. For the physical light quark mass, the crossover occurs for two or three solves, leading to the possibility of using multigrid in our configuration generation work. This is the beginning of a new opportunity for multi-level algorithms for other parts of our code, and will become increasingly important as the quark masses are reduced and lattice sizes increased. In this same spirit, we are exploring and implementing a variety of "deflation" and Schwarz domain decomposition methods [81].

Improved Hybrid Monte Carlo Evolution: Besides the Dirac solvers, the other major consumer of floating point operations in lattice field theory codes is the symplectic integrator for the molecular dynamics equations that arise in the hybrid Monte Carlo algorithms used to generate gauge ensembles. Over the period of the SciDAC grants, a major advance has been the development of the Rational Hybrid Monte Carlo (RHMC) [67], which is implemented in all of our major application codes. It typically results in a two to four times speedup in the generation of gauge configurations. An even higher order symplectic Force

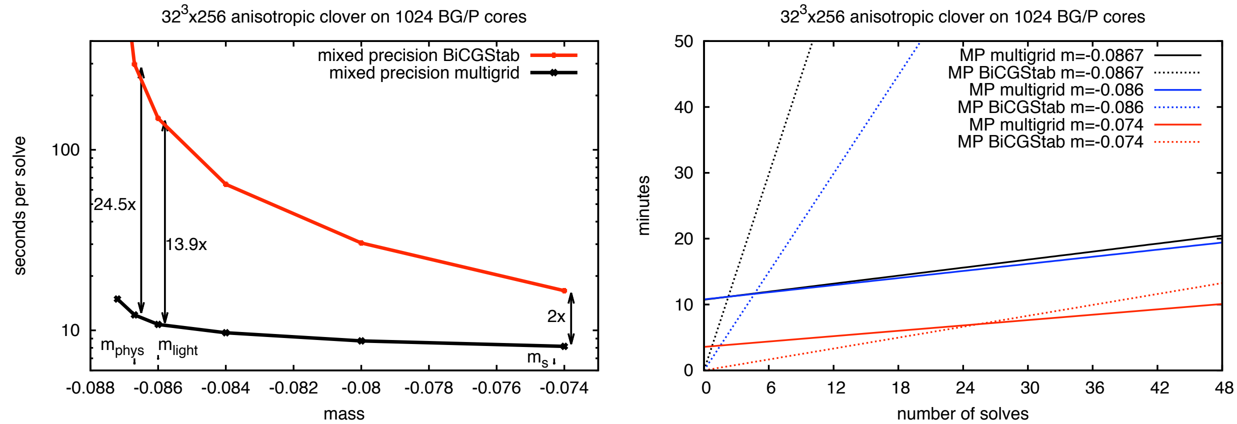

Figure 6: The left panel shows the marginal wall-clock per solve for the multigrid algorithm compared with our best BiCGStab Krylov solver on the Blue Gene/P for a $32^3 \times 256$ lattice with the Wilson-Clover Dirac operator. The right panel shows the total time including the setup for multiple solves on the same configuration by the multigrid and BiCGStab algorithms as a function of the number of solves [77].

Gradient integrator has been designed [82], which promises further improvements in the next generation of gauge configurations on very large lattices.

5 Software and Algorithm Development Plan

In this section we describe the software and algorithm infrastructure that is required to reach the scientific goals set out in Section 2. As is indicated in that section, for the flavor physics work at the intensity frontier, we need codes for the generation of QCD gauge configurations with DWF and HISQ quarks that will run with high efficiency on the computers we expect to use over the next three years. For work on BSM theories at the energy frontier, we again need codes to generate gauge configurations with DWF and HISQ fermions, but in this case for a range of colors, numbers of fermions and fermion group representations. In both cases, the highly optimized Dirac solvers that are critical for configuration generation, will also be of utmost importance for measurement routines run on the configurations. Finally, for some of our QCD measurement routines involving heavy quarks, we will need Wilson-Clover solvers, and for some of our BSM routines, overlap fermion solvers.

The increasingly heterogeneous nature of the computers we expect to use during the grant period, the Blue Gene/Q, the Cray XE/XK, and commodity clusters with GPU accelerators, will require us to make extensive upgrades in the QCD API and in our applications codes. Up to now, we have achieved load balancing on parallel computers by assigning identical subsections of our lattices to each compute core, and using QMP for message passing between cores. For the heterogeneous computers of the petascale and exascale eras, we anticipate using a hybrid programming model in which QMP is used for message passing between nodes, and threaded coding is used on nodes. As indicated in the last section, we have already made a start in this direction, but we must now fully integrate threading into the QCD API and our applications codes. In addition to optimizing our QCD codes for the machines at hand, we must extend the API to cover BSM theories, which, although similar in structure to QCD, in most cases involve different gauge groups and/or fermion representations than QCD. The very small lattice spacings and light quark masses at which we must work to reach the precision demanded by the science, as well as the changing computer architectures, require additional algorithm development. It is important to have improved Dirac solvers for both the evolution of gauge configurations and measurement routines, and improved integrators for the molecular dynamics equations encountered in configuration evolution. Below we discuss each of these subjects. The specific tasks to be undertaken by each participating institution are listed in Section 6.

5.1 Software for Lattice Gauge Theory at the Intensity Frontier

As was indicated in Section 2, our work at the intensity frontier will make use of gauge configurations with the DWF and HISQ quarks. Codes for generating these configurations are contained in the CPS and MILC application suites, respectively. In order to reach the small lattice spacings and light quark masses required by our work, we must optimize these codes for the most capable available leadership class machines. The Blue Gene line of supercomputers has proven particularly useful for lattice QCD studies, and we expect the Blue Gene/Q to enable major advances in our work. We have an allocation on the Argonne Blue Gene/Q, Mira during its Early Science Period, and plan to request substantial time on it thereafter through the DOE's INCITE Program. As mentioned previously, we have gotten an early start on optimizing our DWF code for the Blue Gene/Q because USQCD members at Columbia and BNL, and their international collaborators at Edinburgh, worked with IBM on its design. They already have a highly optimized Dirac solver for DWF quarks, and are working to optimize the remainder of their configuration code using OpenMP on the nodes and MPI for inter-node communication. Completing this development work, and performing similar optimization of the HISQ code are high priority goals of this proposal.

The Cray XE/XK computers also offer major opportunities for our work. We have a PRAC grant from the NSF that will provide early access to its Blue Waters petascale computer, a Cray XE/XK with over 3,000 GPUs. We also plan to apply for time on Oak Ridge's Titan, a Cray XE/XK with approximately 21,000 GPUs. We already have Dirac solvers running on multiple GPUs for all of the formulations of lattice quarks commonly used in QCD studies, except for DWF. Our DWF solver currently runs only on a single GPU, so it must be upgraded to multiple GPUs. A high priority goal of this project is to extend our GPU codes for DWF and HISQ quarks to perform configuration generation. This will require the development of routines for calculating the gauge and fermion forces in the molecular dynamics routines, and, in the case of HISQ quarks, the routine for calculating smeared gauge links. Our current codes have been tested on up to 256 GPUs, but to take full advantage of the Cray XE/XK machines that will become available in the next year will require that the configuration generation codes scale to thousands of GPUs. This will require considerable effort at optimization, and new approaches to eliminating communication bottlenecks. Two of our leading GPU code developers, Mike Clark and Ron Babich, have recently joined NVIDIA's Emerging Applications Group with the charge of working with us on this effort.

Since 2005 the Office of HEP and NP have funded dedicated computers for USQCD, which have been located at BNL, FNAL and JLab. In recent years, USQCD has acquired commodity clusters with GPU accelerators through this program, and we see this direction as growing in importance. So far, the accelerators have been used in measurement routines for which Dirac solvers completely dominate the calculations. The extension of the solvers to larger numbers of GPUs would have a major impact on our measurement calculations, strongly enhancing the return in the investment that the DOE is making in the hardware. Similarly, the development of GPU code for configuration generation would be very useful for exploratory work on moderate-sized lattices that are too small for leadership class computers. In addition, the plans of AMD and Intel to produce commodity processors with large numbers of cores makes the work we propose on hybrid codes very important for the hardware project. As previously noted, we propose to develop high performance code for Intel MIC accelerator cards, which we expect to be strong candidates for use in both our dedicated clusters and in future leadership class computers.

Finally, we note that our colleagues in nuclear physics plan to make use of DWF and HISQ quarks in their studies of the quark-gluon plasma. Exactly the same codes are used to generate zero- and high-temperature gauge configurations, only the ratio of temporal to spatial extent of the lattice changes. They are therefore likely to benefit significantly from the software we develop for configuration generation.

5.2 Software for Lattice Gauge Theory at the Energy Frontier

Like QCD, the theories that have been proposed for physics beyond the standard model have non-Abelian gauge symmetries. However, the number of gluon "colors" N_c, the number of fermion "flavors" N_f, and the dimension of the group representation to which the fermions belong $R(N_c)$, vary from theory to theory. They are all integers, and for QCD they take on the values $N_c = 8, N_f = 6$ and $R(N_c) = 3$. The mathematical equations that define the lattice calculations are algebraically identical to those of QCD once these numbers are specified. But, from a computational perspective, they can dramatically affect the performance of a code because these integers determine the size of data, the sparsity of the Dirac matrix and the ratio of floating point operations performed per byte of data.

We propose to extend our applications codes so they can deal with general values of N_c, N_f and $R(N_c)$. As in the case of QCD, we wish to create a code base that yields high performance on platforms that we anticipate using in the immediate future, and is highly portable. Given that BSM simulations are relatively new, and growing in importance, it is likely that a significant number of physicists new to the field will be taking them up in the near future. It is therefore particularly important that we develop codes that are well documented, straightforward to use by non-experts, and that can be easily modified to incorporate new ideas. As in our QCD studies at the intensity frontier, configurations generation at the energy frontier will involve the DWF and HISQ formulations of lattice quarks. Our existing codes for these formulations have been extended to handle general values of N_f, but accommodating variable N_c or $R(N_c)$ is more complex because it was not anticipated in the original design of the codes. In addition, several optimizations were implemented under the assumption that the target theory would always be QCD. The extension of the existing DWF and HISQ codes to general N_c and $R(N_c)$, while maintaining their high performance and portability, is a high priority goal of this project.

We plan to build a framework for the rapid development and testing of new algorithms for configuration generation, FUEL (Framework for Unified Evolution of Lattices). FUEL will facilitate combining the key routines necessary for lattice generation into a production gauge evolution code. The initial development will focus on code supporting a range of values of N_f, N_c, and R(N_c) that are necessary for simulations of BSM theories. An initial prototype of the framework is being built now on top of the existing SciDAC QOPQDP library of Level 3 routines. The top level control in FUEL is being done using the scripting language Lua [83] , which allows quick and easy modifications to the configuration generation methods and also provides a simple way to configure the application. Lua was chosen because it is an expressive and powerful language with small size, and ease of porting to new platforms. FUEL offers a unique opportunity to accelerate the development of BSM codes because it is designed to work for non-QCD values of N_f, N_c and $R(N_c)$, and to explore the large class of tuning parameters needed to optimize the symplectic integrators. Indeed our first prototype FUEL implementation has already proven useful in providing a convenient framework to tune $N_f = 8$ flavors BSM for the HISQ action. The next target to further test this framework will be the $\mathcal{N} = 1$ Super Symmetric Yang-Mills theory, because it is among the simplest to implement, and yet has all three of these variables different from their QCD values. The successful implementation of this theory requires little modification to also encompass most of the promising composite Higgs theories.

To complement the software effort to produce Dirac solvers and full evolution codes, we propose to develop very large scale and robust eigen-solvers for domain wall, staggered, and Wilson fermions. The main thrust of this effort is to extract the lowest eigenvalues and eigenvectors as required by the investigation of chiral condensates, including the close connection with Random Matrix Theory. BSM applications on the energy frontier will require eigen-solvers for $SU(N_c)$ color gauge groups with $N_c = 2,3,4$ and fermions in the fundamental, adjoint, and two-index symmetric representations. This project will provide a stand-alone package for off-line analysis with interface to the gauge configurations which will be generated in BSM simulations using our newly developed evolution codes.

5.3 Algorithms for New Physics and New Architectures

In our field, as in many others, improvements in algorithms have typically played as large a role in advancing science as improvements in the capability of computers. This has been the case during the period of our SciDAC-1 and SciDAC-2 grants, and we expect that algorithm development will be particularly important in the coming years for two reasons. First, with the advent of petascale computers, we will be able to perform simulations with small lattice spacings at the physical masses of the quarks. This development will significantly increase the precision of our calculations, but to take advantage of it we need to begin to make use of multi-level algorithms that will enable us to separate diverse scales. (QCD simulations involves a wide range of scales from the mass of the u quark, 2 MeV, to the nucleon mass, 10^3 MeV, to the mass of the b quark, 4×10^3 MeV). Second, computers on the immediate horizon will have an even greater ratio of floating point performance to data movement rate than current ones, which means that communications will become an even greater bottleneck than it already is. It is thus important to explore algorithms that minimize data movement.

Our highest priority is to develop improved Dirac solvers, which consume the bulk of the floating point operations in both gauge configuration generation and measurement routines. As indicated above, we now have a very effective multigrid solver for the Wilson-Clover formulation of lattice quarks [1], and we propose to extend this technique to DWF and HISQ quarks. Each of these alternative discretizations of the Dirac differential operator pose new and mathematically interesting challenges to adaptive multigrid methods because of the special properties of the low eigenvectors in the near null space on the lattice. The domain wall operator has an extra fifth dimension with chiral near null vectors confined geometrically to four-dimensional walls, while the staggered operator has four times the size of null space due to the classic "doubling problem". HISQ quarks are particularly challenging for multigrid methods because they involve a local averaging or smearing of the gauge field, which reduces locality and therefore increases communication. An alternative to multigrid, which we plan to explore, is Schwarz domain decomposition approach, which can be specifically designed to minimize communication [81, 84]. As illustrated on the right side of Fig. 5, we have had some success with the simple Block Jacobi version of domain decomposition [77], and expect significant speedups using more sophisticated approaches, such as overlapping blocks, and block Gauss-Seidel. Work on the development of new solvers, and their integration into our production codes will be carried out in collaboration with our partners in the SciDAC FASTMath Institute and the NVIDIA Emerging Applications Group. The goal is to learn how to marry the algorithmic efficiency of multi-scale solvers with the architectural advances of many-core technology. This integration is a challenge but initial software design for porting the Wilson-Clover multigrid algorithm to the GPUs in the weak scaling limit promises a multiplicative cost advantage: O(10) algorithmically and O(10) for GPUs for an estimated total of two orders of magnitude reduction in cost per solve.

Another high priority goal is to improve the symplectic integrators used in the molecular dynamics evolution of the gauge fields. Up to now we have used second order integrators, either leap frog (Verlet) or Omelyan [68]. The cost of maintaining a constant integration error, and therefore a constant Monte Carlo acceptance rate with these integrators grows as $V^{5/4}$. As we move to very large lattices, this super-linear scaling will become increasingly costly. A fourth order integrator (force gradient) has recently been introduced into lattice QCD [85]. Using it, the HMC or RHMC configuration generation algorithms scale as $V^{9/8}$. Numerical experiments indicate that cross-over point at which the force-gradient integrator becomes more efficient than the leap frog or Omelyan ones is $V \approx 50^4$ [85], exactly the region we plan to explore in this project. It is thus important to deploy the force gradient integrator in our codes. However, due to the complicated structure of the force computations, automated code generation is necessary.

In order to facilitate algorithm development, members of USQCD at LLNL and applied mathematicians from the FASTMath Institute at LLNL will provide a framework for fast turn-around exploration and testing of novel algorithmic ideas for lattice field theory that can then be incorporated into the USQCD science codes. In addition this collaboration will be able to test these ideas at scale on LLNLs massively parallel supercomputers as available to the FASTMath and LLNL lattice group for algorithmic development. In particular the focus will be in interfacing lattice QCD kernels with the FASTMath HYPRE applied mathematics effort. HYPRE stands for High Performance Pre-conditioners, and it provides novel scalable linear solvers for sparse linear systems of equations on massively parallel supercomputers. HYPRE represents a large and continuing investment by the applied math community. This collaboration aims to test existing and new algorithms on various lattice QCD kernels in HYPRE leveraging appropriate USQCD infrastructure to access and verify results. The algorithms that give performance improvements will be integrated into the USQCD production software by lattice gauge theorists. Furthermore, extensions to the HYPRE code developed for this project will be incorporated into the main branch, and made available to future users of HYPRE.

6 Tasks and Milestones

Here we present the software tasks and three year milestones for this proposal broken down by the institutions that will carry them out. Much of the work will involve collaborative efforts between participating institutions, as was the case in our previous SciDAC grants. Work of those funded by this grant will be supplemented by that of unfunded senior physicists within USQCD, who will continue to contribute significantly to the software and infrastructure efforts.

Boston University: As software coordinator and Chair of the Software Coordinating Committee, Richard Brower has the responsibility to guide the overall software and algorithm infrastructure project. In addition, Boston University will continue to lead the projects to develop multi-scale algorithms for Dirac solvers and to develop the high performance QUDA library [86] for GPUs. Since their genesis at Boston University under SciDAC-2, these projects have expanded greatly, as has the role of Boston University in co-ordinating activities among participating group, including most prominently the new efforts with the FASTMath Institute and the NVIDIA Emerging Technology group. These developments require a sequence of steps: (i) generalizing the Wilson-clover multigrid algorithm to domain wall and staggered quarks, (ii) implementing these inverters on multi-GPU clusters with appropriate domain decomposition methods and (iii) integrating multi-level solvers (multigrid and domain decomposition) into HMC evolutions codes. Boston University will collaborate on the following tasks:

- *Year 1:*
  – complete multiple–GPU implementation of domain wall solver
  – develop Level 3 code and documentation for the multigrid solver for Wilson-Clover fermions
  – work with NVIDIA to optimize this multigrid solver on multiple–GPU clusters
  – further optimize the Wilson-Clover multigrid in the HYPRE framework with FASTMath partners
  – develop first version of Level 3 multigrid solver for domain wall fermions.
- *Year 2:*
  – develop multigrid solver for Overlap fermions
  – work with NVIDIA to develop multi-precision multigrid solvers for domain wall fermions on multiple-GPU clusters
  – study communication mitigation of hybrid multigrid/domain decomposition solvers with FASTMath collaborators.
- *Year 3:*
  – introduce multigrid/domain decomposition solvers for Blue Gene/Q and cluster architectures
  – test & optimize multigrid/domain decomposition evolution in the FUEL framework in collaboration with ANL

– work with NVIDIA to optimized multigrid/domain decomposition code for GPU systems (Titan, BlueWaters, USQCD Clusters).

Argonne National Laboratory and Syracuse University: Argonne National Laboratory (ANL) and Syracuse University will take the lead in developing a new generation of software required for the study of BSM theories. We propose to support a broader research program for the study of BSM theories. While work on optimizing existing codes will continue, we will develop software to explore theories with a wider range of colors, fermion numbers and fermion representations. A key element of this effort will be the development of the Framework for Unified Evolution of Lattices (FUEL), a tool that will enable the rapid development and testing of new algorithms for configuration generation. In addition ANL will contribute to the multigrid collaboration with FASTMath and to the optimization of codes for the Blue Gene/Q.

- *Year 1:*
  – basic implementation of HMC for SU(2) and SU(3) theories with fundamental representation HISQ, Wilson and domain wall fermions working in FUEL
  – develop optimized code for $\mathcal{N} = 1$ super Yang-Mills for adjoint representation domain wall fermions
  – with others, work on optimizing QMP and QLA for BG/Q
  – investigate HMC integrator improvements for large N_f
  – integrate existing GPU routines into FUEL
  – collaborate with FASTMath and others to get current Wilson-clover multigrid algorithm implemented in HYPRE.

- *Year 2:*
  – investigate Hermitian eigensolver routines for HISQ and domain wall quarks
  – work on adding 2-index symmetric routines to QOPQDP for use in FUEL
  – add nHYP and Stout variations of staggered fermions to QOPQDP – investigate force-gradient integrators in FUEL
  – inclusion of support in Dirac operators for Yukawa interactions necessary for super QCD and general four fermion interactions.

- *Year 3:*
  – add Schrödinger functional support to FUEL
  – develop SU($N_c > 3$) code for Wilson, domain wall and staggered fermions
  – optimize SU(2) and SU($N_c > 3$) codes on BG/Q and clusters
  – collaborate with FASTMath and others on multigrid for staggered and domain wall fermions, and integrate working multigrid codes into FUEL
  – finish support for full super QCD
  – Develop code for $\mathcal{N} = 4$ super Yang-Mills using twisted lattice formulations.

NVIDIA: USQCD has entered into a partnership with the Emerging Applications Group at NVIDIA to develop high-performance lattice gauge theory code for GPU systems (see Sec. A.2). This work will build on the the QUDA library, developed in part with SciDAC-2 support. The first goal, nearing completion, was the development of a set of highly optimized Krylov solvers suitable for running on hundreds of GPUs and supporting all of the most common discretizations of the Dirac operator. The next steps, which will be carried out under this grant in collaboration with members of the NVIDIA Emerging Applications Group, is to develop GPU code for the evolution of gauge configurations using our current algorithms, and then to devise architecture-aware multigrid and domain decomposed algorithms for use on the next generation of capability systems, such as Titan, Blue Waters and clusters with GPU accelerators. The major goal is the development of lattice field theory evolution codes capable of scaling to systems with thousands of GPUs and sustaining hundreds of Tflops or higher.

- *Year 1:*
  – development of adaptive multigrid solvers in QUDA (Wilson-like fermions)
  – refinement of domain decomposition algorithms to improve strong scaling
  – complete deployment of HISQ gauge configuration generation on GPUs
  – develop optimum cache reuse strategies for current GPUs and influence future architectures to improve scaling.
- *Year 2:*
  – develop adaptive multigrid solvers in QUDA (HISQ fermions)
  – develop hybrid domain decomposition and multigrid algorithms
  – complete deployment of Chroma gauge generation code on GPUs.
- *Year 3:*
  – combine multi-scale (multigrid and domain decomposition) methods with HMC evolution code for strong scaling on capability multi-GPU systems (long term R&D).

Brookhaven National Laboratory: BNL scientists, along with members of the RBC and UKQCD collaborations, are strongly focused on zero and finite temperature simulations of DWF QCD, making the ability to efficiently generate DWF gauge configurations on the Blue Gene/Q using CPS a primary goal of their SciDAC work. While an efficient, assembly level DWF solver is already written (Boyle, UKQCD) and threading the remaining evolution code with OpenMP has begun, there is substantial work that still must be done under this proposal, including multi-threading the measurement codes in CPS. In addition, they will work with others in USQCD to implement QMP over SPI for the Blue Gene/Q for efficiency, and to improve QIO for new architectures, including greater control over the number of processes actively writing to the file system. Existing algorithm research, which will focus on disconnected diagrams for chiral fermions, the force gradient integrator and four-dimensional realizations of DWF, is vital to many of the group's physics goals ($K \to \pi\pi$ decays and $(g-2)_\mu$ calculations, for example), and will be continued.

- *Year 1:*
  – finish optimization of DWF evolution code on the Blue Gene/Q
  – begin improvements to QMP and QIO with others in USQCD
  – produce efficient implementation of EigCG for the Blue Gene/Q
  – begin multi-threading of all CPS measurement codes.
- *Year 2:*
  – finalize improvements to QIO and QMP
  – continue to improve CPS, finalizing multi-threading of measurement codes
  – implement deflation techniques in CPS.
- *Year 3:*
  – continue to improve the higher-level organization of CPS
  – implement results of algorithm research in CPS.

University of Arizona, Indiana University and University of Utah: The MILC collaboration code base provides the foundation for the research programs of several groups within USQCD, as well as a number of groups abroad. MILC's current research is focused on the use of staggered (asqtad and HISQ) quarks to study QCD. It developed the SciDAC MPP molecular dynamics algorithms for the HISQ action and continues to support them. As physics goals, algorithms, and machine architectures evolve, the code also evolves. The arrival of radically new multi-core architectures requires some major changes to the code over the next three years. These changes will be built on the SciDAC utilities that MILC and others in this proposal are developing.

- *Year 1:*
  – with Illinois, NVIDIA, and Boston University, develop multiple-GPU support for the four main routines needed for the generation of gauge configurations, including HISQ extensions.
  – integrate the GPU codes into the MILC code suite
  – with Fermilab, test and validate the GPU codes in a production environment
  – develop a QOP QUDA that supports QOP modules required for HISQ and Wilson-Clover fermions and HISQ configuration generation
  – build a trial version of QIO that supports MPI-I/O
  – provide MILC code support for multi-threaded QOP in collaboration with ANL, identify remaining bottlenecks, publish analysis of performance.

- *Year 2:*
  – with Illinois, NVIDIA and Fermilab, optimize QUDA codes developed in year 1 and eliminate the principal bottlenecks
  – eliminate principal bottlenecks in multi-threaded QOP code identified in year 1, and publish analysis of performance improvement
  – with ANL, test the SPI implementation of QMP, publish analysis of MPI *vs* SPI performance.
  – with Fermilab, test a trial Lua version of the MILC code, publish analysis of performance improvements
  – with BU and ANL explore the applicability of the force-gradient integrator and multigrid methods for HISQ and the improved action for heavy-quarks.

- *Year 3:*
  – with Illinois and NVIDIA, develop QUDA versions of successful new algorithms identified in year 2
  – integrate new QUDA modules into the MILC code, test and validate
  – create multithreaded versions of successful new algorithms identified in year 2
  – with Fermilab, expand Lua support for the MILC code, should it prove successful.

U. of Illinois: Guochun Shi of the University of Illinois' Innovative Systems Laboratory, has worked with members of the MILC Collaboration and the Boston University QUDA group for over two years to port key elements of the MILC code to GPUs. He played a major role in porting the Dirac solvers for the asqtad and HISQ formulations of staggered quarks to GPUs. Under this grant he will work with members of MILC and the broader USQCD QUDA group, as well as with colleagues at NVIDIA, to port the remaining routines needed to generate HISQ configurations on multiple GPUs, the gauge and fermion forces, and the smeared links.

- *Year 1:*
  – in coordination with MILC and the rest of the QUDA group, provide the HISQ enhancements needed for running on multiple GPUs
  – port the gauge force routines to multiple GPUs
  – tune the performance of the new routines
  – begin work on tuning for the new NVIDIA Kepler GPU that will be used on Titan and Blue Waters, and very likely on future USQCD clusters.

- *Year 2:*
  – complete tuning of codes for Kepler architecture
  – develop a more exible API that will enable the programmer to indicate when data needs to be moved to the GPU to start the routine.

- *Year 3:*
  – complete implementation of the API developed in Year 2

– implement on GPUs promising new algorithms developed by other members of the USQCD collaboration, FASTMath and NVIDIA.

Lawrence Livermore National Laboratory and FASTMath Institute: Lattice field theorists from LLNL and applied mathematicians from the FASTMath institute at LLNL will provide a framework for rapid exploration and testing of novel algorithms for lattice field theory that can then be incorporated into USQCD codes. In addition, members of this collaboration will test these ideas at scale on LLNL's massively parallel supercomputers which are available to them. In particular, they will focus on interfacing lattice QCD kernels with the FASTMath HYPRE framework.

- *Year 1:*
  – expand HYPRE to include the use of complex numbers
  – expand HYPRE from the current three-dimensional implementation to an arbitrary, user defined dimension
  – implement the Wilson Dirac operator in HYPRE, and develop a solver using HYPRE's multigrid methods
  – compare performance of HYPRE multigrid solver with the existing USQCD Wilson multigrid solver on the Blue Gene/L.
  – develop collaborative work on eigen-solver packages in coordination with the ANL/Syracuse effort.
- *Year 2:*
  – implement and interface the major lattice QCD kernels with HYPRE, including those for the DWF and HISQ formulations of lattice fermions.
- *Year 3:*
  – use the infrastructure built in the first two years to explore new multigrid methods and algorithms for lattice QCD
  – exploit the fast turn-around times to explore a large range of algorithms.

Fermi National Accelerator Laboratory: In 2012 Fermilab will begin operating for USQCD a GPU-accelerated cluster designed to perform optimally on problems requiring large GPU-count parallelization and/or good strong scaling. To inform future USQCD cluster designs and software development for the various leadership-class GPU-accelerated machines, the performance of production applications on this dedicated cluster will be analyzed in depth. Conventional x86_64-based clusters continue to be critical USQCD resources, and new generations of processors from both Intel and AMD will have an impact; specific optimizations for these new processors will be implemented in the QLA library. Fermilab will collaborate with the MILC Collaboration on the application of the Lua scripting language [83] to the MILC code. Fermilab will work with the MILC Collaboration, Boston University, U. of Illinois, and NVIDIA on the development and optimization of software for GPU-accelerated clusters. Fermilab will also provide leadership in the development of optimized software for clusters accelerated with Intel many-core architecture hardware.

- *Year 1:*
  – comprehensive performance analysis of multiple-GPU production runs on the Fermilab USQCD GPU-accelerated cluster
  – performance analysis of multi-TFlop-scale staggered configuration generation on the GPU-accelerated cluster, collaborating with experts from NVIDIA, Illinois, and MILC
  – addition of Intel Sandy Bridge and AMD Interlagos optimizations to QLA library, and measurement of effect on code throughput. Optimizations will include use of AVX and XOP (FMA4) instructions
  – demonstrate Lua-based script version of a MILC production executable for calculation of three-point functions
  – refinement and packaging of the databases and tools developed during SciDAC-2 cluster reliability

sub-project for the monitoring of the batch schedulers deployed on USQCD clusters and the real time analysis of production running on these resources

- *Year 2:*
  – collaboration on the optimization of GPU software and the integration of GPU software into application codes, included performance analysis of production runs on Fermilab USQCD GPU-accelerated clusters
  – leadership in the exploration of the many-core Intel architecture and the guidance of the software optimizations for Intel many-core-accelerated clusters
  – QLA optimizations for Intel Ivy Bridge and AMD socket G2012 processors
  – general Lua-based implementations of MILC applications suitable for production

- *Year 3:*
  – continuing work on software optimizations for GPU-accelerated and Intel many-core-accelerated clusters
  – continued maintenance of QLA for x86_64 processors
  – with the MILC collaboration, further expansion of Lua support for the MILC code, should it prove successful
  – exploration of the use of domain specific languages for automated meta data collection and the direction of analysis workflows

7 Management Plan

The Project Director and lead principal investigator for this effort will be Paul Mackenzie. He will have overall responsibility for the project, and will be the point of contact for the Department of Energy. The co-Director for Computation, Richard Brower, will have overall responsibility for the software and algorithm work, providing direction and coherence to the work, and monitoring progress on all tasks. He will also coordinate work with other collaborators in USQCD, and our partners in the SciDAC FASTMath Institute and in the Emerging Applications Group at NVIDIA. He will provide quarterly reports to the Director on the progress of the software effort, and organize a yearly group meeting that will include our FASTMath and NVIDIA partners. The co-Director for Science, Stephen Sharpe, will track progress towards meeting the scientific goals of the project, and will advise the Director on new scientific opportunities that arise during the course of the project. He will also advise the Director and co-Director for Computation on priorities for software and algorithm development needed to advance the scientific program.

Dr. Mackenzie is Chair of the USQCD Executive Committee, whose other members are R. Brower, N. Christ, F. Karsch, J. Kuti J. Negele, D. Richards, S. Sharpe, and R. Sugar. The USQCD Executive Committee serves as the Advisory Board to the Director, advising him on scientific goals and priorities, the software and algorithm development needed to meet these goals, and the distribution of funds. This procedure will insure that the work performed in this project is coordinated with others efforts of USQCD. The Executive Committee has been been leading the effort to construct computational infrastructure for the U.S. lattice gauge theory community for over twelve years. It holds approximately two conference calls per month, and communicates via email between calls. The Executive Committee has been actively involved in the preparation of this proposal.

Each institution receiving funds under this grant has a local principal investigator, who has first-level responsibility for the work carried out at his institution. The principal investigators will report on progress to the Director and co-Directors during monthly conference calls. At the end of each project year, the Executive Committee will review progress towards meeting the scientic and computational goals of the project, and will advise the Director. Schedule slips of more than two months must be reported to the Executive

A Appendices

A.1 References Cited

[1] R. Babich, J. Brannick, R. C. Brower, M. A. Clark, T. A. Manteuffel, S. F. McCormick J. C. Osborn,and C. Rebbi, *Adaptive multigrid algorithm for the lattice Wilson-Dirac operator*, Phys. Rev. Lett. **105**, 201602 (2010) [arXiv:1005.3043 [hep-lat]].

[2] J. C. Osborn, R. Babich, J. Brannick, R. C. Brower, M. A. Clark, S. D. Cohen and C. Rebbi, *Multigrid solver for clover fermions*, PoS **LATTICE2010**, 037 (2010) [arXiv:1011.2775 [hep-lat]].

[3] R. Babich, J. Brannick, R. C. Brower, M. A. Clark, S. D. Cohen, J. C. Osborn and C. Rebbi, *The role of multigrid algorithms for LQCD*, PoS **LAT2009**, 031 (2009) [arXiv:0912.2186 [hep-lat]].

[4] J. Laiho, E. Lunghi, R. S. Van de Water, *Lattice QCD inputs to the CKM unitarity triangle analysis*, Phys. Rev. **D81**, 034503 (2010). [arXiv:0910.2928 [hep-ph]]; for updates see http:www.latticeaverages.org.

[5] G. Colangelo *et al*., *Review of lattice results concerning low energy particle physics*, Eur. Phys. J. **C71**, 1695 (2011). [arXiv:1011.4408 [hep-lat]].

[6] C. T. H. Davies *et al*. [HPQCD, UKQCD, MILC &FNAL Collabs.], *High precision lattice QCD confronts experiment*, Phys. Rev. Lett. **92**, 022001 (2004) [hep-lat/0304004].

[7] A. S. Kronfeld [FNAL, MILC & HPQCD Collabs.], *Predictions with Lattice QCD*, J. Phys. Conf. Ser. **46**, 147 (2006) [hep-lat/0607011].

[8] J. H. Christenson, J. W. Cronin, V. L. Fitch and R. Turlay, *Evidence for the 2 pi Decay of the k(2)0 Meson*, Phys. Rev. Lett. **13**, 138 (1964).

[9] S. Durr *et al*., *Precision computation of the kaon bag parameter*, Phys. Lett. B **705**, 477 (2011) [arXiv:1106.3230 [hep-lat]].

[10] A. J. Buras and D. Guadagnoli, *On the consistency between the observed amount of CP violation in the K^- and Bd-systems within minimal flavor violation*, Phys. Rev. D **79**, 053010 (2009) [arXiv:0901.2056 [hep-ph]].

[11] *Anatomy of New Physics in $B-\bar{B}$ mixing*, A. Lenz *et al*., Phys. Rev. D **83**, 036004 (2011) [arXiv:1008.1593 [hep-ph]].

[12] E. Lunghi, A. Soni, *Possible evidence for the breakdown of the CKM-paradigm of CP-violation*, Phys. Lett. **B697**, 323-328 (2011). [arXiv:1010.6069 [hep-ph]].

[13] E. Lunghi, A. Soni, *Demise of CKM and its aftermath*, arXiv:1104.2117 [hep-ph];

[14] S. Bethke *et al*., *Workshop on Precision Measurements of α_s*, arXiv:1110.0016 [hep-ph].

[15] S. Aoki *et al*. [PACS-CS Collab.], *Physical Point Simulation in 2+1 Flavor Lattice QCD*, Phys. Rev. D **81**, 074503 (2010) [arXiv:0911.2561 [hep-lat]].

[16] S. Durr *et al*. [BMW Collab.], *Lattice QCD at the physical point: light quark masses*, Phys. Lett. B **701**, 265 (2011) [arXiv:1011.2403 [hep-lat]].

[17] A. Bazavov *et al*. [MILC Collab.], *Simulations with dynamical HISQ quarks*, PoSLATTICE **2010**, 320 (2010) [arXiv:1012.1265 [hep-lat]].

[18] T. Blum *et al*., *The* $K \to (\pi\pi)_{I=2}$ *Decay Amplitude from Lattice QCD*, arXiv:1111.1699 [hep-lat].

[19] T. Blum *et al*., *K to* $\pi\pi$ *Decay amplitudes from Lattice QCD*, arXiv:1106.2714 [hep-lat].

[20] N. H. Christ, *Computing the long-distance contribution to second order weak amplitudes*, PoS **Lattice 2010**, 300 (2010) [arXiv:1012.6034 [hep-lat]].

[21] J. Yu, *Long distance contribution to* K_L*-*K_S *mass difference*, arXiv:1106.2714 [hep-lat].

[22] S. Basak, *et al*., [MILC Collaboration], *Electromagnetic splittings of hadrons from improved staggered quarks in full QCD*, PoS **Lattice 2008**, 127 (2009) [arXiv:0812.4486];

[23] A. Torok, *et al*., [MILC Collaboration], *Electromagnetic splitting of charged and neutral mesons*, PoS **Lattice 2010** 127 (2011).

[24] T. Blum, *et al*., *Electromagnetic mass splittings of the low lying hadrons and quark masses from 2+1 flavor lattice QCD+QED*, Phys. Rev. **D82**, 094508 (2010) [arXiv:1006.1311[hep-lat]].

[25] M. Davier *et al*,, *Reevaluation of the Hadronic Contributions to the Muon g-2 and to alpha(MZ)*, Eur. Phys. J. **C71**, 1515 (2011) [arXiv:1010.4180 [hep-ph]].

[26] T. Blum, *Lattice calculation of the lowest order hadronic contribution to the muon anomalous magnetic moment*, Phys. Rev. Lett. **91**, 052001 (2003) [hep-lat/0212018].

[27] C. Aubin and T. Blum, *Calculating the hadronic vacuum polarization and leading hadronic contribution to the muon anomalous magnetic moment with improved staggered quarks*, Phys. Rev. D **75**, 114502 (2007) [hep-lat/0608011].

[28] P. Boyle, L. Del Debbio, E. Kerrane and J. Zanotti, *Lattice Determination of the Hadronic Contribution to the Muon* $g-2$ *using Dynamical Domain Wall Fermions*, arXiv:1107.1497 [hep-lat].

[29] M. Della Morte, B. Jager, A. Juttner and H. Wittig, *The leading hadronic vacuum polarisation on the lattice*, AIP Conf. Proc. **1343**, 337 (2011) [arXiv:1011.5793 [hep-lat]].

[30] X. Feng, K. Jansen, M. Petschlies and D. B. Renner, *Two-flavor QCD correction to lepton magnetic moments at leading-order in the electromagnetic coupling*, arXiv:1103.4818 [hep-lat].

[31] S. Weinberg, *Implications Of Dynamical Symmetry Breaking: An Addendum*, Phys. Rev. D **19**, 1277 (1979).

[32] L. Susskind, *Dynamics Of Spontaneous Symmetry Breaking In The Weinberg-Salam Theory*, Phys. Rev. D **20**, 2619 (1979).

[33] E. Eichten, K. D. Lane, *Dynamical Breaking of Weak Interaction Symmetries*, Phys. Lett. **B90**, 125-130 (1980).

[34] T. Banks and A. Zaks, *On The Phase Structure Of Vector-Like Gauge Theories With Massless Fermions*, Nucl. Phys. B **196**, 189 (1982).

[35] B. Holdom, *Techniodor*, Phys. Lett. **B150**, 301 (1985).

[36] K. Yamawaki, M. Bando, K. -i. Matumoto, *Scale Invariant Technicolor Model and a Technidilaton*, Phys. Rev. Lett. **56**, 1335 (1986).

[37] T. Appelquist, L. C. R. Wijewardhana, *Chiral Hierarchies from Slowly Running Couplings in Technicolor Theories*, Phys. Rev. **D36**, 568 (1987).

[38] D. D. Dietrich, F. Sannino, *Conformal window of SU(N) gauge theories with fermions in higher dimensional representations*, Phys. Rev. **D75**, 085018 (2007).

[39] S. Catterall, F. Sannino, *Minimal walking on the lattice*, Phys. Rev. **D76**, 034504 (2007).

[40] T. Appelquist, G. T. Fleming and E. T. Neil, *Lattice study of the conformal window in QCD-like theories*, Phys. Rev. Lett. **100**, 171607 (2008).

[41] T. Appelquist, A. Avarkian, R. Babich, R. C. Brower, M. Cheng, M. Clark, S. Cohen, G. Fleming, J. Kiskis, E. Neil, J. Osborn, C. Rebbi, D. Schaich and P. Vranas, *Toward TeV Conformality*, Phys. Rev. Lett. **104**, 071601 (2010).

[42] Z. Fodor, K. Holland, J. Kuti, D. Nogradi and C. Schroeder, *Nearly conformal gauge theories in finite volume*, Phys. Lett. B **681**, 353 (2009).

[43] Z. Fodor, K. Holland, J. Kuti, D. Nogradi, C. Schroeder, K. Holland, J. Kuti and D. Nogradi, *Twelve massless flavors and three colors below the conformal window*, arXiv:1103.5998 [hep-lat] and Phys. Lett. B **703**, 348 (2011).

[44] A. Hasenfratz, *Conformal or Walking? Monte Carlo renormalization group studies of SU(3) gauge models with fundamental fermions*, Phys. Rev. **D82**, 014506 (2010).

[45] X. -Y. Jin, R. D. Mawhinney, *Evidence for a First Order, Finite Temperature Phase Transition in 8 Flavor QCD*, PoS **LATTICE2010**, 055 (2010).

[46] T. DeGrand, Y. Shamir, B. Svetitsky, *Running coupling and mass anomalous dimension of SU(3) gauge theory with two flavors of symmetric-representation fermions*, Phys. Rev. **D82**, 054503 (2010).

[47] J. B. Kogut, D. K. Sinclair, *Thermodynamics of lattice QCD with 2 flavours of colour-sextet quarks: A model of walking/conformal Technicolor*, Phys. Rev. **D81**, 114507 (2010).

[48] T. Appelquist, R. Babich, R. C. Brower, M. Cheng, M. Clark, S. Cohen, G. Fleming, J. Kiskis, E. Neil, J. Osborn, C. Rebbi, D. Schaich and P. Vranas, *Parity Doubling and the S Parameter Below the Conformal Window*, Phys. Rev. Lett. **106**, 231601 (2011) [arXiv:1009.5967 [hep-ph]].

[49] A. Deuzeman, M. P. Lombardo, E. Pallante, *Evidence for a conformal phase in SU(N) gauge theories*, Phys. Rev. **D82**, 074503 (2010).

[50] E. Itou, T. Aoyama, M. Kurachi, C. -J. D. Lin, H. Matsufuru, H. Ohki, T. Onogi, E. Shintani *et al.*, *Search for the IR fixed point in the Twisted Polyakov Loop scheme (II)*, PoS **LATTICE2010**, 054 (2010).

[51] S. Catterall, J. Giedt, F. Sannino, J. Schneible, *Phase diagram of SU(2) with 2 flavors of dynamical adjoint quarks*, JHEP **0811**, 009 (2008).

[52] A. J. Hietanen, J. Rantaharju, K. Rummukainen, K. Tuominen, *Spectrum of SU(2) lattice gauge theory with two adjoint Dirac flavours*, JHEP **0905**, 025 (2009).

[53] A. J. Hietanen, K. Rummukainen, K. Tuominen, *Evolution of the coupling constant in SU(2) lattice gauge theory with two adjoint fermions*, Phys. Rev. **D80**, 094504 (2009).

[54] L. Del Debbio, B. Lucini, A. Patella, C. Pica, A. Rago, *The infrared dynamics of Minimal Walking Technicolor*, Phys. Rev. **D82**, 014510 (2010).

[55] F. Bursa, L. Del Debbio, L. Keegan, C. Pica, T. Pickup, *Mass anomalous dimension in SU(2) with six fundamental fermions*, Phys. Lett. **B696**, 374-379 (2011).

[56] L. Del Debbio, R. Zwicky, *Hyperscaling relations in mass-deformed conformal gauge theories*, Phys. Rev. **D82**, 014502 (2010).

[57] K. A. Intriligator, N. Seiberg and D. Shih, *Dynamical SUSY breaking in meta-stable vacua*, JHEP **0604**, 021 (2006) [arXiv:hep-th/0602239].

[58] J. Giedt, R. Brower, S. Catterall, G. T. Fleming and P. Vranas, *Lattice super-Yang-Mills using domain wall fermions in the chiral limit*, Phys. Rev. D **79**, 025015 (2009) [arXiv:0810.5746 [hep-lat]].

[59] J. Giedt, R. Brower, S. Catterall, G. T. Fleming and P. Vranas, *Domain wall fermion lattice super Yang Mills*, PoS **LATTICE2008**, 053 (2008).

[60] M. G. Endres, *Dynamical simulation of N=1 supersymmetric Yang-Mills theory with domain wall fermions*, Phys. Rev. D **79**, 094503 (2009) [arXiv:0902.4267 [hep-lat]].

[61] J. W. Elliott, J. Giedt and G. D. Moore, *Lattice four-dimensional N=4 SYM is practical*, Phys. Rev. D **78**, 081701 (2008) [arXiv:0806.0013 [hep-lat]].

[62] J. M. Maldacena, *The Large N limit of superconformal field theories and supergravity*, Adv. Theor. Math. Phys. **2**, 231 (1998) [Int. J. Theor. Phys. **38**, 1113 (1999)] [hep-th/9711200].

[63] S. Catterall, D. B. Kaplan and M. Ünsal, *Exact lattice supersymmetry*, Phys.Rept.484:71-130,2009

[64] Simon Catterall, *First results from simulations of supersymmetric lattices*, JHEP 0901:040,2009.

[65] S. Catterall and T. Wiseman, *Black hole thermodynamics from lattice Yang-Mills theory*, Phys.Rev.D78:041502,2008.

[66] S. Duane, A.D. Kennedy, B.J. Pendleton, and D. Roweth, *Hybrid Monte Carlo*, Phys. Lett. **B195**, 216 (1987).

[67] M.A. Clark, A.D. Kennedy and Z. Sroczynski, *Exact 2+1 flavour RHMC simulations*, Nucl. Phys. (Proc. Suppl.) **140**, 835 (2005); M.A. Clark, Ph. de Forcrand and A.D. Kennedy, *Algorithm Shootout: R versus RHMC*, PoS LAT2005, 115 (2005).

[68] I.P. Omelyan, I.M. Mryglod and R. Folk, *Symplectic analytically integrable decomposition algorithms: classification, derivation, and application to molecular dynamics, quantum and celestial mechanics simulations*, Comp. Phys. Comm., **151**, 272 (2003); T. Takaishi and P. de Forcrand, *Testing and tuning symplectic integrators for Hybrid Monte Carlo algorithm in lattice QCD*, Phys. Rev. E**73**, 036706 (2006).

[69] P. Boyle, *The Bagel web site*, *http:www2.ph.ed.ac.uk/ paboyle/bagel/Bagel.html*.

[70] A. Pochinsky, *Writing efficient QCD code made simpler: QA(0)*, PoS **LATTICE2008**, 040 (2008).

[71] M. G. Beckett, B. Joo, C. M. Maynard, D. Pleiter, O. Tatebe and T. Yoshie, *Building the International Lattice Data Grid*, Comput. Phys. Commun. **182** (2011) 1208 [arXiv:0910.1692 [hep-lat]].

[72] See International Data Lattice Grid (ILDG) website: *http://www.usqcd.org/ildg/*

[73] K. Barros, R. Babich, R. Brower, M. A. Clark and C. Rebbi, *Blasting through lattice calculations using CUDA*, PoS **LATTICE2008**, 045 (2008) [arXiv:0810.5365 [hep-lat]].

[74] M. A. Clark, R. Babich, K. Barros, R. C. Brower and C. Rebbi, *Solving Lattice QCD systems of equations using mixed precision solvers on GPUs*, Comput. Phys. Commun. **181**, 1517 (2010) [arXiv:0911.3191 [hep-lat]].

[75] S. Gottlieb, G. Shi, A. Torok and V. Kindratenko, *QUDA programming for staggered quarks*, PoS **LATTICE2010**, 026 (2010).

[76] R. Babich, M. A. Clark and B. Joo, *Parallelizing the QUDA Library for Multi-GPU Calculations in Lattice Quantum Chromodynamics*, arXiv:1011.0024 [hep-lat].

[77] R. Babich, M. A. Clark, B. Joó, G. Shi, R. C. Brower and S. Gottlieb, Proceedings of the 2011 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis (SC '11), 11 (2011), arXiv:1109.2935.

[78] R. C. Brower, H. Neff and K. Orginos, *Moebius fermions: Improved domain wall chiral fermions*, Nucl. Phys. Proc. Suppl. **140** (2005) 686 [arXiv:hep-lat/0409118].

[79] R. C. Brower, H. Neff and K. Orginos, *Moebius fermions*, Nucl. Phys. Proc. Suppl. **153**, 191 (2006) [arXiv:hep-lat/0511031].

[80] J. Brannick, R. C. Brower, M. A. Clark, J. C. Osborn and C. Rebbi, *Adaptive Multigrid Algorithm for Lattice QCD*, Phys. Rev. Lett. **100**, 041601 (2008) [arXiv:0707.4018 [hep-lat]].

[81] M. Luscher, *Solution of the Dirac equation in lattice QCD using a domain decomposition method*, Comput. Phys. Commun. **156**, 209 (2004) [arXiv:hep-lat/0310048].

[82] A. D. Kennedy, M. A. Clark and P. J. Silva, *Force Gradient Integrators*, PoS **LAT2009**, 021 (2009) [arXiv:0910.2950 [hep-lat]].

[83] *The Lua programming language website*, http://www.lua.org

[84] M. Lüscher, *Computational Strategies in Lattice QCD*, arXiv:1002.4232.

[85] M.A. Clark, Bálint Joó, A.D. Kennedy, and P.J. Silva, *Improving dynamical lattice QCD simulations through integrator tuning using Poisson brackets and a force-gradient integrator*, arXiv:1108.1828.

[86] *The QUDA (QCD in CUDA) library and githup repository*, https://github.com/lattice/quda

[87] Websites for *Lattice Meets Experiment* workshops for Flavor Physics at the Intensity Frontier: http://usqcd.fnal.gov/lattice-experiment2010.html and BSM Physics at the Energy Frontier: http://usqcd.fnal.gov/lattice-experiment2010

[88] Website for *QCD Numerical Analysis*workshop: http://super.bu.edu/ brower/qcdna6/

A.2 NVIDIA Letter of Intent

On the next page we reproduce a letter from Dr. David Luebke, Director of Research at NVIDIA Corporation, to Dr. Paul Mackenzie, the Chair of the USQCD Executive Committee, expressing his company's commitment to work with USQCD on the GPU codes discussed in this proposal. Dr. Luebke sets out the joint goals and milestones of NVIDIA and USQCD for this work, and the resources that NVIDIA will contribute to this effort.

December 6, 2012

Paul Mackenzie, Ph.D.
Theoretical Physics Department
Fermilab
MS 106
P.O. Box 500
Batavia, IL 60510 USA

Dear Paul,

On behalf of NVIDIA Corporation, I'm pleased to voice our enthusiastic support for the USQCD collaboration's SciDAC-3 proposals.

As you know, NVIDIA and the LQCD community have had a close but informal collaboration dating back as far as 2006 when the physics community first began researching the use of commodity GPUs for QCD. Since that time the community has created a rich set of open source software assets (e.g. the QUDA library[1]) and an effective software framework for the international community to build application suites that exploit GPU computing (including Chroma, MILC, and CPS).

Although the USQCD community has already achieved a great deal of groundbreaking success using GPUs, we have clearly only scratched the surface. Far greater opportunities to exploit the potential of GPUs lie ahead, with vast implications for research into the fundamental nature of matter. Furthermore, NVIDIA anticipates that the lessons learned in numerical algorithms, software infrastructure, and hardware design, will benefit not only the LQCD community, but the greater scientific community as well.

[1] http://lattice.github.com/quda

Goals

The collaboration between NVIDIA and USQCD will have a number of goals:

1. World-class sustained performance for gauge generation on the order of hundreds of TFLOP/s, scaling to the order of thousands of GPUs.
2. World-class efficiency for the analysis of gauge configurations on the order of tens of TFLOP/s, scaling to the order of hundreds of GPUs.
3. State-of-the-art multigrid-based Dirac inverters that run efficiently on single and multi GPU systems.
4. Complete implementation of GPU support in a domain specific language (e.g. QDP++)
5. Add multi-GPU support for additional fermion formulations, such as domain-wall fermions (DWF).
6. A roadmap for future GPU hardware and software that continues to address the goals of the QCD community.
7. A roadmap for a GPU programming model that continues to address the goals of the QCD community.

Proposed Milestones

Vast GPU installations are becoming increasingly common in NSF and DOE leadership class facilities, notably the Titan system at Oak Ridge National Laboratory, and more recently the Blue Waters system at the National Center for Supercomputing Applications (NCSA). The proposed collaboration between NVIDIA and USQCD should strive to be prepared to exploit these systems as soon as they become available.

NVIDIA and USQCD should target the following high-level milestones:

- 2012 – Refine domain-decomposition algorithms in QUDA & Chroma
- 2012 – Improve scaling through more efficient use of shared memory
- 2012 – Demonstrate Lattice ensemble generation (e.g. full Hybrid Monte Carlo) on GPUs in MILC
- 2013 – Demonstrate adaptive multigrid solver in Chroma (O(10)-fold performance improvement over current QUDA solvers)
- 2013 – Demonstrate adaptive multigrid combined with hybrid Monte Carlo

Resources

NVIDIA is extremely pleased to have the world's two foremost experts in using GPUs for QCD on staff, Dr. Mike Clark and Dr. Ron Babich.

Mike Clark will focus the majority of his time in the near term to supporting the USQCD community to develop software assets that exploit GPUs. He will also work closely with the community to disseminate skills and practical lessons learned so that the community can develop a larger pool of capable resources for this collaboration.

Ron Babich will be responsible for leveraging the lessons learned during this collaboration to guide the design of next-generation GPU architectures from NVIDIA.

Jonathan Cohen manages the Emerging Applications group, an applied research group in GPU computing. Jonathan will support Mike as they explore refinements to domain decomposition algorithms, adaptive multigrid solvers, and scalable hybrid Monte Carlo methods for gauge generation. There are opportunities for adding additional members to the Emerging Applications team, depending on needs, schedule, and suitable candidates.

Finally, Jerry Chen manages NVIDIA's strategic partnership with the QCD community. Jerry will provide program management and strategic alignment between this collaboration and NVIDIA's broader organization.

We look forward to the opportunity to extend our collaboration with the QCD community. Please do not hesitate to contact me with any questions.

Sincerely,

David P Luebke

David Luebke, Ph.D.
NVIDIA Distinguished Inventor
Director of Research
NVIDIA Corporation
dluebke@nvidia.com

A.3 IBM Letter of Intent

On the next page we reproduce a letter from Dr. George Chiu, Senior Manager, Advanced Server Systems, IBM, to Paul Mackenzie, the Chair of the USQCD Executive Committee, expressing his company's plans to continue to work with members of USQCD on the optimization of USQCD codes to run efficiently on the Blue Gene/Q, and on the development of new algorithms that will more effectively exploit this computer.

International Business Machines Corporation

Thomas J. Watson Research Center
P.O. Box 218
1101 Kitchawan Road
Yorktown Heights, NY 10598
gchiu@us.ibm.com
914-945-2436

December 19, 2011

Dr. Paul Mackenzie,
Chair, USQCD Executive Committee
Theoretical Physics Department, MS 106
Fermilab
P. O. Box 500
Batavia, IL 60510 USA

Dear Paul,

Lattice QCD (Quantum Chromodynamics) is an important target application for the Blue Gene Research Project at IBM and for future IBM high performance machines. This is both because of the exciting fundamental science discoveries and the challenges for software and hardware design that must be met to achieve high performance for lattice QCD.

The Blue Gene research group at the Watson Research Center plans to continue its fruitful, 10+ year collaboration with USQCD scientists and particularly the groups at Columbia and Edinburgh Universities and the Brookhaven National Laboratory. We expect to work closely together on many of the research topics addressed by this USQCD SciDAC proposal. This includes both the continued optimization of USQCD codes to run efficiently on the Blue Gene/Q architecture as well as the development of new algorithms which more effectively exploit the heterogeneous architectures built from powerful many-pipe floating point units, embedded in complex multi-thread, many-core nodes which are then replicated tens of thousands of times.

I and my colleagues at Watson look forward to working with those in USQCD on these critical challenges.

Yours truly,

George Chiu
IEEE Fellow,
Senior Manager, Advanced Server Systems, IBM

A.4 Intel Letter of Intent

On the next page we reproduce a letter from Mr. Joseph Curley, Director of Marketing, Technical Computing Group, Intel Corporation, expressing his company's intent to engage in "collaborative interactions" with USQCD on the development of USQCD software for the Intel MIC architecture. This effort will be lead by Dr. Chip Watson of Jefferson Laboratory and Dr. Donald Holmgren of Fermilab. Although not stated in the latter, Fermilab has executed a Restricted Secret information Non-Disclosure Agreement (RSNDA) with Intel regarding the MIC architecture, and expects to receive a MIC accelerator card later this year. Jefferson Laboratory has already received such a card.

Intel Corporation
2200 Mission College Blvd.
Santa Clara, CA 95054

Chip Watson
Group Lead and Deputy CIO, High Performance Computing
Jefferson Lab
12000 Jefferson Avenue
Newport News, VA 23606

January 2, 2011

Dear Chip:

Intel Corporation is pleased to support JLAB and the USQCD grant proposals. We recognize that optimizing and scaling Lattice QCD on new technologies, tools and architectures as we move towards Exascale computing is a significant challenge for the USQCD collaboration team. Intel is excited to be considered such a strategic partner as you tackle Lattice QCD technical challenges through the USQCD collaboration team.

As a strategic investment for Intel, we have accepted JLAB into Intel's early software development program for Intel Many Integrated Core (Intel MIC) Architecture.

We look forward to the opportunity to proceed on collaborative interactions.

Sincerely,

Joseph Curley,
Director of Marketing, Technical Computing Group
Intel Corporation

A.5 Cover Pages and Budgets

In the following pages we reproduce the cover pages and budgets of each participating institution.

A.6 Statements of Work

In the following pages we reproduce the detailed statements of work and milestones of each of the participating institutions.

Argonne National Laboratory Statement of Work

1 Introduction

The multi-institutional USQCD collaboration is submitting a joint proposal to the SciDAC-3 computational high energy physics program. This narrative describes the Argonne National Laboratory contribution to the collaborative effort. Please refer to the full proposal (attached in the appendix) to see how this work integrates with the multi-institutional collaborative effort.

Argonne National Laboratory will be a main contributor to the effort to develop a new framework to more easily integrate code for simulations of theories beyond QCD. In addition we will contribute to the efforts on algorithm development and in optimizations for the IBM Blue Gene/Q.

An important part of this collaborative proposal is to develop codes for the large scale numerical studies of theories of Beyond Standard Model Physics. This includes theories of composite Higgs, for example technicolor, and supersymmetric lattice field theories. Argonne National Laboratory will focus primarily on the composite Higgs work and will involve coordination with the work on lattice supersymmetry lead by Prof. Catterall at Syracuse, as well as collaborations with other members of USQCD.

We also plan to contribute to the efforts to develop better algorithms to perform the large simulations required for Beyond Standard Model physics as well as for QCD. This will include collaborating with the FASTMath institute on using HYPRE to develop and implement multigrid algorithms for solving the Dirac equation, as well as exploring better methods for the HMC integration that are severly needed as one explores theories with a large number (Nf) of light fermions.

The Leadership Computing Facility at Argonne National Laboratory is planning to acquire a large IBM Blue Gene/Q system in 2012. This puts researchers at ANL in a very good position to optimize codes for this platform. We have already started testing LQCD codes on early systems. We plan to continue these efforts and work with other members of USQCD to help optimize a larger set of USQCD codes for BG/Q.

Dr. Osborn has been involved with the USQCD software effort for 10 years, and has contributed to a large number of the LQCD libraries developed under the previous SciDAC grants. He has been a lead developer for the QMP, QLA, QDP/C and QOPQDP libraries and has done extensive work optimizing these libraries for x86 clusters and the IBM Blue Gene platforms. He has also contributed improvements to the QIO library and helped integrate the optimized routines available in QOPQDP into the MILC code. With collaborators at BU and elsewhere he developed an efficient multigrid algorithm for the Wilson-clover Dirac solver and implemented it on top of the QDP/C library.

2 Deliverables: Argonne National Laboratory

- Year 1:
  - Get basic implementation of HMC for SU(2) and SU(3) fundamental HISQ, Wilson and domain wall fermions working in FUEL
  - With Syracuse, work on getting HMC for adjoint doman wall fermions running in FUEL
  - With others, work on optimizing QMP and QLA for BG/Q
  - Investigate HMC integrator improvements for large Nf
  - Integrate existing GPU routines into FUEL
  - Collaborate with FASTMath and others to get current Wilson-clover multigrid algorithm implemented in HYPRE

- Year 2:
  - Investigate Hermitian eigensolver routines for HISQ and domain wall quarks
  - Work on adding 2-index symmetric routines to QOPQDP for use in FUEL
  - Add nHYP and Stout variations of staggered fermions to QOPQDP
  - Investigate force-gradient integrators in FUEL

- Year 3:
  - Add Schrödinger functional support to FUEL
  - Develop SU($Nc > 3$) code for Wilson, domain wall and staggered fermions
  - Optimize SU(2) and SU($Nc > 3$) codes on BG/Q and clusters
  - Collaborate with FASTMath and others on multigrid for staggered and domain wall fermions, and integrate working multigrid codes into FUEL

Arizona, Indiana, Utah Statement of Work

1 Introduction

The multi-institutional USQCD collaboration is submitting a joint proposal to the SciDAC-3 computational high energy physics program. This narrative describes the Arizona, Indiana, and Utah contributions to the collaborative effort. The submitted budget covers only the University of Arizona. Please refer to the full proposal to see how this work integrates with the multi-institutional collaborative effort.

The University of Arizona, Indiana University, and the University of Utah are members of the MILC collaboration, which, in turn, is part of USQCD. The MILC collaboration[1], consisting of approximately eight senior and approximately nine junior members at ten institutions, mostly in the USA,[1] has a long tradition of advancing the state of the art of the numerical simulation of the strong interactions of quarks and gluons using methods of lattice quantum chromodynamics (lattice QCD). (For a recent summary article, see [2].)

We feel we have a strong record of pioneering scientific achievements, made possible through a versatile collaboration code suite. The MILC code, now consisting of approximately 200,000 lines, is constantly evolving in response to new physics challenges, algorithms, and architectures [3]. To achieve maximum efficiency, the code incorporates SciDAC modules developed by us and our USQCD colleagues. This code supports not only our work, but the work of several research groups around the world [1]. We estimate that it is currently responsible for calculations at the rate of over one hundred million core-hours per year at NSF and DOE centers and laboratories in the US and other centers abroad.

Collaboration members, especially at the University of Arizona, Indiana University, the University of Utah, and the American Physical Society, develop and contribute new algorithms and code modules. The integration of these modules and distribution and maintenance of the MILC code is coordinated at the University of Utah [3]. Testing in a production environment is carried out at several collaborating institutions.

The central objective of this project is to advance scientific discovery to a new level of capability by supporting the continued rapid evolution of the community SciDAC code suite for GPU and multicore operation. We will assist in developing CUDA GPU modules of general use to the entire US lattice QCD community, and we will integrate them into the MILC code. Through our active participation in the USQCD SciDAC software committee, we will coordinate with similar efforts on related codes by our colleagues at Argonne National Laboratory, Boston University, Columbia University, Fermilab, Jefferson Laboratory and the Emerging Technology group at NVIDIA. Through our computational science colleagues in the SciDAC institutes, particularly SUPER, we will use state-of-the art performance measures and adaptive optimization. We are confident that such coordination will benefit the entire US lattice gauge theory community.

Three institutional members of the MILC collaboration, namely the University of Arizona, Indiana University, and the University of Utah, are actively involved in the USQCD SciDAC effort and are submitting proposals to the SciDAC-3 high energy physics programs. This narrative details work to be performed at these three institutions. The lead institution is indicated in parenthesis for each item.

[1]A. Bazavov, R.S. Van de Water (Brookhaven), C. Bernard, M. Lightman (Washington U.), C. DeTar, J. Foley, L. Levkova, M. Oktay (U. Utah), J. Kim, D. Toussaint (U. Arizona), S. Gottlieb, R. Zhou (Indiana U.), U.M. Heller (APS), J.E. Hetrick (U. Pacific), J. Laiho (Glasgow U.), J. Osborn (Argonne), R.L. Sugar (UC, Santa Barbara)

[1]Although we do not have a formal registration system, we are aware of users at the following institutions: Boston University, Brookhaven National Laboratory, University of California, San Diego, University of Colorado, University of Kentucky, University of Granada (Spain), University of Gronigen (Netherlands), Ohio State University, Tel Aviv University (Israel), Zhongshan University (China), and, of course, by our collaborating colleagues at Fermilab.

2 Statement of Work: Arizona, Indiana, Utah

Our MILC collaboration code base provides the foundation for the research programs of several groups within USQCD, as well as a number of groups abroad. As our physics goals, algorithms, and machine architectures evolve, the code also evolves. The arrival of radically new multicore architectures requires some major changes to the code. These changes will be built on the general-purpose public-domain SciDAC utilities that we and others in this proposal are developing.

2.1 QUDA modules

Since 2009, we have been working closely with Guochun Shi to develop GPU code for staggered quarks to enhance the MILC code. This work is very timely, since the forthcoming NCSA Blue Waters and the ORNL Titan machines will feature thousands of GPU-enhanced nodes. In the initial stage of this work, an asqtad solver running on a single GPU was developed [4, 5]. Fattening, gauge and fermion force routines for single GPU were developed for the asqtad action [5]. The solver was parallelized by dividing the lattice in the time dimension [6]. Then, the solver was generalized to divide the lattice in all four dimension and scaling to over 100 GPUs on large grids was seen to meet or exceed typical supercomputer speeds [7].

Deliverables, year 1: (Indiana leads) With Illinois and Boston University, multi-GPU support for the four main algorithms needed for molecular dynamics including HISQ extensions. (Utah leads) Integration of the GPU algorithms into the MILC code suite. (Arizona leads) Testing and validation of the modules in a production environment.

Deliverables, year 2: (Indiana) With Illinois and NVIDIA, optimization of codes developed in year 1 and elimination of the principal bottlenecks. (Utah) Integration. (Arizona) Testing and validation.

Deliverables, year 3: (Indiana) Development of QUDA versions of successful new algorithms listed in 2.6 below. (Utah) Integration. (Arizona) Testing and validation.

2.2 Multithreading

In the coming years we propose to coordinate with Argonne National Laboratory in developing multithreading capabilities within the framework of the SciDAC code suite. Currently QLA supports multithreading, but we anticipate that we will need to keep threads active for a longer stretches of the computation. This would entail threading at the QOP level across QDP calls and within the `QDP_loop` constructs. A suitable protocol and thread support utilities need to be developed. Specifically for the Blue Gene Q, we will work with Argonne National Laboratory in selected hand-optimization using Peter Boyle's Bagel tool and the PETSc TAO performance monitoring toolkit [8]. These optimizations will be integrated into the MILC code.

Deliverables, year 1: (Arizona) With Argonne National Lab, MILC code support for multi-threaded QOP. Published analysis of performance and identification of remaining bottlenecks.

Deliverables, year 2: (Arizona) Elimination of the principal bottlenecks identified in year 1. Published analysis of performance improvement.

Deliverables, year 3: (Arizona) Creation of multithreaded versions of successful new algorithms listed in 2.6 below.

2.3 BlueGene SPI

To further improve Blue Gene/Q performance we propose in collaboration with Argonne National Laboratory to develop a version of the QMP (message passing) package replacing MPI support with IBM SPI (System Programming Interface) support.

Deliverables, year 2: (Indiana) With Argonne National Laboratory, testing of the SPI implementation of QMP (message passing). Published analysis of MPI *vs* SPI performance.

2.4 A QOP API for QUDA support

Currently our QOP (Level 3 optimized) routines are built on QDP/C for calculations on conventional clusters. For ease of programming for GPU-based clusters it would be very useful to have a version of QOP routines for HISQ and Wilson clover fermions built, instead, on QUDA. We propose to create one.

Deliverables, year 1: (Utah) A QOP QUDA that supports QOP modules required for HISQ and Wilson clover fermions and HISQ molecular dynamics.

2.5 MPI-I/O version of QIO

Now that MPI-I/O is widely available, it makes sense to build a parallel I/O API based on MPI-I/O, rather than QMP/MPI. The "Lemon" utility developed at Bonn University[9] supports our current binary record formats through MPI-I/O. What remains is to retool the QIO interface so it operates through the Lemon interface, but supports all of our current logical record and file formats.

Deliverables, year 1: (Utah) Trial version of QIO that supports MPI-I/O

2.6 Algorithmic Improvements

A variety of algorithmic developments will benefit the MILC code and the SciDAC code suite: (1) Mixed precision inversion is currently controlled outside the QOP code suite. Moving it into the QOP package makes good sense. (2) The force-gradient integrator improves molecular dynamics evolution. It has not yet been implemented for HISQ fermions. (3) Multigrid methods have been applied successfully to Wilson fermions, but not yet to HISQ fermions. We propose to explore these improvements.

Deliverables, year 3: (Arizona) Formulation of new algorithmic improvements. (Utah) Integration into the MILC code and/or SciDAC code suite of successful algorithmic improvements.

2.7 LUA

LUA is a light-weight scripting language that could greatly accelerate the coding and development of new physics projects. It is the basis for the proposed Argonne National Laboratory Framework for Unified Evolution of Lattices (FUEL) system. Currently, to specify the calculation of a three-point function in the MILC code we write a parameter input file of a few hundred lines (the user interface). This file is read and parsed and the calculations then proceed according to a rigid pattern. With LUA the parameter input file would be replaced by a LUA script that carries out the calculation with a series of procedure calls. Examples are procedures for constructing sources, doing sparse solves (inverters), and tying together propagators. The underlying operations would be mostly our existing MILC procedures, but the management would be done by the C++-like LUA scripting language. We will assist programming experts at Fermilab in the development of the scriptable modules.

Deliverables, year 2: (Utah) With Fermilab, test of a trial LUA version of the MILC code.

Deliverables, year 3: (Utah) With Fermilab, expansion of LUA support for the MILC code, should it prove successful.

2.8 MILC code user interface

We will also continue development of the user interface to the MILC code to increase its flexibility and to mimimize the chances of user errors causing wrong results. The version of the code that is in the final testing stage now implements a powerful and flexible input format for specifying the correlators to be calculated. However, further work will be needed to accomodate new projects and to allow user control of the new low-level features discussed above.

2.9 Miscellaneous

MILC is engaged in a variety of other projects: (1) We have worked with Torsten Hoefler of NCSA to implement derived datatypes in the MILC code with some success. The goal is to require fewer copies of data into message buffers, which will shorten communication times. (2) We are developing an improved heavy quark algorithm, the “OK action”, for Wilson quarks. Continued testing and, if successful, optimization will be required. (3) We have started a new project with the Joint Laboratory for Petascale Computing to try pre-conditioners for the staggered quark inverter. This work will continue in the coming year.

[1] MILC Collaboration Home Page: http://www.physics.indiana.edu/~sg/milc/.

[2] A. Bazavov, C. Bernard, C. DeTar, S. Gottlieb, U.M. Heller, J.E. Hetrick, J. Laiho, L. Levkova, P.B. Mackenzie, M.B. Oktay, R. Sugar, D. Toussaint, and R.S. Van de Water. Nonperturbative QCD simulations with 2+1 flavors of improved staggered quarks. *Reviews of Modern Physics*, 82:1349–1417, 2010.

[3] `http://www.physics.utah.edu/~detar/milc/milc_qcd.html`.

[4] Guochun Shi, Steven Gottlieb, Aaron Torok, and Volodymyr V. Kindratenko. Accelerating Quantum Chromodynamics Calculations with GPUs. In *Proc. Symposium on Application Accelerators in HPC (SAAHPC'10)*. IEEE, 2010.

[5] Steven Gottlieb, Guochun Shi, Aaron Torok, and Volodymyr Kindratenko. QUDA programming for staggered quarks. *PoS*, LATTICE2010:026, 2010.

[6] Guochun Shi, Steven Gottlieb, Aaron Torok, and Volodymyr V. Kindratenko. Design of MILC lattice QCD application for GPU clusters. In *IPDPS*. IEEE, 2011.

[7] Ronald Babich, Michael A. Clark, Bálint Joó, Guochun Shi, Richard C. Brower, and Steven Gottlieb. Scaling Lattice QCD beyond 100 GPUs. In *Proceedings of the 2011 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '11. IEEE Computer Society, 2011.

[8] `http://www.mcs.anl.gov/tao/documentation/installation.html`.

[9] Albert Deuzeman, Siebren Reker, Carsten Urbach, and for the ETM Collaboration. Lemon: an MPI parallel I/O library for data encapsulation using LIME. 2011.

Boston University Project

1 Introduction

The multi-institutional USQCD collaboration is submitting a joint proposal to the SciDAC-3 computational high energy physics. This narrative describes the Boston University contribution to the collaborative effort. Please refer to the full proposal to see how this work integrates with the multi-institutional collaborative effort (see additional attached documents)

Under SciDAC-1 and SciDAC-2, Richard Brower served as the Chair for the Software Coordinating Committee and as a member of the USQCD Executive Committee. In this proposed project, he will continue in this capacity as co-Director for Computation, with overall responsibility for the software and algorithm work, providing direction and coherence to the work, and monitoring progress on all tasks. He will also coordinate work with other collaborators in USQCD, and our partners in the SciDAC FASTMath Institute and in the Emerging Applications Group at NVIDIA. He will provide quarterly reports to the Director on the progress of the software effort. The Software Committee has weekly teleconference to plan and track the projects and annual face to face workshops. These USQCD software workshops will now include partners in FASTMath, NVIDIA, as well as other collaborators in applied mathematics and computer science.

The two main software projects at Boston University under SciDAC-3 will be to further develop multi-scale algorithms for Dirac inverters [1, 2] and the high performance QUDA (QCD in CUDA) library for GPUs [3, 4, 5]. Although these projects had their genesis at Boston University under SciDAC-2, they have now expanded into much larger projects so the role of Boston University will be not only to contribute to them but increasingly to help in co-ordinating developments at other collaborating institutions. Here we summarize the research program at Boston University and list the *Tasks and Milestones* proposed during the 3 years of this proposal.

Beyond the Standard Model (BSM) at the Energy Frontier The driver for all our algorithmic and software projects is to accelerate scientific discoveries through the use of lattice field theory and high performance computing. Boston University has a long history of accomplishments in Lattice Field Theory lead by two senior faculty Prof. Richard Brower and Prof. Claudio Rebbi. With the participation of a series of exceptionally talented postdoctoral fellows supported jointly by SciDAC and NSF funding, they have developed new algorithms and software strategies. For example the Chronological Inverter [6] for evolutions code, the Möbius Domain Wall algorithm [7, 8, 9] with former SciDAC postdoc Hartmut Neff and more recently multigrid with former SciDAC postdocs James Osborn, Mike Clark, Ron Babich and Saul Cohen [10] . To collaborate with the wider applied mathematics community there are semi-annual QCD Numerical Analysis workshops the most recent one at Boston University Sept, 2010 (*QCDNA VI*).

In the last few years the research focus at Boston University has made a transition from lattice QCD to the study of new strongly interacting gauge theories for BSM studies at the TeV energy scale. Nearly five years ago Appelquist, Brower, Fleming, Osborn, Rebbi and Vranas have formed the Lattice Strong Dynamics collaboration (`http://www.yale.edu/LSD`) aimed at exploring non-perturbative scenarios beyond QCD, which may well be part of the new physics discovered at the LHC. A large range of options are described in an early white paper [11] in 2007 and an initial workshop on *Lattice Gauge Theory for LHC Physics* was held at Livermore May 2-3, 2008, followed by a second workshop at Boston University Nov. 6-7, 2009 and the most recent *Lattice Meets Experiment 2010: Beyond the Standard Model* at FNAL on October 14-15, 2011. At present since there is much less experience and no experimental guidance for strongly coupled theories beyond the Standard Model, the risks are greater. To mitigate this risk and minimize software development time, initial projects were chosen in areas close to QCD itself [12]. Boston University's focus has been on the S-parameter [13], which places one of the most stringent constraints on technicolor models and problem of disconnected diagrams, specifically computing the $\bar{s}s$ condensate in the proton needed to estimate the cross section for the direct detection of SUSY neutralino as a possible candidate for dark matter. Both of these project as are limited by the cost of Dirac solvers. To reduce this cost a strong program to develop multigrid methods and cost effective GPU has been pursued over the past few years. Leading the QUDA library project (QCD in CUDA for GPU) were our SciDAC postdoctoral fellows, Mike Clark and Ron Babich, both of who are now employed at NVIDIA to continue this collaborative work with USQCD. Also in collaboration with Argonne Laboratory and Syracuse, Boston University will help to develop the software infrastructure to expand the search for new BSM strongly interacting physics. A key tool will be Framework for Unified Evolution of Lattices (FUEL), a tool that will enable the rapid development and testing of new algorithms for configuration generation based on top level control using the scripting language Lua [14] ,

Multigrid Research After more than 20 years of effort, starting with the early work of Brower, Rebbi, and others, the first success in applying advanced multigrid methods to lattice field theory has been achieved [1, 2]. This breakthrough resulted from combining new insights from applied mathematics and lattice QCD. After nearly four years of effort at Boston University in collaboration with Mike Clark and Ron Babich (both now at Nvidia) and applied mathematicians James Brannick (Penn State) Steve McCormick (Colorado University) and others in TOPS, the first successful multi-grid lattice Dirac inverter was developed. James Osborn after his SciDAC postdoc at BU, now at Argonne National Laboratory, implemented an extension to the QDP API to accommodate multiple lattices, providing a Level 3 multigrid inverter for the Wilson-clover operator, which is demonstrating nearly 25x speed up as one approaches the physical light quark limit in production code on the BlueGene/L. In spite of this significant first step, it is not the end but the beginning of a larger opportunity. Our first multigrid algorithmic advance only addresses one critical component of a single discretization of the Dirac PDE. Further development of multigrid and domain decomposition is proposed under this project.

QUDA: GPU software development In the summer of 2008, Rebbi and Brower enlisted a graduate student in statistical physics, Kipton Barros, to explore the GPU architecture for lattice field theory. In collaboration with our postdoctoral fellows, Mike Clark and Ron Babich, he obtained a performance in excess of 100 Gigaflops on a single 240 core Nvidia GTX280 GPU. This was the beginning of the USQCD software effort at Boston University by Clark and Babich to produce a new SciDAC library called QUDA, targeted at multi-GPU computing. Using multiple precision solvers and a variety of tricks to reduce bandwidth to the device memory on the GPU card, they have obtained an additional factor of more than two in performance for the Wilson-clover inverter. Under SciDAC-2 collaboration has developed in QUDA a full set of inverters for Wilson-clover, Domain Wall, Staggered and Twisted mass discretizations.

Statement of Work: Boston University: The main goal is to continue to explore and develop multi-scale solvers and high performance GPU code with the aim to bring these solvers into the main stream of Lattice Field Theory and to adapt them to the heterogeneous environment in the roadmap to extreme scale computing architectures

During the three years of this proposal these developments require a sequence of steps: (i) generalizing the Wilson-clover multigrid algorithm to Domain Wall and staggered quarks, (ii) implementing these inverters on multi-GPU clusters with appropriate domain decomposition methods and (iii) integrating multi-level solvers (multigrid and domain decomposition) into HMC evolutions codes. Boston University will collaborate on the following tasks: The SciDAC postdoctoral fellow at Boston University will work with another postdoctoral fellow supported by the National Science Foundation under the direction of Brower and Professor Claudio Rebbi, the director of the Boston University Center for Computational Science.

In the first year under SciDAC-3, the Boston University task will be to improve the interface to the Wilson-clover multigrid inverter and collaborate with JLab and ANL to integrate it into the Chroma application suite for wider use. In collaboration with Nvidia, the multigrid inverter will be implemented in CUDA and added to the QUDA software library. The combined reduction in the cost of Dirac solves for a GPU multi-grid solver is projected to reduce the cost of production for analysis by more than a factor of 100.

The next priority is continue to research on multi-scale algorithms for other discretizations of the Dirac operator. This activity is proposed as part of a partnership with FASTMath. A major goal of the FASTMath partnership is to leverage the HYPRE framework to accelerate the design and full scale testing cycle which has proven to be a bottle neck for rapid development of new algorithms under SciDAC-2. The first FASTMath task is to reimplement and further optimize the existing Wilson MG algorithm in a enhance HYPRE framework.

Then the priority in the second year will be take the initial formulation of Domain Wall multigrid algorithm, started last year at Boston University by Saul Cohen (now at Seattle) and collaborate with FASTMath using the HYPRE framework to search for an optimal Domain Wall multigrid inverter.

A final goal at Boston University is to collaborate in developing new evolution codes for Beyond the Standard Model physics, which again will leverage the infrastructure of the FUEL project at Argonne as part of the FASTMath partnership. The first goal is to tune multi-flavor QCD evolution under FUEL and then port the algorithm for high performance for the BG/Q and multi-GPU clusters. This project is already being tested for the QCD staggered action in collaboration with Argonne, Nvidia and MILC. The interest at Boston University in the second and third year is to be able to extend this methodology to new representations for BSM model studies.

We also anticipate testing multi-scale inverters in the context of evolution codes in FUEL and Chroma to understand how best to meet the dual but conflicting requirement of improved algorithmic scaling at small fermion mass and limited communication of increasing importance for heterogeneous architecture, typified but not exclusively by CRAY/Titan and CRAY/BlueWaters GPU enabled machines to be deployed at Oak Ridge and NCSA respectively. The sequence of Tasks and Deliverables are as follows.

- *Year 1:*
  – complete multiple–GPU implementation of Domain Wall solver
  – develop level 3 code and documentation for the multigrid solver for Wilson-Clover fermions
  – work with NVIDIA to optimize this multigrid solver on multiple–GPU clusters
  – further optimize the Wilson-Clover multigrid in the HYPER framework with FASTMath partners
  – develop first version of level 3 multigrid solver for Domain Wall fermions.

- *Year 2:*
  – develop multigrid solver for Overlap fermions
  – work with NVIDIA to develop multi-precision multigrid solvers for Domain Wall fermions on multiple-GPU clusters
  – study communication mitigation of hybrid multigrid/domain decomposition solvers with FASTMath collaborators.

- *Year 3:*
  – introduce multigrid/domain decomposition solvers for Blue Gene/Q and cluster architectures
  – test & optimize multigrid/domain decomposition evolution in the FUEL framework in collaboration with ANL
  – work with NVIDIA to optimized multigrid/domain decomposition code for GPU systems (Titan, BlueWaters, USQCD Clusters).

Bibliography and References Cited: Boston University

[1] J. Brannick, R. C. Brower, M. A. Clark, J. C. Osborn and C. Rebbi, "Adaptive Multigrid Algorithm for Lattice QCD," Phys. Rev. Lett. **100**, 041601 (2008) [arXiv:0707.4018 [hep-lat]].

[2] R. Babich, J. Brannick, R. C. Brower, M. A. Clark, T. A. Manteuffel, S. F. McCormick J. C. Osborn,and C. Rebbi, "Adaptive multigrid algorithm for the lattice Wilson-Dirac operator," Phys. Rev. Lett. **105**, 201602 (2010) [arXiv:1005.3043 [hep-lat]].

[3] K. Barros, R. Babich, R. Brower, M. A. Clark and C. Rebbi, "Blasting through lattice calculations using CUDA," PoS **LATTICE2008**, 045 (2008) [arXiv:0810.5365 [hep-lat]].

[4] M. A. Clark, R. Babich, K. Barros, R. C. Brower and C. Rebbi, "Solving Lattice QCD systems of equations using mixed precision solvers on GPUs," Comput. Phys. Commun. **181**, 1517 (2010) [arXiv:0911.3191 [hep-lat]].

[5] R. Babich, M. A. Clark, B. Joo, G. Shi, R. C. Brower and S. Gottlieb, "Scaling Lattice QCD beyond 100 GPUs," arXiv:1109.2935 [hep-lat].

[6] R. C. Brower, T. Ivanenko, A. R. Levi and K. N. Orginos, "Chronological inversion method for the Dirac matrix in hybrid Monte Carlo," Nucl. Phys. B **484**, 353 (1997) [arXiv:hep-lat/9509012].

[7] R. C. Brower, H. Neff and K. Orginos, Nucl. Phys. Proc. Suppl. **140** (2005) 686 [arXiv:hep-lat/0409118].

[8] R. C. Brower, H. Neff and K. Orginos, Nucl. Phys. Proc. Suppl. **153**, 191 (2006) [arXiv:hep-lat/0511031].

[9] R. Brower, R. Babich, K. Orginos, C. Rebbi, D. Schaich and P. Vranas, PoS **LATTICE2008**, 034 (2008) [arXiv:0906.2813 [hep-lat]].

[10] R. Babich, J. Brannick, R. C. Brower, M. A. Clark, S. D. Cohen, J. C. Osborn and C. Rebbi, "The role of multigrid algorithms for LQCD," PoS **LAT2009**, 031 (2009)

[11] See whitepager at USQCD website: *http://www.usqcd.org/collaboration.html#2007*

[12] T. Appelquist, A. Avarkian, R. Babich, R. C. Brower, M. Cheng, M. Clark, S. Cohen, G. Fleming, J. Kiskis, E. Neil, J. Osborn, C. Rebbi, D. Schaich and P. Vranas, Phys. Rev. Lett. **104**, 071601 (2010).

[13] T. Appelquist, R. Babich, R. C. Brower, M. Cheng, M. Clark, S. Cohen, G. Fleming, J. Kiskis, E. Neil, J. Osborn, C. Rebbi, D. Schaich and P. Vranas, Phys. Rev. Lett. **106**, 231601 (2011) [arXiv:1009.5967 [hep-ph]].

[14] See the programming language Lua website: *http://www.lua.org*

2 Narrative: Scientific Goals and Software Tasks – BNL

To achieve the goals of SciDAC Projects, the High Energy Physics part of Brookhaven National Laboratory (BNL-HEP) conducts various research programs to construct indispensable bridges between high energy experimental data and particle theory. Owing to the dramatic increases in computing resources and theoretical advances over the past several years, BNL-HEP has opportunities to carry out reliable lattice determinations of hadronic weak matrix elements. As a result, significantly better precision has been attained in the determination of fundamental parameters of the Standard Model, in particular, the quark flavor mixing matrix of Cabibbo-Kobayashi-Maskawa (CKM) theory [1, 2, 3, 4], from which, one could test of the Standard Model in the flavored sector at the few-percent level and revealed an approximately $\sim 3\sigma$ tension in the CKM unitarity triangle that may indicate the presence of new physics.

To carry out pion, kaon [5, 6], or B- and D-meson physics [7, 8, 9, 10], dynamics up, down, and strange quarks are computed with lattice quark with chiral symmetry, called Domain-Wall Fermion (DWF). Although it has been well known that the chiral symmetry is useful for various physics applications, dynamical QCD simulation with DWF used to be computationally challenging due to its fictitious fifth dimension that separates the right- and left- chiral modes of quarks. Many breakthroughs both in computations and techniques were made through previous SciDAC 1 and 2 Projects, now DWF QCD is producing many of phenomenologically relevant and precise results. The DWF QCD vacuum configurations, which lie at the core of these computations, have been generated under the SciDAC programs, jointly with RIKEN-BNL-Columbia (RBC) collaboration for two dynamical flavors [11], and RBC and UKQCD collaborations for $N_F = 2+1$ flavors [12]. There are also theoretical developments for the precise renormalization of quantum operators relevant in determinations of quark masses and the electroweak matrix elements using non-perturbative techniques [13, 14, 1, 15].

One new development towards the high precision computation in particle physics is the inclusion of the asymmetry between up and down quarks. Since these two flavors of quarks are similar to each other (isospin symmetry), many lattice QCD computations neglect the small mass difference, $m_{\rm up} \neq m_{\rm down} \sim O(1)$ MeV, and the electric charge difference, $q_{\rm up} = +2/3e, q_{\rm down} = -1/3e$. As the precision of the lattice results are improved lately, we have now an opportunity to investigate accurately effects from isospin symmetry breaking in the hadronic observables. These investigation will allow us to answer important questions such as why the proton is stable (lighter than Neutron), or, what are the precise values of each individual up and down quark [16, 17, 18, 19].

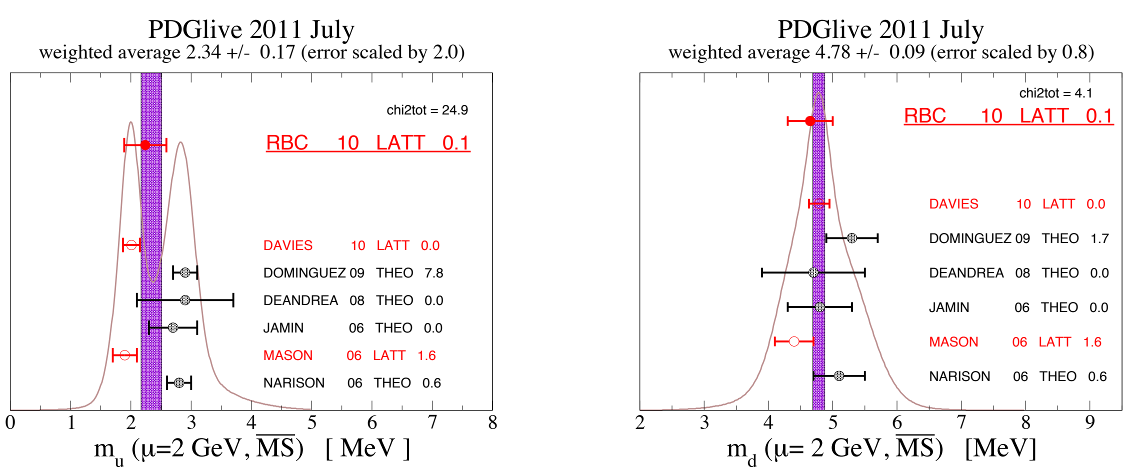

Figure 1: Lattice QED+QCD simulation results: up quark mass (left) and down quark mass

Under SciDAC-3, the main contribution of the team at BNL-HEP to this SciDAC proposal will be in the field of 'Flavor Physics at the Intensity Frontier'. This contribution will be twofold. First, we will improve the errors on quantities for which results with fully-controlled errors exist, but for which the errors are still larger than or comparable to those from other sources. And, the second, we will expand our program of calculations to meet the needs of upcoming intensity-frontier experiments, for example the muon $g2$ experiment at Fermilab, the Project X kaon program at Fermilab, LHC-b, Belle II and Super-B.

Quantities in this first category would be heavy-light decay constants, neutral flavored mesons mixing parameters for K, B_d, B_s and the SU(3) breaking ratio ξ, which leads to yet more precise determinations of the CKM matrix elements and may enable us to confirm or refute the current hints of new physics in the flavor sector. Other likely topics for B- and D-meson electroweak matrix elements pursued at BNL-HEP would include The $B \to \pi l \nu$ and $B \to D^* l \nu$ semileptonic form factors, Beyond-the-Standard Model operator contributions to $B - \overline{B}$ mixing, and short-distance contribution to $D - \overline{D}$ mixing using both Relativistic Heavy Quark action and Static Quark action.

There will be a major efforts towards understanding the Kaon physics, especially for $K \to \pi\pi$ decay, aiming towards understanding the $\Delta I = 1/2$ rule and the direct/indirect CP violation of Kaon system, ε'/ε.

Studies over more than ten years in collaboration with RIKEN, Columbia University, and UKQCD, have led to the decision to compute $K \to \pi\pi$ using the direct method of Lellouch and Luscher [20] exploiting finite-volume effects. (The previously used indirect method of computing Kaon to vacuum and Kaon to single pion amplitudes [21] and then using the leading order chiral perturbation theory was found to have large systematic errors.) The RBC-UKQCD collaboration is now performing a realistic 2+1 flavor computation of the $\Delta I = 3/2$ amplitude on a coarse lattice with close to physical pion masses to obtain both the real and imaginary parts of the $I = 2$ decay amplitude A_2. The electroweak (EW) operators (including the EW penguin operators Q_7 and Q_8) will be matched to the continuum MS bar schemes using the non-perturbative renormalization RI-(S)MOM schemes developed at BNL. The first results are encouraging and suggest that this approach can also be applied to the more difficult $\Delta I = 1/2$ transition amplitude A_0. RBC-UKQCD plans to carry out this calculation on the QCDCQ machine (BNL's prototype BlueGene/Q) by using a G-parity boundary condition in two space directions for the non-zero momentum final state pions. The BNL-HEP group, as part of the RBC-UKQCD collaboration, will also compute the long- distance contributions to neutral kaon mixing, which may shift the location of the B_K constraint on the CKM unitarity triangle by about 5 % [22]. These calculations thus are important for identifying the source(s) of tension in the CKM unitarity triangle. Since on-going efforts are expected to appreciably reduce the error on B_K from the current result of 3.6%, the aforementioned long distance effect on ε_K is becoming important.

For the muon's anomalous magnetic moment, BNL-HEP plans to extend the current QCD+QED simulation to compute the hadronic light-by-light contribution to $(g-2)_\mu$. Since the measurement of $(g-2)_\mu$ by the E821 experiment at BNL, there has been a greater-than-3σ tension between the experimental measurement and theoretical calculations in the Standard Model. This discrepancy relies, however, on a single experimental result that has yet to be confirmed and on theoretical calculations with large hadronic uncertainties. Experiment E989 at Fermilab will improve the measurement of $(g-2)_\mu$ from E821 by more than a factor of four [23]. This will allow a more precise test of the Standard Model provided there are improvements in the theoretical calculations of the hadronic contributions to muon $g-2$. The current QCD+QED simulation under the SciDAC-2 program is planned to be extended to compute the $O(\alpha_{\rm EM}^3)$ light-by-light contribution with a goal of $\sim 10\%$ accuracy. One particular difficulty in the calculation is the necessity to remove unwanted contribution in $O(\alpha_{\rm EM}, \alpha_{\rm EM}^2)$. Under the current SciDAC-2 program, exploratory studies are in progress to subtract the unwanted contribution using methods developed in [24], part of which is already demonstrated to work in [25]. The program will examine related, less involved, quantities to control systematic errors of the $(g-2)_\mu$ calculations on the lattice. These include the quark condensate magnetic susceptibility [26], hadronic vacuum polarizations [27] and $\pi^0 \to \gamma\gamma$ amplitude [28, 29].

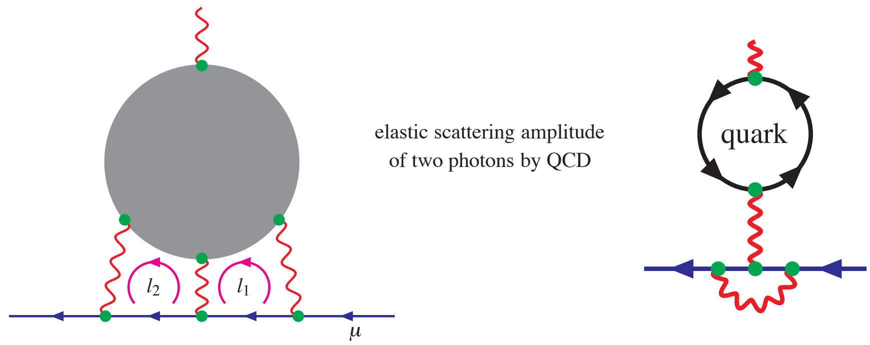

Figure 2: Hadronic light-by-light diagram of muon's anomalous magnetic moment, $(g-2)_\mu$, (right), and an example of unwanted contribution (left).

There are also interesting yet challenging quantities in the flavor-neutral physics besides $(g-2)_\mu$. Long-standing problems for lattice computation include the nucleon's permanent electric dipole moment relevant for the strong CP puzzle and the vacuum angle θ [30, 31, 32], as well as the $U(1)_A$ problem of $\eta' - \eta$ physics [12, 12, 33, 34], or the calculation of the strange-quark content of the nucleon, $\langle N|\bar{s}s|N\rangle$ through reweighting techniques, which is needed to interpret the results of direct dark-matter detection experiments which search for dark matter interacting with nucleons in the detector[35].

These challenging computations, as well as, calculations for $K_L - K_S$ mixing and $\Delta I = 1/2\ K \to \pi\pi$, typically involve disconnected quark diagrams, and thus, they are currently limited by large statistical nose in the Monte Carlo simulations. To efficiently sample the physical observable within reasonable computation resources, various developments both in software, algorithms, and theory will be performed, this is the case even more for DWF simulations due to the larger degrees of freedoms of its five-dimensional structure.

2.1 Software tasks

Future software development at BNL will focus on needs to pursue the physics program at the Intensity Frontier on new leadership class hardware. BNL scientists, along with members of the RBC and UKQCD collaborations, are strongly focused on zero and finite temperature simulations of DWF QCD, making the ability to efficiently generate DWF gauge configurations on the Blue Gene/Q using the Columbia Physics System (CPS) a primary goal of their SciDAC work. While an efficient, assembly level DWF solver is already written (Boyle, UKQCD) and threading the remaining evolution code with OpenMP has begun, there is substantial work that still must be done under this proposal, including multi-threading the measurement codes in CPS. In addition, they will work with others in USQCD to implement QMP over SPI for the Blue Gene/Q for efficiency, and to improve QIO for new architectures, including greater control over the number of processes actively writing to the file system. Existing algorithm research, which will focus on disconnected diagrams for chiral fermions, the force gradient integrator and four-dimensional realizations of DWF, is vital to many of the groups physics goals ($K \to \pi\pi$ decays and $(g-2)_\mu$ calculations, for example), and will be continued.

The HEP group at BNL will be actively involved in software development in coordination with the BNL Nuclear Physics group as well as other members of the USQCD collaboration. Members of the RBC and UKQCD collaborations will also participate in this software development. The physics of the BNL HEP lattice group, as well as much of the RBC and UKQCD collaborations, involves the generation of 2+1 flavor domain wall fermion (DWF) ensembles and the measurement of observables on them. A primary initial goal will be the optimization of codes in CPS (the Columbia Physics System [36]) needed for the generation of DWF ensembles on the Blue Gene/Q computers. This project is well underway, but substantial additional work is needed. Peter Boyle (UKQCD) has written a multi-threaded, highly optimized assembly version of the double precision inverters needed for DWF configuration generation on the Blue Gene/Q which achieves 20%, or higher, of peak already, and it has been integrated into CPS. Other necessary routines, including the force gradient integrator [37] and the gauge and fermion force calculations that are part of the molecular dynamics integrators, are currently being multi-threaded using OpenMP, and substantial efforts are still required to optimize all parts of the evolution code enough to overcome Amdahl's law [38]. Multi-threading the measurement codes in CPS is also a high priority task and we have only just begun.

As part of getting efficient code for Blue Gene/Q, the BNL HEP group will work, in collaboration with ANL and other members of USQCD, to implement QMP over SPI, the native inter-node communication library on the Blue Gene/Q. We have experience in implementing nearest neighbor part of QMP was over SPI for the Blue Gene/P, Peter Boyle has a subset of QMP over SPI working for the Blue Gene/Q, but a full implementation is needed.

The BNL HEP group will also work on improvements to QIO for the new generation of machines, in collaboration with USQCD. On current large supercomputers, we have needed to control the number of nodes doing I/O and use internal machine networks for routing to achieve reliable, and reasonable performance, I/O. Further improvements and functionality are almost certain to be required on the new computer architectures that are appearing now.

Continued research on algorithms will be undertaken by the BNL HEP group and our RBC and UKQCD collaborators. We are actively studying deflation and low mode averaging techniques for DWF, as well as compression techniques to save the large number of eigenvectors that can be required. We will need an efficient implementation of the EigCG algorithm [39] for the Blue Gene/Q to calculate disconnected diagrams for $K \to \pi\pi$ and $(g-2)_\mu$ calculations, among others. Research is ongoing into 4-d realizations of DWF to lower the memory footprint and minimize communications for GPU implementations. We will also continue to pursue reweighting techniques and software supports for $U(1)$ gauge of QCD+QED simulations, relevant for the studies of spectrum studies with electro-magnetism, isospin breaking, and ultimately for $(g-2)_\mu$.

Workplan:

- During the first year, our primary focus will be optimizing DWF calculations on the Blue Gene/Q. We will finish our DWF evolution code and will start ti implement improvements to QMP and QIO. We will produce an efficient implementation of EigCG for the Blue Gene/Q. Furthermore, we will begin to implement multi-threading of all CPS measurement codes.

- During the second year, we will finalize improvements to QIO and QMP and continue to improve CPS. We will finalize multi-threading of measurement codes. Furthermore, we will implement deflation techniques in CPS.

- During the third year, we plan to continue to improve the higher-level organization of CPS and we will implement results of algorithm research in CPS.

3 Literature

[1] Y. Aoki, R. Arthur, T. Blum, P. A. Boyle, D. Brommel, N. H. Christ, C. Dawson and T. Izubuchi *et al.*, Phys. Rev. D **84**, 014503 (2011) [arXiv:1012.4178 [hep-lat]].

[2] D. J. Antonio *et al.* [RBC and UKQCD Collaborations], Phys. Rev. Lett. **100**, 032001 (2008) [hep-ph/0702042 [HEP-PH]].

[3] C. Aubin, J. Laiho and R. S. Van de Water, Phys. Rev. D **81**, 014507 (2010) [arXiv:0905.3947 [hep-lat]].

[4] B. Yoon, T. Bae, Y. -C. Jang, H. -J. Kim, J. Kim, J. Kim, K. Kim and W. Lee *et al.*, PoS LATTICE **2010**, 319 (2010) [arXiv:1010.4778 [hep-lat]].

[5] C. Allton *et al.* [RBC-UKQCD Collaboration], Phys. Rev. D **78**, 114509 (2008) [arXiv:0804.0473 [hep-lat]].

[6] K. -I. Ishikawa *et al.* [PACS-CS Collaboration], Phys. Rev. D **80**, 054502 (2009) [arXiv:0905.0962 [hep-lat]].

[7] J. A. Bailey, C. Bernard, C. E. DeTar, M. Di Pierro, A. X. El-Khadra, R. T. Evans, E. D. Freeland and E. Gamiz *et al.*, Phys. Rev. D **79**, 054507 (2009) [arXiv:0811.3640 [hep-lat]].

[8] R. S. Van de Water and O. Witzel, PoS LATTICE **2010**, 318 (2010) [arXiv:1101.4580 [hep-lat]].

[9] C. Albertus, Y. Aoki, P. A. Boyle, N. H. Christ, T. T. Dumitrescu, J. M. Flynn, T. Ishikawa and T. Izubuchi *et al.*, Phys. Rev. D **82**, 014505 (2010) [arXiv:1001.2023 [hep-lat]].

[10] T. Ishikawa, Y. Aoki, J. M. Flynn, T. Izubuchi and O. Loktik, JHEP **1105**, 040 (2011) [arXiv:1101.1072 [hep-lat]].

[11] T. Blum, T. Doi, M. Hayakawa, T. Izubuchi and N. Yamada, Phys. Rev. D **76**, 114508 (2007) [arXiv:0708.0484 [hep-lat]].

[12] Y. Aoki *et al.* [RBC and UKQCD Collaborations], Phys. Rev. D **83**, 074508 (2011) [arXiv:1011.0892 [hep-lat]].

[13] Y. Aoki, P. A. Boyle, N. H. Christ, C. Dawson, M. A. Donnellan, T. Izubuchi, A. Juttner and S. Li *et al.*, Phys. Rev. D **78**, 054510 (2008) [arXiv:0712.1061 [hep-lat]].

[14] C. Sturm, Y. Aoki, N. H. Christ, T. Izubuchi, C. T. C. Sachrajda and A. Soni, Phys. Rev. D **80**, 014501 (2009) [arXiv:0901.2599 [hep-ph]].

[15] C. Lehner and C. Sturm, Phys. Rev. D **84**, 014001 (2011) [arXiv:1104.4948 [hep-ph]].

[16] T. Blum, R. Zhou, T. Doi, M. Hayakawa, T. Izubuchi, S. Uno and N. Yamada, Phys. Rev. D **82**, 094508 (2010) [arXiv:1006.1311 [hep-lat]].

[17] T. Izubuchi [RBC and UKQCD Collaboration], PoS KAON **09**, 034 (2009).

[18] R. Zhou, T. Blum, T. Doi, M. Hayakawa, T. Izubuchi and N. Yamada, PoS LATTICE **2008**, 131 (2008) [arXiv:0810.1302 [hep-lat]].

[19] T. Ishikawa, Y. Aoki and T. Izubuchi, PoS LAT **2009**, 035 (2009) [arXiv:1003.2182 [hep-lat]].

[20] L. Lellouch and M. Luscher, Commun. Math. Phys. **219**, 31 (2001) [hep-lat/0003023].

[21] C. W. Bernard, T. Draper, A. Soni, H. D. Politzer and M. B. Wise, Phys. Rev. D **32**, 2343 (1985).

[22] A. J. Buras, D. Guadagnoli and G. Isidori, Phys. Lett. B **688**, 309 (2010) [arXiv:1002.3612 [hep-ph]].

[23] New (g - 2) Collaboration, The New $(g-2)$ Experiment: A Proposal to Measure the Muon Anomalous Magnetic Moment to $\pm$ 0.14 ppm Precision, (2010).

[24] M. Hayakawa, T. Blum, T. Izubuchi and N. Yamada, PoS LAT **2005**, 353 (2006) [hep-lat/0509016].

[25] S. Chowdhury, T. Blum, T. Izubuchi, et al., PoS, LATTICE2008:251, (2008).

[26] B. L. Ioffe and A. V. Smilga, Nucl. Phys. B **232**, 109 (1984).

[27] E. Shintani, S. Aoki, H. Fukaya, S. Hashimoto, T. Kaneko, T. Onogi and N. Yamada, Phys. Rev. D **82**, 074505 (2010) [arXiv:1002.0371 [hep-lat]].

[28] E. Shintani, S. Aoki, S. Hashimoto, T. Onogi and N. Yamada, PoS LATTICE **2010**, 159 (2010) [arXiv:1102.5544 [hep-lat]].

[29] M. Knecht and A. Nyffeler, Phys. Rev. D **65**, 073034 (2002) [hep-ph/0111058].

[30] F. Berruto, T. Blum, K. Orginos and A. Soni, Phys. Rev. D **73**, 054509 (2006) [hep-lat/0512004].

[31] E. Shintani, S. Aoki and Y. Kuramashi, Phys. Rev. D **78**, 014503 (2008) [arXiv:0803.0797 [hep-lat]].

[32] S. Aoki, R. Horsley, T. Izubuchi, Y. Nakamura, D. Pleiter, P. E. L. Rakow, G. Schierholz and J. Zanotti, arXiv:0808.1428 [hep-lat].

[33] N. H. Christ, C. Dawson, T. Izubuchi, C. Jung, Q. Liu, R. D. Mawhinney, C. T. Sachrajda and A. Soni *et al.*, Phys. Rev. Lett. **105**, 241601 (2010) [arXiv:1002.2999 [hep-lat]].

[34] K. Hashimoto and T. Izubuchi, Prog. Theor. Phys. **119**, 599 (2008) [arXiv:0803.0186 [hep-lat]].

[35] C. Jung, PoS LAT **2009**, 002 (2009) [arXiv:1001.0941 [hep-lat]].

[36] `http://qcdoc.phys.columbia.edu/cps.html`

[37] H. Yin and R. D. Mawhinney, PoS LATTICE **2011**, 051 (2011) [arXiv:1111.5059 [hep-lat]].

[38] `http://en.wikipedia.org/wiki/Amdahl's_law`

[39] A. Stathopoulos and K. Orginos, SIAM J. Sci. Comput. **32**, 439 (2010) [arXiv:0707.0131 [hep-lat]].

Fermi National Accelerator Laboratory In 2012 Fermilab will begin operating for USQCD a GPU-accelerated cluster designed to perform optimally on problems requiring large GPU-count parallelization and/or good strong scaling. To inform future USQCD cluster designs and software development for the various leadership-class GPU-accelerated machines, the performance of production applications on this dedicated cluster will be analyzed in depth. Specifically, this will involve one major applications in the first year, MILC quark propagator generation. These propagator generation jobs are expected to use to 16 GPUs in parallel, and careful measurement of strong scaling effects will be necessary to adjust GPU counts to optimize physics throughput.

In addition, Fermilab will collaborate with Indiana, Boston, NCSA, and NVidia on the development and optimization of software for GPU-accelerated clusters, with a focus in the first year on gauge configuration generation and analysis of performance of the first versions of this code on the Fermilab GPU-accelerated cluster. This work will involve the asqtad and HISQ actions. For configuration generation, in addition to the Dirac inverter the various force terms are required for a full GPU-implementation of evolution. This analysis of configuration generation will help determine if GPU-accelerated clusters are suitable in the near term for taking on some of the work that to date has required leadership-class supercomputers. These investigations will involve jobs with GPU counts of up to the entire cluster (176 NVidia M2050 GPUs).

Fermilab will also provide leadership in the development of optimized software for asqtad and HISQ actions for clusters accelerated with Intel many-core architecture hardware. This work will start once “MIC” hardware becomes available.

Conventional x86_64-based clusters continue to be critical USQCD resources, and new generations of processors from both Intel and AMD will have an impact on USQCD production in the coming years. Fermilab will provide specific optimizations for these new processors in the QLA library. The areas of optimization that will be applied include the use of the wider SSE units (the AVX instruction set), fused-multiply-add instructions (available now in the AMD XOP extensions, and on Intel processors after Sandy Bridge), and exploitation of the relaxation of memory alignment restrictions. Memory bandwidth typically constrains the performance of LQCD ocdes, but because of the size of SU(3) data structures the 16-byte alignment restrictions in processors prior to Intel’s Sandy Bridge and AMD’s Interlagos prevented the use of certain cache-friendly SSE operations.

LUA is a light-weight scripting language that could greatly accelerate the coding and development of new physics projects. Currently, to specify the calculation of a three-point function physicists using the MILC application write a parameter input file of a few hundred lines. This file is read and parsed and the calculations then proceed according to a rigid pattern. With LUA the parameter input file would be replaced by a LUA script that carries out the calculation with a series of procedure calls. Examples are procedures for constructing sources, doing sparse solves (inverters), and tying together propagators. The underlying operations would be mostly our existing MILC procedures, but the management (currently embodied in our traditional ”setup.c” and ”control.c”) would be done by the C++-like LUA scripting language. We will work with MILC developers at Utah on the development of the scriptable modules.

- *Year 1:*
  – comprehensive performance analysis of multiple-GPU production runs on the Fermilab USQCD GPU-accelerated cluster
  – performance analysis of multi-TFlop-scale staggered configuration generation on the GPU-accelerated cluster, collaborating with experts from Nvidia, NCSA, and Indiana University
  – addition of Intel Sandy Bridge and AMD Interlagos optimizations to QLA library, and measurement of effect on code throughput. Optimizations will include use of AVX and XOP (FMA4) instructions
  – demonstrate Lua-based script version of a MILC production executable for calculation of three-point functions

– refinement and packaging of the databases and tools developed during SciDAC-2 cluster reliability subproject for the monitoring of the batch schedulers deployed on USQCD clusters and the realtime analysis of production running on these resources

- *Year 2:*
  – collaboration on the optimization of GPU software and the integration of GPU softwware into application codes, included performance analysis of production runs on Fermilab USQCD GPU-accelerated clusters
  – leadership in the exploration of the many-core Intel architecture and the guiding the software optimizations for Intel many-core-accelerated clusters
  – QLA optimizations for Intel Ivy Bridge and AMD socket G2012 processors
  – general Lua-based implementations of MILC applications suitable for production

- *Year 3:*
  – continuing work on software optimizations for GPU-accelerated and Intel many-core-acclerated clusters
  – continued maintenance of QLA for x86_64 processors
  – with the MILC collaboration, further expansion of Lua support for the MILC code, should it prove successful
  – exploration of the use of domain-specific languages for automated metadata collection and the direction of analysis workflows

Illinois Statement of Work

1 Introduction

The multi-institutional USQCD collaboration is submitting a joint proposal to the SciDAC-3 computational high energy physics program. This narrative describes the Illinois contributions to the collaborative effort. The submitted budget covers only the University of Illinois. Please refer to the full proposal to see how this work integrates with the multi-institutional collaborative effort.

The central objective of this project is to advance scientific discovery to a new level of capability by supporting the continued rapid evolution of the community SciDAC code suite for GPU and multicore operation. We will assist in developing CUDA GPU modules of general use to the entire US lattice QCD community, and we will integrate them into the MILC code. Through our active participation in the USQCD SciDAC software committee, we will coordinate with similar efforts on related codes by our colleagues at Argonne National Laboratory, Boston University, Columbia University, Fermilab, and Jefferson Laboratory. Through our computational science colleagues in the SciDAC institutes, particularly SUPER, we will use state-of-the art performance measures and adaptive optimization. We are confident that such coordination will benefit the entire US lattice gauge theory community.

2 Statement of Work: Illinois

The Innovative Systems Lab (ISL) at NCSA focuses on the advancement of hardware and software technologies that show potential in early research environments and applies the knowledge of complex computing architectures to scientific domains from a perspective of the underlining science. ISL maintains and operates various machines, including GPUs, Cell/BE, FPGAs and Intel's MIC processors while the staff at ISL is actively porting scientific applications to the emerging processors, by working closely with domain scientists as well as the vendors.

In 2006, Guochun Shi of NCSA's Innovative Systems Laboratory, began experimenting with the MILC code on the Cell/BE processor [1]. In 2009, he started to work with members of the MILC Collaboration and the Boston University to develop the QUDA library, a GPU library for lattice QCD applications. In the next two years the collabration led to the development of GPU code for staggered quarks to enhance the MILC code. This work is very timely, since the forthcoming NCSA Blue Waters and the ORNL Titan machines will feature thousands of GPU-enhanced nodes. In the initial stage of this work, an asqtad solver running on a single GPU was developed [2, 3]. Fattening, gauge and fermion force routines for single GPU were developed for the asqtad action [3]. The solver was parallelized by dividing the lattice in the time dimension [4]. Then, the solver was generalized to divide the lattice in all four dimension and scaling to over 100 GPUs on large grids was seen to meet or exceed typical supercomputer speeds [5]. Under this grant he will work with members of MILC and the broader USQCD QUDA group, as well as with colleagues at NVIDIA, to port the remaining routines needed to generate HISQ configurations on multiple GPUs, the gauge and fermion forces, and the smeared links.

- *Year 1:*
  – in coordination with MILC and the rest of the QUDA group, provide the HISQ enhancements needed for running on multiple GPUs
  – port the gauge force routines to multiple GPUs
  – tune the performance of the new routines
  – begin work on tuning for the new NVIDIA Kepler GPU that will be used on Titan and Blue Waters, and very likely on future USQCD clusters.

- *Year 2:*
  – complete tuning of codes for Kepler architecture
  – develop a more exible API that will enable the programmer to indicate when data needs to be moved to the GPU to start the routine.

- *Year 3:*
  – complete implementation of the API developed in Year 2
  – implement on GPUs promising new algorithms developed by other members of the USQCD collaboration, FASTMath and NVIDIA.

Lawrence Livermore National Laboratory and FASTMath Institute

Lattice field theorists from LLNL and applied mathematicians from the FASTMath institute at LLNL will provide a framework for fast turn-around exploration and testing of novel algorithmic ideas for lattice field theory that can then be incorporated to the USQCD science codes. In addition this collaboration will be able to test these ideas at scale on LLNL's massively parallel supercomputers as available to the FASTMath and LLNL Lattice group for algorithmic development. In particular the focus will be in interfacing Lattice QCD kernels with the FASTMath HYPRE applied math effort. HYPRE provides novel scalable linear solvers for solving sparse linear systems of equations on massively parallel supercomputers. HYPRE represents a large and continuing investment by the applied math community. This collaboration aims to test existing and new HYPRE algorithms on various Lattice QCD kernels with a fast turn-around time. The algorithms that give performance improvements will then be given to USQCD personnel to code them into the main production software stacks of USQCD. Furthermore, extensions to the HYPRE code for Lattice QCD will be incorporated into the main branch and made available to future users of HYPRE.

The yearly deliverables are:

- *Year 1:*
  – Expand HYPRE to include the use of complex numbers. Expand HYPRE from the current three-dimensional implementation to an arbitrary, user defined, dimension.
  – The Wilson Dirac operator will be implemented in HYPRE and solved using HYPRE's multigrid methods. The performance results will then be directly compared with the existing USQCD Wilson Dirac multigrid solver on the BG/L.

- *Year 2:*
  – Implement and interface the major Lattice QCD kernels with HYPRE. In particular the following kernels will be implemented: Domain Wall Fermions, Staggered-based fermions, and Wilson-based fermions.

- *Year 3:*
  – Use the infrastructure built in the first two years to explore new multigrid methods and algorithms for Lattice QCD. The fast turn-around times will enable exploration of a large algorithmic space. The focus will be on multi-scale algorithms that will be needed as Lattice QCD simulations will be done on increasingly larger lattices capable of containing many different physics scales.

Personnel: A postdoc from the field of Lattice QCD with clear ability and interest in algorithmic development will be hired by FASTMath and the LLNL Lattice group to perform the above work under the collaboration and supervision of: R. Falgout (LLNL) 0.1 FTE, R. Soltz (LLNL) 0.05 FTE, and P. Vranas (LLNL) 0.05 FTE.

1 Introduction

The multi-institutional USQCD collaboration is submitting a joint proposal to the SciDAC-3 computational high energy physics program. This narrative describes the Syracuse contribution to the collaborative effort. The submitted budget covers only Syracuse University. Please refer to the full proposal to see how this work integrates with the multi-institutional collaborative effort (see additional attached documents)

The primary goal of the Syracuse effort in collaboration with Argonne National Laboratory will be to lead the development of software suitable for large scale numerical studies of theories of Beyond Standard Model Physics. This includes composite Higgs theories such as technicolor and supersymmetric lattice field theories and fits within the FUEL program being developed at Argonne (Framework for Unified Evolution of Lattices) The Syracuse focus is primarily on lattice supersymmetry and will involve collaboration between Prof. Simon Catterall at Syracuse with other members of the USQCD collaboration, for example Prof. J. Giedt at Rensselaer Polytechnic Institute, Prof. Richard Brower at Boston University, Prof. George Fleming at Yale University Dr. Pavlos Vranas at Lawrence Livermore National Laboratory and Dr. James Osborn at Argonne National Laboratory together with students and postdocs at these institutions.

Prof. Catterall has been a leader for almost a decade in the development of new theoretical approaches to lattice supersymmetry (SUSY) with many highly cited publications on the subject . A good reference to some of this work can be found in the recent monograph written in collaboration with Prof. D. B. Kaplan and Prof. M. Ünsal [1]. He and his collaborators within USQCD have also pioneered numerical simulations of several different supersymmetric lattice theories including the first simulations of $\mathcal{N} = 1$ super Yang-Mills theory using domain wall fermions - the cornerstone of the current proposal [2].

The goal of this part of the current SciDAC proposal is to create a stable, flexible and efficient code base suitable for a high precision study of both this theory and its more phenomenologically interesting companion super QCD. The latter theory plays a very important role in understanding how supersymmetry may be spontaneously broken at low energy - a necessary requirement of any realistic theory incorporating SUSY since the world we inhabit does not exhibit manifest supersymmetry.

The simplest example of a realistic supersymmetric theory incorporating the usual elementary particles is called the Minimal Supersymmetric Standard Model (MSSM). Searching for signals of the MSSM has been one of the primary goals of the Large Hadron Collider at CERN. The consequences of supersymmetry breaking are parametrized in the MSSM by the addition of large numbers of so-called soft breaking parameters. These parameters include, for example, the masses of the superpartners of the Standard model particles. In principle it is generally believed that these parameters are in fact determined by the spontaneous non perturbative breaking of SUSY in a hidden sector; the effects of this breaking being communicated to the MSSM fields through so-called messenger fields. One attractive possibility for this hidden sector theory is super QCD since the latter is known from the work of Seiberg and Intriligator [3] to have long lived metastable SUSY breaking vacua. A small number of non perturbatively determined parameters in this hidden sector super QCD theory then determines, in principle, the large number of soft parameters in the MSSM.

The non-perturbative character of this symmetry breaking has long hindered efforts to compute the low energy features of SUSY breaking. It is a primary goal of this program to develop a quantitative understanding of SUSY breaking in super QCD in order to strongly constrain the parameter space for the MSSM and hence help determine which possible supersymmetric theories of Beyond Standard Model physics are compatible with LHC data.

To study super QCD will require the use of domain wall fermion actions incorporating additional Yukawa terms. Additional fine tuning of the scalar (squark) sector is also needed to recover full supersymmetry in the continuum limit [4]. To access the region of the parameter space where SUSY breaking effects are expected will require the ability to dial both the number of colors and the number of flavors. None of the code for such simulations currently exists; it will be the prime focus of the Syracuse and Argonne efforts to develop a suitable code base. This will piggyback on general BSM kernel development within the SciDAC/USQCD collaboration but will require additional software development to generate efficient full simulation code.

2 Statement of Work: Syracuse

We feel we have a strong record of pioneering scientific achievements in this area and the relevant expertise to push through this program in collaboration with the rest of the USQCD and SciDAC communities. The postdoc supported on the Syracuse grant will focus their energies on developing the relevant application software needed to bridge between the theoretical lattice models that have been constructed and the low level code base being developed elsewhere within this proposal that will support BSM efforts in general. They will also be responsible for a portion of the BSM kernel development being developed at Argonne with James Osborn and with an effort to develop eigensolvers useful for HISQ and DWF actions.

For example, while the base BSM effort will deliver inverters for domain wall fermions in adjoint representation with arbitrary numbers of colors and flavors, additional software will be needed to integrate these inverters into efficient and fully functioning rational hybrid monte carlo code suitable for full dynamical fermion simulations. This will be one of the tasks for the Syracuse group in collaboration with Argonne. Given sufficient time we will also work on developing a first generation code for simulating $\mathcal{N} = 4$ super Yang-Mills using the new formulations which retain exact supersymmetry at non-zero lattice spacing.

3 Deliverables

- Year 1: With Argonne National Lab: Development of inverters for $\mathcal{N} = 1$ super Yang-Mills for adjoint representation domain wall fermions and integration of those inverters into working simulation code.
- Year2: Inclusion of support in Dirac operators for Yukawa interactions necessary for super QCD and general four fermion interactions.
- Year 3: Development of code for full super QCD simulations. Optimization of all elements of code base and simulation codes. Development of code for $\mathcal{N} = 4$ super Yang-Mills using twisted lattice formulations.

A.7 Biographical Sketches

Biographical sketches of senior personnel are presented on the following pages.

Curriculum Vitae
PAUL B. MACKENZIE

Contact Information

Theoretical Physics Department | phone: (630) 840-3347
Fermi National Accelerator Laboratory | fax: (630) 840-5435
PO Box 500 | email: mackenzie@fnal.gov
Batavia, IL 60510

Professional Preparation

Ph.D. Physics, Cornell University, 1981 (Advisor, Peter Lepage)
B.S. Physics and Mathematics, University of Illinois, Summa Cum Laude, 1975

Appointments

2008-present, Fermilab, Scientist III
1991-2008, Fermilab, Scientist I, II
1986-1991, Fermilab, Associate Scientist
1984-1986, Institute for Advanced Study, Member
1981-1984, Fermilab, Research Associate

Visiting Positions

Institute for Theoretical Physics, Santa Barbara, 8/1990-12/1990
Institute for Nuclear Theory, Seattle, 6/1993-7/1993
Center for Computational Physics, Tsukuba, 6/1996-9/1996

Honors and Awards

Fellow, American Physical Society, 1996
Fellow, American Association for the Advancement of Science, 2011

Outside Activities

Member, USQCD Executive Committee, 1999-present; chair, 2009-present
Member, Organizing committee, SciDAC 2010
International Advisory Committee, International Symposia on Lattice Field Theory: Lattice 89, Capri, 1989; Lattice 90, Tallahassee, 1990; Lattice 93, Dallas, 1993; Lattice 96, St. Louis, 1996; Lattice 03, Tuskuba, 2003; Lattice 07, Regensburg, 2007.
Organizing Committee: Lattice 88, Fermilab, 1988; Program on Lattice Gauge Theory, Institute for Theoretical Physics, Santa Barbara, 1990; Aspen Center for Physics Winter Workshop, Aspen, 1995; Aspen Center for Physics Summer Workshop, Aspen, 1997; Lattice 04, Fermilab, 2004.

Collaborators

C. T. H. Davies (Glasgow), G. P. Lepage (Cornell), J. Shigemitsu (Ohio State), C. Bernard (Washington U), C. DeTar (Utah), S. Gottlieb (Indiana), U. M. Heller (APS), J. Osborn (ANL), R. Sugar (UCSB), D. Toussaint (Arizona), M. Di Pierro (DePaul, A. El-Khadra (Illinois), A. S. Kronfeld (Fermilab), R. Van de Water (BNL), J. Laiho (Glasgow), E. Gamiz (Grenada), J. Simone (Fermilab).

Selected Publications

- Jon A. Bailey, C. Bernard, C. DeTar, M. Di Pierro, A.X. El-Khadra, R.T. Evans, E. D. Freeland, E. Gamiz, Steven Gottlieb, U.M. Heller, J.E. Hetrick, A.S. Kronfeld, J. Laiho, L. Levkova, P.B. Mackenzie, M. Okamoto, J. N. Simone, R. Sugar, D. Toussaint, and R.S. Van de Water. "The $B \to \pi \ell \nu$ semileptonic form factor from three-flavor lattice QCD: A Model-independent determination of $|V_{ub}|$," Phys. Rev. D **79**, 054507 (2009) [arXiv:0811.3640 [hep-lat]].
- C. Bernard, C. DeTar, M. Di Pierro, A. X. El-Khadra, R. T. Evans, E. D. Freeland, E. Gamiz, Steven Gottlieb, U. M. Heller, J. E. Hetrick, A. S. Kronfeld, J. Laiho, L. Levkova, P. B. Mackenzie, M. Okamoto, J. Simone, R. Sugar, D. Toussaint, and R. S. Van de Water, "The $\bar{B} \to D^* \ell \bar{\nu}$ form factor at zero recoil from three-flavor lattice QCD: A Model independent determination of $|V_{cb}|$," Phys. Rev. D **79**, 014506 (2009) [arXiv:0808.2519 [hep-lat]].
- C. Aubin, C. Bernard, C. DeTar, M. Di Pierro, E. D. Freeland, Steven Gottlieb, U. M. Heller, J. E. Hetrick, A. X. El-Khadra, A. S. Kronfeld, L. Levkova, P. B. Mackenzie, D. Menscher, F. Maresca, M. Nobes, M. Okamoto, D. Renner, J. Simone, R. Sugar, D. Toussaint, and H. D. Trottier, "Charmed meson decay constants in three-flavor lattice QCD," Phys. Rev. Lett. **95**, 122002 (2005) [arXiv:hep-lat/0506030].
- C. Aubin, C. Bernard, C. DeTar, M. Di Pierro, A. El-Khadra, Steven Gottlieb, E. B. Gregory, U. M. Heller, J. Hetrick, A. S. Kronfeld, P. B. Mackenzie, D. Menscher, M. Nobes, M. Okamoto, M. B. Oktay, J. Osborn, J. Simone, R. Sugar, D. Toussaint, and H. D. Trottier, "Semileptonic decays of D mesons in three-flavor lattice QCD," Phys. Rev. Lett. **94**, 011601 (2005) [arXiv:hep-ph/0408306].
- C. T. H. Davies, E. Follana, A. Gray, G. P. Lepage, Q. Mason, M. Nobes, J. Shigemitsu, H. D. Trottier, M. Wingate, C. Aubin, C. Bernard, T. Burch, C. DeTar, Steven Gottlieb, E. B. Gregory, U. M. Heller, J. E. Hetrick, J. Osborn, R. Sugar, D. Toussaint, M. Di Pierro, A. El-Khadra, A. S. Kronfeld, P. B. Mackenzie, D. Menscher, and J. Simone, "High-precision lattice QCD confronts experiment," Phys. Rev. Lett. **92**, 022001 (2004) [arXiv:hep-lat/0304004].
- S. Hashimoto, A. X. El-Khadra, A. S. Kronfeld, P. B. Mackenzie, S. M. Ryan and J. N. Simone, "Lattice QCD calculation of anti-B $\to$ D l anti-nu decay form factors at zero recoil," Phys. Rev. D **61**, 014502 (2000) [arXiv:hep-ph/9906376].
- A. X. El-Khadra, A. S. Kronfeld and P. B. Mackenzie, "Massive fermions in lattice gauge theory," Phys. Rev. **D55**, 3933 (1997) [hep-lat/9604004].
- G. P. Lepage and P. B. Mackenzie, "On the viability of lattice perturbation theory," Phys. Rev. **D48**, 2250 (1993) [hep-lat/9209022].
- A. X. El-Khadra, G. Hockney, A. S. Kronfeld and P. B. Mackenzie, "A Determination of the strong coupling constant from the charmonium spectrum," Phys. Rev. Lett. **69**, 729 (1992).
- P. B. Mackenzie and G. P. Lepage, "QCD corrections to the gluonic width of the upsilon meson," Phys. Rev. Lett. **47**, 1244 (1981).

Biographical Sketch

Curriculum Vitae

RICHARD C. BROWER

Contact Information

Physics Department	phone: 617 353 6052
Boston University	fax: 617 358 2487
590 Commonwealth Ave	iPHONE: 617 833 5811
Boston, MA 02215	Email: brower@bu.edu

Professional Preparation
Ph.D. Physics, University of California, Berkeley, 1969
M.S. Applied Math, Harvard University, 1964
B.S. Physics, Harvard University, 1963

Appointments and Visiting Positions
1987-present, Professor of Physics and Engineering, Boston University
1973-86, Professor of Physics, University of California, Santa Cruz, CA
1978-86, Research Professor, Santa Cruz Institute of Particle Physics
1972-73, Senior Research Associate, Cal. Tech, Pasadena, CA
1969-72, Research Associate, MIT, Cambridge, MA
2006-present, Visiting Professor of Physics, Brown University, Providence R.I.
1992-present, Visiting Scientist, MIT Center for Theoretical Physics
1980-81, Visiting Associate Professor of Physics, Harvard Cambridge, MA

Honors and Awards
1974-76, A. P. Sloan Research Fellow, SLAC and MIT
1970-71, NSF Postdoctoral Fellow, C.E.R.N., Geneva, Switzerland

Research Interests and Expertise
Professor Brower has worked in several fields of theoretical and computational physics — string theory, hadron phenomenology, Quantum Chromodynamics, lattice formulations of quantum field theory and statistical mechanics and molecular dynamics. He has experience with parallel algorithms starting with data parallel methods on the the Connection Machine to GPU cost effective clusters architectures. Algorithmic research includes in multi-grid Dirac solvers, the Möbius Domain Wall algorithm, chronological inverter and cluster algorithms for bosonic and fermionic systems. He serves on the USQCD Executive Committee and as the National Software Co-ordinator for the SciDAC software infrastructure project.

Synergistic Activities
Organizer of Aspen QCD Workshop 1997 and Strong Dynamics BSM Workshop 2010.

Member, USQCD Executive Committee, 2001-present.
SciDAC software co-ordinator, chair of Software Committee, 2001-present.
International Lattice Data Grid (ILDG)2003-pres; Chair in 2005, 2010.
International Advisory Committee for Lattice Field Theory, Virginia 2008.
Member, Organizing committee, SciDAC 2007.

Thesis, Postdoctoral Advisors and Advisees

Geoffrey Chew (U.C. Berkeley); Francis Low (MIT), Murray GellMann (CalTech); Ronald Babbich (NVIDIA/PSC), Michael Cheng (BU), Mike Clark (NVIDIA/Harvard), Saul Cohen (INT Seattle), James Osborn (ANL), Oliver Witzel (BU)

Collaborators

T. Appelquist (Yale), R. C. Babich (Boston University), James Brannick (Penn. State U), S. Catterall (Syracuse), M. A. Clark (Harvard), S. Cohen (INT Seattle), Marko Djuric (Universidade do Porto), Robert Edwards (Jeff Lab), George T. Fleming (Yale), J. Giedt (Rensselaer), S.Gottlieb (Indiana U.), B. Joo (Jeff Lab), J. Kiskis (U.C. Davis), T. A. Manteuffel (U. of Colorado, Boulder), S. McCormick (U. of Colorado, Boulder) Meifeng Lin (Yale), J. Negele (MIT), E. Neil (FNAL), K. Orginos (William and Mary), J. C. Osborn (ANL), J. Polchinski (Kavali Institute), C. Rebbi (BU), Ina Sarcevic (Arizona U), David Schaich (Boulder Colorado), G. Shi (NCSA), M. Strassler (Rutgers), C-I Tan (Brown), Pavlos Vranas (LLNL)

Publications Most Relevant to Proposal

- R. Babich, M. A. Clark, B. Joo, G. Shi, R. C. Brower and S. Gottlieb `Scaling Lattice QCD beyond 100 GPUs` SuperComputing 2011 paper, arXiv:1109.2935 [hep-lat].
- T. Appelquist, R. Babich, R.C. Brower, M.Cheng, M.Clark, S. Cohen, G. Fleming, J. Kiskis, E. Neil, J. Osborn, R. Claudio, D. Schaich and P. Vranas, `Parity Doubling and the S Parameter Below the Conformal Window` Phys Rev Lett.106.231601 (2011), arXiv: 1009.5967v1 [hep-ph].
- T. Appelquist, A. Avarkian, R. Babich, R.C. Brower, M.Cheng, M.Clark, S. Cohen, G. Fleming, J. Kiskis, E. Neil, J. Osborn, R. Claudio, D. Schaich and P. Vranas, `Toward TeV Conformality` Phys Rev Lett.104.071601 (2010), arXiv:0910.2224 [hep-ph].
- M. A. Clark, R. Babich, K. Barros, R. C. Brower and C. Rebbi) `Solving Lattice QCD systems of equations using mixed precision solvers on GPUs` Comput. Phys. Commun.10.1016 arXiv:2010.05.002 [hep-lat](2010).
- R. Babich, J. Brannick, R. C. Brower, M. A. Clark, T. A. Manteuffel, S. F. McCormick, J. C. Osborn and C. Rebbi , `Adaptive multigrid algorithm for the lattice Wilson-Dirac operator` Phys Rev Lett. 105, 201602 (2010). arXiv: 1005.3043 [hep-lat].
- J. Brannick, R. C. Brower, M. A. Clark, J. C. Osborn and C. Rebbi, `Adaptive Multigrid Algorithm for Lattice QCD` Phys.Rev.Lett.100:041601,2008, arXiv: 0707.4018 [hep-lat].
- R. C. Brower, R. Babich, K. Orginos, C. Rebbi, D. Schaich and P. Vranas `Möbius Algorithm for Domain Wall and GapDW Fermions` LATTICE2001, arXiv: 0910.2813 [hep-lat].

Curriculum Vitae
SIMON CATTERALL

Contact Information

Department of Physics
Syracuse University
Syracuse
NY 13244

phone: (315) 443-5978
fax: (315) 443-9103
email: smc@physics.syr.edu

Professional Preparation

Ph.D. Theoretical Physics, Oxford University, UK, 1989.
B.A. (Hons 1st class) Physics, Christ Church, Oxford University, 1985.

Appointments

2006-present, Professor, Syracuse University.
2000-2006, Associate Professor, Syracuse University.
1994-2000, Assistant Professor, Syracuse University.
1993-1994, CERN Fellow, CERN, Geneva, Switzerland.
1991-1993, SERC Advanced Fellow DAMTP and Junior Research Fellow Trinity Hall, Cambridge University, UK.
1990-1991, Research Associate, Physics Department, University of Illinois at Urbana-Champaign.
1988-1990, Research Associate, DAMTP, Cambridge University, UK.

Visiting Positions

Theory Group, Los Almamos National Laboratory, 1/1998-4/1998.
Niels Bohr Institute, University of Copenhagen 1/2002-7/2002.
Center for Particle Physics Phenomenology, University of Southern Denmark, 2/2011-6/2011.

Honors and Awards

Research Fellow, Trinity Hall, Cambridge University, 1990-1993.
Dixon Graduate Scholarship, Christ Church, Oxford University 1985-88.
Open Scholarship, Christ Church, Oxford 1982-85.

Outside Activities

Member, USQCD Scientific Program Committee, 2010-present.
Chair organizing committee, BalFest, Syracuse Dec 3 2011.
International Advisory Committee, Origins of Mass 2008, University of S. Denmark.
Organizing committee, MRST workshop, Syracuse 1997.
Reviewer for numerous DOE, NSF proposals.
Invited participant in Exascale workshop for High Energy Physics 2008.

Publication Summary
70 peer reviewed publications, circa 40 proceedings, 1 monograph.
10 Top 50/100+ cited papers, 4 single author.
52 invited talks and seminars since 2002. Lectures at 3 summer schools.

Collaborators
J. Hubisz (Syracuse U.), F. Sannino (U. of S. Denmark), L. Del Debbio (U. of Edinburgh), D. B. Kaplan (U. of Washington), M. Ünsal (San Francisco State U.), J. Giedt (Rensselaer Polytechnic Institute), R. Brower (Boston U.), G. Fleming (Yale U.), P. Vranas (Lawrence Livermore National Lab), T. Wiseman (Imperial College, UK)

Selected Publications

- S. Catterall, E. Dzienkowski, J. Giedt, A. Joseph and R. Wells, "Perturbative renormalization of lattice N=4 super Yang-Mills theory," JHEP **1104**, 074 (2011) [arXiv:1102.1725 [hep-th]].
- S. Catterall, D. B. Kaplan and M. Unsal, "Exact lattice supersymmetry," Phys. Rept. **484**, 71 (2009) [arXiv:0903.4881 [hep-lat]].
- J. Giedt, R. Brower, S. Catterall, G. T. Fleming and P. Vranas, "Lattice super-Yang-Mills using domain wall fermions in the chiral limit," Phys. Rev. D **79**, 025015 (2009) [arXiv:0810.5746 [hep-lat]].
- S. Catterall, J. Giedt, F. Sannino and J. Schneible, "Phase diagram of SU(2) with 2 flavors of dynamical adjoint quarks," JHEP **0811**, 009 (2008) [arXiv:0807.0792 [hep-lat]].
- S. Catterall and T. Wiseman, "Black hole thermodynamics from simulations of lattice Yang-Mills theory," Phys. Rev. D **78**, 041502 (2008) [arXiv:0803.4273 [hep-th]].
- S. Catterall, "From Twisted Supersymmetry to Orbifold Lattices," JHEP **0801**, 048 (2008) [arXiv:0712.2532 [hep-th]].
- S. Catterall and F. Sannino, "Minimal walking on the lattice," Phys. Rev. D **76**, 034504 (2007) [arXiv:0705.1664 [hep-lat]].
- S. Catterall, "A Geometrical approach to N=2 super Yang-Mills theory on the two dimensional lattice," JHEP **0411**, 006 (2004) [hep-lat/0410052].
- S. Catterall, "Lattice supersymmetry and topological field theory," JHEP **0305**, 038 (2003) [hep-lat/0301028].
- S. Catterall and S. Karamov, "Exact lattice supersymmetry: The Two-dimensional N=2 Wess-Zumino model," Phys. Rev. D **65**, 094501 (2002) [hep-lat/0108024].

Curriculum Vitae
NORMAN H. CHRIST

Contact Information

Department of Physics	phone: (212) 854-3307
Columbia University	fax: (212) 932-3169
New York, NY, 10027	email: nhc@phys.columbia.edu

Professional Preparation

Ph.D., Columbia University, 1966 (Advisor, T. D. Lee)
B.A., Columbia University (summa cum laude, 1965)

Appointments

Gildor Professor of Computational Physics, Columbia University (1999-present)
Chair, Department of Physics, Columbia University (1994-1999)
Professor of Physics, Columbia University (1974-1999)
Associate Professor, Columbia University (1969-1974)
Assistant Professor, Columbia University (1966-1969)
Instructor, Princeton University (1967-1969)

Honors and awards

Salutatorian, Columbia College (1965)
Sloan Fellowship (1967)
American Physical Society Fellow (1981)
Gordon Bell Prize, QCDSP computer (1998)

Outside Activities

Member Executive Committee, USQCD Collaboration, 1999 - present
IBM contractor, participating in design of LLNL Sequoia System
Leader of QCDOC and QCDSP computer projects
Member International Advisory Committee, Lattice Field Theory Symposia:

Lattice 2012, Cairns, Australia	Lattice 2011, Lake Tahoe, USA
Lattice 2010, Sardinia, Italy	Lattice 2009, Beijing, China
Lattice 2008, William and Mary, USA	Lattice 2007, Regensburg, Germany
Lattice 2004, Fermilab, USA	Lattice 2000, Bangalore, India
Lattice 1996, Saint Louis, USA	Lattice 1995, Melbourne, Australia

Member, Advisory Board, New York Center for Computational Science, 2007–present

Collaborators

Chris Allton (Swansea), Yasumichi Aoki (BNL), Alexei Bazavov (Arizona), Tanmoy Bhattacharya (LANL), Thomas Blum (University of Connecticut), Peter Boyle (Edinburgh), Michael Cheng (LLNL), Saul Cohen (University of Washington), Michael Clark (Harvard), Carleton DeTar (Utah), Shinji Ejiri (Niigata), Jonathan Flynn (Southampton) Steven Gottlieb (Indiana), Rajan Gupta (LANL), Urs Heller (APS), Kay Huebner (BNL), Taku Izubuchi (BNL), Chulwoo Jung (BNL), Andreas Jüttner (Mainz), Frithjof Karsch (BNL), Anthony Kennedy (Edinburgh), Richard

Kenway (Edinburgh), Edwin Laermann (Bielefeld), Ludmila Levkova (Utah), Huey-Wen Lin (University of Washington) Chuan Miao (BNL), Robert Mawhinney (Columbia University), Christopher Maynard (Edinburgh), Shigemi Ohta (BNL), Brian Pendleton (Edinburgh), Peter Petreczky (BNL), Christopher Sachrajda (Southampton), S. Sasaki (Tokyo), Enno Scholz (Regensburg), Christian Schmidt (Bielefeld), Wolfgang Soeldner (GSI), Ron Soltz (LLNL), Amarjit Soni (BNL), Robert Sugar (UCSB), Douglas Toussaint (Arizona), Pavlos Vranas (LLNL), Tilo Wettig, (Regensburg) Takeshi Yamazaki (Tsukuba), James Zanotti (Edinburgh)

Selected Publications Relevant to This Proposal

- P A. Boyle, D. Chen, N. H. Christ, M. Clark, S. D. Cohen, C. Cristian, Z. Dong, A. Gara, B. Joo, C. Jung, C. Kim, L. Levkova, X. Liao, G. Liu, R. D. Mawhinney, S. Ohta, K. Petrov, T. Wettig, A. Yamaguchi, *Overview of the QCDSP and QCDOC computers*, IBM Research Journal, Vol **49**, No. 2/3, 351 (2005).
- M. Cheng, N. H. Christ, M. Li, R. D. Mawhinney, D. Renfrew, P. Hegde, F. Karsch, M. Lin, P. Vranas, *The finite temperature QCD phase transition using 2 + 1 flavors of domain wall fermions at* $N_t = 8$, Phys. Rev. D81:054510, 2010 (arXiv:0911.3450 [hep-lat]).
- N.H. Christ, C. Dawson, T. Izubuchi, C. Jung, Q. Liu, R.D. Mawhinney, C.T. Sachrajda, A. Soni, R. Zhou, *The* η *and* η' *mesons from Lattice QCD* Phys. Rev. Lett. 105:241601, 2010 (arXiv:1002.2999 [hep-lat]).
- Y.Aoki, R. Arthur, T. Blum, P.A. Boyle, D. Brommel, N.H. Christ, C. Dawson, J.M. Flynn, T. Izubuchi, X-Y. Jin, C. Jung, C. Kelly, M. Li, A. Lichtl, M. Lightman, M.F. Lin, R.D. Mawhinney, C.M. Maynard, S. Ohta, B.J. Pendleton, C.T. Sachrajda, E.E. Scholz, A. Soni, J. Wennekers, J.M. Zanotti, R. Zhou (RBC and UKQCD Collaborations), *Continuum Limit Physics from 2+1 Flavor Domain Wall QCD* Phys. Rev. D83:074508 2011. (arXiv:1011.0892 [hep-lat])
- Y. Aoki, R. Arthur, T. Blum, P.A. Boyle, D. Brmmel, N.H. Christ, C. Dawson, T. Izubuchi, C. Jung, C. Kelly, R.D. Kenway, M. Lightman, R.D. Mawhinney, Shigemi Ohta, C.T. Sachrajda, E.E. Scholz, A. Soni, C. Sturm, J. Wennekers, R. Zhou, *Continuum Limit of* B_K *from 2+1 Flavor Domain Wall QCD* Phys. Rev. D, to appear (arXiv:1012.4178 [hel-lat]).
- T. Blum, P.A. Boyle, N.H. Christ, N. Garron, E. Goode, T. Izubuchi, C. Jung, C. Kelly, C. Lehner, M. Lightman Q. Liu, A.T. Lytle, R.D. Mawhinney, C.T. Sachrajda, A. Soni, C. Sturm, *The* $K \to (\pi\pi)_{I=2}$ *Decay Amplitude from Lattice QCD*, (arXiv:1111.1699 [hep-lat])
- Ruud A Haring, Martin Ohmacht, Thomas W Fox, Michael K Gschwind, Peter A Boyle, Norman H Christ, Changhoan Kim, David L Satterfield, Krishnan Sugavanam, Paul W Coteus, Philip Heidelberger, Matthias A Blumrich, Robert W Wisniewski, Alan Gara, George L Chiu, *The IBM Blue Gene/Q Compute Chip*, MicroSI-2011-10-0058, IEEE Micro, to appear.

BIOGRAPHICAL SKETCH December 2011

Carleton DeTar	Telephone: (801)581-7537
	Department of Physics and Astronomy
	115 S 1400 E Rm 201
Citizenship: USA	Salt Lake City, Utah 84112
	e-mail: `detar@physics.utah.edu`

Education:

Undergraduate Studies:	A.B. Chemistry and Physics, 1966, Harvard College
Graduate Studies:	Ph.D. Physics, 1970, University of California, Berkeley
Postdoctoral Research:	Massachusetts Institute of Technology, 1970-1972

Employment:

1985-	Professor, Department of Physics, University of Utah
1998-2005	Associate Chair, Department of Physics, University of Utah
1983-1989	Associate Chair, Department of Physics, University of Utah
1978-1985	Associate Professor (Utah)
1972-1978	Assistant Professor (MIT)

Relevant Publications:

"First Determination of the Strange and Light Quark Masses from Full QCD", (with the MILC, HPQCD, and UKQCD Collaboration), Phys. Rev. D **70**, 031504 (2004).

"Charmed meson decay constants in three-flavor lattice QCD," (with the Fermilab Lattice, MILC, and HPQCD Collaborations) Phys. Rev. Lett. **95**, 122002 (2005).

"Topological Susceptibility in Staggered Chiral Perturbation Theory" (with Brian Billeter and James Osborn) Phys. Rev. D **71**, 034504 (2005).

"QCD equation of state with 2+1 flavors of improved staggered quarks", (with the MILC collaboration), Phys. Rev. D **75**, 094505 (2007).

"QCD Thermodynamics from the Lattice," (with U.M. Heller), Eur. J. Phys. A **41**,405 (2009).

Other Publications:

"A Conjecture Concerning the Modes of Excitation of the Quark-Gluon Plasma", Phys. Rev. D **32**, 276 (1985).

"The Hadronic Spectrum of the Quark Plasma", (with J.B. Kogut) Phys. Rev. Lett. 59, 399 (1987).

"High Precision Lattice QCD Confronts Experiment", (with the Fermilab, HPQCD, MILC, and UKQCD collaborations) Phys. Rev. Lett. **91**, 022001 (2004).

"Lattice Methods for Quantum Chromodynamics", T. DeGrand and C. DeTar, (World Scientific, Singapore, 2006).

"Full nonperturbative QCD simulations with 2+1 flavors of improved staggered quarks," (with the MILC collaboration) Rev. Mod. Phys. **82**, 1349 (2010).

For a current bibliography please see:

`http://www.physics.utah.edu/~detar/biblio.html`

Collaborators:

Claude Bernard (Washington Univ.)
Aida El Khadra (Univ. of Illinois)
Rajan Gupta (Los Alamos National Laboratory)
James Hetrick (University of the Pacific)
Andreas Kronfeld (Fermilab)
Paul Mackenzie (Fermilab)
Robert Sugar (Univ. of Calif., Santa Barbara)
Ruth Van de Water (Brookhaven National Laboratory)
Richard Brower (Boston University)
Steven Gottlieb (Indiana Univ.)
Urs Heller (American Physical Society)
Frithjof Karsch (Brookhaven National Laboratory)
John Laiho (Glasgow University)
James Osborn (Argonne National Laboratory)
Douglas Toussaint (Univ. of Arizona)

Recent Supported Graduate Students and Postdocs:

Tommy Burch (Regensburg)
Justin Foley (Utah)
Bugra Oktay (Utah)
Ludmila Levkova (Utah)

Steven A. Gottlieb

Education:

A.B.; Mathematics and Physics, Cornell University, 1973 *summa cum laude*, with Distinction in All Subjects
Ph.D.; Physics, Princeton University, 1978

Appointments:

Distinguished Professor of Physics, Indiana University, 2008–
Visiting Research Associate, NCSA, University of Illinois, 2009–2010
Professor of Physics, Indiana University, 1992–2008
Frontier Fellow, Fermi National Accelerator Laboratory, 2001–2002
Visiting Physicist, Brookhaven National Laboratory, 1992–1993
Associate Professor of Physics, Indiana University, 1988–1992
Assistant Professor of Physics, Indiana University, 1985–1998
Assistant Research Physicist, UCSD, 1982–1985
Research Associate, Fermi National Accelerator Laboratory, 1980–1982
Postdoctoral Appointee, Argonne National Laboratory, 1978–1980

Five Relevant Publications:

- *Scaling Lattice QCD beyond 100 GPUs*, R. Babich, M.A. Clark, J. Joó, G. Shi, R.C. Brower, S. Gottlieb, Proceeding of the 1011 ACM/IEEE International Conference for High Performance Computing, Network, Storage and Analysis.

  Design of MILC lattice QCD application for GPU clusters, G. Shi, S. Gottlieb, A. Torok, V. Kindratenko, Proc. 2011 IEEE Parallel and Distributed Processing Symposium, (IPDPS '11), Anchorage, AK, May, 2011.

  Acclerating Quantum Chromodynamics Calculations with GPUs, G. Shi, S. Gottlieb, A. Torok, V. Kindratenko, Proc. Symp. on Applications Accelerators in HPC (SAAHPC '10), Knoxville, TN, July 2010.

  QUDA programming for staggered quarks, S Gottlieb, G. Shi, A. Torok, V. Kindratenko, Proceedings of Lattice 2010, Villasimius, Italy, June 2010, Proceedings of Science (Lattice 2010) 026.

  Hybrid-Molecular Dynamics Algorithms for the Numerical Simulation of Quantum Chromodynamics, S. Gottlieb, W. Liu, D. Toussaint, R.L. Renken, and R.L. Sugar, Phys. Rev. D **35**, 2531 (1987).

Five Other Significant Publications:

The QCD spectrum with three quark flavors, MILC Collaboration: C. Bernard, T. Burch, T. A. DeGrand, S. Datta, C. DeTar, S. Gottlieb, U. M. Heller, K. Orginos, R. Sugar, D. Toussaint, Phys.Rev. **D64** (2001) 054506.

High-Precision Lattice QCD Confronts Experiment, The Fermilab Lattice, HPQCD, MILC, and UKQCD Collaborations: C.T.H. Davies, E. Follana, A. Gray, G.P. Lepage, Q. Mason, M. Nobes, J. Shigemitsu, H.D. Trottier, M. Wingate, C. Aubin, C. Bernard, T. Burch, C. DeTar, S. Gottlieb, E.B. Gregory, U.M. Heller, J.E. Hetrick, J. Osborn, R. Sugar, D. Toussaint, M. Di Pierro, A. El-Khadra, A.S. Kronfeld, P.B. Mackenzie, D. Menscher, J. Simone, Phys. Rev. Lett. **92**, 022001 (2004) [arXiv:hep-lat/0304004].

Light pseudoscalar decay constants, quark masses, and low energy constants from three-flavor lattice QCD, The MILC Collaboration: C. Aubin, C. Bernard, C. DeTar, S. Gottlieb, E.B. Gregory, U.M. Heller, J.E. Hetrick, J. Osborn, R. Sugar, D. Toussaint, Phys. Rev. D **70** (2004) 114501 [arXiv:hep-lat/0407028].

Curriculum Vitae
FRITHJOF KARSCH

Contact Information

Department of Physics — phone: (631) 344-8015
Brookhaven National Laboratory — fax: (631) 344-5519
Upton, NY 11973 — email: karsch@bnl.gov

Professional Preparation

Dr. rer. nat. Physics, Bielefeld University, Germany, 1982 (Advisor, H. Satz)
Diplom Physics, Bielefeld University, Germany, 1979 (Advisor, H. Satz)

Appointments

2005–present, Sr. Scientist, Brookhaven National Laboratory, USA
1990–present, Professor for Theoretical Physics, Bielefeld University, Germany
1990–1993, Head of the particle physics group at HLRZ, HLRZ-Jülich, Germany
1986–1990, Staff Member, Theory Division, CERN, Geneva, Switzerland
1984–1986, Research Associate, University of Illinois at Urbana-Champaign, USA
1982–1984, Research Fellow, CERN, Geneva, Switzerland
Visiting Positions: Visiting Scientist, University of Tsukuba, Tsukuba, Japan, May-July 1999; CERN, Geneva, Switzerland, August 1999 - April 2000

Synergetic Activities

Member, Editorial Board of *The European Physical Journal C*
Member, USQCD Executive Committee, 2011–present
Member, USQCD Scientific Program Committee, 2009–present; chair, 2010-present
Member, Advisory Board, New York Center for Comp. Science, 2007–present

International Advisory Committees:
a) International Symposia on Lattice Field Theory: Lattice 2008, William and Mary, USA; Lattice 2009, Beijing, China; Lattice 2010, Villasimius, Italy; Lattice 2011, Squaw Valley, USA
b) International Conference on Ultrarelativistic Nucleus-Nucleus Collisions: Quark Matter 2008, Knoxville, USA; Quark Matter 2011, Annecy, France; Quark Matter 2012, Washington DC
c) International Conference on 'Critical Point and Onset of Deconfinement': CPOD 2009, BNL, USA, CPOD 2010, Dubna, Russia; CPOD 2011, Wuhan, China; CPOD 2013, Berkeley, USA.

Organizing Committee:
Lattice 06, Tucson, 2006; Lattice 94, Bielefeld 1994; Strong and Electroweak Matter 2006, BNL 2006; Critical Point and Onset of Deconfinement 2009, BNL 2009; Extreme QCD 10, Bad Honnef, Germany 2010

Coordination of research programs:
Coordinator of the interdisciplinary research project *Multiscale Phenomena and their Simulation on Massively Parallel Computers*, Center for Interdisciplinary Research, Bielefeld University Germany, 8/1996-12/1998
Coordinator of the European research network *Finite Temperature Phase Transitions in Particle Physics*, 12/1997 - 12/2001

Graduate and Postdoctoral Advisor: Prof. H. Satz, Bielefeld University

Graduate and Postdoctoral Advisees: P. Hegde (BNL), S. Datta (TIFR), H.-T. Ding (BNL), C. Schmidt (Bielefeld), W. Söldner (Regensburg),

Selected Publications: Frithjof Karsch

202 publications in refereed journals; total of 13500 citations (SPIRES)

1.) A. Bazavov, T. Bhattacharya, M. Cheng, N. H. Christ, C. DeTar, S. Ejiri, S. Gottlieb and R. Gupta, U.M. Heller, K. Huebner, C. Jung, F. Karsch, E. Laermann, L. Levkova, C. Miao, R.D. Mawhinney, P. Petreczky, C. Schmidt, R.A. Soltz, W. Soeldner, R. Sugar, D. Toussaint, P. Vranas [HotQCD Collaboration], Equation of state and QCD transition at finite temperature, Phys. Rev. D **80**, 014504 (2009) [arXiv:0903.4379 [hep-lat]] (252 citations).
2.) A. Bazavov, T. Bhattacharya, M. Cheng, C. DeTar, H. T. Ding, S. Gottlieb, R. Gupta and P. Hegde, U. M. Heller, F. Karsch, E. Laermann, L. Levkova, S. Mukherjee, P. Petreczky, C. Schmidt, R. A. Soltz, W. Soeldner, R. Sugar, D. Toussaint, W. Unger, P. Vranas [HotQCD Collaboration], The chiral and deconfinement aspects of the QCD transition, arXiv:1111.1710 [hep-lat] (4 citations).
3.) M. Cheng, N. H. Christ, M. Li, R. D. Mawhinney, D. Renfrew, P. Hegde, F. Karsch and M. Lin, R. D. Mawhinney, D. Renfrew, P. Vranas, The finite temperature QCD using 2 + 1 flavors of domain wall fermions at N(t) = 8, Phys. Rev. D **81**, 054510 (2010) [arXiv:0911.3450 [hep-lat]] (28 citations).
4.) C. R. Allton, S. Ejiri, S. J. Hands, O. Kaczmarek, F. Karsch, E. Laermann, C. Schmidt and L. Scorzato, The QCD thermal phase transition in the presence of a small chemical potential, Phys. Rev. D **66**, 074507 (2002) [hep-lat/0204010] (389 citations).
5.) C. R. Allton, M. Doring, S. Ejiri, S. J. Hands, O. Kaczmarek, F. Karsch, E. Laermann and K. Redlich, Thermodynamics of two flavor QCD to sixth order in quark chemical potential, Phys. Rev. D **71**, 054508 (2005) [hep-lat/0501030] (263 citations).
6.) M. Cheng, P. Hendge, C. Jung, F. Karsch, O. Kaczmarek, E. Laermann, R. D. Mawhinney and C. Miao, P. Petreczky, C. Schmidt, W. Soeldner, Baryon Number, Strangeness and Electric Charge Fluctuations in QCD at High Temperature, Phys. Rev. D **79**, 074505 (2009) [arXiv:0811.1006 [hep-lat]] (84 citations).
7.) O. Kaczmarek, F. Karsch, E. Laermann, C. Miao, S. Mukherjee, P. Petreczky, C. Schmidt and W. Soeldner, W. Unger, Phase boundary for the chiral transition in (2+1) -flavor QCD at small values of the chemical potential, Phys. Rev. D **83**, 014504 (2011) [arXiv:1011.3130 [hep-lat]] (26 citations).
8.) S. Datta, F. Karsch, P. Petreczky and I. Wetzorke, Behavior of charmonium systems after deconfinement, Phys. Rev. D **69**, 094507 (2004) [hep-lat/0312037] (300 citations)
9.) H. -T. Ding, A. Francis, O. Kaczmarek, F. Karsch, E. Laermann and W. Soeldner, Thermal dilepton rate and electrical conductivity: An analysis of vector current correlation functions in quenched lattice QCD, Phys. Rev. D **83**, 034504 (2011) [arXiv:1012.4963 [hep-lat]] (14 citations).
10.) G. Boyd, J. Engels, F. Karsch, E. Laermann, C. Legeland, M. Lutgemeier and B. Petersson, Thermodynamics of SU(3) lattice gauge theory, Nucl. Phys. B **469**, 419 (1996) [hep-lat/9602007] (636 citations)

Collaborators and Co-editors: A. Bazavov (BNL), P. Braun-Munzinger (GSI), T. Bhattacharya (Los Alamos), M. Cheng (Boston U.), N. H. Christ (Columbia U.), C. DeTar (Utah U.), H. -T. Ding (BNL), S. Ejiri (Niigata), J. Engels (Bielefeld U.), A. Francis (Mainz U.), B. Friman (GSI), R. Gupta (Los Alamos), P. Hegde (BNL), U.M. Heller (APS), K. Huebner, C. Jung (BNL), O. Kaczmarek (Bielefeld U.), D. Kharzeev (Stony Brook), M. Kitazawa (Osaka U.), E. Laermann (Bielefeld U.), L. Levkova (Utah U.), C. Miao (Mainz U.), R.D. Mawhinney (Columbia U.), S. Mukherjee (BNL), C. Pica Southern Denmark U.), P. Petreczky (BNL), K. Redlich (Wroclaw U.), H. Satz (Bielefeld U.), B.-J. Schaefer (Giessen U.), C. Schmidt (Bielefeld U.), S. Scheredin (Bielefeld U.), V. Skokov (BNL), R.A. Soltz (LLNL), W. Soeldner (Regensburg U.), R. Sugar (Santa Barbara), K. Tuchin (Iowa State U.), D. Toussaint (Tucson), W. Unger (ETH Zürich), M. Wagner (Bielefeld U.), J. Wambach (Darmstadt U.)

Curriculum Vitae

JULIUS KUTI

Contact Information

Department of Physics, UC San Diego
9500 Gilman Drive, La Jolla, CA 92093-0319
phone:(858)534-6096; fax: (858)-534-0173; e-mail:jkuti@ucsd.edu

Professional Preparation

Ph.D. Physics, Eötvös University, Budapest, 1967
B.S. Physics, Eötvös University, Budapest, Summa Cum Laude, 1963

Appointments

2000–present, Distinguished Professor of Physics
1983-2000, Professor of Physics
1981-1983, Visiting Scientist, KITP, University of California, Santa Barbara
1974–1981, Associate Professor, Eötvös University, Budapest
1973–1974, Visiting Associate Prof., Massachusetts Institute of Technology
1970–1972, Postdoctoral Fellow, Massachusetts Institute of Technology
1968-1970, Assistant researcher, Eötvös University, Budapest

Honors and Awards

1993 Fellow, American Physical Society
1990 Elected Member of Hungarian National Academy
1975 Hungarian State Prize in Science
1972 Novobatzky Prize in Theoretical Physics
1966 Ettore Majorana Fellowship
1965 Gottfried von Herder Fellowship, Vienna University

Outside Activities

Member, USQCD Executive Committee, 2011-present
Co-Director of UCSD Computational Science, Mathematics, and Engineering, 2005-present
UCSD Committee of Academic Personnel, 1990-1992; vice-chair, 1991; chair, 1992
Advisory Committees of international conferences
Associate editor of Phys. Rev. Letters – two year period
Various Department of Energy review panels

Collaborators (*past 5 years*)

Zoltan Fodor (University of Wuppertal)
Kieran Holland (University of the Pacific)
Daniel Nogradi (Eötvös University)
Chris Schroeder (University of Wuppertal)
Ricky Wong(UCSD)

Selected Publications

- *Twelve massless flavors and three colors below the conformal window*
  (with Z. Fodor, K. Holland, D. Nogradi and C. Schroeder), Phys. Lett. B703, 348 (2011).

- *Chiral symmetry breaking in fundamental and sextet fermion representations of SU(3) color*
  (with Z. Fodor, K. Holland, D. Nogradi and C. Schroeder), PoS(Lattice 2010) 060.

- *Fine Structure of the String Spectrum*
  (with K.J. Juge and C.J. Morningstar), Phys. Rev. Letters, Volume 90, 161601 (2003).

- *Ab initio study of hybrid $\bar{b}gb$ mesons*
  (with K.J. Juge and C.J. Morningstar), Phys. Rev. Letters, Volume 82, 4400 (1999).

- *The Equivalence of the Top Quark Condensate and the Elementary Higgs Field*
  (with A. Hasenfratz, P. Hasenfratz, K. Jansen, and Y. Shen), Nuc. Phys. B **79**, 79 (1991).

- *Upper Bound on the Higgs-Boson Mass in the Standard Model*
  (with L. Lin and Y. Shen), Phys. Rev. Lett. 61, 678 (1988).

- *The Deconfining Phase Transition and the Continuum Limit of Lattice Quantum Chromodynamics*
  (with S. Gottlieb, D. Toussaint, A. Kennedy, S. Meyer, B. Pendleton, and R. Sugar), Phys. Rev. Lett. 55, 1958-1961 (1985).

- *Internal Spin Structure of the Nucleon*
  (with V. Hughes), Ann. Rev. Nucl. Part. Sci. 33, 611-644 (1983).

- *Monte Carlo Study of SU(2) Gauge Theory at Finite Temperature*
  (with J. Polonyi and K. Szlachanyi), Phys. Lett. B98, 199 (1981).

- *Inelastic Lepton-Nucleon Scattering and Lepton Pair Productions in the Relativistic Quark-Parton Model*
  (with V.F. Weisskopf), Phys. Rev. D4, 3418-3439 (1971).

Light hadrons with improved staggered quarks: approaching the continuum limit, The MILC Collaboration: C. Aubin, C. Bernard, T. Burch, C. DeTar, S. Gottlieb, E.B. Gregory, U.M. Heller, J. Osborn, R. Sugar, D. Toussaint, Phys. Rev. D **70** (2004) 094505 [arXiv:hep-lat/0402030].

Equation of state and QCD transition at finite temperature, The HotQCD Collaboration: A. Bazavov, T. Bhattacharya, M. Cheng, N.H. Christ, C. DeTar, S. Ejiri, S. Gottlieb, R. Gupta, U.M. Heller, K. Huebner, C. Jung, F. Karsch, E. Laermann, L. Levkova, C. Miao, R.D. Mawhinney, P. Petreczky, C. Schmidt, R.A. Soltz, W. Soeldner, R. Sugar, D. Toussaint, P. Vranas, Phys. Rev. D **80** (2009) 014504 [arXiv:0903.4379 [hep-lat]].

Synergistic Activities:

Chair, Lattice QCD Oversight Committee, 1999–2006
Divisional Associate Editor, Physical Review Letters, 2003–2009
Associate Editor-in Chief, Computing in Science & Engineering, 2007–
Roadrunner Phase Three Assessment, Los Alamos, 2007
International Advisory Committee, *Lattice '04, '03, '95, '94, '91*

Collaborators and Other Affiliations:

Collaborators (Last 48 Months): R. Babich (NVIDIA), J. Bailey (U. Seoul), S. Basak (NISER), G. Bauer (UIUC), A. Bazavov (BNL), C. Bernard (Washington U.), T. Bhattacharya (LANL), R. Brower (Boston U.), T. Burch (Regensburg), M. Cheng (LBNL), N.H. Christ (Columbia), M.A. Clark (NVIDIA), C. Davies (Glasgow), C. DeTar (Utah), M. Di Pierro (DePaul), S. Ejiri (APS), R.T. Evans (Regensburg), E. Freeland (DuPage), W. Freeman (GWU), A. El Khadra (UIUC), Z. Fu (Utah), E. Gamiz FNAL), R. Gupta (LANL), U. Heller (APS), J. Hetrick (Pacific), D. Holmgren (FNAL), T. Hoefler (UIUC), K. Huebner (BNL), R. Jain (UIUC) B. Joo (JLab), C. Jung (BNL), F. Karsch (BNL), V. Kindratenko (UIUC), A. Kronfeld (FNAL), E. Laermann (Bielefeld), J. Laiho (Glasgow), L. Levkova (Utah), P. Mackenzie (FNAL), C. Miao (BNL), R.D. Mawhinney (Colubmia), M. Oktay (Utah), J. Osborn (Argonne), p. Petrecsky (BNL), S. Prelovsek (Ljubljana), D. Renner (DESY), C. Schmidt (Bielefeld, G. Shi (UIUC), W. Soeldener (Darmstadt), R. Soltz (LLNL), R. Sugar (Santa Barbara), J. Simone (FNAL), D. Toussaint (Arizona), A. Torok (Indiana), P. Vranas (LLNL), R. Van de Water (FNAL), R. Zhou (Indiana)

Graduate Advisor: D. J. Gross

Thesis Advisor and Postgraduate-Scholar Sponsor: S. Tamhankar (U. Washington), A. Krasnitz (Cold Spring Harbor Laboratory), K. Rummukainen (Helsinki), L. Levkova (U. Utah), K. Nagata (Hokudai U.), S. Basak (NISER), H. Na (ANL), A. Torok (FEI Co.), R. Zhou (Indiana)

Biographical Sketch of James C. Osborn

Education and Training

University of Virginia	Math/Physics	B.A./B.S. 1994
Stony Brook University	Physics	Ph.D. 1999
Duke University	Theoretical Nuclear & Particle Physics	1999-2001
University of Utah	High Energy Theory	2001-2004
Boston University	Particle Theory	2004-2007

Research and Professional Experience

Argonne National Laboratory	Computational Scientist	2010-
Argonne National Laboratory	Assistant Computational Scientist	2007-2010

Publications

- T. Appelquist, R. Babich, R.C. Brower, M. Cheng, M.A. Clark, S.D. Cohen, G.T. Fleming, J. Kiskis, M. Lin, E.T. Neil, J.C. Osborn, C. Rebbi, D. Schaich and P. Vranas, "Parity Doubling and the S Parameter Below the Conformal Window," Phys. Rev. Lett. 106, 231601 (2011).
- J.C. Osborn, R. Babich, J. Brannick, R.C. Brower, M.A. Clark, S.D. Cohen and C. Rebbi, "Multigrid solver for clover fermions," PoS Lattice2010, 037 (2010).
- R. Babich, J. Brannick, R.C. Brower, M.A. Clark, T.A. Manteuffel, S.F. McCormick, J.C. Osborn and C. Rebbi, "Adaptive multigrid algorithm for the lattice Wilson-Dirac operator," Phys. Rev. Lett. 105, 201602 (2010).
- A. Bazavov, C. Bernard, C. DeTar, W. Freeman, S. Gottlieb, U.M. Heller, J.E. Hetrick, J. Laiho, L. Levkova, M. Oktay, J. Osborn, R.L. Sugar, D. Toussaint and R.S. Van de Water, "Scaling studies of QCD with the dynamical HISQ action," Phys. Rev. D82, 074501 (2010).
- A. Bazavov, C. Bernard, B. Billeter, C. DeTar, S. Gottlieb, U.M. Heller, J.E. Hetrick, J. Laiho, L. Levkova, M.B. Oktay, J. Osborn, R.L. Sugar, D. Toussaint and R.S. Van de Water, "Topological susceptibility with the asqtad action," Phys. Rev. D81, 114501 (2010).
- T. Appelquist, A. Avakian, R. Babich, R.C. Brower, M. Cheng, M.A. Clark, S.D. Cohen, G.T. Fleming, J. Kiskis, E.T. Neil, J.C. Osborn, C. Rebbi, D. Schaich and P. Vranas, "Toward TeV Conformality," Phys. Rev. Lett. 104, 071601 (2010).
- J.C. Osborn, "Initial guesses for multi-shift solvers," PoS Lattice2008, 029 (2008).
- J. Brannick, R.C. Brower, M.A. Clark, J.C. Osborn and C. Rebbi, "Adaptive Multigrid Algorithm for Lattice QCD," Phys. Rev. Lett. 100, 041601 (2008).

- C. Aubin, C. Bernard, C. DeTar, M. Di Pierro, A. El-Khadra, S. Gottlieb, E.B. Gregory, U.M. Heller, J. Hetrick, A.S. Kronfeld, P.B. Mackenzie, D. Menscher, M. Nobes, M. Okamoto, M.B. Oktay, J. Osborn, J. Simone, R. Sugar, D. Toussaint and H.D. Trottier, "Semileptonic decays of D mesons in three-flavor lattice QCD," Phys. Rev. Lett. 94, 011601 (2005).

- C.T.H. Davies, E. Follana, A. Gray, G.P. Lepage, Q. Mason, M. Nobes, J. Shigemitsu, H.D. Trottier, M. Wingate, C. Aubin, C. Bernard, T. Burch, C. DeTar, S. Gottlieb, E.B. Gregory, U.M. Heller, J.E. Hetrick, J. Osborn, R. Sugar, D. Toussaint, M. Di Pierro, A. El-Khadra, A.S. Kronfeld, P.B. Mackenzie, D. Menscher and J. Simone, "High-Precision Lattice QCD Confronts Experiment," Phys. Rev. Lett. 92, 022001 (2004).

Synergistic Activities

- Reviewed papers for PRL and PRD.
- Fellow in ANL/U. Chicago Computation Institute.
- Helped mentor 2 summer students on developing and running benchmark codes on the BG/P.
- Presented tutorial on parallel codes for Lattice QCD at 2007 SciDAC conference.

Collaborators and Co-Editors:

Gernot Akemann (Brunel U.) Tom Appelquist (Yale) Ronald Babich (NVIDIA) Alexei Bazavov (BNL) Claude Bernard (Washington U.) James Brannick (Pennsylvania State U.) Richard Brower (Boston U.) Michael Cheng (Boston U.) Michael Clark (NVIDIA) Saul Cohen (U. Washington) Poul Damgaard (Niels Bohr Institute) Christine Davies (U. Glasgow) Carleton DeTar (U. Utah) George Fleming (Yale) Antonio Garcia-Garcia (Cambridge) Steven Gottlieb (Indiana U.) Urs Heller (American Physical Society) James Hetrick (U. Pacific) Joe Kiskis (U.C. Davis) Jack Laiho (U. Glasgow) Ludmila Levkova (U. Utah) Meifeng Lin (Yale) Scott MacLachlan (Tufts U.) Tom Manteuffel (UC Boulder) Steve McCormick (UC Boulder) Ethan Neil (FNAL) Claudio Rebbi (Boston U.) Dru Renner (DESY Zeuthen) David Schaich (UC Boulder) Kim Splittorff (Niels Bohr Institute) Robert Sugar (UC Santa Barbara) Doug Toussaint (U. Arizona) Jacobus Verbaarschot (Stony Brook U.) Pavlos Vranas (LLNL)

Graduate and Postdoctoral Advisors and Advisees

Richard Brower (Boston U.) Shailesh Chandrasekharan (Duke U.) Carleton DeTar (U. Utah) Claudio Rebbi (Boston U.) Jacobus Verbaarschot (Stony Brook U.) Heechang Na (ANL)

Curriculum Vitae
STEPHEN R. SHARPE

Contact Information

Department of Physics | phone: (206) 685-2395
University of Washington | fax: (206) 685-0635
Seattle, WA 98195–1560 | email: sharpe@phys.washington.edu

Education and Training

Ph.D. Physics, University of California, Berkeley, 1983 (Advisor, Michael Chanowitz)
B.A. Theoretical Physics, Trinity College, Cambridge University, 1978

Research and Professional Experience

1995-present, Professor, University of Washington
1991-1995, Associate Professor, University of Washington
1988-1991, Assistant Professor, University of Washington
1986-1988, Five-year Research Associate, SLAC
1983-1986, Junior Fellow, Harvard Society of Fellows

Visiting Positions

Visiting Professor, University of Marseille, 2008
Visiting Professor, University of Southampton, 2004
Visiting Professor, University of Tsukuba, 1998
Visiting Professor, University of Rome, 1996
Visiting Staff Scientist, CEBAF (now Jefferson Lab), 1991-2

Honors and Awards

Fellow, American Physical Society, 1993
Alfred P. Sloan Foundation Fellow, 1990-1994
DOE Outstanding Junior Investigator, 1989-1991

Synergistic Activities

Member, USQCD Executive Committee, 1999-present; USQCD Scientific Program Committee, 2002-7
Organizer, INT summer school "Lattice QCD and its Applications," 2007
Panel chair, DOE workshop "Extreme Scale Computing,", SLAC, 2008
International Advisory Committees, International Symposia on Lattice Field Theory: Lattice 91, 92, 98, 01, 02, 03, 04, 05, 07 & 09

Collaborators in last 48 months

S. Aoki (Tsukuba), T. Bae (Seoul), O. Bär (Berlin), C. Bernard (Wash. U.), B. Bringoltz (IIAR, Israel), M. Golterman (SFSU), Max Hansen (UW), Y.-C. Jang (Seoul), C. Jung (BNL), Jangho Kim (Seoul), Jongjong Kim (Arizona), K. Kim (Seoul), H.-J. Kim (Seoul), M. Koren (Krakow), W. Lee (Seoul National), A. Lytle (Southampton), Y. Shamir (Tel Aviv), B. Yoon (Seoul)

Graduate Students and Postdoctoral Advisees in last 5 years
B. Bringoltz (ITAR), A. Lytle (Southampton), M. Hansen (UW)

Selected Publications

- M. T. Hansen and S. R. Sharpe, "Constraint on the Low Energy Constants of Wilson Chiral Perturbation Theory," arXiv:1111.2404 [hep-lat].
- T. Bae, Y. -C. Jang, C. Jung, H. -J. Kim, J. Kim, K. Kim, W. Lee, S. R. Sharpe and B. Yoon, "B_K using HYP-smeared staggered fermions in $N_f = 2+1$ unquenched QCD," Phys. Rev. D **82**, 114509 (2010) [arXiv:1008.5179 [hep-lat]].
- B. Bringoltz and S. R. Sharpe, "Non-perturbative volume-reduction of large-N QCD with adjoint fermions," Phys. Rev. D **80**, 065031 (2009) [arXiv:0906.3538 [hep-lat]].
- S. R. Sharpe, "Rooted staggered fermions: Good, bad or ugly?," PoS LAT **2006**, 022 (2006) [hep-lat/0610094].
- C. h. Kim, C. T. Sachrajda and S. R. Sharpe, "Finite-volume effects for two-hadron states in moving frames," Nucl. Phys. B **727**, 218 (2005) [hep-lat/0507006].
- S. R. Sharpe and N. Shoresh, "Physical results from unphysical simulations," Phys. Rev. D **62**, 094503 (2000) [hep-lat/0006017].
- W. -J. Lee and S. R. Sharpe, "Partial flavor symmetry restoration for chiral staggered fermions," Phys. Rev. D **60**, 114503 (1999) [hep-lat/9905023].
- S. R. Sharpe and R. L. Singleton, Jr, "Spontaneous flavor and parity breaking with Wilson fermions," Phys. Rev. D **58**, 074501 (1998) [hep-lat/9804028].
- S. R. Sharpe, "Quenched chiral logarithms," Phys. Rev. D **46**, 3146 (1992) [hep-lat/9205020].
- M. S. Chanowitz and S. R. Sharpe, "Hybrids: Mixed States of Quarks and Gluons," Nucl. Phys. B **222**, 211 (1983) [Erratum-ibid. B **228**, 588 (1983)].

Curriculum Vitae
Guochun Shi

Contact Information

3050D
1008 NCSA 1205 W.Clark
Urbana, IL 61801
phone: (217) 244-6950
email: gshi1@illinois.edu

Professional Preparation

M.C.S. Computer Science, University of Illinois, Urbana-Champaign, 2002
B.S. Physics, Peking University, Beijing, China, 1998

Appointments

2002-present, National Center for Supercomputing Applications (NCSA), UIUC, Research Programmer

Collaborators

S. Gottlieb, Indiana University
R. Brower, Boston University

Selected Publications

- R. Babich, M.A. Clark, B. Joó, G. Shi, R.C. Brower, S. Gottlieb, "Scaling Lattice QCD beyond 100 GPUs," In Proceedings of the 2011 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis (SC'11) [arXiv:1109.2935v1 [hep-lat]].
- F. Pratas, P Trancoso, L. Sousa, A. Stamatakis, G. Shi, V. Kindratenko, "Fine-grain Parallelism using Multi-core, Cell/BE, and GPU Systems," Parallel Computing, 2011.
- G. Shi, S. Gottlieb, A. Torok, V. Kindratenko, "Design of MILC lattice QCD application for GPU clusters," In Proceedings of International Parallel and Distributed Processing Symposium (IPDPS) (2011), IEEE, pp. 363371.
- S. Gottlieb, G. Shi, A. Torok, V. Kindratenko, "Quda programming for staggered quarks," In Proc.XXVIII International Symposium on Lattice Field Theory (Lattice 2010).
- J. Enos, C. Steffen, J. Fullop, M. Showerman, G. Shi, K. Esler, V. Kindratenko, J. Stone, J. Phillips, "Quantifying the Impact of GPUs on Performance and Energy Efficiency in HPC Clusters ," The Work in Progress in Green Computing(WIPGC), 2010
- J. Stone, D. Gohara, G. Shi, "OpenCL: A Parallel Programming Standard for Heterogeneous Computing Systems," Computing in Science and Engineering (CiSE), vol. 12, no. 3, May/June 2010.
- G. Shi, I. Ufimtsev, V. Kindratenko, T. Martinez "Direct Self-Consistent Field Computations on GPU Clusters," 24th IEEE International Parallel and Distributed Processing Symposium (IPDPS), 2010.
- V. Kindratenko, J. Enos, G. Shi, M. Showerman, G. Arnold, J. Stone, J. Phillips, W. Hwu, "GPU Clusters for High Performance Computing," Workshop on Parallel Programming on Accelerator Clusters (PPAC), IEEE Cluster 2009.
- G. Shi, V. Kindratenko, S.Gottlieb, "The bottom-up implementation of one MILC lattice QCD application on the Cell blade," International Journal of Parallel Programming (IJPP), 2009.
- G. Shi, V. Kindratenko, "Implementation of NAMD molecular dynamics non-bonded force-field on the Cell Broadband Engine processor," In Proc. 9th IEEE International Workshop on Parallel and Distributed Scientific and Engineering Computing (PDSEC'08), 2008

Curriculum Vitae

Robert L. Sugar

Education and Training:

A.B. Physics, Harvard University 1960, *summa cum laude*
Ph.D. Physics, Princeton University 1964
Postdoctoral Research in Theoretical Physics, Columbia University, 1964-66

Professional Experience:

Research Professor, University of California, Santa Barbara, 2004-present
Chair, Department of Physics, University of California, Santa Barbara, 1994-97
Deputy Director, Institute for Theoretical Physics, Santa Barbara, 1979–81; 1983–85.
Professor, University of California, Santa Barbara, 1973-2003
Associate Professor, University of California, Santa Barbara, 1969-73
Assistant Professor, University of California, Santa Barbara, 1966-69

Honors and Awards:

Phi Beta Kappa
Fellow of the American Physical Society
Fellow of the American Association for the Advancement of Science
Co-founder, Institute for Theoretical Physics, Santa Barbara

Selected Publications:

Variants of fattening and flavor symmetry restoration, Kostas Orginos, Doug Toussaint and R.L. Sugar, Phys. Rev. D **60**, 054503 (1999) [arXiv:hep-lat/9903032].

High-Precision Lattice QCD Confronts Experiment, The Fermilab Lattice, HPQCD, MILC, and UKQCD Collaborations: C. T. H. Davies, E. Follana, A. Gray, G. P. Lepage, Q. Mason, M. Nobes, J. Shigemitsu, H. D. Trottier, M. Wingate, C. Aubin, C. Bernard, T. Burch, C. DeTar, Steven Gottlieb, E. B. Gregory, U. M. Heller, J. E. Hetrick, J. Osborn, R. Sugar, D. Toussaint, M. Di Pierro, A. El-Khadra, A. S. Kronfeld, P. B. Mackenzie, D. Menscher, J. Simone, Phys. Rev. Lett. **92**, 022001 (2004) [arXiv:hep-lat/0304004].

Light hadrons with improved staggered quarks: approaching the continuum limit, The MILC Collaboration: C. Aubin, C. Bernard, T. Burch, C. DeTar, Steven Gottlieb, E.B. Gregory, U.M. Heller, J. Osborn, R. Sugar, D. Toussaint, Phys. Rev. D **70**, 094505 (2004) [arXiv:hep-lat/0402030].

First determination of the strange and light quark masses from full lattice QCD, The HPQCD, MILC and UKQCD Collaborations: C. Aubin, C. Bernard, C. Davies, C. DeTar, S. Gottlieb, A. Gray, E. Gregory, J. Hein, U.M. Heller, J. Hetrick, G. Lepage, O. Mason, J. Osborn, R. Sugar, D. Toussaint, H. Trottier M. Wingate, Phys. Rev. D **70** 031504(R) (2004) [hep-lat/0405022].

Semileptonic decays of D mesons in three-flavor lattice QCD, The Fermilab Lattice, MILC, and HPQCD Collaborations: C. Aubin, C. Bernard, C. DeTar, M. Di Pierro, A. El-Khadra, Steven Gottlieb, E. B. Gregory, U. M. Heller, J. Hetrick, A. S. Kronfeld, P. B. Mackenzie, D. Menscher, M. Nobes, M. Okamoto, M. B. Oktay, J. Osborn, J. Simone, R. Sugar, D. Toussaint, H. D. Trottier, Phys. Rev. Lett. **94**, 011601 (2005) [arXiv:hep-ph/0408306].

Charmed meson decay constants in three flavor lattice QCD, The Fermilab Lattice, MILC, and HPQCD Collaborations: C. Aubin, C. Bernard, C. DeTar, M. Di Pierro, E.D. Freeland, Steven Gottlieb, E.B. Gregory, U.M. Heller, J.E. Hetrick, A.X. El-Khadra, A.S. Kronfeld, L. Levkova, P.B. Mackenzie, F. Maresca, D. Menscher, M. Nobes, M. Okamoto, D. Renner, J.N. Simone, R.L. Sugar, D. Toussaint, H.D. Trottier, Phys. Rev. Lett. **95** 122002 (2005) [arXiv:hep-lat/0506030].

The $B \to D * l\nu$ form factor at zero recoil from three-flavor lattice QCD: A model independent determination of $|V_{cb}|$, The Fermilab Lattice and MILC Collaborations: C. Bernard, C. DeTar, M. Di Pierro, A.X. El-Khadra, R.T. Evans, E.D. Freeland, E. Gamiz, Steven Gottlieb, U.M. Heller, J.E. Hetrick, A.S. Kronfeld, J. Laiho, L. Levkova, P.B. Mackenzie, M. Okamoto, J. Simone, R. Sugar, D. Toussaint, R.S. Van de Water, Phys. Rev. **D79**, 014506 (2009) [arXiv:0808.2519 [hep-lat].

The $B \to \pi l\nu$ semileptonic form factor from three-flavor lattice QCD: A model-independent determination of $|V_{ub}|$, The Fermilab Lattice and MILC Collaborations: Jon A. Bailey, C. Bernard, C. DeTar, M. Di Pierro, A. X. El-Khadra, R. T. Evans, E. D. Freeland, E. Gamiz, Steven Gottlieb, U. M. Heller, J. E. Hetrick, A. S. Kronfeld, J. Laiho, L. Levkova, P. B. Mackenzie, M. Okamoto, J. N. Simone, R. Sugar, D. Toussaint, R. S. Van de Water, Phys. Rev. **D79**, 054507 (2009) [arXiv:0811.3640 [hep-lat]].

Full nonperturbative QCD simulations with 2+1 flavors of improved staggered quarks, The MILC Collaboration: A. Bazavov, C. Bernard, C. DeTar, Steven Gottlieb, U.M. Heller, J.E. Hetrick, J. Laiho, L. Levkova, P.B. Mackenzie, M.B. Oktay, R. Sugar, D. Toussaint, and R.S. Van de Water, Rev. Mod. Phys. **82**, 1349-1417 (2010) [arXiv:0903.3598 [hep-lat]].

Scaling studies of QCD with the dynamical HISQ action, MILC Collaboration: A. Bazavov, C. Bernard, C. DeTar, W. Freeman, Steven Gottlieb, U. M. Heller, J. E. Hetrick, J. Laiho, L. Levkova, M. Oktay, J. Osborn, R.L. Sugar, D. Toussaint, R.S. Van de Water, Phys. Rev. D**82** 074501 (2010) [arXiv:1004.0342]

Selected Synergistic Activities:

Member, USQCD Executive Committee, 1999-present, Chair 1999-2008
Member, SciDAC 2005 Organizing Committee, 2005
Co-Chair, NSF Computational Physics Steering Committee, 2001-2002
Member, NSF Advisory Committee for the Directorate of Computer and Information Science and Engineering, 1992-96.
Chair, User Advisory Committee, National Computational Science Alliance, 1997-2004

Collaborators (Last 48 months):

A. Bazavov (BNL), C. Bernard (Washington U.), R. Brower (BU), N. Christ (Columbia), M. Creutz (BNL), C. DeTar (Utah), M. Di Pierro (DePaul), E. Freeland (U. Illinois), A. El Khadra (Illinois), S. Gottlieb (Indiana), U. Heller (APS), J. Hetrick (Pacific), F. Karsch (BNL), A. Kronfeld (FNAL), J. Kuti (UCSD), J. Laiho (Glasgow), L. Levkova (Utah), P. Mackenzie (FNAL), J. Negele (MIT), J. Osborn (ANL), C. Rebbi (BU), D. Richards (JLab), S. Sharpe (U. Washington), J. Simone (FNAL), D. Toussaint(Arizona), R. Van de Water (BNL).

Co-Editors (Last 24 months): None.

Graduate Advisor: R. Blankenbecler (Stanford Linear Accelerator Center).

Postdoctoral Advisor: T.D. Lee (Columbia).

Thesis Advisor and Postgraduate-Scholar Sponsor:

My most recent graduate students were Christopher Martin, now an associate professor at Oberlin College, and Robert Sedgewick, now at Google-Pittsburgh. I have had a total of seven graduate students and sponsored over twenty postdoctoral scholars.

DOUG TOUSSAINT
Biographical Information
December 23, 2011

Education:

Univ. of North Carolina	Physics	BS 1974
Princeton University	Physics	Ph.D. 1978
Univ. of Calif. Santa Barbara	High Energy Physics	1978-1980
Institute for Theoretical Physics	High Energy Physics	1980-1983

Appointments:

July 1994 - present
Professor, Department of Physics
University of Arizona, Tucson, AZ 85721

May 1997 - Aug 1997
Visiting Professor, Center for Computational Physics
University of Tsukuba, Tsukuba, Ibaraki 305, Japan

Aug. 1988 - June 1994
Associate Professor, Department of Physics
University of Arizona, Tucson, AZ 85721

Jan. 1988 - May 1988
Visiting Scientist, Fermi National Accelerator Laboratory
PO Box 500, Batavia, IL 60510

July 1983 - June 1988
Assistant Professor, Physics Department, B-019
University of California at San Diego, La Jolla, CA 92093

Selected Publications:

1. S. Gottlieb, W. Liu, R. Renken, R. L. Sugar, and Doug Toussaint, "Hybrid Molecular Dynamics Algorithms for the Numerical Simulation of Quantum Chromodynamics," Phys. Rev. D **35**, 2531-2542 (1987).

2. Light pseudoscalar decay constants, quark masses, and low energy constants from three-flavor lattice QCD, MILC Collaboration: C. Aubin, C. Bernard, C. DeTar, Steven Gottlieb, E.B. Gregory, U.M. Heller, J.E. Hetrick, J. Osborn, R. Sugar, D. Toussaint, hep-lat/0407028 , Phys. Rev. D **70**, 114501 (2004).

3. Full nonperturbative QCD simulations with 2+1 flavors of improved staggered quarks, Authors: A. Bazavov, C. Bernard, C. DeTar, Steven Gottlieb, U.M. Heller, J.E. Hetrick, J. Laiho, L. Levkova, P.B. Mackenzie, M.B. Oktay, R. Sugar, D. Toussaint, R.S. Van de Water , arXiv:0903.3598, Rev. Mod. Phys. **82**, 1349-1417 (2010).

4. Scaling studies of QCD with the dynamical HISQ action, MILC collaboration: A. Bazavov, C. Bernard, C. DeTar, W. Freeman, Steven Gottlieb, U. M. Heller, J. E. Hetrick, J. Laiho, L. Levkova, M. Oktay, J. Osborn, R.L. Sugar, D. Toussaint, R.S. Van de Water, [arXiv:1004.0342], Phys. Rev. D **82**, 074501 (2010).

5. The strange quark condensate in the nucleon in 2+1 flavor QCD, D. Toussaint and W. Freeman [MILC collaboration], [arXiv:0905.2432], Phys. Rev. Lett. **103**, 2009, 122002.

6. B- and D-meson decay constants from three-flavor lattice QCD, A. Bazavov *et al.*, [arXiv:1112.3051].

Synergistic Activities

Local Organizer for Lattice-2006 conference
SciDAC Lattice Gauge Theory Science Committee Member

Collaborators and Co-editors (past 48 months)

C. Aubin (William and Mary)	J. Bailey (FNAL)	S. Basak (Orissa U)
A. Bazavov (Arizona)	C. Bernard (Washington U.)	T. Bhattacharya (LANL)
C.M. Bouchard (FNAL)	M. Cheng (LLNL)	C. Detar (Utah)
H.T. Ding (BNL)	M. Di Pierro (DePaul U.)	A. El Khadra (Illinois)
R.T. Evans (U. of Illinois)	E. Freeland (Benedictine Univ)	W. Freeman (George Wash. Univ.)
E. Gamiz (U. de Granada)	S. Gottlieb (Indiana)	R. Gupta (LANL)
P. Hegde (BNL)	J. Hetrick (U. Pacific)	U. Heller (APS)
R. Jain (U. of Illinois)	F. Karsch (BNL)	A. Kronfeld (Fermilab)
E. Laermann (U. Bielefeld)	J. Laiho (Washington U)	L. Levkova (Utah)
P. Mackenzie (Fermilab)	S. Mukherjee (BNL)	E.T. Neil (FNAL)
M. Oktay (Utah)	J. Osborn (ANL)	P. Petrezcky (BNL)
C. Schmidt (U. Bielefeld)	J. Simone (Fermilab)	W. Soeldner (BNL)
R. Soltz (LLNL)	R. Sugar (UC Santa Barbara)	W. Unger (BNL)
R. Van de Water (BNL)	P. Vranas (LLNL)	

Graduate and Postdoctoral Advisors

Thesis advisor	Franck Wilczek	MIT
Postdoctoral supervisor	Bob Sugar	UCSB

Students and Postdocs supervised, last 5 years

Master's degree supervisor	Steve Bildstein	Colorado Springs, CO
Thesis director	Walter Freeman	George Washington Univ.
Postdoctoral Supervisor	Dru Renner	JLAB
Postdoctoral Supervisor	Alexei Bazavov	Brookhaven Nat'l Lab
Postdoctoral Supervisor	Jongjeong Kim	Univ. of Arizona

Curriculum Vitae

PAVLOS M. VRANAS

Contact Information

Physical Sciences Directorate
Lawrence Livermore National Laboratory
7000 East Ave., L415
Livermore, CA 94550

phone: (925) 422-4681
fax: (925) 422-6594
email: vranasp@llnl.gov

Education and Training

1985 B.Sc. in Physics, University of Athens, Greece.
1987 M.Sc. in Theoretical Physics, University of California, Davis.
1990 Ph.D. in Theoretical Physics, University of California, Davis.
1990-1994, Supercomputer Research Institute, Post Doctoral Research Scientist
1994-1998, Columbia University, Post Doctoral Research Scientist
1998-2000, University of Illinois, Urbana-Champaign, Research Scientist

Appointments

2007-present, LLNL, Staff Physicist, Lattice Group Leader
2000-2007, IBM, T.J. Watson Research Laboratory, Senior Research Scientist

Honors and Awards

Gordon Bell Special Achievement Award 2006 for QCD on BG/L
Gordon Bell Award 1998 for the QCDSP supercomputer
LLNL PAT Directorate Award 2007 for Lattice QCD on the BlueGene
IBM Outstanding Innovation Award 2006

Synergistic Activities

Lattice 2011 International Conference Chair
INT Extreme Computing and NP, organizer, 2011
Organizer "Lattice Gauge Theory for LHC Physics" workshops 2008, 2009, 2011
USQCD software executive committee member
Core team architect and designer of the IBM BlueGene/L/P/Q supercomputers
Lead architect and designer of the Columbia Physics System (CPS)

Collaborators

M. Buchoff (LLNL), T. Appelquist (Yale), R.C. Brower (BU), M. Cheng (BU), N.H. Christ (Columbia), C. DeTar (UoU), G.T. Fleming (Yale), M. Golterman (SFSU), J. Greensite (SFSU), S. Gottlieb (IU), R. Gupta (LANL), W. Haxton (UCB), U.M. Heller (APS), P. Hegde (BNL), K. Holland (UoP), K. Juge (UoP), C. Jung (BNL), F. Karsch (BNL), J. Kiskis (UCD), J. Kuti (UCSD), L. Levkova (UoU), M. Li (Columbia), M.F. Lin (Yale), T. Luu (LLNL), R. D. Mawhinney (Columbia), C. Miao (BNL), E.T. Neil (FNAL), J.C. Osborn (ANL), P. Petreczky (BNL), C. Rebbi (BU), D. Renfrew (Columbia), D. Schaich (BU), R. (LLNL), B.

Sugar (UCSB), S. Syritsyn (LBNL), T. Tomboulis (UCLA), D. Toussaint (UoA), A. Walker-Loud (LBNL), J. Wasem (LLNL)

Selected Publications

- T. Appelquist, R. Babich, R.C. Brower, M. Cheng, M.A. Clark, S.D. Cohen, G.T. Fleming, J. Kiskis, M. Lin, E.T. Neil, J.C. Osborn, C. Rebbi, D. Schaich, P. Vranas, "Parity Doubling and the S Parameter Below the Conformal Window", Phys. Rev. Lett. **106**, 231601, (2011), [arXiv:1009.5967]
- T. Appelquist, A. Avakian, R. Babich, R.C. Brower, M. Cheng, M.A. Clark, S.D. Cohen, G.T. Fleming, J. Kiskis, M. Lin, E.T. Neil, J.C. Osborn, C. Rebbi, D. Schaich, P. Vranas, "Toward TeV Conformality", Phys. Rev. Lett. **104**, 071601, (2010) [arXiv:0910.2224]
- M. Cheng, N. H. Christ, P. Hegde, F. Karsch, Min Li, M. F. Lin, R. D. Mawhinney, D. Renfrew, P. Vranas, "The finite temperature QCD using 2+1 flavors of domain wall fermions at $N_t = 8$", Phys. Rev. **D81** 054510 (2010), [arXiv:0911.3450]
- A. Bazavov, T. Bhattacharya, M. Cheng, N.H. Christ, C. DeTar, S. Ejiri, S. Gottlieb, R. Gupta, U.M. Heller, K. Huebner, C. Jung, F. Karsch, E. Laermann, L. Levkova, C. Miao, R.D. Mawhinney, P. Petreczky, C. Schmidt, R.A. Soltz, W. Soeldner, R. Sugar, D. Toussaint, P. Vranas, "Equation of state and QCD transition at finite temperature", Phys. Rev. **D80** 014504 (2009), [arXiv:0903.4379]
- J. Giedt, R. Brower, S. Catterall, G.T. Fleming, P. Vranas, "Lattice super-Yang-Mills using domain wall fermions in the chiral limit", Phys.Rev. **D79** 025015 (2009), [arXiv:0810.5746]
- M. Blumrich, D. Chen, A. Gara, M. Giampap, P. Heidelberger, V. Salapura, J. Sexton, R. Soltz, P. Vranas, "Massively parallel QCD", IBM Journal of R&D **Vol. 52, 1/2** (2008)
- T. Luu, R. Soltz and P. Vranas, "Simulating the birth of the universe on a Petaflop", Computing in Science and Engineering, November 2007
- P.M. Vranas "Gap Domain Wall Fermions", Phys. Rev. **D74** 034512 (2006), [arXiv:hep-lat/0606014]
- A. Gara M. A. Blumrich D. Chen G. L.-T. Chiu P. Coteus M. E. Giampapa R. A. Haring P. Heidelberger D. Hoenicke G. V. Kopcsay T. A. Liebsch M. Ohmacht B. D. Steinmacher-Burow T. Takken P. Vranas, "Overview of the Blue Gene/L system architecture", IBM Journal of research and development, **Vol 49, 2/3** (2005)
- G. Fleming, J. Kogut, P. Vranas, "Super Yang-Mills on the lattice with domain wall fermions", Phys. Rev. **D64** 034510 (2001), [arXiv:hep-lat/0008009]

A.8 Facilities and Resources

The USQCD Collaboration operates its own dedicated clusters for lattice QCD at its partner laboratories, Fermilab and JLab, and will have access to dedicated hardware at its third partner laboratory, Brookhaven. At Fermilab, it operates 1277 cluster nodes with 20,320 total cores, and it will soon have 76 GPU-accelerated cluster nodes with 152 total GPUs (all NVIDIA Tesla M2050). It has 540 TBytes of dedicated disk storage, and access to Fermilab's robotic tape mass storage system. At JLab, it operates 544 dual-socket quad-core Xeon nodes with 4352 Intel cores, as well as 500 NVIDIA Fermi GPUs. It has dedicated disk capacity of 300 TBytes for Lattice QCD, and access to JLab's tape mass storage system. At BNL, it will have dedicated access to one of BNL's three racks of Blue Gene/Q hardware. USQCD will use this access for software development to prepare for upcoming large Blue Gene/Q installations at the Argonne Leadership Computing Facility and Lawrence Livermore National Laboratory.

USQCD also has significant grants at the DOE's leadership computing centers, 50 M core-hours at the Argonne Leadership Computing Facility and 46 M core-hours at the Oak Ridge Leadership Computing Facility for the current year. In addition, the partner institutions taking part in this proposal also have access to significant local resources which are described in their respective sections in the appendix on Budgets and Statements of Work.