SciDAC3 Project Title: A MultiScale Approach to Nuclear Structure and Reactions: Forming the Computational Bridge between Lattice QCD and Nonrelativistic Many-Body Theory ("CalLAT")

Principal Investigator:	Wick Haxton, Nuclear Science Division, LBNL
Co-Director for Science:	Tom Luu, Physical Science Directorate, LLNL
Co-Director for Computing:	Esmond Ng, Computational Research Division, LBNL

Co-PIs and Institutions:

Michael Clark	NVIDIA
Bronis R. de Supinski and Robert Falgout	Center for Applied Scientific Computing, LLNL
Juan Meza, Samuel Williams, and Chao Yang	Computational Research Division, LBNL
Pavlos Vranas	Physical Sciences Directorate, LLNL
André Walker-Loud	Nuclear Science Division, LBNL

Scope of Work

The project's goal is to develop procedures by which a non-relativistic effective theory of nuclear physics can be linked to the exact theory of the strong interaction, quantum chromodynamics (QCD), by connecting the low-energy constants of the former to lattice QCD calculations of nucleon-nucleon (NN) scattering parameters. The envisioned first applications are to tests of symmetries in the standard model, specifically the strangeness-conserving hadronic weak interaction and nucleon and nuclear electric dipole moments.

The effective theory (HOBET) is formulated in a discrete harmonic oscillator basis, which will allow us to exploit numerical tools for diagonalizing large sparse matrices. This numerical step addresses the mid-scale nuclear physics problem, and numerically provides adequate scale separation between omitted short-distance (ultraviolet) and long-distance (infrared) physics to ensure HOBET's rapid convergence. The approach is designed to exploit the multi-scale character of the nuclear physics problem, dividing it into three regions that can be attacked with distinct physical and numerical techniques. It also circumvents the standard effective interactions problem, as the appropriate soft interaction is constructed directly from data (experiment or LQCD input), with no need for a high-momentum NN potential. Because HOBET is formulated in a large but discrete basis of antisymmetric states, the fermion sign problem that would arise in direct LQCD calculations of few-nucleon systems is absent. Instead, LQCD results from two- and three-nucleon calculations can be related to short-range operators in HOBET, with HOBET then faithfully embedding that physics in more complicated systems. The scheme depends on one numerically testable assumption, convergence in the number of nucleons that are allowed to interact at short range at one time.

The LQCD work will initially focus on calculations of two-nucleon (NN) phase shifts needed to constrain NN effective range parameters that we then relate to the many-nucleon Hamiltonian. This research aspect augments, as well as benefits from, other SciDAC efforts, such as USQCD and NUCLEI, and will provide an opportunity for CalLAT to impact the broader computational nuclear physics program.

The target applications are to parity nonconservation (PNC) in the hadronic weak interaction and to neutron and nuclear electric dipole moments. The PNC work builds on a recent result from our group, the first calculation of the parity-violating pion-nucleon coupling h_π^1. One project goal is to improve the result in calculations with nearly physical pion masses and including disconnected diagrams. A second goal is to calculate the isotensor coupling h_ρ^2, as the absence of a reliable estimate of this contribution reduces the impact of the field's best measurement, the $\vec{p}+p$ asymmetry.

Work on electric dipole moments will focus on the the QCD theta parameter, and specifically the scalar pion-nucleon coupling this standard-model coupling induces. A lattice QCD calculation of this coupling provides the starting-point for interpreting the limit on the edm of ^{199}Hg, the field's most sensitive result.

Both applications are important to experimental programs supported by the Office of Nuclear Physics. Our LQCD calculations of NN PNC couplings will provide essential information for the global analysis of PNC that will be possible once the NPDGamma experiment finishes in 2013. In addition, HOBET will allow us to relate these results to those from light nuclei. Our efforts on CP violation will better define the potential sensitivities of the edm searches proposed by the nedm, Storage Ring edm, and Hg collaborations.

Mathematics, Computer Science, and Data/Visualization Issues

The project will be carried out as a partnership between LBNL, LLNL, and NVIDIA physicists and their applied mathematics/computer science colleagues who have developed relevant tools under the FASTMath and SUPER Institutes. The project's computational needs, such as LQCD calculations of NN effective range parameters for nearly physical pion masses, can be addressed with the 10-20 Pflop Office of Science and NNSA computing facilities that are coming available in 2012. The collaboration has significant hardware and software experience on leadership-class machines, as well as expertise with large GPU clusters and other heterogeneous architectures that can be efficiently utilized in specific phases of LQCD calculations. (The initial work on h_π^1 required approximately six months of nearly dedicated running on LLNL's Edge GPU cluster.)

The HOBET diagonalization stage for the application to PNC will require Lanczos methods for sparse matrices of dimension up to $\sim 1.4 \times 10^{10}$ with complex Hamiltonians requiring, in their initial stage, ~ 65 TB of storage, and ultimately (four-body operators) 200 PB of storage. The initial-stage calculations are about a factor of two beyond a no-core shell model calculation for ^{12}C recently performed on 100,000 cores on JaguarPF, the largest nuclear calculation of this type so far attempted. The final calculations we envision are well beyond state-of-the-art, but could likely be done on platforms in the 20-100 Pflop range. Ng and Yang have extensive experience integrating PARPACK with solvers designed for nuclear physics problems because of their involvement in the UNEDF collaboration. HOBET could benefit from a second, novel application of Lanczos methods to the kinetic energy operator, in which moments methods are used to generate the Green's function. The kinetic energy operator governs the infrared behavior discussed above.

The lattice QCD studies will exploit two LQCD codes, CPS and Chroma, in generating ensembles of lattice configurations and making measurements. Much of the computation involves the application of a sparse linear matrix to a column vector. The spatial lattice is mapped directly onto the hardware mesh, while the time dimension is either on-node or spread across multiple designated fourth-dimension nodes. Improving the resiliency and self-checks, while minimizing impacts to the computational throughput, is a major collaborative performance research goal of the proposal. A key mathematics research goal is the design of effective multi-grid methods for the null space of the system matrix, which in the case of the Dirac operator in LQCD is not known a priori, and thus must be "discovered" by an adaptive procedure. The *hyper* linear solver library will be used as a platform for the efficient inversion of the Dirac operator, for the cold LQCD applications relevant to this project.

As degradation of the computations associated with the quark-field contractions step is well documented if performed on the CPU, it is important to move such computations to GPUs to remove this bottleneck. These calculations should be performed on the thread or SIMD level. The quark-field contractions then become dense-linear problems, well suited to the high arithmetic intensity that GPUs typically feature. The development of optimized quark-field contractions will be done within the QUDA framework, and would be available to the general LQCD community.

An important research project will be autotuning the performance and resilience of the three major codes, Chroma and CPS for LQCD, and the Lanczos-based code for the HOBET shell-model-like diagonalization. We will design roofline models of the major kernels of all of these codes, and develop a workflow model of the overall process. These efforts will guide manual code changes to optimize performance for specific architectures, including GPUs, and guide our investments of time in seeking alternative implementations of critical kernels.

Though the algorithms developed under this SciDAC program are needed to fulfill the science goals outlined in the Scope of Work, their applicability is very general and will benefit multiple domain science applications. Development of these algorithms on both homogeneous and heterogeneous architectures ensures that the impact of this research remains relevant as the progress toward exascale platforms continues.