Probing MATTER at

Scientists funded under the SciDAC National Computational Infrastructure for Lattice Gauge Theory are using high-end computing to understand the nature of matter and the Universe. Significant success has been achieved by these computational science techniques that map particles onto a space-time lattice, providing powerful tools to understand matter at subnuclear scales.

Understanding Hadrons

The role of high-end computing in researching macroscopic systems such as evolving stars, fusion plasmas, or global climate patterns may be familiar to many. Perhaps less known is the efficacy of computational science tools in empowering the study of micro-cosmos in the deep interior of matter and its nuclei. Advanced computing has played an important role in some of the major successes that have been achieved in unveiling fundamental aspects of how the internal ingredients of subatomic particles like protons and neutrons interact to hold their particle homes together. In an expanded context, this advances science by enhancing the holistic understanding of the nature of the Universe.

One of the various classification schemes of subatomic particles classifies hadrons as particles that undergo strong interactions. Quantum chromodynamics (QCD) is the theory that describes the structure of hadrons in terms of the strong interactions of the quarks and gluons that comprise them (figure 1; figure 4, p39). Because hadrons are essential building blocks of matter and account for most of the visible matter in the Universe, the stakes are very high for understanding QCD and solving it quantitatively. The interaction strengths, or coupling, between quarks and gluons moving at high energies deep inside hadrons is relatively weak and can be handled by conventional theoretical methods. However, the coupling strength increases dramatically at lower energies, and also when the particles try to move apart. The only viable way to deal with this regime is to use advanced numerical techniques and high-end computing, essentially solving QCD on a four-dimensional space-time lattice—a method known as lattice QCD.

Quantum chromodynamics is the theory that describes the structure of hadrons in terms of the strong interactions of the quarks and gluons that comprise them.

Three Pillars of Discovery

The comprehensive study of the complexities of hadron structure and binding, as well as the related physics of what are called "strong forces," necessitates a true synthesis of the three pillars of science—experiment, theory, and computation ("Scientific Discovery: Powerful, Unpredictable, and Aesthetic," *SciDAC Review*, Spring 2006, p8). The search for order and structure at length scales on the order of fractions of a Fermi (one femtometer: 10^{-15} meter) is undertaken experimentally via huge accelerators (figure 2, p38) that are traditionally part of Big Science in particle physics ("Designing Accelerators: Precision Probes for Scientific Discovery," *SciDAC Review*, Spring 2006, p12). These big machines accelerate minuscule subatomic particles to giga-electron volts (GeV) of energy, or more, and smash them into nuclei or other particles. The plethora of new particles created in accelerator experiments yields information about the nature and behavior of matter at the hadronic and subhadronic levels. The interpretation and analysis of the explosion of experimental data resulting from these experiments, and indeed the ideas and projections that fuel them, follow from theoretical descriptions and models, such as QCD, that are developed in the language of mathematics. Some of the major difficulties that complicate analysis arise when dealing with the critical energy and space-time regimes that dominate the recombination of quarks and gluons to form bound hadronic states. The structure of these hadrons cannot be understood rigorously and correctly without lattice QCD and the tools of high-end computing. The path to discovery is thus strongly dependent on an approach in which theory, experiment, and computational science reinforce each other.

SciDAC Lattice QCD Projects

Today, the confluence of theory, numerical algorithms, and computer technology enables theorists to quantitatively solve lattice QCD for a range of fundamental problems. This research helps determine the limits of validity of the Standard

Subnuclear SCALES

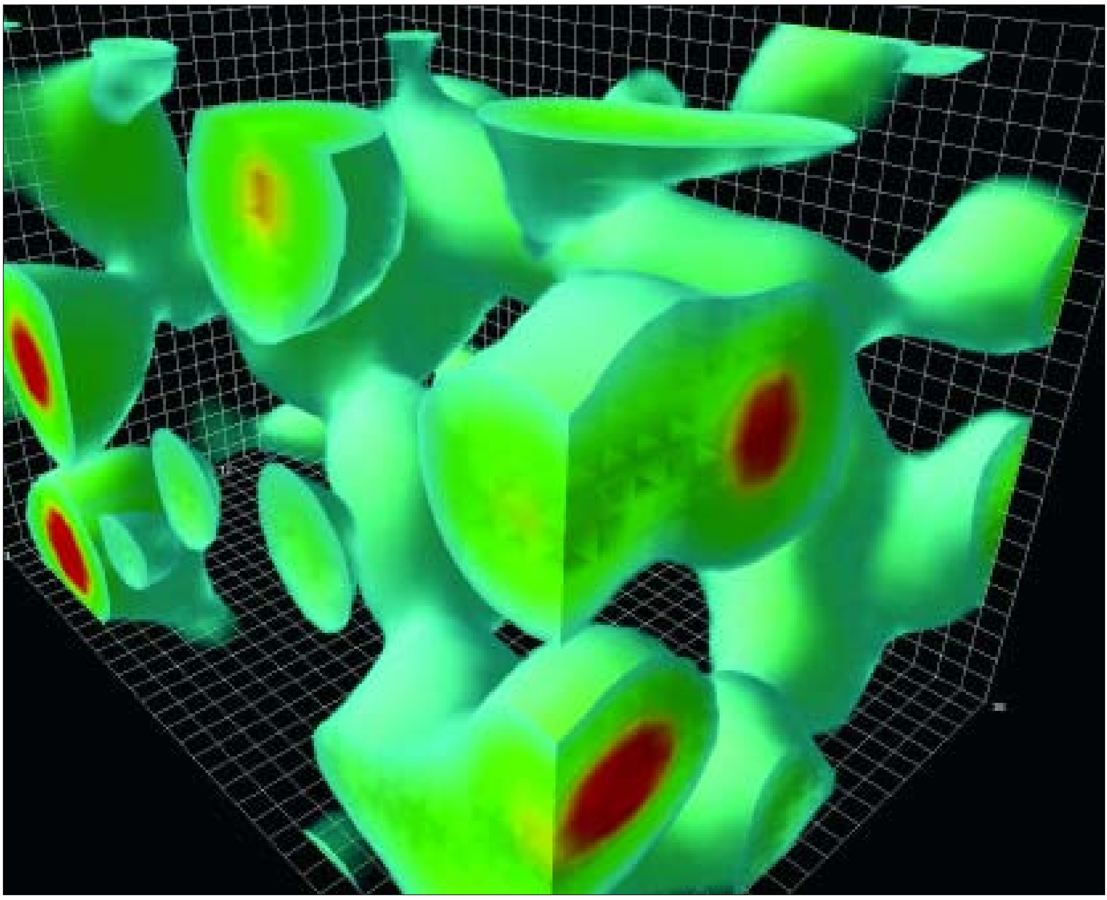

Image courtesy of Derek Leinweber, CSSM, University of Adelaide

Figure 1. *A frame from an animation illustrating the typical four-dimensional structure of gluon-field configurations used in describing the vacuum properties of QCD. The animation was featured in the 2004 Nobel Prize lecture.*

Model of high-energy physics, and also allows scientists to explore the physics beyond it. Lattice QCD also facilitates understanding of the internal structure of hadrons. These efforts, along with exploration of the properties of strongly interacting matter at extreme temperature and density, are having a decisive impact on some of the major science goals of the DOE Office of Science (SC).

Catalyzed by SciDAC support, the U.S. lattice QCD community—approximately 85 senior theorists and their associated post-docs and students—is working under the National Computational Infrastructure for Lattice Gauge Theory project to develop the hardware and software infrastructure necessary for world class lattice QCD research and to utilize this infrastructure for scientific discovery. This SciDAC-2 project ("SciDAC-2: The Next Phase of Discovery," *SciDAC Review*, Spring 2007, p16) will build upon the progress of the SciDAC alumni project, National Infrastructure for Lattice Gauge Computing. Notable achievements include insights into the masses, properties, and decay characteristics of specific hadrons, and the elucidation of the behavior of high-energy-density (HED) matter produced in relativistic heavy ion collisions.

Introducing the Hadrons and their Hidden Residents

One way to introduce the world of hadrons—and the quarks and gluons that reside in their interiors—is to first discuss the interactions that govern the particles, and then explore beyond the hadronic world to determine connections to the large-scale behavior of the Universe itself. The fundamental interactions of nature as they are known today are the gravitational, electroweak, and strong interactions. Gravitation, which holds the Earth in orbit and determines motion at the

FERMILAB

Figure 2. *An aerial view of the Fermilab Tevatron accelerator, where experiments have been conducted to constrain the fundamental parameters of the Standard Model.*

largest astronomical scales, is the weakest in strength of the fundamental forces. At low energies, the electroweak force decouples into the electromagnetic interaction that holds atoms and molecules together and the weak force responsible for decay of particles and nuclei. The star of the QCD story—the strong interaction—is responsible for holding hadrons together while its residual extensions bind protons and neutrons in the nucleus. Taking part in strong interactions—something only some particles can do—is the characteristic that defines the class of particles called hadrons.

The word "hadron" derives from the Greek "hadros," meaning "strong," and the best known examples of hadrons are the familiar constituents of atomic nuclei—protons and neutrons. Protons and neutrons are in fact the lightest of a whole class of similar hadrons called baryons.

After a formal introduction to the hadron group and its subclasses, the baryons and mesons, it is time to probe deeper and meet their residents. The main constituents of baryons are quarks (figure 3). Quarks were named by Murray Gell-Mann, winner of the 1969 Nobel Prize in Physics. The particles were named after "three quarks for Muster Mark," a line from *Finnegan's Wake,* the novel by James Joyce. Quarks are an unusual breed of particles. They carry fractional electric charge, are the primary participants of strong interactions, and always remain hidden inside hadrons. Along with the quarks, hadrons also host gluons, particle carriers of the strong interaction force field. The role of gluons in strong interactions is thought of as analogous to the photons of the electromagnetic field.

Taking part in strong interactions—something only some particles can do—is the characteristic that defines the class of particles called hadrons.

The interiors of hadrons are scenes of incredible activity with quarks exchanging gluons, gluons exchanging gluons, quark–antiquark pairs rapidly materializing from and annihilating back into the vacuum, and, sometimes, transient quarks and antiquarks playing catch with gluons. The mass of a hadron is determined not only by its main quarks but in large part by the gluon fields that hold the quarks within the nucleus. By spin classification, quarks are fermions and gluons are bosons (sidebar "Meeting the Particles").

Asymptotic Freedom

QCD is the theory describing how quarks and gluons interact by exchange of gluons to form protons, neutrons, and other strongly interacting particles (figure 4). In structure it is a field theory, drawing on similarities with the analogous field theory, quantum electrodynamics (QED), which describes the interactions of electrically charged particles via the photon field. However, by behavior, and by the physics it represents, QCD is strikingly different from QED (sidebar "QCD: Colorful and Strong," p40). Unlike the neutral photons of QED, the gluons of QCD are charged, and so gluons can couple to each other, in addition to coupling to quarks in the way photons couple to electrons. The charges responsible for QCD interactions are called "color charges," hence the colorful "chromo-" component of the theory's name. However, it is important to understand that color charges have no connection with visual color perception, or to the electrical charges of QED.

A quantitative analysis of the interaction between quarks by gluon exchange reveals the remarkable property of asymptotic freedom, a

Meeting the Particles

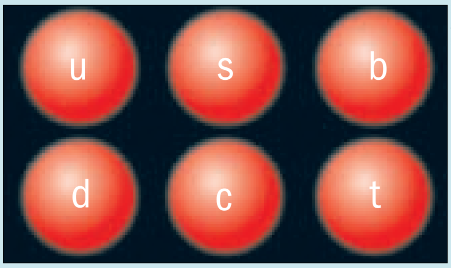

A. TOVEY

Figure 3. *The six types of quarks are up (u), down (d), strange (s), charm (c), beauty/bottom (b), and top (t). In this visual representation, the lightest quark flavors are shown on the left, heavier quarks in the center, and the heaviest flavors on the right. Each quark can carry one of three "color charges."*

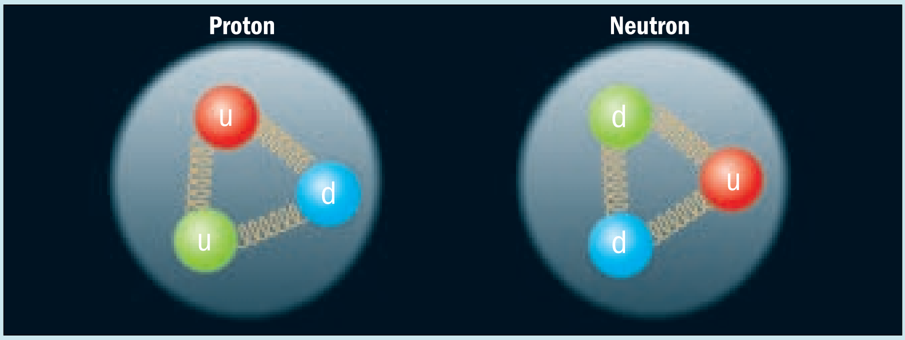

A. TOVEY

Figure 4. *Hadrons, such as protons and neutrons, are composed of quarks. Protons are made up of two up (u) quarks and one down (d) quark; neutrons have one up quark and two down quarks. The colors in the image symbolize the conventional use of red, blue, and green for the quark color charges. The additional gluons, their interactions, and the quark–antiquark sea are not visibly represented.*

The hadrons—baryons and mesons—of QCD can be represented as composites of their ultimate building blocks, the quarks and gluons. The six quarks (figure 3) are called: up (u), down (d), strange (s), charm (c), beauty/bottom (b), top (t). Each quark can carry one of three color charges, called red, green, and blue (figure 4). Quarks also carry fractional electric charges (multiples of 1/3 e), positive or negative, and quarks have half-integral spin. The eight gluons (also with color charges) have zero electric charge and have integral spin.

Baryons contain three main quarks and gluons. Baryons have half-integral spin and are thus classified as fermions. They also define the particle property of baryon number, as all baryons have a baryon number of unity. Protons contain a combination of two up quarks and one down quark; neutrons are composed of one up quark and two down quarks (figure 4). Other baryons are made up of different combinations of the six quarks.

Mesons, on the other hand, contain a quark and an antiquark, plus gluons. They have integer spin and a baryon number of zero. Examples include: pion (u,d); kaon (u/d,s); D (c,u); D_s (c,s); charmonium (c,c); upsilon (b); B_c (b,c); B_s (b,s); and many more variations.

A party with the particles would not be complete without a hello to the non-hadrons, or particles that do not take part in strong interactions. These are called leptons and, to date, retain the status of being elementary and without constituents. The most well known lepton is the electron, which needs no introduction. Leptons have no color charge, can have plus, minus, or zero electric charge, and have half-integral spin. Baryon number is also zero, as they are non-baryons and non-hadrons.

Quarks and leptons also fall under three flavor groups, or generations, characterized by electron, muon, and tau flavors, respectively for the leptons. The electron, muon, and tau are the charged leptons, while their three neutrino counterparts are electrically neutral. For the quark family, the lightest quarks, u and d, are first generation, c and s the next flavor doublet, and the heaviest quarks, t and b, make up the heavy flavor states (figure 3).

Weak interactions transform quarks and leptons between flavor states and the rates for these are important quantities for physics. In the case of quarks, their decays occur inside hadrons and the resulting hadronic decay parameters are measurable by experiment. Quark flavor changes and decays are characterized matrices called the Cabibbo–Kobayashi–Maskawa (CKM) matrix elements.

The particle building blocks and their interactions as described above are in accordance with the Standard Model of particle physics that has passed many experimental tests though efforts have been made to make these tests more stringent. To probe physics beyond its predictions continues to be an exciting goal for current and future research.

discovery honored by the 2004 Nobel Prize in Physics ("Quantum Chromodynamics," *SciDAC Review*, Fall 2006, p5). Asymptotic freedom states that the interactions mediated by gluons become weaker as the collision energy increases, or as the quarks move closer to each other. Typically, these are situations where the quarks are deep inside the hadrons where quarks are believed to move about relatively freely. In contrast, this force increases greatly as the quarks move further apart, corresponding to lower energy regimes. As quark separations grow to those approaching hadron dimensions, the gluon forces strengthen to permanently confine the quarks. Thus, unlike electrons, which can be removed from atoms, a single quark cannot be removed from a proton. Since quarks are confined inside a hadron, scientists have to learn how to understand an unprecedented composite system—one that has constituents which cannot be removed for study in isolation.

QCD on the Lattice

The good news is that QCD has been studied quantitatively in high energy collisions where the interactions are weak, providing overwhelming

QCD: Colorful and Strong

Why is it that quarks cannot leave hadrons? Why must they always stay confined inside? To answer these questions it is necessary to revisit the color force of QCD, and the "color charge" property that particles must possess to be participants in QCD interactions. At the internal level of hadrons, the quarks and gluons carry the color charge that makes them players in the QCD game. In addition, quarks also carry fractional electric charge, and hence can take part in electromagnetic interactions, too, while gluons are inert to these interactions.

There are six quarks (sidebar "Meeting the Particles," p39) and each variety can come in any of three color charges, conventionally called red, blue, and green—analogous to familiar notions of visual color.

One of the reasons free quarks are not observed is because only color neutral (or, in more technical terms, color singlet) states of quarks and gluons are observable states. Combinations of three quarks of suitable color or quark–antiquark pairs of same color generate these physically realizable bound states of baryons and mesons, respectively.

Delving deeper, QCD processes occurring by gluon exchange are different from the analogous quantum electrodynamics (QED) processes by photon exchange due to the fact that gluons themselves carry color charge, unlike the photons, which do not carry electric charge. Thus, gluons can interact with each other or with quarks, whereas photons cannot interact with each other, but only with charged particles, such as electrons.

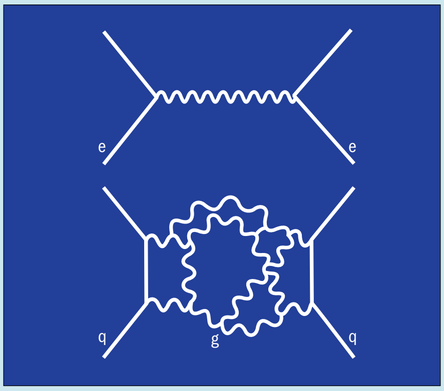

Figure 5. *Two examples of Feynman diagrams, a practical visual depiction of the interaction between particles. The upper diagram shows the interaction between two electrons in quantum electrodynamics (QED); the lower diagram depicts the interaction between two quarks (q) in quantum chromodynamics (QCD).*

A. Tovey

For example, figure 5 shows a process in QED in which an electron, denoted by the left-hand solid line, propagates in space-time, and at some point emits a photon, denoted by the wavy line. This photon, in turn, is later absorbed by another electron, represented by the right-hand solid line. Writing the formula from quantum field theory for this process and summing over all the places the photon could be emitted and absorbed yields the one-photon exchange contribution to the interaction between two electrons. One can similarly enumerate all the diagrams in which two photons are emitted from one electron and absorbed by the other electron in all possible ways. Because photons have no charge, they cannot interact with each other, and the final result is rapidly convergent. In the case of atomic physics, the lowest-order one-photon exchange dominates and yields the familiar Coulomb interaction plus tiny corrections. All the features of atoms follow from this, including the qualitative fact, often taken for granted, that the force between two electrons approaches zero at large distances, and this allows an electron to be removed from an atom.

With this introduction, it should be clear why QCD is so strikingly different. Unlike the neutral photons of QED, the gluons of QCD are charged so that gluons can couple to each other in addition to coupling to quarks in the way photons couple to electrons. Hence, as shown in figure 5, a typical Feynman diagram has an immense tangle of gluons interacting with each other, as well as with the emitting and absorbing of quarks. As the separation between quarks increases, it turns out that the proliferation of these gluons causes the interaction to increase without bound. Thus, unlike electrons that can be removed from atoms, a single quark cannot be removed from a proton.

Since quarks are confined inside a hadron, scientists have to learn how to understand an unprecedented composite system—one that has constituents which cannot be removed for study in isolation.

evidence for its validity. The bad news is that at the large distances and low energies governing the binding and structure of hadrons, the interactions become so strong and nonlinear that, unlike the case of QED, no known analytical technique can quantitatively describe them. Theoretical physicists have turned to large-scale computation to solve QCD numerically in order to unlock the secrets of the strong interaction. The inspiration and techniques come from Feynman's path integral formulation of field theory that sums over all possible time histories contributing to a process, with appropriate weighting. To implement this idea in a practical calculation, scientists replace the space-time continuum with a discrete lattice in space and time that possesses a finite lattice spacing, a, and a finite volume. Quark fields are defined on the sites of the lattice, and the gluon fields are defined on the links connecting lattice sites. A discrete action is written for this lattice with the property that it approaches the QCD action in the limit as a goes to zero. The integral over the gluon fields is evaluated by Monte Carlo methods, so that as the number of time histories, N, sampled with a weight given by the discrete QCD action increases, the exact path integral is approximated with controlled errors that decrease as the reciprocal of $\sqrt{N}$.

To illustrate how this process works, figure 8 shows a typical time history contributing to a lattice calculation of the propagation of a meson between two points in space-time, denoted by the solid black dots. Quarks and antiquarks are represented by directed lines propagating forward and backward in time, respectively. The blue lines represent one possible path for a quark–antiquark pair, starting at one dot and ending at the other. The green lines represent a dynamical vacuum polarization process in which a quark–antiquark pair is excited out of the vacuum. The set of red

QCDSP, QCDOC, and BlueGene

The roots of the QCDOC machine (figure 6) extend back to an earlier QCD machine, QCDSP, which was constructed from digital signal processor (DSP) chips. Two versions of this pioneering machine were constructed: an 8,192-node (0.120 teraflops sustained) DOE-funded machine at Columbia University and a 12,288-node (0.180 teraflops sustained) RIKEN-funded machine at BNL. The BNL machine won the Gordon Bell prize in 1998 for its $10/Mflops price performance.

After completion of the QCDSP machine in 1998, one of its architects, Dr. Alan Gara, joined the IBM Watson Research Center. Dr. Gara brought with him an enthusiasm for this low-power, tightly-integrated, mesh-based approach to massively parallel supercomputing. This quickly led to QCDOC, a joint project between Columbia, the RIKEN BNL Research Center, the University of Edinburgh, and IBM. As this project progressed the group at IBM recognized commercial possibilities for this approach to supercomputing.

BNL

Figure 6. *The QCDOC machine at BNL.*

LLNL

Figure 7. *The BlueGene/L supercomputer at LLNL.*

Recruiting three additional scientists from the original Columbia QCDSP team—Dr. Dong Chen, Dr. James Sexton, and Dr. Pavlos Vranas—IBM embarked on a more ambitious project which built upon the QCDSP/QCDOC concept with additional memory, enhanced mesh communications with cut-through routing, a fast tree collective network, and more than four times the floating point capability. The result was the IBM BlueGene/L computer (figure 7), now a highly successful commercial product which is installed in national laboratories, universities, and supercomputer centers around the world.

The largest BlueGene/L installation at LLNL (figure 7; "High-Performance Hardware," *SciDAC Review*, Fall 2006, p40) with a Linpack performance of 280.6 teraflops is expected to occupy for some time the first position on the top 500 list, a ranking of the most powerful computer systems based on the Linpack performance benchmark. This close relation between the QCDSP and QCDOC machines and the IBM BlueGene computers provides an excellent example of technology transfer from a DOE-supported university project to a breakthrough commercial product.

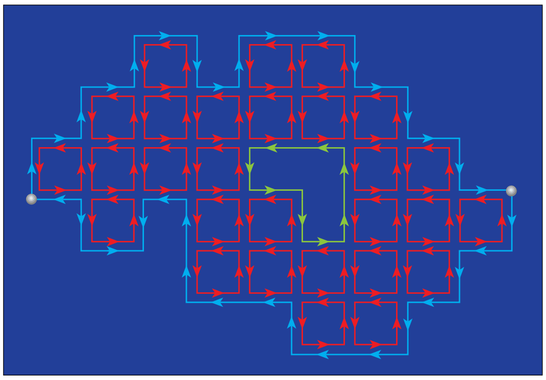

A. Tovey

Figure 8. *One time history contributing to the propagation of a meson in lattice QCD between the two points denoted by the gray dots. The blue path shows the trajectory of a quark–antiquark pair, the green path shows the excitation of a quark–antiquark pair from the vacuum, and the red lines indicate gluon fields.*

squares, or "plaquettes," tiling the region between quarks and antiquarks represents the gauge field configuration. A corresponding diagram for a nucleon would have three blue quark lines connecting the gray dots, along with any number of dynamical green quark–antiquark excitations, and a corresponding gluon tiling. For a lattice with 64 lattice sites in each of the three spatial directions and 128 sites in the time direction, there are roughly 1.2 billion gluon variables. The integration over the quarks can be done analytically, but leads to the problem of solving a huge system of linear equations, the well known Dirac equation, a discretized partial differential equation (PDE) with a 400 million × 400 million matrix. The main computational challenge is to solve this lattice Dirac equation, which must be done tens of thousands of times to generate an appropriate statistical sampling of the QCD action via a variant of the molecular dynamics technique.

Computers for Lattice QCD

Over the past twenty years computers specifically optimized for lattice QCD have provided platforms for frontier physics calculations and achieved record price performance. From the Caltech Cosmic Cube to the Quantum Chromodynamics on a Chip (QCDOC; figure 6) and the SciDAC-1 clusters, technological opportunities have been exploited to provide economical, large-scale computational capabilities. These national efforts and similar activities carried out in Germany, Italy, Japan, and the United Kingdom have also played an important role in the development of commercial supercomputers. For example, the QCDSP and QCDOC machines developed and demonstrated the architecture that is now the basis for the highly successful IBM BlueGene computers (sidebar "QCDSP, QCDOC, and BlueGene").

Lattice QCD calculations have a simple, regular structure. Most of a lattice QCD calculation is

Vital to both the cluster and QCDOC machines is the common software environment created under this SciDAC project.

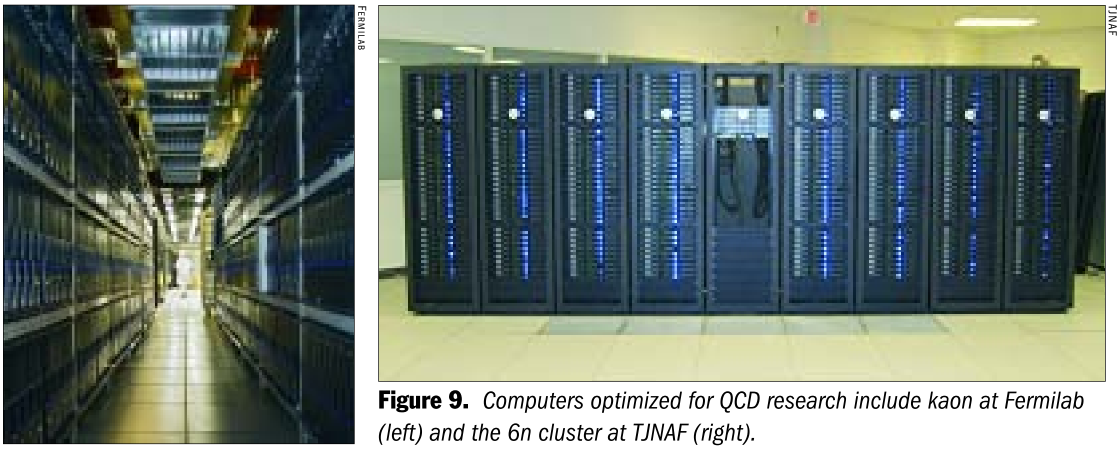

Figure 9. *Computers optimized for QCD research include kaon at Fermilab (left) and the 6n cluster at TJNAF (right).*

spent using a conjugate gradient routine to invert the Dirac difference operator. For this inversion the largest computational load comes from performing the 3 x 3 matrix-times-vector multiplies needed to parallel transport the quark spinors between neighboring sites. This is a calculation that is easy to parallelize, with the individual nodes of a massively parallel machine assigned small subcubes in the larger volume. Nearly all of the required communication is between nodes that contain physically adjacent sub-volumes. This implies that a mesh or hierarchical switched network will work well. High floating point performance for these matrix-times-vector multiplies, substantial memory read bandwidth, and very little input/output requirements are further characteristics.

The collaboration participating in this QCD SciDAC effort has pursued two hardware optimization strategies. For the first, specialized workstation clusters have been developed with mother board components and inter-node communications fabric optimized especially for lattice QCD. The second approach involves designing a customized QCD computer from the ground up.

The workstation cluster effort began with extensive tests of a variety of commercial processors and networks. The original SciDAC-1 grant supported experiments with both individual mother boards and network cards but also with entire clusters constructed from a few hundred nodes. This allowed significant tests of complete systems including hardware (processors, memory bus, and network cards) and software (buffering strategies, network protocol drivers, and the SciDAC software) under "battle field" conditions. The resulting test-bed machines provided the experience being exploited in the multiteraflop clusters now being constructed as part of the DOE Lattice QCD Computing Project, which started in October 2005, and the hardware that was used to obtain some of the physics results presented here. These clusters include the 3g, 4g, and 6n clusters (figure 9) at the Thomas Jefferson National Accelerator Facility (TJNAF) and the qcd, pion, and kaon clusters (figure 9) at the Fermi National Accelerator Laboratory (FNAL). The largest of these clusters, 6n and kaon, contain 280 and 600 nodes, respectively.

The second approach undertakes the design and construction of the entire QCD computer. The QCDOC project began in 1999 and achieved a price-per-sustained performance on QCD of \$1/Mflops for calculations on multiteraflops partitions comprised of thousands of computing nodes. This approach allows a choice of not only the processor and network fabric, but also cooling, packaging, power consumption, diagnostics, and reparability. The resulting QCDOC computers have a typical size of 12,000 nodes with each node made up of a single Applications Specific Integrated Circuit (ASIC) and commercial memory module. The ASIC contains a standard PowerPC processor with an integrated floating point unit and fast serial communications to 12 nearest neighbors in a 6-dimensional mesh network. These machines were designed and constructed by a collaboration of Columbia University, the RIKEN BNL Research Center (RBRC), the University of Edinburgh, and IBM. The QCDOC project resulted in three large computers, two 12,288-node machines installed at Brookhaven (one funded by the DOE and the other by the RBRC). A third PPARC-funded, 13,302-node QCDOC computer is installed in Edinburgh.

Vital to both the cluster and QCDOC machines is the common software environment created under this SciDAC project. For example, given the importance of high-bandwidth, low-latency communications it is essential that both the cluster network cards and the direct-memory-access QCDOC communication circuitry be accessible to standard application code in a way that permits the special features of each implementation to be exploited. The QCD Message Passing (QMP) communications

A. TOVEY

Figure 10. *The SciDAC QCD API. SciDAC-1 components are shown in white lettering, and the new SciDAC-2 components in yellow lettering.*

layer provides just this capability. Using this and other QCD-optimized software components, application codes from many groups within the SciDAC collaboration are now running efficiently in production on both the QCDOC and cluster platforms. The clusters and QCDOC machines achieve comparable cost performance and run a mixture of jobs allocated to exploit their somewhat different characteristics.

SciDAC Software

The rapid transition from terascale to petascale facilities in less than a decade puts extraordinary demands on nimble software development. Inferior application software not only wastes state-of-the-art hardware facilities, but also diverts scarce human resources from the scientific mission.

To meet this software challenge the U.S. lattice gauge theory community under SciDAC has created a unified programming environment. The core is a layered QCD Applications Programming Interface (QCD API) that isolates the machine-dependent routines for communication and algebraic primitives (figure 10, Level 1) from an application-specific data parallel QCD interface (Level 2). At the top (Level 3) the most critical algorithms are recoded directly to each architecture to obtain maximal overall performance. All of the fundamental components of the QCD API have been implemented and are in use on the U.S. QCDOC hardware at BNL, on both the switched and mesh architecture Pentium 4 clusters at FNAL and TJNAF, and on a number of general purpose supercomputers. The QCD API is being used by a growing number of physicists in the U.S. and abroad. The libraries and documentation can be found at the USQCD (see Further Reading, p47) and made available worldwide. This ensures that all involved physics students and the increasingly mobile postdoctoral researchers can move freely along their career paths without fear of losing their investment in software experience and tools to continue their research at the next research institution.

All components are initially written in highly portable C or C++ code over a Message Passing Interface (MPI) so applications can be run and developed on all available architectures. However, extreme performance depends on specific implementation of the libraries on Level 1 & 3 native to the target architecture. The following sections describe how each component and current work are bringing high performance to the next generation of hardware, specifically the BlueGene/P and Cray XT3.

Message Passing

QMP defines a uniform subset of MPI-like functions with extensions that partition the QCD space-time lattice and map it onto the geometry of the hardware network, providing a convenient abstraction for the Level 2 data parallel API, QCD

The rapid transition from terascale to petascale facilities in less than a decade puts extraordinary demands on nimble software development. Inferior application software not only wastes state-of-the-art hardware facilities, but also diverts scarce human resources from the scientific mission.

Data Parallel (QDP). The functions also contain specialized routines designed to access the full hardware capabilities of novel networks such as that on the QCDOC or BlueGene.

Linear Algebra

All lattice QCD calculations make use of a set of linear algebra operations in which the basic elements are the three-dimensional complex matrices required by the QCD symmetries. These operations are local to lattice sites or links and do not involve inter-processor communications. The direct C implementation has about 19,000 functions generated in Perl, with a full suite of test scripts, whereas the C++ implementation makes considerable use of templates, and so contains only a few dozen templated classes. The required specific classes are generated on demand by the compiler.

Data Parallel Interface

QDP (Level 2) contains data parallel operations that are built on QMP and QCD Linear Algebra (QLA). The re-engineering of the fundamental application code on top of QDP is both simplifying the code and bringing increased performance. For example, the MILC code—a large, publicly available suite of applications, carefully optimized over its fifteen year lifetime—has shown significant improvement. Chroma, an entirely new application code base, has been written *di novo* in the C++ implementation of QDP. 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, and other elements are hidden from the user.

Level 3 Subroutines

A very large fraction of the resources in any lattice QCD simulation go into a few computationally intensive subroutines, to describe the space-time trajectory of the quarks. Computationally this is a set of specialized linear algebra operations on sparse matrices on the order of 100 million rows and columns. These subroutines are optimized for each architecture. Typical assembly codes for these Dirac matrix inverters achieve 30% to 45% of peak. The result is to bring current simulations on commodity clusters to well under $1.00/ Mflops sustained performance.

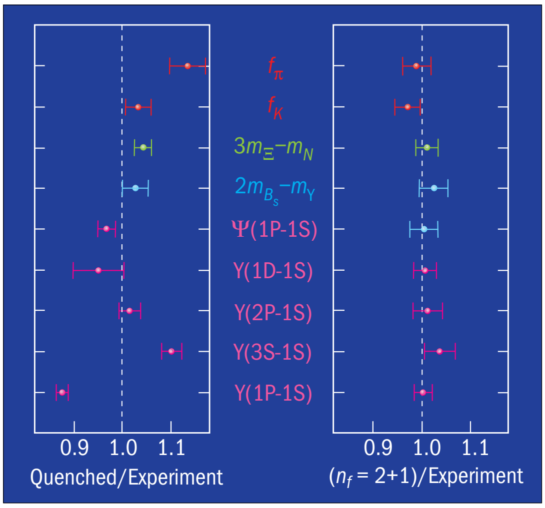

A. Tovey

Figure 11. *Lattice QCD predictions for meson decay constants and hadron mass splittings compared with experiments. The graph on the left shows the results omitting the effects of vacuum polarization due to light quark loops; the graph on the right shows the results including these effects.*

QMC: Threaded Libraries for Multicore Processors

Perhaps the most immediate change in programming requirements is the transition to multicore commodity microprocessors. Intel, AMD, and IBM have started the move toward multicore processors in the last two years. In the short run, lattice QCD application codes can take advantage of multicore microprocessors by simply treating the cores as independent processors. However, to fully exploit the multicore capability the QMC interface is being developed to provide the data parallel code portable access to thread optimized components. It will provide an abstraction for threads that can be implemented for portability in the Portable Operating System Interface (POSIX) standard, but will have native implementations for highest performance and for unconventional multicore architectures.

Other areas of active research and development include QCD specific graphics and visualization tools, workflow and data analysis tools to carry out analysis campaigns on shared configuration database, and performance analysis tools to rapidly assess the quality of compilers and the need for specialized implementation for the low-level API core functions.

Science Accomplishments

Weak Decays and Tests of the Standard Model

Quark–gluon interactions play a role beyond hadron binding and structure and can influence the way quarks inside hadrons decay by weak interactions. The Standard Model (sidebar "Meeting the Particles," p39), which has been remarkably successful in describing nature via the fundamental interactions, also provides a classification of quarks and leptons via what we call their "flavor" categories. The weak decays of quarks result in flavor transitions that are governed by the parameters of a 3 x 3 matrix, the

Two essential goals for lattice QCD are the validation of lattice calculations by comparing lattice results with experiment, and the use of lattice calculations to extract from experiment the fundamental parameters of the Standard Model.

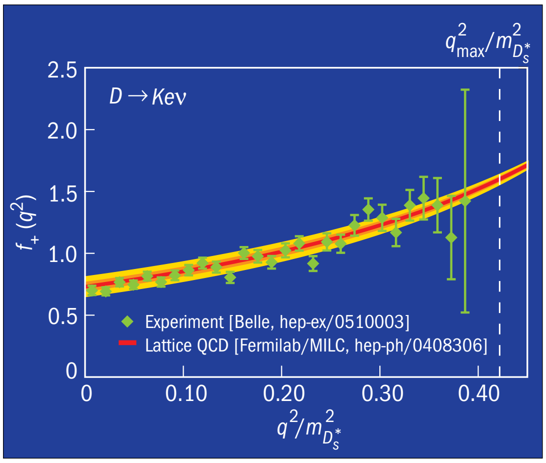

A. Tovey

Figure 12. *When a D meson, made up of a charmed quark and an up antiquark, decays into a kaon, made up of a strange quark and an up antiquark, plus a neutrino and an anti-electron, the form factor $f_+(q^2)$ describes the dependence on the momentum transfer q^2 of the decay. The gold bands show the shape of this form factor as predicted by lattice calculations, and the green diamonds show the subsequent confirmation by the Belle experiment.*

Cabibbo–Kobayashi–Maskawa (CKM) matrix, characterizing mixing of quark states.

Two essential goals for lattice QCD are the validation of lattice calculations by comparing lattice results with experiment, and the use of lattice calculations to extract from experiment the fundamental parameters of the Standard Model, such as the masses of quarks, the strong coupling constant, and the matrix elements characterizing the weak interactions of quarks. By comparing various determinations of these parameters, researchers test the Standard Model itself. The greatest challenge to accurate calculations is the proper treatment of vacuum polarization due to the light quarks. Significant progress has been made in meeting this challenge during the past five years through the use of improved formulations of QCD on the lattice and through rapid growth in the computing resources available to the field. Among the notable results have been calculations of the leptonic decay constants of pions and kaons, and mass splittings in the charmonium and bottomonium spectra to an accuracy of 3% or better. Figure 11 shows some of these quantities compared with experiment, omitting the vacuum polarization of light quarks (left), and including it (right). These calculations and others are now enabling the determination of the fundamental parameters of QCD, including the light quark masses, the strong coupling constant, and CKM matrix elements such as V_{us}.

The lattice gauge theory community has moved from the validation of techniques through the calculation of quantities that are well known experimentally to the successful prediction of quantities that had not previously been measured. Three cases in which predictions were subsequently confirmed by experiment were the calculations of the leptonic decay constant and semi-leptonic form factors of the D meson (figure 12), and the mass of the charmed B meson, represented symbolically as B_c. The successful calculations for D mesons provide important validation of similar calculations for B mesons, which play important roles in tests of the Standard Model. These results indicate that the research community is in a position to make very significant progress over the next five years.

Thermodynamics of Hot and Dense Matter

The study of properties of strongly interacting matter under extreme conditions, such as high temperature and/or density, provides the basis for a deeper understanding of the early evolution of the Universe, as well as for understanding the properties of compact stellar objects. Moreover, by analyzing the response of matter to extreme conditions scientists also learn about basic properties of QCD (sidebar “Matter Under Extreme Conditions,” p46).

At very high temperatures and/or densities, one expected observation is a dramatic change from ordinary strongly interacting matter to a plasma of quarks and gluons. Along with the generation of new degrees of freedom, which remain invisible in ordinary matter, there occurs a drastic change of thermal properties. The energy density and pressure in a hot and dense medium of quarks and gluons are expected to gradually approach that of an ideal gas. Large density and charge fluctuations should characterize the occurrence of the transition. Among the issues that can uniquely be addressed by lattice QCD calculations are the nature of the transition, the temperature at which it occurs, the properties of the plasma, and, in particular, the equation of state which relates the energy density and pressure of matter under extreme conditions to its temperature.

Indeed, lattice calculations provided the best estimates of the temperature at which the transition from ordinary matter to a quark–gluon plasma occurs. It is now known from lattice calculations that the energy densities at the transition temperature and beyond can be generated in modern particle accelerators such as the Relativistic Heavy Ion Collider (RHIC; figures 13 and 14, p46) at BNL. Lattice calculations also show that the equation of state of hot and dense elementary particle matter is significantly different from that of an ideal gas. Unlike in an ideal gas, the pressure of the medium is significantly less than one third of its energy density in the entire temperature regime accessible to experiments at RHIC. The

Indeed, lattice calculations provided the best estimates of the temperature at which the transition from ordinary matter to a quark–gluon plasma occurs.

Matter Under Extreme Conditions

The properties of strongly interacting matter under extreme conditions are studied in large-scale numerical calculations on today's most modern computers. Numerical results deduced from such calculations are confronted with experimental findings from collisions of nuclei accelerated to almost the speed of light. In fact, a primary motivation for the construction of the Relativistic Heavy Ion Collider (RHIC) at BNL was to recreate matter under extreme conditions hypothesized to exist milliseconds after the Big Bang. The collision of two heavy nuclei creates a "little bang." The dense matter formed in each little bang leaves its traces in the form of thousands of newly generated particles, which are observed by sophisticated detectors.

An important experimental finding of recent years is that the dense matter created at RHIC quickly equilibrates and behaves almost like a perfect fluid. This is in accordance with lattice QCD results. Calculations of the transition temperature to the high-temperature and high-density regime of QCD indicate that this new state of matter is reached in experiments at RHIC. Moreover, the numerical studies of the dependence of energy, density, and pressure on the temperature show that the equation for the state of hot and dense matter deviates significantly from that of a non-interacting gas. Thus, there is sufficient room for interactions that can lead to a liquid-like form of hot and dense matter.


BNL

Figure 13. *A section of the RHIC accelerator that creates colliding beams of relativistic heavy ions at BNL.*

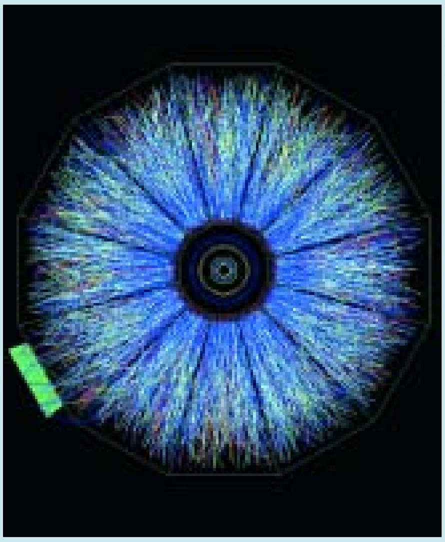
STAR Collaboration

Figure 14. *Tracks from a relativistic heavy ion collision at BNL's RHIC, showing the thousands of particles produced. Analyses of these collisions indicate that the dense matter created in these collisions equilibrates quickly and forms a nearly perfect fluid.*

reduction of the pressure hints at the presence of strong attractive forces in a quark–gluon plasma. This is in accordance with experimental findings at RHIC, which suggest that the dense matter created in heavy ion collisions equilibrates quickly and behaves almost like a perfect fluid.

Today's numerical techniques are well suited for studies of matter under extreme conditions that existed in the early Universe and that can now be reproduced in heavy ion experiments performed at BNL. This matter is characterized by having a vanishing net baryon number, meaning it contains as many baryons as antibaryons. Its study through lattice calculations only requires sufficient computational resources to reach high precision results. Numerical calculations at a non-vanishing baryon number, however, are much more subtle; the numerical techniques used for these calculations are at a much earlier stage of development, and new algorithm developments will be required. Matter with a non-zero net baryon number content is expected to show a rich new phase structure that will also be explored in upcoming experiments at RHIC, as well as at new heavy ion facilities being built in Europe. A major effort in refining lattice studies of matter at non-zero baryon number is thus expected to be performed in the next few years.

Hadron Structure

How do the observed structures of nucleons and other hadrons arise from QCD? What is the spatial distribution of charge and current, and how does the total spin arise from the spin and orbital motion of quarks and gluons? What is the fraction of the momentum carried by the quarks and gluons in a rapidly moving nucleus? How do the heavier, strange quarks contribute to all these quantities? What is the spectrum of meson and baryon resonances of QCD, and are there states in which the gluonic degrees of freedom are manifestly exposed, or states with exotic quantum numbers? And finally, how do hadron–hadron interactions arise from QCD? These questions go to the heart of understanding strong interactions, and all of them can now be addressed by lattice QCD.

The combination of new DOE computational resources and sharing a common set of dynamical quark configurations between high-energy and nuclear physics has enabled theorists to study hadron structure at unprecedented low values of the up and down quark masses. These quark masses are light enough that the results can be compared with experiment using chiral perturbation theory, and the initial successes are impressive. The axial charge of the nucleon, a fundamental property of the strong interaction governing the decay of the neutron, agrees with

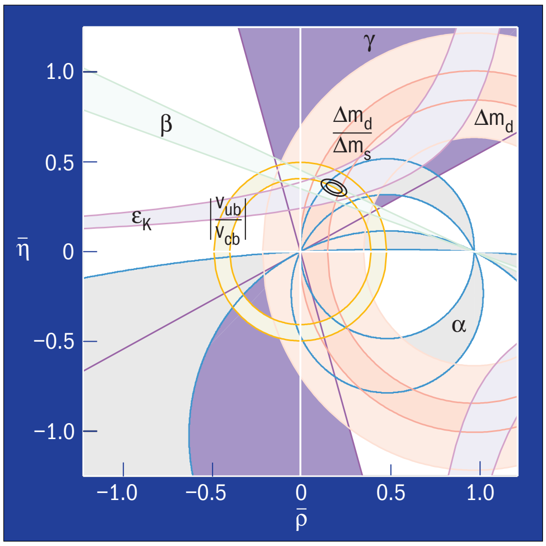

A. Tovey

Figure 15. *Constraints on the least well known parameters of the CKM matrix. If the Standard Model is correct, all measurements should give consistent values. At present, a consistent solution is found: the ellipses.*

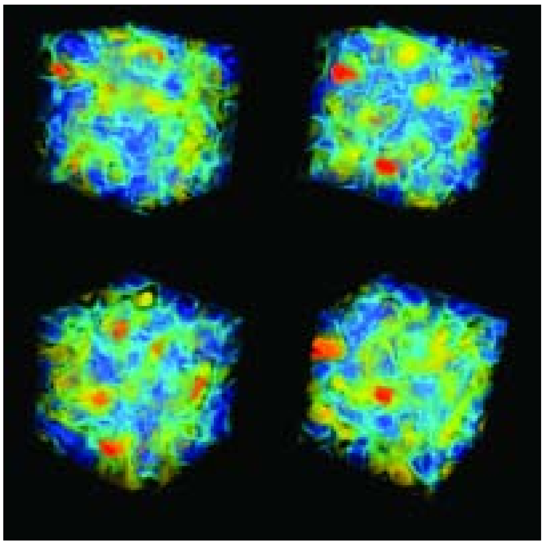

Visualization: M. McGuigan & S. Tomov Data: T. Izubuchi

Figure 16. *Four 3D slices of topological charge density of a lattice QCD gauge configuration. This image was generated at the RIKEN BNL Research Center, which operates two QCDOC supercomputers.*

experiments, as do form factors characterizing the distribution of charge and current in the nucleon, low moments of quark momentum distributions, and the pion scattering length, which characterizes the low energy scattering of two pions.

Building on these successes, ambitious calculations are under way to calculate the origin of the nucleon spin, the transverse spatial distribution and longitudinal momentum distribution of quarks in a rapidly moving nucleon, the contributions of strange quarks to nucleon properties, the distribution of gluons in the nucleon, the spectrum of low-lying mesons and baryons, the expected production rate for mesons having gluonic excitations, and the nucleon scattering length. These calculations will be valuable in both guiding and interpreting experiments at the present TJNAF electron accelerator and its 12 GeV upgrade, in the RHIC spin program, and at a future electron–ion collider.

Future Challenges

The successes to date in using lattice QCD to extract CKM matrix elements from the decays of heavy mesons, to probe the physics of the Standard Model and beyond (figure 15), to study matter under extreme conditions, and to understand the structure of hadrons are exciting, and the continued development of computational infrastructure and algorithms promises important further improvements. However, there is still a host of significant QCD problems that cannot yet be calculated well on the lattice, and hence require new ideas and theoretical breakthroughs. Transport coefficients, such as the viscosity of the quark–gluon plasma, are essential for understanding relativistic heavy ion collisions. Near-cancellation of positive and negative contributions associated with fermion antisymmetry currently limit calculations of strongly interacting matter to low baryon density, so fundamentally new algorithms are needed to understand the high density matter in neutron stars and the full phase diagram of QCD. Much of the momentum and spin of the nucleon arises from gluons, so it is important to calculate the gluon as well as the quark structure of hadrons. Finally, although no one knows what physics beyond the Standard Model will be revealed by the Large Hadron Collider (LHC) and the subsequent linear collider, explorations are already underway to invent lattice formulations of the supersymmetric field theories that may have to be solved at the next high-energy frontier. These and other scientific challenges, partnered with advanced high-end computing capabilities, will continue to energize lattice field theory, keeping it intellectually vibrant and productive in terms of scientific discovery. ●

Contributors: Dr. Richard Brower, Boston University; Dr. Michael Creutz, BNL; Dr. Norman Christ, Columbia University; Dr. Frithjof Karsch, BNL; Dr. John Negele, Massachusetts Institute of Technology; Dr. Paul Mackenzie, Fermi National Accelerator Laboratory; Dr. Jeffrey Mandula; Dr. Claudio Rebbi, Boston University; Dr. David Richards, TJNAF; Dr. Stephen Sharpe, University of Washington; and Dr. Robert Sugar (PI), University of California–Santa Barbara

Further Reading

USQCD
http://www.usqcd.org/usqcd-software