

Lattice Quantum Chromodynamics Comes of Age

Quantum chromodynamics is the elegant but notoriously intractable theory of the strong interactions. Recent advances in numerical computer simulation are beginning to reveal, in impressive detail, what the theory predicts.

Carleton DeTar and Steven Gottlieb

The strength of the electron–photon interaction is characterized by the fine-structure constant $\alpha \approx 1/137.036$. Because α is small, quantum electrodynamics (QED), the theory of interacting electrons and photons, can be solved to very good approximation with the traditional technology of pencil and paper. By contrast, quantum chromodynamics (QCD), the generally accepted theory of strongly interacting quarks and gluons, has proven to be remarkably resistant to that approach. But in recent years, advances in computer technology and algorithms have brought the ab initio, numerical simulation of QCD to a level of credibility that will have a significant impact on scientific discovery.

Lattice QCD

In formulation, QCD and QED are strikingly similar. Both are gauge-invariant quantum field theories. The key difference is that photons in QED are neutral; so they can't interact directly with each other. The gluon is the QCD analog of the photon; it carries the strong force between quarks. But quite unlike photons, gluons do carry color charge, the analog of electric charge. So gluons interact directly with each other as well as with quarks. (See the article by Frank Wilczek in *PHYSICS TODAY*, August 2000, page 22.)

That seemingly innocent change has dramatic consequences for phenomenology. It is the root of QCD's daunting complexity. Electrons, positrons, and photons can be separated and isolated at macroscopic distances. Quarks, antiquarks, and gluons cannot. This prohibition, called color confinement, assures that all the elementary particles (the hadrons) composed of quarks, antiquarks, and gluons come in precise color-neutral combinations. Loosely speaking, this means that they come either in quark–antiquark pairs (the mesons) or in triplets of quarks (the baryons). Several recently discovered “pentaquark” baryons appear to combine a quark triplet with a quark–antiquark pair (see page 19 of this issue.)

Why only color-neutral combinations? In QCD, quarks can have three colors. Conventionally, they are labeled red, blue, and green, but of course they have nothing to do with optics. Antiquarks have the corresponding anticolors. Triplets of quarks containing equal portions of the three

colors are color neutral.

Try to pry loose one of the three valence quarks in a proton. Before going much farther than the radius of the proton (about 1 fm or 10^{-13} cm), you've done enough work to create a new quark–antiquark pair. Pairs promptly appear, choose new partners, and you find a meson in one hand

and a proton or neutron in the other. No isolated quarks!

At distances an order of magnitude smaller than 1 fm or, equivalently, at interaction energies or momentum transfers in the multi-GeV range, α_s , the energy-dependent QCD analog of the fine-structure constant, is effectively weak. In that limited regime, perturbation theory works, and pencil-and-paper methods succeed. But for the larger distances and softer interactions, where confinement is the dominating process, α_s is effectively large and we must resort to computerized numerical simulation.

In 1974, Kenneth Wilson at Cornell University formulated a version of QCD on a discrete spacetime lattice (see the left panel of figure 1) and, with pencil and paper, used it to provide a plausible, but not rigorous, argument for color confinement.¹ Wilson argued that, on a coarse spacetime lattice, the potential energy of separation of a quark and an antiquark must rise linearly with distance. In 1979 at Brookhaven National Laboratory, Michael Creutz, Laurence Jacobs, and Claudio Rebbi demonstrated the feasibility of doing meaningful numerical simulations with Wilson's formulation on a Control Data Corp 7600 computer.² Shortly thereafter, Creutz obtained numerical results for the confinement potential that supported Wilson's conclusions. That success launched a new branch of computational physics, called lattice gauge theory or lattice QCD.³ The right panel of figure 1 shows a modern lattice-QCD result for the quark–antiquark potential.⁴

High-precision calculations

For two decades after Creutz's pioneering 1979 calculations, refinements in algorithms and computing power brought steady gains in precision and consistency. But only in the past four years have powerful algorithmic and theoretical improvements launched us into the age of high-precision lattice QCD—at least for some key hadronic quantities.

By the standards of the strong interactions, “high precision” means 1 or 2%. The impact of this new precision extends beyond the strong interactions. Determining key features of the *weak* interactions of hadrons—for example, the Cabibbo-Kobayashi-Maskawa (CKM) parameters—requires correcting measured weak decay rates for strong-interaction effects (see box 1). The uncertainties in our knowledge of such fundamental parameters limits the precision with which the standard model of the elementary particles can be tested and probed for new realms of physics.

The most important theoretical advance in recent

Carleton DeTar is a professor of physics at the University of Utah in Salt Lake City. Steven Gottlieb is a professor of physics at Indiana University in Bloomington.




Figure 1. Lattice quantum-chromodynamic calculation. (a) In the Wilson lattice formulation, the gluon field operator U links the quark fields q at adjacent lattice sites in the discretized four-dimensional spacetime. In the 3D internal color space of QCD, q is a 3-vector and U is a 3×3 matrix. (b) A recent lattice-QCD calculation of the potential energy of separation between a quark and an antiquark.⁴

years has been the development of improved actions, that is, improved methods of formulating QCD on the lattice. As in classical field theory, the QCD action is the integral, over space and time, of the Lagrangian density. In lattice calculations, this four-dimensional integral is approximated by summing over discrete lattice points in spacetime.

With substantial computational resources at NSF and DOE national centers during the past three years, lattice gauge theorists have used an improved “staggered fermion” (ISF) action to generate, and make publicly available, a large set of gauge-field configurations (see box 2).⁵ Staggered fermion actions, introduced by John Kogut and Leonard Susskind in 1976, are so called because the algorithm spreads the fermion spins over adjacent lattice points.

The newly available gauge-field configurations include the vacuum-polarization (quark loop) effects of u , d , and s quarks. Several lattice-QCD collaborations, working together,⁶ have recently used these configurations to determine a variety of hadronic quantities to an unprecedented accuracy of 3%. All of those quantities had been measured previously in the laboratory. Figure 2 plots the ratio of the simulated value to the experimental one for each observable. The only inputs were a few experimentally known hadron masses that were used to determine the lattice spacing and the masses of five of the quark flavors. The t quark is too heavy to contribute. The rest is pure prediction.

The left panel shows the result from a widely used quenched approximation (omitting quark-loop effects). The right panel shows what happens when quark loops are included in the calculations. The results show that quark-loop effects are essential; when they are included, the agreement with experiment is encouraging.

The quantities shown in figure 2 demonstrate the predictive capabilities of lattice QCD. The two hadronic decay parameters f_π and f_K , which describe the strong-interaction contribution to the weak decays of the π and K mesons, measure their quark–antiquark wavefunctions at the origin. The particular mass-difference combinations shown in the figure—involving the nucleon, the doubly strange Ξ

baryon, and the ground states of the B_s (strange and bottom flavored) and Y meson families—were chosen because those linear combinations are rather insensitive to a variety of systematic errors.

The flavor-neutral ψ (charmonium) and Y (bottomonium) meson families, quark analogs of positronium, are, respectively, $c\bar{c}$ and $b\bar{b}$ bound states. The figure shows level splittings between different orbital states of these “quarkonium” families. The c and b quarks, with masses of a few GeV, are a thousand times heavier than the u and d .

For the hadronic quantities in figure 2, extrapolation to the physical u and d masses is well under control. But many other important quantities, such as the nucleon mass itself, present greater difficulties. We expect, however, to achieve comparable precision with the nucleon mass, for example, once we have developed the extrapolation procedure for it to the same level of sophistication we already have for the π and K mesons.

As a byproduct of these calculations, one gets a new value for the color fine-structure coupling α_s by combining a nonperturbative lattice determination of energy scales with lattice perturbation theory. This quantity is traditionally calculated at very high energies, where perturbation theory applies. Although the starting energy scale of the new determination is two orders of magnitude lower than that of the perturbative calculation, the two values turn out to agree reassuringly well. For α_s at 91 GeV (the mass of the Z^0 weak boson, the conventional point of comparison), the current lattice QCD calculation gives 0.121 ± 0.003 . The world average from other determinations is 0.117 ± 0.002 .

What advances do we foresee?


► **CKM matrix elements.** An immediate scientific objective is to determine the weak decay constants f_D for the charm-flavored D mesons and f_B for the analogous bottom-flavored B mesons. One needs those constants to extract the CKM matrix elements (box 1) from the experimentally measured meson decay rates. As we have done with f_π and

Box 1. The CKM Matrix Elements

Quantum chromodynamics (QCD) is the strong-interaction part of the standard model that summarizes our current understanding of the most fundamental interactions and particles in nature. Quarks come in six “flavors,” the lightest of which, u (up) and d (down), have masses of just a few MeV and make up the proton and neutron. The other quark flavors, in order of increasing mass, are s (strange), c (charm), b (bottom), and t (top). Each quark flavor comes in the three colors that are the QCD analogs of electric charge. All hadrons are color-neutral combinations of quarks and antiquarks (denoted by overbars).

The fundamental parameters of QCD are the quark masses and the coupling strength α_s . Lattice formulation adds arbitrary computational parameters: the grid spacing a and the total lattice volume. In numerical simulations we can vary all the parameters, but the objective is to choose values that approach those in nature, including, of course, taking the lattice spacing to zero while keeping the volume large enough (a few fm on a side) that it doesn’t distort the physics.

In the standard model, weak interactions, unlike the strong interactions, can change quark flavors. For example, the figure at right shows a bottom-flavored B meson decaying to a charmed D meson (plus positron and neutrino) with the metamorphosis of a $\bar{b}$ to a $\bar{c}$. Alternatively, the B meson can decay by the weak annihilation of a $\bar{b}$ with a u. The fundamental weak-interaction parameters include the measured coupling amplitudes for such conversions of one quark flavor into another (the red dots) with the virtual emission of the W boson that mediates the weak interactions. These amplitudes are the elements of the unitary 3×3 Cabibbo-Kobayashi-Maskawa (CKM) matrix.



One of the big questions is whether there are still more quarks to be discovered. If so, the true CKM matrix has to be bigger than 3×3 , and we would find that the incomplete matrix describing the known quarks is not unitary. An intense experimental and theoretical effort is under way to measure the CKM matrix elements precisely enough to detect such departures from unitarity. But to accomplish that, one has to correct for initial- and final-state strong interactions between quarks. That is where lattice QCD calculations become essential. The gluons that mediate the strong interaction are represented in the figure by wavy lines, sometimes sprouting loops—virtual quark–antiquark loops. The weak decay process occurs in an instant, but strong interactions set the stage and dress the hadronic actors.

f_K , we expect to determine the heavy-flavored meson decay constants to high precision.

Another important strong-interaction parameter, B_K , characterizes the influence of the strong interactions on the remarkable and well-known quantum mixing phenomenon in which a neutral kaon oscillates between the positive strangeness K^0 and negative strangeness $\bar{K}^0$ states. The analogous but less well-measured process involving the B^0 and $\bar{B}^0$ meson states is under current study

at the SLAC BaBar collider and the Belle collider at KEK in Japan (see PHYSICS TODAY, May 2001, page 17). Both measurements also determine elements of the CKM matrix. Study of B-meson mixing and decay is an important part of the Fermilab Tevatron program.

► **The quark–gluon plasma.** The cores of sufficiently dense stars are expected to contain an intriguing phase of matter, predicted by QCD, in which hadrons are so crowded that quarks and gluons are “liberated” and move as if they were in a kind of deconfined plasma. The universe was very likely such a quark–gluon plasma (QGP) for a brief moment after the Big Bang. Such a phase is anticipated to manifest itself at extremely high temperature or density. Experiments under way at Brookhaven’s Relativistic Heavy Ion Collider (RHIC) and at CERN bring heavy nuclei into high-energy collision in hopes of creating the QGP for brief instants (see the article by Thomas Ludlam and Larry McLerran in PHYSICS TODAY, October 2003, page 48).

Although lattice-gauge simulations do not describe the kinematics of





Figure 2. Comparing lattice QCD with experimental values for a variety of hadronic quantities. The ratios of simulated to measured values are shown for lattice calculations with (right) and without (left) quark-loop effects. The quantities are π - and K-meson hadronic decay parameters; mass-difference combinations between the Ξ baryon and the nucleon and between the ground-state B_s and Y mesons; and energy-level splittings between various bound states of the charmed quark and its antiquark (the ψ mesons) and between bound states of the bottom quark and its antiquark (the Y mesons). The quenched approximation (left) is much less expensive to compute, but including quark-loop effects brings substantial improvement. The yellow bands indicate 1% departure from perfect agreement between calculation and measurement. (Adapted from ref. 6.)

Figure 3. Predicted fluctuation of the net strangeness number N_s as a function of temperature T in a hot, flavor-neutral ensemble of quarks and antiquarks. The mean strangeness $\langle N_s \rangle$ vanishes, because there are as many $\bar{s}$ as s quarks. But the mean-square value $\langle N_s^2 \rangle$, which measures the random strangeness fluctuation, does not. The inflection point near 180 MeV in this lattice-QCD calculation is thought to signal the crossover to a quark–gluon plasma. Scaled by T^3 , the mean-square number of strange quarks per unit volume V , plotted here in natural units ($\hbar = c = 1$), is dimensionless. (Adapted from ref. 7.)



a heavy-ion collision, they can predict the equilibrium properties of the QGP: the phase diagram, the equation of state, and fluctuations in particle number. Figure 3 contains the result of a recent lattice-gauge calculation that shows how the fluctuation of net strangeness density increases with temperature in a hot, flavor-neutral ensemble of quarks and antiquarks.⁷ The inflection point near 180 MeV in the figure is taken to indicate the transition to the QGP from a lower-temperature phase of quarks and antiquarks confined in hadrons. Such calculations provide crucial input for phenomenological models of QGP formation and decay.

In a flavor-neutral ensemble, with equal populations of quarks and antiquarks, the net baryon number vanishes. (A quark carries baryon number $+1/3$. For an antiquark, it's $-1/3$.) A long-standing algorithmic challenge has been to find a way of simulating QCD at nonzero net

baryon-number density—for example, in stellar interiors or the “nuclear fragmentation” regions of phase space in relativistic heavy-ion collisions.

Introducing a nonzero chemical potential to push the baryon density away from zero makes the determinant of the quark-action matrix (box 2) complex. Then the weight-

Box 2. How Lattice Calculations Are Done

The successful numerical treatment of QCD uses the Feynman path-integral technique to quantize the field theory.³ For gluons, the numerical problem is then reduced to carrying out a massive multidimensional integration. Quarks, however, present a greater challenge. The essential problem is that quarks are fermions. In the Feynman path integral, they are represented by anticommuting quantities called Grassmann numbers. Fortunately, the fermion integration can be done by hand, leaving only an integration over the ordinary numbers that describe the gluon degrees of freedom. But the quarks make the gluon integrand more complicated. After the quarks are integrated out, any quantity of physical interest—for example, the mass of an elementary particle—is obtained from observables $\mathcal{O}(U)$ that are functions of the gluon field U . Its expectation value is given by

$$\langle \mathcal{O} \rangle = \frac{\int dU \mathcal{O}(U) \exp[-S(U)] \det[M(U)]}{\int dU \exp[-S(U)] \det[M(U)]},$$

which is just a weighted average of the observable function over the multidimensional integration volume. The weight is determined by the purely gluonic action $S(U)$ and the determinant of the quark-action matrix $M(U)$ that describes the motion of the quarks and their interaction with the gluons. In the language of perturbation theory, the quark determinant generates the closed quark loops in Feynman diagrams. These loops represent a polarization of the QCD vacuum. They produce the quark–antiquark pairs whose creation makes it impossible to pry one quark loose from a proton.

The so-called quenched approximation drops the quark determinant in the numerator and denominator of the preceding equation, thereby omitting the quark loops. The determinant can be included, but at a cost that increases by a few orders of magnitude when one decreases the masses of

the u and d quarks toward their experimental values. Thus even when vacuum polarization effects are included, it is standard practice to carry out computations with unphysically large u and d masses and then extrapolate to their physical values.

One of the larger lattices in current practice has a regular grid of 28 points in each of three spatial dimensions and 96 in time, with lattice spacing $a \approx 0.1$ fm and time intervals of a/c . With eight gluon color combinations and four spacetime link directions at each lattice site, that comes to 67×10^6 integration variables needed to describe the gluon field. That's definitely not the place to use Simpson's rule! Instead, one uses a variety of importance-sampling techniques.

The form of the weighted integral in the previous equation lends itself naturally to importance sampling. If the net baryon density is zero and we are careful with the number of quark flavors involved, the weight will be positive definite and it can be treated as a probability. If we sample points U_i for $i = 1, 2, \dots, N$ in the multidimensional integration space according to that probability, the expectation value of $\mathcal{O}(U)$ becomes a simple average of the values that $\mathcal{O}$ assumes on each point in the sample:

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

The statistical error on this expectation value decreases like $1/\sqrt{N}$.

To specify one sample point U_i in the integration space requires giving the value of the gluon field on each of its four lattice links. With the large lattice grid described above, it takes 600 megabytes to specify the complete gauge-field configuration. A common computational strategy is to generate and archive a large sample of such configurations. They then become a resource that can be used subsequently to “measure” a wide variety of different observables.




Figure 4. Circuitous paths linking quarks at neighboring lattice sites improve QCD simulations by effectively smoothing the sharp corners of the four-dimensional spacetime lattice. The diagrams show different terms linking quarks in the ISF action. The arrows represent gluon color matrices on links between adjacent sites. The traditional unimproved action used only term **a**. A chain of arrows is a product of such matrices. The five-link diagram (**c**) involves three directions and the seven-link diagram (**d**) is meant to represent all four spacetime directions. Some paths, like **f**, link next-nearest neighbors.

ing factor in the first equation in box 2 can no longer be interpreted as a probability, and the usual importance-sampling techniques lose their effectiveness. (A similar complication plagues condensed-matter physics calculations for systems in which the conduction-electron occupancy deviates from half filling.) Good progress has been made recently in obtaining results for small net baryon density.⁸

► **Hadronic structure.** The internal structure of the nucleon is fundamental to nuclear physics. A principal scientific goal is to determine quantitatively how quarks and gluons produce the binding and spin of the nucleon. To that end, vigorous experimental programs are in progress at MIT, Thomas Jefferson National Accelerator Facility, SLAC, Fermilab, DESY, and CERN. (See page 9 of this issue.) RHIC will soon be joining the enterprise. The nucleon is easy to study by lattice QCD, because it is the lightest of the three-quark baryons. And determining the static properties of hadrons is natural for lattice QCD. We hope to understand the nucleon's structure through a combination of numerical simulation and experiment.

What we can and can't calculate

Lattice-gauge theorists use the Feynman path-integral technique to quantize the field theory (see box 2). The Feynman approach actually leads us to carry out calculations with an imaginary time coordinate. That feature is standard in statistical quantum mechanics. In fact, the Feynman integration over alternative paths determines the partition function for an ensemble of interacting gluons, quarks, and antiquarks in thermal equilibrium. In lattice-QCD calculations, the temperature is inversely proportional to the duration of the whole lattice volume in imaginary time.

If we keep the imaginary-time duration small, we can study high-temperature features such as the QGP. But if we make the duration large enough, we're simulating a temperature close to zero. Quantities of interest at zero temperature include the masses of a wide variety of hadrons, their decay amplitudes, the quark-antiquark potential, and various static properties of hadrons, such as the internal distributions of charge and magnetization.

Lattice QCD near zero temperature also addresses the complicated structure of the vacuum. The vacuum state of QCD, its zero-temperature ground state, is remarkably rich in structure.⁹ The gluon field fluctuates with twists and turns, tracing out topological knots called instantons. Understanding the ground state is fundamental to understanding QCD.

We cannot, however, calculate everything. Because of its close relationship to statistical thermodynamics, lattice

QCD in its current formulation is unsuited for simulating real-time processes such as multiparticle scattering and the nonequilibrium behavior of the QGP. For such processes we rely on phenomenological models to extrapolate from the domain where lattice QCD does work.

The principal computational challenges faced by lattice QCD are reducing discretization errors and extrapolating down to the small physical masses of the u and d quarks:

► **Discretization errors.** Representing spacetime by a regular grid of discrete points introduces artifacts that become small as the lattice spacing a is decreased. However, the computational cost (the number of requisite computer operations) grows very steeply with decreasing a —something like a^{-7} or a^{-8} —when the computation includes quark-loop effects. In the end, one must extrapolate to zero lattice spacing. To improve the accuracy of that extrapolation, we place a high premium on finding improved algorithms that reduce discretization artifacts. Currently, large-scale computations with improved action algorithms work with a lattice spacing as small as 0.09 fm.


► **Light-quark masses.** Computational cost, at fixed a , grows approximately as the inverse square of the quark masses in question. That makes it too expensive to let the u and d quarks be as light as they are in nature. (The proton is a hundred times heavier than the sum of its three valence quarks.) In the limit of vanishing quark masses, QCD has a special "chiral" symmetry from which one can define a perturbative expansion in the small quark mass and thus anchor the extrapolation down from the unphysically high masses used in the lattice calculations. So if we can simulate QCD in the regime where chiral perturbation theory applies, we can extrapolate from lattice QCD to the small u and d masses with some confidence. In current large-scale lattice simulations, it is feasible to take the u and d masses as low as three times their physical values. That's well within the range of chiral perturbation theory and well below the quark-mass values that had to be used in lattice calculations before the ISF improvement.

How improvement is accomplished

The key to the recent advances in lattice-gauge theory has been the development of improved lattice actions for describing the motion and interaction of quarks and gluons. The improvements refine the discretization of the quark action. The ISF action is the most extensively exploited of these algorithmic improvements.⁵

Let us examine one of the steps in the improvement process. In a lattice simulation, the simplest term describing the interaction between quarks and gluons involves the inner product of an antiquark field at one lattice site and the quark field at a neighboring site. To maintain gauge invariance, however, one also has to include in that product the gluon matrix U on the link joining the two lattice sites.

Figure 5. Calculated nucleon mass as a function of lattice spacing a for a variety of lattice quark-action algorithms. The smaller the slope, the better the algorithm. The blue points indicate results with the unimproved staggered-fermion action devised in 1976 and used well into the 1990s. The red points indicate an improved version of Kenneth Wilson's original quark action used in the late 1990s. The ISF action, represented by the green points, clearly shows the least dependence on lattice spacing. The nucleon mass, in the limit of vanishing a , comes out about 300 MeV too high because all these calculations, for simplicity of comparison, used unphysically high u and d masses and ignored quark-loop effects. (Adapted from ref.10.)



The improvement schemes use not only the shortest path to connect adjacent sites but also a combination of products of gluon matrices along longer, more circuitous paths between immediate and more distant neighbors, as shown in figure 4. With the correct linear combination of such paths, we can reduce discretization errors, in effect, by smoothing the sharp corners of the lattice. We remove errors proportional to a^2 . So we're left with errors that scale like a^4 and $a^2\alpha_s$. The added complexity of the circuitous paths increases the computational cost by a factor of two or three, but the greater cost is handsomely repaid in better accuracy at modest lattice spacing.

One way to measure the effectiveness of the improved formulation is to observe how calculated quantities vary with lattice spacing. For example, figure 5 compares the lattice-spacing dependence of the ISF-action calculation of the nucleon mass with calculations that use older action algorithms.¹⁰ To avoid the added complication of extrapolation, all of these lattice calculations used unrealistically


heavy quarks. Therefore the nucleon mass, in the limit of vanishing a , comes out about 300 MeV too heavy. But what matters in this comparison of quark-action algorithms is the sensitivity to lattice spacing. The smaller the slope, the better the improvement. Clearly, the ISF action does better in this test.

Given sufficient computational resources (see box 3), the future prospects for high-precision lattice calculations are excellent. The ISF action is only one of several improved actions currently being investigated.¹¹ Others make even further improvements.¹² Their formulation is more

Box 3. Special-Purpose Computers for Lattice QCD

The rate-limiting computational problem in lattice QCD is the solution of large linear systems. These systems are well-suited for massively parallel computers with fast communication, high local memory bandwidth, and relatively small memory per processor. Thus it becomes cost-effective to consider special-purpose machines for lattice QCD calculations. These can be custom-built machines like the QCDOC computer designed by researchers at Columbia University, Brookhaven, the Japan Institute of Physical and Chemical Research (RIKEN), and in Britain. Alternatively, they can be specially configured PC clusters built mostly from commercial parts.

The photo at right shows a QCDOC daughter board with two nodes, each with a gigaflop-per-second double-precision IBM 440 PowerPC unit, 4 megabytes of memory, and 12-way bidirectional communication.¹⁵ Sustained 10-Tflop/s QCD calculation requires 20 000 such nodes. European examples include the custom-built Italian apeNEXT computer.¹⁶ Currently, special-purpose lattice-QCD computers can be built at a cost of about \$1 per Mflop/s.



Practically the entire US lattice gauge theory community, a group of about 50 scientists, has been working together for the past four years to develop the computational infrastructure it needs to study the wide variety of high-energy and nuclear physics problems that require numerical QCD simulation. In addition to the QCDOC, large PC clusters are under development at Fermilab and Jefferson Laboratory. The US Department of Energy's SciDAC advanced computing program is providing funds for the development of community software to facilitate efficient use of these powerful computers.¹⁷

complicated; their importance-sampling algorithms are not yet ready for large-scale computation. But we expect significant gains in the near future.

An altogether different approach emphasizes a more accurate treatment of chiral symmetry at nonzero lattice spacing that provides much better control of the extrapolation to the physical u and d quark masses.^{13,14} That approach is considerably more expensive than the ISF action, but some promising work is in progress. It's important to pursue alternatives to ISF for two reasons. First, the implications of some of the approximations used in the ISF simulation are not completely understood. Second, systematic uncertainties in the chirally accurate methods and the ISF methods are sufficiently different that they provide good cross checks.

A host of interesting fundamental physics questions await application of the new tools. These include stringent tests of the standard model of particle physics, full characterization of the quark–gluon plasma, quantitative determination of the structure of the nucleon, and the prediction of masses and decay channels for observed and conjectured exotic hadronic states, including pentaquarks, purely gluonic particles called glueballs, and quark–gluon hybrids (see PHYSICS TODAY, September 2003, page 19).

References

1. K. G. Wilson, *Phys. Rev. D* **10**, 2445 (1974).
2. M. Creutz, L. Jacobs, C. Rebbi, *Phys. Rev. D* **20**, 1915 (1979); M. Creutz, *Phys. Rev. Lett.* **43** 553 (1979) [Erratum **43**, 890 (1979)].
3. I. Montvay, G. Münster, *Quantum Fields on a Lattice*, Cambridge U. Press, New York (1994); J. Smit, *Introduction to Quantum Fields on a Lattice*, Cambridge U. Press, New York (2002). For links to lattice QCD collaborations, see http://www.ph.ed.ac.uk/ukqcd/community/lattice_links.html.
4. C. Bernard et al. (MILC collaboration), *Nucl. Phys. Proc. Suppl.* **119**, 598 (2003).
5. S. Naik, *Nucl. Phys.* **B316**, 238 (1989); C. Bernard et al. (MILC collaboration), *Phys. Rev. D* **58**, 014503 (1998); J. F. Lagaë, D. K. Sinclair, *Phys. Rev. D* **59**, 014511 (1999); G. P. Lepage, *Phys. Rev. D* **59**, 074502 (1999); K. Orginos, D. Toussaint, R. L. Sugar (MILC collaboration), *Phys. Rev. D* **60**, 054503 (1999). The ISF action is also known as the Asqtad action.
6. C. T. H. Davies, E. Follana, A. Gray, G. P. Lepage, Q. Mason, M. Nobes, J. Shigemitsu, H. D. Trotter, 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, <http://arXiv.org/abs/hep-lat/0304004> and *Phys. Rev. Lett.* (in press).
7. C. Bernard et al. (MILC collaboration), <http://arXiv.org/abs/hep-lat/0309118>.
8. Z. Fodor, S. D. Katz, *J. High Energy Phys.* **0203**, 014 (2002); C. R. Allton et al., *Phys. Rev. D* **66**, 074507 (2002); P. de Forcrand, O. Philipsen, *Nucl. Phys. B* **642**, 290 (2002).
9. C. Bernard et al. (MILC collaboration), *Nucl. Phys. Proc. Suppl.* **119**, 769 (2003).
10. C. Bernard et al. (MILC collaboration), *Phys. Rev. D* **61**, 111502 (2000) and refs. therein.
11. U. M. Heller, F. Karsch, B. Sturm, *Phys. Rev. D* **60**, 114502 (1999).
12. A. Hasenfratz, F. Knechtli, *Phys. Rev. D* **64**, 034504 (2001); C. Gattringer et al. (BGR collaboration), <http://arXiv.org/abs/hep-lat/0307013>.
13. R. Narayanan, H. Neuberger, *Phys. Lett. B* **302**, 62 (1993).
14. D. Kaplan, *Phys. Lett. B* **288**, 342 (1992); V. Furman, Y. Shamir, *Nucl. Phys. B* **439**, 54 (1995).
15. P. A. Boyle et al. (QCDOC collaboration), <http://arXiv.org/abs/hep-lat/0309096>.
16. F. Bodin et al., eConf **C0303241**, THIT005 (2003), available at <http://arXiv.org/abs/hep-lat/0306018>.
17. See <http://www.lqcd.org>. ■

Circle number 30 on Reader Service Card