Physics Today

QCD Made Simple

Frank Wilczek

Citation: Physics Today **53**(8), 22 (2000); doi: 10.1063/1.1310117
View online: http://dx.doi.org/10.1063/1.1310117
View Table of Contents: http://scitation.aip.org/content/aip/magazine/physicstoday/53/8?ver=pdfcov
Published by the AIP Publishing

QCD MADE SIMPLE

Quantum chromodynamics is conceptually simple. Its realization in nature, however, is usually very complex. But not always.

Frank Wilczek

Quantum chromodynamics, familiarly called QCD, is the modern theory of the strong interaction.[1] Historically its roots are in nuclear physics and the description of ordinary matter—understanding what protons and neutrons are and how they interact. Nowadays QCD is used to describe most of what goes on at high-energy accelerators.

Twenty or even fifteen years ago, this activity was commonly called "testing QCD." Such is the success of the theory, that we now speak instead of "calculating QCD backgrounds" for the investigation of more speculative phenomena. For example, discovery of the heavy W and Z bosons that mediate the weak interaction, or of the top quark, would have been a much more difficult and uncertain affair if one did not have a precise, reliable understanding of the more common processes governed by QCD. With regard to things still to be found, search strategies for the Higgs particle and for manifestations of supersymmetry depend on detailed understanding of production mechanisms and backgrounds calculated by means of QCD.

Quantum chromodynamics is a precise and beautiful theory. One reflection of this elegance is that the essence of QCD can be portrayed, without severe distortion, in the few simple pictures at the bottom of the box on the next page. But first, for comparison, let me remind you that the essence of quantum electrodynamics (QED), which is a generation older than QCD, can be portrayed by the single picture at the top of the box, which represents the interaction vertex at which a photon responds to the presence or motion of electric charge.[2] This is not just a metaphor. Quite definite and precise algorithms for calculating physical processes are attached to the Feynman graphs of QED, constructed by connecting just such interaction vertices.

In the same pictorial language, QCD appears as an expanded version of QED. Whereas in QED there is just one kind of charge, QCD has three different kinds of charge, labeled by "color." Avoiding chauvinism, we might choose red, green, and blue. But, of course, the color charges of QCD have nothing to do with physical colors. Rather, they have properties analogous to electric charge. In particular, the color charges are conserved in all physical processes, and there are photon-like massless particles, called color gluons, that respond in appropriate ways to the presence or motion of color charge, very similar to the way photons respond to electric charge.

FRANK WILCZEK *is the J. Robert Oppenheimer Professor of Physics at the Institute for Advanced Study in Princeton, New Jersey. Next month he moves to Cambridge, Massachusetts, to take up the Herman Feshbach Chair of Physics at the Massachusetts Institute of Technology.*

Quarks and gluons

One class of particles that carry color charge are the quarks. We know of six different kinds, or "flavors," of quarks—denoted u, d, s, c, b, and t, for: up, down, strange, charmed, bottom, and top. Of these, only u and d quarks play a significant role in the structure of ordinary matter. The other, much heavier quarks are all unstable. A quark of any one of the six flavors can also carry a unit of any of the three color charges. Although the different quark flavors all have different masses, the theory is perfectly symmetrical with respect to the three colors. This color symmetry is described by the Lie group SU(3).

Quarks are spin-1/2 point particles, very much like electrons. But instead of electric charge, they carry color charge. To be more precise, quarks carry *fractional* electric charge ($+2e/3$ for the u, c, and t quarks, and $-e/3$ for the d, s, and b quarks) in addition to their color charge.

For all their similarities, however, there are a few crucial differences between QCD and QED. First of all, the response of gluons to color charge, as measured by the QCD coupling constant, is much more vigorous than the response of photons to electric charge. Second, as shown in the box, in addition to just responding to color charge, gluons can also change one color charge into another. All possible changes of this kind are allowed, and yet color charge is conserved. So the gluons themselves must be able to carry unbalanced color charges. For example, if absorption of a gluon changes a blue quark into a red quark, then the gluon itself must have carried one unit of red charge and minus one unit of blue charge.

All this would seem to require $3 \times 3 = 9$ different color gluons. But one particular combination of gluons—the color-SU(3) singlet—which responds equally to all charges, is different from the rest. We must remove it if we are to have a perfectly color-symmetric theory. Then we are left with only 8 physical gluon states (forming a color-SU(3) octet). Fortunately, this conclusion is vindicated by experiment!

The third difference between QCD and QED, which is the most profound, follows from the second. Because gluons respond to the presence and motion of color charge *and* they carry unbalanced color charge, it follows that gluons, quite unlike photons, respond directly to one another. Photons, of course, are electrically neutral. Therefore the laser sword fights you've seen in *Star Wars* wouldn't work. But it's a movie about the future, so maybe they're using color gluon lasers.

We can display QCD even more compactly, in terms of its fundamental equations (figure 1). You should not necessarily be too impressed by that. After all, Richard Feynman showed that you could write down the Equation of the Universe in a single line: $U = 0$, where U, the total *unworldliness,*[3] is a definite function. It's the sum of contributions from *all* the laws of physics:

$$U = U_{\text{Newton}} + U_{\text{Gauss}} + \ldots,$$

where, for instance, $U_{\text{Newton}} = (\mathbf{F} - m\mathbf{a})^2$ and $U_{\text{Gauss}} = (\nabla \cdot \mathbf{E} - \rho)^2$.

So we can capture all the laws of physics we know, and all the laws yet to be discovered, in this one unified equation. But it's a complete cheat, of course, because there is no useful algorithm for unpacking U, other than to go back to its component parts. The equations of QCD, displayed in figure 1, are very different from Feynman's satirical unification. Their complete content is out front, and the algorithms that unpack them flow from the unambiguous mathematics of symmetry.

A remarkable feature of QCD, which we see in figure 1, is how few adjustable parameters the theory needs. There is just one overall coupling constant g and six quark-mass parameters m_j for the six quark flavors. As we shall see, the coupling strength is a relative concept; and there are many circumstances in which the mass parameters are not significant. For example, the heavier quarks play only a tiny role in the structure of ordinary matter. Thus QCD approximates the theoretical ideal: From a few purely conceptual elements, it constructs a wealth of physical consequences that describe nature faithfully.[4]

Describing reality

At first sight it appears outrageous to suggest that the equations of figure 1 or, equivalently, the pictures in the box, can describe the real world of the strongly interacting particles. None of the particles that we've actually seen appear in the box, and none of the particles that appear in the box has ever been observed. In particular, we've never seen particles carrying fractional electric charge, which we nonetheless ascribe to the quarks. And certainly we haven't seen anything like gluons—massless particles mediating long-range strong forces. So if QCD is to describe the world, it must explain why quarks and gluons cannot exist as isolated particles. That is the so-called confinement problem.

Besides confinement, there is another qualitative difference between the observed reality and the fantasy world of quarks and gluons. This difference is quite a bit more subtle to describe, but equally fundamental. I will not be able to do full justice to the phenomenological arguments here, but I can state the essence of the problem in its final, sanitized theoretical form. The phenomenology indicates that if QCD is to describe the world, then the u and d quarks must have very small masses. But if these quarks do have very small masses, then the equations of QCD possess some additional symmetries, called chiral symmetries (after *chiros*, the Greek word for *hand*). These symmetries allow separate transformations among the right-handed quarks (spinning, in relation to their motion, like ordinary right-handed screws) and the left-handed quarks.

But there is no such symmetry among the observed strongly interacting particles; they do not come in opposite-parity pairs. So if QCD is to describe the real world, the chiral symmetry must be spontaneously broken, much as rotational symmetry is spontaneously broken in a ferromagnet.

Clearly, it's a big challenge to relate the beautifully simple concepts that underlie QCD to the world of observed phenomena. There have been three basic approaches to meeting this challenge:

▷ The first approach is to take the bull by the horns and just solve the equations. That's not easy. It had better not be too easy, because the solution must exhibit properties (confinement, chiral-symmetry breaking) that are very different from what the equations seem naively to suggest, and it must describe a rich, complex phenomenology. Fortunately, powerful modern computers have made it possible to calculate a few of the key predictions of QCD directly. Benchmark results are shown in figure 2, where the calculated masses[5] of an impressive range of hadrons

QED and QCD in Pictures.

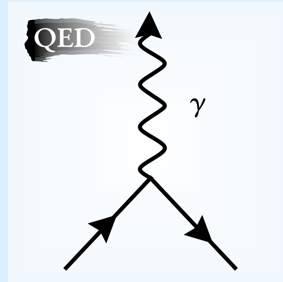

The physical content of quantum electrodynamics is summarized in the algorithm that associates a probability amplitude with each of its Feynman graphs, depicting a possible process in spacetime. The Feynman graphs are constructed by linking together interaction vertices of the type at left, which represents a point charged particle (lepton or quark) radiating a photon. To get the amplitude, one multiplies together a kinematic "propagator" factor for each line and an interaction factor for each vertex. Reversing a line's direction is equivalent to replacing a particle by its antiparticle.

Quantum chromodynamics can be similarly summarized, but with a more elaborate set of ingredients and vertices, as shown below. Quarks (antiquarks) carry one positive (negative) unit of color charge. Linear superpositions of the 9 possible combinations of gluon colors shown below form an SU(3) octet of 8 physical gluon types.

A qualitatively new feature of QCD is that there are vertices describing direct interactions of color gluons with one another. Photons, by contrast, couple only to electric charge, of which they carry none themselves.

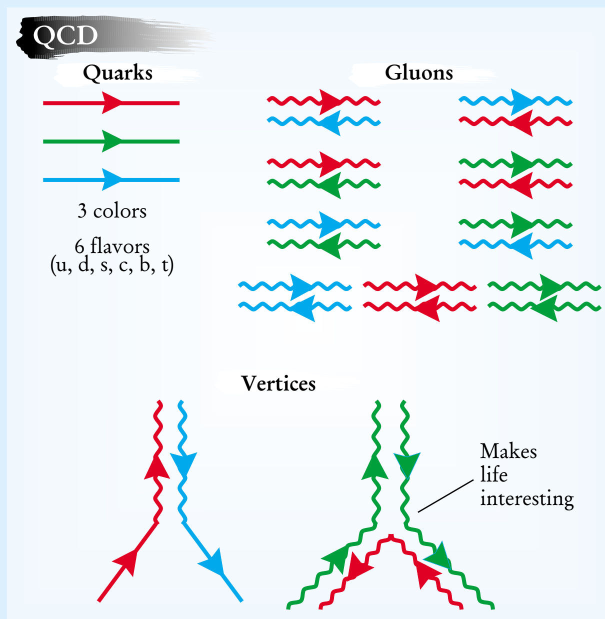

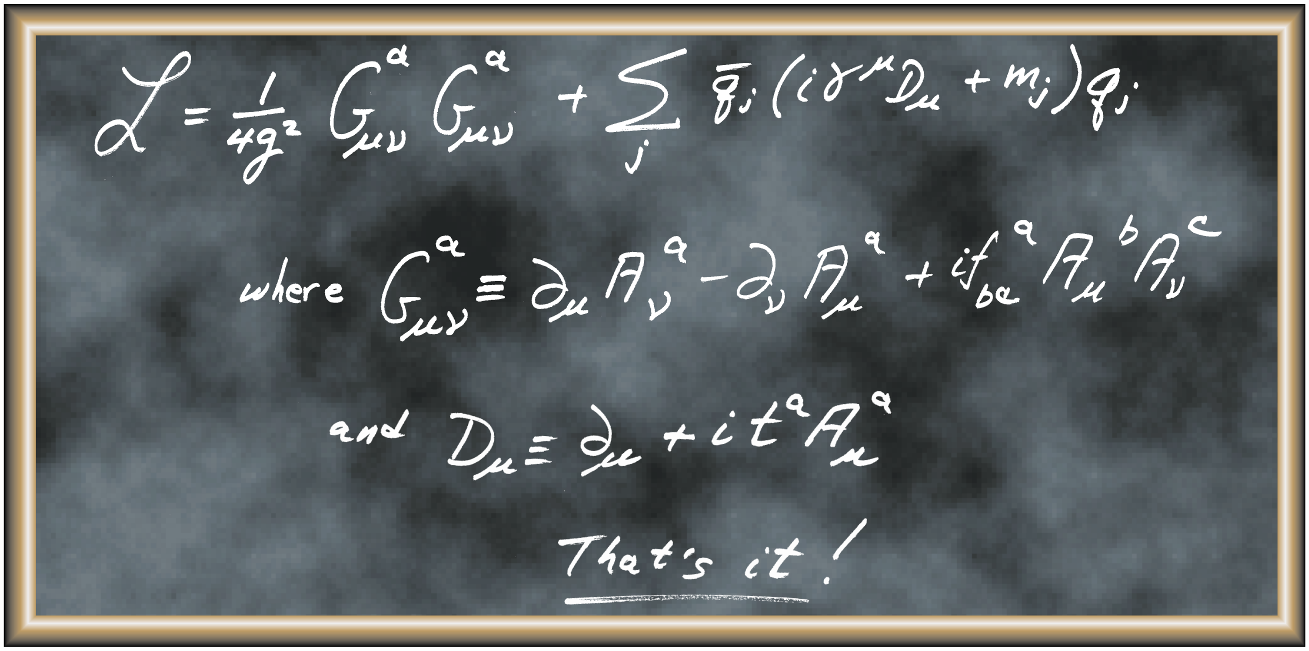

FIGURE 1. THE QCD LAGRANGIAN $\mathcal{L}$ displayed here is, in principle, a complete description of the strong interaction. But, in practice, it leads to equations that are notoriously hard to solve. Here m_j and q_j are the mass and quantum field of the quark of jth flavor, and A is the gluon field, with spacetime indices μ and ν and color indices a, b, c. The numerical coefficients f and t guarantee SU(3) color symmetry. Aside from the quark masses, the one coupling constant g is the only free parameter of the theory.

are compared with their measured values. The agreement is encouraging.

Such calculations clearly demonstrate that confinement and chiral-symmetry breaking are consequences of solving the equations of QCD. The calculations show us no massless gluons, nor any fractionally charged particles, nor the enlarged multiplets that would indicate unbroken chiral symmetry. Just the observed particles, with the right properties—neither more nor less.

While these and other massive numerical calculations give impressive and useful results, they are not the end of all desire. There are many physically interesting questions about QCD for which the known numerical techniques become impractical. Also, it is not entirely satisfying to have our computers acting as oracles, delivering answers without explanations.

▷ The second approach is to give up on solving QCD itself, and to focus instead on models that are simpler to deal with, but still bear some significant resemblance to the real thing. Theorists have studied, for example, QCD-like models in fewer dimensions, or models incorporating supersymmetry or different gauge groups, and several other simplified variants. Many edifying insights have been obtained in this way. By their nature, however, such modelistic insights are not suited to hard-nosed confrontation with physical reality.

▷ The third approach, which is the subject of the rest of this article, is to consider physical circumstances in which the equations somehow become simpler.

Extreme virtuality

The most fundamental simplification of QCD is illustrated in figure 3. There we see, on the left, the jet-like appearance of collision events in which strongly interacting particles (hadrons) are produced in electron–positron annihilations at high energy. One finds many particles in the final state, but most of them are clearly organized into a few collimated "jets" of particles that share a common direction.[6] In about 90% of these hardron-producing events, there are just two jets, emerging in opposite directions. Occasionally—in about 9% of the hadronic final states—one sees three jets.

Compare those multiparticle hadronic events to collisions in which leptons, say muons, are produced. In that case, about 99% of the time one observes simply a muon and an antimuon, emerging in opposite directions. But occasionally—in about 1% of the muonic final states—a photon is emitted as well.

If history had happened in a different order, the observation of jet-like hadronic final states would surely have led physicists to propose that they manifest underlying phenomena like those displayed on the right-hand side of figure 3. Their resemblance to leptonic scattering and QED would be too striking to ignore.

Eventually, by studying the details of how energy was apportioned among the jets, and the relative probabilities of different angles between them, the physicists would have deduced directly from experimental data that there are light spin-1/2 and massless spin-1 objects lurking beneath the appearances, and how these covert objects couple to one another. By studying the rare 4-jet events, they could even have learned about the coupling of the spin-1 particles to each other. So all the basic couplings we know in QCD might have been inferred, more or less directly, from experiment. But there would still be one big puzzle: Why are there jets, rather than simply particles?

The answer is profound, and rich in consequences. It is that the strength with which gluons couple depends radically on their energy and momentum. "Hard" gluons, which carry a lot of energy and momentum, couple weakly; whereas the less energetic "soft" gluons, couple strongly. Thus, only rarely will a fast-moving colored quark or gluon emit "radiation" (a gluon) that significantly redirects the flow of energy and momentum. That explains the collimated flows one sees in jets. On the other hand, there can be a great deal of soft radiation, which explains the

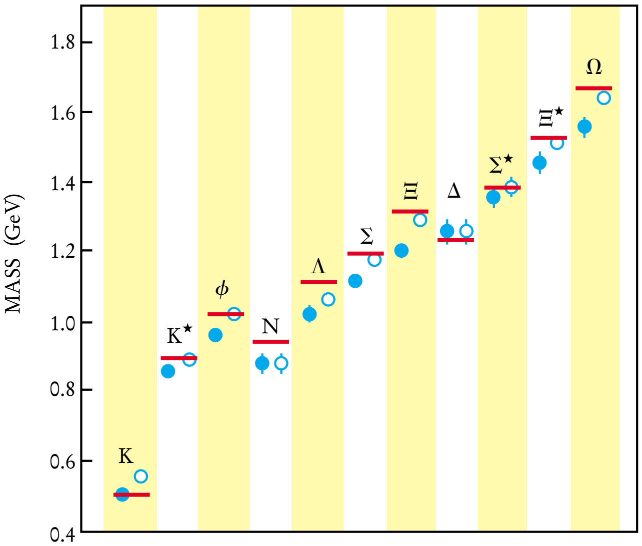

FIGURE 2. MASS SPECTRUM of mesons and baryons, as predicted by QCD and intensive number-crunching.[5] Only two adjustable parameters went into this calculation: the coupling strength and the strange-quark mass (by way of the K or ϕ meson, both of which incorporate strange quarks). The up- and down-quark masses can be approximated by zero. Solid dots are for calculations tied to the measured K mass; open dots are tied to the ϕ mass. The agreement with the measured masses (red lines) is at the 10% level. Such calculations are improving as computer power and techniques get better.

abundant particle content of the jets. So, in a rigorous and very tangible sense, we really do get to see the quarks and gluons—but as flows of energy, not individual particles.

We refer to the phenomenon of weak coupling for hard gluons but strong coupling for soft gluons as "asymptotic freedom."[7] Despite its whimsical name, the concept is embodied in precise equations. It allows us to make quantitative predictions of how often hard-radiation events occur in strong-interaction processes of many different kinds, at different energies. As we see in figure 4, there is by now a plenitude of direct evidence for the central prediction that the coupling strength of gluons *decreases* with increasing energy and momentum.[8] Note that several of the individual points in the figure summarize hundreds of independent measurements, all of which must be—and are—fitted with only one adjustable parameter (the quark–gluon coupling measured at the Z-boson mass).

The actual history was different. The need for asymptotic freedom in describing the strong interaction was deduced from much more indirect clues, and QCD was originally proposed as the theory of the strong interaction because it is essentially the unique quantum field theory having the property of asymptotic freedom.[9] From these ideas, the existence of jets, and their main properties, were predicted before their experimental discovery.[5]

High temperature QCD

The behavior of QCD at high temperature is of obvious interest. It provides the answer to a childlike question: What happens if you keep making things hotter and hotter? It also describes the behavior of matter at crucial stages just after the Big Bang. And it is a subject that can be investigated experimentally with high-energy collisions between heavy nuclei. (See PHYSICS TODAY, May, page 20.) Brookhaven National Laboratory's Relativistic Heavy Ion Collider, where experiments are just getting under way, will be especially devoted to this kind of physics. (See figure 5.)

To avoid confusion, I should state that, when I discuss high-temperature QCD in this article, I'm assuming that the net baryon density (quarks *minus* antiquarks) is very small. Conversely, when I discuss high-density QCD, I mean a high net density of quarks at low temperature, but well above the ordinary quark density of cold nuclear matter. Temperature and net baryon density are generally taken as the two independent variables of the phase diagram for hadronic matter.

Asymptotic freedom implies that QCD physics gets *simpler* at very high temperature. That would seem implausible if you tried to build up the high-temperature phase by accounting for the production and interaction of all the different mesons and baryon resonances that are energetically accessible at high temperature. Hoping to bypass this forbidding

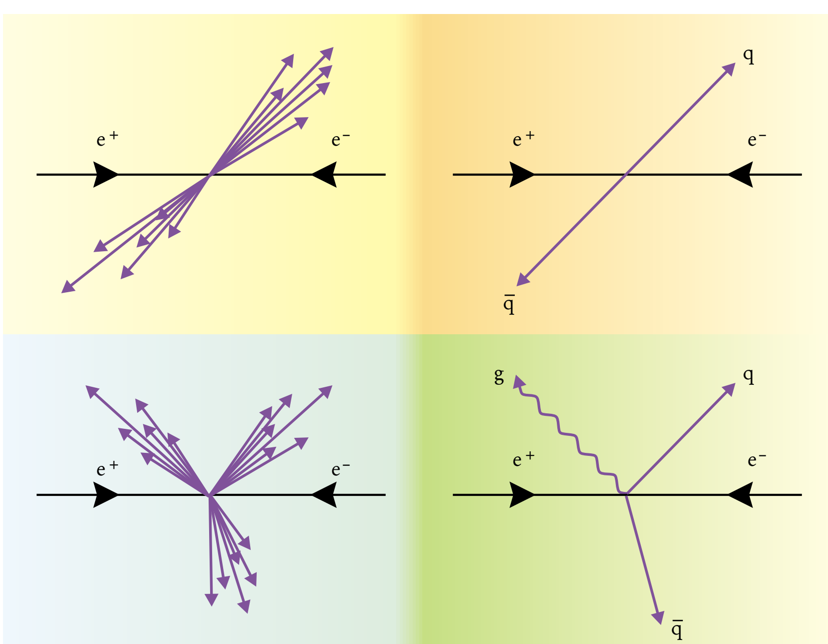

FIGURE 3. IN HIGH-ENERGY e^+e^- annihilations into strongly interacting particles, the many-particle final state is observed (left) to consist of two or occasionally three (or, very rarely, four or more) "jets" of particles leaving the collision in roughly the same directions. QCD predicts their production rates and angular and energy distributions by assuming that (right) a single primary quark or gluon underlies each jet. The jets are explained by asymptotic freedom, which tells us that the probability is small for emitting a quark or gluon that drastically alters the flow of energy and momentum.

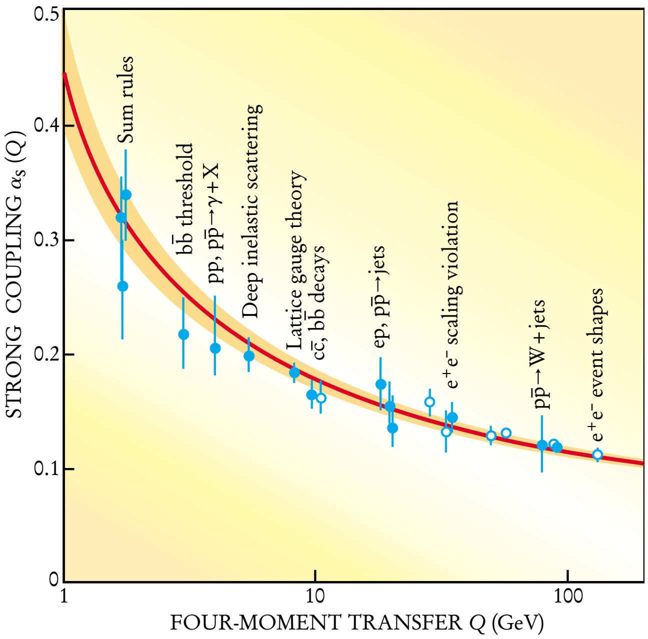

Figure 4. The running coupling "constant" α_s for the strong interaction is predicted by QCD to decrease with increasing energy and momentum. That's asymptotic freedom. The red curve is the predicted dependence of α_s on Q, the magnitude of the four-momentum transfer at a QCD vertex. An empirical input is the measured coupling of a quark pair to a virtual gluon at the Z boson mass; the orange swath reflects its uncertainty. The theory yields excellent agreement with a great variety of experiments,[14] shown by the data points and labels. The open points are results based on the general shapes of many-particle final states in momentum space.

mess, we invoke a procedure that is often useful in theoretical physics. I call it the Jesuit Stratagem, inspired by what I'm told is a credal tenet of the Order: "It is more blessed to ask forgiveness than permission." The stratagem tells you to make clear-cut simplifying assumptions, work out their consequences, and check to see that you don't run into contradictions.

In this spirit we tentatively assume that we can describe high-temperature QCD starting with free quarks and gluons. In an ideal (noninteracting) gas of quarks, antiquarks, and gluons at high temperature, most of the energy and pressure will be contributed by particles with large energy and momentum. How do interactions affect these particles? Well, significantly deflecting such a particle requires an interaction with large momentum transfer. But such interactions are rare because, as asymptotic freedom tells us, they are governed by rather weak coupling. So interactions do not really invalidate the overall picture. To put it another way, if we treat the hadron jets generated by quarks, antiquarks, or gluons as quasiparticles "dressed" in hadronic garb, then we have a nearly ideal gas of quasiparticles. So it seems that ignoring the interactions was a valid starting point. The stratagem has succeeded.

Remarkably, the thermodynamic behavior of QCD as a function of temperature is another one of those things that can be calculated directly from the equations, using powerful computers.[10] Figure 6 shows the qualitative expectations dramatically vindicated. At "low" temperatures ($\lesssim 150$ MeV or 1.5×10^{12} K), the only important particles are the spinless pi mesons: π^+, π^-, and π^0. They represent 3 degrees of freedom. But from a quark–gluon description we come to expect many more degrees of freedom, because there are 3 flavors of light spin-1/2 quarks, each of which comes in 3 colors. If you then include 2 spin orientations, antiquarks, and 8 gluons, each with 2 polarization states, you end up with 52 degrees of freedom in place of the 3 for pions. So we predict a vast increase in the energy density, at a given temperature, as you go from a hadron gas to a quark–gluon plasma. And that is what the calculations displayed in figure 6 show.

What about real experiments? Unfortunately our only access to the quark–gluon plasma is through the production of tiny, short-lived nuclear fireballs, of which we detect only the debris. Interpreting the data requires complicated modeling. In the quest for evidence of the quark–gluon plasma, there are two levels to which one might aspire. At the first level, one might hope to observe phenomena that are very difficult to interpret from a hadronic perspective but have a simple qualitative explanation based on quarks and gluons. Several such effects have been observed by the CERN heavy-ion program in recent years.[11] But there is a second, more rigorous level that remains a challenge for the future. Using fundamental aspects of QCD theory, similar to those I discussed in connection with jets, one can make quantitative predictions for the emission of various kinds of "hard" radiation from a quark–gluon plasma. We will not have done justice to the concept of a weakly interacting plasma of quarks and gluons until some of these predictions are confirmed by experiment.

High density QCD

The behavior of QCD at large net baryon density (and low temperature) is also of obvious interest. It answers yet another childlike question: What will happen when you keep squeezing things harder and harder? It is also interesting for the description of neutron star interiors. But perhaps the most interesting and surprising thing about QCD at high density is that, by thinking about it, one discovers a fruitful new perspective on the traditional problems of confinement and chiral-symmetry breaking.

Why might we hope that QCD simplifies in the limit of large density? Again we use the Jesuit Stratagem. Assume we can neglect interactions. Then, to start with, we'll have large Fermi surfaces for all the quarks. (The Fermi surface bounds the smallest momentum-space volume into which you can pack all those fermions, even at zero temperature.) This means that the active degrees of freedom—the excitations of quarks near the Fermi surface—have large energy and momentum. And so we might be tempted to make essentially the same argument we used for the high-temperature, low-density regime and declare victory once again.

On further reflection, however, we find this argument too facile. For one thing, it doesn't touch the gluons, which are, after all, spin-1 *bosons*. So they are in no way constrained by the Pauli exclusion principle, which blocks the excitation of low-momentum quarks. The low-momentum gluons interact strongly, and because they were the main problem all along, it is not obvious that going to high density really simplifies things very much.

A second difficulty appears when we recall that the Fermi surfaces of many condensed-matter systems at low temperature are susceptible to a pairing instability that drastically changes their physical properties. This phenomenon underlies both superconductivity and the superfluidity of helium-3. It arises whenever there is an effec-

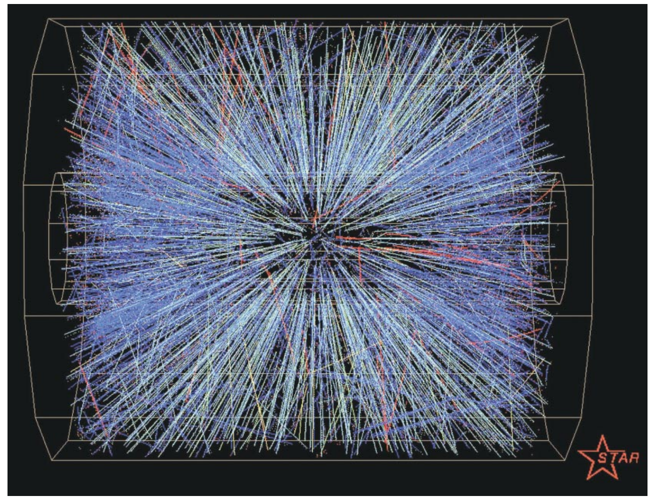

FIGURE 5. ONE OF THE FIRST COLLISIONS between high-energy gold nuclei at Brookhaven's new Relativistic Heavy Ion Collider was recorded by the Star detector facility in June. In this reconstructed side view of the detector, the two 28 GeV-per-nucleon beams of gold nuclei enter from left and right, and collide at the center. About a thousand charged tracks were recorded emanating from this one collision. Significantly higher multiplicities are expected as RHIC works up to its design beam energy of 100 GeV-per-nucleon.

tive attraction between particles on opposite sides of the Fermi surface. As elucidated by John Bardeen, Leon Cooper, and Robert Schrieffer, even an arbitrarily weak attraction can, in principle, cause a drastic restructuring of the ground state.

A nominally small perturbation can have such a big effect because we're in the realm of degenerate perturbation theory. Low-energy excitation of pairs of particles on opposite sides of the Fermi surface, with total momentum zero, can be scattered into one another. By orchestrating a coherent mixture of such excitations, all pulling in the same direction, the system gains an energy advantage.

In condensed-matter physics, the occurrence of superconductivity is a difficult and subtle affair. That's because the fundamental interaction between electrons is Coulomb repulsion. In the classic metallic superconductors, an effective attraction arises from subtle retardation effects involving phonons. In the cuprate superconductors, the cause is still obscure.

In QCD, by contrast, the occurrence of what we might call "color superconductivity" is a relatively straightforward phenomenon.[12] That's because the fundamental interaction between two quarks, unlike that between two electrons, is already attractive! One can see this by a group-theoretical argument: Quarks form triplet representations of color SU(3). A pair of quarks, in the antisymmetric color state, form an antitriplet. So when two quarks are brought together, the effective color charge is reduced by a factor of two compared to when they were separated. The color flux emerging from them is reduced, lessening the energy in the color field. That implies an attractive force. So we should consider carefully what color superconductivity can do for us.

Two of the central phenomena of ordinary superconductivity are the Meissner effect and the energy gap. The Meissner effect is the inability of magnetic fields to penetrate far into the body of a superconductor. Supercurrents arise to cancel them out. Electric fields are, of course, also screened by the motion of charges. Thus electromagnetic fields in general become short-range. Effectively it appears as if the photon has acquired a mass. Indeed that is just what emerges from the equations. We can therefore anticipate that in a color superconductor, gluons will acquire mass. That's very good news, because it removes our problem with the low energy–momentum gluons.

The energy gap means that it costs a finite amount of energy to excite electrons from their superconducting ground state. That's quite unlike what we had for the free Fermi surface. So the original pairing instability, having run its course, is no longer present.

Now with both the sensitivity to small perturbations (pairing instability) and the bad actors (soft gluons) under control, the remaining effects of interactions really are small and under good theoretical control. Once again, the Jesuit Stratagem has served us well.

Color–flavor locking

The simplest and most elegant form of color superconductivity is predicted for a slightly idealized version of real-world QCD in which we imagine there are exactly three flavors of massless quarks: u, d, and s. The strange quark is in fact much lighter than c, b, or t. And anyway, neglecting quark masses is an excellent approximation at extremely high density.

Here we discover the remarkable phenomenon of color–flavor locking.[13] Ordinarily the perfect symmetry among different quark colors is quite distinct and separate from the imperfect symmetry among different quark

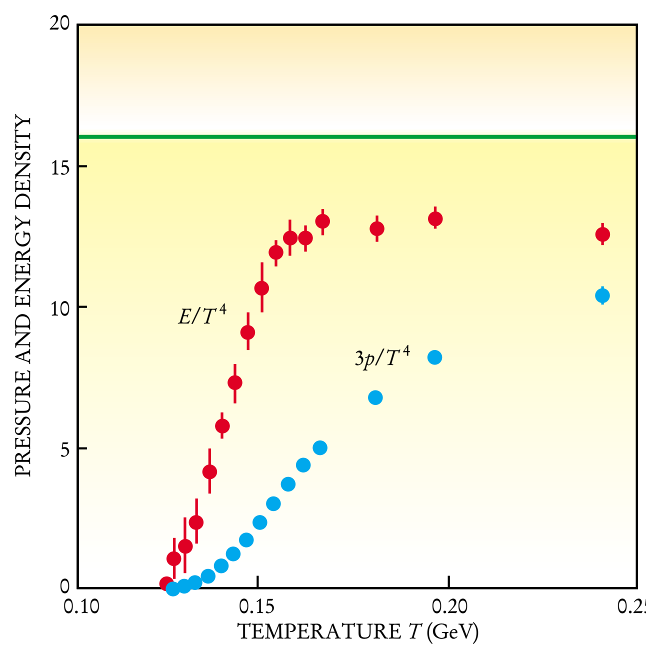

FIGURE 6. STEEP RISE of pressure p (blue points) and energy density E (red points) with increasing temperature T above 130 MeV indicates the opening of many quark–gluon degrees of freedom in this lattice-gauge QCD calculation of the thermodynamics of very hot nuclear matter.[15] For simplicity, the calculation assumes only two quark flavors. Normalized to T^4, both variables become dimensionless (in natural units) and asymptotically approach the green line if the quark masses are zero.

flavors. But in the imagined color–flavor locked state they become correlated. Both color symmetry and flavor symmetry, as separate entities, are spontaneously broken, and only a certain mixture of them survives unscathed.

Color–flavor locking in high-density QCD drastically affects the properties of quarks and gluons. As we have already seen, the gluons become massive. Due to the commingling of color and flavor, the electric charges of particles, which originally depended only on their flavor, are modified. Specifically, some of the gluons become electrically charged, and the quark charges are shifted. The electric charges of these particles all become integral multiples of the electron's charge!

Thus the most striking features of confinement—the absence of long-range color forces, and integer electric charge for all physical excitations—emerge as simple, rigorous consequences of color superconductivity. Also, because both left- and right-handed flavor symmetries are locked to color, they are also effectively locked to each other. Thus chiral symmetry, which required independent transformations among the left- and the right-handed components of the quarks, is spontaneously broken.

Altogether, there is a striking resemblance between the *calculated* properties of the low-energy excitations in the high-density limit of QCD and the *expected* properties—based on phenomenological experience and models—of hadronic matter at moderate density. This suggests the conjecture that there is no phase transition separating them.

Unfortunately both numerical and direct experimental tests of this conjecture seem out of reach at the moment. So it is not certain that the mechanisms of confinement and chiral-symmetry breaking we find in the calculable, high-density limit are the same as those that operate at moderate or low density. Still, I think it astonishing that these properties, which have long been regarded as mysterious and intractable, have been simply—yet rigorously—demonstrated to occur in a physically interesting limit of QCD.

I have tried to convince you of two things: first, that the fundamentals of QCD are simple and elegant, and second, that these fundamentals come into their own, and directly describe the physical behavior of matter, under various extreme conditions.

References

1. For a wide-ranging survey, see, H. Kastrup, P. Zerwas, eds., *QCD 20 Years Later*, World Scientific, Singapore (1993).
2. R. Feynman, *QED: The Strange Theory of Light and Matter* Princeton University Press, Princeton, N. J. (1985).
3. R. Feynman, *The Feynman Lectures on Physics*, Addison-Wesley, New York (1964) vol. II, p. 25-10.
4. F. Wilczek, Nature **397**, 303 (1999).
5. S. Aoki *et al*., http://xxx.lanl.gov/abs/hep-lat/9904012
6. G. Hanson *et al*., Phys. Rev. Lett. **35**, 1609 (1975).
7. D. Gross, F. Wilczek, Phys. Rev. Lett. **30**, 1343 (1973). H. Politzer, Phys. Rev. Lett. **30**, 1346 (1973).
8. For a recent review, see I. Hinchcliffe, A. Manohar, http://xxx.lanl.gov/abs/hep-ph/0004186.
9. D. Gross, F. Wilczek, Phys. Rev. D **9**, 980 (1974).
10. For a recent review, see F. Karsch, http://xxx.lanl.gov/abs/hep-lat/9909006.
11. U. Heinz, M. Jacob, CERN preprint nucl-th/0002042 (2000).
12. For a recent review, including extensive references, see T. Schäfer, http://xxx.lanl.gov/abs/hep-ph/9909574.
13. M. Alford, K. Rajagopal, F. Wilczek, Nucl. Phys. B **537**, 443 (1999).
14. M. Schmelling, http://xxx.lanl.gov/abs/hep-ex/9701002.
15. S. Gottlieb *et al*., Phys. Rev. D **47**, 3619 (1993). ■