have a dramatic impact on several subfields of nuclear science, but its effect on nuclear structure studies will be extraordinary. This bold new concept will define and map the limits of nuclear existence; make possible the exploration of the exotic quantal systems that inhabit these boundaries; and isolate, amplify, or reveal new phenomena, new types of nucleonic aggregations, and key interactions in ways that stable beams cannot. RIA will provide new foundations for the understanding of nuclei. It offers the promise to guide the development of a unified theory of the nucleus in which both the familiar properties and excitation modes of the nuclei at or near stability and the exotic structures far from stability may be encompassed in a single theoretical framework.

RIA is a project of such a scale that it is likely to operate toward the end or after the period covered by this Plan (fiscal years 2002–12). However, because of its importance to the future of the field, its influence will be felt immediately. For a timely start of construction, it is essential that the necessary technical developments be carried out expeditiously during the coming years. It is also necessary to develop the theoretical concepts and the experimental techniques required to fully realize the discovery potential of RIA. Most importantly, the scientific community that will use RIA must be nurtured at universities and national laboratories.

In the short term, the keys to progress in understanding nuclear structure are linked to the continued, vigorous exploitation of existing stable-beam facilities and of the first-generation exotic-beam facilities that are just now coming on-line. These facilities are critical to the pursuit of current exciting initiatives and promising physics themes. Stable and exotic beams are essential for the training of young scientists who will work in this field now and at RIA in the future, and who will meet national needs in high-technology areas such as medicine, stockpile stewardship, and energy production. Vigorous exploitation of these facilities requires funding at a level necessary to operate efficiently. While every effort must be made by the laboratories and universities to maximize productivity, it is clear that funding is currently inadequate to fully operate these very cost-effective facilities. The *Facilities Initiative,* which addresses this issue, is an essential component of this Plan.

Instrumentation is an area where new developments will enhance ongoing scientific productivity and pave the way for RIA. The success of Gammasphere has demonstrated that state-of-the-art instrumentation has a dramatic effect on the rate of scientific progress. New instrumentation initiatives should be encouraged, including but not limited to a *4π Gamma-Ray Tracking Array,* presented as part of this Plan.

Equally important to the vitality of nuclear structure research is increased support for theoretical investigations. The development of new concepts and methods is essential as an inspiration for the experimental program and for the interpretation of the fascinating and unexpected observations that will surely emerge. Exceptional progress has been made in recent years with very limited resources, both human and fiscal, thanks to radically new approaches and to the power of modern computing techniques. To take full advantage of the exciting science opportunities, especially in the context of RIA, a theory initiative is needed. A number of excellent suggestions to strengthen the nuclear structure theory program have been put forth. They are presented in the *Nuclear Theory Initiative* and the *Large-Scale Computing Initiative.*

This section has discussed the most basic question facing nuclear scientists: What is the structure of the nucleus? While many facets of the nucleus have been uncovered, much remains to be done, and the pace of discovery is rapid. For this reason, a coordinated framework for nuclear science research in the U.S. must maintain a vigorous program in this fundamental area. The 1996 long-range plan challenged the nuclear science community to develop a cost-effective proposal for an exotic-beam accelerator. The RIA concept not only meets this challenge, but also exceeds the performance expectations set at that time. RIA will provide the U.S. with the opportunity to assume a leadership role in nuclear structure physics for the coming decades. With RIA, the outlook for nuclear structure research is indeed a bright one.

QCD at High Energy Densities:

Exploring the Nature of Hot Nuclear Matter

Overview: In Pursuit of the Quark-Gluon Plasma

One of the fundamental tasks of modern nuclear physics is to understand the structure of the vacuum and the long-distance behavior of the strong interaction, one of the four basic interactions of nature. Quantum chromodynamics (QCD) is the current theory of the strong interaction in the context of the Standard Model of particle physics. However, many of the primary features of our universe are not easily understood from the form and symmetries of the Standard Model. In fact, a rather startling view of the vacuum arises when one examines the long-distance behavior of QCD. The vacuum—rather than being empty—is composed of a quark condensate that fills all of space, breaking the symmetries of the Standard Model and giving rise to the large masses of hadrons as compared with the light quarks (see "Chiral Symmetry, Mass, and the Vacuum," pages 46–47). It is a remarkable fact that the proton, made primarily of three light quarks, has a mass two orders of magnitude larger than its quark constituents; most of the proton's mass comes from its coupling to the quark condensate that fills the QCD vacuum.

Can we change the properties of the vacuum experimentally? Lattice QCD calculations argue that this is possible. When the vacuum "melts" at temperatures exceeding 170 MeV, the underlying symmetries of QCD will be restored. At these temperatures, matter should behave as a plasma of nearly massless quarks and gluons—the quark-gluon plasma, a state that existed in the first few microseconds after the Big Bang.

The vacuum, and in fact any system of quarks and gluons, is likely to have a complex phase structure, similar to that of many other bulk materials. The QCD phase diagram, as envisioned in current theory, is shown in Figure 2.14. The phase structure of the vacuum state is along the vertical axis where the baryon density is zero. There are actually two related phase transitions. The first is that of deconfinement, in which the quarks and gluons become free of their bondage in baryons and mesons. The second is that of chiral symmetry restoration, in which the masses of the quarks are reduced to their bare quark values. At temperatures above the chiral phase transition, the sum of the masses of the three quarks that make up the proton would be very small. It is still an open question as to whether these two phase transitions occur under exactly the same conditions of pressure and temperature.

The collisions of heavy ions produce matter with a range of energy densities and baryon densities corresponding to the two axes of the phase diagram in Figure 2.14. We expect

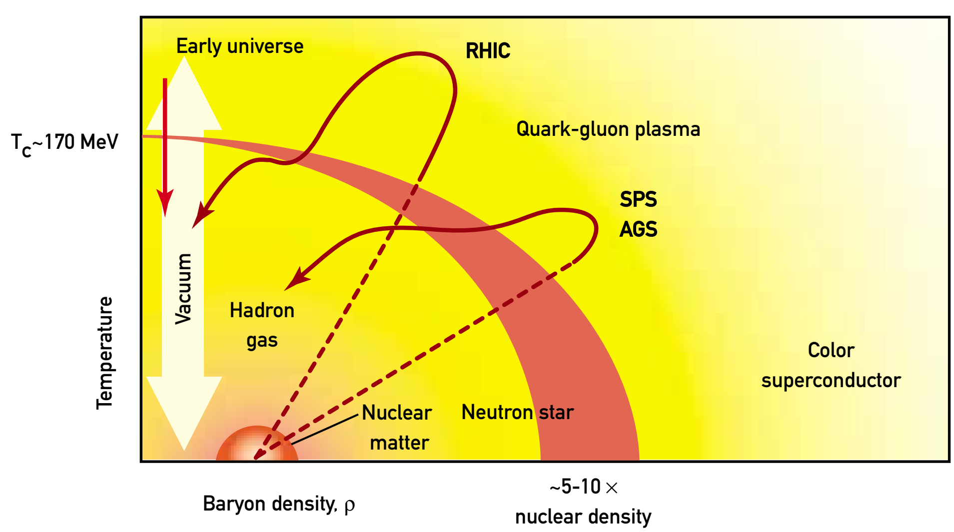

Figure 2.14. The phases of QCD. The QCD phase diagram, with baryon density as the abscissa and temperature as the ordinate. The red band indicates the range of temperatures and baryon densities at which phase transitions (both chiral and deconfinement) are thought to occur. The vacuum is the region near the *y*-axis, with a net baryon density near zero. Trajectories are also shown for systems created at the AGS and SPS, with center-of-mass energies of approximately 4–20 GeV per nucleon pair, and at RHIC, with a center-of-mass energy of 200 GeV per nucleon pair. The region of cold, high-density color-superconducting matter is shown to the far right.

both the deconfinement phase transition, and the chiral phase transition to occur in regions of either high energy density or high baryon density. Theoretical calculations on the lattice predict the location of the phase transition at low baryon densities to be at about 1 GeV fm^{-3}, or 170 MeV, as mentioned above. Theoretical calculations for nonzero baryon densities are much less certain, and only general estimates can be made. In experiments at lower energies, nuclei collide and essentially stop, producing systems of high baryon density. At the much higher energies provided by colliders such as RHIC, the nuclei pass through each other leaving behind high-energy-density debris with almost zero baryon density—essentially a highly heated vacuum.

In the context of this general picture, we have posed a number of fundamental questions:

- In relativistic heavy-ion collisions, how do the created systems evolve? Does the matter approach thermal equilibrium? What are the initial temperatures achieved?
- Can signatures of the deconfinement phase transition be located as the hot matter produced in relativistic heavy-ion collisions cools? What is the origin of confinement?
- What are the properties of the QCD vacuum and what are its connections to the masses of the hadrons? What is the origin of chiral symmetry breaking?
- What are the properties of matter at the highest energy densities? Is the basic idea that this is best described using fundamental quarks and gluons correct?

Achievements in relativistic heavy-ion physics. As implied in these questions, we are seeking to uncover many of the secrets of QCD with relativistic heavy-ion collisions (see “Anatomy of a Heavy-Ion Collision,” pages 48–49). The U.S. program in relativistic heavy-ion physics, until recently based on fixed-target machines, has a long history. The first such machine was Berkeley Lab’s Bevalac, which operated with a center-of-mass energy of about 2.5 GeV per nucleon pair. This was followed by the AGS at Brookhaven (~5 GeV per nucleon pair) and the SPS at CERN (~20 GeV per nucleon pair), where a substantial contingent of U.S. nuclear scientists collaborated. A new frontier has now been opened at RHIC, a colliding-beam machine with a center-of-mass energy up to 200 GeV, producing matter in a regime where lattice calculations are expected to be reliable. QCD has yielded its secrets grudgingly. Even in the perturbative regime where theoretical calculations are straightforward, the physics community did not immediately accept experimental evidence for gluons. A number of expected signals for the formation of the quark-gluon plasma have been observed in fixed-target experiments at CERN (and some at the AGS), but the evidence is not unambiguous. RHIC, with its broad range of accessible phenomena, can be expected to resolve many of these uncertainties—and indeed, early findings are tantalizing. But the flow of results has just begun, and conclusions are still tentative.

The most important recent achievements pertinent to the questions raised above include the following:

- Studies of particle abundances and spectra—as well as Bose-Einstein correlations, which give information about the space-time evolution of the collision—indicate that, in a nucleus-nucleus collision, the system undergoes rapid expansion and is close to both chemical and thermal equilibrium. Thermal equilibrium is thought to be reached very rapidly, but standard hadronic cross sections have difficulty accounting for the rapid rate at which this thermalization occurs. However, interaction cross sections arising from colored quarks and gluons are expected to be larger and could be driving this rapid thermalization.
- A state of matter in which quarks and gluons are mobile is expected to show a strong enhancement of strangeness production, particularly antistrange particles whose yield would ordinarily be suppressed by their relatively large masses. Experiments at CERN have seen enhanced strange antibaryon production, with increasing enhancement for each additional unit of strangeness. Recently, a similar strangeness enhancement has also been observed at RHIC. Experiments at the AGS, which have been able to detect only the $\bar{\Lambda}$, have seen a strong enhancement in the $\bar{\Lambda}$-to-$\bar{p}$ ratio.
- In 1985 charmonium ($c\bar{c}$) production was suggested as a probe of a deconfined medium created in relativistic heavy-ion collisions. Quarks and gluons in the medium would screen the strong interaction between charm and anticharm quarks and thus cause the ($c\bar{c}$) pair created by hard nucleon-nucleon scatterings to “melt.” This occurrence depends on the energy density of the medium and the species of charmonium being consid-

ered, with the less tightly bound χ and ψ' states breaking up at lower energy densities than the J/ψ. Just such a trend has been observed in experiments at CERN. However, uncertainties remain in the interpretation, because it is difficult to separate the contribution of charmonium suppression from other processes that produce similar effects.

- Signatures that may be interpreted as evidence of chiral symmetry restoration have also been seen. At CERN, excess electron-pair yields have been observed at invariant masses between 200 and 800 MeV, which can be explained as a mass shift of the ρ meson due to the onset of chiral symmetry restoration. Competing interpretations of the data as arising from collision-induced resonance broadening are also possible.

- Beginning in the summer of 2000, the first data were collected at RHIC in a run lasting three months, during which the machine reached 10% of its design luminosity, as planned. The center-of-mass energy during this run was 130 GeV per nucleon pair in gold-gold collisions. About 10^7 events were collected among the four RHIC detectors (see Figure 2.15 and "First Results at RHIC," pages 50–51). Much of this data has now been analyzed and published. During the fall of 2001, the full design luminosity was achieved, at the full energy of 200 GeV per nucleon pair, and 50–100 times as many events were collected.

- The multiplicity and transverse energy of particles produced in collisions at RHIC have been measured. These measurements provide an estimate of the energy density achieved, which was at least 20 times that of nuclear matter. The multiplicity was also measured as a function of the number of nucleons participating in the interactions. This provides a probe of the initial conditions for the collisions, which are believed to be a very high density of coherent gluon fields.

- One of the early, unexpected results at RHIC was the strong elliptic flow signal seen at relatively high momenta. Flow is a measure of the degree to which a group of particles moves collectively. Collective behavior can occur only if there is a strong degree of thermal equilibration. A strong elliptic flow indicates that this equilibration developed at very early times when the pressure was very large. Such an early thermalization, combined with the measured transverse energy, implies an energy density substantially higher than that required for the phase transition, as indicated by theoretical lattice calculations.

- The emission of hadrons with high transverse momenta is expected to be suppressed in a quark-gluon plasma, owing to the energy loss of partons. Such an effect was not observed in lower-energy collisions at CERN. Indeed, the opposite was found. However, at RHIC, the yield of high-transverse-momentum hadrons, measured in central nucleus-nucleus collisions, was

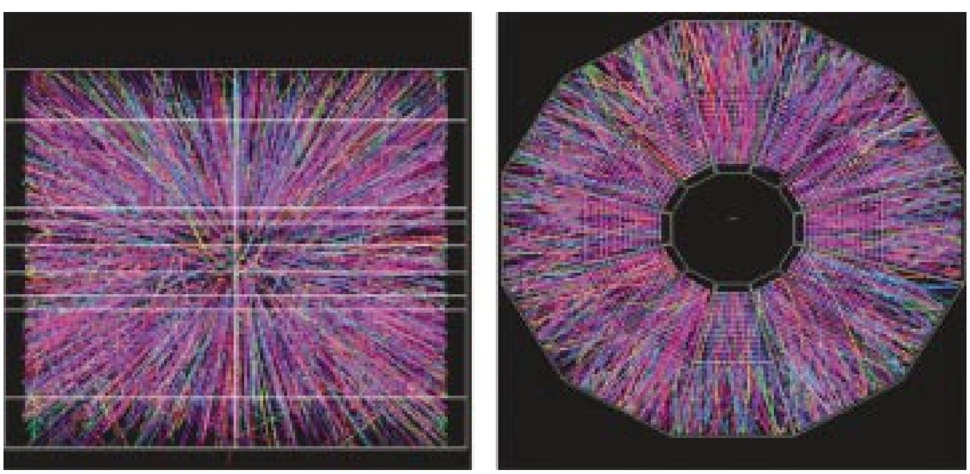

Figure 2.15. Golden glow. A central gold-gold event as seen in the STAR detector at RHIC. The side view and end view of the time-projection chamber are shown. Each colored radial line emanating from the center corresponds to the track of a particle produced in the collision. Such a central collision typically produces about 6000 particles.

substantially reduced as compared with that from ordinary proton-proton collisions.

- A major theoretical advance has also been made in a very different region of the phase diagram. In the very dense but very cold environment at the far right of Figure 2.14, quark matter is predicted to display many characteristics more familiar to a condensed-matter physicist than to a plasma physicist: Cooper pairs form, and the quark matter becomes a color superconductor, characterized by Meissner effects and gaps at the quark Fermi surfaces. Such cold quark matter may exist in the centers of neutron stars.

The years ahead: The RHIC era. Fixed-target experiments have clearly shown that heavy-ion collisions create high energy and high baryon densities. The density of hadrons is so large that there is simply not enough room for them to coexist as a superposition of ordinary hadrons. The observed signatures are not readily explainable by standard hadronic models. It is also clear that a great deal remains to be done in the energy regime accessible at the SPS. For example, a new experiment (NA60) is now under construction at CERN to measure the charm-production cross section, necessary to clarify the interpretation of the dilepton results.

In addition, we hope it will become possible to use astrophysical observations of neutron stars to learn more about the region of the QCD phase diagram occupied by dense, cold quark matter. Ultimately, information from astrophysical observations, data from the lower-energy experiments, and data on the hot quark-gluon plasma to be gained from experiments at RHIC must be pieced together into a coherent, unified phase diagram for QCD.

Notwithstanding the continuing promise of fixed-target experiments and astrophysical observations, the completion of RHIC at Brookhaven has ushered in a new era. Studies are now possible of the most basic interactions predicted by QCD in bulk nuclear matter at temperatures and densities great enough to excite the expected phase transition to a quark-gluon plasma. As the RHIC program matures, experiments will provide a unique window into the hot QCD vacuum, with opportunities for fundamental advances in the

Chiral Symmetry, Mass, and the Vacuum

Chiral symmetry is the symmetry between right-and left-handed objects, that is, between things that rotate clockwise and things that rotate counterclockwise. Physicists believe that the underlying rules governing the strong interaction is left-right—that is, chirally—symmetric. (The strong interaction is the force responsible for binding the atomic nucleus together.) The handedness of a particle is defined by the direction of the spin relative to the direction of motion. If one looks along a particle's direction of travel, a clockwise spin is defined as right-handed, a counterclockwise spin as left-handed. The flaw in this definition is that one can transform the coordinate system and change the definition of the spin, even while the intrinsic characteristics of the particle remain unaltered. Imagine an observer moving faster than the particle itself. He would see a "right-handed" particle moving in the *opposite* direction and would thus believe the particle to be left-handed.

In this case, it would be possible to change a right-left symmetric universe, one where half the particles are left-handed and half right-handed, into a universe in which all particles are right-handed—provided only that one had a very fast rocket ship. This would then spoil the chiral symmetry. How might this situation be avoided, so as to preserve chiral symmetry? One possibility is for all particles to be massless. It turns out that all massless particles move at the speed of light. Since nothing can move faster than the speed of light, no spin-redefining transformation is possible, and thus a universe of massless particles would be chirally symmetric. But of course, this doesn't match the universe that we see. Where then does mass come from?

Physicists believe that particles are, in their basic nature, massless and that they acquire mass through their interactions with the vacuum. This is the process of chiral symmetry breaking. QCD has the property that the lowest energy state is not empty space, but rather is a vacuum filled with a "condensate," which is itself composed of quarks. A computer simulation of this condition is shown at right. In turn, the interactions of quarks with this quark condensate conspire to make the quarks behave as if they have mass. This state of the universe depends on the temperature. If the understanding of quark confinement, chiral symmetry breaking, and, very possibly, new and unexpected phenomena in the realm of nuclear matter at the highest densities.

By colliding beams of ions from protons to gold, with center-of-mass energies from 20 to 200 GeV per nucleon pair, RHIC will create conditions favorable for melting the normal vacuum and creating states of matter unknown in the universe since the Big Bang. With these unique capabilities, RHIC addresses all of the fundamental questions posed on page 44. The U.S. thus now possesses the premier laboratory in which to study these questions. It is likely that an initial understanding of high-density QCD matter and its associated phase transitions will be achieved in the next several years. However, much will remain to be done after these initial discoveries.

By their very nature, phase transitions introduce a host of unusual phenomena—for example, critical phenomena leading to the formation of large-scale fluctuations. In addition, since there are actually two phase transitions, the chiral transition and the deconfinement transition, a great deal will remain to be understood about the relationship between the two. Furthermore, new puzzles will undoubtedly present themselves as we begin to understand more about the bulk behavior of QCD.

The following paragraphs discuss in more detail some of the opportunities that lie ahead.

Mimicking the Big Bang: Thermalization and Equilibration

The highly compressed, then rapidly expanding, nuclear matter created at RHIC, which has many of the characteristics of the early universe shortly after the Big Bang, is the system now available for answering the questions on page 44. Among them are questions of thermalization and equilibration: How do the systems created in these collisions evolve? Does the matter approach thermal equilibrium? What are the initial temperatures achieved? Experiments will probe these questions by focusing on hadrons—on single-particle distributions and on correlations among parti-

temperature is high, as in the early universe—or in a relativistic heavy-ion collision—the vacuum state is such that the quarks once again exhibit their massless nature.

Although this seems to contradict the natural belief that space is empty, this is, in fact, an integral part of the Standard Model of particle physics—a model that has proved remarkably successful. Relativistic heavy-ion collisions offer the possibility of observing the effects of the vacuum directly, by heating it up and changing its characteristics, that is, by "melting" the vacuum.

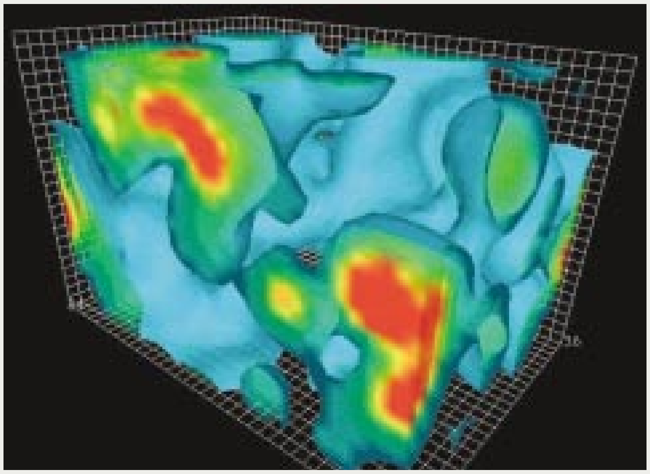

Simulating the vacuum. According to the Standard Model, all space is filled with the QCD condensate. The interaction of particles with this background condensate gives rise to most of the mass that makes up "ordinary" hadronic matter. The computer simulation shown here is a snapshot of the gluon field that binds quarks together to make up particles such as protons and neutrons. The red color indicates areas of intense "action" in the gluon field, associated with winding of the field lines. The green and blue colors correspond to weaker gluon field strengths. *Image courtesy of D. B. Leinweber, CSSM, University of Adelaide.*

Anatomy of a Heavy-Ion Collision

Relativistic heavy-ion collisions, such as those produced by Brookhaven's RHIC, provide physicists a chance to study very hot, dense matter similar to that which existed a few microseconds after the Big Bang. In a typical gold-gold collision, when viewed in the laboratory frame, the two nuclei initially appear as flat pancakes—a result of relativistic contraction (see the figure below). In the early preequilibrium phase, many "hard" collisions occur between the quarks and gluons of the nuclei, producing thousands of other quarks and gluons in an enormous cascade. The next stage of the collision is the one of primary interest. These secondary quarks and gluons equilibrate into a hot cauldron of matter, the quark-gluon plasma. Because of the low number of baryons (protons, neutrons, and their kin), this plasma is essentially a high-temperature vacuum. In the final stages, the plasma cools and condenses into ordinary particles, which are then seen by the detectors.

Head-on "central" collisions are the most violent sort, producing the largest number of "participants" and the largest volume of hot matter. More glancing "peripheral" collisions produce little, if any, hot matter, as suggested in the figure to the near right. However, these peripheral collisions are particularly important, since they can be used for comparison. The centrality of a collision can be monitored by detecting the cold "spectator" material.

In all these collisions, one of the most important questions is, How can one see if a plasma is made? One probe is provided by

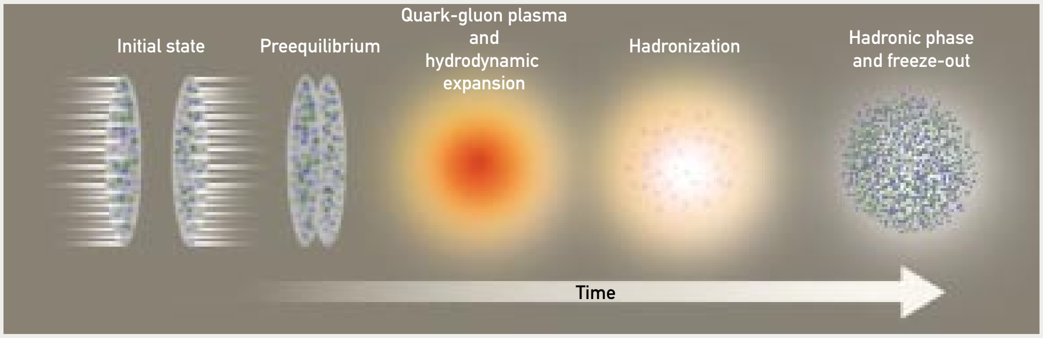

A Little Bang. Relativistic heavy-ion collisions replicate in the laboratory some of the conditions thought to exist a few microseconds after the Big Bang. In the schematic illustration here, two gold nuclei give rise to thousands of other quarks and gluons, which then equilibrate into a hot cauldron of matter, the quark-gluon plasma. As this plasma cools, it condenses into the ordinary particles seen by the detectors.

cles—and by detecting penetrating probes, which interact only electromagnetically and therefore escape the dense system relatively unperturbed. Theory, aimed at matching models and experimental data, will be crucial in understanding these measurements.

Extensive studies of lower-energy heavy-ion collisions, at the AGS and at CERN's SPS, have shown that the analysis of the distributions and correlations of soft hadrons yields the temperature and dynamics at the time the hadrons cease to interact, or "freeze out." The space-time evolution thus measured is crucial to understanding the collision dynamics, and it lends confidence to back-extrapolations to the early, hottest phase of the collision. Systematic study of the conditions under which the hadrons freeze out, as a function of initial temperature and collision volume, will help separate signatures of new physics from the underlying hadronic processes.

Momentum and flavor distributions of the hadrons provide information on the degree of thermal and chemical equilibration when the colliding system becomes dilute enough that hadronic strong interactions cease. When combined with other experimental information, such as thermal radiation, the space-time evolution of the entire collision can be inferred. An important goal at RHIC will be to determine whether equilibration occurs early in the collision, or only later, in the cooler hadronic phase. The complementary capabilities of the suite of RHIC detectors will be invaluable for this study, since they will allow us to combine measurements high-momentum particles. In the preequilibrium phase of the collision, some of the quarks acquire a very large momentum and thus appear as "jets" of particles. About half of the energy is carried by a single leading particle, which yields information about the momentum of the original quark as it left the collision region and before fragmenting into the particles that compose the jet. Fast quarks can traverse a region of ordinary hadronic matter with little hindrance (as illustrated schematically to the right); however, if the central hot region were a quark-gluon plasma, the fast quark would lose a great deal of energy, whereby the momentum of the leading particle would be greatly reduced. This leads to a "softening" of the momentum spectrum. Although it is too early to conclude that a quark-gluon plasma has been seen at RHIC, preliminary results of this sort suggest such a possibility and signify a spectacular beginning to the RHIC scientific program (see also "First Results at RHIC," pages 50–51).

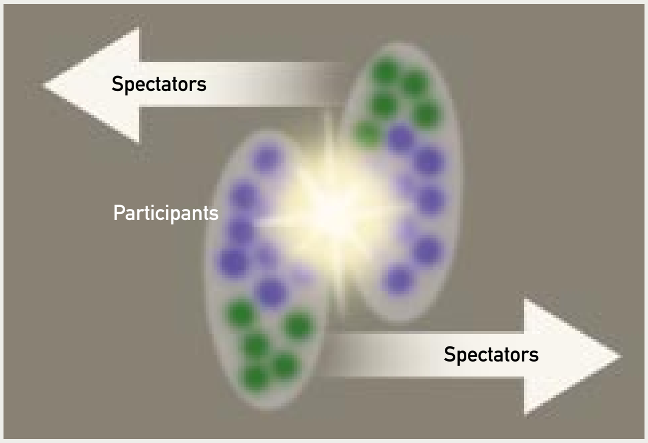

A glancing blow. Not all collisions are head-on, or "central," collisions; a "peripheral" collision is shown here. The centrality of a collision can be determined experimentally by measuring the number of "spectator" particles.

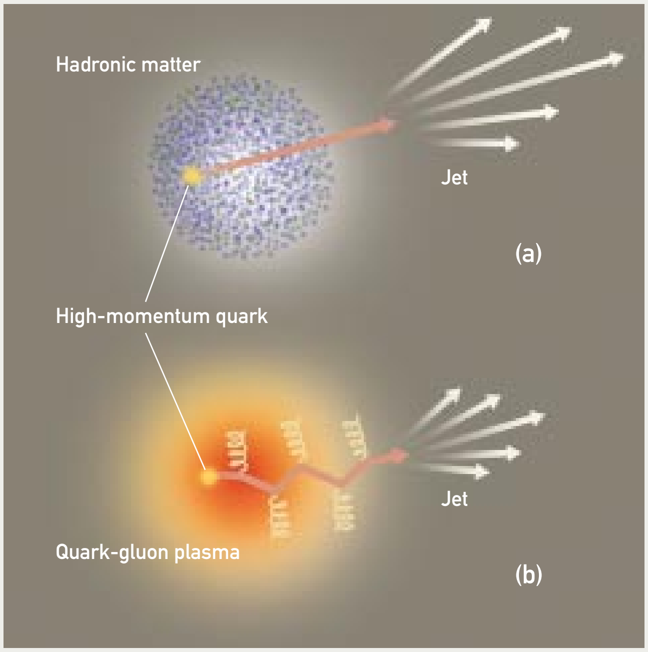

Quarks and jets. Jets arise from high-momentum quarks that fragment into particles after they leave the collision region. If no quark-gluon plasma has been formed, as in (a), the quark passes though nuclear matter with little resistance. However, in the presence of a quark-gluon plasma, as in (b), high-momentum quarks lose a great deal of their initial energy, and the detected particles in a jet have considerably less momentum.

of hadronic observables, collective behavior reflecting early conditions, and thermal emission of virtual and real photons.

Early results from RHIC on some of these topics have already indicated that the system freezes out at a lower baryon density and at a somewhat higher temperature than at the SPS or AGS. As already mentioned, flow measurements indicate that the degree of thermalization is high; hence, the concepts of temperature and pressure have meaning in the system under study. More information will be coming as physicists refine their measurements.

Real and virtual photons, materializing from quark-antiquark annihilations as electron or muon pairs, are radiated from the hot, dense QCD matter. Although such radiation is emitted at all times during the collision, for photons above about 100 MeV, the reaction dynamics significantly favors emission from the hottest part of the colliding system. (Detectors at RHIC are not currently sensitive to very low-energy photons.) Thus, measurements of the distribution of the black-body thermal radiation will yield the initial temperature. The background to such a signal is formidable, however, since photons and electrons are copiously produced from other sources as well, such as π^0 decay. Upgraded detectors designed to reject such backgrounds will be necessary in the future. Systematic analysis and variation of the initial conditions will also be required to solidify the interpretation.

First Results at RHIC

Construction began on RHIC in 1991 and was completed at the end of 1999. The first data-taking run commenced shortly thereafter. The second run, currently under way, is scheduled to conclude in early 2002. The four RHIC detectors were only partially instrumented for the first run but are now substantially complete. STAR is a large-acceptance detector built around a central time-projection chamber (TPC) in a solenoidal magnetic field. Inside the TPC is a silicon vertex tracker for detecting secondary vertices. The PHENIX detector is composed of four spectrometers optimized for detecting and identifying electrons, muons, photons, and hadrons. Multiple detector subsystems are used in the two central arms, yielding good momentum resolution and particle identification. Of particular note is redundancy in electron identification capabilities, giving a total e/π rejection of better than 10^{-4}. PHOBOS, one of two smaller detectors, is composed primarily of silicon and is optimized for large event rates. The second smaller detector, BRAHMS, specializes in measuring the fragmentation region of the collisions.

The early data have now been analyzed, and much has already been learned. Prior to the start of the RHIC experiments, very little was known about collisions of heavy ions at very high energies; for instance, predictions of the number of particles that would emerge from such a collision varied by a factor of four. Among the early findings are the following:

- The particle density in the hottest region in central gold-gold collisions was about a factor of two higher than previously achieved at CERN. A measurement of the transverse energy shows a similar result. Furthermore the yield per participant was found to increase with centrality, indicating the importance of hard processes and multiple collisions, which, assuming an early thermalization (as implied by the strong flow signal), would lead to energy densities greater than that required for the phase transition, as indicated in QCD lattice calculations.
- The azimuthal asymmetry of particle production in peripheral and semicentral collisions, known as elliptic flow, was found to be surprisingly large (as shown in the figure at the left)—evidence of a high degree of thermalization early in the collision, with a buildup of high pressure followed by a violent explosion. The magnitude of the flow agrees well with the predictions of hydrodynamic model predictions for a wide range of momenta and particle types. This important result supports the thermalization hypothesis.
- The antibaryon-to-baryon ratio was found to approach unity, in contrast to previous measurements at CERN, where the ratio was more than tenfold smaller. Fits to a thermal model of various particle yields lead to a very low baryo-chemical potential (consistent with a high antibaryon-to-baryon ratio), again substantially different from that observed at CERN. While not entirely baryon-free, these measurements indicate that at RHIC experiments are now in a region where lattice calculations are reliable and where the observed system should exhibit the characteristics of a vacuum at high temperature.
- One of the most intriguing results from the first run comes from a measurement of the transverse momentum spectrum for neutral pions, shown in the figure to the right. A fast-moving colored parton (quark or gluon) is a useful probe of hot nuclear matter. In normal nuclear matter, a quark would experience only a small energy loss; hence, the

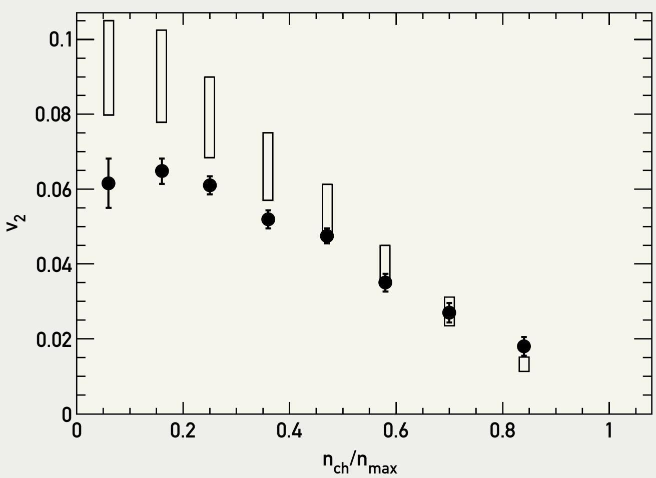

Nuclear flow. The figure plots elliptic flow v_2 (solid points) as a function of centrality, defined as n_{ch}/n_{max}, as measured by the STAR detector. The open rectangles show a range of values expected for v_2 in the hydrodynamic limit, scaled from ε, the initial space eccentricity of the overlap region. The lower edges correspond to 0.19ε and the upper edges to 0.25ε. The startling feature of these data is that, for central events ($n_{ch}/n_{max} > 0.5$), the flow signal appears to be as strong as allowed by hydrodynamics, implying an early equilibration of the system.

resulting jet would carry essentially all of the energy originally imparted to the parton. In a quark-gluon plasma, however, the deconfined color fields would slow the parton down considerably; energy losses can be as great as 10 GeV fm^{-1}. This would lead in turn to a reduced yield of high-momentum particles—exactly the result observed at RHIC. Whether this is a definitive signal of a quark-gluon plasma, however, has not yet been determined. Future data will provide more statistics and higher transverse momenta, as well as proton-nucleus data for comparison.

The interpretation of all these findings, in terms of temperature, entropy production, and ultimately, the existence of a phase transition, will take some time. Early results on charmonium suppression, dilepton spectra, and multistrange antibaryons will require data taken during the 2001–02 run. In any case, this has been a spectacular beginning for the RHIC experiments.

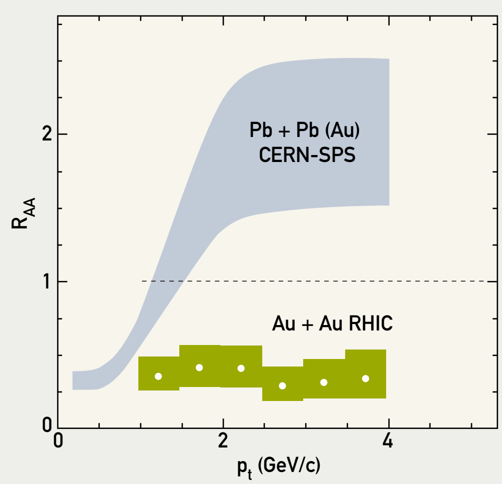

Jet suppression. Plotted here as a function of transverse momentum is the ratio of the measured yield of neutral pions in nuclear collisions to the yield that would be expected based on extrapolation from proton-proton collisions. Results are plotted for gold-gold collisions at RHIC (lower data points) and for lead-lead and lead-gold collisions at CERN at lower energies (broad upper band). The colored bars around the RHIC data points indicate the level of systematic error. The results are qualitatively different: At CERN energies, the yield at large transverse momentum is enhanced, whereas at RHIC energies, it is depleted. Such a depletion was predicted on the basis of the expected energy loss of partons in a quark-gluon plasma. The results, therefore, provide intriguing indications that this state may have been formed in collisions at RHIC. Further experiments will be needed to confirm this interpretation.

Quarks and Gluons Unbound: Deconfinement

A second suite of unanswered questions is tied to the phenomenon of deconfinement. For example, Can signatures of the deconfinement phase transition be located in the cooling debris of relativistic heavy-ion collisions? And what is the origin of confinement? The fundamental degrees of freedom in QCD are quarks and gluons. However, free quarks and gluons have never been observed, and the physical spectrum of particles contains only hadrons—"colorless" bound states of quarks, antiquarks, and gluons. The origin of this "confinement" is linked to the properties of the vacuum. Numerical calculations on the space-time lattice have shown that, at high temperatures and densities, the vacuum structure of QCD may "melt," leading to a novel form of QCD matter in which quarks and gluons move freely—the quark-gluon plasma. Further progress in the theory, including both new analytical methods and improved simulations on the lattice, is imperative to confirm these predictions. Heavy-ion collisions create a hot and dense environment in which this transition may be induced experimentally. Connecting observations from these experiments to "signatures" of deconfinement is now a prime goal of the field.

One such signature is the suppression of high-momentum particles. Measurement of the hard-scattering processes via high-transverse-momentum hadrons and heavy-flavor distributions will indicate to what extent the fast particles lose energy in the dense medium. This energy loss results in energy transfer from fast particles to the medium and thus drives thermalization. Furthermore, this energy transfer multiplies the number of gluons and drives particle production, increasing the density of the medium further. In fact, some theoretical models predict that matter may reach the stage of gluon saturation, in which case the physics is determined by interactions in a dense gluon gas, calculable using perturbative QCD, with subsequent hydrodynamic expansion. Measured particle yields, spectra, and correlations to transverse momenta of at least 10 GeV/*c* are needed to determine whether such predictions are correct. As noted earlier, RHIC may already have revealed hints of this thermalization phenomenon in the π^0 and charged-hadron transverse momentum spectra. Further measurements, in the second year of data-taking, have just begun and should extend the spectra to a transverse momentum of 10 GeV/*c*. In addition, important comparison data will be taken in the coming years for proton-proton and proton-nucleus collisions.

At higher luminosities, it will be possible to directly measure the photons produced opposite the high-transverse-momentum hadrons. Since these photons recoil against the quark jets, and since they do not suffer energy loss in the deconfined medium, they serve as indicators of the initial transverse momenta of the jets. Such observations will provide a means to make careful, quantitative measurements of the energy loss. One interesting possibility is to "flavor tag" the high-transverse-momentum hadrons. A leading K^{-} with no valence quarks is more likely to come from a gluon jet. This would allow us to measure the difference in the energy loss between gluon and quark jets. Gluon jets are expected to lose energy at twice the rate of quark jets in a deconfined medium. Later in the decade, the LHC will be able to make similar measurements at 30 times the center-of-mass energy available at RHIC, where the lifetime of the quark-gluon plasma is expected to be several times longer and jet cross sections at high transverse momentum are two to three orders of magnitude greater.

J/ψ suppression is another well-known proposed signature of deconfinement. RHIC will be able to measure J/ψ production in both the muon and electron channels. One of the crucial measurements that must accompany the measurement of the J/ψ is that of open charm production. To make this possible, specialized vertex detectors must be constructed, with the position resolution needed to discriminate between the charm vertex and the original event vertex. R&D programs focused on such an upgrade are under way.

The J/ψ is but one of the vector mesons in the charm family. The excited states of the J/ψ, as well as the Υ family (bound states of $b\bar{b}$), should all exhibit some degree of suppression. The suppression of the associated states, χ_b and χ_c, can also be observed, since they decay to the detectable vector mesons. Each of these states should melt at a different temperature. In fact, the Υ will be used as a control, since it should not be suppressed at all at RHIC energies. By varying the temperature and volume of the system by means of changes in beam energy and species, we can change the pattern of suppression of the various states. Not only would this be a convincing signature of a phase transition, but it would also give a good measure of the actual energy density. This will require a higher luminosity than currently available at RHIC, as in the proposal for RHIC II. In addition, when the LHC begins heavy-ion operation, the Υ family will be produced and detected at rates two to five time higher than at RHIC and will be easy to analyze.

Looking into the QCD Mirror: Exploring Symmetries

Investigations of chiral symmetry breaking respond directly to questions about the most fundamental properties of the QCD vacuum and its connections to the masses of the hadrons. The challenge for RHIC experiments is to search for evidence of in-medium mass changes among the low-mass vector mesons associated with the restoration of chiral symmetry. Direct, in-medium measurements of the masses of light vector mesons such as the ρ, ω, and ϕ are possible, since they decay rather rapidly within the fireball before hadronization. The decay to di-electrons is particularly interesting, since electrons should not be rescattered in the medium, and their invariant mass should reflect the mass of the vector meson in the altered vacuum state. Since some fraction of the vector mesons decay outside the medium (in the case of the ω, some 70–80% do so), these can be used as a calibration point for the measurement. The fraction exhibiting a shifted mass should change as a function of the transverse momentum and the size of the central fireball. This shift would be a particularly dramatic signature of the altered vacuum. Higher luminosities and improved detectors will be needed to reject background for detection of the ρ, the shortest lived, and hence the broadest, of the vector mesons. Observation of the ρ will be important, since it decays entirely within the fireball, and its spectrum may yield a history of the thermal evolution of the system.

The presence of a phase transition is also expected to cause inhomogeneities, which may survive the hadronic phase as fluctuations in particle number and type. Fluctuations and droplet formation are of particular interest, since similar processes may account for much of the large-scale structure of the universe and the inhomogeneities observed in the cosmic microwave background. Several fluctuations have been proposed as signatures of a phase transition. If the transition is first order, the growth of hadronic droplets and the shrinking of quark-gluon droplets may yield a lumpy final state and large fluctuations in particle number. Different scenarios may lead to other signatures, such as abnormal ratios of charged to neutral pions, or enhancements of pions at low momenta. Experiments will search for such phenomena and correlate their appearance with other signatures of the quark-gluon plasma.

The theory of chiral symmetry breaking and restoration is under active development. Progress requires the development of new analytical tools and further advances in lattice calculations. In order to investigate chiral symmetry on the lattice, we must be able to perform calculations with realistic quark masses. This places severe constraints on the size of the lattice and requires new methods (for example, "domain wall fermions") and new and more powerful computers.

Weak interactions violate both parity (P) and combinations of charge conjugation and parity (CP). By contrast, the strong interactions appear experimentally to preserve both of these symmetries under normal conditions. QCD as a theory does not require this. Therefore, it would be of great interest to learn whether CP-violating processes occur in the strong interactions under extreme conditions of high temperature and density. Theoretical progress in this area is linked to the understanding of topological effects in gauge theories at finite temperatures. This requires improvements in both analytical tools and lattice simulations. Clever experimental signatures have been devised for CP-violating bulk phenomena in heavy-ion collisions at RHIC. In theory, since CP is conserved in ordinary strong interactions, the signature of the altered CP state should be preserved during the evolution of the collision and may be quite distinct.

High-Density Matter

A final realm of investigation is the nature of matter at the highest energy densities. The behavior of QCD at the high-energy frontier is not yet understood theoretically. The simplest and most fundamental questions are still unanswered: Why do hadron cross sections rise at high energies? How are particles produced? What is the wave function of a high-energy hadron? RHIC will help find the answers to such questions by providing detailed data on particle production over a wide range of atomic numbers and energies. Progress in understanding high-energy behavior in QCD will, in turn, allow the reconstruction of the initial conditions in heavy-ion collisions, a crucial prerequisite to theoretical descriptions of the entire process. Such strides will require continuing development of theoretical tools, as well as large-scale, real-time Monte Carlo numerical simulations.

Gluonic interactions may be expected to dominate the first few fm/*c* of RHIC collisions, immediately following the initial nucleon-nucleon interactions as the nuclei penetrate one another. Gluon fusion processes dominate the production of charm and bottom quarks at energies attainable at RHIC. Consequently, measurements of open charm and bottom decays will likely be the most important ways to study the gluon fields inside heavy nuclei and their excitations in heavy-ion collisions. Of particular interest are distributions at low x, where x is a measure of the momentum fraction of a nucleon carried by an individual quark or gluon. These measurements require greater luminosity and detector efficiency than is currently available at RHIC. The RHIC II initiative addresses these two issues and will lead to a 40-fold increase in luminosity. RHIC II will be an invaluable tool to study the evolution of the quark structure functions to small x inside heavy nuclei (measurements of proton-nucleus collisions will yield this information), as the parton distributions evolve during a heavy-ion collision.

The dependence of the multiplicity upon the number of nucleons participating in the collisions was measured in the RHIC experiments. This dependence reflects the nature of the initial distribution of gluon fields inside the colliding nuclei. The early results provide support for a picture in which these fields are very dense and highly coherent, and in which the typical density scale of these fields inside nuclei is significantly greater than that inside a single nucleon.

Nuclear shadowing is another important process in understanding the initial stages of RHIC collisions. This can be measured directly via Drell-Yan and other hard processes in proton-nucleus collisions. Experiments must measure, with sufficient statistics, the dimuon distributions at high mass and the hadron spectra at high transverse momentum (at or above 10 GeV/*c*) to determine the extent of shadowing in kinematic regions accessible at RHIC. The results feed back, of course, into understanding the initial conditions in nucleus-nucleus collisions. However, they also probe the gluon field properties directly. If the gluon and quark densities can saturate, this will affect the gluon distribution deep inside a heavy nucleus, as well as the dynamics of the early stage of a heavy-ion collision. Measuring the intrinsic transverse momentum of the quarks within the nucleon via hard probes, and observing how this depends on x, as well as the volume of dense matter, can address the issue of saturation. All of these measurements require the capabilities of RHIC II. The LHC should have excellent capabilities to study this physics as well, since the apparent density of low-x virtual gluons will almost certainly be at saturation there.

A future electron-ion collider will make significant contributions to these measurements. Because of the Lorentz contraction, the nucleus will effectively amplify the parton densities seen by the incoming electron by a factor of the thickness of the ion, $\sim A^{1/3}$. Hence, lower center-of-mass energies are adequate for the observation of various phenomena such as gluon saturation. (In technical terms, one is able to reach thresholds for these interesting phenomena at higher values of x, the fraction of the nucleon momentum carried by the parton.)

Outlook

RHIC has just begun its task of uncovering the secrets of QCD. The next few years will yield a wealth of new information, and we have an outstanding opportunity to revolutionize our understanding of matter at the highest energy densities. Accordingly, the highest priority for the relativistic heavy-ion community is to utilize RHIC to its fullest potential. Sufficient running time is required to realize the physics promise of RHIC and to reap the rewards of our investment in RHIC's construction. This priority is also recognized in the first recommendation of this Plan and in the *Facilities Initiative* that supports it. Certain short-term upgrades are also essential, as well as R&D aimed at major upgrades to the machine luminosity and to the detectors. In the more distant future, significant upgrades of the collider and the experiments will be needed. An upgrade program such as the *RHIC II* initiative, which increases luminosity and adds new capabilities to the experiments, will allow in-depth pursuit of the most promising observables characterizing the deconfined state. Timely completion of the technical R&D is essential so that a detailed plan and schedule can be developed.

As discussed earlier, many open questions remain in the study of QCD at the high-energy frontier. Electron-nucleus collisions can provide complementary information to that obtained at RHIC. One of the technical options of the *Electron-Ion Collider* initiative would add this capability to the facility. This is an extremely exciting opportunity for the long term, since it allows access to a new regime within QCD and should shed light on the initial conditions for heavy-ion collisions at RHIC. As with RHIC II, R&D is essential in the near term so that a full scientific proposal can be developed.

Finally, the CERN heavy-ion program will be starting soon at the LHC. It would be wise to make a modest investment of manpower and money so that some U.S. participation is possible. This program should focus on those aspects of relativistic heavy-ion physics not easily addressed at RHIC. This includes jet and photon production at transverse momenta above 20 GeV/c, Υ family vector meson production, and W and Z production in heavy-ion collisions. The LHC offers jet and photon production rates one to three orders of magnitude larger than those at RHIC, for transverse momenta in the range 20–100 GeV/c, opening the door to detailed studies. The LHC also has a kinematical reach some 25 times better than RHIC's, extending into the realm of very soft gluons, where we may expect saturation effects to matter.

While the general structure of QCD is now firmly established, its properties are not yet fully understood. Many fundamental problems remain unsolved and are thus at the forefront of modern theoretical physics. One of the most important tools for making progress is lattice gauge theory, which allows us to solve complex nonlinear field theory problems using a computer. Such problems are among the most complex in computer science and require enormous computing power. New computational capabilities will be needed in order to make progress, both in interpreting experimental results and in furthering fundamental theoretical understanding. The necessary capabilities are contained within the *Large-Scale Computing Initiative*.