NUCLEAR MATTER IN EXTREMIS SOUGHT WITH ULTRARELATIVISTIC HEAVY-ION BEAMS

Last summer's Quark Matter '87 conference, at Schloss Nordkirchen near Dortmund, was the sixth in a series of quasiannual gatherings of physicists interested in extended nuclear matter under extreme conditions of density and temperature. But it was the first that had significant data to chew on. At sufficiently high energy densities, it is confidently believed, we will reach the Holy Grail—the phase transition from ordinary nuclear matter to the "quark–gluon plasma." (See PHYSICS TODAY, March 1985, page 40.) The experimental method of choice is to bang heavy nuclei together at the highest possible collision energies. Before September 1986 the best one could do (aside from examining a handful of spectacular cosmic-ray events) was to accelerate uranium ions to a modest 1 GeV per nucleon at the Berkeley Bevalac and fire them at a uranium target.

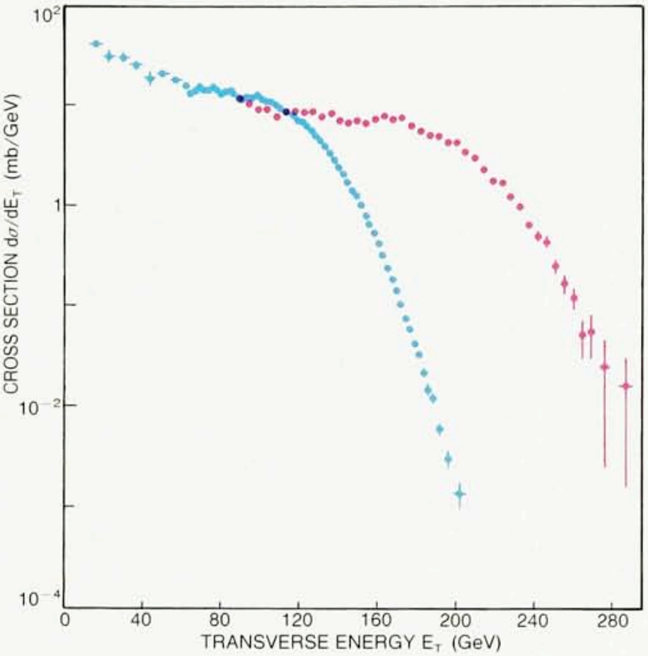

Sulfur nuclei, being twice as heavy as oxygen, produce a much broader distribution in transverse energy in 200-GeV-per-nucleon heavy-ion collisions at CERN. Cross section, as measured by the HELIOS collaboration, is plotted against transverse energy for 1986 oxygen-beam run (blue) and 1987 sulfur-beam run (red).

A new era of high-energy heavy-ion physics was inaugurated in the autumn of 1986. The Super Proton Synchrotron at CERN began providing experimenters with beams of oxygen nuclei at 3.2 TeV. That comes to 200 Gev *per nucleon*, roughly half the energy to which the SPS traditionally accelerates its proton beams. At about the same time, the venerable Alternating Gradient (proton) Synchrotron at Brookhaven offered its first oxygen beams at 14.5 GeV per nucleon. This latter may seem pale compared with the 200 GeV per nucleon provided by the much younger and larger CERN machine. But the Brookhaven energy is closer to what theorists expect to be the optimal stopping energy for colliding heavy nuclei. At much higher energies, nuclei are more transparent to one another. Somewhere between the AGS and SPS energies, however, one expects colliding nuclei to lose enough energy to fuse into a single "fireball," providing a glimpse of nuclear matter with perhaps the highest baryon densities one can get. At higher collision energies one is interested primarily in a "central region" of low baryon density but very high energy density.

The harvest of experimental results from these first high-energy heavy-ion runs at the SPS and AGS began in earnest at Quark Matter '87. The more expected results were accompanied by significant surprises. Physics results from the April 1987 AGS run with a silicon beam are just now coming in, and the analysis of the autumn 1987 SPS run with a 200-GeV-per-nucleon sulfur beam has just begun.

The heavier the projectile nucleus, the greater is the energy density one can expect in a collision. With their present injector systems neither the SPS nor the AGS can go beyond sulfur beams. But with the new AGS booster, now under construction, one will be able to accelerate gold nuclei to 12 GeV per nucleon; and a new injector system awaiting external funding at CERN would provide lead nuclei at 180 GeV per nucleon in the SPS. Both these accelerators, however, are restricted to *fixed-target* operation. Although the SPS ring serves as a collider for countercirculating proton and antiproton beams, there are as yet no heavy-ion colliders. (Proton–antiproton colliders can make do with a single ring of magnets because the beams are of *opposite* charge. An ion collider would require a double magnet ring.) In their center of mass, therefore, a pair of nuclear protons collide with a total energy of only 20 GeV at the highest SPS beam energy for heavy-ion running, significantly less than what is needed, the theorists tell us, to give assurance of reaching the quark–gluon plasma.

Therefore, in spite of the experimental *tours de force* at the SPS and AGS, all eyes are still on the vacant 4-km-circumference tunnel at Brookhaven originally constructed to house a proton–proton collider (Isabelle) that was never built. In 1986 DOE approved the construction of RHIC, the Relativistic Heavy Ion Collider, in the Isabelle tunnel. With its countercirculating 100-GeV-per-nucleon heavy-ion beams (with masses up to gold), RHIC would provide 10 times the center-of-mass energy per nucleon one now has at the SPS. Quantum-chromodynamics calculations and the data already in hand indicate that RHIC energies should easily suffice for making the quark–gluon plasma. But construction has not yet begun. The fiscal-1989 DOE budget provides only research and development money for RHIC, but no construction funds as yet.

The push for the new $20 million lead injector at CERN reflects, to some extent, the international community's disappointment at this postponement of RHIC. Carlo Rubbia, who becomes director general of CERN next January, foresees a use for this new injector beyond the confines of the SPS. If the Large Hadron Collider is eventually built in the 27-km LEP tunnel at CERN, he points out, this injector would make it possible to produce 1600-TeV lead–lead collisions from the very start of LHC operation.

The experiments

The 7-km-circumference SPS ring at CERN remains principally a $p\bar{p}$ collider. Heavy-ion beams are invited in for occasional short runs. The heavy-ion program was initiated in 1983 with an agreement among CERN, the Lawrence Berkeley Laboratory and the Gesellschaft für Schwerionenforschung in Darmstadt. GSI was to provide the electron-cyclotron-resonance ion source, and LBL undertook to build the rf quadrupole that would focus and bunch the ions for injection into the linac stage that precedes the PS and SPS synchrotrons. After the linac stage the ions are stripped of their last remaining electrons. The maximum ion energy per nucleon reached in the SPS is only half what one gets for protons because the charge/mass ratio of middleweight nuclei is about half that of the proton.

Worldwide participation in the CERN heavy-ion program has grown rapidly from its inception. The 1986 oxygen-beam run involved 62 institutions from 18 countries. Four large detector systems collected the bulk of the data in this run: WA80 is a collaboration of LBL, GSI, Oak Ridge and the universities of Lund, Münster and Tennessee.[1] Its detector configuration begins with the Plastic Ball, which was lovingly crated and flown to CERN (see cover of this issue) after long service at the Berkeley Bevalac. Its 650 plastic scintillator modules surround the target, except for a hole in the forward direction that lets small-angle collision products proceed to the additional counter and calorimeter arrays and spectrometer units further downstream. The NA35 experiment[2] (the letter designations refer to north area, west area and so forth), a collaboration of 14 European institutions and LBL, features a new streamer chamber built by a team from the Max Planck Institute in Munich. The streamer chamber, with the target just inside its upstream end, provides the most visual display of these ultrarelativistic nuclear collisions. The streamer chamber is operated at lower-than-usual voltage to resolve the many hundreds of charged-particle tracks emerging from a single collision.

NA34, next door to the streamer-chamber experiment, is an experiment of the HELIOS collaboration, 16 institutions from western Europe, North America, the Soviet Union and Israel. This survey experiment[3] surrounds the target almost completely with hadronic and electromagnetic calorimeters that tell us where the collision energy is going. The HELIOS detector also incorporates magnetic spectrometers that allows one to sample momentum spectra and measure the invariant-mass spectrum of muon pairs emerging from the collisions. NA38, a French, Portugese, Spanish, CERN collaboration, is a "beam-dump" experiment that looks only at the penetrating muons produced in a heavy-ion collision, measuring the dimuon spectra by means of a magnetic spectrometer.

These large experiments at the SPS are augmented by several small-scale undertakings involving visual, non-electronic detectors—photographic emulsion stacks and plastic sheets in which traversing charged particles leave permanent tracks.[4] Two more large detector systems joined the others in time for last fall's sulfur run: NA36, featuring a time-projection chamber built under the direction of LBL; and WA85, which uses the huge Omega Spectrometer, a veteran of the old Proton Synchrotron that preceded the SPS.

The relativistic heavy-ion program at Brookhaven was made possible by the upgrading of the tandem Van de Graaff accelerator and the construction of a 2000-foot transfer line from the Van de Graaff to the AGS. The data from the inaugural 1986 oxygen-beam run at 14.5 GeV per nucleon were taken by the E802 group,[5] an Argonne, Brookhaven, Columbia, MIT, Tokyo, University of California collaboration. With its venerable "Henry Higgins" dipole magnet, which dates back to the old Cambridge Electron Accelerator, the E802 multiparticle spectrometer is particularly good at particle identification. For the silicon-beam run last spring, E802 was joined by the E814 detector, operated by a collaboration of Brookhaven, Stony Brook and eight other institutions. The first results[6] gathered with the E814 detector's hadron calorimeter and wall of sodium iodide modules were presented at Quark Matter '87.

Bread and butter

Summarizing the early experimental results at Quark Matter '87, Berkeley theorist Miklos Gyulassy divided the more important findings into "bread and butter" results and exciting new surprises. "The bread and butter is the basic physics we need to know so we can make estimates for RHIC energies—to tell us we're on the right track," Gyulassy explains. The first of these issues was to estimate the energy densities the collisions produce. The estimate begins with dE_T/dy, the observed distribution of "transverse energy" E_T per unit "rapidity" y. One measures the energies of collision products as they are absorbed in the various calorimeters arrayed around the target. Before the collision, of course, all the energy is longitudinal—along the beam axis. The interesting energy is that which the collision diverts from the longitudinal to the transverse direction—E_T, given by a sum of contributions $E\sin\theta$, where θ is the scattering angle from the beam axis to the particular calorimeter that absorbs the energy E.

The rapidity variable y is a relativistically convenient measure of longitudinal velocity, given by $\tanh^{-1}(P_L/E)$, where P_L is the longitudinal momentum component of a particle of energy E. This definition is chosen so that a distribution of rapidities is simply shifted to the left or right by a Lorentz boost along the beam axis; its shape is invariant. The rapidity variable naturally divides the aftermath of a high-energy nuclear collision into three regions: Particles departing with rapidities near that of the projectile (target) nucleus are said to come from the "projectile (target) fragmentation region." If the collision energy is high enough, so

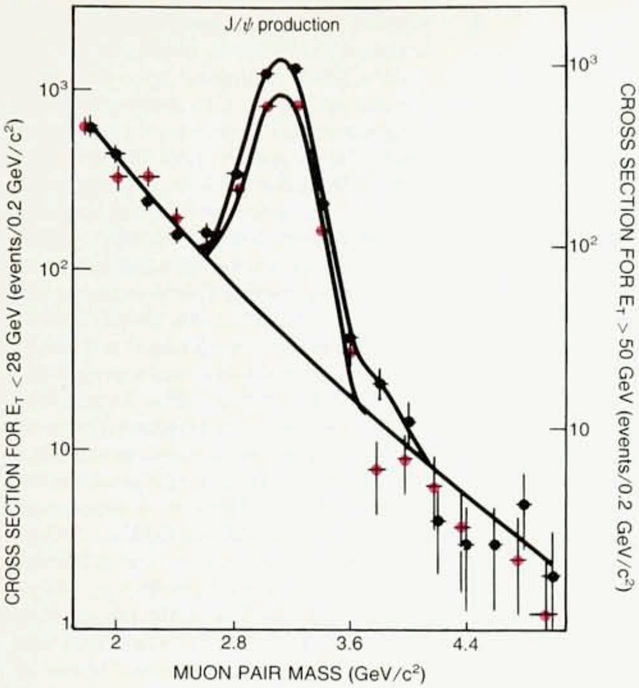

Suppression of J/ψ production in head-on, ultrarelativistic heavy-ion collisions is thought to indicate approach to quark–gluon plasma or, at least, very high energy density. Cross section for producing the charm–anticharm J/ψ meson in collisions of 200 GeV-per-nucleon oxygen nuclei on a uranium target at CERN[8] is plotted against the invariant mass of the $\mu^+\mu^-$ lepton pair into which the J/ψ decays. Head-on collisions ($E_T > 50$ GeV) are plotted in red; peripheral collisions ($E_T < 28$ GeV) are plotted in black. J/ψ production is indicated by resonant peak at 3.1 GeV/c². The two curves, with separate cross-section axes at right and left respectively, are normalized so that their backgrounds overlap in the figure. Relative to background, production of J/ψ mesons in head-on collisions is only 64% of its frequency in peripheral collisions.

that the two fragmentation regions are well separated in rapidity, a "central region" emerges between them. Particles with such intermediate rapidities are thought to come from a region of high energy density and low (net) baryon density, left behind near the center of mass by the departing nuclear fragments.

In the CERN 200-GeV-per-nucleon runs, the rapidity interval between incident nuclei is about 6. As each of the colliding nuclei slows and fragments, it spreads out over at least 2 units of rapidity, leaving perhaps 2 units of rapidity for the central region. Detector elements in these experiments are arrayed to cover all three rapidity regions with varying completeness. At AGS energies, where the initial beam–target rapidity difference is only 3, there is really no central region. The overlap of the two fragmentation regions reflects the fact that the colliding nuclei are just about stopping one another. At the highest RHIC beam energies the rapidity interval between the incident nuclei would exceed 10.

First results

Measuring the transverse energy deposited in their central-region calorimeters, the three large-detector groups taking data in the 200-GeV-per-nucleon oxygen-beam run found that dE_T/dy peaks with a value of about 60 GeV at central rapidities in head-on collisions. To convert this measurement into a determination of the maximum central energy density attained in these collisions, they invoke an argument originally put forward by James Bjorken (Fermilab), which attempts to relate the observed rapidity distributions to the space–time structure of the collisions. The oxygen beam is directed at a variety of target nuclei, all of them much larger than oxygen. One therefore imagines that the small projectile nucleus bores a cylindrical hole through the larger target nucleus. Target nucleons outside this impact cylinder are regarded as mere "spectators." The collision energy density in a particular rapidity region is then given by dE_T/dy divided $\pi R_B{}^2\tau$, where R_B is the radius of the oxygen nucleus and τ is an estimate of the proper time interval during which the collision energy is deposited. In essence, $\pi R_B{}^2\tau dy$ is the presumed volume in which an energy E_T has been deposited in the region of rapidity dy.

From the Bjorken formula the CERN experiments arrive at a conservative estimate of about 1 GeV per cubic fermi for the maximum central-region energy density attained in the 200-GeV-per-nucleon oxygen run. That doesn't sound like much if one recalls that the rest mass of a nucleon is about 1 GeV and its diameter is about 1 fermi (10^{-13}cm). But it should be noted that undisturbed nuclei are quite loosely packed; the normal nuclear density is only about 0.14 GeV/fm³. So what's already been seen in the CERN heavy-ion collisions is perhaps seven times normal energy density in a region from which the incident nucleons are assumed to have fled. That's still presumably not enough to reach the quark–gluon plasma, "but we think we may be halfway there," Gyulassy told us. "We've probably reached a mixed phase, in which the overlap of the hadrons (nucleons and mesons) is considerable—perhaps a partial meltdown of the hadronic matter into little blobs of quark matter." With beam nuclei heavier than oxygen one expects to do considerably better, and 3 GeV/fm³ should be well within the reach of the lower end of the RHIC energy scale. At the highest RHIC energies and nuclear masses, central energy densities on the order of 10 GeV/fm³ are foreseen.

Would one have predicted a central energy density as high as 1 GeV/fm³ in these initial experiments? Suppose one simply extrapolates from proton–proton data, after the manner of Roy Glauber (Harvard)—superposing the interactions of the 16 oxygen nucleons with the roughly 50 "participating" target nucleons in the impact cylinder and ignoring the possibility of nonlinear collective effects. This exercise does indeed yield roughly 1 GeV/fm³. But recent detailed calculations of dE_T/dy in this spirit, based on earlier theoretical work by Larry McLerran (then at Stanford) and coworkers, consistently yield predictions roughly 20 percent below the heavy-ion data. "The order of magnitude is under control," says Gyulassy, "but there seems to be a systematic tendency for all the oxygen data to be a bit higher than linear extrapolations from proton–proton data." He warns, however, that the magnitude of this apparent tendency is small compared with the uncertainties inherent in the application of the

Glauber–McLerran phenomenology. One ignores, for example, the spectator nucleons outside the cylinder supposedly bored by the projectile through the much larger target nucleus. But, in fact, we see at these energies that the spectators behave with surprising violence. "At CERN energies the spectators behave like the *participants* at Bevalac energies," say Arthur Poskanzer, a Berkeley member of the WA80 collaboration.

Pion interferometry

After the departing fragmented nuclei leave behind a hot central region of seven times normal energy nuclear density (and net baryon number close to zero), what happens next? The hot region expands and cools before its final products (mostly several hundred pions) "decouple" from the hadronic cauldron and go their separate ways. In particular one wants to know how much the system expanded before decoupling. Was there really any thermalization, and is some of this heat converted into collective hydrodynamic motion?

The NA35 streamer-chamber group undertook to measure the actual size of the central region at the decoupling time by means of pion interferometry. This interferometric effect of source size on pion correlations was first observed by Gerson and Shulamith Goldhaber, Wonyong Lee and Abraham Pais at Berkeley in 1962. The idea goes back to the method of intensity (as distinguished from amplitude) interferometry developed in the 1950s by Robert Hanbury Brown (Manchester) and Richard Twiss (Sydney) for measuring the angular diameters of stars.

Because pions, like photons, are bosons, they tend to be positively correlated in momentum space. The two-particle correlation function in momentum space peaks when the momentum difference $\Delta\mathbf{P}$ between two identical pions goes to zero. The rate at which the correlation function falls as $\Delta\mathbf{P}$ increases is a measure of the source size: the smaller the source, the larger the $\Delta\mathbf{P}$ at which one still sees correlation.

The streamer chamber is particularly well suited for this sort of measurement because one can pick out negatively charged tracks rather easily, and one makes a mistake of only about 10 percent by assuming these tracks are all pions. At Quark Matter '87 Thomas Humanic (GSI Darmstadt) reported the results of the NA35 group's "probing of the space-time geometry of ultrarelativistic heavy-ion collisions" by pion interferometry. The first surprise was that the transverse radius of the central pion source region appears to grow from an initial 3 fm (the radius of the oxygen nucleus) to more than 7 fm by the time the pions decouple. This implies a 20-fold decrease in central energy density from the initial post-collision maximum to the final "freeze-out" decoupling.

Under these circumstances one would surely expect to have seen some evidence for nuclear hydrodynamic flow. Such flow would manifest itself in the anomalous departure of transverse-momentum distributions from the simple exponentials one sees in proton–proton collisions at these energies. But no clear evidence of this kind has been seen in these heavy-ion experiments. "How can it be that we have a 20-fold expansion without seeing evidence of collective flow in the transverse-momentum distributions," Gyulassy asks. "If both these results are confirmed, we'll have the first real puzzle posed by the bread-and-butter data." The experimental situation, however, is unclear. Protons, being generally slower than pions, should manifest collective flow more clearly. But it is difficult to identify protons among the far more numerous positive pions. Therefore the transverse-momentum spectra thus far available are less than ideal for detecting collective flow.

The bread-and-butter studies have uncovered yet another surprise, which has turned up in investigations of nuclear stopping power. How much is an incident oxygen nucleus slowed down by passing through the larger target nucleus? From high-energy proton–nucleus scattering experiments we know that a single energetic proton traversing a nucleus in a head-on collision slows down by about two units of rapidity. This maximum Δy is more or less independent of the incident proton energy. One wanted to know whether things would be very different when the incident projectile is a nucleus.

A Δy of 2 implies that an oxygen nucleus at AGS energies (10 or 15 GeV per nucleon) can be completely stopped in a target nucleus, and indeed stopping seems to be confirmed by the first results of the E802 and E814 groups at Brookhaven. The E802 group deduces the stopping of incident oxygen nuclei from the observation that the maximum transverse energy deposited in their central-region calorimeters is independent of the size of the target nucleus. If passage through a gold nucleus produces no more E_T than passage through a copper nucleus only half as wide, it would seem that the oxygen projectile doesn't get all the way through.

At CERN energies $\Delta y = 2$ is not enough to stop the projectile; the target nucleus is effectively transparent. The surprise here comes from the WA80 and NA35 experiments. Looking at protons emerging from 200-GeV-per-nucleon oxygen collisions at rapidities appropriate to target fragmentation, the groups see far more of them than one would expect from the naive geometrical picture of a collision cylinder surrounded by passive spectators. "The target nucleus explodes," as Gyulassy puts it. Though this came as a surprise, Gyulassy told us, "it's easy to understand after the fact." If each incident oxygen nucleon deposits a GeV or two in the target, that comes to more than 100 MeV per target nucleon, "which is gigantic on the scale of nuclear binding energies." Understandable though it may be, the effect tends to obscure the question of stopping power at CERN energies. "You've blown the hell out of the target, scattering the spectators so widely in rapidity that the distribution of 'real' participants is buried underneath."

Less charm, more strangeness

Departing from the bread-and-butter investigations, we come to two quite new observations that have excited considerable attention. In 1986 Tetsuo Matsui (MIT) and Helmut Satz[7] (Bielefeld and Brookhaven) suggested that the onset of a quark–gluon plasma should be signalled by a significant suppression of J/ψ meson production as a result of "color screening." The famous J/ψ meson, discovered in 1974, is the lightest (3.1 GeV) and most spectacular of the $c\bar{c}$ bound states of the charmed quark and its antiquark. Just as a sufficiently dense ordinary plasma can prevent the formation of atoms by electromagnetic screening of nuclei from electrons (the "Mott transition"), so, they argue, would a quark–gluon plasma screen the color (gluon-exchange) force between quarks. Ordinary mesons, being bound states of light, abundantly produced quarks, would not suffer such suppression. But the heavy charm–anticharm pairs necessary to make J/ψ mesons are produced only rarely, in very hard quark collisions that can happen only in the initial instant of nuclear collision. If such a pair is eventually to emerge as an intact J/ψ, the two charmed quarks must stick together through an extended interval of space and time. If this involves passage through a quark–gluon plasma, relatively few $c\bar{c}$ pairs will survive.

Whether lesser excitations, short of a full-blown plasma, would cause significant J/ψ suppression, is a theoretical question much discussed in recent months.

The NA38 beam-dump experiment is designed specifically to look for J/ψ mesons manifested as a peak at 3.1 GeV in the invariant-mass spectrum of emerging muon pairs. Initial results from the NA38 oxygen-beam run at 200 GeV per nucleon indicate that J/ψ production is indeed suppressed by about one-third in head-on collisions.[8] (See figure on page 19.) One estimates this suppression by comparing the observed J/ψ peak with its underlying muon-pair background in head-on and glancing collisions. Although uncertainties remain in its interpretation, the one-third suppression is thought to be quite significant, both experimentally and theoretically. "This result is perhaps the most intriguing so far," Gyulassy told us. "It's much more than we expected at this stage, when the energy densities are still quite modest, and we don't think we have a quark–gluon plasma yet."

A similar surprise comes to us from the E802 detector at Brookhaven, running with a silicon beam at 15 GeV per nucleon. In proton–proton collisions at this energy one sees about one strange K^+ meson produced for every ten or twenty π^+. But when they collide silicon nuclei with gold, the group finds a K^+/π^+ of $(20 \pm 5)\%$—*twice* as high as the average ratio in pp collisions. The E802 detector is particularly good at distinguishing kaons from pions. If one looks only at head-on collisions, signalled by unusually large numbers of collision products, the K^+/π^+ ratio is higher still. (The K^-/π^- ratio, which is always lower, has not been observed to differ significantly between pp and heavy-ion experiments.)

This result is clearly interesting, but it is not well understood. One popular explanation goes as follows: At AGS energies, with beam and target stopping together, one is presumably seeing extraordinarily high baryon densities. Under these circumstances, up quarks and down quarks, the ordinary constituents of nucleons, are much more abundant than their antiparticles. (In a baryon-free central region at high energies, by contrast, one expects them all in roughly equal numbers.) Therefore when an $s\bar{s}$ pair of strange quarks is produced in the collision, the $\bar{s}$ strange antiquark should have a particularly easy task of finding an up quark to form a K^+. This effect is referred to as "K^+ distillation." The observation is greeted with enthusiasm as evidence that one can reach very high baryon densities at these relatively low energies.

"These experimental runs have been a great success," comments theorist Gordon Baym (University of Illinois). "These relatively light ions have achieved unprecedented energy densities over extended volumes. Our predictions and hopes for RHIC appear to be on track. Though much of the data remains to be analyzed, the experiments show that we understand heavy-ion collisions rather well."

—BERTRAM SCHWARZSCHILD

References

1. R. Albrecht *et al.* (WA80 collaboration), Phys. Lett. B **199**, 297 (1987).
2. A. Bamberger *et* al. Phys. Lett. B **184**, 271 (1987).
3. T. Akesson *et al.* (NA34 HELIOS collaboration), Z. Phys., to be published (1988).
4. P. L. Jain, K. Sengupta, G. Singh, Phys. Rev. Lett. **59**, 2531 (1987). G. Gerbier, W. Williams, P. B. Price, R. Guoziao, Phys. Rev. Lett. **59**, 2535 (1987).
5. T. Abbott *et al.* (E802 collaboration), Phys. Lett. B **197**, 285 (1987).
6. B. Bassalleck *et al.* (E814 collaboration), in Proc. Quark Matter '87, to be published in a special issue of Z. Phys. (1988).
7. T. Matsui, H. Satz, Phys. Lett. B **178**, 416 (1986).
8. M. C. Abreu *et al.* (NA38 collaboration) in Proc. Quark Matter '87, to be published in a special issue of Z. Phys. (1988).
9. T. Abbott *et al.* (E802 collaboration), in Proc. Quark Matter '87, to be published in a special issue of Z. Phys. (1988).

NUCLEON CORRELATIONS SEEN IN PION DOUBLE CHARGE-EXCHANGE REACTIONS

The position of one nucleon in a nucleus is assumed to affect the location of another, but such spatial correlations are hard to demonstrate. Many hoped they would appear in pion double charge-exchange reactions once the meson factories were built in the 1970s. Unfortunately, any correlations were masked by other complex interactions at the normal operating energies of the meson factories. Only recently—in new data taken at energies as low as 35 MeV—have nucleon correlations begun to surface. They show up in pion double charge exchange because these reactions necessarily involve two nucleons: The ingoing positive pion has successive charge-exchange scatterings with two neutrons in a nucleus, changing them into protons in the process. At low energies the cross sections for pion double charge-exchange transitions to certain states are larger than expected, and they are peaked in the forward direction. Theorists have explained these features and others in terms of nucleon correlations. Their work indicates that the pion interacts with two valence nucleons outside a closed shell that are within one femtometer (about a proton radius) of each other.

The simplest case of a pion double charge-exchange reaction is one in which the nucleus consists of an even number of valence neutrons outside a closed shell. The pion interaction just changes two valence neutrons to protons and leaves the nucleus otherwise unchanged. The initial and final states are members of the same isospin multiplet; the final state is the double-isobaric-analog state of the ground state. In these transitions the reaction is essentially elastic.

Experiments on carbon-14

The first pion double charge-exchange reactions to show convincingly the importance of nucleon correlations were performed below 80 MeV—an energy region well away from the possible interfering effects of the strong pion–nucleon interaction known as the delta resonance. William Gibbs (Los Alamos National Laboratory) told us that single charge-exchange reactions indicate that the nucleus is rather transparent to the pion at these energies; that is, the pion essentially sees only the nucleons with which it interacts and is not strongly affected by other nuclear interactions. Thus nucleon correlations should show up most clearly and with least distortion here.

In 1984 a team of Canadian, American and Israeli physicists, working at the Tri-University Meson Facility in Vancouver, measured[1] the double charge-exchange reaction cross section on carbon-14 at 50 MeV. They used TRIUMF's time projection chamber, with its large solid angle, to