CERN EXPERIMENT CLARIFIES ORIGIN OF CP SYMMETRY VIOLATION

Experimental and theoretical floodgates were opened in 1957 by the discovery that parity inversion (*P*) and charge conjugation (*C*) are not among nature's elite of inviolate symmetry transformations. By contrast, the even more surprising discovery of *CP* violation seven years later has been something of a disappointment. In the quarter-century since James Cronin, Val Fitch, James Christenson and Rene Turlay at Princeton discovered that the decay of neutral K mesons is not precisely invariant under the combined operations of parity inversion and charge conjugation, *CP* violation has been seen in no other physical system. Nor have the experimenters been able to clarify the mechanism underlying this exotic asymmetry of the elementary particles.

When it was first observed, *CP* violation in K^0 decay had no natural place in the theory of the day. Lincoln Wolfenstein (Carnegie Mellon University) soon put forward his *ad hoc*, phenomenological "superweak theory," attributing *CP* violation to a new interaction, so much weaker than the ordinary weak interactions that one could hope to see it only in the peculiar neutral-kaon system, with its bizarre coupling between particle and antiparticle. But now, of course, we have a full-blown "standard model" that purports to be a complete description of the strong, electromagnetic and weak interactions of the elementary particles. This standard model incorporates quantum chromodynamics and the electroweak theory with three generations of quarks and leptons. What does it have to say about *CP* violation?

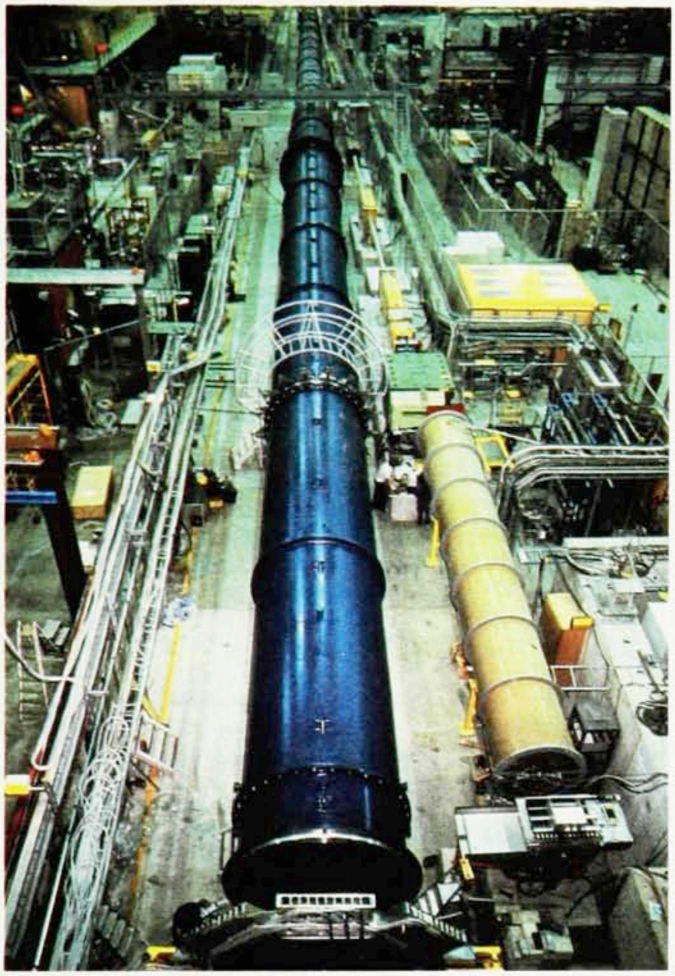

Evacuated decay region for neutral K mesons produced at the CERN Super Proton Synchrotron. The blue vacuum vessel is more than 100 meters long and two meters wide. A beam of high-energy long- or short-lived neutral kaons is allowed to decay in the upstream (top of picture) half of the vessel. Decay products (mostly pions and photons) proceed through the vessel to an array of detector planes beyond the downstream end. The first wire-chamber plane is partially visible at the bottom center of the picture. Short-lived kaons are produced at a target on rails that moves from place to place inside the vessel. Long-lived kaons are produced far upstream of the decay region.

It turns out that the standard-model predictions for *CP* violation in K^0 decay differ from those of the old superweak theory. But the differences are so subtle that no experiment—until this year—has ever detected them. Every experimental result since 1964 was consistent with the superweak theory. None was sensitive enough to detect the departures from the superweak theory demanded by the standard model.

The superweak theory is not the issue. It has never been regarded as more than a phenomenological stopgap pending a more comprehensive theory. Now that the standard model has become a serious candidate for such a theory, its proponents have been eagerly awaiting the delicate confirmation of its predictions in *CP*-violating K^0 decay. The basic question is: can *CP* violation be understood as a subtle manifestation of the well-known weak interactions, or must we invoke a new and separate regime of superweak physics, with quanta much heavier than the weak intermediate vector bosons?

First evidence

Under the provocative title "First Evidence for Direct *CP* Violation," a recent issue of *Physics Letters B* carries the report of an experiment at CERN yielding the first observation of a significant departure from superweak predictions in *CP*-violating K^0 decays.[1] A CERN-Dortmund-Edinburgh-Mainz-Orsay-Pisa-Siegen collaboration, headed by Heinrich Wahl (CERN), has measured the *CP*-violating parameter ratio ε'/ε to be $(+3.3 \pm 1.1) \times 10^{-3}$, in good agreement with the prediction of the standard model. The superweak theory, by contrast, requires that ε' be identically zero. Experimenters had been looking for a nonvanishing ε' for two decades. But before this new result no experiment has had sufficient sensitivity to establish a significant departure from zero.

Why is this new result described as *direct* *CP* violation, when the *CP*-violating two-pion decay of K_L, the longer lived of the two neutral kaons, has been known since 1964? The two *CP* eigenstates of the neutral kaon, K_1 and K_2, are the obvious linear superpositions of K^0 and its antiparticle $\overline{K}^0$, the eigenstates of strangeness, which is not conserved in the weak interactions.

$$K_1 = (K^0 + \overline{K}^0)/\sqrt{2}$$
$$K_2 = (K^0 - \overline{K}^0)/\sqrt{2}$$

K_1 and K_2 are, respectively, the $+1$ and -1 eigenstates of *CP*. Before the discovery of *CP* violation, it was simply assumed that they were also the short- and long-lived mass eigenstates of definite lifetime, K_S and K_L, respectively. Fitch and Cronin's surprising observation that about one K_L in every 300 will decay into two pions (which *must* be in a $+1$ state of *CP*) implied that K_L carries a small admixture of the "wrong" *CP* eigenstate K_1. This small *CP*-violating admixture is described by the parameter ε.

$$K_L \approx K_2 + \varepsilon K_1$$

In the superweak theory, this mixing of *CP* eigenstates in the decay eigenstates is the *only* source of *CP* violation. The small mixing parameter ε tells the whole story. This mixing implies *CP* violation in the weak coupling between K^0 and $\overline{K}^0$, particles differing in strangeness (S) by two units. The superweak theory attributes all *CP* violation to such unusual $\Delta S = 2$ transitions: K_L occasionally decays to the "forbidden" two-pion state only because a *CP*-violating $\Delta S = 2$ coupling has given it a small admixture of K_1. The actual final decay mechanisms of the component *CP* eigenstates K_1 and K_2, which involve only the usual $\Delta S = 1$ transitions, are taken to be rigorously *CP* conserving in the superweak scenario.

In the standard model the *CP*-violating mixing parameter ε is also due to the $\Delta S = 2$ coupling between K^0 and $\overline{K}^0$. But beyond this mixing, the standard model predicts an additional "direct" *CP* violation in the $\Delta S = 1$ decays of the components K_1 and K_2 themselves. This very small effect, which distinguishes the standard model from the superweak prediction, is parametrized by ε'.

In terms of the *CP*-forbidden and allowed decay amplitudes A of the short- and long-lived neutral kaons, this parametrization gives

$$A(K_L \to \pi^+\pi^-)/A(K_S \to \pi^+\pi^-) = \varepsilon + \varepsilon'$$

$$A(K_L \to \pi^0\pi^0)/A(K_S \to \pi^0\pi^0) = \varepsilon - 2\varepsilon'$$

Both ε and ε' are complex numbers. The measured value of ε is about $0.0023\, e^{i\pi/4}$. This has sufficed to describe all the *CP*-violating data reported before 1988, with ε' taken to be zero in the spirit of the superweak theory. The standard model, on the other hand, predicts that ε' should be a few hundred times smaller than ε, with roughly the same phase. The detailed prediction depends on the mass of the as-yet-unseen top quark, and on hadronic complications that are difficult to calculate with assurance.

Box diagrams and penguins

The standard model, as it has evolved from the experimental discoveries of the 1970s, asserts that the hadrons are made from six species of quarks paired into three generations, each heavier than its predecessor. The quotidian world of ordinary matter is accounted for by the up and down quarks. K mesons and other carriers of strangeness or charm require the strange and charmed quarks of the second generation. The third generation is confidently presumed to encompass the bottom and top quarks. But vain searches for the top quark and the recently observed high level of mixing between the bottom meson and its antiparticle (PHYSICS TODAY, August 1987, page 17) tell us that the top quark is at least 50 times heavier than the proton.

The six quarks, as they manifest themselves in the spectroscopy of the hadrons, are not quite the eigenstates of the weak interaction. The 3×3 intergenerational mixing matrix that gives the weak eigenstates as superpositions of the hadronic quark eigenstates is named for Makoto Kobayashi and T. Maskawa, who developed this formalism at the University of Nagoya in 1973, before there was any

Contributions to CP Violation in K^0 Decay

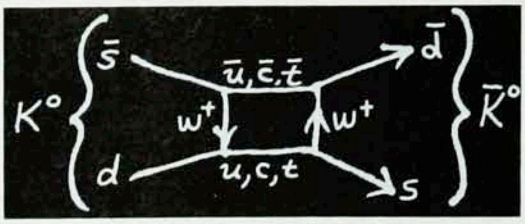

Box diagram describes fourth-order weak transition from K^0 to its antiparticle, responsible for the small admixture of the "wrong" *CP* state in the long-lived neutral kaon. The down and antistrange quarks making up the initial K^0 exchange a weak vector boson W. The intermediate state, before a second W is exchanged, involves all three generations of charge $+\frac{2}{3}$ quarks: up, charmed and top.

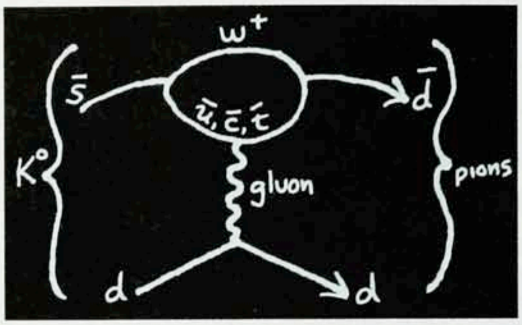

Penguin diagram introduces *CP* violation *directly* into the transition from K^0 to pionic states. The antistrange quark in the decaying K^0 becomes an antidown quark via a loop involving a W and all three quark generations. The exchange of a gluon couples the loop to the "spectator" down quark. A quark can alter its identity only by coupling to the weak intermediate vector boson W. This diagram generates the *CP*-violating parameter ε', which distinguishes the predictions of the standard model from those of the superweak theory.

experimental hint of a third generation. With admirable prescience they pointed out that if there are at least three quark generations, this mixing matrix must contain an irreducible phase angle that would account naturally for *CP* violation.

Therefore *CP* violation in the standard model arises from weak-interaction Feynman diagrams in which all three quark generations can be exchanged. The principal contributor to the mixing parameter ε is the "box" diagram in the inset on page 18, which describes the $\Delta S = 2$ transition between K^0 and $\bar{K}^0$. All three quark generations can participate in the intermediate state, thus bringing into play the *CP*-violating phase angle of the Kobayashi–Maskawa matrix. This fourth-order weak $\Delta S = 2$ coupling is an explicit mechanism for Wolfenstein's phenomenological superweak *CP*-violating interaction in the modern language of the standard model.

But unlike the *ad hoc* superweak theory, the standard model presents us with an additional *CP*-violating mechanism, described by the "penguin" diagram in the inset. This second-order contribution to the $\Delta S = 1$ transition of a K^0 (or $\bar{K}^0$) to pions makes possible the "direct" *CP*-violating transition $K_2 \to 2\pi$ because the loop once again involves all three quark generations. More specifically, because the penguin diagram contributes only to the production of pionic final states with isotopic spin zero, it generates a phase difference between K^0 transition amplitudes to different charge states of two pions. It is this phase difference, which would be forbidden if *CP* were conserved, that gives a nonvanishing ε'.

The experiments

In terms of the two-pion decay rates counted by the experimenters, the real part of ε'/ε is given to good approximation by

$$\mathrm{Re}\,(\varepsilon'/\varepsilon) = \tfrac{1}{6}(1 - R)$$

R is the experimental "ratio of ratios" given by

$$R = \frac{\Gamma(K_L \to 2\pi^0)/\Gamma(K_L \to \pi^+\pi^-)}{\Gamma(K_S \to 2\pi^0)/\Gamma(K_S \to \pi^+\pi^-)}$$

The numerator is the ratio of decay rates (Γ) for the neutral and charged modes of the *CP*-forbidden decay of the long-lived neutral kaon. The denominator is the corresponding ratio for the *CP*-allowed decay of the short-lived neutral kaon. Both ratios are close to 46%, giving an *R* so close to unity that finding a nonvanishing ε' is a daunting experimental undertaking.

About 99 percent of all K_S mesons decay to two pions, whereas more than 99 percent of K_L mesons decay into negative-*CP* three-body final states. Less than 0.3 percent of all K_L decays are to the *CP*-forbidden two-pion states. The lifetime of the K_L is 5×10^{-8} seconds, roughly 600 times that of the K_S. This great disparity of lifetimes is due largely to the very skimpy phase space available for the three-body decay of a particle that weighs little more than three pions.

Each of the four decay rates in the ratio of ratios presents its own peculiar experimental problems. The various experiments that have sought a nonzero ε' in recent years have grouped the four modes into pairs to be recorded simultaneously, so that the various experimental uncertainties would cancel one another as much as possible in the determination of *R*. The CERN collaboration, for example, undertook to measure the neutral and charged decays simultaneously. Ordinarily one measures the momentum of charged pions by tracking their curvature through a magnetic field. Neutral pions cannot be measured this way; in fact each π^0 decays into two photons before it can travel a macroscopic distance. Seeking to confront the charged and neutral decay modes with essentially the same detector geometry, the CERN collaboration has chosen to do without a magnetic field, so that the raw data will not have to be subjected to elaborate, mode-dependent geometric acceptance corrections.

An advantage of this design is that the necessary geometric corrections are quite small, not requiring the demanding Monte Carlo computer simulations needed to interpret the data from a magnetic detector. Magnetic fields, on the other hand, not only facilitate the measurement of charged tracks; they also serve to sweep junk out of the way. In early days there was some concern that the proliferation of unwanted background tracks in the absence of a large magnet might render the CERN design unfeasible.

Because they measure the charged and neutral dacay modes at the same time, the CERN experimenters run their K_S and K_L beams at different times. The opposite strategy has been chosen by a group carrying out a rival experiment at Fermilab. This University of Chicago-Fermilab-Princeton-Saclay collaboration, led by Bruce Winstein (Chicago), has been measuring neutral and charged decay modes at separate times while running parallel K_S and K_L beams into their detector simultaneously, a few inches apart. The first results of the Fermilab collaboration, based on a very small, early fraction of their total data sample, were reported[2] in April, three weeks before the CERN group's paper. The preliminary Fermilab result:

$$\varepsilon'/\varepsilon = (+3.2 \pm 2.8\,[\text{statistical}] \pm 1.2\,[\text{systematic}])\times10^{-3}$$

is certainly consistent with the CERN result, but the analyzed data sample is as yet too small for one to call it a confirmation.

Long- and short-lived beams

In both of these experiments a high-energy proton beam produces neutral K^0 and $\bar{K}^0$ mesons, among a profusion of other hadronic debris, by impinging upon a metal target. After the debris is removed by sweeping magnets and collimators, one is left with a beam of neutral kaons with a fairly broad energy spectrum centered near 100 GeV. At these energies a K_S will decay, on average, after traveling 6 meters, while a K_L typically lasts for several kilometers. Therefore, when the CERN group wants to run a K_L beam into its detector complex, it puts the production target into the proton beam some 100 meters upstream of the fiducial decay volume, so that less than one K_S in a million survives the journey.

During the eight-hour runs that call for a K_S beam, the upstream K_L production target is replaced by a movable K_S production target inside the 50-meter-long decay region of the detector complex. Because the mean K_S decay length is much shorter than the fiducial decay volume of the detector, a stationary K_S target would result in a concentration of K_S decays near the target. But to minimize biases one wants the distribution of K_S decay vertices in the detector to approximate the almost flat longitudinal distribution of the K_L decays. Therefore, during the K_S runs the production target is moved on a track from one position to another along the 50-meter axis of the fiducial decay volume.

The beam produced at the K_L target far upstream has a K_S contamination of less than 1 in 10^6 by the time it reaches the detector. But when the K_S target is in place, fully one or two percent of the decays observed in the detector are actually K_L decays. This would be a problematic contamination in an experiment looking for exceedingly small effects, except that 99.7 percent of all the unbidden K_L decays in the K_S beam are three-body decays that the detector is explicitly designed to reject.

The principal function of the CERN collaboration's detector complex is to identify neutral kaon decays to two pions, either charged or neutral. The chief problem is to catch and eliminate *CP*-allowed three-body K_L decay modes masquerading as two-pion decays. One might, for example, fail to detect a neutral pion in the common decays $K_L \to 3\pi^0$ or $K_L \to \pi^+\pi^-\pi^0$, or one might misidentify an electron or a muon as a second charged pion in the semileptonic decay modes $K_L \to \pi^+ e^- \nu$ or $K_L \to \pi^+ \mu^- \nu$.

The evacuated decay region of the CERN detector is followed by two planes of wire chambers in a helium atmosphere, 25 meters apart, to register the passage of charged decay products and permit their extrapolation back to their point of origin. After the wire planes comes an electromagnetic calorimeter consisting of liquid argon interspersed with thin sheets of lead, whose principal role is to identify and measure the positions and energies of electrons and photons. This electron–photon calorimeter is followed by a hadron calorimeter, made of alternating planes of iron and plastic scintillator, that does the same for charged pions. Next comes an iron plate too thick to be penetrated by pions or electrons, followed by muon veto counters that assume any charged particle able to reach them must be a muon. The sides of the decay region are also lined with veto counter rings to warn of extra decay products coming off at angles too wide to reach the calorimeters.

One identifies a two-pion decay essentially by determining that the sum of the four-momenta of the decay products is consistent with that of a beam kaon. The absence of a magnetic bending field complicates this process somewhat for the charged-pion decay modes. It turns out, however, that measuring the angle between the charged-pion trajectories and the energy that each deposits in the hadron calorimeter is sufficient to distinguish $K_L \to \pi^+\pi^-$ from background, if one discards events in which the two pions lose very different amounts of energy in the calorimeter—a cut that pares away roughly half of all $\pi^+\pi^-$ decays.

When a kaon decays to $2\pi^0$, the detector sees only the four photons from the decays of the two neutral pions. They are, of course, not seen by the wire chambers. But their energies and positions are measured well enough in the liquid argon calorimeter to permit the complete reconstruction of the event. This reconstruction is facilitated by the extra requirement that the four-momenta must match up for the two $\pi^0 \to 2\gamma$ decays as well as for the initial kaon decay. The severest background plaguing the measurement of the *CP*-forbidden decay $K_L \to 2\pi^0$ comes from the much more abundant $K_L \to 3\pi^0$ decays, when one π^0 evades the clutches of the detector. This happens most often when a $3\pi^0$ decay occurs near the downstream end of the detector's fiducial decay volume.

The data sample encompassed in the CERN collaboration's report of its first significant evidence for a nonzero ε' included more than three million K_S decays and almost half a million *CP*-violating K_L decays. This enormous accumulation of two-pion decays yielded a value of

$$R = 0.980 \pm 0.004 \text{ [statistical]} \pm 0.005 \text{ [systematic]}$$

for the double ratio. This is the first experiment that has claimed a value of R significantly different from unity. Of the four decay rates that make up R, only the $K_L \to 2\pi^0$ rate required a background subtraction exceeding 1 percent.

At this point, the group's measurement of ε'/ε is only three standard deviations away from zero. Such a portentous result cries out for confirmation. The CERN collaboration has in fact begun a new run, which will double their statistics. More important are the improvements introduced to beat down the systematic errors. The experimenters have augmented the shielding and collimation downstream of the production targets to reduce the charged-particle debris that contaminates the two-pion decay signal. This lets them accept decays occurring closer to the entrance of the detector. Thus they can, without losing statistics, reject $K_L \to 2\pi^0$ events at the downstream end, where the $3\pi^0$ contamination is worst. The collaboration has also installed a transition-radiation detector just before the liquid argon calorimeter, thereby achieving a tenfold improvement in the rejection of electrons posing as pions.

At Fermilab

The Fermilab experimenters, having already finished taking their data, believe they have on tape a somewhat larger sample of usable decay events than the CERN group will have gathered. Early next year they plan to report on about 20 percent of this final sample. In the Fermilab experiment, a regenerator is flipped back and forth between two adjacent K_L beams, so that one of the beams acquires a fresh K_S component. Thus K_L and K_S decays are measured at the same time. In the early phases of the experiment, neutral and charged decay modes were gathered in separate runs. The problem is to determine the position of decay vertices well enough to know unambiguously from which beam a decay originated. To this end the Fermilab experiment employed a thin lead sheet in its neutral-pion runs to convert decay photons to electron–positron pairs for better determination of trajectories. Eventually the group found they could dispense with the lead sheet and still have adequate vertex resolution. This simplification greatly increased their data-taking rate and made it possible to collect the neutral and charged-pion modes simultaneously, with all the attendant cancellation of systematic errors.

To predict ε'/ε from the standard model, one must calculate the matrix element of the penguin diagram sandwiched between real hadronic initial and final states. This brings in all the uncertainty of how quarks are distributed in physical particles. The first detailed calculation in this spirit was carried out in 1979 by Fred Gilman and Mark Wise at Stanford. With regard to a theoretical lower limit, one can at present say only that the ε'/ε should not be much smaller than 1×10^{-3}.

Because the top quark plays a dominant role in the imaginary part of the penguin diagram, the theoretical upper limit on ε'/ε depends on the mass of that still elusive sixth quark. The prevalent opinion among theorists and experimenters that the top quark is heavier than 50 GeV translates into an upper limit of about 1×10^2 for ε'/ε. Turning this around, one can argue that the first nonzero measurement of ε' has now strengthened the belief that the top quark is relatively heavy. In any case, the CERN collaboration's result takes away much of the mystery of *CP* violation and leaves the standard model of the elementary particles a little more secure.

"After almost 25 years of searching for additional information about the origin of *CP* violation," comments Gilman, "we now have the first evidence. It points to an origin of the phenomenon in physics at the scale of the weak interactions rather than some new superweak scale."

—Bertram Schwarzschild

References

1. H. Burkhardt *et al.* (CERN, Dortmund, Edinburgh, Mainz, Orsay, Pisa, Siegen collab.), Phys. Lett. B **206**, 169 (1988).
2. M. Woods *et al.* (Chicago, Fermilab, Princeton, Saclay collab.), Phys. Rev. Lett. **60**, 1695 (1988). ■