The Tau Neutrino Has Finally Been Seen

Twenty-five years after the discovery of the tau lepton ($\tau^{\pm}$), its neutrino (ν_{τ}) has at long last been detected. At a Fermilab seminar in July, Byron Lundberg of the international DONUT (Direct Observation of Nu Tau) collaboration reported that five months of exposing a photographic-emulsion target to an intense neutrino beam from the Tevatron had netted four identifiable ν_{τ} interactions.[1]

> The existence of the third neutrino was hardly in doubt. But it took the sub-micron precision and electronic sophistication of a modern photo-emulsion target to find the decay of its telltale collision product.

The existence of the tau neutrino is in no sense a surprise. Quite the contrary. If it *didn't* exist, the spectacularly successful "standard model" of particle theory would be in grave trouble. But detecting a ν_{τ} interaction is unusually difficult; and finding a reliable way to do it has important implications for major experiments now on the drawing boards.

Like the other two neutrino varieties (ν_e and ν_{μ}), the tau neutrino rarely interacts with matter. But even when it does, the experimenter's problem is greatly compounded by the fact that the τ, whose appearance is the telltale sign that a nucleus was struck by a tau neutrino, has a lifetime of only 3×10^{-13} s. So, at Tevatron energies, one needs to ferret out charged particle tracks that decay within about a millimeter of their production point.

That's where photographic emulsion comes in. A charged particle traversing photographic emulsion dissociates silver-salt grains much as light exposes ordinary film. Emulsion can provide three-dimensional tracking of particle trajectories and their decay kinks with sub-micron precision. Nuclear physicists were using it half a century ago. But traditional emulsion scanning techniques are far too labor intensive for the modern study of very rare events. The pioneering DONUT experiment incorporated electronic tracking elements that identified candidate events and pinned down their vertex positions to within a few millimeters. After weeks or months of exposure, emulsion layers were removed from the target, and the regions pinpointed by the electronic tracking were scanned and digitized by automated microscopes.

Perhaps as important as DONUT's verification of the tau neutrino's existence is the experiment's demonstration that modern emulsion targets are a feasible means of searching for the metamorphosis of ν_{μ} into ν_{τ} in accelerator-based neutrino-oscillation experiments. (See PHYSICS TODAY, August 1999, page 9.)

One such experiment, called CHORUS, ran at CERN's SPS accelerator from 1994 to 1997 without finding any neutrino oscillation to ν_{τ}. The CHORUS emulsion target, built by Kimio Niwa's group at the University of Nagoya, was a precursor of the DONUT target. Niwa now heads DONUT's Japanese contingent, which is responsible for the experiment's emulsion components.

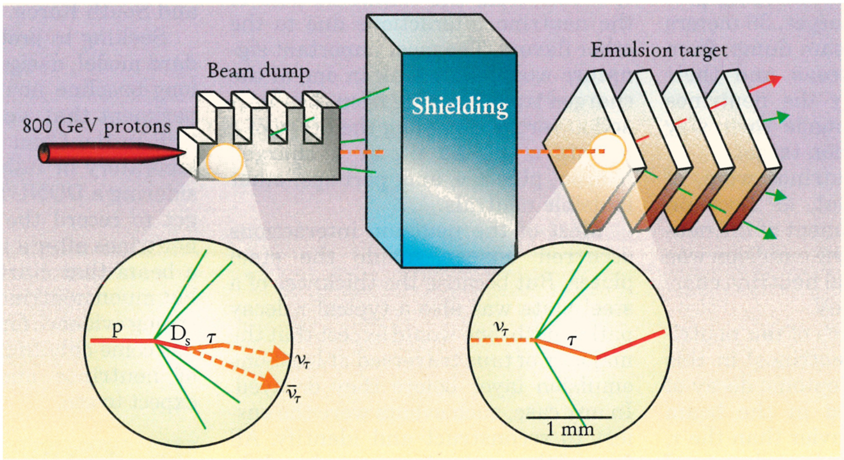

FIGURE 1. THE DONUT EXPERIMENT at Fermilab. The intense shower of particles created by ramming an 800-GeV proton beam into the tungsten beam dump is mostly absorbed in the iron shielding or in the beam dump itself. Only neutrinos get through to the photographic emulsion target. About 8% of them are tau neutrinos, produced mostly by the decay of D_s mesons in the beam dump. A tiny fraction of the tau neutrinos collide with nuclei in the emulsion target, creating telltale charged tau leptons that decay after about a millimeter.

Three charged leptons

In 1962 we learned that neutrinos come in at least two distinct flavors associated, respectively, with the electron and the muon. Except for the still unexplained fact that the muon ($\mu^{\pm}$) is more than 200 times heavier than the electron (or positron), the two charged leptons appear to be almost identical twins. One can tell their neutrinos apart by the charged leptons they turn into: A nuclear collision can turn a ν_e into an electron or a ν_{μ} into a muon—but never the other way around.

In 1975, Martin Perl and colleagues at Stanford discovered the τ, which turned out to be 17 times heavier than the μ. The decay phase space afforded by this great mass accounts for the fact that the tau's lifetime is 10 million times briefer than the muon's.

There's good reason to believe that the τ is the last of the charged leptons. The standard model asserts "lepton universality." That is to say, except for properties that depend explicitly on mass, the three charged leptons are presumed to have precisely equivalent interactions. From this— and from a lot of indirect experimental evidence—it follows that the τ must have a neutrino of its own. And now we've finally seen it.

The experiment

The DONUT emulsions were exposed at the Tevatron in the spring and summer of 1997, and the arduous data analysis is now nearing completion. The experimental setup, shown schematically in figure 1, began with an intense pulsed beam of 800-GeV protons terminating in a tungsten "beam dump" target and producing a myriad of assorted hadrons. Deflect-

© 2000 American Institute of Physics, S-0031-9228-0010-320-5

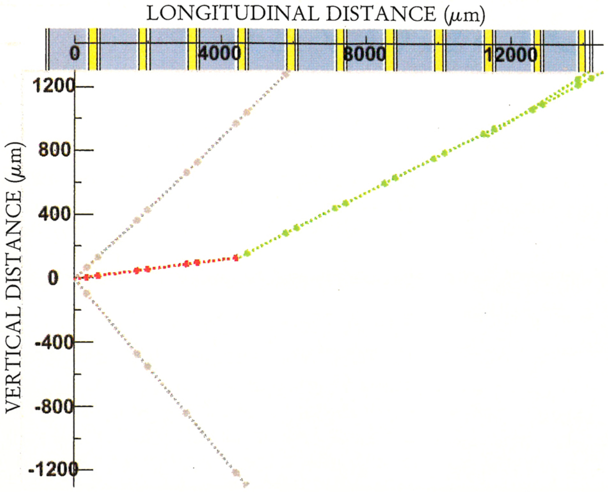

FIGURE 2. A TAU NEUTRINO EVENT, one of the four identified in DONUT's emulsion target. Three charged tracks emerge from a collision between a ν_τ and a nucleus in one of the steel plates: The 4.5-mm track shown red is a tau lepton that decays into a roughly 4-GeV electron (green) plus two unseen neutrinos. The two gray tracks are presumably hadrons. The striated bar indicates the steel (blue), plastic (yellow), and emulsion (white) layers of the target. Dots indicate tracks found in the 100-μm-thick emulsion layers.

ing magnets and iron shielding protected the emulsion target, 36 meters downstream of the beam dump, from this shower of hadrons and their decay products. Only the neutrinos (plus a few muons) made their way through to the emulsion target.

Most of these neutrinos were, of course, ν_e and ν_μ. But, at Tevatron energies, the ν_τ component of the neutrino flux traversing the emulsion was as high as 10%. Typical neutrino energies were about 50 GeV.

The principal source of tau neutrinos (and their antineutrinos) in this experiment was the leptonic decay of $D_s^\pm$ mesons produced in the beam dump. About 10% heavier than the τ, the short-lived D_s is a bound state of a charmed quark and a strange antiquark—a combination that favors leptonic decay. Although π and K mesons were by far the most abundantly produced hadrons in the beam dump, most of them were reabsorbed in the long, dense tungsten target before they could decay.

The emulsion target was a row of four 7-cm-thick modules interleaved with scintillation-fiber tracking planes that electronically recorded the coordinates of charged tracks. Additional electronic tracking, timing, and particle identification were accomplished by a downstream spectrometer complex of wire chambers, muon detectors, and a lead glass calorimeter. A typical emulsion module was a stack of emulsion-coated plastic plates alternating with 1-mm-thick steel plates. Each emulsion coating layer was only 100-μm thick. The steel was there to provide most of the mass for neutrino collisions with nuclei.

In five months of running, with a total of 4×10^{17} protons bombarding the beam dump, the electronic triggering and tracking signaled about 700 candidate neutrino interactions in the emulsion target. After microscopic digitization of the emulsion regions localized by the scintillation-fiber tracking data, 203 candidates survived as neutrino interactions—or plausible impostors.

A few good kinks

Now the group had to pick out the few ν_τ events from the estimated 95% of the neutrino interactions due to the other flavors. The most important signature would be a kink in one of the charged tracks coming from the interaction vertex, signaling the decay of a $\tau^\pm$ into a $\mu^\pm$, an $e^\pm$, or a charged hadron, plus a ν_τ and, perhaps, other invisible neutrals.

Most of the neutrino interactions occurred unseen inside the steel plates. But because the thickness of a steel plate was also a typical τ decay path length, one would expect that the majority of taus traversed at least one emulsion layer before they decayed. In any case, such events were the easiest to reconstruct and identify. All four of the ν_τ events reported in July were of this kind. Figure 2 shows one of them. A fifth strong candidate event, in which the τ appears to decay before it leaves the steel plate in which it was born, is still under consideration.

Even though the primary interaction vertex is rarely seen, one can reconstruct its position, as well as the telltale kink vertex, quite precisely from the tracks recorded in the emulsion layers. In a true ν_τ event, the other primary tracks accompanying the τ should all be hadrons. So the group discarded all events in which one of the unkinked tracks was unmasked as a lepton by the downstream muon detector or the calorimeter. That left fewer than 30 candidate events with well-documented kinks.

The largest remaining group of impostors were ν_e or ν_μ events in which one of the hadron tracks was kinked by a nuclear collision. But such scattering, especially without a detectable recoil track, would almost always impart a transverse momentum kick of less than 250 MeV/c to the hadron track. Imposing a 250 MeV/c transverse-momentum cut separated the four surviving ν_τ events cleanly from an obvious background of hadron scattering events piled up at very low transverse momenta. Any surviving background due to hadronic scattering and other things like charmed-meson decays or random superpositions, the group concluded, could account for less than half an event.

"So we finally have direct evidence that the tau neutrino really exists," Lundberg told us, "and, as far as we can tell, that its interactions with matter are what the standard model predicts." Lundberg (Fermilab) and Vittorio Paolone (University of Pittsburgh) are the two spokesmen for the DONUT collaboration, which, in addition to its US and Japanese contingents, includes groups from Greece and South Korea.

Seeking to probe beyond the standard model, designers of the proposed long-baseline neutrino-oscillation experiment that would link the CERN SPS to the Gran Sasso underground laboratory in Italy are seriously considering a DONUT-type emulsion target to record the appearance of tau neutrinos after a journey of 700 km in a beam that starts out with nothing but muon neutrinos. Until now, all the strong evidence for neutrino oscillation has come only from the *disappearance* of neutrinos one would otherwise expect to see.

Has the Higgs also been seen?

As LEP, the large electron–positron collider at CERN, approached its scheduled final shutdown at the end of September, several of the LEP detector groups were reporting tentative evidence for a Higgs boson with a mass near 114 GeV/c^2. Combining the data of all four groups, one gets an enhancement above background of less than 3 standard deviations. To increase statistics at the collider's highest beam energy, the CERN directorate decided on 14 September to reprieve LEP until 2 November.

We've now seen all the fundamental fermions (the 3 charged leptons, their neutrinos, and the 6 quarks) required by the standard model; and all the spin-1 gauge bosons that mediate their interactions. Of all that the standard model demands, the only ingredient still to be found is the spin-0 Higgs boson. But as we go to press, tantalizing hints of a low-mass Higgs are being reported at CERN. (See the box on page 18.)

Bertram Schwarzschild

Reference

1. See http://fn872.fnal.gov/.

Dendritic Signals Start Out Stronger When They Have Farther to Go

The trillion or so neurons in our brains are input–output devices. Feeding each neuron with information is a tree-like body of fine filaments called dendrites, whose trunk originates at the neuron's cell body or soma. Also attached to the soma is the axon, which carries the processed information to other cells—mostly the dendrites of other neurons via junctions called synapses.

Dendritic arbors (as the tree-like structures are known) can be quite extensive, up to a millimeter long and connected to more than 10 000 axons. Somehow—it's not clear exactly how—dendrites integrate the inputs they receive to produce a single command for the neuron: Signal! (or not, as the case may be).

Inevitably, because of the arbor's shape, some dendritic terminals are farther from the soma than others (see the cover). If signals moved along dendrites without attenuation, then the synapse-to-soma distance wouldn't matter. But that's not the case. Lacking the insulating layer of myelin that covers axons, dendrites act like leaky wires. Without a method for boosting their signals, distant synapses would not be heard over the din of their closer cousins.

Apparently, nature has either (a) built a system with ineffectual components or (b) provided a method for boosting signals. The answer is *b*. In last month's *Nature Neuroscience*,[1] Jeff Magee (Louisiana State University) and Erik Cook (Baylor College of Medicine) confirm that many signals arrive at the soma with the same strength regardless of how far away in the dendritic arbor the signals originate. They also found out how nature pulls off this trick: Rather than being topped up along the way with a system of repeaters (analogous to the devices used in long-distance optical and electric cables), the signals actually start out stronger.

Because such a method requires that the synapse know how far away the soma is, Magee and Cook's result provides further evidence of "backward" information flow in neurons—

> Without a way to compensate for leaking charge, neurons wouldn't be able to receive their most distant inputs.

that is, from the soma to the dendrites. This route, discovered six years ago, is a key ingredient of models that seek to determine the cellular basis of learning and memory.

Propagation primer

In general, the signals that travel through neurons consist of temporary changes in the voltage across the cell membrane. These voltage changes occur in response to changes in the permeability of the membrane to various ions, principally potassium and sodium. At rest—that is, when the neuron isn't signaling—specialized membrane transport proteins enforce the ionic status quo: a higher K^+ concentration inside the cell than outside; a lower Na^+ concentration inside than outside; and a voltage difference of about −70 mV between the cell's exterior and interior.

To send a signal, neurons disrupt this status quo in one of two main ways. What's termed active propagation begins when Na^+ channels are triggered to open by a signal at the soma. Propelled by both their concentration gradient and the membrane's negative potential, Na^+ ions enter the cell, reducing the voltage difference and causing more Na^+ channels, which are voltage sensitive, to open. Na^+ ions rush into the cell until the polarity of the membrane is reversed. At that point, K^+ ions come into play. Voltage-sensitive K^+ channels open to let K^+ ions out of the cell and restore the original membrane potential.

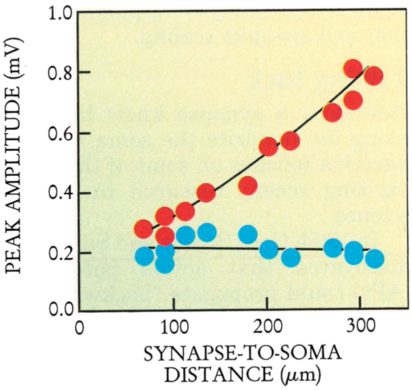

FIGURE 1. THE STRENGTH of an excitatory postsynaptic potential (EPSP) measured simultaneously at the soma (blue) and synapse (red) for various synapse-to-soma distances. The figure demonstrates that distant EPSPs start off stronger. (Adapted from ref. 1.)

This to-and-fro of ions doesn't happen along the entire length of a neuron simultaneously. Rather, it propagates along the neuron as a self-sustaining pulse of voltage known as an action potential. In effect, the original trigger's role is taken by the pulse itself. Because the signals regenerate themselves, active propagation is the only effective way to transmit signals down axons, which, in humans, can be as long as a meter.

Passive propagation is much simpler. Ion channels, open at one end of the neuron, let in a packet of charge that simply pushes itself along the neuron by its own Coulomb repulsion. The signal gets no help from voltage-sensitive ion channels because the voltage never climbs high enough to open them. Consequently, passive propagation is unsuitable for axons, but it plays the major role in dendrites. The passive signals that originate at synapses and flow down dendrites to the soma are known as EPSPs (excitatory postsynaptic potentials).

Patch-clamp recording

The basic principles behind active propagation were worked out in the 1940s by Alan Hodgkin and Andrew Huxley. The pair experimented on the common squid's "giant" axon, whose 1-mm diameter is large enough to accommodate wires.

Studying signal propagation in dendrites is more challenging because the diameter of a typical dendrite in the mammalian brain is at most 25 μm. For their experiments, Magee and Cook exploited a technique developed in the 1970s by Erwin Neher and