A thousand phototubes fixed to the walls stare unblinkingly into the oil, watching for the tiny flashes of light hat signal the presence of a neutrino.

A Thousand Eyes

The story of LSND

Bill Louis, Vern Sandberg, and Hywel White as told to David Kestenbaum

There is no original truth, only original error.
Gaston Bachelard (1884–1962)

If Las Vegas took bets on physics results, the odds on LSND’s 1995 announcement would have been very, very long. While no one could quite say what was amiss, few believed the experiment had detected a neutrino mass. When the LSND team made the rounds, giving talks at universities, conferences, and the national labs, it was a tough sell.

“The community was . . . intrigued, o say the least,” recalls Ion Stancu, an SND (liquid scintillator neutrino etector) collaborator from the Univerty of California, Riverside. Stancu ave one of the first talks that January o a room packed past fire codes. The uestion period ran longer than the talk self. “Some thought it was complete ubbish, others were very excited,” he ays. “In the end, all I could say was his is the data . . . take it or leave it.’” Many preferred to leave it, thinking it would just go away. Previous expernce suggested it might.

For one, the Standard Model of parcle physics states unequivocally that eutrinos are massless, and the Stanard Model had yet to be proved wrong. Although the notion that neuinos might have mass was not new—measurements of atmospheric- and olar-neutrino rates pointed to a similar onclusion—the LSND result didn’t oincide with many physicists’ expecations. Most theoretical models had to e stretched quite a bit to accommodate ll three sets of data; there just didn’t eem to be room for yet another posive result. (See the article “The Evience for Oscillations” on page 116.)

Pressed to pick the wrong results om the lineup, many in the field susected LSND’s. Rather than wait for eutrinos from the heavens, LSND manufactured its own with a kilometer-ong particle accelerator. In principle, his afforded greater control over the xperiment. In practice, it had been a esson in humility. Similar experiments ad checkered histories—their claims eemed to flit in and out of existence ke the neutrinos themselves.

Outliving this legacy would be hard, nd rumors that several LSND collaboators were questioning the results did ot help. Neither did the fact that the esults had appeared on the front page f the *New York Times* before they were made public to the physics community, all at a time when the experiment was under the budget axe. But what, then, to make of the results? The SND team argued the odds were only 1 in 1,000 that their results were wrong. Still, no one was placing bets.

LSND—A Walking Tour

Los Alamos, with its long history of defense work, may seem an odd place for the delicate task of weighing the neutrino. Above ground, in the foothills of the Jemez Mountains, there is little to betray the intricate machinery. The accelerator lies buried under a kilometer-long mound of dirt. Seen from an airplane, it recalls the inhumanly large constructions of the ancient Mayas, designed to catch the eyes of the gods. At its far end, metal blocks, planks, and bricks are stacked several meters high like the abandoned toys of a giant toddler. Most are recycled relics from the cold war—iron from magnets at Oak Ridge Lab, steel from chopped-up battleships, and counterweights from missile silo doors. In their retirement, they shield a giant underground neutrino detector from cosmic rays and an occasional rattlesnake seeking refuge in the cracks.

The detector is so well shielded that it is all but impossible to get to. In a small shack nearby, a ladder leads down a narrow metal pipe into the dry rocky ground. Climb down, crawl through another sewer-size pipe, flip the light switch, and you’ll find yourself in a small, dusty room fondly known as the Black Hole, silent except for the whir of fans cooling the electronics. The 6-meter-diameter circle of steel that forms one wall of the room is yet more shielding—the front face of a monstrous archlike shell called the cosmic-ray veto shield. Like the shielding blocks, the veto shield is a relic, this time from a previous neutrino experiment. Nestled inside it is the main detector: an enormous tank filled with 52,000 gallons of mineral oil. This is the heart of LSND, where over a thousand phototubes fixed to the walls of the tank stare unblinkingly into the oil, watching for the tiny flashes of light that signal the presence of a neutrino. If a phototube fails, it’s left for dead. Once the tank has been stuffed into the underground tunnel, only neutrinos and cosmic rays can get inside.

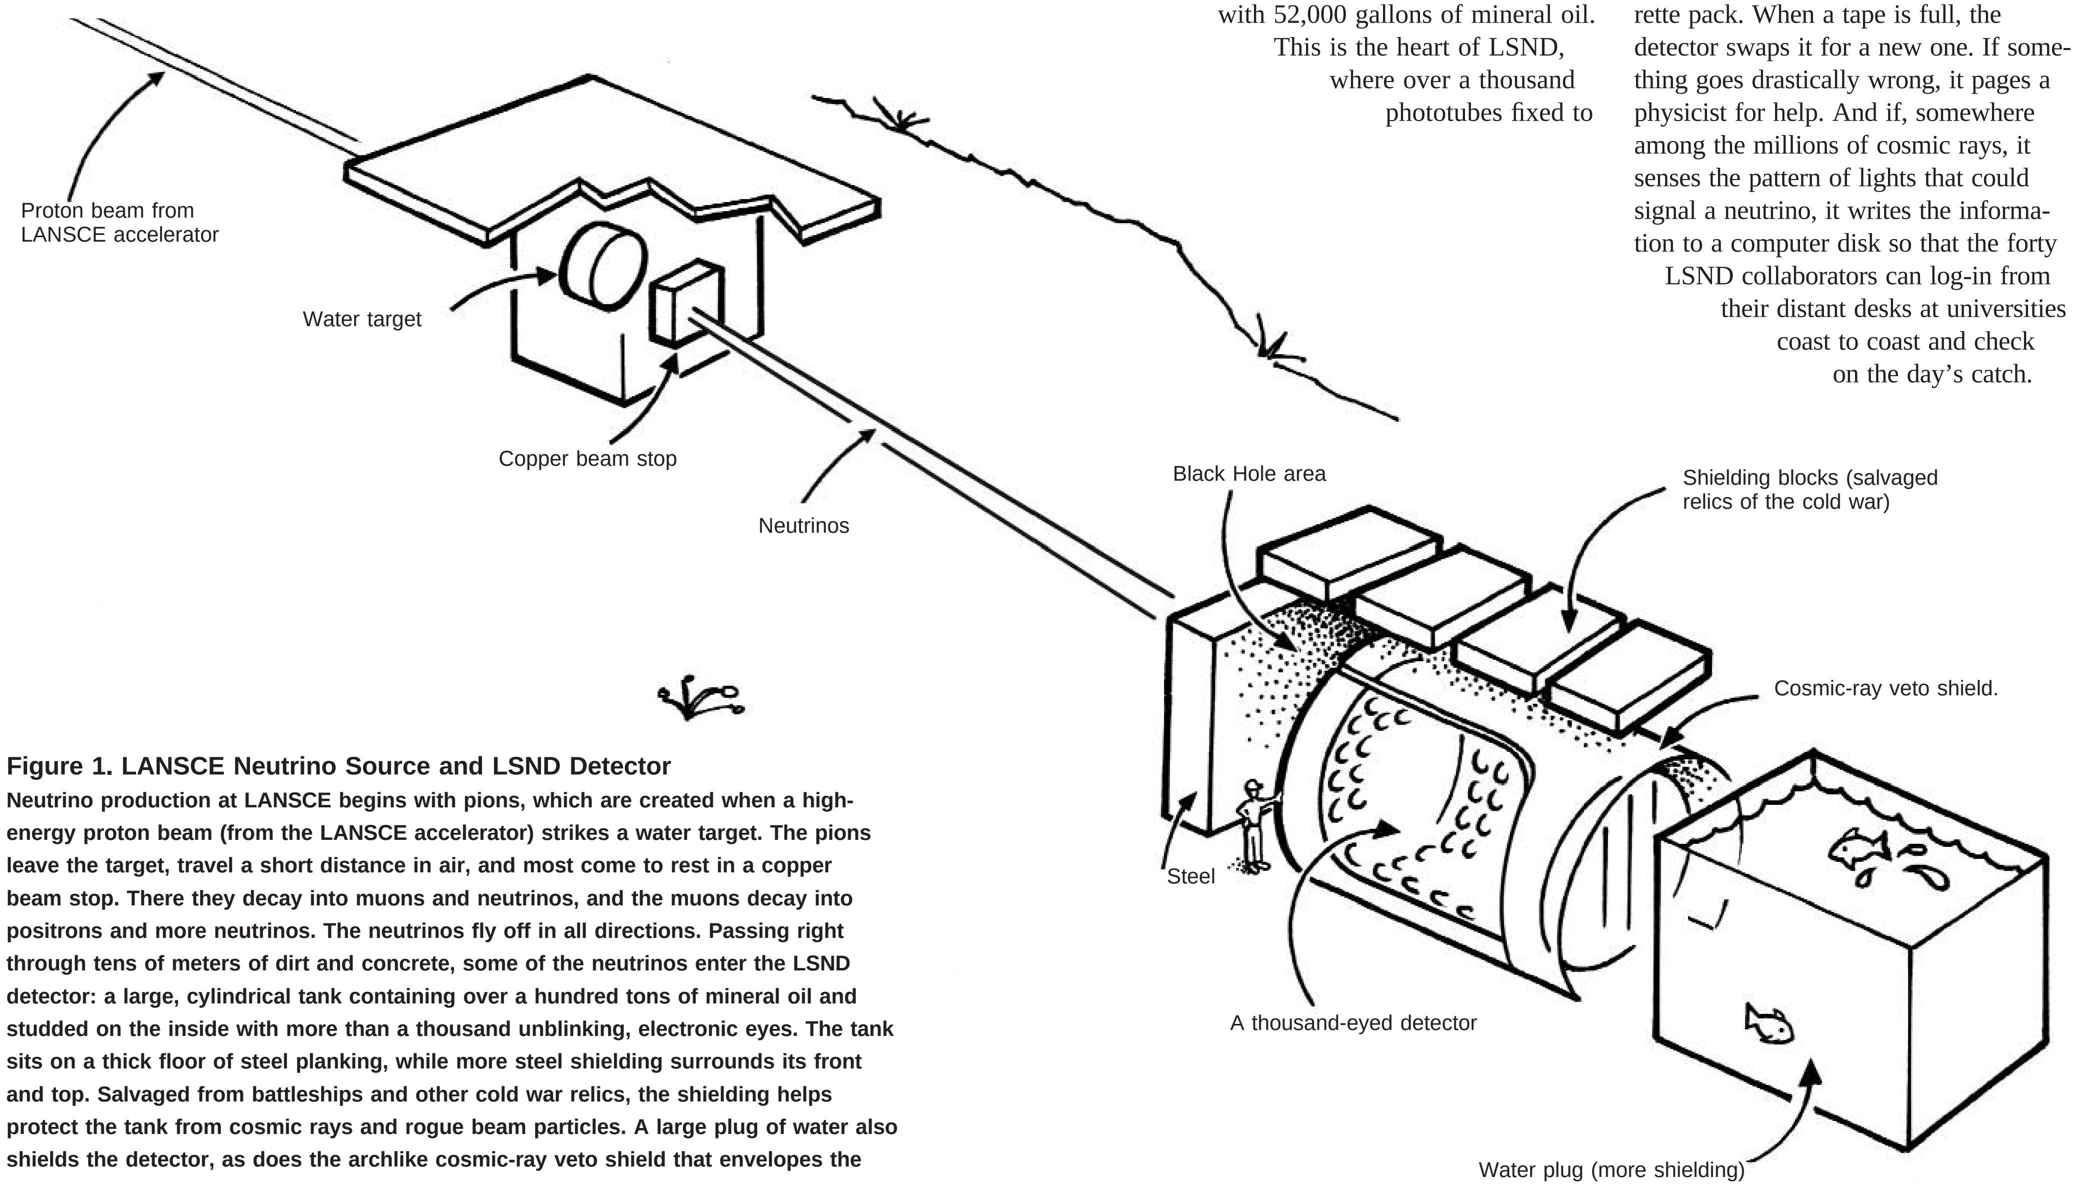

Figure 1. LANSCE Neutrino Source and LSND Detector

Neutrino production at LANSCE begins with pions, which are created when a high-energy proton beam (from the LANSCE accelerator) strikes a water target. The pions leave the target, travel a short distance in air, and most come to rest in a copper beam stop. There they decay into muons and neutrinos, and the muons decay into positrons and more neutrinos. The neutrinos fly off in all directions. Passing right through tens of meters of dirt and concrete, some of the neutrinos enter the LSND detector: a large, cylindrical tank containing over a hundred tons of mineral oil and studded on the inside with more than a thousand unblinking, electronic eyes. The tank sits on a thick floor of steel planking, while more steel shielding surrounds its front and top. Salvaged from battleships and other cold war relics, the shielding helps protect the tank from cosmic rays and rogue beam particles. A large plug of water also shields the detector, as does the archlike cosmic-ray veto shield that envelopes the tank. The veto shield clues experimenters when something other than a neutrino enters the tank.

Fortunately, after years of troubleshooting by humans, the detector can essentially take care of itself. Most of the time, it clicks happily away, analyzing the electronic pulses from the phototubes with an elaborate array of hardware and then writing the data to magnetic tapes half the size of a cigarette pack. When a tape is full, the detector swaps it for a new one. If something goes drastically wrong, it pages a physicist for help. And if, somewhere among the millions of cosmic rays, it senses the pattern of lights that could signal a neutrino, it writes the information to a computer disk so that the forty LSND collaborators can log-in from their distant desks at universities coast to coast and check on the day’s catch. Figure 1 shows a sketch of the underground accelerator and neutrino detector.

How to Weigh a Neutrino

If neutrinos have mass, it is so slight that it hardly impedes their motion. Were it possible to produce a neutrino that stood perfectly still, the tiniest tap would suffice to send it fleeing to the ends of the universe at, or close to, the speed of light. And because the neutrino is electrically neutral, it cannot be grasped with electric or magnetic fields the way electrons or protons can. The only possibility of detecting neutrinos at all is through the weak force, which is roughly one hundred million times feebler than the electromagnetic force.[1] The weak force is the agent behind all neutrino behavior—how they produce flashes of light in the tank, how they are made, and even, perhaps, how they are “weighed.”

Since it is impossible to sit a neutrino on a scale or to determine its mass by running it through the magnetic fields of a spectrometer, the neutrino can only be weighed indirectly. LSND, like atmospheric and solar experiments, looks for neutrinos to “oscillate,” a strange behavior that can betray their mass.

Neutrinos come in three varieties—the electron neutrino, the muon neutrino, and the tau neutrino. Each neutrino also has an antimatter counterpart, called the electron, muon, and tau antineutrinos. When neutrinos (or antineutrinos) oscillate, they undergo a kind of identity crisis. An electron neutrino made in the Sun, for instance, may transform enroute to the earth and present itself instead as a muon neutrino or a tau neutrino. The probability of observing one neutrino type or another varies periodically as the neutrino travels, hence the term oscillation. Oscillations can occur only if neutrinos have mass. Definitive observation of neutrino oscillations would settle the

---

[1] Neutrinos feel gravity’s tug, but too weakly to be of use experimentally.

igure 2. A Neutrino Oscillation xperiment

neutrinos oscillate, the probability of bserving a given neutrino type varies ith distance. An idealized oscillation xperiment would consist of a neutrino ource that produces only one type of eutrino, here muon antineutrinos (blue) om pion decay. Enroute to the detector, he muon antineutrino transforms into an ectron antineutrino (red), which can teract with a proton in the detector to reate a positron and a neutron. These wo particles are taken to be the signaure of the electron antineutrino.

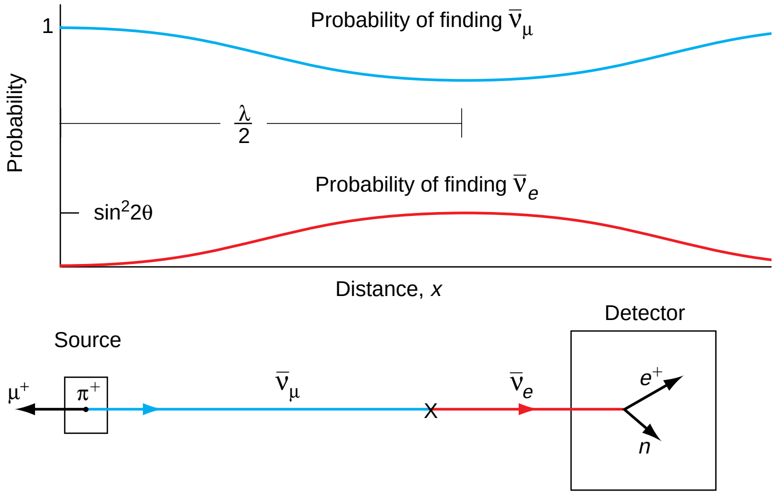

ecades-old question of whether “massess” neutrinos are really massless.

Assume for the moment that neuinos have mass, and furthermore, each eutrino would have a specific mass. he electron neutrino would weigh ome amount m_1, the muon neutrino would weigh m_2, and the tau neutrino would weigh m_3—simple, elegant, and asy to explain to students. But in the aradoxical realm of quantum mechancs, a neutrino can be several things at nce. In all likelihood, each neutrino as a split personality and possesses, in sense, three masses. Put an electron eutrino on a scale, and it might read m_1, m_2, or m_3.

The three masses are like three hostly neutrinos (called “mass states”) hat inhabit the electron, muon, and tau eutrinos. Mathematically, they can be een as three ingredients which, ombined in various proportions, form he electron, muon, and tau neutrinos.

To see how oscillations occur, magine a muon neutrino produced in he decay of a muon. The muon neuino can be viewed as a mixture of he three ghostly neutrinos, each with different mass. The particular mixure (the recipe) that defines the muon eutrino is dictated by three numbers alled “mixing angles.” If other subtomic particles are any guide (that is, he quarks), one “ghost” will dominate ach neutrino type so that, for instance, a muon neutrino might be 90 percent m_1, 9 percent m_2, and 1 percent m_3 (see the article “The Oscillating Neutrino” on page 28).

The muon neutrino may oscillate into another neutrino type, because the mix of its ingredients can change as it travels. In the quantum picture, the neutrino is described by a “wave function” that can be seen as the sum of three separate waves, one for each of the mass states. As the neutrino travels, each mass-state wave “vibrates” at a frequency that depends on the neutrino’s mass, so that the neutrino is like a three-note chord with each note beating against the others. The relative amounts of each mass state change with the rise and fall of one wave against another until the neutrino arrives at some detector designed to measure its type. Depending on how the three waves are synchronized at the detection point, the particle will have some probability of appearing as an electron, muon, or tau neutrino.

The mathematics of oscillations can fit on a single page (see the box “Derivation of Neutrino Oscillations” on page 52), but the trickier problem of why there is a primordial mixup of masses remains unsolved. The related problem of why particles have the masses they do also presents a conundrum that only a few broad-minded theorists have dared to tackle. Were physics a religion, mass would have its own creation myth.

Blueprints

All neutrino oscillation experiments follow the same conceptual blueprints (see Figure 2). At the level of a sketch one might make on a napkin, there are only two components: a "source", which like a pitching machine, hurls out neutrinos of a known type, and a "detector", which like a catcher’s mitt, absorbs and counts the neutrinos. The game is simple. If neutrinos have mass, they can oscillate as they travel, changing their identity back and forth as they go. Any difference between what the source throws out and what the detector observes can be chalked up to oscillations.

But most neutrinos fly straight through the detector, so oscillation experiments, like baseball, are a kind of long-attention span sport, with extended periods of thumb twiddling between bits of action. The slow pace is a reflection of the fact that the neutrinos must interact through the weak force. Because the weak force is so feeble, fewer than one in one trillion neutrinos will leave a mark in the tank. Even with LSND's high-intensity source, it is often an hour between neutrino catches.

Still, the weak force runs the show. Like any other force, it can push and pull, but it can also create and absorb particles. In the source, the weak force gives birth to the neutrinos (here muon antineutrinos produced through the decay of a muon), and at the other end, it is again responsible for their demise.

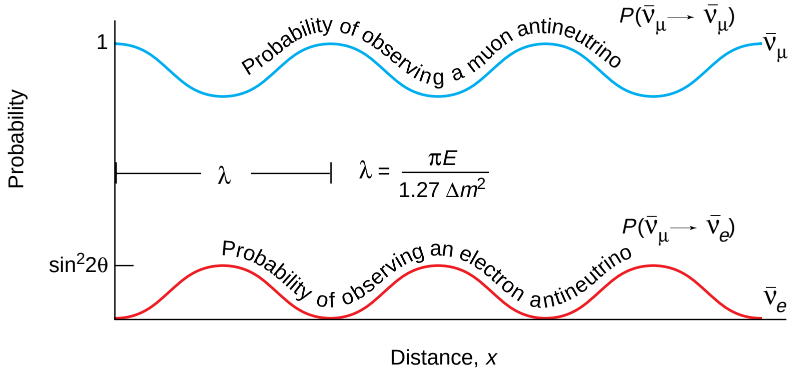

Figure 3. The Rise and Fall of the Electron Antineutrino

The probability of finding a muon antineutrino (blue curve) goes down as the probability of finding an electron antineutrino (red curve) goes up. Both probabilities oscillate with the same wavelength λ, and the sum of the two is always equal to one.

The death is swift, if rare. The neutrino disappears, replaced by new particles that generate telltale patterns of light in the oil. The "disappearance" is really just a change in identity. The weak force tugs on the neutrino and the protons in the tank, and can occasionally "transfer" the electric charge of one proton to the neutrino. In the process, the proton becomes an (electrically uncharged) neutron, and the electron antineutrino becomes a positively charged electron (called a positron).[2] Both the neutron and the positron generate flashes of light in the oil which draw the attention of a roomful of electronics.

Thus the weak force anchors the experiment at both ends, creating and then providing the means of detecting the neutrinos. But when the neutrino flies between the source and the detector, the weak force is, in a sense, left behind. In the interim, the ghostly hand of quantum mechanics takes over and starts the neutrino oscillating.

Feeling Around in the Dark

At the heart of every oscillation experiment lies a single all-important equation that gives the probability that a neutrino beginning its journey as one type will be observed as another type. For simplicity, assume that the tau neutrino is out of the picture and that oscillations take place only between a muon antineutrino and an electron antineutrino. (LSND searches for the latter particle.) In that case, the probability for a muon antineutrino to transform into an electron antineutrino is

$$P\left(\bar{\nu}_\mu \rightarrow \bar{\nu}_e\right) = \sin^2 2\theta \, \sin^2\left(\frac{1.27\Delta m^2 x}{E}\right) \; .$$

Here θ is the mixing angle (the “recipe” for how to combine two mass states to make a muon antineutrino), Δm^2 is the difference between the squares of the masses of the two mass states (that is, $\Delta m^2 = \left| m_2^2 - m_2^1 \right|$) and is in units of electron volts squared (eV^2), E is the neutrino energy in million electron volts (MeV), and x is the distance between creation and detection in meters. Despite its complicated appearance, all the equation really means is that the probability of observing an electron antineutrino goes up and down like a sine wave as the distance to the source or the energy of the neutrinos is changed. (See Figure 3.) Indeed, since the mixing angle is a constant of the world (albeit unknown), and since the energy of neutrinos is typically fixed for a particular experiment, the oscillation probability is often written as just

$$P\left(\bar{\nu}_\mu \rightarrow \bar{\nu}_e\right) = \mathrm{A}\,\sin^2\left(\frac{\pi x}{\lambda}\right) \; ,$$

where the oscillation “wavelength” is given by $\lambda = (\pi E)/(1.27\Delta m^2)$ and the size of the oscillation is $\mathrm{A} = \sin^2 2\theta$.

Given the difficulty of detecting neutrinos, every experimenter would love to place the detector where the probability of seeing an electron antineutrino peaks. The first place would be one-half of a wavelength from the source (see Figure 2). Unfortunately no one knows exactly how long a wavelength is.

Ideally, theoretical calculations could predict the neutrino masses that determine the wavelength, but so far such predictions lie beyond the scope of even the most far-reaching theories. From a theoretical perspective, searching for oscillations is like digging for buried treasure without a map. The oscillation wavelength could be 2 millimeters or 2 light-years.

The quest is made still more difficult because an experiment essentially only measures one number (here the number of electron antineutrinos) but seeks information about two quantities (Δm^2 and $\sin^2 2\theta$). If, after a year, an experiment saw nothing, it could mean that the mixing angle is very small (and hence A, which functions like a “volume” control, squelches the probability of observing an electron antineutrino) or it could mean that the detector happened to sit at a distance where the oscillation probability was low (a distance very much smaller than the wavelength for example). Then, the best one can do is to rule out the values of $\sin^2 2\theta$ or choices of Δm^2 that would have given an observable number of electron antineutrinos. Figure 4 shows the values that had been searched and ruled out by experiments before LSND began taking data in 1993.

Similarly, if after a year, the catcher gazes into the glove and miraculously sees a few electron antineutrinos, it is

---

[2] The neutrino names, in fact, derive from this pairing. The weak force transforms an electron neutrino into an electron, a muon neutrino into a muon, and a tau neutrino into a tau.

mpossible to sort out how much of the scillation was due to A and how much o Δm^2. In fact, it wouldn't even be lear that oscillations had produced the lectron antineutrinos. The appearance f electron antineutrinos could also be nterpreted, perhaps more interestingly, s evidence for a new, bizarre decay of he muon, forbidden by the accepted aws of physics.[3]

This was the strange limbo that eset LSND following the 1995 nnouncement. Although LSND results howed the appearance of electron ntineutrinos in a flood of muon antieutrinos, the measurement had been made at only one distance. Without eeing the number of detected electron ntineutrinos rise and fall periodically s a function of the distance x (or as function of the energy E), few were willing to write massive neutrinos into he textbooks.

The LSND collaborators themselves greed that one point did not make an scillation. To address this , LSND had een designed to see oscillations in a econd way, by looking for the transformation of muon neutrinos into electron eutrinos, the "matter" counterpart of its rimary antimatter analysis. This second nalysis would later provide an invaluble cross-check on LSND results, but n 1995 it was not yet complete.

Like all neutrino oscillation experiments, LSND was a shot in the dark. If he experiment had indeed observed scillations, it would have been a lucky appenstance that the source-to-detector istance was right.

Or of course, it could have been a mistake. The best-designed detectors re imperfect and can be duped by eleconic noise, by other particles from the eam, or by the unrelenting rain of cosmic rays. Understanding these "backrounds" formed the linchpin for the SND experiment and was the focus of the questions that filled the air whenever LSND researchers presented their results. While a seminar audience could not judge in an hour what had taken years to put together, many feared that the experimenters, too, had somehow been duped.

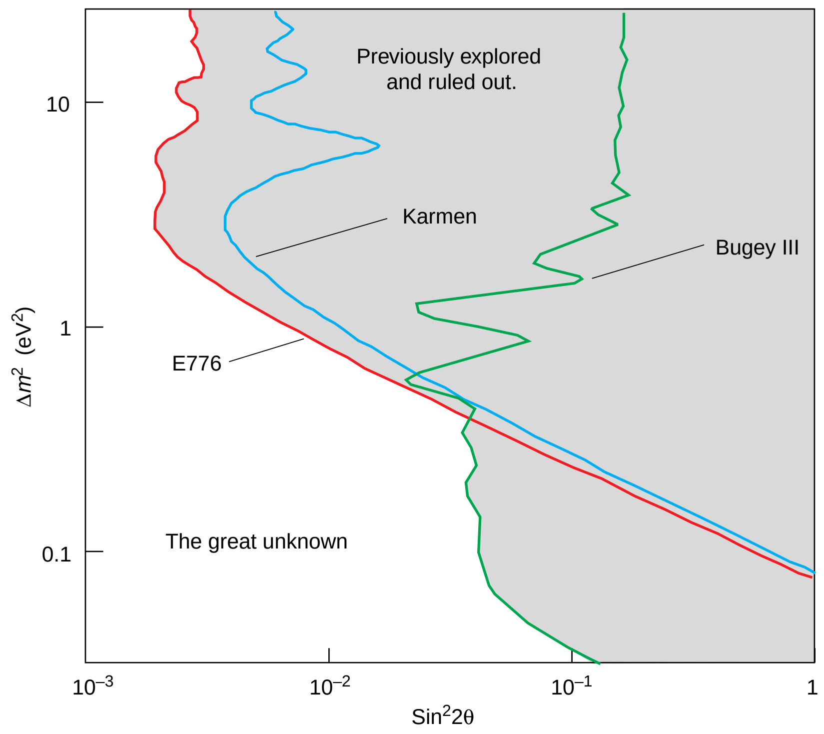

Figure 4. Values of Δm^2 and Sin$^2 2\theta$ Explored Before LSND

Several experiments had searched for neutrino oscillations and reported negative results. The values of Δm^2 and sin$^2 2\theta$ probed and ruled out are shown in grey: E776 (red) was an accelerator-based experiment at Brookhaven National Laboratory; Karmen (blue), an accelerator-based experiment at Rutherford Appleton Laboratory in England; and Bugey III (green), a French reactor-based experiment. When LSND began taking data in 1993, it was hoped that the experiment would have enough sensitivity to probe "the great unknown" region that remained.

Celestial vs Terrestrial Neutrinos

LSND is the fifth in a series of neutrino experiments at Los Alamos, and its design draws heavily on the experience of its predecessors. Most importantly, it has inherited a decade-old high-intensity neutrino source based on the Los Alamos Neutron Science Center, or LANSCE, accelerator. An old warhorse, LANSCE is still the highest-intensity proton accelerator in the world for its energy. Physicists are as familiar with its behavior as a soloist is with a well-rehearsed piece of music.

Earthbound neutrino sources like LANSCE outshine their celestial counterparts in many respects. From their far-flung birthplaces, celestial neutrinos must travel a long way to visit terrestrial detectors—thousands of light-years from supernovae, a hundred million kilometers from the Sun. As messengers, they are invaluable, bringing news of distant events such as the death of stars and reactions in the heart of the Sun; but as sources for oscillation experiments, their remote and uncertain origins make them less than ideal. There is no user's manual for the Sun that states exactly what kind of neutrinos are being produced, no dials to turn to adjust its output, and no switch to turn it off. Particle accelerators provide a way to construct an intense neutrino source closer to home, a kind of sun-on-earth that can be tuned and tested.

Figure 5. LANSCE Production of Decay-at-Rest (DAR) Neutrinos and Their Energy Spectra

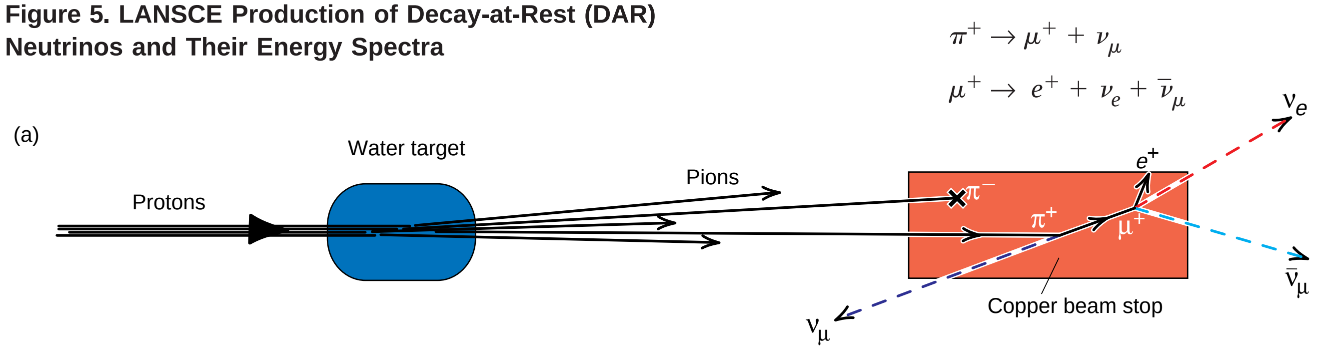

(a) LANSCE accelerates protons and shoots them into a water target, where they produce pions. Most of the pions travel a half meter in air before striking a copper beam stop and coming to rest. Nearly all of the negative pions (π^-) are absorbed by copper nuclei before they decay, but each positive pion (π^+) decays at rest to a muon (μ^+) and a muon neutrino (ν_μ). The muons also come to rest in the beam stop and decay to a muon antineutrino ($\bar{\nu}_\mu$), a positron (e^+), and an electron antineutrino ($\bar{\nu}_e$). The neutrinos fly off in all directions as from a miniature sun. (b) Energy spectra for the muon neutrinos (purple), muon antineutrinos (blue), and electron neutrinos (red) created in the beam stop. Although electron antineutrinos from negative pion decay have energies similar to those of electron neutrinos, these antineutrinos are produced so infrequently that their spectrum would appear as only a small dot on the plot.

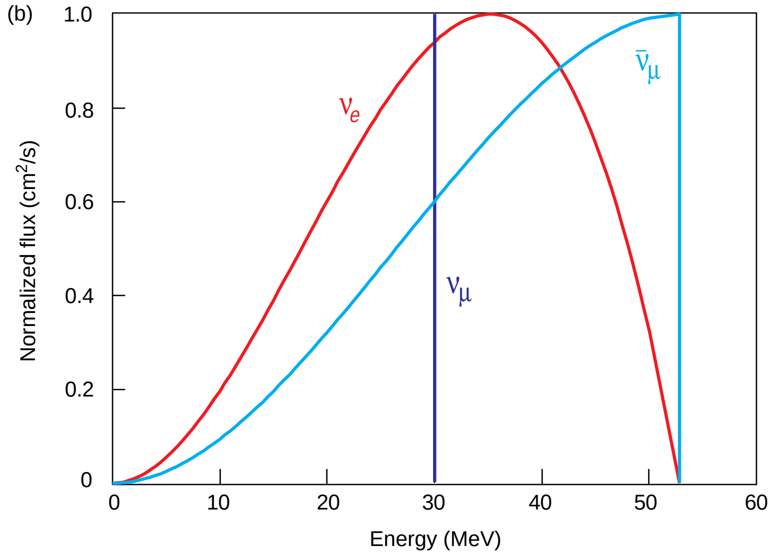

But even a custom-made neutrino factory will have flaws. The ideal source would produce only one neutrino type so that oscillations could be easily identified. Neglecting the tau neutrino and its antiparticle still leaves four other types (electron neutrinos, electron antineutrinos, muon neutrinos, and muon antineutrinos) and, unfortunately, LANSCE makes them all. The goal, then, is to get rid of one (electron antineutrinos in this case) to clear a channel so that oscillations can be detected. The trick is preventive medicine—to stop electron antineutrinos before they are made. Miraculously, all this takes is a block of copper.

Neutrino production begins with a burst of protons from the kilometer-long LANSCE accelerator. As seen in Figure 5, the protons strike a water target, producing pions. Virtually all pions will produce both muon neutrinos and antineutrinos when they decay. The positively charged pions will also make electron neutrinos, while the negatively charged pions will make electron antineutrinos.

The pions fly through the air and plow into a copper block, called the beam stop, where they slow and come to rest. Fortunately, the negative pions, because of their charge, are absorbed by the positively charged copper nuclei before they can decay, so that only a very few electron antineutrinos are produced. The positive pions, however, hang around until they decay into muons, which in turn decay to produce abundant neutrinos. Thus the copper block filters out the negative pions and keeps the source relatively free of electron antineutrinos. In the end, the other neutrino types produced outnumber electron antineutrinos by a factor of roughly 10,000 to 1.

It should be noted that a small number of pions decay in flight before reaching the copper block. These pions produce higher-energy neutrinos than the pions that decay at rest in the copper block. The "decay-in-flight" (DIF) neutrinos thus become a second, separate source riding piggyback on the first and can be used to cross-check the results from the decay-at-rest (DAR) neutrinos. In fact, the DIF neutrinos are the subject of the second analysis alluded to earlier.

LANSCE loses out to neutrino sources like the Sun, however, in one important category. The problem isn't physical, it's fiscal. The sun shines for

---

f some positively charged muons made by the celerator decayed into electron antineutrinos, it ould (falsely) appear as if oscillations had ocrred. See the article "The Nature of Neutrinos Muon Decay and Physics Beyond the Standard odel" on page 128.

ee, but LANSCE eats up about $20 million in electric bills every four months. LANSCE serves several xperiments in addition to LSND, but n 1993, the Department of Energy DOE) threatened to shut down the ower-hungry accelerator in order to eed other research efforts.

At the time, LSND researchers were just preparing to take their first ata. If LANSCE died, LSND would ollow, buried in its underground rypt. At the last minute LANSCE was spared, and the researchers, crambling to finish LSND construc-on, took a total of six weeks of data uring September and October. Even n this small amount of data, there were signs that a few extra electron ntineutrinos had reared their heads. he researchers found eight electron ntineutrino-like events[4] when they ad expected only 0.9 ± 0.2.

The results, which were published in conference proceedings the next year Louis et al. 1994), drew considerable nterest, but more data was needed to onfirm the excess of events. If people new one thing about neutrino physics, ecalls Richard Imlay, “it was that it was hard. Many people wanted to wait nd see if the results held up.” Despite he cloudy funding picture, LSND was leared to take another three and a half months of data beginning in August f 1994.

At the time, *Scientific American* evoted a page to the LSND story. Missing Matter Found?” it asked. Not et. “We feel we have a high burden f proof,” it quoted Hywel White, one f the Los Alamos collaborators, as aying, because if neutrinos have mass, it’s very important.” After the upcom-ng LSND run, the article quipped, “the eam at Los Alamos should be able to erify—or otherwise—their nonclaim” Mukerjee 1994).

Detecting the Electron Antineutrino

Neutrinos produced in the LANSCE beam stop travel at velocities near the speed of light through steel shielding, earth, and concrete, finally reaching the detector 30 meters downstream. Over this short distance, the muon antineutrinos may change their identity and oscillate into electron antineutrinos. If the weak force, then, should chance to connect an antineutrino to a proton in the detector’s oil, a single positron and neutron will emerge, each producing light as it passes through the oil.

The mineral oil that fills the tank is composed of carbon and hydrogen (chains of 30 or so CH_2 molecules). The electron antineutrinos are detected when they react with the hydrogen atoms (which are essentially free protons) through a process called “inverse beta decay,”

$$\bar{\nu}_e + p \rightarrow n + e^+ \ ,$$

so named because it represents the reverse of the normal “beta decay” process common in radioactive nuclei.

The creation of the neutron and positron heralds an electron antineutrino event. Both particles lead to the production of light in the detector, and it is by observing that light that LSND knows an event has taken place. By human standards the light is invisibly faint, but to the 1,220 phototubes that can see a single photon and measure its arrival time to a nanosecond, the light shines like a miniature pyrotechnic display.

The positron generates light through two mechanisms. When created, the positron has a velocity greater than the speed of light in the oil and produces an electromagnetic shock wave analogous to the wake of a speedboat or to the sonic shock wave of a Concorde breaking the sound barrier. The light, called Cerenkov radiation, forms a cone that expands along the positron’s trajectory like the headlight of a tiny car. The cone has a 47-degree opening angle and forms a shell rather than a solid. Projected onto a flat surface, the cone leaves a telltale ring.

The positron also produces light with the help of a small amount of scintillator that is added to the oil. As the positron travels, it loses energy by inducing atomic excitations in the oil; by a secondary process, the scintillator also gets excited. When the scintillator de-excites, it produces blue light that is detected by the phototubes. The excitation process takes a little time and delays the scintillation light by about 15 nanoseconds relative to the Cerenkov light. Also, because the positron typically travels only about 25 centimeters before it wears itself out, the scintillation light appears as an almost spherical cloud. The ratio of scintillator to mineral oil is selected to give roughly four scintillation photons to every Cerenkov photon. All told, a typical positron produces enough photons to trigger 450 phototubes.

Unlike the positron, which leaves a bright trail of light, the neutron goes quietly, wandering randomly away from the neutrino collision until it comes close enough to a proton to be captured through the reaction

$$n + p \rightarrow \mathrm{D} + \gamma \ ,$$

producing a deuteron (D) and a 2.2-MeV gamma ray (that is, a 2.2-MeV photon). On average, the capture takes 186 microseconds, so the 2.2-MeV gamma ray emerges somewhat after the light from the positron. The gamma ray also generates some scintillation light, which fires between 20 and 50 phototubes.

The light generates small electric pulses in the phototubes that then travel along some of the one thousand cables connecting the detector tank to the crates of electronics near the tunnel entrance. Like a brain processing a visual image, the electronics sift through the light signals, trying to assemble a picture of an electron antineutrino.

The purpose of the data collection electronics is to distinguish electron antineutrinos from the endless stream of cosmic rays that penetrate the detector’s shielding and enter the tank. The probability that an electron antineutrino (or any other neutrino) will interact with matter is unimaginably small. Even with 167 metric tons of mineral oil, 99.999999999 percent of the neutrinos will pass through the tank unhindered and unnoticed. In a day, only 25 neutrinos will leave their mark. In the same time period, 300 million cosmic rays will also pass through. (For more details on the data collection electronics, see “From Tank to Tape—The LSND Data Acquisition System” on page 112.)

The scenario is summarized in Figure 6. The electron antineutrino penetrates the tank and strikes a proton, giving rise to a neutron and a positron. The positron generates a cloud of blue scintillation light and the characteristic cone from Cerenkov radiation. Then all is quiet for roughly 186 microseconds, after which a tiny 2.2-MeV gamma ray signals the presence of a neutron. Taken together, these signals constitute the signature of an electron antineutrino.

The central difficulty in analyzing the data is how to pick out the real electron antineutrinos. After all,

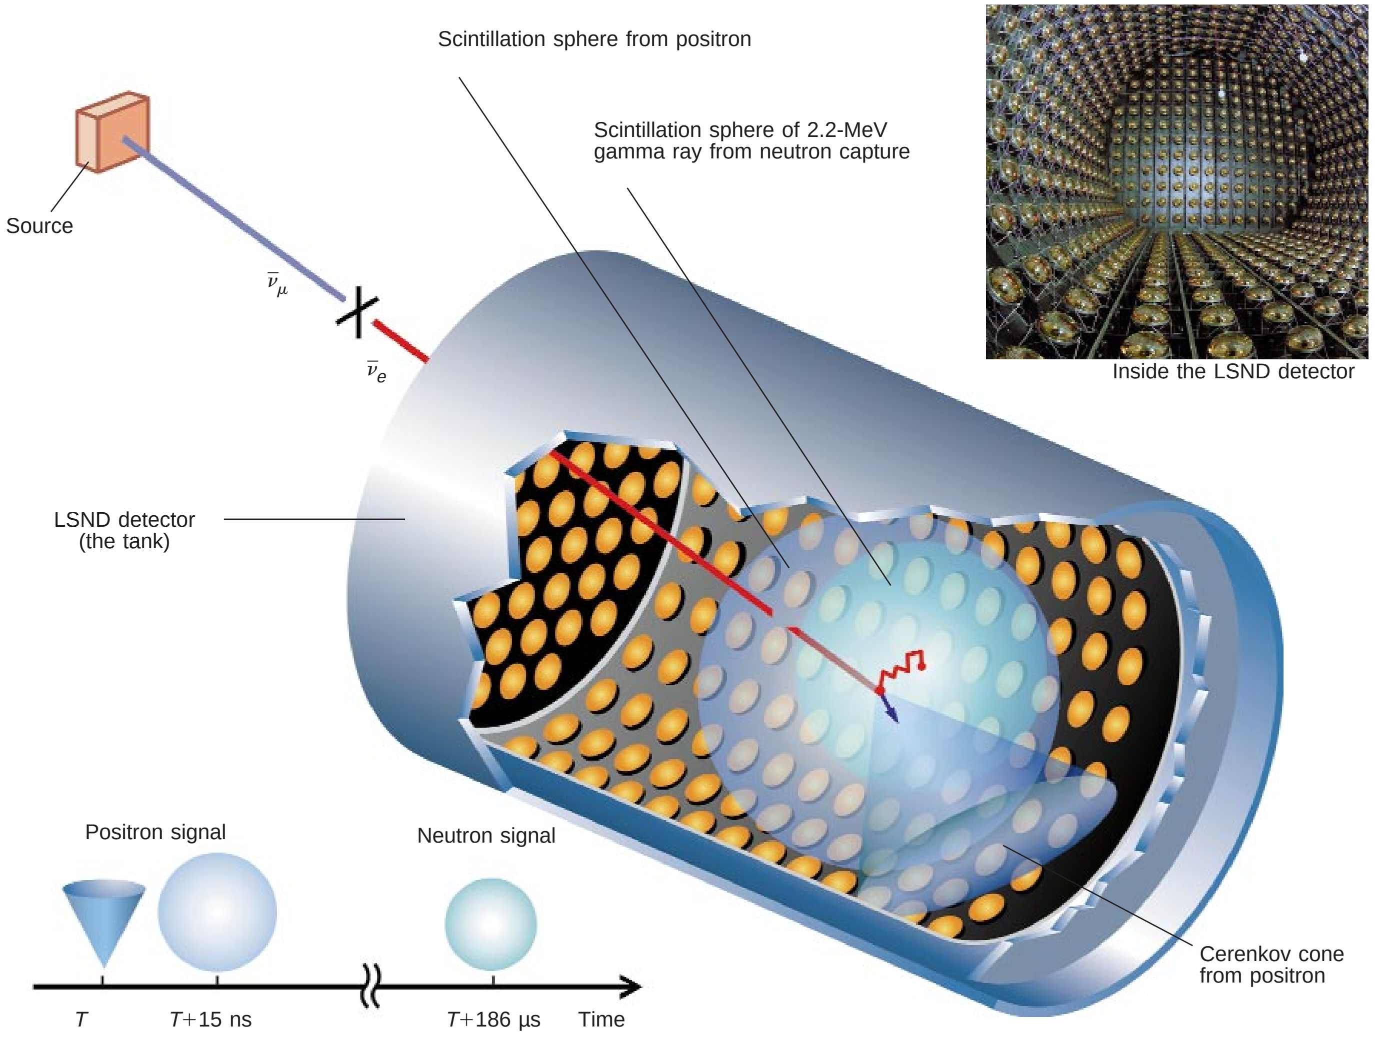

Figure 6. Signature of a Neutrino Event in LSND

In an oscillation, a muon antineutrino (blue) produced in the beam stop oscillates enroute to the detector and appears as an electron antineutrino (red). The neutrino strikes a free proton in the oil, creating a positron and a neutron. The positron travels faster than the speed of light in the oil and so produces a Cerenkov cone. As it loses energy through collisions with atoms in the oil, the positron also produces a sphere of scintillation light. The neutron survives about 186 microseconds and wanders 100 centimeters before it is absorbed by a nucleus, emitting a 2.2-MeV gamma ray that also produces a sphere of scintillation light. This succession of events—the apex of a Cerenkov cone centered on a sphere of scintillation light followed by emission of a 2.2-MeV gamma ray—is the signature of an electron antineutrino.

This was a faint heartbeat. If correct, it implied at the oscillations were small indeed, since SND would have seen 5000 events if all the uon antineutrinos produced had fully oscillated to electron antineutrinos.

igure 7. Identifying a Positron

**n energetic particle moving through the
nk creates both Cerenkov and scintilla-
on light, but the ratio ρ of Cerenkov to
cintillation light varies depending on the
article's type and energy. Positrons (or
ectrons) tend to have a very consistent
atio of about 0.5, which is evident in this
figure as a large spike. Neutrons and pro-
ons tend to produce very little Cerenkov
ght, and thus they have small values of
(the rounded hump around 0.15). This
ronounced difference in ρ is used to
elp distinguish positrons from other
articles that leave trails in the tank.**

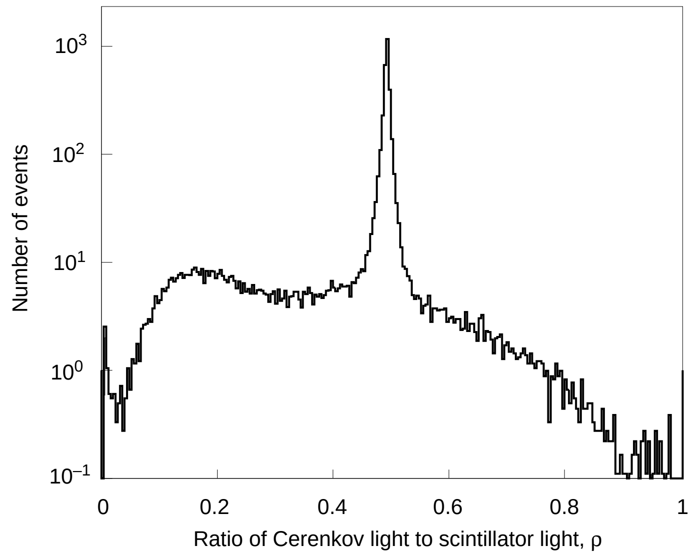

igure 8. Accidental Photons

**hotons due to background processes—
ccidental photons—should appear at
andom positions in the tank. However,
hese plots reveal that most accidental
hotons appeared at the bottom of the
nk. Each point represents the position
f a photon projected onto a two-
mensional plane. The heavy concen-
ation along the bottom was likely due to
piece of steel shielding that has some
inute amount of radioactivity or to a
mall hole in the shielding through which
he electronic cables pass.**

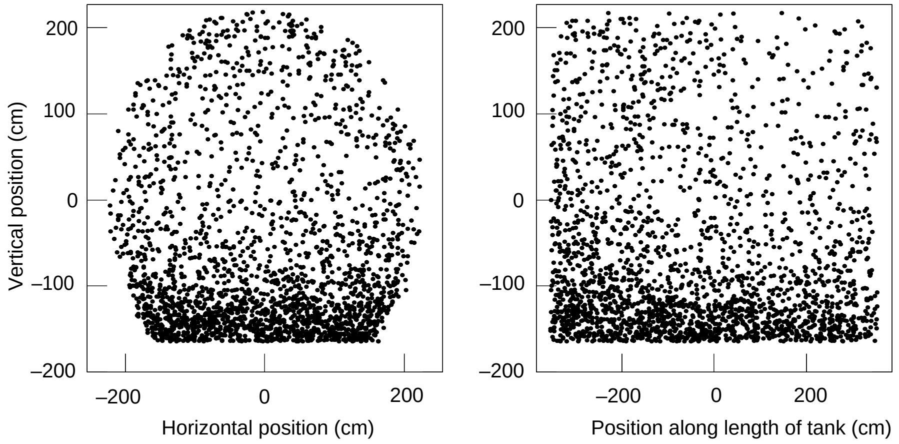

eutrinos aren't the only particles
haking up the atoms in the oil. Even
with a carefully designed detector,
vents are not always what they seem.
ull the plug on the accelerator, and the
etector will still record events that
ook indistinguishable from electron
ntineutrinos. As in any physics experiment, the researchers began by playing
evil's advocate and drawing up an
xhaustive list of possible impostors,
r backgrounds. Topping the list were
wo potential showstoppers. First, what
the positron wasn't really a positron
ut was a cosmic ray instead? And
econd, what if the putative photon
rom neutron capture really came from
omewhere else? To separate the real
lectron antineutrinos from the fakes,
SND researchers had to develop a
ariety of tools.

Positrons, Photons, and Impostors

The first problem was how to distinguish a positron from cosmic-ray muons and their by-products. Because most muons betray their identity by leaving a signal in the cosmic-ray veto shield (the outer shell that encloses the main detector tank), they are easy to track. If a muon is energetic enough to make it through the veto shield and into the tank, it tends to make a flashy entrance, producing so much ionization that all 1,220 phototubes light up. Looking more like a than a positron, these muons are easy to identify.

Occasionally, however, a cosmic ray can pull off a more convincing impersonation, passing near the tank and knocking a neutron free in the shielding. The neutron can pass undetected through the veto shield and into the tank. There it may strike and propel a proton through the oil. If the proton is mistaken for a positron and a neutron is captured nearby, the combination would be a dead ringer for an electron antineutrino. (See "Other Things in the Tank: Backgrounds" on page 104.)

But forging a positron signature is not so easy. A typical 45-MeV positron lights up 450 phototubes, each of which measures the time and number of the photons it receives. This is an enormous amount of information from which the positron's position, trajectory, and energy can be culled. The positron's position can be determined to about 25 centimeters by finding the point in the tank that is equidistant in

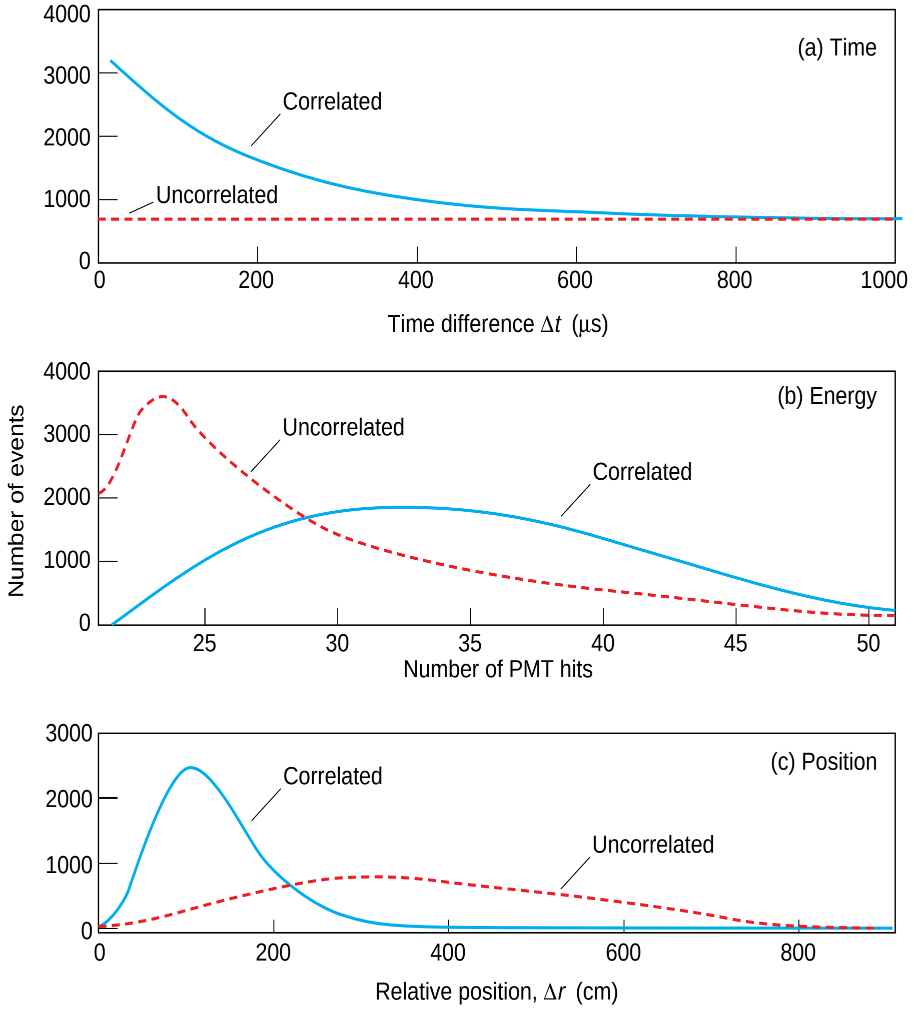

Figure 9. Identifying the Photon from Neutron Capture

Although the photon from neutron capture cannot be distinguished from a random photon, statistically, it has a distinct signature in time, energy, and position. In these graphs, blue lines show the distributions of photons that are correlated with neutrons, while red lines show the distributions for photons that are uncorrelated. (Cosmic-ray neutrons leave a telltale trace of light when they stop in the tank. They are used to determine the correlation functions.)

(a) On average, a neutron is captured 186 microseconds after it is created. Uncorrelated photons show no time correlation to the initiating event.

(b) Background photons have less energy and thus produce less light in the tank. On average, the 2.2-MeV photon lights up about 35 photomultiplier tubes (PMTs).

(c) The neutron wanders about 100 centimeters before it is captured. Photons arising from random processes can occur anywhere in the tank. These distributions are used in the likelihood function *R* to assess whether a photon is associated with a neutron created from a neutrino event.

time from all of the phototubes that detected photons. The trajectory can be determined by finding the ring of Cerenkov light that is superimposed on the uniform scintillation light. The trajectory runs through the center of the ring and can be determined to within 12 degrees. The positron's energy is simply proportional to the total charge from all of the phototubes that were hit and can be determined to within 6 percent.

Protons produced by cosmic rays, by contrast, will rarely be traveling fast enough to emit a Cerenkov cone (see Figure 7). Requiring a well-defined cone and an accompanying sphere of scintillation light removes 99.9 percent of all cosmic-ray events, while 80 percent of the real positrons pass the selection criteria.

The second problem was to sift out the photons (the 2.2-MeV gamma rays) that truly came from neutron capture. The steel of the cosmic-ray veto shield around the main detector tank absorbed most photons coming from the outside, but there was a chink in its armor. The shield covered the tank like an arch, stopping at the tunnel floor and leaving the underbelly of the detector exposed to the concrete. To extend the shielding, the LSND collaborators had laid down a 15-centimeter-thick floor of steel planks underneath the detector.

Even with this flooring, however, accidental photons turned out to be a bigger problem than expected. When the LSND physicists looked for photons in their 1993 data, they found many of them clustered suspiciously at the bottom of the tank toward one end (see Figure 8). Whether due to a bit of slightly radioactive steel or to the small hole in the shield where the electronic cables exited the tank, the extra photons threatened to swamp the real electron antineutrino signature.

To filter out these accidental photons, Richard Imlay and his group from Louisiana State University came forward with a new technique based on a quantity called *R*, which determined the likelihood that the "positron" and "neutron" signals were correlated. If the two were the true signature of an electron antineutrino, the photon (from neutron capture) should have 2.2 MeV of energy and appear slightly after the positron but at roughly the same location in the detector tank (see Figure 9).

This correlation was in contrast to what you expected for an unrelated or

Other Things in the Tank: Backgrounds

Despite the shielding that surrounds it, the detector tank remains a bubbling cauldron of activity. Every second, thousands of accidental photons and cosmic rays stream through the oil, leaving trails of scintillation light in their wake. Occasionally, the endless background signals combine to look just like the signature of an oscillation event—a positron followed some 186 microseconds later by a photon. The LSND experimenters have taken great pains to understand how background signals could mimic real electron antineutrino events so they can estimate the number of "false positives" that would make it through the data analysis. Only by finding an excess of oscillation over background events can they state with confidence that they have observed neutrino oscillations.

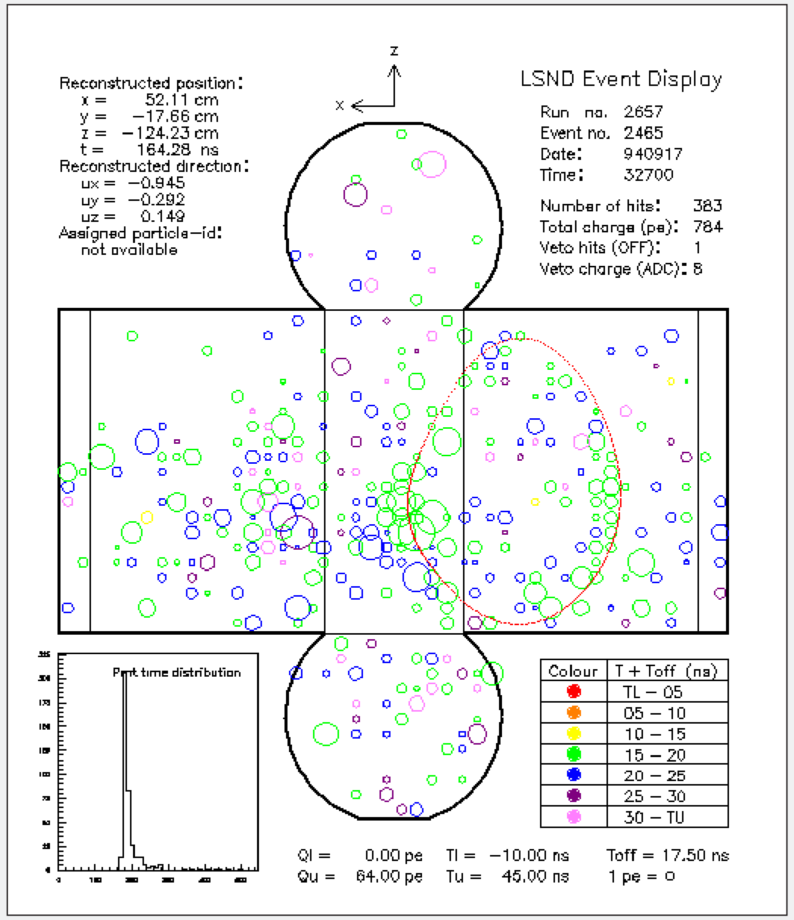

econstructing an LSND Event. In this flattened ew of the tank's inner surfaces, the colored rcles identify 383 photomultiplier tubes that ave been hit by photons. A circle's diameter orresponds to the number of photons that ruck the tube; its color indicates when those rikes occurred (see color key at the lower ght). This data is used to reconstruct the erenkov ring (which looks oval-shaped in this ojection) and the scintillation sphere (not hown). The position and trajectory of the parti- e associated with the photons (see upper left) e then calculated. In this case, the particle as an energetic electron produced from the ecay of a muon.

Non-Beam-Related Backgrounds: Cosmic Rays and Accidental Photons. Cosmic rays are the largest source of background. If every muon antineutrino were to oscillate to an electron antineutrino, cosmic rays interacting in the tank would still outnumber electron antineutrinos by a factor of more than 100,000. Most cosmic rays that reach the earth are muons, produced from the decay of pions created when high-energy protons strike nuclei in the upper atmosphere.

Surrounding most of the detector tank is a cosmic-ray veto shield that is LSND's main line of defense. The archlike shield has double walls. The inner wall is a 15-centimeter-thick layer of lead shot that absorbs accidental photons and a significant number of cosmic rays. But this dense layer is not enough. Every second, approximately 4,000 cosmic-ray muons pass straight through the lead and enter the tank. Thus, the outer wall of the veto shield is studded with 292 photomultiplier tubes looking inward, and the space between the walls is filled with mineral oil and liquid scintillator. On their way into the tank, cosmic rays leave a trail of scintillation light in the veto shield. Removing events in which more than a few of the veto-shield phototubes fire eliminates 99.999 percent of all cosmic-ray-induced events.

Being very energetic, most cosmic-ray muons that pass through the shield also pass straight through the tank. About 10 percent, however, don't make it. They stop in the tank and decay, with the positive muons producing positrons and the negative muons producing electrons. (The two particles are treated the same in the data analysis, since they are indistinguishable as seen through the eyes of the detector.) Although the muon is detected by the veto shield, the positron that is born typically 2 microseconds later is not. It appears to come from nowhere, exactly like the positron created from a electron antineutrino event. Thus, in this background process, the muon acts as a kind of Trojan horse for the positron, in effect sneaking it past the defenses of the veto shield.

Most of these positrons, however, are ignored by the data acquisition system (see the box "From Tank to Tape—The LSND Data Acquisition System" on page 112). The system requires that there be no activity in the detector or veto shield for a period equivalent to about 7 muon lifetimes before a positron appears. If any activity occurs during this "all quiet" time, the positron is rejected as signaling a potential neutrino event.

Although not all cosmic-ray background events can be detected, their number is relatively easy to estimate: it is simply measured when the accelerator is off. The accelerator does not produce protons continuously but has a regular heartbeat, pumping out 600-microsecond bursts of protons 120 times a second. The beam is on only about 7 percent of the time, and the non-beam-related backgrounds can be studied during the relatively long rest periods that compose the remaining 93 percent. In fact, the data acquisition system does not initially distinguish between beam-on and beam-off events. It only uses that information later, when it assesses which events represent true neutrino interactions. In their 1996 paper, the LSND collaborators estimated that they expected 2.5 ± 0.4 events from beam-off sources to look like electron antineutrinos (Athanassopoulos et al. 1996).

Beam-Related Backgrounds: Neutrinos. The beam contains equal numbers of muon neutrinos, muon antineutrinos, and electron neutrinos. If, for instance, 1 percent of the muon antineutrinos oscillate to electron antineutrinos, they will be outnumbered 300 to 1 by the other neutrino types. Although the other neutrinos cannot easily imitate electron antineutrinos, they can still lead to background events that must be estimated.

A muon antineutrino interacting with a proton produces a positive muon (instead of a positive electron) and a neutron. This process is a potentially dangerous background; the muon can decay in the tank and produce a positron, which, when combined with the neutron, would convincingly mimic the oscillation signature of an electron antineutrino. Fortunately, the newly created muon produces Cerenkov and scintillation light in the tank, so the muon is observed by the data acquisition system. The "all quiet" requirement removes these muon decay events. In addition, because the muon weighs a hefty 105 MeV, only the relatively few decay-in-flight muon antineutrinos have enough energy to produce muons. (All decay-at-rest neutrinos have energies below 55 MeV.) In the end, muon antineutrinos constitute a small background that can be reliably calculated.

Electron neutrinos are a background because they can change a carbon atom in the detector's mineral oil into a nitrogen atom. An electron is also produced in the process:

$$\nu_e + \mathrm{C} \rightarrow e^- + \mathrm{N} \quad .$$

However, since no neutron is produced in this reaction, the process is only a problem in the unlikely event that it coincides with an accidental photon.

The largest source of beam-related backgrounds is the electron antineutrinos present in the beam itself. They arise from negative muons that decay in the beam stop before being absorbed. These decay-product neutrinos are essentially indistinguishable from electron antineutrinos produced from oscillations. Fortunately, the decay-product flux is very well known, and this background can also be calculated with confidence. In their 1996 paper, the LSND collaborators estimated the background from electron antineutrinos in the beam to be 1.1 ± 0.2 events and the total background from beam-related events to be 2.1 ± 0.4 events. Thus, all told, backgrounds accounted for a grand total of 4.6 ± 0.6 events (Athanassopoulos et al. 1996).

"accidental" photon, which could come from the small amounts of radioactive elements (thorium, for instance) present in the concrete and earth surrounding the detector or in components used to build the phototubes. Accidental photons tended to have energies below 2.2 MeV and, by definition, appeared uncorrelated in position and time with the positron. The data plots shown in Figure 9 reveal some of the differences between accidental photons and those associated with neutron capture.

The *R* correlation not only picked out events in which the positron and photon were correlated, but it also rejected events in which the photon looked accidental. Mathematically speaking, *R* was a "relative likelihood," but it worked like a magic box. Feed in the number of phototubes fired by the photon and the distance and time between the photon and the positron signals, and out popped a number. Real electron antineutrino events tended to have high values of *R*, while accidental photon events piled up at low *R*. Unfortunately, although *R* could remove 99.4 percent of the accidentals, it did so at a high price: *R* also removed 77 percent of the real electron antineutrinos.

The final source of background events was neutrinos produced in the beam stop. The electron antineutrinos that contaminated the source contributed a few background events, as did the other neutrino types that could, on occasion, leave what appeared to be the signature of an electron antineutrino.

Still, when the LSND collaborators ran through the data from both 1993 and 1994, they found further evidence for electron antineutrinos. Their improved analysis let fewer impostors slip through the cracks than before. This time they had nine events that looked like electron antineutrinos. The expected number of background events came to only 2.1 ± 0.3. The odds, they calculated, that background could account for the nine events were roughly 1 in 300.

One Experiment, Two Interpretations

These were heady times at LSND, ut they were also times of trouble: had driven a wedge between some members of the group. Alfred Mann, ne of the original collaborators and a rofessor at the University of Pennsylania, withdrew from the collaboration ver what he saw as a dangerous disreard for the scientific method. Mann eared that the group had lost objecvity. The experimenters wanted to see n excess of events, he believed, so hey had unconsciously shaped their nalysis to find one. Mann was an auhority, having worked on many experiments designed to search for the as-yet nobserved, and he preferred simpler methods to the "unnecessarily complex" . "My whole experience says that if ou're going to find something new, it enerally rises up out of the data and okes you in the eye," he said.

At the same time, one of Mann's raduate students, James Hill, had eturned to Philadelphia from Los Alamos and was hard at work on his wn analysis of the LSND data. Hill ecided to restrict himself to the 1994 ata and to dispense with *R*. It made more sense, he argued, to simply cut ut photons that appeared in the bottom f the tank or near the tank walls. hese cuts, however, reduced the data et by a factor of 3. Hill also used less-estrictive selection criteria to pick out hotons. In the end, he found five vents that looked like electron antieutrinos. He calculated the background o be 6.2 ± 1.6 events. By his estimate, he excess was a mirage.

But the two analyses weren't necesarily contradictory. If there were only few oscillation events, it was entirely ossible that they would stand out in ne analysis but not in another—specially, the advocates of *R* argued, the other analysis cut out two-thirds f the data and took a simple-minded pproach to selecting photons. Mann, n the other hand, thought Hill's work was "entirely sensible." It was, he said, the conservative and hence safe approach. The debate was amicable, Mann maintains: "We just differed in our interpretation of the data."

When January 1995 rolled around, LSND again found itself on the losing side of a budget war. DOE had begun to draft a five-year plan that, in the words of one collaborator, "slit the throat of LSND." The collaborators decided it was time to show their hand. They would go to a nuclear and electroweak physics meeting at Berkeley later that month and announce that they had what looked like a hint of oscillations. Since it would be impolite, they reasoned, not to present the results first at Los Alamos, they scheduled an on-site colloquium in advance. Bill Louis, the LSND spokesman, would give the Los Alamos colloquium on Thursday. Hywel White would give the Berkeley talk the following Sunday.

Fit to Print

With preparations underway for the Thursday colloquium, the phone rang with a call that would change everything. It was the *New York Times*. Someone, perhaps at a recent astrophysics conference, had tipped *Times* reporter John Wilford to the LSND results. Wilford called John Gustafson at the Los Alamos public affairs office, and Gustafson, pleased with the prospect of a *New York Times* article, approached the LSND group with the idea of letting the *Times* report on the results. The collaborators on hand at Los Alamos were hesitant about speaking to the press before informing their peers, but in the end they decided that it was better to talk to Wilford than to let the *Times* run a story that could be wrong or overblown. White recalls a sense of helplessness: "It's just like going on a rubber raft," he remembers. "Once you decide to get on, jumping off doesn't make any sense, so you hang on as best you can."

The ride was long and rough, and in retrospect, many wished they had kept to higher, drier ground. The Tuesday before the Los Alamos colloquium, the *Times* ran its story—not buried behind the fashion page in the Science Times section but on page one, just below the fold, making it look like a definitive discovery. "Cosmos's Missing Mass: Wispy Particle Weighs In," the headline read (Wilford 1995). Looking at the article now, White says it seems balanced and accurate, but at the time, it made him swallow hard. In the physics community, there is nothing as close to a sin as "publishing in the *Times*." Colleagues want to have a crack at reviewing new results before they hit the press. "They really screwed the pooch," one physicist remembers thinking after reading the article at breakfast. The incident gave critics a peg to hang their skepticism on.

In the following weeks, LSND collaborators found themselves apologizing for their misstep with the media. "If I could do it again," says Louis, "I would just say 'no comment.'"

At the same time, word of Hill's contrary analysis began to circulate, fueling doubt in the physics community as to the credibility of the results. Louis recalls, "They seemed to be saying 'we don't know what you've got there, but it's not oscillations.'" One physicist at a Fermilab colloquium said White seemed slick, "like a lawyer who knew his client was guilty." Vern Sandberg, an LSND experimenter, says that under the circumstances, he could sympathize with the skeptics: "I wouldn't have believed us either."

Looking back, Imlay thinks the results were a hard sell because few in the audience had the expertise to understand them. And it was true, LSND sat at the intersection of two, vast fields—high-energy physics and nuclear physics—fields that were like adjacent neighborhoods that spoke different languages. "Neutrino physics is a niche," Imlay says, "it's not the sort of thing where you can walk in, hear a talk, and understand it." Things would have gone easier, he suspects, if the group had first published some conventional results to establish credibility. But there had been no time for such niceties.

Doubt even began to trouble the collaborators as they prepared their paper for publication. The process, which by some estimates should have taken two weeks, took two months. Right up until the final edit, Louis recalls getting a flood of comments by E-mail. Had we checked this? Could we change this word here? Hill also wanted to publish his results, and by this time, the relationship between Hill, Mann, and the rest of the LSND group had become so strained that the two factions were unable to consolidate their results into a single paper. In April, they submitted two papers to *Physical Review Letters*, which ultimately ran them back to back. One bore the names of 39 collaborators; the other had a single author, James Hill (C. Athanassopoulos et al. 1995, Hill 1995).

In politics, debates can drag on indefinitely, without hope of resolution. In physics, there is no arguing with nature. With more data, the truth would out. Eventually.

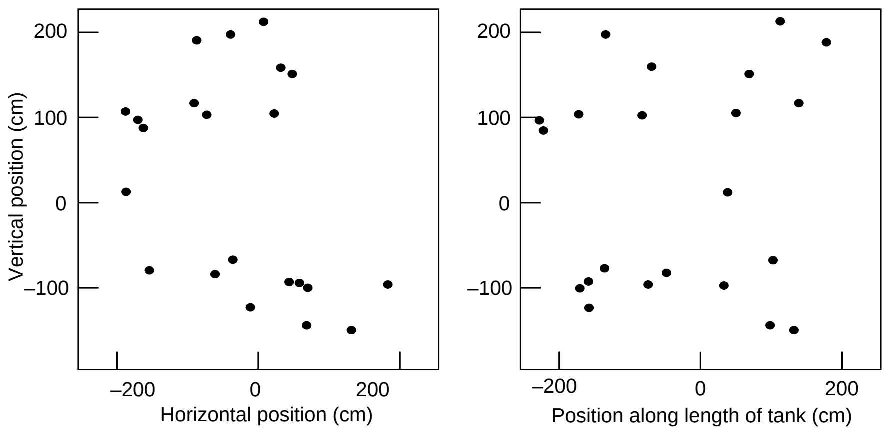

Figure 10. Evidence for Oscillations

The likelihood function *R* was used to sift through the 1995 LSND data to determine if a photon was likely to have come from neutron capture and was correlated with a positron event. In all, there were 22 events when the beam was on that had *R* values >30 (that is, that had a high level of correlation). Only 4.6 ± 0.6 events were expected from background processes. The 22 events appeared scattered across both tank cross sections.

Back Out to Sea

Fortunately, it looked like there would be more data. Though the nuclear physics division of DOE had orphaned LANSCE (which at the time was called LAMPF for Los Alamos Meson Physics Facility) by cutting its funding in 1995, the defense projects division had arranged to assume custody. The new overseers renamed the accelerator LANSCE to reflect what would be its new focus on neutron physics but agreed to allow neutrino production to continue on the side.

Gerry Garvey, an LSND collaborator who was watching events from a temporary post at the White House's Office of Science and Technology, attributed LANSCE's stay of execution to a law of nature. "Good things have a way of continuing," says Garvey, who had been LAMPF's director for five years; "names may change, people will come and go, but where there is will, research will persevere."

Drawing upon outside sources and Los Alamos discretionary funds, LSND cobbled together enough money to run for another four months. The much-anticipated run began in August 1995. That first month, Louis checked every day to see if the detector had recorded any electron antineutrino events. They expected less than one per week. It was like being on a long fishing expedition, staring into the dark waters and waiting for something to bite. But every day the nets came up empty.

Louis began to fear that they might have been wrong. Anxiety woke White at 2 A.M. many mornings to think the experiment through again. Insomnia caught on like the flu, leaving many weary and frustrated. Some questioned the electronics that stood between the collaborators and the neutrinos. Questions popped up like "We've got 1,500 channels here, could there be some mistake? A dirty connection? A glitch in the trigger memory? An electronic hallucination?" But Sandberg, the technological hero who had engineered the data acquisition system, had a parental faith in its performance. Every day he dove into the raw data and made sure the hardware was doing what it should. Still, Sandberg remembers having his doubts, too: "We were worried we might end up in the Journal of Irreproducible Results."

Then on the last day of August, a single event came in. In September there were a few more. At the end of the run, looking at the entire data set, they had a grand total of 22 events with a predicted background of 4.6 ± 0.6 (see Figure 10). "I learned a new appreciation for what low statistics means," recalls Imlay, salvaging a lesson from the nail-biting experience. The group estimated that the odds that all 22 events were background were less than 1 in 10 million. Working backwards, the collaborators calculated the possible regions of Δm^2 and $\sin^2 2\theta$ that could explain the oscillations. Many had been ruled out by previous experiments, but a few small regions stretched tantalizingly out into the unprobed region.

Burned by the spotlight once, the collaborators stonewalled their curious colleagues and took a full five months

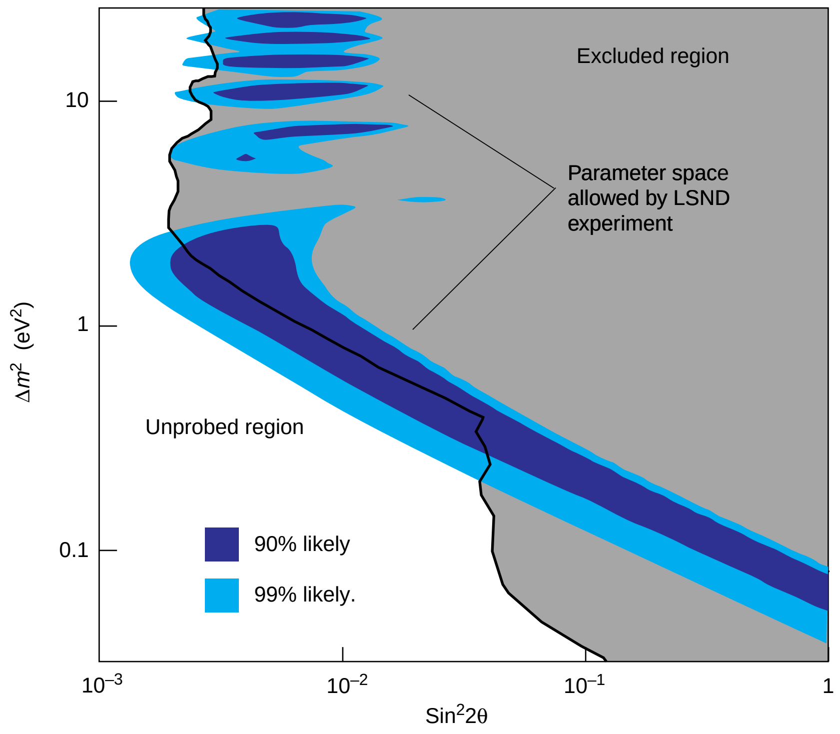

igure 11. Parameter Values from the Initial Analysis of the LSND Data

uch of the area favored by the LSND search (blue regions) had been ruled out by pre-ous searches (the grey region). A small strip and the edges of a few of the “islands” ppeared to be the most likely values of Δm^2 and $\sin^2 2\theta$ that could have generated the scillations. The true parameter values are 90 percent likely to lie in the dark blue egion and 99 percent likely to lie in the light blue region. The shape of the allowed egions is caused by the fact that two terms in the oscillation probability expression ($\bar{\nu}_\mu \rightarrow \bar{\nu}_e$) can vary independently.

reparing their next paper. “People would say, ‘Hello! How are you? Haven’t seen you in a while, so . . . WHAT ARE THE RESULTS?’” andberg remembers, “but we made hem wait.” The idea was to write a aper that would hang together with he certainty and logic of a mathematical proof. Publicity had forced intense ntrospection, and this time the group roduced a 24-page tome entitled Evidence for Neutrino Oscillations rom Muon Decay at Rest” C. Athanassopoulos et al. 1996). A olleague praised it as one of the most xtensive and exhaustive descriptions f an analysis ever published. Around SND, it is simply referred to as “the ig paper.”

The big paper spelled out the analy-s in excruciating detail, at the level of a graduate thesis, and went a long way toward restoring LSND’s credibility. Paranoia had paid off.

In the paper, the collaborators also tried cutting out photons in the bottom of the tank as Hill had done. They found six events with a background of 1.7 ± 0.3 events. The odds were about 1 in 100 that the six events could all be background. Responding to Mann’s allegations that they had willed the excess into existence, they repeated their analysis, varying the requirements for photons and positrons. In each case, they reported an excess of events.

But the physics community reasonably demands a high level of proof before declaring victory. Today Mann, while impressed by the excess, cautions that the results are nothing to yell Eureka about: “If you knew your house had a 1 percent probability of burning down, you’d be out buying insurance,” he says.

Mapping out the Territory

Nothing makes a physicist happier than data, and having observed a statistically significant excess of events, the LSND collaborators proceeded to map out the regions of Δm^2 and $\sin^2 2\theta$ that could have led to the oscillations.

Given the detection efficiency, they could calculate the number of electron antineutrinos that had passed through the tank without getting caught. By comparing this number with the number of muon antineutrinos that emanated from the source, they estimated that only 0.31 ± 0.13 percent of the muon antineutrinos had oscillated. To unfold the Δm^2 and $\sin^2 2\theta$ information from the data, the collaborators performed a likelihood fit to all events that contained a positron. The fit took into account the positron’s energy and direction, the photon likelihood R, and the distance to the source for each event. It also took into account the expected distributions of these quantities for electron antineutrinos from oscillations and for background processes.

As shown in Figure 11, the fit does not pinpoint particular values of Δm^2 and $\sin^2 2\theta$ but rather carves out regions of more-likely values. The shapes of the regions are a consequence of the fact that the oscillation probability is the product of two terms, one relating to Δm^2, x, and E, and the other to $\sin^2 2\theta$. The spots that spread out like a chain of small islands arise because some oscillation events at relatively high energy tend to exclude Δm^2 near integral multiples of 4.3 eV^2. (Those values of Δm^2 make multiples of the oscillation wavelength approximately equal to the source-to-detector distance. Hence, $\sin^2(\pi x/\lambda)$ is almost zero.) The longer “island” corresponds to smaller values of Δm^2, where $\sin^2(\pi x/\lambda)$ no longer oscillates but slowly approaches zero and where the data must be accounted for entirely by increasing the allowed values of $\sin^2 2\theta$.

Figure 12. LANSCE Production of Decay-in-Flight (DIF) Neutrinos and Their Energy Spectra

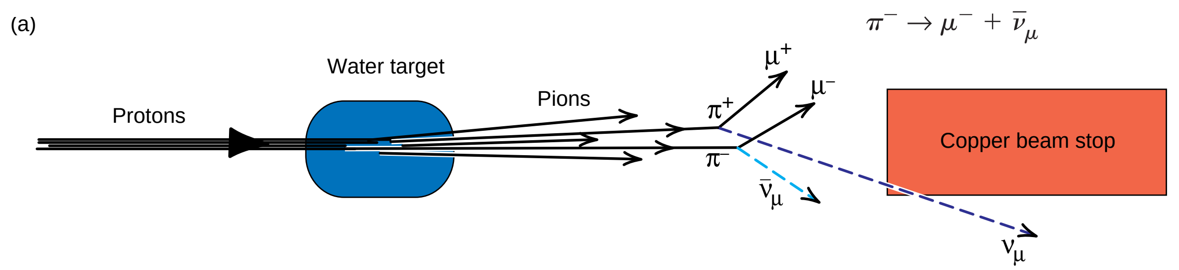

$$\pi^+ \rightarrow \mu^+ + \nu_\mu$$

$$\pi^- \rightarrow \mu^- + \bar{\nu}_\mu$$

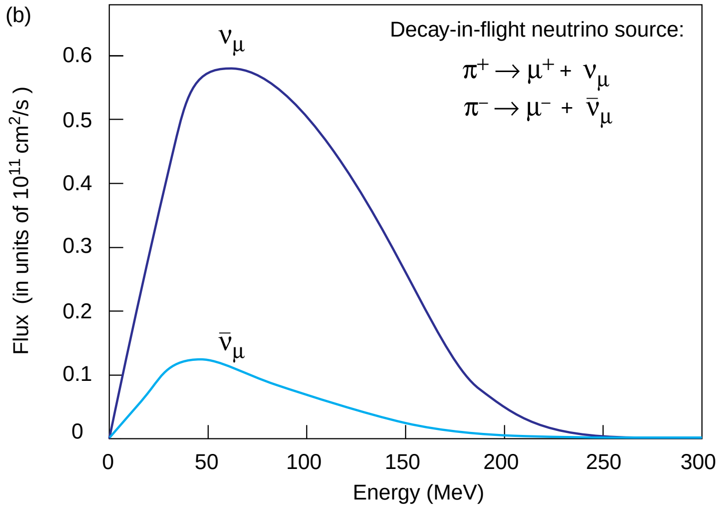

(a) Most pions emerging from the water target strike the copper beam stop and produce neutrinos when they decay at rest. Occasionally, however, a pion emerging from the target will decay in flight before reaching the beam stop. This decay produces both muon neutrinos and muon antineutrinos. The muon neutrinos are the source for the DIF analysis. (b) Because the pions are in motion when they decay, the resulting neutrinos can have energies many times those of the DAR neutrinos (compare with Figure 5). Muon neutrinos outnumber muon antineutrinos because the target produces 8 times fewer negative pions than positive pions. Very rarely (0.001 percent of the time), a positive muon also decays in flight and gives birth to an electron, a muon antineutrino, and an electron neutrino. Roughly 1 electron neutrino is produced for every 1,000 muon neutrinos.

Many of the regions allowed by the LSND data, however, had already been explored and found barren (see Figure 4). An experiment at Brookhaven National Laboratory (E776) and a French reactor-based experiment (BUGEY III) had essentially ruled out all of LSND’s preferred areas except for a narrow strip that stretched from Δm^2 = 0.2 to 20 eV^2 and from $\sin^2 2\theta$ = 0.03 to around 0.001. Although the LSND results taken alone allowed a variety of interpretations, when combined with the previous null results, the parameter space for neutrino masses and mixing angles was quite limited.

Decay in Flight—The Second Analysis

While the LSND collaborators were drawing up Figure 11 and putting the finishing touches on the big paper, they began work on a second analysis that would either confirm or disprove all their previous work. Knowing that an excess of electron antineutrinos was not enough to establish that oscillations had occurred (the electron antineutrinos could be coming from some equally surprising source, such as an exotic type of muon decay), the collaborators had built in a second method to look for oscillations. Instead of using the neutrinos from decay-at-rest (DAR) pions and muons, this method looked to a smaller sample of neutrinos produced by the pions that decayed in air on their way to the beam stop. These decay-in-flight (DIF) neutrinos had a much higher energy than the DAR neutrinos and so could be easily distinguished in the detector.

In the production of DAR neutrinos, pions produced in the target travel about a half meter through open space before striking and coming to rest in the copper beam stop (refer to Figure 5). When the still pion decays, it can give at most 53 MeV, or about half of its mass, to one of the resulting neutrinos. Thus all neutrinos used in the DAR analysis had energies below 53 MeV. By contrast, a pion that decays in flight before reaching the beam stop passes on some of its kinetic energy to the resulting neutrino, giving it as much as 300 MeV of energy. By considering only neutrinos that had at least 60 MeV of energy, LSND collaborators could essentially tap a second neutrino source, for free. Figure 12 shows this DIF source and the energy spectrum of the neutrinos produced by it.

Unlike the DAR neutrinos, DIF neutrinos come mainly from pion decay, since there is rarely time in the half-meter journey to the beam stop for the muon produced by the pion to also decay. (Roughly 3 percent of the pions

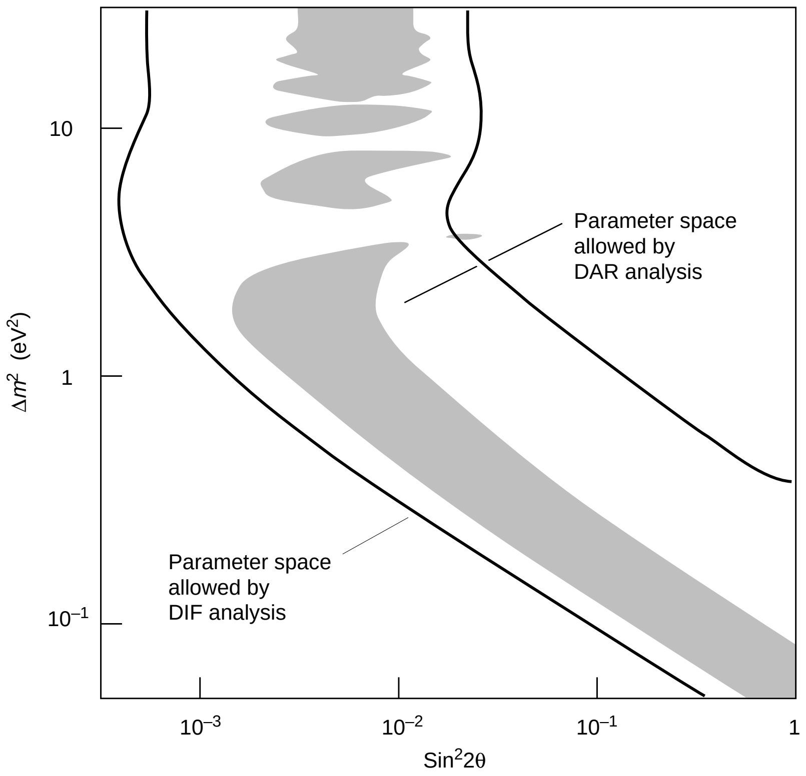

igure 13. Parameter Values Allowed by DAR and DIF Neutrino Analyses
he DIF analysis also indicated neutrino oscillations: a small excess of electron eutrino events was detected above background events. The parameter values allowed y the DIF analysis lie between the two solid lines, whereas the values allowed by the AR analysis are shown in grey. Although there were greater uncertainties associated ith the DIF analysis, the allowed values of Δm^2 and $\sin^2 2\theta$ for the two analyses verlapped.

ecay in flight, but only 1 in 100,000 muons do.) As a result, the DIF neu- inos are mostly muon neutrinos and ontain relatively few contaminating lectron neutrinos. The DIF analysis ooked for muon neutrinos to oscillate nto electron neutrinos—the "matter" ounterpart of the DAR analysis. If the DAR analysis observed oscillations, so oo should the DIF analysis.

Alfred Mann thought the DIF analy- s so important, and so integral to the mission of LSND, that he had urged the ollaborators to keep a low profile until hey finished it. But the LSND group ad gone out on a limb with their DAR nalysis, and it remained to be seen whether the DIF data would support he earlier analysis.

While the DIF analysis sought to observe the same kind of oscillations as the DAR analysis, it did so in a way that was completely independent. Even the way the detector observed the oscillations was different. The electron antineutrinos in the DAR analysis struck a free proton in the oil and produced a low-energy positron and a neutron. By contrast, the electron neutrinos in the DIF analysis interacted with a neutron in a carbon atom from the oil, producing a high-energy electron and transforming the carbon into another atom (X), typically nitrogen:

$$\nu_e + \mathrm{C} \rightarrow e^- + \mathrm{X} \quad .$$

Unlike in the DAR reaction, no neutron is produced, so the DIF analysis boiled down to the difficult task of separating electrons from background sources such as cosmic rays.

Since the detector is essentially charge-blind, identifying electrons is just like picking out positrons in the DAR analysis. But unlike the positrons, the DIF electrons can have considerable energy, so much so that they may travel for half a meter before stopping in the oil. As a result, the "sphere" of scintillation light the electron produces in the tank becomes stretched out, looking instead like the superposition of spheres from a string of electrons.

The LSND group developed two methods for selecting electrons. Both made careful study of the amount and timing of the light expected to hit each phototube. Both also looked for a Cerenkov cone and scintillation light and discriminated against cosmic rays, although in slightly different ways. Finally, both methods offered improved position and direction resolution over the positron method used in the DAR analysis. In the end, each method had its own strengths and weaknesses, so the collaborators decided to use both. Taken together, the two methods were expected to pick out roughly 17 percent of the electrons produced by electron neutrinos.

The collaborators identified 40 events in the data set that seemed to contain a high-energy electron. They had expected roughly 11 events from backgrounds such as cosmic rays and 10 events from electron neutrinos present in the beam, leaving an unexplained 19.2 ± 7.8 events. They calculated the probability that backgrounds could account for the excess to be less than 1 percent. More importantly, when they worked backwards to see what neutrino masses and mixing angles could have generated the oscillations, the results overlapped quite nicely with the earlier DAR results (see Figure 13). The two analyses seemed to point to the same conclusion—that oscillations and massive neutrinos were behind the excess events.

Epilogue

At the time of this writing, LSND had just begun to make its DIF results public. This time around, reactions are more enthusiastic than condemning, and question-and-answer sessions no longer run an hour. Whereas before, LSND results were often downplayed in neutrino talks, they now take center stage along with the atmospheric and solar data.

Neutrinos defied detection for nearly 25 years after Pauli first proposed them, and today, near the 70-year mark, physicists still disagree over whether neutrinos have mass. While there seems to be a growing suspicion in the physics community that neutrinos do indeed have mass, many are still waiting for the day when some experiment sees the cyclic rise and fall of the number of neutrinos from oscillations as the neutrino energy or the distance between the source and the detector is gradually altered. It would be impossible to stare those results in the face and deny that neutrinos have mass.

That day would mark the end of one of the longest quests in the history of particle physics, one that currently stretches over half a century and spans generations of physicists. It would also reserve a place in the history books for LSND, solar, and atmospheric experiments. But even those intimately involved in neutrino work are uncertain exactly when all this might come to pass. White has a page from a word-a-day calendar tacked to the wall of his office, just beside a copy of the troublesome *New York Times* article. The word is obscure: *Greek calends*, defined as a time that will never arrive, the next blue moon, or when pigs fly. With the mounting evidence, it may be that pigs are preparing for takeoff. ■

Further Reading

Athanassopoulos, C., et al. 1995. Candidate Events in a Search for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ Oscillations. *Physical Review Letters* **75**: 2650.

Athanassopoulos, C., et al. 1996. Evidence for Neutrino Oscillations from Muon Decay at Rest. *Physical Review C* **54**: 2685.

Athanassopoulos, C., et al. 1997. Evidence for $\nu_e \rightarrow \nu_\mu$ Oscillations from Pion Decay in Flight. Submitted to *Physical Review C.* LA-UR-97-1998.

Hill, J. E. 1995. An Alternative Analysis of the LSND Neutrino Oscillation Search Data on $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$. *Physical Review Letters* **75**: 2654

Louis, W. C., et al. 1994. Neutrino Oscillation Studies at LAMPF. *Nuclear Physics B* (Proc. Suppl.) **38**: 229.

Mukerjee, M. 1994. Missing Matter Found? *Scientific American* **271** (2): 22.

Wilford, J. 1995. Cosmos's Missing Mass: Wispy Particle Weighs In. *The New York Times*. January 31, 1995.

William C. Louis is currently the scientific spokesman of the liquid scintillator neutrino detector (LSND) experiment at LANSCE. Following a postdoctoral appointment at Rutherford Laboratory and a faculty appointment at Princeton University, Louis joined the Laboratory in 1987 as a staff member of the Medium Energy Physics Division. In 1993 he joined the Physics Division's Subatomic Physics Group. Louis earned his Ph.D. in physics from the University of Michigan in 1978, studying high-energy neutrino interactions at the then recently opened Fermi National Accelerator Laboratory. Louis' research at Los Alamos has centered on weak interactions, neutrino reactions in nuclear and particle physics, and neutrino oscillations.

Vern D. Sandberg received his Ph.D. in 1975 from the University of Utah, where he studied strong gravitational fields. Following a postdoctoral appointment at the Center for Relativity Theory at the University of Texas at Austin, Sandberg accepted a postdoctoral position at CalTech, where he participated in the discovery of quantum non-demolition techniques that circumvented the "quantum limit" for the readout of gravitational wave detectors. In 1979, he joined LAMPF (now LANSCE), working with Darragh Nagles to develop instruments for precision experiments. The high intensity of LAMPF's accelerator beam opened up possibilities of experiments with well-defined beams of neutrinos, and soon thereafter he collaborated on the neutrino-electron elastic scattering experiment that measured the interference between the weak neutral current and the charged neutral current in the scattering of electrons by neutrinos. This experiment led to a program to search for neutrino oscillations, which culminated with the LSND experiment. Sandberg is currently involved in the neutrino oscillation physics program, consulting on gravitational wave detectors and building a quantum computer.

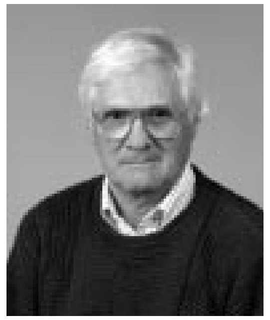

D. Hywel White received his B.Sc. in mathematics and physics from the University of Wales in 1953, and earned his Ph.D. in experimental particle physics from Birmingham University, England, in 1956. After a stint on the faculty at Birmingham University, White became an assistant professor at the University of Pennsylvania. In 1964 he was appointed associate professor, and in 1968 professor, at Cornell University. White joined Broookhaven National Laboratory in 1978 as head of experimental facilities on the ISABELLE project, where he began experiments in neutrino physics. In 1986 he came to the Laboratory as group leader of Nuclear and Particle Physics Research at LAMPF to continue work in experimental neutrino physics. In 1994 the neutrino group observed a signal interpreted as neutrino oscillations, an event that has since dominated their research.

From Tank to Tape—The LSND Data Acquisition System

Neutrinos interact in the LSND tank at a rate of approximately one an hour. Go to lunch, and if you're lucky, there will be a new neutrino event written to tape when you get back. But in that same hour, nearly 15 million cosmic rays will also have left their marks in the tank. That's a staggering number to contemplate. "You have all this background from cosmic rays," says Vern Sandberg, the principal designer of the LSND data acquisition system, or DAQ. "But you want to be absolutely sure that a positron, which is the primary thing we look for, is isolated from anything to do with cosmic rays. You need some way to separate the wheat from the chaff, so to speak. The only way you can convincingly sort it all out is by keeping track of what happened before the positron showed up and after it was detected."

The DAQ keeps track. It is an array of electronics (one circuit board is shown in the photo) that can be thought of as a kind of brain with the single-minded task of identifying potential neutrino events whose signature is a positron followed by a 2.2-MeV gamma ray. When the DAQ identifies a promising positron, it grabs from its short-term memory everything that happened in the tank for 6 microseconds (µs) before the positron was detected. It also records all gamma-like activity that occurs within the next 1 millisecond (ms). Armed with this information, the DAQ tries to make sense of what it saw by correlating, in space, time, and energy, the positron signal with a gamma ray. If the correlation matches the profile of a neutrino interaction, it writes the information to long-term memory (a magnetic tape).

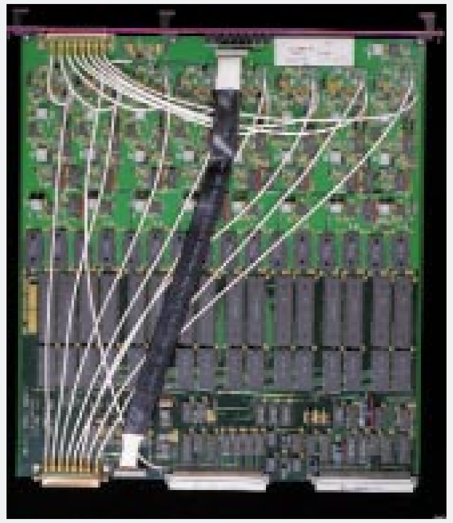

his circuit board is one of 210 that make p the LSND data acquisition system. A oard contains eight individual channels, nd each channel consists of analog rcuitry for processing signals from one hotomultiplier tube and digital circuitry hat stores and/or sends that data to the omputer. The board is built from mostly ff-the-shelf components and is thus irly robust and inexpensive. To achieve flexible yet highly reliable system, each oard was designed to be used in a andard VME electronics crate. This eans that the DAQ can be easily dapted to meet the needs of other xperiments.

The human brain, with its exceptional pattern recognition ability, evolved over hundreds of millions of years. The DAQ used by LSND was designed and assembled in less than 1 year by Sandberg and a team of students and visiting staff as they scrambled for funding and raced to finish before LANSCE started producing neutrinos. By all accounts, the final system has been a smashing success. "It's a unique system," says Darryl Smith, who was involved in developing the DAQ. "No other data collection system has this look-back capability, to see what was happening in the detector *before* the triggering signal occurred. If in the future someone asks, 'Did you check for this, or look for this oddball correlation,' we can go back to the data and see."

How the DAQ searches for a neutrino signature is shown in the diagram on pages 114 and 115. In the text that follows, the letter callouts correlate with those on the figure.

The Front End—Selecting Promising Signals. **(*A*)** The DAQ perceives the world through 1,220 photomultiplier tubes (PMTs) that line the inside of the detector tank. Like huge eyes, the phototubes watch for the brief pulses of light produced when energetic particles pass through the oil. Outside the tank, 292 more phototubes sit within the cosmic-ray veto shield to signal the arrival of cosmic rays. When any tube is hit by a photon, it sends a tiny current pulse to a digital version of short-term memory called a circular buffer. There is one buffer per PMT.

(*B*) Conceptually, the circular buffer is analogous to an office Rolodex filled with 2,047 electronic "cards." At every time T, where T is the system time [measured in 100-nanosecond (ns) units, that is, $(T) - (T-1) = 100$ ns], the Rolodex is "turned" and the value of the electric charge q in the current pulse and the precise time t that the pulse occurred (accurate to about ±0.5 ns) are written to a card. The card is stamped with the system time T. If a tube was not hit, $q = 0$ and $t = 0$, but the data are still written to the card, which still receives a time stamp. Because a tube is hit on average only once every 200 µs, the data on most cards are just zeroes.

At 100 ns per card and 2,047 cards, each circular buffer maintains a 204-µs history of its phototube. The DAQ can access any portion of that history by asking the buffer to "dump" the information contained on specific cards, for example, the cards stamped $T-1$ through $T-60$. Because it has access to data from all 1,512 circular buffers, the DAQ in principle can construct detailed, 204-µs histories of everything that occurred in the tank and veto shield.

But a full history comes at a steep price tag. There are over 15 million pieces of data coming in per second, far too much data for a computer to mull over in detail. The system needs to quickly cull promising positron signals from the heavy traffic of cosmic rays and other "stuff" that cruises through the tank. Thus, the DAQ applies some rules of thumb so that it can react "instinctively." **(*C*)** In addition to producing an analog current pulse, each phototube when hit by a photon sends a digital pulse to a summing circuit. The digital pulse simply indicates that the tube fired in the preceding 100 ns, and the sum of the pulses gives the total number of tubes that fired in the tank and the veto shield within that time interval. (Tank and veto shield sums are kept separate.) The information goes to a programmable "trigger" (called the signal flagger) that crudely identifies the signal. For example, if at least 21 phototubes fire in the detector and less than 4 fire in the veto shield, the signal is flagged as a gamma-ray candidate. If at least 150 tubes fire in the detector and less than 4 fire in the shield, the signal is flagged as a positron candidate. Cosmic-ray muons typically light up the tank like a Roman candle, setting off 250 to 1,000 or more tubes. Flagged signals, and the times that they occurred, are passed onto a trigger computer charged with the task of selecting promising ones.

(*D*) The trigger computer monitors and stores signals as they roll in, always on the lookout for a positron candidate. When it sees one, it reviews the signal roster, checking back in time to see if there was any activity in the tank within a 15-µs "all-quiet" period before the positron appeared. If there was activity, it ignores the signal and keeps looking. The reason for the all-quiet condition is to weed out positrons that come from muons.

"Cosmic-ray muons that decay in the tank are the bugaboo of this experiment," says Sandberg. "They produce a Michel electron that is identical to the positron produced by electron antineutrinos. The average lifetime of the muon is about 2 µs, so if there's any suspicious activity within about 7 muon lifetimes before a positron signal, we don't want to waste time checking that positron. We're not interested."

Cosmic rays, beta particles from radioactive decay, neutrons, and muons create a riotous background of activity in the tank, and flagged signals are sent to the trigger computer at an average rate of one every 60 µs or so. Positrons account for about a third of that rate. With the all-quiet condition, the rate of signals that will be examined further is cut down to about one every 10 ms. That rate is slow enough that a large, mainframe "analysis" computer—the cerebrum of the DAQ brain—can examine the selected signals in detail. Because the decision to flag a signal is based on low-level information (the total number of tubes that fired in the tank and veto shield), it can be made quickly. **(*E*)** Within about 400 ns of finding a clean positron with no prior history, the trigger computer sends a message to all circular buffers to dump their detailed information.

The Back End—Reconstructing Events. **(*F*)** If the positron signal occurred at time T, the trigger computer tells the circular buffer to dump the data stored on the card stamped T into a temporary storage bin called a FIFO (an acronym for first in, first out). There is one FIFO per photomultiplier. In addition, the trigger tells the circular buffer to dump into the FIFO the 60 pieces of (q,t) data corresponding to the 6 µs before T. This "look-back" information is a double check on the signal. Anything that happened in the 6 µs before the positron made itself known in the tank will be looked at in detail. Finally, all gamma-like signals that occurred within 1 ms after T are also placed in the FIFO.

The FIFO allows the analysis computer, which has the tough job of trying to figure out what happened, to leisurely collect the detailed information from each tube. **(*G*)** The analysis

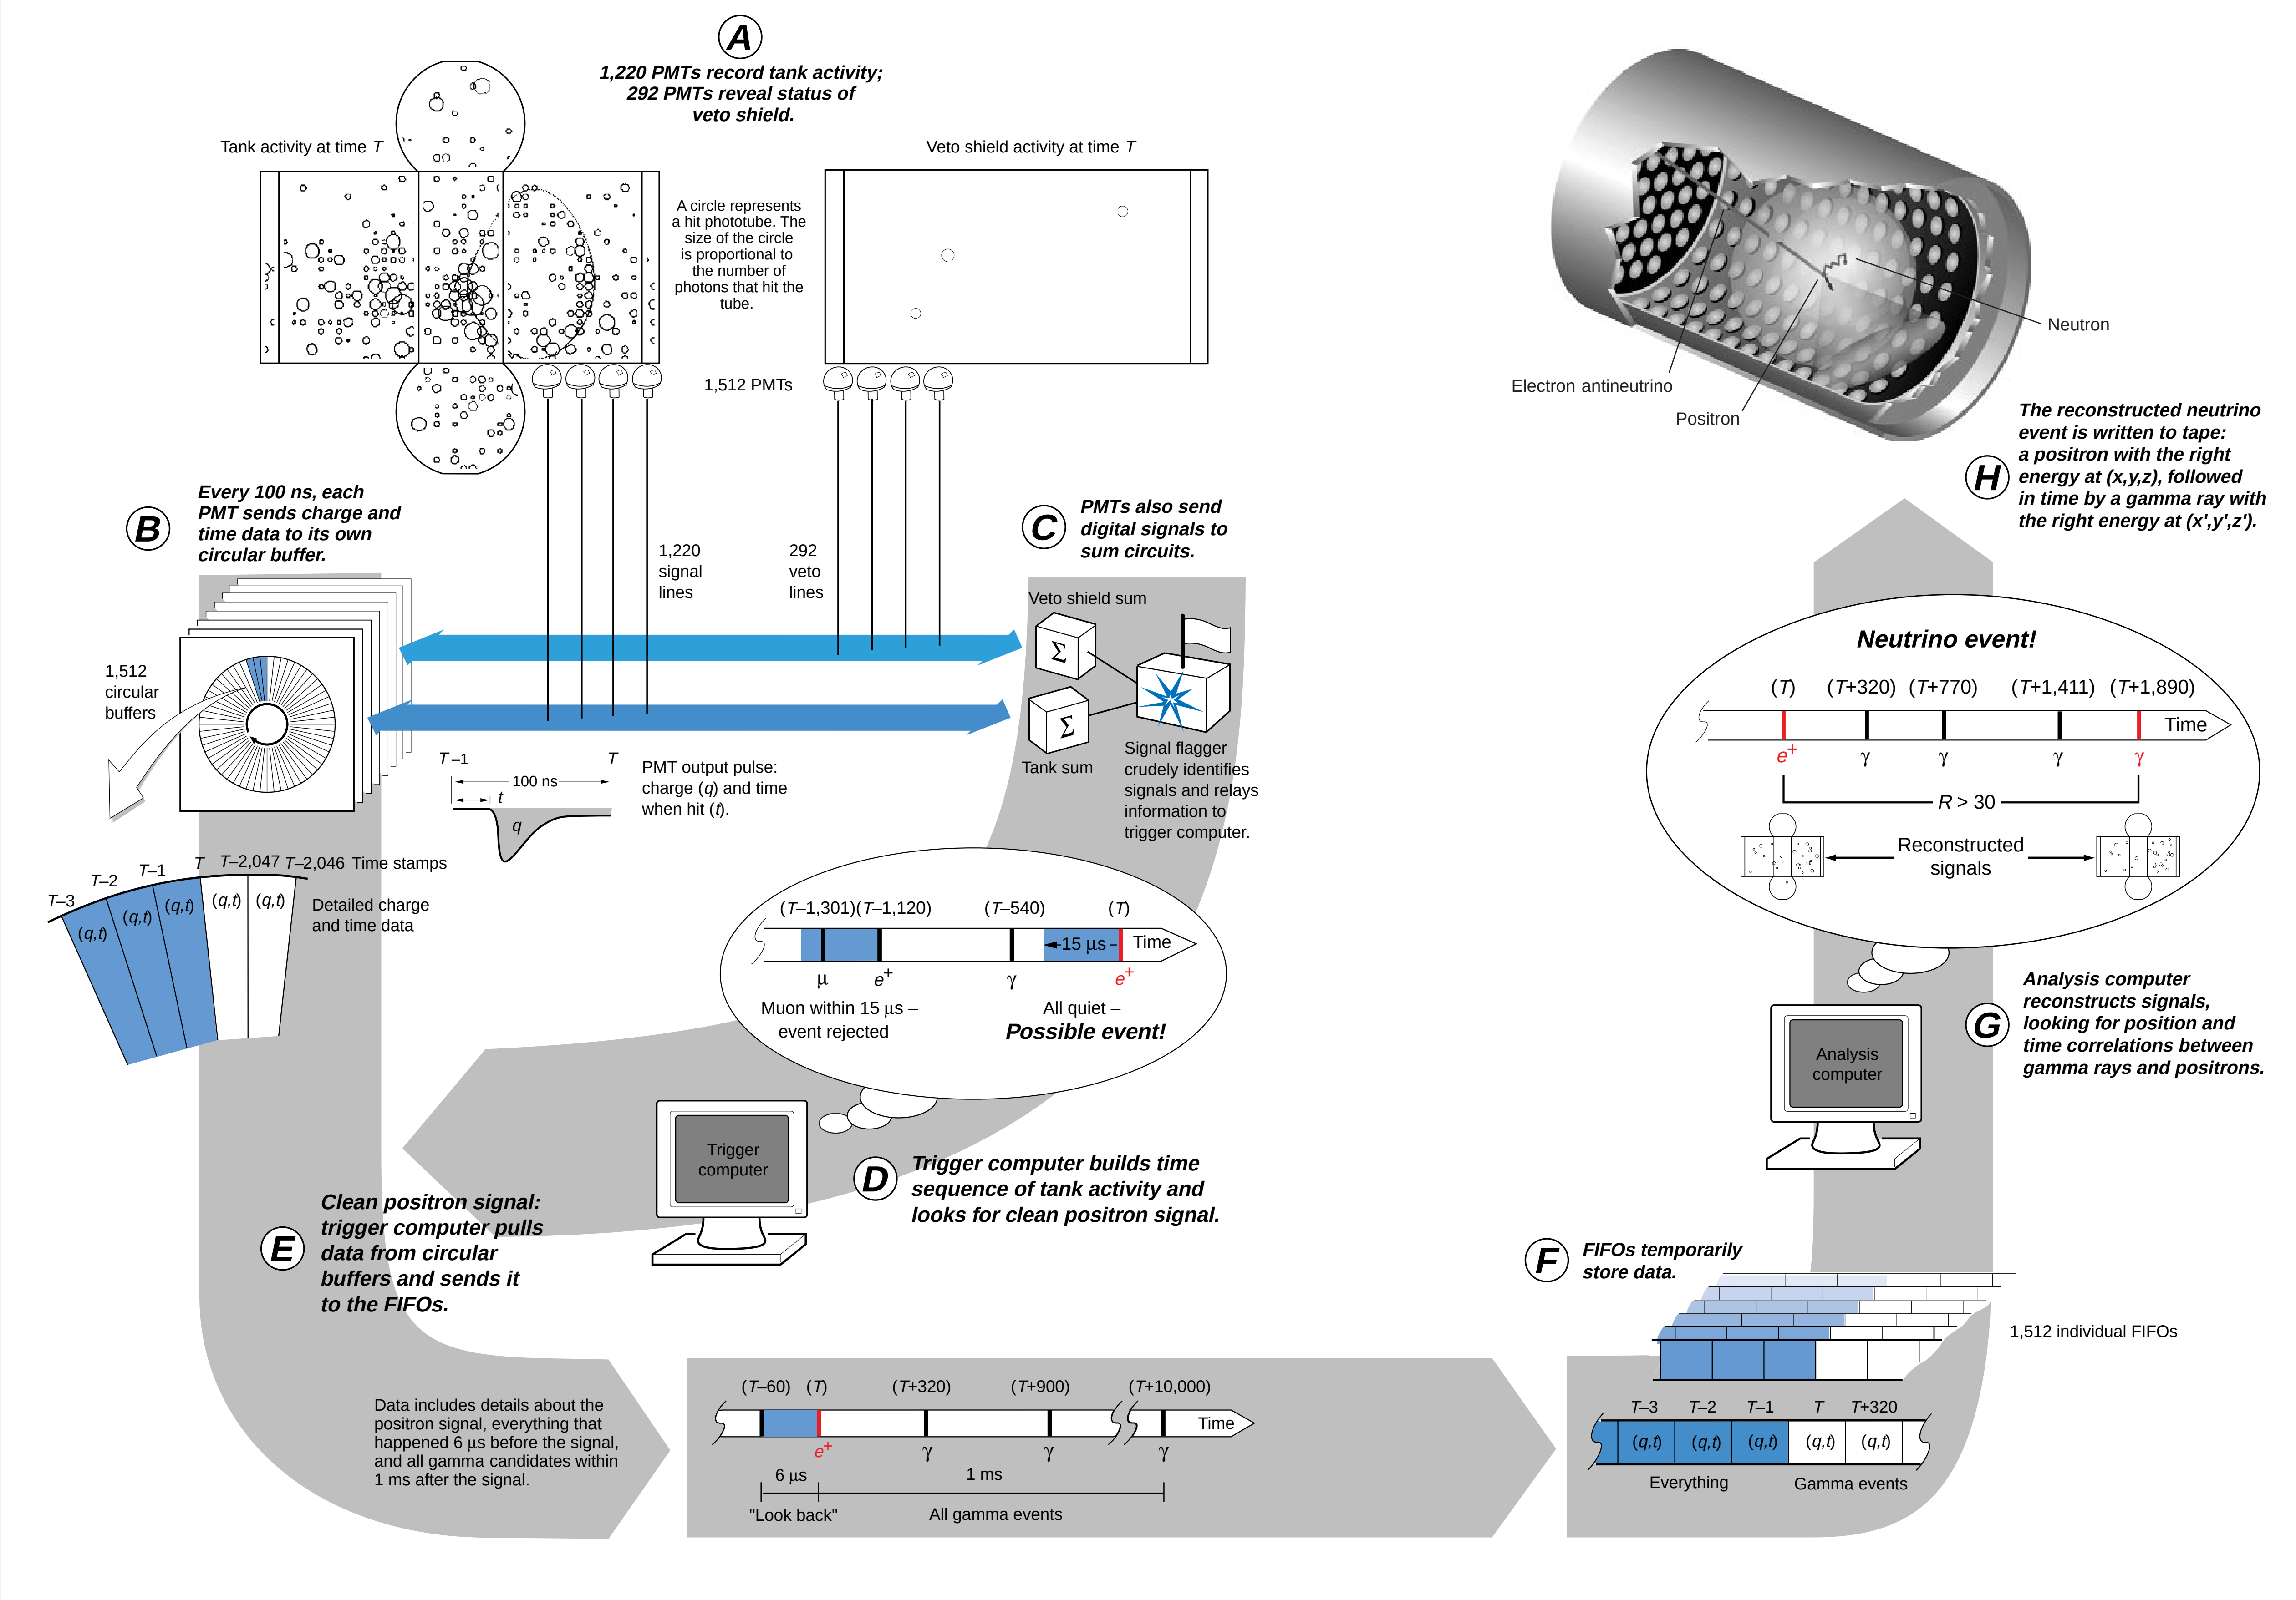

computer gathers the detailed photomultiplier data from all 1,512 FIFOs and reconstructs the signals. First, it finds the Cerenkov cone and the sphere of scintillation light and determines the positron's trajectory and position. Next, it finds the position and energy of the gamma rays (by reconstructing their spheres of scintillation light) and then uses the likelihood function *R* to find a gamma that has the right time, energy, and position to have come from neutron capture. If *R* is high (>30), the sequence of signals—a positron followed by a gamma ray from neutron capture—is taken as the signature of a neutrino event. ***(H)*** The reconstructed neutrino events are written to tape.

Since 1993, the trigger computer has looked at half a billion flagged signals, of which twenty-two were identified as the signature of electron antineutrinos.* It was a significant challenge to design a system that could handle such a low, "asynchronous" event rate. The DAQs used in most particle physics experiments operate on a clock that is synchronized with the particle beam, so that the electronics know exactly when to pay attention. By contrast, neutrino oscillation events in the LSND experiment appear at almost random times. The DAQ has to operate continuously, look at all events, but select only a tiny subset. "Traditional experiments threw stuff away," says Sandberg. "We couldn't afford to do that. But we also couldn't afford to keep it all." The front-end selection of signals and the look-back capability of the DAQ helped solve the problem.

"Because neutrino oscillations have come and gone from decade to decade," says Smith, "we needed as much credibility as we could bring to bear on the problem. This DAQ is totally solid. It's built almost entirely from off-the-shelf components. It can live in an enormously noisy environment. It can look back in time. Because it's built using a standard VME architecture, any graduate student could plunk it into an experiment and get it up and running. It works." ■

---

* Through a separate analysis, nineteen electron neutrino events were also identified.