The universe is just swimming in neutrinos, and if these guys have even the most infinitesimal mass imaginable, it might be enough to account for the elusive "dark matter"—the 90 percent or so of the mass of the universe that we can't see but know must exist, or else galaxies would fly apart from their own centrifugal force.

The Secret Life of Neutrinos

by Douglas L. Smith

The Palo Verde Nuclear Generating Station is the largest nuclear power plant in the U.S. The plant's three reactors (red arrows) are spaced along a circular arc. The white arrow points to the neutrino detector.

Sit quietly, and count off 10 seconds to yourself. Roughly 200 trillion neutrinos from the sun, from cosmic rays, and from distant supernovas have just passed through you, but you'd never know it. Neutrinos are the ghostliest of subatomic particles. They have no electrical charge, so they're not subject to electromagnetic forces. They're immune to the strong nuclear force, which binds atomic nuclei together. In fact, you could shoot your average neutrino through a light-year's worth of lead bricks before anything would happen to it. These few interactions are a result of the weak nuclear force—a wimpy excuse for a fundamental force that causes neutrons to turn into protons via a process called beta decay, and whose effective range is less than the diameter of the decaying neutron. And until recently, everybody thought neutrinos were massless, like photons of light.

Or are they? The biggest physics news of 1998 was that a Japanese experiment called Super Kamiokande intimated that these evanescent creatures might have just a whisper of mass after all. This set theorists abuzz, because a glimpse of phenomena beyond the so-called Standard Model is the sort of thing that can lead to a Grand Unified Theory of Everything, and eventually to a Nobel Prize. And it set cosmologists abuzz, because the universe is just swimming in neutrinos, and if these guys have even the most infinitesimal mass imaginable, it might be enough to account for the elusive "dark matter"—the 90 percent or so of the mass of the universe that we can't see but know must exist, or else galaxies would fly apart from their own centrifugal force. So now a collaboration headed by Felix Boehm, Valentine Professor of Physics, Emeritus, and including people from Stanford, the University of Alabama, Arizona State University, and the Arizona Public Service Company, hopes to find out how much mass neutrinos have, using the Palo Verde Nuclear Generating Station 60 miles west of Phoenix, Arizona. Nuclear power plants are dandy neutrino sources, pumping out huge fluxes of them in accurately calculable amounts at precisely known energies.

Super Kamiokande didn't measure the neutrino's mass directly—that's not yet possible—but instead measured the difference between the masses of two types of neutrinos. Neutrinos come in three "flavors" mirroring the three kinds of particles produced along with them during beta decay—the electron, the muon, and the tau. When a neutron, which has no electrical charge, turns into a proton (charge +1), the quantum accountants force it to emit an electron (charge −1) as well, in order to preserve the overall charge (zero) of the decay reaction. But the books also have to balance on the number of particles and antiparticles, and now there's a net gain of one particle on the right-hand side of the ledger. Thus an electron antineutrino gets created to keep the auditors happy and the sum of the particles and antiparticles unchanged; the same goes for decays featuring muons and taus. (So whenever we're talking about neutrinos here, we really mean antineutrinos, but that's just too much of a mouthful to keep repeating.) If neutrinos were truly massless, like photons, that would be the end of the story. But if the neutrino's mass is not exactly zero-point-zero-zero-zero-zero-zero-zero-zero-zero to as many decimal places as you care to go, a seldom-invoked clause in the laws of quantum mechanics says that there's a small but calculable possibility that the neutrino will eventually change flavors, due to an overlap in the wave functions that describe them. (The degree of overlap is called the "mixing angle.") The farther a neutrino travels, the greater the odds that it will have changed; but the new particle also has a probability of changing flavors. If you followed the career of this neutrino long enough, you'd see it oscillating between flavors, and the wavelength of the oscillation depends on the difference between the flavors' masses. Super Kamiokande,

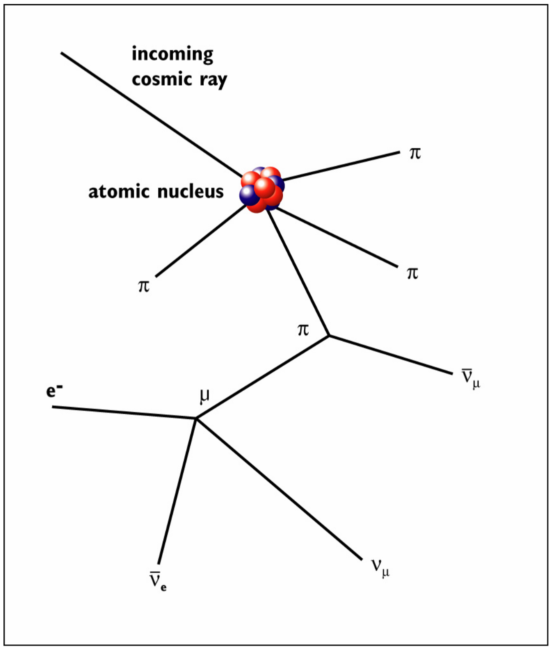

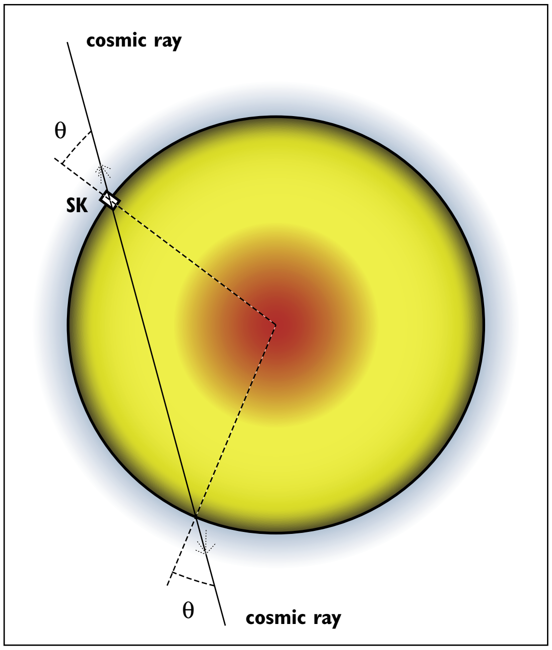

Top: An incoming cosmic ray (usually a high-energy proton, but sometimes something heavier) slams into the nucleus of an air molecule in the upper atmosphere, creating a shower of pions (π) and other particles. Each pion quickly decays into a muon (μ) and a muon antineutrino ($\bar{\nu}_\mu$); the muon in turn decays into a muon neutrino (ν_μ), an electron (e^-), and an electron antineutrino ($\bar{\nu}_e$). Thus for each electron neutrino detected, there should be two muon neutrinos.

Bottom: Cosmic rays bombard Earth from all directions, but shallow-angle particle showers may not reach the surface. Still, equal ratios of neutrinos should be seen in diametrically opposite directions, because the angle of incidence (θ) into the atmosphere is the same. Super Kamiokande (SK) found the expected 1:2 ratio of electron to muon neutrinos coming down from above, but only about 1:1.3 coming up from below.

which was watching for the electron and muon neutrinos created when cosmic rays collide with atomic nuclei in Earth's upper atmosphere, counted equal numbers of electron neutrinos coming from all directions, but fewer muon neutrinos coming up through the planet than down from above. The explanation was that the muon neutrinos taking the longer road through Earth had had more time to change into tau neutrinos. (In 1992, the original Kamiokande experiment had revealed that only about half as many muon neutrinos were making their way into the detector as physicists calculated there should be, but it lacked the directional sensitivity to say from whence the shortfall came.) Super Kamiokande's results indicated a difference between the masses of the muon and tau neutrinos of one ten-millionth of the mass of an electron or less, which is so small it makes one's brain hurt just thinking about it. But if there's a measurable difference in the neutrino masses, then obviously they *have* mass, and that's the point.

And here's where the story starts getting tricky. You might wonder why, if the muon neutrinos were becoming tau neutrinos, Super Kamiokande didn't see tau neutrinos. Well, that's because electrons, muons, and taus differ in their masses but are otherwise identical—they undergo the same reactions and have the same properties. And the mass is the key—because $E = mc^2$, before a particle can spring into existence in a nuclear reaction, there has to be enough energy available to transmute into the particle's mass. Electrons are the lightest of the three, so they are the easiest to create. Muons are 200 times more massive than electrons, but still only one-fifth the mass of a neutron—a piece of cake to make at cosmic-ray energies. Taus, at 3,500 times the mass of an electron, are just too darned heavy. They're only found in the most powerful particle accelerators. But neutrino detectors work by running the beta-decay reaction backward, creating particles that are easier to see. So when the muon neutrinos oscillate into tau neutrinos, they become cloaked in invisibility—they won't register in the detector because they don't have enough energy to create tau particles. Until they oscillate back again at some point further on, they have, for all practical purposes, vanished. It's all in the bookkeeping. We could have used some particle physicists to balance the federal budget back in the '80s.

Nuclear power plants are relatively low-energy systems—at least compared to cosmic rays and particle accelerators—and don't have enough oomph to make muons, let alone taus. Consequently, the Palo Verde Neutrino Detector only sees electron neutrinos. Therefore, what Boehm et al. are really looking for is not the neutrinos, but their disappearance as they change flavors. A straightforward set of calculations based on the reactor's power level and operating characteristics shows how many neutrinos the plant is cranking out. Calibrations done when the reactor is shut down for refueling tell the researchers what the background levels are, and how efficient the detector is. (These calibrations aren't as easy as they might be, because the Palo Verde power plant has three reactors, located at varying distances from the detector, and they're shut down in rotation so that only one, at most, is off-line at any given time.) So if the collaborators see every neutrino their calculations tell them they're entitled to see, then the neutrinos aren't oscillating to an appreciable degree. But if the detector records fewer neutrinos than predicted, it shows that they're oscillating. And the visible neutrinos, taken as an aggregate, should have a specific

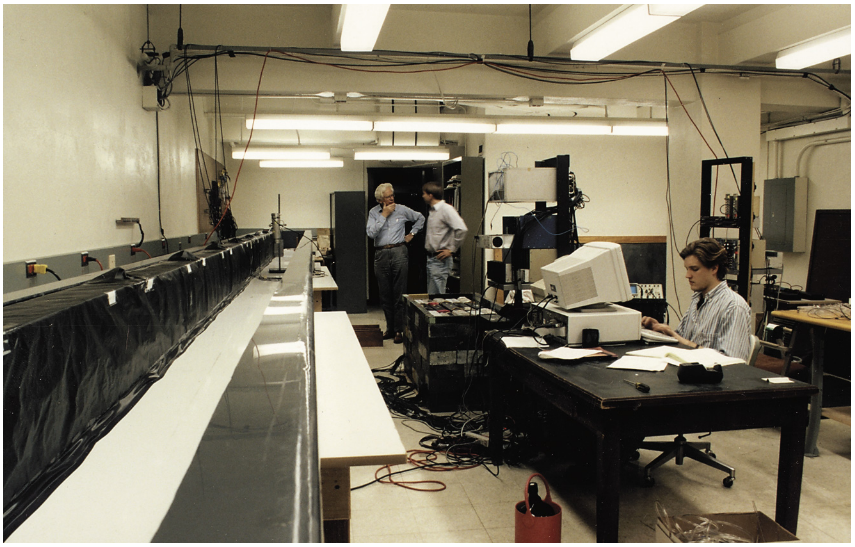

All the detector components were tested individually at Caltech before being shipped to Arizona. Here, Brian Cook (at computer), tests the light-transmission properties of a nine-meter-long scintillator cell (at left), swaddled in black plastic to keep the room lights out. Felix Boehm (left) and Andreas Piepke (right) confer in the background.

Detecting something that has built a career out of not interacting with matter in any form is, shall we say, a bit of a challenge.

energy distribution that was predicted by Super Kamiokande, because neutrinos with different energies should disappear at different rates.

If the electron neutrino is changing flavors, part of the trick to watching it disappear is to position the detector in the trough of its probability wave. The smaller the mass difference, the longer the oscillation's wavelength is going to be. On the other hand, you don't want to put the detector much beyond the wavelength, because the total neutrino flux falls off with the square of the distance from the source. The farther away you get, the bigger and more expensive your detector has to be. Early results from Kamiokande had shown a slight surplus of electron neutrinos as well as the famous muon-neutrino deficit, suggesting that a small percentage of muon neutrinos were going to electron neutrinos and that a site one kilometer or so from the neutrino source would be a good distance from which to watch this happen. Therefore, the Palo Verde detector was built 890 meters from Reactors 1 and 3, and 750 meters from Reactor 2. Says Petr Vogel, senior research associate in physics and the house theorist for the project, "This is 10 times farther away from the reactor than any previous such experiment, which means that our flux is 100 times less. We have to really push the detector technology in order to see any neutrinos at all."

Boehm and Vogel have been chasing neutrinos for 20 years, starting with an experiment in Grenoble, France, in 1979. (In fact, they wrote the book on massive-neutrino physics.) The field was launched that year at Caltech by Murray Gell-Mann, then Millikan Professor of Theoretical Physics, and postdocs Harald Fritzsch and Peter Minkowski, who did the first calculations of neutrino oscillations. Says Vogel, "Mixing is the hottest issue in particle physics today. Since 1992, four or five other experiments have confirmed that the muon deficit exists. Nobody doubts that neutrinos have mass any more, so the question now is what the mass is and what the mixing angle is. That will be the program for the next decade, to explore this parameter space."

The Palo Verde project is about five years old. It took three years for grad students Brian Cook (MS '93, PhD '96), now at JPL, and Mark Chen (PhD '94); Humboldt Fellows Ralf Hertenberger and Andreas Piepke; postdocs Nick Mascarenhas and Vladimir Novikov; and staff engineer John Hanson to design, develop, and test the detector elements, while member of the professional staff Herb Henrikson, who got his BS in mechanical engineering at Caltech in 1953 and has been a project engineer here ever since, did the nuts-and-bolts design. At the same time, Boehm had to find a site for the project, line up money and collaborators, bid out the construction contracts, and so forth. A year's worth of ground was lost to a competing experiment, subsidized by the French nuclear-power industry, when the initial plan to use the San Onofre reactor, about an hour's drive south of Caltech, fell through—endangered gnatcatchers were nesting on the proposed excavation site. Assembling the detector apparatus and building the underground chamber that houses it took another year, followed by a six-month shakedown period. The detector has been fully operational and taking data since October 1998 under postdoc K. B. Lee and colleagues from Caltech, Stanford, and the University of Alabama.

Detecting something that has built a career out of not interacting with matter in any form is, shall we say, a bit of a challenge. You have to rely on indirect evidence: in this case, the flashes of light produced when a neutrino hits a proton, creating a positron (or anti-electron) and a neutron—as mentioned earlier, the neutron-decay reaction run backward. To maximize the collision rate, the detector contains 12 tons of proton-rich mineral oil, whose average molecular formula is $C_{22}H_{46}$. The oil is heavily laced with pseudocumene, a benzene derivative that has half a dozen easily

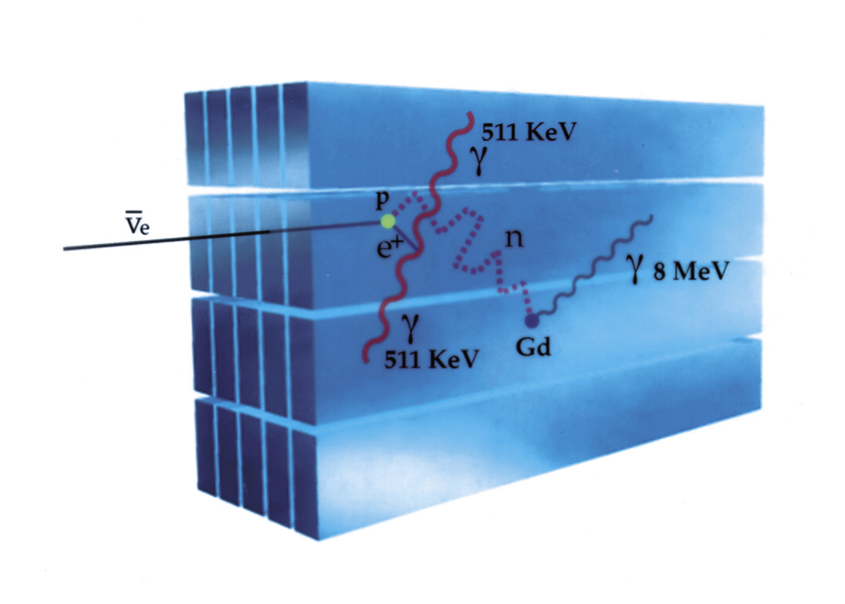

How to catch a neutrino: This schematic shows a portion of the detector array five cells wide by four cells high. An electron antineutrino ($\bar{\nu}_e$) from the reactor hits a proton (p), creating a positron (e^+) and a neutron (n). The positron quickly annihilates an electron (not shown), creating a pair of gamma rays (γ). (KeV stands for thousand electron volts.) The neutron wanders off and is eventually sucked up by a gadolinium atom (Gd), which emits another gamma ray—a whole slew of them, in fact. (For clarity, only one gamma ray is shown.)

"Ah. Triple pulse with the right energy distribution. That's a keeper."

excitable electrons per molecule. The positron jangles these electrons as it screams by with an average kinetic energy of three million electron volts (MeV). In a process called scintillation, the excited electrons emit flashes of blue light that are recorded by photomultipliers—light detectors capable of sensing a single photon—and the energy measurement of each flash is sent to a computer. The positron travels about two centimeters, losing energy with every electron it twangs. But to slow down is to die—eventually (within about 30 billionths of a second, that is) it no longer has enough zip to get by its mortal enemy. The last electron it runs into annihilates it, producing two gamma rays at 0.5 MeV, which is the energy equivalent of the mass of an electron or positron. These gamma rays also jangle the pseudocumene's electrons, causing two more pulses of light.

In order to chart the particles' paths, the scintillator oil is parceled out into 66 cells—acrylic-walled rectangles nine meters long by 12 centimeters wide and 25 centimeters high, with a photomultiplier on each end. The cells are wrapped in copper foil so that a flash in one won't trip the photomultiplier in a neighbor. But the gamma rays normally fly through several cells before petering out, so the computer continuously digitizes the arrival time and energy of all the flashes picked up by all the photomultipliers in the array, and scans the lot for "coincidences"—signals from photomultipliers in blocks of up to 15 adjoining cells at the same time—and says, "Ah. Triple pulse with the right energy distribution. That's a keeper."

Meanwhile, the neutron plows through the oil, gradually losing steam until it gets absorbed by an atom of gadolinium, which soaks up neutrons like a sponge. (Persuading gadolinium to dissolve in mineral oil is no small feat—like most metals, it's soluble in acids, but uninterested in oil. Some pretty harsh things used to have to be done to the gadolinium to get it into solution, and the result was a dark, nasty liquid that blotted out all light passing through it within half a meter or so. The solution also went bad in just a few months, meaning that the detectors were constantly in the shop for an oil change. So Piepke and Novikov, in collaboration with Bicron, a leading manufacturer of radiation detectors, developed a new recipe for dissolving gadolinium that results in a fluid as clear as water that remains stable for at least two or three years. Bicron now sells the stuff, which has become the industry standard.) Upon catching a neutron, the gadolinium atom emits a fresh cascade of gamma rays at energies of up to 8 MeV. Because these gamma rays are so hopped up, the computer looks for them in coincidences of up to 35 cells at once. A couple of hundred microseconds (millionths of a second) separates the positron's demise and the neutron's capture, and the three-one flash pattern with its set of characteristic energies and delay times is the unmistakable fingerprint of a neutrino.

But lifting that print is not trivial. A bazillion other particles are also zipping through the detector, and they all leave their mark. Says Boehm, "Our detector registers 20 neutrino interactions a day, but we get about 2,000 hits per second from the cosmic-ray flux, plus other background radiation, so it's a very difficult experimental problem. We have to use lots of clever tricks." The Super Kamiokande detector is buried a kilometer deep in a zinc mine to screen out as much background radiation as possible. "Unfortunately," says Boehm, "the Arizona desert has no commercial mineral deposits, so there are no deep mines." Instead, the Palo Verde Neutrino Detector is buried about 25 meters (82 feet) deep—as far down as Caltech could afford to dig. In lieu of a kilometer of rock, the scintillator cells

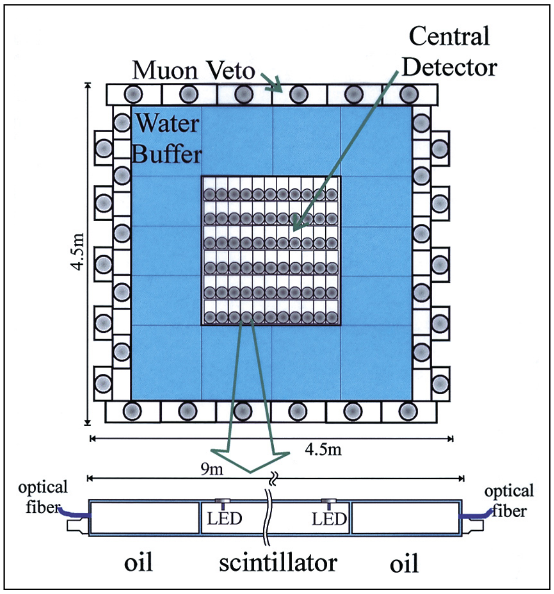

Above: An end-on view of the detector array (main drawing). The gray circles are the photomultipliers. The muon veto detectors are plastic boxes filled with a liquid scintillator sans gadolinium—muons make gamma rays just like positrons do. There's about an inch of headspace over the liquid in each box, so there are two layers of overlapping boxes along each side wall to ensure complete coverage. At the bottom is a side view of a single scintillator cell. The last half meter on each end of the cell is partitioned off and filled with pure mineral oil to act as an additional buffer. A series of light-emitting diodes (LEDs) down the length of the cell provide standardized pulses of light to calibrate the photomultipliers and, with the optical fibers, keep tabs on the scintillator oil's clarity.

Right: The flat gray boxes are the muon detectors and the canisters are their photomultipliers; the scintillator cells are hidden behind the festoons of red cable.

"Our detector registers 20 neutrino interactions a day, but we get about 2,000 hits per second from the cosmic-ray flux, plus other background radiation, so it's a very difficult experimental problem. We have to use lots of clever tricks."

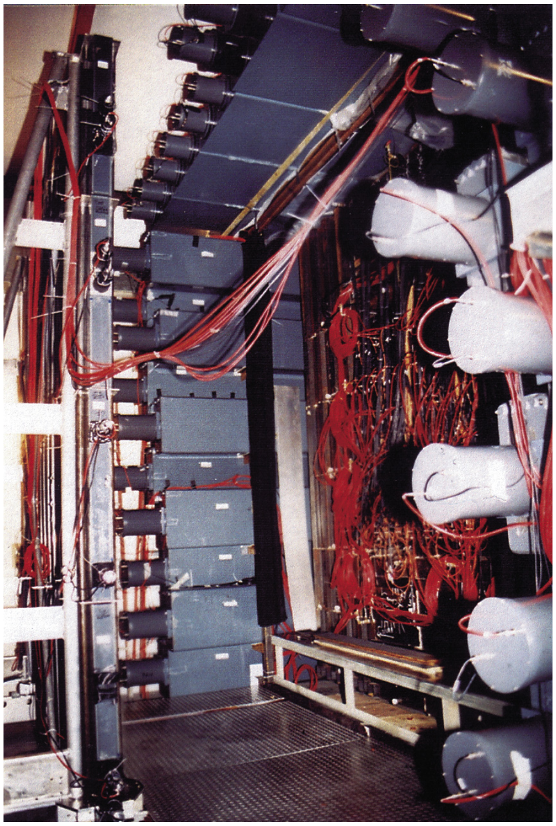

are surrounded by a bank of muon detectors that register cosmic-ray hits. Also called veto detectors, the muon detectors when they go off tell the computer, "Any data you are getting right now is from a cosmic-ray shower. Ignore all inputs for the next 10 microseconds." To help keep costs down, the muon detectors were spares from the MACRO (Monopole, Astrophysics and Cosmic Ray Observatory) project, lent by Linde Professor of Physics Barry Barish and then-Division Chair Charles Peck (PhD '64). Between the veto and neutrino detectors, a 100-ton, one-meter-thick wall of water absorbs neutrons, the other chief byproduct of cosmic rays. And, finally, the computer filters the data through screening programs that reject flashes that aren't energetic enough or otherwise don't look promising.

And there are other subtleties. Samples of all the construction materials had to be vetted by an exquisitely sensitive radiation detector in the subbasement of Caltech's Bridge Laboratory. Ordinary rebar contains a trace amount of radioactive cobalt 60, added to help monitor the production process. "This cobalt 60 is weak for all practical purposes," says Boehm. "But not for neutrino detectors! We had to request special batches of low-cobalt steel to be shipped to us. And concrete is always slightly radioactive, and there's nothing you can do about it. Normally, it contains about one part per million of uranium and thorium, which would have emitted enough gamma rays to choke our detector." These trace elements are found naturally in Earth's crust, so that when you crush rock into gravel, or quarry limestone for cement, they come along. In fact, the product from the local gravel plant was particularly bad. The rock was volcanic, so that it had lots of heavy elements from Earth's interior, and relatively young, so that the hot stuff hadn't had much time to decay.

"The USGS helped us find a marble deposit near

Left: The detector chamber awaits reburial. The access road will soon become a tunnel.
Above: To save money, the tunnel was made from the eight-foot-diameter corrugated steel sewer pipe normally seen on highway drainage projects.

Phoenix that was 10 times radio-cleaner than the local stone. It was trial and error: they sent us lots and lots of rock samples, and we tested them here at Caltech." The Phoenix gravel plant crushed this marble for them, adding about 6 percent to the construction costs. Marble is a soft rock—that's why sculptors use it—and it had never been used in concrete before, so nobody had any idea whether the stuff would be strong enough to support the thousands of tons of dirt that was going to be backfilled onto the roof of the detector chamber. A lot of special testing had to be done at the cement plant before the first batch was trucked to the site. These two factors make the Palo Verde Marble Mix probably the most expensive concrete ever poured—with the possible exception of the night that Jimmy Hoffa disappeared—and certainly the fanciest.

Even with all these precautions, the computer records some 600 megabytes' worth of flashes per day. The data is stored on hard drives at the site, and gets shipped once a day over a fast, dedicated Internet connection to computers at Caltech and Stanford that tease out the fingerprints of the 20 neutrinos a day the collaboration is hoping not to see. These computers also reconstruct the neutrinos' trajectories and energy distributions.

The analysis of the first 70 days' worth of data is now complete. The results are bad news for the neutrinos-as-dark-matter folks, says Vogel. "To be blunt, we do not see oscillations, so the mixing angle cannot be large. And we have moved the mass parameter by a factor of 10 toward smaller masses." These results have almost completely closed the window in parameter space that Super Kamiokande had allowed for electron neutrinos. Thus, it appears that muon neutrinos may mix, but electron neutrinos don't—at least, not to within Palo Verde's detection limits.

But electron neutrinos must mix, because of another long-standing conundrum called the solar neutrino problem. For decades, people have been

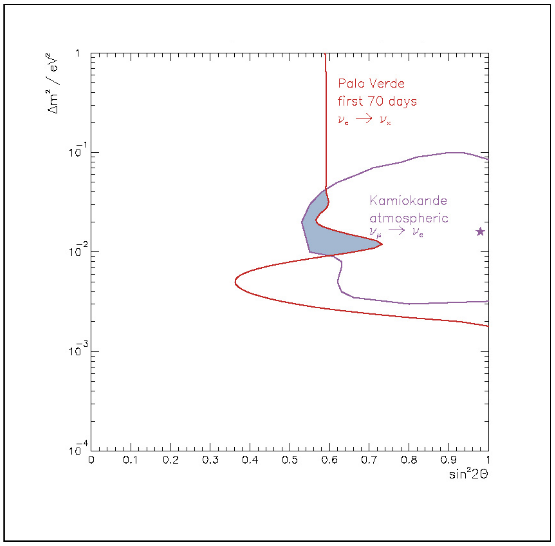

The purple star marks the preferred mixing angle and mass difference derived from Super Kamiokande's results, within the area of experimental uncertainty enclosed by the purple line. However, the Palo Verde results exclude the area lying above and to the right of the red line, so only the blue area remained to be explored. (As *E&S* went to press, the accumulation of additional data had excluded this area as well.) The mixing angle is plotted horizontally, with 0 being no mixing, and 1 being the maximum possible mixing, i.e., a 45° mixing angle. The mass difference is plotted vertically in logarithmic units of electron volts squared.

measuring the electron-neutrino flux from the most powerful nuclear reactor in our neighborhood—our friend, Mr. Sun. These measurements are only coming up with about half as many neutrinos as the solar physicists say should be produced. Either we don't understand the nuclear reactions going on inside the sun as well as we think we do, which is highly unlikely, or else electron neutrinos are disappearing en route to Earth. With a flight path of 93 million miles, even a very tiny mass and minuscule mixing angle would show an effect.

The Palo Verde collaboration will continue to run the experiment through the end of 1999 in order to refine the statistical accuracy of their numbers tenfold—down to the residual uncertainty left in the calculations of the reactor's flux and detector's efficiency. Then the detector gets dismantled. "It's expensive to run," says Boehm. "We have to pay rent to the utility. We have to keep somebody on site to maintain the complex electronics, do all the calibrations, change computer disks, and so on. That person also has to reset all the detectors whenever there's a thunderstorm—we get power outages all the time." You'd think that, being on the premises of a nuclear power plant, they'd have an uninterruptible source of electricity, but no—all their amps come by wire from Phoenix.

The next step, says Vogel, is to explore longer wavelengths. The Caltech group is collaborating on a proposal to build a new detector, called Kamland, down in that Japanese zinc mine. The mine is located near the city of Kamioka, which lies some 40 kilometers north of Osaka near the center of the main island of Honshu. Japan gets about one-third of its electricity from nuclear power, and Kamland will use the 16 nuclear plants on the island as its neutrino source. (If calibrations with three reactors at Palo Verde were tough, calibrating this detector is going to be a real bear!) The plants lie from 100 to 300 kilometers away from the detector, which will contain 1,000 tons of scintillator oil. But even with a detector that size, the collaboration expects to see only about a thousand neutrinos a year, because of the distances involved. Still, this very long baseline will make the experiment sensitive to mass differences 1,000 times smaller than either Super Kamiokande or Palo Verde could see.

The Palo Verde project has been very fruitful, says Boehm. "We have clearly shown that, unlike atmospheric (muon) neutrinos, reactor (electron) neutrinos do not oscillate at these wavelengths. We explored a promising set of wavelengths, and answered a challenging question in neutrino physics while advancing the state of the art in scintillator technology." And they did so for a bargain-basement price: the whole shebang only cost about $2.5 million to build, which is peanuts as particle physics goes. Although the Department of Energy and the collaborating institutions have helped finance the project, Caltech put up a substantial contribution out of the provost's discretionary funds, says Boehm. "Both Jennings and Koonin felt this was an important opportunity, and have been very supportive. We certainly appreciate it." ☐