

Results from the Palo Verde neutrino oscillation experiment

F. Boehm,³ J. Busenitz,¹ B. Cook,³ G. Gratta,⁴ H. Henrikson,³ J. Kornis,¹ D. Lawrence,² K. B. Lee,³ K. McKinny,¹ L. Miller,⁴ V. Novikov,³ A. Piepke,^{1,3} B. Ritchie,² D. Tracy,⁴ P. Vogel,³ Y-F. Wang,⁴ and J. Wolf¹

¹*Department of Physics and Astronomy, University of Alabama, Tuscaloosa, Alabama 35487*

²*Department of Physics and Astronomy, Arizona State University, Tempe, Arizona 85287*

³*Division of Physics, Mathematics and Astronomy, Caltech, Pasadena, California 91125*

⁴*Physics Department, Stanford University, Stanford, California 94305*

(Received 16 March 2000; published 7 September 2000)

The $\bar{\nu}_e$ flux and spectrum have been measured at a distance of about 800 m from the reactors of the Palo Verde Nuclear Generating Station using a segmented Gd-loaded liquid scintillator detector. Correlated positron-neutron events from the reaction $\bar{\nu}_e p \rightarrow e^+ n$ were recorded for a period of 200 d including 55 d with one of the three reactors off for refueling. Backgrounds were accounted for by making use of the reactor-on and reactor-off cycles, and also with a novel technique based on the difference between signal and background under reversal of the e^+ and n portions of the events. A detailed description of the detector calibration, background subtraction, and data analysis is presented here. Results from the experiment show no evidence for neutrino oscillations. $\bar{\nu}_e \rightarrow \bar{\nu}_x$ oscillations were excluded at 90% C.L. for $\Delta m^2 > 1.12 \times 10^{-3} \text{ eV}^2$ for full mixing and $\sin^2 2\theta > 0.21$ for large Δm^2 . These results support the conclusion that the observed atmospheric neutrino oscillations do not involve ν_e .

PACS number(s): 13.15.+g, 14.60.Lm, 14.60.Pq

I. INTRODUCTION

Results of a long baseline study of $\bar{\nu}_e$ oscillations at the Palo Verde Nuclear Generating Station are reported here. The work was motivated by the observation of an anomalous atmospheric neutrino ratio ν_μ/ν_e reported in several independent experiments [1–3] that can be interpreted as ν_μ - ν_e oscillations requiring large mixing. The mass parameter suggested by this anomaly is in the range of $10^{-2} < \Delta m^2 < 10^{-3} \text{ eV}^2$ for two flavor neutrino oscillations.

The quantity Δm^2 , defined as the difference between the square of the masses of the mass eigenstates, and the mixing parameter θ are related to the transition probability P for two-flavor $\nu_a \rightarrow \nu_b$ oscillations (see, for example, [4]) by

$$P_{\text{osc}}(\nu_a \rightarrow \nu_b) = \sin^2 2\theta \sin^2 \left(\frac{1.27 \Delta m^2 L}{E_\nu} \right), \quad (1)$$

where E_ν (MeV) is the neutrino energy, L (m) is the source-detector distance, and Δm^2 is measured in eV^2 .

Exploring Δm^2 down to 10^{-3} eV^2 requires that the quantity L/E_ν (m/MeV) have a value of around 200. For reactor neutrinos ($E_\nu \sim 5 \text{ MeV}$), a baseline of $L \sim 1 \text{ km}$ is adequate. Reactor experiments are generally well suited to study $\bar{\nu}_e$ oscillations at small Δm^2 ; however, they are restricted to the disappearance channel $\bar{\nu}_e \rightarrow \bar{\nu}_x$.

Reactor antineutrinos have been used for oscillation studies with ever increasing Δm^2 sensitivity since 1981 [5,6]. All of the experiments are based on the large cross section inverse beta decay reaction $\bar{\nu}_e p \rightarrow e^+ n$. The correlated signature, a positron followed by a neutron capture, allows significant suppression of backgrounds. As the reactor $\bar{\nu}_e$ yield and spectra are well known [5], a “near detector” is not required. It is, however, important to control well the detector efficiency and backgrounds.

The considerations mentioned have led to the design of the Palo Verde and Chooz [7] experiments, which have similar Δm^2 sensitivities. While both experiments have pursued their goal of exploring the unknown region of small Δm^2 , recent data from Super-Kamiokande [8] favor the $\nu_\mu \rightarrow \nu_x$ oscillation channel over $\nu_\mu \rightarrow \nu_e$. This paper reports in greater detail results presented earlier [9] and describes the detector calibration, background subtraction, and data analysis techniques used to extract results on neutrino oscillations.

II. EXPERIMENT

A. Detector

The Palo Verde Nuclear Generating Station in Arizona, the largest nuclear power plant in the U.S., consists of three identical pressurized water reactors with a total thermal power of 11.63 GW. The detector is located at a distance of 890 m from two of the reactors and 750 m from the third at a shallow underground site. The 32-meter of water-equivalent overburden entirely eliminates any hadronic component of cosmic radiation while reducing the cosmic muon flux to $22 \text{ m}^{-2} \text{ s}^{-1}$. In order to reduce the ambient γ -ray flux in the laboratory all materials in and surrounding the detector were selected for low activity. The laboratory walls were built with an aggregate of crushed marble, selected for its low content of natural radioisotopes. Concentrations of 170, 750, and 560 ppb for ^{40}K , ^{232}Th , and ^{238}U were measured in the concrete, resulting in a tenfold reduction of γ -ray flux when compared with locally available aggregate. A low ^{222}Rn concentration of about 20 Bq/m^3 in the laboratory air was maintained with forced ventilation. Temperature and humidity were controlled to ensure stable detector operation.

The segmented detector, shown in Fig. 1, consists of a 6×11 array of acrylic cells dimensioned at $900 \text{ cm} \times 12.7 \text{ cm} \times 25.4 \text{ cm}$ and filled with a total of 11.34 tons of liquid




FIG. 1. The Palo Verde detector. One of the 66 target cells with PMTs, calibration light emitting diodes (LEDs), and optical fiber flashers is shown lengthwise at the bottom.

scintillator. A 0.8 m long oil buffer at the ends of each cell shields the central detector from radioactivity originating in the photomultiplier tubes (PMTs) and laboratory walls. The cells were made by cutting and bonding large 0.62 cm thick acrylic sheets. The total acrylic mass in the detector is 3.48 tons. Each cell is individually wrapped in 0.13 mm thick Cu foil to ensure light tightness and is viewed by two 5-in. low activity PMTs [10], one at each end, housed in mu-metal boxes. The target cells are suspended on rollers held in place by thin sheet metal hangers. All structural materials were dimensioned as lightly as possible to minimize dead material between cells. Each cell can be individually removed from the mechanical structure for maintenance. The detector is oriented such that the $\bar{\nu}_e$ flux is perpendicular to the long axis of the cells.

The liquid scintillator is composed of 36% pseudocumene, 60% mineral oil, and 4% alcohol, and is loaded with 0.1% Gd by weight. This formulation was chosen to yield long light transmission length (11.5 ± 0.1 m at 440 nm), good stability, high light output, and long term compatibility with acrylic. Details of the scintillator development have been published elsewhere [11].

The central volume is surrounded on the sides by a 1 m buffer of high purity de-ionized water (about 105 tons) contained in steel tanks which, together with the oil buffers at the ends of the cells, serve to attenuate gamma radiation from the laboratory walls as well as neutrons produced by cosmic muons passing outside of the detector. The low Z of water minimizes the neutron production by nuclear capture of stopped muons inside the detector and has a high efficiency for neutron thermalization.

The outermost layer of the detector is an active muon veto counter, providing 4π coverage. It consists of 32 12-m-long




FIG. 2. Schematic of the front-end electronics. The charge and timing of PMT pulses are read out by two banks of ADCs and TDCs. The timing information is discriminated with two thresholds to avoid crosstalk and retain dynamic range. Trigger discriminator signals from the two PMTs in each cell are OR'ed and input to the trigger circuit.

PVC tanks (from the MACRO experiment [12]) surrounding the detector longitudinally, and two endcaps. The endcaps are mounted on a rail system to allow access to the central detector. The horizontal tanks are read out by two 5-in. PMTs at each end; the vertical tanks are equipped with one 8-in. PMT at each end while the endcaps use 3-in. PMTs. The liquid scintillator used in the veto is a mixture of 2% pseudocumene and 98% mineral oil, with a light attenuation length at 440 nm in excess of 12 m.

A schematic of the central detector's front-end electronics is shown in Fig. 2. Each channel can be digitized by either of two identical banks of electronics. The dual bank system allows both parts of the sequential inverse beta decay event to be recorded with no deadtime by switching between banks. Because of the large dynamic range of energy in the data of interest (40 keV to 10 MeV or 1–250 photoelectrons typically), each PMT has both a dynode and anode output connected to analog to digital converters (ADCs), as well as three discriminator thresholds for the trigger and time to digital converters (TDCs). The higher TDC threshold serves to avoid crosstalk from large signals in adjacent channels while the lower threshold allows timing information to still be available at the single photoelectron level. The relative time of arrival from each end of a cell is used to reconstruct longitudinal position. The measured PMT pulse charge at each end, corrected for light attenuation based on the distance traveled in the cell, allows energy reconstruction.

Each cell is connected to the trigger via the OR of the discriminated signals from the two PMTs. Signals are tagged according to two thresholds: a *high* threshold corresponding to ~ 600 keV for energy deposits in the middle of the cell and a *low* threshold corresponding to ~ 40 keV, or one photoelectron at the PMT. The *low* trigger threshold also serves as the lower TDC threshold. The trigger, which has a decision time for each event of around 40 ns, uses a field programmable gate array to search for patterns of energy deposits in the central detector, and can be reprogrammed easily to

change trigger conditions as needed for calibrations [13].

A veto signal disables the central detector trigger for 10 μ s following the passage of a muon to avoid most related activity. Typical veto rates are ~ 2 kHz. With each event, the time and hit pattern of the previous muon in the veto counter is recorded along with information as to whether or not the muon passed through the target cells. The veto inefficiency was measured to be $(4 \pm 1)\%$ for stopping muons (one hit missed) and $(0.07 \pm 0.02)\%$ for through-going muons (two hits missed). We note that the small size of this second quantity with respect to the first is due to correlations between incoming and outgoing muons as confirmed by a simple Monte Carlo model.

B. $\bar{\nu}_e$ signal

The $\bar{\nu}_e$ signal is detected via the reaction $\bar{\nu}_e p \rightarrow n e^+$ as illustrated in Fig. 17 further below along with the dominant backgrounds. Signal events consist of a pair of time-correlated subevents: (1) the positron kinetic energy ionization and two annihilation γ 's forming the prompt part and (2) the subsequent capture of the thermalized neutron on Gd forming the delayed part. By loading the scintillator with 0.1% Gd, which has a high thermal neutron capture cross section, the neutron capture time is reduced to ~ 27 μ s from ~ 170 μ s for the unloaded scintillator. Furthermore, Gd de-excites by releasing an 8 MeV γ cascade, whose summed energy gives a robust event tag well above natural radioactivity. In contrast, neutron capture on protons releases only a single 2.2 MeV γ .

Background is rejected at trigger level using the detector segmentation by looking for coincidences of energy deposits matching the pattern of inverse beta decay. Each of the subevents of a $\bar{\nu}_e$ signal is triggered by scanning the detector for a pattern of three simultaneous hits in any 3×5 subset of the cell array. This threefold coincidence, called a *triple*, must consist of at least one *high* trigger hit, due to either the positron ionization or neutron capture cascade core, and at least two additional *low* trigger hits, resulting from either positron annihilation γ 's or neutron capture shower tails. The use of identical trigger requirements for the two *triples* is found to give rise to close to an optimal signal to noise ratio. Five μ s after finding an initial *triple*, the trigger begins searching for a delayed *triple*. The blank time suppresses possible false signals from PMT afterpulsing. If two *triples* are found within 450 μ s of each other, the candidate $\bar{\nu}_e$ event is digitized for offline analysis.

C. Expected $\bar{\nu}_e$ interaction rate

In order to calculate the expected $\bar{\nu}_e$ interaction rate in the detector, the status of the three reactors is tracked daily, and the fission rates in the cores are calculated based on a simulation code provided by the manufacturer of the reactors. This code uses as input the power level of the reactors, various parameters measured in the primary cooling loop, and the original composition of the core fuel elements.

The output of the core simulation has been checked by measuring isotopic abundances in expended fuel elements in




FIG. 3. $\bar{\nu}_e$ flux, inverse beta decay cross section, and $\bar{\nu}_e$ interaction spectrum at the detector.

the core; errors in fuel exposure and isotopic abundances are estimated to cause $< 0.3\%$ uncertainty in the $\bar{\nu}_e$ flux estimate. Of the four isotopes— ^{239}Pu , ^{241}Pu , ^{235}U , and ^{238}U —whose fissions produce virtually all of the thermal power as well as neutrinos, measurements of the neutrino yield per fission and energy spectra exist for the first three [14,15]. The ^{238}U yield, which contributes 11% to the final $\bar{\nu}_e$ rate, is calculated from theory [16]. When the same theoretical method was used to calculate the spectrum from the other three isotopes, the theory agreed with experimental results within 10%. The contribution of ^{238}U fission to the overall uncertainty in $\bar{\nu}_e$ rate is therefore expected to be $\sim 1\%$.

This calculated $\bar{\nu}_e$ flux is then used to compute the expected rate of $\bar{\nu}_e$ candidates $N_{\bar{\nu}_e}^-$ at the detector as a function of the oscillation parameters Δm^2 and $\sin^2 2\theta$:

$$N_{\bar{\nu}_e}^- = n_p \int dE_{\bar{\nu}_e}^- \sigma(E_{\bar{\nu}_e}^-) \eta(E_{\bar{\nu}_e}^-) \times \sum_{i=1}^3 \frac{\mathcal{I}_{\bar{\nu}_e, i}^-(E_{\bar{\nu}_e}^-) [1 - P_{\text{osc}, i}(\Delta m^2, \sin^2 2\theta, L_i, E_{\bar{\nu}_e}^-)]}{4\pi L_i^2} \quad (2)$$

where $\sigma(E_{\bar{\nu}_e}^-)$ is the inverse beta decay cross section [17], $\eta(E_{\bar{\nu}_e}^-)$ is the (energy dependent) detector efficiency, n_p is the number of target free protons, and $\mathcal{I}_{\bar{\nu}_e, i}^-$ is the source strength of reactor i at distance L_i with oscillation probability $P_{\text{osc}, i}$. In Fig. 3 we show the energy spectrum of the $\bar{\nu}_e$'s emitted by a reactor, the $\bar{\nu}_e$ (energy) differential cross section in the detector and the actual interaction rate in the de-




FIG. 4. The calculated $\bar{\nu}_e$ interaction rate in the detector target. The two long periods of reduced flux from reactor refuelings were used for background subtraction. The decreasing rate during the full power operation is a result of the changing core composition as the reactor fuel is burned.

detector target before detector efficiency corrections, referred to here as $R_{\bar{\nu}_e}^-$ [obtained by setting $\eta(E_{\bar{\nu}_e}^-)=1$]. The energy spectrum actually measured in the detector is the energy of the positron created by the inverse beta decay. This spectrum is approximately $E_{\bar{\nu}_e}^- - 1.8$ MeV, slightly modified by the kinetic energy carried away by the neutron (10–50 keV).

Previous short baseline experiments which measured the rate of $\bar{\nu}_e$ emission by reactors have found good agreement between calculated and observed neutrino flux by using largely the same method of calculation. A high statistics measurement at Bugey [6], in particular, found excellent agreement both in spectral shape ($\chi^2/N_{\text{DF}}=9.23/11$) and in absolute neutrino yield (agreement better than 3%, dominated by systematic errors). These previous generation experiments prove that the reactor antineutrino spectrum, i.e., the $\bar{\nu}_e$ flux at the distance $L=0$, is well understood.

The expected $\bar{\nu}_e$ interaction rate in the whole target, both scintillator and the acrylic cells, is plotted in Fig. 4 for the case of no oscillations from July 1998 to October 1999. Around 220 interactions per day are expected with all three units at full power. The periods of sharply reduced rate occurred when one of the three reactors was off for refueling, the more distant reactors each contributing approximately 30% of the rate and the closer reactor the remaining 40%. The short spikes of decreased rate are due to short reactor outages, usually less than a day. The gradual decline in rate between refuelings is caused by fuel burnup, which changes the fuel composition in the core and the relative fission rates of the isotopes, thereby affecting slightly the spectral shape of the emitted $\bar{\nu}_e$ flux.




FIG. 5. Time difference Δt as function of $1/Q$ for near- and far-end LEDs of a cell, illustrating the *time-walk* effect. For values of Q close to the discriminator threshold the *time-walk* correction is kept at a constant (maximum) level.

III. CALIBRATION

In order to maintain constant data quality during running, a program of continuous calibration and monitoring of all central detector cells is followed. Blue LEDs installed inside each cell are used for relative timing and position calibration. Optical fibers at the end of each cell, also illuminated by blue LEDs, provide information about PMT linearity and short term gain changes. LED and fiberoptic scans are performed once a week. Radioactive sources are used to map the light attenuation in each cell, for absolute energy calibration, and to determine detection efficiencies for positrons and neutrons. A complete source scan is undertaken every 2–3 months.

A. LED and optical fiber calibrations

As seen in Fig. 1, every cell of the central detector has two LEDs, one at each end at a distance of 90 cm from the PMTs. These blue LEDs, which provide fast light pulses with a rise time comparable to scintillation light, are used for timing calibrations needed for position reconstruction along the cell's axis.

The difference in pulse arrival time between the two PMTs of a cell Δt is described as a function of the position z with an effective speed of light c_{eff} , an offset z_0 and a small nonlinear correction $f(1/Q)$:

$$\Delta t = (z - z_0)/c_{\text{eff}} + f(1/Q_n, 1/Q_f). \quad (3)$$

The correction $f(1/Q_n, 1/Q_f)$, a function of both near and far PMT pulse charge Q , describes the dependence of the pulse height due to a *time walk* in the leading edge discriminators used in the front-end electronics. To extract these calibration parameters and compensate for the *time-walk* effect, a third order polynomial is fit to Δt versus $1/Q$ (see Fig. 5). The

intercepts at $1/Q=0$ for the two LED positions provide c_{eff} and z_0 , while the slopes are used to parametrize the *time-walk* correction.

In order to check the suitability of longer wavelength 470 nm LED light to measure timing properties of ~ 425 nm scintillation light, data taken with a ^{228}Th source at several longitudinal positions were reconstructed with the LED timing calibration parameters. Comparing the reconstructed positions with the actual source positions, the effective speed of light measured with the LED system was found to be on average 3.6% lower than that with the sources. A simulation of the light transport in a cell with various indices of refraction and attenuation lengths of the scintillator suggested that the small discrepancy in Δt between LED and scintillation light was due to the difference in attenuation length. The correction factor was found to be constant over several months. Weekly LED scans are therefore used to correct for short term variations in Δt and a constant correction factor is applied to the effective speed of light.

The fiberoptic system includes 15 blue LEDs, each illuminating a bundle of 12 fibers. The light output of each LED is measured in two independent reference cells with PMTs checked to be linear over the whole dynamic range of the LEDs. By taking a run which scans through all light intensities and mapping each PMT's response relative to the reference cells, the nonlinear energy response of the PMTs is calibrated. Low intensities are used to determine the single photoelectron gain of each PMT, which is used to correct for changes from the nominal gain setting of 4×10^7 .

B. Scintillator transparency and energy scale calibration

In addition to weekly LED and fiberoptic calibrations, the energy response of the scintillator is measured every 3 months using a set of sealed radioactive sources. Eighteen 2.4 mm diameter tubes run along the length of the detector, allowing insertion of the sources adjacent to any cell at any longitudinal position. The response of each PMT as a function of longitudinal position is measured by recording the Compton spectrum from the 2.614 MeV γ of a ^{228}Th source at seven different locations along each cell.

Monte Carlo simulation found that the half maximum of a Gaussian function fitted to the Compton spectrum is relatively independent of resolution; this point is therefore used as the benchmark of the cell response. The response versus distance from the PMT, shown in Fig. 6 for one cell, is then fit to the phenomenological function $\exp(p_0 + p_1 z) + \exp(p_2 + p_3 z)/z$, where z is source longitudinal distance from the PMT. The effective attenuation length of the scintillator (including multiple total reflection on the acrylic walls) is generally between 3 and 4 m and over a year was found to change on average ~ 1 mm/day, demonstrating that the Gd scintillator was remarkably stable.

The overall energy scale was determined from the position of the 1.275 MeV peak of a ^{22}Na source, and then verified by taking data with several γ sources in different energy ranges, ^{137}Cs (0.662 MeV), ^{65}Zn (1.351 MeV), ^{228}Th (2.614 MeV), and the capture of neutrons (8 MeV) from an Am-Be source. The gamma cascade from neutron capture was mod-




FIG. 6. Effect of aging on Gd loaded scintillator. Top: ^{228}Th Compton edge position at seven different longitudinal locations along a typical cell. The two curves are from calibrations taken a year apart. The curves are normalized at the location nearest to the PMT. Bottom: Effective attenuation lengths for all 66 cells from the two calibrations. The quantity μ denotes the averaged effective attenuation length AL_{eff} .

eled according to measurements of the emitted spectrum [18]. In contrast to homogeneous detectors which measure total absorption energy peaks, 25% of the detector target mass consists of the inert acrylic of the cell walls, which absorbs some energy. The Monte Carlo simulation was therefore used to find the correct final distributions of energy detected from single and multiple scattering of the γ 's. The total energy reconstructed for data and the Monte Carlo simulation for each source is plotted in Fig. 7. The data were matched with a Monte Carlo simulation for the ^{22}Na spectrum in Fig. 12, below, to find the overall energy scale and to the spectra in Fig. 7 to assure that the scintillator response is linear over the energies of interest. The light yield after PMT quantum efficiency was found to be ~ 50 photoelectrons per MeV in the center of the cells. The agreement for three of the four sources in Fig. 7 is good, the exception being ^{228}Th , in which the data have a consistently higher Compton scattering peak than the Monte Carlo simulation predicts. This discrepancy is consistent across all the data taken and therefore does not affect the scintillator transparency calibration.

C. Monte Carlo simulation

The ν_e efficiency of the detector is a relatively strong function of event location in the detector and, to a lesser extent, of time due to scintillator aging. A further complication comes from the trigger efficiency being a function of threshold (voltage) while only energy (charge) is measured. For this reason a Monte Carlo model which included a detailed simulation of the detector response, including the PMT pulse shape, is used for an estimate of the overall efficiency




FIG. 7. The total energy reconstructed for various sources compared for data and Monte Carlo simulation.

for $\bar{\nu}_e$ detection. A variety of measurements was performed to cross-check that the Monte Carlo accurately models the detector response.

The physics simulation program is based on GEANT 3.21 [19]. This code contains the whole detector geometry and simulated the energy, time, and position of energy deposits in the detector. Hadronic interactions are simulated by GFLUKA [20] and the low energy neutron transport by GCALE [21]. Scintillator light quenching, parametrized as a function of ionization density, is included in the simulation [22].

The event reconstruction program reads the output of this physics simulation and then applies the second step of the Monte Carlo method, the simulation of the detector response as PMT pulses which are then converted into time and amplitude digitizations and trigger hits. A logical scheme of this detailed detector simulation is shown in Fig. 8.

The calibrations discussed above empirically provide the scintillator light yield (photoelectron/MeV) and attenuation function for each cell, which in turn provide the number of photoelectrons, $\bar{N}_{pe}$, expected for a simulated energy deposit. The total charge of the pulse then follows from sampling a Poisson distribution with mean $\bar{N}_{pe}$ and folding the number of simulated photoelectrons, N_{pe} , with the PMT's nominal gain, first stage gain variance ($10N_{pe}$), and cell-to-cell energy scale calibration uncertainty (10%).

To simulate the pulse shape, an arrival time is assigned to each photoelectron, and individual photoelectron pulses (whose shape is derived from real data) are summed into a final pulse. The calculated arrival time of each photon is a combination of two processes, scintillator de-excitation and propagation along the cell. The latter distribution is parametrized by the distance traveled to the PMT, larger distances giving larger variances, using a light transport simulation of




FIG. 8. Schematic of the Monte Carlo detector simulation program. In order to convert energy deposited into accurate threshold simulation, calibration information from the detector was used to reconstruct PMT pulses.

2×10^7 photons. The resulting pulse is then analyzed to extract TDC and trigger hits.

The Monte Carlo threshold simulation, position reconstruction, and positron and neutron efficiency predictions were checked using calibration data. The trigger threshold simulation for each cell was compared to data taken with a ^{22}Na β^+ source near the center of each cell. The trigger conditions were loosened for these data, a single *low* hit producing a trigger and the event tagged if a *high* threshold was crossed. By plotting the reconstructed energy for each event versus the efficiency for a *high* trigger tag, an effective *high* trigger threshold in MeV for that location in the cell was determined. The *low* threshold was measured similarly. The Monte Carlo pulse shape parameters were tuned to these data. A typical cell's trigger threshold efficiency as a function of energy is shown in Fig. 9 for both data and the Monte Carlo simulation. The trigger threshold, defined as the energy at 50% efficiency, is also plotted for all 66 cells. On average, the thresholds agree to within 1%.

TDC thresholds were checked by the same algorithm, plotting the threshold hit efficiency versus reconstructed energy. A more direct check of the TDC simulation, however, compares the position reconstruction for data and Monte Carlo simulations. Figure 10 shows the longitudinal position of the third largest energy deposit in each event for a ^{22}Na calibration run, representing the position reconstruction of the energy deposited by one of the two positron annihilation γ 's. Since these energy deposits tend to be small (~ 100




FIG. 9. A comparison of the trigger thresholds from data and Monte Carlo simulation. The data were taken with a ^{22}Na source at the center of each cell. The top portion shows the efficiency of the trigger threshold for a typical cell as a function of energy deposited; the bottom shows the energy of the 50% efficiency threshold for all 66 cells.

keV), some fraction of them have one or both PMT's responses below the TDC low threshold. These events constitute the tails of the distribution in Fig. 10 since only the relative signal amplitude was used for position reconstruction.




FIG. 10. A comparison of the reconstructed Z position of an annihilation gamma energy deposit from a ^{22}Na run, between data and Monte Carlo simulation. The long tails are due to lower energy deposits which do not have a TDC for position reconstruction, in which case relative signal amplitudes are used, resulting in worse resolution. Statistical errors in the plot are negligible.




FIG. 11. Comparison of data and Monte Carlo detection efficiency for ^{22}Na and Am-Be source runs at various locations. Locations at the edges of the detector tend to have lower efficiencies.

tion. The narrower central peak is populated by events with TDC information available. The simulation and data agree well, in both resolution and relative frequency of the two cases.

D. $\bar{\nu}_e$ detection efficiency

The absolute efficiency of the detector for positron annihilations and neutron captures was verified using ^{22}Na and Am-Be sources respectively. The ^{22}Na source emits a 1.275 MeV primary γ which is accompanied 90% of the time by a low energy positron which annihilates in the source capsule. The primary γ can mimic the positron ionization of a low energy $\bar{\nu}_e$ event. This deposit, in conjunction with the positron's annihilation γ 's, closely approximates the positron portion of a $\bar{\nu}_e$ event near the trigger threshold.

In two rounds of data taking, 10 months apart, the ^{22}Na source was inserted into the central detector at 35 locations chosen to provide a sampling of various distances from the PMTs and edges of the fiducial volume. The source activity is known to 1.5%, allowing determination of an absolute efficiency. After applying the offline selections used for $\bar{\nu}_e$ prompt triples and correcting for detector DAQ downtime, the measured absolute efficiency was compared with the Monte Carlo prediction; the results are summarized in the top portion of Fig. 11. Good agreement is seen in the average efficiency over all runs, and run by run agreement was 11%.

The energy spectra predicted by the simulation and measured in the data for the ^{22}Na runs were compared. The total energy seen in all cells and the energy detected in the three most energetic hits are plotted in Fig. 12. The trigger thresholds can be seen in the spectra: the high trigger threshold is the rising edge at around 0.5 MeV in the spectrum of the most energetic hit (E_1), and the low trigger threshold is the




FIG. 12. The Monte Carlo simulation and the data compared for the spectra of total energy and first, second, and third most energetic hit (E_{total} , E_1 , E_2 , and E_3) digitized in the ^{22}Na calibration runs.

rising edge at around 50 keV of the third most energetic hit (E_3).

A similar procedure was used to check the neutron capture detection efficiency. The Am-Be neutron source is attached to one end of a thin (7.5 mm) NaI(Tl) detector, which tagged the 4.4 MeV γ emitted in coincidence with a neutron. The NaI(Tl) tag forces the digitization of the 4.4 MeV γ as the prompt part of an event and opens a 450 μs window for neutron capture; this is the same coincidence window used in the $\bar{\nu}_e$ runs.

All neutron cuts used for the $\bar{\nu}_e$ data selection were applied, and the resulting detection efficiency was corrected for detector deadtime and a small random coincidence background. On average, the Monte Carlo efficiency predictions agree well over the 25 locations tested with a run by run agreement of better than 4%, as shown in the bottom of Fig. 11.

As with the ^{22}Na runs, the energy spectra predicted by the simulation and measured in the data were compared. The total energy seen in all cells and the energy detected in the three most energetic hits are plotted in Fig. 13. Note the small peak in E_{total} at ~ 2 MeV arising from neutrons being captured on hydrogen. The differences in data versus Monte Carlo spectra for ^{22}Na and Am-Be were taken into account in estimating systematic errors.

The Am-Be source emits neutrons with kinetic energies up to 10 MeV, creating proton recoils in the detector scintillator in coincidence with the NaI(Tl) induced trigger. By digitizing any energy deposits seen during the neutron release, the high ionization density of these recoiling protons was used in setting the parameters which control scintillator light quenching in the simulation.




FIG. 13. The Monte Carlo simulation and the data compared for the spectra of total energy and first, second, and third most energetic hit (E_{total} , E_1 , E_2 , and E_3) digitized in the Am-Be calibration runs.

The above cross-checks verify our ability to accurately generate the events, model the detector response, reconstruct the events, and correctly calculate the livetime of the data acquisition (DAQ) system. Taken together these procedures complete the task of estimating our $\bar{\nu}_e$ efficiency.

The Monte Carlo simulation for $\bar{\nu}_e$ events models the expected interactions throughout the entire target, including the acrylic walls of the cells, since there is significant efficiency for inverse beta decay originating in the acrylic. The Monte Carlo simulation yields an average efficiency over the entire detector as a function of $\bar{\nu}_e$ energy. The efficiency from the simulation is folded with the incident $\bar{\nu}_e$ spectrum (which may be distorted by oscillations depending on the hypothesis tested), to get the effective efficiency.

E. Independent reconstruction and Monte Carlo simulation

A parallel and independent event reconstruction and the simulation of the detector response has been developed. This second version follows the same general outline of detector calibration, event reconstruction, and simulation described above, but differs in the algorithms and parametrizations used. Major differences include

(i) The functional form for the scintillator light attenuation is the sum of two exponentials $\exp(p_0 + p_1 z) + \exp(p_2 + p_3 z)$, without z^{-1} in the second term.

(ii) The cell response benchmark is the 70% maximum rather than half maximum of the fitted Compton scattering spectrum.

(iii) A different parametrization is used for the linearity correction of the dynode signals.

(iv) The low threshold parameters are tuned to the TDC hit efficiencies rather than trigger efficiencies as discussed above.




FIG. 14. The ratio of predicted to observed efficiency of ^{22}Na and Am-Be for the two reconstruction codes (1) and (2).

(v) An alternate algorithm for simulating the PMT pulse shape was developed and tuned to observed PMT pulse characteristics.

These differences manifest themselves as slightly different $\bar{\nu}_e$ efficiency predictions and $\bar{\nu}_e$ candidate rates in the data. Tests with radioactive sources have been performed to evaluate the quality of the second data reconstruction. The ^{22}Na and Am-Be efficiency runs shown in Fig. 11 were reconstructed by the second analysis to test its efficiency prediction throughout the detector. The ratio of predicted to observed efficiencies over all the e^+ and neutron runs for the first reconstruction (1) and the second reconstruction (2) are plotted in Fig. 14. While the results presented in the analyses below come from the first reconstruction code described above, the development of a second simulation and event reconstruction offers a useful cross-check of the systematic uncertainties of the results. The differences between the two analyses were used to corroborate the estimate of systematic errors.

IV. $\bar{\nu}_e$ SELECTIONS AND BACKGROUNDS

A. $\bar{\nu}_e$ selection

The trigger rate for time-correlated events (two *triples* occurring within 450 μs) is ~ 1 Hz. Most of those events are random coincidences of two uncorrelated *triple* hits, which occur individually at a rate of ~ 50 Hz, mostly from natural radioactivity. In order to select neutrino events, the following offline cuts are applied:

(i) The energy reconstructed in both prompt and delayed *triples* has at least one hit with $E > 1$ MeV and at least two additional hits with $E > 30$ keV. No single hit was allowed to be greater than 8 MeV.

TABLE I. Summary of efficiency and data rates after trigger, selections, and livetime. The entry “Live $_{\mu\text{veto}}$ ” refers to the deadtime induced by the 150 μs blanking window that accompanies each muon detected in the veto counter.

Cut	Cumulative efficiency		Data rate (d^{-1})	
	1998	1999	1998	1999
Trigger	0.271	0.328	69k	106k
Selection cuts	0.149	0.177	1k	1.2k
Live $_{\mu\text{veto}}$	0.102	0.121	<50; see	
Live $_{\text{DAQ}}$	0.075	0.112	Table V	

(ii) The prompt *triple* is required to resemble a positron, i.e. annihilation γ 's each less than 600 keV, and together less than 1.2 MeV. (This cut is the only one which treats the two *triples* asymmetrically.)

(iii) At least one of the two *triples* in the event has more than 3.5 MeV of reconstructed energy for rejection of γ backgrounds.

(iv) The prompt and delayed portions of the event are correlated in space and time (within 3 columns, 2 rows, 1 m longitudinally, and 200 μs).

(v) The event started at least 150 μs (~ 5 neutron capture times) after the previous veto tagged muon activity.

The trigger and selection efficiencies are summarized in the first two columns of Table I.

In addition to corrections for selection cut efficiency and trigger efficiency, detector livetime is a substantial correction to the number of neutrinos seen and deserves some comment. Deadtime comes from two sources, the DAQ and the muon veto. DAQ livetime is the ratio of the number of *triples* the DAQ was available for digitizing to the total number of *triples* the trigger saw. These numbers are available from trigger scalars. The trigger livetime was measured to be $>99.9\%$. The DAQ livetime varies with the *triple* rate, and for the four data periods was determined to be 73.2%, 74.4%, 92.3%, and 91.8% for 1998 full power, 1998 refueling, 1999 full power, and 1999 refueling, respectively. The higher livetime in 1999 is due to improvements made in the trigger conditions.

The muon deadtime can be further divided into two contributions: 150 μs of deadtime caused by each muon, which at $R_\mu = 1990$ Hz left the detector live 74.2% of the time, and muons which interrupted a neutrino event between the positron and the neutron capture, which estimated from the fit parameters of the Monte Carlo capture time left 92.5% of events uninterrupted. The total uncertainty in the calculation of detector deadtime is less than 1%.

B. Backgrounds

Backgrounds can be separated into two types: correlated and uncorrelated. Uncorrelated background events are due to unrelated *triple* hits which randomly coincided in the time window allowed. Although most of the events collected were random coincidences, almost all of this type of background is removed by requiring at least one subevent to have more




FIG. 15. The time elapsed between the prompt and delayed portion of neutrino candidate events for Monte Carlo simulation and data. The Monte Carlo simulation is fit to two exponentials. The data are fit to three exponentials, constrained to have the time structure of the simulation and an additional contribution due to uncorrelated background (dashed line).

than 3.5 MeV of reconstructed energy. These events do not have a time correlation between prompt and delayed sub-events (*inter-event* time) characteristic of neutron capture. They have instead a longer time correlation determined by the probability that the veto detected no muon between the prompt and delayed random *triples*. At a 2 kHz muon rate, this background is seen as a 500 μs tail under the normal neutron capture distribution. By looking at the *inter-event* times of the candidate $\bar{\nu}_e$ events at longer time scales, this background can be measured.

The *inter-event* time distribution after all neutrino selections (except the time correlation cut) is shown in Fig. 15. The Monte Carlo simulation for a pure neutron capture sample is empirically fitted to the sum of two exponentials. There are two time constants due to the inhomogeneity of the target: neutrons which remain in the scintillator have a 27 μs capture time, whereas those which enter the acrylic have a longer capture time due to the absence of Gd. The data *inter-event* time distribution is fitted to a function of three exponentials with fixed time constants consisting of the Monte Carlo fit τ 's multiplied by a third time constant of 500 μs . Integrating the resulting 500 μs exponential of the uncorrelated background in the signal region gives an estimate of 4.1 ± 0.2 events per day, or 9% of the $\bar{\nu}_e$ candidates being uncorrelated background events.

To measure the uncorrelated background in smaller parts of the data set, the statistical accuracy of the three exponential fit method becomes unacceptably poor. A simpler method is therefore used in conjunction with the above fit. For *inter-event* times longer than 200 μs , the $\bar{\nu}_e$ candidates are dominated by uncorrelated backgrounds. The integrated




FIG. 16. The uncorrelated background rate as measured in approximately 1-month-long segments of the data. Overlaid is a fit to a constant within each year's data; the fit quality is consistent with a stable background.

number of candidates from 200–400 μs is scaled to estimate the number underneath the signal region ($<200 \mu\text{s}$). Using the scaling from the fit of the entire data set shown in Fig. 15, the uncorrelated background was measured in approximately 1-month-long intervals as shown in Fig. 16. For both the 1998 and 1999 data sets, the rates are found to be stable within statistical errors.

Correlated backgrounds have the neutron capture *inter-event* time structure of the $\bar{\nu}_e$ candidates. These events come mainly from cosmic muon induced fast neutrons from spallation or muon capture, as shown schematically in Fig. 17. These fast neutrons can either (1) induce more neutrons via spallation, two of which can be captured in the detector with one capture mimicking a positron signature, or (2) they can cause proton recoil patterns in the central detector which appear as a positron signature and then get captured. Spallation neutrons originate from muons passing through the walls of the laboratory without hitting the veto detector or from muons passing through the detector shielding undetected by the veto. Muon capture neutrons mainly originate from muons stopping in the water buffer without registering in the veto.

To illustrate some properties of correlated background, Fig. 18 shows the time elapsed since the previous veto hit for $\bar{\nu}_e$ candidates, with all selection cuts applied except that of the previous muon timing. This distribution is fit to a three exponential function analogous to that used for the *inter-event* time fits. The two time constants for neutron capture are not identical to those for $\bar{\nu}_e$ events, but tend to be smaller since after passage of a muon there are often more than one neutron in the detector to be captured. The third exponential time constant is again constrained to 500 μs as expected in a random sampling of events unrelated to the previous muon.




FIG. 17. A schematic drawing of the detector showing a signal event and two examples of interactions which contribute to the correlated background. Neutrons which cause correlated backgrounds come mainly from muons undetected by the veto and captured in the water buffer (left), and spallation from muons in the lab walls (right). These neutrons in turn can either induce more neutrons via spallation (as shown at right) or cause three coincident proton recoils as they thermalize (as shown at left).

Since at very short times there are other contributions such as muon decay, times less than $15 \mu\text{s}$ are excluded from the fit. Muon-induced neutron backgrounds dominate the candidates in the first $150 \mu\text{s}$ after the previous tagged muons, motivating the selection cut on μ timing.

In order to show that the correlated background was constant in time, the previous muon time cut was disabled and a plot was made of the $\bar{\nu}_e$ candidate rate versus time, as shown in Fig. 19. When fit to a constant for each year, a χ^2/N_{DF} of $382.7/371$ is obtained, which has a 33% likelihood, indicating that the detector efficiency for correlated backgrounds was stable during each year's data taking.

Aside from the detector efficiency for background, however, a loss of veto efficiency could also cause a fluctuation in background. (The rates fit in Fig. 19 are with the muon timing selection disabled, and hence do not vary with veto inefficiency.) To track veto efficiency, the veto hit patterns recorded with each event are used. If a muon hit was recorded only on the bottom of the veto, where only exiting muons are seen, then the muon must have entered the veto without recording a hit. By measuring the percentage of these events a *one hit missed* veto inefficiency of $(4 \pm 1)\%$ is found as mentioned above. The through-going (*two hits missed*) veto inefficiency is measured to be $(0.07 \pm 0.02\%)$ by looking at the rate of μ tracks triggered on in the central detector. These inefficiencies were tracked in time to assure their stability.

C. Neutron- $\bar{\nu}_e$ direction correlation

The neutrons produced in the inverse beta decays will have momenta slightly biased away from the source, whereas no correlation is expected for background. This effect is the consequence of momentum conservation which requires that the neutron should always be emitted in the forward hemisphere with respect to the incoming $\bar{\nu}_e$. Such a correlation has been observed already in the Gösigen experiment [23] and again at Chooz [24]. The theoretical treatment of the effect can be found in [17].

The signal to background ratio can be independently verified using this effect. The $\bar{\nu}_e$ source is to the left of the detector in Fig. 1. The relative horizontal location (relative column in the target cell array) of neutron capture cascade cores versus positron ionizations for data and the simulation of the $\bar{\nu}_e$ signal are plotted in Fig. 20. Defining the asymmetry $A_{\text{data}} = (R - L)/(R + L)$ in terms of the number of neutrons captured one column away from the source R and one column toward the source L , a slight asymmetry 0.050 ± 0.017 is found in the data, at 2.9σ significance. Using the Monte Carlo simulation which gives $A_{\text{MC}} = 0.134$ to estimate




FIG. 18. The time between the previous muon traversing the veto counter and a neutrino candidate event trigger. The fit is to a sum of three exponentials. The dashed portion denotes events uncorrelated to the previous muon; the rest are correlated background events induced by cosmic muons. The neutron capture time constants, shorter than the simulation prediction for $\bar{\nu}_e$ events, imply multiple neutron production.




FIG. 19. The number of candidates per day in the 373 runs used in this analysis, fit to a constant for each year. All neutrino selections except the previous muon time cut are used, leaving a data set dominated by muon-induced neutron background. The different rates in 1998 and 1999 are due to trigger and livetime improvements.

the portion of the data consisting of $\bar{\nu}_e$ signal and assuming the background to be symmetric in this variable, an effective signal to noise ratio

$$\frac{S}{N} = \frac{A_{\text{data}}}{A_{\text{MC}} - A_{\text{data}}} = 0.6^{+0.4}_{-0.3} \quad (4)$$

is found. This value agrees well with the ratio of 0.81 ± 0.03 found with the *swap* analysis method described below.

V. ANALYSIS

The data set presented here was taken from July 1998 to September 1999 in 373 short runs, each on average about 12 hours long. In 1998, 35.97 days of data were taken with the three reactors at full power and 31.35 days with one of the reactors at a distance of 890 m off for refueling. The detector was then taken offline in January to February of 1999, when DAQ improvements were made to increase livetime, and the *high* trigger thresholds were lowered by 30% to increase trigger efficiency. The 1999 data set includes 110.95 days with all three reactors at full power and 23.40 days with the 750 m baseline reactor off for refueling. Thus, the entire data set has four distinctly different periods, with three different baseline combinations and neutrino fluxes.

After all selection cuts there is still substantial background in the remaining data set. The correlated background, coming mainly from muon induced neutrons, is difficult to predict and subtract. The yield and spectrum of neutron spallation is a function of muon flux and energy, which in turn is a function of depth. While some measurements of fast neu-




FIG. 20. The relative horizontal location of the neutron capture and positron ionization in the detector target. The kinematics of the inverse beta decay causes a slight bias in the momentum of the neutron from inverse beta decay, seen here as an asymmetry in the populations of neutrons captured one column away or towards the $\bar{\nu}_e$ source. Backgrounds should not exhibit such asymmetry. The Monte Carlo prediction for the pure $\bar{\nu}_e$ signal, normalized to unity, is also plotted.

tron spectra and fluxes have been done in the past, there is no model which can consistently predict the fast neutron production. Below we present three methods used to extract the $\bar{\nu}_e$ signal from data.

A. Analysis with the *on-off* method

The conceptually simplest method of subtracting background is to take advantage of periods of reduced power levels of the reactor source. Ideally all three reactors would be down at once allowing for a direct measurement of the background. However, in practice only one of the three Palo Verde reactors was refueled at any given time. These reduced power periods occurred twice annually for about a month. Each year's data set is treated independently, subtracting 1998 *off* from 1998 *on* and 1999 *off* from 1999 *on*, since the efficiency of the detector changed between the two years. By subtracting these data taken at reduced flux from the full flux data, a pure neutrino sample is retrieved, albeit containing the statistical power of only a small portion of the potential data set: the subtraction is limited by the 2 months of refueling time and treats the $\bar{\nu}_e$ flux from the two reactors still at full power as background.

The primary concern arising from use of this method, aside from the loss of statistics, is guaranteeing that the background rates during the *on* and *off* periods were stable. Both correlated and uncorrelated backgrounds were carefully tracked to ensure stability as discussed above.

TABLE II. Results for the simple *on-off* background subtraction analysis, showing candidate rates in the data, N , and efficiency corrected ν_e interaction rates R observed and calculated. The data sets for each year are considered independently here. Uncertainties are statistical only. Systematic errors are estimated to be 10%.

	1998	1999
L (m)	890	750
ON N_{cand} (day^{-1})	38.2 ± 1.0	52.9 ± 0.7
OFF N_{cand} (day^{-1})	32.2 ± 1.0	43.9 ± 1.4
ON-OFF N_{cand} (day^{-1})	6.0 ± 1.4	9.0 ± 1.6
Total efficiency ON (OFF)	0.0746 (0.0772)	0.112 (0.111)
R_{obs} (day^{-1})	95 ± 19	77 ± 14
R_{calc} (day^{-1})	63	88

The numerical results of this analysis of the total rate are summarized in Table II. After correcting for efficiency (for the no-oscillations scenario) and livetime, the data sets were subtracted to find observed neutrino interaction rates in the detector. No significant deviation from the expected neutrino interaction rates was found at either baseline distance.

The results from the alternate reconstruction (2) for this analysis are shown in Table III for comparison. This analysis selects about 5% more candidates, but also gives a correspondingly higher efficiency. For this analysis, the uncorrelated background was measured and removed from the data before the subtraction.

In order to test the results for oscillation hypotheses in the two flavor Δm^2 - $\sin^2 2\theta$ plane, a χ^2 analysis is performed comparing the calculated $R_{\text{calc},ij}$ and observed $R_{\text{obs},ij}$ spectra divided into 1 MeV bins j for each year i . The spectra used are the prompt energies of the two subtracted data sets. At each point in the oscillation parameter plane, taking into account the changes in detector efficiency due to distortions of the neutrino spectrum, the quantity

$$\chi^2 = \sum_{i=1}^2 \sum_j^{E_{\text{bins}}} \frac{(\alpha R_{\text{calc},ij} - R_{\text{obs},ij})^2}{\sigma_{ij}^2} + \frac{(\alpha - 1)^2}{\sigma_{\text{sys}}^2} \quad (5)$$

is computed, where α accounts for possible global normalization effects due to systematic uncertainties (discussed be-

TABLE III. Results from the alternate reconstruction (2) for the simple *on-off* background subtraction analysis. Uncorrelated background was accounted for in the data before subtraction. Uncertainties are statistical only. Systematic errors are estimated to be 10%.

	1998	1999
L (m)	890	750
ON N_{cand} (day^{-1})	37.3 ± 1.2	49.3 ± 0.7
OFF N_{cand} (day^{-1})	31.6 ± 1.2	38.6 ± 1.6
ON-OFF N_{cand} (day^{-1})	5.7 ± 1.7	10.7 ± 1.8
Total efficiency ON (OFF)	0.0809 (0.0838)	0.121 (0.121)
R_{obs} (day^{-1})	85 ± 20	89 ± 15
R_{calc} (day^{-1})	63	88




FIG. 21. 90% C.L. limits of the *on-off* analysis. (a) On-off subtraction with normalization constraint and (b) using only the spectral shape.

low) across both periods and σ_{ij} is the statistical uncertainty in each bin. Systematics which can affect spectral shape, mainly energy scale uncertainty, are negligible relative to the statistical uncertainties in the analysis. The function is minimized with respect to α . The point in the physically allowed parameter space with the smallest chi square χ_{best}^2 was found, which represents the oscillation scenario best fit by the data.

The 90% confidence level (C.L.) acceptance region is defined according to the procedure suggested by Feldman and Cousins [25] by

$$\Delta \chi^2 = \chi^2(\Delta m^2, \sin^2 2\theta) - \chi_{\text{best}}^2 > \Delta \chi_{\text{crit}}^2(\Delta m^2, \sin^2 2\theta) \quad (6)$$

where $\chi^2(\Delta m^2, \sin^2 2\theta)$ is the minimized fit quality at the current point in Δm^2 - $\sin^2 2\theta$ space and $\Delta \chi_{\text{crit}}^2$ is the C.L. χ^2 cutoff. Because of the sinusoidal dependence of the expected rates on the oscillation parameters and the presence of physically allowed boundaries to those parameters, the cutoff is not simply the $\Delta \chi^2$ one would analytically find for a three parameter minimization but has to be calculated for each point in the plane. To find the $\Delta \chi_{\text{crit}}^2$ for a point, the experiment is simulated 10^4 times under the assumption that the oscillation hypothesis represented by that point is true. For each simulated data set, a χ_{best}^2 is extracted and a $\Delta \chi^2$ found for the point. These $10^4 \Delta \chi^2$, the simulations' fit qualities to the hypothesis, are then ordered. The $\Delta \chi^2$ of which 90% of the simulations are a better fit is a 90% C.L. and therefore that oscillation hypothesis' $\Delta \chi_{\text{crit}}^2$.

The region excluded by the analysis is shown in curve (a) of Fig. 21. The results of the χ^2 analysis, including the oscillation parameters' best fit to the data, are summarized in




FIG. 22. Total energy of the prompt subevent ($\sim E_{\nu_e}^- - 1.8$ MeV) in 1 MeV wide energy bins (observed and expected for no oscillation) from the *on-off* analysis, summed over 1998 and 1999. The spectra are in agreement with each other, with a $\chi^2/N_{\text{DF}} = 8.5/8$. Uncertainties are statistical only.

the first column of Table IV further below. For the *on-off* analysis a best fit preferring the no-oscillation hypothesis was found.

In addition to the analysis of the absolute $\bar{\nu}_e$ rates observed, one can analyze the shape of the spectrum of neutrinos seen independently of the absolute normalization, thereby relieving the result of most systematic uncertainties. The χ^2 is calculated at each point in the oscillation parameter plane as in Eq. (5), with no constraint on normalization ($\sigma_{\text{syst}} \rightarrow \infty$). The same procedure as before is followed in defining a 90% C.L. region in the $\sin^2 2\theta - \Delta m^2$ plane. At large Δm^2 where $\bar{\nu}_e$ of all energies are oscillating many times within the baseline, the energy spectrum of the incident flux is affected only in magnitude. As a result, the region excluded in the plane does not extend to large Δm^2 , as shown in Fig. 21(b).

For the spectrum analysis, when the normalization α is left free, the minimum χ^2 is obtained for $\alpha \sim 2$ (see Table IV) and maximum mixing. Such a value of α is “compensated” for by the maximum value assumed by the mixing parameter that reduces the neutrino flux to 1/2 of the non-oscillation value for large Δm^2 . This is clearly due to the fact that this particular fit ignores completely the independent constraint given by the efficiency calibration of the detector. Such calibration fixes α to be 1 within better than 10%. In addition this result has no effect on the exclusion plot in Fig. 21 because, as shown in Table IV, below, the no-oscillation hypothesis has actually better χ^2/N_{DF} than the minimum. Also the exclusion plots based on Eq. (6) and either χ^2_{best} or $\chi^2_{\text{no osc}}$ are found to be virtually identical. Furthermore, changing the bin size from 1 MeV to 0.5 MeV does not appreciably change the exclusion plot, either.

Since the analyses reported above and in the following sections find no evidence for neutrino oscillations, the spectra of the two years are added and the summed spectrum is plotted in Fig. 22 along with the Monte Carlo expectation.

B. Analysis with the reactor power method

Part of the statistical limitations of the direct subtraction of the preceding analysis is a result of the separation of the data set by year. By correcting the four periods for efficiency and then subtracting the respective reduced flux from full flux periods, the subtraction is forced to treat the $\bar{\nu}_e$ flux from two of the reactors as background. A second χ^2 analysis was performed which effectively uses the full $\bar{\nu}_e$ flux of the refueling periods.

To use the 1998 and 1999 data sets together, the change of both signal and background efficiencies are accounted for. The $\bar{\nu}_e$ efficiency difference is found through the detector Monte Carlo simulation. The efficiency change for background from 1998 to 1999, which is not necessarily the same as for $\bar{\nu}_e$, is extracted via the high statistics fits of correlated background shown in Fig. 19. The 27% increase in background efficiency observed roughly corresponds to what the Monte Carlo simulation predicts for a background composed mainly of double neutron captures. Uncorrelated background accounts for less than 10% of the candidate set; this background’s efficiency changed by a similar amount within the measurement statistics seen in Fig. 16.

The combined data sets are analyzed for oscillation hypotheses by calculating the χ^2 summed over runs i (runs with less than ten candidates are combined with adjacent runs):

$$\chi^2 = \sum_i \frac{[(\alpha N_{\text{calc},i} + b) - N_{\text{cand},i}]^2}{\sigma_i^2} + \frac{(\alpha - 1)^2}{\sigma_{\text{syst}}^2} \quad (7)$$

where $N_{\text{cand},i}$ is the total $\bar{\nu}_e$ candidate rate, $N_{\text{calc},i}$ is the calculated rate, α is the overall normalization as before, and b is the background rate. The background, b , is scaled as appropriate for the year but is otherwise assumed to be constant. The function is minimized at each point with respect to b and α . We found no evidence for oscillations and the 90% C.L.

TABLE IV. The best fit to the data for the oscillation parameters of the two flavor model for each of the analyses presented here: *on-off*, *spectral shape*, *reactor power*, and *swap*. The χ^2 of the no-oscillation hypothesis is also tabulated for each analysis. For the *reactor power* analysis all runs with less than $10\bar{\nu}_e$ candidates were combined with adjacent runs, leaving $327N_{\text{DF}}$. In all cases the χ^2 obtained for the no-oscillation hypothesis has essentially the same value as χ^2_{best} . The different analysis techniques are described in the text.

Analysis	On-Off	Spectrum	Power	Swap
$\chi^2_{\text{best}}/N_{\text{DF}}$	17.9/13	16.9/13	317.6/325	0.9/1
α_{best}	1.00	1.99	1.02	1.03
$\chi^2_{\text{no osc}}/N_{\text{DF}}$	17.9/15	17.9/15	317.7/327	0.9/3




FIG. 23. 90% C.L. limits of oscillation parameters for (a) the reactor power analysis, which fits for background based on varying power levels, and (b) the swap analysis, which directly subtracts background from the data set. The Kamiokande $\nu_e - \nu_\mu$ atmospheric neutrino 90% C.L. (dashed line) and best fit (star) are also shown.

plot, shown in Fig. 23, curve (a), is then constructed around the χ^2_{best} as before by comparing $\Delta\chi^2$ with $\Delta\chi^2_{\text{crit}}$ at each point. The predictions of signal and background from this fit for the no-oscillation hypothesis are shown in Table V. The no-oscillation likelihood and best fit results of the χ^2 analysis with this method are summarized in the third column of Table IV.

C. Analysis with the swap method

A third analysis is used which has the potential of using the full statistical power of the neutrino data set by subtracting background directly. The method, discussed in more detail elsewhere [26], takes advantage of the asymmetry of the prompt (positron) and delayed (neutron capture) subevents of the neutrino signal. The data selection and trigger treat the

TABLE V. The rate of candidates N_{cand} , signal N_{detected} , and background b found by the reactor power analysis, using a χ^2 analysis of the data set run by run. The efficiency-corrected total interaction rates R are also listed. Errors are statistical only. The systematic uncertainty is estimated to be 10%.

Events (day ⁻¹)	1998 ON	1998 OFF (890 m)	1999 ON	1999 OFF (750 m)
N_{cand}	38.2 ± 1.0	32.2 ± 1.0	52.9 ± 0.7	43.9 ± 1.4
b	19.5 ± 1.7		26.3 ± 2.2	
N_{detected}	18.7 ± 2.0	12.7 ± 2.0	26.6 ± 2.3	17.6 ± 2.6
R_{obs}	225 ± 24	140 ± 22	216 ± 19	140 ± 21
R_{calc}	218	155	218	130

two portions of the event identically with the exception of two cuts designed to isolate events with annihilation-like γ 's in the prompt *triple*.

The candidates remaining after the selection cuts can be written as

$$N = B_{\text{unc}} + B_{nn} + B_{pn} + S_\nu \quad (8)$$

where B_{unc} , B_{nn} , and B_{pn} are uncorrelated, two-neutron, and proton-recoil–neutron-capture backgrounds respectively, and S_ν is the neutrino signal. Applying the same neutrino cuts with the positron cuts reversed, or *swapped* (such that the positron cuts are now applied to the delayed *triple*), gives

$$N' = B_{\text{unc}} + B_{nn} + \epsilon_1 B_{pn} + \epsilon_2 S_\nu. \quad (9)$$

Since the uncorrelated background and two neutron capture backgrounds are symmetric under exchange of the prompt and delayed *triples*, their efficiencies with the reversed cuts remain the same. The parameters ϵ_1 and ϵ_2 denote the relative efficiency change for proton recoils and neutrino signal under the swap, respectively.

The positron cuts are highly efficient for positron annihilation events but have poor efficiency for neutron captures. The Monte Carlo simulation is used to estimate $\epsilon_2 = 0.159$. Subtracting Eq. (8) from Eq. (9) leaves the majority of the neutrino candidates and only proton recoil background:

$$N - N' = (1 - \epsilon_1)B_{pn} + (1 - \epsilon_2)S_\nu. \quad (10)$$

To estimate $(1 - \epsilon_1)B_{pn}$, it is noted that the proton recoil spectrum extends beyond 10 MeV, well above the positron energies of the neutrino signal and other sources of background. These measured high energy events can be used to normalize the B_{pn} background in the signal using the Monte Carlo ratio:

$$r \equiv \frac{B_{pn}^{\text{MC}}(E_{1,e^+} < 8 \text{ MeV})}{B_{pn}^{\text{MC}}(E_{1,e^+} > 10 \text{ MeV})}, \quad (11)$$

where $B_{pn}^{\text{MC}}(E_{1,e^+} < 8 \text{ MeV})$ is the fraction of simulated B_{pn} events passing the normal $\bar{\nu}_e$ selections, and $B_{pn}^{\text{MC}}(E_{1,e^+} > 10 \text{ MeV})$ are the fraction of simulated events in the high energy background region. Multiplying the ratio r by the measured high energy proton recoil rate gives the B_{pn} background contribution:

$$B_{pn} = r B_{pn}^{\text{data}}(E_{1,e^+} > 10 \text{ MeV}). \quad (12)$$

The neutrons which cause the proton recoil background are created either by muon capture or spallation in the laboratory walls, or by muons entering the veto counter undetected. The spectrum of the fast neutrons from spallation is not well understood. However, such a spectrum can be decoupled somewhat from the resulting proton recoil spectrum. The expected backgrounds were simulated for various possible fast neutron spectra and the resulting ϵ_1 and r for neutrons created in the lab walls were calculated. The same calculation was performed for neutrons created in the passive detector shielding by untagged muons; in this case, the expected yield is much smaller, being only a few percent of

TABLE VI. Results for the *swap* analysis, including the various background estimates. Uncertainties are statistical only.

Period	1998 ON	1998 OFF (890 m)	1999 ON	1999 OFF (750 m)
Time (days)	35.97	31.35	110.95	23.40
$\bar{\nu}_e$ overall efficiency (%)	7.46	7.72	11.2	11.1
$B_{pn}^{\text{data}}(E_{1,e^+} > 10\text{MeV})(\text{day}^{-1})$	8.79	9.09	13.52	13.29
$(1 - \epsilon_1)B_{pn}(\text{day}^{-1})\mu$ spallation	-0.88	-0.91	-1.35	-1.33
$(1 - \epsilon_1)B_{pn}(\text{day}^{-1})\mu$ capture	0.58	0.58	0.86	0.86
$N(\text{day}^{-1})$	38.2 ± 1.0	32.2 ± 1.0	52.9 ± 0.7	43.9 ± 1.4
$N'(\text{day}^{-1})$	24.6 ± 0.8	21.2 ± 0.8	32.3 ± 0.5	31.7 ± 1.2
$N_\nu(\text{day}^{-1})$	16.5 ± 1.4	13.5 ± 1.4	25.1 ± 0.9	15.0 ± 1.9
Total background $B_{\text{unc}} + B_{nn} + B_{pn}(\text{day}^{-1})$	21.7 ± 1.0	18.7 ± 1.0	27.8 ± 0.6	28.8 ± 1.3
$R_{\text{obs}}(\text{day}^{-1})$	221 ± 19	174 ± 17	225 ± 8	137 ± 17
$R_{\text{calc}}(\text{day}^{-1})$	218	155	218	130

that from the walls. The simulated spectra of spallation neutrons are chosen to span the wide range of predictions quoted in literature.

A value for ϵ_1 of 1.14 ± 0.07 is found after averaging over spectra, implying that the spallation proton recoil background is essentially symmetric like the other backgrounds. Upon simulating the possible spectra, the quantity $(1 - \epsilon_1)r = 0.1 \pm 0.05$ is found to vary little.

The yield and spectrum of neutrons from muon capture are reasonably well understood. Since these neutrons tend to be lower in energy, only those created in the vicinity of the detector have any efficiency for creating background. Knowing the veto inefficiency to miss stopping muons $[4 \pm 1]\%$, the capture rate in water surrounding the detector and its contribution to the background can be estimated using Monte Carlo simulation. Overall this proton recoil background appears to be symmetric as well, $\epsilon_1 = 0.77 \pm 0.32$, meaning that the subtraction also strongly rejects this background. The uncertainty of the residual background $(1 - \epsilon_1)B_{pn}$ is conservatively estimated to be about 160%, corresponding to $\sim 4\%$ error in $N_{\nu_e}^-$.

The results of this analysis are summarized in Table VI. Overall

$$\frac{R_{\text{obs}}}{R_{\text{calc}}} = 1.04 \pm 0.03(\text{stat}) \pm 0.08(\text{syst}). \quad (13)$$

The background estimates returned by the reactor power analysis in Table IV compare well with the results of the *swap* analysis. The 90% C.L. region for this analysis follows the same χ^2 formula, Eq. (7), as for the reactor power analysis but uses the background estimated by the *swap* method subtraction instead of minimizing the function with respect to background. Again, we find no evidence for neutrino oscillations and the excluded region for this analysis is shown in Fig. 23, curve (b).

D. Systematic uncertainties

The systematic uncertainties have three sources: the prediction of expected ν_e interactions, the efficiency estimate,

and, for the *swap* analysis, the B_{pn} estimate. The expected $\bar{\nu}_e$ uncertainty is dominated by the conversion of fission rates into neutrino fluxes, which relies on direct empirical measurements of β spectra emitted by the isotopes. The Bugey experiment [27], which directly measured the neutrino flux and energy spectrum emitted by a reactor at short baseline, found agreement within 3% using the same methods; the 3% value is used here as the estimated uncertainty.

The efficiency uncertainty can be further subdivided into that arising from direct comparisons of Monte Carlo e^+ and neutron efficiency from calibration measurements and that arising from the selection cuts themselves. The calibration runs taken with the positron and neutron sources, when compared with Monte Carlo simulations, show overall agreement across all locations of better than 1% in the efficiency predictions. However, the run-by-run agreement was at a level of 4% for neutrons and 11% for positrons. Since the ^{22}Na source is similar to the inverse beta decay signal with the e^+ close to detector threshold, the positron efficiency uncertainty over the entire $\bar{\nu}_e$ spectrum was estimated to be closer to 4% in any particular location. These run-by-run variations are then used as our systematic uncertainties in the $\bar{\nu}_e$ efficiency.

To test the robustness of the event selection, each cut is varied within a reasonable range and variations of the ratio between data and Monte Carlo simulation are examined. In order to take into account correlations all cuts were varied simultaneously by randomly sampling a multidimensional *cut space*. The rms of the resulting ratio of observed/expected is given as the selection cut uncertainty.

The *swap* method analysis has a somewhat smaller uncertainty for the selection cuts variation as the subtraction tends to cancel out systematics. However, the *swap* analysis uses a Monte Carlo estimate of the proton recoil background. Because of limited Monte Carlo statistics and the uncertainty in the fast-neutron energy spectrum, a 4% uncertainty is assigned to the neutrino signal. All of the systematic uncertainties are summarized in Table VII. The total systematic uncertainty is obtained by adding the individual errors in quadrature.

TABLE VII. Summary of the systematic uncertainties.

Error source	On minus off (%)	Swap (%)
e^+ efficiency	4	4
n efficiency	3	3
$\bar{\nu}_e$ flux prediction	3	3
$\bar{\nu}_e$ selection cuts	8	4
B_{pn} estimate	—	4
Total	10	8

The development of a second simulation and event reconstruction proved to be helpful in understanding systematic uncertainties of the analyses due to the algorithms chosen. For comparison the results for the *on-off* analysis from both reconstructions are shown in Tables II and III. An independent analysis of systematic errors was performed for the second reconstruction, similar to the method described above, giving comparable results.

VI. CONCLUSION

In conclusion, the data taken thus far from the Palo Verde experiment show no evidence for $\bar{\nu}_e \rightarrow \bar{\nu}_x$ oscillations. This result, along with the results reported by Chooz [7] and Super-Kamiokande [8], excludes two family ν_μ - ν_e mixing as being responsible for the atmospheric neutrino anomaly as originally reported by Kamiokande [1]. Later results of Super-Kamiokande, in particular data on the zenith angle distribution of muons and electrons, suggest that muon neutrinos ν_μ strongly mix with either ν_τ or with a fourth flavor of neutrino sterile to weak interaction. Clearly it is becoming important to include at least three neutrino flavors when studying results from oscillations experiments.

The most general approach would involve five unknown parameters, three mixing angles and two independent mass differences. However, an intermediate approach consists of a simple generalization of the two flavor scenario, assuming that $m_3^2 \gg m_1^2, m_2^2$ (i.e. $\Delta m_{13}^2 = \Delta m_{23}^2 = \Delta m^2$, while $\Delta m_{12}^2 \approx 0$). In such a case the mixing angle θ_{12} becomes irrelevant and one is left with only three unknown quantities: $\Delta m^2, \theta_{13}$, and θ_{23} . With this parametrization the $\bar{\nu}_e$ disappearance is governed by

$$P(\bar{\nu}_e \rightarrow \bar{\nu}_x) = \sin^2 2\theta_{13} \sin^2 \frac{\Delta m^2 L}{4E_\nu}, \quad (14)$$

while the $\nu_\mu \rightarrow \nu_\tau$ oscillations in this scenario responsible for the atmospheric neutrino results are described by

$$P(\nu_\mu \rightarrow \nu_\tau) = \cos^4 \theta_{13} \sin^2 2\theta_{23} \sin^2 \frac{\Delta m^2 L}{4E_\nu}. \quad (15)$$

A preliminary analysis of the atmospheric neutrino data based on these assumptions has been performed [28] and its




FIG. 24. Exclusion plot showing the *allowed* region of θ_{13} and Δm^2 based on the Super-Kamiokande preliminary analysis (the region *inside* the dotted curve). The region *excluded* by the neutrino reactor experiments are to the right of the corresponding dashed and continuous curves.

results are shown in Fig. 24 for the $\bar{\nu}_e$ disappearance channel. One can see that while the relevant region of the mass difference Δm^2 is determined by the atmospheric neutrino data, the mixing angle θ_{13} is not constrained very much. Here the reactor neutrino results play a decisive role.

We plan to continue taking data through summer of 2000, which will provide two additional reduced flux refueling periods.

ACKNOWLEDGMENTS

We would like to thank the Arizona Public Service Company for the generous hospitality provided at the Palo Verde plant. The important contributions of M. Chen, R. Hertenberger, K. Lou, and N. Mascarenhas in the early stages of this project are gratefully acknowledged. We thank K. Scholberg for illuminating discussions on the Super-Kamiokande three flavor analysis. We are indebted to J. Ball, B. Barish, R. Canny, A. Godber, J. Hanson, D. Michael, C. Peck, C. Roat, N. Tolich, and A. Vital for their help. We also acknowledge the generous financial help from the University of Alabama, Arizona State University, California Institute of Technology, and Stanford University. Finally, our gratitude goes to CERN, DESY, Fermilab, LANL, LLNL, SLAC, and TJNAF who at different times provided us with parts and equipment needed for the experiment. This project was supported in part by the Department of Energy. One of us (J.K.) received support from the Hungarian OTKA fund and another (L.M.) from the ARCS Foundation.

- [1] K. S. Hirata *et al.*, Phys. Lett. B **205**, 416 (1988); Y. Fukuda *et al.*, *ibid.* **335**, 237 (1994).
- [2] R. Becker-Szendy *et al.*, Phys. Rev. Lett. **69**, 1010 (1992).
- [3] E. Peterson *et al.*, Nucl. Phys. B (Proc. Suppl.) **77**, 111 (1999).
- [4] F. Boehm and P. Vogel, *Physics of Massive Neutrinos*, 2nd. ed. (Cambridge University Press, Cambridge, England, 1992).
- [5] G. Zacek *et al.*, Phys. Rev. D **34**, 2621 (1986).
- [6] Y. Declais *et al.*, Nucl. Phys. **B434**, 503 (1995), and references therein.
- [7] M. Apollonio *et al.*, Phys. Lett. B **466**, 415 (1999); M. Apollonio *et al.*, *ibid.* **420**, 397 (1998).
- [8] Y. Fukuda *et al.*, Phys. Rev. Lett. **81**, 1562 (1998).
- [9] F. Boehm *et al.*, Phys. Rev. Lett. **84**, 3764 (2000).
- [10] Model 9372 Electron Tubes Inc., Ruislip, UK.
- [11] A. G. Piepke, S. W. Moser, and V. M. Novikov, Nucl. Instrum. Methods Phys. Res. A **432**, 392 (1999).
- [12] S. P. Ahlen *et al.*, Nucl. Instrum. Methods Phys. Res. A **324**, 337 (1993).
- [13] G. Gratta *et al.*, Nucl. Instrum. Methods Phys. Res. A **400**, 456 (1997).
- [14] A. A. Hahn *et al.*, Phys. Lett. B **218**, 365 (1989).
- [15] K. Schreckenbach *et al.*, Phys. Lett. **160B**, 325 (1985).
- [16] P. Vogel *et al.*, Phys. Rev. C **24**, 1543 (1981).
- [17] P. Vogel and J. F. Beacom, Phys. Rev. D **60**, 053003 (1999).
- [18] L. V. Groshev *et al.*, Nucl. Data, Sect. A **5**, 1 (1968).
- [19] R. Brun *et al.*, GEANT 3, Report No. CERN DD/EE/84-1 (revised), 1987.
- [20] P. A. Aarnio *et al.*, “FLUKA user’s guide,” Report No. TIS-RP-190, CERN, 1990.
- [21] T. A. Gabriel *et al.*, Report No. ORNL/TM-5619-mc, 1977.
- [22] R. L. Craun and D. L. Smith, Nucl. Instrum. Methods **80**, 239 (1970).
- [23] G. Zacek, Ph.D. thesis, Technical University Munich, 1984.
- [24] M. Apollonio *et al.*, Phys. Rev. D **61**, 012001 (2000).
- [25] G. J. Feldman and R. D. Cousins, Phys. Rev. D **57**, 3873 (1998).
- [26] Y.-F. Wang *et al.*, Phys. Rev. D **62**, 013012 (2000).
- [27] Y. Declais *et al.*, Phys. Lett. B **338**, 383 (1994).
- [28] Preliminary result, SuperKamiokande Collaboration; see also K. Okumura, Ph.D. thesis, University of Tokyo (1999).