

IMB-3: a large water Cherenkov detector for nucleon decay and neutrino interactions

R. Becker-Szendy ^a, R.M. Bionta ^o, C.B. Bratton ^h, D. Casper ^s, R. Claus ^q, B. Cortez ^l, S.T. Dye ^d, S. Errede ^m, G.W. Foster ⁿ, W. Gajewski ^a, K. Ganezer ^t, M. Goldhaber ^c, T.J. Haines ^k, P.G. Halverson ^a, E. Hazen ^d, T.W. Jones ^f, D. Kielczewska ^{a,g}, W.R. Kropp ^a, J.G. Learned ^e, J.M. LoSecco ^l, S. Matsuno ^e, J. Matthews ^b, G. McGrath ^e, C. McGrew ^a, R. Miller ^j, M.S. Mudan ^f, H.S. Park ^o, L. Price ^a, F. Reines ^a, J. Schultz ^a, S. Seidel ^p, D. Sinclair ^b, H.W. Sobel ^a, J.L. Stone ^d, L.R. Sulak ^d, R. Svoboda ⁱ, G. Thornton ^r, J.C. Van der Velde ^b and C. Wuest ^o

^a The University of California, Irvine, Irvine, California 92717, USA

^b The University of Michigan, Ann Arbor, Michigan 48109, USA

^c Brookhaven National Laboratory, Upton, New York 11973, USA

^d Boston University, Boston, Massachusetts 02215, USA

^e The University of Hawaii, Honolulu, Hawaii 96822, USA

^f University College, London, WCI E6BT, UK

^g Warsaw University, Warsaw, Poland

^h Cleveland State University, Cleveland, Ohio 44115, USA

ⁱ The University of Notre Dame, Notre Dame, Indiana 46556, USA

^j Louisiana State University, Baton Rouge, Louisiana 70803, USA

^k The University of Maryland, College Park, Maryland 20742, USA

^l AT & T-Bell Labs, Holmdel, New Jersey 07733, USA

^m University of Illinois, Urbana, Illinois 61801, USA

ⁿ Fermi National Accelerator Laboratory, Batavia, Illinois 60510, USA

^o Lawrence Livermore National Laboratory, Livermore, California 94550, USA

^p The University of New Mexico, Albuquerque, New Mexico 87131, USA

^q Stanford Linear Accelerator Center, Stanford, California 94309, USA

^r University of Adelaide, Adelaide, Australia

^s CERN, Geneva, Switzerland

^t California State University, Dominguez Hills, California, USA

Received 23 June 1992

The IMB experiment, a large water Cherenkov detector which began data collection in September 1982, has undergone several upgrades to improve light collection, on-line processing power, data throughput and buffering, calibration, and operating efficiency. The current device, known as IMB-3, enjoys a factor of four light collection advantage over its precursor. Since May 1986, it has been used to search for such diverse phenomena as nucleon decay, dark matter, neutrino oscillation, and magnetic monopoles, and to study stellar collapse and cosmic rays. Due to its large size and long exposure time IMB presents unique challenges. The design and operation of the IMB-3 detector are described in detail.

1. Design motivation and history

The IMB detector [1], which was built by the Irvine–Michigan–Brookhaven collaboration, was designed to observe nucleon decay at a lifetime predicted by the grand unified theory Minimal SU(5). This theory gave a partial lifetime of the proton into the distinctive, low-background mode $p \rightarrow e^+ \pi^0$ of

$$\tau/\beta = 4.5 \times 10^{(29 \pm 17)} \text{ yr}$$

[2]. To study this lifetime required a considerable leap in sensitivity beyond the previous generation of detectors [3]: an exposure of about 10^{32} proton-years. To accomplish this in a reasonable time (a few months) required about 10^{33} protons.

The most economical way of observing this quantity of matter, 3.3 kton, is via the water Cherenkov technique. In water, charged particles with a velocity $\beta > 0.752$ produce Cherenkov radiation in a cone about the

particle direction. The opening angle of the cone approaches 42° as $\beta \rightarrow 1$. Sufficiently purified water is highly transparent (absorption length > 40 m) between 350 and 500 nm. This allows Cherenkov light to be collected by photomultiplier tubes (PMTs) around the periphery of a large volume. Hence water provides both an inexpensive source of matter and an efficient detecting medium for decay products.

3.3 kton of liquid water fills a cube ~ 15 m on a side. Because products of two-body proton decay and heavy neutral particles travel about 2.5 m in water, an additional 2 m veto region is needed outside the active volume. The full detector measures $22.5 \times 17 \times 18$ m³, which at 8 kton makes it the largest water Cherenkov detector in the world.

The amount of photosensitive area used in the original (IMB-1) detector was determined by the $p \rightarrow e^+ \pi^0$ signature and 10% energy resolution. Two 500 MeV particles create ~ 500 cm of Cherenkov-radiating track. Each centimeter of track radiates ~ 250 photons in the range of sensitivity of the PMTs. Taking a rough estimate of 10% PMT efficiency, 12 500 of the radiated photons should be detectable. Energy resolution of 10% requires detection of 100 or more photons, so the PMTs must cover about 1% of the detector's surface area.

The IMB-1 detector used 2048 13 cm hemispherical PMT [4] made by EMI Ltd. Between October 1982 and May 1984, about 417 days of data were accumulated using these tubes. By that time, the proton decay predictions of Minimal SU(5) had been decisively refuted, but newly developed theories based on supersymmetry favored lower light-level proton decays such as $p \rightarrow \mu^+ K^0$ and $p \rightarrow \nu K^+$. To increase the detector's

sensitivity to these modes involving more massive particles (and hence less available kinetic energy to radiate Cherenkov light), a light collection upgrade was necessary. It was decided to enhance the 13 cm PMTs by coupling them to wavelength-shifting plates (see section 4), and to add an additional 2048 20 cm PMTs to the detector.

In late 1983, before this upgrade began, the 13 cm PMTs began failing at an increasing rate. In the dominant failure mode, "flashing", a hairline crack at the junction of two different kinds of glass near the pins of the PMT allowed air to enter the PMT bulb. When a critical pressure was reached, the tube began irregular sparking and became a light source. By this time, it had long ceased to function as a light detector. A tube could flash so infrequently as to be unnoticeable, or glow almost continuously. Since a discharging PMT emitted enough light to trigger the detector, a high-frequency flasher rendered the detector insensitive to interesting physics. The problem was only correctable by determining the PMT responsible and disconnecting its power. The bulb failures were likely caused by residual stress in the glass stem aggravated by a tightly fitting base and possibly moisture.

IMB-2 [5] was an intermediate stage using wavelength-shifting plates and 13 cm PMTs. Because of flashing PMTs and problems with the high voltage power supplies, only 69 days of usable data were collected from September 1984 to June 1985. Further operation became impossible due to the rapidly rising PMT failure rate. Installation of 2048 20 cm PMTs (without waveshifters) was already planned as a second stage upgrade, but instead these tubes were used as replacements for the 13 cm tubes. The waveshifters




Fig. 1. Layout of the underground laboratory.

were remachined to fit 20 cm tubes and returned to the tank. A new PMT base, designed to lessen stress on the PMTs' pins, was used for the 20 cm tubes. This configuration of the detector, IMB-3 [6] with 2048 20-cm PMTs coupled to waveshifters, operated from May 1986 through March 1991, and is the focus of this paper.

2. Detector construction

High-energy muons made by cosmic-ray showers in the atmosphere penetrate tremendous distances through matter. To shield against them it was essential to locate our experiment deep underground. Even at this depth (~ 600 m, 1570 meters water equivalent), it

is traversed by 2.7 muons/s. A higher rate would have decreased the time the detector is able to accept events (such as proton decay) originating within its volume and raised the rejection factor needed to extract contained events from background. A deeper site would have been more difficult to find and more costly to excavate and operate.

The IMB detector is located about 65 km east of Cleveland, Ohio (41.72° N, 81.27° W) in a salt mine operated by Morton International (formerly Morton Thiokol). Its location in the active Fairport Mine ensures safety and regular access. A salt mine is particularly attractive due to the ease of excavation, stability, and low background from radioactivity in the surrounding rock. Several earlier experiments had been performed at this site with successful results [7].

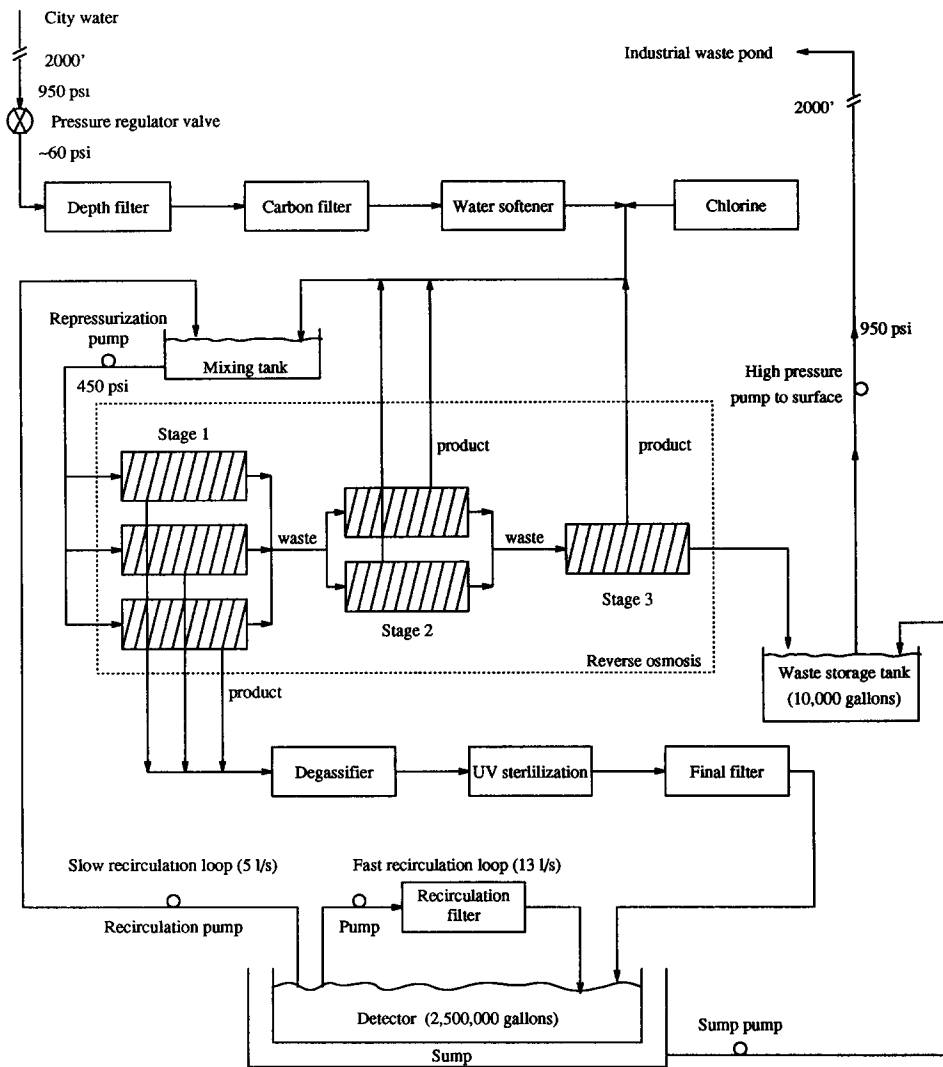


Fig. 2. Water filtration system.

In addition to a large cavity to hold the detector's mass, areas to house the readout electronics, on-line computer, and water filtration system were also required. The layout of the underground site is shown in fig. 1. The cavity was excavated by a continuous mining machine digging a long ramp down to the eventual floor of the detector pool. The ramp was then filled in to create a shelf-like platform for the laboratory, leaving a rectangular volume for the pool. The pool area and laboratory are separated by an opaque wall, with light-tight baffles to allow electrical and optical cables to pass to and from the detector. A light-tight door allows access to the pool through a clean room.

Excavated faces are covered with chain-link fence and are rock-bolted to enhance structural stability. The cavity is lined with two layers of 2.5 mm thick high-density nonreflective black polyethylene (manufactured and installed by Shlegel), separated by a plastic drainage grid which allows water to flow between the sheets. A sump at the bottom of the cavity pumps away any water which may leak through the inner liner. This leak rate (typically less than 5 l/min) is monitored by the on-site personnel, and periodic maintenance to patch leaks is performed by divers. A layer of low density ($\rho = 1 \text{ g/cm}^3$) concrete was poured between the outer liner and the cavity walls as the detector was filled with water. This served to support the liner and eliminate any irregular spaces in the excavation into which the liner might stretch to the point of tearing. Above the water line in the detector area, and throughout the rest of the laboratory, a suspended plastic dome seals the site from the hostile salt environment.

Five large I-beams, set in niches mined out near the top of the detector cavity, support a catwalk above the water level. The catwalk is reached via the light-tight door to the pool area and circuits the entire detector, allowing easy access to the PMTs and other equipment.

3. Water system

The detector pool is filled with a total of about 8000 ton of water. Water is continuously recirculated through one of two filtration systems to remove particulate, ionic, and biological contamination (see fig. 2).

The primary system takes water from the municipal supply and pretreats it using depth filters, carbon filters, water softeners, and chlorine. The depth filters remove particulates down to about 10 μm ; the carbon filters remove organic materials and chlorine; the water softeners remove magnesium and calcium salts; and the chlorine, injected at 0.1 ppm, prevents biofouling. The prepared water then goes to a multistage reverse osmosis (RO) system, where particulates larger than 0.001 μm are removed and total dissolved solids are

rejected with an efficiency of 95%. Using the Painesville, Ohio municipal supply, the RO produces water with about 4 ppm of total dissolved solids.

Conventional RO systems produce about 30–50% waste water. Because the detector is located in a deep mine, it is expensive to dispose of large amounts of waste water. Consequently a multistage RO system was designed which produces less than 10% waste.

A degassifier is used to remove dissolved carbon dioxide as a precaution against small bubble formation on the phototube surfaces. UV sterilization is followed by a final 1 μm filtration stage. About 5 l/s can be processed.

Water purity in the pool deteriorates due to constant leaching of materials from underwater components. The capacity of the water system is matched to this effect. This recirculation of the pool takes about 30 days for one complete volume. Since the quality of filtration product depends on the input water quality, when high purity water from the pool is recycled, output with less than 0.3 ppm dissolved solids results.

A parallel recirculation system, employing 1–5 μm filters, is used to quickly remove any suspended solids which are introduced into the pool. This system can process 13 l/s and recirculate the detector volume in 10 days.

Covering the pool surface is a black "floating floor", constructed from Hypalon (manufactured by Burke Rubber Co.) sheeting and closed-cell polyethylene foam planks for buoyancy. The floor can support the weight of several people. Hinged flaps at the edges can be raised to allow strings of PMTs to be removed from the detector and openings in a number of places allow light sources and test devices to be lowered into the pool. In addition to reducing evaporation from the pool, the floor prevents dust and other impurities from reaching the water. The most important function of the floating floor is to eliminate an air/water interface which would cause total internal reflection of Cherenkov photons. Because exhaust air from divers repairing leaks in the liner is introduced into the pool periodically, short drapes of Hypalon are hung vertically from the underside of the cover.

In order to preserve water clarity, all underwater components were carefully selected to be nonpolluting. Metal parts were avoided wherever possible. Plastic components, including the PMT high voltage/signal cables, were baked to eliminate contaminants which might eventually be drawn into the water. The Hypalon floating floor and the high-density polyethylene liners were chemically cleaned during installation as an added precaution.

Although the water in the detector is extremely transparent to light in the region of 350–500 nm, it is not sterile. Within several months of the initial filling, evidence of a slimy biological growth was noticed. The

major portion of the slime was identified as *Beggiatoa*, although other slime producing organisms such as *Pseudomonas* spp. were found as well. The organism was difficult to eliminate. Because it does not adversely affect water clarity, its presence is tolerated.

4. Light collection

The 2048 20-cm Hamamatsu R1408 PMTs which comprise the sensitive element of the IMB-3 detector are arrayed along the walls of the detector, facing inward on a ~ 1 m lattice, as shown in fig. 3. Since visible Cherenkov photons are few in number, PMTs must be able to detect single photons with high efficiency. The quantum efficiency of producing a photoelectron (pe) is a function of the color of the incident light, as plotted in fig. 4, depending mainly on the alkali photocathode and the borosilicate glass envelope. The potential difference between the photocathode (which is held at ground since it is in contact with water) and anode is generally ~ 2000 V. At this voltage, the efficiency for collecting pes is estimated to be roughly 60%. Timing information from the PMTs is vital to the spatial reconstruction of events. For these PMTs, the timing resolution at single pe illumination is about 13 ns (FWHM). The resolution improves with increasing light-level, reaching about 5 ns at 10 pes.

Each tube is housed in a neutrally buoyant gray PVC assembly (see fig. 5). A PMT is glued with epoxy onto the beveled edge of the housing, and the joint is covered with a silicone rubber water-tight sealant. The back end of the housing is sealed by an O-ring and a round transparent acrylic plate which is attached by stainless steel screws. The clear window allows visual inspection of the seal and base without opening the housing. An opaque plastic cover is fitted over the Plexiglas window to prevent any light from electrical discharge within the PMT from entering the detector. A black plastic disk is fitted around the PMT neck to block light leakage into the bulb itself. The cable carrying high voltage to and signals from the PMT enters the housing through a water-tight feedthrough. Two PVC cylinders filled with lead shot are attached to the housing to make the assembly as a whole neutrally buoyant.

PMTs arrayed along the faces of the detector are supported by two methods. Tubes along the sides are suspended vertically, in strings of 16. These PMTs are hung by two parallel nylon-coated stainless steel cables attached at one end to the catwalk and at the other end to a hanging weight near the bottom of the pool. The two stainless cables pass through slots on either side of the PMT housing where they are secured by PVC screws. PMTs on the top and bottom are mounted on a 1 m long rectangular parallelepiped beam section,




Fig. 3. PMT grid arrangement.

constructed of PVC pipe. A total of 16 beam sections are then joined by stainless steel screws into an entire beam. Beams are lowered vertically, swung into a horizontal position, and then attached at both ends to the catwalk.

When the decision was made to upgrade IMB-1, several different methods of light collection enhancement were studied. The most economical and effective of these possibilities was waveshifter technology [5,8]. Although waveshifters have a variety of applications in particle detection, they had never before been used on so large a scale to enhance the light collection of PMTs. The fundamental idea is quite simple: a large area plate containing a wavelength shifting compound is optically coupled to the bulb of a hemispherical PMT. The PMT continues to record light striking its photocathode directly. In addition, light missing the




Fig. 4. PMT quantum efficiency.




Fig. 5. Wavelength-shifter plate and PMT.

PMT but striking the waveshifter is absorbed and re-emitted at a longer wavelength. Although re-emission is isotropic inside the plate, a fraction of the re-emitted light is trapped by total internal reflection. Some

smaller fraction eventually reaches the PMT/waveshifter interface and is recorded by the PMT.

Of utmost importance are the spectral characteristics of the waveshifting floor. Ideally, its absorption

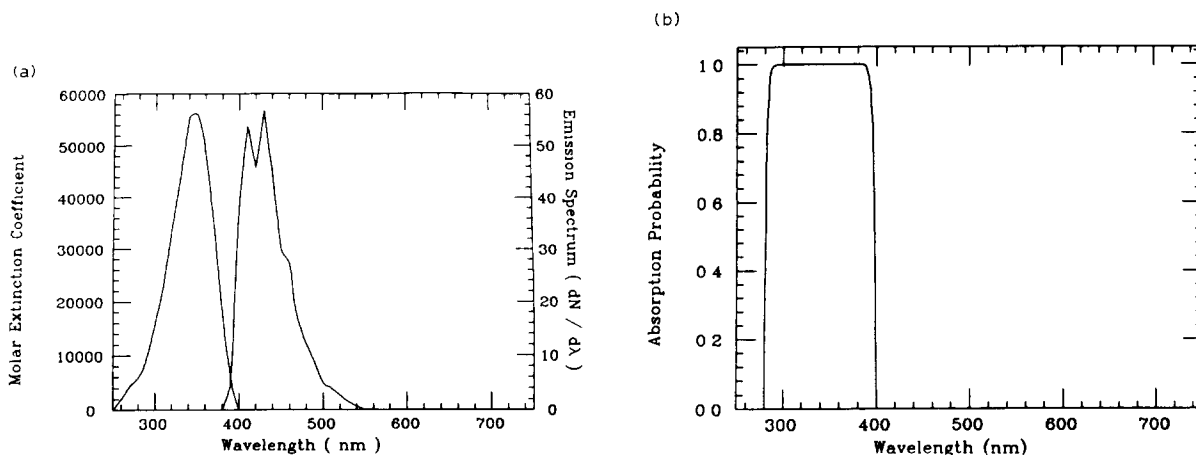


Fig. 6. (a) Bis-MSB optical properties. (b) Wavelength-shifter absorption probability.

spectrum would cover the full range of Cherenkov wavelengths in water. Cherenkov light obeys a $1/\lambda^2$ spectrum, but the light reaching the photocathode is reduced in the near UV, below 300 nm due to scattering in water and absorption in the PMT's glass envelope. The ideal fluor would also have a re-emission spectrum closely matched to the blue-green sensitivity of the bialkali photocathode (see fig. 4).

Based on these considerations, the fluor bis-MSB (p-bis (o-methyl styryl) benzene, or $C_{24}H_{22}$), was selected. Its spectral characteristics are shown in fig. 6a. The absorption curve plotted is actually the molar

extinction coefficient (μ), which is related to the absorption probability by

$$P_{\text{abs}} = 1 - 10^{-\mu\rho d},$$

where ρ is the fluor's concentration in mol/l, and d is the thickness of the material. The remainder of the parameters, namely the concentration (50 mg/l), the plate thickness (1.3 cm), mounting position on the PMT (30° up from the PMT equator), and method of coupling to the PMT (simple physical contact) were determined by a series of computer simulations and actual tests in the detector. A small gain in light

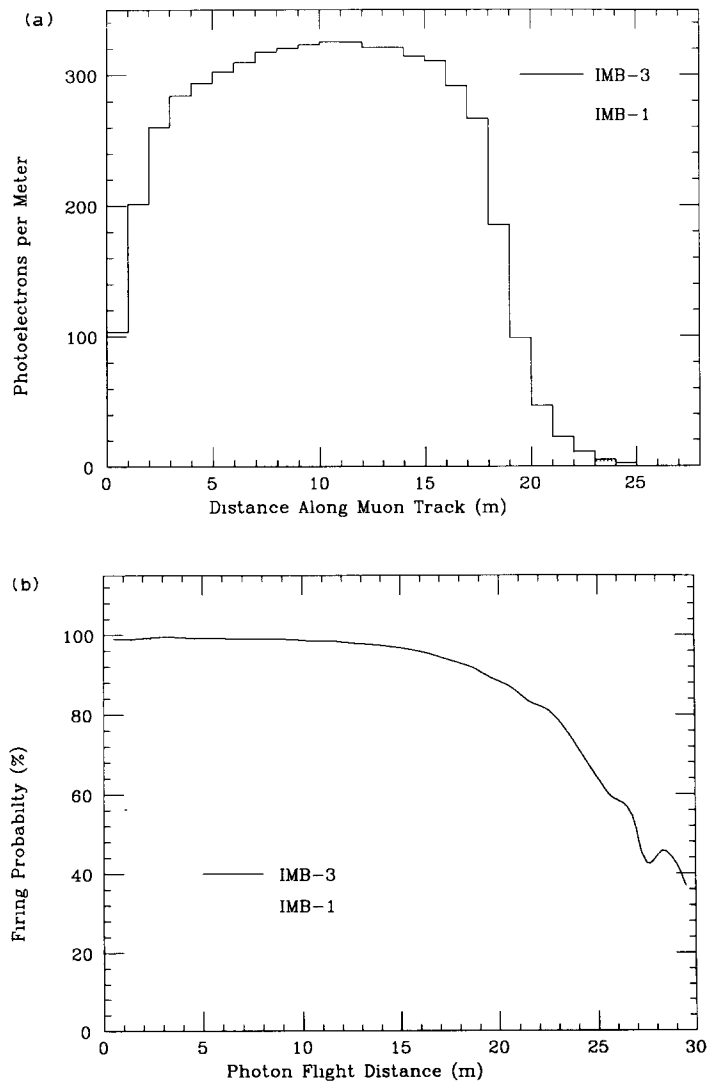


Fig. 7. (a) Cosmic-ray muon corrected pulse-height, IMB-1 and IMB-3. (b) Cosmic-ray muon occupancy vs distance, IMB-1 and IMB-3.




Fig. 7. Continued. (c) Muon decay coincidence levels, IMB-1 and IMB-3.

collection could be achieved by fixing the plates to the PMT bulb with epoxy, but for IMB, the benefits of the easy-to-maintain flexible contact fit are greater.

Plate performance is relatively insensitive to small (factor of two) variations in fluor concentration and plate thickness, since these values are chosen to give nearly 100% absorption over a wide range of wavelengths. The absorption efficiency of 50 mg/l, 1.3 cm thick plates is shown in fig. 6b. Since the decay constant of the waveshifting fluor (~ 4 ns) is small compared to the inherent timing resolution of the PMTs, the plates do not degrade timing information significantly. At single pe illumination, the time resolution is 15 ns (FWHM) with plate, versus 13 ns without.

The 60×60 cm² plate of UV transmitting acrylic, doped with wavelength-shifting compound, is mated to the bulb of the PMT, as shown in fig. 5. The outside edges of the plates are masked with reflective tape to prevent light from escaping. The central hole is tapered and rounded to make snug contact with the PMT surface. Plates on side PMTs are held in place by 4 PVC clips attached to each plate, with grooves through which the stainless steel support cables are strung under tension, pulling the plate against the PMT. Plates on top and bottom PMTs are clamped to PVC holders on the beam sections by eight large PVC washers and four PVC screws.

Three different comparisons of the detector's light collection abilities for the IMB-1 and IMB-3 configurations are presented in figs 7a–7c. Cosmic-ray muons are one light source which is constant through time. Fig. 7a shows the number of pes collected per meter along a through-going muon track. Fig. 7b shows PMT firing probability versus the photon pathlength from a through-going muon track. Another time-independent

light source is the spectrum of electrons from the decay of cosmic-ray muons which stop in the detector. Fig. 7c shows the number of PMTs fired by electrons from muon decay. Each of these quantifications of light collection has its limitations, but it is clear that roughly a factor of four increase in sensitivity over IMB-1 was achieved by the new PMTs augmented with waveshifter plates. About one-half of this increase was provided by the waveshifter plates.

5. Power supplies and associated hardware

High voltage (HV) is provided to each tube by a specially ordered 93 Ω coaxial cable (RG59) with an outer jacket double the normal thickness. The same cable is used to read out the PMT's signal. Cables are joined together into flat ribbons of eight for ease of handling. Thus, 256 8-conductor cables are required to serve the 2048 PMTs in the detector. Cables are guided through trays on the catwalk, through light-tight baffles, and into the electronics laboratory. Two different cable lengths are used, depending on a PMT's distance from the laboratory. It is necessary to compensate for these two different lengths (and hence two different signal propagation times).

Voltage provided to the tube is divided between dynodes and anode by a custom-designed PMT base. The base also sets the various focusing potentials inside the tube. The entire base assembly is dipped in an insulating coating to prevent arcing between HV conductors inside the PMT housing. The base circuit is back-terminated by a 91 Ω resistor to limit reflections.

Several sources of power have been used over the course of the experiment. The units inherited from

IMB-1 were extremely unreliable and required constant maintenance. A nonprogrammable custom-built supply was used as a stopgap until the HV system was finally replaced by a LeCroy 1440 system. The supplies are current-limited to prevent damage to power supplies and PMTs should water in a tube housing cause a short-circuit or should HV be left on with lights turned on in the detector area. For additional protection, a magnetic interlock trips off power to all tubes when the door to the detector is opened. The HV system's controller unit communicates with the on-line computer via an RS-232 terminal line. Under normal circumstances, the computer can monitor and set the voltage and turn on or off any of the 128 HV channels in the detector.

Groups of 16 PMTs selected to have similar gains are operated at a common voltage from one HV channel. A total of 128 channels of HV are required. Voltage from a given channel is passed across the readout electronics card (see below) and then through a "cable paddle" to the 8-conductor cable for each of two 8-tube "Strings". The cable paddle has a socketed resistor for each channel, which allows an individual PMT to be (manually) turned off or have its voltage trimmed to optimize performance. It also contains hardware to decouple the ac PMT signal from the dc HV level and feed the low-voltage pulse into the readout electronics. Cable paddle circuits are painted with an insulating coating to prevent arcing and shocks, and to restrict cross-talk. Nonetheless cross-talk proves a persistent problem when one channel has a pulse height in the range of hundreds of pes.

6. Detector trigger

Detector triggers are formed from analog sums of individual PMT discriminator pulses (see fig. 8). The amplitude and width of a PMT pulse are related to the number of pes reaching the first dynode. A single pe pulse from one of the PMTs is typically 50 mV high and 30 ns wide. The PMT signals are passed from the cable paddles to custom discriminator/analog-to-digital converting circuit boards called "8-channel cards".

Several different forms of information are gleaned from the pulse of a tube by the 8-channel cards. The most basic information is whether the tube has seen any light at all. Since thermal and electronic noise constantly produce small signals, a discriminator monitors the PMT and fires only when a pulse exceeds ~ 30 mV. Although discriminator thresholds can be individually set for all 2048 tubes, in practice this is not necessary. A single threshold value, chosen well below the single pe level, but above the background of thermal noise, is used for all PMTs. Digitally encoded PMT threshold values are stored in memory on the 128-channel digitizing cards (see below). Corresponding threshold voltage is provided to discriminators on the 8-channel cards.

When a PMT's discriminator fires, a number of processes are initiated. Each 8-channel card creates an analog sum of its 8 discriminator outputs, forming a gate 55 ns long and ~ 70 mV deep per PMT. The levels of 8 cards corresponding to an 8-by-8 PMT array or "patch" are in turn added together on the back-

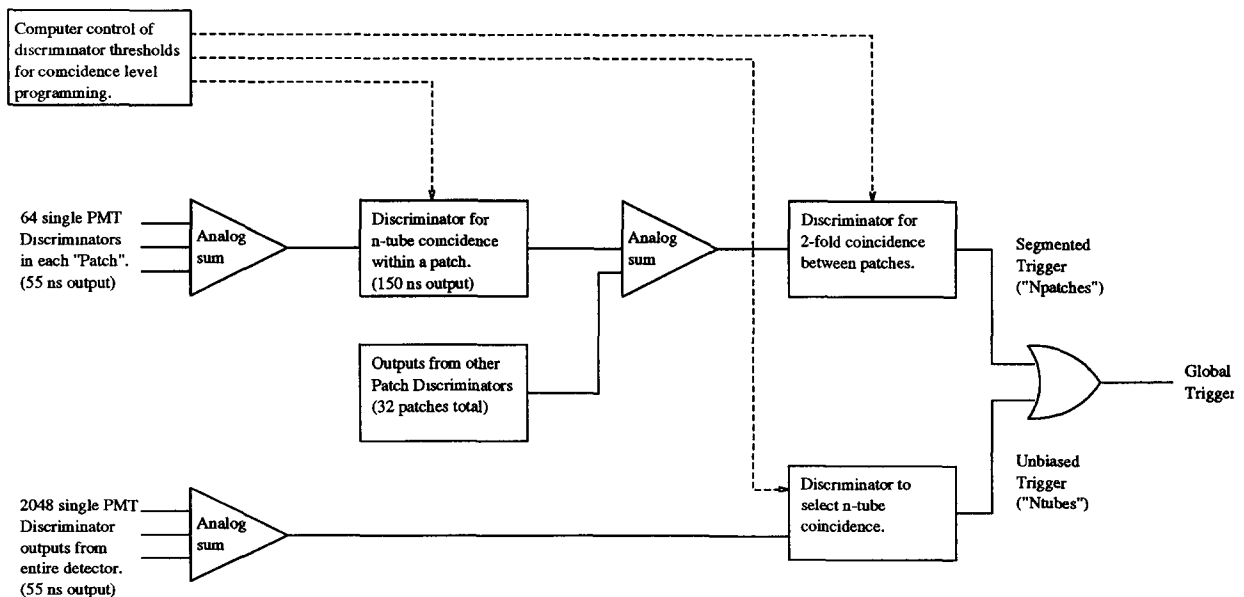


Fig. 8. Trigger system.

plane of a custom electronics crate. A total of 32 such analog discriminator sums or “patch OR-outs” (one for each patch) are used to form the detector trigger. Firing of a tube’s discriminator also activates two analog storage circuits called the “ \mathcal{F} 1 ramp” and “ \mathcal{E} ramp”, which record the time and amplitude of the tube pulse and are explained in section 7.

Patch OR-outs from patches with shorter (52 m) cables run through cable delays so their signals reach the trigger hardware with the same timing as those from patches with longer (70 m) cables. These signals are then checked against two different trigger criteria.

The unbiased or “ N_{tubes} trigger is tested by summing all 32 patch OR-outs into a single global sum for the entire detector. This N_{tubes} analog sum is proportional to the number of PMTs fired in the last 55 ns. A programmable discriminator set under computer control monitors this sum, and when it exceeds the specified number of PMTs firing in coincidence, the discriminator fires, creating an N_{tubes} trigger. Until the beginning of 1989, this level was about 20 PMTs. After reducing PMT noise and cable paddle cross-talk through tuning of voltages, it was reduced to about 10 PMTs.

A second, segmented or N_{patches} trigger condition is also tested. All patch OR-outs are individually discriminated against a single computer-selectable threshold level. Patch discriminators generate a 150 ns logic gate when the level corresponding to three tubes is reached. Outputs of all 32 patch discriminators are then added into an analog sum of the number of patches firing in

150 ns coincidence. This signal is finally discriminated against a computer-selected level corresponding to two patches, and when that condition occurs, an N_{patches} trigger is created. Hence, the N_{patches} trigger requires at least three tubes firing in 55 ns coincidence in at least two different patches, with the two threefold coincidences occurring within 150 ns of each other. The 55 ns window for coincidence within a patch is comparable to the light travel time across a patch, and the 150 ns window for coincidence between two patches is comparable to the time for light to cross the diagonal of the detector.

When either an N_{tubes} or N_{patches} trigger occurs, a detector trigger signal (global trigger) is fanned out to a variety of places. A WWV clock is read into a CAMAC input register, freezing the absolute time of the trigger to an accuracy of a millisecond. Global Trigger is also routed to the custom electronics read-out supervisor/CAMAC interface, initiating the data acquisition process. The supervisor produces two logic signals: Digitize Gate, which is fanned-out to the back-planes of all 16 custom electronics crates (see below); and CPU Busy, which indicates that an event is in process and further triggers cannot be accepted. The Digitize Gate signal must be cable-delayed en route to the half of the detector with longer PMT cables, in order to preserve the timing of the event. The timing of this cable-delayed signal is monitored by CAMAC to preclude any drift in timing. Finally, the supervisor asserts the Look-At-Me (or LAM) signal on its own CAMAC slot to alert the on-line computer that a

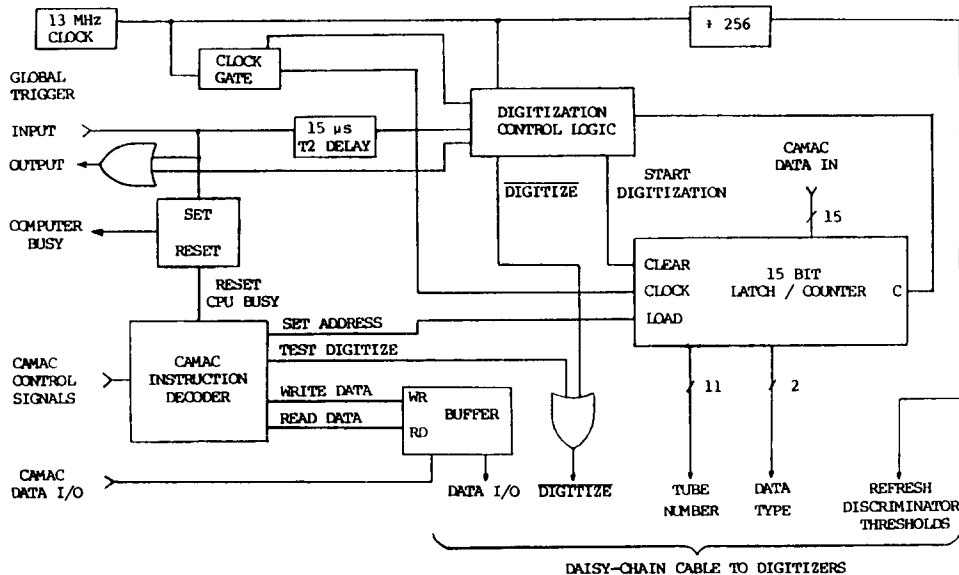


Fig. 9. Supervisor block diagram.

trigger has occurred. The supervisor's functions are summarized in fig. 9.

7. PMT readout electronics

Readout of the detector's 2048 PMTs is accomplished by a hierarchy of electronics, performing the following functions: < 1 pe pulse discrimination on each PMT, two-hit TDC, charge ADC. The bonded cable set shared by 8 tubes is bolted and soldered to a cable paddle, which is itself plugged into an 8-channel readout card. Sixteen paddle/readout card assemblies share one custom crate. The crate distributes low voltage power for the readout circuits, and is controlled by one 128-channel custom digitizing card. Each crate controls two 8×8 arrays of tubes called patches (see fig. 10). Sixteen such crates (or equivalently, 32 patches) comprise the entire detector. All crates are in turn controlled by one custom readout supervisor, which communicates with them via two 60-conductor flat cables. A diagram of the readout hierarchy is shown in fig. 11.

When activated by a given tube's discriminator firing, the PMT's $\mathcal{F}1$ ramp circuit on its 8-channel card begins charging a capacitor at a constant rate. The




Fig. 10. Segmentation of the detector into patches.

same tube's \mathcal{C} ramp" circuit simultaneously charges a different capacitor at a variable rate, integrating the charge of the PMT pulse. The $\mathcal{F}1$ capacitor charges until the Digitize Gate signal is received via the crate backplane, or until ~ 600 ns elapse. Arrival of Digitize Gate indicates that the detector has triggered and freezes the charges stored on the $\mathcal{F}1$ and \mathcal{C} capacitors. The $\mathcal{F}1$ capacitor thus contains a charge proportional to the time elapsed between the discriminator firing and the arrival of the trigger signal. The \mathcal{C} capacitor stores a charge related to the number of pes collected by the PMT. Should 600 ns elapse after a PMT's discriminator fires without arrival of the Digi-

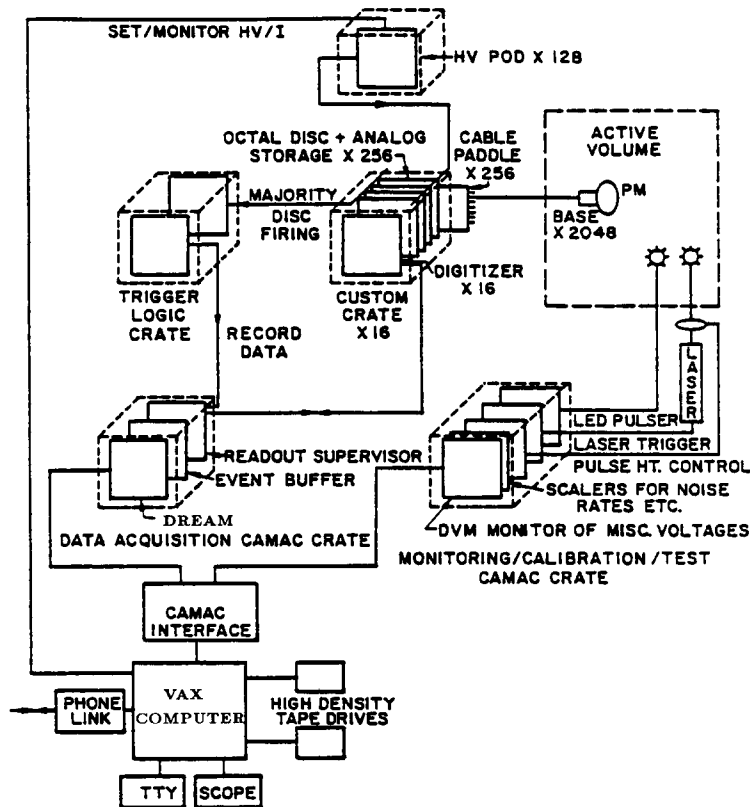


Fig. 11. PMT readout electronics hierarchy.

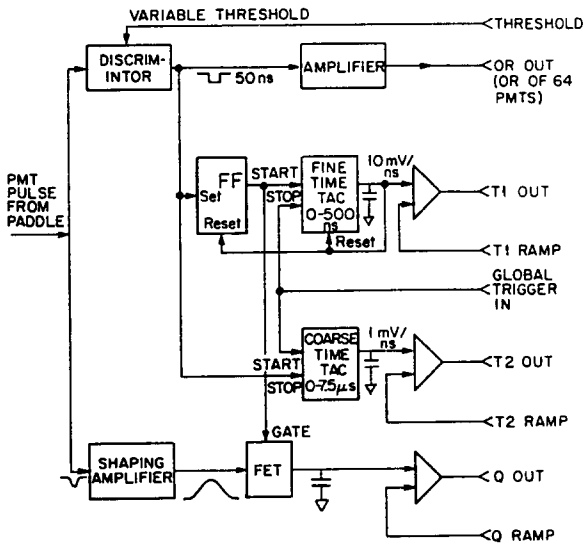


Fig. 12. 8-channel readout card.

tize Gate trigger signal (because the PMT firing was isolated dark noise or events below the detector threshold), both the $\mathcal{F}1$ and \mathcal{Q} capacitors are discharged and the discriminator is reset to accept another pulse from the tube.

After firing once, a given tube's discriminator cannot immediately fire again. Only after the Digitize Gate signal arrives or the $\mathcal{F}1$ and \mathcal{Q} ramps are reset is the discriminator re-enabled. The Digitize Gate signal begins a third analog circuit called the $\mathcal{F}2$ ramp, which also charges a capacitor at a constant rate. Should the discriminator fire after receipt of Digitize Gate, the charge on the $\mathcal{F}2$ capacitor is frozen, storing a charge proportional to the time elapsed between arrival of the trigger and firing of the discriminator. Both $\mathcal{F}1$ and $\mathcal{F}2$ values can exist for a single tube. If the discriminator fires before the trigger arrives, the \mathcal{Q} ramp is associated with the pretrigger pulse. If the tube does not fire before the trigger, the charge measurement is associated with the post-trigger pulse. A block diagram of the 8-channel card is shown in fig. 12.

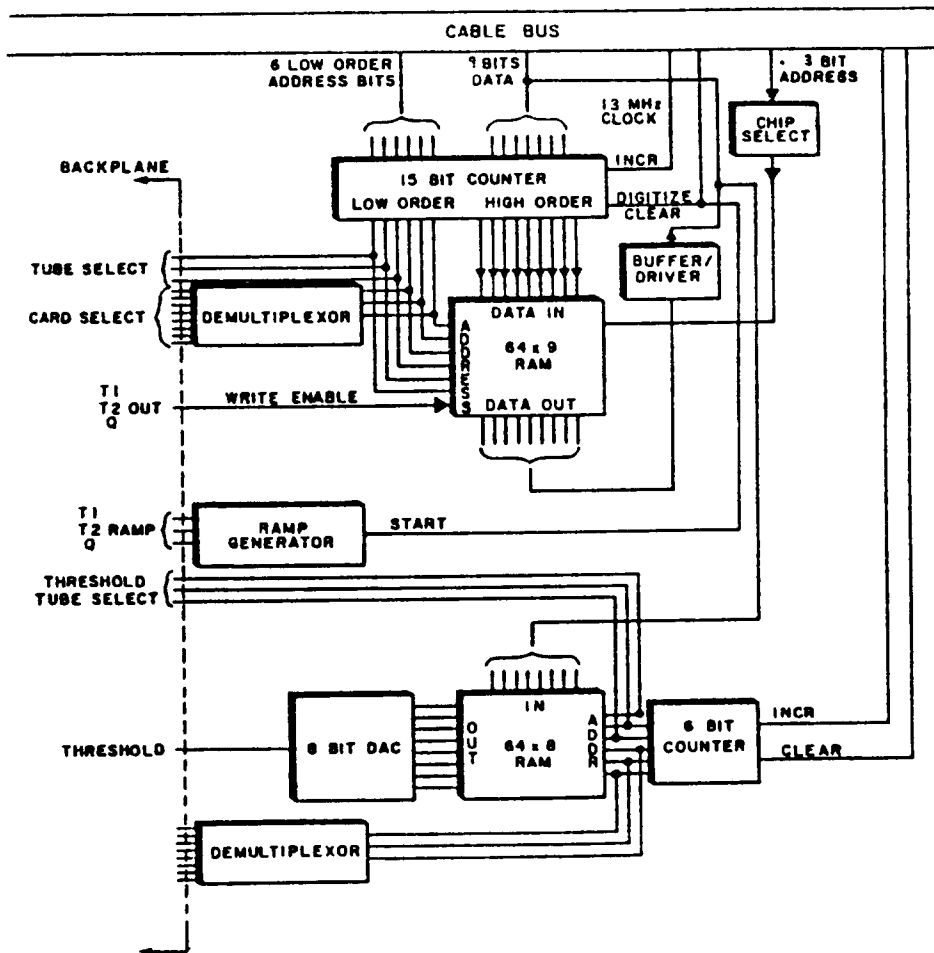


Fig. 13. 128-channel digitizing card.

About 15 μ s after receipt of Global Trigger, the supervisor sends a Digitize signal over two flat ribbon cables to all 16 digitizer cards, instructing them to begin converting the analog $\mathcal{F}1$, $\mathcal{F}2$, and \mathcal{Q} voltages into digital values. Simultaneously, all crates digitize the information on their 16 8-channel cards. The supervisor also provides a 13 MHz clock signal, which increments a counter on each digitizer a total of 2^{15} times. The lowest 6 bits in this counter address one of 64 PMTs in either half of the crate. The upper 9 bits are used by each digitizer to reference a linearly rising voltage ramp, which is fed to all 8-channel cards via the crate backplane. The $\mathcal{F}1$, $\mathcal{F}2$, and \mathcal{Q} voltages stored on the capacitors of a selected tube are simultaneously compared with this reference voltage on the 8-channel card. When the reference ramp reaches the value stored on a capacitor, an associated voltage comparator fires, latching in the digitizer's memory a 9-bit digital value corresponding to the level of the refer-

ence ramp. Fig. 13 shows a block diagram of the digitizer board.

Thus, the $\mathcal{F}1$, $\mathcal{F}2$, and \mathcal{Q} data for each tube, stored as analog information on each 8-channel card, are converted to a digital value between 0 and 511 in a digitizer memory. Capacitor charging rates on the 8-channel cards and reference ramp slopes on the digitizers are selected to make 1 count of $\mathcal{F}1$ correspond to 1 ns, 1 count of $\mathcal{F}2$ correspond to 15 ns, and 1 count of \mathcal{Q} correspond to 0.1–0.25 pes. Hence the $\mathcal{F}1$ time scale lasts for about 511 ns before the trigger and the $\mathcal{F}2$ scale lasts for about 7.5 μ s after the trigger, while the Q scale overflows at about 50–100 pes.

After 2^{15} cycles of the 13 MHz clock (or about 2.5 ms), the supervisor turns off both the Digitize signal on the flat cables and the Digitize Gate NIM signal fed to all crates. The end of Digitize instructs a custom buffer memory module to collect all the digital PMT data ($\mathcal{F}1$, $\mathcal{F}2$, and \mathcal{Q}) from the entire detector. A buffer




Fig. 14. Event acquisition timing.

memory module has room to store data from sixteen events. 9-bit data are transferred in parallel from digitizer RAM to the buffer memory via the two flat cables, a process requiring about 1 ms. The end of Digitize Gate causes the $\mathcal{T}1$, $\mathcal{T}2$, and \mathcal{Q} capacitors on all 8-channel cards to discharge and all tube discriminators to reset themselves to prepare for subsequent pulses. The 1 ms required for data transfer to the buffer memory allows ample time for all capacitors to be discharged. Fig. 14 illustrates the timing of the various steps in the trigger-digitization-transfer process.

8. On-line computer / trigger interface

A VAX minicomputer controls and monitors the experiment. In 1990, the 11/750 model was replaced with a more powerful and reliable MicroVAX II. A

Kinetic Systems model 2051 (2061 after installation of the MicroVAX II) parallel highway branch driver connects the computer to the CAMAC system. The sophisticated VMS operating system, availability of high level languages such as FORTRAN, and the convenience of performing CAMAC operations by FORTRAN-callable subroutines make the VAX system easy to operate and diagnose. Full compatibility with computers used for off-line analysis is also an asset.

A custom data acquisition program runs on the VAX as a detached process. The main body of this program is a loop which reads from the buffer memory, applies a variety of real-time algorithms and cuts, formats, and writes on magnetic tape the PMT data from previous events (see below). When a detector trigger occurs and the supervisor turns on its LAM signal, the branch driver relays this information to the VAX. Loop execution is momentarily suspended and a special interrupt servicing routine called an AST (for

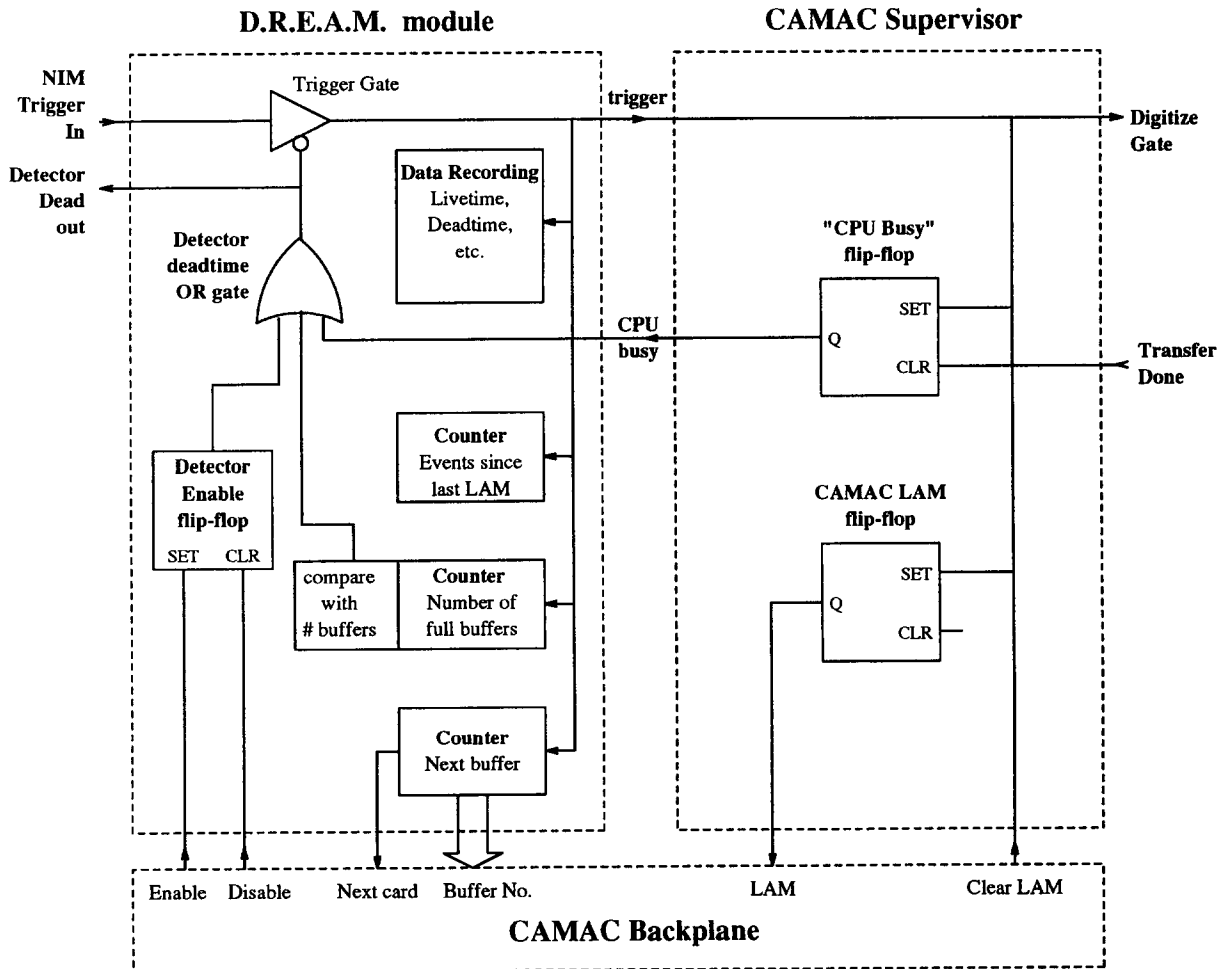


Fig. 15. Dead-time reduction event acquisition module and supervisor.

asynchronous system trap) is invoked. The AST routine reads CAMAC modules containing, among other things, the absolute time of the event, the type of trigger responsible for the event, the total time elapsed in the run, and the cumulative amount of time during which the detector has been “dead” or unable to accept triggers. Miscellaneous monitoring information such as the water level in the pool and air temperature are read out as well. The AST notes the fact that the buffer memory now contains one more event than it had before the trigger. If room is left in the buffer for another event, the final step taken by the AST is to clear the LAM of the supervisor (by CAMAC instruction), which also removes the CPU Busy detector veto signal. The apparatus is now able to accept another trigger. Should no room remain in the buffer memory for another event, the detector is left disabled until space is created by the main program loop reading a previous event’s PMT data.

Re-enabling the detector by computer command takes a total of about 30 ms from the occurrence of a trigger, and more when the buffer memory is full due to a rate fluctuation. This delay is almost 10 times as long as the inherent deadtime of the custom electronics, and is mostly due to relatively slow processing of interrupts by the VMS operating system. At the usual trigger rate of 2.7 s^{-1} , the detector would be live between 85% and 90% of the time.

9. Deadtime reduction electronics

To improve the detector’s ability to record rapid bursts of events and also eliminate the 10–15% deadtime for recording the ambient data stream, modifications were made to the data acquisition system in mid-1988. These changes were threefold: an increase (256 vs ~ 8) in the number of CAMAC buffer memories available to receive PMT data from the digitizing cards, installation of a custom module to perform the control of buffer memory cards and trigger timing readout which previously required computer intervention, and modification of the supervisor to re-enable triggering, digitization, and transfer without CAMAC reset.

The goal of this upgrade was to eliminate the delay caused by waiting for the VAX to process one trigger before accepting another, without major changes to the existing electronics, CAMAC interface, or computer system. Reducing the deadtime per event to the minimum 3.5 ms required for the actual recording, digitization, and transfer to buffer memory allows high trigger rate bursts up to 250 s^{-1} to be collected, limited only by the capacity to buffer these triggers in CAMAC. Using Supernova 1987A as a guide, and estimating the likely distance to the next stellar collapse as 10 kpc

rather than 50 kpc, the ~ 10 interactions from 1987A scale to ~ 250 expected from a similar event in our galaxy, thus determining the necessary buffering capacity.

The heart of the system is a module christened the deadtime reduction event acquisition monitor (or DREAM for short). A block diagram of the DREAM/supervisor system is shown in fig. 15. The DREAM performs two general functions: first, keeping account of how many buffer memory slots remain available and which is to be the next filled, and second, recording and storing the relative times of, and the specific trigger condition (N_{tubes} and/or N_{patches}) responsible for, each trigger.

This upgrade decreased the detector’s dead-time to $\sim 1\%$, corresponding to the 3.5 ms needed to record, digitize and transfer an event, multiplied by the 2.7 s^{-1} ambient trigger rate. Steady-state data rates up to 5 s^{-1} can be recorded without buffer overflow. Using the timing circuitry, the relative times of any two triggers are measured to within about $4 \mu\text{s}$. The system was designed to be expandable to over 2000 event memories, using a separate DREAM module for each crate of 16 memory cards.

10. On-line data processing

In addition to servicing trigger interrupts, the on-line computer must process the data from events which have already occurred, decide which events are potentially interesting and should be written on tape for subsequent analysis, format these events for output, and transmit them to a magnetic tape drive. on-line analysis and cuts are designed to reduce the quantity of data which must be analyzed off-line.

Simple algorithms such as finding the number and “center of mass” of PMTs which fired during the $\mathcal{S}1$ scale and averaging the raw \mathcal{E} values recorded for $\mathcal{S}1$ hits are performed first. Events having between 450 and 1800 $\mathcal{S}1$ hits are then subjected to three more sophisticated routines which correct the raw times and pulse heights on a channel-by-channel basis and attempt to determine the direction of the presumed muon track. This real-time reconstruction ability can be used to track sources through the sky. A histogram of reconstructed zenith-angle cosines for cosmic-ray muons is plotted in fig. 16. Off-line checks estimate the accuracy of these track fitting programs as 7° to 10° . Events with less than 450 hits are typically muons which either stop inside the detector or barely clip one corner. Events with over 1800 tubes firing are bundles of parallel muons or very energetic showers. Both of these types of events do not fit the geometry of a through-going muon, and consequently the fitting programs’ results are unreliable.




Fig. 16. Cosmic-ray muon z -direction cosine reconstructed on-line.

Approximately 250 000 triggers are recorded each day, almost all of which are downward-going cosmic-ray muons. A histogram of the number of PMTs firing during the $\mathcal{F}1$ timescale is shown in fig. 17. The simplest cut applied to the raw data is a cut on the number of $\mathcal{F}1$ hits. All events radiating less than about 2 GeV of energy inside the detector (including all proton decays, and low and medium energy neutrino interactions) should have 900 or fewer PMTs illuminated. This cut includes 45% of the raw data, since many cosmic rays with less than ~ 10 m path-length in the detector also satisfy this criterion. Events reconstructed on-line at greater than 70° from the zenith (i.e., upward or nearly upward-going) are saved, since they may be produced by neutrinos outside the detector. This cut saves only 2% of the raw data. Events with greater than 1850 $\mathcal{F}1$ hits are saved for studies of high-energy phenomena (e.g., interactions of high-energy neutrinos from active galactic nuclei). This cut saves roughly 3% of the raw data. Finally, one-




Fig. 17. NT1 distribution of raw data.

eighth of the raw data, selected randomly, is saved as an unbiased data sample. The combination of all these cuts saves somewhat more than half of the detector triggers.

Events accepted by on-line cuts are packed together for output to magnetic tape. PMT data are zero-suppressed to conserve tape. In addition to PMT data, pertinent CAMAC information (including the time of the event), results of the on-line algorithms, and the reason the event was saved are stored with the event. The data format on tape has been described elsewhere [6]. Storage of one event typically requires about 4000 bytes. Until May 1989, data were recorded in 32 Kbyte blocks on Aviv 6250 bpi magnetic tape drives. One 715 m tape reel stored about one-half day of cut data. In March 1990, Ebyte 8 mm video tape drives were introduced to reduce data handling and storage costs. These permit every trigger to be written on tape and allow data-taking to continue during periods when access to the laboratory is unavailable. One 120 min 8 mm tape stores two days of uncut data (or four days of cut data).

Although full pulse-height and timing information for all cosmic-ray muons was not saved on tape before installation of the 8 mm tape drives, condensed summaries, containing reconstructed directions, event times, and other useful quantities recorded or calculated on-line, were written out for all events having between 600 and 1800 PMT hits. These “muon summary records” allow searches for muon sources, anisotropies, and periodicities while taking up less than 1% of the space needed to record the full data of the events.

In November 1986, a second on-line computer (a MicroVAX II) was added to the system to perform further data reduction on-site. This machine performed the first stage of (formerly) off-line analysis (Pass0) on the data stream broadcast over an Ethernet connection from the data acquisition computer.

Since the inception of the experiment, two independent analysis chains have been used to maximize efficiency and identify possible systematic errors. After November 1986, the first stage analyses of both groups were carried out underground, with each group providing their own algorithms and calibrations. An event rejection factor of about 20 is achieved at this stage, with $> 80\%$ efficiency for retaining contained events and upward-going muons. In addition to routine analysis, several algorithms which search for bursts of low-energy events were installed, and the system was used for real-time muon astronomy at certain times. The Pass0 results were written on a separate tape which was subjected to further off-line analysis by each group. Tapes of unprocessed data were simultaneously written by the 11/750 as backups.

This powerful yet somewhat awkward arrangement was finally consolidated by the upgrade of the analysis

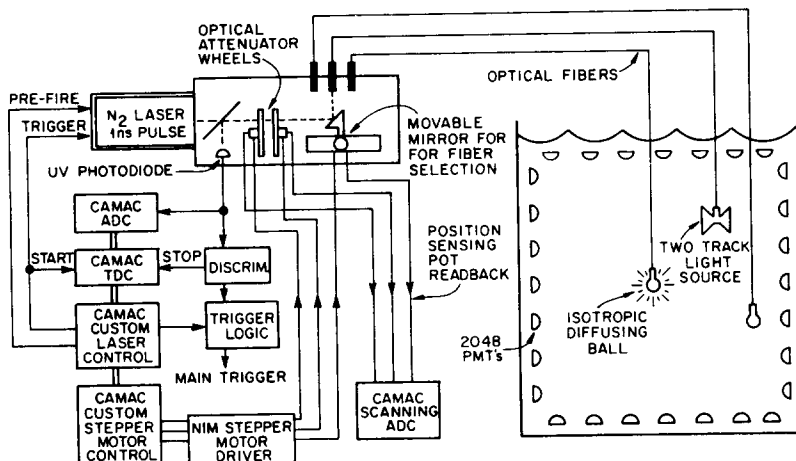


Fig. 18. Calibration hardware.

computer to a VAXstation 3200 in March 1990. The more powerful processor of the 3200, coupled with the more efficient mass storage of the 8 mm tapes, allowed replacement of the data acquisition computer by the MicroVAX II, the recording of every trigger on 8 mm tape, and addition of a second stage of (previously) off-line analysis in situ providing an additional factor 10–20 rejection in real-time.

11. Calibration

Although the same types of PMTs and electronics are used in all the detector's 2048 channels, the apparatus cannot be so finely tuned that all channels perform exactly the same. Differing PMT gains, operating voltages, signal propagation times, alignments in the Earth's magnetic field, and tolerances in electronic components introduce 10–20% channel-to-channel variations in timing and pulse-height response. It is necessary to compensate for these variations to make reliable analysis and interpretation of the data possible.

The IMB-1 calibration system [9] was built around a nitrogen laser and used largely unchanged through IMB-3. A dye laser was installed in 1989 to allow measurements to be made at a variety of wavelengths. Fig. 18 shows a schematic of the calibration hardware, while figs. 19 and 20 illustrate the laser and filter housing assemblies in greater detail. The dye (nitrogen) laser provides a ~ 3 ns (< 1 ns) pulse of 385 nm (337 nm) UV light. Part of the laser beam immediately splits off into a fast photodiode. This photodiode is monitored by a discriminator to verify that the laser has fired when requested, and by a 1 ns least-count TDC (time-to-digital converter), to allow precise mea-

surement of the time delay between firing request and discharge of the laser.

Light from the laser passes through two eight-filter wheels. The two wheels rotate independently under computer control, selecting one of 64 levels of attenuation, varying the light eventually reaching the PMTs over four orders of magnitude. The filter wheels are aligned slightly out of parallel with each other to prevent reflected light from following the primary beam. Attenuated laser light reflects into a 0.6 mm diameter quartz optical fiber, which winds through the same light-tight baffles as the PMT cables, and into the pool area. The fiber passes through an opening in the




Fig. 19. N₂ laser.




Fig. 20. Filter wheel assembly.

pool's floating floor, and then to a diffusing ball near the detector's center.

The diffusing ball (see fig. 21) is a 500 ml flask filled with a solution of Ludox (manufactured by Du Pont), and sealed from the water of the detector. Two parallel




Fig. 21. Diffusing ball.

mirrors and the suspended silica particles (21 nm diameter) in the Ludox solution combine to scatter the beam leaving the fiber into an omni-directional light pattern. The light output of the diffusing ball varies by about a factor of two depending on the angle from the fiber beam direction, but the calibration algorithms do not depend on knowledge of the exact distribution of light. The optical fiber is sheathed in opaque black tubing to protect the fragile strand from damage and to prevent light leakage into the detector before the fiber reaches the diffusing ball.

The laser operates in a remote trigger mode, which allows it to be fired under computer control by a custom-built CAMAC pulser. Two separate delays of up to 82 μ s may be programmed, in increments of 20 ns. A prompt pulse initiates precharging of the laser in preparation for firing. A delayed pulse, 40 μ s later, quickly (~ 150 ns) initiates a discharge. During normal operation, two different types of diagnostic triggers are available. In a "fixed" trigger, the detector is triggered by the discriminated and cable-delayed output signal of the laser photodiode. Hence, the laser light is always recorded by the PMTs in the pool at a fixed time (about 250 ns) before the global trigger. By contrast, a "ramped" trigger allows a second programmed output from the pulser to trigger the detector. Varying the time between the "fire" and "trigger" signals under computer control allows laser light to be recorded by the PMTs anytime before (during $\mathcal{T}1$) or after (during $\mathcal{T}2$) the detector trigger. The N_{tubes} and N_{patch} triggers used for data-taking are disabled during laser firings to prevent cosmic-ray muons from being recorded.

A set of more than 20 000 laser firings over a wide range of filter settings and trigger/firing delays is

recorded during a calibration run. The attenuation level is continually varied throughout the run to average any variation of the laser light output over all filter settings. The light output of the laser (in arbitrary units) is calculated during each run, and variation of laser intensity between runs (due to N_2 pressure variation, cleaning of laser head, or electrode wear) does not affect the calibration results.

PMT data taken during the calibration are analyzed on-line to determine timing intercepts, slopes, and pulse-height-dependent slewing for the $\mathcal{F}1$ and $\mathcal{F}2$ time-scales. Pulse-height readings and occupancy are used to establish the translation from digital counts to Poisson-distributed pe's. Cosmic-ray muon events recorded at the end of a calibration run are used to determine the conversion from pes to energy.

12. Operation and monitoring

The detector was originally designed to operate for 10 years, since it searches for very rare processes. Continuous maintenance of the physical plant has been performed throughout the experiment to insure the safety and stability of the experimental site. The electronic components of the detector have also been constantly monitored and maintained.

Diagnostic tools are required to allow occasional failures to be detected and fixed. An updating general status display provides most important information about the current run at a glance. Information in graphic and numeric form on every PMT and its electronic channel in the detector is available from the on-line calibration program. A custom real-time display package allows any calculated or recorded quantity associated with the data to be immediately displayed as a user-defined histogram or continuously updated screen display. Events in the data stream can also be displayed by a color plotter showing the position and timing of hits in perspective view identical to off-line displays. Results of on-line muon trajectory fits can be displayed to assess fitter performance and verify data quality. The status of the entire data handling pipeline can be displayed for debugging or study. All displays are available either on video terminals or hard-copy printout.

Hardware diagnostic programs exist to test most components of the data acquisition system, including 8-channel cards, digitizers, buffer memories, and the supervisor. Routines to diagnose the laser, trigger, and CAMAC systems also exist. Finally, there are generalized programs to allow individual operator-specified CAMAC instructions to be sent to any module in the system and the module response monitored.

When necessary, most operation and diagnosis of the experiment can be carried out remotely using a

modem. During the week, technicians and one or more physicists are on-site eight hours per day, to maintain and monitor the experiment. They are also assigned the task of overnight and weekend monitoring, and are available to enter the lab and correct problems which may develop during these times. All detector operations and displays are initiated by short mnemonic commands which then effect the desired result without further human intervention. In this way, detailed knowledge of the on-line code is not required for routine operation.

Ideally, the experiment would collect data continuously. In practice, of course this is not possible for many reasons. Some causes of downtime are merely detector calibration and maintenance, without which the data would be worthless. Most downtime is due to malfunctions which occur overnight, during holidays, or at other times when the laboratory is inaccessible. Power shutdowns by the mine management and tapes running out while the lab is inaccessible are causes of downtime over which we have no control. Other problems may not prevent the detector from running, but may render the data collected unusable. Typically, an electronic component whose failure is not immediately apparent is involved. The duty factor for the collection of high quality data achieved by the IMB-3 experiment is 0.47.

13. Conclusion

At the end of March 1991, the data-taking with the IMB-3 phase was terminated, due to a water leak which could not be stemmed without emptying the detector. Options for resumed operations after a further upgrade of the detector's capabilities are being studied. Notwithstanding a possible "IMB-4" at some time in the future, it seems appropriate to conclude with a review of the past 10 years' accomplishments.

IMB was the world's first very large water Cherenkov detector. At the time of its inception, it was the only device capable of testing newly formulated theories of grand unification: It was made possible by rapidly advancing electronics technology, particularly large area photomultipliers. Other advances in mining engineering, geotextiles, and water filtration helped keep the cost of construction down.

A review of the initial proposal reveals that most physics goals were met or exceeded. Roughly 2.5 live-years of useful data were collected during the IMB-3 stage. When combined with 1.3 live-years from IMB-1 and IMB-2, sensitivity to proton decay into $e^+\pi^0$ and other low-background modes will reach the level of 10^{33} years. For lower-energy modes and those with greater background, the sensitivity will be a factor of 3–10 less. Data from IMB-1 [10] and 2 [11], plus the

first half of IMB-3 which has been fully analyzed [12], give no compelling evidence that nucleon decay has been observed. Approximately 2000 neutrino interactions (both contained and outside the detector) spanning neutrino energies from 20 MeV to ~ 1 TeV, have been collected. Neutrino oscillations have been studied in regions of parameter space inaccessible to traditional accelerator and reactor experiments [13]. The added sensitivity of IMB-3 has also allowed classification of events by neutrino flavor with good efficiency [14], the fortuitous appearance of SN1987 allowed detection of neutrinos from stellar collapse [15], and with better than expected capability. This observation contributed to many unique results in both particle physics and astrophysics. Limits on $n-\bar{n}$ oscillation [16], magnetic monopoles [17], dark matter [18], galactic supernovae [19], and neutrino [20] and muon [21] point sources have also been set.

In summary, IMB has brought into successful operation the first and largest of a new class of detectors, solving many of the problems of operation, calibration, maintenance, and particularly detector simulation, display of events, and data analysis. The method of imaging Cherenkov rings and interpreting them has created a new paradigm for visualization of particle tracks. Since September 1982, this technique has been refined and applied to a wide variety of problems by Kamiokande, HPW, DUMAND, SNO, and other instruments under test or proposal for operation in mines, lakes, oceans, and Antarctic ice. The large-scale water Cherenkov approach seems likely to remain the method of choice for nucleon decay and astrophysical neutrino detection for some time to come.

Acknowledgements

It is a pleasure to acknowledge our hosts at Morton International's Fairport Mine, who have always gone out of their way to accommodate the unusual needs and schedules of our group. Also worthy of special praise are the on-site technicians who have operated and maintained the detector these ten years: Joe Reese, Ted Darden, Robert Render, Cheryl Gleason Mailander, John Hise, and Joe Bastulli. We also wish to thank Karl Luttrell for reliable technical support over the lifetime of the experiment. This work was supported in part by the United States Department of Energy.

References

- [1] M. Goldhaber et al., Proposal for a Nucleon Decay Detector, unpublished (1979);
G.W. Foster, Harvard University Ph.D. Thesis (1983);
B.G. Cortez, Harvard University Ph.D. Thesis (1983);
C.R. Wuest, University of California, Irvine Ph.D. Thesis (1983),
E.L. Shumard, University of Michigan Ph.D. Thesis (1984);
H.S. Park, University of Michigan Ph.D. Thesis (1985);
R. Svoboda, University of Hawaii Ph.D. Thesis (1985);
G. Blewitt, California Institute of Technology Ph.D. Thesis (1986);
T.J. Haines, University of California, Irvine Ph.D. Thesis (1986)
- [2] H. Georgi and S.L. Glashow, Phys. Rev. Lett. 32 (1974) 438;
P. Langacker, SLAC-PUB-2544 (1984).
- [3] H.S. Gurr et al., Phys. Rev. 158 (1967) 1321.
- [4] C.R. Wuest et al., Nucl. Instr. and Meth. A239 (1985) 467.
- [5] S. Seidel, University of Michigan Ph.D. Thesis (1987).
- [6] S.T. Dye, University of Hawaii Ph.D. Thesis, (1988);
M. Mudan, University College, London Ph.D. Thesis (1989);
D. Casper, University of Michigan Ph.D. Thesis (1990);
R. Becker-Szendy, University of Hawaii Ph.D. Thesis (1991).
- [7] M. Crouch et al., Phys. Rev. D18 (1978) 2239.
- [8] R. Claus et al., Nucl. Instr. and Meth. A261 (1987) 540.
- [9] R.M. Bionta et al., Proc. of the Fermilab Workshop on Calorimeter Calibration (1983) unpublished.
- [10] H.S. Park et al., Phys. Rev. Lett. 54 (1985) 22;
G. Blewitt et al., Phys. Rev. Lett. 55 (1985) 2114;
T.J. Haines et al., Phys. Rev. Lett. 57 (1986) 1986.
- [11] S. Seidel et al., Phys. Rev. Lett. 61 (1988) 2522.
- [12] R. Becker-Szendy et al., Phys. Rev. D42 (1990) 2974.
- [13] R.M. Bionta et al., Phys. Rev. D38 (1988) 768;
R. Becker-Szendy et al., Phys. Rev. Lett. 69 (1992) 1010.
- [14] D. Casper et al., Phys. Rev. Lett. 66 (1991) 2561;
R. Becker-Szendy et al., to be published in Phys. Rev. D. (November 1, 1992).
- [15] R.M. Bionta et al., Phys. Rev. Lett. 58 (1987) 1494;
C.B. Bratton et al., Phys. Rev. D37 (1988) 3361.
- [16] T.W. Jones et al., Phys. Rev. Lett. 52 (1984) 720.
- [17] S. Errede et al., Phys. Rev. Lett. 51 (1983) 245.
- [18] J.M. LoSecco et al., Phys. Lett. B188 (1987) 388.
- [19] S.T. Dye et al., Phys. Rev. Lett. 62 (1989) 2069
- [20] R. Svoboda et al., Astrophys. J. 315 (1987) 420.
- [21] R.M. Bionta et al., Phys. Rev. D36 (1987) 30;
R. Becker-Szendy et al., Phys. Rev. D43 (1990) 1413