FERMILAB-PUB-19-624-E-ND

The Effects of Thermal Cycling on Two Light Guide Technologies and Two Readout Technologies in Liquid Argon

S. Mufson[a], B. Adams[a], B. Baugh[a], B. Howard[a], C. Macias[a], G. Cancelo[b], E. Niner[b], D. Totani[b]

[a]*Indiana University, Bloomington, Indiana 47405, USA*
[b]*Fermi National Accelerator Laboratory, Batavia, Illinois 60510, USA*

December 13, 2019

Abstract

In this investigation two light guide technologies and two readout technologies are tested for the effects of thermal cycling in liquid argon (LAr). The two light guide technologies are the technology developed at Indiana University [1] and the technology developed at MIT and implemented at Fermilab [2, 3]. The two readout technologies are passively and actively ganged 12 Hamamatsu MPPCs. The experiment, mounted in the TallBo dewar facility at Fermilab, measured the scintillation light from through-going, minimum ionizing cosmic muons. Thermal cycling of the experiment's components occured when the dewar was drained after five weeks of running, refilled, and then run again for an additional four weeks. The results show that the MPPCs were unaffected by thermal cycling and there is good evidence that both ganged readout technologies were similarly unaffected. Since the experiment showed that there was a clear degradation of the cosmic signal when the LAr was drained and refilled, it is concluded that it was both light guide technologies that were degraded. Quenching of the scintillation light by contaminants and path length differences in the cosmic tracks through the LAr were considered as causes and but did not play a role.

Keywords:

[☆]FERMILAB-PUB-?, arXiv: 1912.05987

liquid argon scintillation, neutrino detectors, photon detection

1. Introduction

As charged particles traverse large liquid argon (LAr) detectors, the scintillation light they generate provides important information for investigations of neutrinos and dark matter. To collect these scintillation photons for analysis, different technologies have been developed for the single phase far detector of the Deep Underground Neutrino Experiment (DUNE) [4]. Two of these light guide technologies were studied in this experiment – one developed at Indiana University [1] and one developed at MIT and implemented at Fermilab [2, 3]. Both of these technologies make use of multiple silicon photomultipliers (SiPMs) in their readout. Since the individual readout of multiple SiPMs is quite costly, two different ganging schemes were also studied in this experiment – one in which the SiPMs were ganged by passively connecting the photodetectors in parallel and one in which the SiPMs were actively summed electronically. The potential value of an active summing scheme over passively ganging SiPMs is reduced capacitance and noise.

The question addressed here asks what effects, if any, does thermal cycling have on this subset of technologies developed to detect the scintillation light in LAr. The experiment took place in the TallBo liquid argon dewar at Fermilab and studied single through-going, minimum ionizing cosmic muons. Thermal cycling resulted when the dewar was drained after five weeks of running to remount the detectors, refilled, and then run again for an additional four weeks. Since scintillation light can be significantly quenched by contaminants in the LAr [5–9], the contaminants were carefully monitored during the experiment to assess their effects on the outcomes.

2. The TallBo Experiment

The experiment took place in the liquid argon dewar facility “TallBo”, housed in the Proton Assembly Building (PAB) at Fermi National Accelerator

Laboratory (FNAL). The experiment ran from August 3, 2018 until October 17, 2018. Data collection took place in 4 separate runs (runs 1-4), which were defined by the detector configuration read out by the data acquisition.

2.1. Experimental Design

Fig. 1 shows a schematic of the experimental design in runs 2 and 3. Two

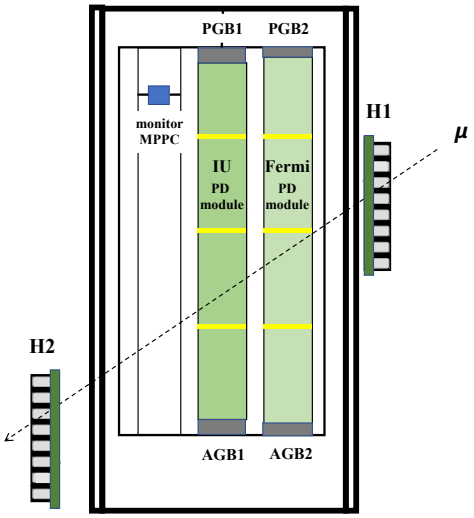

Figure 1: A schematic of the experimental apparatus in the TallBo dewar during run 2 and run 3. Two photon detector (PD) modules, "IU" and "Fermi", are mounted in a custom frame suspended from the dewar lid. Each end of the PD modules was read out by 12 (MPPC) photodetectors ganged together into one readout channel. At the top, the MPPCs were ganged passively (PGB1, PGB2); at the bottom the MPPCs were actively summed (AGB1, AGB2). A monitoring MPPC was mounted in the position shown. Four-fold coincidence triggers were provided by the adjacent hodoscope trigger paddles H1 and H2 that flank the outside of the dewar. A representative triggered cosmic muon trajectory is superposed.

photon detector (PD) modules, one developed at Indiana University [1] and one developed at MIT and implemented at Fermilab [2, 3], are mounted in a custom frame suspended from the dewar lid. Each end of the two PD modules was read out by 12 Hamamatsu SiPM (MPPC) photodetectors ganged together into one readout channel. Two schemes were used to gang the MPPCs, passive and

active. The passive ganging boards read out the modules from the top of the dewar and the active ganging boards read out the modules from the bottom. In the final data-taking period of the experiment, the IU and Fermi modules were exchanged. The readout boards, however, remained in the configuration shown during all runs. A single MPPC, mounted in the position shown, was included on the frame to monitor photodetector performance characteristics during all runs in the experiment. Four-fold coincidence triggers that define single minimum ionizing cosmic muons passing through the detectors were provided by adjacent hodoscope trigger paddles (H1 and H2) that flank the outside of the dewar. A representative single cosmic ray muon track that triggers the readout is also displayed in Fig. 1.

2.2. Light Guide Technologies

Light guide technologies developed at Indiana University [1] and Fermilab/MIT [2, 3] were used in the experiment. Both technologies have been previously described.

In the IU light guide technology [1], scintillation photons from liquid argon at 128 nm strike one of four acrylic plates with the wavelength shifter TPB embedded in their surfaces. The TPB in the struck plate converts VUV scintillation photons to visible photons, typically in the range 420 – 450 nm. These photons are transmitted through the plate and are subsequently absorbed by a commercial light guide made by Eljen Technology. In the Eljen light guide, photons are again wavelength shifted to the range 480 – 510 nm and channeled to the Hamamatsu MPPCs at the end, where they are detected.

The Fermi light guide technology is described in [2, 3]. These light guides are manufactured from cast acrylic bars that have wavelength shifter (TPB) embedded into their surface using a dip-coating technology. Before the wavelength shifter was applied, the acrylic bars were annealed. Scintillation photons from liquid argon are absorbed by the TPB and wavelength shifted to the range 420 – 450 nm. These waveshifted photons are channeled to the MPPCs on the ends of the light guides, where they are detected.

2.3. Photodetectors

The SiPMs used in the experiment, including the monitor MPPC, were Hamamatsu 6x6mm MPPCs (S13360-6050VE[1]). These MPPCs have 50μm pixels in a TSV package and are coated with epoxy resin. Fig. 2 shows the dark

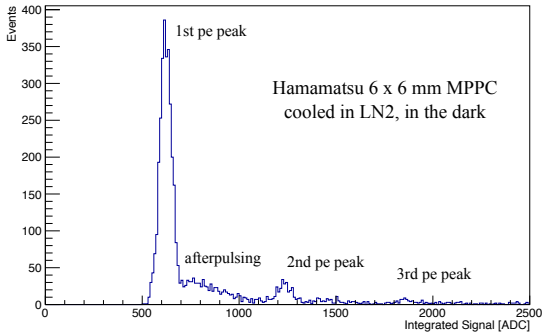

Figure 2: The dark spectrum of a previously uncooled MPPC in LN2 that has been biased at 44.5 V as measured in the lab. The single p.e. peaks are marked. Also marked is the signal due to afterpulsing.

spectrum of a previously uncooled MPPC in LN2 that has been biased at 44.5 V. The first, second, and third photoelectron (p.e.) peaks are clearly visible. Also labeled is the signal from afterpulsing. Afterpulsing is a second avalanche in the same pixel as the primary avalanche. They are smaller than a standard avalanche because they occur before the cell can fully recover. Afterpulsing can seriously blur primary event signals when it becomes significant.

The MPPCs in this experiment were biased at 44.5 V, a compromise that minimizes afterpulsing while maximizing the amplitude of the first p.e. peak. This bias voltage is close to the 3 V over breakdown recommended by the Hamamatsu data sheets for operation at room temperature. The breakdown voltage for the MPPCs is 42 V, determined by plotting the "pseudo-gain", as a function of bias voltages ranging from 43V - 47.5V, as shown in Fig. 3 for 4 previously uncoold MPPCs. The pseudo-gain is the difference in the number of

[1] https://www.hamamatsu.com/resources/pdf/ssd//s13360-2050ve_etc_kapd1053e.pdf

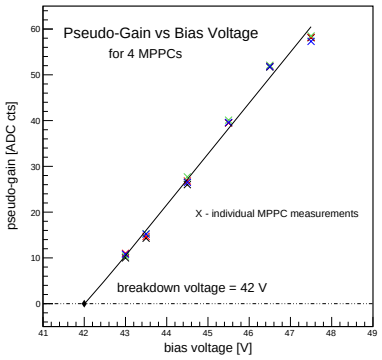

Figure 3: The pseudo-gain, or the difference in the number of ADC counts between the first and second p.e. peaks, as a function of bias voltage for 4 previously uncooled MPPCs. The x intercept of a least-squares straight line fit to the 4 pseudo-gain measurements determines the breakdown voltage.

ADC counts between the first and second p.e. peaks, a quantity that is related to the MPPC gain by a multiplicitive constant. The breakdown voltage was determined as the x intercept of a least-squares straight line fit to the mean of the pseudo-gain measurements for the 4 MPPCs as a function of bias voltage.

Laboratory measurements of the the performance characteristics of 12 previously uncooled Hamamatsu MPPCs immersed in LN2 are shown in Fig. 4. The MPPCs were biased at 44.5 V. Superposed on these histograms are 3 independent laboratory measurements of the TallBo monitor MPPC (c.f., Fig. 1) that were made after the experiment was completed. The description of the measured characteristics are with reference to Fig. 2. The dark noise (a) was computed as the sum of the ADC counts in the dark spectrum divided by the data acquisition time. The MPPCs run quietly in LN2 with rates in the range of 5-25 Hz, compared with their MHz dark rates at room temperature. The first p.e. peak (b) is the number of ADC counts at the peak of the first p.e. peak. This calibration of the individual MPPCs showed variations of approximately 15%. The "pseudo-gain" (c) is the difference in the number of ADC counts between the first and second p.e. peaks, a quantity that is related to the gain by a

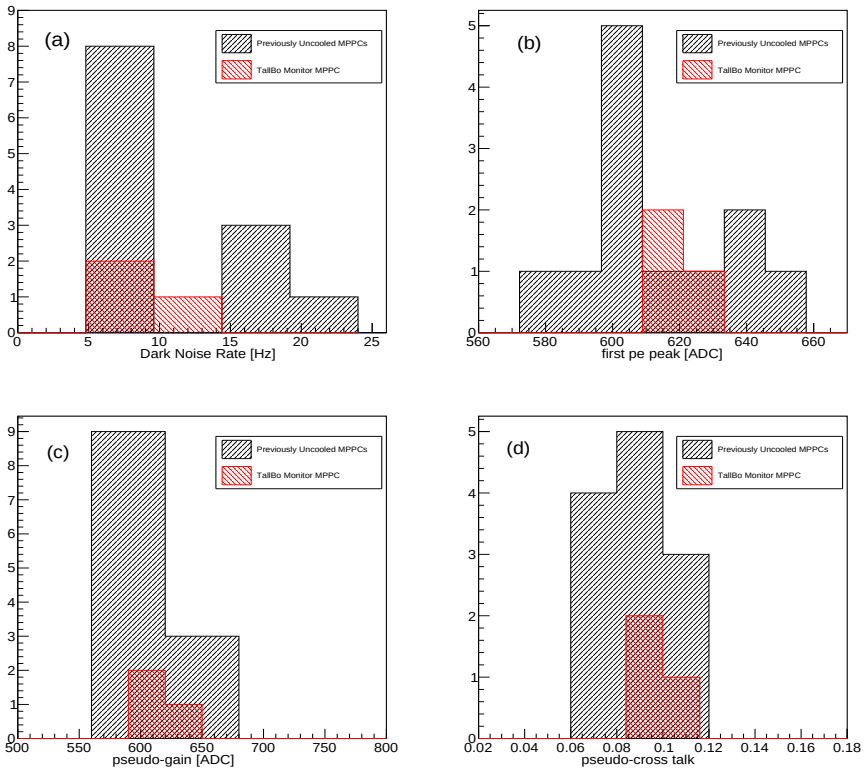

Figure 4: Laboratory measurements of the the performance characteristics of 12 previously uncooled Hamamatsu MPPCs immersed in LN2 biased at 44.5 V. Superposed on these histograms are 3 independent laboratory measurements of the TallBo monitor MPPC (c.f., Fig. 1). (a) The dark noise, characterized as the sum of the ADC counts divided by the data acquisition time; (b) the number of ADC counts at the peak of the first p.e. peak; (c) the pseudo-gain or the number of ADC counts between the first and second p.e. peaks; and (d) the pseudo-cross talk determined as the ratio of the number of ADC counts in the peak of the second p.e. peak to the number of ADC counts in the peak of the first p.e. peak.

multiplicitive constant. Cross talk events occur when a photon emitted during the electron avalanche in one pixel is re-absorbed by another pixel elsewhere on the SiPM and induces a second avalanche in immediate coincidence with the first. The "pseudo-cross talk" (d) is the ratio of the number of ADC counts in the peak of the second p.e. peak to the number of ADC counts in the peak

of the first p.e. peak and estimates the contribution of cross-talk to the event signal. It is typically in the 10% range.

One important conclusion from Fig. 4, as discussed below, is that the characteristics of the monitor MPPC are consistent with the characteristics of the test devices that had never been cooled. This implies that the monitor MPPC was not significantly affected by the environmental conditions it experienced during the experiment.

2.4. Readout

There are 12 ganged MPPCs on each end of the light guides. By ganging the MPPCs into one readout channel, the number of electronics channels and associated cables will clearly be significantly reduced, which in turn implies large reductions in electronics costs, an important consideration in large detectors. There were two different ganging board designs for the MPPCs tested in the experiment – passive and active.

The signals from the readout boards were processed by a 12-channel SiPM Signal Processor (SSP) module that was designed and built by the HEP Electronics Group at Argonne National Laboratory (ANL). The SSP has been described in detail in [10].

2.4.1. Passive Ganging Boards

The passive ganging boards, PGB1 and PGB2 in Fig. 1, connect the 12 MPPCs in parallel. MPPCs have a large output capacitance ($\sim$1 nF for a 6mm x 6mm array) and when operated in Geiger mode require a bias voltage in the tens of volts. Parallel ganging increases the output capacitance proportional to the number of MPPCs connected in parallel but leaves the MPPC bias voltage unchanged. Although the larger capacitance reduces the signal to noise ratio of the output, the integrated signal from the 12 MPPCs on the board remains constant. The larger capactance also decreases the signal to noise on the waveforms during the increased signal collection time, particularly important for weak events. Panel (a) in Fig. 5 shows the baseline subtracted mean of 25

waveforms from typical single track muon events crossing TallBo as read out by the passive board PGB1 during run 2. The waveforms from the readout board PGB2 on the Fermilab technology are qualitatively very similar. The increased capacitance is seen in the long exponential tail on the waveform of ≈ 6.5 μs. Single MPPCs have an exponential tail more typically ≈ 500 ns. The artifact at the waveform minimum is thought to be the result of mismatches in impedance at the board-SSP coupling.

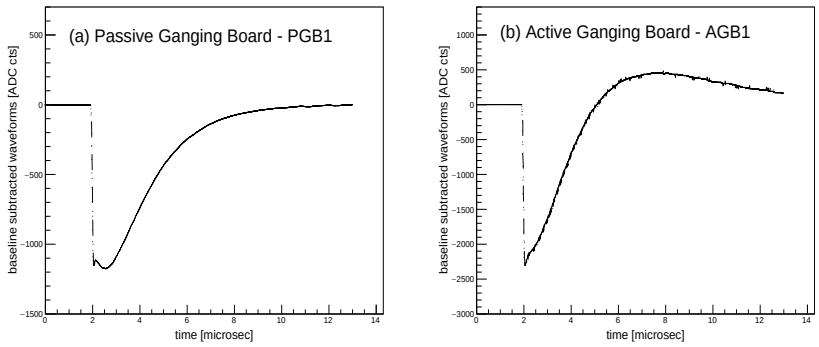

Figure 5: The baseline subtracted mean of 25 waveforms from typical single track muon events crossing TallBo and which pass all analysis cuts as read out by (a) passive ganging board PGB1 and (b) active ganging board AGB1 (Fig. 1) during run 2. In these waveforms each bin is 6.67 ns wide and the waveforms shown are $\sim$10.7 μs in length. The waveforms from the readout of the Fermilab technology are qualitatively very similar.

2.4.2. Active Ganging Boards

One method to add the signals from the 12 MPPCs without increasing the capacitance as significantly is to use an active ganging board that implements an active summing node like the circuit shown in the *top* panel of Fig. 6. In this circuit the negative input of the OpAmp provides a virtual ground that decouples the signals coming on each input. The gain of the amplifier can then be adjusted by the ratio between the feedback resistor and Ra. There are several challenges to designing this active summing node: it must work in the cold (87K); it must not boil the LAr by dissipating too much heat; it must be

high bandwidth to keep the MPPC signal integrity; and it must not contribute significantly to the total noise budget. The circuit used in this experiment employed a Texas Instruments THS4131ID[2] OpAmp which has 150 MHz of bandwidth for unitary gain and 1.3 nV/$\sqrt{(\mathrm{Hz})}$ of noise. In this circuit the Op Amp summed 2 rows of 6 MPPCs each. The active summing board is shown in the *bottom* panel of Fig. 6.

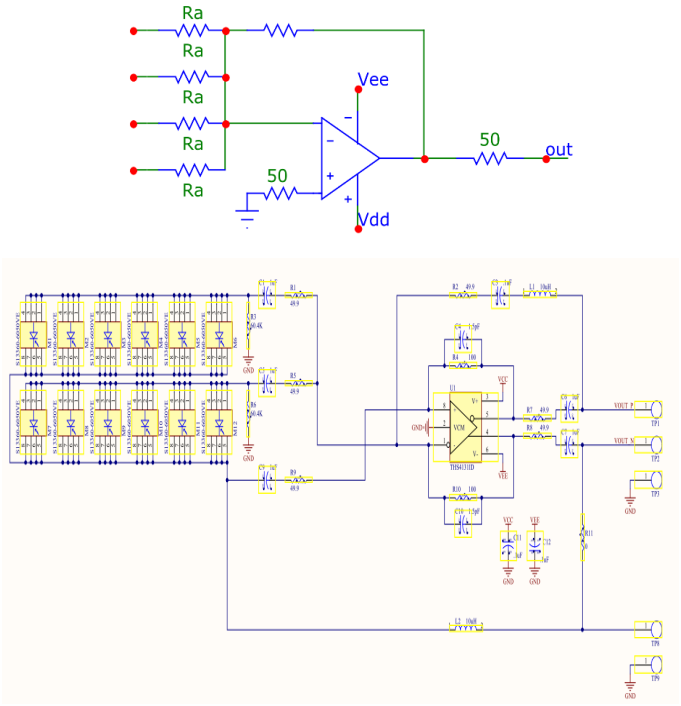

Figure 6: *top*: The active summing node used in the design of the active ganging board; *bottom*: the active ganging board.

Panel (b) in Fig. 5 shows the mean of 25 waveforms from typical single track muon events crossing TallBo as read out by the active board AGB1 (Fig. 1) during run 2. The waveforms from the readout of the Fermilab technology are qualitatively very similar. Clearly this board introduces a significant overshoot

[2]https://www.ti.com/store/ti/en/p/product/?p=THS4131ID

in the waveforms. The overshoot is the most problematic artifact when analyzing the summed waveforms. Its origin is unclear and it does impact the accuracy of the results determined by the data collected by these active ganging boards. An artifact is also seen at the waveform minimum from the active ganging board and is again the likely result of mismatches in impedance at the board-SSP coupling.

2.5. Hodoscopes

Two hodoscope modules were installed on opposite sides of the TallBo dewar to select single-track cosmic-ray muons passing through the LAr volume [10]. Each hodoscope module consists of 64 2-inch diameter barium-fluoride crystals, coated with TPB and arranged in an 8×8 array. Each crystal is monitored by a 3″ PMT. Since the hodoscope modules were originally designed to detect bremsstrahlung photons in the CREST balloon flight experiment, they are very sensitive to extraneous photon activity around our experiment. To reject this γ ray activity, a pair of plastic scintillator panels covering the entire face of a hodoscope module were placed between each hodoscope module and the TallBo dewar. That is, there were 4 total scintillator panels, two on each side of the TallBo dewar. These 4 panels were individually read out by PMTs salvaged from the QuarkNet program at FNAL. The SSP readout was triggered by four-fold coincidence logic that required at least one hit in each hodoscope module on opposite sides of the dewar, as well as one hit in the adjacent scintillator plane. Typically this trigger indicates that each event contains at least one charged particle passing through the LAr. Events were further filtered offline to reject showers by requiring one and only one hit on a PMT in each hodoscope module. Single-track events crossing from one side of the frame to the other were rejected in order to exclude any tracks that could pass through a light guide.

After their installation adjacent to the TallBo dewar, one of the PMTs on one of the scintillator panels failed which made the scintillator module inoperable. The PMT could not be replaced because spares were not available. Therefore, the experiment collected tracks on either the “front side” or the “back side” of

the detector frame by employing only one scintillator panel on either side of the detector frame in the four-fold coincidence trigger during separate runs. The designations "front side" and "back side" are purely descriptive and are not meant to imply any asymmetries in the two sides of the detector.

3. Operations

To prepare for a run, the TallBo dewar was first evacuated by a turbo pump to help reduce contamination from residual gases. It was then back-filled with gaseous argon. The gaseous argon was next replaced with ultra-high-purity (UHP) LAr that passed through a molecular sieve and a copper filter that had been regenerated just prior to the run. The volume of LAr in TallBo was approximately 460 liters.

Once filling was complete, the TallBo dewar was sealed and subsequently maintained at a positive internal pressure of 8 psig to assure no contamination from the outside. Gaseous argon from the ullage was recondensed to liquid argon with a liquid nitrogen condenser and then returned to the dewar.

3.1. Run Rates and Selection

The experiment consisted of recording the scintillation light from single track minimum ionizing cosmic ray muons traversing the LAr in the TallBo dewar.

There were four periods of running, defined by the detector configuration read out. The dates of these runs are given in Table 1. The dewar was actually filled on 7/5/18 but problems with the new LAr filling system prevented data taking from beginning until 8/3/18. The first two runs differed only by cable switches. There were actually two monitor photosensors on the frame of two different types: the Hamamatsu MPPC shown in Fig. 1 and a SensL SiPM like those used in our previous experiments at PAB [1, 10]. During data analysis, it was concluded that the SensL SiPM did not contribute useful information but this was not established during the experiment. At the outset of data-taking, one of the channels assigned to the monitor SiPMs failed. Between run 1 and

Table 1: Experimental Runs

Run	Tracks Analyzed	Center Detector	Run Start	Run End
1	front	IU	Aug 3, 2018	Aug 19, 2018
2	front	IU	Aug 20, 2018	Aug 27, 2018
3	back	IU	Aug 28, 2018	Sept 9, 2018
drain/refill				
4	front	Fermi	Sept 18, 2018	Oct18, 2018

run 2 the readout channel for the two different monitor SiPMs were exchanged, which resulted in a cable switch. Both of these runs read out front side tracks. For run 3 the hodoscope scintillator paddles were switched so that back side tracks were read out. This switch was made to ascertain whether front side and back side tracks gave the similar results, as was expected. In addition, the monitor SiPM cables were again exchanged.

For run 4, the positions of the IU and Fermi PD modules were exchanged. The purpose of this switch was to determine whether the geometric placement of the light guides affected the results. There was a concern that the scintillation signal from cosmics would be affected by the position of the light guides with respect to the reflecting dewar walls. To prepare for run 4, the dewar was first drained, the light guide positions were switched, and the dewar then was refilled. The readout boards remained in their original positions to keep the light guide response independent of the board response. In addition, the hodoscope scintillator paddles were switched back to read out front side tracks. Draining and refilling the dewar was the point at which thermal cycling occurred.

The event statistics are given in Table 2. Column [2] gives the total number of hours of run time for the runs. The number of four-fold coincidence triggers per hour for each run are given in column [3]. Column [4] gives the number of four-fold coincidence triggers per hour in which there was one and only one hit

on a PMT in each hodoscope module. These single track events were the ones analyzed. Table 2 shows clearly that there was an abrupt change in the trigger rates between run 1 and run 2, run 3, and run 4. This change in the rates is discussed below.

Table 2: Run Statistics

Run	Run Time [hr]	4-fold Coinc. Rate [Hz]	Single Track Rate [Hz]
1	343	0.085	0.024
2	333	0.062	0.015
3	203	0.056	0.013
drain/refill			
4	612	0.064	0.016

Each run in Table 2 consisted of a series of 24-hour subruns. Fig. 7 shows the superposition of a histogram for the single track rates for the subruns in run 1 and a histogram for the single track rates for the subruns in runs 2, 3, and 4.

Fig. 7 shows that the individual subrun rates for run 1 were significantly higher than those for runs 2, 3, and 4. The increase in the rates seen in Table 2 is not the result of a few hot runs. There was nothing obviously different in the environment in PAB during run 1. Since the rates change was seen after cable switches and a DAQ restart, probably something varied in the electronics. But truthfully the cause is not known.

3.2. Contaminants

The contaminants that most affect LAr scintillation light are O_2, N_2, and H_2O, which can both quench scintillation light and decrease the argon transparency at 128 nm [5–9]. The LAr delivered to TallBo is UHP LAr and is typically delivered with low levels of these contaminants. The UHP LAr is

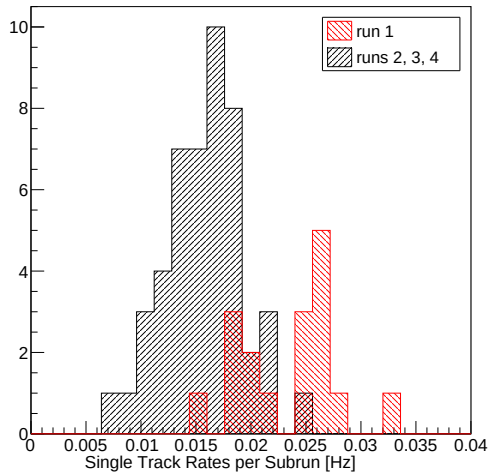

Figure 7: The superposition of a histgram for the single track rates for the subruns in run 1 and a histogram for the single track rates for the subruns in runs 2, 3, and 4.

pumped into TallBo through filters which are very effectve at removing O_2 and H_2O. N_2 is not removed by the filters. Before this experiment, the filters were regenerated, thus becoming very efficient at further removing the O_2 and H_2O from the UHP LAr delivered.

The O_2, N_2, and H_2O concentrations were monitored during the run. The concentrations of these contaminants are shown in Fig. 8. In addition, Fig. 8 shows the liquid argon level relative to the top of the detector frame. Only during brief excursions was the top of the frame above the liquid argon level. The high N_2 readings during run 4 are likely due to the fact that the LAr was delivered with much higher than typical N_2 contaminataion, although its concentration falls within specs. The O_2 were high in run 3 and the filters were regenerated after the dewar was drained. As shown below the O_2 contamination does not affect the results.

The probability that a scintillation photon is absorbed by a contaminant along its path is given by $P_{abs} = A \exp(-L/\lambda_{abs})$, where A is a normalization constant, L is the track length to the intersection, and λ_{abs} is the absoption

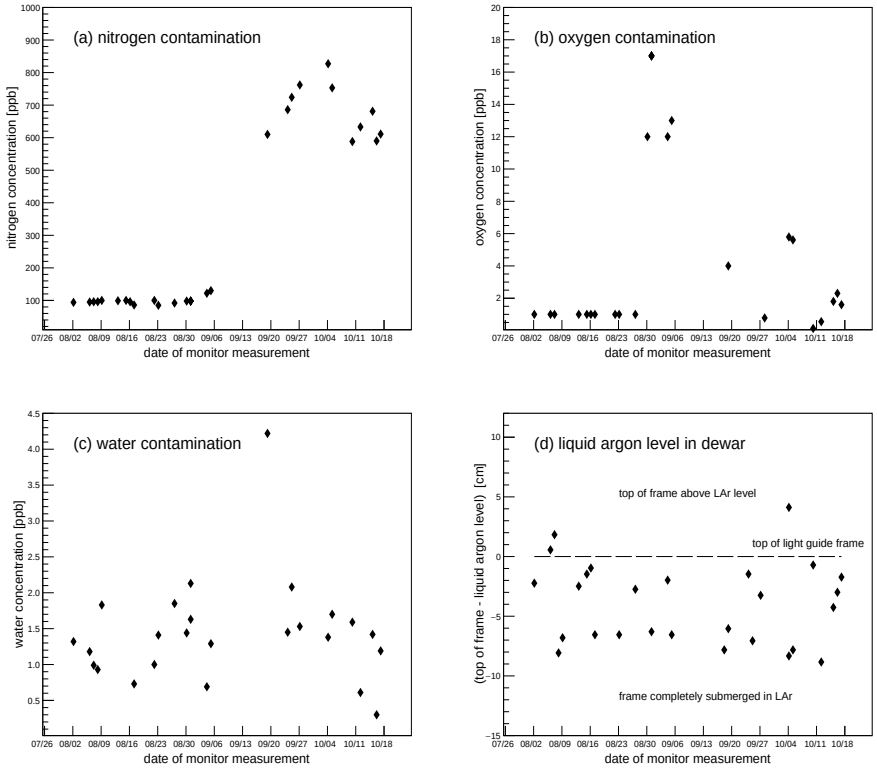

Figure 8: The contaminants in the liquid argon and the LAr level in the dewar: (a) The nitrogen contamination; (b) the oxygen contamination; (c) the H_2O contamination; and (d) the LAr level in the dewar.

length [1]. The absorption length is calculated by

$$\lambda_{abs} = 1/n_c\sigma = 1/[< \chi > \cdot n_{LAr} \cdot \sigma], \tag{1}$$

where n_c = the number density of contaminants, σ is the cross section for absorption of 128 nm VUV photons, $< \chi >$= mean fractional number of contaminant molecules, and n_{LAr} = the number density of LAr atoms/cm^3 = (1.396 g/cm^3)/(39.948 g/mol) $\times$ (6.02×10^{23} atoms/mol) [11] = 2.1037×10^{22} LAr atoms/cm^3. Table 3 gives σ for the contaminants [1]:

Table 4 gives $< \chi >$ and λ_{abs} for the contaminants measured for the different runs. The means were computed from the measurements shown in Fig. 8. The

Table 3: Absorption cross sections for contaminants.

contaminant	σ [cm^2]	Ref.
N_2	7.1×10^{-21}	[5]
O_2	2.8×10^{-19}	[7]
H_2O	8.0×10^{-18}	[9]

one outlier H_2O measurement is well outside the range of all others and is thought to be due to a glitch in the monitor at an LAr topoff. This level of H_2O was not repeated during the experiment and was not used in computing $< \chi >$ for H_2O.

Table 4: Concentration and absorption length for contaminants.

contaminant	run 2	run 3	run 4
N_2:			
$< \chi >$ [ppb]	95.0	109.2	727
λ_{abs} $[m]$	7.0×10^2	6.1×10^2	9.2×10^1
O_2:			
$< \chi >$ [ppb]	0.8	14.2	4.0
λ_{abs} $[m]$	2.0×10^3	1.2×10^2	4.2×10^2
H_2O:			
$< \chi >$ [ppb]	1.3	1.4	1.3
λ_{abs} $[m]$	4.7×10^1	4.1×10^1	4.5×10^1

At these concentrations, the N_2 is not expected to affect the scintillation light [5, 12]. At concentrations <100 ppb, the effects of O_2 contamination on scintillation light in LAr are thought to be negligible [6]. The effects of H_2O contamination on LAr scintillation light are not well studied. However, studies in gaseous argon suggest the H_2O concentration at this level does not affect our

results [8]. Simulations described below confirm these expectations.

4. TallBo Simulation

The track simulation through TallBo follows closely the simulation described in [1] and details can be found there. Briefly, in this simulation cosmic muons are treated as traveling along straight paths with end points fixed at the centers of the two triggered PMTs in the hodoscopes on either side of the dewar. It is assumed that the cosmic muons are minimum ionizing particles with an energy deposition in LAr of 2.105 MeV/cm [11] and that 40,000 photons/MeV [13–16] are generated along their tracks. There are consequently 8.42×10^4 photons/cm scintillation photons emitted by the muons as they traverse the LAr.

The starting positions of the scintillation photons were distributed uniformly along the track segment passing through the LAr volume and the photons' momentum vectors were distributed uniformly in solid angle. Photons were tracked along straight line paths until they intersected with either an IU or Fermi photon detector, the dewar wall, or were lost. Along their tracks, photons could also undergo a Rayleigh scattering or reflect off the walls of the dewar. Photons could be lost if they were absorbed by a contaminant, or the dewar wall, or if they hit the top of the liquid level. Every track in the TallBo data set that struck two single hodoscope PMTs and passed data selection cuts was simulated.

The path length distribution of all the scintillation photons that strike the IU or Fermi PD and the positions of all the scintillation photons that strike the IU or Fermi photon PD were stored in histograms. These histograms were used for two purposes. (1) The histogram of the photon path length distribution was used to determine whether the contaminants in Table 4 quenched the number of photons striking the detectors in any significant way. (2) The histograms of the positions of the photons on the IU and Fermi PDs were used to evaluate how the geometric position of the PD in the dewar (whether in the center or closer to the wall) affected the number of scintilation photons striking the detector.

The contamination calculation (1) was done independently for the three data runs included in the analysis. A photon path length was first drawn from the distribution of all photon path lengths simulated for that run. Three absorption lengths for that photon for the three contaminants were drawn as exponential deviates using ROOT's TRandom3 package and λ_{abs} from in Table 4. If any of the 3 simulated absorption lengths were shorter than the photon path length, the photon was assumed lost to absorption before it reached the detectors. The calculation was repeated 10^7 times for each run. The probability of absorption $\mathbf{P}_{abs}$ by any of the three contaminants in each run are given in Table 5. These

Table 5: Photon absorption probability by contaminants in Table 4.

absorption probability	run 2	run 3	run 4
$\mathbf{P}_{abs}$	0.008	0.009	0.011

calculations confirm that the level of contaminants in TallBo shown in Table 4 have only a minor effect on quenching the scintillation light before it reaches the detectors, as expected from the literature [5–9]. As important, the quenching is approximately equal during all 3 runs, implying this effect will not introduce any run-by-run differences in the results.

5. Results

5.1. Data Analysis

All events selected for analysis were four-fold coincidence triggers with at least one hit in each hodoscope module on opposite sides of the dewar. Each trigger was further required to have one and only one hit on a PMT in each hodoscope module and hits in its adjacent scintillator plane. If a straight line track between the PMT centers crossed the detector plane, the event was discarded. If a track was on the back side for runs 2 and 4 or the track was on the front side for run 3, the event was discarded. The sample of events selected

was assumed to be dominated by single minimizing muons passing through the liquid argon.

The data analysis began by filtering the waveforms of the selected events with an 11-point running mean. This procedure smooths out fluctuations along the waveform that can result from the decreased signal to noise due to ganging. As seen in Fig. 5, the readout retains $\sim$2 μs of data before the trigger and the mean of the pre-trigger data was used to calculate the waveform baseline. A cut was then applied on the standard deviation of the pre trigger samples about the baseline to remove anomalous waveforms with spurious shapes or unusually large fluctuations. These anomalous waveforms were found to correlate strongly with large values of the standard deviation. Fig. 9 shows the standard deviations of all waveforms from boards PGB1 and AGB1 from run 2. For both distributions, the

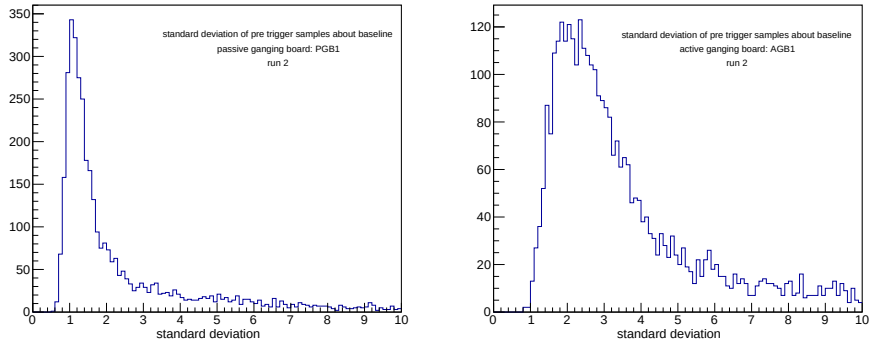

Figure 9: The distributions of the standard deviation of the pre trigger samples about the baseline for the passive readout board PGB1 and the active board AGB1 from run 2.

cut was made at a standard deviation of 2, which was found to eliminate the vast majority of anomalous waveforms. PGB2 and AGB2 have similar distributions and the cut was also made at a standard deviation of 2. Qualitatively similar histograms were found for the readout boards in runs 3 and 4, and the cut was again made at a standard deviation of 2.

One characteristic of the histograms in Fig. 9 is apparent. Comparing the two distributions, a larger fraction of the waveforms from AGB1 have large

standard deviations. These differences are also seen when comparing PGB2 and AGB2. In runs 3 an 4, the differences are often even more striking. This suggests that the active boards are often shaping and distorting the waveforms in significant ways, which is likely the result of amplifier noise . It also means that the standard deviation cut removes a far larger fraction of waveforms read out by the active boards than from the passive boards.

The selected waveforms were then integrated. The start point for the integration is the trigger at 2 μs. The end point for the integration, however, depends on whether the waveform was read out by a passive board or an active board. As shown in Fig. 5, the waveforms from the passively ganged boards and the actively ganged boards are distinctly different. For the passively ganged boards, the integral was extended out to 10.7 μs to collect all the charge information. It is not so clear, however, how far to extend the integral for the active boards. The main concern was whether to include in the integral the piece of the waveform after the waveform overshoots the baseline.

In this analysis the choice was made to integrate the waveforms from AGB1 and AGB2 only out to the overshoot. This choice reflects the fact that the overshoot was found to occur at a position (time) along the waveform that is approximately independent of signal strength, suggesting that the overshoot is an artifact of how the active ganging board shapes the waveform and the piece of the waveform after the overshoot does not add to the charge information in the waveform. A lab test of this hypothesis was made at IU after the experiment was completed. In these tests the boards were immersed in LN2 and then flashed with LED pulses of variable width. The assumption was made that the number of photons falling on the boards was directly proportional to the LED pulse width. The range in pulse widths reflect the transit times of a muon assuming the light is dominated by early light [10]. The waveforms from the passive boards were integrated out to 10.7 μs. The waveforms from the active boards were integrated out to the stable overshoot point. Since all charge information is contained in the integral of the passive waveforms, it should be linear with LED pulse width. If the charge information in the integrated waveforms from

the active boards is a constant fraction of the total contained in the integrated wavefroms from the passive boards, it should also be linear with LED pulse width and have the same slope. Fig. 10 shows this is indeed the case. It was

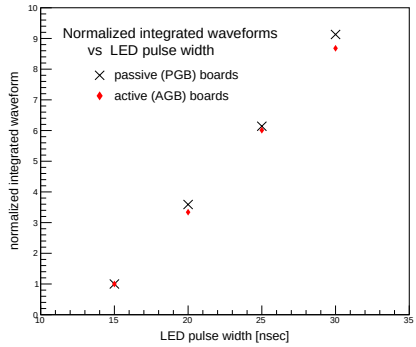

Figure 10: The comparison of the integrated waveforms from passive and active readout boards from the TallBo experiment. The boards were tested in LN2 in the lab by flashing them with an LED having different pulse widths. The results obtained from the two passive and the two active boards were averaged. The mean integrated waveforms were normalized to their values are 15 ns.

not feasible to determine whether all charge information is contained in the integral out to the overshoot. The test dewar was small and reflections off the stainless steel sides made the results dependent on the exact placement of the boards in the dewar. The exact position of the boards in the dewar from test to test proved impossible to reproduce exactly. For this reason, the comparison of the passive boards and the active boards was made with the means of the boards and the results were normalized to their values at 15 ns. Nevertheless, Fig. 10 does make the case that the integrated waveforms from passive boards and the active boards out to the overshoot yields consistent information about the response of the detectors.

The integrated waveforms from the passive and active boards on the IU light guide and the Fermi light guide were put into separate histograms for runs 2, 3, and 4. Examples of these histograms for the passive boards on the IU light guide in runs 2 and 4 are shown in Fig. 11. Two cuts were made on the histograms.

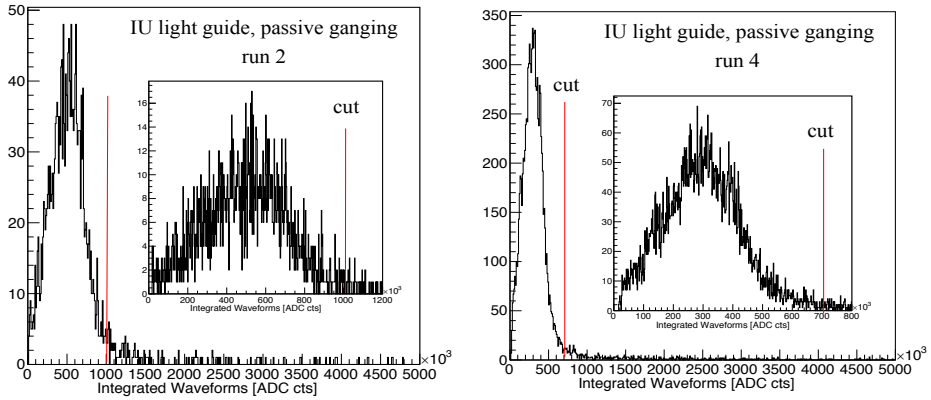

Figure 11: The integrated waveforms from the passive boards on the IU light guide in runs 2 and 4. The long high tail in the integrated are apparent. The muon "bump" for each run is shown in the inset. The cut on the tail is marked.

(1) Integrated waveforms with very low ADC counts were considered noise and cut. (2) As seen in Fig. 11, there is a large peak and a long, high energy tail in the histograms of the single track events. The distribution of events around the peak ("bump") was associated with single, minimum ionizing cosmic muons. The long tail are likely a mix of high energy muons and noise events. These events were also cut from the analysis. For each run in the analysis, the cuts were made independently for the passive boards and the active boards. For each run the waveform samples for PGB1 and PGB2 included the same tracks, and the waveform samples for AGB1 and AGB2 included the same tracks. The passive and active samples, however, were not the same.

The results of the analysis are given in Table 6. The number of tracks in the muon bump for each readout board in each run is given in column [3]. Although the track sample for the passive boards in each run and the track sample for the active boards in each run are the same, the number of tracks in each sample used in the analysis differs somewhat. This is the result of the cut on the high energy tail affecting the two modules somewhat differently. Further, it is also quite apparent that there are significantly fewer tracks in the active board samples than in the passive board samples. This is the result of cutting events

Table 6: Results.

	readout board	#tracks	bump integral [ADC]	geometry correction	[ADC]/track
IU light guide	passive				
run2	PGB1	2,391	1.18×10^9	1.00	4.94×10^5
run3	PGB1	2,322	1.14×10^9	1.00	4.91×10^5
run4	PGB2	10,160	3.03×10^9	1.04	2.98×10^5
	active				
run2	AGB1	683	2.63×10^8	1.00	3.54×10^5
run3	AGB1	997	1.69×10^8	1.00	3.70×10^5
run4	AGB2	1,217	2.86×10^8	1.04	2.35×10^5
Fermi light guide	passive				
run2	PGB2	2,380	6.06×10^8	1.04	2.65×10^5
run3	PGB2	2,249	6.35×10^8	1.02	2.82×10^5
run4	PGB1	10,187	2.10×10^9	1.00	1.98×10^5
	active				
run2	AGB2	728	1.67×10^8	1.04	2.29×10^5
run3	AGB2	886	9.79×10^7	1.03	2.30×10^5
run4	AGB1	2,311	5.71×10^8	1.00	2.47×10^5

with standard deviations >2, which impacts the waveforms in the active board sample more strongly, as shown in Fig. 9. Column [4] gives the total number of ADC counts in the waveforms in the bump.

There was a concern that the differing proximity of the light guides to the reflecting dewar walls could systematically affect the number of scintillation photons striking the light guides. The simulation described in §4 was used to address this issue. Every track in the TallBo data set that passed analysis cuts was simulated 10 times [1]. For each of these 10 simulated tracks, the positions of all photons that stike either the IU or Fermi PD modules were collected

into a histogram. The resulting 10 histograms for each PD were then averaged into a single histogram to determine the expected number of photons striking either the IU or Fermi PD the light guide for that track. The individual track histograms were then summed for all tracks in the passive or active waveform samples in each run to find the number of scintillation photons striking either the IU or Fermi PD. The ratio of the number of photons striking the PD modules, or the correction for the position of the light guides in the dewar for each run (normalized to PBG1 or AGB1) is given in column [5] of Table 6. Finally column [6] gives the mean # of ADC counts per track in the sample corrected for the position of the light guides in the dewar.

There are two results that can be inferred from column (6) of Table 6. First, the switch from reading out front side tracks in run 2 to reading out back side tracks in run 3 had little effect on the mean # of ADC counts per track. This was the expected result because there was no change in the readout boards and the light guides see scintillation photons symmetrically from both sides. The mean variation in the # of ADC counts per track between the front side and the back side is ≈4%. Since the track trajectories in the samples of front side and back side tracks have somewhat different geometries, it would not be expected that the match would be exact. The efficiency of the TPB plates (IU) or coating (Fermi) are also expected to be somewhat different.

On the other hand, there was no clear expectation on how the mean # of ADC counts per track would differ from runs 2 and 3 to run 4, after the dewar was drained and refilled, and the readout boards for the IU and Fermi light guides, by staying in place on the frame, became switched on the light guides. Table 6, however, shows that there was a significant fall off in the # of ADC counts per track after the light guides were exchanged for the IU and Fermi technologies when read out passively and for the IU technology when read out actively. Calculating the percent change as

$$(< \text{run } 2, \text{run } 3 > - \text{ run } 4)/(< \text{run } 2, \text{run } 3 > + \text{ run } 4) \times 100$$

shows that the # of ADC counts per track fell by 25% for the IU light guide

read out passively, by 18% for the Fermi light guide read out passively, and by 21% for the IU light guide read out actively. Only the Fermi light guide read out actively does not follow this trend and shows no fall-off at all.

To test whether the results of Table 6 are robust, the analysis was repeated with each data set broken into two. The same analysis routines were used with the same cuts. For passive boards in runs 2 and 3, the (# ADC counts)/track for both pieces remain close to their values for the whole run. For the active boards, where there are fewer tracks, the results show more variation but they remain consistent with Table 6. The same drop-off is also seen between runs 2 and 3 and run 4,. For both pieces of the run 4 data set for the active board on the Fermi light guide, no drop-off is seen, consistent with Table 6.

A careful visual inspection of both light guides once the experiment was completed did not show any noticeable evidence for crazing, which could scatter photons out of the light guide. Inspection of both light guides under UV light also showed no evidence for crazing. There was no visible degradation in the waveshifting plates used in the IU technology.

There is one additional result that can be deduced from column (6) in Table 6. Comparing the # of ADC counts per track for the IU and Fermi light guides for runs 2 and 3, before the dewar was drained, there is plausible evidence that the IU light guide technology has a higher efficiency than the Fermi technology. An estimate of the relative efficencies of the IU and Fermi technologies is best obtained from the passive boards, calculated using the ratio of the mean # of ADC counts in run 2 and run 3 for the two technologies. This estimate implies that the IU technology is $\sim$16% more efficient than the Fermi technology for the same sample of cosmic tracks. It is not possible, however, to quantify the absolute efficiencies of these technologies in this experiment. Multiple dark measurements of the passive boards before and after the experiment like those shown in Fig. 2, as well as calibration runs during the experiment, show no evidence at all for the single p.e. peaks needed for the quantification of the absolute efficiency.

6. Discussion

There are several possible explanations for the significant drop-off in the efficiency of the IU and Fermi light guides after refilling the dewar for run 4 when read out passively and the IU light guide when read out actively. One possibility is that contamination quenched scintillation light in run 4. The discussions in §3.2 and §4, however, show that contamination is not likely to be the cause. The literature on quenching in LAr and the TallBo simulations are both inconsistent with contamination being responsible for the drop-off. Further, this explanation would fail to explan why the Fermi technology, when read out actively, does not follow the trend.

A second possibility is that the MPPC response degraded due to the refilling of the dewar prior to run 4. Fig. 12 makes the case that this explanation is also

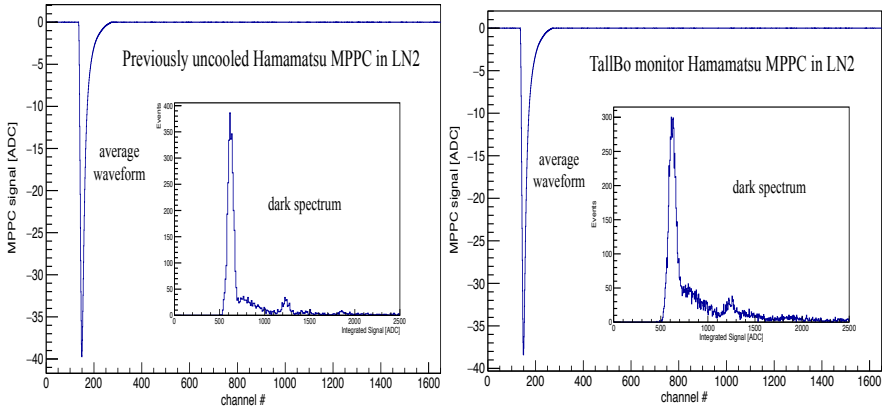

Figure 12: The comparison of the dark spectrum of a previously uncooled MPPC in LN2 biased at 44.5 V with the dark spectrum of the monitor MPPC also biased at 44.5 V. Both spectra are superposed on an average of 25 dark waveforms. These MPPCs demonstrate behavior that is quite similar.

unlikely. Fig. 12 shows the comparison of the dark spectrum of a previously uncooled MPPC in LN2 biased at 44.5 V (cf., Fig. 2) compared with the dark spectrum of the monitor MPPC also biased at 44.5 V. These spectra are superposed on an average of 25 dark waveforms. Both MPPCs demonstrate very similar waveforms and dark spectra. There does not seem to be some evidence

for degradation in the dark spectra of the monitor MPPC, and by extension the MPPCs used on the readout boards. The first p.e. peak peak is somewhat reduced in the monitor MPPC and the afterpulsing has increased. The integrated spectra are estimated to differ by less than 10%. Further evidence is found in Fig. 4, where the dark noise, position of the first p.e. peak, the pseudo-gain, and the pseudo-cross talk of the monitor MPPC is seen to fall well within the distributions of these characteristics in a sample of 12 previously uncooled MPPCs. The 48 individual MPPCs in the 4 readout boards were removed once the expeirment was completed and tested electrically by measuring their resistance and continuity. Their properties were similar to what they were before the experiment. They were not individually tested in LN2.

A third possibility is that the readout boards degraded after the refill. Fig 13

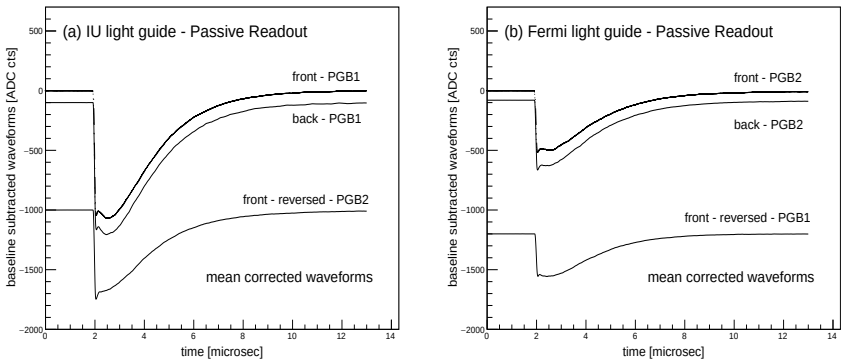

Figure 13: The baseline subtracted mean of 25 waveforms from single track muon events crossing TallBo as read out by a passive ganging board in runs 2 (front), run 3 (back), and run 4 (front-reversed). The waveforms are offset from one another for clarity and are labeled with the readout boards that collected them. The waveforms in the figure have been multiplied by the geometry correction given in Table 6.

suggests this explanation is unlikely for the passive readout boards. This figure shows the baseline subtracted mean of 25 waveforms from single track muon events crossing TallBo as read out by a passive ganging board in runs 2 (front), run 3 (back), and run 4 (front-reversed). The waveforms are offset from one

another for clarity. Since the positions of the light guides were exchanged from run 3 to run 4, but the readout boards remained in place, the waveforms are labeled with the readout boards that collected them. The waveforms were multiplied by the geometry correction given in Table 6. The 25 waveforms used in computing the mean were selected near the peak of the muon bump (Fig. 11).

The reduced signal strength responsible for the drop-off in efficiency, as seen in the depth of the waveform response, could be a result of degradation in either the light guide response or the board response. The readout boards remained on the same DAQ channels. Fig. 4 and Fig. 12 imply that it is unlikey to be due to a degradation in the individual MPPC responses for the 12 MPPCs ganged in parallel. Fig. 13 shows the structure of the waveforms from the boards are all similar, both before and after the refill, and that the exponential tails remain comparable. This suggests that the board response has also not degraded. It seems most likely that the it is the light guides and not the readout boards that have degraded due to thermal cycling in the data for the passively ganged technologies.

The light guides read out by the active boards show more complicated behavior. Fig 14 shows the baseline subtracted mean of 25 waveforms from single track muon events crossing TallBo as read out by an active ganging board in runs 2 (front), run 3 (back), and run 4 (front-reversed). The waveforms are offset from one another for clarity and are labeled with the readout boards that collected them. The waveforms were multiplied by the geometry correction given in Table 6. The 25 waveforms used in computing the mean were selected near the peak of the muon bump (Fig. 11).

For the IU light guide there is again reduced signal strength as seen in the depth of the waveform response. Again it is unlikely that this is due to a degradation in the individual MPPC responses. The waveforms all have very similar structure before and after the refill, with the overshoot occurring at approximately the same point on the waveform. As for the passive boards, this suggests that the board response has also not degraded. For the IU light guide it is probable that the light guide and not the readout board has degraded due

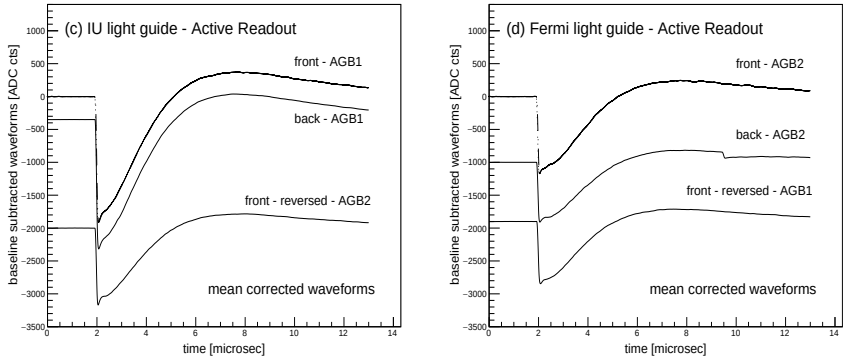

Figure 14: The baseline subtracted mean of 25 waveforms from single track muon events crossing TallBo as read out by an active ganging board in runs 2 (front), run 3 (back), and run 4 (front-reversed). The waveforms are offset from one another for clarity and are labeled with the readout boards that collected them. The waveforms were multiplied by the geometry correction given in Table 6.

to thermal cycling.

For the Fermi light guide the situation is different. Although the waveforms have very similar forms in the runs, with the overshoot occurring at approximately the same point on the waveform, signal strength as seen in the depth of the waveform response has not fallen off. The most likely explanation for this behavior is that the response of AGB1 is greater than the response of AGB2. The signals are summed electronially on these boards and a difference in board response is not improbable. Lab tests of this hypothesis on the AGB boards at IU once the experiment was completed were not successful. As discussed above, these tests were made in a small dewar and reflections off the stainless steel sides made the results dependent on the exact placement of the boards in the dewar. It was not possible to reproduce the exact position of the boards in the dewar from test to test. The tests were therefore unable to determine the absolute response of the AGB boards. It seems plausible that the Fermi light guide was also degraded in its response due to thermal cycling as seen in the data from the passive readout boards.

Another possibility, in which TPB leaches out of the detectors while submerged in LAr, has been suggested in [17]. This explanation is unlikely for the results discussed here, however, since the response dropped precipitously from run 3 to run 4 but showed almost no change from run 2 to run 3. If the TPB had leached out of the detectors into the LAr, the drop off would be expected to be cmore or less ontinuous from run 2 to run 4.

The final possibility to be discussed is that the drop off in signal from run 2 to run 4 is due to a difference in the path length distributions in LAr of the tracks in the two runs. Here the assumption is made that the tracks passing the cuts are minimum ionizing muons that give off 8.4×10^4 sintillator photons/cm as they traverse the LAr. The difference in the [ADC]/track from run 2 to run 4 could result from a difference the number of photons striking the light guides due to different path length distributions in LAr for the two runs. A cosmic track between fixed PMT positions on the hodoscopes can be characterized by its zenith and azimuthal angles and every track with these angles passes through a specific path length of LAr. Consequently, the zenith angle and azimuthal angle distributions for the tracks correlate with the path length distributions of the cosmics in LAr and the scintillation photons given off. It is possible that different track geometries in run 2 and run 4 could conspire to have the same zenith angle and azimuthal angle distributions. But this seems unlikely. The zenith angle and azimuthal angle distributions for the tracks passing the cuts from the passive readout boards in runs 2 and 4 are shown in Fig. 15. As can be seen, these distributions are quite similar. The differences can be attributed mostly to the stochastic nature of cosmic rays and the different statistics in the two runs. Fig. 15 suggests that the drop off in signal seen from run 2 to run 4 are unlikely to be from differences in the path length distributions between the two runs.

A second way to test whether the fall off in the signal/track between the two runs is due to differences in the track length distributions from runs 2 and 4 is to use the simulations that compute the geometry correction factors in Table 6. These simulations show that an average of 3.73×10^5 photons/track fall on the

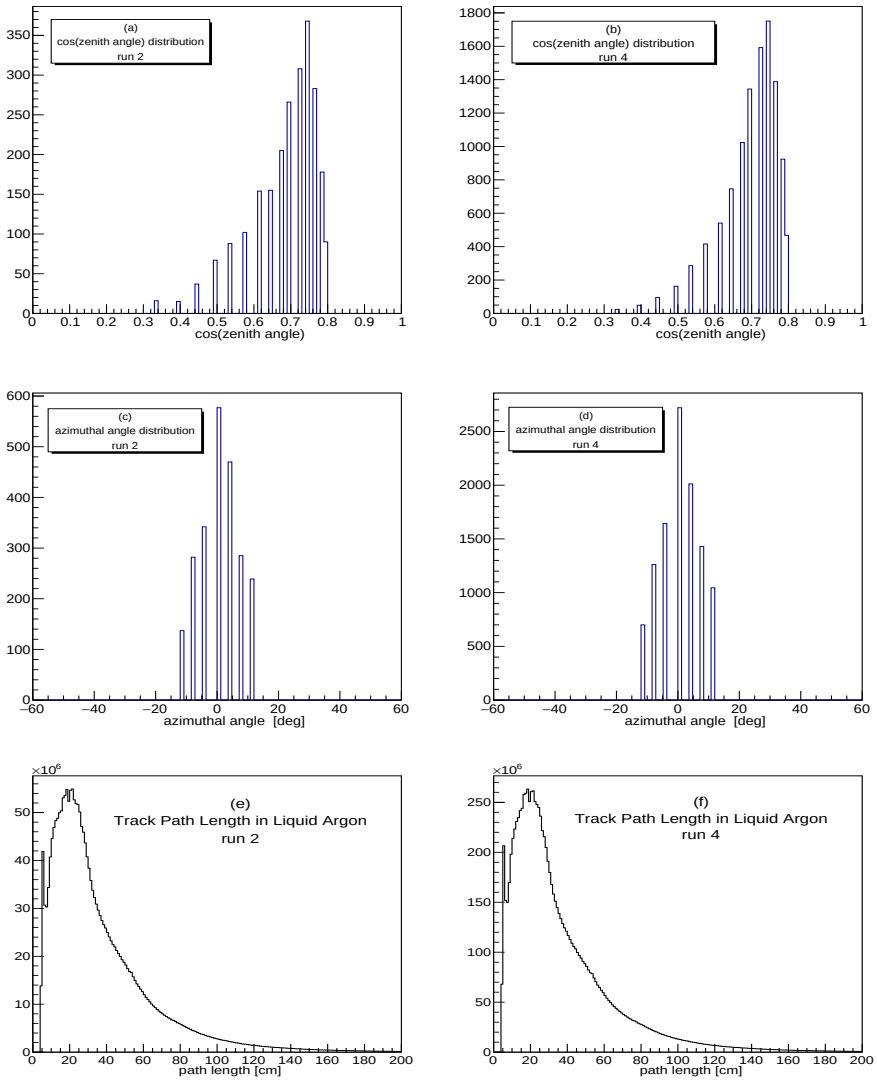

Figure 15: The zenith angle and azimuthal angle distributions for the tracks recorded by the passive readout in runs 2 and 4: (a) the zenith angle distribution in run 2; (b) the zenith angle distribution in run 4; (c) the azimuthal angle distribution in run 2; a(d) the azimuthal angle distribution in run 4; (e) the track path length distribution in run 2; and the path length distribution in run 4.

IU light guide in run 2 and an average 3.76×10^5 photons/track in run 4. These

results incorporate the geometry corrections in Table 6. Calculations for the Fermi technology are redundant since the simulation uses the same tracks for both light guides. These calculations again suggest that the drop off in signal seen from run 2 to run 4 are unlikely to be from differences in the path length distributions between the two runs

The conclusion drawn from the TallBo experiment described here is that the response of the two light guide technologies tested in this investigation were degraded as the result of thermal cycling after the TallBo dewar was emptied and refilled. All measurements of the light guide response made with the passive and active readout boards were consistent with no significant change from run 2 to run 3. Both measurements of the light guide response with the passive readout boards showed a significant decrease from run 3 to run 4. The IU light guide also showed a significant decrease when read out by the active board. The Fermi light when read out by the active board did not. The bulk of evidence, therefore, favors the concluson that the response of the two light guide technologies were degraded due to the effects of thermal cycling.

Acknowledgements

This work was supported in part by the Trustees of Indiana University, the DOE Office of High Energy Physics through grant DE-SC0010120 to Indiana University, and grant #240296 from Broookhaven National Laboratory to Indiana University. The authors wish to thank the many people who helped make this work possible. At IU: M. Gebhard, M. Lang, J. Urheim. At Fermilab: R. Davis, A. Hahn, B. Miner, T. Nichols, B. Ramson. At ANL: G. Drake. At MIT: J. Conrad. At Eljen Technology: C. Hurlbut.

This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

References

[1] B. Howard, et al., A novel use of light guides and wavelength shifting plates for the detection of scintillation photons in large liquid argon detectors, NIM A907 (2018) 9.

[2] L. Bugel, et al., Demonstration of a lightguide detector for liquid argon TPCs, NIM A640 (2011) 69.

[3] Z. Moss, et al., Improved TPB-coated light guides for liquid argon TPC light detection systems, JINST 10 (2015) P08017.

[4] R. Acciarri, et al., Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE) Conceptual Design Report vol.1: The LBNF and DUNE Projects, arXiv:1601.05471.

[5] B. Jones, et al., A measurement of the absorption of liquid argon scintillation light by dissolved nitrogen at the part-per-million level, JINST 8 (2013) P07011.

[6] R. Acciarri, et al., Oxygen contamination in liquid argon:combined effects on ionization electron charge and scintillation light, JINST 5 (2010) P05003.

[7] R. Goldstein, F. Mastrup, Absorption coefficients of the O2 Schumann–Runge continuum from 1270 Å - 1745 Å using a new continuum source, JOSA 56 (1966) 765.

[8] K. Mavrokoridis, et al., Argon purification studies and a novel liquid argon re-circulation system, JINST 6 (2011) P08003.

[9] K. Watanabe, M. Zelikoff, Absorption coefficients of water vapor in the vacuum ultraviolet, JOSA 43 (1953) 9.

[10] D. Whittington, S. Mufson, B. Howard, Scintillation light from cosmic-ray muons in liquid argon, JINST 11 (2016) P05016.

[11] C. Patrignani, et al., The Review of Particle Physics, Chin. Phys. C 40 (2017) 100001.

[12] R. Acciarri, et al., Effects of nitrogen contamination in liquid argon, JINST 5 (2010) P06003.

[13] M. Miyajima, et al., Average energy expended per ion pair in liquid argon, Phys.Rev.A 9 (1974) 1438.

[14] T. Doke, et al., LET dependence of scintillation yields in liquid argon, NIM A 269 (1988) 291.

[15] T. Doke, et al., Estimation of absolute photon yields in liquid argon and xenon for relativistic (1 mev) electrons, NIM A291 (1990) 617.

[16] T. Doke, et al., Absolute scintillation yields in liquid argon and xenon for various particles, Jpn.J.Appl.Phys. 41 (2002) 1538.

[17] J. Asaadi, et al., Emanation and bulk fluorescence in liquid argon from tetraphenyl butadiene wavelength shifting coatings, JINST 14 (2019) P02021.