Available online at www.sciencedirect.com
ScienceDirect
Nuclear Instruments and Methods in Physics Research A 591 (2008) 453–466
NUCLEAR INSTRUMENTS & METHODS IN PHYSICS RESEARCH Section A
www.elsevier.com/locate/nima

Testing of photomultiplier tubes for use in the surface detector of the Pierre Auger observatory

D. Barnhill[a,*], F. Suarez[b,**], K. Arisaka[a], B. Garcia[c], J.P. Gongora[d], A. Lucero[c], I. Navarro[c], T. Ohnuki[a], A. Risi[d], A. Tripathi[a]

[a]*Department of Physics and Astronomy, UCLA, Los Angeles, CA 90095-1547, USA*
[b]*INFN, Universitá degli Studi di Torino, Italy*
[c]*UTN-FRM Mendoza, Argentina*
[d]*UTN-FRSR San Rafael, Argentina*

Received 13 July 2007; received in revised form 15 January 2008; accepted 23 January 2008
Available online 5 February 2008

Abstract

In the array of water Cherenkov detectors of the Pierre Auger Observatory, 4800 large photomultiplier tubes (PMTs) will be used. Before being deployed, each PMT is evaluated to check that various parameters, such as the linearity, dark noise, and gain, fall within a specified range. The large scale test system, designed and constructed for this purpose, is capable of testing multiple large PMTs simultaneously. The test system and the results of the tests for the first 3964 PMTs are presented in this paper.

PACS: 07.05.Hd

Keywords: PMTs; Photodetectors; Water Cherenkov; Cosmic rays; Astroparticle physics; Pierre Auger

1. Introduction

The Pierre Auger Observatory, currently being constructed in the province of Mendoza in Argentina, is designed to measure the energies, directions, and composition of the highest energy cosmic rays arriving at the earth. To accomplish this task, the observatory consists of two detectors: a surface detector, which is an array of 1600 water Cherenkov detectors deployed over $\sim$3000 km^2, and 24 fluorescence telescopes grouped into four sites, which overlook the surface detector [1]. Being able to measure ultra high energy cosmic rays using both techniques will provide unprecedented information about the nature and origin of these particles.

However, the Auger Observatory can only operate in "hybrid" mode (or using both the fluorescence and surface detectors together) $\sim$10% of the time [2]. The remaining 90% of the time, the surface detector operates alone. It is critical, then, that the surface detector is well understood as it is the foundation for all of the data taken at the Auger Observatory.

The surface detector is made up of an array of water Cherenkov detectors, or stations, which are cylindrical water tanks, 3.6 m in diameter and 1.2 m deep. They are filled with purified water and have a reflective interior that is fitted with 3×9 in. Photonis XP1805 photomultiplier tubes (PMTs) which look down into the station. The PMTs are equipped with a resistive base and a local High Voltage (HV) module [3]. When a relativistic particle enters the station, it emits Cherenkov radiation that propagates through the water, being reflected at the station walls until it is either absorbed or detected by the PMTs.

When completed, the surface detector will have 1600 stations, which will employ the use of 4800 PMTs. Including spares, the Auger Observatory will receive more than 5000 PMTs for use in the surface detector, and it is the

*Corresponding author.
**Also for correspondence.
E-mail addresses: david.scott.barnhill@gmail.com (D. Barnhill), fedesuarez@auger.org.ar (F. Suarez).

0168-9002/$ - see front matter © 2008 Published by Elsevier B.V.
doi:10.1016/j.nima.2008.01.088

testing and characterization of these PMTs that is addressed in this paper.

The layout of the paper is as follows: first, in Section 2, the test system will be described, including the hardware that was designed specifically for this system. Then, the tests run on the PMTs along with the results of testing will be presented. Finally, in Section 3, the monitoring of the test system will be described with a discussion of the results.

2. Testing of PMTs

The purpose of testing each PMT before being deployed is two-fold. The tests are to verify that each PMT is within the specifications given to Photonis, specifications that are designed to ensure that only PMTs of the desired quality are used in the surface detector, resulting in uniform behavior across the array of stations. Secondly, we are able to give valuable feedback to the company regarding the performance of the PMTs which they can use to improve their product.

The specifications regarding the performance of the PMTs to be used in the surface detector is driven by the physics that is being done. It is desirable to have PMTs with a large dynamic range, good linearity, low counting rate, and low background. Because the calibration of the surface detector is done using atmospheric muons rather than depending on a knowledge of the absolute gain of each PMT, the desired energy resolution is not strictly specified. In Table 1 the specifications are listed to determine which PMTs will be used in the surface detector and are adapted from the original specifications given in Tripathi et al. [4], where PMTs from different companies were compared and analyzed as to which would best suit the needs of the Auger Observatory.

To illustrate the relationship between the specifications and the physics done with the Auger Observatory, we consider the calculation of the primary energy of a cosmic ray. As a first step, the energy deposited in each station involved in the air-shower must be known. The calibration of each station is done using single muons which are constantly passing through it [5], while a cosmic ray air-shower may cause tens of thousands of particles to enter a given station during a period of up to 1 μs. Therefore, it is desirable to have PMTs which have a linear response over this large dynamic range. The non-linearity test specification (less than 6% non-linearity below 50 mA, see Table 1), is designed to reject PMTs which deviate from linearity over this range. Related to this issue is the afterpulse measurement, as any afterpulsing in the PMTs may lead to a miscalculation of the energy deposited in a station.

Table 1
Specifications to determine if a PMT passed or failed a given test

Test	Specification
SPE peak to valley	>1.2
Gain versus voltage	10^6 gain with $V<2000$ V
Dark pulse rate	<10 kHz at $\frac{1}{4}$ pe threshold
Non-linearity	$<6\%$ below 50 mA peak current
Dynode to anode ratio	Between 25 and 40
Afterpulse ratio	$<5\%$

SPE stands for Single Photoelectron.

To cover the dynamic range of physical signals, from single muons to tens of thousands of particles, the output of the last dynode before the anode is tapped and amplified [3]. The amplified signal from the dynode makes it possible for the station to be triggered by single muons passing through a station, while the signal from the anode is used for the detection of any signal that saturates the readout electronics of the amplified dynode. Therefore, the ratio of the signal from the dynode to the signal from the anode must be known and the overall gain of the dynode chain must fall within a specific range. The absolute gain as a function of input voltage is measured using a single photoelectron (SPE) spectrum. The PMTs and bases were designed to operate with a gain between 2×10^5 and 10^6, and are currently being operated at $\sim3\times10^5$ gain [5]. All tests are explained in more detail in Section 2.2, but first, the design of the test system itself is described.

2.1. The test system

The system is designed to test 16 PMTs in a single run (see Figs. 1 and 2). However, to monitor the stability of the system, there are four permanent PMTs located at the corners of the test stand. These PMTs monitor the stability of the light source as well as the readout electronics and the performance of the system overall. Each test run lasts about 5 h and is completely automated. This makes it possible to do two test runs per day resulting in 24 PMTs per day being tested.

The entire test system is controlled with the data acquisition (DAQ) computer, see Fig. 1. This computer controls the voltages delivered to the 16 PMTs via a HV Box, controls the intensity and the firing of four LEDs (three blue and one UV) through the Light Control Box and LED Pulsers, and controls the signal which triggers the CAMAC data acquisition system.

2.1.1. The light control box

The light sources consist of four LED Pulsers and are controlled through a National Instruments multi-function IO PCI card (6025E) connected to the DAQ computer and the Light Control Box. The PCI card is used to control the intensity of the light through its analog outputs, as well as to generate the TTL signals to use as triggers for the LED pulsers. The Light Control Box is used to supply power for the LED Pulsers, to enable the triggers, and to buffer the analog output of the PCI card to control the brightness of the light sources. Thus, it is divided into three sections: the power supply, the trigger lines, and the analog buffer lines. The power supply section is made with two commercial linear power supplies which provide the ±12 and $+5$ V

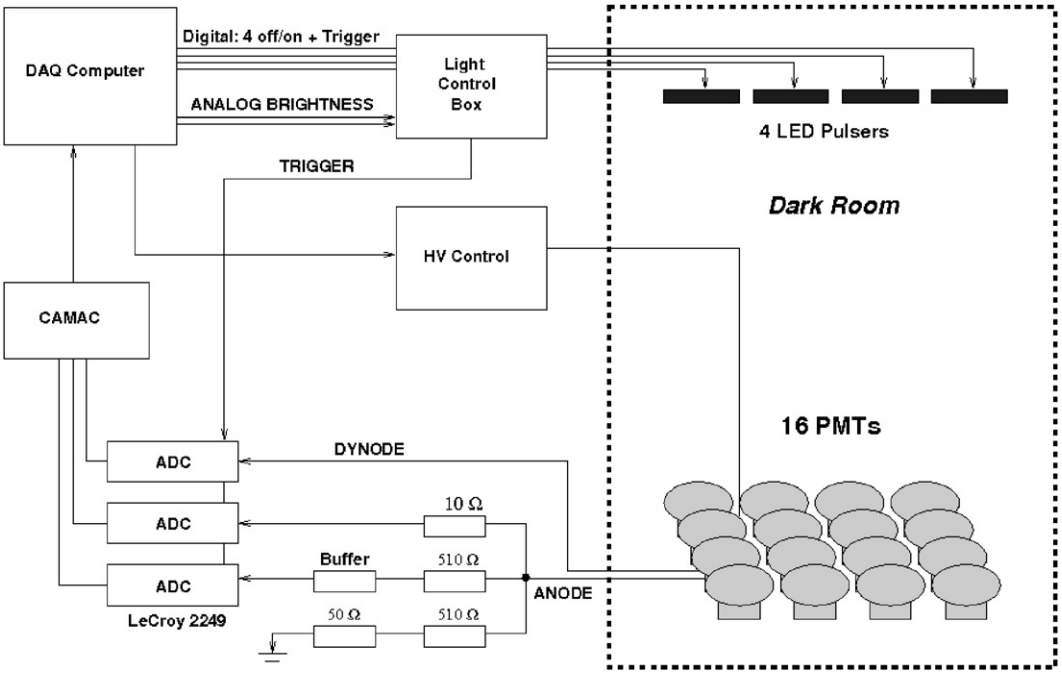

Fig. 1. Layout of the PMT data acquisition system.

required by the LED pulsers. As shown in Fig. 3, the trigger lines are made with five digital signals to enable the four light sources (one enable line per light source) and to trigger the acquisition gate for the integrating ADCs in the CAMAC crate. Thus, the signal lines for the light sources are considered only when the gate signal is present. The analog buffer lines are needed because there are only two analog outputs available on the PCI card, therefore two light sources are controlled by each output (See also Fig. 4).

2.1.2. The LED pulsers

The light sources are four LED Pulsers attached to the ceiling of the test room. Three of the four light sources use blue LEDs (Nichia NSPB520, 470 nm peak output and 30 nm FWHM) and one uses an UV LED (Nichia NSHU550E, 370 nm peak output and 15 nm FWHM). Two of the blue LEDs are used for non-linearity and gain-voltage measurements, and the third one is more attenuated and used for SPE and quantum efficiency (QE) measurements. The UV light source is only used for QE measurements in the UV range. Each one of the pulsers consists of a pulse generator, an analog buffer, the LED, and a diffuser (Fig. 4). The pulse generator (74F121) is controlled by the trigger signal coming from the Light Control Box, and generates a pulse with a width of approximately 10 ns. The analog buffer takes the analog signal coming from the Light Control Box and adjusts its level to correctly feed the LED. Thus, the signal width applied to the LED is controlled by the pulse generator (fixed) and the light intensity is controlled by the DAQ computer (variable). Since the LEDs are insufficiently diffuse, diffusers are used. The same diffuser is used for the SPE and QE light sources, but a less opaque one is used for the other two LEDs because it is necessary to have a more intense light source for the other measurements.

2.1.3. HV control box

The PMTs used in the Auger Observatory have their own HV power supply on their bases. The purpose, then, of the HV Control Box is to supply all the voltages needed by the base as well as the reference signal to control the HV applied to the PMT. Each PMT is fed with a +12 V signal generated with a LM78L12 regulator, ±3.3 V signals generated with LM317 and LM337 regulators, and a variable signal (called the reference line) which is between 0 and 2.5 V. The reference line is an input to the HV power supply module located on the base of each PMT which converts this 0–2.5 V signal to a voltage between 0 and 2000 V. The reference line voltage originates in a United Electronics Industries (UEI) PowerDAQ PD2-AO 32 PCI analog output card, and passes almost directly through the HV control box having just a 2.6 V zener diode in parallel to ensure that there is not more voltage applied than permitted to the HV module on the base of the PMT.

This box also provides an interlock system that requires the door of the dark room to be closed and the key to be inserted into the HV Control Box, which also automatically turns off the lights of the test room. If the key is not in the HV Control Box, no voltage is supplied to the PMTs. This is needed because the PMTs can be seriously damaged if HV is applied when ambient light is present.

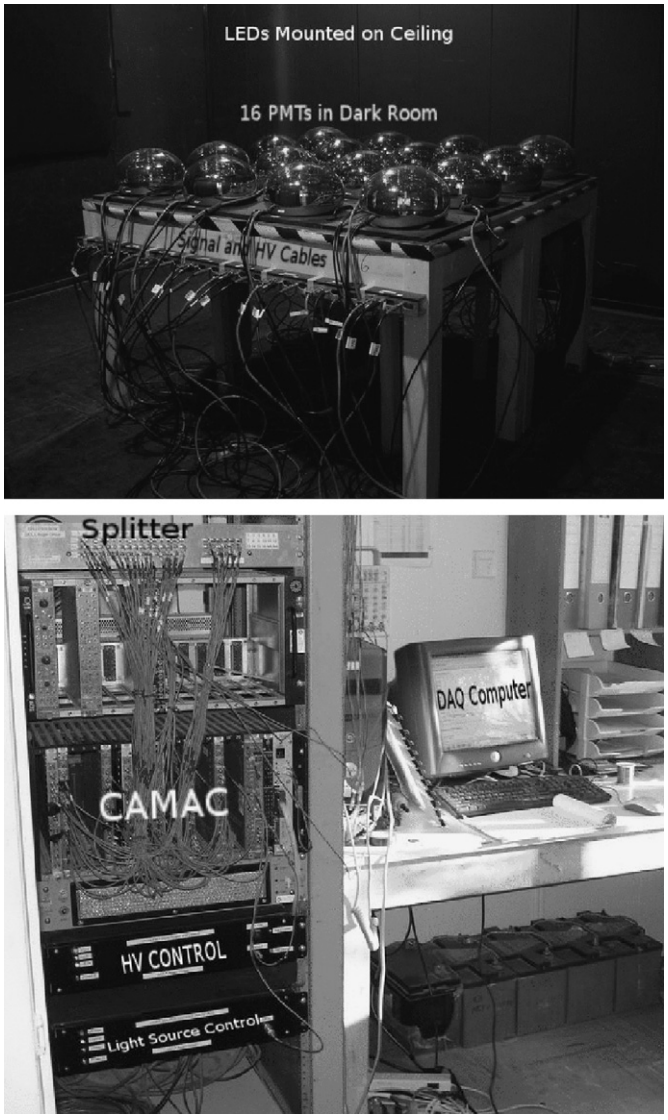

Fig. 2. Top: PMTs in the dark room. Bottom: The data acquisition system with the CAMAC crate and custom electronics in the rack on the left.

2.1.4. The splitter

For certain measurements, the dynamic range of the charge-integrating ADC modules is insufficient for the desired range of signals. Therefore, the signal from the anode is divided into three signals through a resistive splitter (see Fig. 5). Thus, two of the outputs from this splitter have approximately 8.5% of the charge of the original signal (called the attenuated anode signal), and the other output has the remaining 83% (called the anode signal). The attenuated anode signal is used when the anode signal saturates the ADCs. The remaining attenuated anode output is used for monitoring purposes, oscilloscope connections, or other measurements.

2.1.5. PMT signals and DAQ

The signals of the PMTs come from two places, the anode as is customary, but also from an amplified tap from the last dynode in the amplification chain of the PMT. This is done to extend the dynamic range of the PMTs when they are in the detectors taking data. To further extend the dynamic range of the system for testing purposes, the signal from the anode is broken into three components, as explained in the previous section, see Fig. 1. Again, this is done to extend the range over which the PMTs can be tested.

The signals from the PMT are then put into a charge-integrating ADC in a CAMAC crate (LeCroy 2249A and 2249W) where they are measured and then read out by the DAQ computer. The gate for charge integration is triggered by the Light Control Box, but the gate and width are controlled by a CAMAC gate-and-delay generator (LeCroy 2323A). All the analysis is then done on the DAQ computer.

2.2. Tests and results

Using this test system, 3964 PMTs have been tested out of the 5000 needed for the surface detector. The tests run by the system are SPE spectrum, gain as a function of voltage, dark pulse rate at $\frac{1}{4}$ photoelectron (pe) threshold, non-linearity, dynode to anode ratio, excess noise factor (ENF), and afterpulse ratio. Each test will be described in greater detail in the following sections.

2.2.1. SPE spectrum

To obtain the absolute gain of the phototube at a certain voltage, a single pe spectrum is measured. This is done by setting the PMT to a gain of $\sim 2 \times 10^6$, according to measurements done at Photonis, and flashing the LED at an intensity such that 90% of the time there are no pes at the first dynode in the PMT. From Poisson statistics

$$P(n) = \frac{e^{-v} v^n}{n!} \quad (1)$$

where $P(n)$ is the probability to see exactly n pe, if there are 0 pe 90% of the time, then $P(0) = e^{-v} = 0.9$, so that $v = 0.105$. Then, the probability of seeing 1 pe is: $P(1) = ve^{-v} = 0.095$, and the probability of seeing more than 1 is 0.005, or there is a $\sim$0.5% contamination of events caused by 2 or more pe in any given SPE spectrum. The signal then is dominated by SPE events, $P(1)/P(n>1) = 21$.

To determine a suitable intensity for the light source to achieve the necessary ratio of events with no pes at the first dynode of the PMTs, an offline calibration is performed. During this calibration procedure the light intensity is varied and the fraction of 0 pe events are determined at several locations in the test stand. When the fraction of 0 pe events at each monitored location in the test stand is above 90%, the intensity of the light source is fixed. The fraction of 0 pe events is monitored during testing and if the ratio becomes less than 90%, the calibration procedure is repeated and the intensity of the light source is altered.

A typical SPE spectrum can be seen in Fig. 6. To calculate the gain from this spectrum, the pedestal (or the

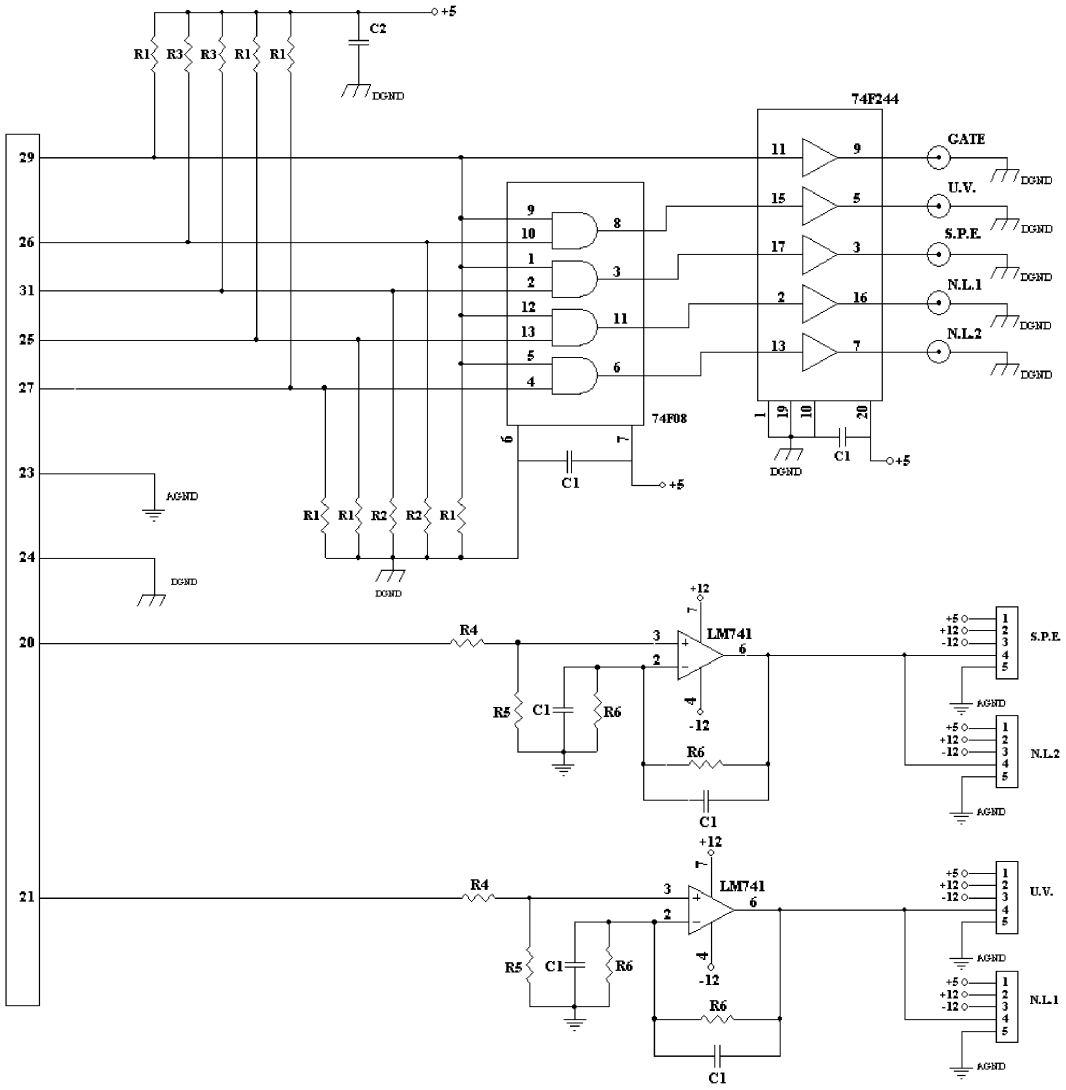

Fig. 3. Schematic of the light control box.

signal deposited with no light) and the standard deviation of the pedestal is established in a measurement directly preceding the SPE measurement. Once this is known, data are taken with the light source such that 90% of the events are 0 pe events. The resulting spectrum contains events from 0 pe, 1 pe, 2 pe, etc. To calculate the gain, it is necessary to find the mean of the single pe distribution, correcting for 2 pe contamination and compensating for events that are under the pedestal.

To compensate for events under the pedestal, a simple extrapolation of the behavior near the pedestal is assumed. As mentioned above, the mean and standard deviation of the pedestal is determined immediately before the SPE measurement. Using the first six non-pedestal bins in the histogram (starting with the bin that is $+3\sigma_{\text{pedestal}}$ away from the pedestal mean) the average number of events per bin is calculated. This average is then used as the estimated number of events in each bin between the pedestal mean and the bin that is $+3\sigma_{\text{pedestal}}$ away from the mean. Using this extrapolation and ignoring the events in the pedestal, the mean of the non-zero pe distribution is calculated.

To account for 2 pe events, a correction is made utilizing the Poisson nature of the light source. The correction is calculated knowing that the mean of the non-zero pe distribution is

$$\overline{x} = \frac{\mu_1 P(1) + \mu_2 P(2) + \mu_3 P(3) + \cdots}{P(1) + P(2) + P(3) + \cdots} \tag{2}$$

where μ_n is the mean of the n pe distribution and $P(n)$ is the probability to have exactly n pe. Ignoring any events from 3 or more pe and knowing that $\mu_2 = 2\mu_1$, Eq. (2) is solved for μ_1, the true mean of the SPE spectrum, in terms of $\overline{x}$, the mean of the measured distribution:

$$\mu_1 = \frac{1 + v/2}{1 + v}\overline{x} \tag{3}$$

using v from Eq. (1).

To quantify the resolution of the SPE spectrum, the peak to valley ratio is used. The peak to valley ratio is the ratio

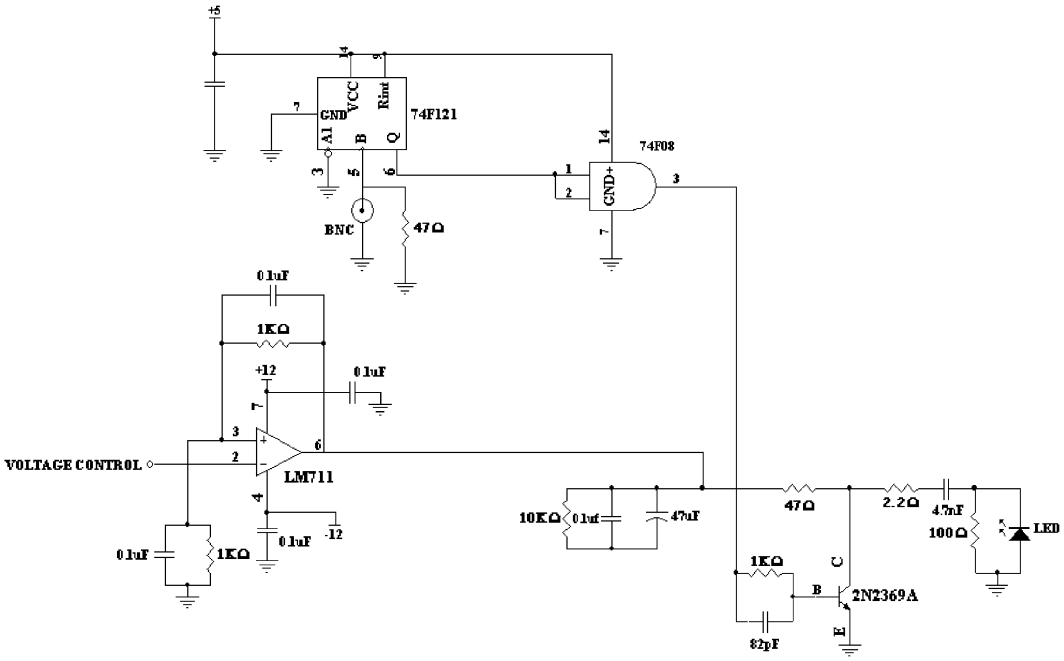

Fig. 4. Schematic of the LED pulsers.

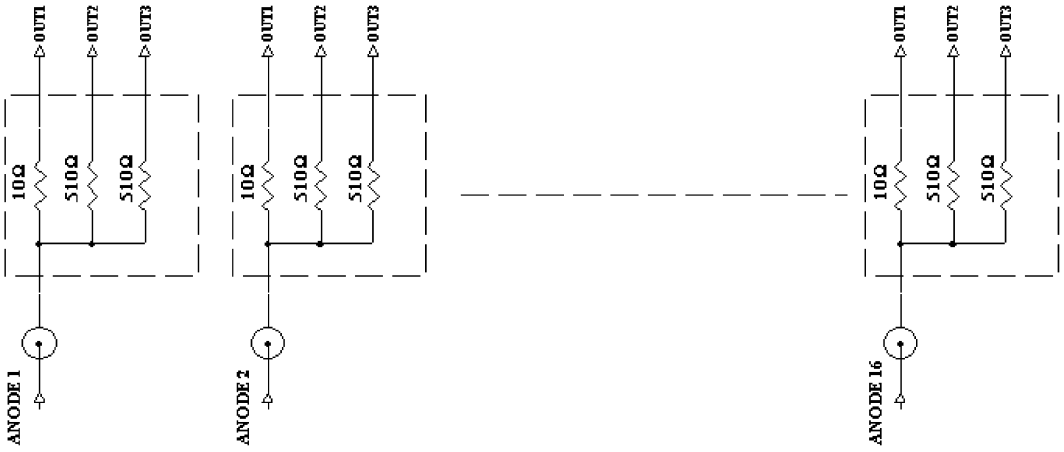

Fig. 5. Schematic of the splitter.

of the maximum value of the histogram of the SPE spectrum to the minimum value between the pedestal and the maximum. To calculate this number, smoothing is done on the distribution by creating an array where each element is the average number of events in six consecutive bins of the histogram:

$$a[n] = \frac{x_{n-2} + x_{n-1} + \cdots + x_{n+3}}{6}. \quad (4)$$

One then just steps through the smoothed array (a) and finds the minimum value (between the pedestal and the mean of the SPE distribution) and the maximum value of the histogram. The peak to valley ratio of the PMT shown in Fig. 6 is 1.47. In Fig. 7, the distribution of peak to valley ratios for PMTs tested in the system is shown. In the figure, the results are shown for 3463 PMTs for which a peak to valley ratio is successfully determined. While all PMTs (3964) are tested, the algorithm to find a peak to valley ratio may return a null result if there is no discernible peak and valley structure. No PMTs are rejected at this stage if either the peak to valley ratio is not determined (501 PMTs) or the ratio is less than 1.2 (55 PMTs).

2.2.2. Gain as a function of voltage

Once the gain is calculated from the SPE spectrum, the absolute gain at that voltage is known. However, it is necessary to have several points to determine the gain as a function of the input voltage. The relationship between the gain and input voltage can be described accurately as a power law:

$$G = kV^{\beta} \quad (5)$$

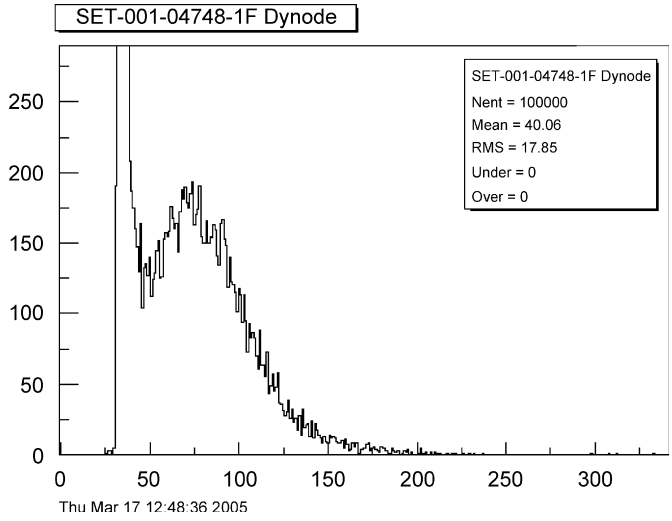

Fig. 6. Typical single photoelectron spectrum. The first tall peak is the pedestal and second peak is due to 1 pe events. The peak to valley ratio in this plot is 1.47.

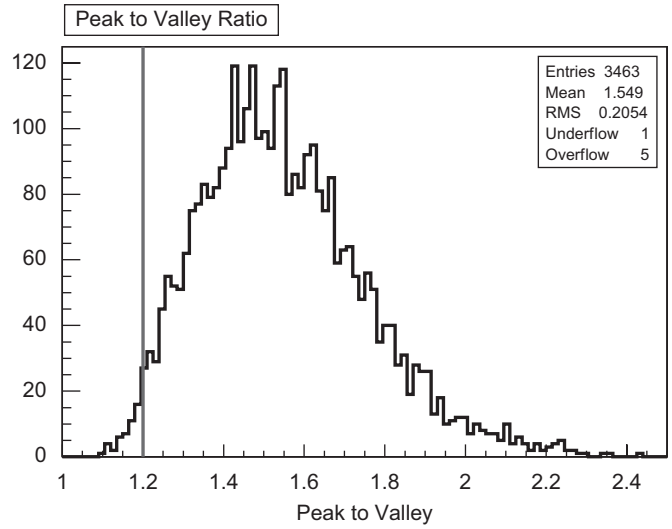

Fig. 7. Distribution of peak to valley ratios for 3463 PMTs tested in the system. The grey line represents the specification given to Photonis for the minimum desired peak to valley ratio and 55 PMTs fall below this line. Note that the algorithm to find peak to valley can fail if there is no discernible peak and valley structure, which returns a null result. This explains the discrepancy in the total number of PMTs tested (3964) and the number of entries in this histogram.

$$\log G = \gamma + \beta \log V \quad \text{with } \gamma = \log k \tag{6}$$

with γ and β being parameters to be determined from measurements.

To determine this relationship, the LED is pulsed with a constant intensity and the PMT is set at different voltages. The resulting measurements are used to calculate the slope of the line (β) in Eq. (6). The value of the gain at the voltage used by the SPE measurement is then used as a point through which this line must pass. This uniquely determines the y-intercept of the line (γ) and one can calculate the gain of the PMT at any input voltage. An

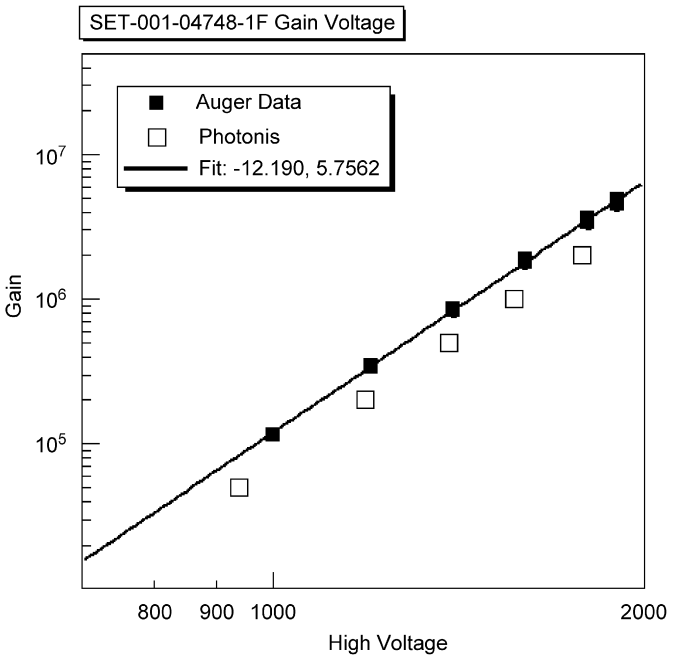

Fig. 8. A typical curve for the gain as a function of voltage for a PMT tested in the system, shown with the curve obtained by Photonis.

example of the gain versus voltage curve is given in Fig. 8, and for this particular PMT $\gamma = -12.190$ and $\beta = 5.7562$.

In Fig. 8, the measurements taken at Photonis are shown on the same plot to illustrate the difference between two methods of measuring the gain. At Photonis, the measurements are made using a constant light source and measuring the current of the photocathode and the anode. The gain is then the ratio of the anode current to the photocathode current. In this method, the collection efficiency (or the percentage of pes from the photocathode that reach the first dynode) is included, whereas, in the SPE method it is not. As a result, one can take the ratio of the gain as calculated by the SPE method and the gain calculated using the method just described and obtain an estimate of the collection efficiency at a given voltage. Calculating the collection efficiency is not necessary in any physics done at the Auger Observatory, but it is a nice bonus feature of the different measurements. The difficulty arises in understanding the systematic uncertainties when comparing two measurements from different systems. However, knowing that the systematics between the systems are not taken into account, the mean collection efficiency estimate is calculated as an exercise and the result is ~70%. This is calculated at the voltage necessary for a gain of 10^6 using 3964 PMTs.

Fig. 9 is a plot of the voltage necessary to get a gain of 10^6 as determined by the SPE method (axis labeled Auger Voltage) and using the photocathode current to anode current ratio method (axis labeled Photonis Voltage). The specification given to Photonis is that the optimal operating gain should be between 2×10^5 and 10^6, which should be possible with less than 2000 V applied. Again, this is to match the PMT output of physical signals, from single

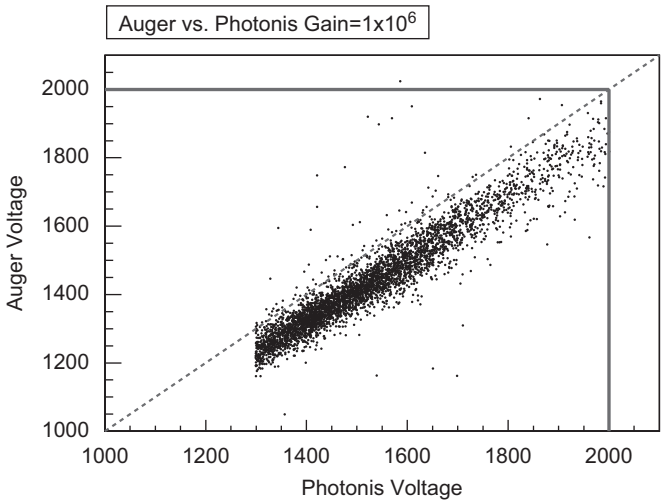

Fig. 9. Voltage to get a gain of 10^6 as measured by our test system (axis labeled Auger Voltage) and Photonis. The grey line is the specification given to Photonis and the dashed line is $y = x$. Three PMTs failed to meet the specification (0.08%).

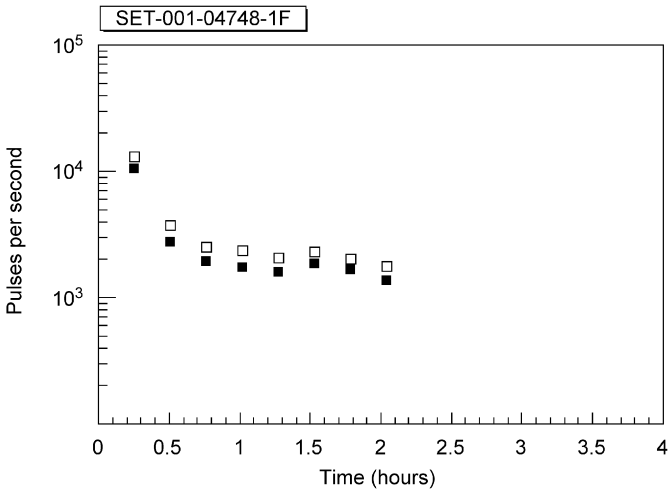

Fig. 10. Plot of the dark pulse rate versus time spent in the dark room. Open boxes are $\frac{1}{4}$ pe threshold while solid boxes are $\frac{1}{2}$ pe threshold.

muons to thousands of particles, to the dynamic range of the readout electronics of the PMT. Of the 3964 PMTs tested, three did not meet this requirement (0.08%).

2.2.3. Dark pulse rate

When a PMT is exposed to a lot of light, like when it is exposed to the light in a room or daylight, it needs a chance to sit in the dark to reduce the dark pulse rate. The dark pulse rate is the rate at which signals above a certain threshold are observed in a given PMT with no light incident on the photocathode. It is analogous to the dark current of a PMT when discussing DC measurements. Dark current is caused mainly by the leakage current (or ohmic leakage), thermionic emission, and regenerative effects which are only a concern when operating at very HVs [6].

In the test system, before other tests are run, the PMTs are allowed to sit in the dark room for a period of 2 h during which time the dark pulse rates above $\frac{1}{4}$ pe and $\frac{1}{2}$ pe threshold are monitored. Fig. 10 is a plot of a typical dark pulse rate versus time.

Of course, to know the magnitude of the $\frac{1}{4}$ pe threshold for a PMT, the absolute gain of the PMT must be known. Therefore, after the 2 h period in the dark, the gain of the PMT is calculated and the measurement of the dark pulse rate is made again. The distribution of dark pulse rates for the PMTs is in Fig. 11. The specification given to Photonis was that after two hours in the dark the dark pulse rate above $\frac{1}{4}$ pe would be below 10 kHz and 12 PMTs (0.30%) failed to meet this requirement. This limit is set because there is a correlation between dark pulse rate and the lifetime of a PMT, and the less noise in the PMT, the longer it will last. The Auger Observatory is intended to run for the next 20 years, so longevity is important.

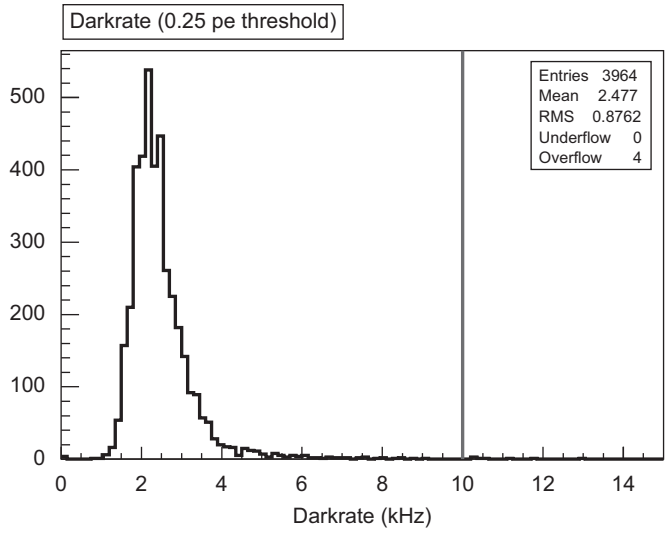

Fig. 11. Distribution of dark pulse rates (1/4 pe threshold) after 2 h in the dark room. The grey line represents the specification given to Photonis and 12 PMTs (0.30%) fail to meet the requirement.

2.2.4. Non-linearity

The validity of applying the calibration of the stations in the surface detector to the extensive air-showers that are detected relies on the linearity of the response of the PMTs. The calibration of the PMTs is done with single muons passing through the station whereas large showers can have thousands of particles entering the station [5]. To extend the calibration from single muons to thousands of particles, the linearity of the response of a PMT is a concern. Non-linearity normally occurs when the current gets high enough to cause a space–charge effect around the last dynode. This space–charge effect is caused by an excessive amount of electrons which change the electric field in that region, causing the normal trajectory of the electrons to be skewed. Thus, the amount of electrons

arriving at the last dynode, and hence the anode, is smaller than expected. This causes a negative non-linearity. In many of the PMTs from Photonis, there is a positive non-linearity which is due to the design of the dynode chain. It is designed to collect electrons more efficiently at higher currents which causes signal to be lost at lower currents. This appears as a positive non-linearity due to the definition of non-linearity (see Eq. (7)).

The method to measure non-linearity uses two LEDs. LED A is fired, LED B is fired, LED A and B are fired simultaneously, then no LED is fired (to obtain the baseline or pedestal). The non-linearity is then defined as

$$\mathrm{NL}(\%) = 100 \times \frac{Q_{AB} - (Q_A + Q_B)}{Q_A + Q_B} \quad (7)$$

where Q_A is the signal from firing LED A alone, Q_B is the signal from LED B alone, and Q_{AB} is the signal from firing LEDs A and B simultaneously (all are baseline subtracted). This sequence is repeated at several light intensities to map out the non-linearity as a function of peak anode current. A typical non-linearity curve is shown in Fig. 12. In this figure, the positive non-linearity feature is evident as well as the following negative non-linearity due to the space–charge effect.

To illustrate the properties of the non-linearity in the PMTs from Photonis, a plot of the maximum non-linearity versus the non-linearity at 50 mA is shown in Fig. 13. The reason 50 mA is chosen is because the specification given to Photonis was that the non-linearity be less than ±6% with a peak anode current of less than 50 mA. In the original design of the Pierre Auger Observatory, this was estimated to be the peak current at 1000 m from the core of an air-shower initiated by a cosmic ray with an energy of 10^{21} eV [7].

In Fig. 13, it is evident that almost all the PMTs have a positive non-linearity, and the maximum positive non-linearity often occurs around 50 mA of peak anode current. It should be noted that the maximum positive non-linearity

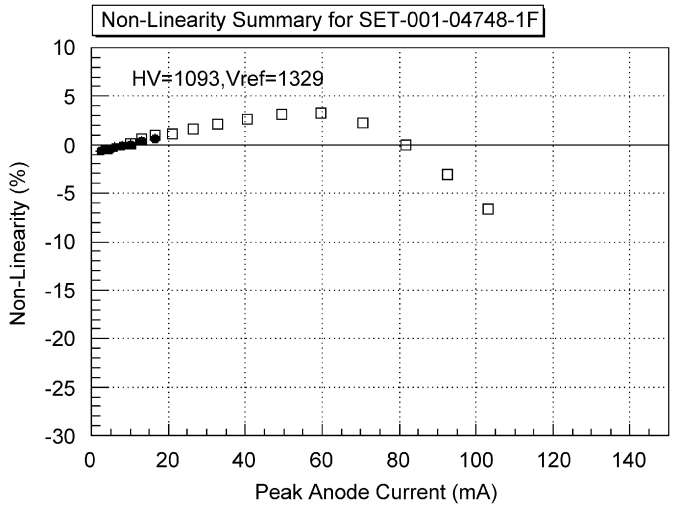

Fig. 12. Non-linearity versus peak anode current for a typical PMT.

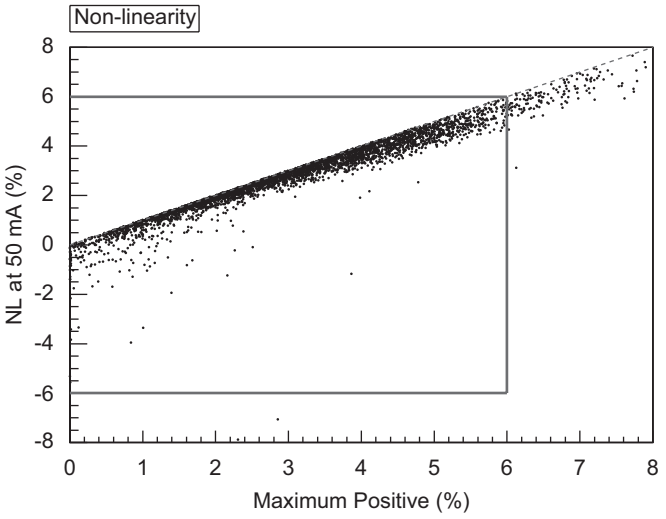

Fig. 13. Non-linearity at 50 mA versus maximum positive non-linearity. The grey lines represent the 6% limit given to Photonis and the dashed line is $y = x$. One hundred and fifty six PMTs (3.93%) fail to meet the non-linearity requirement.

is defined as the maximum non-linearity with a peak anode current less than 50 mA. If non-linearity continues to increase after a current of 50 mA, that is not considered in the definition of maximum positive non-linearity, which is why there are no points above the dashed line in Fig. 13.

Of the 3964 PMTs tested, 156 (3.93%) fail to meet the non-linearity requirement. This has proven to be the most critical cut in the testing process as the bulk of the rejected PMTs are due to excessive positive non-linearity. Recently, however, Photonis has improved the linearity of the PMTs to eliminate the positive non-linearity effect. Results from the first batches of the new PMTs indicate that the mean value of the maximum positive non-linearity is less than 1%, and this change does not affect the acceptable behavior of the other characteristics of the PMTs.

2.2.5. Dynode to anode ratio

To enable the detection of small signals in the water Cherenkov detectors and extend the dynamic range of the detector, the signal from the last dynode is extracted and amplified. The amplification is fixed via the electronics on the base of the PMT to be a factor of 40. There are then two signals from the PMT, the amplified dynode and the signal from the anode. As stated before, the calibration is done with single muons which give a small signal and are recorded using the amplified signal from the dynode. Once a large number of particles pass through a station from an air-shower, the dynode reaches the maximum dynamic range of the PMT readout electronics and the signal from the anode is used in the analysis instead. To be able to extend the calibration using muons, which is measured using the signals from the amplified dynode, to the signals from an air-shower, recorded using signals from the anode, the amplification of the signal from the dynode when compared to the signal from the anode must be known.

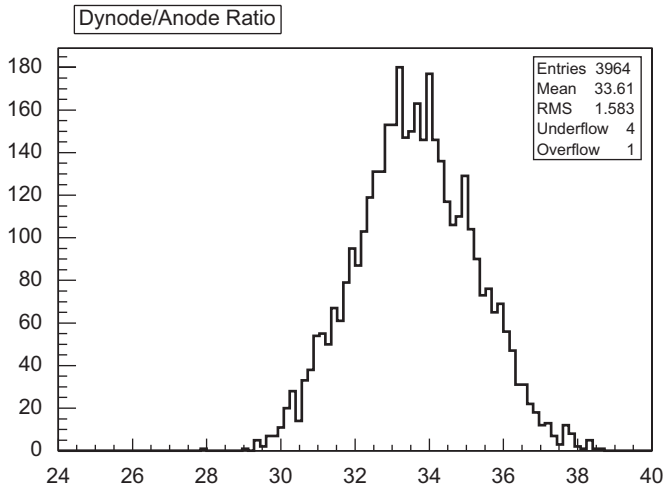

Fig. 14. Distribution of dynode to anode ratios at a gain of 10^6 for 3964 PMTs. 5 PMTs (0.13%) failed to meet the requirement that the value be between 25 and 40.

This is known as the dynode to anode ratio, and it depends on the gain of the last dynode since the amplification of the dynode signal is fixed to a value of 40.

$$D/A = \alpha = 40\frac{\delta - 1}{\delta}. \tag{8}$$

In Eq. (8), the relationship between the dynode to anode ratio (α) and the gain of the last dynode (δ) is defined. The factor of 40 is the value of the gain of the amplifier. The factor of $(\delta - 1)/\delta$ represents that for every electron that hits the last dynode, δ electrons leave. This gives a signal of $1 - \delta$ on the last dynode and δ on the anode (in arbitrary units). The amplifier inverts and amplifies the signal from the last dynode to make it the same polarity as the anode.

To measure the dynode to anode ratio, the PMTs are set to a fixed gain and the light source is flashed at varying intensities. The signal of the dynode versus the signal of the anode is plotted and the slope of the resulting line gives the dynode to anode ratio. The distribution of dynode to anode ratios at a gain of 10^6 is shown in Fig. 14. From Eq. (8), the dynode to anode ratio depends on the gain of the last dynode, and the gain of the last dynode is dependent on the voltage applied to the PMT. A plot of the gain of the last dynode as a function of input voltage for a typical PMT is shown in Fig. 15, where offset and slope refer to γ and β from Eq. (6). The mean gain of the last dynode when applying the voltage necessary for an overall gain of 10^6 is ~6, meaning that the nominal value of the dynode to anode ratio is ~33, and accepted values range from 25 to 40. Of the 3964 tested PMTs, 5 (0.13%) failed to meet this requirement.

2.2.6. Excess noise factor

The ENF of a PMT is a useful quantity because it is related to the signal to noise ratio. It arises from the statistical nature of the multiplication process in a PMT. Light amplification is achieved via a series of dynodes

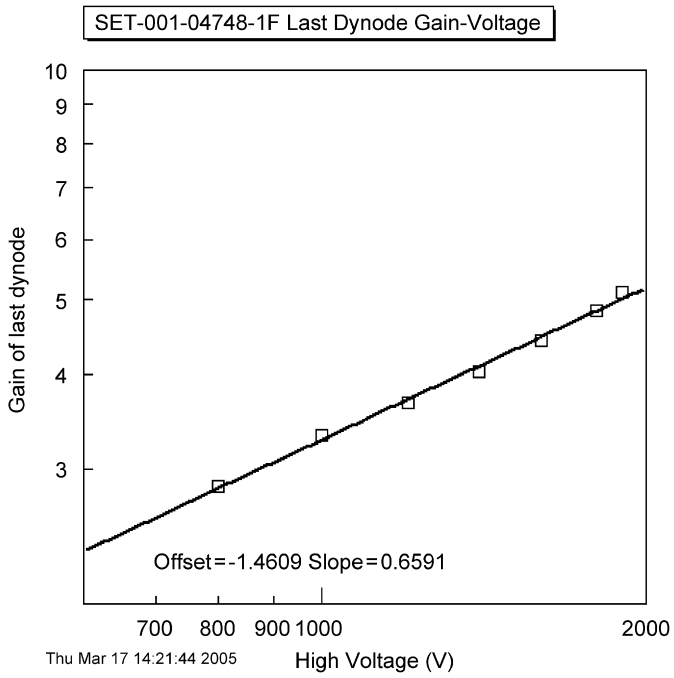

Fig. 15. Gain of the last dynode versus input voltage for a typical PMT.

which multiply the number of incoming electrons. Each dynode stage has a mean secondary emission, δ, and a variance, which results in the output signals being more spread than the distribution of the incoming photons.

Following the arguments as presented in Teich et al. [8], the ENF is:

$$\mathrm{ENF} = 1 + \frac{\sigma_G^2}{G^2} \tag{9}$$

where G is the mean amplification, or gain of the PMT, and σ_G^2 is the variance in the gain. Using an expression for the variance of the gain from Teich et al. and assuming that the secondary emission of each dynode stage is a Poisson process, the ENF can be expressed as

$$\mathrm{ENF} = 1 + \frac{1}{\delta_1} + \frac{1}{\delta_1\delta_2} + \cdots + \frac{1}{\delta_1\delta_2\cdots\delta_N} \tag{10}$$

for a PMT with N dynode stages, where δ_n is the mean secondary emission of the nth dynode.

To measure the ENF, it is useful to note that if the distribution of incoming pes follows a Poisson distribution, then the output variance will be the input variance multiplied by the ENF:

$$\left(\frac{\sigma_{\mathrm{out}}}{S_{\mathrm{out}}}\right)^2 = \mathrm{ENF}\left(\frac{\sigma_{\mathrm{pe}}}{N_{\mathrm{pe}}}\right)^2 \quad \text{or} \quad \mathrm{ENF} = N_{\mathrm{pe}}\left(\frac{\sigma_{\mathrm{out}}}{S_{\mathrm{out}}}\right)^2 \tag{11}$$

where S_{out} is the mean of the output signals, σ_{out} is the spread of the output signals, and N_{pe} is the mean number of incoming pes.

In the test system, the PMTs are set to a fixed gain and the LED is pulsed multiple times at an intensity such that each PMT receives ~100 photoelectrons. The mean and variance of the output signals is computed, and since the

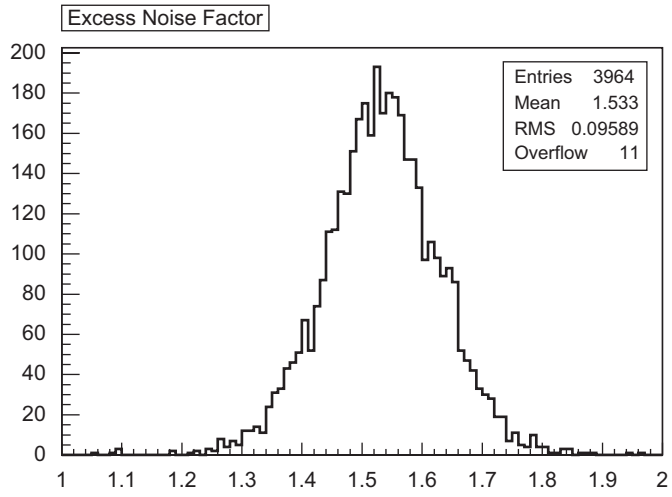

Fig. 16. Distribution of excess noise factors at a gain of 2×10^6 for 3964 PMTs.

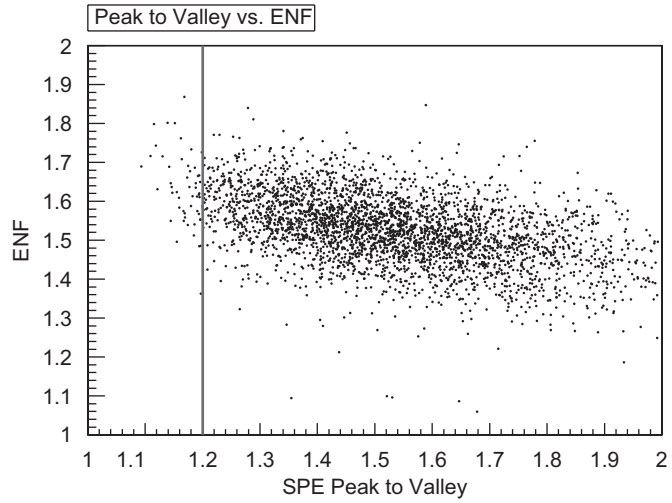

Fig. 17. ENF versus peak to valley ratio showing an anti-correlation. The grey line represents the peak to valley specification given to Photonis (>1.2).

gain is known, it is possible to determine the mean number of pes for a given PMT using only the mean of the output signals (S_{out}):

$$N_{\text{pe}} = kS_{\text{out}}/G \tag{12}$$

where k is a constant related to the DAQ electronics. It is then straightforward to compute the ENF using Eq. (11). The distribution of ENFs at a gain of 2×10^6 is presented in Fig. 16.

The ENF is related to the peak to valley ratio of the single pe spectrum. The larger the ENF, the broader the distribution of output signals will be for the same input. Therefore, we expect an anti-correlation between ENF and peak to valley ratio. There is no specification for the ENF of a PMT, but because it is related to the peak to valley ratio, any excessive noise will cause a failure in the peak to valley requirement. The ENF versus the peak to valley ratio is shown in Fig. 17.

2.2.7. Afterpulse

One concern with PMTs is contamination of gases. The PMT is made of a glass envelope around a dynode structure with a vacuum inside the glass tube. If there are molecules of gas inside the glass envelope, as the pes pass through the gas the molecules will ionize and these ions will travel back to the glass where they will eject more electrons. This will cause a pulse proportional to the initial pulse delayed in time anywhere from hundreds of nanoseconds to microseconds, depending on the gas. Ultimately, this could cause a miscalculation of the energy deposited in a surface detector.

In the PMTs tested, there is no significant afterpulsing, indicating that the vacuum is free from gases. There is, however, a systematic negative value for the afterpulse measurement (see Fig. 18). The negative value in the afterpulse is caused by a shift in the baseline after a large signal, as is the case in this test, and is a property of the DAQ system, not the PMT. The baseline is shifted to a value that is smaller after a large signal, and when integrating over $\sim$5 μs the result is that there appears to be a negative value for the afterpulse because the afterpulse percentage is defined as

$$\text{AP}\ (\%) = 100 \times \frac{Q_{\text{after}} - Q_{\text{baseline}}}{Q_{\text{signal}} - Q_{\text{baseline}}} \tag{13}$$

where Q_{signal} is the charge deposited while the LED is flashing, Q_{after} is the charge after the initial pulse, integrating over 5 μs, and Q_{baseline} is measured before the LED flashes. It should be noted that the shift in the baseline is negligible when integrating for less than 500 ns, as is the case in all other measurements.

The specification requires that the afterpulse ratio be less than 5% when integrating for up to 5 μs. Of the 3964 PMTs tested, 12 (0.30%) failed to meet this requirement.

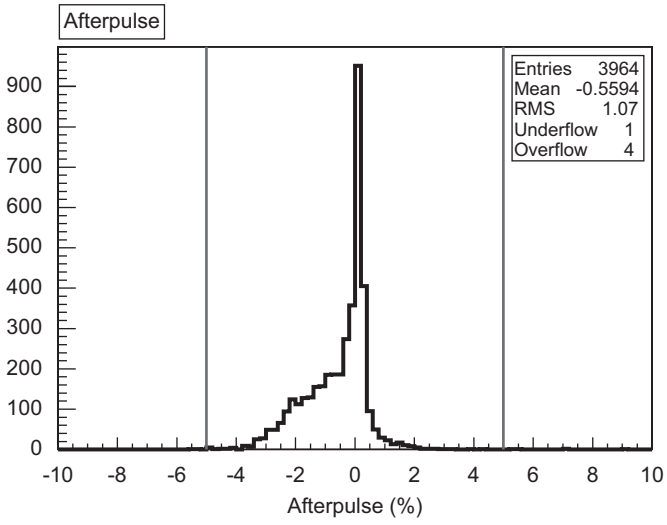

Fig. 18. Afterpulse distribution for 3964 PMTs, with the grey lines representing the specification that the afterpulse be less than 5% of the total signal. Twelve PMTs (0.30%) failed to meet the requirement.

3. Test system performance

The test system, as it is operated, has four PMTs that are left permanently in the test stand. The results from these PMTs are used to monitor the performance of the system as a whole, from the LEDs to the DAQ electronics. The permanent PMTs provide a comparison for tests run currently versus tests run when the system was first commissioned, to see any systematic shifts or anomalous behavior. In addition, checking the spread of the measurements of a given parameter for the four permanent PMTs indicates the resolution of the test system, correcting for any time and temperature effects. It is worth noting that the temperature inside the dark room is recorded for each test, and that there are no noticeable temperature effects for any test in the temperature range from 12 to 27 °C.

Fig. 19 is a plot of the voltage necessary for a gain of 10^6 for the four permanent PMTs during a period of 1000 days. There is no noticeable drift of this value with time, and any temperature effects are lost in the spread of the voltages. One PMT is taken as an example, PMT 892, to show the spread of the measurements over this same time period (see Fig. 20). For this PMT, the standard deviation is 10.5 V with a mean of 1525 ± 0.5 V. Fixing the voltage to the mean value, the standard deviation of the voltage corresponds to a standard deviation of less than 4% on the gain. Repeating this process for all the tests, the resolution of the system is determined for each test and is reported in Table 2.

To illustrate the monitoring capabilities of the four permanent PMTs, two examples are given. The first is shown in Fig. 21. In this figure, the dark pulse rate at $\frac{1}{4}$ pe threshold is shown versus time, for 1000 days. For PMT 879, there is a period of a steady decrease in the dark pulse rate, followed by a period of stability. This behavior is demonstrated solely by this PMT, therefore it is not necessary to make any corrections to the PMT test results.

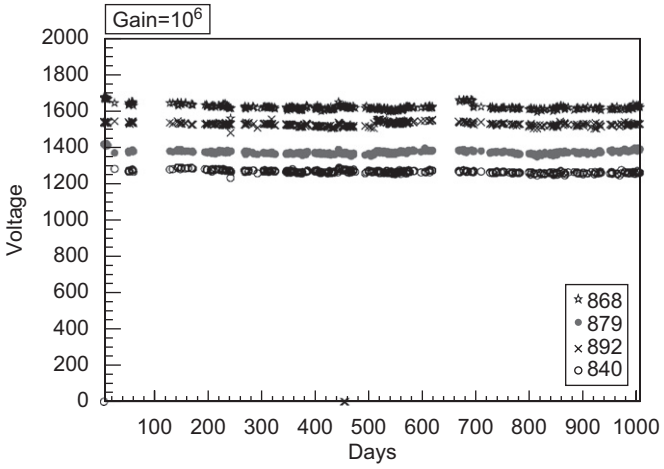

Fig. 19. Voltage to get a gain of 10^6 for the four permanent PMTs over a 1000 day period. The consistency of the value for a given PMT is what is monitored.

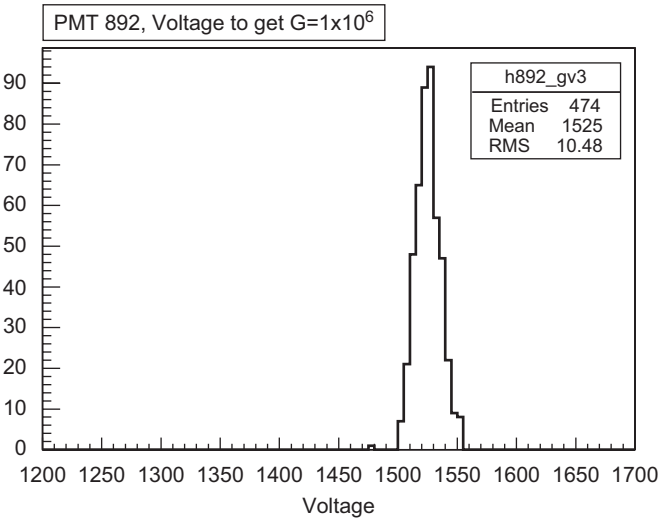

Fig. 20. Distribution of the voltage for a gain of 10^6 for PMT 892 over a 1000 day period.

Table 2
Resolution of the PMT test stand based on the standard deviation of the measurements of the specified tests for the permanent PMTs

Test	868	879	892	840	Average
SPE Peak to Valley ratio (%)	4.5	4.7	5.5	5.7	5.1
$G = 2 \times 10^5$ Voltage (%)	0.64	0.48	0.62	0.57	0.58
$G = 1 \times 10^6$ Voltage (%)	0.66	0.46	0.69	0.61	0.61
$G = 2 \times 10^6$ Voltage (%)	0.72	0.48	0.48	0.66	0.59
Dark pulse rate (Hz)	897	963	643	764	817
Dynode to anode ratio	0.34	0.40	0.41	0.31	0.37
Excess noise factor	0.07	0.06	0.06	0.07	0.07
Non-linearity at 50 mA (%)*	0.35	0.39	0.49	0.38	0.40
Max. positive non-linearity (%)*	0.31	0.39	0.48	0.38	0.39
Afterpulse (%)*	1.5	2.0	2.0	1.2	1.7

For the peak to valley and the gain voltage values, the percentage is given of the standard deviation to the mean value. For all other tests, the value given is the raw value of the standard deviation of the measurements because the spread is independent of the mean value of the observable. Values with an asterisk (*) denote that the value is the raw spread in the variable which is a measurement of a percentage and are not to be confused with the percentage of the standard deviation to the mean value.

In the second example, however, in the non-linearity measurements there was a drift detected in the system starting around day 300 and recovering from day 500 to 600, see Fig. 22. Since all four PMTs experienced the same change in behavior (with PMT 879 experiencing an abrupt shift independently), the drift can be attributed to something that is happening in the test system itself and the results for the tested PMTs can be corrected for this behavior. The cause of this drift is unknown, but it can be monitored and the results can be adjusted accordingly.

To determine any systematic effect associated with location in the test stand, test results are also plotted as a function of position in the test stand, see Fig. 23. Each position in the test stand is in a fixed location, meaning the

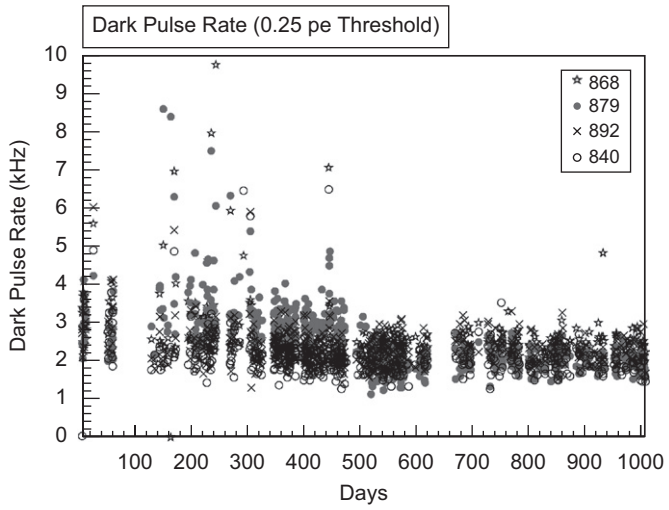

Fig. 21. Dark pulse rate at $\frac{1}{4}$ pe threshold for the four permanent PMTs. PMT 879 experiences a steady decrease over the first 500 days.

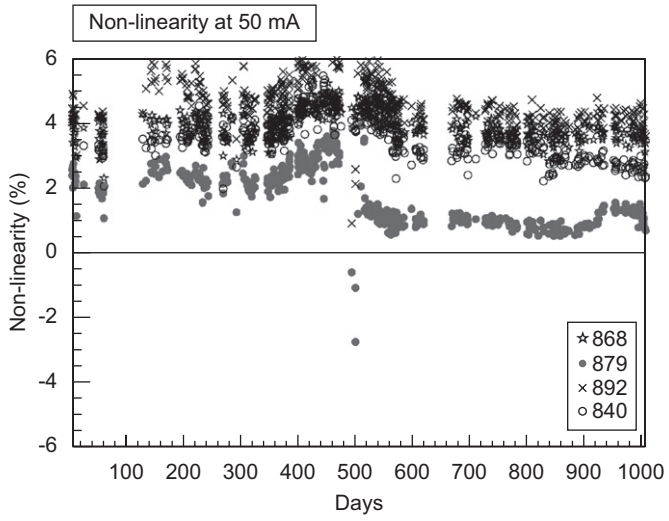

Fig. 22. Non-linearity at 50 mA for the four permanent PMTs. A system wide drift is noticed starting at day 300, peaking at day 500, and recovering at day 600.

orientation with respect to the LEDs is constant. Each position is also associated with a fixed channel in the data acquisition electronics. In the plots in Fig. 23, each data point for the given location in the test stand has anywhere from 200 to 350 PMTs to compute the average. In the bottom plot, it is shown that the mean dark pulse rate does not depend on the location in the test stand, whereas non-linearity measurements (top plot) vary with the location in the test stand by around 1%. Any systematic effect due to the location in the test stand can be corrected in the final results.

4. PMT testing conclusions

Out of the 3964 PMTs tested, 179 (4.51%) have been rejected for failing to conform to the specifications given to

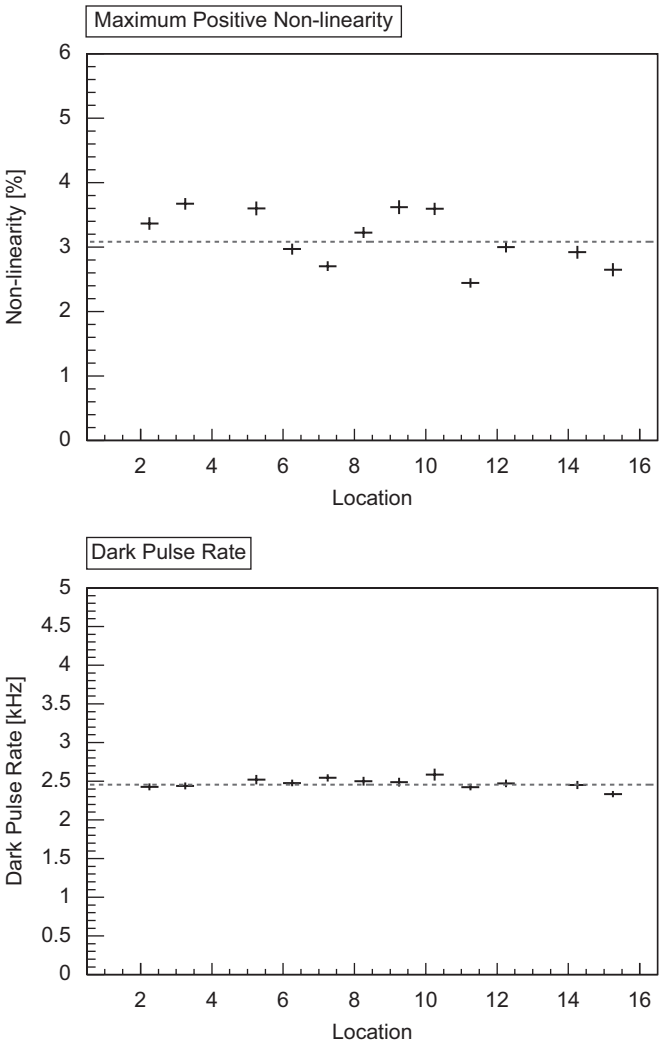

Fig. 23. Plots as a function of position in the test stand. In each, the dashed line is the mean value of that parameter for all tested PMTs. Top: Maximum positive non-linearity. Bottom: Dark pulse rate.

Table 3
Rate of failure for tested PMTs

Test	Specification	Failed	Failed (%)
SPE peak to valley[a]	>1.2	55	1.59[b]
Gain versus voltage	10^6 gain with V<2000 V	3	0.08
Dark pulse rate	<10 kHz at $\frac{1}{4}$ pe threshold	12	0.30
Non-linearity	<6% below 50 mA peak current	156	3.93
Dynode to anode ratio	Between 25 and 40	5	0.13
Afterpulse ratio	<5%	12	0.30

[a]No PMTs are rejected for failing SPE peak to valley specification.
[b]The percentage is calculated using the number of PMTs for which the peak to valley ratio was successfully determined (3463).

Photonis, see Table 3. This does not include the PMTs that failed the peak to valley requirement as no PMTs are rejected for failing this specification.

Studying Table 3 will show that the most critical cut is the requirement that the non-linearity be less than 6%, as 156 PMTs failed this requirement, or 3.93% of the PMTs. This failure rate is an order of magnitude larger than the next most frequent failures, the dark pulse rate and afterpulse ratio, which are at 0.30% each. However, as discussed earlier, Photonis has improved the non-linearity of the PMTs while not altering the performance in the other areas of concern, so the failure rate due to non-linearity should decrease by an order of magnitude according to preliminary studies.

What is not expressed in the table is the frequency of PMTs failing more than a single test. Of the PMTs that failed to meet specifications (ignoring the peak to valley results), eight failed more than a single test. Of those eight PMTs, one failed three tests: gain voltage, dynode to anode ratio, and afterpulse ratio. The most frequent combination of failed tests is non-linearity with dark pulse rate, with three PMTs failing both. The other test combinations were non-linearity with afterpulse ratio and non-linearity with dynode to anode ratio where two PMTs each failed both tests respectively. The frequency of PMTs that failed more than a single test implies that test failures are correlated. If the assumption were made that test failures were independent and the probability to fail a particular test is given by the fractions in Table 3, then the probability to see eight out of 3964 PMTs fail more than one test would be less than 0.01%.

References

[1] J. Abraham, P. Auger Collaboration, et al., Nucl. Instr. and Meth. A 523 (2004) 50.
[2] Pierre Auger Collaboration, Proceedings of the 29th International Cosmic Ray Conference, Pune, vol. 7, 2005, p. 369.
[3] B. Genolini, et al., Nucl. Instr. Meth. A 504 (2003) 240.
[4] A. Tripathi, et al., Nucl. Instr. Meth. A 497 (2003) 331.
[5] X. Bertou, Pierre Auger Collaboration, et al., Nucl. Instr. Meth. A 568 (2006) 839.
[6] R.W. Engstrom, RCA Photomultiplier Handbook (PMT-62), RCA Electro Optics and Devices, Lancaster, PA, 1980.
[7] Pierre Auger Collaboration, Pierre Auger Project Design Report, second ed., November 1996. Revised March 14, 1997.
[8] M.C. Teich, K. Matsuo, B.E.A. Saleh, IEEE J. Quant. Electron. 22 (1986) 1184.