SNO-STR-99-039

The Gain and the Pain: A Study of Detector Response and Instrumental Backgrounds After the High Voltage Raise

James Cameron, Bill Frati, Rich Helmer, Chris Jillings, Josh Klein, Neil Mccauley, Gordon Mcgregor, Mark Neubauer, Richard Van de Water, Peter Wittich

Oxford University, Queens University, TRIUMF, University of Pennsylvania

December 21, 1999

Abstract

We describe here a study of both the gain increase from the recent high voltage raise and the increase in high-voltage related instrumental backgrounds. We find that, as anticipated, the average PMT gain increase is roughly 35%, which puts our average discriminator threshold down at roughly 1/4 photoelectron, and this results in an overall detector efficiency increase of 8-9%. In addition, the increase in gain has lead to an expected flattening of the occupancy vs. threshold curve, thus reducing the overall differences in occupancy tube-to-tube. There has been no significant increase in wet-end breakdown events, however there has been an increase in both sharkfins and flashers. The flasher rate showed an initial increase well above expectations, but has since settled down to levels just slightly above the original prediction.

1 Introduction

The recent raise in high voltage was undertaken to increase both the detector efficiency and stability. The increase was done over period of five days, in discrete steps: +40 V, + 80 V, and +100 V, with the exception of the crate at the highest voltage (Crate 9) which was raised only by +50 V. As anticipated, the increase did not come without a price: high voltage-related phenomena such as wet-end breakdown and flashers are known to be high-voltage dependent, and in the case of flashers the increase at the new voltages was easily observable. The most dangerous of the effects from a detector standpoint is the wet-end breakdown, but from the standpoint of neutrino backgrounds it is likely to be the flashers.

In two previous memos [1, 2], we presented the state of the PMT gain and relative efficiency at the previous nominal voltage as well as a full accounting of the high-voltage instrumental backgrounds in the detector. In this memo, we try to combine these two goals, documenting both the increase in gain and the increase in some of the high-voltage related backgrounds.

2 Gain

There are several ways to measure the increase in gain with high voltage. The most relevant, ulimately, is simply the increase in overall detector energy scale in response to a source (like the ^{16}N). However, to ensure that the increase in overall response conforms to our expectations of PMT behavior, it is important to study the PMT response itself to measure (at least on average) the increase in tube gain with high voltage. Seeing the effect at the PMT level goes beyond simply looking at the change in the charge spectrum, however, since part of the original claim was that the increased gain would make the differences in tube occupancies as a function of discriminator threshold smaller. This flattening means that our energy resolution will be less affected by the various threshold settings in the detector, ultimately leading to a narrower resolution function.

Run Number	Date	Run Type
4369	July 9	Gated Charge
5478	October 12	Standard PCA/Laser Occupancy (7% occ.)
5527	October 15	Standard PCA/Laser Occupancy (7% occ.)
5683	October 24	Standard PCA/Laser Occupancy (7% occ.)
5644	October 22	Gated Charge
5646	October 22	Gated Charge, Laser Blocked
10237	November 28	Standard Neutrino

Table 1: Runs for Study of Gain Effects of High Voltage Raise.

2.1 Data Sets

The data sets were used for the studies described in Section 2 are shown in Table 1.

2.2 PMT Gain Change

To look directly at the PMT gain change at the new high voltage, we have examined both standard laser discriminated charge spectra as well as 'gated' charge spectra in which each channel is read out regardless of whether its discriminator fires. The discriminated data has the advantage that this is the way we normally take data, acquiring these data sets can be done either during laser PCA calibrations or even with standard NHIT data, and also allow one to obtain high tube statistics with little data overhead. The statistics argument also means that we can take these data sets at low occupancies and still produce a reasonable measurement, so the spectra are more guaranteed to be dominated by single, rather than multiple, photon hits. On the other hand, the gated charge runs, although they 'look' like high NHIT data as far as the acquisition system is concerned because a large fraction of the channels are fired by the gate on each event, actually take some time to accumulate because few of the fired channels actually contain any data. The gated charge data, however, has the advantage that we get to see the entire charge spectrum without any hardware threshold imposed. This allows us to get an idea of the position of the true peak of the single photoelectron spectrum, and allows us to impose software thresholds (in charge, not

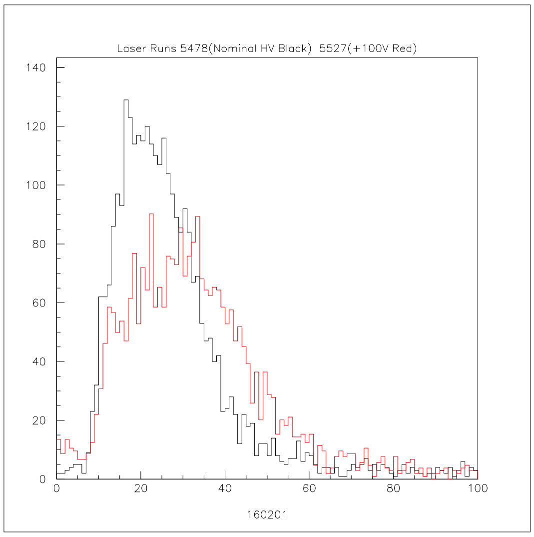

Figure 1: Discriminated single PE spectrum for a single channel in crate 16, shown before and after high voltage was raised by +100 V.

voltage) which are guaranteed to be uniform channel-to-channel.

Figures 1 and 2 show the gain change for two different channels using the standard laser discriminated data. Figure 3 shows the gain change for a channel in Crate 9, which, being the crate with the highest nominal voltage, was left at only 50 V above nominal. It is clear that, at least for this channel, the change in gain for Crate 9 was marginal. For the other channels shown, however, it is clear that there is a substantial change in gain. Figure 4 shows the ratio of gains before and after the high voltage raise. For this plot, the gains were determined using the 'high half point'—the point on the high side of the single pe charge spectrum where the number of counts is half that in the peak—of the (QHS) charge distribution. This method has greater robustness against changes in the shape of the distribution due to the varying thresholds channel-to-channel than one which uses either the mean or the peak itself. From the Figure, the gain change (excluding crate 9) looks to be about 35% between the two voltages. For Crate 9, Figure 5 shows a gain change of about 15%.

The results from the gated charge data are in reasonable agreement with the discriminated data, though they tend to point to a somewhat smaller change in gain. Figures 6 and 7 show the changes in the spectrum of sample single channels. Figure 6 is the same channel shown

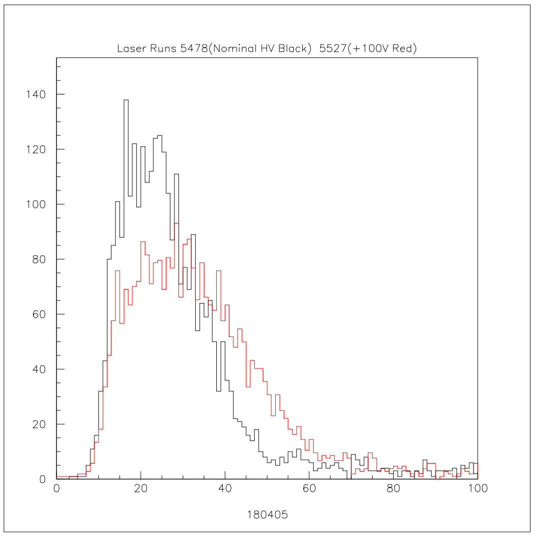

Figure 2: Discriminated single PE spectrum for a single channel in crate 18, shown before and after high voltage was raised by +100 V.

in the previous memo [1] on threshold efficiencies, and Figure 7 is the same channel as is shown in Figure 1. Comparing Figures 1 and Figure 7, one can see that (for this channel) the 'high half point' method works well, as the gated spectrum has the same high-side tail as the discriminated one. Although only this one channel is shown here, looking at many others verifies the general conclusion. Comparing the gated and discriminated data also allows us to see the effects of the higher occupancy of the gated charge data (taken at $\sim 20\%$ rather than the discriminated data's 7% or so occupancy). From the Figures (and looking at many other spectra not shown here) we can see that the contribution from multiple photon hits in the gated charge data is not a significant effect on the high side of the charge spectrum.

Figure 8 shows a gain comparison between the nominal high voltage charge spectrum and the raised high voltage spectrum. In this plot, the gated charge spectra (QHL) for all channels from the even crates (even and odd crates are done separately to lower the data rate) have been summed together. An important caveat is that for the raised high voltage spectrum, a cell-by-cell correction to the standard ECA pedestals has been done using gated charge data in which the laser light was blocked. For the nominal data, a global 4.8 count offset has been subtracted, and the residual smearing of the pedestal can be seen compared to

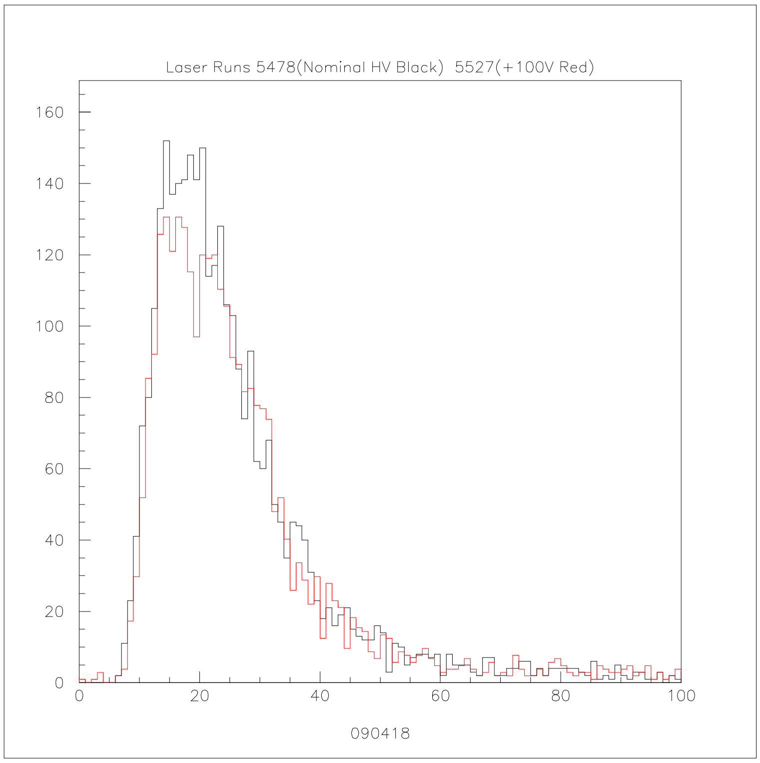

Figure 3: Discriminated single PE spectrum for a single channel in crate 9, shown before and after high voltage was raised by +50 V.

the raised high voltage spectrum (the reason that the standard ECA pedestals do not agree with the gated charge pedestals is still unknown). The rms of the pedestal for the raised high voltage data is less than two counts (the scale and occupancy make it look larger) and the additional smearing in the nominal high voltage data is probably on the order of two or three counts. Nevertheless, the difference in gain is quite clear, and the peak of the spectrum for the raised high voltage stands out clearly. It is not fair, however, to infer a peak/valley ratio from this data, since that will be somewhat occupancy dependent. The relative normalization for the two spectra was done using the above-pedestal peaks of the two distributions, and the occupancies for both sets were roughly the same.

In Figure 9, both the gated charge spectrum and the discriminated spectrum for same channels is superimposed. The discriminated spectrum was taken from a subsequent run,

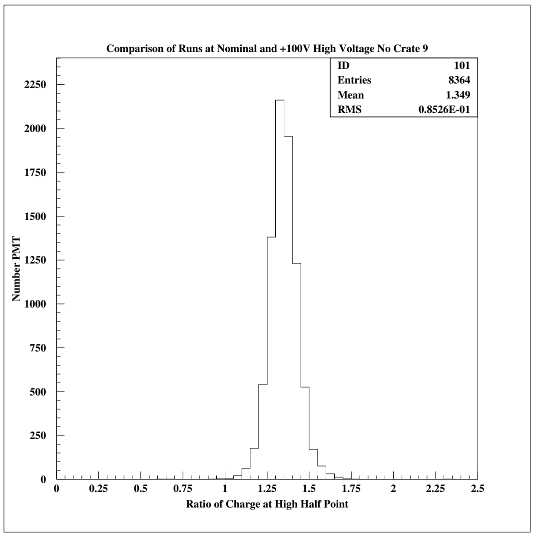

Figure 4: Ratio of gains determined using ‘high half point’ of charge spectrum, shown before and after high voltage was raised by +100 V, for all channels except those in Crate 9.

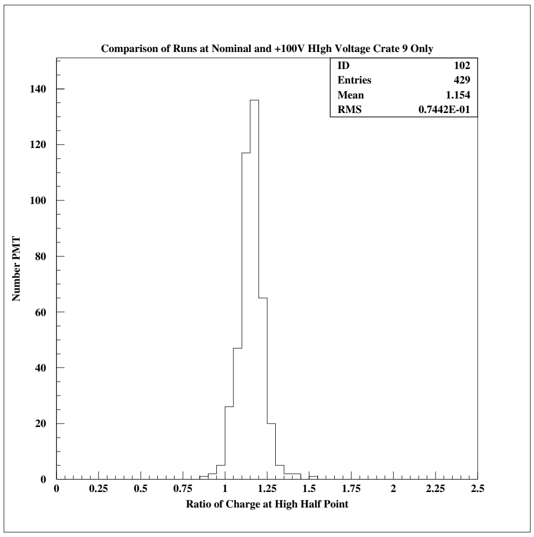

Figure 5: Ratio of gains determined using ‘high half point’ of charge spectrum, shown before and after high voltage was raised by +50 V, for channels in Crate 9.

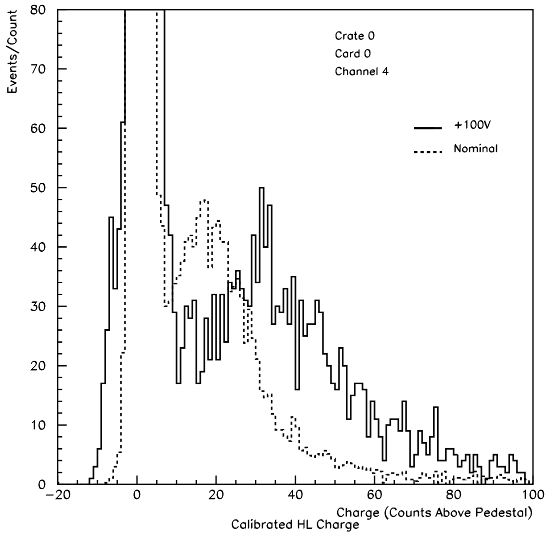

Figure 6: Gated single PE spectrum for a single channel, shown before and after high voltage was raised by +100 V.

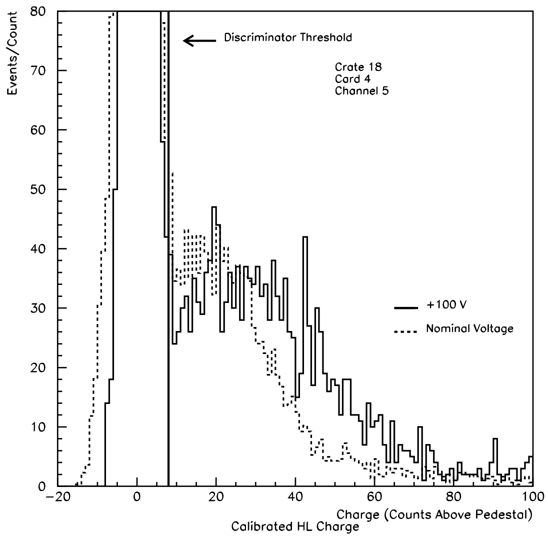

Figure 7: Gated single PE spectrum for a single channel, shown before and after high voltage was raised by +100 V.

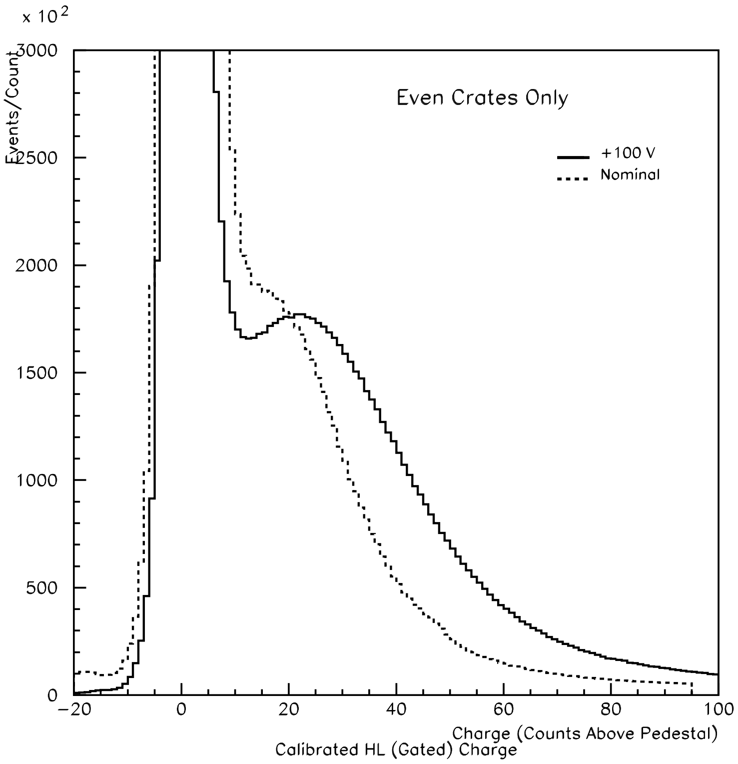

Figure 8: Gated single PE spectrum for all channels in even crates, shown before and after high voltage was raised by +100 V. Note that the broadening of the pedestal in the nominal voltage data is partly a result of the fact that the standard pedestals were used rather than pedestals derived from the gated charge data itself.

with all run parameters the same, and the same cell-by-cell calibration determined from the lightless gated charge data has been applied. The normalization here was done to ensure that the high charge tails agreed. In addition, for illustration, the mean discriminator threshold (around 8.5 counts above zero) is indicated on the plot, under the assumption (supported by evidence in the previous document on threshold efficiencies) that ADC charge counts and DAC threshold counts have a one-to-one correspondence. It is clear from this Figure that we are doing quite well now in the fraction of photoelectrons we are keeping—the mean threshold is at the bottom edge of the valley (for this occupancy) and the discriminated spectrum is very close to the gated one (compare this to Figure 7 in the previous threshold efficiency memo [1]).

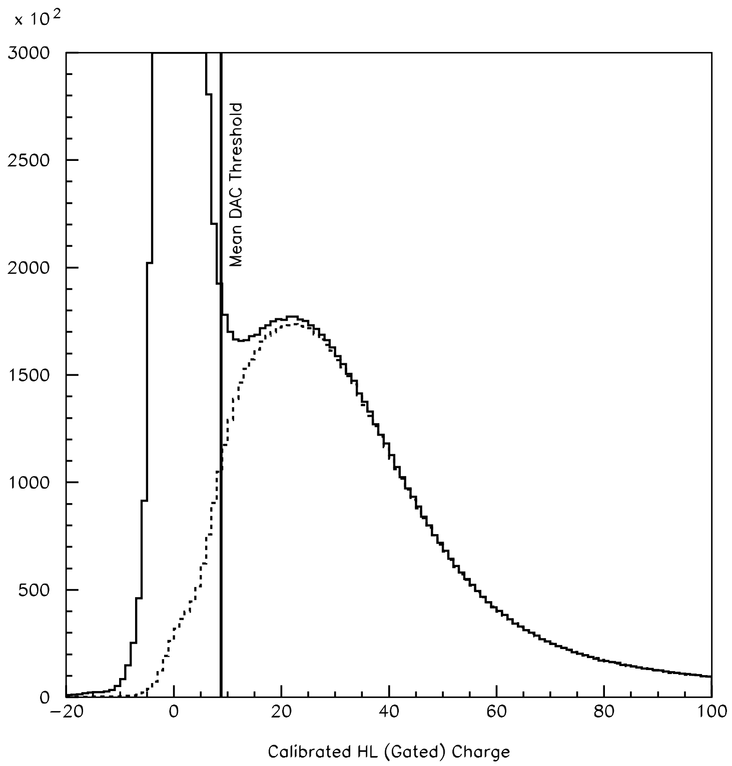

Figure 9: Gated single PE spectrum for all channels in even crates, compared to discriminated single pe spectrum taken at the same occupancy. The 'gated pedestals' have been used to pedestal correct both sets of data.

As stated earlier, one of the advantages of the gated charge data is that we can impose uniform software thresholds, which then allow us to calculate statistics on the charge spectrum without any concern for discriminator effects. Fitting the peak, calculating a mean or a median, are all made easier with the full spectrum. Of these statistics, the most robust for determining relative gain is the median, as fitting the peak of a distribution with a poorly known shape can be a dicey operation, and the mean is too easily pulled by changes in the tails. The drawback of the median in this case, however, is that the software threshold does force the median higher for the lower gain data. Figure 10 shows the relative gains of the nominal and +100 V data, using the medians calculated for both cases above a software threshold of 12 counts above pedestal (and below 100 counts). The 30% gain seen here is lower than what was measured with the discriminated spectrum, and may partly be a result of effects of the software threshold pushing the median, or perhaps from the difference in the pedestal calibrations used in the nominal and +100 V data.

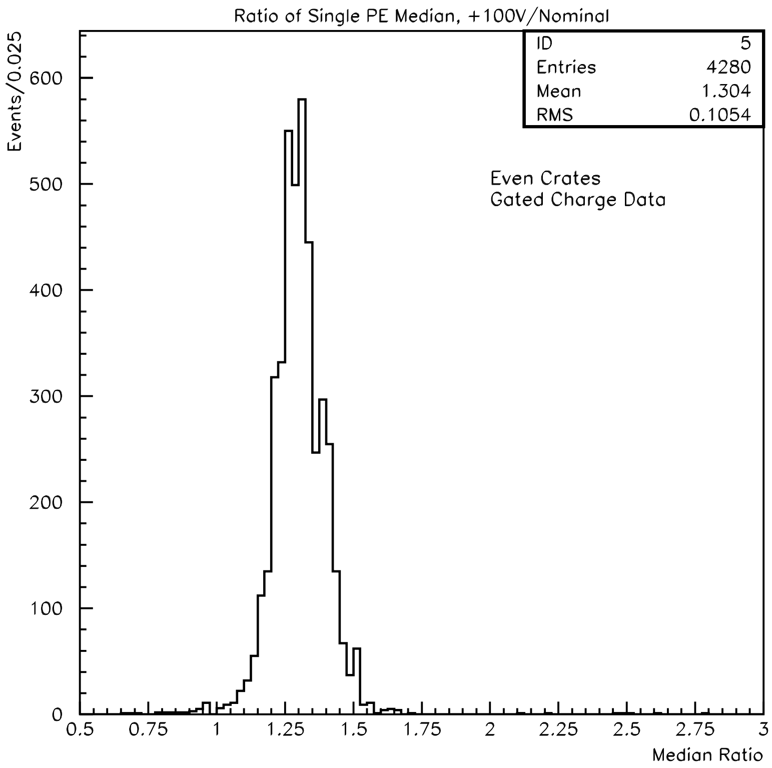

Figure 10: Comparison of median distribution of gated charge spectra, before and after +100 V raise. The median is measured between 12 and 100 counts above pedestal.

We can also use the gated charge data to produce a measure of where the discriminator thresholds sit in 'photoelectrons'. The definition of single photoelectron charge is somewhat arbitrary—whether one uses the mean or median over some interval, or the peak, for example, can change the numbers drastically. The Queens' measurements of the single pe charge used the mean as the definition, taken between 1/4 pe and 2.5 pe. In Figure 11 we show the ratio of the discriminator threshold to the mean of the gated charge distribution taken between 12 and 100 counts above pedestal (roughly 1/3 to 3 pe). This ratio assumes that DAC counts have a one-to-one correspondence with ADC counts, something which was crudely verified in the first memo on threshold efficiencies [1]. As can be seen in the Figure, the average threshold (for the even crates shown here) is a bit less than 1/4 pe, and the median threshold is somewhat lower.

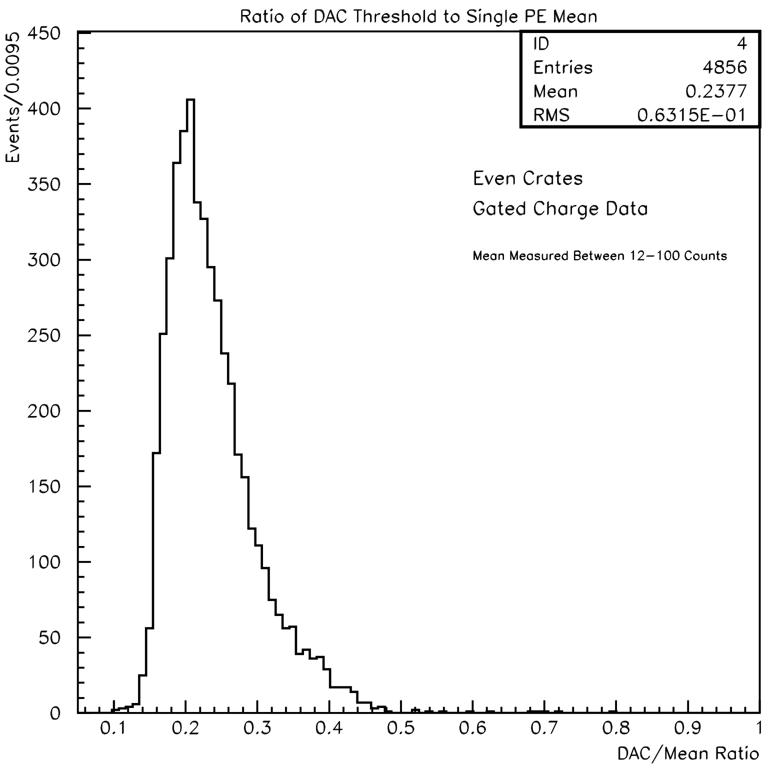

Figure 11: Ratio of discriminator thresholds to mean of the single pe spectrum, under the assumption that DAC counts and ADC counts have a one-to-one correspondence.

2.3 Occupancy Spreads

One of the predicted advantages of the higher gain was a reduction in the slope of efficiency loss as a function of threshold. We expect that as the gain is increased, the effective thresholds move lower and the absolute efficiency of each channel increases. With the absolute efficiency higher, the loss of efficiency with slightly higher thresholds must become lower—there isn't that much efficiency left to lose down in the single pe 'valley'. The ultimate consequence of this is that it reduces the systematic effect of having different threshold settings for different channels.

Figure 12 shows plots of the relative occupancy for channels at each threshold setting for both the nominal and the +100 V data. The difference in slopes is apparent in the plot but obviously not a very big effect. In fact, fitting to the first six points in the two plots (the region that looks linear) shows that the slope has changed from about $3.33 \pm 0.14\%$ per DAC count to $2.86 \pm 0.13\%$ per DAC count. Including all the points shown, the change is somewhat greater, from $3.24 \pm 0.08\%$ to $2.62 \pm 0.14\%$.

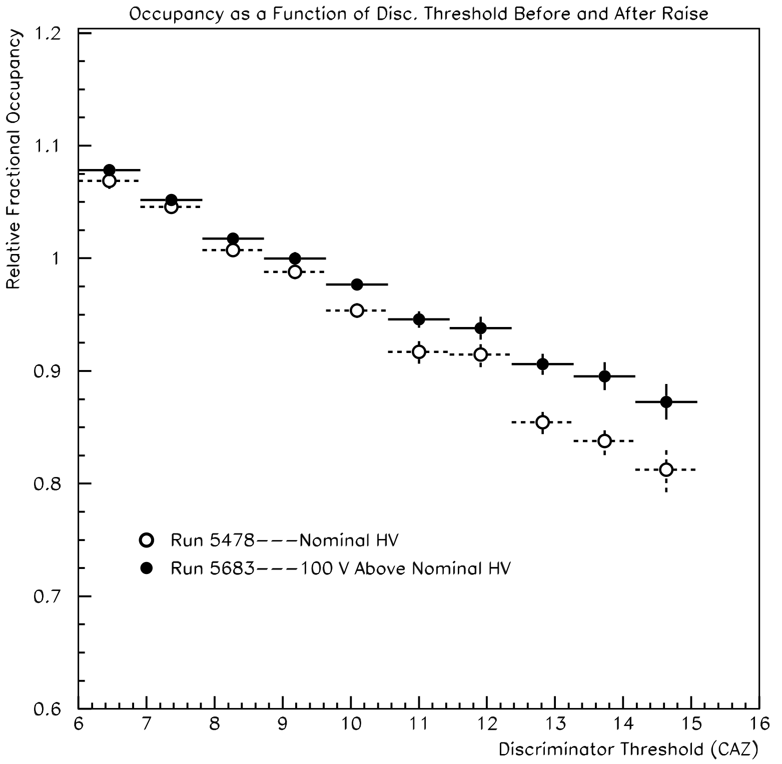

Figure 12: Comparison of the relative occupancy as a function of discriminator threshold, before and after +100 V raise.

We can also look at the degree to which, at our new gain settings, the occupancy is a function of gain (rather than threshold). Figure 13 shows the dependence of occupancy on gain for two different sets of discriminator thresholds, one below 9 counts above zero (CAZ) and one for those channels 9 CAZ and above. As expected, for those channels with high thresholds, the occupancy is a steep function of gain (measured here by the median of the gated charge distribution, taken between 12 and 100 counts above pedestal). However, for those channels with low thresholds (recall, the mean threshold setting is around 8.5 counts above zero, the median is between 7 and 8 CAZ), we see that the occupancy is relatively flat, tailing up slightly for high gains and down steeply (as expected) at very low gains.

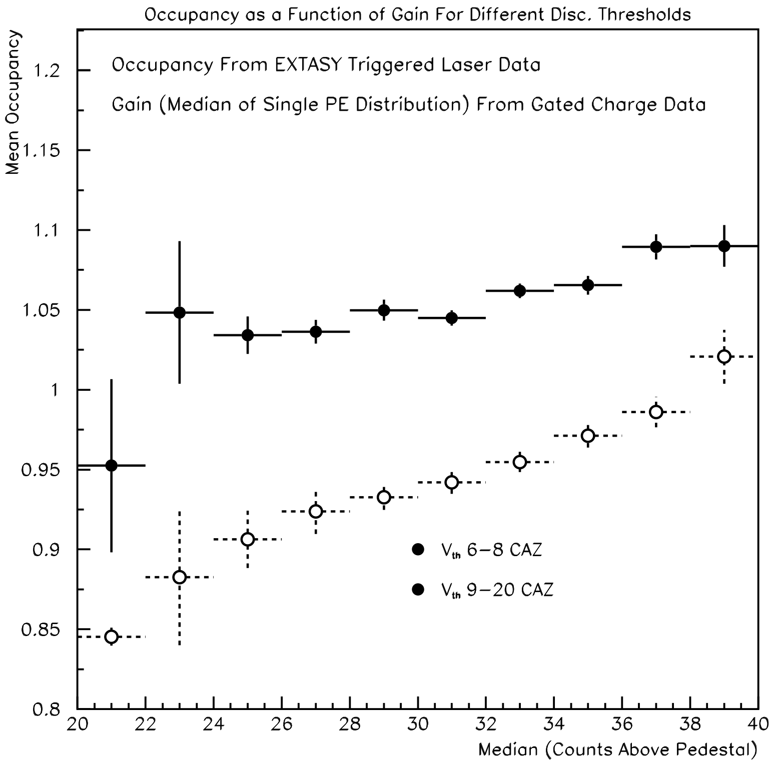

Figure 13: Comparison of the relative occupancy as a function of gain (measured with median of QHL gated charge spectrum) for two sets of thresholds, those below 9 CAZ and those with 9 CAZ and higher.

We can turn this around, of course, and ask 'What does the occupancy look like as a function of threshold, for different sets of tube gains?'. Figure 14 answers this, showing two different sets of tube gains, separated somewhat in the range of their medians. The results here are less spectacular than those in Figure 13—the difference in slopes is noticeable but

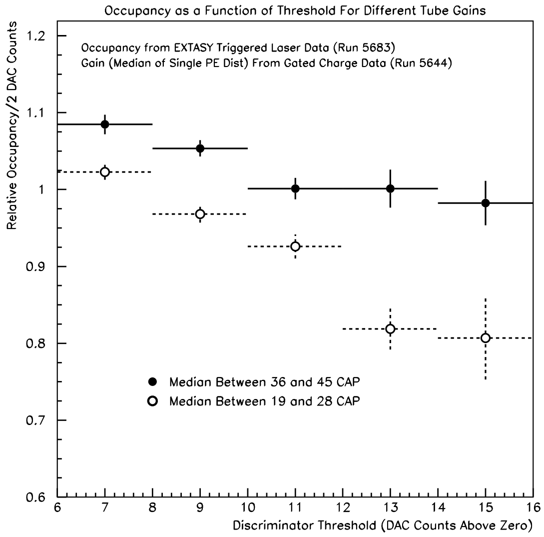

Figure 14: Comparison of the relative occupancy as a function of threshold at two different gains. Gains were determined by the median of the gated charge distribution, taken between 12 and 100 counts above pedestal.

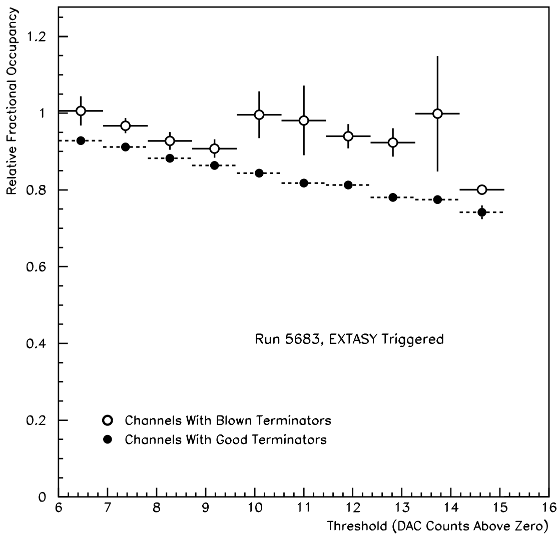

Figure 15: Comparison of the relative occupancy as a function of threshold for channels with blown to those with non-blown terminators.

not overwhelmingly so.

Another way of looking at this information was what was done in the previous document on the relative efficiencies [1]. In that document we plotted occupancy versus threshold for blown and non-blown terminators, concluding that at low thresholds there was not much difference but at higher thresholds the difference was large. Figure 15 is the same plot now for the raised high voltage data. We can draw the same conclusion here as we did then, though there is clearly a systematic effect at higher thresholds which tends to discount this. In fact, it is somewhat disappointing that at the lowest threshold the difference in occupancy is higher now that it was before the high voltage raise. It is possible that this is simply a result of a higher noise rate for the doubled-gain tubes (the thresholds for these tend to be sitting at close to 1/8 pe) compared to before, but the structure at high thresholds indicates that there is probably more going on here. Figure 16 shows both the occupancy as a function for Pulse_GT data (bottom plot) and the CMOS rate versus threshold, for both channels with blown back terminations and channels with good terminations. The Figure hints at a noise-correlation, but disentangling increase noise rate from just the bias of each channel's discriminator is not possible. It is interesting to note that for the channels with good back terminations, the slope of the occupancy versus threshold plot is somewhat steeper ($\sim 3\%$) for the Pulse_GT data than for the laser triggered data.

Despite the flattening out of the response as a function of threshold at the newer high voltage, it is clear that the overall spread in occupancy is still dominated by other factors. Figure 17 compares the occupancy distribution for all operational channels, before and after the high voltage raise. The gaussian fits indicate that there is essentially no change in the spread in the two cases (and that a gaussian is clearly not a good model).

One last effect to check is the difference in occupancy of Crate 9, which was raised on 50 V, to the rest of the crates which were raised by 100 V. Figure 18 shows the comparison of the occupancy distributions of Crate 9 and all other crates. As can be seen from the Figure, although Crate 9 does seem to have a slightly broader distribution, the mean is not significantly different.

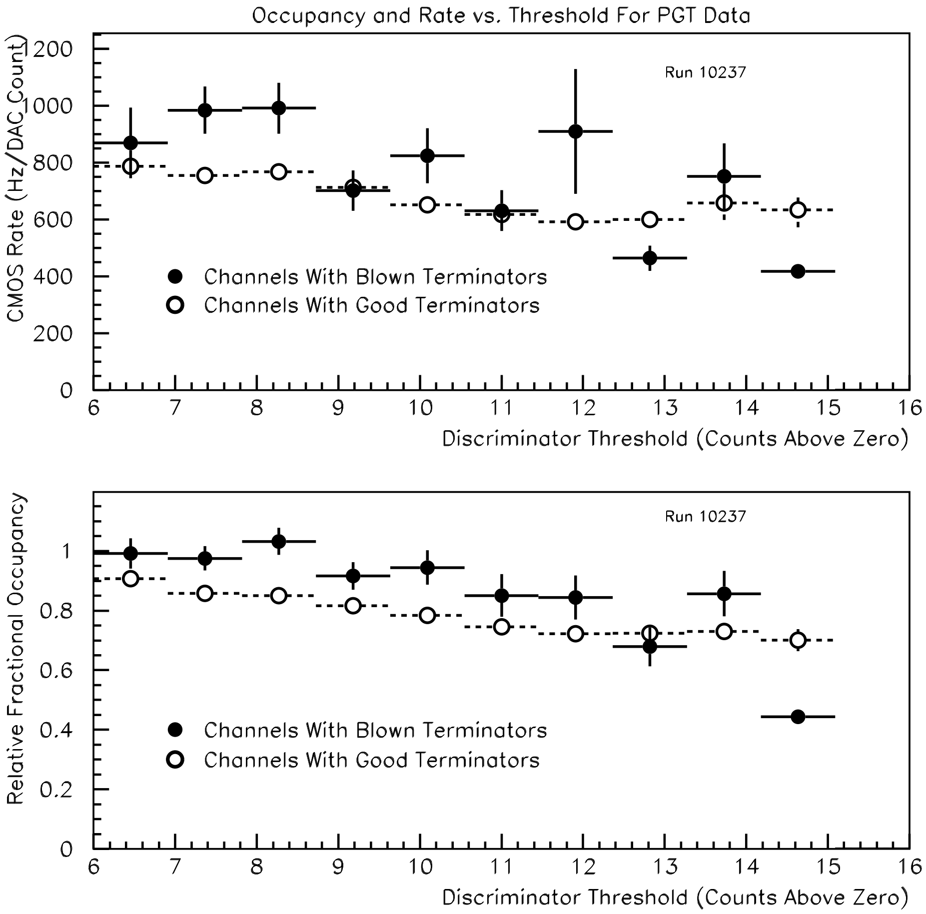

Figure 16: Comparison of the CMOS rate as a function of threshold (top plot) and the occupancy as a function of threshold for PULSE_GT data.

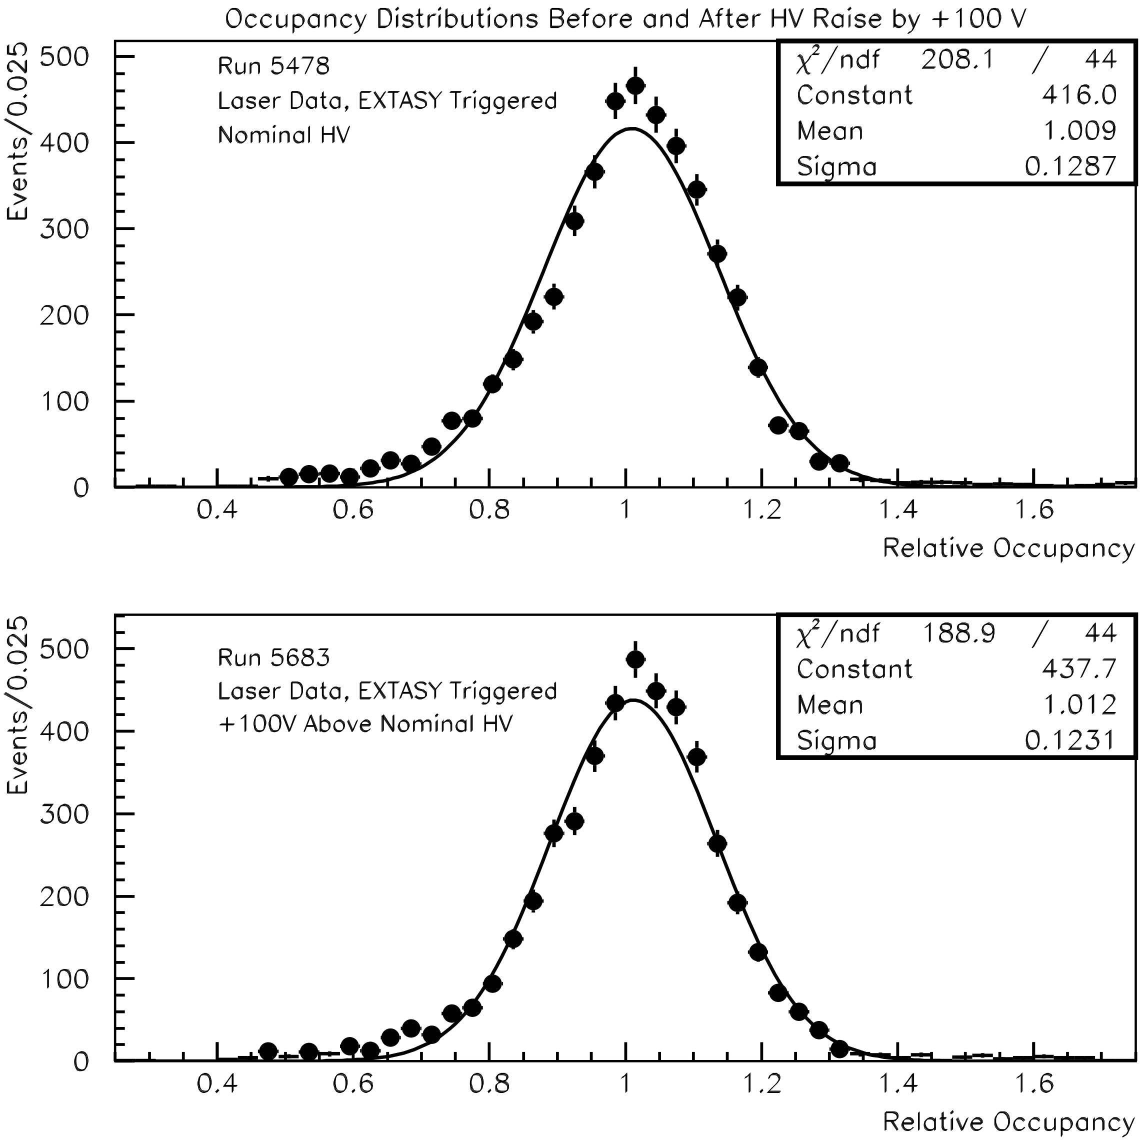

Figure 17: Comparison of the occupancy distributions for the whole detector, before and after the high voltage raise.

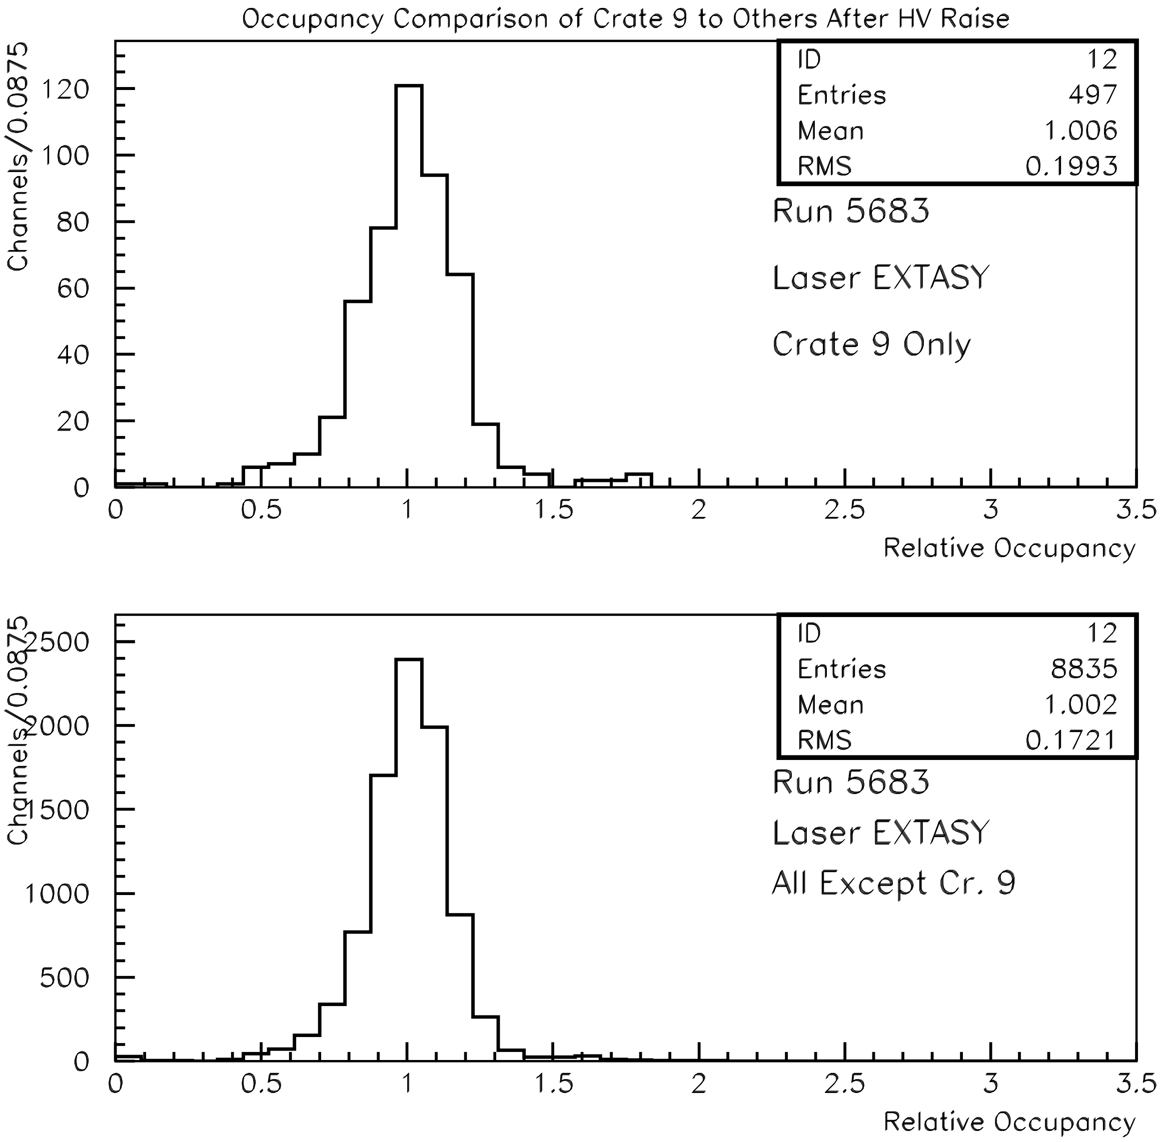

Figure 18: Comparison of the occupancy distributions for Crate 9 to the rest of the detector, after the high voltage raise.

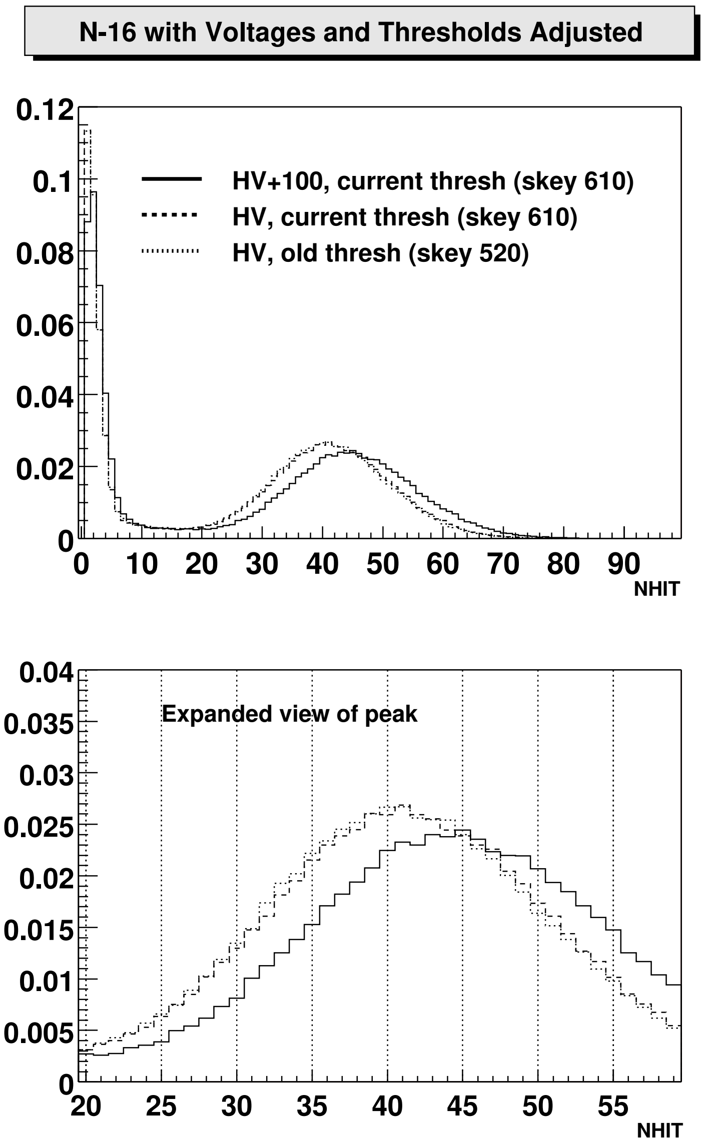

Figure 19: Shift in the detector response to cerenkov events with change in thresholds (dashed curve) and change in high voltage (solid curve).

2.4 Detector Response to Cerenkov Events

By far the most important gain to measure with the raised high voltage is the increase in the overall detector efficiency. Using the ^{16}Nsource, we can see this increase clearly. Figure 19 shows the shift in the NHIT distribution before and after the high voltage raise. The Figure also shows the shift in response due to the recent lowering of thresholds on the exterior slots of each crate. The shift in the mean due to the high voltage raise is roughly 8%, determined by fitting a gaussian over a narrow range near the peaks of the distributions. This increase is slightly lower than the expectation of about 10%. The fact that Crate 9 was only raised 50 V rather than the 100 V that the rest of the detector was raised may have reduced the expected change somewhat, though that would at most be a 5% reduction (that is, reducing the change from 10% to 9.5%). In addition, the lowering of thresholds may have lowered the potential change, as more channels were sitting in the PMT valley with less efficiency to gain by stretching out the single pe spectrum. Overall, the change including the threshold changes was roughly 9%.

3 Pain

The detector stability and performance before and after the +100V increase in nominal high voltage has been carefully monitored and studied. A detailed report of the expected effects of the HV increase has been written [2]. It details and records the baseline HV behavior of the detector, so that subsequent effects of the HV increase, made between October 12th and 14th, 1999, can be monitored and compared.

To ensure that we are being sensitive to potential causes of damage to the detector, five HV sensitive phenomena have been carefully tracked before, during, and after the +100V increase. They are flashers, sharkfins, breakdown, blown terminators, and tube failure rates. Some of these phenomena are expected to increase moderately. However, if any showed major changes beyond what was predicted, then the higher HV experiment would have been terminated and the HV returned to nominal values. The HV was increased in steps of +40V, +40V, and then +20V, each over one-day intervals. This careful ramp up allowed the gentle conditioning of the hardware to the new, higher HV (Table 2), and also allowed close monitoring of detector behavior and the ability to detect and react to any unforeseen

Crate	HV	Crate	HV	Crate	HV
0	2110	1	2240	2	2075
3	2160	4	2043	5	2170
6	2170	7	2170	8	2060
9	2435	10	2240	11	2370
12	2220	13	2270	14	1970
15	2025	16	1995	OWL	2500
17	1945	18	2010		—

Table 2: New crate nominal HV setpoints. All crates were raised +100V, except crate 9, which was raised only +50V, and the OWL's, which were not increased.

problems immediately.

3.1 Flashers

Of all the high voltage related processes, PMT flashers are the biggest potential contributers to neutrino background. These are light emissions inside the PMT due to some type of high voltage discharge, of which the exact mechanisms are not understood at present. Flasher rates are easily tracked, and much has been learned about their behavior. Linear scaling with HV was observed, and an overall rate increase of about 30% was predicted for a 100V increase [2].

Figure 20 tells a slightly different story from our idealized prediction. During the three step HV increase, the flasher rate shot up dramatically, much more than was anticipated, to about four or five times the nominal rate. Since conditioning has been observed before, it was decided to maintain the increased HV and closely monitor the flasher rates, to see if they eventually settled down. Indeed, after four weeks, the rate has come down significantly, to a threshold uncorrected level about 60% above the pre-HV raise. There are significant fluctuations, but the overall trend is still downwards, inching towards the original prediction.

Even though the rate increase was at first worrisome, it has been a useful exercise, and the observations made might eventually help us to understand this phenomena in more detail, and how it impacts on detector stability and physics backgrounds. Certainly, the concept of HV conditioning has been proven, and also, a curious feature, after the two power failures/shutdowns, a rate dip was observed. This might be due to a charging or settling time constant of the part of the internal PMT structure that is causing the flashing.

3.2 Sharkfins

The second effect closely monitored are sharkfins. These were originally named for small HV micro-discharging effects in the PMT connector [2] that produce characteristically wide and large pulses. This rather large charge pulse (up to hundreds of millivolts) fires the energy sum trigger and shows up in the detector as a single channel with a medium to large charge, a few other associated pickup channels, and no light hits. We make here the conservative estimate that all these large ESUM pulses are sharkfins, which gives us an upper limit on their rate.

Because of the excellent stability of the ESUM trigger with the additions of the AC/F's, the sharkfin rate can be reliably monitored. From inspection of Figure 21, the rate of large charge events before the HV increase was steady at about 1.3Hz, afterwards it climbed to 4.5Hz. Then at run 5723, the ESUM threshold was turned up from 100mV to 120mV, to offset the increase in charge gain. At that point, the sharkfin rate dropped trigger gain correction is probably not exact, we should rely on the down to 2.3Hz, where it has essentially remained to date.

The sharkfin rate can also be monitored using PULSE_GT data. With these data, effects of the ESUM threshold are essentially removed from the analysis, and so sharkfins of lower amplitude can be monitored and there will be no artificial rate changes resulting from changes in threshold. Little of the data have been analyzed in this way - the PULSE_GT live time is only 7.2 ms in each hour. Before the high voltage raise the sharkfin rate was typically 600/hr; after the raise it was about 2000/hr (statistics of 6amplitudes therefore tracked the change observed with those large enough to cross the ESUM threshold. The higher rate for sharkfins of all amplitudes has not shown any evidence of subsiding with time, however, as the flasher rate has done. The rate will continue to be monitored using both methods.

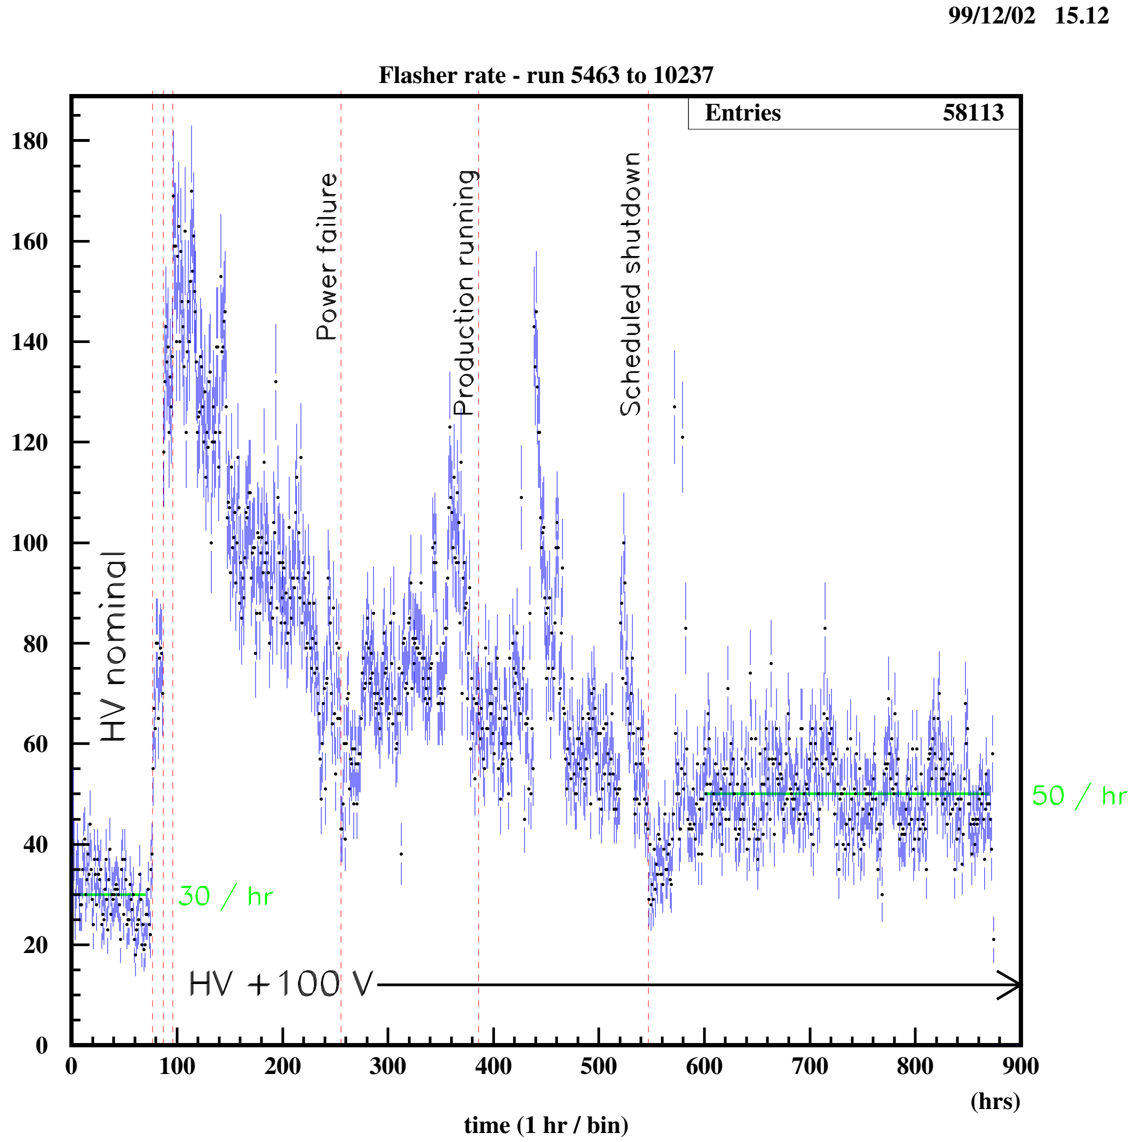

Figure 20: Rate of Flashers as a function of time.

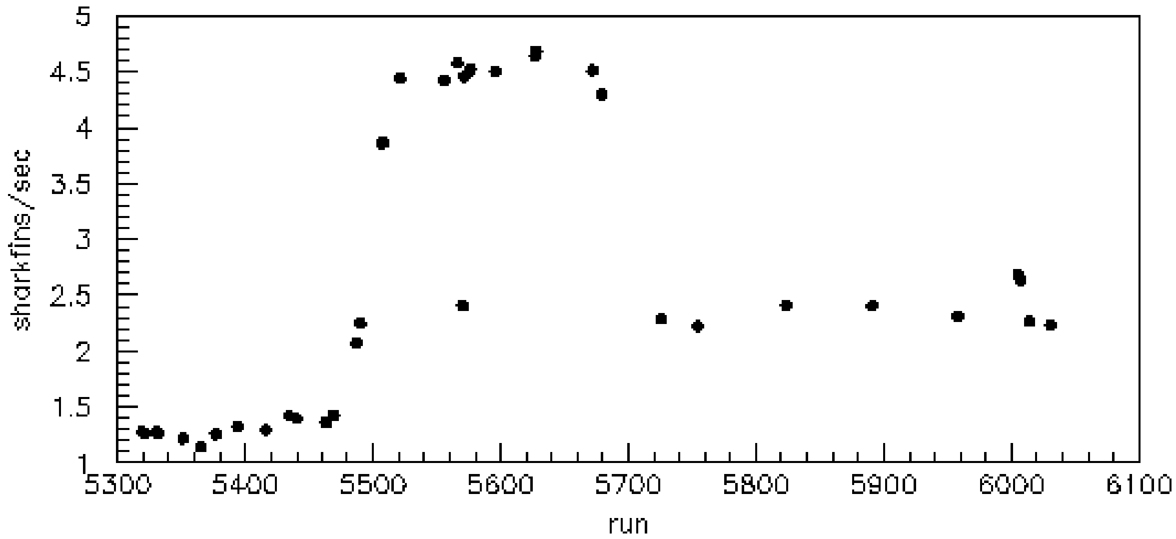

Figure 21: Rate of Sharkfins as a function of run number. The HV increase started at run 5487, and the increased ESUM threshold set at run 5723.

3.3 High Voltage Breakdown

Potentially the most serious effect to be monitored and studied is breakdown in both the wet-end (WEB), and dry-end (DEB), components of the HV chain. There are many online as well as offline detector tools for monitoring breakdown events, which by their nature, are easy to identify [2]. Given that the nominal crate to crate HV ranges from about 1800V to 2400V, and that the WEB/DEB was never found to correlate strongly with this HV range, it was anticipated that a 100V increase would have little impact on breakdown.

Before the HV increase there was virtually no DEB observed, after the increase online monitoring showed there to be some DEB activity, which after a few days, decayed away to previous low levels. This is to be expected as the dry-end connections settle into their new HV setpoint.

Close online and offline monitoring of WEB showed there to be no activity for the first four weeks after the HV increase. However, during the fifth week two separate cases of classical WEB episodes were observed in the detector. They were bad enough that the channels, 2/6/Ch21 and 1/4/ch3, completely shorted and died. Subsequent measurements showed that the PMT base in both cases were essentially shorted (low impedance), which may indicate that water leaked into the connectors or base. It is hard to say if these events were due to the HV increase, or was just part of the expected PMT/connector in water failure rate. A few cases of this has already been observed over a previous four-month period before the HV increase. A more detailed discussion of tube failure rates is given below.

3.3.1 Blown Terminators

After a small number of repeated wet (or dry) end HV breakdowns, the PMT base back-terminator resistor can blow open [2], not, by itself, a large loss. There are two methods to look for the blown terminator. The first, more direct and reliable, is to use the on-board individual channel charge injection circuit to look for a reflection pulse from the termination mismatch [2]. A scan was made before the HV raise and compared directly with a scan made three weeks afterwards, which showed there to be no new blown terminators. Note, this contradicts the statement above that WEB was observed in two channels, since it should show evidence of a blown terminator. The terminator test was done after a scheduled INCO power outage (Nov 19, 1999), which was approximately a week before the two channels blew. The next time a reflection test is made, these two channels should show blown terminators.

A second, less robust, but more up to the minute on-line method is to look for channel gain changes in the neutrino data [2]. Since the HV increase, no new clearly identified blown terminators have been found with this technique. This includes the two recent WEB channels which were missed because they shorted out and had their HV turned off, and hence, do not produce data. Recently, three marginal candidates have been found (00/13/ch27, 04/09/ch17, 10/11/ch24). Since no clearly identified WEB has been observed from these channels, they are most probably not real blown terminators, but some other detector artifact. However, they will be closely monitored and checked at the next available time with the more robust reflection test.

3.4 Pulse GT

The detector pulsed global trigger (PULSE_GT) samples the detector at a continuous 5Hz, which makes it an excellent non biased monitor of detector stability and background rates. Before the HV increase, the mean PULSE_GT nhit was measured to be 1.89 ± 0.05. After

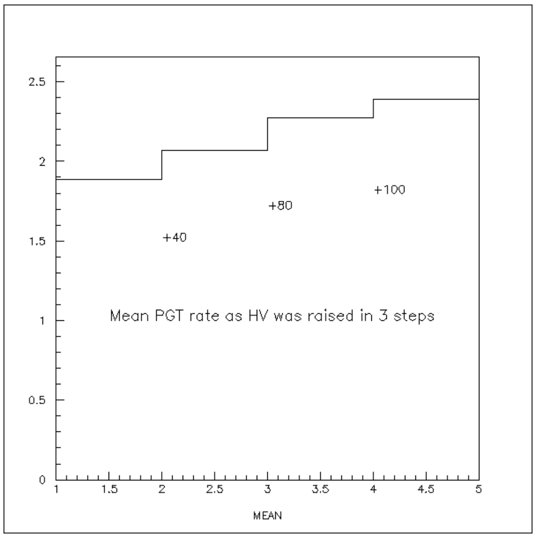

Figure 22: Mean PULSE_GT nhit as a function of HV steps.

the increase, it was found to be 2.35 ± 0.05 (see Figure 22), and shown to be very stable. This increase in mean nhit can be attributed to the expected higher tube noise rates. Given that the mean nhit in the $400nsec$ trigger time window is low, the even smaller delta increase is insignificant and easily tolerated. Figure 23 shows a typical example of the PULSE_GT mean nhit versus time for a standard neutrino run at the HV increase. Notice the excellent stability.

3.5 Occupancy Changes

Another means of checking the detector health is to monitor the tube death rate. Through continual laser occupancy runs, PMT rates and base current measurements, the overall PMT health and death rate have been closely monitored. As reported before, our typical loss rate through the summer has been low at about 0.4% per year, or about 1 tube per week.

Since the HV increase, two tubes have shorted and died, showing an open PMT base (clear evidence of connector failure due to water leakage). Three new tubes were found to have low data occupancy (10% or less) and are probably a result of damage to the PMT base or channel HV distribution resistors. This gives us five new dead channels in a period

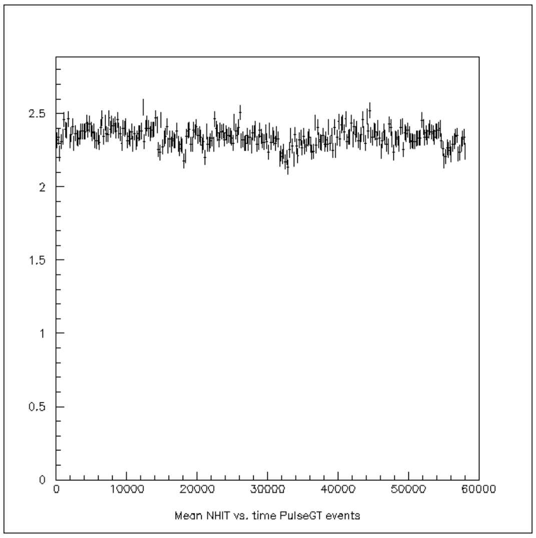

Figure 23: PULSE_GT mean nhit rate over a period of 16 hours for neutrino run 10142.

of 6 weeks since the HV increase, which is consistent with the previously measured tube death rate. In addition, no overall degradation in tube performance (rates, basecurrents, occupancies) has been observed. Figure 24 plots the number of dead tubes versus time.

3.6 Future HV Monitoring Plans

Overall, no significant unforeseen ill effects have been observed with the increase in HV. After the initial six weeks of close monitoring and analysis, it has been decided that the HV increase is a success! The +100V over previous nominal, are now defined as the new detector nominal HV operating setpoints (Table 2). The numerous tools developed for monitoring the HV related phenomena have been refined and now incorporated into the online/offline data stream as automated processes. The output from these analyses will continue to be monitored over the long term, and any sudden changes or deviations, will be quickly noted and action taken if required.

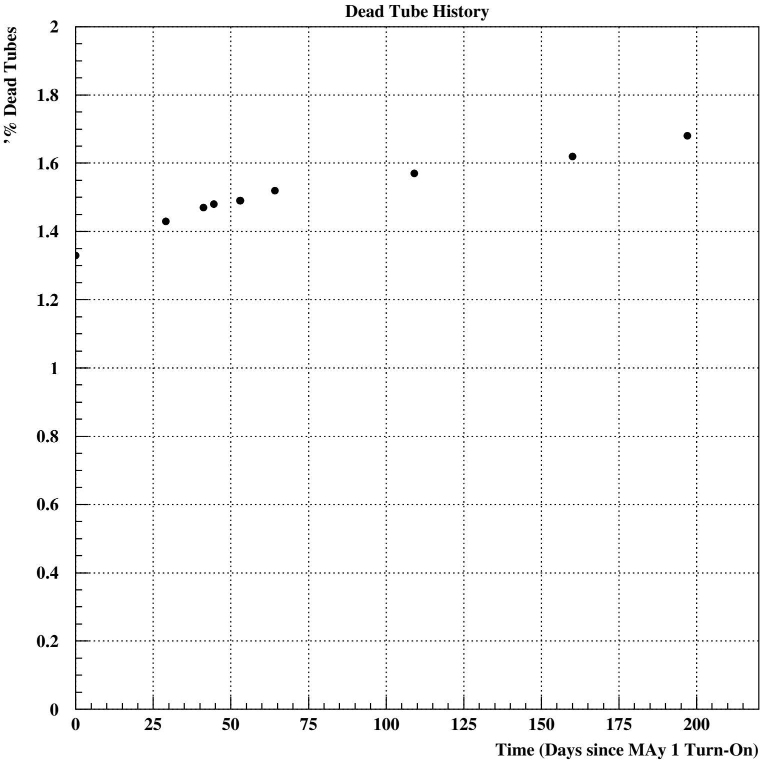

Figure 24: Tube deaths as a function of time since full detector turn-on. HV increase was around day 155.

References

[1] Frati et al. *Investigation of Relative Channel Efficiencies*, SNO-STR-99-021.

[2] Frati et al. *Study of High Voltage-Dependent Phenomena in the SNO Detector*, SNO-STR-038.