SNO-STR-99-038

Study of High Voltage-Dependent Phenomena in the SNO Detector

James Cameron, Bill Frati, Rich Helmer, Josh Klein, Neil Mccauley, Gordon Mcgregor, Mark Neubauer, Richard Van de Water, Peter Wittich
Oxford University, Triumf, University of Pennsylvania

December 21, 1999

Abstract

We summarize here studies done on some of the non-physics phenomena in SNO at the current high voltage settings, focussing in particular on the relationship between the high voltage settings and severity of the problem. This study also serves as a demonstration that the tools we have are sufficient for diagnosing any changes that may occur in these phenomena either over time or with a change in the high voltage setting. Specifically, we have looked at the rates of flashers, 'sharkfins', wet-end breakdown, and tube failure, and the correlation of these with high voltage, with (new and old) blown back terminations, and with each other. We find that both flashers and sharkfins may increase by $\sim$25-30% for a global 100 V increase, but see no evidence for an increase in tube loss from breakdown. Overall, the tube death rate is now roughly 0.4%/year, about half of which may be attributed to classical wet-end breakdown problems.

1 Introduction

It has been pointed out recently that a small increase (80-100 V) in the high voltage settings on the SNO phototubes could raise the overall photon detection efficiency by up to 10%—a significant change in that it increases both the overall energy and reconstruction resolution, allowing better separation of signal and background. A 10% increase could put SNO's energy scale into the regime (9-10 tubes/MeV) in which much of the original Monte Carlo studies of SNO physics were done.

A change in the high voltage settings does not come without a price, of course. There are many (now) well-known high voltage related phenomena which are both unpleasant (to put it mildly) and which may worsen with the voltage increase. In particular, the rate of flashers, so-called sharkfin pulses, and wet-end breakdown events may all increase if the high voltage is raised. While the flashers appear fairly easy to eliminate, a significant change (factor of two or three) would probably be beyond our comfort level. The sharkfins (or what we are calling sharkfins) are not known to cause any harmful effects in the current detector, though before regassing they were strongly correlated with wet-end breakdown. The wet-end breakdown itself, of course, is something which is clearly dangerous, and we are unlikely to tolerate any increase which may to contribute significantly to long-term tube loss.

We describe here a study of all of these phenomena, paying particular attention to their dependence on high voltage. We feel it is important to point out before beginning that because of the wide spread of crate voltage settings (500 V), in only one case does an increase of even 100 V explore a new voltage regime: Crate 9, which is at the highest voltage setting (2385 V) is more than 100 V higher than the next highest crate. Therefore, studies of the differences between crates and channels at different voltages give us some idea of what we can expect if each one is raised.

In addition to giving us a baseline for comparisons, this document also serves to demonstrate that we have sufficient tools to detect and diagnose any changes that we see. In undertaking any global high voltage raise (even the $\sim$4% raise we are considering) we must be vigilant; a significant increase in high voltage related problems must be noticed quickly and a decision made whether the problem is severe enough for us to return to the lower voltage settings.

2 Flashers

One of the potential consequences of changing the detector high voltage is a change in the rate of flasher PMT's. Although it is not known what the exact mechanism is that causes the flashers to flash, the assumption has been that the light emission is inside the PMT itself, due to some type of high voltage discharge. As is well known, the characteristics of flashers are a large charge in a single tube, accompanied by pickup in surrounding channels, and light (typically in the 40-200 NHIT region) on the opposite side of the detector. Since the water has gone in, the pattern of light has taken on a distinctive, asymmetric shape.

There are several questions regarding flashers which we will try to address here: Do flashers 'condition' away? What is the correlation between flasher frequency and high voltage? How many tubes exhibit this behavior? Understanding these before we change the detector high voltage will allow us to decide whether any change in rate is indicative of a new problem or consistent with our predictions.

Tracking flashers with good efficiency is relatively easy, though some work has to be done to prevent contamination from other types of events. The algorithm used to identify flashers in this memo looks for a combination of a cluster of pickup hits with significant light on the opposite side of the detector. The clusters that are searched for have six or more consecutive electronics channels hit, and the total amount of light in the rest of the detector is required to be 20 or more tube hits. In addition, no more than six 'small clusters' (three or more consecutive channels hit) are allowed, unless the small clusters are part of a big (six-hit or more) cluster. This last requirement eliminates non-flasher events like 'manipulight' and muon retriggers. While this type of ID will tag only a fraction of the total flasher events in a data set, it is good enough as a relative measure for us to be able to draw conclusions about the time and high voltage dependence of the flashers.

Figure 1 shows the rate of flashers for roughly 700 hours of running. As shown in the Figure, the average rate is $\sim$ 28 flashers/hour. The spikes in the plot are bursts of flashers almost certainly due to seismic activity in the mine. While there are certainly large changes in the average rate—it drops in the first 120 hours or so by nearly a factor of two and then recovers with roughly the same time constant—there nevertheless does not appear to be any long-term conditioning on this kind of timescale. An important point to make is that the

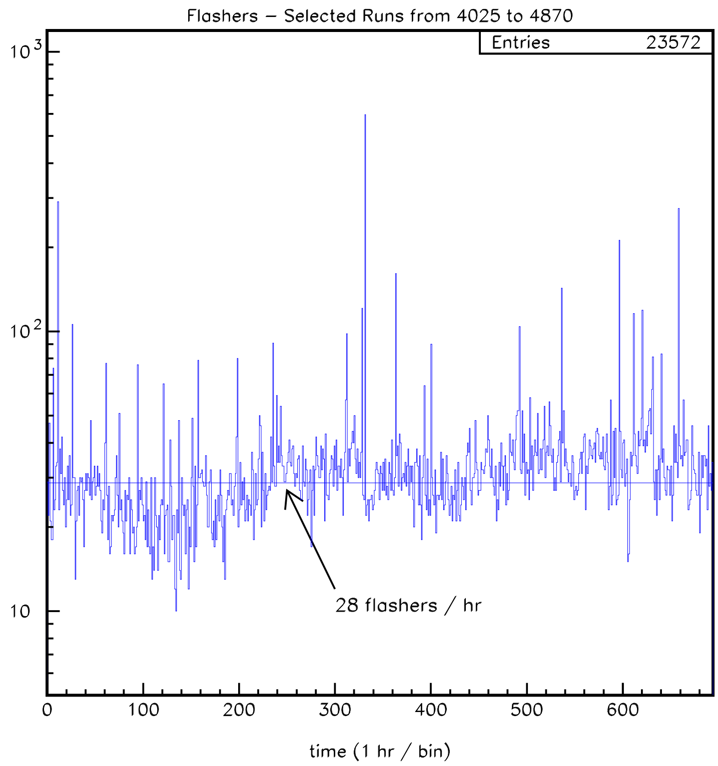

Figure 1: Rate of Flashers vs. Time for $\sim$ 1 month of running.

time variation of the flasher rate (both the slow variations and the bursts) will make it much easier to determine their contamination in our final data samples, assuming that the ones that leak through our cuts vary in the same way in time as the ones we identify easily.

The total number of flashers and their distribution in the detector is also interesting—if we were dominated by just a few tubes we would very likely turn those off. Figure 2 shows the channel distribution of flashers, with some of the 'worst offenders' labelled. The total of 4385, as shown in the Figure, indicates that in this time period nearly half the tubes in the detector exhibited flasher behavior. Although worst offenders can be identified, it is not obvious that they remain at the top of the list over time due to the "bursty" nature of the flashers.

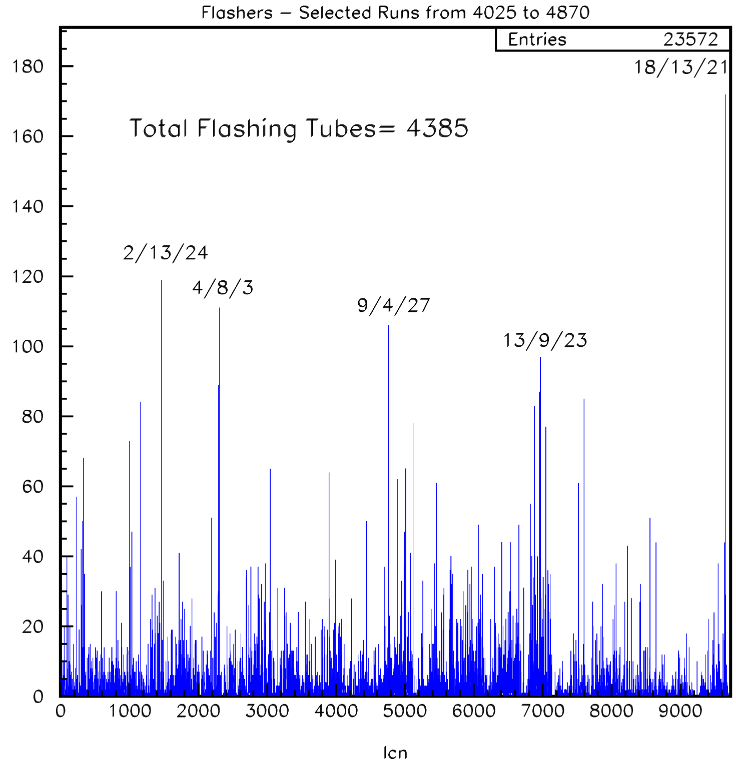

Figure 2: Channel Distribution of Flasher PMT's.

More interesting in relation to the question of a high voltage change is the correlation of flashers with their high voltage setting. Many have noted this correlation based upon the different crate high voltages. In Figure 3 we show the correlation with the *channel* high voltage setting, which shows the effect even more clearly. The slope of the line yields a factor of two increase in rate with roughly a 300 V increase in voltage. A change in voltage of 100 V produces a $\sim$ 25-30% increase in the rate. The turnover near 2200 V is interesting, but more work would need to be done to decide if it is a real effect.

In addition to the flasher rate as a function of high voltage, it is interesting to look at the flasher intensity. Figure 4 shows the average NHIT of flasher events as a function of channel high voltage. Interestingly enough, there is a clear correlation which starts after a flat region and becomes steeper at higher voltage. This may be an insight into the flasher

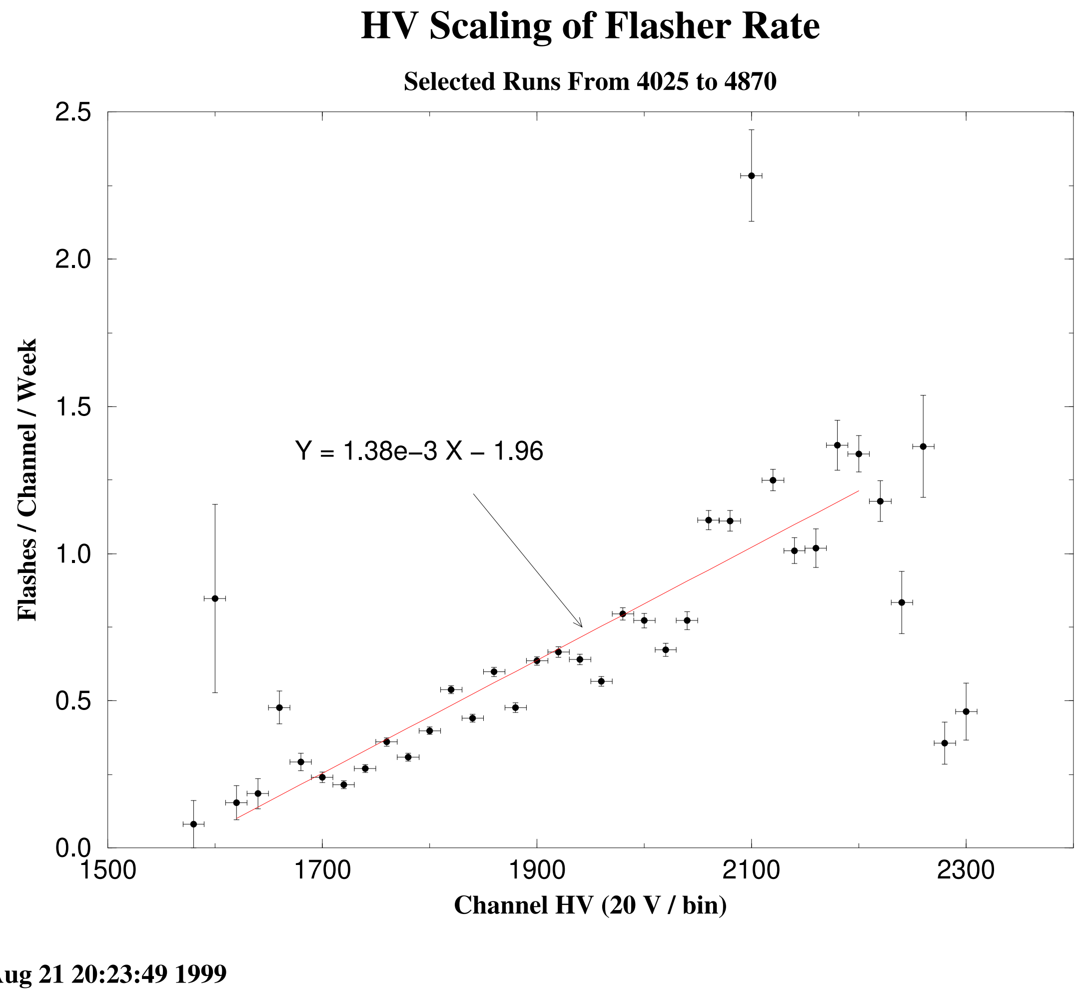

Figure 3: Correlation of Flasher Rate with Channel High Voltage

mechanism itself, but would have to be studied in more detail.

3 Sharkfins

So-called 'sharkfin' pulses, present in large numbers before the re-gassing of the light water, were shown to precede wet-end breakdown events. In fact, it is possible that all breakdown events are precipitated by sharkfin pulses. The sharkfins themselves, however, occured most often on their own, and in some channels occasionally at high rate ($\sim$ kHz).

The signature of a sharkfin pulse (as seen on the ESUMHI analog sum) was a risetime (leading edge) longer than a typical PMT risetime ($\sim$ 10-20 ns) followed by a falltime much longer—up to hundreds of nanoseconds. In addition, the pulses were larger than a typical PMT pulse—over 100 mV on the ESUMHI trace, sometimes saturating the channel's ESUMHI output (this corresponds to a pulse height in excess of 200 photoelectrons). Of course, there may have been many smaller sharkins that we were unable to see due to our own trigger biases. The sharkfins may be some kind of corona discharge, or micro-discharge on the connector surfaces, or perhaps something else entirely.

Since re-gassing, the characteristic long rise and longer fall seems to have disappeared, but large pulses of 100-200 mV persist, with no light evident in the detector. Whether these pulses are sharkfins or not—that is, some type of high-voltage related discharge associated with the wet-end connectors and not the PMT itself—is not clear. If they are sharkfins, we must assume that the possibility exists that they can cause damage to the connector, and therefore their rate and their dependence on high voltage is worth investigating before we consider a change.

With no robust criterion now with which to identify sharkfins, we choose the most conservative approach and define *any large phototube pulse not associated with light in the detector* to be a sharkfin. This definition is implemented by simply looking for clusters of hits (five or more on a paddlecard) with fewer than twelve hits in the detector as a whole. No charge requirement is applied, in an attempt to minimize bias. However, the cluster requirement itself is a charge measurement, as pickup to four channels or more does not occur until the charge in a single channel becomes larger than many tens of pe. Although this definition is certainly not an optimal way of identifying sharkfins, it is good enough as a relative measure

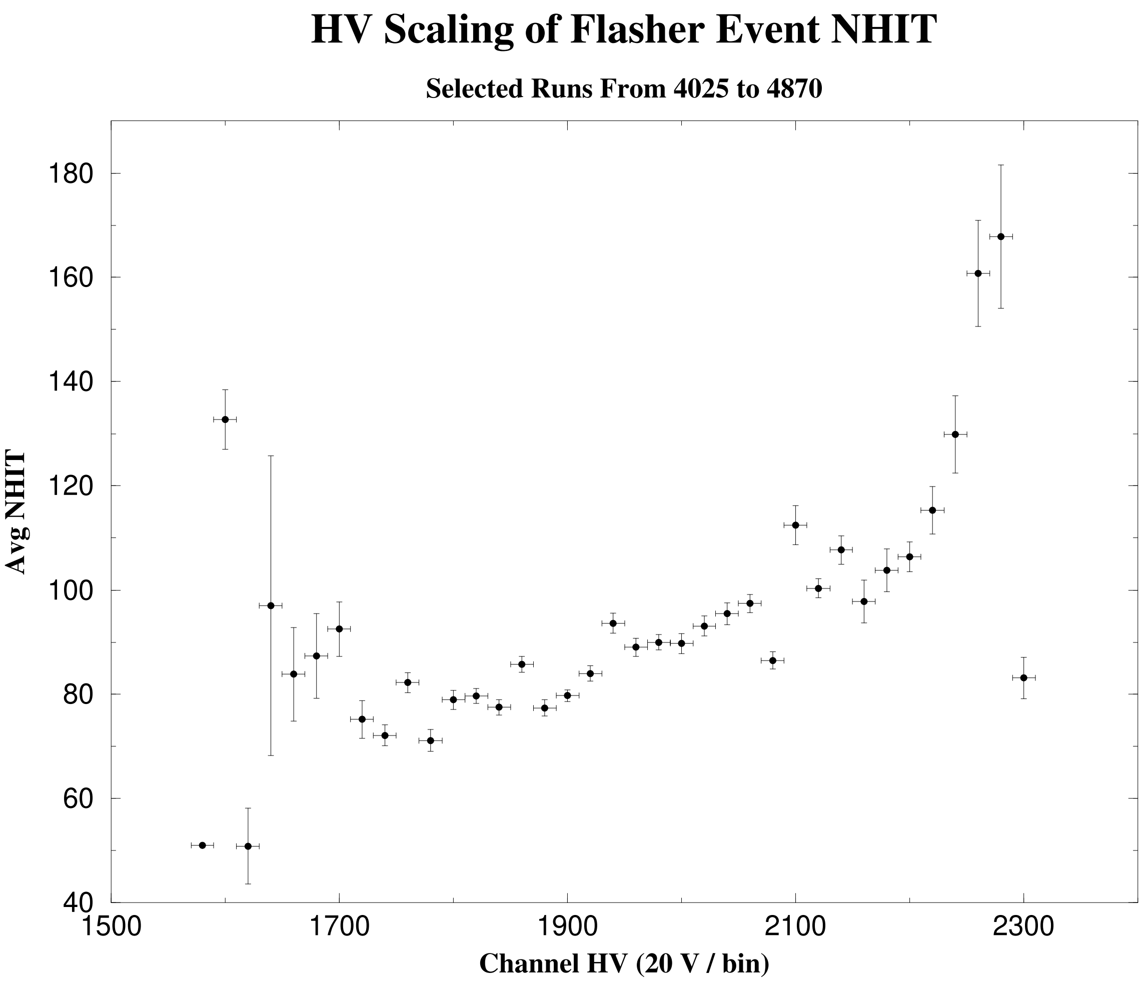

Figure 4: Correlation with Flasher NHIT with High Voltage.

of their rates and their dependence (if any) on high voltage.

Two types of triggers were considered in this study: PULSE_GT events and ESUMHI triggers. The PULSE_GT events have the great advantage that they are unbiased—we can sample sharkfin pulses down to a size as small as the cluster requirement allows. In this way we avoid seeing time-dependent changes which are a result of differences in triggering. The drawback to the PULSE_GT analysis is that the probability of an accidental coincidence between a sharkfin and a PULSE_GT trigger is very small until the sharkfin rate becomes extremely high.

For most of the data taken so far, The ESUMHI triggers have the distinct disadvantage that the threshold has not been stable. Temperature changes, sources of pickup and sag, the whims of the detector operators— all combined to produce a situation in which the rate of sharkfins measured using the ESUMHI triggers tell us more about the threshold than about the sharkfins. The installation of the ESUM-AC/F's stablized the threshold considerably, allowing the ESUMHI threshold to be lowered and making it much more useful for sharkfin studies.

Figure 5 shows 1000 hours of sharkfins using the PULSE_GT data. The average coincidence rate shown in the Figure of 4 sharkfins/hour corresponds to a detector rate for sharkfins of $\sim$ 250 Hz. As can be seen from the Figure, there is no noticeable rate change over the course of the 1000 hours, though the statistics are obviously poor.

Figure 6 shows the distribution of channel numbers producing the sharkfins. There are definitely a few that appear to have higher rates than others, though overall the statistics are not high even for the 1000 hours.

Looking at the same 1000 hours of data with the ESUMHI trigger as shown in Figure 7, shows the kind of structure we would anticipate due to the varying trigger threshold over time, and therefore it is hard to say very much about the changes in rate. Figure 8 shows a similar plot for 48 hours of data after the ESUM AC/F's were installed. The rates here are much more stable, and also higher due to the lowering of the threshold from 200-250 mV to 100 mV. In addition, we have enough statistics with these data to look at individual channel rates and correlations with blown terminations and channel high voltage.

Figure 9 shows the distribution of sharkfins as a function of channel number using the ESUMHI triggered data (after the AC/F's were installed). There are clearly some channels

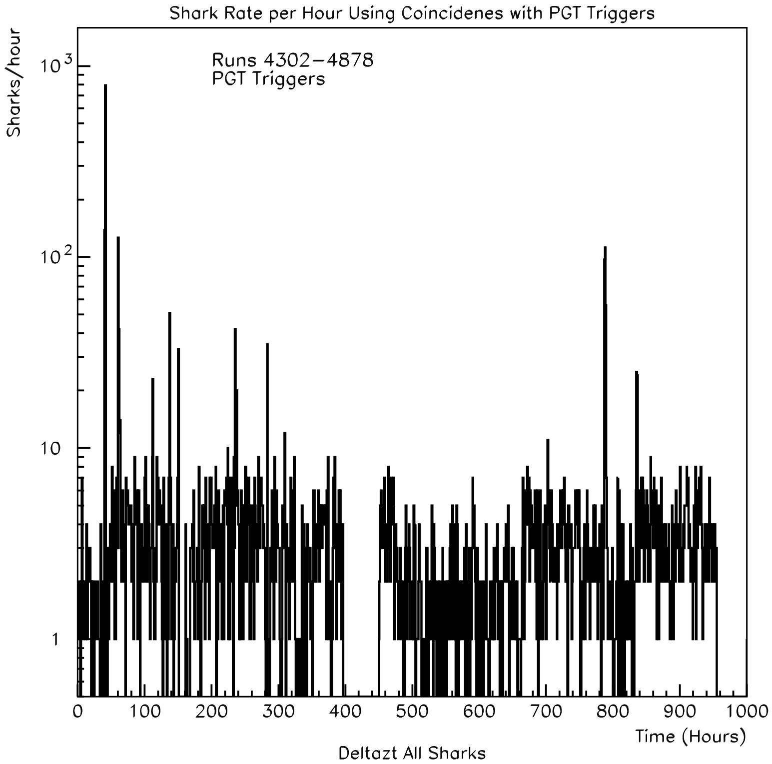

Figure 5: One Thousand Hours of Sharkfins, identified with PULSE_GT triggers (the detector was off during the period around 400 hours).

much worse than others in this sample. Plotting the number of sharkfins vs. channel high voltage as in Figure 10, we can see that there is, in fact, a strong correlation. The only way to produce such a correlation with high voltage for normal PMT pulses is if the gains of the tubes are not well-matched. In this case, all we are seeing is the ESUMHI trigger bias. (It is interesting to note that Figure 10 and Figure 3 have very similar slopes— perhaps flashers are the tail of sharkfins?). Figure 11 shows the distribution of low gain charge for sharkfin events identified using the ESUMHI trigger, and the steep fall of the spectrum past the bite from the ESUMHI threshold setting indicates that only a small change in gain can result in a big change in apparent rate. If the gain settings are correlated with the high voltage settings (the design intention was that the gains would be matched), then we could see such a correlation. (Recent work looking at occupancies shows that such a strong gain correlation

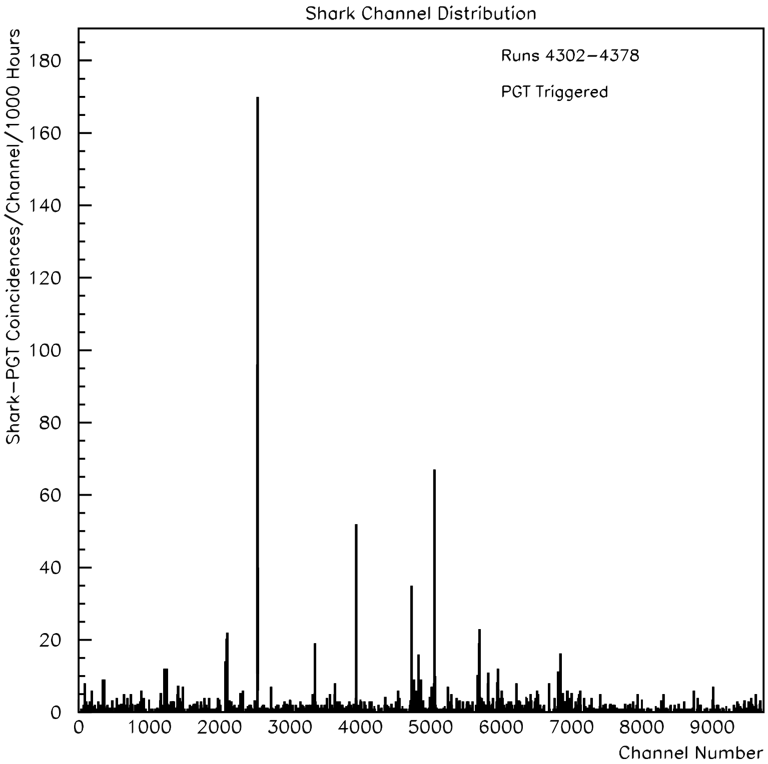

Figure 6: Channel Distribution of Sharkfins, using PULSE_GT triggers.

with high voltage is not likely).

Some of the bias is eliminated by looking back at the PULSE_GT triggered data. Figure 12 shows the same correlation plot as Figure 10 with the 1000 hours of PULSE_GT data. Here the correlation is not very obvious, potentially an indication that the correlation seen with the ESUMHI data is a pulse height (either sharkfin or normal PMT) correlation, not a rate correlation.

We can also look at the correlation between sharkfins and channels with blown terminators. We have the same problem here as we had earlier, however—if the sharkfins are just the tail end of the PMT pulse height spectrum, then channels with blown terminators will be much more likely to cross the ESUMHI threshold and we will be biased to see these channels. The same is true of sharkfins, if there is a spectrum of sharkfin pulseheights—doubling the gain with a blown termination will result in an apparent rate increase. Figure 13 shows the

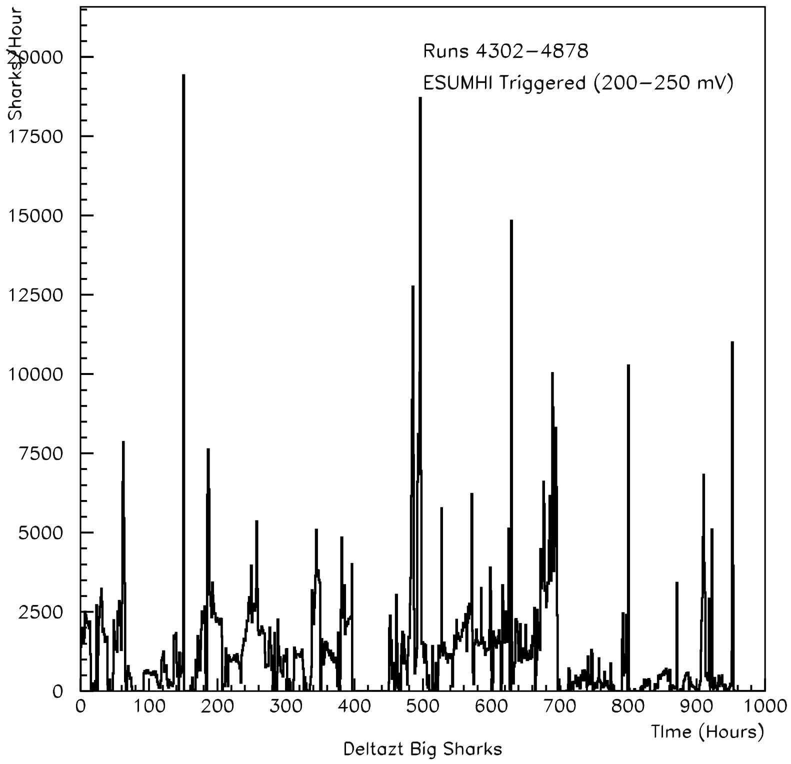

Figure 7: One Thousand Hours of Sharkfins, using ESUMHI triggers, before installation of ESUM-AC/F's.

total count of sharkfins for Runs 4988-4991 for channels with blown and non-blown terminators. It is clear that those with blown terminators are tagged as showing sharkfins much more that those without.

We can use the distribution of sharkfin rates to get an idea of whether sharkfins eventually cause tube problems. The 10 channels with the highest sharkfin rate (measured using the ESUM triggered data) do not show occupancies that are signficantly different than average tubes. These 10 tubes have sharkfin rates that are two or three orders of magnitude larger than the vast majority of the phototubes, and therefore the fact that these tubes do not become worse with time indicates that there is no reason to believe the others with much smaller rates will either.

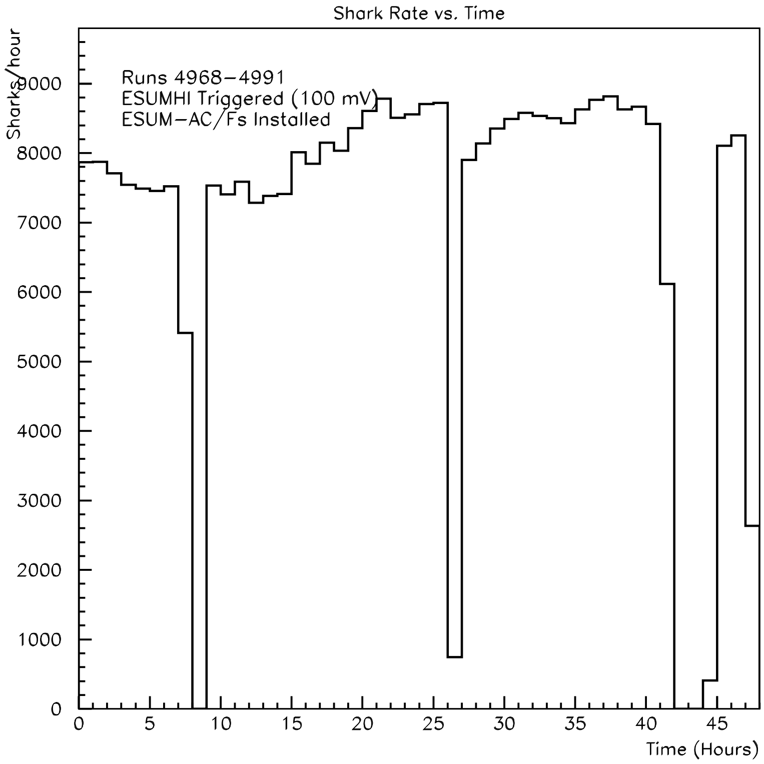

Figure 8: Forty-eight hours of sharkfins, using ESUMHI triggers after installation of ESUM-AC/F's.

4 High Voltage Breakdown

The most serious consequence from raising the high voltage could be an increase in wet-end breakdown problems. There are several ways of tracking the degree of the problem that we have had over the past few months of running. Here, we have tried to be comprehensive by using several algorithms: looking for wet-end breakdown events directly, looking at the number of blown back-terminations as a function of time, looking for changes in tube gains (and correlating these with the blown back terminations), and looking for tube deaths as a function of time using laser occupancy runs.

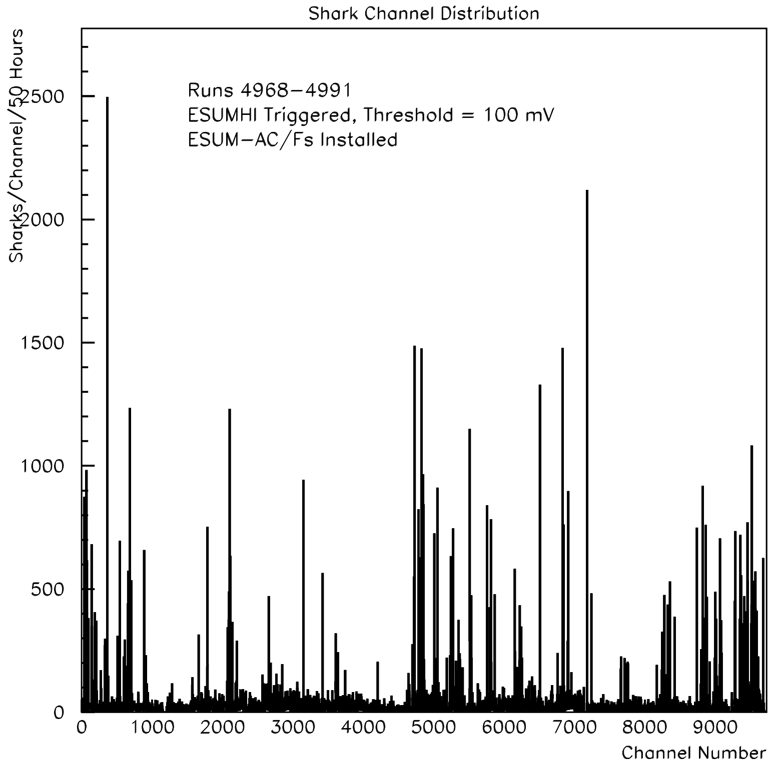

Figure 9: Channel Distribution of Sharkfins, using ESUMHI data. Note the vertical scale is roughly an order of magnitude larger than for the same plot using PULSE_GT data.

4.1 Wet-End Breakdown Rate

The rate of identifiable connector wet-end breakdown has dropped to such a low level that it has become hard to track. Other high NHIT events—muons, so-called 'bubblers' and 'neck events' can look (to a simple algorithm) like wet-end breakdown. Nevertheless, we have used an algorithm that was demonstrated to track breakdown well before the re-gassing with nitrogen and applied it to all of the relevenat data taken between July 10 to August 18, 1999.

The algorithm was intended to be a way of identifying bad channels (rather than simply bad paddlecards), and is therefore somewhat particular about what it calls wet-end breakdown. The general outline of the algorithm is that it looks for a card with at least one negatively railed QLX and ten channels with unusual QHS. If, as was often the case before

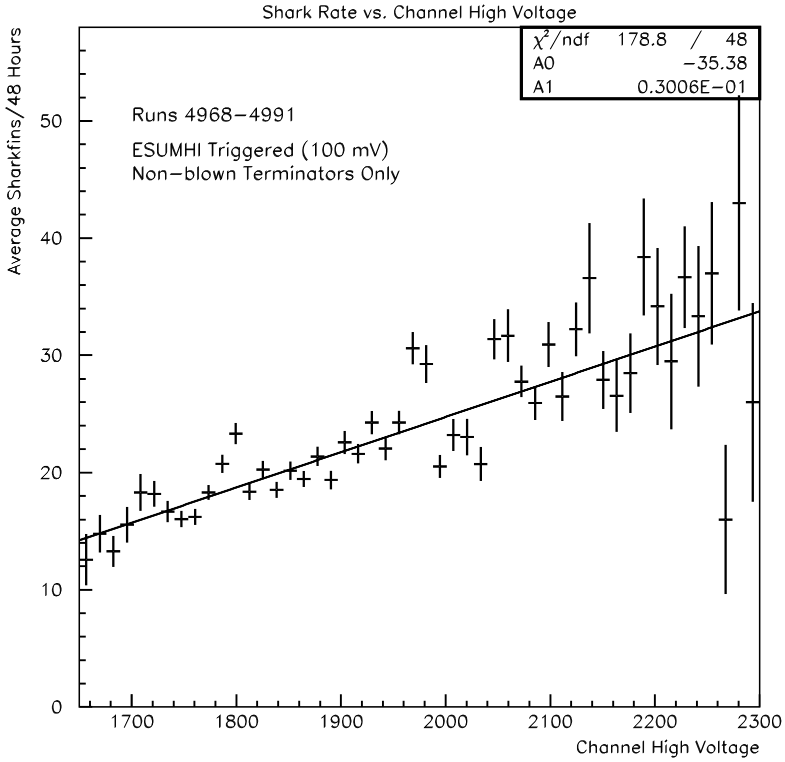

Figure 10: Sharkfins vs. Channel High Voltage.

re-gassing, there is more than one channel with negatively railed QLX then the paddle card with the highest QLX 'score' (essentially the sum) is considered to be the bursting paddle card. The tube considered to be the one causing the breakdown is the one which is most often missing during the subsequent bursts.

Using this code, we find that a total of five different paddlecards were identified as showing breakdown over the month of data analyzed. In all cases, the total number of breakdown events from these paddlecards was 1 or 2, and therefore no channels could be identified. On none of the paddlecards identified are there any channels which became non-operational during the running time, nor are there any which subsequently were found to have blown back-terminations. In addition, the two crates which were identified as showing problems (Crates 8 and 10) are crates with relatively low voltages (1970 V and 2140 V, respectively).

Although it is possible that some level of wet-end breakdown still remains, it may now

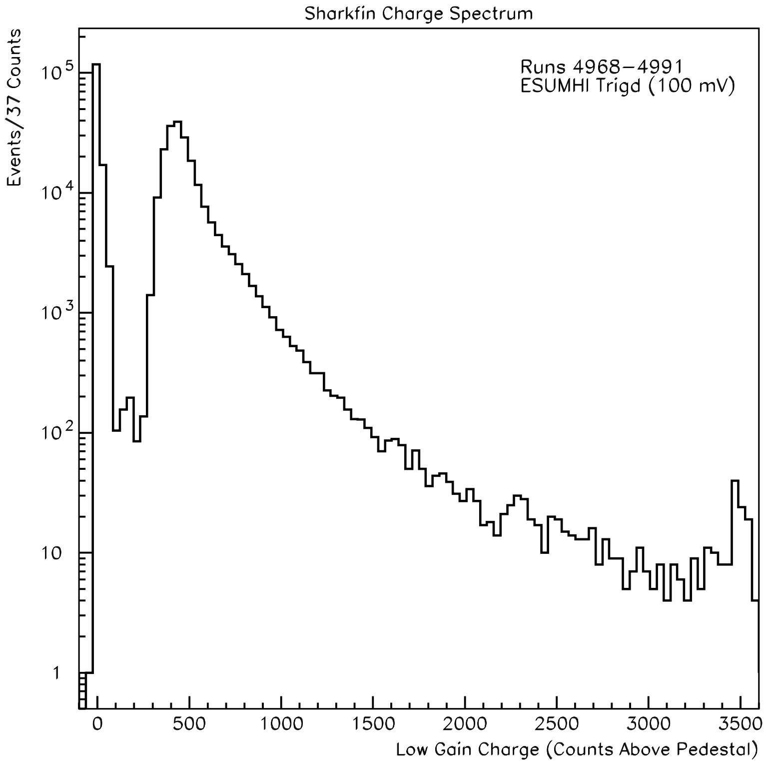

Figure 11: Low gain charge spectrum of sharkfins.

have a character different enough that our algorithms are no longer appropriate. Building new algorithms will depend upon us being able to identify the problems in an independent way. It also likely that, with or without a change in high voltage, some connectors will fail. We must keep this in mind if and when we see a change after a high voltage raise.

4.2 Blown Terminations

Perhaps the simplest way to track tubes suffering from breakdown is to look for the destruction of the base back-termination. The wet-end breakdown has enough energy to blow the 75 Ohm back-terminations in a fairly short amount of time—perhaps as few as 30 or 40 events. The blown back-termination is then a 'permanent record' of the breakdown, insensitive to analysis biases and inefficiencies (though sensitive to its own measurement errors).

The most direct way of determining whether a back-termination has been blown is by

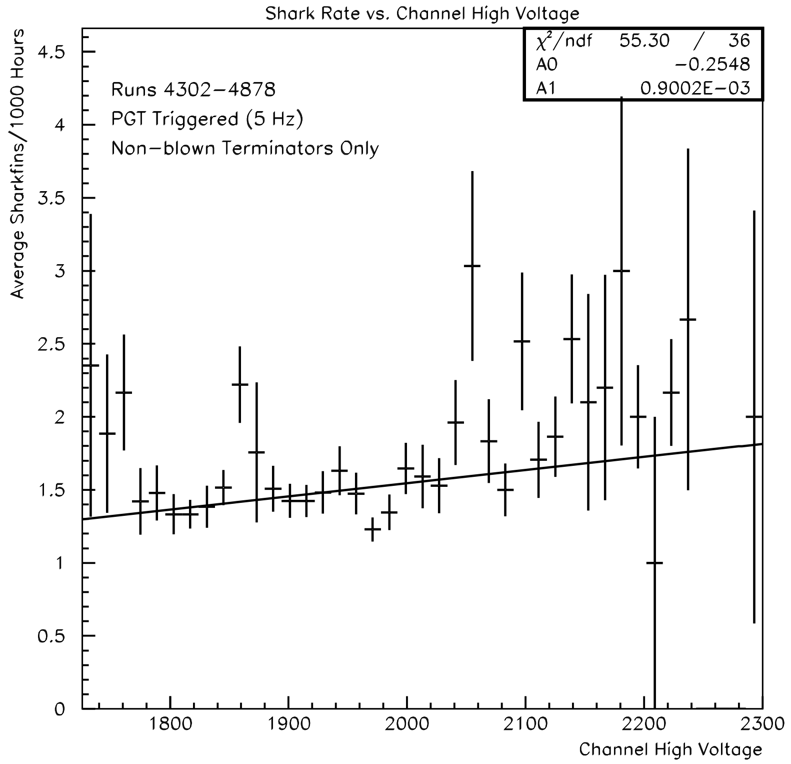

Figure 12: Sharkfins vs. Channel High Voltage for PULSE_GT Data.

injecting a pulse down the PMT cable and looking for the reflection. This is done most often 'by hand'—visually inspecting a scope trace for sign of the reflected pulse.

The total number of blown terminations currently known is 450, the vast majority of which occured before re-gassing. Figure 14 shows the correlation of blown terminators with high channel voltage, using only the crates which were on for substantial times before re-gassing. The plot has been made simply by histogramming the number of blown terminators in bins of 20 V, and dividing by the total number of channels in each bin. As can be seen from the Figure, despite the very wide range of voltages—over 500 V—there is no indication of a correlation.

Since the turn-on in May, two complete reflection tests were done (one at the end of July, and one in early September). In the July run, 30 previously unreported blown terminators were found, with a relatively uniform distribution over the crates in the detector. Of

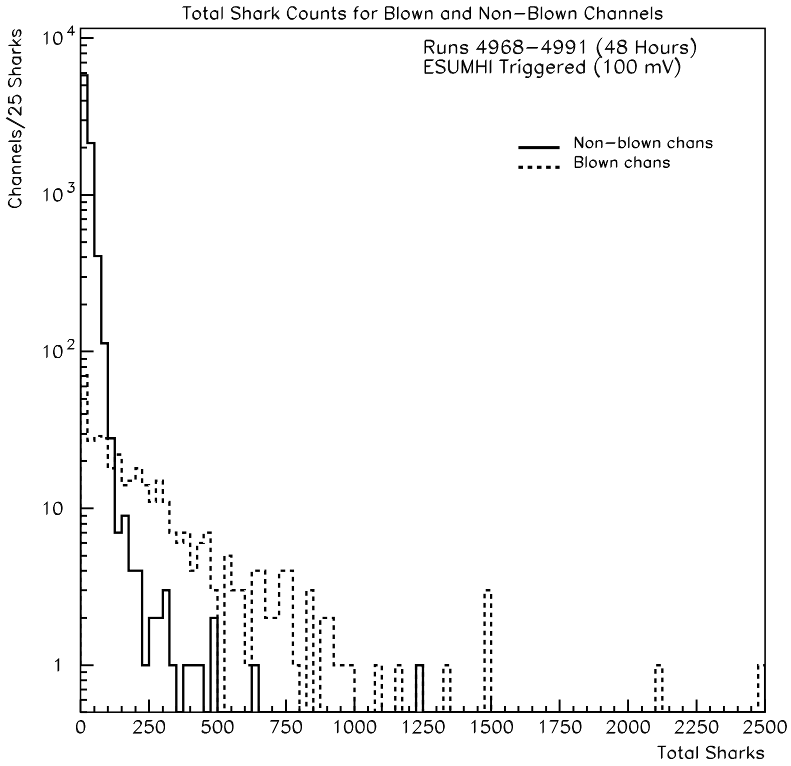

Figure 13: Comparison of sharkfin rate distributions for channels with blown and non-blown back terminations.

these 30, 7 had 'small' reflections which may not be indications of damage from breakdown. Subsequent analysis looking for the expected gain doubling of the blown terminations using laser data (see next Section) showed that, of the 23 remaining channels, only four showed clear gain changes since May, indicating that the rest had either been blown before and mis-recorded or were not, in fact, blown at all. Of these four, one was known to have participated in a breakdown event in early May, and was shown to no longer have any signal at all (it was completely dead). The remaining three had also all blown before June, one of which was also now dead.

In the second reflection test done in September, three previously unreported blown back-terminations were found. In this test, a different method was used, looking at the full ESUM (rather than the output of the crate sums themselves). This makes the reflection

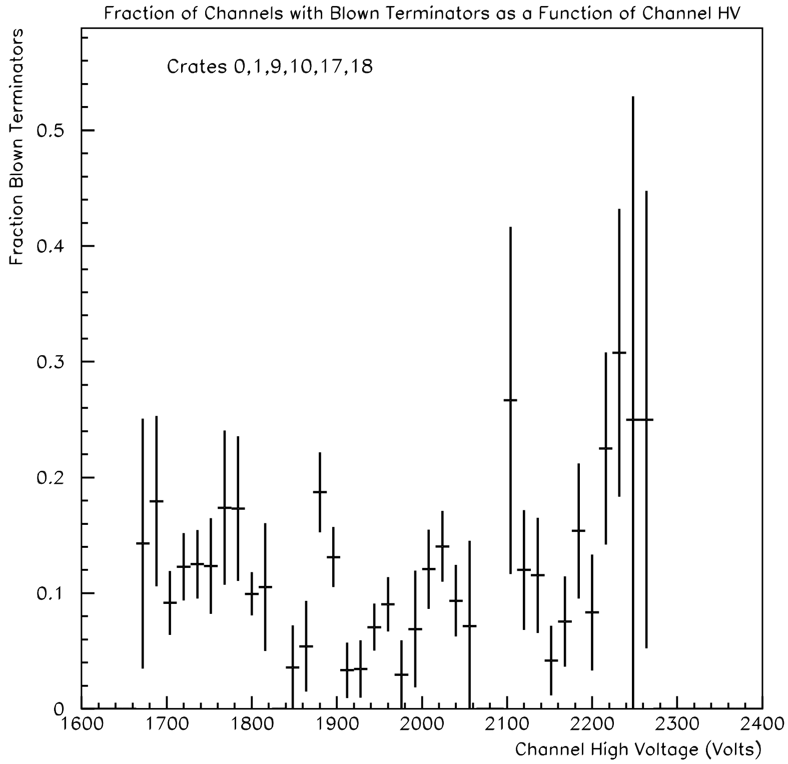

Figure 14: Number of blown 75 Ohm back-terminations as a function of channel high voltage, normalized to the number of channels at each high voltage.

test somewhat faster, but has the disadvantage of being less sensitive. Based on the number of known terminators that were missed, the sensitivity loss is roughly 3%. Of the three new blown terminators, later analysis showed one which was clearly blown long before and was somehow missed in previous reflection tests. The two others also did not show any gain change (see next Section), but their gains did not look obviously doubled as we would expect from the loss of the back-termination.

In addition to the 'manual' reflection tests, an automated tool has been developed which uses the difference between the short and long integrate charges to look for reflections. That tool was not used in this study as it has its own inefficiencies, but will be considered for more long term use during production running.

4.3 Gain Changes

In addition to looking for the blown terminations with the reflection tests, we can look for the resulting gain doubling in the data instead. The advantage here is that the test can be performed offline, without any interruption in detector running. This method can be used with either laser runs or standard NHIT runs, with the former case having somewhat better sensitivity but the latter being more frequently available.

For the NHIT analysis, the runs chosen are at least 24 hours long, and only events with NHIT100 triggered events are considered. The size of the events is required to be between 20 and 200 hits. In addition, some low level data cleaning cuts are applied, eliminating orphans, bursts, clusters, and events with railed charges. Individual PMT's are rejected if their charge pedestals or time slopes are outside broadly defined nominal limits, if the TAC is within the curl region, or if it fails the walk correction. With these data, histograms of the high gain, short integrate charge spectra are made for each qualifying tube, in units of counts above pedestal.

A KS test is then applied to the histograms from two runs of interest, with the first 20 counts above pedestal zeroed to restrict the comparison to the high charge region. In addition, occupancy tables from the two runs are generated for each channel.

Channels are then flagged as bad if they have either a change in occupancy greater than 25%, or bad KS probabilities ($< 10^{-5}$), or KS probabilities that are zero. For each of the bad channels, overlays of the charge histograms are made from each run, and all are hand scanned.

Figure 15 compares the distributions for a channel with its gain doubled due to a blown back-termination to one without. The change over the interval between two of the laser runs is very clear in this case, and is typical of those which the reflection tests show have blown terminators. Figure 16 shows four distributions compiled from two NHIT runs which were flagged as being bad. The distribution in the upper left hand corner of the Figure, for Crate 16, Card 9, Channel 14, is the only one of these which the hand scanning tagged as being a candidate for a blown termination. The others are simply shifts due to changes in pedestal for those channels. A subsequent reflection test showed that Crate 16, Card 9, Channel 14 did, in fact, have a doubled gain but that there was no accompanying reflection, indicating that the gain change was caused by some other mechanism not likely to have been

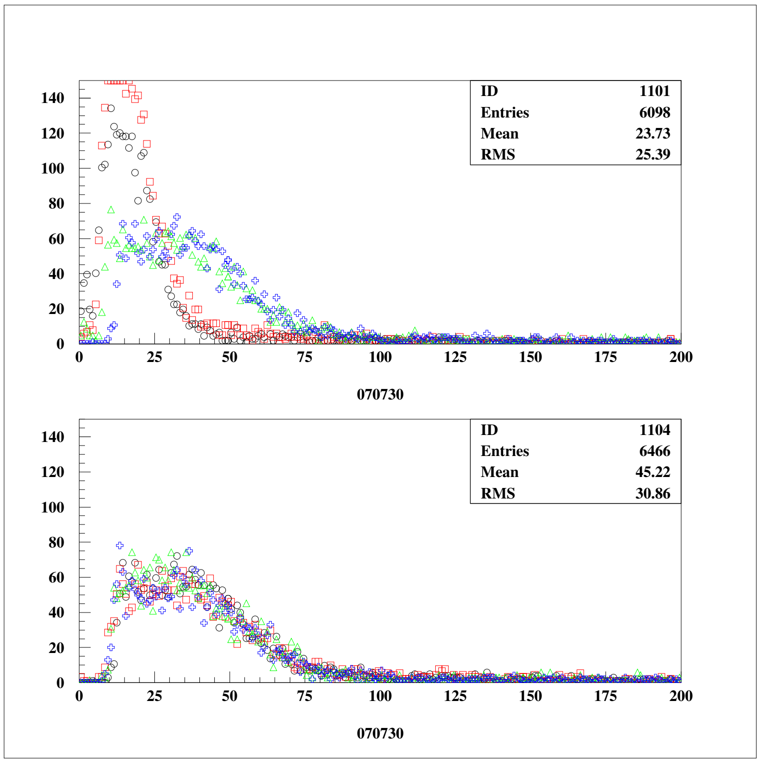

Figure 15: Comparison of channel identified as blowing its back termination to one that did not, using laser data.

breakdown.

In addition to 16/9/14, a total of four other channels were identified as having potentially blown terminations based on the NHIT analysis, all before June (and none since then). These were all considered 'marginal' and subsequently showed no reflection nor a doubled gain.

4.4 Occupancy Changes

The most important measure of the negative effects of high voltage on the detector is the rate of tube death. Not all tubes that die are due to high voltage problems, but if the rate of overall tube death is low then we are safe in assuming that the rate of death from high voltage related phenomena must be even lower.

Measuring the rate of tube death is not as easy as it might appear at first, since the

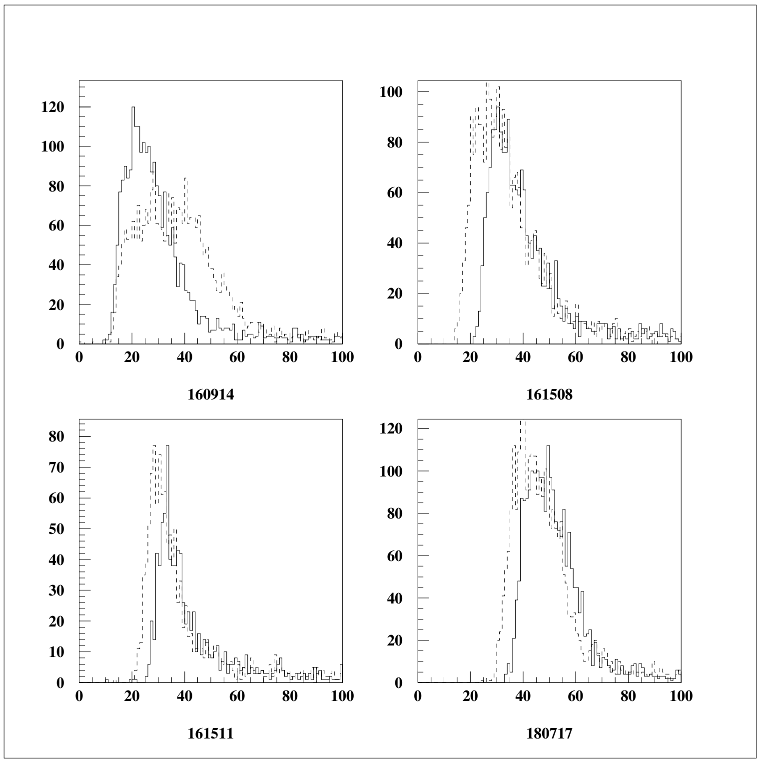

Figure 16: Channels flagged as potentially having a doubled gain. Only 060914 was identified with hand-scanning as having its terminator blown, and was subsequently found to have a doubled gain but no reflection.

detector running over the past few months has been characterized by periods of shutdowns and repairs which can artificially increase the death rate, and often tubes which appear dead can be easily fixed by replacing or fixing some part on the dry end. However, using data taken with the laser and correlating these with measurements of the tube channel noise rates and the tube base current measurements, we can come up with the history of tube failures shown in Figure 17. The first point, which was the first measurement made after turn-on, shows the total number of dead tubes with which we began running. These tubes are those which never worked (due to some unknown problem when the detector was still dry), plus those damaged due to the wet-end breakdown problem. The second point, which also shows a fairly steep slope, is likely dominated by those tubes which were missed the first time the

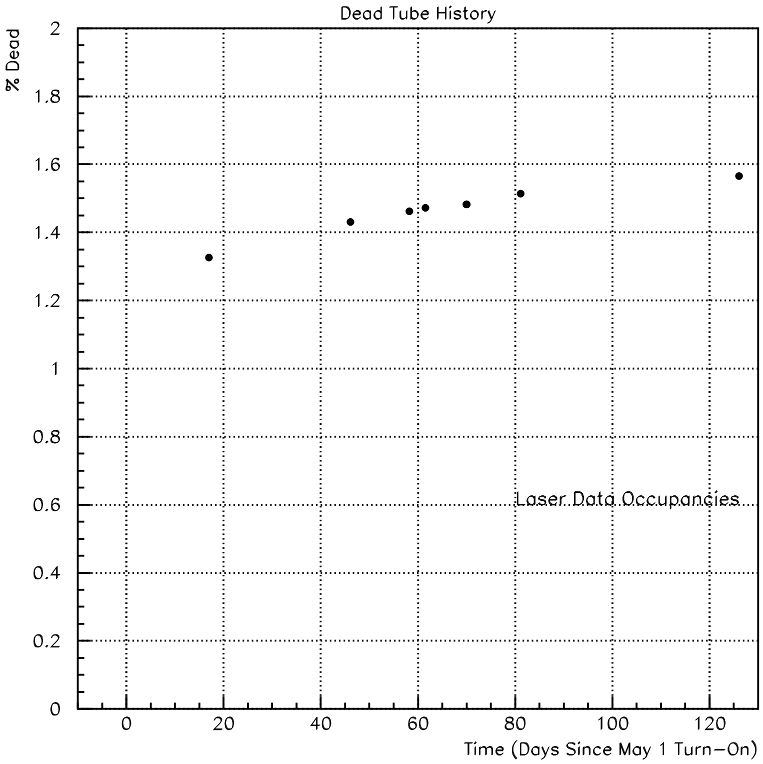

Figure 17: History of Tube Failures

measurement was made, and are not likely to be representative of the true rate of tube death. After that, the slope flattens out, and by the last two points on the plot (those taken during the most stable running period, except for the INCO power outage) we see a tube loss rate of roughly 0.4%/year.

Table 1 details the faults for each of the currently non-operational tubes. In the Table, the only classes which can be attributed to having been caused by high voltage related problems are those channels explicitly listed as having failed due to WEB (wet-end breakdown) and due to WEB studies (the cables or phototubes were removed to help understand the breakdown itself), plus some fraction of those with 'bad bases', 'other PMT problems', and 'zero occupancy'. At most, this number is 91, and, more importantly, of the five tubes which failed between the last two data points in Figure 17, only two are identifiable as having possibly failed due to wet-end breakdown. The others showed symptoms of base impedances

Failure Type	Total
WEB	13
Bad Base	8
Bad PMT/base	26
Zero Occupancy	34
Other PMT Problems	18
WEB Studies	18
Dry End	33

Table 1: Details of Tube Failure Modes

becoming near shorts, perhaps due to water leakage into a connector, cable, or base. The 0.4%/year of loss may only be a 0.2%/year loss due to high voltage-related breakdown.

5 Verification Plans

If a high voltage raise is undertaken, clearly it will have to be done with caution, and with attention paid to the gain (in both the specific and general sense) achieved.

To determine the change in detector gain, we will have to run a voltage scan with the ^{16}N or laser source, and in addition look at the spread in occupancies crate-to-crate to determine whether additional tuning is necessary.

The tools we have describe here will all be applied to monitor the rate of (and any change in) high voltage problems after a high voltage raise. Some of the tools (sharkfin and flasher rates) have already been incorporated into the automated SNOMAN processing, and therefore we will know almost as soon a each subrun is completed whether these rates are changing substantially. There are tools within the automated processing which now can help us identify high rates of breakdown (NHIT spectra, etc), but an alogrithm which looks explicitly for WEB events will be added very soon.

Initially, at the time the raise is done, we will rely immediately on detector operator monitoring, which has always been the fastest way to see problems. Any noticeable increase in the WEB rate (that is, anything non-zero) during the first 24 hours will be cause to bring the detector back to nominal voltage. During this first 24 hours, we will (as usual) be looking at the output of the automated processing routines to see if there is any evidence of increase high-voltage breakdown, or an unexpectedly large increase flashers or breakdown.

Long term tests will essentially be a repeat of what has been done in this memo, with the possible exception of the WEB algorithm which may no longer be appropriate. We propose that a first study of the increase in flasher and sharkfin rates be done on the timescale of a week after turn-on, and that a search for gain changes and an occupancy study done with calibration source data as well. (The gain change can also be done with NHIT data at this time or earlier). Between 4 and 6 weeks after the high voltage raise (and assuming there are no obvious problems otherwise) a reflection test and a second comparison of laser occupancies and gain measurements will be done. If we see no significant deviation from what we have seen in the past (for example, a doubling of flasher or sharkfin rates, an increase in the rate of blown terminations, etc.) then we will stay at the new voltage and continue to monitor the situation primarily through the automated processing, the occupancy studies (which will become automated themselves), and the gain change tool, with reflection tests as appropriate during detector downtime.

6 Summary

We find that while both the flasher and sharkfin rates may increase by 25-30% with a 100 V raise in the high voltage, the level of breakdown is not likely to change. In the case of flashers, the overall rate already varies by factors of two (not including the seismic burts, and so the additional 30% change is not likely to present a new problem. So far, our tube death rate is overall about 0.4%/year, with perhaps half of that from high-voltage related phenomena. We have also shown that sharkfins do not appear to have any effect on channel occupancies, and therefore may be benign.

In addition to providing a quantitative baseline of rates of the various phenomena, this document has also demonstrated the various tools we have for alerting us to any changes or new problems, and, in the light of this, we have outlined a proposal for how to watch the detector after the high voltage is raised. Many of the tools described here have already been added to the automated processing, and hopefully more will be added soon.