SNO-STR-99-021

Investigation of Relative Channel Efficiencies

Bill Frati, Josh Klein, Richard Van de Water, Peter Wittich
University of Pennsylvania

September 24, 1999

Abstract

We describe here an investigation of the SNO relative channel efficiencies. Based on threshold scans and gated charge runs with the laser source, we find that the loss of relative efficiency with each threshold DAC count increase is $\sim 3\%$. Using scans in high voltage (+40 V, +80 V, +100 V) with the laser, we find that the increase in efficiency is roughly 10% for a 100 V increase.

1 Introduction

The determination (or estimation) of the efficiency of light detection for each channel in the SNO detector is critical for a complete understanding of SNO's energy scale and resolution. What we would like, of course, would be an absolute measurement of the efficiency: if a photon is incident upon a given photocathode, how often do we register a hit? Such a measurement depends upon the photocathode quantum efficiency, the tube gain and the channel (relative) threshold—and it is unlikely we will be able to measure these each individually. However, what we *can* measure is the *relative* efficiency of each channel and, perhaps more importantly, the relative efficiency as a function of channel threshold or tube high voltage.

The relative efficiency as a function of channel threshold can tell us how much we stand to lose or gain by changing thresholds to improve, for example, detector stability. In principle, with a knowledge of the absolute efficiency at any given threshold (perhaps specified as a fraction of a photoelectron) we can also know what the average absolute efficiency is in the detector as a whole and compare that with our physics needs. In addition, knowledge of the relative effiencies may be useful for inclusion in the detector simulation, which will allow us to understand the systematic effects on the energy scale and resolution.

In the event that we find we need to increase the detector's energy scale to achieve our physics goals, a measurement of the relative efficiency as a function of high voltage may tell us how small a change we can make and still get a significant improvement. Obviously, the smaller that change the less likely the detector is to suffer from other high-voltage influenced problems (breakdown, sharkfins, flashers...).

Currently, the threshold settings in the SNO detector are determined by the point at which the channel starts to suffer from readout-induced pickup. In many cases (as we will see) this level is far above the noise level of the phototubes or any other uncorrelated noise source. Lowering these thresholds will only be possible if the pickup problem itself is solved—and this may require significant modifications to parts of the system such as the backplane. We therefore expect that, without either a solution to the pickup problem or a change in the channel gain, we will remain somewhat below our optimum channel efficiency.

2 Data Sets

The data sets used in this study are summarized in Table 1.

There are essentially four classes of runs:

- **Δ Threshold Scan**— Varying the thresholds on all channels relative to nominal for either laser or ^{16}N data. This lets us learn the slope of the efficiency change near where we actually run the detector, as well as allow us to see how the spread in occupancies is affected by the different thresholds from channel to channel. (The threshold differences are for the most part a result of different levels of noise and pickup on individual channels).

Run Number	Date	Run Type
3467	May 14	Standard PCA/Laser Occupancy (3% occ.)
3526	May 16	ΔThreshold Scan V$_{th}$ =Nominal
3527	May 16	ΔThreshold Scan V$_{th}$ =Nominal + 3
3528	May 16	ΔThreshold Scan V$_{th}$ =Nominal - 2
3563	May 17	Gated Charge
3548	May 17	Uniform Threshold Scan V$_{th}$ = 3/Laser
3547	May 17	Uniform Threshold Scan V$_{th}$ = 4/Laser
3545	May 17	Uniform Threshold Scan V$_{th}$ = 6/Laser
3552	May 17	Uniform Threshold Scan V$_{th}$ = 7/Laser
3544	May 17	Uniform Threshold Scan V$_{th}$ = 8/Laser
3553	May 17	Uniform Threshold Scan V$_{th}$ = 9/Laser
3543	May 17	Uniform Threshold Scan V$_{th}$ = 10/Laser
3542	May 17	Uniform Threshold Scan V$_{th}$ = 12/Laser
3554	May 17	Uniform Threshold Scan V$_{th}$ = 16/Laser
3555	May 17	Uniform Threshold Scan V$_{th}$ = 20/Laser
4261	July 6	ΔThreshold Scan V$_{th}$ =Nominal
4262	July 6	ΔThreshold Scan V$_{th}$ =Nominal+6
4263	July 6	ΔThreshold Scan V$_{th}$ =Nominal+3
4264	July 6	ΔThreshold Scan V$_{th}$ =Nominal+1
4265	July 6	ΔThreshold Scan V$_{th}$ =Nominal-1
4369	July 9	Gated Charge
4370	July 9	Gated Charge (no light)
4377	July 9	HV Scan, Nominal + 40 (4 crates)
4378	July 9	HV Scan, Nominal + 80 (4 crate)
4380	July 9	HV Scan, Nominal + 100 (4 crate)
4538	July 16	Δ Threshold Scan V$_{th}$ =Nominal
4541	July 16	Δ Threshold Scan V$_{th}$ =Nominal - 1
4542	July 16	Δ Threshold Scan V$_{th}$ =Nominal - 2
4543	July 16	HV Scan, Nominal + 80
4544	July 16	HV Scan, Nominal + 100

Table 1: Runs for Threshold Efficiency Studies.

- **Uniform Threshold Scan** — Setting the thresholds for all channels to the same value, and varying these up and down from run to run, for either laser or ^{16}N data. From these runs we learn how well we could do if all thresholds were the same, and at what point we gain more (uncorrelated) noise than signal by lowering the threshold. We can also compare these to the ΔThreshold scans to see how the spread in occupancy narrows when we have all thresholds the same.

- **Gated Charge** — Gating the integrator with the laser, so that the discriminators are removed from the process (as if we had set the channel thresholds to zero). The pulse height spectra from these runs shows us the complete PMT charge distribution, from pedestal upward. We therefore know whether the spectrum has a valley, a good single photoelectron peak, etc. In principle, placing a software threshold on these spectra should reproduce the occupancy fractions we see with the discriminator thresholds cutting on the spectrum.

- **High Voltage Scan** — Varying the high voltage on some of the crates to see how the gain and efficiency change.

3 Occupancy Studies

Perhaps the simplest (and maybe the most simplistic) way of viewing the relative channel efficiencies is just to look at how the fractional detector occupancy changes with threshold during laser runs. As the threshold is lowered, more tubes are hit with each event and once the contribution from uncorrelated tube noise is subtracted (as measured by PULSE_GT data taken simultaneously) and the change in laser intensity from run to run is corrected for, we get a measure of the increase in relative efficiency.

Figure 1 shows this type of occupancy plot, for a Uniform Threshold Scan. In the Figure, the mean NHIT for events which have triggered the EXT_ASYNC trigger used by the laser and those that have set the PULSE_GT trigger bit are plotted separately as a function of discriminator threshold. Although the occupancy increases steeply once the threshold gets below about 6 or 7 Counts Above Zero (CAZ), it is clear that the mean NHIT for the PULSE_GT data (that is, the noise) is responsible for most of this increase. In addition, the spread in occupancy event-to-event is also significantly larger below 6 CAZ, indicating that there may be some instability at those thresholds. One important point to keep in mind in this type of analysis, however, is the possibility that the increase in the PULSE_GT occupancy at low threshold (and the noise part of the EXT_ASYNC occupancy) may be due primarily to correlated sources of noise, such as pickup from the laser or event readout, and therefore the 'limit' of 6 CAZ may not reflect the point at which we are dominated by PMT noise. If that is the case, solving some of the pickup problems may allow us to move the average threshold below 6 CAZ.

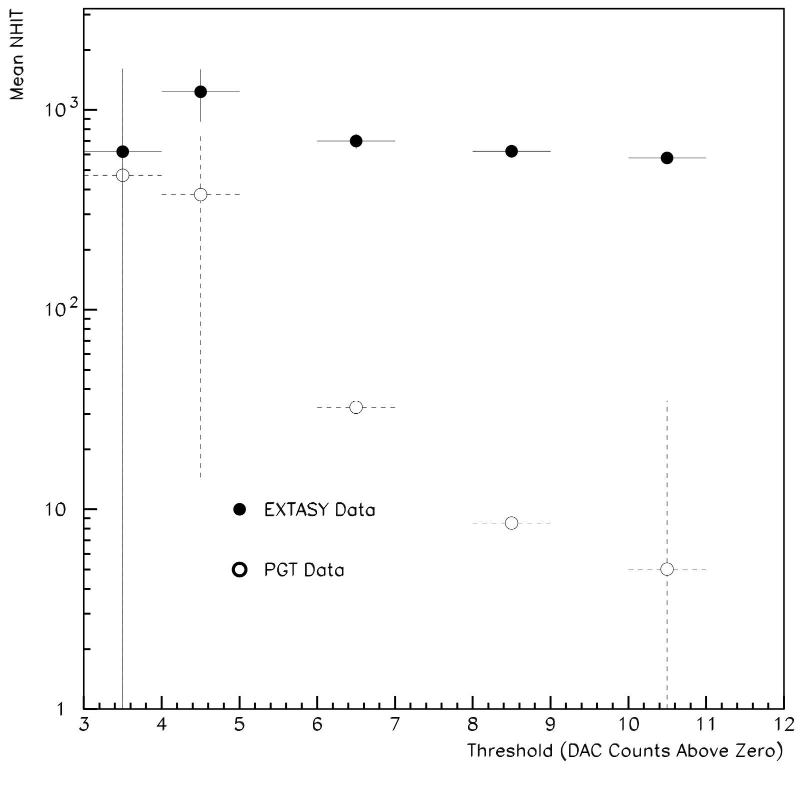

Figure 1: Occupancy Plot for a Uniform Threshold Scan, both PULSE_GT and EXT_ASYNC Triggered Data are Shown.

Although the Uniform Threshold Scan provides a map of the occupancy (and, in turn, relative efficiency) as a function of threshold, it does not tell us what we would stand to gain by globally varying the thresholds currently set. The threshold settings now (depending upon which set is being used) sit at an average of around 9.5 CAZ, but because the pickup from readout is not uniform with crate geography (slot 0 and slot 15 having the worst problems) there is a large spread in threshold channel-to-channel (mostly slot-to-slot). Figure 2 shows the distribution of threshold settings in the detector for the 'skey 520' threshold set extracted from the DQXX banks written out by the DAQ, for all operational channels and for only those channels which are not on the exterior of the crate (slots 0 and 15 are excluded as are all channels 0 and 31). We are likely to change thresholds by moving them relative to these non-uniform settings, and Figure 3 shows an occupancy plot for such a ΔThreshold Scan. As in Figure 1, both the data from EXT_ASYNC laser triggers and PULSE_GT triggers are shown. In Figure 4, the same data is shown with the PULSE_GT occupancy subtracted, and with the overall occupancy normalized to the occupancy at nominal (ΔThreshold = 0) threshold. (Note that because of laser instabilities at the time, both Figure 3 and Figure 4 have been corrected for the intensity changes by using slot 2 of every crate whose thresholds were unchanged). The error bars in these two plots show the spread in the occupancy, not the uncertainty on the mean occupancy.

Figure 4 shows us that lowering all the thresholds by 1 count relative to nominal (nominal here was the 'skey 520' threshold set) provides a gain of about 3% above the increase in the PULSE_GT noise, and raising the thresholds by 1 count is a loss of about the same. Lowering the thresholds globally by more than 1 count is not likely, as that puts many of the channels around 4 CAZ, well into the regime where noise dominates (either from readout pickup or from PMT noise itself).

Taken together, the Uniform Threshold Scans and the ΔThreshold scans show us that the lowest thresholds we could run at are near 6 or 7 CAZ, and the fact that our average threshold sits near 9.5 CAZ means we are running at a relative efficiency up to $3\% \times 3.5 = 10.5\%$ lower than what we could do were readout pickup not a problem. Of course, we do not really know how this compares to our expected absolute efficiency—that depends on the details of how that absolute efficiency was estimated for the Monte Carlo or assumed in our earlier calculations.

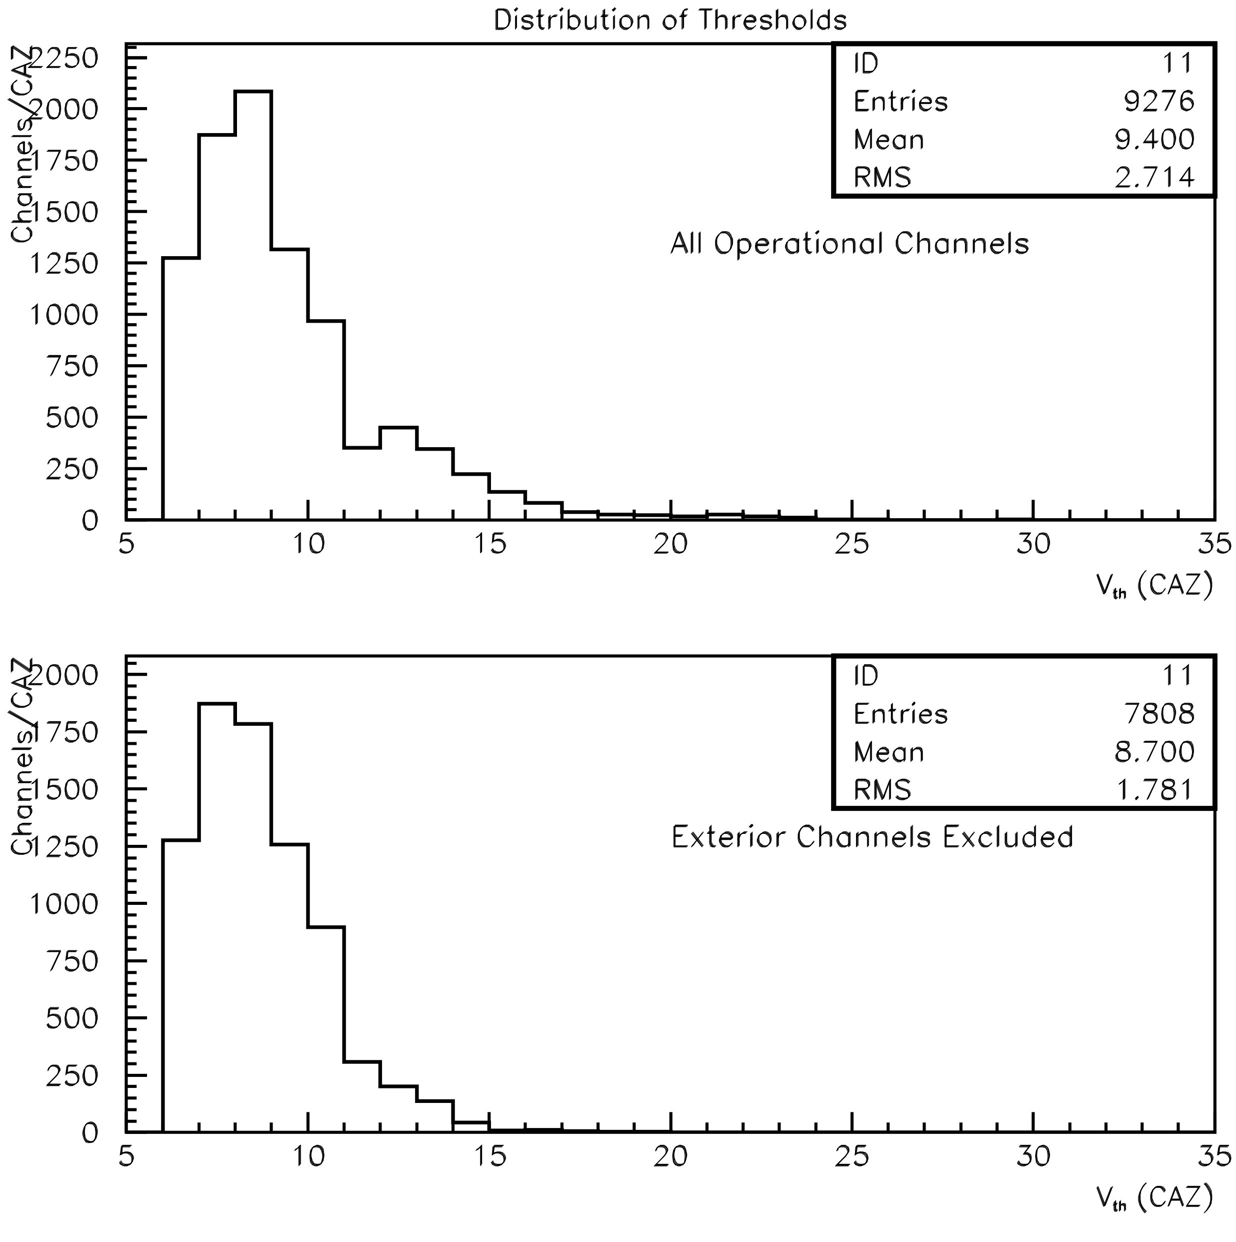

Figure 2: Distribution of Channel Thresholds for all channels (upper plot) and those not in exterior channels (lower plot).

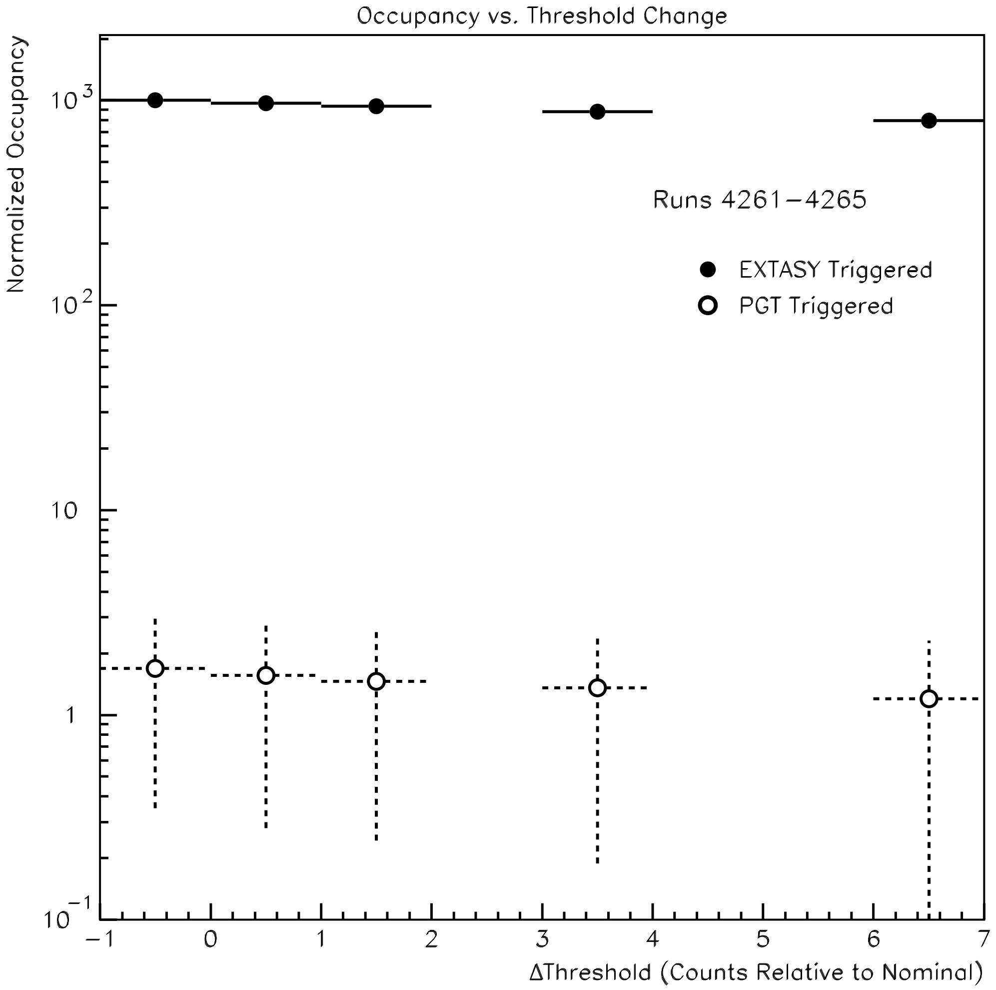

Figure 3: Occupancy Plot for a ΔThreshold Scan. Both laser (EXT_ASYNC) and PULSE_GT data are shown.

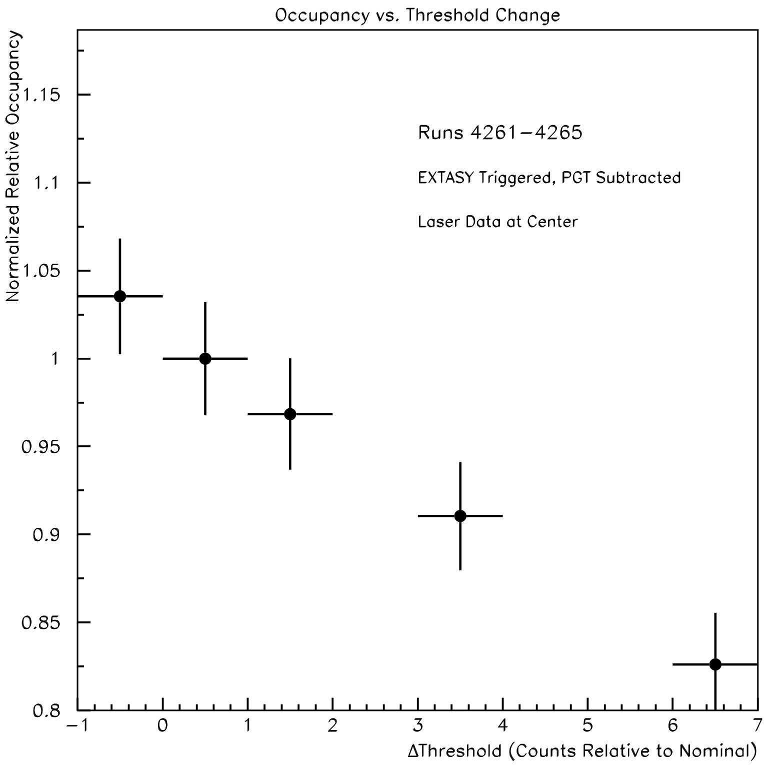

Figure 4: Occupancy Plot for a ΔThreshold Scan. The contribution from PULSE_GT has been subtracted, and the occupancies have been normalized to the nominal threshold setting.

4 Relative Efficiency on a Channel Level

Although the threshold scans tell us how the detector as a whole behaves under global changes in threshold, they do not necessarily tell us about the variation channel-to-channel under normal conditions, nor do they explain why the relative efficiency vs. threshold curve looks as it does. To get a more channel-level view, we can produce occupancy vs. threshold curves which use the average channel (rather than detector) occupancies, and plot those against the various channel threshold settings.

Figure 5 is plot of the relative occupancies of channels at various thresholds, for a standard run (Run 3467) and a run in which all the thresholds were lowered by two counts (Run 3528). The normalization was done relative to the occupancy for channels at the lowest threshold in the two runs (4 CAZ). Each point is an average over all channels at the same threshold,

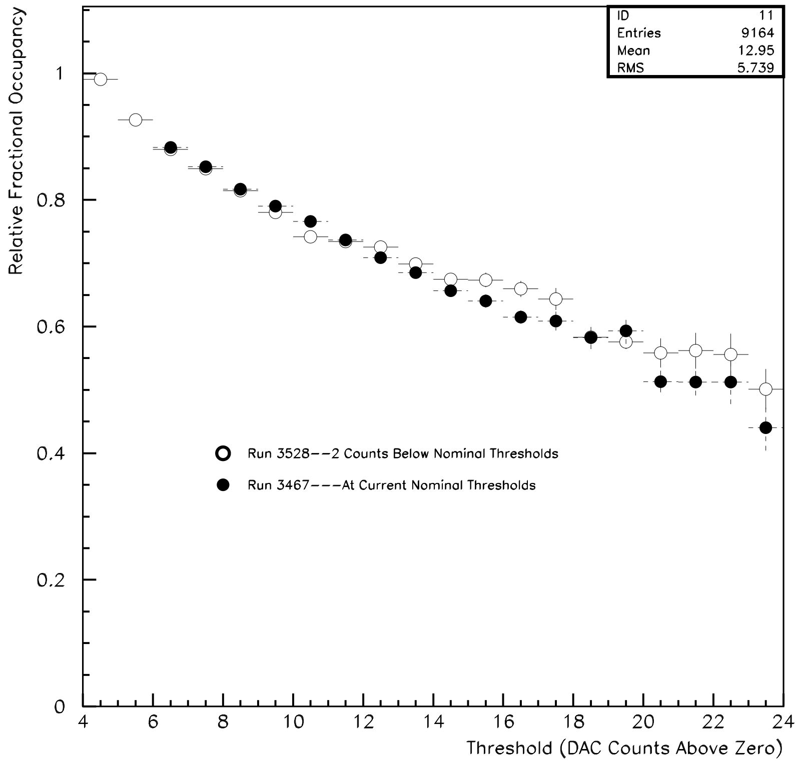

Figure 5: Occupancy vs. Channel Threshold for a run at nominal threshold and a run at nominal-2 CAZ.

and channels which are believed to be non-operational for any reason (based on the DQXX banks) are not included in the average.

By averaging over all channels at each threshold, we reduce effects on the channel occupancy (like optics) that have nothing to do with the relative threshold efficiency. Not all of these effects are removed, however, and in particular the fact that most of the channels with high threshold ($>$ 16 CAZ) are in slot 15 of each crate and that these tubes tend to be geographically grouped in the detector may be the reason that there is some structure in the plot at those high thresholds. In addition to being a way to understand the relative threshold efficiencies, Figure 5 can be used in the Monte Carlo simulation to adjust the relative efficiencies and thus more accurately reproduce the detector behavior.

Overall, however, what we see here is very similar to what we saw in the occupancy plots for the threshold scans: near 9 or 10 CAZ, the change in relative efficiency with each count in threshold is $\sim$3%. At low thresholds (below 6 CAZ) this slope becomes steeper, agreeing with what we found before for the increase in noise. One disappointment, however, is the lack of a plateau region: we would hope that as the threshold becomes lower and falls into the 'valley' of the single photoelectron spectrum, that the increase in relative efficiency would take on a much shallower slope until the noise peak is hit and the curve becomes steep again. Figure 5 does not show any evidence of such a plateau region, though as stated above, the inflection as the curve becomes steeper does look to occur around 6 CAZ. The lack of a plateau may be a result of the fact that we are dominated by pickup (and so we begin to fire on correlated noise sources when we are in the valley of the charge spectrum) or simply because of the averaging we are doing—if the valley occurs in slightly different places for different tubes because of differences in gain, we will smear out the plateau.

One way to investigate the curve shape is to look at the same occupancy plot for tubes with higher gain. We would expect that if the efficiency is already fairly high, doubling the gain (for example) should make a small difference in the relative occupancy. The set of channels which have blown back terminations (roughly 400) provide a fair sample of tubes with doubled gain, though their distribution is not particularly uniform in the detector. Figure 6 compares the relative efficiency curves for channels with blown back terminations to those with normal terminations, and it is clear that the sample with blown back terminations has a significant amount of structure in it, perhaps because of the non-uniform distribution. Looking at the points at the lowest thresholds, however, we see what looks to be roughly a 5% increase in the relative efficiencies of those channels with blown terminators. Some of this increase is probably due to the increased PMT noise hits with the doubled gain, but overall this is roughly what we would expect if our efficiency was already fairly high ($> 90\%$). It also appears (with the use of some imagination) that the change in efficiency is much greater for those with high thresholds, as one would expect. However, given the systematic structure in the curve, we still need something else to tell us why there is no plateau region, and why the falloff is 3% with each CAZ.

The Gated Charge runs provide us with a way of understanding the relative efficiency vs. threshold curve in terms of the charge spectrum of the channels themselves. By removing the discriminator entirely from the data chain, we get to see what we're missing. Figure 7 shows

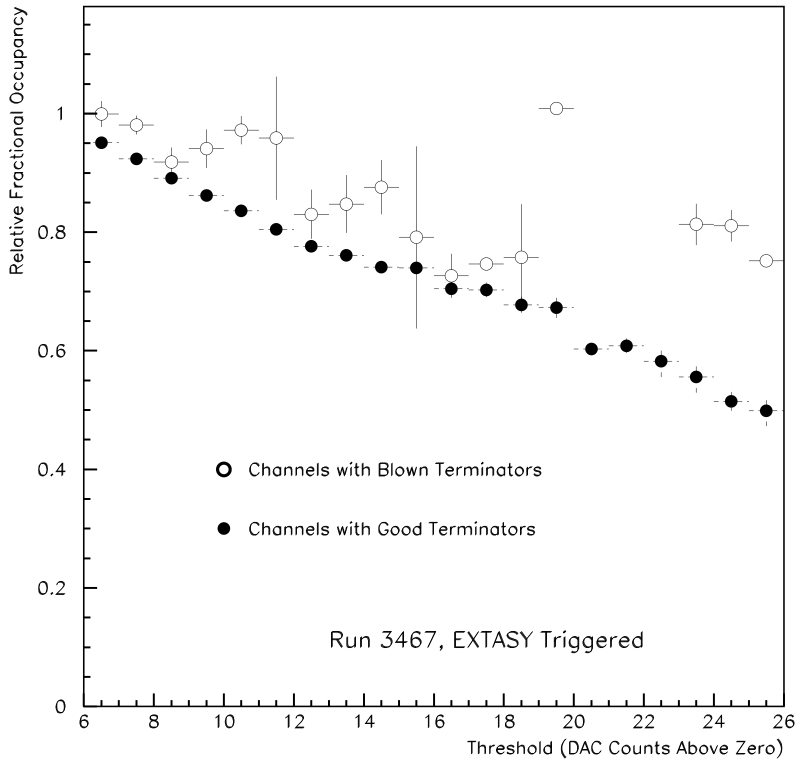

Figure 6: Occupancy vs. Threshold for normal channels and those with blown 75 Ohm back-terminations.

the gated charge spectrum for half of the detector, compared to the discriminated charge spectrum for the other half using the high gain, long integrate charge (QHL). Interestingly (and perhaps fortuitously) the two halves of the detector agree well with one another—no normalization has been done in the plot, they line up naturally. Although the data has been pedestal corrected, there is still a $\sim$ 5 count shift which as yet is unexplained (the shift occurs regardless of whether new- or old-style ECA constants are used). From this average gated charge spectrum, we do not see any obvious valley in the spectrum. This may partly be a consequence of the varying gains of the tubes (no gain correction has been made) so that the spectrum gets smeared out. The spectrum of a single tube will typically have a valley in it, however, as shown in Figure 8.

Figure 7 clearly shows the average bite the discriminator takes out of the full spectrum.

We can theoretically reproduce the occupancy vs. threshold curves we obtained with the discriminator threshold scans by simply taking the gated charge spectrum and applying a software threshold. Integration from the software threshold to infinity, normalized to the total integral, then gives us a map of the relative efficiency as a function of the software threshold. In principle, this is like taking data 'off-line'—we are applying the channel discrimination after the data has already been taken, and therefore this is a measure of relative efficiency just as the occupancy plots are. If the software threshold (which cuts on charge) is the same as the discriminator threshold (which cuts on mV), the two curves should look the same. Figure 9 shows the full efficiency curve using a software threshold cut on the gated charge spectrum, and Figure 10 compares this curve to the curve obtained earlier by varying the discriminator threshold. In Figure 10, the channels responsible for the structure out at high threshold (the channels on the exterior slots of the crate) have been removed.

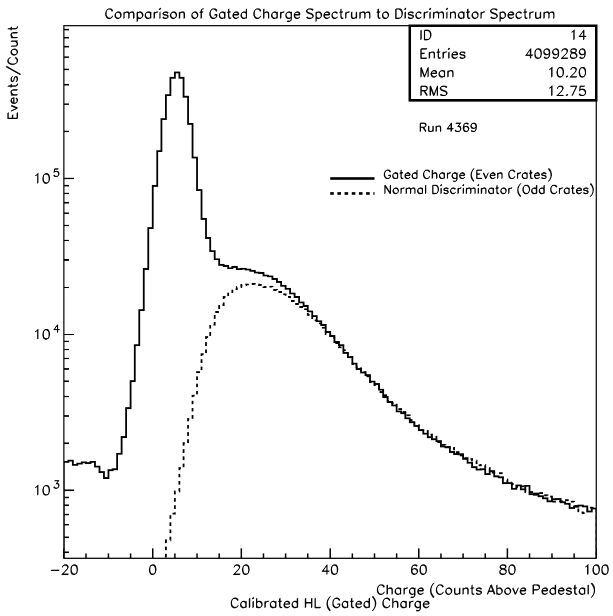

Figure 7: Comparison of gated (HL) charge spectrum for all even crates and discriminated spectrum for odd crates.

The close agreement between the two curves in Figure 10 illustrates our understanding of the effects of the discriminator—what we are seeing in relative efficiency changes as we vary

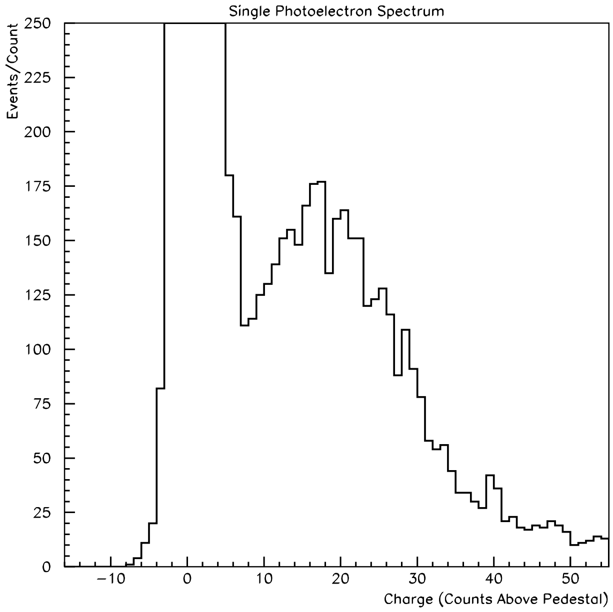

Figure 8: Gated Charge spectrum for a single channel, crate 0, card 0, channel 4. The additional 5 count offset of the pedestal seen in the gated charge data has not been accounted for here.

the discriminator threshold makes sense in terms of the charge spectra that the channels are seeing. If we were dominated by pickup at the low thresholds, we would expect the two curves to disagree: the disciminator 'sees' pickup that the integrator does not, because the integral of the pickup is generally zero. By cutting on charge offline and seeing agreement with the discriminator's cut on mV on-line, we have shown that the noise increase at 6 or 7 CAZ is primarily due to sources with non-zero integrals (like PMT noise). Although this means that the pickup is not so bad that it is affecting the charge spectrum, this says nothing about how pickup effects our triggering, since it occurs as a correlated source which fires many tubes at once. The lack of a plateau in this plot looks like a consequence of the averaging we are doing (or, more correctly, as a consequence of the average channel behavior). A single channel does show some plateau as shown in Figure 11, which is a comparison of the integral of the gated charge spectrum for a single channel as a function of software threshold to that for the sum over all the even crates in the detector and the channel shown in Figure 8. The

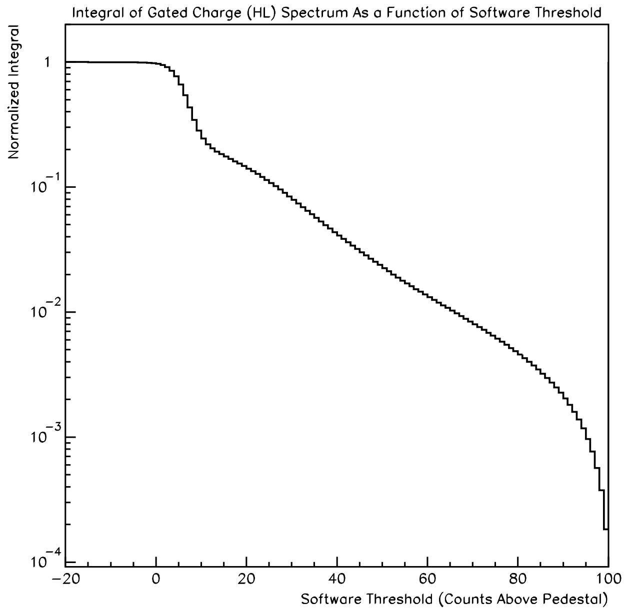

Figure 9: Integral of Gated Charge Spectrum as a function of software charge threshold, normalized to total integral of charge spectrum.

flattening of the dotted curve in Figure 11 occurs just where the (offset-corrected) valley occurs in Figure 8. It would be interesting to see if, for this single channel, there is a plateau in a occupancy vs. threshold curve using, for example, a set of Uniform Threshold Scans.

What we conclude, therefore, is that the 3% change with each CAZ near our nominal threshold setting determined in the previous Section makes sense in terms of the average channel behavior. The steep increase around 6 CAZ looks to be a consequence of the actual noise seen by each channel (the noise peak in the average gated charge spectrum hits the single photoelectron spectrum there), and the discriminators themselves are behaving as we expect. What is missing here, of course (intentionally), are the effects of correlated noise such as readout pickup. Obviously from this analysis we would conclude that we should lower the thresholds so that we are running closer to 6 or 7 CAZ. However, the pickup problem prevents us from doing that on most of the channels. In that case, the best way to increase the relative efficiency is to change the gain of the tubes.

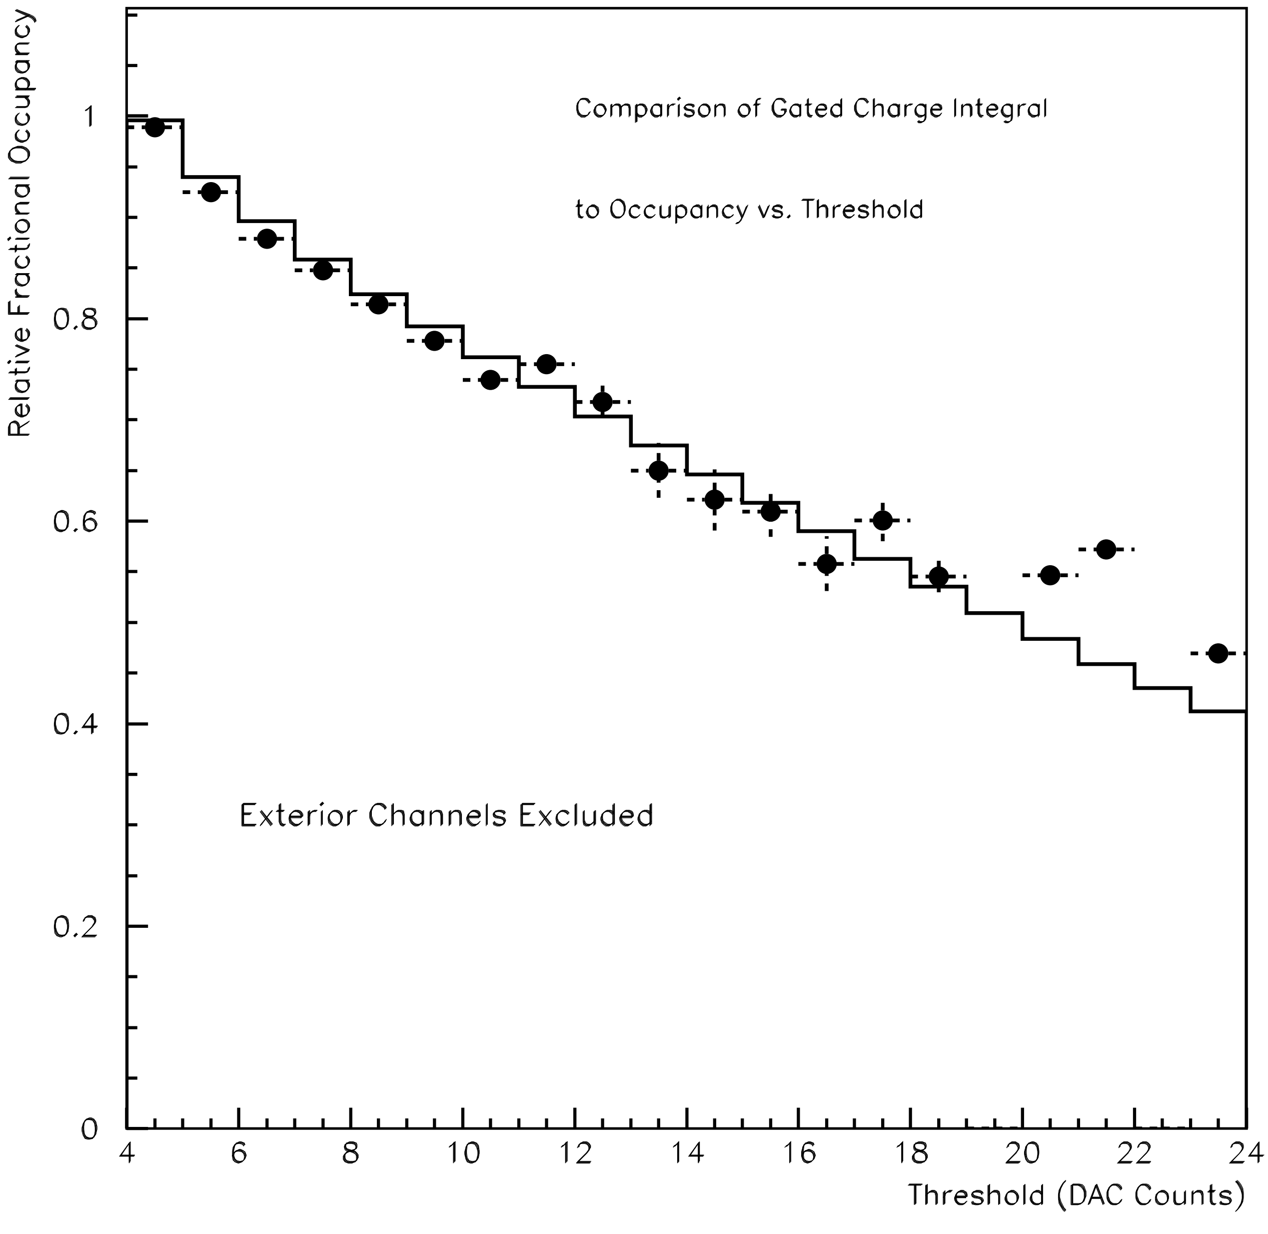

Figure 10: Comparison of Relative Efficiency Derived from Software Threshold Applied to Average Gated Charge Spectrum (solid line) and Hardware Discriminator Thresholds (dots).

5 High Voltage (Gain) Scaling

The previous sections have shown that we would like to be running our thresholds closer to 6 or 7 CAZ—the point at which further lowering of the thresholds adds more uncorrelated noise than signal, and that by doing so we stand to gain roughly 10% relative efficiency. Figure 10 shows that were we able to move the peak of single photoelectron distribution further away from the pedestal peak, the bite of the discriminator would be far less severe, and we would be in a situation perhaps where we were running at a threshold at which the channel efficiencies effectively looked like our ideal 6 or 7 CAZ.

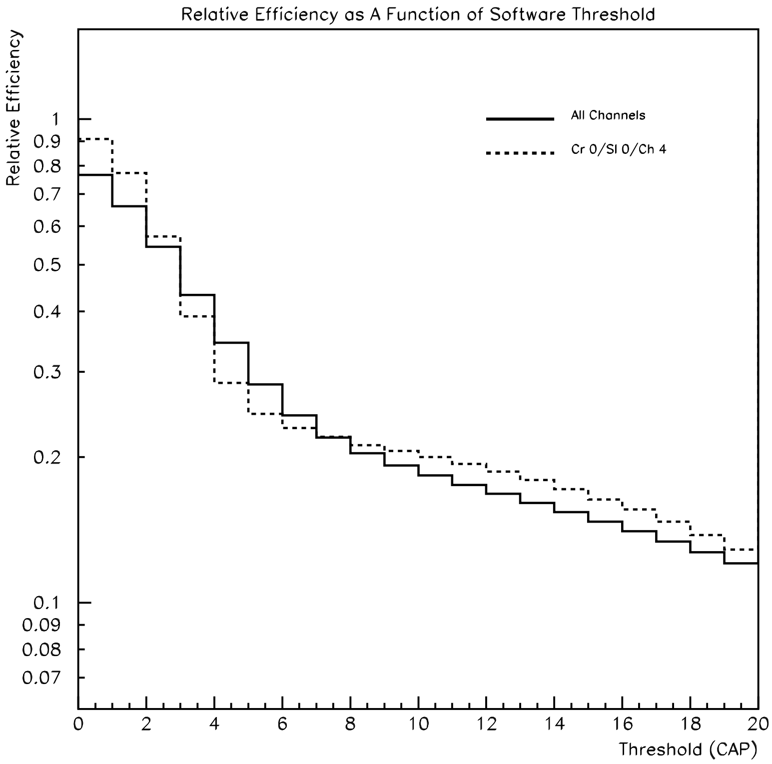

Figure 11: Comparison of Normalized Integrated Gated Charge Spectrum for both the average over all channels in the even crates and a single channel (Crate 0, Slot 0, Channel 4), in units of Counts Above Pedestal (CAP)— laser data.

Beyond the question of efficiency, however, are two equally (and perhaps more) important advantages of increasing the PMT gain: energy resolution and detector stability. As can be seen in Figure 12, the spread in occupancies is somewhat broader for our actual threshold settings (where different channels have different thresholds based on the point at which they suffer from readout pickup) than for a threshold setting where all tubes are the same. The additional smearing in occupancy results in a broadening of the detector's energy resolution. Although in principle such a broadening can be reproduced in the Monte Carlo simulation (and therefore incorporated into a spectral fitting algorithm) or unfolded through calibration (we would need to adjust the inferred energy for each event dependent upon *which* tubes were hit, not just how many) there will inevitably be some loss of information. Detector stability can be enhanced by raising the tube gain as well, because problem channels can have their thresholds raised with a smaller marginal hit in efficiency than at lower gain.

5.1 Relative Efficiency as a Function of High Voltage

To explore the increase in efficiency as the high voltage is increased, several sets of data were taken with both the laser and the ^{16}N. In these runs, four crates had their high voltage increased in steps from +40 V above nominal to +100 V above. The crates were selected to span as large a range as possible in nominal voltage settings, to account for differences in the slope of the gain curve at different voltages.

Figure 13 shows the increase for each crate in relative occupancy for each of the voltage steps in the laser data set (Crate 12 inadvertently did not have its voltage raised to the full 100 V in the last step). The varying intensity of the laser was corrected for by using the rest of the detector for normalization. What we see is roughly a 1% increase in occupancy for each 10 V increase in high voltage, for all crates except Crate 11, the crate with the highest nominal voltage in the sample (2270 V), which had only a 0.8% rise for each 10 V. (The average crate voltage is about 2050 V). In addition to the linear increase in occupancy with voltage, there is also an interesting correlation in relative occupancy with the nominal high voltage of the crates—the crate with the highest high voltage appears to have the highest relative occupancy, the one with the lowest nominal HV the lowest relative occupancy, and the other two somewhere in between. In fact, if we plot the relative occupancies of all the crates against their nominal HV settings as in Figure 14, we see that with the exception of Crates 0 and 1, the correlation is fairly strong. This may be an indication that either the HV settings on the crates is wrong or the gains are not correctly matched. However, looking at

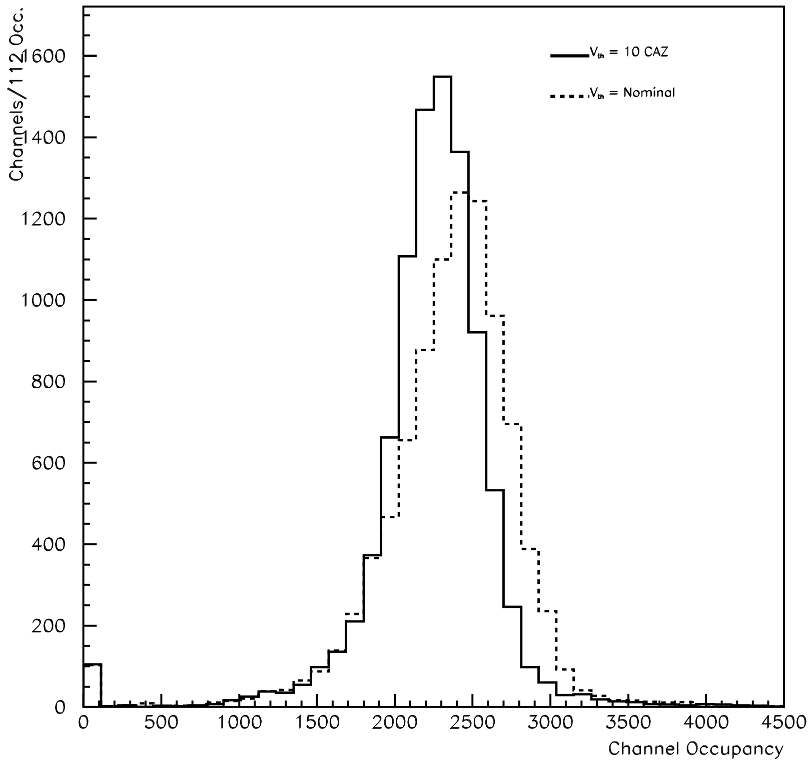

Figure 12: Comparison of the spread in occupancy for the real threshold distribution (dotted line) and a uniform distribution set equal to the average of the real distribution (solid line). The two runs were taken for equal times, so the absolute occupancy has been used here (laser data).

several runs shows the correlation is not consistently present in all the data, or at least not as strongly as is shown in Figure 14. Figure 15 shows plots similar to Figure 14 (although with the ordinate and abscissa reversed). Runs 4376-4380 all appear to have some correlation, as indicated both by their appearance and the correlation coefficient R (note that the two 'anomalous' crates, 0 and 1, were included in the correlation test), but Runs 4070 and 4603 do not. In any case, we will likely want to tune voltages on at least a crate-by-crate basis to get as flat an occupancy response as possible as a function of high voltage, whether we undertake a global change or not.

To better understand both the occupancy change and the differences from crate to crate,

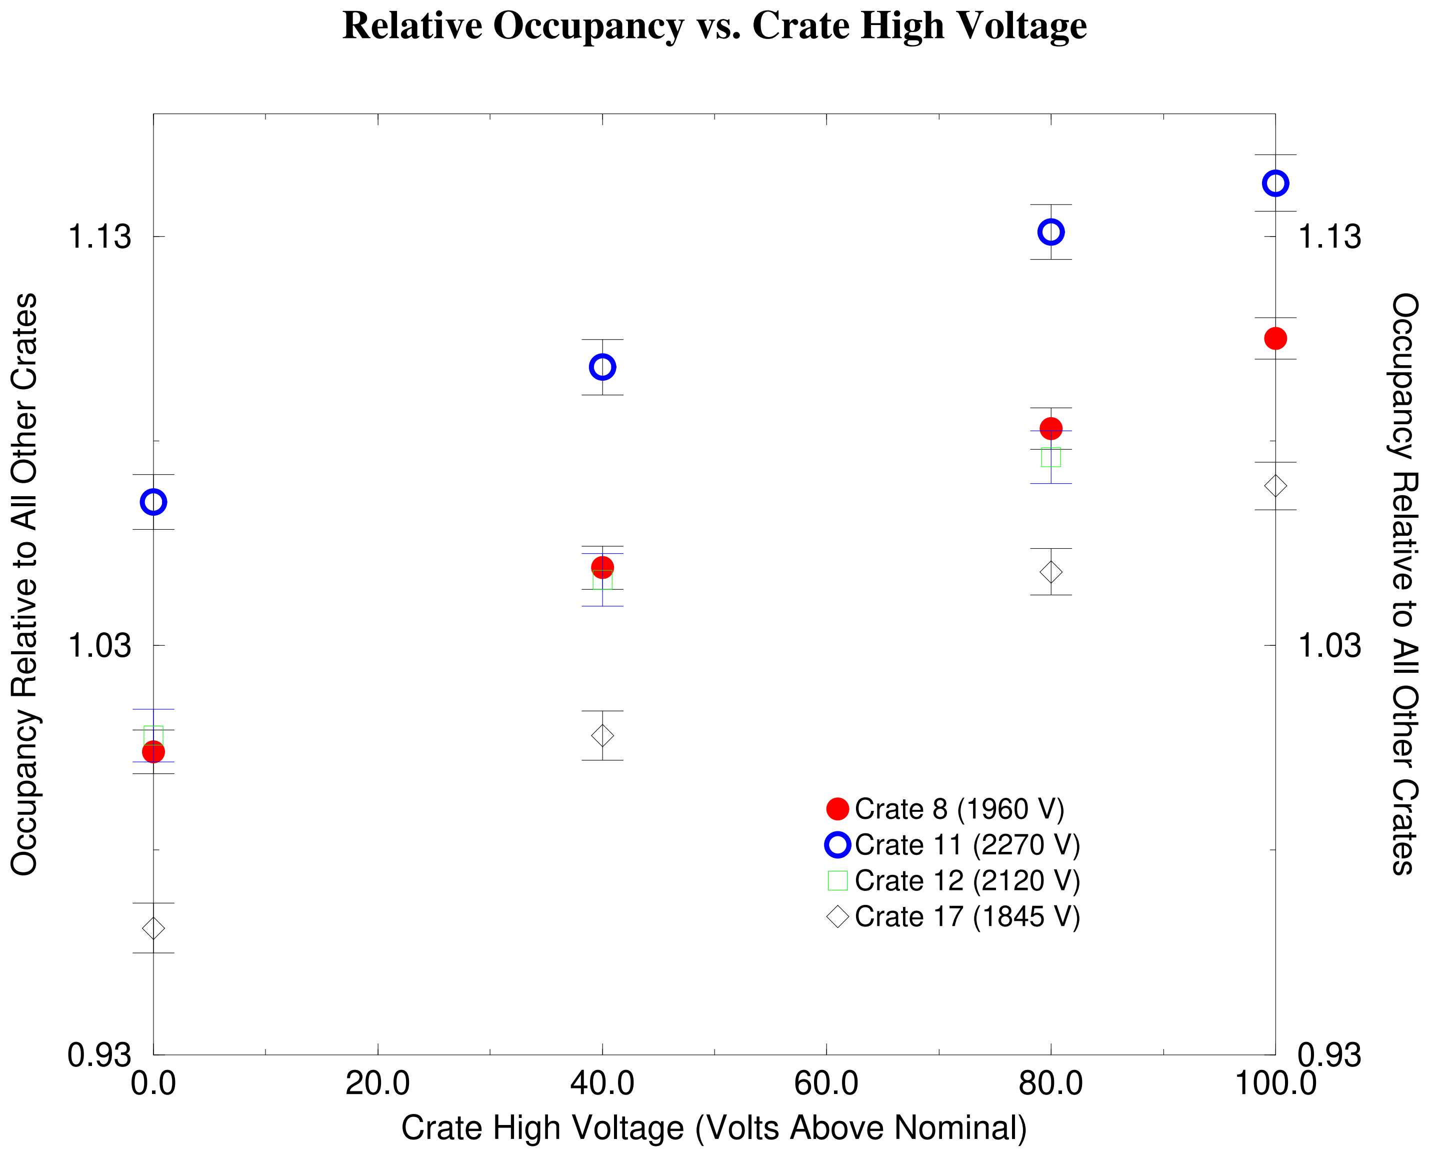

Figure 13: Crate Relative Occupancy vs. Voltage for High Voltage Scan, laser data.

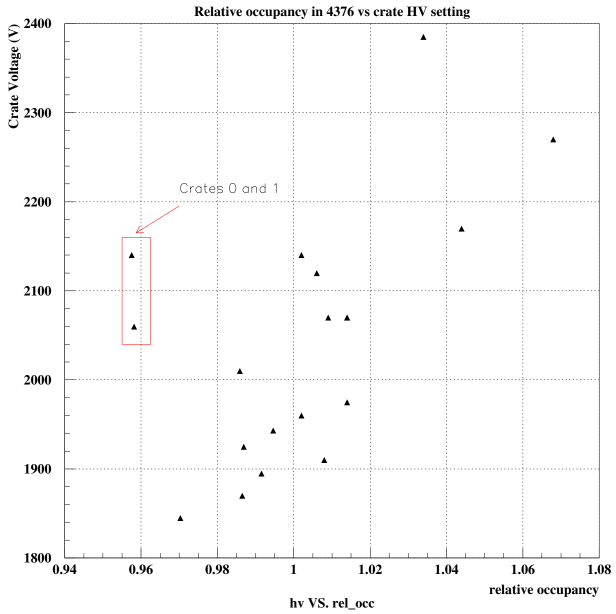

Figure 14: Crate Relative Occupancy vs. Crate Nominal Voltage, laser data.

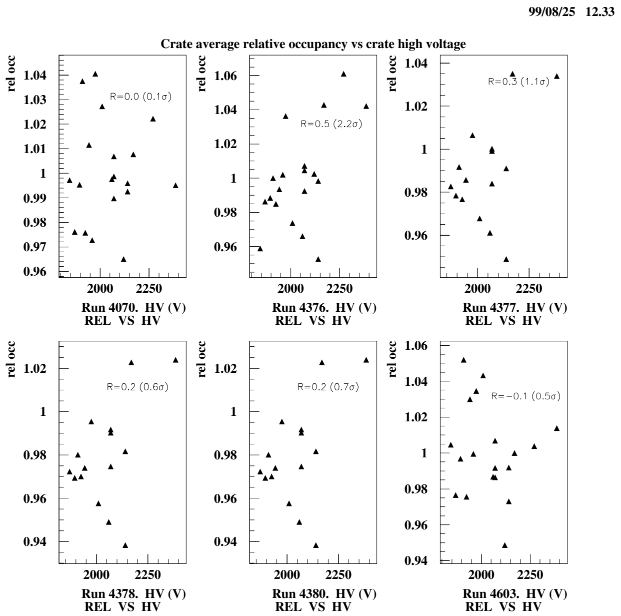

Figure 15: Crate relative Occupancy vs. Crate Nominal Voltage, laser data, for six different runs.

we have also looked at the gain change for some of the tubes with high voltage. The basic idea is that the pulse height (PH) shape does not change with the change in conditions. If the thresholds were set at zero then a simple scaling of PH by the gain change would leave the PH spectrum identical under different conditions. The discriminator not being zero however, means that only the upper portion of the PH distribution should scale. Under these conditions the shape of the PH spectrum at the fall off is the most sensitive parameter to the PMT changing its gain. It is this parameter that can be used to determine the gain change.

Fig 16 shows plots made by Richard Ford during air fill of the PH distribution at nominal HV and nominal plus 200 volts, using gated charge data as well as discriminator scan data. We can see from the Figure that the high half point increases by a factor of 1.6 when the HV is increased by 200 volts.

To see somewhat more clearly that we can see a gain change even with the discriminator thresholds active, Fig 17 shows the PH distribution of one PMT whose termination resistor is currently blown (resulting in a gain change of about 2). The top figure overlays data taken on 05/05(black circle), 05/14(red square), 06/15 (green triangle) and 06/27(blue cross).

The bottom figure shows the runs taken on 06/27(back circle), 07/05(red square), 07/09 (green triangle) and 07/20(blue cross). The gain change here is obvious (and it is obvious the termination resistor blew between 05/14 and 06/15). The high half point increases from 25 to 50 indicating a change of a factor of two as one would expect.

Fig 18 has the overlay of a PMT (in a crate whose HV = 2270 V) at nominal and +100 Volts. The gain change see here is 1.2. From the air fill run one would expect the gain change to be $\sqrt{(1.6)} = 1.26$.
Fig 19 is a similar plot for a PMT sitting in a crate whose HV is 1875 V. Here the gain change is a bit higher at 1.4 Note that the PMT sitting at the lower nominal voltage has the larger increase in gain.

5.2 Occupancy Spread as a Function of High Voltage (Gain)

We expect that, as the gain of the tubes is increased, those channels which have the highest thresholds will see the largest increase in efficiency, since the discriminator settings are sitting near the peak of the single photoelectron distribution. The result of this is that the

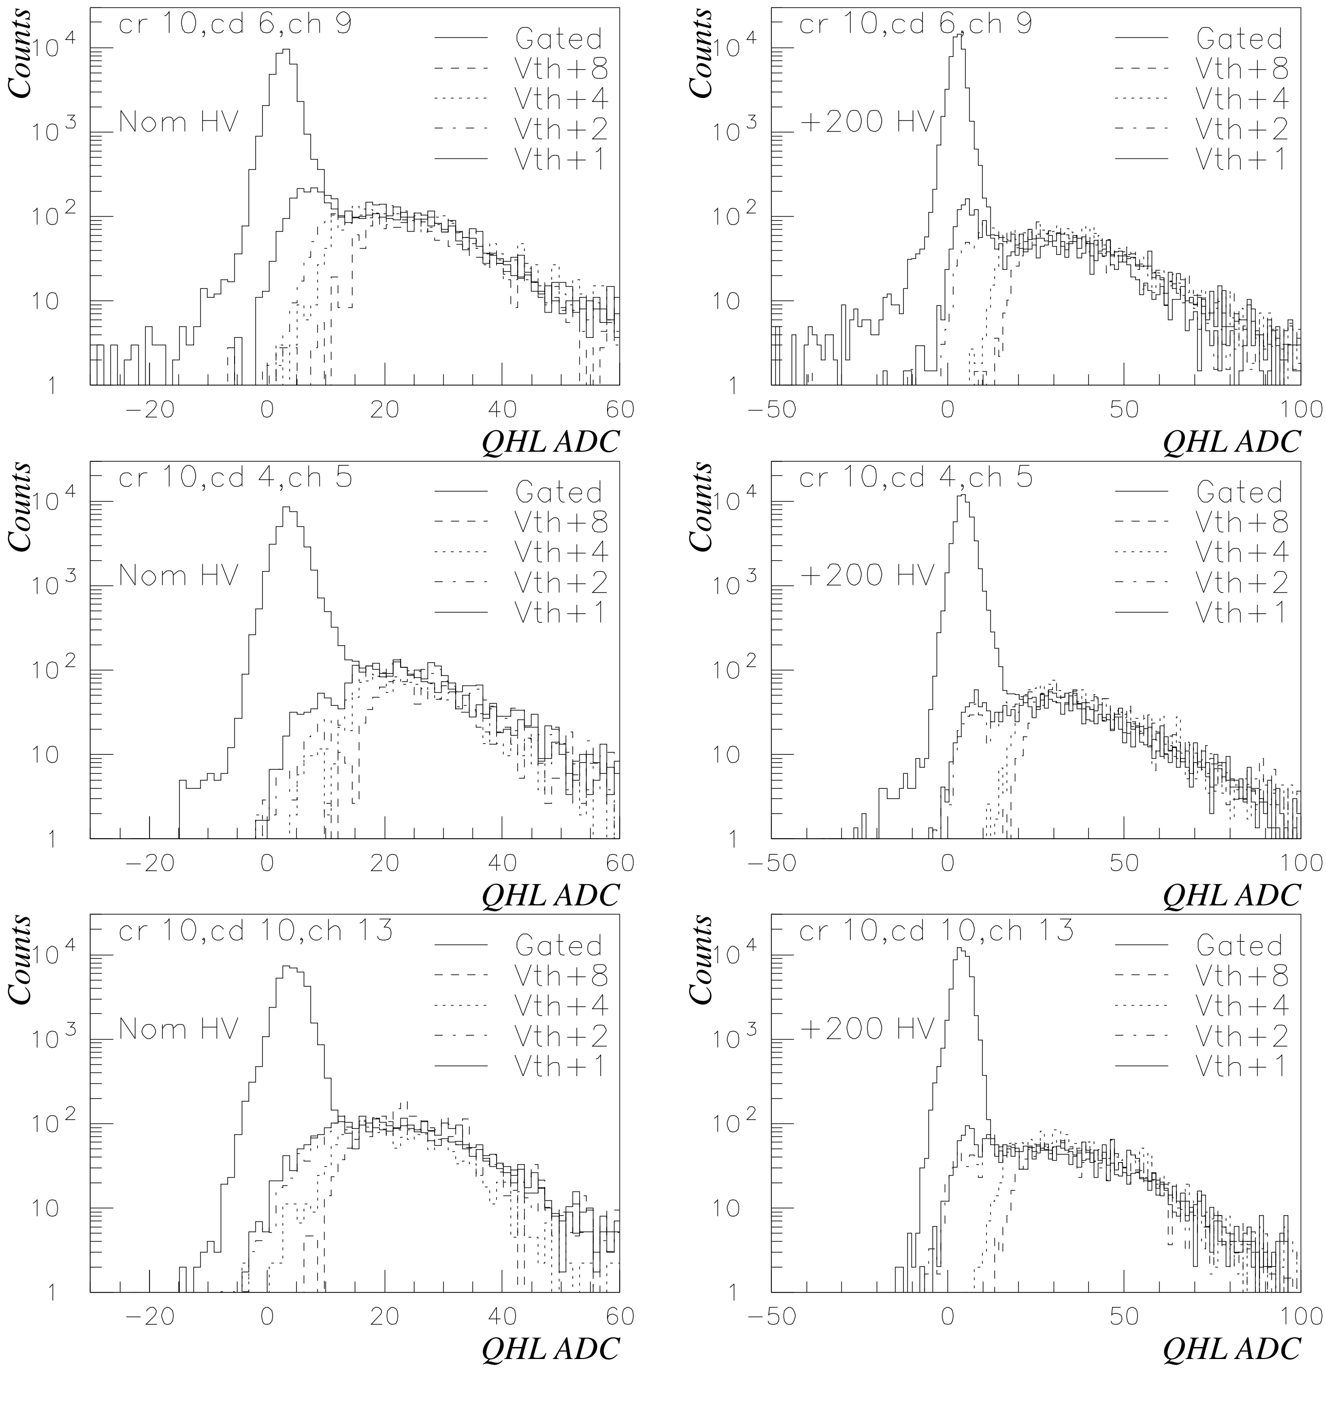

Figure 16: Richard Ford's Plots of Charge Distribution with HV at Nominal and at +200 V, for both Gated Charge and Discriminator Scans (air fill laser data).

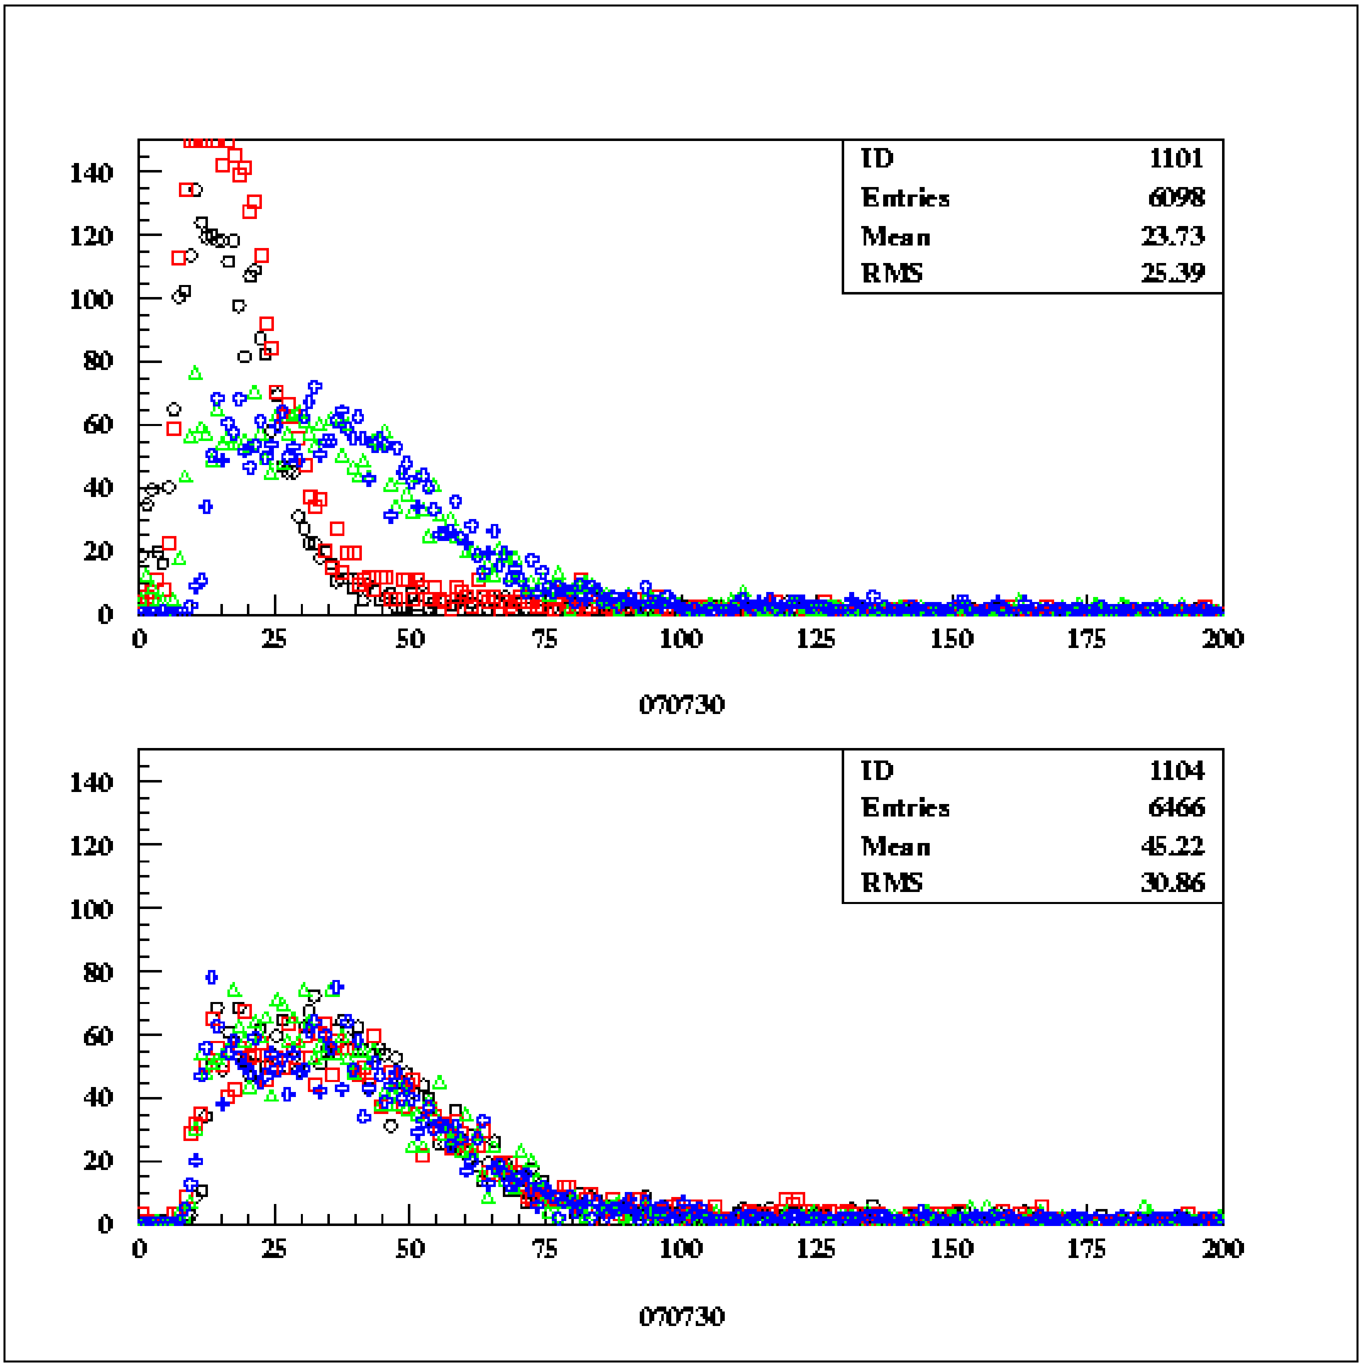

Figure 17: Change in Charge Distribution Between May 05 and July 20, for Crate 7, Card 7, Channel 30 (laser data).

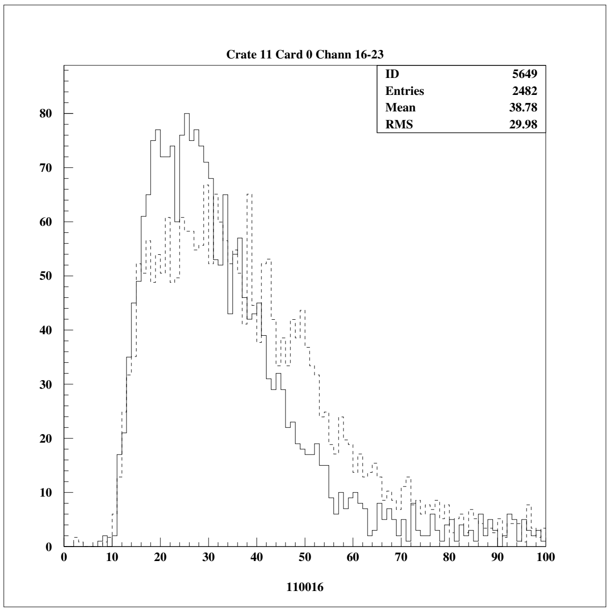

Figure 18: Charge Distribution with HV at Nominal and Nominal+100 V, for a channel in a crate whose nominal voltage is 2270 V (laser data).

variation in occupancy should be narrowed somewhat. Figures 20 23 plot the distributions in occupancy for each crate as the high voltage is increased. The fit to the distributions gives the best measure of the change in spread and is somewhat disappointing. Although there is some narrowing for Crate 8 and Crate 17 (the two with the lowest nominal high voltages), it is not particularly large. This may indicate that the spread in occupancy is being driven mainly by something other than the varying channel thresholds.

We can see the effects due to threshold alone when we plot the relative efficiency as a function of threshold for the different voltages. In this case, we do see a change in slope from 3% to $\sim$ 2-2.5% per CAZ as shown in Figures 24- 27 by the time the voltage is increased by +100 V. The fact that the overall distributions shown in Figures 20- 23 do not become significantly narrower may mean that a large fraction of the variation in occupancy is due to factors other than threshold. This does not mean, however, that reducing the spread due to the threshold variation is not a significant effect—factors like optics that can lead to variations in occupancy will be explicitly measured with calibration sources. We would do

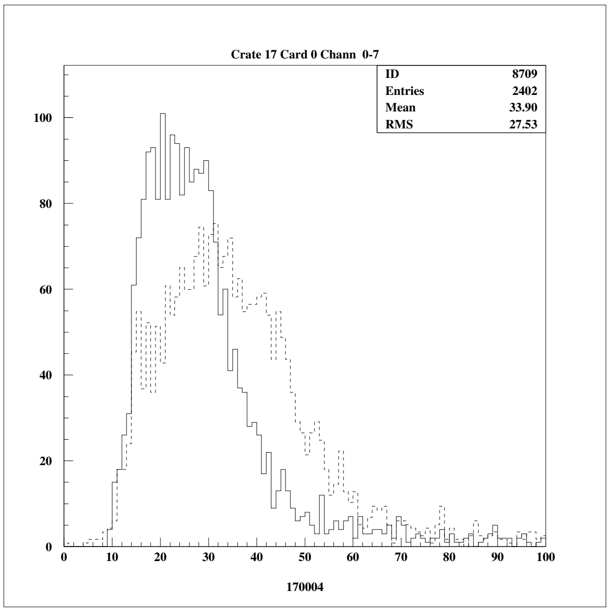

Figure 19: Charge Distribution with HV at Nominal and Nominal+100 V, for a channel in a crate whose nominal voltage is 1875 V (laser data).

well to limit the number of factors for which we need such calibration and correction.

There is an additional advantage to reducing the variation per CAZ to 2-2.5% from 3%. As we know, some channels (those in slots 0 and 15 as well as channels 0 and 31 on all slots) are more sensitive to readout pickup than others, and it is generally these channels which can lead to detector instabilities which force runs to stop, and which add artifacts in the data stream such as the 'rings of fire' which are events lighting up only these channels. One way to reduce the number of these events in the data stream as well as to enhance the detector's running stability is to raise the thresholds on these channels. Currently, raising the thresholds on these channels is a 3% hit in efficiency for each CAZ (actually, it may be somewhat more since these channels already have high thresholds). With a higher voltage setting, the penalty (as shown in Figures 24- 27) is reduced to roughly 2%. (Note that the chi-squareds shown in the fits are not meaningful because the spread in the occupancy was used rather than the uncertainty on the mean occupancy). With an increase of (say) 10% overall for a +100 V change, reducing the efficiency of some channels by 2-2.5% with the

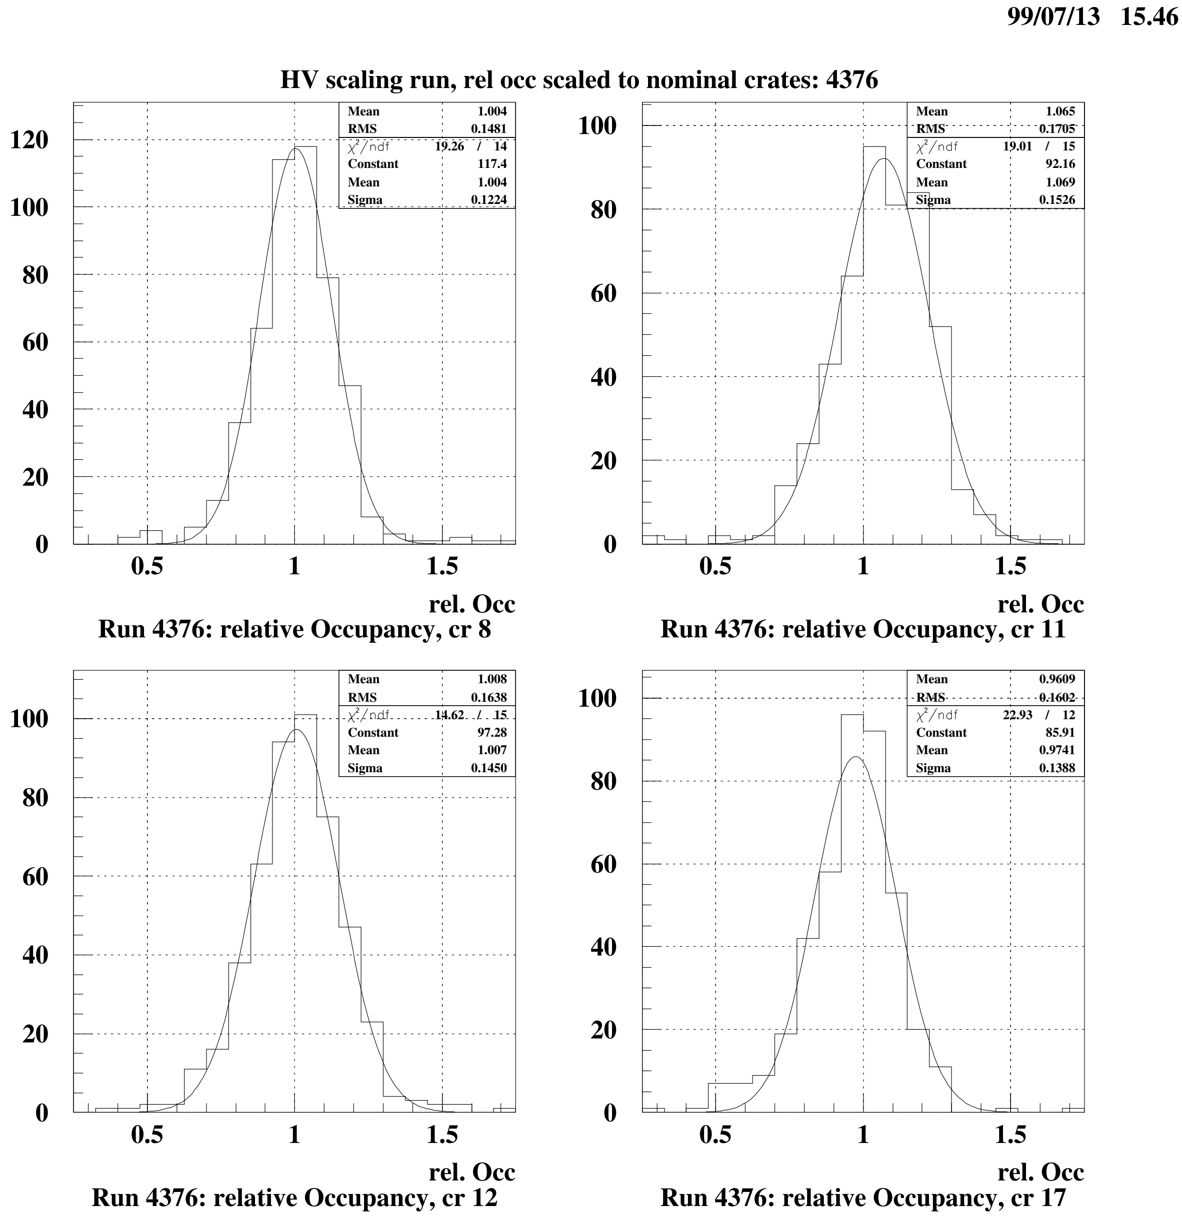

Figure 20: Occupancy Spread for Crates 8,11,12 and 17 at Their Nominal High Voltage Settings (laser data).

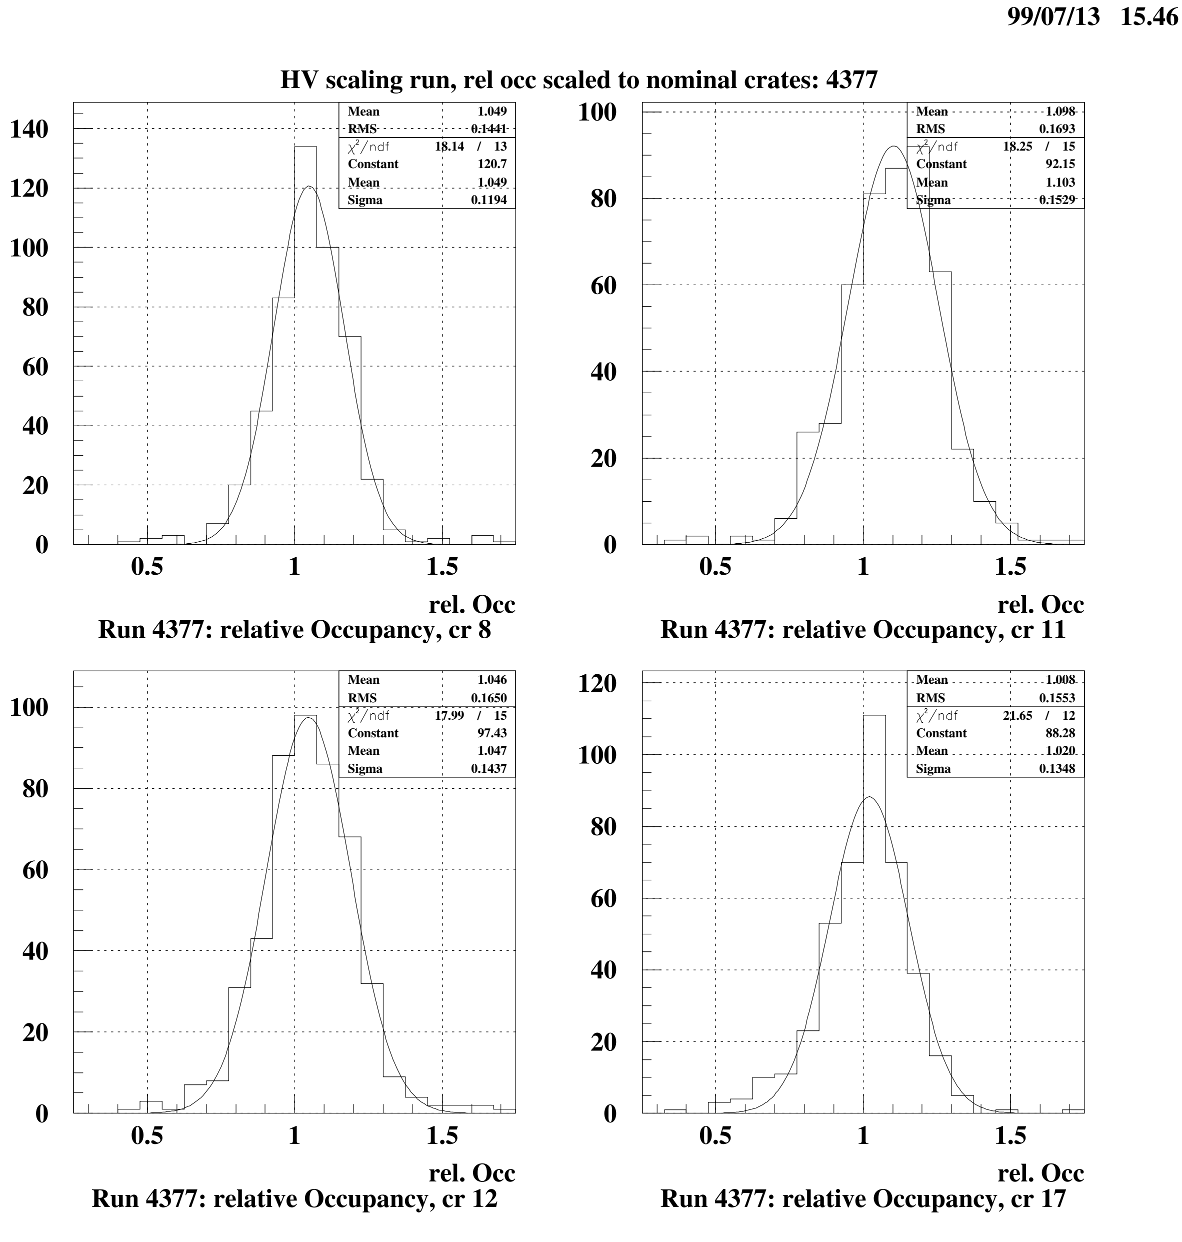

Figure 21: Occupancy Spread for Crates 8,11,12 and 17 at 40 V Above Their Nominal High Voltage Settings (laser data).

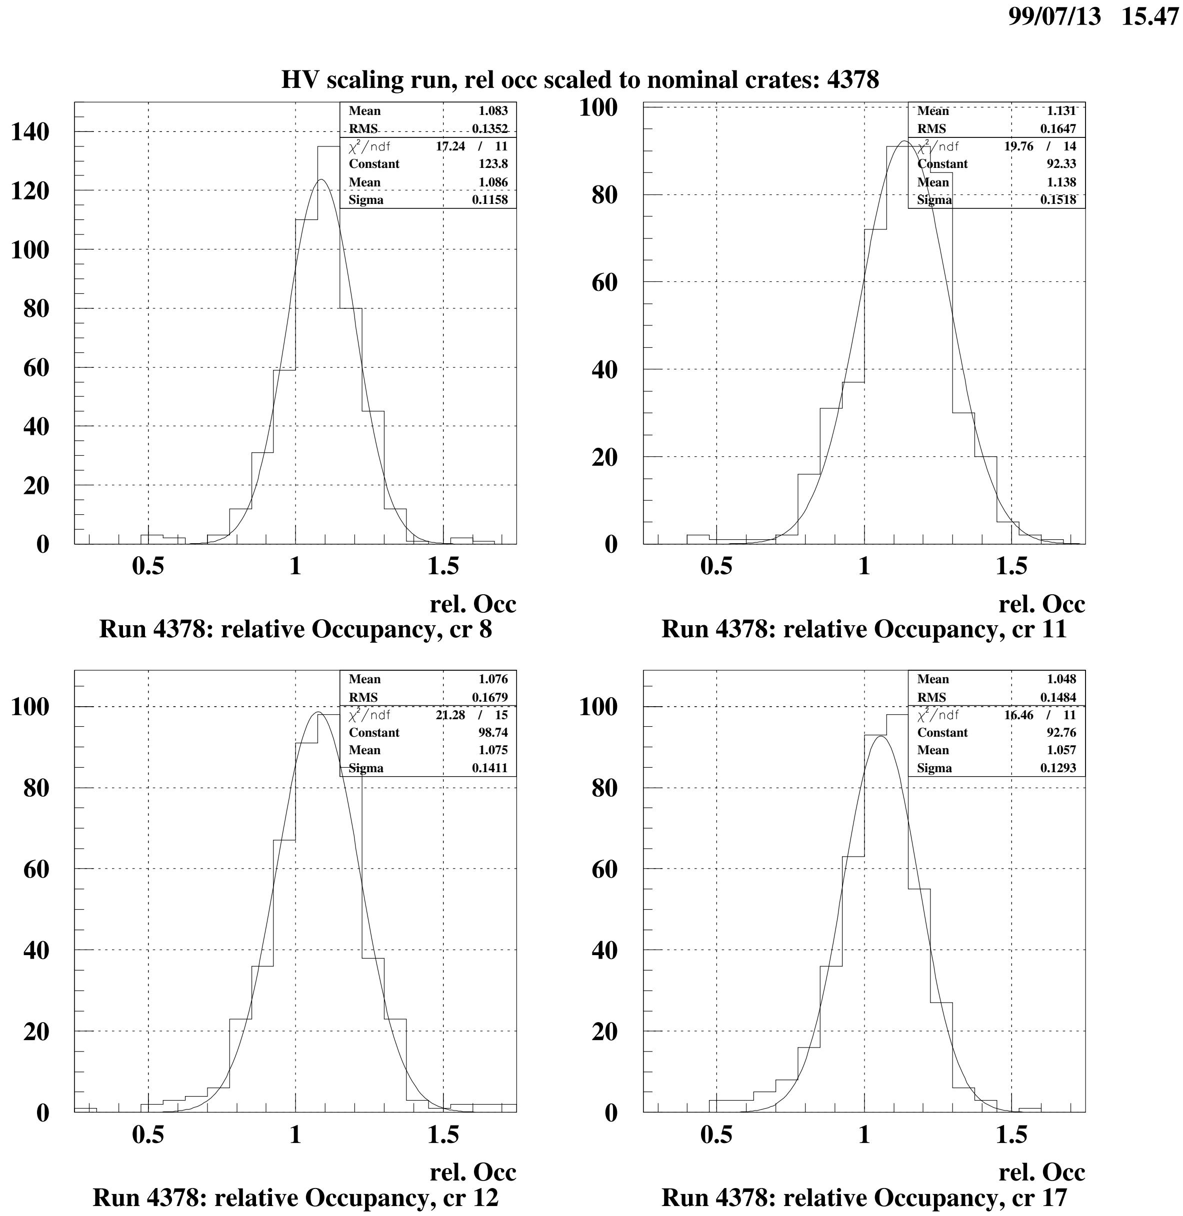

Figure 22: Occupancy Spread for Crates 8,11,12 and 17 at 80 V Above Their Nominal High Voltage Settings (laser data).

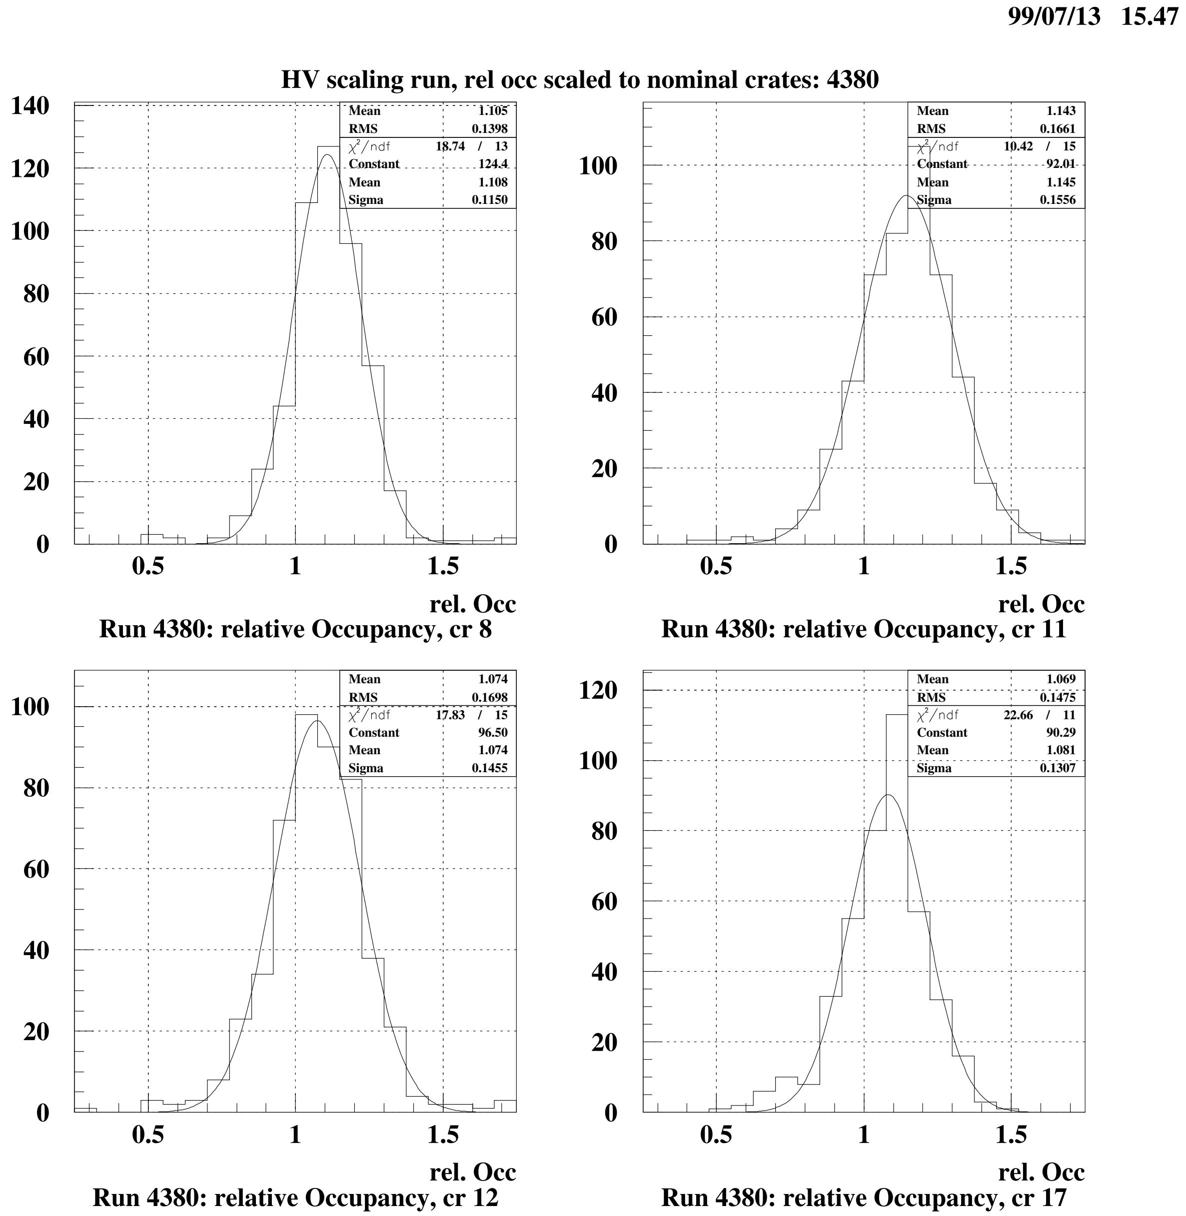

Figure 23: Occupancy Spread for Crates 8,11,12 and 17 at 100 V Above Their Nominal High Voltage Settings (Note that Crate 12 was still only 80 V above for these plots) (laser data).

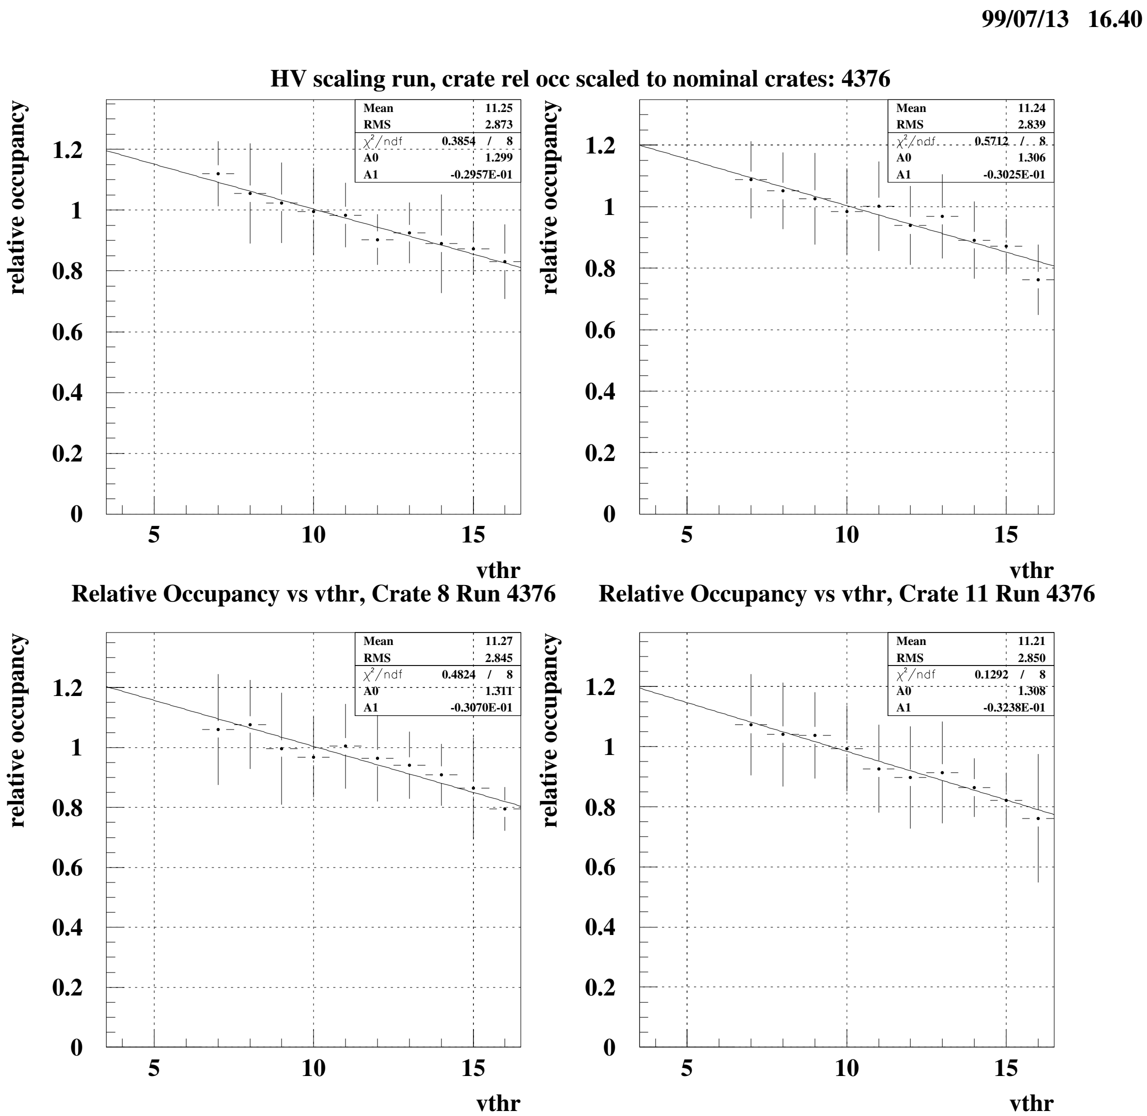

Figure 24: Occupancy vs. Threshold at Nominal HV for Crates 8,11,12, and 17 (laser data).

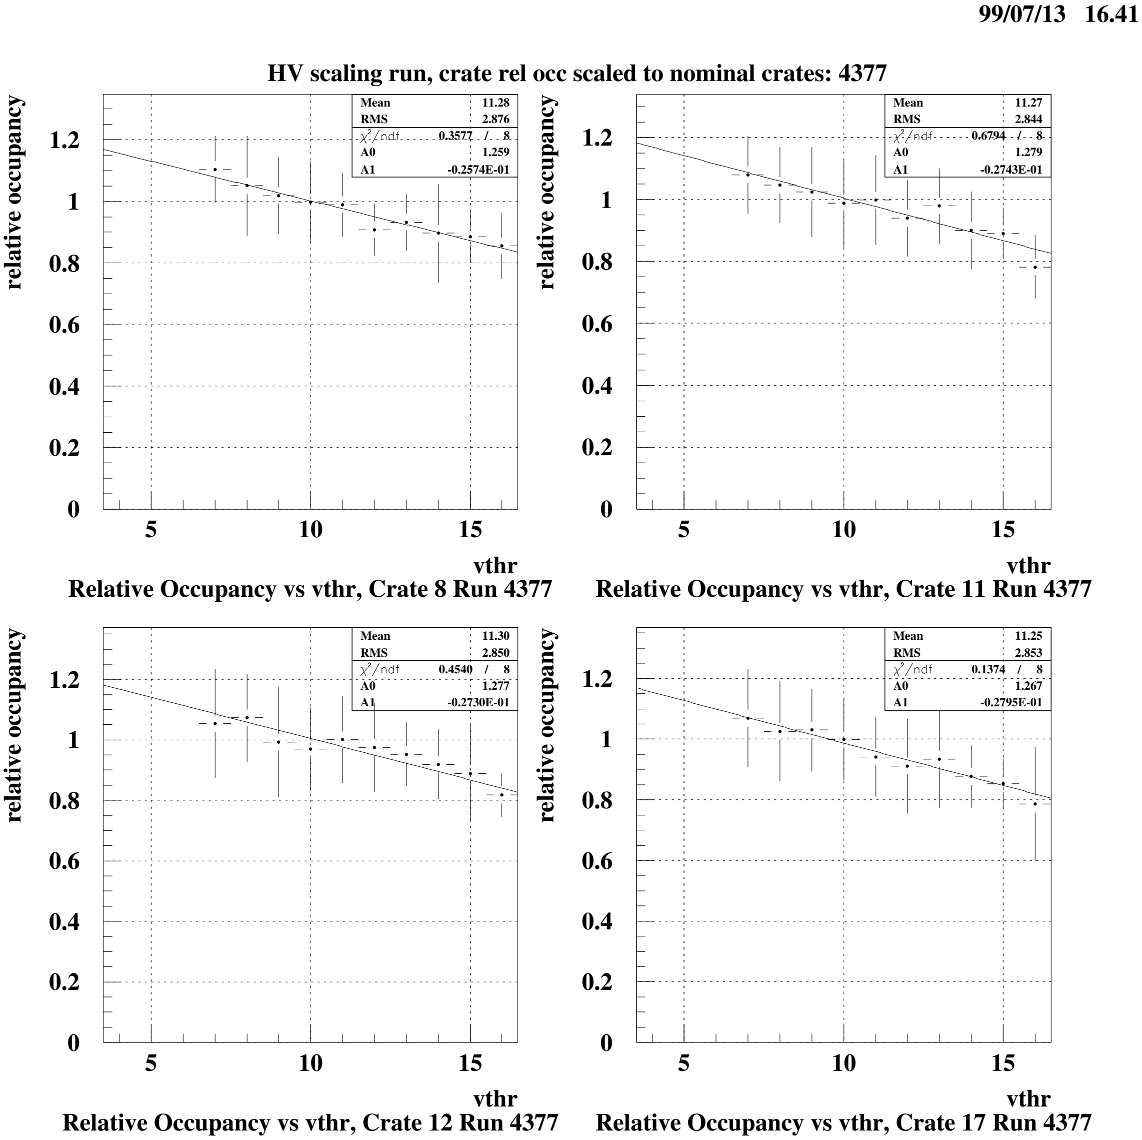

Figure 25: Occupancy vs. Threshold at 40 V Above Nominal HV for Crates 8,11,12, and 17 (laser data).

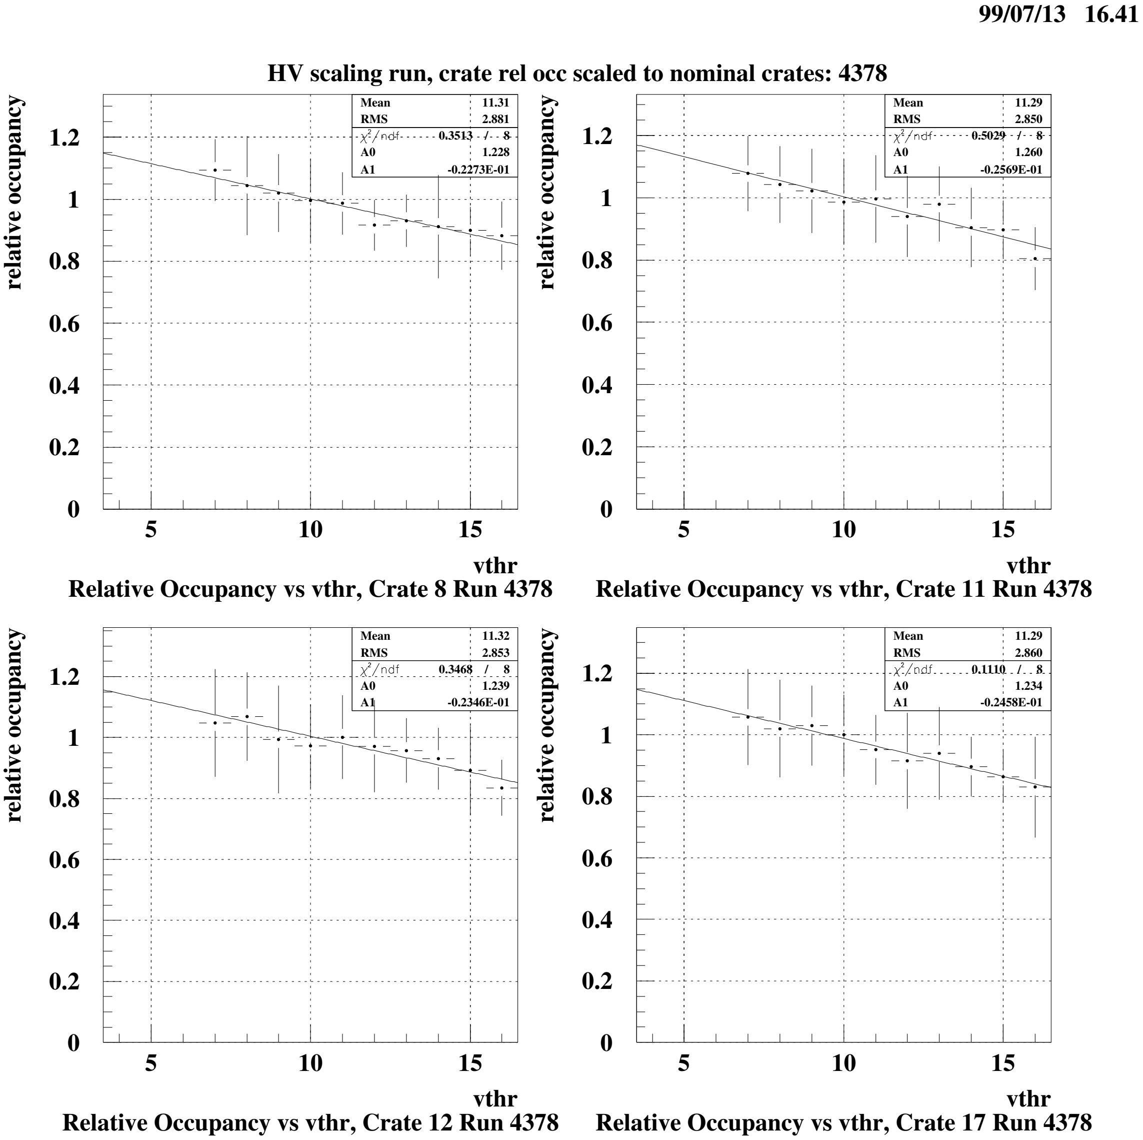

Figure 26: Occupancy vs. Threshold at 80 V Above Nominal HV for Crates 8,11,12, and 17 (laser data).

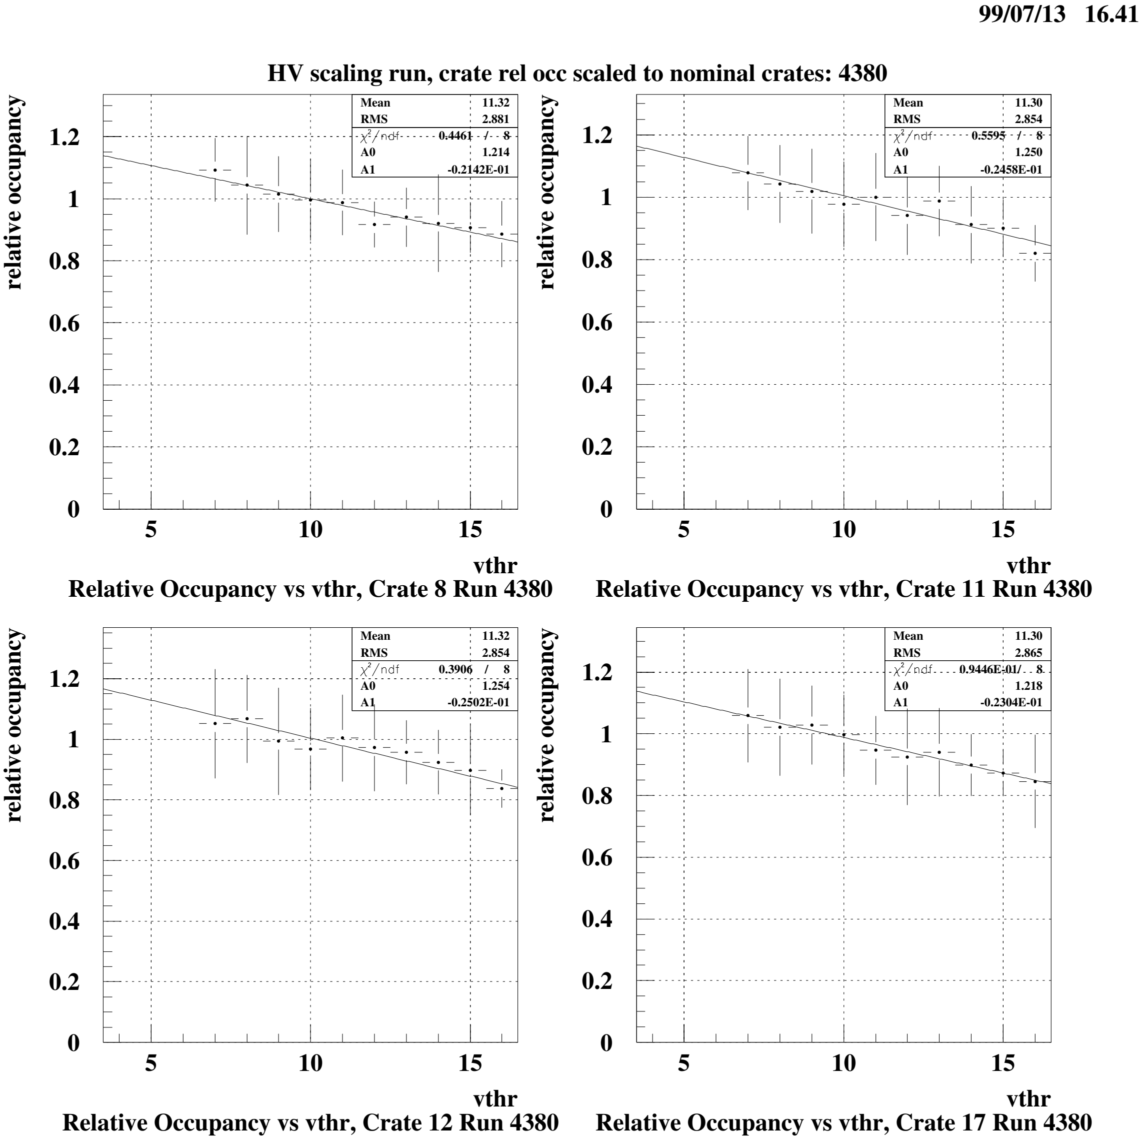

Figure 27: Occupancy vs. Threshold at 100 V Above Nominal HV for Crates 8,11,12, and 17 (laser data).

reward of a more stable detector and a cleaner data stream certainly seems like a worthwhile move.

5.3 The Effects of Raised High Voltage on PMT Noise

In addition to the obvious potential penalties for a global raise in the high voltage (higher breakdown rate, more sharkfins and flashers), there is the possibility that the noise rate in each PMT (for the same relative threshold) will go up, adding more noise hits into each event. The addition of noise will make reconstruction more difficult, as the ability to throw out out-of-time tubes in a given fit is compromised as the number of those tubes rises.

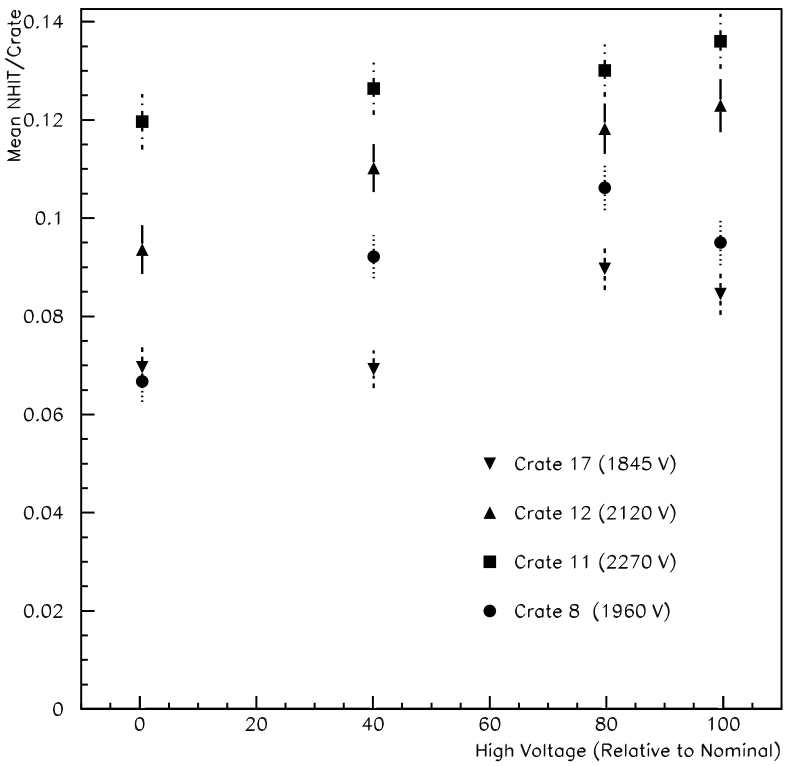

Figure 28: Mean NHIT for PULSE_GT Data for each crate as a function of high voltage increase relative to nominal.

At the current threshold settings, the mean NHIT of the PULSE_GT events (the best measure of the effects of the tube noise rate) sits near 2 tubes. To see how the high voltage affects this noise rate, we ran PULSE_GT triggers simultaneously with both the laser and ^{16}N source. Figure 28 shows the change in average NHIT for the crates whose voltage was raised, and Figure 29 compares the sum of those crates to the rest of the detector. Looking at the crates individually, we can use the mean occupancy of each channel as a measure of the tube noise rate. As can be seen, there is a small change with the raised high voltage, but it is not very significant. As can be seen in Figure 29, this increase translates into an additional 0.2 tubes/event were all the crates to be raised by +100 V.

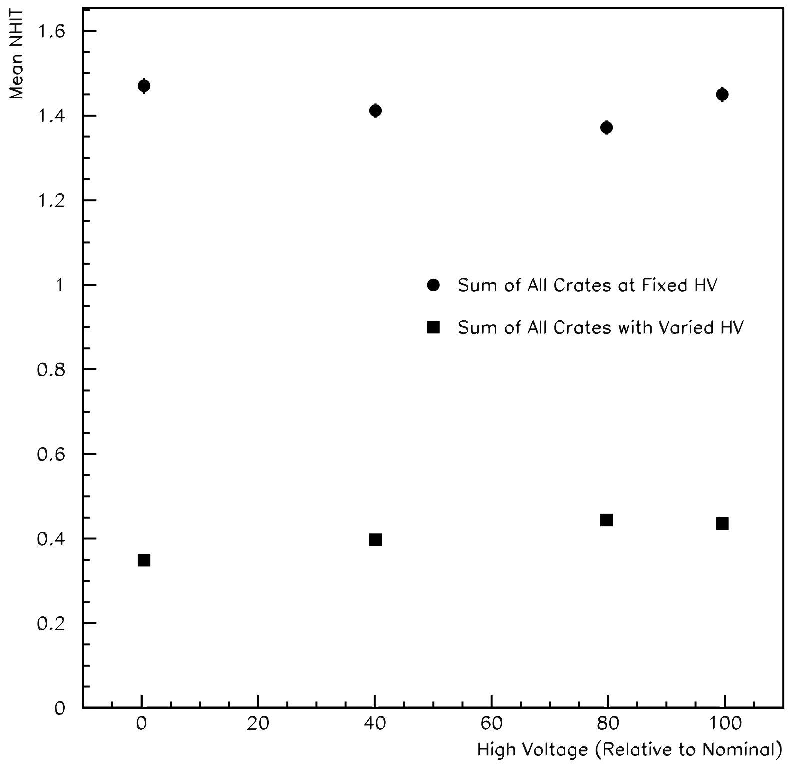

Figure 29: Comparison of Mean NHIT for PULSE_GT events for crates with changed voltages and the rest of the detector.

The numbers used in this Section are summarized in Table 2.

6 Conclusions

From the studies of the threshold scans, it looks fairly convincing that the loss in relative efficiency at the current crate voltage settings is about 3% per CAZ. In addition, there is some evidence that we could not run the detector with an average threshold lower than about 6 or 7 CAZ even in the absence of pickup from readout. With an average threshold of about 9.5 CAZ, we are $\sim$ 10% lower than where we could run were pickup not a problem.

Raising the high voltage by +100 V results in a gain of roughly 10% detector-wide, as well as a reduction in the slope of the efficiency vs. threshold curve to 2-2.5%, with no signficant increase in the number of noise tubes in each event. With these numbers, we gain not only through the increase of our energy scale, but in an increase in our energy resolution and in an opportunity to increase detector stability. In addition, there is some evidence that the gains crate-to-crate are not matched as well as they could be, and this may indicate that we should raise some of the crates more than others to flatten out their relative occupancies. Clearly, we must consider carefully any change, for not only does it carry the known risks of an increase in breakdown, or sharkfins, or flashers, but also the unknown risk that any global change in detector running conditions may bring.

Crate	+0 V			+40 V			+80 V			+100 V		
	Occ.	σ_{occ}	N_{pgt}	Occ.	σ_{occ}	N_{pgt}	Occ.	σ_{occ}	N_{pgt}	Occ.	σ_{occ}	N_{pgt}
8	1.004	0.1224	0.066	1.049	0.1194	0.092	1.083	0.1158	0.106	1.105	0.1150	0.094
11	1.065	0.1526	0.120	1.098	0.1529	0.124	1.131	0.1518	0.129	1.143	0.1556	0.135
12	1.008	0.1450	0.094	1.046	0.1437	0.109	1.076	0.1411	0.116	1.074	0.1455	0.120
17	0.961	0.1388	0.070	1.008	0.1348	0.069	1.048	0.1293	0.090	1.069	0.1307	0.085

Table 2: Summary of Results from HV Scaling Runs.