FINAL:
Data Cleaning: The Cluster Cut
SNO-STR-2000-001

Peter Wittich
University of Pennsylvania

21 October 1999

Abstract

The cluster cut takes advantage of crosstalk in the electronics to tag backgrounds from the SNO data. Effects of the cluster cut on physics data and efficiency for removing events with crosstalk such as flashers and bursters is discussed, based on the standard benchmarking data sets assembled by Steve Brice and the standard cluster cut available in SNOMAN.

1 Introduction

The primary goal of the SNO detector is to make a sensitive measurement of the solar neutrino flux from the sun. This measurement requires a relatively clean data sample of neutrino events, with an understanding of how much background is contaminating the data and how many physics events have been inadvertently eliminated in the data reduction process. In addition, the data reduction process might also have systematic effects in eliminating data of a given type or a given energy or N_{Hit} range. These systematic effects need to be understood. In order to make a good measurement of the solar neutrino flux and be confident in one's results, SNO must very clearly characterize any cuts applied globally to the data in order to understand and account for these biases. Global cuts are particularly dangerous as the non-global cuts made by individuals will be cross-checked in independent individual analyses, whereas most analyses will use the global cuts without question. The verification of the global cuts must occurr outside the normal analysis path.

For this reason, the data cleaning group's job is to come up with a set of standard cuts and quantify their effects on the data. Here, the cluster cut is described.

2 Crosstalk and the cluster cut

The cluster cut is a simple cut based on an imperfection in the electronics. When a channel registers a PMT signal that is much larger than a typical single photo-electron signal (equivalent to 15 - 25 photo–electrons), adjacent channels on the front end electronic cards (FECs) tend to fire in sympathy though their PMTs are not seeing light. This effect is called *crosstalk*. Physics events that are interesting for solar neutrino analyses typically do not have tubes with high charge in them, and hence no events with crosstalk. On the other hand, flashers, one of SNO's typical background events, exhibit huge signals in a single tube, with a pulse many tens of volts in amplitude at the FEC. As a result, many channels near the flashing channel fire at the same time. This crosstalk is one of the hallmarks of a flasher event. Crosstalk at this level was first examined with flashers, and the cluster cut was developed with flashers in mind.

The cluster cut is easily described: for each paddle card (crate) in the system, the total number of hits in an event is counted. Each paddle card (crate) with more than a predefined number of counts is flagged and called a "cluster". Events with clusters are tagged for exclusion or further inspection.

The cut is simple as it does not require any calibration to be applied to the data to work – problems with missing or incorrect calibrations cannot confuse this cut. On the other hand, if tubes are omitted from the data stream in some way[1], crosstalk can be masked out. In this case, channels that did fire are not read out, and the simple counting above fails. Additionally, channels that are not instrumented[2], but are still being read out can still see this pickup. Failing a smaller cluster size will tag more clusters, however, more physics data can be tagged as well.

In each cluster, there is a tube which caused the other tubes to fire. This tube will have abnormal charge[3], as the integrator is being presented with a huge pulse. The charge cut[4] can be combined with this cut to effectively determine which tube in a cluster is causing the crosstalk. Alternatively, the existence of a tube with abnormal charge could be used to extend the simple counting version of the charge cut to a more sophisticated 'cluster–charge' cut, allowing one to flag smaller clusters as of instrumental origin and extending the sensitive range of the cluster–like cuts.

An additional characteristic of the crosstalk channels is that they fire at the same time. A 'cluster–time' cut (also called 'cluster–density') cut was developed to take advantage of this correlation to tag clusters with fewer tubes but tight time correlation as suspicious[5]. However, simpler cuts are easier to understand

[1]For instance, they have been explicitly disabled by the data acquisition system (DAQ).
[2]I.e., Electronics channels that do not have photo-tubes attached.
[3]Abnormal in this case means a charge which is inconsistent with a single photo–electron event.
[4]The charge cut is described elsewhere in a similar document.
[5]At the time of this writing, the a description of the cluster-density cut is available in my

and their effects are easier to quantify, so we will try to stick with the simple cluster cut until we are forced to consider something with more muscle.

3 The Cluster Cut

3.1 Code available in SNOMAN

SNOMAN, in versions past 3.01, comes with a standard cluster cut. This cut has two criteria which an event must fulfill in order to pass[6]:

- In a given event, no single paddle card may contain *more* than N hits, where N is a tuneable parameter.
- In a given event, no single crate may contain *more* than M hits, where M is again a tuneable parameter.

In SNOMAN version 3.0188 (a development release of SNOMAN 3.02), the parameters above default to $N = 6$ and $M = 200$. The ranges of the cuts are determined by the quantities in question: a paddle card and a crate have eight and 512 channels, respectively. So the ranges on can investigate N and M are

- $1 \leq N \leq 7$
- $1 \leq M \leq 511$

Additionally, the cluster cut will tag orphans, as I discovered.

The first thing one can do is apply the cluster cut while varying the relevant parameters on different standard data sets. I will first investigate the paddle card based cut, and then the crate based cut. First, however, a discussion of the standard data sets is presented, along with what one might hope to achieve with the cluster cut with these data sets.

3.2 The data sets

The benchmarking data sets are a mixture of physics data, calibration data, and known junk. A complete list is presented below, with a brief discussion of each run.

- **Laser calibration runs:** Laser ball runs can provide single photo–electron events whose time distribution mimics that of a neutrino interaction. At higher intensities, many events will fail the cluster test. For the runs in question, the source position is fixed at $x = y = 0$ cm with a changing z co–ordinate.

part of the Penn analysis web page (`http://www.hep.upenn.edu/SNO/`).).

[6]Currently, a data reduction cut is considered *passed* if the event or set of events is *not* cut, i.e., it is not thrown away as undesirable.

- DC_laser_4085.zdab, $z = 0$ cm
- DC_laser_4086.zdab, $z = -550$ cm
- DC_laser_4092.zdab, $z = +400$ cm

- ^{16}N **calibration runs:** The ^{16}N source is currently our only calibration source for Cherenkhov light. Here too, the source position is fixed at $x = y = 0$ cm with a changing z co–ordinate.

  - DC_n16_4505.zdab, $z = -11$ cm
  - DC_n16_4512.zdab, $z = -581$ cm
  - DC_n16_4534.zdab, $z = +389$ cm

  This data is our only source of "pure" physics events. As the γ events from this source look like neutrino-induced cherenkhov electrons, any useful cut should not cut many of these events. The effect of moving the source close to the PMTs is a test for the cluster cut, as the spatial proximity of tubes that are hit will be mirrored in the proximity of channels that have data.

  Additionally, using the as yet undeployed pT source as a high-energy γ source would provide an additional test for this cut at the high end of the solar neutrino spectrum. Unfortunately, we don't have this data.

- **Various assorted crap:** These files contain "hand-assembled" collections of background events which we are hoping to eliminate in an automated fashion.

  - DC_golden_flashers_4025_5112.zdab , a set of golden flashers as collected by Mark Neubauer. Both this data set and the following flasher burst should be effectively cut with the cluster cut, as it was initially developed with this in mind.
  - DC_flash_burst_3887.zdab , a (rock burst–induced?) burst of flashers. These flashers should be cut by the cluster cut.
  - DC_rings_of_fire_4782.zdab , a collection of so–called "ring of fire" events, where the outermost channels in crate–space in a given crate fire, as collected by Josh Klein. Most of the ones I have looked at by hand seem to be more 'slots of fire' with just one slot triggering, mostly slot 15 or zero. There is also some flasher contamination, and at least one ring event that might be a decay e^-.

  This data one expects to be cut by the cluster cut. However, one must be concerned about self-selection effects. Depending on how these data were assembled, it is possible that we will be misled. If the data were assembled with a cluster cut, it will not be surprising that a cluster cut tags them.

- ν **data:** N_{Hit}-triggered data, meant to mimic a standard neutrino run. This data has been studied intensely, and its many peculiarities are well-known. Given what we learn from the above, one might hope to extract the above bad data from the ν data sets.
  - DC_nu_run_3887.zdab
  - DC_nu_run_5125.zdab

  Calling this neutrino data doesn't imply there are only neutrinos in it. In fact, the amount of garbage in this run hopefully exceeds to a large fraction the amount of garbage in future stable production running. Rather, the neutrino runs are just data taken in the same configuration as we will be taking our production data in the future.

3.3 Cut tightly or loosely?

The goal of first pass data cleaning (and hence the cluster cut, as a part of first pass data cleaning) is to eliminate obvious junk from the data stream. The cuts at this level are meant to be *loose* cuts, i.e., non-controversial cuts that are designed to preserve signal at the expense of letting background through. More sophisticated cuts, such as cutting on the results of reconstruction or an event classifier, are expected to be able to remove the remainder of the background that has been passed through this level of cuts, and therefore it is not in the goal of first pass to eliminate *all* junk from the data, but rather to remove most of the junk while preserving signal.

4 Paddle Card cluster cut

Figure 1 shows the results of applying the cluster cut to the benchmarking data with $2 \leq N \leq 7$. In this section, the crate cluster cut is set to the impossible value of $M = 600$ in order to insure that we are just looking at the paddle card test. The figure shows the amount of data cut in each of the benchmarking sets as a function of the cut parameter N. Additionally, tables of the numerical data are appended in Appendix A, and a distribution of maximum number of counts on a paddle card for a given event can be seen in Figure 12 for each data set.

In the following paragraphs, a detailed breakdown of the cluster cut behavior can be seen for different runs.

4.1 Golden Flashers

The golden flasher data (figure 2) holds no surprises. The cluster cut was developed for flashers, and hence one is strongly tempted to tune the cluster parameter to this cut and investigate what the effects on other data types then result. The cluster cut effectively cuts all of this data for low values of N,

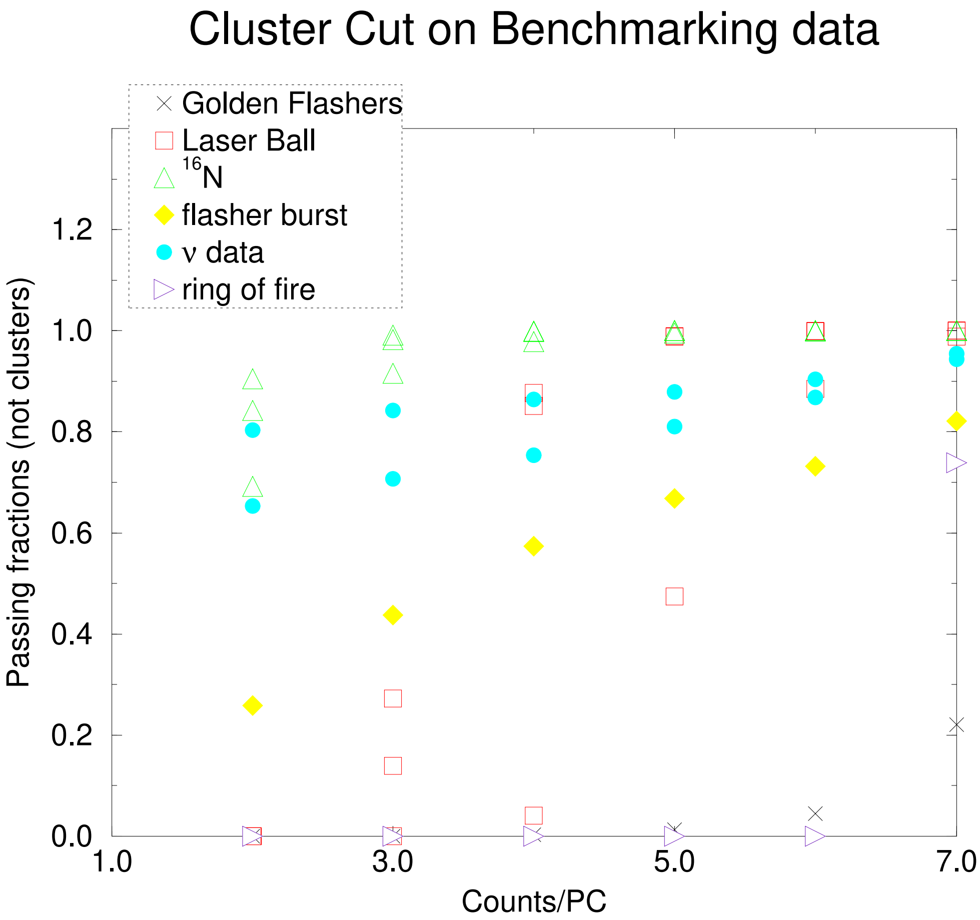

Figure 1: Effect of the cluster cut on the set of standard SNO data cleaning benchmarks. At values of $N = 6$, the cluster cut is highly effective at differentiating flashers from other types of data. (There are multiple entries for a given run type at each data point as there were multiple runs with differing run conditions.) Note that few events from the flasher burst are tagged.

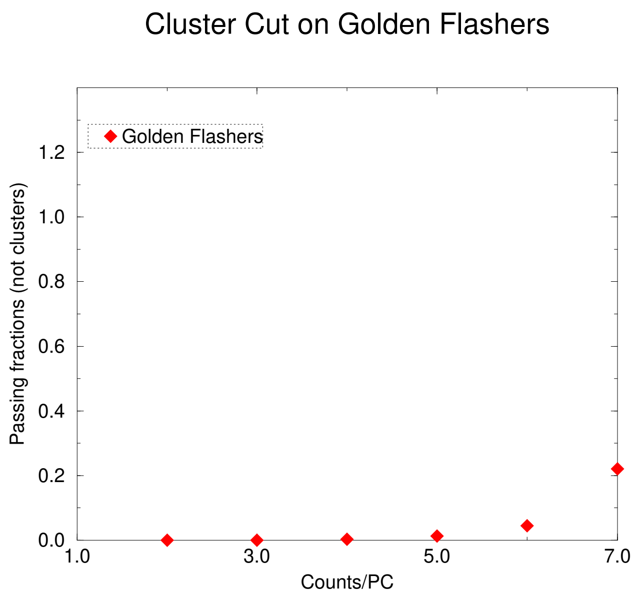

Figure 2: Effect of the cluster cut on golden flashers. For $N \leq 6$, more than 95% of data is flagged.

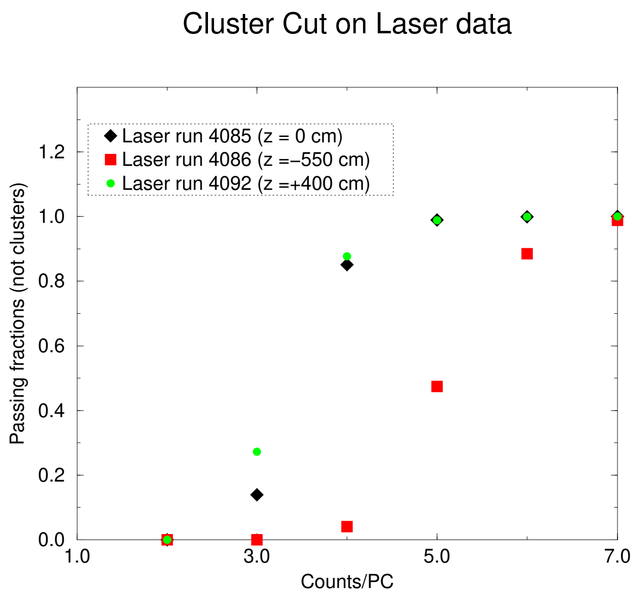

Figure 3: Effect of the cluster cut on laser data. The high intensity of the laser run and the anisotropy for the off-center runs ($\langle\langle N_{\text{Hit}}\rangle \simeq 670$ in runs 4085, 4086 and $\langle N_{\text{Hit}}\rangle \simeq 450$ in run 4092) make a large fraction of these events get flagged by the cluster cut. None the less, at $N = 6$, fewer than 12% of events in run 4086 are flagged.

and even at the highest value $N = 7$, only 20% are missed. Since we want to conserve signal, we will be lenient and focus on the value $N = 6$, where about 5% of the golden flashers are not tagged by this cut.

4.2 Laser

For the laser data, a non-zero fraction of the events do get flagged by the cluster cut, as can be seen in Figure 3. The high intensity of the laser run and the anisotropy due to the off-center source position ($\langle\langle N_{\text{Hit}}\rangle \simeq 670$ in runs 4085, 4086 and $\langle N_{\text{Hit}}\rangle \simeq 450$ in run 4092) are to blame for 15-85% of the data being tagged. For neutrino data, the detector occupancies are in a much lower range: $N_{\text{Hit}} \leq 200$ is a conservative range for neutrino data. With occupancies three times higher, the ability of this data to be used for the cluster cut is misleading. Additionally, the $z \neq 0$ source positions mean concentrated light at the upper and lower ends of the detector. This *is* the sort of data we need to look at for

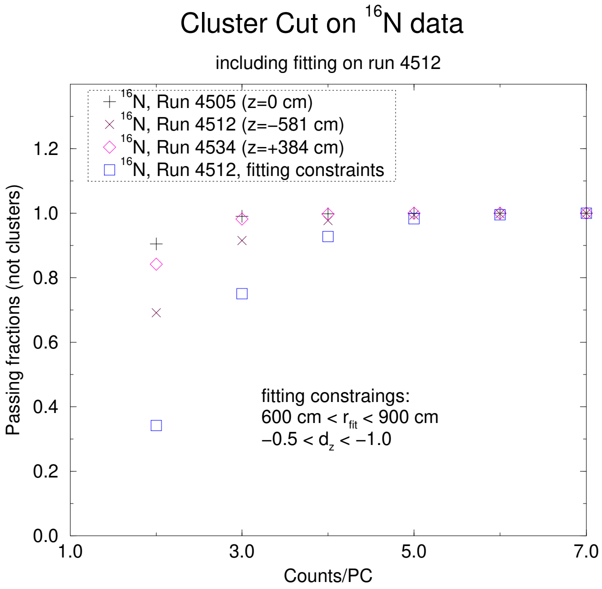

Figure 4: Effect of the cluster cut on ^{16}N data. For $N \geq 3$, less than 9 % of calibration data is flagged. At $N \geq 6$, less than 1% is flagged. When selecting data from run 4512 going radially outwards ("fitting constraints"), the worst-case scenario is tested.

the cluster cut, as e^- in water are directional. However, it would be better to look at data in more neutrino-like energy ranges.

None the less, for $N = 5$, runs 4085 with $z = 0$ cm and 4092 with $z = +400$ cm, less than 2% of all data is flagged as a cluster. For run 4086, with $z = -550$ cm, more than 50% of all events are marked as clusters. This is due to the increase in solid angle subtended from the source by the tubes with the extreme z position. For $N = 6$, the solid angle effect is mitigated, and 12% of this laser run is tagged.

4.3 ^{16}N

The ^{16}N data, as shown in figure 4, is the most crucial test for this cut and the other cuts in the data cleaning. As a sample of roughly 5 MeV electrons, this looks a lot like the solar neutrino data we are trying to preserve. With the value $N = 6$ we got from the golden flasher data, we see that here we get roughly 0.1% cut in the worst scenario (run 4512, corresponding to a source position of $z = -581$ cm). This seems to be acceptable.

While run 4512 has a large z source position, a better test would be to reconstruct the data from this run and see if we can't grab the γ rays that are going downwards (i.e., radially outwards from center). These γs are the ones we really want to look at, as the Compton electrons from them are generated close to and pointing at the PSUP and hence have a tight light pattern on the PSUP, illuminating adjacent photo tubes. As explained in the description of the runs (section 3.2), the spatial proximity of tubes that are hit will be mirrored in the proximity of channels that have data, causing more clusters. To grab this data set, I reconstructed the data using the standard SNOMAN FTT fitter[7]. For all events that reconstructed between 600 cm $\leq r_{\text{fit}} \leq$ 900 cm, I check the z component of the reconstructed direction cosine, d_z. For $-0.5 \leq d_z \leq -1.0$, we declare these events radially outward looking data and can do the same test, looking at the behavior of the cluster cut on this data set[8]. After these cuts, we are left with roughly 14000 events out of a total of 123000. (DO WE KNOW IF THE SOURCE EMITS ISOTROPICALLY?) The fit results are summarized in Figure 5. The overflow on the distribution of z fit positions shows that the r_{fit} cut is not grabbing some pathelogical data set, but events at the bottom of the detector. Hand-scanning a small sample of the fit results agrees with this.

The cluster cut now grabs some more data than it did before (see the data points in figure 4, and compare tables 5 and 6). However, even with this stringent test, at $N = 6$, $\frac{4}{10}$% of the data is cut. As expected, this is more than is cut without the fitting constraints, where about 0.1 percent were tagged. None the less, the results are acceptable.

4.4 Flasher Burst

The flasher burst data shows a small amount of data that is cut via cluster cut. This is suprising if one assumes that the flasher burst is just a collection of flashers in rapid succession. It turns out that this is not the case. A quick look into the viewer shows that a lot of the data are follower events after a normal flasher, separated by times with $\delta t \sim t_{\text{lockout}}$, where t_{lockout} is the minimum time between events as determined by the electronics . Slicing the data into four parts shows that the data changes over time (see table 1). Figure 7 shows that the bulk of the data in this run occur in the last quarter of time (all but the last 10%, actually), and that in these times, the δt distribution goes down sharply. Figure 8 shows the cumulative distribution of δt, the time to the previous event, for cluster and non-cluster data from this data set. The non-cluster distribution reaches 90% of the data with $\delta t \simeq 3\mu s$, while the distribution of clusters barely reaches 70% within $10\mu s$. The cluster cut therefore will catch some of the flasher burst data, but this job is clearly better left to a real burst cut.

[7] Calibration is done with ECA only, as no PCA is available in the data base.

[8] A better test might have been 600 cm $\leq z \leq$ 900 cm, to get around data reconstructing at largely wrong areas. It turns out that both results are the same within statistics - very few events are at large radii in the wrong position

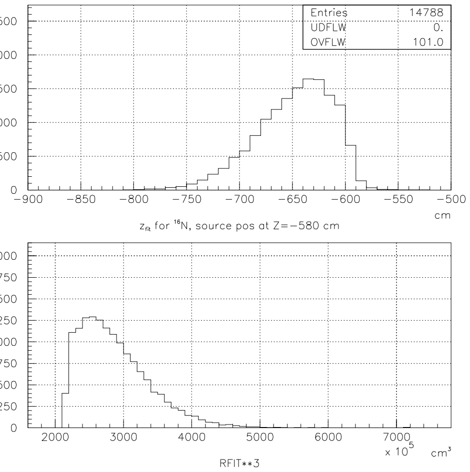

Figure 5: Reconstruction of ^{16}N with constraints as described in section 4.3. The constraint on the fit is on r_{fit}, but checking the z postion shows the bulk of the events in the desired region. Recall $z_{\mathrm{source}} = -581$ cm for this ^{16}N run.

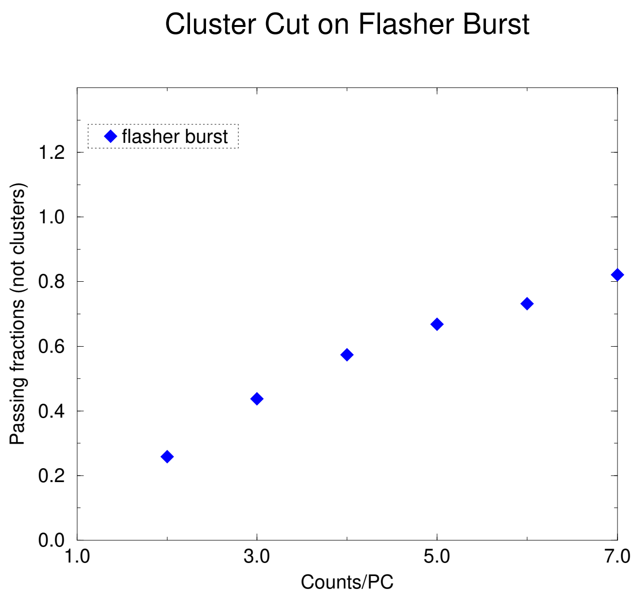

Figure 6: Effect of the cluster cut on flasher burst data. Suprisingly little data is cut, if you assume this is just a succession of flashers.

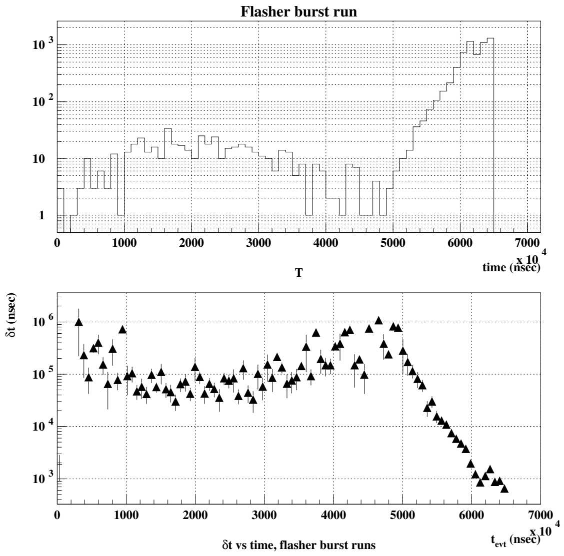

Figure 7: Relative time distribution in flasher burst data.

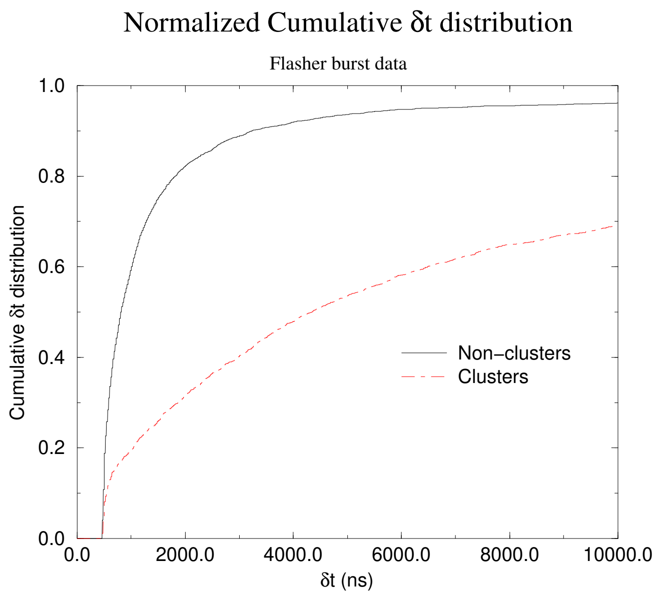

Figure 8: Cumulative integrals of δt, the time to the previous event, for cluster and non-cluster data in the flasher burst data set. The distribution for non-clusters has a much larger fraction at lower values of δt, indicating that it is mostly retriggers that do not get tagged by the cluster cut.

File part	Fraction tagged
First quarter	70%
Second quarter	23%
Third quarter	25%
Fourth quarter	14%

Table 1: Flasher burst data in four parts

A different question is whether these sorts of followers only exist in flasher burst or if they exist in flashers themselves, albeit at a lower level. It would be interesting to do a study of time correlation between flasher events and those that follow immediately after to see if we can see this effect in normal flashers.

4.5 Rings of Fire

The 'rings of fire' data set is clearly something that the cluster cut is ideal for and the cut bears this out. Figure 9 shows for $N \leq 6$, every event is tagged. This is largely self-selection (the data set was assembled with a cluster cut), so using this data set for a cluster cut is mostly meaningless.

4.6 Other data sets

The cluster cut suffers a bit from its own success: a lot of the data sets used in benchmarking were, in fact, developed with a cluster cut and hence are hard to use for benchmarking the cluster cut. Therefore, we need other data sets to at least convince ourselves that the cluster cut is not crazy. Flashers are the obvious choice, since the cut is used mostly as a flasher cut. To this end, two sets of flashers were grabbed from the first three subruns of run 5377 using the AMB and Q vs T cuts available in SNOMAN. The data was then hand-scanned and seen to be made up mostly of flashers. Figure 10 shows the distribution of maximum number of counts per event on a paddle card from these two data sets, and one can see that the success of the cluster cut on the golden flashers is not just a selection effect.

5 Crate cluster cut

A similar set of analyses is performed here for the crate cluster cut. The paddle card cut is turned off by requiring the unphysical $N = 10$. M is varied with $50 \leq M \leq 450$, in steps of 50. The results of this can be seen in Figure 11. All data is presented together, except the laser data. Figure 11 is rather remarkable

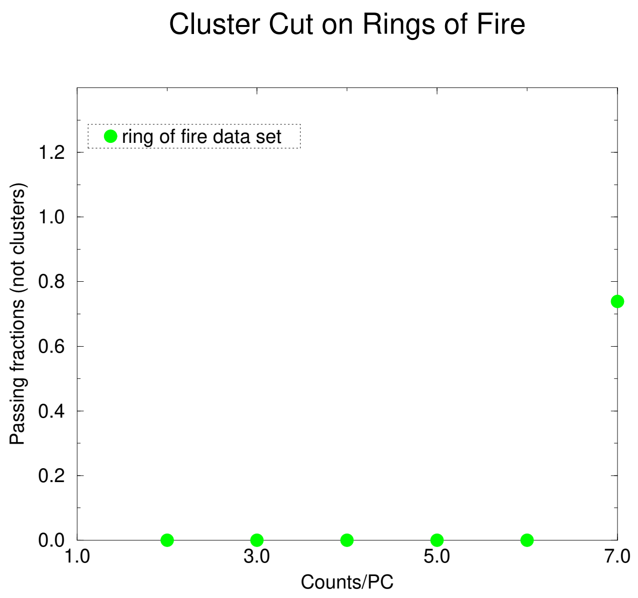

Figure 9: Effect of the cluster cut on rings of fire data. For $N \leq 6$, 100% of the data is flagged.

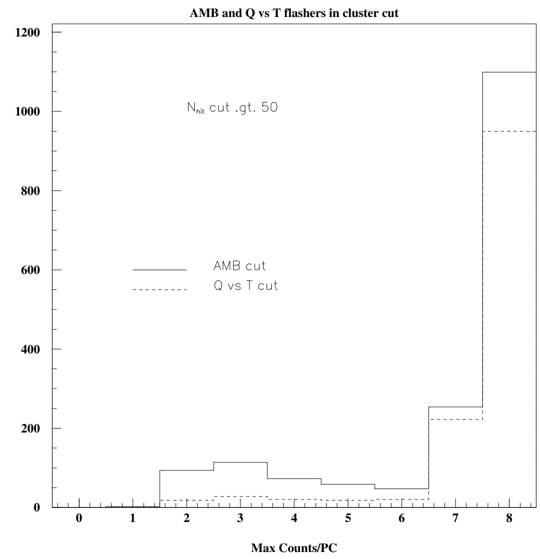

Figure 10: Distribtion of maximum number of counts on a paddle card for flasher data sets from run 5377, snagged with the AMB and Q vs T cuts. An additional $N_{\rm Hit} > 50$ cut was applied. At $N \geq 6$, the bulk of the data is grabbed.

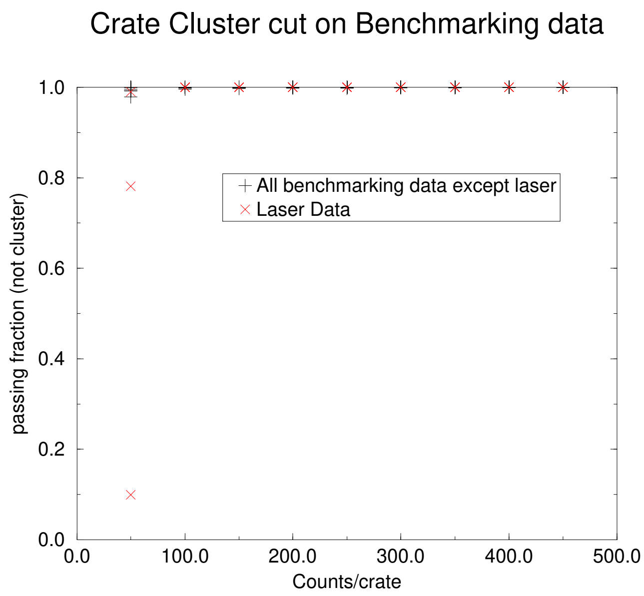

Figure 11: Crate cluster cut on all data. Only flasher burst and laser data are shown seperately. All other data sets do not get flagged by the crate cluster cut.

in that varying the parameter makes no difference for $M \geq 100$. Even at lower values of M, only the laser runs (with $\langle N_{\mathrm{Hit}} \rangle \simeq 450 - 670$) are tagged.

This cut does not mark any data which would otherwise be used for solar neutrino analyses. It could be useful to find bursters or dry end breakdown, though specific burst cuts have been developed for those events. The crate cluster cut is, when compared to the other cuts being developed, not useful and probably just a historic artifact of the earlier days of data cleaning.

6 Conclusions and Recommendations

The paddle part of the cluster cut is an effective cut for removing simple flashers and ring of fire events. Given the default value of $N = 6$ for this cut, a good compromise of preserving signal while tagging garbage is achieved[9]. Even for γ data radially outward at the periphery of the detector, less than one percent of the data is tagged with the paddle card cluster cut.

[9] A skeptic might ask, since I end up with the same number I started out with, why did it take me 20-odd pages to get there?

I also recommend that the crate cluster cut be removed from the cluster code in SNOMAN, or the default value set such that the test cannot inadvertently flag an event. For solar neutrino analysis, this cut does not appear to be useful. It might be useful for other analyses, such as a muon or atmospheric analysis, but it's not clear to me how this would be the case.

An additional recommendation is that the cluster cut *not* fail orphans. Ideally, the cluster cut would never be passed an orphan event, and this discussion would be moot. Applying a cluster cut to an orphan event is non-sensical, and an ideal solution might be a return value of "not applicable" when presented nonsense data. There is little harm in tagging orphans as clusters, except in the possible danger of misestimating your contaimination of flashers or bursters when you assume that all tagged events are in these catagores. Orphans should be tagged in a specific orphan cut.

Finally, an important caveat is that the cluster cut as currently implemented is paddle-card based. It's not clear to me that we know that crosstalk occurs on the paddle card or on the daughter board - it might well occur in other places, such as in the PMTIC or on the mother board, where the eight channel boundries are not as pronounced, or even in the cables. A better cut (though a bit more complicated programatically) would be to strictly count the number of adjacent channels that fired, ignoring paddle card boundries. While I am confident it would be the same most of the time, an inspection by hand of the events that fail the cluster cut with $N = 6$ show that a lot of them would have been caught in a method that just looked for adjacent channels, rather than channels on a paddle card. The algorithm probably should be rewritten to include this update.

A Complete results for the cluster cut

For completeness, the data displayed graphically in figures 1 and company are displayed below in tables 2 and following. Also, figures a and b in 12 show the distribution of the maximum number of counts on a paddle card for a given event. The proposed cut value is at $N = 6$.

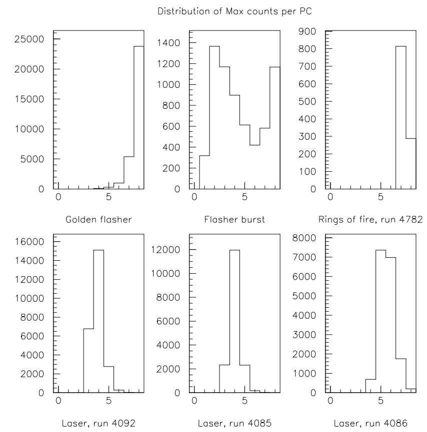

(a)

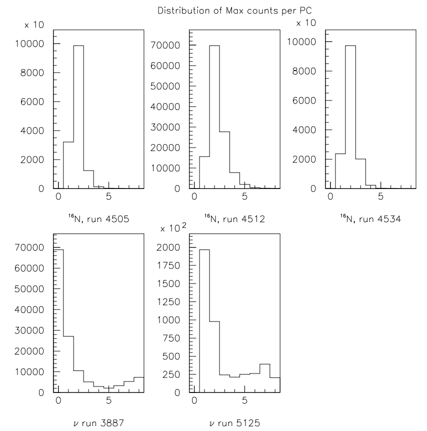

(b)

Figure 12: Distribution of maximum number of counts on a paddle card for a given event. Note that the ν run 3887 has lots of $N_{\rm Hit} = 0$ events. Note also that the maximum for the ^{16}N is *not* at 1 but at 2, and that for run ^{16}N 4512 ($z = -581$ cm) there are more counts at $N = 3$ than $N = 1$.

Events Passed	Events Failed	Passing Fraction	Cluster Window Size	Run information
0	30526	0.000	2	DC_golden_flashers_4025_5112.zdab
2	30524	0.000	3	
99	30427	0.003	4	
390	30136	0.013	5	
1379	29147	0.045	6	
6744	23782	0.221	7	

Table 2: Complete results for the golden flasher data.

Events Passed	Events Failed	Passing Fraction	Cluster Window Size	Run information
0	16795	0.000	2	DC_laser_4085.zdab
2332	14463	0.139	3	
14286	2509	0.851	4	
16606	189	0.989	5	
16775	20	0.999	6	
16793	2	1.000	7	
0	17000	0.000	2	DC_laser_4086.zdab
4	16996	0.000	3	
701	16299	0.041	4	
8065	8935	0.474	5	
15045	1955	0.885	6	
16803	197	0.988	7	
4	24936	0.000	2	DC_laser_4092.zdab
6775	18165	0.272	3	
21877	3063	0.877	4	
24653	287	0.988	5	
24924	16	0.999	6	
24940	0	1.000	7	

Table 3: Complete results for the laser data.

Events Passed	Events Failed	Passing Fraction	Cluster Window Size	Run information
1684	4848	0.258	2	DC_flash_burst_3887.zdab
2853	3679	0.437	3	
3750	2782	0.574	4	
4363	2169	0.668	5	
4783	1749	0.732	6	
5365	1167	0.821	7	

Table 4: Complete results for the flasher burst data.

Events Passed	Events Failed	Passing Fraction	Cluster Window Size	Run information
130722	13784	0.905	2	DC_n16_4505.zdab
143146	1360	0.991	3	
144327	179	0.999	4	
144487	19	1.000	5	
144504	2	1.000	6	
144505	1	1.000	7	
85286	38048	0.692	2	DC_n16_4512.zdab
112917	10417	0.916	3	
120667	2667	0.978	4	
122710	624	0.995	5	
123188	146	0.999	6	
123304	30	1.000	7	
120827	22665	0.842	2	DC_n16_4534.zdab
140869	2623	0.982	3	
143179	313	0.998	4	
143451	41	1.000	5	
143478	14	1.000	6	
143483	9	1.000	7	

Table 5: Complete results for the ^{16}N data.

Events Passed	Events Failed	Passing Fraction	Cluster Window Size	Run information
5059	9729	0.342	2	DC_n16_4512.zdab, with fitting constraints. See section 4.1
11111	3677	0.751	3	
13725	1063	0.928	4	
14537	251	0.983	5	
14735	53	0.996	6	
14781	7	1.000	7	

Table 6: Complete results for the ^{16}N, run 4512, with fitting constraints.

Events Passed	Events Failed	Passing Fraction	Cluster Window Size	Run information
106497	26050	0.803	2	DC_nu_run_3887.zdab
111602	20945	0.842	3	
114465	18082	0.864	4	
116565	15982	0.879	5	
119820	12727	0.904	6	
125169	7378	0.944	7	
294315	156351	0.653	2	DC_nu_run_5125.zdab
318521	132145	0.707	3	
339693	110973	0.754	4	
364893	85773	0.810	5	
391145	59521	0.868	6	
430269	20397	0.955	7	

Table 7: Complete results for the ν data.

Events Passed	Events Failed	Passing Fraction	Cluster Window Size	Run information
0	1102	0.000	2	DC_rings_of_fire_4782.zdab
0	1102	0.000	3	
0	1102	0.000	4	
0	1102	0.000	5	
0	1102	0.000	6	
814	288	0.739	7	

Table 8: Complete results for the "ring of fire" data.