Grid Fitter for Reducing Tails in Spatial Distributions

M.E. Moorhead
SNO-STR-95-042

October 23, 1995

1 Motivation

There is a growing number of fitters whose aim is to reconstruct the vertex of any given event with as small an error as possible. It is especially important to reduce the long tail in the spatial distributions of these fitted vertices, which can cause much more frequent background events to 'spill over' from external regions (PMTs, H_2O and acrylic vessel) into the relatively background-free D_2O . In general terms, the motivation for reducing these tails is provided by the pursuit of three separate goals listed below in rough order of increasing difficulty.

1. To minimize the number of PMT β-γ's which reconstruct inside the D_2O so that they do not determine the threshold for analysing neutrino signals, i.e. establish a 'neutrino window', free from PMT β-γ's, in which the threshold is determined by the irreducible background from D_2O β-γ's. Current wisdom suggests that this should be easy but the huge ratio of PMT β-γ decay rates to D_2O β-γ decay rates ($\sim 10^6$) is a concern.

2. To minimize the number of PMT, acrylic vessel (AV) and H_2O β-γ's which reconstruct inside a central volume of D_2O of approximately 4m radius and within an Nhits range of about 30 $\leq$Nhits$<$ 40. This 'Th/U monitoring window' is interesting because it may be possible to measure the ^{214}Bi and ^{208}Tl concentrations in the D_2O *independently* by using the same hit pattern recognition techniques that are being developed for NC/CC separation. This is because above 30 Nhits ^{214}Bi events are mostly (90% of the time)

single electron events (CC-like) whereas ^{208}Tl events are always one 2.6 MeV gamma + an electron (1.8 MeV endpoint) + one or more lower energy gammas (NC-like). As a prerequisite for this method to have any chance of being successful the Th/U monitoring window must be relatively free from events other than D_2O β-γ's.

3. Another 'Th/U monitoring window' has been proposed by B. Frati which is similar to the one above except that its enclosed volume (radius) would be smaller and the Nhits threshold lower (down to about 15 hits). The reason for looking this low in Nhits is that below about 25 Nhits the spectral shape of ^{214}Bi and ^{208}Tl events begin to differ and one can use this difference to determine their concentrations independently. Thus there are two proposed Th/U monitoring windows and in order to differentiate them let us call them the 'Th/U monitoring high window' and the 'Th/U monitoring low window' where high and low refer to the Nhits threshold.

The second goal is part of a new proposal for determining the in-situ neutral current background in SNO without relying on radio-chemical methods. It is inspired by the recent success of S. Brice in separating NC and CC events using neural network techniques. It also seems to be a reasonable goal to pursue because it is not as daunting as the third goal and yet sufficiently challenging to be a good test of vertex fitters. Its feasibility has been investigated using a fitter developed by the author and preliminary results are given in Section 3.

This new fitter is called the grid fitter and is described in Section 2. However, before jumping into the details of the grid fitter, it is important to discuss a few general consideration which apply to any fitter that is written to pursue goal 2. It is difficult to estimate the raw trigger rate in SNO at a threshold of 30 Nhits but a *very* rough guess is a few events per minute. This rate provides a lower limit for the fitting rate, but in practice a higher fitting rate is desirable. This is because it is necessary to fit many thousands of events, per event class, to study the tails spatial distributions. It will be necessary to perform these studies both to optimize the fitters and to study systematic effects such as the dependence on the physical input parameters of the Monte-Carlo. Thus a fitting rate of about one event per second is better suited to the task.

Another somewhat obvious general consideration is that the fitter should not be optimised for any one particular class of events, to the detriment of its ability to fit other classes of events. Otherwise one can end up with large rejection factors for the optimized class but poor

performance for other classes.

2 Description of Grid Fitter

The grid fitter is in essence a hybrid fitter which combines elements of the box fitter of T. Anderson and the maximum likelihood fitter of J. Klein. The author is greatly indebted to T. Anderson, J. Klein and all the participants of the last SNO fitter's meeting whose ideas have been recycled. Many thanks are due to S. Brice for writing the extremely helpful 'Adding a New Fitter to SNOMAN' guide.

Only timing information is used in the grid fitter, i.e. angular distributions are ignored. This choice was made in order to keep the grid fitter as general as possible, so that it can simultaneously fit different classes of events which can have very different angular distributions.

The first step in the fitting procedure is to search a coarse 3D grid of approximately one meter spacing across the entire detector volume, i.e. a sphere of approximately 10 meter radius. At each 3D grid point (fixed x,y and z coordinates) a log-likelihood function is minimized with respect to time, the only remaining free parameter of the vertex. Then, after the grid search has been completed, the best grid point (lowest log-likelihood) is used as a start point for minimizing the same log-likelihood function, but this time in 4D, i.e. with x,y,z, and t (of the vertex) all free parameters. The initial grid search is the main workhorse of the fitter, since it uses most of the cpu, and it is an effective way of greatly reducing the susceptibility of the fitter to local minima. A possible improvement in the method, which has not been investigated, would be to perform several 4D minimizations starting from, say, the ten best grid points or even a coarse grid itself.

The log-likelihood function is calculated in the dimension of residual time, t_{res}, which, for the ith PMT hit, is equal to the hit time, $t_{hit}(i)$, minus a fitted time, t_{fit} , and minus the time it takes light to travel from a fit vertex position, r_{fit} , to the ith PMT position, $r_{pmt}(i)$, at speed c/n:

$$t_{res}(i) = t_{hit}(i) - t_{fit} - |r_{fit} - r_{pmt}(i)|n/c \qquad (1)$$

Thus, t_{res} depends on the x,y and z values of r_{fit} and on t_{fit} . If one fixes r_{fit} at the true vertex position and fixes t_{fit} at zero then the t_{res} distribution, averaged over all PMT hits and over many events, is a gaussian centered on the true event time, with a slightly larger sigma than the PMT transit time spread and with long non-gaussian tails due to scattering, reflections and

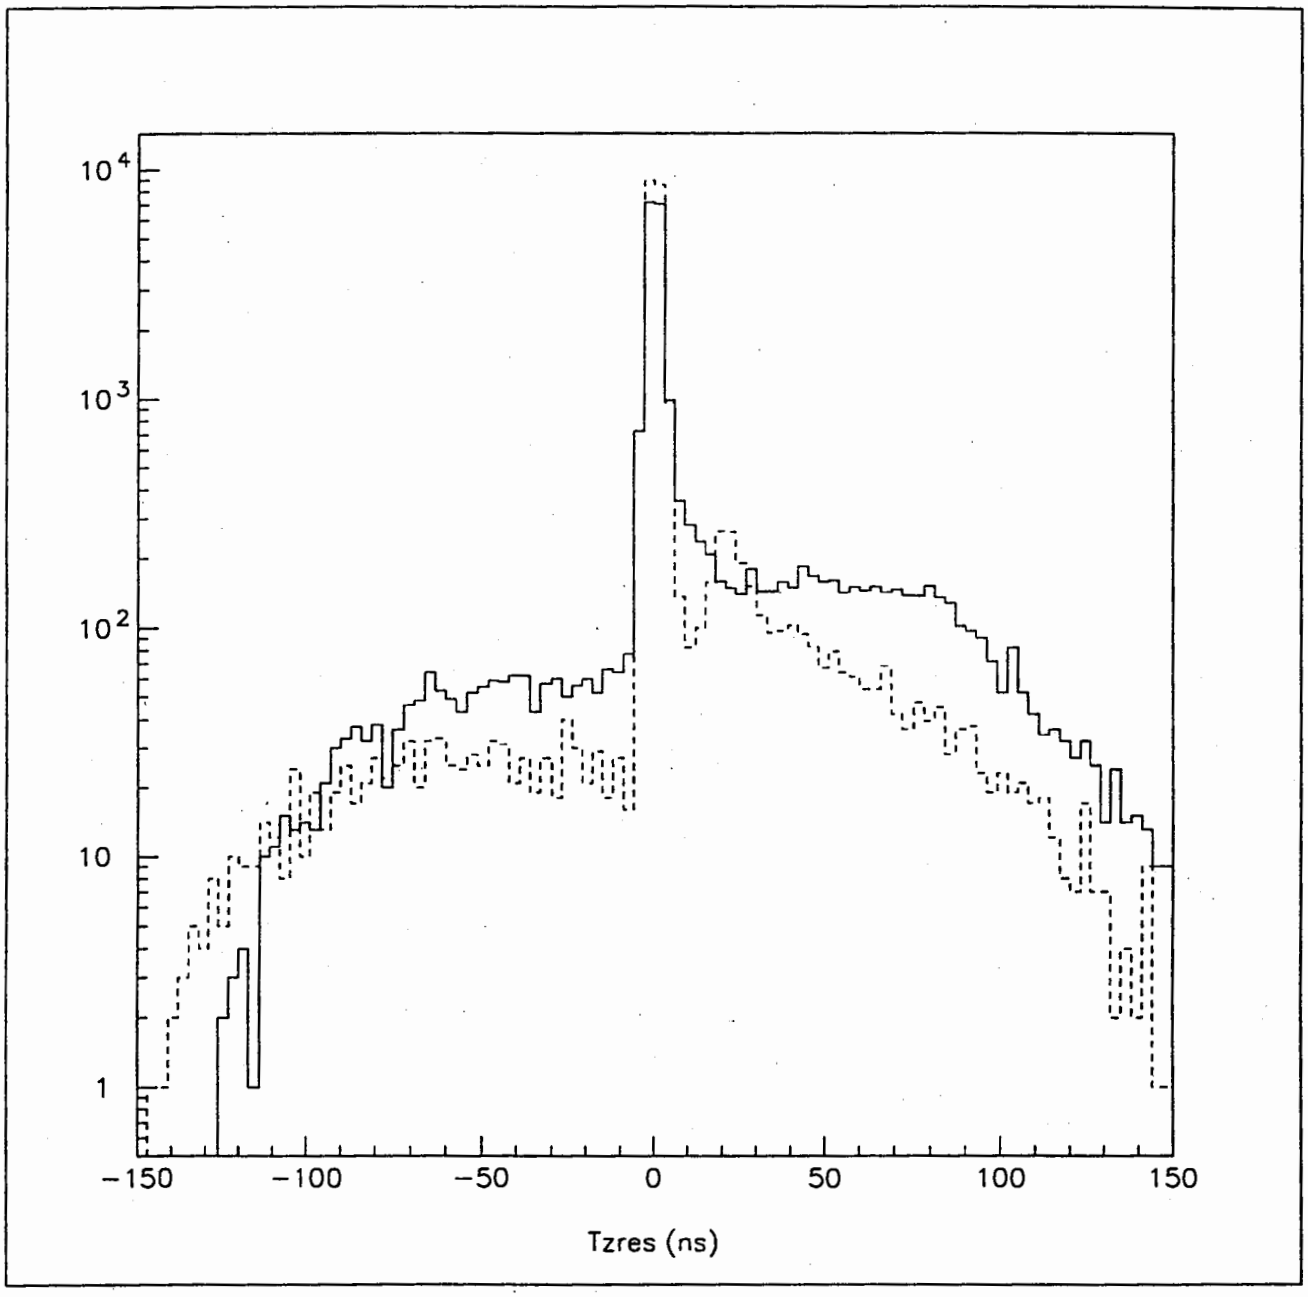

Figure 1: t_{res} distributions for PMT ^{208}Tl events (solid line) and AV ^{208}Tl events (dashed line), both after 1D minimization of t_{fit} on an event-by-event basis. Both histograms comprise 23,000 PMT hits.

PMT noise. If one allows t_{fit} to be fitted in 1D (see procedure below) for each event separately, whilst keeping r_{fit} at the true vertex position, an almost identical distribution in t_{res} is obtained apart from a shift in t_{res} (it is now centered on $t_{res} = 0$) and a very slight sharpening of the gaussian peak. This distribution is shown in Fig. 1 for two different classes of events: PMT and AV ^{208}Tl decays.

From Fig. 1 it is clear that the t_{res} distribution has large tails which can be different for different classes of events. The tail on the early side ($t_{res} < 0$) is entirely due to PMT noise as all scattered and reflected light can only arrive later (from the true vertex). The difference between PMT and AV events in the amplitude of the early side tail is due to the fact that the *relative* probability for getting noise hits is inversely proportional to the total number of hits per event and Nhits was on average greater for the AV events than for the PMT events. On the late side of the t_{res} distributions there are differences due to this same effect, but also due to the fact that the effects of reflection and scattering are dependent on the location of the event.

In order to obtain a probability distribution which is not too event class dependent, the *average* of the t_{res} distributions of Fig. 1 was fitted to a smooth functional form with six free parameters. The best fit of this functional form is shown in Fig. 2, together with the average t_{res} distribution. Also shown are two other t_{res} distributions which have been generated, for each of the event classes, by placing r_{fit} not at the true vertex but instead at a location 2 meters radially inwards from the true vertex. These 'false' t_{res} distributions show how the t_{res} distribution begins to change as one moves away from the true vertex: the central peak becomes broader and more events appear in the tails.

All of this behaviour was expected, apart from the early side secondary peak in the false t_{res} distribution for PMT events. This peak is caused by the 'near' PMT hits (those which occur in the PMT where the event originated and/or a few neighbouring PMTs). As one moves r_{fit} radially inwards into the detector, the best t_{fit} time will be dominated by the more numerous 'far' PMT hits which occur at the other side of the detector (due to the focussing effect of the concentrators). As a result the best value for t_{fit} is ***increased*** by about 8 ns, since r_{fit} is now 2 meters *closer* to the far PMT hits. This in turn decreases the t_{res} values for the near PMTs by about 8 ns (see Eq. 1). In addition, a further 8 ns decrease in t_{res} occurs since r_{fit} has moved 2 meters *away* from the near PMT hits (see Eq. 1). Thus the overall effect is a secondary peak at about -16 ns. This peak is quite helpful in reducing the number of PMT events that reconstruct inside the D_2O since it occurs on the early side where the amplitude of the probability distribution is lowest. It may even become more helpful, once the grid fitter is installed on a version of SNOMAN where the PMT hits can be weighed by ADC values, since

the near hits are often several photoelectrons each.

Having obtained a functional form for the probability distribution $P(t_{res})$ one can fit t_{fit} in 1D (keeping r_{fit} fixed) or t_{fit} and r_{fit} simultaneously in 4D by minimizing the log-likelihood function

$$\log - \text{likelihood} = \Sigma_{i=1}^{\text{Nhits}} - \log(P(t_{res}(i))) \tag{2}$$

where t_{res} is given by Eq. 1 and $P(t_{res})$ is the function shown in Fig. 2, normalized so that $P(0) = 1$ and modified so that the PMT noise component (a flat background across the entire distribution) is inversely proportional to the Nhits value of the fitted event. Both types of minimization (1D and 4D) are performed using the CERNLIB MINUIT package routines, and in particular the routine MIGRAD with the strategy set to maximum reliability which also means slowest performance.

A couple of tricks which greatly help the speed of the initial grid search may be interesting to some readers. To find a good start value for t_{fit} before the 1D minimization, the average t_{res} value (over all PMT hits) was calculated with $t_{fit} = 0$. However, this is not always a great starting value since there tend to be a few very late hits which skew the average considerably to the late side. To remove these late (and sometimes early) hits an iterative cutting and re-averaging procedure was used where the time window of accepted hits was successively narrowed until it was only ± 5 ns away from the final average. If, at any point during this process the number of accepted hits fell below a prescribed fraction of the available hits, then the grid point was automatically discounted and the search moved on to the next grid point without performing the 1D minimization. Setting this fraction at 0.55 greatly increases the speed of the grid search and has no effect on the fitter efficiency and reliability for all classes of events except PMT events where the efficiency is slightly reduced, but no effect on reliability is observed.

It should be noted that the probability distribution of Fig. 2 is very close to a truncated gaussian. In fact, if one replaces it with a simple truncated gaussian with flat tails then one obtains almost as good performance. Perhaps the only real advantage of using the complex probability distribution on Fig. 2 is that it incorporates the difference between the late side and early side tails, and this difference is physically motivated (scattered and reflected light can only arrive late, from the true vertex).

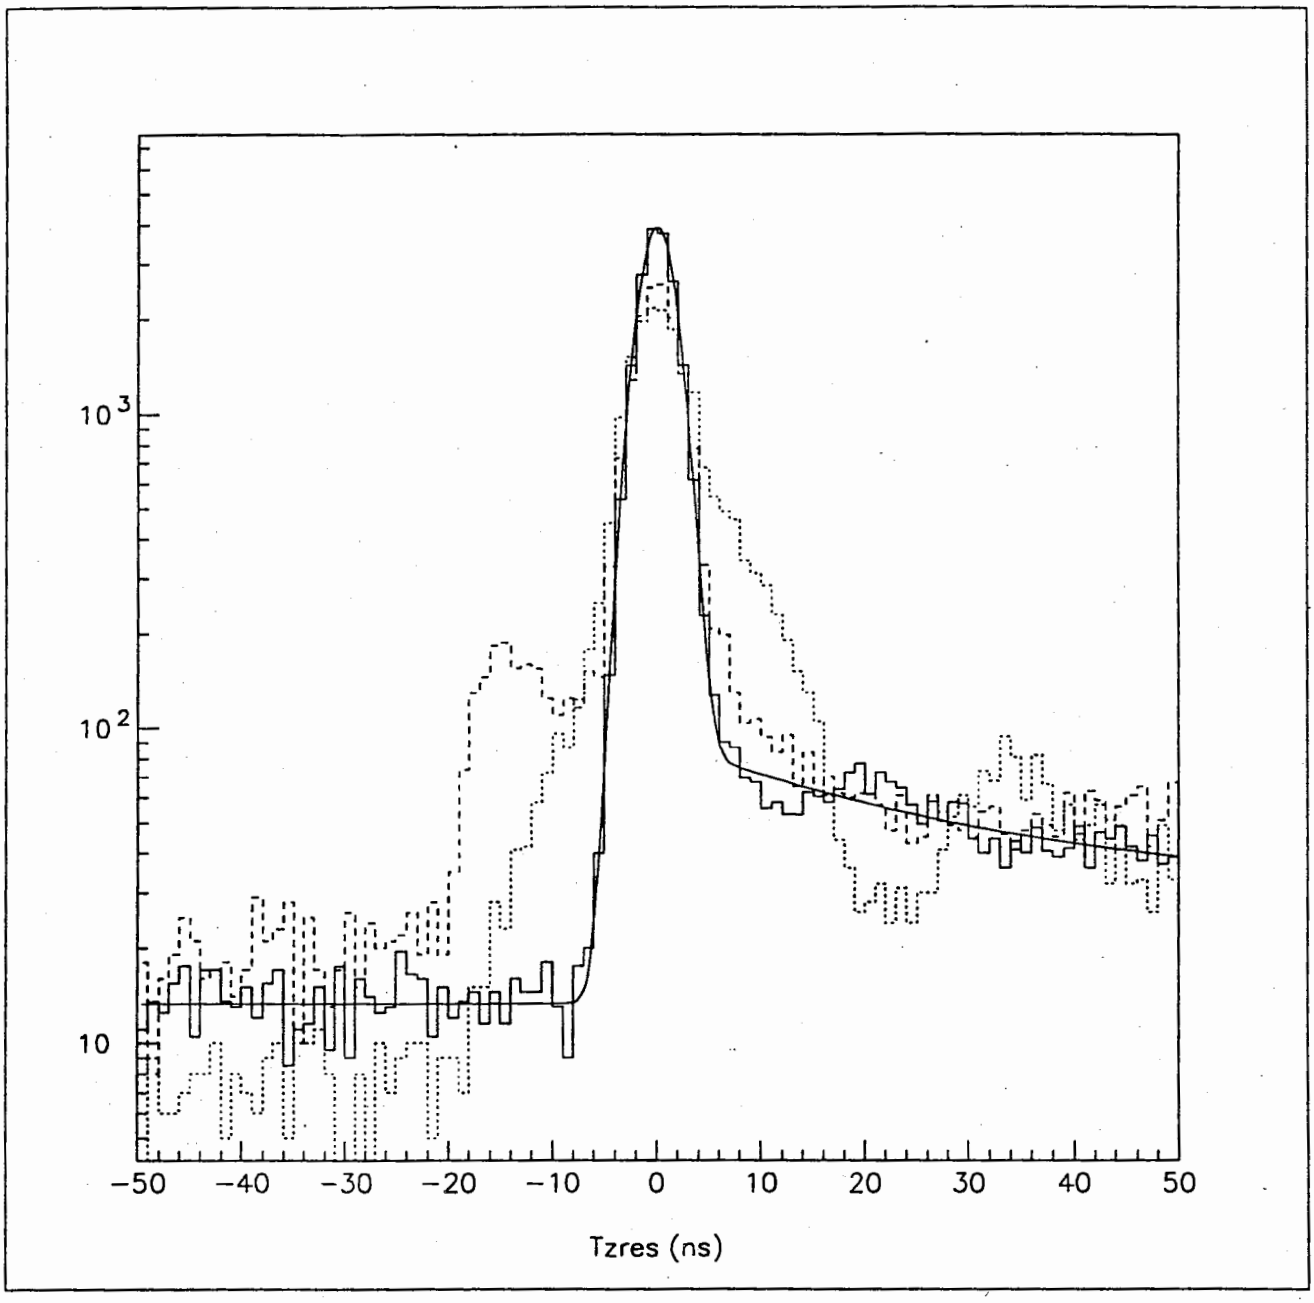

Figure 2: average t_{res} distribution for AV and PMT ^{208}Tl events (solid histogram) fitted to a 6 parameter functional form, $P(t_{res})$ (smooth curve). False t_{res} distributions for PMT (dashed histogram) and AV (dotted histogram) ^{208}Tl events, where r_{fit} was positioned 2 m radially inwards from the true vertex. Notice the change in t_{res} scale from Fig. 1.

3 Preliminary Results

Many classes of events were investigated with the grid fitter, installed on SNOMAN 2.07, and a brief summary of the results will be presented. First the speed of the fitting procedure varies between 1-5 seconds per event depending on the event class and the number of PMT hits. This is about an order of magnitude slower than the speed of the Monte-Carlo generation of the events (in the circumstance of 100% trigger efficiency). However, if one considers an 'analysis threshold' set way above the nominal hardware threshold of 10 Nhits then only a small fraction of the low energy events will pass this analysis threshold and require fitting, e.g. for PMT ^{208}Tl events only $\sim 10^{-4}$ of generated events would pass an analysis threshold of Nhits$\geq$ 30, which means that the grid fitter would consume only $\sim$ 0.1% of the cpu time.

The efficiency of the grid fitter is $>$ 99.9% for all events with Nhits$>$ 25 (CC, NC, ^{208}Tl and ^{214}Bi events in either the D_2O , AV or H_2O). The only exception are PMT events where the efficiency drops to 98% if a speed optimizing choice is made, but 99.9% can be achieved for 20% slower performance.

The ultimate performance indicator of a vertex fitter is the spatial difference, Rdev, between the true event vertex and the fitted vertex. Distributions in Rdev have been generated for all important classes of events using both the grid fitter and the standard SNOMAN time fitter. For every event class tested the relative performance of the two fitters follows the same pattern. For events with Rdev$<$ 200 cm the Rdev distributions are practically identical in shape (with the same mean) but the grid fitter manages to fit considerably more events within Rdev$<$ 200 cm, whereas the SNOMAN time fitter has much larger tails beyond Rdev$=$ 200 cm. These differences are clearly shown in Fig. 3 where the Rdev distributions for 13,000 AV ^{208}Tl events with Nhits$>$ 26 are shown for both the grid fitter and the time fitter. Unfortunately, comparisons with the quad fitter were not made because the author did not understand how to correctly run the quad fitter installed on SNOMAN 2_07.

Having smaller tails in Rdev than the standard time fitter is not too surprising. The real question is whether the grid fitter, or any other of the new generation of vertex fitters, are able to accomplish the goals listed in the introduction. Unfortunately, it is very difficult to investigate the first goal (obtaining a negligible number of PMT events in the 'neutrino window') because of the cpu time to generate a sufficiently large sample of PMT events. Nevertheless, 300,000 PMT ^{208}Tl events were generated which corresponds to roughly 1/2 an hour of real time in SNO. Of these events, 1,400 passed a 20 Nhits trigger and the radius of their fitted vertices,

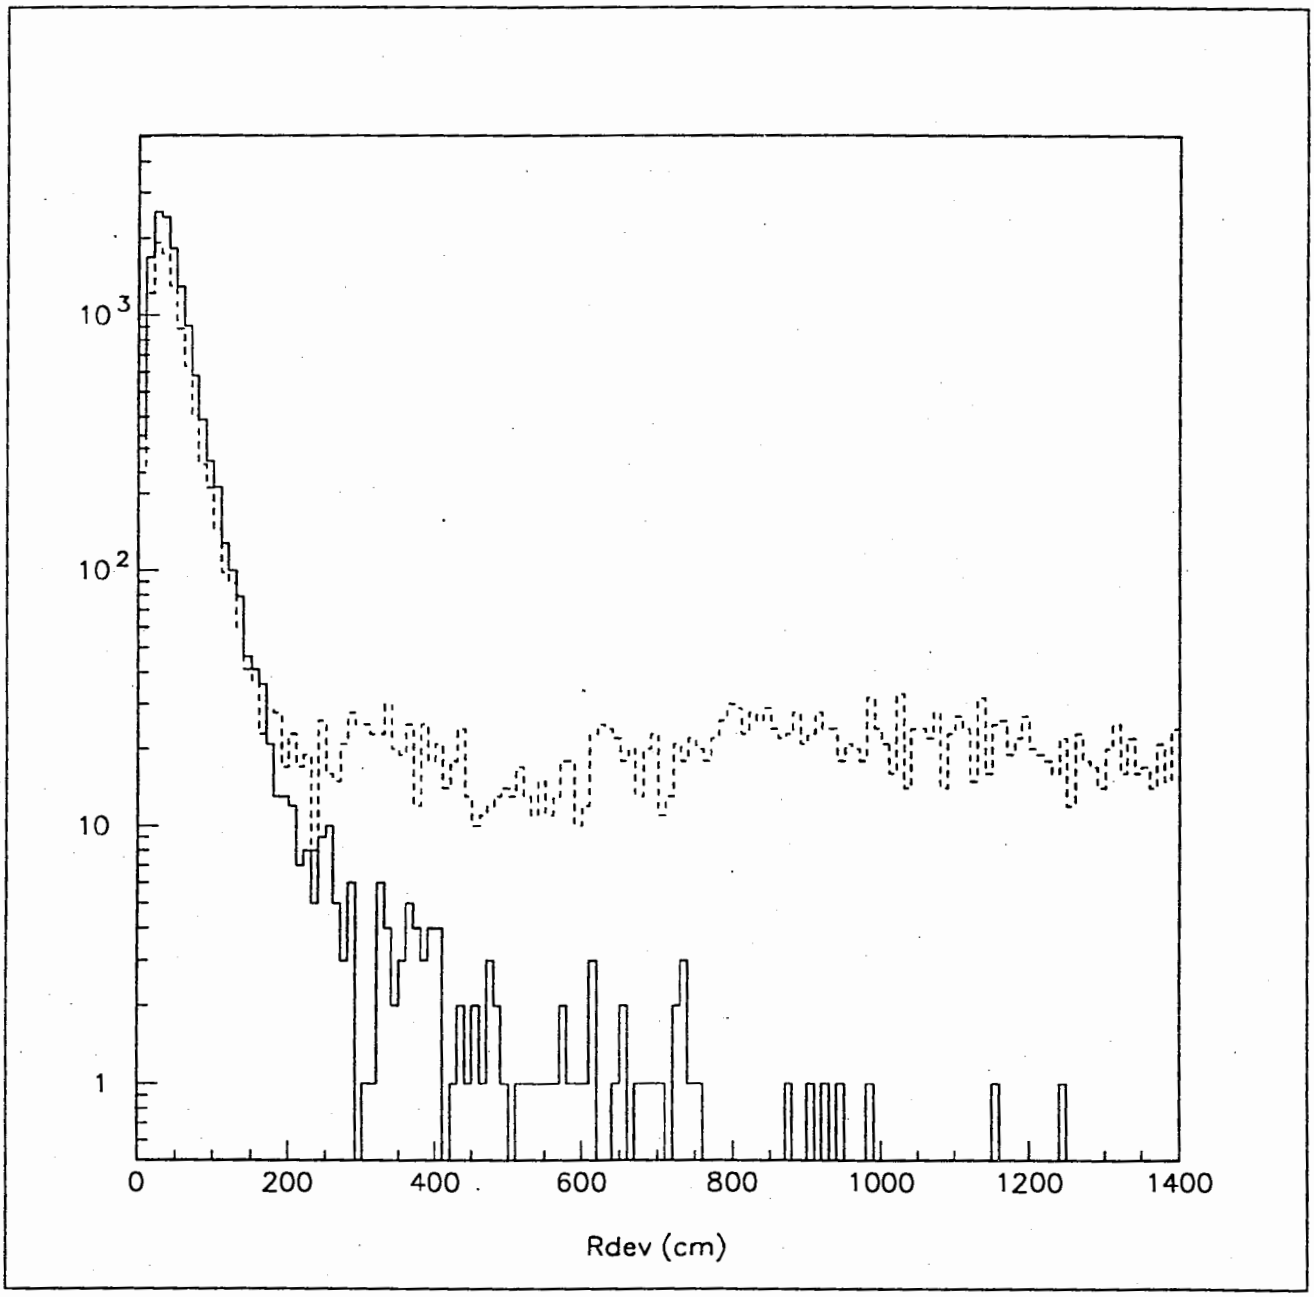

Figure 3: Rdev distributions for 13,000 AV ^{208}Tl events, with Nhits> 26, using the grid fitter (solid line) and the SNOMAN time fitter (dashed line).

region	event	assumption	#/month	# generated	scaling	raw cnts	adj. cnts
D_2O	^{208}Tl	10^{-14}g/g	38,600	30,000	1.285	499	641
D_2O	^{214}Bi	10^{-14}g/g	322,000	200,000	1.61	324	521
D_2O	CC	1/3 SSM	500	5,000	0.1	101	10
D_2O	NC	full SSM	300	5,000	0.06	72	4
AV	^{208}Tl	10^{-12}g/g	116,000	100,000	1.16	23	27
AV	^{214}Bi	10^{-12}g/g	967,000	250,000	3.87	9	35
H_2O	^{208}Tl	3×10^{-14}g/g	193,000	200,000	0.965	18	17
H_2O	^{214}Bi	3×10^{-14}g/g	1,610,000	200,000	8.05	5	40
PMT	^{208}Tl	10^{-8}g/g	3×10^8	300,000	10^3	0	< 50

Table 1: The raw counts gives the number of the events which the grid fitter placed in the Th/U monitoring high window: Rfit$<$ 400 cm and 30 $\leq$Nhits$<$ 40. These numbers have been scaled in the last column to account for the difference between the number of events generated and the correct number of events/month for the radiopurity or ^{8}B flux levels which have been assumed.

Rfit, are shown in Fig. 4 for both the grid fitter and the SNOMAN time fitter. The grid fitter shows a marked improvement, but much larger data sets are required before reliable conclusions can be drawn.

The second goal (a background-free 'Th/U monitoring high window') is easier to investigate. Table 1 shows the number of events, for all relevant classes, which the grid fitter places into the Th/U monitoring high window. These numbers have been scaled to one month of SNO real time according to a given set of assumptions for radiopurity purity levels and for the ^{8}B neutrino flux (see Table 1). Clearly, the 'background' from CC and NC events is well below the expected 'signal' from D_2O β-γ's. However, the sum of all external backgrounds is about 10% of the expected signal, and the radiopurity assumptions for the AV and the H_2O are perhaps a little optimistic. So it is not yet clear whether this window is sufficiently clean to perform the suggested ^{208}Tl /^{214}Bi separation by event topology. Nevertheless, the grid fitter improves the external background rejection by about an order of magnitude wrt. the time fitter, and it has not yet been fully optimised. For example, the t_{res} distributions for AV ^{214}Bi and H_2O ^{208}Tl or ^{214}Bi events were not used in making the smoothed probability distribution, $P(t_{res})$, and this may help to reject these backgrounds. Another improvement would be to include more Nhits dependence in $P(t_{res})$. Also, the fitting procedure could be improved by performing more 4D minimizations, using several sets of start points for r_{fit} and t_{fit} .

A natural unit for illustrating the distribution in Rfit for the above event classes is Rcube = (Rfit/600)^3. This unit expresses linearly the enclosed volume at a given radius Rfit in units where the entire D_2O volume equals 1 and the volume of the Th/U monitoring window is 0.3.

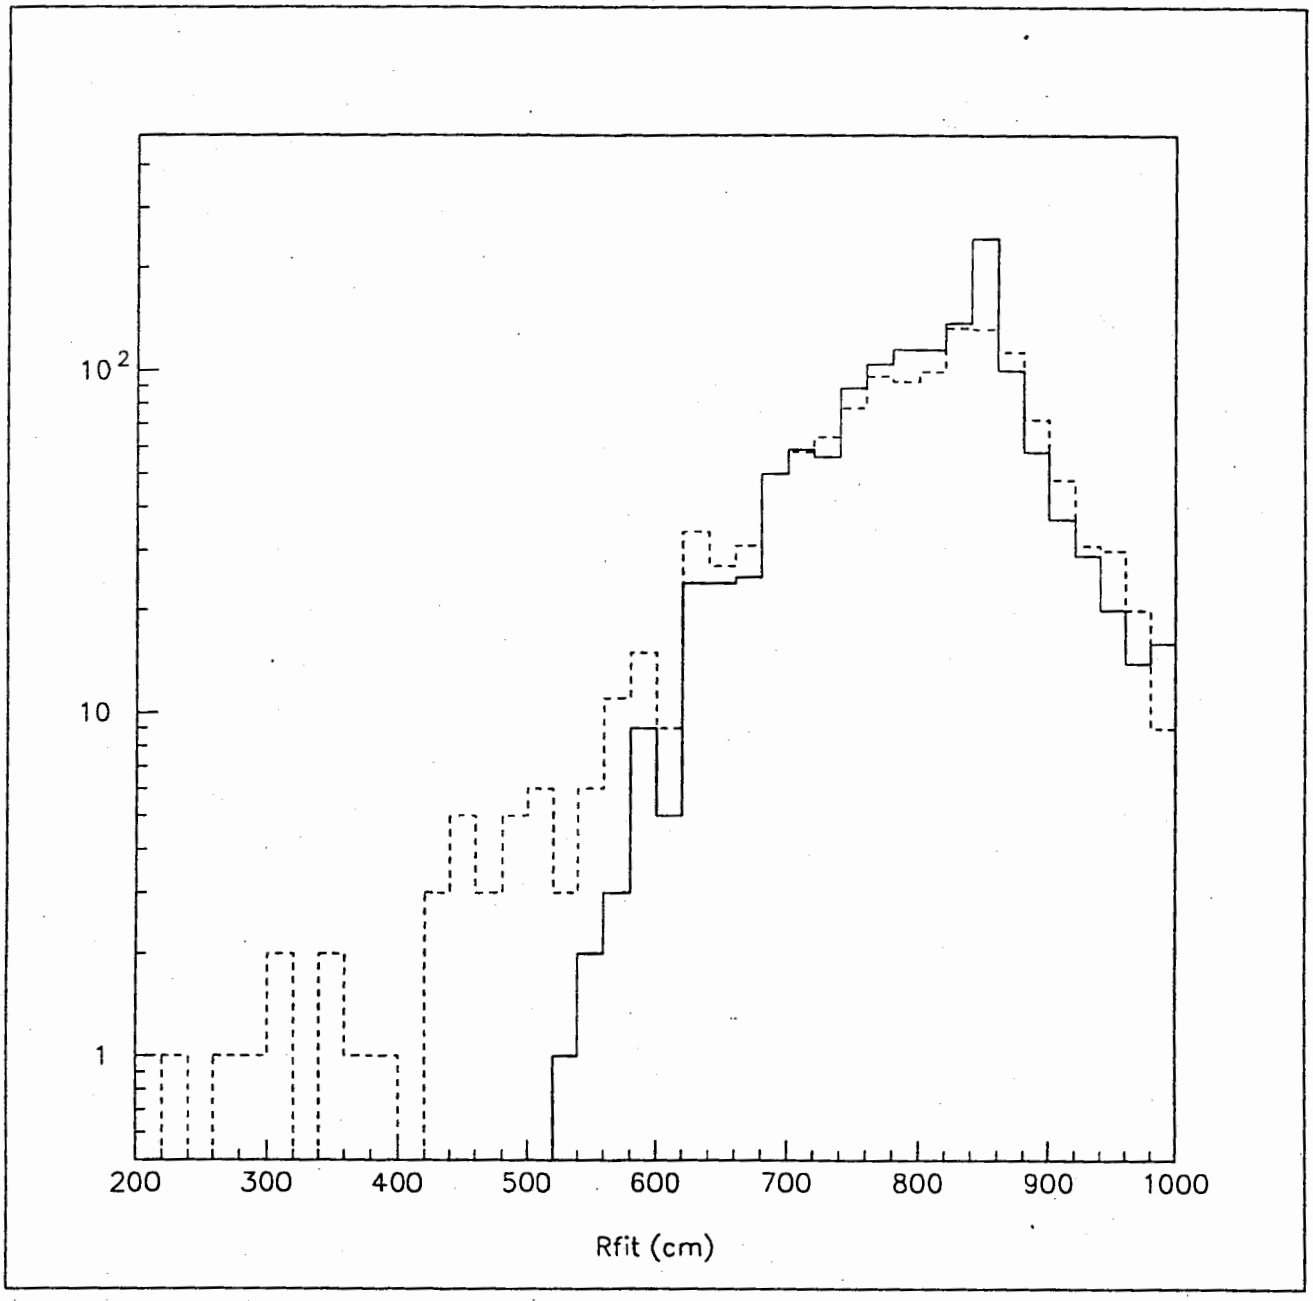

Figure 4: Rfit distributions for 1,400 PMT ^{208}Tl events, with Nhits$\geq$ 20, using the grid fitter (solid line) and the SNOMAN time fitter (dashed line).

The distributions in Rcube for the event classes of Table 1 are shown in Fig. 5, scaled to one month SNO real time and with 30 $\leq$Nhits$<$ 40.

In the above discussion the PMT β-γ background has been treated in an overly simplistic manner. Only ^{208}Tl activity in the glass bulb has been considered and the way in which it has been scaled is very crude indeed. Of the 300,000 events generated only 25 had Nhits$\geq$ 30, whereas 1400 had Nhits$\geq$ 20. So, the distributions in Rfit were taken from the Nhits$\geq$ 20 sample under the assumption that it was a fair representation for events with Nhits$\geq$ 30. Most likely, this will turn out to be a conservative assumption, but it needs to be investigated with much larger samples of PMT β-γ's, before any reliable conclusions can be drawn..

4 Conclusion

The grid fitter is able to reconstruct event vertices with higher efficiency and much smaller tails in Rdev (the spatial error) than the standard time fitter and at a relatively small cost in cpu time (1-5 seconds/event). A preliminary investigation shows that, using the grid fitter, a Th/U monitoring high window (Rfit$<$ 400 cm and 30$\leq$Nhits$<$ 40) has a 1% background from CC and NC events and a 10% background from AV, H_2O and PMT β-γ's. Ideally, this latter background from external events should be reduced to the same level (1%) as the internal CC and NC rate. This may be within reach of either the grid fitter (after some further optimization) or other fitters such as the quad fitter.

One possible role of the grid fitter could be as a 'first-pass' fitter since it has such high efficiency and is relatively fast. Then, events within given windows defined by Rfit radii and Nhits ranges could be passed on to slower fitters such as the quad fitter which (probably) have even higher rejection factors for external backgrounds. In this way, the highest efficiency for 'internal' events and minimum contamination from 'external' events could be obtained for any window of interest. Obviously, the window in the first-pass fitting must be larger, in Rfit, than the final window of subsequent fitting in order to avoid edge effects.

It is hoped that large data sets, for the event classes discussed in Section 3, could be generated using SNOMAN 2_08 and distributed amongst SNO institutions so that valid fitter comparisons can be made. In particular, there is an urgent need for sufficient numbers of PMT β-γ events to be able to draw solid conclusions about the possible contamination, by PMT β-γ events, of the three windows discussed in Section 1.

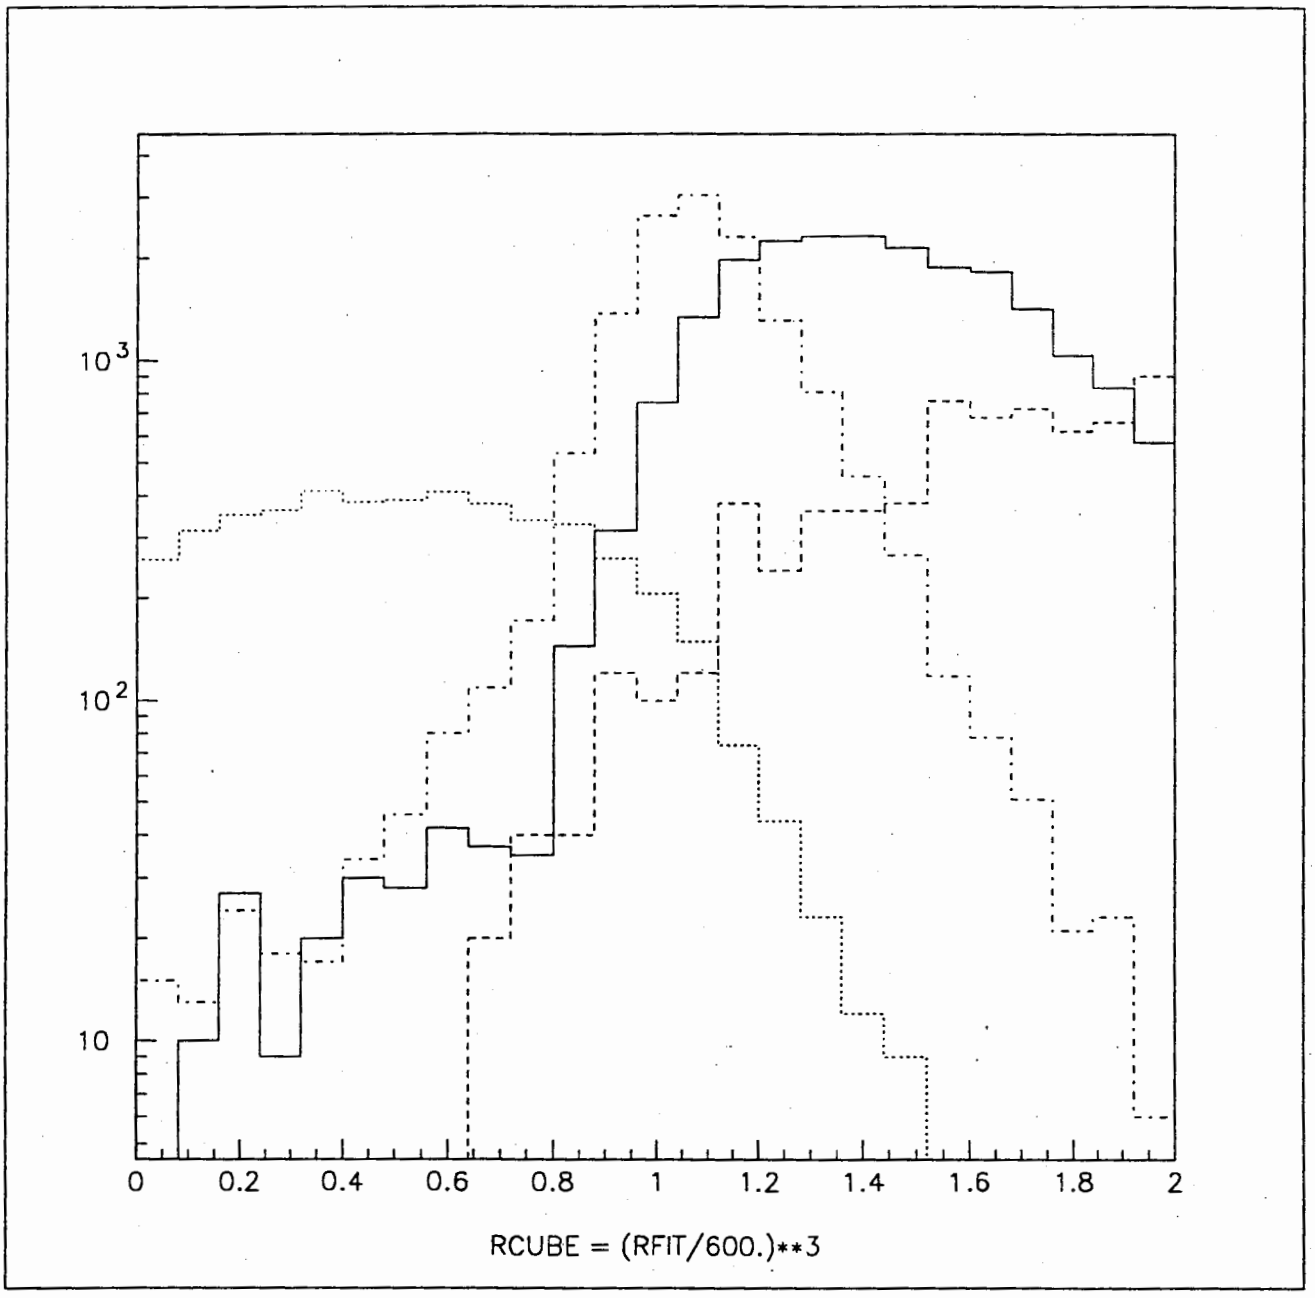

Figure 5: Rcube distributions for D_2O β-γ's (dotted histogram), AV β-γ's (dash-dot histogram), H_2O β-γ's (solid histogram) and PMT ^{208}Tl events (dashed); all with 30 $\leq$Nhits$<$ 40 and scaled to one month SNO real time according to the radiopurity assumptions of Table 1.