MINERvA Technical Note No. 037
DocDB document 9131

Birks' Law parameter from MINERvA testbeam data

Rik Gran, University of Minnesota Duluth

December 14, 2013

Using the testbeam sample of protons stopping at the end of their range, we have obtained a more accurate value for the Birks' law quenching parameter for the Monte Carlo simulation. Birks' quenching affects the observed photon yield for high dE/dx particles, and is especially interesting for the end point of proton and nuclear fragments as well as vertex activity near a hadronic interaction on a pion or proton track or from a neutrino nucleus interaction. Mis-modeling this parameter would be visible in proton/pion PID and energy estimations using the dE/dx profile or calorimetry. This is one of several core results from the MINERvA testbeam experiment's 2010 data, Fermilab experiment T977. The data prefer a Birks parameter of 0.0905 ± 0.012 mm/MeV, just beyond the one-sigma limit of the old default value of 0.133 ± 0.040 mm/MeV. This shift to less suppression causes an 8% increase in the MC's apparent energy in the strip where the proton stops. This analysis uses a smaller Geant4 and MCHit step size than the default MINERvA in a way that is intrinsically correlated with this parameter. A simulation using the larger step size (which saves 20% file size) should use 0.0934 ± 0.012 mm/MeV.

In addition to Birks' parameter, this proton sample is sensitive to several other aspects of our detector simulation, which become systematic uncertainties on Birks' parameter and might be interesting for other special situations in MINERvA neutrino interaction analysis. The largest effect (which we take as a systematic but might not be) is the Geant4 and MCHit step size. Other significant physics uncertainties include PMT nonlinearity and the material assay, though analysis design and event selection are not negligible.

Contents

1 Form of Birks' suppression compared to true proton dE/dx

In our software, Birks' suppression is applied by us in the OpticalModel in the following way for every MCHit coming from Geant4, given the hit's dE and the hit's dx:

$$[\text{Light yield factor in (photons/MeV)}]\ \frac{dE}{1.0 + \text{Birks Constant} \times (dE/dx)}$$

which is the simple form captured in the Particle Data Group (eq. 31.3 of the 2012 version of Particle Detector section). The PDG itself wisely does not give a value, it suggests that if your experiment is sensitive to it, you should measure it for yourself and your situation. The effect of Birks' quenching is illustrated in Fig. 1 for a proton in the equivalent of the last MINERvA plane or so of its energy loss. In this 22mm piece of range the Birks' suppressed response averages to 65% of the true energy loss. The difficult to see low value for energy loss at the 22 mm point is 1.3 MeV/mm, already well above the muon minimum ionizing energy loss of 0.2 MeV/mm, and the Birks' quenching there reduces the response to only 85% of the true energy loss. This figure is constructed for a proton using 0.1 mm steps using the Bethe-Bloch calculation in our software framework.

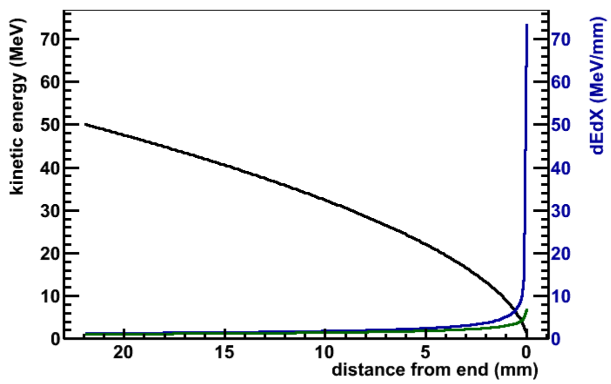

Figure 1: Kinetic energy and energy loss at the very end of a proton track, from the Bethe-Bloch calculation. This illustrates the last 50 MeV or 22 mm of energy loss, about a single MINERvA plane's worth. The green line shows how much of this energy is available to be reconstructed using our nominal Birks' constant, integrated over the 22 mm range shown, the value is 65% of the true energy loss.

The provenance of our original value 0.133 mm/MeV for Birks' constant is difficult to figure, it probably came from MINOS, for whom some say it came from MACRO, which is an uncertain trail to follow via the internet. Below are some values obtained from creative search terms and time devoted to a name-brand internet search engine.

The real world does what it does, but applying Birks' constant to the MC is affected by the step size used in the simulation, which is discussed in the systematics section. It can also depend on the

0.126 mm/MeV	Wikipedia, quoting Leverington 1970 for polystyrene
0.133	MINOS, Dave Petyt's thesis, maybe quoting MACRO, and separate work by N. Tagg around.
0.1	MINOS, quoting Doug Michael
0.116	MACRO for liquid scintillator
0.208 ± 0.003	SciBar testbeam T551 at KEK, protons stopping in polystyrene
0.135 ± 0.02	nuclear recoils 0.3 to 0.85 MeV, L. Reichhart arXiv:1111.2248
0.09 ± ?	alphas, from above ref 15. Tretyak Astropart. Phys. 2010
0.151	says CALICE from their own fit to some data.
0.079	says Hirschberg, IEEE 1972, quoted by CALICE for a while.
0.103 ± 0.012	the final best value from this work.

Table 1: Some values of Birks' parameter found in the wild.

property of Geant4 to deposit in one spot all the remaining energy of a low-energy (below energy cutoff) particle or a secondary particle that is not to be simulated. In our OpticalModel implementation, this last effect is probably caught because a tiny path of less than 0.001 mm will have its Birks' law quenching calculated based on a path of 1mm instead. A separate effect that has similar properties is the (lack of) PMT non-linearity in the simulation, though it affects high dE per strip, not only high dE/dx.

There are also granularity effects in using data to constrain Birks' parameter. We effectively average over finite step size, the crossing length of a triangular scintillator bar, or more traditionally a doublet of triangles spanning the thickness of the plane.

For a muon, the quenching is mild, because at the canonical 2 MeV/cm (or 0.2 MeV/mm), the quenching is between 2 and 3%. In the plane before the end of the track, a proton in MINERvA will deposit around 1 MeV/mm, so the quenching should be correspondingly higher, more like 10 to 15%. In the last plane (one or two pieces of triangular scintillator) the proton crosses, it will go through a changing range of dE/dx from 1 MeV/mm to 10 MeV/mm, with the last 10 MeV dumped in a very small space at the end. The apparent quenching in the last physical plane depends quite specifically on how far it penetrates into the detector.

Converge through iterations The starting point for the Birks analysis was variations of Birks' parameter of 30% from its nominal 0.133 value. The analysis proceed with iterations with a parameter values of 0.133, 0.120, 0.103, 0.0934, 0.0905 mm/MeV and increasingly smaller one-sigma uncertainties. The converging iterations allowed for bugfixes, improvements to the analysis technique, improvements to the calibrations, improvements to the MC, and better understanding of systematics at each step, so a comparison between iterations does not simply test only two things. Some elements of the analysis involve interpolating through effects that are not necessarily linear, and every next iteration allowed for a validation that improvements from the previous step took effect. Plots from different steps appear in the following discussion if they best illustrate the nature of the fit or systematic better than the final iteration, in particular many of the systematics estimates were evaluated using the iterations that produced the 0.103 and 0.0934 results.

2 Proton selection

The selection tries to make a sample of protons which stop at the end of their range ("range out") between modules 11 and 19 inclusive of the 20TRAK20ECAL detector and rejects protons that likely experienced an interaction or are contaminated in some way. First the sample passes the beamline and beam-detector match quality selection criteria and events were selected because they were protons.

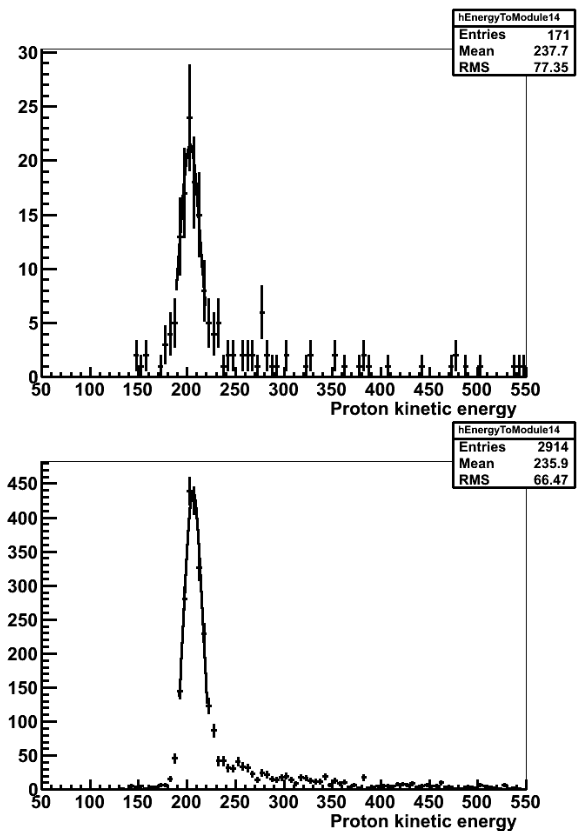

Figure 2: Distribution of beamline proton kinetic energy of particles that appeared to stop in plane 14 for data (left) and MC (right). Events in the peak actually did stop without interacting, the tail are protons that interacted. The Gaussian fit to locate the peak is shown. When this profile is used to select protons that likely ranged out, I use a range of Kinetic energies surrounding the Gauss fit mean and within ± 1.5 sigma, further to the base of the distribution than the roughly one-sigma shown in this plot.

When the sample is analyzed, the code dynamically finds the energy spectrum for protons that appear to stop in each module. This spectrum has a peak which corresponds to protons that actually range out, significant high-side tail of protons which had too much energy to stop in that plane and must have interacted. An example for plane 14 is shown in Fig. 2. There is a smaller low-side tail of protons too, in the MC these are protons whose simulated energy is fluctuated up from the beamline energy plus a small fraction that interacted upstream whose products just happened to land near the right plane. In both the data and the MC, the events in this low-side tail also visually look like good protons, except for one data event which suffers from a second particle entering the front of the detector.

This peak is fitted with a Gaussian and the mean and sigma are recorded. Visual inspection of the distribution indicates that the MC reproduces the peak and tails quite well, and quantitatively the means and sigmas are within sample fluctuations of each other. This selection also forms the basis for the separate analysis of proton KE vs. range, for example docdb:7984.

To select protons in this peak and not in the tail, for each plane we select protons that are within 1.5 sigma of the Gaussian mean. This selection can be made narrower or wider to investigate possible systematics. To emphasize, this is done dynamically and separately for the data and the MC; if it were the case that MC protons travelled less far into the detector for a given energy compared to data, the selection would still return the subset of protons in that plane that were ranging out in that plane. As it happens, there is quite good agreement between data and MC for energy vs. range, 1.3% in the WilliamClayFord processing.

In addition, three other quality selections are made. There is a subset of events that have a gap of two or more planes with energy less than 0.5 MeV along the path of the proton. In AnaTuple post-processing, a simple pattern recognition algorithm recognizes when there is such a gap in activity per plane disregarding any tracking information. The selection shown in Fig. 2 uses the last module before a gap as its criteria for inclusion. After this additional selection, only events without a gap are kept, which eliminates about 2% of the events in the MC but 14% of the events in the data. Most of the data in this category have either a few stray hits downstream or evidence of muons or other particles that passed the standard cleaning selections, with activity often occurring beyond plane 30. In principle, much of this lost 14% could be harvested back, but the statistical power of the sample is already adequate. In the MC, the events that are cut say they underwent some kind of interaction, usually in or near the plane before the gap, pushing some activity further downstream by 15 planes or less.

This analysis has almost no dependence on tracking. The cluster energy per plane is not from track nodes but rather is the sum of clusters greater than 0.3 MeV within 60mm distance of the point where the original proton direction suggests it would have crossed that plane. If there are two or more planes missing anywhere between the last and first plane, the event is not included. This could be simple inefficiency, or a proton that kinked enough that in one view it was no longer near the original trajectory by the end of its range.

The energy quantity at the center of this analysis is energy per plane. It is NOT formed directly

Selection	data	MC
last hit planes 11 to 19	1342	26097
consistent with range out	790	16096
no gap	682	15727
at least 10 nodes on track	667	15680
last track node is really at end	645	15274

Table 2: [probably should update this for the final WCF processing and MC.] Selection of range-out protons in the data and 20x data-driven MC. The only one that is severely different is the no-gap selection, which is sensitive to beam pileup effects we do not simulate.

using tracking information, and track cleaning and other cluster breaking routines are turned off. The direction of the incident proton is projected onto each plane. All clusters which are within 60mm from this point in the plane and with an energy greater than 0.3 MeV are summed to be the total proton energy per plane. At the point where this quantity is formed, the analysis tool can optionally apply a correction to each cluster to make the MC PMT response non-linear, and/or apply our canonical LPos effect energy tweak.

Because the analysis is using clusters, if a track was formed, the clusters will have attenuation corrections applied from the default center of strip to the tracker's estimate of the position along the strip. The protons in this analysis are reliably tracked, though sometimes the tracker stops short and does not include the cluster in the actual last plane of the event.

The result is a sample of 645 protons that range out between tracker planes 11 and 19. [Update ?]

There is a scan of a subset of the selected events by a UMD undergraduate Alexander Lovelin to check that the sample is as clean and minty as hoped.

2.1 Protons that are lost

A potentially interesting question, though not directly for this analysis, is how many protons that should have stopped between planes 11 and 19 didn't actually make it? An analysis of a sample of low energy protons (not necessarily range-out protons) will have a choice between energy reconstruction by range, by passive-material corrected calorimetry, or by full calorimetry. The first two should essentially agree, while the last one will be as significant as the fraction of protons that interact and especially are absorbed or undergo a neutron knockout reaction. I asked the MC to answer this question.

The protons that range out according to the criteria above are roughly between 160 and 260 MeV kinetic energy. Actually, the protons at 160 MeV are also consistent with ranging out upstream in planes 8, 9, and 10. Of the protons in the MC sample in this energy range, 5.3% don't even make it to plane 8 by the above criteria, 16.9% appear to stop in planes 8, 9, and 10, and 6.8% stop a little long, in planes 20, 21, or 22. Of the ones that do make it into the range of planes 11 to 19, another 18.3% don't make the strict requirement of being in the peak, some of which should be considered to stop significantly early, some of which are simply caught by a modestly tight cut.

So, I infer that between 5.3% and 20% of protons in the range $160 < KE < 260$ MeV would have a low energy estimate if reconstruction by range was used. Depending on the nature of the sample, a calorimetric reconstruction might have benefits.

3 The dE/dx profile

This section describes a simplified analysis using a Gaussian shape fit which informed the development of the final Birks' parameter fit and binning. It is by far the best place to describe how the fit works and what physics it is sensitive to and how much the Birks tuning affects the MC model; the discussion that follows will refer constantly to these two plots. Though it looks simple once its plotted, several analysis decisions have already been wrapped up in the profile plots in Fig. 3. This set of plots is from two iterations Birks' tuning process, one with the original Birks' value of 0.133 ± 0.040 mm/MeV and one just prior to the final tuning with 0.120 ± 0.024 mm/MeV on the way to the final fit value.

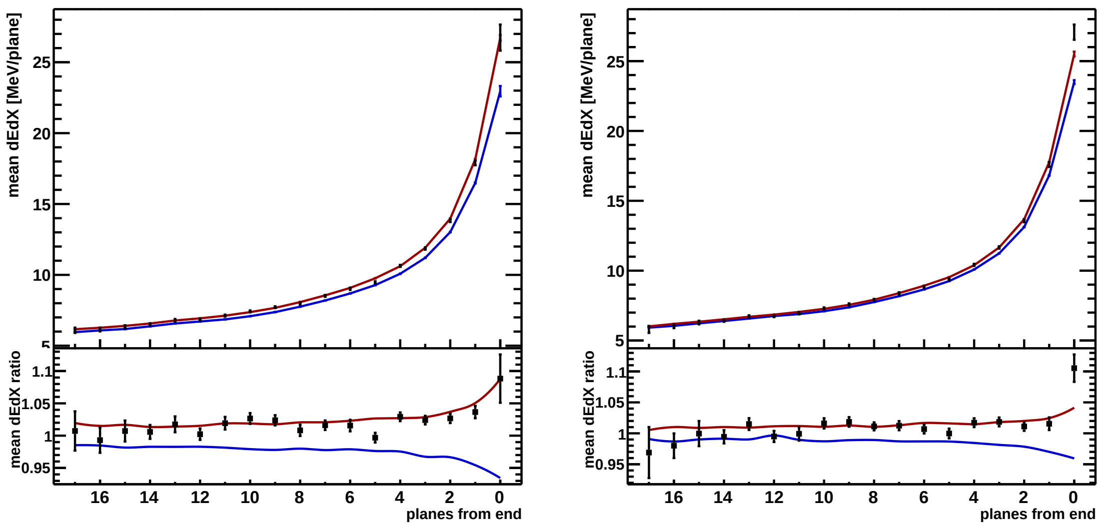

Figure 3: (In color) A pair of iterations in the dE/plane profiles for the data compared three different MC samples. The left plot is MC generated with Birks' parameters 0.133 ± 0.040 mm/MeV and the final processing of the proton data. The second sample is MC generated with parameters 0.120 ± 0.024 mm/MeV and a self consistent (not final) processing of the proton data. The two MC have less Birks' quenching (red) and more (blue) by tweaking the parameter down and up respectively. Each point is a Gaussian fit to the peak of the dE/plane distribution, however the plane-from-end = 0 point is not at all Gaussian so its fit is not robust and has a large uncertainty, see text for demonstration of how this is constructed. The default MC is not shown explicitly, but it is the denominator used to form the ratio in the bottom portion of each plot. There are more variations between the MC and data samples (MEU constant, geometry tuning, etc.) than just the Birks constant, but the main trend visible, especially near the right side of each plot, is the need to tune the Birks parameter much closer to the red line.

We have the dE/plane for clusters from plane one to the end point, and renumber the planes so that all events contribute to the plane-from-end = 0, which is the observed end point, and so up to plane-from-end = 17 = the second plane from the front for the few events that travelled all the way to planes 19. The way this is constructed, the planes-from-end dE/plane distributions are averaging over several X,U,V planes including a full set each of east and west readout. All planes have a well defined peak in the dE/plane distribution. None of them are perfectly Gaussian. The spectrum of the plane-from-end=0 is especially sculpted because some protons stop at the very front of the plane and appear to have a low dE while others stop at the very back of the plane and have a higher apparent dE.

There are three MC sets active in the plots in Fig. 3. The default MC is not shown directly, but the data is bracketed by the two MC with tweaked Birks' constants. The red has less quenching because the Birks' parameter has been turned down to the -1σ value, the blue has more suppression because its constant is at the $+1\sigma$ value. The figure illustrates the original lack of agreement in the proton profile with the Birks' parameter 0.133 ± 0.040 mm/MeV, and the improved agreement with the starting point for the next iteration and Birks' parameter 0.120 ± 0.024 mm/MeV.

The ratio shows the data and the two tweaked MC's relative to the same default MC. There is an overall offset because the Birks' law suppression affects the whole range of dE/dx including the muon MIP peak and most parts of the proton tracks. If the default MC had the wrong Birks parameter, about half of this offset would naturally be absorbed in the MEU factor, and anyway there is an additional uncertainty in the MEU tuning between data and MC. There is a trend upward at the end of the track, which is where the fit will have the most power to constrain the parameter.

Changes between these iterations (and the final result) include more than just tuning the Birks parameter, they include modifications to the geometry, smearing, and MEU tuning, so some of the offset seen in the ratio plot is due to these changes. Nevertheless, the Birks' parameter is the most significant feature visible, and the iterative nature of the analysis mitigates this tangle of effects.

3.1 Spectra for individual planes

Each data point in the above plot is obtained from a Gaussian fit to locate the peak of the energy loss spectrum in each plane. Examples of these fits are shown in the following figures. The Birks' parameter extraction is done by fitting these distributions individually, interpolating between the Birks shifted distributions for each bin, computing a χ^2, and finding the Birks parameter that best describes all of the distributions together, so these plots are really the data and model for the real fit described in Section 4.

In the middle, like plane-from-end 11

For planes in the middle of the profile, the distribution is quite Gaussian with good statistics, and well modeled by the MC. The events are averaging over many planes, including planes with both east and west readout, and X,U,V orientations. The Gaussian is constrained to only use the data from within $\sim$1 RMS of the simple mean value of the distribution, and the Gaussian mean and its error are transferred to the full profile plot. The Birks fit is similarly constrained, using the bins from X to Y. What that full fit is done, bins that cover most of the peak are used, partly to avoid mis-modeling there, and partly so energy scale shifts never produce a MC bin with zero entries. In the case of plane 11, bins from X to Y are included in the fit.

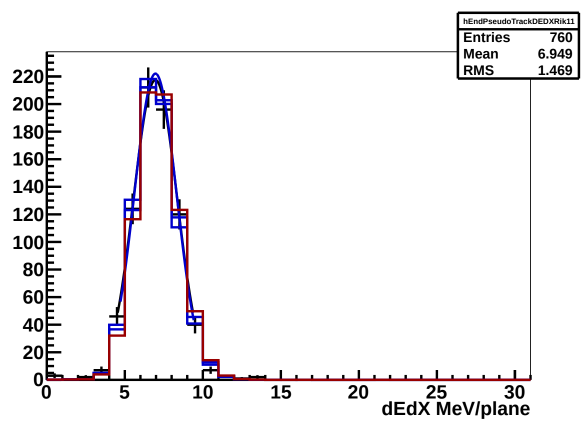

Figure 4: (In color) The energy loss distribution for the point 11 planes from the end.

At the front of the detector, like plane-from-end 16

For the planes far from the end, the statistics are necessarily low, the data is restricted to come specifically from the very front planes of the detector, and they specifically average over very few physical planes. The data at planes-from-end = 16 contains events crossing physical planes 1, 2, and 3 only. Despite this, they look okay. The planes-from-end = 17 looks mangy but is included in the profile plot, and the planes-from-end=18 defeats the Gaussian fit and is not included anywhere.

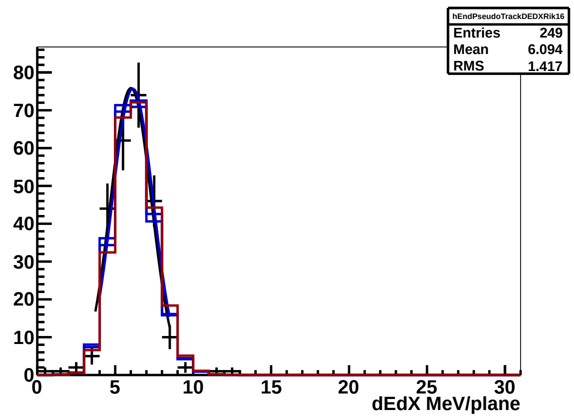

Figure 5: (In color) The energy loss distribution for the point 16 planes from the end.

In a way that doesn't affect the conclusions, there are misbehaviors are consistent with low statistics in the bins far from the end. They have 64, 122, 197, and 265 events. The width of the range used for the Gaussian fit is expanded ad-hoc to nearly twice as wide as the simple RMS because otherwise it would not be using enough of these relatively coarse bins. The fitting procedure has also made the bins even more coarse for the planes-from-end 17 data point. In any case, bins from plane-from-end 15 and higher are not used in the Birks fit.

Near the end of the track, like plane-from-end 1 and 2

The distributions close to, but not right at the end of the track exhibit larger energy loss and larger fluctuations. In addition, these planes especially show a source of fluctuation in the data that places more events up in the high reco energy tail than is observed in the MC. The method of fitting to the peak insulates this analysis from that feature. Finally, the plane-from-end 1 has a high side tail modeled by the MC that is somewhat non-Gaussian, though not nearly as radical as the plane at the very end of the track.

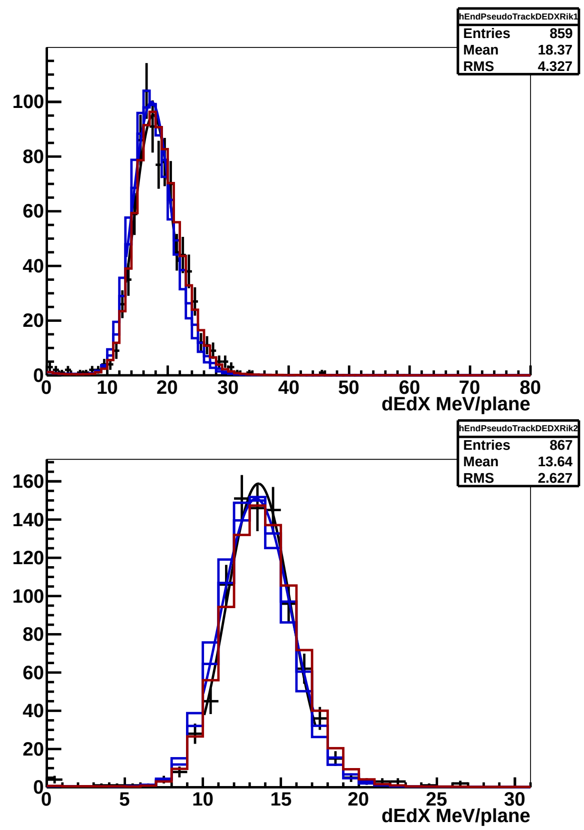

Figure 6: (In color) The energy loss distribution for the points 2 and 1 planes from the end. Caution: the horizontal scale is significantly different between the two plots!

At the end of the track

The plane with the end of the track is particularly interesting. This is not at all Gaussian because the observed energy loss depends on how far into the plane the proton went before it stopped. Additional fluctuations are possible if the proton stopped near or in the gap between triangles. There is obviously no simple Gaussian-like peak to fit. An ad-hoc choice has been made to locate where the shoulder falls off the edge, with some by-hand tuning of the horizontal range for the Gaussian fit. It is not even clear that one can interpret a simple shift in the "peak" as evidence for needing a tweaked Birks' parameter. The actual Birks' fit interpolates between the two Birks' extremes for range from 10 to [where], and so captures how the Birks parameter (and other systematic parameters) affect the whole shape of this distribution, not just the shoulder.

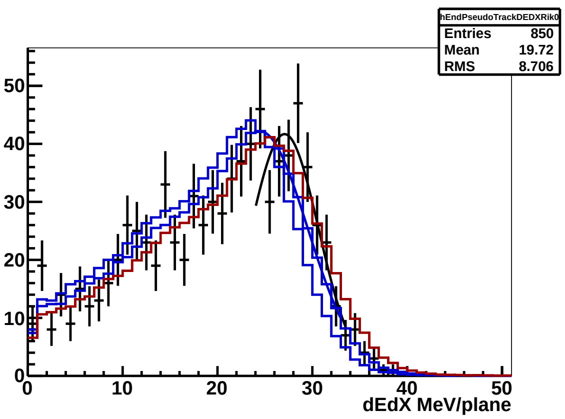

Figure 7: (In color) The energy loss distribution the last plane.

4 Birks parameter fit

The fit is done using the per plane profiles in the previous section, limited to bins covering most of the central peak, and is χ^2 scan in 2D space of Birks' parameter and energy scale, rather than an optimized minimization routine. The procedure for the χ^2 scan is to choose a value for the energy scale, and apply it to the energy per plane right before it is put into the histograms, and write out a set of histograms for nominal, high, and low Birks' parameter values. In this sense, the energy scale is has an effect identical to a modified value for the MEU factor. With each set, a Birks parameter in steps between $\pm 2\sigma$ are used to generate a value based on the high and low parameter entries for every relevant bin from every plane which is compared to the data for that bin. For every such bin, a contribution to the χ^2 is added, including both data and MC statistical uncertainties.

The best fit parameter is different than the default 0.103 mm/MeV of that MC by 1.04 sigma, which translates to 0.0905 mm/MeV. From the fit alone, the uncertainty (using $\Delta\chi^2 = 1$) is $+0.52\sigma - 0.38\sigma$.

This treats the best fit energy scale of 0.9875 as an unconstrained parameter to marginalize over and not a parameter itself to be measured. The χ^2 surface is shown in Fig. 8.

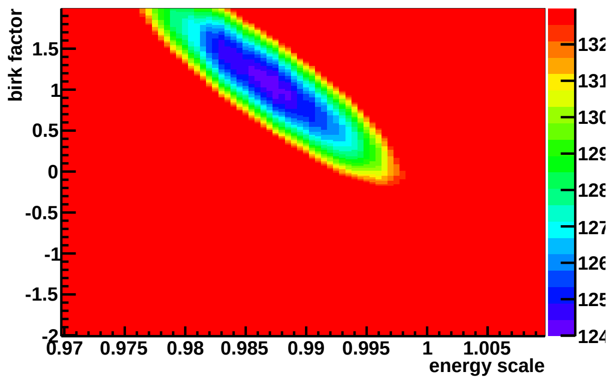

Figure 8: (In color) chisquare surface that produces the final fit result. The vertical axis is NOT the Birks' parameter, it is a scale factor between nominal, one, and two sigma ($1\sigma = \pm 0.012$ mm/MeV), with higher response (lower Birks' parameter, less suppression) at the top of the figure and lower response at the bottom.

The parameter in the fit is not the Birks' parameter, it is the scale factor between nominal and one sigma, and is defined such that +1 sigma means higher light yield and response and lower value for the Birks' parameter. In this iteration from penultimate to final fit value, $1\sigma = 0.012$ mm/MeV, and all uncertainties described below will be expressed in terms of this, and/or in terms of an absolute quantity in mm/MeV. In a few cases, a systematic was evaluated with a previous iteration that had Birks parameter 0.120 ± 0.024 mm/MeV iteration, where 1σ is twice as large, in which case the outcome of the systematic study is obviously also scaled.

4.1 Subsets of planes in the fit

Figure 9 reinforces how the full range of planes in the proton sample yield a constraint on the Birks parameter. Fitting each sample individually shows the natural correlation between the two fit parameters, and the combined best fit lands where the individual slope correlations converge. This set is with the small Geant4 and MCHit step size and the best cluster-based smearing of 6%.

The lower right plot is from planes-from-end 5 through 13, and its primary role is to constrain the energy scale in the way the MEU factor does. It has a mild preference for a high energy scale and high Birks constant because there is a mild downward trend in the data/mc ratio which would prefer a little more Birks suppression (blue line in Fig. 3) and the positive energy scale accommodates it, but fundamentally the planes closer to the end will dominate the choice of Birks factor and this subset will dominate the choice of energy scale.

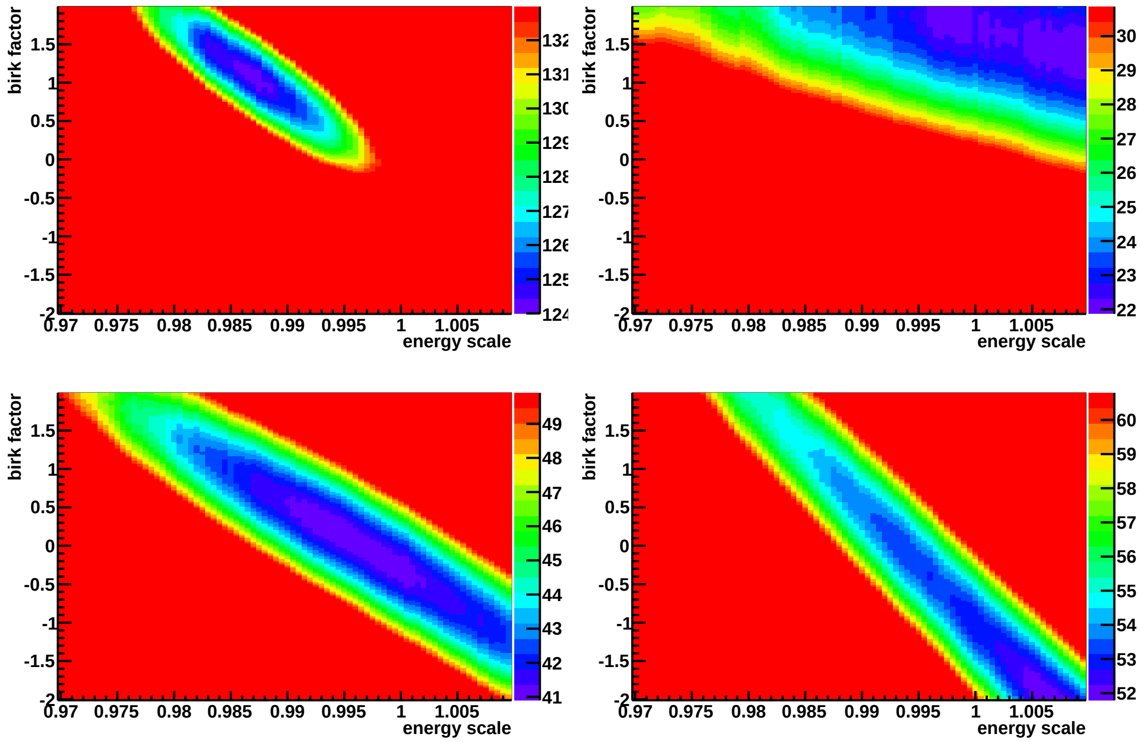

Figure 9: (In color) chisquare surface for subsets of the data. The vertical axis is a scale factor between nominal, one, and two sigma ($1\sigma = \pm 0.012$ mm/MeV). Upper left is the same as previous figure, upper right is a fit to the last plane only, lower left is the fit to planes-from-end 1 through 4, and lower right is fit to planes-from-end 5 through 13. See text for discussion.

The lower left plot is planes-from-end 1 through 4. These planes have modest sensitivity to both Birks parameter and the energy scale. If the Birks parameter was far off in the MC, the data/MC should show a significant trend away from flat 1.0. In the iteration shown here, these data points prefer the MC's default Birks parameter, but with a large uncertainty. The magnitude of the Birks effect for these planes relative to a simple energy scale is different, so the slope of the correlation is different. If plotted plane by plane, there would be a consistent counterclockwise trend in the slope that describes the correlation.

The upper right is the end plane, and has only one plane of data in it, but that plane has a non-trivial dE/plane profile. Along the line of correlation between the two parameters, the chisquare does not distinguish very well, even as the trend continues off the top of the plot. But the best fit everywhere in this range of energy scale wants a higher Birks response.

The best fit obviously balances these three, and there is some tension. It is the plane at the end that gives up the most χ^2, about $\Delta\chi^2 = 8$ out of 10 total, to achieve the best fit. In some sense, this best fit Birks' parameter does not magically, perfectly describe the data. The discussion of systematics will touch on, but not resolve this observation.

5 Systematics

The following table summarizes the effects of systematics on the MC description of the samples and the Birks fit parameter. There is modest tension in the fit, and the systematics involved are interesting in their own right, so the description presented here is more rich than simple shifts in the fit parameter. These are presented relative to the final result, and are expressed below as a change in the dE/plane response. The two most significant modifications of the MC, relative to the default result, produce one-sided changes.

source of effect	end plane	next plane from end	calorimetry	Birks response
± 0.012 Birks' parameter	$\mp$ 3.1%	$\mp$ 1.6%	$\mp$ 1.0%	1σ
coarse G4+MCHit step size	+3.7%	<0.1%	-0.4%	0.24σ
coarse G4+MCHit via mock data				0.7σ
nonlinearity 0.5 x Howard's line	-1.8%	-0.8%	0.8%	0.6σ
$\pm$0.5% cluster energy smear	$<$ 0.1%	$<$ 0.1%	$<$ 0.1%	0.2σ
MEU factor 1%	1%	1%	1%	0.0σ
data - mc discrepancy at best fit	1.25%	0.3%		
uncertainty from the fit				$+0.52\sigma$
				-0.38σ

Table 3: Effect on response due to different tunable parameters and sources of uncertainty. The first three fields are simple shifts in the average dE/plane, calorimetry is the shift in energy for a sample that includes all testbeam protons, and the Birks response is the shift in the resulting fit parameter itself relative to $\pm 1\sigma = \pm 0.012$.

The two most significant effects have a definite direction. Reducing the Geant4 and MCHit aggregator step sizes from very fine to the MINERvA default will only affect plane-from-end 0, and will push the fit to want a lower response (best fit toward lower in the plot). Adding PMT non-linearity (there is none by default) can only be a suppression, and so can only push the fit to counteract with a higher response through less Birks suppression to describe the same data.

5.1 Geant4 small step and MCHit aggregator

The default MINERvA simulation uses Geant4's adaptive step size, which looks ahead to the geometry boundary and to the mean free path corresponding to interaction processes, but not necessarily to the distance that will bring a proton to rest. Secondly, if there are multiple MC steps, under many circumstances our simulation will aggregate them together, up to the strip size, to save 20% file size when writing the MC truth information, and saving a small amount of processing time. Both these have the effect that the dE at the end of the proton will be averaged over too large a step, and less Birks suppression will be applied than the better simulation does. The MC response is slightly too high in this case. As of this writing, the testbeam simulation uses the finer step size, and the default MINERvA step size is evaluated as a systematic.

This systematic is intrinsically, but not perfectly correlated with the Birks parameter. In the last plane, the difference between the default MC and the MC with ± 0.012 mm/MeV Birks parameter is $\mp 3.1\%$, with a smaller effect in the few planes upstream from the end. The difference between the fine step used for the result above and the coarse step size used in the rest of MINERvA is -3.7%, with negligible effect in the planes upstream of the end. If this was all the information available, then switching between the two cases should probably cause a $\sim 1\sigma$ shift downward in the fit Birks' parameter, but the situation is more subtle, the parameter shifts only -0.24σ, to 0.0934 mm/MeV.

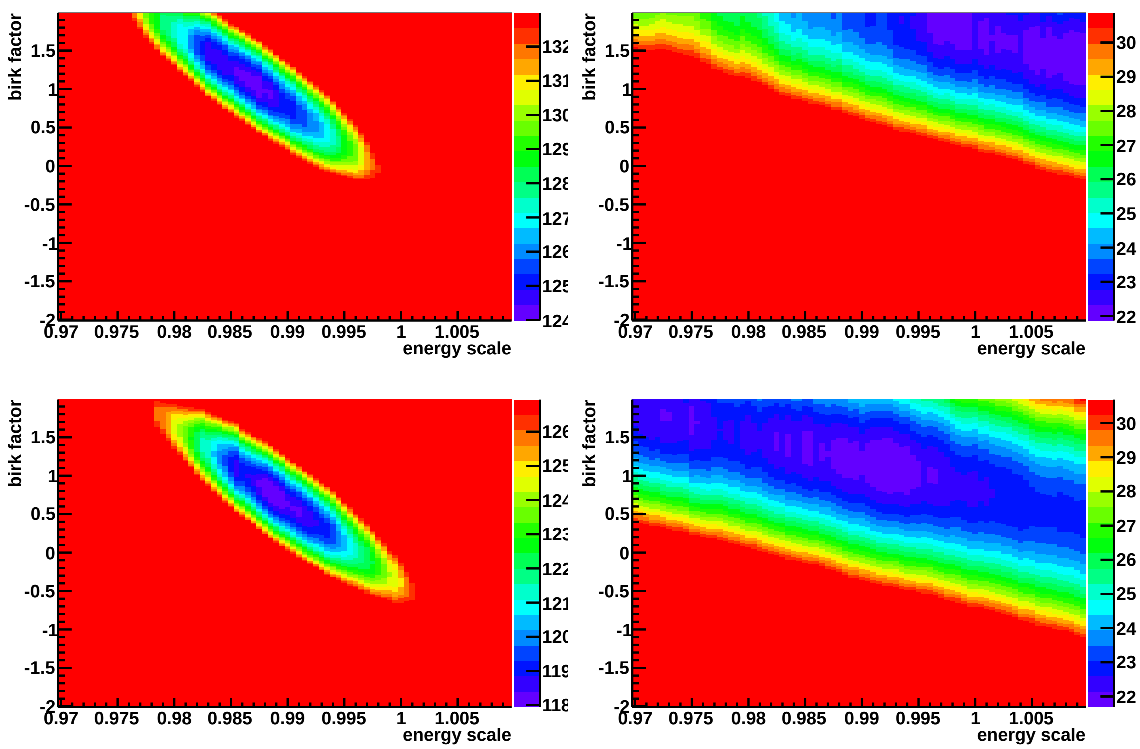

Figure 10: (In color) chisquare surface for the default fit with tiny step size (upper plots) and the fit for the larger MINERvA default step size (lower plots). The full fit is on the left, the plane-from-end = 0 fits are on the right. See text for discussion.

Because more Birks' suppression is needed for the plane from end with the larger step size, the best fit response moves down by 1σ in the plane at the end (right figures), but that also lessens the tension between the subsamples, so the combined fit only moves down by 0.25σ. Making this change back to the coarse step size, reducing tensions, also reduces the χ^2 for the combined fit from 124 for 123 data points (121 dof) to 118. Some tension still remains even in this case, but not much, the individual components at their best individual fits add up to 115 (123 data points, but 117 dof) for both these cases.

Another estimate of this effect is to use a sample of the coarse step MC as mock data and fit it with the fine step size mock data. This gives a shifted result of -0.7σ.

Inconvenient that a supposed improvement to the modeling of protons causes the MC to do a worse job describing the data, but it is testbeam after all. The other systematics described below do not magically cause better agreement. Either there is a systematic or method issue we're not thinking of, or we are at the limit of the Birks' Law approximation. The uncertainty in the fit result should be between the two cases above, and can be evaluated by fitting a mock data sample, we will use the mock data result of 0.7σ as the contribution to the total uncertainty. This only a little larger than the uncertainty in the fit (from statistics and tension), as well as other errors, so it does not destroy the error budget.

5.2 PMT non-linearity

Our simulation assumes that the PMT response is perfectly linear, but in the extreme this is not true, there is a saturation effect which scales with the instantaneous current reaching the anode. Because the testbeam sample has intrinsically 50% more light, we are 50% closer to non-linear effects. Because we have a sample constructed explicitly to have the 25 MeV end of the proton sample instead of the typically 3 MeV muon tracks, we are 8x closer to non-linear effects. Howard, in docdb:7652 and docdb:9311 modeled the saturation non-linearity, and the best model is shown in the green line of the left figure. The green line does not take into account that the light reaching the PMT includes both direct and reflected pulses separated in time, so at most we assume the non-linearity effect would be half what is illustrated in the green line. We do not have our own direct measurement of the non-linearity from either bench test or in-situ analysis. For the testbeam at the end of a proton track, 25 MeV corresponds to about 180 PE which corresponds to about 15 pC.

To study non-linearity effects, the green line is used to calculate a correction for each hit in a cluster, and by combining those we get a correction for the whole cluster. This correction can be applied whole, or a fraction of it can be applied. Because clusters are made up of multiple hits which share the energy in different ways, clusters of the same energy get wildly different amounts of correction. These are illustrated in the right plot in Fig. 11 for a sample of all protons, some which stop and some which interact. This effect is specific to high dE/plane, not high dE/dx, so large pulse heights in the core of EM showers and large response from a particle travelling along a strip will have a suppressed response too, not just the end of stopping protons.

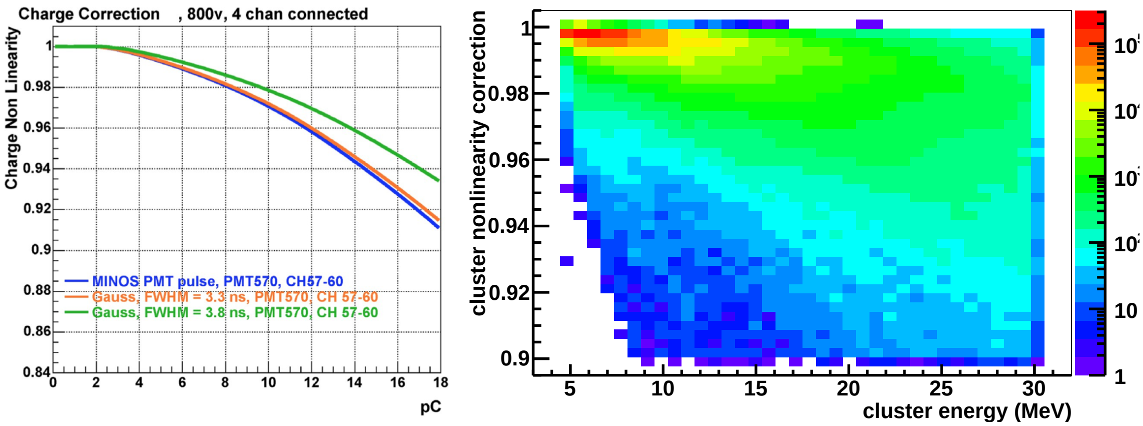

Figure 11: (In color) simulation of the nonlinearity effect by Howard from docdb:7652 and docdb:9311 (left), and implementation of half of the green curve from Howard's plot propagated to the full (interacting and stopping) testbeam proton sample. At 25 MeV, the center of the distribution is 1.5% suppressed, with large fluctuations from 0 to 5%. For energy deposits higher than is seen for stopping protons, the effect can be several percent for some clusters.

Adding non-linearity at the level of 0.5 times the green line causes the response in plane-from-end 0 to decrease by 1.8% and a decrease of 0.8% in plane-one-from-end, and produces a 0.6σ shift in the fit Birks parameter. This is the direction that we think would be a better simulation of the detector response, however we do not know how much is the right amount. Adding this to the simulation (or alternatively using this to correct the non-linearity out of the data) would cause the best fit to move toward higher response (want a smaller Birks parameter in the MC).

In all iterations where this test was done, the χ^2 increased by about 1 for an amount of 0.2 x Howards green line; the data are consistent with zero or small non-linearity, with shifts of 0.2σ. Because this effectively marginalizes over this parameter, take the 0.2σ as the uncertainty contribution to the final Birks' parameter in the context of this nuisance parameter. Nominally this also looks like a constraint on the amount of non-linearity. However, we have not systematically explored uncertainties in this case, and we are experiencing mild tension in the fit. At this time, we plan to use the full 0.5 times Howard's line when non-linearity is explored for calorimetry or other analyses.

5.3 channel to channel smearing

The nature of the fit to the peak of the dE/plane distributions is naturally sensitive to having the channel to channel smearing correct. In fact, the results described above are all done with a best amount of smearing by scanning through a range of anomalous smearing, and this fit method depends strongly on doing that step. This scan is not done with the same strip-level channel-to-channel smearing extracted from the MEU tuning, because that requires a full simulation. Instead, it is done by adding Gaussian smearing at the cluster level; the best χ^2 is obtained with $6.0 \pm 0.5\%$ additional smearing. Moving 0.5% away reveals fluctuations of $+0.08\sigma - 0.38\sigma$, either because the fit is snapping in to the right smearing, or because the smooth effect is hidden by statistical fluctuations in the fit. For the estimate of what to include in the uncertainty on the extracted parameter, I have used half of half of the total spread from $\pm 1\%$ smearing, which gives a symmetric $\pm 0.2\sigma$.

Another estimate using these same data and a linear interpolation between full simulations with 0% and 10% channel to channel smearing suggest an amount of 5.75% would be appropriate for that parameter. In this case, the effect is not perfectly linear, so that estimate carries some uncertainty. Comparing the two methods with similar sized parameters reveals that they do not have identical effects on the widths of the dE/plane distribution, especially toward the end of the proton. In principle this might indicate the kind of modeling error we suspected in the section describing the step-size effect, but as of this writing we have not rerun a test sample with 5.75% smearing to see if the fit and the χ^2 change.

This 5.75% amount of direct c2c smearing is similar to the MINERvA detector, but smaller than the 10% amount that is apparently needed for the cosmics MEU tuning sample. This feature we interpret as revealing an additional fluctuation in the cosmics sample specific to cosmics muons or more likely to the fact that they sample the entire detector while the proton sample samples a much smaller fraction of the front-center of the detector. Or it could be related to the protons coming in horizontally rather than the higher angles typical of our cosmic proton sample. Recall also, the plane-from-end distributions also sample many physical planes and strips, except at the very front of the detector, which we do not use in the analysis.

5.4 Other small systematic effects

5.4.1 MC material model

Changing the MC material model by 2% will change the points along the dE/plane curve that the MC is sampling, affecting the fit energy scale and interestingly the rise at the end. To be specific, if I add planes to the material, then the rise at the end will happen in fewer planes in the modified MC, affecting the fit. Fitting a MC mock data sample with this feature using the default MC yields a shift of 0.2σ in the fit Birks parameter when tested against the Birks 0.120 $\pm$0.024 sample, so 0.4σ as reckoned in the final analysis.

5.4.2 Analysis setup

Bins in fit the fit doesn't use all bins in the dE/plane distribution, only ones that are in the peak. Eliminating a bin at the high or low edge of the peaks changes the answer negligibly, consistent with statistical fluctuations. Adding a bin at the high edge produces a $+0.4\sigma$ shift in the Birks' parameter, adding a bin at the low edge causes some MC bins to have zero entries for particular choices of energy scale. Since the default fit does not include those extra high-side bins, we suppose we are right not to include them and take half this amount as the uncertainty.

Physical planes each plane-from-end is made up from several physical planes. A scan systematically eliminating one physical planes from the distribution and refit. Some have an effect as large as 0.2σ when tested against the Birks 0.120 $\pm$0.024 sample, so 0.4σ as reckoned in the final analysis.

Event selection there are parameters that tweak the proton event selection that have not been extensively studied yet. There is a little ongoing work to see more precisely how clean this super-clean proton sample really is.

5.5 MEU nuisance parameter

The best fit energy scale parameter comes out to 0.9875, or to put it another way, the MC energy is scaled down 1.25%. We do not go change the MEU scale because of this. Describe what this means.

5.6 Total uncertainty

source	+	-
fit	0.52	0.38
G4 + MChit step size	0	0.7
nonlinearity	0.25	0.25
energy smear	0.2	0.2
material assay	0.4	0.4
bins included	0.2	0.0
physical planes	0.4	0.4
total	0.9	1.0

Table 4: Summary and total of systematic uncertainties on the final Birks parameter. The total uncertainty is not very asymmetric, even though its components are, so we will quote a symmetric uncertainty. A new parameter of $+1.04\sigma \pm 1.0\sigma$ corresponds to 0.0905 ± 0.012 mm/MeV.

6 Conclusion

We have used a special high precision proton sample to measure the Birks' parameter for MINERvA scintillator. The value is near the -1σ boundary of the 0.133 ± 0.040 mm/MeV value that is the original default parameter in the MINERvA simulation. The new value is 0.0905 ± 0.012 mm/MeV.

MINERvA analysis using the old parameters, such as the Titan/Resurrection v10r6pXX releases might be able to simply use the special Birks shifted samples to evaluate the uncertainty, and consider that the negative shifted parameter is very close to the best one. Or if the effect on a particular analysis is negligible, then nothing more needs to be done.

Finally, any MINERvA analysis that uses the current default coarse Geant4 and MCHit step sizes (and we might continue to do this), would be better off using the central fit done with MC generated in that configuration, 0.0934 ± 0.012 mm/MeV.