

Magnetic field map for T977 tertiary beam

R. Gran, University of Minnesota Duluth, 10 December 2012, updated 20 June 2016
this is MINERvA technical note TN077 docdb:8067 , Fermilab pub FERMILAB-TM-2628-ND

This note describes the final magnetic field model and a simplified momentum reconstruction being used for the analysis of the 2010 data for T977 (MINERvA test beam experiment) at the Fermilab Test Beam Facility. The results from analysis of those data are published in “MINERvA neutrino detector response measured with test beam data,” Nuclear Instruments and Methods A, volume 789 (2015), pages 28-42, which has open access availability, and is also available as arXiv:1501.06431. The beamline used for this experiment was designed by MINERvA in collaboration with experts at the Fermilab Test Beam Facility. Support for constructing the beamline was provided by Fermilab, including support to MINERvA personnel and also and a wide range of technical help from across Fermilab.

The field map is based on an ab-initio calculation performed by Bob Wands using the ANSYS finite-element software. The map has been compared to data from a campaign to map the field and the model describes the data well; the uncertainties intrinsic to the mapping data coordinate system dominate the comparison. Several systematic studies related to scale factors and the position survey of the magnets have been evaluated and are described. For particles traveling through the good field region of the magnets, the momentum is reconstructed accurately, picking up a 1% error from the field model and 1% error from the wire chamber.

The results in this note were obtained from an early, simple tracking and stepping algorithm and primarily illustrate the properties of the magnetic field map. The followup note, MINERvA TN017 docdb: Fermilab Pub describes the final beamline reconstruction and uncertainties. On the other hand, these results are an improvement over a field map generated directly from a mapping measurement exercise. The direct use of the calculated field map, constrained by the data, is equivalent to using the calculation as the correct (superior) interpolation between coarsely sampled data. These two technical notes are primarily intended as a citeable resource to followup experiments (e.g. LArIAT) who are using variations of this beamline.

1 The magnet field map

Bob Wands made two baseline maps. A final map was created with 2 mm size elements and 2 mm grid spacing for the field points. The calculation assumed a nominal current of 100 Amperes and the BH curve based on the one for MINOS steel. The field map extends from -1500 mm to +800 mm along Z and from -195 mm to +195 mm horizontally, and from 0 mm to 66 mm from center vertically. Only the top half, above the plane of symmetry, is saved to the data file. Because the upstream magnet is rotated and offset relative to the downstream, illustrated in Fig.1, this box doesn't perfectly capture the field of the upstream magnet; it is too narrow in the transverse direction, but it does capture essentially all trajectories of interest.

Some of the comparisons in this document are based on an initial calculation with 10 mm size elements and 5 mm grid spacing for the field points. The coarse grid map cuts off at the same ± 195 mm horizontally, at 0 and 65 mm vertically, and extends from -1215mm to 500 mm along a Z axis that runs perfectly down the center axis of the downstream magnet.

The following three modifications must be applied to bring the as-calculated map to a form usable for reconstructing events in the data.

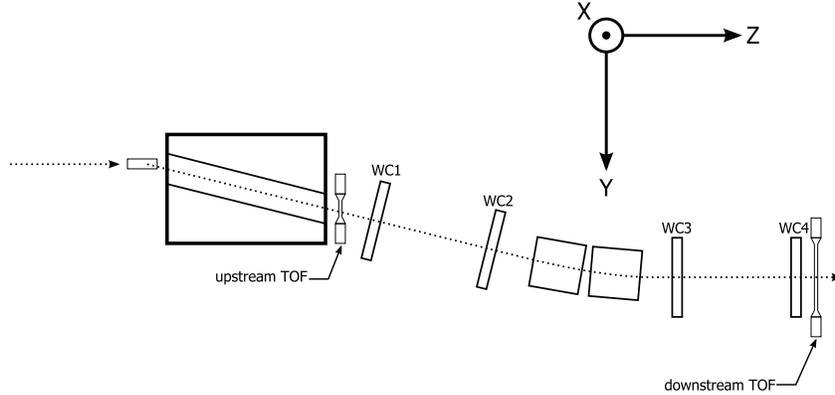


Figure 1: A map of the two magnets (boxes in center), four wire chambers, two time-of-flight systems, and the upstream target and collimator (left).

- A scale factor of 0.9942 bring the model field in line with the measured (during the mapping campaign) field in the center of the magnets. Then an additional scale factor of 1.003 accounts for beam running at a slightly higher current than we used for the mapping data. These changes in current are low enough that small uncertainties in the BH curve should produce negligible effect after this empirical scaling.
- Technically, a padding of zeros around the edge of the field map gives the fit a smoother interpolation to work with. Also, the map has an error code of 1.0 to most of the horizontal and vertical edges. This error code is set to 0.0 at the front and back face, plus the middle region of the horizontal edge. The presence of this error code shows up in the quantity called the error integral as a non-zero value, and is now used to tag the 10% events which had a valid fit but whose best fit trajectory strayed outside the edge of the map. These events will not be used. No error code is written to the horizontal edges outside -1100 and +410 mm in Z (about 40 cm from the magnet center), which is where the principal component of the field has fallen to 5% of its maximum. Especially for the upstream magnet, this mitigates the fact that the magnet is rotated but the strip of field map is not. Particles can sneak in from the side and cross this boundary, though it was just four events out of the whole 20ECAL20HCAL positive sample.
- When the map is read in to the Reco software, a swap of xy convention is made and the vertical structure of the map is reflected around the vertical origin with the appropriate sign flips. Persons reading the reco software code should take care that they know what side of the xy convention swap their favorite bit of code is on.

2 Field integral and reco momentum rule of thumb

The leading order effect on the reconstructed momentum is the field integral from the field map. It always takes me a minute to remember how this goes.

Rule of thumb: the field map (or field scale) with a larger field integral produces a larger estimate of the momentum.

To see this qualitatively, consider that for every event we have measured the bend. A smaller

momentum gives a larger bend, likewise a larger field gives a larger bend. If we increase the field in the reco, we need to also increase the momentum to yield the same bend and fit the data. Imagine fuzzy-thinking faculty members go through this monologue every time, even daily, when he is thinking about the beamline reconstruction.

The following Fig. 2 illustrates the field integral of the first two pions that pass the cuts from the 20ECAL20HCAL Pi+ run; they take two very different paths through the magnetic field.

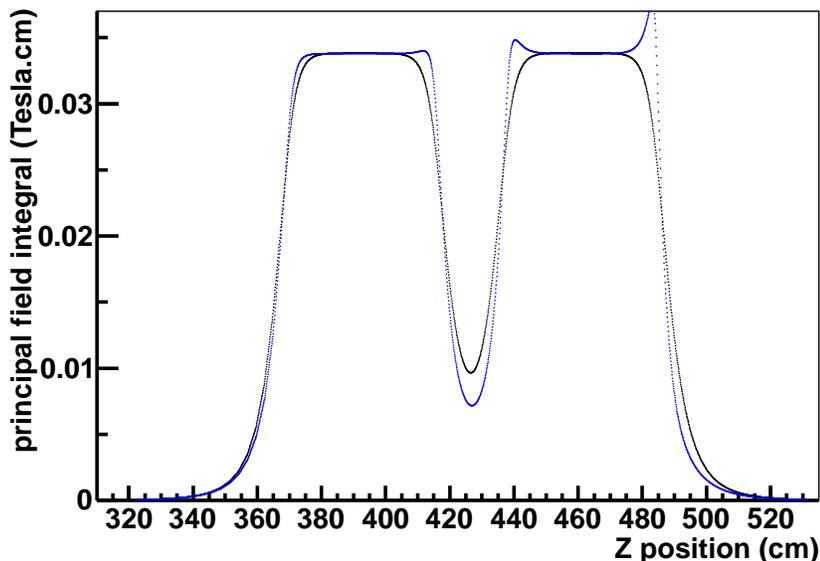


Figure 2: (In color) A demonstration of paths through the magnets for the first two pions in run 209_0001. The vertical axis is the field integral for each mm steps through the field along the path of the field, or multiply by ten to convert to Tesla at that Z position. The blue one (the one with the devil horns on the downstream magnet is event event 209/1/6 momentum 683.4 reco mass 103.8) took a rising path through the magnets while the black one (209/1/1 momentum 585.1 MeV/c reco mass 143.8) took a relatively level path through the middle. Despite the obviously different field experienced along the path, the field integral is the same within 0.3% for the two paths.

The blue one, with what I call the “devil’s horns” is typical for a particle that took a path that was vertically quite distant from the center, about 6 cm by the time it got to WC3 in this example. These are the places in that region where the field changes dramatically, producing spiking high field regions at the longitude edge of the magnet (as in this example) and spiking low field regions at the transverse edge. It was these trajectories through fast changing field regions that were not well modeled with with the spotty, coarse data from the mapping campaign.

2.1 Angle distribution

Lower momentum particles bend much more, and come off their original 16 degree trajectory and end up hitting the detector within a couple degrees of normal incidence, suggested by the dotted line in Fig. 1. High momentum particles don’t bend nearly as much, and squeak in at the edge of the magnet aperture around 10 to 11 degrees from normal incidence, heading to the lower right (+Y +Z) in the view of Fig. 1. These plots illustrate this basic feature of the beamline, and you can imagine the horizontal angle axis is as viewed by the particle, with positive angles going slightly

east toward Chicago.

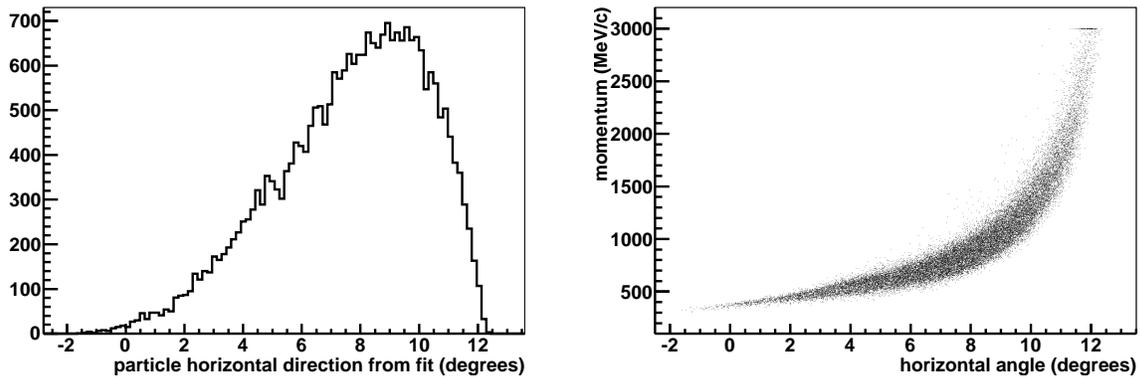


Figure 3: The angle distribution with respect to the beamline axis (practically the same as the detector axis). Left is the distribution itself, right shows how it correlates with the momentum of the particle.

2.2 Field integral is mostly constant

The quantity that matters most for the reconstruction of the momentum is the integral of the principal field for typical particle trajectories through the field. As suggested in the two pion example above, it is nearly constant, even for particles that traverse apparently quite different regions of the field. There is a slight dependence on the momentum of the particle, shown for pions here. The contour scale is adjusted to show the peak, and does not show the long tail, a feature which is illustrated toward the end of this note.

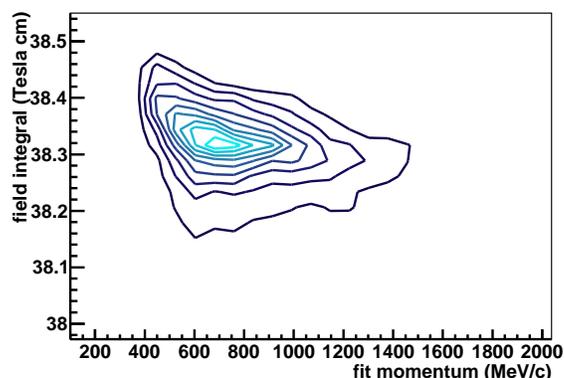


Figure 4: (In color) A demonstration of how the field integral and path through the magnet depends slightly on the momentum of the particle being reconstructed. The lowest momentum particles experience the most significant bend and have slightly longer paths through the magnets. The contour scale is adjusted to show the peak.

The integral is strongly peaked because most trajectories produced by valid triggers go through the middle part of the magnets and because the principal component of the field changes slowly until you get pretty close to the edge of the magnet.

Especially low momentum particles preferentially accrue about 0.2% more field integral due in

large part to their larger bend and therefore longer path through the two magnets. Other variation comes from trajectories being slightly higher or lower or left or right. Later in this note there is a discussion of making a quality cut based on the field integral.

The non-principal components change dramatically in fractional terms, but their absolute size is small, they don't have an effect on the momentum fits, have only a slight effect on the goodness of the fit in the vertical direction, and are of negligible interest for this analysis.

3 Systematics studies

Most of the systematic studies were performed in the following way. Reconstruct the data with the default field. Separately reconstruct the data, changing the field or some other aspect of the reconstruction like WC or magnet positions. Take events that pass the quality cuts in both cases (ignore those that migrate across the quality cuts) and form the difference in *momentum* (new-default) or the fractional difference in momentum (new-default)/default on an event by event basis. Histogram that quantity, and analyze its mean and RMS. Do this separately for low and high momentum, protons and pions. Draw conclusions and write them down. Rinse. Repeat. There is an example plot of this with the first example comparison; the plots for the others are not shown but the resulting bias is presented and discussed.

3.1 Comparison to the old data-driven field map

The new field produces approximately 2% lower momentum than the old data-driven field. Unfortunately, it is not simply a difference between the contents of the data-driven field and the calculated field model. There was a mistake in the old data-driven field downstream magnet that made the field too wide in Z, incorrectly increasing the field integral and therefore the momentum estimates. There is a graphical description in MINERvA docdb:7889 of this error, and a comparison where the new field is compared to the old version constructed from two copies of the upstream magnet, which did not have this error. Sigh. The error is of historical interest, but other aspects of this illustration are important.

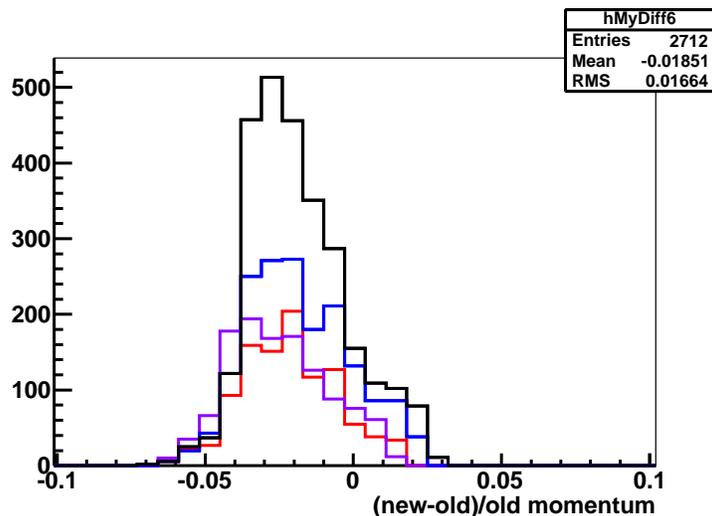


Figure 5: (In color) A histogram of the event-by-event fractional difference between Bob Wands' field (new) and the data-driven field we used through June 2012. The stats box refers to pions with momenta around 600 MeV/c, the other four colors go in increasing momentum up to a bin centered at 1200 MeV/c.

The width of this distribution is the widest of all such systematic comparisons made in this note; the reconstruction resolution depends on the accuracy of the field model, not just multiple scattering and wire pitch. Bob Wands' field is much better. We suppose there is a remnant contribution to the resolution from the intrinsic non-uniformity of the field map.

A second reason the new field is lower: the old data driven field used a simple superposition which gives a higher estimate for the field in between the two magnets, but Bob Wands' field

considers the steel and the coils of two magnets together in one calculation, see next subsection.

3.2 Use superposition of one magnet

This is the first of several special purpose field maps I asked Bob Wands to produce. He made a map of a single magnet, and the fit code is willing to make a superposition of it, just like the way it worked in 2010-2011 with the data-driven maps. The superposition fails to account for the effect of the steel of the other magnet, and produces a distinctly different shape in and near the gap of the magnet. That slightly higher field integral yields momenta that are 0.5% higher than the default map.

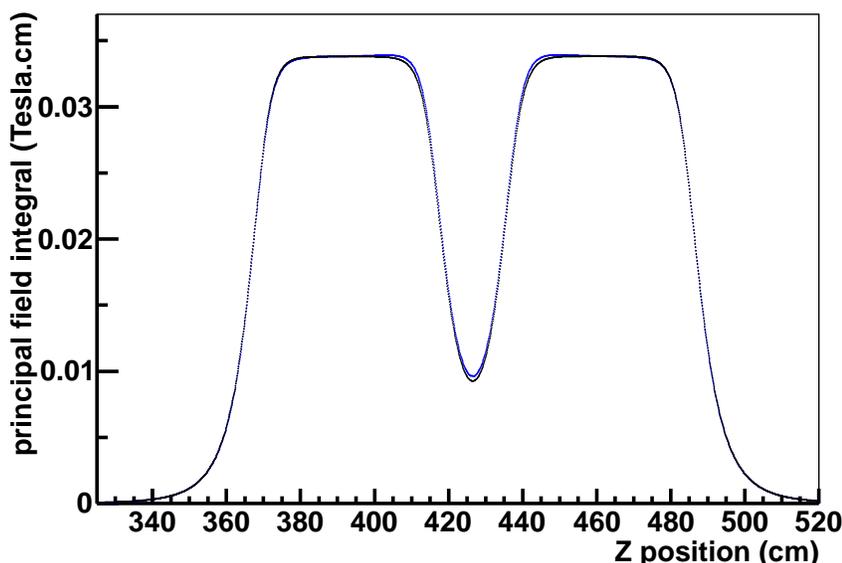


Figure 6: (In color) A demonstration of the difference between superposing two copies of one magnet (blue) and the default map which has the field generated for two magnets placed at that specific geometry. There is an obvious difference in the field in the gap between the two magnets, though it goes in different directions at different places, giving a net 0.5% higher field integral, on average.

3.3 Narrow the magnet 5mm horizontally

This is another special purpose field map, Bob narrowed the horizontal dimension of the magnet, and with this map I re-reconstructed the positive polarity 20ECAL20HCAL sample. The result was a bias of +0.24% (+0.19% at high momentum).

But with this geometry, the field in the middle of the magnet also increased by 0.25% to 0.33988 Tesla. If followed to completion, my procedure would have adjusted the field from the field map downward to match the as-measured field in the middle of the magnet. Therefore, the actual bias from a possible geometry error this large is only -0.01% (-0.06% at high momentum).

3.4 Narrow the magnet 5mm vertically

For this special purpose field map, Bob narrowed the vertical dimension of the magnet. Remember, this is the shortest dimension, so 5mm is significant. Out of the box it has quite a large effect, producing a 2.8% bias (more like 3.0% for high momentum particles), but a major part of the reason is that the field strength in the center of the magnet is higher by 3.1% with this geometry, 0.34943 Tesla. Since the procedure adjusts for this using the as-measured field in the center, the resulting bias would actually be -0.3% if the geometrical error was actually as large as 5mm.

3.5 Extend the magnet 5mm in Z

For this map, Bob took the geometry and current that he used and repeated it making each magnet 5mm longer (i.e. add 2.5mm on each side of center). This represents an overestimate of the uncertainty coming from how well the technical drawings match the as-built magnets, but we've produced maps that exaggerate this source of error to make it easier to identify the possible effects on the reconstruction.

By itself, making this change shifts the momentum higher by 0.9%, simply because it increases the field integral by that much. The field at the very center of the downstream magnet for this modified map is 0.33922 Tesla instead of 0.33903, but scaling the central field to the as-measured field is part of my procedure, so in fact I would also have scaled the field (and the field integral) down by this tiny 0.06%, which is so small I would not have bothered.

In the three examples above, we generated maps with 5mm shifts, though that magnitude is larger than we think the actual error could possibly be. But this is one we have suitable measurements to test, see next subsection.

3.6 Longitude extent vs. Mapping data

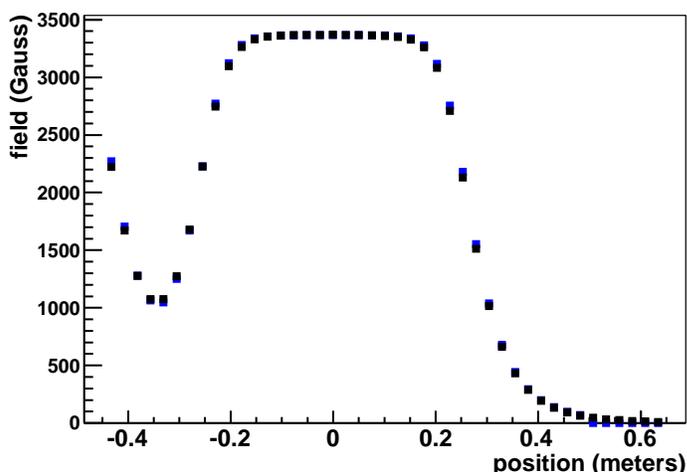


Figure 7: (In color) A demonstration of possible discrepancies in the longitude dimension of the magnet. This is from the 1H zip of the downstream magnet. Its hard to see, but at the right and left downgoing edges of the magnet, the blue (Wands' field) is higher than the black (data) by a few tens of Gauss, corresponding to between 2 and 3 mm of additional width in the blue field.

This figure shows the same kind of comparison I made with the two calculated field maps, but this time made with data. Its hard to see, at the edges of the field, the points on either side around 1500, 2200, and 2700 Gauss, the blue point (Wands' field) is higher than the data point, suggesting that the calculated field is slightly too wide. I quantified the implied increase in width for three zips, two near the center transverse (zips 1H and 3H) and one at the edge (zip 1L). Averaging the results from the center zips yields the conclusion that Wands' field is 2.6mm longer in the transverse direction total than the as-built magnet's field

Combined with the previous subsection, this yields an error that Wands' field is 0.5% too high. This result can also be obtained with pencil and paper by assuming that the extra width is effectively in the middle of the magnet with the max 0.3375 Tesla field, thus $0.26\text{cm} \times 2 \text{ magnets} \times 0.3375 \text{ Tesla} / 38 \text{ T.cm} = 0.5\%$.

On the other hand, the result for the 1L suggests that Wands' field is 1mm too narrow. But the fact that it contradicts two zips in the center, that the discrepancy at the edge could just as easily be due to a transverse offset, and that the pegboard alignment is imperfect leads me to trust the center transverse data for this particular systematic.

I've repeated this analysis for the other, upstream magnet and find negligible discrepancy. I also reanalyzed the data with modest simple shifts in the vertical and horizontal direction and can bring the measurement above about a factor of two closer to agreement. I conclude that the manufacturing tolerances and the uncertainty in shifts and rotations of the pegboard during mapping allow me to say that the two magnets are consistent with their technical drawings within 2.6mm, and we should take the 0.5% as an uncertainty but not a correction with a smaller uncertainty.

3.7 Shift the magnet closer in Z by 5mm

A final special purpose field map I asked Bob Wands to produce. He took the geometry, based on the alignment survey of the locations of the magnets, and purposely shifted the upstream magnet 5mm closer to the downstream magnet. The result is a shift of -0.17%. As above, this is an over estimate of the possible error, the alignment survey in all other respects appears to be good to 1mm or so.

3.8 Comparison to the coarse grained Wands field map

The fine grid map, our default, returns negligibly different values than the coarse grid map. The mean reconstructed momentum shifts by 0.01% rms 0.15% with the fine grid yielding slightly lower momentum estimates. This is tiny and consistent with the addition of 30 cm of very small field at the very edge of the map. There is a bit of an ambiguity, not only did the output grid get smaller, but the element mesh size also got smaller to match. To the extent that it matters at all, probably its the element mesh size that made the difference.

The fine grained field map doesn't take much longer to process though it requires more memory, about 2 GB per instance, to hold the field map as a ROOT TH3F. This is easily within the capabilities of even a recent laptop computer. Unfortunately, this exceeds the memory limit for grid running at Fermilab. I run these interactively in a screen-based shell, and when split into three sets they take less than 5 hours total. I didn't mention it until just now, but all the comparisons here are based on the 20ECAL20HCAL positive polarity data set. I spot checked a couple with the 20TRAK20ECAL and protons separate from pions, with little difference.

3.9 Use a different steel B vs. H curve

A final pair of special purpose field map was made, one using a different B vs. H curve. All the above field maps were produced using what Bob Wands calls the “kjs” curve measured for the MINOS detector steel. This one requires only a small correction to yield the measured field at the center of the magnet (in fact, this is why it was chosen). Bob reran the simulation with the B vs. H curve measured by the University of Wisconsin for the CMS endwall steel.

A technical matter, a fresh map of each was produced with the same fine grid points for the output, but to save computation time a coarser mesh for the FEM input. Because they were both done with the same coarse mesh, its an apples to apples comparison, and as described above the coarse mesh makes very little difference.

The CMS B vs. H curve produces a field whose center is 2.4% below the kjs BH curve, and whose field integral is 2.75% below. Since I adjust the field to match the measured field at the center of magnets, the choice of B vs. H curve has only 0.35% effect. Equivalently, the magnet steel is not substantially in saturation given these currents.

3.10 100 Amp current, stability, and other small errors

A number of these are discussed in internal MINERvA document docdb:5676. The instrument that measures the magnet current and reports it to us via ACNET indicates that the system is extremely stable. We had a current of 100.3 Amps during mapping, and took data at 100.6 or 100.7 amps, both numbers reported with the same instrument. We correct the mapping data upward by 0.3% when we do the reco, and in this sense there is negligible error, either from the calibration or from drift.

The field in the center of the two magnets is identical to a fraction of a percent. The uncertainty in the principal component due to rotation of the probe is measured to be negligible, though the uncertainty in the non-principal components is more significant. The Hall probe plus the ADC plus the DAQ computer as a combination was calibrated in an NMR machine for MIPP and the data Doug Jensen sent to me suggest it is known to be 209.7 ± 0.2 ADC counts. Is there time drift in the Hall probe? I hope not.

3.11 Non-principal components of the field

We are now including the non-principal field components in the field map and the fit. Because they are so weak, they do not affect the fit very much. Setting the non-principal component to zero changes the fit less than 0.1% bias with an rms of 0.05% to 0.1%. The fit chisquare, which includes the distance between fit line and measured points in both X and Y, degrades when I do not include the non-principal components; average by 0.2329 while the rms of the event-by-event chisquare change is 1.6, more for high momentum events. Ordinarily, the fit chisquare distribution averages 1.8, but is only 1.2 for events with momentum greater than 1 GeV/c.

3.12 Field cutoff

The code assumes the field has dropped to zero by the location of WC3. This is close, but not exactly accurate. Physically, WC3 is a distance of 72cm along the Z axis from the center of the downstream magnet. In Bob Wands’ field calculation, we went out to 80cm along the magnet’s Z axis which is 79cm along the physical Z axis. At 0.72 along the magnet Z axis, Wands calculates

the field is 6 Gauss and falls to 3 Gauss by 80 cm along the magnet Z axis, integrating 0.00364 Tesla cm in that span. Compared to the typical 38 Tesla cm for a path through the magnets, this is 0.01% and is totally negligible. If I include a rough estimate for the additional field beyond 80cm not included in Wands map, that adds another 0.003%.

4 Fit technique, some variations, and more systematics

We have tried several variations in techniques for fitting out the momentum of the particle in order to be sure we are not missing something that matters. The default simplified fit minimizes 2D distance between the observed WC hits and the fit path, but does so weighted with terms that convolute a wire pitch resolution with multiple scattering estimates for which WC4 has a worse resolution than WC3. Two variations on this are to use equal weights for WC3 and WC4 and to use only the horizontal distance between hit and path, not the 2D distance. A second major version of the fit will use a Kalman Filter implemented by Will Bergan and Josh Devan at William and Mary is described in a separate note

The first variation is equivalent to doing a simple linear fit to two points with equal errors and treats neither point as special. Compared to the default multiple scattering error terms, the simple fit changes the momentum for pions between 0.1% at low momenta and 0.3% at high momenta, with an rms between 0.2% and 0.3%. That is delightfully small.

The variation with just the horizontal constraint gets a second change: the multiple scattering term in the chisquare calculation is reduced by $\sqrt{2}$ which is just a guess and may not be exactly the right thing, but even with this reduction, the chisquare comes out much lower, an average of 0.5 instead of 1.84. I think this reflects actually that the degrees of freedom are such that we fit out the horizontal portion of the multiple scattering and the remaining piece in the chisquare must be describing the mm-level offsets of the wire chambers. The fit momentum itself is changed less than 0.01%.

4.1 More about the chisquare

The chisquare for this fit looks like the chisquare for two degrees of freedom. That seems approximately right; its good that its average is not 10 nor 0.1. But its not super obvious to me that it is exactly right, considering the chisquare-like quantity I am forming includes both horizontal and vertical positions. On the other hand, the freedom in the fit to push the vertical trajectory around is very slight. What I think is a reasonable guess is that we should have the chisquare for one degree of freedom as if we were fitting a slope and intercept to three data points, plus another degree of freedom from the vertical components which are not fit. The combination of issues below, plus mm-level WC misalignments degrade the chisquare distribution.

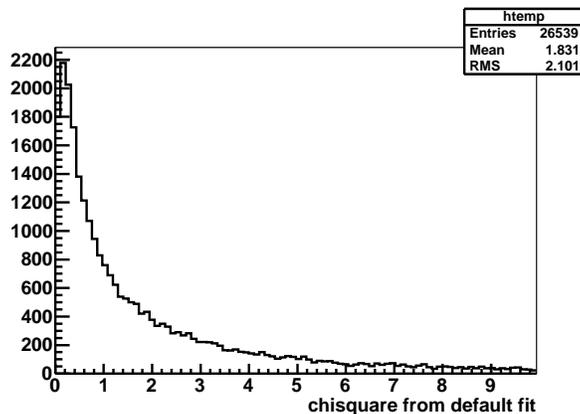


Figure 8: The chisquare distribution for the default fit. See text for discussion.

The chisquare is a little better for high-momentum particles, and also a little better for pions

compared to protons, but a big part of this is artificial. The multiple scattering estimate being used to construct the chisquare is not energy or (especially) species dependent. We could relatively easily implement the energy dependence, with more difficulty also the species dependence, using the seed (based on the pT kick) estimate of the momentum and TOF. Instead we'll let the Kalman Filter effort run its course, and back-patch this feature if needed.

Finally, I can check these things with the high resolution range-out proton sample. There I can implement the expected resolution from Will Bergan docdb:8002 (with its energy and species dependence mentioned in Section 4) and I see that I need to add roughly 1% in quadrature to smear the otherwise perfect MC to look like the data. There are a couple good reasons. I'm not sure the WC wire pitch is included in Will's estimates, but it might be. The non-uniform field can add significant smearing on top of multiple scattering; going from the old field to Wands' field reduced this additional quadrature piece from about 2.5% to 1% for example. Finally, Will's estimates more properly apply to his Kalman-filter style reconstruction, not the simple reconstruction described here; the latter may in fact require a different resolution function.

If I really want to mess with these issues, including the Energy and Mass dependence and a toy model for the actual dof in the face of perfect alignment but correct fluctuations, and a toy model or full MC that propagates and reconstructs events to derive the resolution function that properly goes with this field map and fit technique. This seems like a low priority.

4.2 Field step size

Our standard stepper goes in 1mm steps and completes the job in a comfortably short amount of time; it takes a couple hours to complete the processing on four cpu cores. This step size is also well matched to both the 2mm elements and 2mm grid used in the ANSYS computation of the field and is reasonably smaller than the smallest changes in the field. Changing this from 1.0mm to 0.2mm has a negligible effect on the fit, momentum bias $-5.532e-07$ rms $3.413e-05$. We probably have room in the error budget to increase the step size, but the economics of the computation don't require it.

4.3 WC offset

The fit we use to extract the momentum does not treat the WC equally in a couple ways. It insists that the line defined by the activity in WC1 and WC2 is perfect despite the presence of multiple scattering, and it emphasizes the position at WC3 more than WC4 by an amount consistent with multiple scattering. Compare this to the two alternative, a simplest linear fit with a Pt kick and two slopes in the horizontal plane, or a full Kalman Filter that knows multiple scattering at each WC.

Below is a comparison of the resulting fit momenta in response to a +1mm shift (east, toward Chicago) for each WC in turn. It is expressed as a fraction (new-old)/old. Survey errors in WC2 or WC1 have the largest effects.

These values can also be used to back-of-the-envelope the effect due to multiple scattering in the context of this momentum fit technique, turning the "resolution" in horizontal into a resolution in momentum.

Memo to myself (not done) understand why do + shifts in both WC3 and WC4 move the momentum fit the same direction? That's not the behavior I would get if I used the Pt slope technique. I guess they both wag a line from the middle of the magnets. A similar wonder applies to the point I make next.

Momentum	WC1	WC2	WC3	WC4
600 MeV/c	0.7	-1.1	0.2	0.1
800 MeV/c	0.9	-1.3	0.3	0.15
1000 MeV/c	1.1	-1.6	0.4	0.2
1200 MeV/c	1.5	-2.1	0.5	0.25

Table 1: Biasing effect on fit momentum of a +1mm transverse (east, toward Chicago) shift of each wire chamber. With this style of fit, a shift in WC2 and WC1 have the largest effect.

In addition to changing the fit momentum slightly, these shifts affect the apparent offset between fit and actual positions at WC3 and WC4.

offset (mm)	nominal	WC1	WC2	WC3	WC4
mean at WC3	-0.43	-0.71	+0.18	-1.01	-0.23
mean at WC4	+0.75	+1.31	-0.49	1.87	0.35

Table 2: Effect of a +1mm transverse (east, toward Chicago) on the apparent horizontal position discrepancy at WC3 and WC4. Some illustrations of this are in docdb:7947. If it were possible, we would want to construct an adjustment of shifts to the WC that take out the offset in the nominal configuration.

As expected, shifting the WC by a millimeter causes the apparent discrepancy between fit and actual position at WC3 and WC4 to decrease or increase. Its not a trivial millimeter because the fit doesn't treat them uniformly; the effect on WC4 is just over a mm and WC3 is just over half-millimeter if we shift WC2 or WC3. That decreases to half and quarter if we shift WC1 and WC2, the ends wag the fit less than the middle. But fundamentally, we are supposed to imagine what offsets to apply to WC1-4 in order to reduce the nominal discrepancy to zero. The answer is between half and two if we thought exactly one WC was off, but on average a combination of shifts around 1mm would be adequate to explain it. Neither this, nor the previous table seem to allow us to infer if a particular WC is off significantly.

4.4 Shift the whole magnet system in XYZ

The result is a bias in the momentum of less than 0.2% in all three cases. Boring.

5 Some other observations

Most of the studies and conclusions in this document were performed during and after I had made substantial progress with both pion calorimetry and range-out proton analyses, plus other progress the group has made with calibrations and software infrastructure for testbeam analysis. Several aspects of the beamline reconstruction stand out because they seem to be relevant (or show themselves to be negligible) for some aspect of analyzing the data in the detector.

5.1 Particle direction as it exits the field

We have two solid estimates for the direction of the particle as it exits the field region since the field integral is negligible beyond WC3. One is directly from the data: the line that connects the hits in WC4 and WC3. The other is the direction of the trajectory at the best fit momentum. These are naturally in good agreement. The typical dy/dz slopes range from 0 to 0.2 (equivalently from 0 to 10 degrees), and the disagreement in those slopes has a bias of -0.0006 and is a pretty-good Gaussian with an RMS or sigma of 0.0014.

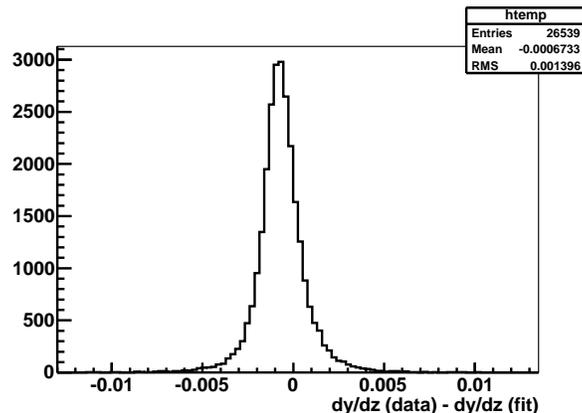


Figure 9: The difference between the dy/dz slope from the measured points in WC3 and WC4 compared to the value from the fit. Imagine that $\text{atan}(dy/dz)$ is the angle in degrees whose distribution was shown earlier.

This comes into play when we project the particle trajectory onto the detector. The particle goes approximately one meter beyond WC4 to the detector, so if had a slope of 0.1 then it would travel 10 cm horizontally before it hit the detector. An error of 0.001 (coming from the choice of slopes or the apparent resolution available) corresponds to an error of 1mm in the projection. For all conceivable projection purposes, this is plenty accurate, it doesn't matter which one we should choose, though the measured slope is presumably more correct.

5.2 Pileup at +19ns

There is a band of pileup at $\text{TOF} = \text{nominal} + 19\text{ns}$ which appears on our standard momentum vs TOF plots. This occurs when one particle arrives in an earlier bucket and triggers TOF1 before the particle that triggered the rest of the beamline and TOF2. Because we put in tight proton mass cuts, most of this background is not included in our analysis, except for the little bit that coincidentally overlaps the rest of the triggers between momentum of 500 and 600 MeV/c. If we think its needed for a particular proton analysis, we can try to perform a pion subtraction to these protons (the pions have that same momentum).

5.3 Best and loose mass cuts

The beamline fit used here has several indications of fit quality, beyond a much simpler MINUIT fit-is-valid criteria used previously. If the best fit value went into the field error region, it might not be a good fit. If the chisquare is bad, meaning the hits in WC3 and WC4 are quite far from the best fit line, it might not be a good fit. Finally, the field integral typically takes on a value of 38 Tesla.cm, with some variation; if the best fit trajectory carves out a significantly different field integral, it might not be one of the best trajectories. I should mention, as implemented the error integral is strangely defined, because its a mashup of two different purposes. The values here are appropriate to the current implementation.

Tight cuts: chisquare less than 10.0, error integral less than 0.0002 (practically zero), field integral between 36 and 40 Tesla.cm.

Loose cuts: chisquare less than 20.0, error integral less than 0.01, field integral between 34 and 42 Tesla.cm.

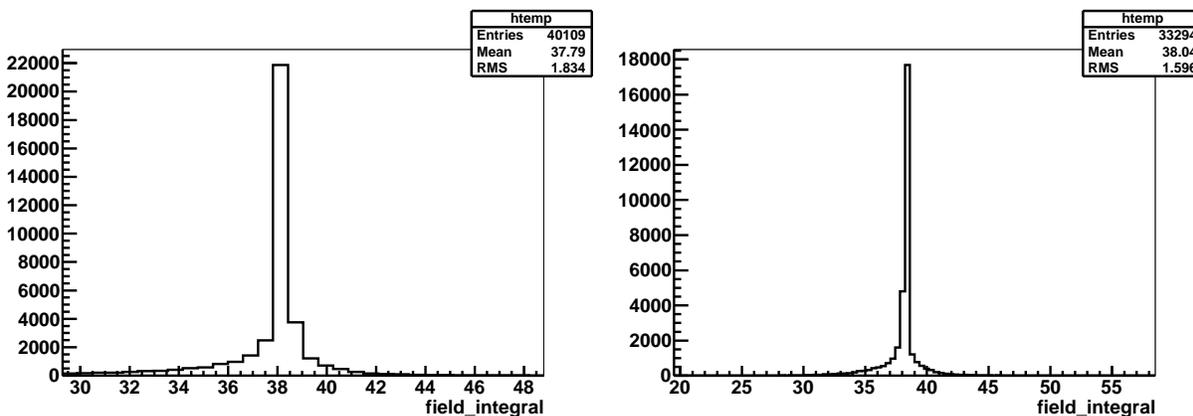


Figure 10: The field integral distribution (units of Tesla.cm), before other quality cuts (left) and the field integral distribution with loose cuts on the error integral and fit chisquare (right, in effect the N-1 cuts distribution). The axis range is different, but note the RMS improves, largely because the left-side tail is absent with the fit quality cuts.

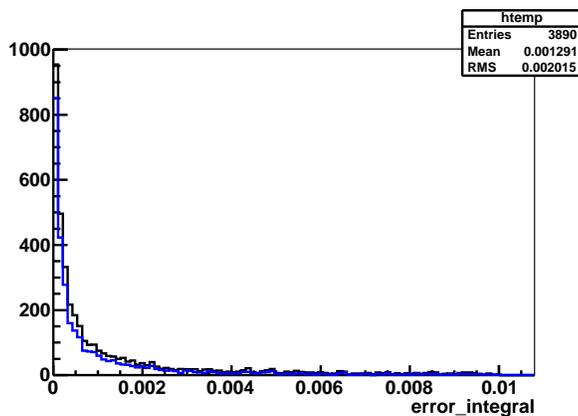


Figure 11: The error integral distribution, before other quality cuts (black) and the error integral distribution with loose cuts on the field integral and fit chisquare (blue, N-1 distribution).

In the context of analyzing detector data where there are additional clean event and pileup cuts, choosing the tighter cuts reduces the sample size by about 15%. For most analyses I recommend we start with the loose cuts, and apply tighter cuts as part of your systematic error studies. If an analysis is shown to be really sensitive to this, the tight cuts can easily become a new default. Based on several quantities in the range-out proton and pion calorimetry analyses, the loose cuts seem to produce a stable result.

6 Summary of systematic uncertainties

Source	error and units	modifies what quantity
WC alignment low-p	1.3%	momentum
WC alignment high-p	2.6%	momentum
longitude magnet dimension	0.5%	momentum
other uncertainties in the field map	0.5%	momentum
All other systematics	<0.5%	momentum
above in quadrature p <800 MeV/c	1.5%	momentum
above in quadrature p >1000 MeV/c	2.7%	momentum

Table 3: Summary of the errors incorporated into the analysis of the data.

An earlier version of this table discussed the mass in the beamline as a systematic related to the beamline. As of this writing, we start the data-driven ParticleCannon at a point one meter behind WC3, which corresponds roughly to the middle of the magnets. In this way, we are taking the beamline momentum estimate as if it were the momentum at this point and trusting GEANT4 and the mass model for the air and WC windows to account for the energy loss. Further, when we analyze the data, we don't take this particle energy (momentum), we take a momentum and species dependent Bethe-Bloch corrected energy (momentum) at the front face of the detector. So this part of the uncertainty is no longer considered a beamline uncertainty, and is not either considered to be 2 MeV big.

Parting comment on systematic uncertainties. We are probably okay with these errors, and they might improve some with the effort to put a Kalman Filter in place. But they could be reduced if we had (or a future user of the tertiary beamline would) designed the magnet off data in such a way that we could constrain the location of the wire chambers to half-millimeter or better.

The only error related to the magnet field in this table comes from the longitude dimension of the magnet. I should think more carefully if there is some additional metric that suggests an uncertainty in the field map. None of the studies in this document suggest one. Pulling another error out of the comparison of Wands' field and the mapping data could be tried again. In my experience trying to tweak coordinate offsets and rotations, I found I could arbitrarily match any set of zips at the edge of the magnet aperture, but not simultaneously match all such edge measurements. The quality of the data and its alignment don't argue for a serious discrepancy or permit a good constraint.

A future effort to map the magnet more seriously would be better than guessing. However, this applies to the few events that travel the edges of the detector. In a previous analysis docdb:6891 I reported that only 4% of all events that passed the cuts at that time travelled through the poorly known field regions, and they picked up a potential momentum bias of 1% to 4%, which translates to a bias of 0.2% averaged over all events. The conclusions of that analysis carry to this new field without change, except that we suppose Wands field is more accurate than the one extracted from

the mapping data, plus now we are implicitly making tighter cuts on the sample through the field integral constraint.