The 3m arc dipole (BA) - modeling and comparison with sparse experimental data

Jay Benesch

Abstract

The 3m arc dipole is the most widely used in CEBAF, comprising arcs 7, 8, 9 and halls A and C. Pascal Vernin and Mike Tiefenback made extensive measurements of central field and BdL on the "ninth dipole" which is part of the hall A energy measurement system. Beam-based measurements were made of average focusing (quadrupole) term provided by the magnets in the arcs. CEBAF dipoles were specified in TN91-095. Total field error, systematic quadrupole and random quadrupole were specified. Jefferson Lab does not have the equipment used by high energy labs to measure dipole multipoles, usually called "moles".

Simulation of the 6 GeV and 12 GeV CEBAF accelerators require knowledge of the multipole terms of the magnets. These drive emittance growth, x-y coupling and halo formation among other things. Detailed modeling of the BA dipole has been undertaken with Vector Fields Opera version 12 with the intent of learning the limitations of the software and demonstrating that the models can accurately reproduce the sparse experimental data. Such agreement makes credible the use of the multipoles derived from such models.

The effects of manufacturing tolerances on multipoles is NOT considered in this paper. A graduate student, Nicolas Ruiz, is working on this under my direction. He will prepare a paper on this work in partial fulfillment of his masters degree.

The Software

Vector Fields Opera comprises two pre-processors, TOSCA for magnetostatic calculations, and a post-processor for evaluating the TOSCA solution. Version 11 introduced the Modeller as the second pre-processor option. This program produces 3D models directly rather than "extruding" them as done with the old pre-processor. It is based on the ACIS kernal (http://www.spatial.com/products/acis.html) and so is likely written in a modern language. TOSCA itself is written in Fortran. Portions of it likely date back to punch cards. Commands are space-sensitive so it is difficult to read and write command files. Text output is all caps. Commands are not case sensitive, so input files need not shout.

The Modeller, unlike the original pre-processor, has no limit on the number of nodes and elements. Since it is now possible to buy a quad core workstation with TB disks and 16 GB RAM for less than \$6000, much less than the software license price, models which are much more detailed than previously possible can be created and evaluated in reasonable calendar time. The tradeoff between time creating the mesh for maximum efficiency and minimum size versus computation time has changed dramatically in favor of simple dense meshes and 4-16 hour computation time versus days refining mesh to keep computation time under a week. It is generally more productive to spend about four hours creating a model with simple dense mesh and then run several variations on it rather than refine meshes for days as I had to in the Eighties. The Modeller provides many facilities for mesh refinement in areas in which fields change

rapidly, and I use these. I simply don't go to coarse mesh in the bulk of the pole face, for instance, but retain a moderate size there, typically 2.5mm, because it's not worth my time.

This paper will document my learning process about the software as well as results on final models.



🛛 Vector Fields 🔣

Figure 1. Basic BA magnet model. Air surrounding steel is 3 times X direction, 2 times Y direction and 1.3 times Z direction. Tangential field is set to zero on the boundary. 8cm mesh on the boundary.

This model was run at 30A intervals from 30 to 300A. Central fields and field integrals were computed using two methods, nodal interpolation and source integration. Vector Fields claims the latter is more accurate. This turns out to be true only for regions with low field far from the magnet proper. Between the poles nodal interpolation is much closer to reality than integration. As one moves the boundary farther away from the steel and conductor the "integration" value

approaches the nodal interpolation value between the poles. To keep the number of mesh elements acceptable the mesh size on the boundary must increase to \sim 75 cm.



Figure 2. Same mesh as figure 1, looking in Z direction from magnet midplane. Dense mesh in the air between the poles and in the first cm of the pole face where saturation effects are important on the field in the volume the beam occupies. Less dense mesh in the area occupied by the coils and still less dense in the remainder of the steel. Mesh size of 2.5mm was chosen for the air between the poles from the sampling theorem. Multipoles will be evaluated on a 1cm radius circle. 2.5mm mesh gives ~25 quadratic elements on the circumference, adequate for the allowed decapole error component. "Allowed" refers to pole symmetry: dipole, sextupole and decapole all have positive pole at 12 o'clock and negative pole at 6 o'clock. Quadratic elements are used in the steel and in the air between the poles to improve accuracy.

At the top of page 2 I wrote of the effect of boundary location on field values generated by nodal interpolation and integration. In table 1 I show the differences in central field and straight line JBdl for models with different background sizes. These were run with 90A in the BA to speed

computation. The model has 10 million elements and 7.6 million nodes. Straight line integrals JBdl have 350 points on a line 25 cm longer than the steel, 175cm total length. The integration method of evaluating fields takes about one hundred times longer than nodal interpolation. The last column will only be shown for two cases for that reason.

x range	y range	z range	cond	B nodal	B integration	BdL nodal	BdL integ.
-51,43.3	0,28.5	0,221	full	-1763.154	-1757.562	-268072.060	
-82.4,74.8	0,28.5	0,221	full	-1763.252	-1757.979	-268086.220	
-90.2,82.6	0,35.6	0,212	full	-1763.201-	-1759.982	-268076.440	
-90.2,82.6	0,35.6	0,212	half	-1763.085	-1759.562	-267893.373	
-161,153.4	0,56.9	0,340	half	-1763.015	-1762.189	-267874.648	
-161,153.4	0,142.2	0,340	half	-1762.984	-1763.350	-267866.504	-267948.71
-161,153.4	0,142.2	0,340	full	-1762.991	-1763.465	-268032.356	-268090.02

Table 1. Effects of background size and conductor area on calculated magnetic field. Cond(uctor) full and half refer to area, where full is the actual area of the coil pack and half an artificial coil pack with half the area and twice the current density. Vector Fields support personnel recommended putting more elements between the conductor and the steel. The easiest way to do this without expanding the number of elements is to halve the area. Models with smaller coils solve a bit faster. The 66 ppm difference in "nodal" central field between the third and fourth lines of the table (excluding header) are not significant for known purposes.

I conclude from the table above that using "integration" for field evaluation is inappropriate in most cases. The computation penalty far outweighs any gain in accuracy which might be obtained. Putting the boundary condition far from the sources is desirable in all cases and doesn't impose a computation penalty if one increases the mesh size to match. The difference between the nodal values in the first and last line in the table is small enough that I will continue to use models with "close" boundaries rather than repeat 1000+ hours of computation. Models built after this investigation have far boundaries, of course.

To check on the effect on multipoles of the two post-processing methods I am running in background a five day job using the integration method to evaluate the field on 900 circles along the beam orbit in the 4m arc 10 dipole model. There is noise on skew multipoles in the nodal-evaluation results, in the 5-50 mG range, which may be numerical or mesh derived.

Comparison with ninth dipole data

On the next page are two graphs. Figure 3 compares calculated values to ninth dipole measurements with and without a correction for hysteresis loop effects. The BH curve provided by ME does not have the effects of the "setup" hysteresis loop used in the machine. Attempts to modify the BH curve to match the measurements have failed. Simple left (right) shifts of the curve decrease (increase) the central field while raising (lowering) the end field, leaving JBdl unchanged. Figure 4 shows the ad hoc correction applied to the model results. The flat portion of the curve is the EPICS magnet map value at 0 amps. The linear decline which begins half-way through the current range was chosen to improve the model/measure fit at the associated currents. 180A is the first model which has any steel volume in the non-linear piece of the BH

curve. The amount of non-linear steel increases through 300A. One would expect the "memory" of the steel to fade as it approaches the maximum current.



Tosca vs ninth dipole measurements





Remanent field correction

Figure 4. Remanent field correction applied to blue points in figure 3 to yield red points.

The model/measure ratio for the uncorrected 300A model is 0.9999. The corresponding field is 5.83 kG. Almost all the dipoles in the 12 GeV CEBAF will be operating at or above this field even when final energy is reduced. At 210A the uncorrected ratio is 0.9981 - still good. No data exists above 300A so agreement at higher fields is not assured. The hall A power supply maximum is 300A.

Multipoles

Normal quadrupole and sextupole terms calculated at 2.5mm intervals along a beam arc approximately 175cm long are shown in figures 5 and 6. Integrating the values and dividing to get G per cm of length, one arrives at 0.63 G. Beam measurements by Johan Bengtsson and Mike Tiefenback gave 0.3 G at 80% of this dipole fields. *Yves Roblin result: Quadrupole field map errors of ~115G would account for the difference*. Sextupole has not been measured. Skew quadrupole may be inferred from the typical values of two skew quads in the machine, in 8S and 9S. These correct both the residual from the linacs and skew from ~550m of dipoles. The total needed in 8S and 9S to null x-y coupling ranges from 500-1000G. Skew quadrupole due to dipole roll and manufacturing tolerances may thus be comparable to the normal quadrupole term. Better data will be available in fall 2008 after linac skew quads have been upgraded so full compensation may be accomplished for the C50 modules running at high gradient in the linacs. Skew quads in 8S and 9S will then compensate only for the dipoles.







Figure 6. Same as figure 5 except for vertical axis change so one may see magnitude in the body of the dipole.

Yves also looked at the multipole results I have for the BE magnet in arc 3 at one field. Calculated quadrupole term is -3.5G per cm of length versus -1G from beam based measurements in the Optim models. A change of 115G in the arc QAs again restores the optics after the TOSCA value is substituted for dipoles. The QA field map is known inaccurate at this tolerance. The beam-based measurements were made in a manner which was intended to greatly reduce dependence on QA field map, so this may be a coincidence. Other sources of the discrepancy might be feeddown from sextupole or orbit offset in the dipole in experiment vs model.

Magnet Line integrals

Dave Douglas referred me to TN97-019 about body and end gradients. I ran the Opera "table" command to generate By and Bmod on a grid of points in the midplane, x=[-5,5] y=0 z= [0,174.75] at 2.5mm intervals. It took about ten minutes to compute the table with nodal interpolation. Fearing time requirement for integration, I cut x to [-2.5,2.5]. It took over 1500 minutes (>25 hours) to compute that table. I ran these on a 150A BA model because that current is cited in a handwritten note on one of the TN 97-019 figures. It may be for a BE magnet instead of a BA - I can't quite read that. I also ran a model at 201.3A, corresponding to 3645 MeV in arc 9. The values for 201.3A are all precise multiples of those for 150A - the model is perfectly linear in this range. I've sent Dave all the numbers for his perusal.

Here I show straight line integrals for the 150A case, normalized to value at x=0, y=0. Figure 7 is a coarse view of nodal data. One can see some non-physical variations at the 50ppm level around x=-1.5 and 2 cm. In figure 8 I show nodal and integration results. The integration result is much smoother - at the cost of over a day's delay. Not worth it.



Figure 7. Straight line integrals on 150A BA magnet, normalized to value at x=0, y=0



Figure 8. Comparison of nodal interpolation and integration results at finer scale for inner 5cm in which the beam generally resides. Sagitta is 3.7 cm. Greatest discrepancy is ~150ppm.

Default Warnings

While some of the defaults, like the background multiplier 10, proved to be good advice, others must be changed and/or are misleading.

The default surface mesh spacing is 1cm. This must be increased, to \sim 75cm, for the large volumes created by large background multipliers.

The default "absolute" tolerance for the Newton-Raphson non-linear iteration is 0.001. I have been unable to find any information on the units. I generally tighten it to 0.0002, forcing one or two more iterations. I increase the maximum number of iterations from 21 to 40 because I got bitten once by a model which required 22 iterations and had to be re-run, losing 15 hours. Since increasing it, I've had several more models require 22-24 iterations. The higher the field and therefore the more non-linear the steel, the more iterations.

The misleading default I know about is the tolerance for current density. There may be others. The default value is zero. When one declares the units to be cgs in creating the analysis database, this is changed to 10G without notification. Iterations on coil field calculations will stop when the differences are less than 10G. This seemed huge, so I sent a flame to VF. I was told that the effect on calculations would be small. I reran several BA cases with 0.1G tolerance and found this to be the case, as shown in table 2 below. Still, I now use 0.1G.

current (A)	B(0,0,0), 10G tol.	B(0,0,0), 0.1G tol.	ratio, 10G/0.1G
30	-587.74	-587.73	1.000018
60	-1175.47	-1175.45	1.000018
90	-1763.20	-1763.15	1.000025
120	-2350.87	-2350.83	1.000018
150	-2938.39	-2938.33	1.000019
180	-3525.48	-3525.41	1.000020
210	-4111.72	-4111.63	1.000023

Table 2. Comparison of central field values with different current density tolerances.

Advice to the magnet modeler and his/her supervisor

- 1. Use a big workstation: at least four CPU cores, 16GB RAM and two TB disks, mirrored so you're unlikely to lose data. I didn't buy enough disk space. I'm storing output files on a 1 TB USB drive and hoping that it doesn't fail. I'm keeping the input files mirrored so I can rerun the jobs if it does. If you've \$20K, two 3.2GHz quad cores, 64GB and four TB disks.
- 2. Run version 12, 64 bit, so RAM, disk and patience are the only limits on model size. This requires a network license. JLab's individual licenses can be changed at no cost.
- 3. Buy two network licenses for the Opera suite (TOSCA, post-processor) and one Modeller license for each user. This will allow the user to run two calculation jobs in background while creating or post-processing a third in the foreground. With the fourth CPU core one can run long post-processing jobs, like my multipole job, in the background too. This will cost twice what the workstation in (1) does but a lot less than the annual salary and benefits of the user. The productivity increase will amortize the cost in under a year.
- 4. Accept the default background size for X and Y, 10. Adjust Z background ratio so the Z boundary is about the same distance from the coil as the transverse boundary is from the steel. If the solved field shows more than a few gauss at the boundary, expand it and resolve.
- 5. Mesh the air between the poles at ~2.5mm maximum extent due to sampling theorem limit if you want multipoles to icosapole.
- 6. Split the pole piece in two so one can mesh the bottom cm or so more densely than the rest. 2.5mm is a good value for that bottom cm. 1cm is OK for the rest of the pole. The return steel can be meshed at 2-4 cm without change in field between the poles. If one is interested in saturation of the return steel, 1cm mesh.
- 7. Mesh the 1cm high faces of the pole piece at 1mm using "face properties". The mesh generator grades to the body value of 2.5mm over two elements. This provides more elements where the field changes are steepest and saturation is most likely without increasing the number of elements unduly. Meshing the bottom pole face at 1mm may also be desirable to reduce noise in multipoles but the number of elements doubles. Calculation time becomes excessive (weeks). One is better off meshing the steel and air discussed in (5) and (6) at 2mm than face meshing the bottom of the pole at 1mm. Errors in skew multipoles at the 10⁻⁶ level of bending field can be obtained with the 2mm mesh. With 2.5mm, about 6*10⁻⁶.
- 8. Layering is another option for faces, but I've found it less than useful because one can layer in only one direction. The transverse mesh size is taken from the value in the adjacent face. It doesn't help to put a few layers at 200 microns each in the pole if the transverse dimensions

of the volumes created are still 2.5mm. One still gets unphysical field jumps between first and second tetrahedron from an edge because of the large transverse dimensions. Face meshing as discussed in (7) works better for me; YMMV.

- 9. Mesh the volume with the coils so there's at least one full element between coil and steel, preferably two. Face meshing on the rest of the steel at 2mm would help but the number of elements gets too large to tolerate.. Coil size is just as easy to adjust and as seen in table 1 does not have significant effect on the result.
- 10. Cutting planes every 20-25cm in Z help the mesh generator keep on track.
- 11. The "check" function is worthless. It has never found an error in a model which surface meshed but failed on volume mesh. When the volume meshing fails, a new characteristic is assigned to the offending volume. This may be located by "picking by property Error". One can then play with the maximum mesh for that region to see if a slight change will allow volume meshing to succeed.
- 12. Table 1 and other results not included herein convinced the author that one needs one model per steel configuration, evaluated at multiple current densities, rather than models for each coil pack and each core. This reduces the number of models needed by almost a factor of two. It makes sense physically: the field in these magnets is defined by the pole configuration and its saturation, not the direct field from the conductor. This is true even at 14 kG in the 10cm pole width, 98 cm long dipoles, where the pole edges are saturated for almost a cm.
- 13. Don't bother with integration mode except far from all sources. It's not even useful for crosstalk evaluation because the magnets which affect each other in the spreaders and recombiners are no more than 50cm apart - not far enough.
- 14. The next few items assume models are built with poles parallel to the X axis and the beam is bending in that plane. I place the origin at the center of the pole in X and Z and the midplane between the poles in Y. Models must have some elements in positive (x,y,z) space for the background element to work. I've generally put the original C return in -x.
- 15. The descriptions used by Opera for model symmetry definition are ones only someone who has just finished Jackson could love. If one wants reflection across the Z=0 plane because one is meshing only half the length of the magnet, one sets tangential field boundary (=0) for XY plane. If one wants reflection across the Y=0 plane, one sets normal field boundary for the ZX plane. This forces all skew terms to zero, BTW. If one has an H magnet and can also reflect at the X=0 plane, one sets tangential boundary on YZ plane.
- 16. For circles in XZ plane with negative X origin, rotation angle around Z 90 degrees and "around new Y" 90 degrees. Circle starts at -90 degrees and moves more negative.
- 17. For circles in XZ plane with positive X origin, rotation around Z -90, around new Y -90, and circle starts at -90 degrees and moves less negative, e.g. -90 to -85.8 degrees.
- 18. Buy a copy of Amperes, 3D hybrid electromagnetic software from Integrated Engineering Software, <u>http://www.integratedsoft.com/</u>, so one can compare the results with Vector Fields. It is called hybrid software because it mixes two methods of calculation. Within sources, coils and steel, it uses finite elements. It uses boundary elements on the boundaries of the sources to calculate the fields in free space. One does not have to mesh air out to ten times the magnet extent to get consistent results. Initial solution time is a bit longer than a pure finite element code like TOSCA but post-processing takes the same time as Opera nodal interpolation with the accuracy promised, but not delivered, with Opera integration mode. I suspect integration mode was added to Opera to respond to the competitive threat of hybrid

codes. Hall D collaboration is comparing TOSCA to the ANSYS electromagnetic module; that result will be interesting.

Acknowledgments

Dr. Klaus Hoeffner of Vector Fields UK and Robert D. Pillsbury, Ph.D. of Sherbrooke Consulting, Inc. (Arlington, VA) provided technical support on Opera/TOSCA, the latter under contract to the former.

Mike Tiefenback suggested modeling the BA magnet in the first place because he had accurate experimental data available for it. Yves Roblin suggested the method used to derive the remanent field compensation as well as making the cited calculations. Many conversations with Mike and Yves improved this work. Nicholas Ruiz had questions which led me to clarify three sections.