Three-Way RFOFO Ring Field Comparison -- Part 1
Steve Bracker
December 14, 2002
RFOFO Field Map 28.doc

In this writeup I compare six field maps for the RFOFO Cooling Ring. For all six field maps, the following ring configuration parameters remain the same:

24 solenoids centered on the Coil Placement Circle, +-6 degrees from cell boundaries
Solenoid length = 0.5 meters
Solenoid inner radius = 0.77 meters
Solenoid outer radius = 0.88 meters
Solenoid current-turns = 5239850 amp-turns
Solenoid tilt angle = 53 milliradians
Coil Placement Circle 0.10 meters larger in radius than Nominal Beam Trajectory
Circumference of the Nominal Beam Trajectory = 33.0 meters

All six fields are generated for the same pseudo-random distribution of 5000 test points within the active area of the RFOFO ring, which for these purposes is defined to be all points within 0.25 meters of the Nominal Beam Trajectory. The location of the test points are recorded in a disk file; see RFOFO Field Map 27.doc for details.

The differences:

Field Map A: Generated with the Ole Miss Biot-Savart field generator. Each coil has 50 x 11 = 550 1 cm x 1 cm turns. Radial layers of turns are centered at R = 0.775, 0.785 ... 0.875 meters from the coil axis. 300 current segments per turn. The Biot-Savart field is relatively straightforward to code; only location points must be transformed from tilted coil coordinates to ring coordinates. The fields appear directly in ring coordinates (which is the form our GEANT simulation wants them), and inasmuch as the ring coordinate system is cartesian, it is straightforward to compute curls ande divergences for ensuring that the field conforms to Maxwell's Equations. The primary disadvantage of the Biot-Savart field generator is that is slow. It would takes weeks of computing time to generate an adequately fine-grained field map grid, for example. However, it is perfectly practical to apply it to small ensembles (thousands) of points, and to check the performance of other field map generators against it.

Field Map B: Generated with the Ole Miss coordinate rotations, but finding the field for each coil using 11 current sheets (radii as above) and a local version of BSHEET / ELLTHREE / RC modified from one provided to us by Rick Fernow. BSHEET is a much more difficult routine to understand; there are behaviors around certain "magic regions" in a coil that have to be handled with some care. (Magic regions include the sheet center and any point on the axial extension of the sheet cylinder, even well beyond the sheet itself.) There are fairly complex issues of numeric precision and convergence criteria that must be looked after. Assessing the accuracy of the field calculated is not entirely straightforward. However, BSHEET is very much faster than the Biot-Savart field generator; it can compute an eleven-sheets-per-coil field grid for a cell (1 cm lattice) in just a few hours. That fact has made the 11-sheet BSHEET field our standard "workhorse" field generator. A field map grid based on it is used in our GEANT simulations, and it is the generator of choice whenever more than a few thousand field points must be calculated. By comparing Field Map B with Field Map A, we can determine whether the "tuning" of BSHEET to the task at hand has been successful.

Field Map C: Generated by interpolation within a 1 cm x 1 cm x 1 cm grid of points, where the field at each grid point was computed a version of BSHEET very similar (but not quite identical) to that used to generate Field Map B. Now mostly of historical interest, this field map is included here only because it remains the field map used by GEANT, but will be replaced within a few days by . . .

Field Map D: Generated by interpolation within a 1 cm x 1 cm x 1 cm grid of points, as in Field Map C. However in this case the version of BSHEET used to generate the fields at the grid points is exactly the same as that used in Field Map B. The field at an arbitrary point is interpolated using multi-dimensional linear interpolator FINT. In GEANT, the field grid for one cell (1/12 of the ring) is loaded into memory from a binary disk file by routine ***LoadField*** (in a few seconds). As particles are stepped around the ring, the needed field values are computed by routine ***FindFieldAnywhere***, which rotates the particle location into an equivalent position in Cell 1, calls ***FindField*** to compute the field there, and then rotates the field components back to the cell where the particle actually is. ***FindFieldAnywhere*** is very fast, about 30 times faster than calling BSHEET directly, but since the grid is fairly coarse and the interpolator is only linear, the question naturally arises as to whether Field Map D is sufficiently accurate. By comparing Field Map B with Field Map D at a suitable sample of test points, we can answer that question.

Field Map E: Identical to Field Map B except that each coil is represented by only three current sheets instead of the eleven used in Field Map B. By comparing Field Maps B and E and their respective discrepencies from Field Map A (Biot-Savart) we can assess whether it is appropriate to save some execution time, when generating large field grids, by reducing the number of sheets per coil.

Field Map F: Identical to Field Map B (and E) except that each coil is represented by 22 current sheets instead of the eleven used in Field Map B. It is appropriate to wonder whether the field map is actually converging as the sheet count goes up. Comparing Field Maps B, E and F helps answer that question.

Scaling the Fields

Each field may be viewed as a ***field shape*** times a ***field scale***. In a working ring, if the field ***scale*** were found to be slightly different than expectations, one would fix this by increasing or decreasing the coil currents. (Remember that there is no iron in this design.) Finding that the fields had a ***shape*** consequentially different from expectations might be much harder to fix. It might require anything from realigning the coils to introducing correction elements to facing catastrophe.

All field generating procedures discussed here are approximations. In the Biot-Savart generator, the simulated turns are larger than real turns, and the circular arcs of the coils are approximated by many short straight segments. In the BSHEET generator, there are fewer sheets per coil than there would be layers of winding, and the calculation of the elliptic integrals is done iteratively. Hence it should not be surprising that there are small differences between the fields generated by the two methods and their variants. To the degree that these differences are only small differences of field ***scale***, they may be ignored. To the degree that they are differences of field ***shape***, they are potentially more serious.

In the comparisons that follow, I have not hesitated to apply a scaling correction whenever it is helpful to minimizing the discrepency between two field maps. What remains is the ***field shape discrepency***, which is what we care about. However, I emphasize that the scaling corrections applied are very small; the largest one used below is 1.00015. It is a measure of the substantial consilience of the various field generating methods that adjustments this small remove discrepencies that otherwise would completely dominate those relating to field shape.

If a field is scaled, then all three components must be scaled by the same amount.

How Good is Good Enough?

To what accuracy must the fields be known? At this point, the quick answer is "well enough so that remaining errors in the field map do not materially affect the performance measures of competing ring designs". Yes, but what does that mean in terms of actual errors in the field? Is "no more than a 1% error at 95% of the points and no more than 3% error at any point" good enough? Should we strive for errors ten times smaller? One hundred times smaller?

At some point in time, if a cooling ring is ever built, there will have to be field sensitivity studies done. Magnet builders cannot wind coils or position coils to perfect accuracy, and if the cooling ability of the ring is jeopardized by minute variations in the coils, comparable to those expected from imperfect winding technology, surveying errors or temperature fluctuations, then the ring is too fragile to be built, regardless of its performance in an ideal world. The same kind of sensitivity studies will shed light on the accuracy required of the simulation. But at the moment, they do not exist, as far as I know. Hence for now it is reasonable to hold the simulated fields to a high standard, something like "accurate enough so that, lacking sensitivity studies, there seems to be a vanishingly small likelihood that field map errors will be a major contributor to uncertainties in the assessment of the ring's performance." I have somewhat arbitrarily taken this point to be per-field-component standard deviation of 0.0001 tesla (1 gauss) for a sample of points randomly distributed over the active (particle-accessible) volume of the ring.

I am willing to be argued with about this; I imagine (perhaps wrongly) that most people will consider this too stringent a requirement, and its pursuit a waste of time. However, as we shall see below, we are probably already at or near that level, without having done anything terribly heroic. In any case, if the project (as we hope) moves from conceptual design toward engineering, we shall have to place this matter on firmer footing than considered opinion, even the opinions of those who have considered such matters for a lot longer than I have.

Inputs and Outputs of the Comparison Process

All code is written in Fortran and is located in C:\Magnet and its subdirectories. All data are written in C:\Magnets and its subdirectories. Most data are in the form of "vanilla text files" but field grids are converted to binary files.

Project ***DefineTestPoints*** generates a data file of standard test points. The standard test point file is ***C:/Magnets/StandardTestPoints/StdTestPoints1.txt*** containing locations of 10,000 test points. The first 5000 test points are used in the comparisons that follow. For each point, the location of the point is given in Ring (cartesian) coordinates X, Y, Z, and in Accelerator coordinates S, b, X' and Y'. (S and b are redundant but it's convenient to have both.) There is one line of text per test point.

All field-generating program read the standard test point file, compute the field at each point, append the field components to the input line, and write it to an output file. At Ole Miss the fields are in Ring Coordinates and units of tesla, but other coordinates and scales may also be used; the field components are converted to Ring Coordinates before being submitted to the comparison program.

The results of test runs of lasting interest are given a letter (as in A-F above) and stored in ***C:/Magnets/StandardTestPoints/TestFieldReportA.txt*** etc. In some cases there is an accompanying ***TestFieldDiagA.txt*** with certain diagnostic information pertinent to the reported results.

The comparison program is run on pairs of results files. The output of the comparison file is a ***report*** which contains histograms of unscaled field component discrepencies and a ***diagnostic*** file containing point-by-point location / field value / field differences, one line for each point. The diagnostic file (which has proven to be the more useful of the two) is imported into a standard Excel worksheet, where interactive scaling is performed if needed, discrepency scatterplots are produced, and a few basic statistics are computed.

I have found it most useful to characterize the discrepencies component-by-component in the form of scatterplots, such as BX difference vs BX. This makes it easy to distinguish scale-dependent discrepencies (BXdif correlated with BX) from others, and to ferret out any "magic regions" in which the discrepencies are unusually high. As the overall measure of discrepency I use the standard deviation of BXdif for the whole ensemble of test points.

The Biot-Savart Field A vs the Eleven-Sheet BSHEET Field B

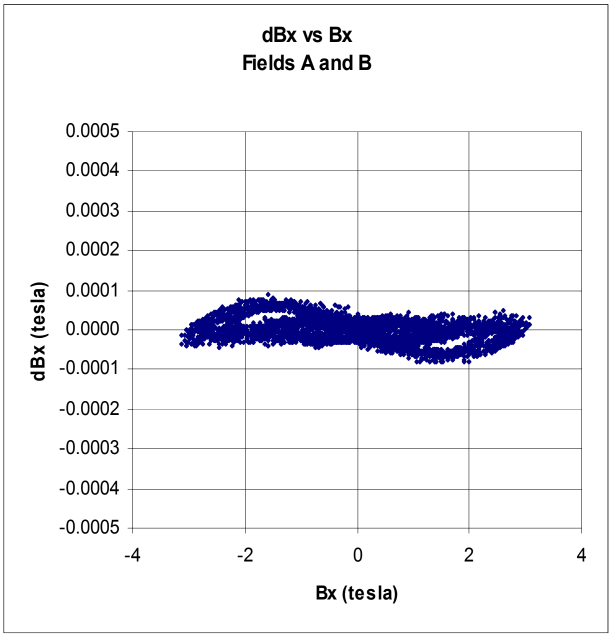

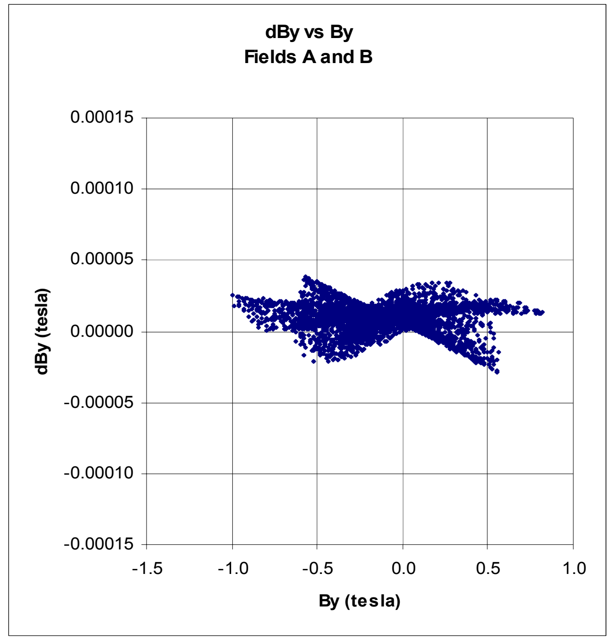

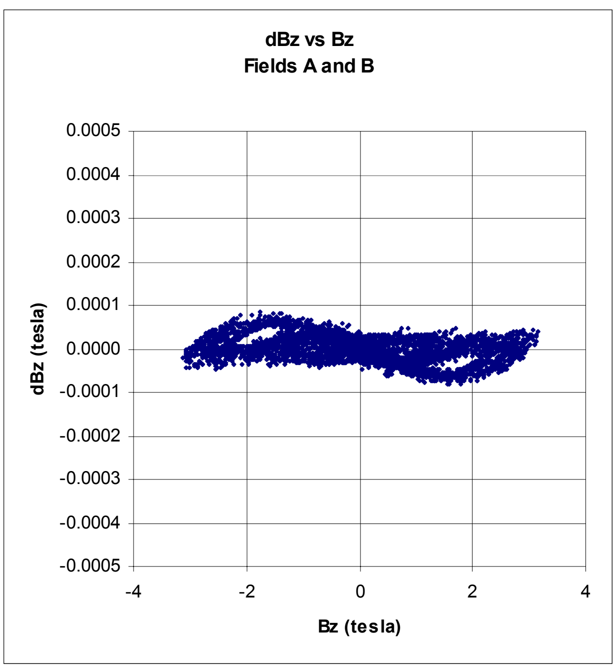

0.0000285	StDev dBx
0.0000092	StDev dBy
0.0000287	StDev dBz

The scaling factor is 1.00011. There is some remaining structure to the field differences, but the discrepencies are very low -- these two very different methods of generating the field for points around the entire ring, neither exact, nonetheless give highly consistant results.

The Eleven-Sheet BSHEET Field B vs the Interpolated Field D

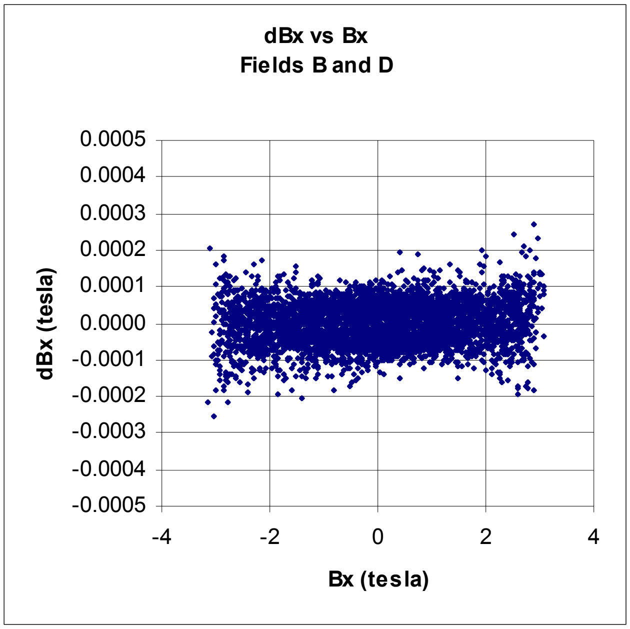

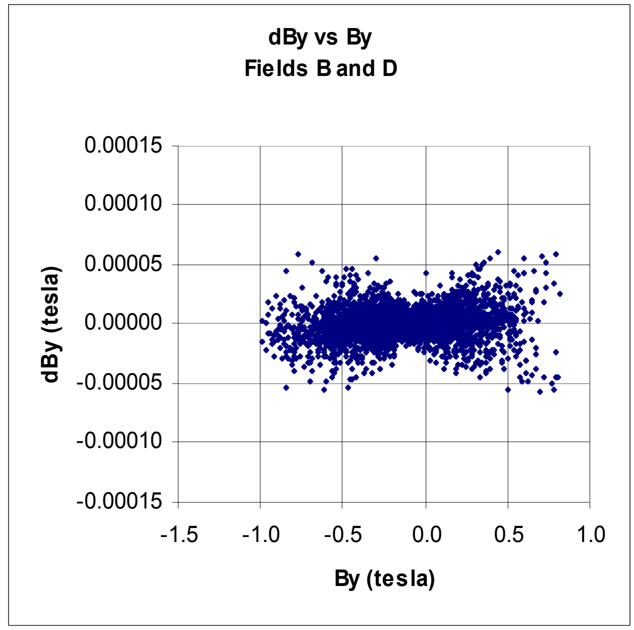

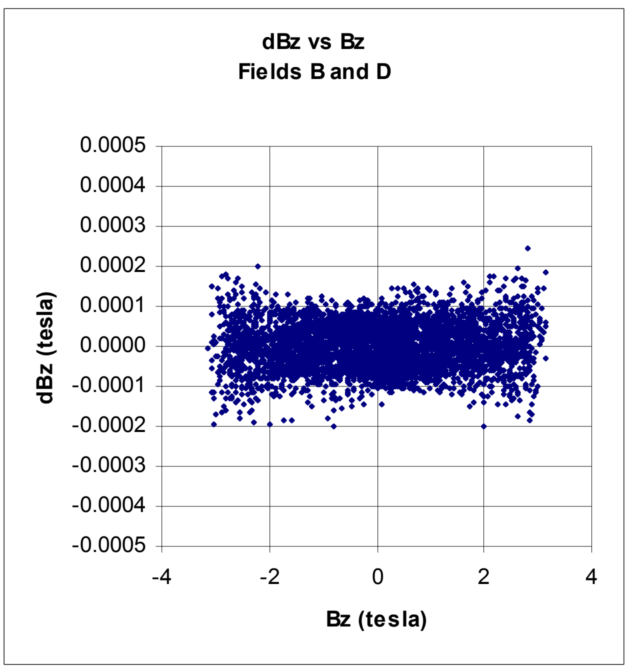

Bx1-SBx2
-0.00000091 Average
0.00005770 Std Dev
0.00013577 FWHM

By1-SBy2
-0.00000026 Average
0.00001161 Std Dev
0.00002733 FWHM

Bz1-SBz2
0.00000037 Average
0.00005757 Std Dev
0.00013547 FWHM

How does the field interpolated from the BSHEET-derived grid compare to the BSHEET field itself? Again, the results are very encouraging; very little accuracy is lost by resorting to the (much faster) interpolation-in-grid. The memory required to hold the grid is substantial, however.

The Three-sheet BSHEET Field vs the Eleven-sheet BSHEET Field

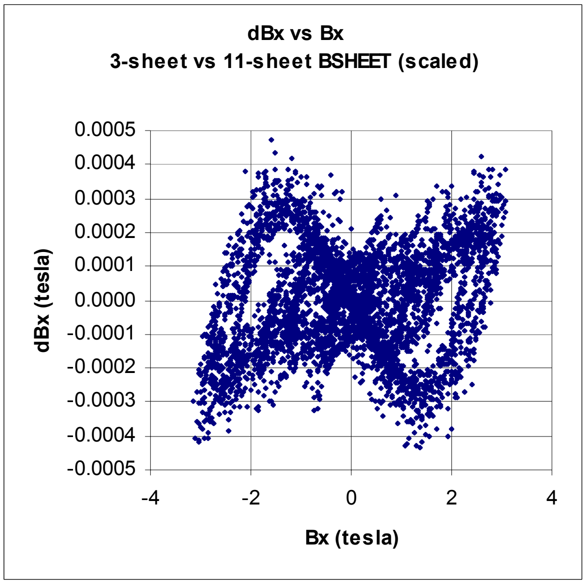

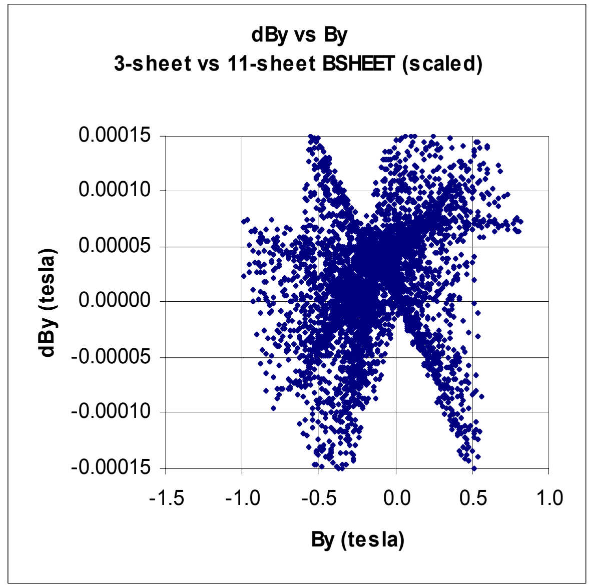

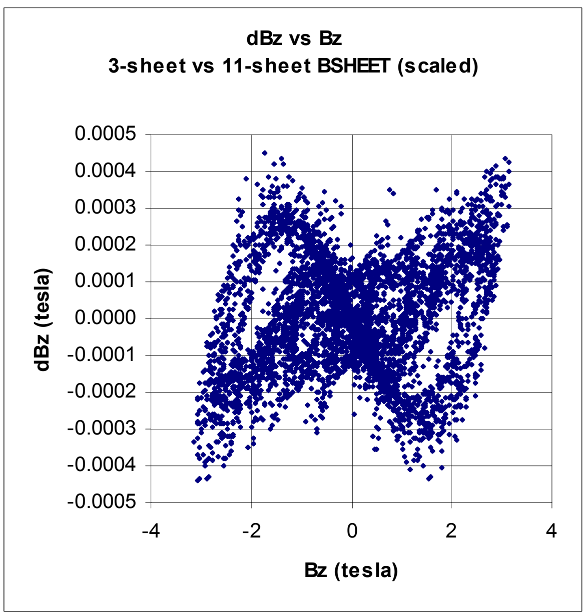

Bx1-SBx2
-0.00000536 Average
0.00015235 Std Dev
0.00035847 FWHM

By1-SBy2
0.00001749 Average
0.00005715 Std Dev
0.00013447 FWHM

Bz1-SBz2
0.00000097 Average
0.00015506 Std Dev
0.00036487 FWHM

If we assume that the 11-sheet field is a lot closer to the real field than is the 3-sheet field, then the 3-sheet field is not really an acceptable approximation (given the criteria for acceptability stated above). In our case, generating a grid with an 11-sheet approximation takes only about 2 hours longer than generating with a 3-sheet approximation. For this comparison the scaling factor is 1.00015.

The Eleven-sheet BSHEET Field vs a Twenty-two Sheet Field

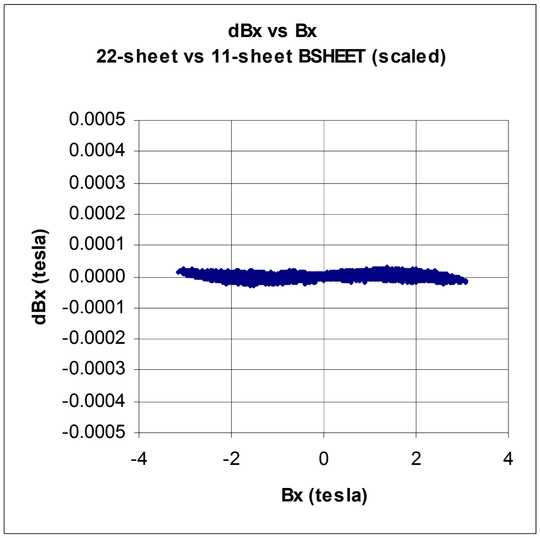

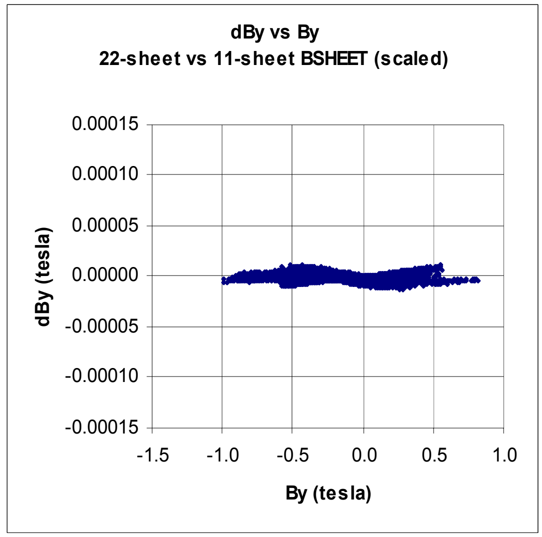

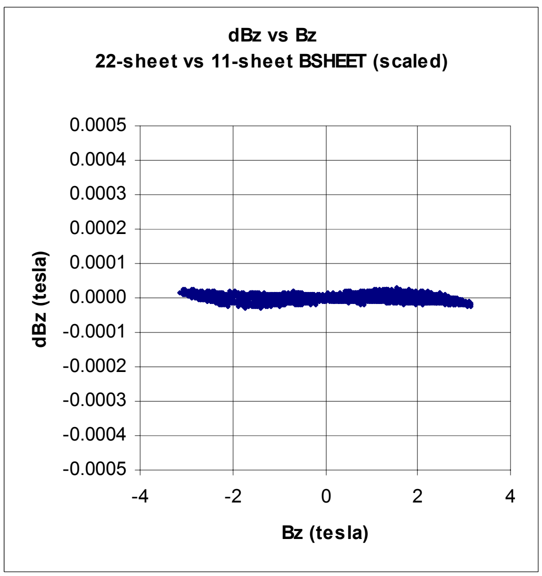

Bx1-SBx2
0.00000032 Average
0.00000905 Std Dev
0.00002128 FWHM

By1-SBy2
-0.00000118 Average
0.00000339 Std Dev
0.00000798 FWHM

Bz1-SBz2
-0.00000004 Average
0.00000917 Std Dev
0.00002158 FWHM

As expected, the 11-sheet field and the 22-sheet field are very close to each other; maybe we are even converging to the Truth. The scale factor here is 0.999990.

Next Steps

A second dataset is being prepared which has 1000 primary test points scattered around the ring, and for each primary, six nearby secondaries which enable the taking of derivatives that go into calculating the curls and divergence of the field. A very similar test has already been done to our fields A and B, but the new scheme will formalize the process in a way that makes it easy to check other fields as well.

Perhaps the most important thing to do is to perform the same tests, using the same ensemble of test points, on fields generated elsewhere, e.g. BNL. This is a very powerful consistency check, especially because BNL and Ole Miss have developed the transforms used to place and rotate coils almost completely independently That is one area where mistakes have been found in the past and problems could still be lurking. Such consistency checks as we have performed have been done mostly for small ensembles of points taken along the Nominal Beam Trajectory or the Coil Placement Circle. The tests being reported here are far more rigorous.

I understand from Rick that he is reworking a parameterization of his field map, an alternative method to interpolating within a field map grid. When that task has been completed, I would welcome the opportunity to plug that field into the comparison process as well. We are of course happy to participate in comparison processes that other devise, or to code and carry out comparison protocols suggested by others.

How sensitive is a ring's cooling performance to details of the magnetic field map? It is important to answer that question eventually, and perhaps not premature to start doing some serious work on it now. As soon as we have identified the parameters of a satisfactory (not necessarily optimal) cooling ring, it is fairly straightforward to alter the coil configurations and corresponding field maps and test the impact on the cooling effectiveness. What kind of alterations to the field are important to make? Presumably not random fluctuations on a point-by-point basis; these will be mostly averaged out as the particle passes thousands of field points. It is larger-scale alterations -- changes analagous to moving a coil or changing a current -- that are of the most interest.