Simulation and Fitting of Magnetic Fields of Superconducting Multipole Magnets at FRIB

Michael C. Wentzel[1]
[1]3rd Year Undergraduate, Department of Physics, University of Michigan, Ann Arbor, MI 48109, USA

Abstract. Essential to the beam delivery system at the Facility for Rare Isotope Beams (FRIB) is the precise calibration of steering and focusing magnets along the beam line. To properly calibrate multipole magnets used in the beam delivery system, the field throughout the bore of the magnet must be known. Specifically, when modeling or measuring multipole magnets with both large apertures and short lengths the fringing field at the edges of the field distribution must be accounted for. Moreover, the magnets in the beam line are not pure multipoles and thus exhibit behavior associated with dipole, quadrupole, sextupole, and other harmonics. Here we present an analysis of the mathematical and computational methods used to simulate and study the magnetic field within multipole magnets.
Keywords: Magnetic Field, Quadrupole, Modeling, FRIB, Curve Fitting, Enge Function.

1. INTRODUCTION

We carry out a numerical implementation of the results presented by Hiroyuki Takeda et al. [1] in the paper on the extraction of 3D field maps from 2D surface measurements. The process is acutely useful in the case of superconducting magnets due to their predilection for yoke saturation. We discuss here the specific case of short-length, large-aperture superconducting quadrupoles (SQs) which are used extensively in the beam delivery system at the Facility for Rare Isotope Beams.

An Enge function product was used to model the on-axis gradient for the superconducting quadrupoles (denoted as $b_{2,0}$) [1]. The model was then used to simulate the on-axis gradient at multiple currents, since the shape of the field varies significantly over the operating range of the quadrupole. We tested the accuracy of four similar fitting algorithms by fitting the parameters of the Enge function product to the simulated data and observing the deviation from the previously known parameters used in the simulation.

2. MODELING THE ON-AXIS GRADIENT FUNCTION

An on-axis gradient function is a mathematically convenient way of defining the magnetic field at any point inside the bore of a multipole magnet using a function of a single variable. The magnetic field inside the bore is determined by Laplace's equation of the magnetic scalar potential, $\nabla^2\Phi = 0$, with Φ defined in cylindrical coordinates owing to the geometry of the magnets. Solving for Φ and

applying the conditions that the scalar potential be single-valued and the length of the magnet be finite, we can achieve the following expression for Φ [1].

$$\Phi(r,\theta,z) = \sum_{n=1}^{\infty}\sum_{m=0}^{\infty}\left[\frac{b_{n,m}(z)r_0}{n+2m}\left(\frac{r}{r_0}\right)^{n+2m} sin(n\theta) + \frac{a_{n,m}(z)r_0}{n+2m}\left(\frac{r}{r_0}\right)^{n+2m} cos(n\theta)\right] \tag{1}$$

Here n represents the multipole harmonic of the magnet (i.e. dipole, quadrupole, etc.). We can see that the functions $b_{n,m}(z)$ and $a_{n,m}(z)$ (on-axis gradient functions) can be used to describe the behavior of the normal and skew magnetic field within a multipole. In fact, it has been shown that it is sufficient to find only the $m = 0$ term from which all further terms can be determined. [1, 2] For our purposes, we are interested in the normal component of a superconducting quadrupole magnet, and thus we focus on $b_{2,0}(z)$. While an analytic solution for $b_{2,0}(z)$ does exist, we model the on-axis gradient at the centers of the SQs using the product of two fifth order Enge functions as shown in Eq. (2).

$$b_{2,0}(z) = \frac{B_0}{(1+exp(\sum_{n=0}^{5} c_{i_n}(\frac{z-L_{eff}/2}{2r_0})^n))} \times \frac{1}{(1+exp(\sum_{m=6}^{11} c_{e_m}(\frac{-z-L_{eff}/2}{2r_0})^m))} \tag{2}$$

As in the solution for the magnetic potential Φ, cylindrical coordinates were used such that multipole fields about the longitudinal axis can be used for focusing charged particle beams. The longitudinal axis, here labeled z, refers to the position along the reference beam axis relative to the center of the magnet, and r_0 is the reference radius at which the magnetic field is measured (often treated as the radius of the pole-tips). Here c_{i_n} and c_{e_m} represent the incident and exit Enge function coefficients respectively, and the constant B_0 is a factor corresponding to the maximum on-axis gradient value. The effective length of the SQ, L_{eff}, is a parameter associated with the length of the magnet which can be defined in a number of ways provided the definition is consistent for each calculation. Throughout our tests of the fitting process, we used two methods of determining effective length.

The first method of effective length determination was to simply set the length to a fixed constant. When using experimental SQ data, the analogous step simply would be to probe for an effective length at a particular current and set L_{eff} to that value for all other field and current calculations. As explained in the following section, this was the method used to simulate the data. The second method for determining SQ effective length accounted for a variable effective length as a function of current. The formula used to define effective length is shown in Eq. (3) below where $b_{2,0}$ is a measurement of the on-axis gradient. In Eq. (3) $b_{2,0}(z = 0)$ is the value of the on-axis gradient at the center of the z axis of the magnet (presumably the maximum gradient value). It is possible in experiment for the maximum value to stray slightly from the center at $z = 0$; however, we chose to

set the measurement at a fixed position. Note that in practice, the integral would turn into a discrete Riemann sum evaluated over the region of data acquisition.

$$\frac{L_{eff}}{b_{2,0}(z=0)} = \int_{-\infty}^{\infty} b_{2,0}(z)dz \tag{3}$$

To further adjust the model to account for variable current, another variation was accounted for. The Enge coefficients (c_{i_n} and c_{e_n}) vary with excitation current, I, and can usually be described by a second or higher degree polynomial relation. We chose a second order polynomial for the purpose of testing different techniques of fitting the data to the model. Therefore, each coefficient is related to current by Eq. (4).

$$c_{i_n}, c_{e_n} = p_0 + p_1 I + p_2 I^2 \tag{4}$$

Applying this relation, we simulated test data for a known set of polynomial coefficients. We then implemented various fitting approaches to fit the data to the models and determine which methods gave the smallest residual errors.

3. DATA SIMULATION

We simulated the on-axis gradient of a quadrupole to behave as the model in Eq. (2). Each Enge coefficient (c_{i_n} and c_{e_n}) was modeled as a second order polynomial in current I as in Eq. (4). The parameters p_n were chosen to mimic similar behavior observed in an SQ designed for the FRIB separator. Additionally, the parameters were set such that for each n, the behavior for the polynomials c_{i_n} and c_{e_n} matched. At ten currents spaced evenly between 18 and 180 Amps, we simulated gradients along z with z set to vary between -1 and 1 m and the step size set to $5mm$. In the simulation we used a constant effective length set to $0.7m$. The length approximately marks where the on-axis gradient drops to half its maximum value for any given current. The final simulated on-axis gradient data were therefore a function both of z position and current. Fitting methods for the solution to the polynomial coefficients could then be tested against known values.

4. FITTING METHODS

We tested four different methods to fit the polynomial coefficients, p_n, and then compared the residuals from each of the three fits. All the methods were performed using the ROOT software with the Minuit minimization package. The methods differed in the technique used to determine effective length and global or local fitting of the p_n parameters with respect to the dimension of current. Local (or indirect) fitting was performed by fitting the Enge coefficients, c_{i_n} and c_{e_n}, for each current individually and then using the results of each coefficient for all ten currents to reach the objective of fitting the polynomial coefficients, p_n, for each c_{i_n} and c_{e_n}. Global (or direct) fitting was performed by fitting the polynomial coefficients in Eq. (4) for each c_{i_n} and c_{e_n} while considering the data over all currents.

First, we performed a local, indirect fit of the polynomial coefficients while treating L_{eff} as a constant value of $0.7m$, as was simulated. Since our data were generated using that set length, the fit using a constant length was predicted to be more accurate. However, it is important to note that in an experimental setting, one would not have knowledge of this value and would use Eq. (3) to determine its value. Note that the normalization of Eq. (3) is such that one does not necessarily get the same value of L_{eff} as was used in simulating the data. This has implications on the resulting fitted coefficients as will be demonstrated below.

Using the set effective length, we individually fit each of the generated data sets for the ten currents valued at $18n$ amps with $n \in 1, ..., 10$. Each fit generated values for the Enge coefficients (c_{i_n} and c_{e_n}). These values were plotted with the y-axis corresponding to the coefficient value and the x-axis corresponding to current. Finally, each coefficient was fit to a second degree polynomial to determine each p_n.

Second, we performed a local, indirect fit using a calculated effective length according to Eq. (3). The integral was evaluated numerically over the interval $z \in [-1, 1]$. Beyond this domain, the on-axis gradient becomes negligible. This calculated effective length was then substituted into Eq. (2), and the remainder of the fitting procedure followed as in the first method.

Third, we used the global fitting method to directly fit the polynomial coefficients. We started by plotting the data sets corresponding to all ten currents, yielding a three-dimensional plot. The on-axis gradient was plotted as a function of current and z position. The combination of Eq. 2 and Eq. 4 was then fit to the three-dimensional data by varying the polynomial coefficients $p_0, ..., p_{37}$ and B_0. For this test, we returned to a constant effective length of $0.7m$ for all data sets. In this method, we avoided two different fitting processes by circumventing the fits for the Enge coefficients (c_{i_n} and c_{e_n}) and fitting the polynomial coefficients directly. We postulated this method to be the most accurate due to requiring only one fit and using a constant effective length.

Fourth, we performed the fit using the same global fit method described above, only changing the effective length determination method. Here, we determined the effective length for each data set numerically as in Eq. (3). Using these calculated effective lengths, we determined a formula for the effective length as a function of current by fitting the effective lengths for the ten data sets to a second degree polynomial as in Eq. (5).

$$L_{eff} = p_0 + p_1 I + p_2 I^2 \tag{5}$$

Eq. 4 and Eq. 5 were then substituted into Eq. (2) for the Enge coefficients and effective length respectively. The resultant function was then fit to the data from all ten data current sets, keeping the L_{eff} polynomial coefficients fixed. As before, we expected this method to be slightly less accurate than the method using a constant effective length. It does, however, result in a more reliable method for analysis of real SQ data.

In each fitting method, the final parameter values are heavily dependent on the initial values. Particularly, when there are many parameters to vary in the fitting process, the number of parameter sets which constitute satisfactory fits is large. Therefore, in the interest of obtaining consistent results

between all four fitting methods, a common set of initial values was set. For the local methods, the initial values of the coefficients were all set to constants. In the global methods, these initial values were set as the constant terms of their respective coefficient polynomials. The remaining polynomial coefficients were all initially set to zero.

5. TEST RESULTS: FITTING

Using each of the four fitting methods described, we fit an Enge function product to the simulated data for ten currents. The goal of the test was to observe the difference in the errors of the fitting processes as measured by the sum of the squared residuals. Of course, there is some systematic error in the fitting process. By setting the initial values of the global fit with constant effective length to the values used in the simulation, a minimum residual on the order of 10^{-9} was determined and attributed to systematic error. In each test, the total residual was recorded, and the results are shown in Table 1.

As expected, the global fitting methods proved to be more accurate, producing residuals roughly ten times as small as their local counterparts. An explanation for the higher accuracy of the global fits is twofold: they are limited to a single fitting procedure, and they presume the second-degree-polynomial behavior of the Enge coefficients with current. The local fits use two fitting procedures which naturally increases systematic error; and in the first fit, they do not assume that each c_{i_n} and c_{e_n} behaves as a second degree polynomial meaning each c_{i_n} and c_{e_n} is fit individually. That is, there is no inherent co-variance between the Enge coefficients and current when, according to the simulation, there is.

Shown in Fig. 1 are plots showing the results of each of the four fitting methods along with the expected behavior of the first three incident and exit coefficient polynomials. From the plots, we can see than none of the polynomials matches the expected polynomial behavior. This is likely due to the large number of parameters in the fit which allows for a plethora of possible parameter sets that minimize the squared residuals while not matching the parameters used in simulation. Note, however, that in simulating the data, the polynomial behavior for each c_{i_n} and c_{e_n} was set such that for all $n \in 0, ..., 5$ and $I \in \mathbb{R}$, $c_{i_n} = c_{e_n}$. We therefore expected the plots of c_{i_n} and c_{e_n} to match for each fit. We do indeed observe this behavior in Fig. 1.

Additionally, Fig. 1 exhibits a plot of the polynomial used to describe effective length as a function of current. This same polynomial was used in both the local and global fits employing a calculated effective length.

In Fig. 2, the difference in expected and fitted behavior produced by the global fit with constant effective length is displayed. The difference is measured as $\delta = \frac{b_{2,0_{fit}}}{b_{2,0_{exp}}} - 1$. We can see that in the majority of the fitted domain, the value of δ is approximately zero as we expect. However, nearing high z and I, δ increases. Likewise, in the current range approaching zero, the fitted function appears to undervalue the expected behavior. Given the final χ^2 value on the order of 10^{-2}, we can conclude that a majority of the error comes from the boundaries of the fit. Fitting all p_n coefficients provides

a completely general fit for any current or z position, allowing for ease of access to on-axis gradient data for an entire fixed radius.

To further test the range of the two global fitting techniques, we implemented a function to recursively fit the model in Eq. 2 to the on-axis gradient data (the local fits were not tested using this method due to the different number of initial and final parameters). Following each fit, the final parameters were recorded and subsequently used to create new initial parameters for the next minimization process. Shown in Fig. 3 are graphs of χ^2 plotted against the number of fit iterations. The minimum χ^2 value for each global fit is shown in Table 1. For both global fits, there was a initial decrease in the χ^2 value followed by a sharp drop and plateau. There were cases where the residuals increased, and in plotting, any increase in χ^2 was ignored and treated as a repeated value.

The relation between change in the residual squared and number of fits allows for an optimization of computing time. For the global fit with a calculated effective length, pursuing more than five fits will yield little increase in accuracy. In fitting globally with a fixed effective length, the number of useful fits increases to near ten, but results in a final decrease of only 1.2% of the initial χ^2. Final results for the maximum number of iterations are shown in table (1). Interesting to consider is the fact that after multiple fits, the χ^2 value for the calculated effective length drops below that for the constant effective length. Furthermore, there was a considerably larger percentage decrease applying the iterative algorithm to calculated effective length, also displayed in Table 1.

It is important to note that with respect to the time of each fitting process, there are increasing returns to scale. As the residuals decrease, the fitting procedure makes fewer calls and thus returns a final parameter set more quickly than in the previous fit. Therefore, although a higher number of iterations reaches a plateau in $\Delta\chi^2$, it is possible for the return to be great enough at a number of iterations past the plateau. We tested this possibility.

Shown in Fig. 3 are graphs of the decrease in residual, $\Delta\chi^2$, over the CPU time of operation plotted against the iteration number. In fact, we still see a sharp decrease and plateau of $\Delta\chi^2/CPU_{time}$. More specifically, in both the constant and calculated effective length fits, the decreases and the plateaus occur at the same points as the decreases in their corresponding plots of χ^2 versus iteration number. The correlation indicates increasing the number of iterations past the point where significant decreases in the squared sum of residuals occur is disadvantageous.

Fit Method	Local: L_{eff} fixed	Local: L_{eff} calc.	Global: L_{eff} fixed	Global L_{eff} calc.
χ^2 Single Fit	2.07×10^{-1}	1.30×10^{-1}	5.83×10^{-2}	5.97×10^{-2}
χ^2 Recursive Fit	——	——	5.76×10^{-2}	5.11×10^{-2}
$\%$ Decrease	——	——	1.71	14.4

Table 1. Residual sums squared for the four tested fitting methods. These are calculated as $\chi^2 = \sum(b_{fit} - b_{sim})^2$. The third row shows the resultant χ^2 values after the maximum number of minimizations as given by Fig. 3; thirty minimizations were used with L_{eff} calculated and L_{eff} constant. The final row displays the percentage decrease in the χ^2 value from the single fit to the recursive fit.

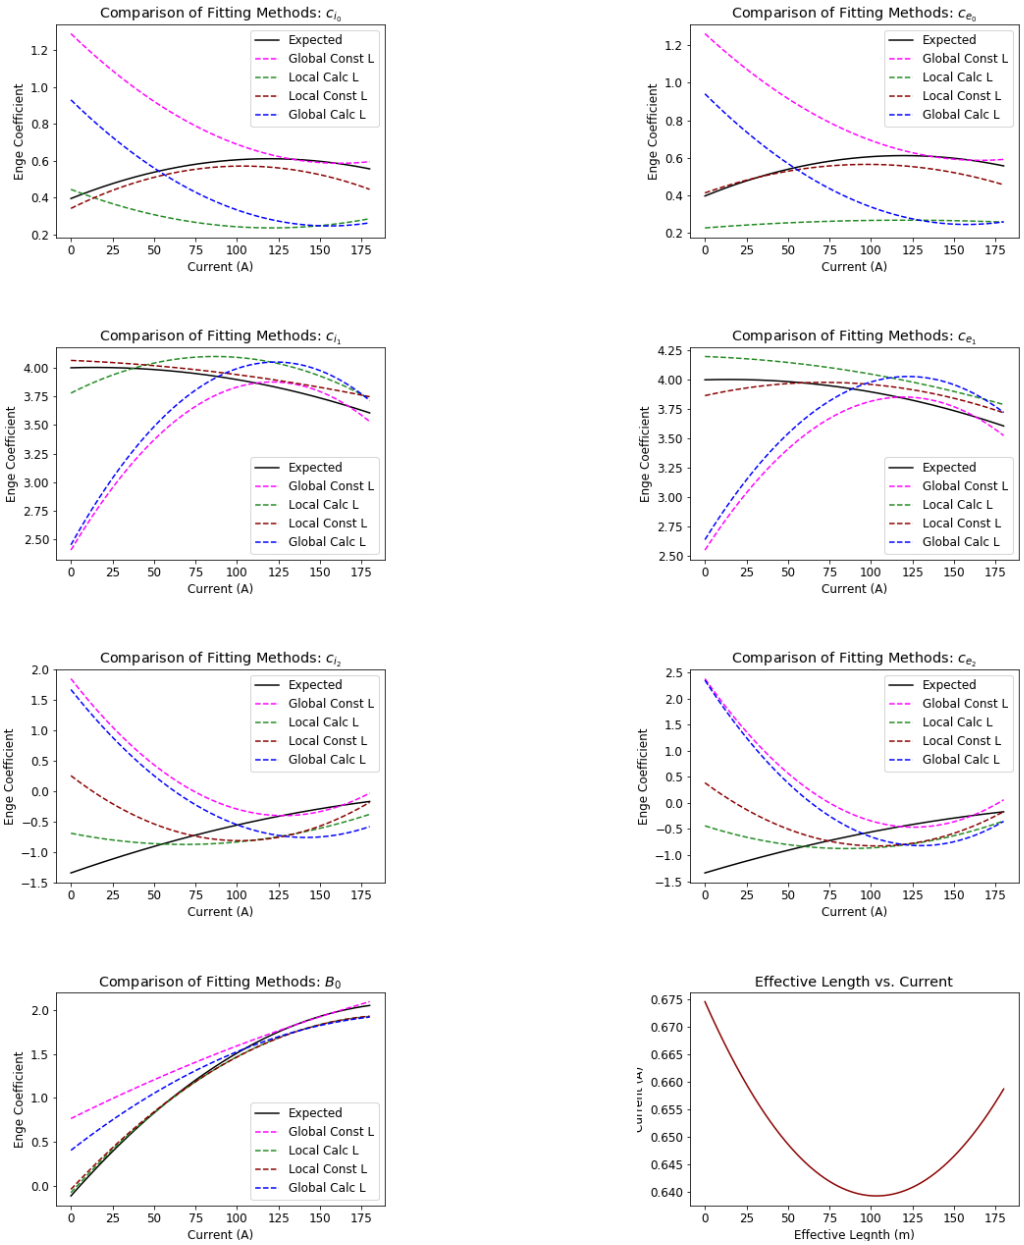

Figure 1. Expected versus fitted behavior for the first three incident and exit Enge coefficients,c_{i_n} and c_{e_n}, and the field gradient constant, B_0. Coefficients were fitted using all four fitting methods described above. While the local fit with constant effective length appears to yield a closer fit to the expected behavior, in fact the χ^2 for the global fit with constant effective length is on the order of ten times smaller than that for the local fits. Also shown is a plot of effective length versus current.

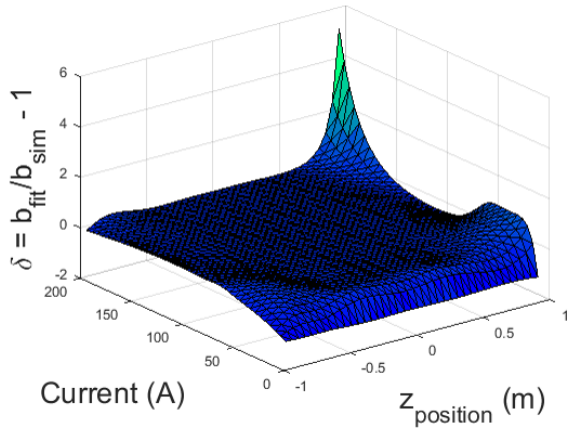

Figure 2. Difference in the fitted and expected behavior for the global fit with constant effective length. The difference is measured as $\delta = \frac{b_{2,0_{fit}}}{b_{2,0_{exp}}} - 1$. For a perfect fit, the result would be a plane at $\delta = 0$.

Figure 3. Plot showing χ^2 values and $\Delta\chi^2/CPU_{time}$ plotted against the number of fitting iterations for constant and calculated L_{eff} respectively.

6. CONCLUSIONS

Although the fixed effective length global fit yields the highest accuracy fit, it is not feasible to apply it in an experimental setting. The effective lengths of the SQs will not be known explicitly without calculation, and that variance must be accounted for in the fitting procedure. Therefore, we use this method to determine the maximum accuracy of the fit which is determined to be a minimum χ^2 on the order of 10^{-9}.

The optimal method to use in experiment is the global fit using a calculated effective length. Although the residual sum is higher than that for the constant effective length, it accounts for the natural behavior of the SQs and the unknown L_{eff}. Furthermore, under an iterative algorithm it has potential to make additional reductions in the χ^2 value.

We establish the optimal number of minimzations to perform to be approximately five based not only on the return on $\Delta\chi^2$ but also on the CPU performance time for each iteration of the fit. Increasing the iteration number beyond this yields little to no return in $\Delta\chi^2$.

In the future, we intend to test the fitting methods against data for operating SQs. Implementing the methods described here will decrease the time and cost of magnetic field mapping for SQs while simultaneously maximizing the accuracy. Moreover, the same procedure can be applied to superconducting or room temperature magnets of higher order poles with ease.

7. ACKNOWLEDGMENTS

The author gratefully acknowledges the incredible guidance and mentorship from Dr. Mauricio Portillo of the Facility for Rare Isotope Beams in the research leading up to this paper (supported by REU).
This material is based upon work supported by the U.S. Department of Energy Office of Science under Cooperative Agreement DE-SC0000661, the State of Michigan, and Michigan State University. Michigan State University designs and establishes FRIB as a DOE Office of Science National User Facility in support of the mission of the Office of Nuclear Physics.

References

[1] Takeda H., Kubo T., Kusaka K., Suzuki H., Inabe N., and Nolen J. A., 2013, "Extraction of 3D Field Maps of Magnetic Multipoles from 2D Surface Measurements with Applications to the Optics Calculations of the Large-Acceptance Superconducting Fragment Separator BigRIPS", Nuclear Instruments and Methods in Physics Research B 13, **798**

[2] Walstrom P. L., 2004, "Soft-Edged Magnet Models for Higher-Order Beam-Optics Map Codes", Nuclear Instruments and Methods in Physics Research A 519, **216**