

Mapping the OLYMPUS Magnetic Field

Axel Schmidt

August 19, 2013

1 Introduction

During the spring and summer of 2013, a major effort was undertaken to map the OLYMPUS magnetic field and to incorporate that data into the simulation and other analysis applications. Knowledge of the magnetic field is necessary to simulate the trajectories of charged particles, a major step in the simulation's propagator and in the track fitter. Up until this point, the analysis software relied on a magnetic field map produced for the BLAST experiment. Since the BLAST magnet was rebuilt as the OLYMPUS magnet, the BLAST field map was certainly a good approximation to the OLYMPUS magnetic field. However, for a final analysis, an accurate map of the OLYMPUS field was needed.

Between March and May, measurements were taken using a Hall probe that was moved by translation tables through the volume of the toroid. The position of the Hall probe was determined by means of optical survey. Between May and August, the data from the field measurements the survey were compiled in order to create an updated field map. In addition, the software for reading and interpolating this map was improved and unified into a single software library, for all applications to use.

1.1 The OLYMPUS Magnet

The OLYMPUS Magnet was previously used as the magnet for the BLAST experiment and was described in a NIM article [1]. The magnet was made up of eight coils in a toroid configuration, with the beamline running through the center of the toroid axis. The coils divided the region around the beamline into eight wedge-shaped sectors. The two sectors in the horizontal plane were instrumented with detectors.

Each coil was approximately 4.3 m long, and 1.9 m wide, and was formed by 26 windings of square copper bars, arranged in two layers of 13 windings. The coils had an irregular shape; they were narrower upstream to accommodate the target, and wider downstream, to extend the region of high magnetic field to smaller scattering angles. During normal OLYMPUS operations, the coils were driven with 5000 A of current, which produced a magnetic field of about 0.3 T in the high-field region.

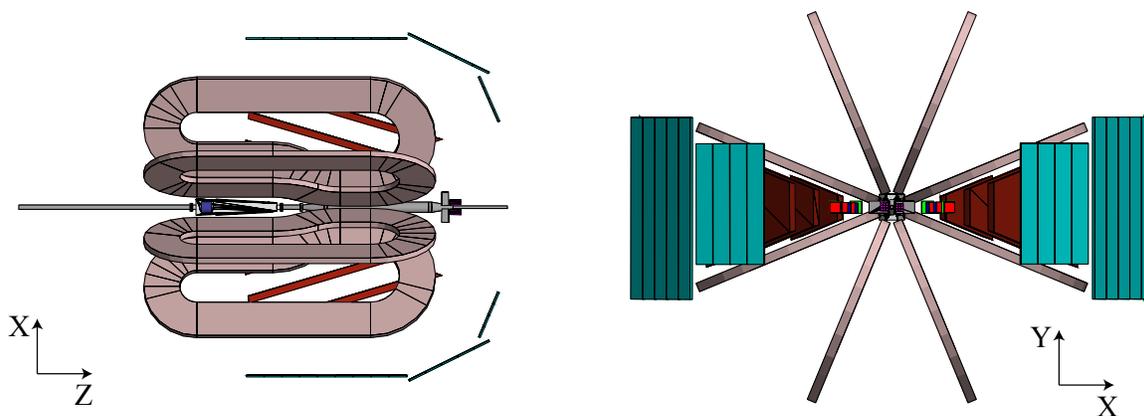


Figure 1: The OLYMPUS Magnet and detectors are shown from above (left) and from downstream (right). The standard OLYMPUS coordinate axes are also shown.

1.2 The OLYMPUS Coordinate System

This note makes frequent mention of positions and orientations using the standard OLYMPUS coordinate system. By this convention, the Z axis points in the direction of the beam, the Y axis points upwards, and the X axis points to the left of the beamline, forming a right-handed system. During the measurement of the magnetic field, a few other coordinate systems were commonly used to describe, for example, the positions of translation tables. In this note, all positions are described using the standard OLYMPUS system. Figure 1 shows the standard coordinate axes in relation to the OLYMPUS Magnet.

2 Measurements

2.1 Overview

The measurements of the magnetic field were taken with a 3D Hall probe. The probe was moved by translation tables through the volume of the toroid. The motors in the tables could not operate within a strong magnetic field, so a long rod connected the probe to the tables, which sat outside the magnet. To cover the full volume of the magnet, several rods with different lengths were used, and the set-up had to be recalibrated several times. Measurements were performed first on the left sector, before the set-up was taken apart and reassembled for measurements on the right sector. During all of these measurements, the probe position was determined by optical survey with a variety of equipment. These procedures are detailed below.

2.2 Translation Tables

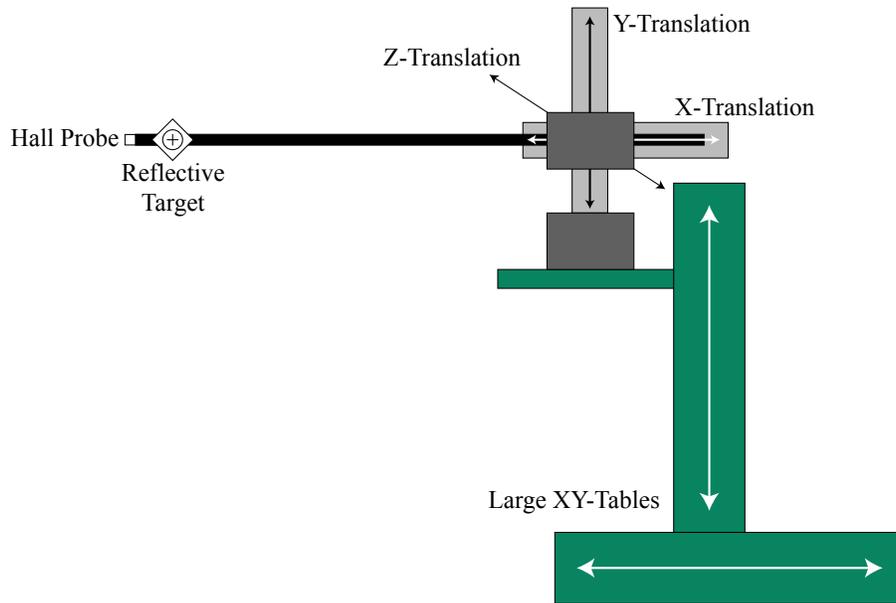


Figure 2: The rod was mounted on a system of three translation tables to provide motion in all three axes. This viewing angle approximately matches the photo in figure 3.

To move the Hall probe through the toroid volume, the probe was attached to one of several carbon fiber rods, which attached to a system of translation tables. The translation system sat away from the magnet so that the magnetic field would not interfere with the stepping motors. The rod pointed inwards toward the beamline, parallel to the X axis. The rod was moved in the X and Y directions by two 0.2 m tables, and in the Z direction by a long 6 m table. This apparatus was bolted to a long aluminum I-beam, which sat parallel to the Z axis. The I-beam rested on an additional pair of XY tables which could add an additional meter of translation in the X and Y directions.

Measurements were made at points that formed a grid in the toroid volume. A grid spacing of 50 mm was chosen for the region within 1 m of the beamline because the field gradients in that region were large. Outside

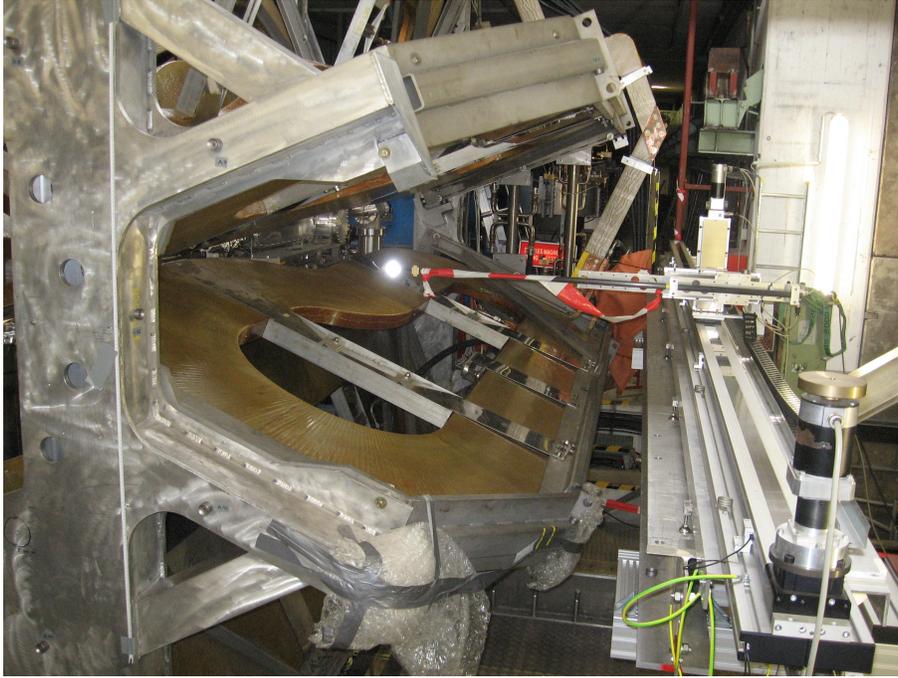


Figure 3: The apparatus is set up to measure the left sector. The rod is midway down the long Z table.

this region, a grid spacing of 100 mm was chosen. The measurements were divided into scans, each at a particular position in X and Y . During a scan, the probe stepped along the Z -axis, taking measurements of the field at the designated grid points. At the end of each scan, the probe was reset to a new X and Y position for the next scan. In this way, a regular grid was mapped throughout the toroid volume. In all, 703 scans were taken, totalling more than 36,000 measurement points. The positions of these scans in X and Y is shown in figure 4.

The scans nominally extended from $Z = 3600$ mm to $Z = -400$ mm except in regions where the full range was obstructed by the magnet supports. The boundaries of the grid were wedge shaped in X and Y to fit between the coils surrounding the tracking volume, but extended from the beamline, all the way to $|X| = 2500$ mm in both sectors. To achieve such a large coverage, three different rods were used, each with a different length. Additional coverage in Y was achieved by mounting a vertical extension to the brackets that mounted the rod on the translation tables.

2.3 Surveying

During the set-up of the translation tables, the DESY Survey Group used a laser tracker apparatus to measure the positions of survey targets on the I-beam to align the the tables with the axes of the OLYMPUS coordinate system. The Survey Group also measured the position of a target on Z -axis stepping motor over the course of a scan. The results indicated that the target did not follow a straight line, but in fact followed a trajectory which wiggled several millimeters up and down over the range in Z . The shape of the trajectory was found to be repeatable, and so could be corrected, if the start and end positions of every scan were measured. The decision was made to survey the probe position at the beginning and end of every scan.

The DESY survey group provided a Leica Wild T3000 Total Station, and a Kern E2 Theodolite to allow the OLYMPUS measurement personnel to perform their own survey during the measurement scans. Both devices can measure the polar and azimuthal angles to a point in space, as determined by a telescope and cross-hairs. The total station can additionally measure distance to a set of special reflective targets. The theodolite was positioned on the downstream concrete block on the opposite side from the measurement, looking down the rod at the probe tip when in its nominal start position of $Z = 3600$ mm. The total station was positioned in the downstream tunnel, and aimed at targets that were attached to the measurement rod. The total station was used to confirm that deviations in trajectory had the same qualitative shape for all of the measurements in a given sector, but the amplitudes were a function of the specific rod being used. The decision was made to survey the start and end points of each scan, so that the trajectory in between could be inferred.

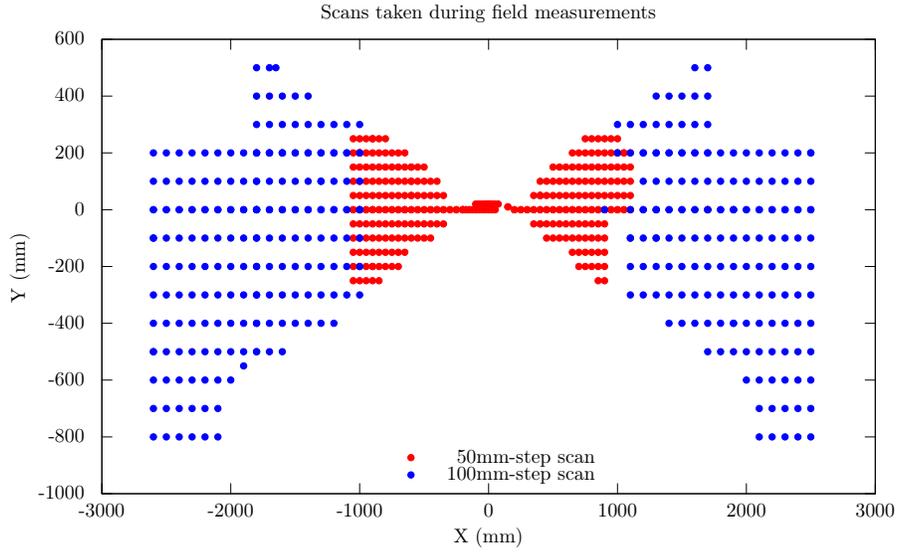


Figure 4: 703 scans were taken, covering more than 36,000 grid points.

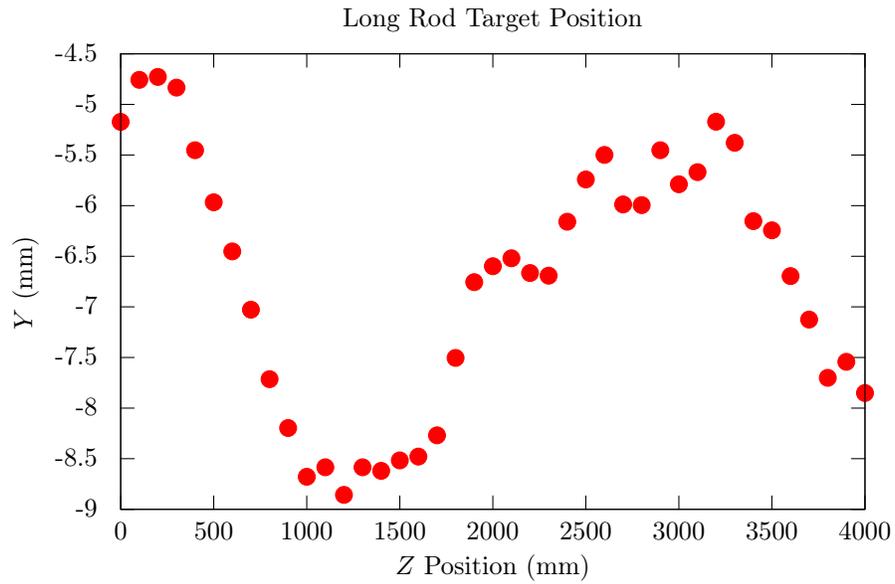


Figure 5: The Hall probe did not follow a straight line over the course of a scan. This position deviation was corrected using survey data

2.4 Procedure

During data taking, a crew of three people operated equipment to make measurements. One person controlled the translation tables from a computer near the shielding wall. A second person operated the theodolite to survey the probe tip start position. The third person used the total station to survey the reflective targets at each scan's start and end position.

All of the scans for a given rod were performed at once, to minimize the number of times that a rod had to be exchanged. First, the medium rod measurements were taken (for the region $950 \text{ mm} < X < 1700 \text{ mm}$). Next, the short rod was installed for the region $1800 \text{ mm} < X < 2500 \text{ mm}$. Then the long rod was installed for the region $150 \text{ mm} < X < 900 \text{ mm}$. The brackets attaching the long rod were then adjusted to extend the long rod to cover the Møller region ($|x| < 100 \text{ mm}$). Finally, the vertical extension was added and medium rod reattached to cover the area above the beamline, near the upper edge of the tracking volume. After these scans were completed on the left sector, the apparatus was taken apart and reassembled to measure the right sector.

2.4.1 Daily Calibration

A daily calibration procedure was designed and followed during data taking. Each morning, the theodolite was calibrated by measuring the positions of tick marks on rulers mounted on the shielding wall and the beamline support. If the position deviated significantly from the nominal aligned position, the theodolite was repositioned. The total station was calibrated by measuring the positions of all of the visible OLYMPUS survey marks. In this way, the position and orientation of the total station could be determined. Other daily tasks included checking the positions of the large XY tables to make sure the readings on the manual controls matched their physical position, taking photos for documentation, and using a Mu-metal shield to reset the zero point on the Hall probe. The beginning of the measurements began with a test run with the magnet off. Then the magnet was activated and real measurements began for the day.

2.4.2 Changing Rods

A set of procedures were followed in order to change the rod. The work of mounting the probe in the rod, and attaching the rod with brackets to the translation tables was performed by a technician in the DESY group MEA 1. No easy way was found to align the probe in the rod, so the alignment was checked by using the theodolite, and then adjusted by hand. Once the new rod was installed, a reflective total station target was attached tightly with a cable tie to the rod, 51 mm from the tip. Since the probe was 1 mm from the tip, the target was nominally 50 mm from the probe. The angles between the tip of the rod and target were measured with the total station. To check the trajectory of the probe over the course of the scan, the next task would be to step the probe through the grid points of a scan in Z , measuring the target position at each point with the total station.

When the long rod was used, the full range in Z was not available because the K-beam supporting the coils sat in the path of the rod. Scans were made only in the region upstream of the K-beam, and frequently in the start position, the K-beam would obscure the total station's view of the reflective target. The solution to this problem was to add a second target further up the rod, so at least one of the targets was always visible.

3 Analysis of the Survey Data

3.1 Overview

The first task for analyzing the magnetic field scans was the determination of the spatial position and orientation of the probe at each measurement point. The scanning procedure was designed to map out a grid in space, but early on it was determined that the scanning mechanism introduced significant deviations from that grid. For example, the long translation table parallel to the Z direction introduced vertical deviations on the order of several millimeters. To account for these grid deviations, survey equipment was used to measure the positions of targets attached to the scanning rods at the beginning and end of each scan. The analysis of this survey data divided into three principle tasks:

- Determining the positions of the targets at the beginning and end of each scan
- Determining the probe positions relative to these targets
- Modeling the trajectory of the probe along the length of the scans.

Each of these topics are covered in the following sections.

3.2 Determining Target Positions

For each scan, the scanning rod was surveyed using a Leica Wild T3000 Total Station, and a Kern E2 Theodolite. These devices are similar; they both have a telescope with crosshairs and can report the polar and azimuthal angles of the crosshairs. The total station can also measure the distance to special reflective targets using a laser. This additional distance measurement, combined with the measurement of the two angles uniquely constrains the target point in three dimensional space. Thus the total station data was far more useful than that from the theodolite, which was not used in the analysis.

Both the total station and the theodolite were calibrated daily. Neither the total station nor the theodolite were rigidly fixed, and so were vulnerable to vibrations and the occasional bump. To calibrate the theodolite, rulers were fixed to the walls and floor, and the position of the theodolite was readjusted relative to marks on the ruler. The total station was calibrated by measuring survey marks on the walls, quadrupoles, and magnet frame. Given that the positions of these survey marks were known from the survey results, the position and orientation of the total station could be established, and a transformation to the OLYMPUS coordinate system produced. This method proved successful.

After transformations were produced from each of daily calibrations, they were applied to all of the total station measurements to produce a list of target positions in the global OLYMPUS coordinate system. The next task was to comb through this list for accuracy. The measurements of the total station had to be transcribed by hand and this transcription had errors with a rate of a few dozen over the approximately 700 scans. It was endeavored to correct as many of these errors as possible.

To search for these errors, the position data was plotted in several ways to look for inconsistencies. One way was to plot the difference between the target position and the nominal scan position. If the position data were correct, then they would appear in this plot to be clustered around a constant value: the difference between the target position and the probe position. Positions that were very far from the cluster, were likely to have transcription errors. Another useful plot was the difference between the start and end positions of a scan. In the X and Y directions, this difference clustered around zero. Outliers were likely to have errors.

To correct transcription errors, the electronic log was first checked against any paper records that were kept. If that failed to turn up the error, then the total station values were compared to those of scans of nearby positions. Often the error would be a single digit written incorrectly, or two neighboring digits transposed. If such a correction could be made to move the total station position into the clusters of the plots, then it was adopted, and recorded in the electronic log book. After examining every measurement, all of the positions were well-clustered to within a few mm or better.

One observation that was made during this process was that during the left sector measurements, the Z axis of the translation table came out of alignment with the OLYMPUS Z axis. This occurred between the measurement of the end point of scan 0093 and the start point of scan 0094. Between these scans, the large XY tables were adjusted, so it is possible that during this movement, the tables made contact with part of the magnet. This misalignment posed no problems because the survey data allowed the probe positions for scans after 0094 to be corrected.

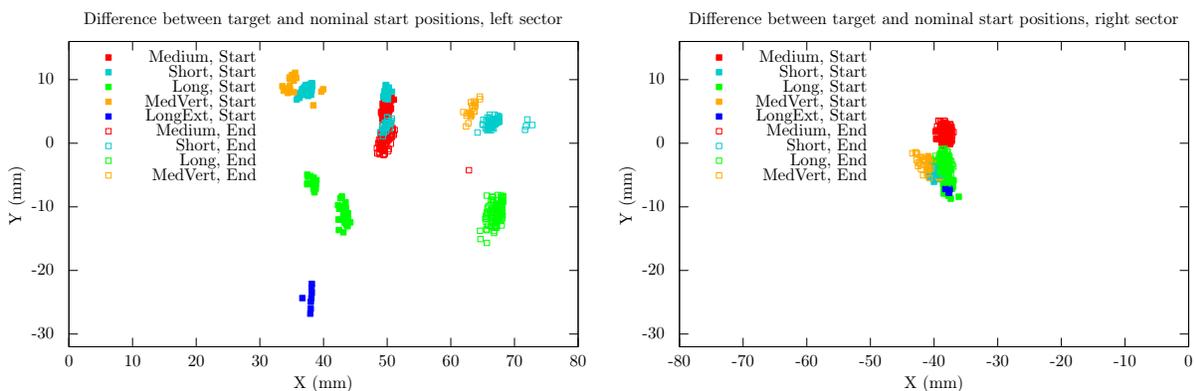


Figure 6: The start and end points of the scans should be clustered tightly around their nominal values. A few outlier points probably have transcription errors, but they could not be identified. The left sector points are widely spread because of the bump in the tables that occurred after scan 0094.

3.3 Determining the Probe Positions

The next task was to determine the position of the Hall probe relative to the positions of the total station targets. The reflective targets were attached to the measurement rods by hand using cable ties, so were fixed in position, but this position was not precisely aligned. A set of offsets were needed to establish the position of the Hall probe relative to the positions of the targets on each rod.

Each rod had a target that was attached 51 mm away from the tip. This was done by hand and checked with calipers. According to the probe's specification sheet, the probe was located 1 mm from the tip of the rod, so the thinking was that the target would be a round 50 mm from the probe. When each rod was changed, some total station measurements were made of the angles between the target and the tip of the rod (no distance measurement was possible because the tip of the rod was not a reflective target). These angle measurements were a useful check of both the 51 mm distance between the target and rod tip, and as well that the target and the probe had the same position in Y . However, the precision was not enough to make a convincing statement about small deviations in the distance between the target and the probe. It was therefore assumed that the probe was 50 mm away along the X axis from the total station target and that there was no deviation in Y . The deviation in Z was calculated based on caliper measurements of the diameter of each rod.

The problem was complicated by the fact that when the long rod was used, the target was not always visible from the total station. Therefore a second target was added to the rod, much farther away from the tip. At least one of these targets was always visible, and on many scans, both were visible. In these later scans, it was possible to extract a relative position between the two targets. This relative position could be used when the primary target was not visible.

3.4 Scan Trajectories

Early on, during the set up of the measurement apparatus, it was noticed that the translation tables do not maintain a consistent vertical position over the course of the full range in the Z direction. Instead, the table height fluctuates several millimeters over the course of a scan. Corrections needed to be made to the scan positions to accurately reflect the true trajectory of the Hall probe over the course of a scan.

To measure these trajectories, a test scan was made for each rod, in which the rod was stepped in increments of Z and the target positions were measured with the total station. From these measurements, it was found that the overall shape of these fluctuations was consistent for each set-up of the apparatus, i.e. all scans on the left sector had similarly shaped trajectories, but these were different from the shape of the trajectories on the right sector. It was also found that the magnitude of these fluctuations depended predictably on the length of the rod being used. In fact it was sufficient to parameterize the trajectories of the Hall probe in the following way.

$$x_{r,s}(z) = x_s(z) + L_r \cos(\theta_s(z)) + \text{linear term} \quad (1)$$

$$y_{r,s}(z) = y_s(z) + L_r \sin(\theta_s(z)) + \text{linear term} \quad (2)$$

Index r represents the rod being used, while index s represents the sector (left or right). The trajectory could be expressed in terms of three functions $x_s(z)$, $y_s(z)$, and $\theta_s(z)$ that were consistent across each sector, as well as the rod length L_r which was measured. The three functions were extracted using a fit to the calibration data, and the residuals were found to be sufficiently small. A linear term was used to match the trajectory to each scan's start and end points. An example of these corrections being applied to fit the left sector calibration data is shown in figure 7.

3.5 Software

The application of these corrections was performed by a program in the OLYMPUS software repository called "DataPrep." The DataPrep program reads in all of the magnetic field data files as well as all of the survey data. For each magnetic field scan, it first transforms the survey data into positions in OLYMPUS coordinates by using the proper calibration. Next, it generates the Hall probe start and end positions from the available target positions. Then, it uses the start and end positions to map out a trajectory over the course of the scan. The final step is to write out, for every magnetic field measurement, the corrected positions, the measured magnetic field, and the magnet current used for that measurement.

The DataPrep program is found in the 'bin' directory of the repository. The executable can be called from any directory, since the required input files can all be found in the user's '.olympus/shared' directory. The results are written to a file called 'fieldgrid.txt.'

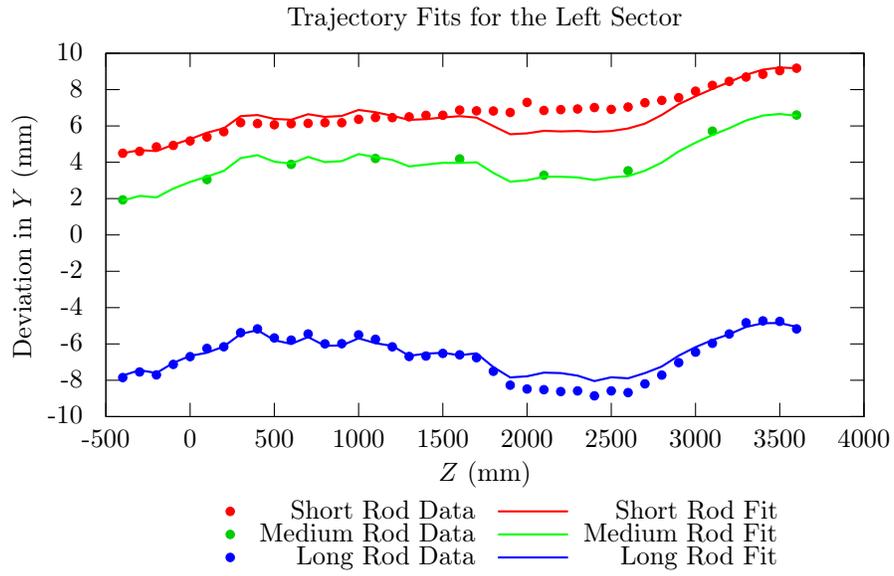


Figure 7: The same trajectory correction is able to fit the deviations for all of the different rods. These Y fits follow equation (2).

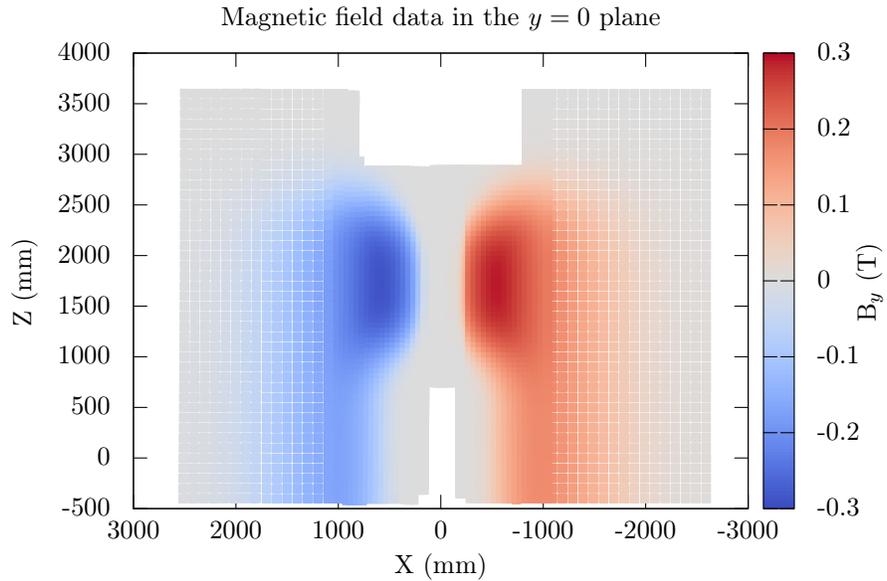


Figure 8: This plot shows the Y component of the magnetic field data as measured, after position corrections were applied.

4 Generating a Field Map

4.1 Overview

The magnetic field measurements are limited to a discrete set of points. In the OLYMPUS simulation, knowledge of the magnetic field is needed at every point in space. This requires determining the magnetic field in between nearby measurement points (interpolation) and at points that are outside the spatial extent of the measurements (extrapolation). One can achieve this by making a model of the magnetic field with only a few assumptions. This model can be validated by how well it reproduces the correct magnetic field at the measurement points. One can improve the model by fitting its parameters to find the point of best agreement with the data.

A model of the magnet was designed for this purpose. The model consisted of filaments of current in the shape of the copper conductors in each coil. The magnetic field was then calculated using the Biot-Savart Law. Parameters in this model were fit to the magnetic field data. These included the positions and orientations of the coils, the current in the coils, and geometrical deformations.

The Biot-Savart calculation is of course too slow to serve within the OLYMPUS simulation, which needs to query the field at thousands of positions in order to simulate a single particle's trajectory. A further approximation scheme is needed to speed up the process. The field and its derivatives were calculated using the Biot-Savart law over an evenly spaced grid of approximately 400,000 points over the full volume of the experiment. Spline interpolation could then be performed to return a magnetic field value at any point in the volume.

4.2 Magnet Model

The OLYMPUS magnet was modeled as a collection of current segments that traced the nominal shape of the coils. The model was designed to balance ease and speed of field calculations with accuracy. The nominal OLYMPUS coil shape is irregular but can be constructed of straight sections and circular arcs, so expressions for the field created by current in those shapes was necessary. This is easy to do if the current loops are filaments, but not so easy to do if the current has some cross-sectional extent. Therefore the conductors, which in reality, are square copper tubes, 1.5 inches across, were approximated by filaments.

The magnetic field calculation was performed using the Biot-Savart Law:

$$d\vec{B} = \frac{\mu_0 I}{4\pi} \cdot \frac{d\vec{l} \times (\vec{p} - \vec{c})}{|\vec{p} - \vec{c}|^3} \quad (3)$$

where $d\vec{B}$ is the differential magnetic field at measurement point \vec{p} produced by a differential current segment located at point \vec{c} , with differential length $d\vec{L}$. If the segment is a straight line centered at \vec{c} and length vector \vec{L} , then Biot-Savart Law integrates to

$$\vec{B} = \frac{\mu_0 I}{4\pi\beta^2} \left[\frac{(\alpha + |\vec{L}|/2)}{|\vec{c} - \vec{p} + \vec{L}/2|} - \frac{(\alpha - |\vec{L}|/2)}{|\vec{c} - \vec{p} - \vec{L}/2|} \right] \cdot (\vec{c} - \vec{p}) \times \vec{L}, \quad (4)$$

$$\alpha \equiv \hat{L} \cdot (\vec{c} - \vec{p}), \quad (5)$$

$$\beta^2 \equiv (\vec{c} - \vec{p})^2 - \alpha^2. \quad (6)$$

An analytic expression also exists for the integration of a circular arc-shaped segment but it involves the use of Appell functions and the attempts to implement this were not successful. Instead, arc shaped segments were approximated by many line segments. The proper placement of these line segments is a non-trivial exercise. If the start and end points lie on the path of the arc, then the area of the arc is diminished and the dipole moment decreased. The segments can be spread out to preserve the area of the arc, but this means that the end points first and last segment will no longer lie on the end points of the arc. If the arc is a portion of a current loop, then the loop will not be continuous at the start and end points of the arc. To balance these concerns, it was decided to position the segments so that the first and last segment began and ended, respectively, on the endpoints of the arc, preserving continuity of any current loops. The other segments' endpoints would be expanded outward to a larger radius R' to maintain the area of the arc. The expression for R' was derived using geometry, and the result is

$$R' = R \left(\sqrt{\frac{1}{(N-2)^2} + \frac{\theta_a}{(N-2)\sin(\theta_a/N)}} - \frac{1}{N-2} \right), \quad (7)$$

where N is the number of segments used to approximate the arc, and θ_a is the angle subtended by the arc.

All eight magnetic coils in the model were identical in shape and composition, but could be positioned and oriented independently. Each coil was composed of 26 current loops arranged in two rows of 13. Each loop was positioned at the center of the nominal position of each of the copper windings. The model current loops were built of three straight sections, two 180° arcs (approximated by 180 segments) and two arcs of approximately 38.58° (approximated by 40 segments).

4.3 Fitting the Magnetic Field Data

The parameters in the magnetic field model were chosen after fitting various combinations of parameters to the magnetic field data. A program, titled “coilfit” was written to perform these fits and has been included in the OLYMPUS software repository. The quality of a fit could be judged by the degree to which the model reproduced the measured data and by the degree to which the resulting parameters were believable. After experimenting with many degrees of freedom that could be included in the magnetic field model, a final set of parameters were chosen, so that the magnetic field could be calculated at any point in space.

The “coilfit” program was written as a standalone executable to perform magnetic field fits. The program reads in the magnetic field data as written by the “DataPrep” program, scaling any data taken at -5000 A current by -1 to match the positive polarity data. Using the model, the program can also calculate the magnetic field at the exact same spatial points. Using least squares criteria, parameters in the model can be fit to achieve the best agreement with the data. The minimization is achieved by the Levenberg-Marquardt algorithm as implemented by “cminpack.” Several techniques were chosen to increase the speed of the fit. First, the model was constructed of a single coil, of which eight copies were made and positioned independently. The field of each coil was calculated and cached independently. This meant that field recalculations only needed to be performed for coils which moved since the last cache. Secondly, each iteration of the fit required calculating the magnetic field at each measurement point. These calculations were performed in parallel, and this was implemented using the OpenMP API.

There proved to be no obvious set of parameters to use for the fit, and a great deal of experimentation had to be performed to discover the discriminating power of the measurements. A balance had to be struck between giving the fit too much or too little freedom. With fewer free parameters, the model was less able to reproduce the data. With too many free parameters, especially with parameters that were not well constrained by the measurements, the best-fit values were too different from their nominal values to be believable. The goal became to add well-constrained free parameters to minimize the residuals.

One natural parameter set were the three positions and three rotations for all eight coils. However the final positions for the top two and bottom two coils were wildly different than what one would expect. The measurements, which were never close to those coils, could not provide the constraint that was needed for their positions and rotations to serve as useful parameters. Instead, it was more productive to give the four “tracking volume” coils more freedom than the top pair and bottom pair.

The final parameter set had 35 free parameters. The four “tracking volume” coils were given 24, the full freedom of position and rotation, while the remaining four coils were collectively given eight: three positions for the toroid origin, three rotations for the toroid axis, an collective in-plane rotation, and a collective in-plane expansion. The three other parameters in the fit were added to compensate for the filament approximation in the magnet model. The coils were collectively allowed to shrink or expand with two scale factors. The third parameter was total current in the coils. The final values of these parameters are shown in table 1.

Param.	Value	Param.	Value	Param.	Value	Param.	Value	Param.	Value
I	5007 A	S_x	0.986	S_y	0.998	θ_{coil}	-0.05°	$\Delta\rho$	12.35 mm
x	3.357 mm	$\Delta\rho_0$	26.58 mm	$\Delta\rho_3$	24.64 mm	$\Delta\rho_4$	-4.962 mm	$\Delta\rho_7$	-1.494 mm
y	-21.7 mm	$\Delta\phi_0$	0.01°	$\Delta\phi_3$	-0.006°	$\Delta\phi_4$	-0.006°	$\Delta\phi_7$	0.002°
z	4.633 mm	Δz_0	8.214 mm	Δz_3	1.429 mm	Δz_4	0.332 mm	Δz_7	22.74 mm
α_x	1.368°	$\Delta\theta_{x0}$	-0.040°	$\Delta\theta_{x3}$	-0.065°	$\Delta\theta_{x4}$	0.12°	$\Delta\theta_{x7}$	0.547°
α_y	0.017°	$\Delta\theta_{y0}$	-0.939°	$\Delta\theta_{y3}$	1.141°	$\Delta\theta_{y4}$	1.1°	$\Delta\theta_{y7}$	-0.891°
α_z	-0.163°	$\Delta\theta_{z0}$	0.402°	$\Delta\theta_{z3}$	0.635°	$\Delta\theta_{z4}$	-0.846°	$\Delta\theta_{z7}$	-0.492°

Table 1: The final fit had 35 free parameters. Coils 0, 3, 4, and 7 were given full freedom because they surrounded the tracking volume and were well constrained by the measurements.

The fit to these parameters had an average residual of 18 G per point per direction. The residuals are smallest in the high-field region in the tracking volume, and are largest in the gradient region between the toroid axis (where

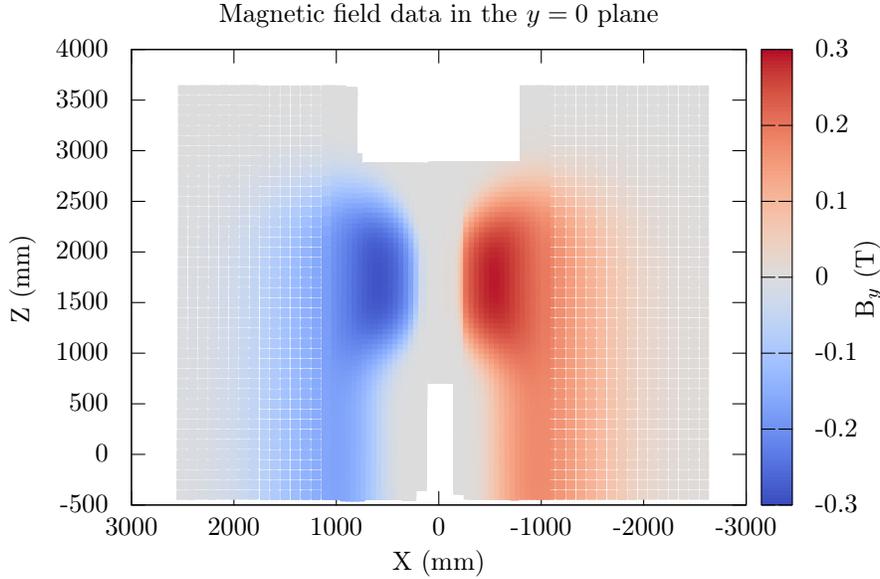


Figure 9: The magnet model with the parameters in table 1 was used to create this field map.

the field is small) and the tracking volume (where the field is large). Unfortunately, this is the location of the 12° detectors. Attempts to find improvements in this region were not successful. Errors in this region might be mitigated by the fact that tracks have to pass through regions of both too-high and too-low field. Further study is needed to assess how much of a systematic effect is created by these field residuals.

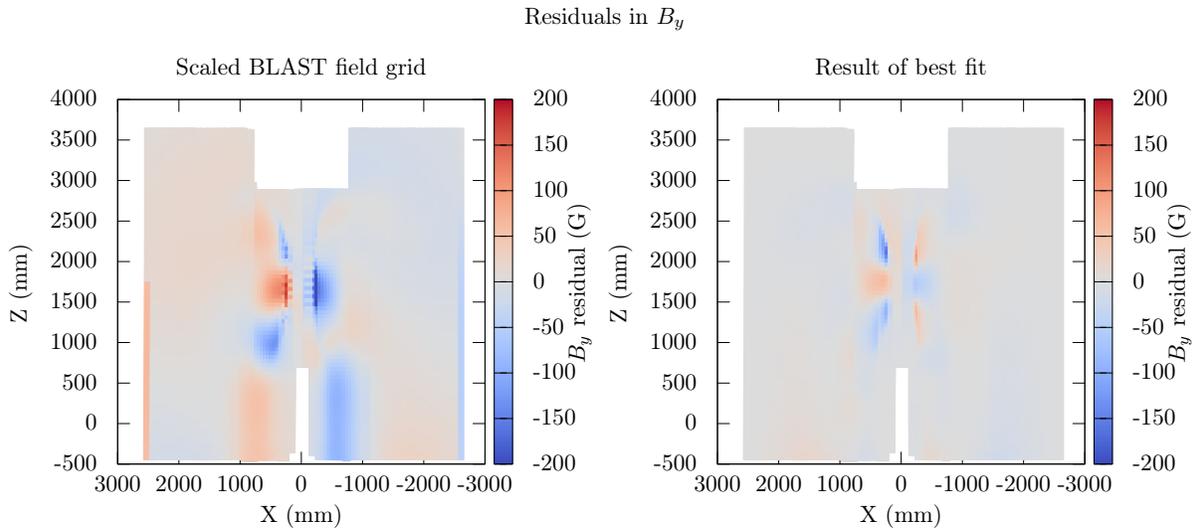


Figure 10: The new field map is a big improvement from the scaled BLAST grid that was being used.

4.4 Spline Interpolation

The Biot-Savart calculation of the magnetic field using the model was fast enough to perform a fit in a few minutes, but not nearly fast enough for simulating track propagation. The OLYMPUS Geant4 Simulation needs to query field values at several thousand positions to simulate a single track. To achieve the speed to make this possible, the magnetic field was calculated via Biot-Savart on a discrete regular grid, and used as an input for a fast spline interpolation, which could return an approximate field value at any intermediate position. The method of

interpolation on a grid had been used up until this point using the grid produced for the BLAST analysis, however the interpolation was only linear, and thus had discontinuous first derivatives.

Interpolation would have been possible directly on the measurement data. There were three reasons for not pursuing this approach. Primarily, the measurement data only covered a limited amount of space, and extrapolation was needed as well as interpolation. However in addition to this, the algorithms for interpolation on a non-uniform grid are much slower than those on a uniform grid. Lastly, the presence of a fitting step in between the measurements and interpolation served to smooth the data, and remove errors that might be associated with faulty survey measurements.

The grid points were chosen to be the same as those in the BLAST grid. The range in X was ± 2.5 m, the range in Y was ± 1.0 m, and the range in Z was from $Z = -0.5$ m to $Z = 3.5$ m. The grid spacing was 50 mm. The grid had a total of 418,241 points.

Several big improvements were added to the interpolation scheme. The first was to use an algorithm [2] that uses both information about the function and its derivatives at the grid points. Secondly, the splines were computed using basis functions chosen so that the polynomial coefficients were simply the function and derivative values at the grid points, reducing the amount of information which needed to be stored in memory. A method in the “coilfit” program called “printOLYMPUSGrid” can be used to write out a binary data file in the same format needed by the magnetic field class in the analysis repository.

4.5 Implementation

To make the magnetic field information available for multiple applications, a new static library was written for the Magnetic Field. The library contains the “Magnetic_Field,” class, which takes care of reading the binary data file produced by “coilfit.” The class only needs to be passed the name of the data file to be read and the value of current (in Amperes) in the coils.

The binary data file with the field grid from this analysis has been included in the repository as the file “olympus_v1.deriv.bin.” If improvements are made, they will be released with incremented version number. The files are installed in the directory “.olympus/shared/MagneticField.”

The Magnetic_Field class inherits from the Geant4 “G4MagneticField” class and matches the expected Geant4 format so that it naturally fits into the Simulation. To get the magnetic field at a point, one simply calls the “GetFieldValue” method, which takes two arguments. The first is an array of four doubles: the (x, y, z, t) spacetime point. (Since our magnetic field is not time-varying, t is ignored). The second argument is an array of six doubles into which the results are written in the form $(B_x, B_y, B_z, E_x, E_y, E_z)$. (The electric field components will always be returned as 0.0.)

5 Summary

Over the past few months the OLYMPUS magnetic field was measured, and the data analyzed to provide an updated field map. Between March and May, the Magnet was measured at DESY using a 3D Hall probe attached to a system of translation tables. Surveying was an integral part of the process so that the probe position was well known at every measurement point. The magnetic field and survey data were analyzed during the past few months at MIT. The data were fit with a model whose field was calculated by the Biot-Savart law. Once a suitable set of parameters were found, the magnetic field was calculated on a grid of points, which were passed as data for a spline interpolation. This was integrated into a new magnetic field library, which can be used by any analysis application.

References

- [1] K. A. Dow, T. Botto, A. Goodhue, D. K. Hasell, D. Loughnan, and others, “Magnetic field measurements of the BLAST spectrometer,” *Nucl.Instrum.Meth.*, vol. A599, pp. 146–151, 2009.
- [2] F. Lekien and J. Marsden, “Tricubic interpolation in three dimensions,” *International Journal for Numerical Methods in Engineering*, vol. 63, no. 3, pp. 455–471, 2005.