ARGONNE NATIONAL LABORATORY

23 July 1992 LIGHT SOURCE NOTE: LS–

INTERPOLATION OF HALL PROBE CALIBRATION DATA

David W. Carnegie

Advanced Photon Source
Argonne National Laboratory
9700 S. Cass Ave., Argonne, IL 60439-4814

ABSTRACT

Calibrated Hall-effect magnetic-field sensors will be used to map the magnetic field in insertion devices. Typical calibration data give the magnetic field as a function of measured signal and temperature on a two-dimensional grid. We need to calculate the magnetic field from the two measured signals. As an example, this work uses the calibration data supplied with a Hall-effect measurement system from Group 3 Technology Ltd. Detailed field-versus-signal data are given for three calibration temperatures for each of four gain settings.

Two methods for performing the interpolation are presented for a fixed gain setting. The first method fits the three field-versus-signal data sets to three polynomials minimizing the sum of squares of the errors and then interpolates for the temperature. The second method uses a bivariate interpolation routine. In this method, there are no residual errors at the calibration points. The two methods are compared. The selection of the method used will depend on what errors are present in the calibration data.

1. Introduction

The Magnetic Measurement Facility of the APS ID Group will map the magnetic field in insertion devices before they are installed on the Storage Ring. Calibrated Hall-effect magnetic-field probes will be used. These devices are intrinsically non-linear, and corrections need to be applied to their output. Several systems on the market perform digital correction to the signals but are too slow for our purposes. We plan to use Hall sensors with temperature monitoring and analog interfaces to increase the reading rates.

In this report, we will review a particular system manufactured by Group 3 Technology Ltd. [1] that was calibrated by the vendor. The calibration data consist of tables of numbers relating the output voltage to the field at three temperatures. Our problem is to make the best use of this data. The field mapping system will record the field and temperature signals. Thus, we must perform a two-dimensional interpolation to calculate the field.

In the first section, a description of the system and calibration data is presented. The second section describes an interpolation scheme that fits the data to a polynomial for a particular range and temperature. Three such fits would be used for a particular gain setting of the system. The temperature effects are calculated by a linear interpolation between two of the fit polynomials. This is followed by a section that describes a bivariate interpolation utilizing splines. The report concludes with a comparison of the two methods and recommendations for our own field sensor calibration system.

2. Description of the Group 3 Hall Probe System

The Group 3 Hall probe system [1] consists of a Hall sensor and thermocouple mounted together on an aluminum plate. This unit is connected to an analog interface that supplies current to the Hall plate and amplifies the field and temperature signals in which we are interested. Specifically, this work refers to probe model 11000007 (serial number 07018,) and Model HPI Probe Interface (serial number 01500013). There is one additional probe (serial number 07019) calibrated to the same interface. The Hall sensors are not linear to the accuracy we require, so they are commonly calibrated. Temperature effects are also taken into account in the calibration of these units. The analog interface has four gain settings summarized in Table 2.1 below. The first column refers to the relay that selects the range.

Table 2.1: HPI Field Ranges

Relay	Range
K3	3,000 Gauss
K2	6,000 Gauss
K1	12,000 Gauss
None	30,000 Gauss

The maximum output field signal in each range is about 3 volts. This output is, of course, bipolar corresponding to the bipolar nature of the measured field. The output temperature signal ranges from about 0 volts for 0°C to about 2.5 volts at 50°C. In the following, the K1 range was selected because it corresponds to the field measuring range of interest for Undulator A.

For each range selected, the manufacturer calibrated the probe at three temperatures: 0°C, 25°C, and 50°C. Thus, there are a total of twelve calibration scans of magnetic field versus field output signal. For a particular range, three scans are involved in the interpolation. Table 2.2 shows the probe temperature calibration data for probe serial number 7018. The plot in Fig. 2.1 shows the probe temperature as a function of the interface signal.

Table 2.2: Group 3 probe serial number 7018 calibration temperature data.

Temperature (Celsius)	Signal (volts)
0.03	-0.03021
24.94	1.24291
49.99	2.53146

Fig. 2.1: Plot of the calibration temperature as a function of the signal from the HPI interface.

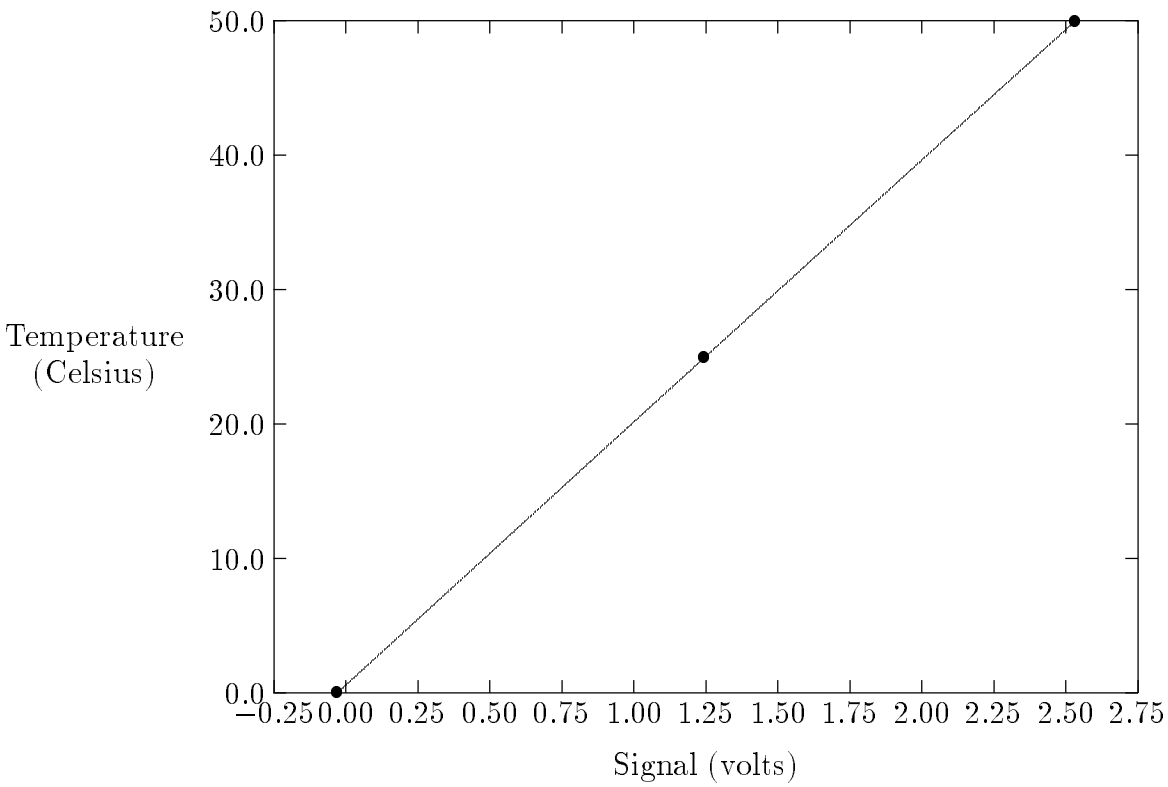

Fig. 2.2 shows the actual calibration data for this probe at 25°C. It is difficult to see the differences between the plots for the three temperatures on a single graph. So, to get an estimate of the effects of temperature on the calibration of the probe, Table 2.3 contains the coefficients from linear fits for each of the calibration temperatures. Note that there are temperature dependencies in both the intercept and the slope of the fits. These numbers are useful in estimating errors in measurements due to temperature changes during field mapping. We can estimate that the overall scale factor or slope changes at 0.15 Gauss/volt $\cdot^{o}$ C and that the intercept shifts by about -0.03 Gauss/°C. We will see later that a linear fit is not accurate enough for our purposes.

Fig. 2.2: Plot of the calibration data for probe serial number 7018 at 25°C.

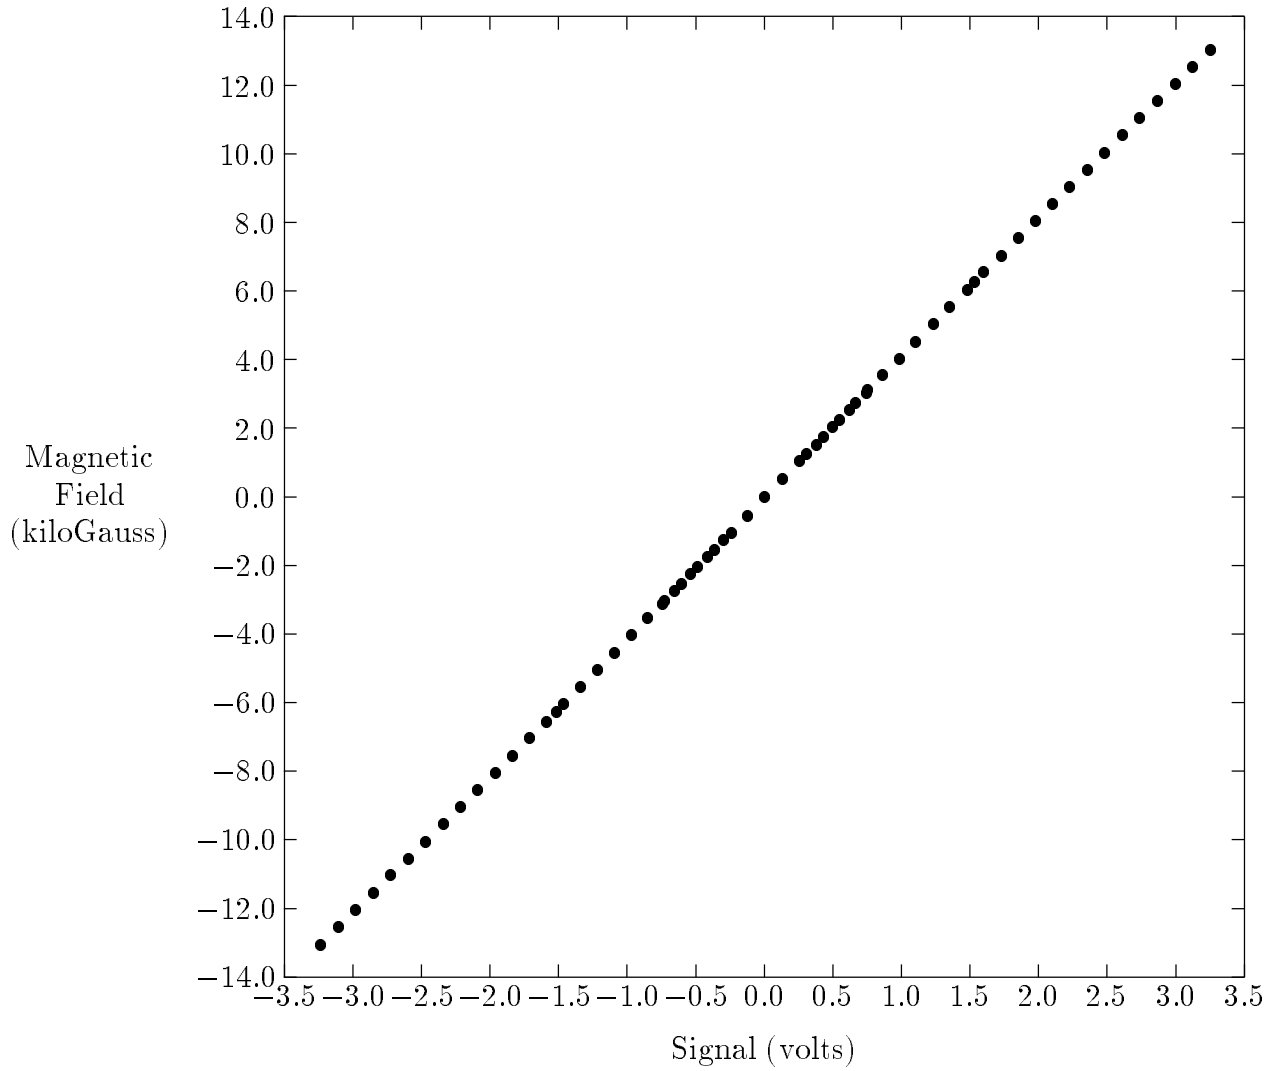

Table 2.3: Results of linear fits to the calibration data for probe 7018 in the K1 range.

Temperature (Celsius)	Slope (Gauss/volt)	Intercept (Gauss)
0.03	4055.47	-46.30
24.94	4059.01	-46.83
49.99	4063.01	-47.79

3. Polynomial Fits to the Calibration Data

The first method considered fits each of the calibration scans to a polynomial in the field signal. Fitting high-order polynomials to data in a straight-forward manner is prone to errors because of the ill-conditioned matrices generated in the analysis. An alternative method fits the data to a series of orthogonal polynomials avoiding the ill-conditioning problem. Routines from IMSL [2] that perform such a fit and decode the coefficients were used. Routine RLFOTH does the fitting, RLDOPM decodes the coefficients, and RLOPDC computes the predicted values from the fit model. A description of the algorithm used is contained in Ref. 3. Fits were performed for orders 1 through 15 with the results printed out for each order. The most useful results for evaluating a fit for interpolation are the residuals and the error sum of squares. The latter quantity is shown in Fig. 3.1, where it is plotted as a function of the order of the fit. This information can be used in selecting the highest order fit that should be used. The step structure occurs because of the odd symmetry of the original data. This symmetry can be observed in Fig. 3.2, a plot of the residuals from a first order (linear) fit. There is little improvement in the fit when going from an odd order to an even order. A large improvement occurs when the next odd order is added to the fit. The last significant improvement occurs between order 10 and 11. There is little improvement in the higher order fits. The 11th order fit is the best fit achievable with this data. Higher order fits may introduce "wiggles" in the function that do not represent the device being calibrated. When we look at the residuals for the 11th order fit (Fig. 3.3), there is no clear pattern as there is in the 1st order fit. This lack of pattern represents errors in the measurements themselves. We see a tight distribution about zero for all but four extreme points. It is interesting that these "problem" points are near zero field where unipolar power supplies used in laboratory magnets are reversed with relays. There are also remanent fields in the electromagnet that may not be as uniform as those at higher fields. We would probably obtain a better representation of the actual behavior of the Hall sensor if the fit were repeated without these four extreme points.

Without these four points, the magnitude of the residuals is less than about 0.1 Gauss. The maximum field in the K1 range is 12,000 Gauss. Therefore, the residuals are smaller than about one part in 10^5 when compared to the maximum field. The standard deviation of the residuals changes from 0.099 Gauss with the four points included to 0.068 Gauss with them removed for the 11th order fit. In practice, it is a good idea to select the order of the fit from a plot like Fig. 3.1 and then examine the residuals for extreme points. The extreme points should be removed from the input data and the fit repeated.

An actual calculation for arbitrary field and temperature is done in two steps. The measured probe temperature signal will be in one of two possible domains. Roughly speaking, it's either between 0°C and 25°C or between 25°C and 50°C. In the first case, the field signal is used as input to both polynomials for 0°C and 25°C. A linear interpolation is used on the output to correct for temperature. A similar procedure is used for the other temperature range.

The effects of temperature on the computed field can be seen if the output field is calculated from the *same* input field signal in each of the three polynomials. Table 3.1 shows the results of these calculations for three different input signals spanning the field signal range. Note that the zero field signal does not correspond to zero field. The zero offset must be determined by placing the probe in a zero-Gauss chamber and measuring that signal.

The plots of the zero field signal and positive field signal from Table 3.1 are given in Fig. 3.4 and Fig. 3.5. Note the discontinuities in the slopes at 25°C in each of these plots. This is an artifact of the interpolation scheme. It is likely that the actual behavoir of the Hall sensor is smooth in this region. Without more detailed calibration near room temperature, it is difficult to evaluate the error from this effect. A quadratic function could be used for the temperature interpolation instead of the piecewise linear function. However, the difference between a quadratic function and the piecewise linear interpolation is less than the 0.01-Gauss resolution of the supplied calibration data.

Fig. 3.1: Plot of the error sum of squares for fits from order 1 through 15. The last significant improvement in the fit occurs between order 10 and 11. We should therefore use the order 11 fit.

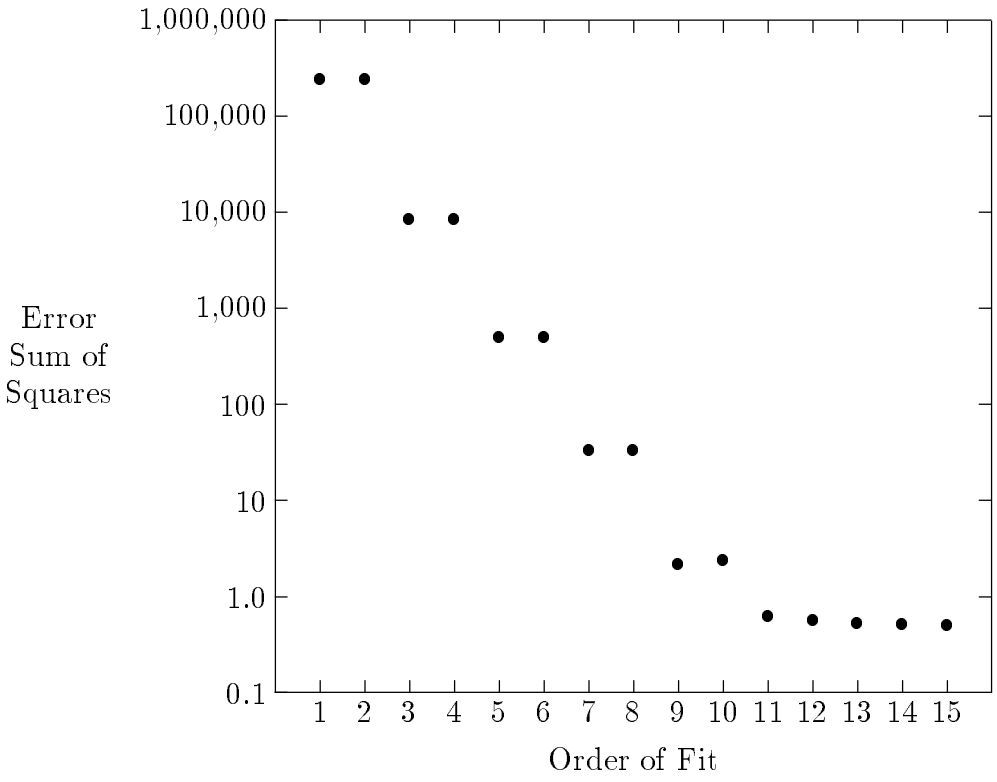

Fig. 3.2: Plot of the residuals from a first order fit for probe serial number 7018, K1 range at 25°C.

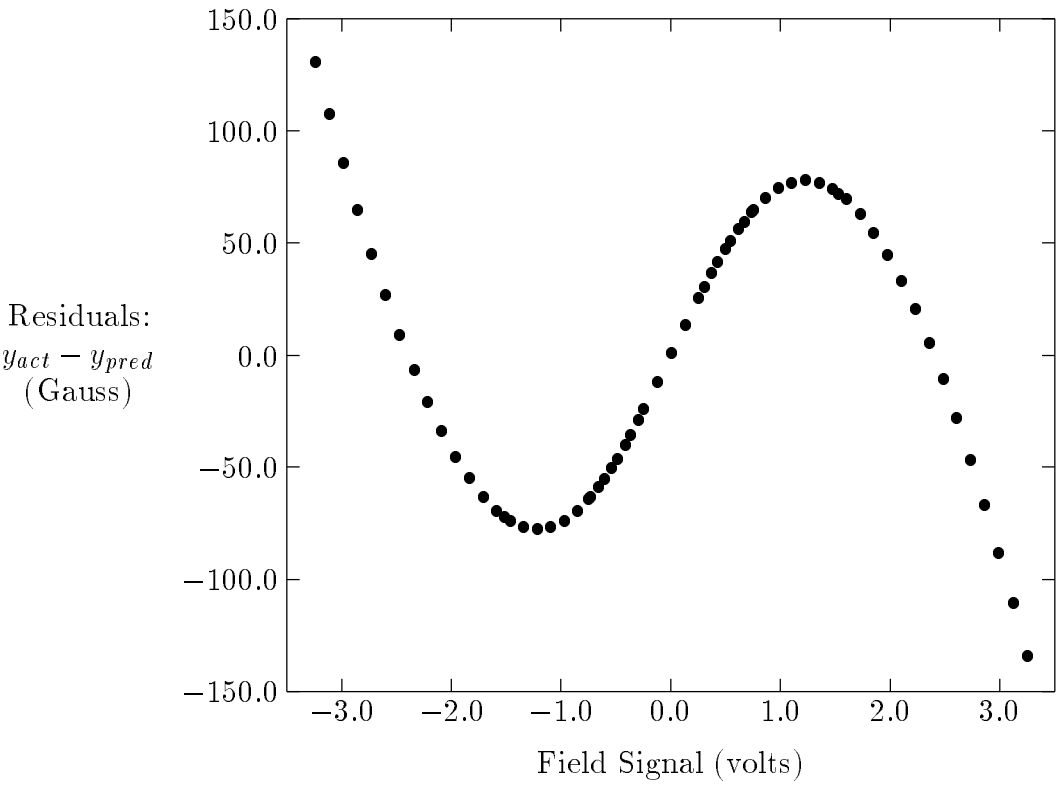

Table 3.1: Computed fields for fixed input signal as a function of temperature. For each input signal, the output was computed at three calibration temperatures.

Field Signal (volts)	Temperature Signal (volts)	Temperature (Celsius)	Computed Field (Gauss)
-3.23314	-0.03021	0.03	-13025.31
-3.23314	1.24291	24.94	-13039.28
-3.23314	2.53146	49.99	-13055.09
0.0	-0.03021	0.03	-46.60
0.0	1.24291	24.94	-47.23
0.0	2.53146	49.99	-48.16
3.25704	-0.03021	0.03	13026.37
3.25704	1.24291	24.94	13039.56
3.25704	2.53146	49.99	13053.76

Fig. 3.3: Plot of the residuals from an eleventh order fit for probe serial number 7018, K1 range at 25°C. Note the four extreme residuals that are candidates for elimination from the data set.

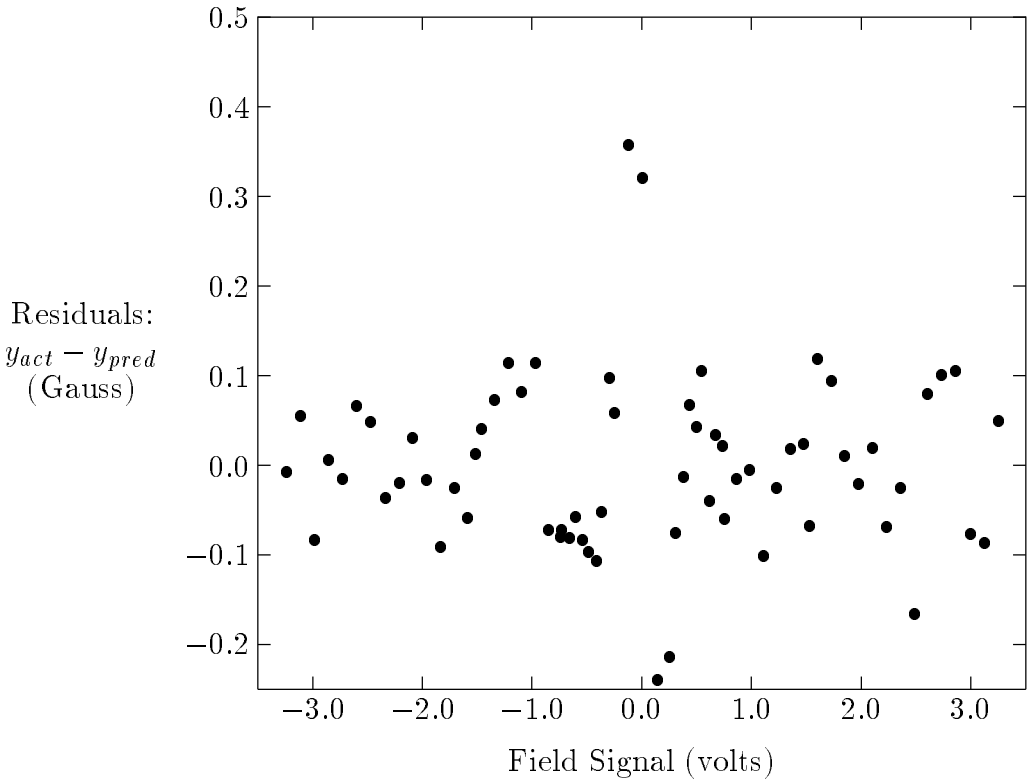

4. Bivariate Interpolation for Irregularly Distributed Data Points

There are a variety of interpolation methods available for one- and two-dimensional data. The routines in the IMSL [2] library will perform spline interpolation for points on a regular rectangular grid. In this application, the calibration points are not on such a grid. Algorithm 526 from the ACM [4] will perform this interpolation. The main feature of this type of interpolation is that there are no residuals at the calibration points. That is, the calibration points are reproduced exactly. This is an advantage if the data are tabulated values of an exact function. In our case, the data points are experimentally measured values with errors. This technique is still useful if the accuracy of the calibration data is much higher than the desired accuracy of the interpolated points. This will be described in more detail in the concluding section.

Fig. 3.4: Plot of the computed field values for zero input signal at the calibration temperatures.

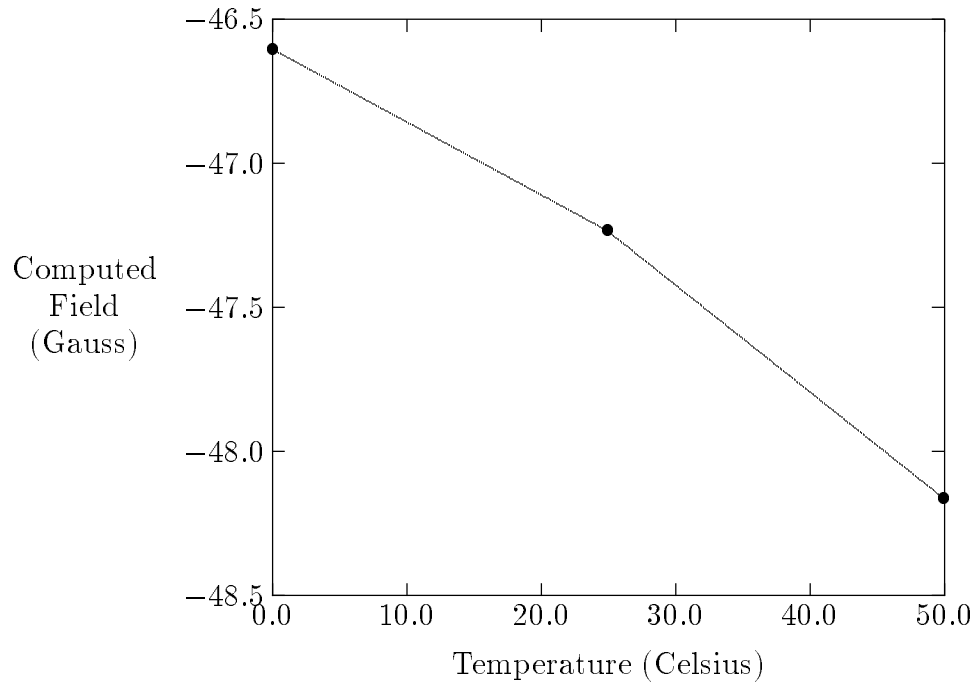

Fig. 3.5: Plot of the computed field values for the high input signal at the calibration temperatures.

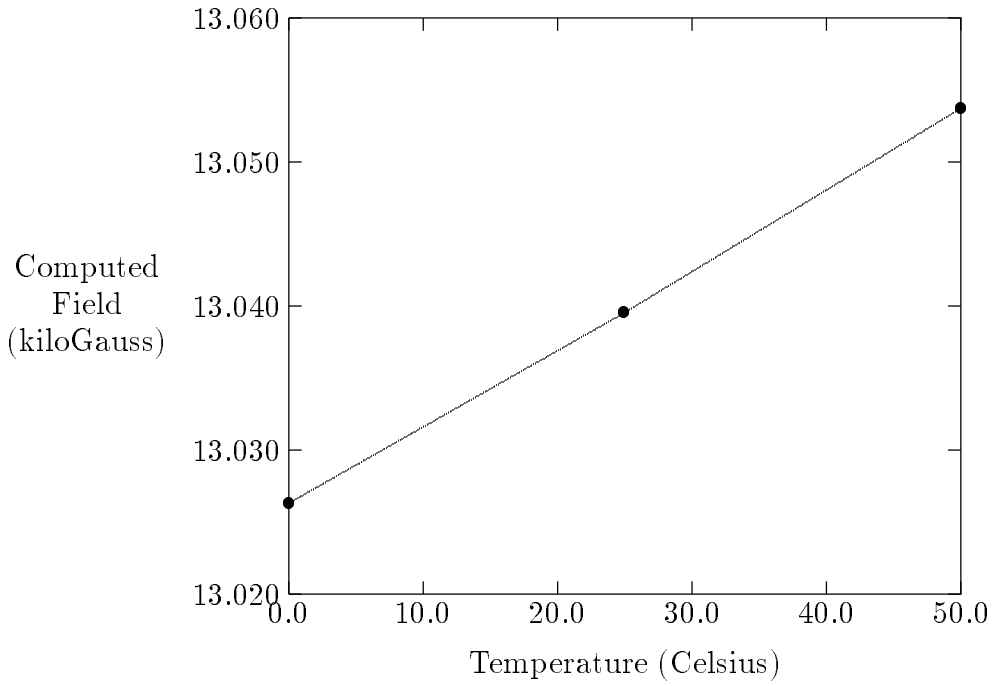

Algorithm 526 uses in the calibration data for all three temperatures together. The plane defined by the independent variable of the calibration data is divided into triangles. The data to be interpolated are entered, and the calculations performed. The algorithm locates in which triangle each input data point is located and performs the interpolation. Table 4.1 contains the output from the algorithm for three input signals as a function of input temperature. Fig. 4.1 is a plot of the output for 0.0 volt input signal and a range of temperatures, and Fig. 4.2 plots the output at the maximum field for the same temperatures. The main difference between these plots and those shown in the previous section is the smooth interpolation around 25°C. The algorithm used here results in almost linear interpolation away from the central calibration temperature by constraining the slopes at the extreme temperature values to equal the slope of the lines connecting the end points to the central point.

Fig. 4.1: Plot of the output of Algorithm 526 for 0 volt input as a function of temperature.

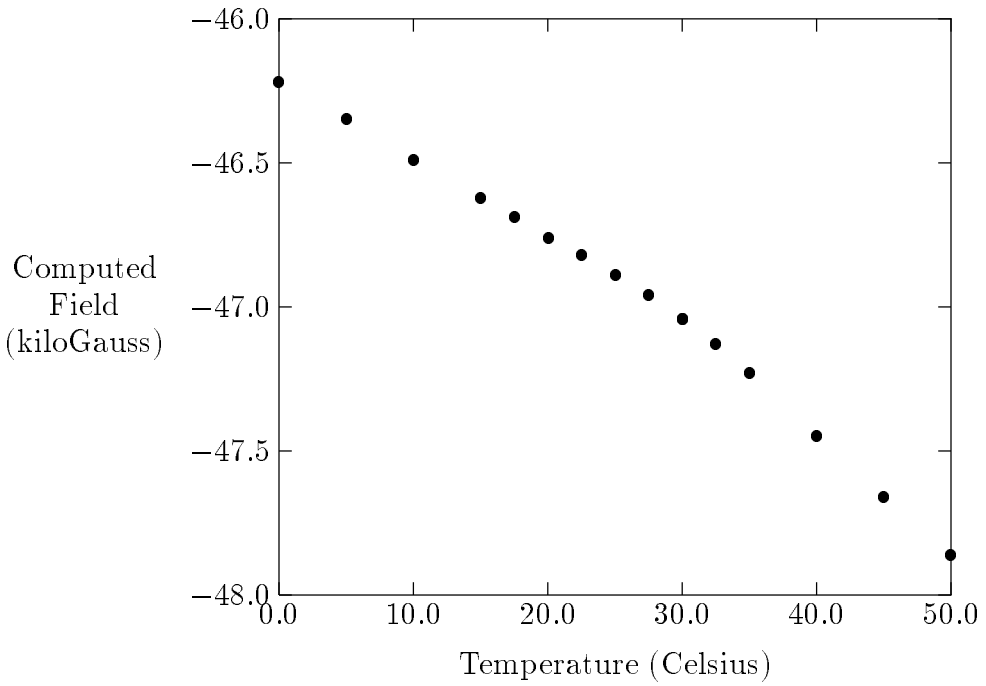

Table 4.1: Output from Algorithm 526 for three input signals as a function of temperature.

Temperature (Celsius)	0.0 volt Signal	-3.23314 volt Signal	3.25704 volt Signal
0.03	-46.22	-13025.26	13026.38
4.99	-46.35	-13028.05	13029.01
10.01	-46.49	-13030.87	13031.67
15.02	-46.62	-13033.70	13034.34
17.53	-46.69	-13035.11	13035.67
20.04	-46.76	-13036.52	13037.00
22.55	-46.82	-13037.94	13038.34
25.05	-46.89	-13039.35	13039.67
27.55	-46.96	-13040.77	13041.01
30.04	-47.04	-13042.50	13042.37
32.53	-47.13	-13044.30	13043.77
30.04	-47.04	-13042.50	13042.37
35.02	-47.23	-13045.83	13045.19
40.01	-47.45	-13048.90	13048.09
44.99	-47.66	-13051.99	13050.99
49.96	-47.86	-13055.11	13053.83

Fig. 4.2: Plot of the output of Algorithm 526 for 3.25704 volts input as a function of temperature.

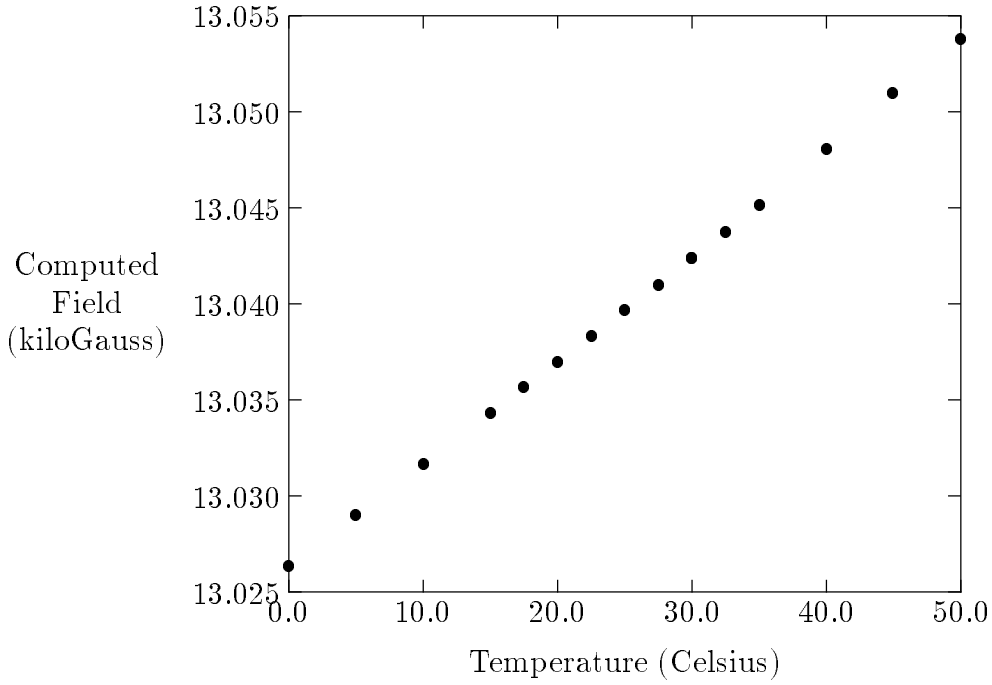

5. Comparison of the Two Methods

This report presents the results of an exercise to determine a reliable interpolation method. The results can be used in the interpolation of the data from the two Group 3 probes now on hand for our R&D work. In the future, the facility will have a calibration electro-magnet equipped with an NMR field measurement instrument for calibrating Hall probes. The complete system will be able to run calibration curves at several temperatures. The selection of interpolation polynomials will be a part of the software controlling the calibration system.

The 11th order fit described in the third section of this report has, with the four extreme points removed, an rms residual of about 0.07 Gauss. This is 70 times the resolution of the NMR system and 7 times the resolution of the calibration data supplied by the vendor. In the plot of the residuals, this looks like noise in the experimental data. The polynomial fit smooths out this noise.

The interpolation scheme described in the fourth section reproduces the calibration points exactly. Noise in the calibration data is passed on to the output data. This noise would appear as an elevated noise level in the Fourier analysis of the periodic field of the insertion device.

By looking at the error sum of squares and the plots of the residuals, we are able to select the best fit order for the calibration data. The remaining residuals are randomly distributed about zero and appear to be random errors in the calibration data. These errors are filtered out by the fit procedure. This filtering does not take place in the algorithm described in section four. Thus, our preference is to use the polynomial fits to smooth out the measurement errors and noise present in the calibration data.

The measurements on the insertion devices will take place in a room with the temperature controlled to better than $\pm 0.5^{\circ}$C. The average temperature will be on the order of 25°C. We will not require the wide temperature range supplied with the vendor calibration data. To better characterize the behavior of the Hall probe around 25°C, we should perform calibration scans at three to five temperatures closer to 25°C. A recomendation for initial work on calibration would be scans at 20°C, 22.5°C, 25°C, 27.5°C, and 30°C, which would supply the detail necessary to evaluate the type of interpolation required for the temperature signal. Without this additional detail, it is difficult to estimate the magnitude of the errors from the temperature signal.

Acknowledgment

I would like to acknowledge many helpful conversations with Fred Lopez of ASD. He spent time working on the calibration of similar Hall probe systems for the Magnetic Measurement Group in his division. In particular, he tried out Algorithm 526 on his data when I first came across it. His experience with this package is similar to mine.

References

1. Group 3 Technology Ltd., P.O. Box 71-111 Rosebank, Auckland 7, New Zealand. The U.S. distributor is GMW, P.O. Box 2578, Redwood City, CA 94064.
2. *User's Manual: IMSL Library Edition 9.2*, IMSL, Inc., 1984.
3. "Generation and Use of Orthogonal Polynomials for Data-Fitting with a Digital Computer," George E. Forsythe, J. Soc. Indust. Appl. Math., **5**(2)74-88 (1957).
4. "Algorithm 526: Bivariate Interpolation and Smooth Surface Fitting for Irregularly Distributed Data Points," Hiroshi Akima, ACM Transactions on Mathematical Software, **4**(2)160-164 (1978).