



Terms and Conditions of Use of Digitised Theses from Trinity College Library Dublin

Copyright statement

All material supplied by Trinity College Library is protected by copyright (under the Copyright and Related Rights Act, 2000 as amended) and other relevant Intellectual Property Rights. By accessing and using a Digitised Thesis from Trinity College Library you acknowledge that all Intellectual Property Rights in any Works supplied are the sole and exclusive property of the copyright and/or other IPR holder. Specific copyright holders may not be explicitly identified. Use of materials from other sources within a thesis should not be construed as a claim over them.

A non-exclusive, non-transferable licence is hereby granted to those using or reproducing, in whole or in part, the material for valid purposes, providing the copyright owners are acknowledged using the normal conventions. Where specific permission to use material is required, this is identified and such permission must be sought from the copyright holder or agency cited.

Liability statement

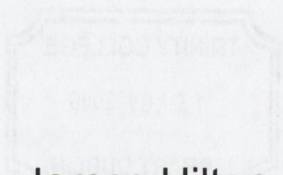
By using a Digitised Thesis, I accept that Trinity College Dublin bears no legal responsibility for the accuracy, legality or comprehensiveness of materials contained within the thesis, and that Trinity College Dublin accepts no liability for indirect, consequential, or incidental, damages or losses arising from use of the thesis for whatever reason. Information located in a thesis may be subject to specific use constraints, details of which may not be explicitly described. It is the responsibility of potential and actual users to be aware of such constraints and to abide by them. By making use of material from a digitised thesis, you accept these copyright and disclaimer provisions. Where it is brought to the attention of Trinity College Library that there may be a breach of copyright or other restraint, it is the policy to withdraw or take down access to a thesis while the issue is being resolved.

Access Agreement

By using a Digitised Thesis from Trinity College Library you are bound by the following Terms & Conditions. Please read them carefully.

I have read and I understand the following statement: All material supplied via a Digitised Thesis from Trinity College Library is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of a thesis is not permitted, except that material may be duplicated by you for your research use or for educational purposes in electronic or print form providing the copyright owners are acknowledged using the normal conventions. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone. This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.


Computational Modelling of Novel Permanent Magnetic Designs



James Hilton
Supervisor: Dr Sara McMurry

A thesis submitted to the University of Dublin
In application for the degree of Doctor of Philosophy

Department of Pure and Applied Physics
University of Dublin
Trinity College
September 2005



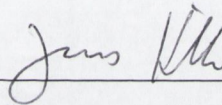
Thesis
8066

Declaration

This thesis is submitted by the undersigned for the examination for the degree of Doctor of Philosophy to the University of Dublin. It has not been submitted as an exercise to any other university.

Apart from the advice and assistance mentioned in the text, this thesis is entirely my own work.

I agree that this thesis may be freely lent or copied.

A handwritten signature in cursive script, appearing to read "James Hilton", is written above a horizontal line.

James Hilton

September 2005

Summary

Modern permanent magnetic materials provide unique design opportunities as their low susceptibility makes them effectively transparent to external magnetic fields. This also simplifies the calculation of magnetic fields in devices incorporating them. We have developed two software programs to simulate the magnetic field from arrays of such permanent magnets. The first is based on the magnetic charge model, and the second uses a magnetostatic finite element method. We have used this software to develop novel designs for three particular types of application:

1. A design to improve the homogeneity within short length Halbach cylinders. We compared this new design to existing methods of increasing the field homogeneity. Our design achieves an approximate twofold increase in the homogeneity in central plane of a Halbach cylinder over the next best design.
2. Three designs producing a constant magnetic field gradient using only permanent magnetic material. In two of these designs the gradient is parallel to the field, and in one the gradient is orthogonal to the field. These designs can produce around 30 T/m in a working region of length 1cm.
3. Three designs producing a constant product of magnetic field (B) and magnetic field gradient (dB) using only permanent magnetic material. In two of these designs the gradient is parallel to the field, and in the other it is orthogonal to the field. We commissioned two scalings of the design with gradient orthogonal to the field to be manufactured, and the field measured within the design is shown to be in excellent agreement with our simulations. The second design with the gradient parallel to the field can achieve a constant BdB of magnitude $20T^2/m$ over a working area of length 1cm, and the design with the gradient orthogonal to the field can achieve a constant BdB of magnitude $30T^2/m$ over a working region of length 1cm.

Acknowledgements

Acknowledgements seem to be the most tricky part of writing a thesis, as there are so many people who have contributed both their time and effort to help me complete the work of the past few years, all of which have been condensed into this volume.

I would like to start by thanking my supervisor, Dr. Sara McMurry, to whom I am extremely grateful for giving me the opportunity to undertake this PhD in Trinity college, and whose guidance and encouragement throughout the past few years I could not have done without. I would also like to thank Enterprise Ireland for their strategic grant (ST/2000/025) for funding over the course of the project.

I would like to thank all of my friends in the computational lab (including the chemists!): Gary and Aengus for having to sit beside me during particularly difficult debugging sessions, Paul for keeping us all amused throughout the day, Gemma for being her kindly self and having to put up with us all, and Aron for bringing his version of normality to the lab.

To my friends outside the lab, who have seen me through the past few years, Sarah and Johannes; thank you for everything, better friends I could not ask for. I am also grateful to all my other friends, who have made my time here so memorable and enjoyable.

Finally, I would like to thank my family for their constant love and support over these past few years.

Contents

Declaration	2
Summary	3
Acknowledgements	4
Contents	5
1 Introduction	9
1.1 Magnetism	9
1.1.1 Introduction	9
1.1.2 Magnetic Materials	14
1.1.3 Designs Exploiting Modern Materials	18
1.1.4 Designs Producing Constant Force on Dipoles	23
1.1.5 Project Outline	27
2 Theoretical Background	28
2.1 Magnetic Fields	28
2.1.1 Maxwell Equations	28
2.1.2 Magnetic Vector Potential	29
2.1.3 Magnetic Boundary Conditions	30
2.1.4 Magnetic Dipoles	32
2.1.5 Forces and Torques on Magnetic Dipoles	33
2.1.6 Magnetisation	33
2.2 Magnetostatic Field Calculations	35
2.2.1 Charge Model	35
2.2.2 Current Model	37
2.2.3 Green Function Expansion	38
2.3 Magnetic Structures	40
2.3.1 One-Sided Flux	40
2.3.2 Infinite Length Halbach Cylinder	44
2.3.3 Finite Length Halbach Cylinder	47
3 Numerical Methods for Magnetic Field Calculations	51
3.1 Surface Charge Model	51
3.1.1 2D Surface Charge Model	51
3.1.2 2D Transformations	53

3.1.3	2D Computational Realisation	54
3.1.4	3D Surface Charge Model	56
3.1.5	3D Transformations	57
3.1.6	3D Computational Realisation	60
3.2	Finite Element Method	61
3.2.1	Weighted Residual Methods	62
3.2.2	Finite Element Method	66
3.2.3	Application to Magnetostatic Systems	71
3.2.4	Solution of the Matrix System	76
3.3	Automated Mesh Generation	86
3.3.1	Overview	86
3.3.2	Delaunay Incremental Insertion Algorithms	88
3.3.3	Geometric Primitives	90
3.3.4	Fast Point Location	93
3.3.5	Ruppert Algorithm	94
3.3.6	Shewchuk Algorithm	97
3.3.7	Boundary Recovery	100
3.3.8	Implementation	101
3.3.9	Computational Realisation	102
4	Improvement of Homogeneity in Halbach Cylinders	105
4.1	Introduction	105
4.1.1	Effects of Cylinder Length on Field Homogeneity	105
4.1.2	Increasing Field Homogeneity	108
4.2	Split Halbach Cylinder	110
4.2.1	Design Overview	110
4.2.2	Effect of Split Width on Homogeneity	113
4.3	Halbach Cylinder With Non-Uniform Inner Radius	115
4.3.1	Design Overview	115
4.3.2	Effect of Increasing Central Bore Radius	118
4.4	Halbach Cylinder With Soft Iron Insert	119
4.4.1	Design Overview	119
4.4.2	Effect of Insert Thickness on Homogeneity	120
4.5	Comparison of the Methods of Improving Homogeneity	122

5	Designs Producing a Constant Field Gradient	128
5.1	Introduction	128
5.2	Halbach Cylinder With Linearly Varying Inner Radius	128
5.2.1	Design Overview	128
5.2.2	Design Summary	132
5.3	Halbach Cylinder With Offset Bore	134
5.3.1	Design Overview	134
5.3.2	Design Summary	135
5.4	Rolled Flux Sheet	137
5.4.1	Design Overview	137
5.4.2	Design Summary	140
5.5	Comparison of Designs Producing a Constant Field Gradient	143
6	Designs Producing a Constant Product of Field and Field Gradient	145
6.1	Introduction	145
6.2	Uniformly Magnetized Cylinder With Linearly Varying Inner Radius	146
6.2.1	Design Overview	146
6.2.2	Effect of Segmentation	149
6.2.3	Design Summary	149
6.3	Halbach Cylinder With Linearly Varying Inner Radius	151
6.3.1	Design Overview	151
6.3.2	Design Summary	153
6.3.3	Performance of Magnets Built According to Design	156
6.4	Rolled Flux Sheet	162
6.4.1	Design Overview	162
6.4.2	Design Summary	164
6.5	Comparison of Designs Producing a Constant Product of Field and Field Gradient	166
7	Conclusions	168
	Appendix	171
A.1	Derivations	171
A.1.1	Rectangular Charge Sheet	171
A.1.2	Triangular Charge Sheet	173
A.1.3	Cylindrical Halbach With Linear Inner Bore	177

A.1.4	Field From a Conical Magnet	180
A. 2	Effect of Varying Parameters For Designs In Section 5	182
A.2.1	Design of Section 5.2	182
A.2.2	Design of Section 5.3	185
A.2.3	Design of Section 5.4	186
A. 3	Effect of Varying Parameters For Designs In Section 6	190
A.3.1	Design of Section 6.2	190
A.3.2	Design of Section 6.3	192
A.3.3	Design of Section 6.4	195
References		198

1 Introduction

1.1 Magnetism

1.1.1 Introduction

Magnets, although familiar and extremely common, continue to exert fascination due to their seemingly magical properties. Since ancient times magnets, in the form of lodestone, have been attributed supernatural powers and were often used in divination



Figure 1.1 – Alive with magnetism.

and fortune telling. The Greek philosopher Thales even postulated that the magnet had a soul of its own. Even so, the historical significance of magnets would probably be very little had their use in navigation never been discovered. It is believed that around the 12th century both the Chinese and (possibly independently) the Europeans discovered that lodestone tends to align itself in a preferred direction. It was discovered soon after that an iron needle rubbed with loadstone exhibited the same effect. The discovery spread and was being used by the Arabs around the 13th century and the Scandinavians by around the 14th century. Magnets were therefore of great

importance to travel and trade and so instrumental to the spread of knowledge and ideas in the early stages of civilisation.

Magnets that immediately come to mind are the horseshoe magnet or fridge magnet, both examples of permanent magnets (with which this project is concerned). Of course, magnetic fields can also be generated by electric currents and time varying magnetic fields can generate electric currents; this is commonly put to use in devices such as electromagnets and electric motors, for either small electromagnetic actuators or larger industrial applications. Electromagnets currently provide the only source of very high magnetic fields, although resistive heating of the current carrying wire becomes a problem past a certain point. Although designs such as the Bitter solenoid (in which a small frog was famously levitated at the Nijmegen magnet laboratory¹) go some way to alleviating this problem the only practical way in which to achieve

extremely high magnetic fields is by using superconducting material (although much higher fields than these can be created over very short timescales using destructive methods).

The main advantage of permanent magnets over electromagnets is that they provide a constant flux output for no energy expenditure (along with no need for maintenance, control circuitry and possibly heat dissipation systems). They are also extremely compact and can be engineered to any shape or size depending on the design. Furthermore, as Furlani points out², if an electromagnet is shrunk by a factor L the current density of the electromagnet must be increased by $1/L$ to give the same field strength at a rescaled observation point, which is not the case for a permanent magnet. Therefore, for small high flux systems, permanent magnets provide the only viable field source.

In this project we looked at designs that produce high uniform fields and designs that produce constant field gradients to provide a constant force on magnetic dipoles, constructed entirely from modern high-energy ferromagnetic materials. It may seem somewhat strange that permanent magnets can be arranged in ways to produce a uniform field as, for example, a bar magnet can be approximated as a magnetic *point dipole* with field \mathbf{H} (given in polar co-ordinates relative to the dipole)³:

$$H_r = \frac{2m \cos(\theta)}{4\pi r^3}, H_\theta = \frac{m \sin(\theta)}{4\pi r^3} \quad [1.1.1]$$

Where m is the magnetic moment of the dipole. However, if the length of the bar magnet is much larger than both the height and width of the magnet, it can be represented by a *line dipole* with field of the form³:

$$H_r = \frac{m \cos(\theta)}{2\pi L r^2}, H_\theta = \frac{m \sin(\theta)}{2\pi L r^2} \quad [1.1.2]$$

Where L is the length of the line dipole. It can be seen that the line dipole obeys an inverse square rule as opposed to the inverse cube rule for the point dipole, but the crucial difference is that the factor of two in the radial component disappears in the expression for the line dipole. This means that for the line dipole, \mathbf{H} is orientated at a constant angle 2θ to the dipole, whereas for the point dipole the orientation of \mathbf{H} relative to the dipole is not constant (in fact, the orientation is $\theta + \tan^{-1}(\frac{1}{2} \tan(\theta))$). Also, the magnitude of \mathbf{H} is independent of θ for the line dipole, whereas the point dipole has a $\cos^2\theta$ dependence. The fact that the orientation of the field due to a line

dipole is simply related to the orientation of the dipole itself can be then exploited in magnetic design. For example, a set of line dipoles can be arranged with orientation such that the components of \mathbf{H} from the line dipoles cancel in all directions apart from one, giving rise to a uniform field.

Some examples of the application of high uniform magnetic fields include Faraday rotators, Magnetic Resonance Imaging (MRI) applications and spin-alignment in sputtering systems. Uniform fields are also used to rotate or guide electron beams in magnetrons and klystrons.

Faraday rotators

Faraday rotators rotate the polarisation angle of plane polarised light and are used in detectors and optical isolators. Plane polarised light can be expressed in terms of coherent left and right-handed circularly polarised light of equal amplitude; if one of these two components has a higher phase velocity than the other in the material (due, in this case, to an applied magnetic field) the plane of polarisation will rotate. This is called the Faraday or magneto-optical effect and was discovered by Faraday in 1845⁴.

The rotation angle, β , of the plane of the incident polarised light can be expressed as:

$$\beta = \nu B d \quad [1.1.3]$$

Where ν is the Verdet constant for the material and d is the distance the light travels through the material. Modern optical materials such as terbium gallium garnet have high Verdet constants, making rotation of 45° or more possible, which is crucial for optical isolators.




Figure 1.2 – A Faraday rotator rotates the polarization angle of plane polarized light by an angle β .

An optical isolator protects optical components from optical feedback, for instance in laser cavities. The rotation of the plane of polarization of the light beam is independent of propagation direction, so a Faraday rotator with a vertical polarizer at the input end and a 45° polarizer at the output end will prevent any feedback; light

coming back through the output will be polarised at 45° to the vertical, rotated by a further 45° into the horizontal plane and then blocked by the vertical polarizer at the input end. A high field, **B**, is important to rotators as the length, *d*, should ideally be kept as short as possible for economic reasons.

Magnetic Resonance Imaging

High uniform fields are also used in Magnetic Resonance Imaging (MRI)⁵. This is a non-invasive imaging procedure that uses the principles of Nuclear Magnetic Resonance⁶ to gather information about the molecular structure of a given sample. A strong and highly homogeneous magnetic field is used to align nuclear magnetic moments within the sample. Some moments will align parallel to the field (the ‘low’ energy state) and some will align anti-parallel to the field (the ‘high’ energy state). Boltzmann statistics tells us that the number of nuclear spins in the lower energy state will slightly outnumber those in the higher energy state, according to:

$$\frac{N_{higher}}{N_{lower}} = e^{-E/kT} \quad [1.1.4]$$

Where *E* is the energy difference between the spin states, *k* is the Boltzmann constant and *T* is the temperature in Kelvin. This difference in spin population gives a net magnetization, **M**, to the sample which is initially aligned with the applied field **B**. The nuclear spins precess around the applied field direction at the Larmor frequency ν ,

$$\nu = \gamma |\mathbf{B}| \quad [1.1.5]$$

Where γ is the nuclear gyromagnetic ratio. Once the spins are aligned, the sample is given a pulse of radio frequency radiation tuned to the Larmor frequency to produce a resonant transition from the lower energy state to the higher energy state. The pulse duration is long enough to ensure that the spin population will approximately balance (i.e. $N_{higher} \approx N_{lower}$). This effectively rotates the net magnetization vector, **M**, from aligned with the applied field, **B**, into the plane orthogonal to the applied field. The applied field is defined to be in the z-direction, so that the orthogonal plane is the x-y plane. When the applied radio frequency field is then turned off, **M** relaxes back into alignment with **B**. Microscopically this involves the nuclei losing energy $h\nu$, with frequency ν , and switching back into their original energy state. Macroscopically this involves the vector **M** spiralling upward out of the x-y plane until it is fully re-aligned

with \mathbf{B} . The recovery of \mathbf{M} to equilibrium is given by an exponential function, the parameter of which varies depending on the type of molecule and is known as the time constant $T1$.

$$M_z = |\mathbf{B}|(1 - e^{-t/T1}) \quad [1.1.6]$$

Furthermore, the individual microscopic spins that make up \mathbf{M} (initially all in phase) gradually de-phase in the x-y plane due to microscopic interactions between the individual spins. This de-phasing is also given by an exponential function with parameter $T2$.

$$M_{xy} = |\mathbf{B}|e^{-t/T2} \quad [1.1.7]$$

Positional information in MRI is determined by superimposing static gradient fields on top of the uniform alignment field. Usually one field gradient parallel to an axis is used to select a slice within the sample whilst the sample is being radiated. Two further field gradients along the other axes are used, once the electromagnetic radiation has been turned off, to give positional information within the sample slice. The measured frequency responses can then be turned into an (2D slice) image by extracting positional data using Fourier transform methods. For medical MRI the nucleus under investigation is usually Hydrogen due to its abundance and large gyromagnetic ratio. The $T1$ and $T2$ parameters for hydrogen are different depending on the molecule of which the hydrogen is a component and can be used to differentiate, for example, healthy tissue from diseased tissue.

The most common magnets in MRI machines are usually superconducting magnets, as conventional electromagnets require far more power to produce the required fields. A good alternative to superconducting magnets is provided by permanent magnets as they require no power, no maintenance once set up, have high homogeneity and also (in the case of Halbach design) produce little stray fields outside the region of interest. However, permanent magnets used on this scale are extremely heavy and provide only limited field strength and uniformity. Also, ambient temperature fluctuations cause field instability which can have a large effect on the sensitive MRI measurements. They are, though, a viable option for MRI as inhomogeneities in the field can usually be removed by careful placement of small coils ('active' shimming) and placement of small amount of soft ferromagnetic ('passive shimming') material to tune the field.

Spin Alignment

Uniform fields are also of importance in controlling spin alignment in devices which exploit the spin of electrons (known as *spintronic* applications). Modern spintronic materials require precise control of spin-alignment during fabrication. This can be controlled using a uniform field source during the deposition (or sputtering) process as the spins align parallel to the applied field. Spintronics is a rapidly developing field and devices employing spin effects, such as the giant magnetoresistive effect (GMR), are just starting to be commercially exploited. Hard drive read heads now employ GMR tips for extreme sensitivity to the magnetic field of the disc, resulting in a great increase in storage capacity⁷. Magnetic RAM (MRAM)⁸ is also currently on the horizon, promising a non-volatile, and therefore more energy efficient, alternative to conventional DRAM.

1.1.2 Magnetic Materials

Magnetisation in materials is caused by magnetic dipoles on an atomic scale. These dipoles are not necessarily fixed and constant; they can appear in response to an applied field (diamagnetism), be present but only weakly coupled and only align in response to an applied field (paramagnetism) or be present and strongly coupled giving rise to a bulk magnetisation (in, for example, ferromagnetism)². Ferromagnetism is a special case in which individual microscopic magnetic *domains* within the material couple to each other. On the macroscopic scale it is not necessary to consider which mechanism gives rise to the presence of these dipoles; we can simply take into account the bulk effect of the individual microscopic contributions from *all* the dipoles by defining the *magnetisation vector*, \mathbf{M} , as the net magnetic dipole moment per unit volume (with units amperes/metre). \mathbf{H} is related to the magnetic field \mathbf{B} and magnetization vector \mathbf{M} by the definition:

$$\mathbf{H} \equiv \frac{1}{\mu_0} \mathbf{B} - \mathbf{M} \quad [1.1.8]$$

Where \mathbf{B} has units of Tesla and \mathbf{H} units of A/m. Therefore, in a magnetic material \mathbf{B} and \mathbf{H} are not necessarily parallel as the components of \mathbf{B} normal to boundaries must be continuous (to obey the first Maxwell equation) whereas \mathbf{H} has to obey the above relationship. In free space this relationship reduces to:

$$\mathbf{B} = \mu_0 \mathbf{H} \quad [1.1.9]$$

The magnetic polarisation vector \mathbf{J} , is defined by:

$$\mathbf{J} = \mu_0 \mathbf{M} \quad [1.1.10]$$

and is measured in Tesla. Working in terms of the polarisation vector is sometimes more convenient as it has units of Tesla, the same as \mathbf{B} .

In all ferromagnetic materials \mathbf{M} depends on the history of the sample under consideration due to the effect of individual domains of magnetization within the sample, giving rise to the phenomenon of *hysteresis*. This can be visualised by plotting the applied field (\mathbf{H}) against the magnetization (\mathbf{M}), and takes the form of a hysteresis loop (see figure 1.3, below).




Figure 1.3 – A typical hysteresis loop of a ferromagnetic material.

Initially, the domains within the sample are randomly orientated giving rise to zero bulk magnetic field (point A). As the applied field is increased the domains experience a torque that tends to align them parallel to the applied field. Within the domains each of the individual dipoles constituting the domains tend to remain parallel to each other, but at the boundaries between the domains dipoles aligned less favourably with the applied field tend to rotate to match those aligned more favourably with the applied field. The net effect of this is the movement of the domain boundaries; domains parallel to the applied field enlarge whilst others decrease in size. This causes a rise in the bulk magnetization parallel to the applied field (dotted line). Eventually most of the domain dipoles will be aligned with the applied field (point B). The value of \mathbf{M} at this point is called the *saturation value* \mathbf{M}_s . If the applied field is then reduced the domains relax, but only to a point; when the applied field is zero some residual magnetization remains and the ferromagnet is now a permanent

magnet (point C). The value of \mathbf{M} at this point is called the *magnetic remanence* \mathbf{M}_r . If the applied field is decreased past zero the domains will eventually relax to the point where the net magnetization over the sample is zero (point D). However, even though the magnetization is zero, the domain structure is *not* the same as the domain structure of the material at point A. The value of \mathbf{H} at this point is called the *coercivity* \mathbf{H}_c . Reducing the field further still results in saturation in the negative sense (point E). Increasing the applied field from this point causes the material to go through negative remanence (point F) and positive coercivity (point G). The loop is complete when saturation in the positive sense is achieved again (point B).

Magnetic materials are classified as *hard* or *soft* depending on the value of their permeability and coercivity. The *permeability* of a material is a measure of its response to an applied field (defined in section 2.1.6). A material with a large permeability is easy to magnetize and de-magnetize and is called a soft material. These typically have a coercivity \mathbf{H}_c approximately less than 10^3 A/m. Hard materials, which constitute all permanent magnets, have a low permeability and a high coercivity, typically greater than 10^4 A/m. This makes them hard to magnetize and demagnetize, and these materials can be assumed (to a first order approximation) to be unaffected by any external magnetic fields. This property, known as *magnetic transparency*, greatly simplifies the design of magnetic structures incorporating only hard material, as the influence of the magnetization of any component of the design on itself and other hard components within the design need not be taken into account. On the other hand, if soft materials are included in a magnetic design, their influence on themselves and other soft components must be calculated, making the calculation of the field from such a design a much more difficult task.




Figure 1.4 - Schematic hysteresis curves for both hard and soft materials.

Figure 1.4. shows schematic hysteresis curves for typical hard and soft materials². Ideally, a soft material should have a curve as narrow as possible (making it as easy to magnetize and de-magnetize) and a hard material should have a curve as wide as possible (giving a constant M_r as possible over the largest range of applied fields).

- Soft materials are used as magnetic flux conduits, as flux amplifiers to boost flux within a given region and as flux shields to prevent magnetic fields from entering a given region (this is the magnetic analogue of the shielding effect of a conducting surface in electrostatics). They are commonly used in transformers and electromagnets.
- Hard materials are employed as permanent magnetic flux sources in a wide variety of applications such as data storage devices, biomedical equipment, computers and telecommunication equipment. Their main advantages over current-driven flux sources are their constant flux expenditure for no power consumption, their compact size (useful, for example, in data storage equipment) and a lack of need for cooling systems (which are necessary for high power electromagnets). The development of magnetic material since the 1950s has followed a path of increasing energy product $B \cdot H_{max}$ (a measure of how much energy is stored in the material and hence roughly a measure of maximum achievable field), increasing remanence J_r (the intrinsic magnetic polarization of the material, see section 2.1.6) and increasing coercivity (which is a measure of the difficulty with which the material can be magnetized and demagnetized). Some examples of hard magnetic materials are shown in table

1.1 below. At time of writing the material with highest coercivity and remanence is Neodymium-Iron-Boron (NdFeB) with a remanence of around 1.2 Tesla, discovered in the 1980s. In this project only magnetostatic designs assembled from this particular material were investigated.

Date	Name	Composition	BHmax (kJm ⁻³)	J _r (T)
~1930	Alnico*	Al, Ni, Co, Fe (Cu)	~35	~0.75
~1950	Ferrites	XO·6(Fe ₂ O ₃); X=Ba/Sr/Pb	~30	~0.4
~1960	Samarium-Cobalt	Sm, Co	~190	~1.0
~1980	Neodymium-Iron-Boron	Nd, Fe, B	~250	~1.2

Table 1.1 Chronology of Magnetic Materials

**Although the development of Alnico dates from the 1930s and has impressive characteristics, the alloy unfortunately has poor physical qualities (it is very hard and brittle)². By contrast the other materials listed here can be ground into powder form and synthetically bonded, making working the material much easier.*

1.1.3 Designs Exploiting Modern Materials

Cylindrical Designs

A ferromagnetic cylinder with magnetization given by:

$$\mathbf{M} = M_r \{ \sin(k\phi)\hat{\rho} - \cos(k\phi)\hat{\phi} \} \quad [1.1.11]$$

Where k is an integer and M_r the magnetic remanence of the material, gives a cylinder with a magnetic field confined entirely to the bore of the cylinder and zero external field for positive k , or field entirely external to the cylinder and zero field within the bore for negative k ⁹ (see figure 1.5, below, for positive k flux distributions).




Figure 1.5 – One sided flux cylinder arrangements showing positive k magnetization patterns and flux.

For a cylinder of infinite length and $k = 1$ there is a uniform field within the bore given by:

$$\mathbf{H} = M_r \ln\left(\frac{R_o}{R_i}\right) \hat{\mathbf{y}} \quad [1.1.12]$$

where the inner and outer cylinder radii are R_o and R_i , respectively. \mathbf{H} is in the y direction (figure 1.5). This is the simplest form of the *Halbach cylinder*¹⁰, and it can be seen that if the ratio of outer to inner radii is greater than e the flux inside the bore actually exceeds the remanence of the magnetic material used to create the cylinder. Ideally, these structures would be created from a infinite length cylinder of magnetic material with direction of magnetization continuously varying. The magnetic flux produced by this ideal design would be perfectly uniform and be entirely confined to the bore of the cylinder. However the difficulty of manufacturing a cylinder with a continuously varying magnetization usually leads to segmentation of the ideal design. Furthermore, the ideal case of infinite length is not realisable and in practice the finite length of the cylinders produce ‘end effects’ which introduce non-uniformities in the field within the bore.

These cylindrical structures are used in devices such as brushless AC motors, magnetic couplings and high field cylinders. Brushless motors and magnetic coupling devices are not discussed in this thesis; high field structures are, however, covered in detail. Both brushless motors and coupling devices use multipole field arrangements. Brushless motors typically use cylindrical designs in which all the flux is confined to the centre of the bore (such as $k = 3$ above, a six pole rotor) with the AC coils also contained within the bore. Such self-shielding motors designs are more efficient and produce higher torque than conventional motor designs¹¹. Magnetic coupling devices transmit torque through magnetically *transparent* barriers (that is the barrier is non-magnetic *or* is magnetic but is not affected an by applied magnetic field), for instance between sealed containers or pressurised vessels. The optimal torque couplings consists of a pair of coaxially nested cylinders with opposite $+k$ and $-k$ flux magnetization patterns, as $-k$ magnetization patterns produce fields entirely *external* to the cylinder. In the lowest energy state, the outer flux of the inner cylinder exactly matches the internal flux of the outer cylinder. Rotating one cylinder relative to the other from this state results in a restoring torque¹².

If the two dimensional magnetic distribution pattern of the Halbach cylinder is extended to three dimensions, the result is the *Halbach sphere*³. These designs have an extremely uniform field within the interior of the design, as they are not affected by the ‘end effects’ prevalent in the finite length cylinder design. The magnitude of the uniform field for a sphere also increases to 4/3 the amount for a the ideal cylindrical design with the same inner and outer radii. However, being spherical, access to the region of uniform field is usually restricted to a narrow hole at the top and bottom of the design. Higher fields are possible by optimising the spherical design to take account of the fact that it is composed of point dipoles (and not line dipoles). This results in the stretching of the sphere to an elliptical shape and having a non-uniform distribution of magnetization over the component parts of the sphere. Using this method, as well as soft pole pieces within the design, 4.5T in a working volume of 20mm³ was achieved by Bloch *et al*¹³ in 1998 and this was increased further to 5T in 2000, although over a smaller working area of 0.05mm¹⁴. As hard materials are temperature dependent, refrigeration of the entire magnet array can increase the field within the working area further as shown by Kumada *et al*¹⁵. This group also reported development of a 5.16T Halbach *dipole* cylinder in 2003¹⁶.

Other Uniform Field Designs

The cylindrical design is only one class of design which produces a uniform field inside a cavity within an array of permanent magnets³. Other classes of design include *wedge designs*, proposed by Abele and Jensen¹⁷ in which wedges of magnetized material are arranged to provide uniform field within cavities inside the design as shown in figure 1.6 (A).

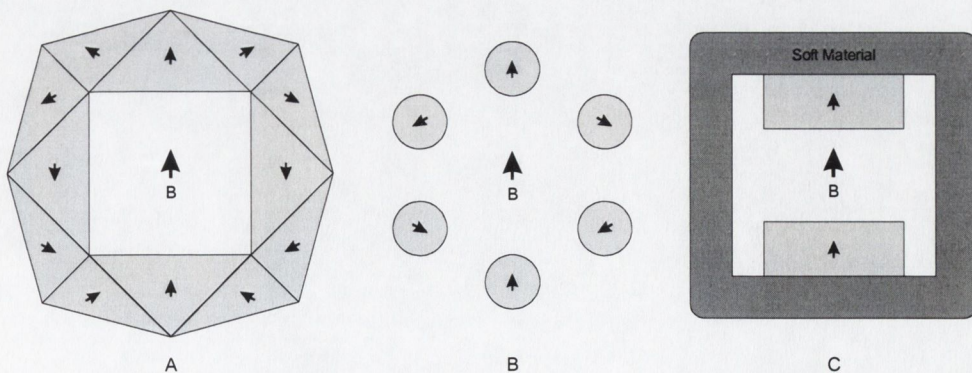


Figure 1.6 – Three designs producing uniform magnetic fields within their central air gap.

The direction of magnetization of the wedges can be calculated using a set of rules given by Abele¹⁸, and allows for great freedom in the shape of the cavity. Another

class of design is the *magnetic mangle*, proposed by Coey and Cugat¹⁹, in which uniformly magnetized rods are arranged such that their magnetization matches that of a Halbach cylinder, as shown in figure 1.6 (B) for a six rod design. This design greatly increases access to the region of uniform field, at the expense of the volume of uniform field being smaller than in the cylindrical designs (although this area can be made larger by increasing the number of component rods). Rotating the rods relative to each other results in many possibilities including a dynamically variable field and various dipolar configurations. Very simple designs for a uniform fields include designs with soft iron return paths, as shown in figure 1.6 (C).

One Sided Flux

Discovered by Mallinson in 1973²⁰, one sided flux structures were initially described as a ‘curiosity’ by Mallinson although he recognised at the time the potential for significant improvements in magnetic tape technology. A one sided flux sheet is a sheet of magnetic material magnetised such that all magnetic flux from the sheet is confined to one side of the sheet. Although this seems somewhat counter-intuitive to those familiar with, for example, bar magnets or solenoids, the reason for this flux distribution can be easily visualised and the theory is straightforward (see section 1.5). Mallinson shows that the criterion for the magnetization of the infinite length sheet is simply that the components of the magnetization in the plane of the flux sheet (x) and normal to the flux sheet (y) are $\pi/2$ out of phase with each other. The mathematical transform which shifts the phase of all components of some function by $\pi/2$ is called a Hilbert transform²¹; the components of the magnetization vector can therefore be any Hilbert transform pair (the simplest of which is $\sin\{x\}$, $\cos\{y\}$, shown in figure 1.7).




Figure 1.7 – Magnetic fields from uniformly magnetized sheet (left) and one-sided flux sheet (right).

The Halbach cylinder (which was found later and independently by Halbach⁹) can be thought of as a ‘rolled up’ finite length flux sheet with sinusoidal magnetization. We have investigated the properties of semi-rolled flux sheets in this project.

The advantages of one sided flux distributions are twofold: First, the field is effectively twice as large on the side on which the flux is confined and secondly, no stray field is produced (in the ideal, infinite length case) on the opposite side. This helps with field confinement, usually a problem in the design of magnetic structures. Although one sided flux distributions may seem somewhat abstract, they have a surprising number of applications ranging from the humble fridge magnet through industrial applications such as the brushless AC motor and magnetic coupling, to high-tech applications such as wiggler magnets used in particle accelerators and free electron lasers. Furthermore, one-sided flux theory could also, surprisingly, explain the lack of magnetic field on celestial bodies such as the moon and Mars²², due to the cooling process after the formation of such planets ‘freezing in’ a spherical internal one sided flux distribution giving rise to zero field external to the planet.

The simplest example of a one sided flux magnet is a fridge magnet. These are usually composed of powdered ferrite in a binder such as plastic or rubber. The extruded magnet is exposed to a rotating field giving the ferrite particles in the magnetic compound a magnetization resulting in a one-sided flux distribution. This distribution increases the holding force of the magnet when placed on a permeable surface, compared to the holding force from, say, a uniform magnetization of the magnetic compound.




Figure 1.8- The magnetization and flux distribution of a fridge magnet.

Scaling up this design and adding a top sheet gives a wiggler magnet, used in synchrotrons²³ and free electron lasers. Wiggler magnets ‘wiggle’ or oscillate an electron beam perpendicular to the magnetic field. As the electrons are undergoing acceleration they radiate electromagnetic energy and these photons can be trapped between two parallel mirrors that form a resonant cavity similar to than of a conventional laser²⁴.




Figure 1.9- A wiggler magnet, used in a free electron laser, oscillates an electron beam creating electromagnetic radiation at a specific wavelength.

The design shown above is usually know as a Halbach wiggler. The magnetization vectors in the magnetized sheets rotate in the opposite senses to each other; above, the top sheet's magnetization vector rotates clockwise and the bottom sheet's magnetization vector rotates counter-clockwise. This design is chosen so that the x -components of the magnetic fields from the sheets cancel and the y -components reinforce so that the field is given by $H_y \propto \cos(kx)$ where k is the 'wavenumber' of the magnetic sheet given by the spacing between magnetic blocks with the same magnetization vector. If the electron velocity approaches the speed of light the radiation wavelength produced can be approximated by²:

$$\lambda_l = \frac{\lambda_w}{2\gamma^2} \quad [1.1.13]$$

where $\lambda_w = 2\pi/k$ and γ is given by:

$$\gamma = \frac{1}{\sqrt{1-\beta^2}} \quad [1.1.14]$$

and β is the fraction of the speed of light at which the electron is travelling, $v = \beta c$.

1.1.4 Designs Producing Constant Force on Dipoles

Magnetic field gradients produce a force on magnetic dipoles, and are used in a variety of applications. Magnetic latching, in which a magnet attaches to another magnet or a soft magnetic material (in which case the magnetic dipoles are induced in the soft material), and magnetic suspension are examples of forces produced by magnetic gradients. However, gradients such as these are often gradients created by magnetic *dipoles* and we looked at designs producing *uniform* gradients in this

project. The force on a permanent dipole with moment \mathbf{m} in a magnetic field \mathbf{B} is given by²⁵:

$$\mathbf{F} = \nabla(\mathbf{m} \cdot \mathbf{B}) \quad [1.1.15]$$

For a dipole with a constant moment, the field gradient then gives a force:

$$\mathbf{F} = m_x \nabla B_x + m_y \nabla B_y + m_z \nabla B_z \quad [1.1.16]$$

One application of a uniform magnetic gradient has been discussed in MRI, although MRI typically uses electromagnets to provide the field gradient. Uniform field gradients are also used for pure research purposes where a controlled magnetic gradient is necessary for a uniform force on permanent magnetized samples under investigation. They are also used in commercial devices such as magnetic separators which, for example, can separate ferromagnetic and paramagnetic materials from ore slurries.

Devices producing a constant field times field gradient, in which the product of the field, \mathbf{B} , and the gradient of the field, dB , is uniform over a given region of space produce a constant force on paramagnetic and diamagnetic materials. This has a number of applications, including Faraday balances and diamagnetic levitation devices. The foam physics group at Trinity college has shown that the bubble size in a ferrofluid foam is proportional to BdB in the region where the foam is created²⁶.

Faraday Balance




Figure 1.10 – A Faraday balance, used to measure magnetic moments; the vertical cylinder diameter is around 2cm.

Picture taken from ²⁷; Dr. J. C. P. Klaasse, Van der Waals-Zeeman Institute

A common device producing a region of uniform field-field gradient product (BdB) is the Faraday balance²⁷, used to measure the magnetic moments of samples. These are

typically based on electromagnetic designs, such as the balance shown above, and we show that the same strength BdB ($\sim 10 \text{ T}^2/\text{m}$) can be created using a novel permanent magnetic design that requires no power input and occupies far less space than the conventional design. The balance shown has the field \mathbf{B} in the x-direction only and the field gradient dB in the y direction only. This means for small samples (where we assume \mathbf{m} is constant throughout the sample volume) in a region of constant BdB , the force on the sample is given by:

$$F_y = m_x \frac{\partial B_x}{\partial y} \quad [1.1.17]$$

Hence the magnetic moment of the sample can be obtained directly from the y component of the force on the sample if we know the gradient. However, if BdB is constant we can measure the susceptibility of diamagnetic samples as, for isotropic samples, the magnetic moment is directly proportional to the applied field. For small sample size:

$$m_x = M_x V \quad [1.1.18]$$

Where M_x is the x-component of the sample magnetization and V is the sample volume. The magnetization vector \mathbf{M} for linear, homogeneous and isotropic materials is given in terms of \mathbf{B} and the magnetic susceptibility χ_m by:

$$\mathbf{M} = \frac{1}{\mu_0} \left(1 + \frac{1}{\chi_m} \right)^{-1} \mathbf{B} \quad [1.1.19]$$

So χ_m can be determined from the force using:

$$\chi_m = \left[\frac{1}{\mu_0} \frac{V}{F_y} \left(B_x \frac{\partial B_x}{\partial y} \right) - 1 \right]^{-1} \quad [1.1.20]$$

Magnetic Levitation




Figure 1.11 - Stable diamagnetic levitation of a water droplet in the bore of a Bitter solenoid.

Picture taken from²⁸; HFML, University of Nijmegen

Earnshaw's theorem²⁹ forbids stable magnetic levitation of a magnet by any arrangement of other magnets but does not forbid *diamagnetic* levitation as the magnetic moment of the diamagnetic sample appears in response to the applied field. Paramagnetic levitation is, however, unobtainable as no energy minimum exists for the paramagnetic case³⁰. The theory behind diamagnetic levitation is similar to the principles behind the Faraday balance; essentially, we seek to make F_y large enough to counteract gravitation forces. The gravitational force is given for an object of density ρ by:

$$F_g = \rho V g \quad [1.1.21]$$

Which must be balanced by the diamagnetic force:

$$F_y = V \frac{1}{\mu_0} \left(1 + \frac{1}{\chi_m} \right)^{-1} \left(B_x \frac{\partial B_x}{\partial y} \right) \quad [1.1.22]$$

So our field gradient product must obey:

$$\left(B_x \frac{\partial B_x}{\partial y} \right) > \mu_0 \rho g \left(1 + \frac{1}{\chi_m} \right) \quad [1.1.23]$$

Which means for $\chi_m \approx 10^{-7}$ and $\rho \approx 10^3 \text{ kg/m}^3$, BdB needs to be around $100 \text{ T}^2/\text{m}$ to achieve levitation; it can be noted that the expression does not involve the mass of the sample, only the density.

Of course, this does not mean stable levitation can necessarily be achieved; the design must ensure that the levitation point is at a point where the gravitational and magnetic energy is a local energy minimum. Unfortunately, the strength of BdB necessary for

diamagnetic levitation is unachievable in designs involving only permanent magnets, but can be achieved in high power solenoid designs¹.

1.1.5 Project Outline

In this project, we have used computer modelling to investigate novel designs for three particular types of application.

- We present designs with improved homogeneity in the bore of Halbach cylinders. Finite length cylinders are subject to ‘end effects’ in which the stray field from the ends of the cylinder reduces the uniformity of the field in the centre of the cylinder. We present an overview of the problem, and compare our new design to two known methods of increasing the homogeneity.
- We investigate designs for producing a uniform field gradient using structures composed entirely from permanent magnets, in which the field gradient is either orthogonal or parallel to the field.
- We investigate designs producing a region in which the product of the field and field gradient is constant (as, for example, used in a Faraday balance). Again, the field gradient can be either orthogonal or parallel to the field. We have also had two “field times field gradient” designs constructed for two different scalings of the design and compared its performance with our computational model.

We developed two computer programs for the modelling of magnetic fields from magnetostatic designs. The first uses a method of determining the field by applying the magnetic charge model to three dimensional polygonal simplexes. This is only suitable for designs composed entirely of modern high-strength permanent magnetic materials, as it uses their property of magnetic transparency. The second uses a finite element method² and allows more flexibility in designs as it can incorporate non-linear ‘soft’ materials. Both of these applications are intended to be used for future work in magnetic design.

2 Theoretical Background

2.1 Magnetic Fields

2.1.1 Maxwell Equations

Magnetic fields are created in some reference frame by electric charges moving relative to the frame. These electrodynamic relations can be summarised using the Maxwell equations, the first of which is:

$$\nabla \cdot \mathbf{B} = 0 \quad [2.1.1]$$

This shows that the *magnetic induction*, \mathbf{B} , does not diverge, i.e. no magnetic monopoles can exist in conventional electromagnetic theory. The second equation gives the *magnetic field* \mathbf{H} , in the case of steady electric current density \mathbf{j}_f and electrostatic fields, as:

$$\nabla \times \mathbf{H} = \mathbf{j}_f \quad [2.1.2]$$

which takes into account the magnetization of any materials present; \mathbf{j}_f is the free electric current within the material. [2.1.2] is the differential form; the integral form can be obtained by applying Stokes' theorem to the above, and shows that the line integral of \mathbf{H} around a loop is dependent only on the free electric current enclosed by the loop. In vacuum \mathbf{B} is simply related to \mathbf{H} by:

$$\mathbf{B} = \mu_0 \mathbf{H} \quad [2.1.3]$$

However, if a material has an intrinsic magnetization, \mathbf{M} , the total magnetic induction, \mathbf{B} , is then due to the magnetic field, \mathbf{H} , plus the magnetization, \mathbf{M} :

$$\mathbf{B} = \mu_0 (\mathbf{H} + \mathbf{M}) \quad [2.1.4]$$

Taking the curl of both sides of [2.1.4] and using [2.1.2] gives:

$$\nabla \times \mathbf{B} = \mu_0 (\mathbf{j}_f + \nabla \times \mathbf{M}) \quad [2.1.5]$$

If we consider the magnetization of the material as being composed of a large number of microscopic current loops, the macroscopic effect of all these current loops can be thought of as giving rise to a fixed surface and volume current within the material. This *bound current* \mathbf{j}_b can be defined as:

$$\nabla \times \mathbf{M} = \mathbf{j}_b \quad [2.1.6]$$

This bound current is simply an imaginary concept; there is no true bulk current flow due to the magnetization at the surface of the material. However, the bound current

can be treated in exactly the same way as a real electric current for calculation purposes; this is the basis of the current model, covered in section 2.2.2.

With this definition of the bound current, the right hand side of [2.1.5] can now be expressed as the summation of two separate ‘current’ parts. This leads to the introduction of the *total current density* \mathbf{j} , composed of the *free electric current* within the material, \mathbf{j}_f , and the *magnetic bound current*, \mathbf{j}_b , arising due the intrinsic magnetization of the material:

$$\mathbf{j} = \mathbf{j}_f + \mathbf{j}_b \quad [2.1.7]$$

The curl of the magnetic induction can then be expressed in terms of this total current density:

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{j} \quad [2.1.8]$$

2.1.2 Magnetic Vector Potential

Since the divergence of $\mathbf{B}(\mathbf{x})$ is zero everywhere it can be expressed as the curl of some vector field. We can therefore introduce the magnetic vector potential $\mathbf{A}(\mathbf{x})$ ³¹:

$$\mathbf{B} = \nabla \times \mathbf{A} \quad [2.1.9]$$

\mathbf{A} is, however, not unique; any function with zero curl added to \mathbf{A} will leave \mathbf{B} unchanged. If we express \mathbf{A} in terms of a second vector field, \mathbf{A}' , and a scalar field, $\Lambda(\mathbf{x})$:

$$\mathbf{A} = \mathbf{A}' + \nabla \Lambda \quad [2.1.10]$$

and take the curl of this, we have:

$$\nabla \times \mathbf{A} = \nabla \times \mathbf{A}' \quad [2.1.11]$$

as the curl of the gradient of an arbitrary scalar field, $\nabla \times [\nabla \Lambda(\mathbf{x})]$, is always zero. This ability to add an arbitrary gradient to \mathbf{A} is known as *gauge invariance*, and the act of adding an arbitrary gradient is known as a *gauge transformation*. Taking the divergence of [2.1.10] gives:

$$\nabla \cdot \mathbf{A} = \nabla \cdot \mathbf{A}' + \nabla^2 \Lambda \quad [2.1.12]$$

So, we can chose the divergence of \mathbf{A} by choosing an appropriate scalar field. If we pick:

$$\nabla^2 \Lambda = -\nabla \cdot \mathbf{A}' \quad [2.1.13]$$

This gives a particularly simple expression for the divergence of \mathbf{A} :

$$\nabla \cdot \mathbf{A} = 0 \quad [2.1.14]$$

Which is called the *Coulomb gauge*. [2.1.13] is simply a Poisson equation for $\Lambda(\mathbf{x})$, so a vector potential satisfying [2.1.14] can be found provided a solution to [2.1.13]

exists. For the magnetostatic situations we are considering, solutions of [2.1.13] exist and we can use the Coulomb gauge. The usefulness of this gauge becomes evident when we substitute [2.1.9] into [2.1.8], giving:

$$\nabla \times (\nabla \times \mathbf{A}) = \mu_0 \mathbf{j} \quad [2.1.15]$$

Using the vector identity for the curl of a curl gives:

$$\nabla(\nabla \cdot \mathbf{A}) - \nabla^2 \mathbf{A} = \mu_0 \mathbf{j} \quad [2.1.16]$$

Substitution of the Coulomb gauge into this expression reduces it to:

$$\nabla^2 \mathbf{A} = -\mu_0 \mathbf{j} \quad [2.1.17]$$

which reduces to the magnetostatic Poisson equation in free space. The Poisson equation can be solved using Green's theorem²⁵

$$\int d^3 x (\phi \nabla^2 \psi - \psi \nabla^2 \phi) = \oint dS \left(\phi \frac{\partial \psi}{\partial \hat{n}} - \psi \frac{\partial \phi}{\partial \hat{n}} \right) \quad [2.1.18]$$

where ϕ , ψ are arbitrary scalar fields and \hat{n} is the surface normal. Taking:

$$\psi = \frac{1}{|\mathbf{x} - \mathbf{x}'|} \quad [2.1.19]$$

and using the property of the Dirac delta function $\delta(\mathbf{x} - \mathbf{x}')$:

$$\nabla^2 \left(\frac{1}{|\mathbf{x} - \mathbf{x}'|} \right) = -4\pi \delta^3(|\mathbf{x} - \mathbf{x}'|) \quad [2.1.20]$$

Gives:

$$A_x = \frac{\mu_0}{4\pi} \int d^3 \mathbf{x}' \frac{j_x(\mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|} + \frac{1}{4\pi} \oint dS' \left[\frac{1}{|\mathbf{x} - \mathbf{x}'|} \frac{\partial A_x}{\partial n'} - A_x \frac{\partial}{\partial n'} \left(\frac{1}{|\mathbf{x} - \mathbf{x}'|} \right) \right] \quad [2.1.21]$$

where we have taken, for example, the x component of \mathbf{A} , $\phi = A_x(\mathbf{x}')$. If the potential tends to zero as the surface becomes infinite (i.e. the source distribution is localised), this reduces to the familiar result:

$$A_x = \frac{\mu_0}{4\pi} \int d^3 \mathbf{x}' \frac{j_x(\mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|} \quad [2.1.22]$$

2.1.3 Magnetic Boundary Conditions

Boundary conditions on \mathbf{B} and \mathbf{H} are important for constructing solutions to the field equations where one or more material interfaces lie within the domain of interest. Applying the divergence theorem to [2.1.1] and Stokes' theorem to [2.1.2] gives the integral form of the magnetostatic Maxwell equations²⁵:

$$\int_S \mathbf{B} \cdot d\mathbf{a} = 0 \quad [2.1.23]$$

$$\oint_c \mathbf{H} \cdot d\mathbf{l} = \mu_0 \int_S \mathbf{j}_f \cdot d\mathbf{a}$$

Applying these to a Gaussian pillbox and an Amperian loop straddling an interface, we have:

$$(\mathbf{B}_2 - \mathbf{B}_1) \cdot \hat{\mathbf{n}} = 0 \quad [2.1.24]$$

$$(\mathbf{H}_2 - \mathbf{H}_1) \times \hat{\mathbf{n}} = \mu_0 \sigma_s$$

Where $\hat{\mathbf{n}}$ is the outward normal to the interface, the subscripts 1 and 2 refer to the inner and outer fields respectively and σ_s is the surface current density. If there is no free conduction current \mathbf{j}_f , these reduce to the simple expressions:

$$\mathbf{H}_1 \times \hat{\mathbf{n}} = \mathbf{H}_2 \times \hat{\mathbf{n}} \quad [2.1.25]$$

$$\mathbf{B}_1 \cdot \hat{\mathbf{n}} = \mathbf{B}_2 \cdot \hat{\mathbf{n}}$$

Showing that the *normal* component of \mathbf{B} is continuous and the *tangential* component of \mathbf{H} is continuous at an interface.

The scalar magnetic potential is introduced in section 1.4.1, but we will review its properties at interfaces within this section. If there is no free current density, [2.1.2] reduces to

$$\nabla \times \mathbf{H} = 0 \quad [2.1.26]$$

Allowing us to represent the magnetic field by a scalar magnetic potential φ_m :

$$\mathbf{H} = -\nabla \varphi_m \quad [2.1.27]$$

The condition on \mathbf{H} leads to the simple condition on φ_m :

$$\hat{\mathbf{n}} \times \nabla \varphi_{m1} = \hat{\mathbf{n}} \times \nabla \varphi_{m2} \quad [2.1.28]$$

At material interfaces this gives:

$$\varphi_{m1} = \varphi_{m2} \quad [2.1.29]$$

The condition on \mathbf{B} leads to a slightly more complex expression, as from [2.1.3]:

$$\mathbf{B} = \mu_0 (-\nabla \varphi_m + \mathbf{M}) \quad [2.1.20]$$

Taking the normal in, for example, the z direction gives:

$$\mathbf{B} \cdot \hat{\mathbf{n}} = \mu_0 \left(M_z - \frac{\partial \varphi}{\partial z} \right) \quad [2.1.31]$$

So that:

$$M_{1z} - \frac{\partial \varphi_{m1}}{\partial z} = M_{2z} - \frac{\partial \varphi_{m2}}{\partial z} \quad [2.1.32]$$

2.1.4 Magnetic Dipoles

The term ‘dipole’ is something of a misnomer when applied to magnetism as [2.1.1] shows that no isolated magnetic monopoles are allowed to exist in electromagnetic theory. Also, a magnetic dipole cannot be thought about in the same terms as an electric dipole, i.e. two isolated point charges of opposite sign separated by some small distance. Nevertheless, if we take a multipole expansion of an arbitrary current distribution in [2.1.22] the vector potential of the *dipole* term in the expansion is given by:

$$\mathbf{A}_{\text{dipole}} = \frac{\mu_0}{4\pi} \frac{\mathbf{m} \times \mathbf{x}}{|\mathbf{x}|^3} \quad [2.1.33]$$

Where \mathbf{x} is the vector from the dipole to the observation point and \mathbf{m} , the dipole moment, is given by:

$$\mathbf{m} = \frac{1}{2} \int d^3 \mathbf{x}' [\mathbf{x}' \times \mathbf{j}(\mathbf{x}')] \quad [2.1.34]$$

where $\mathbf{j}(\mathbf{x}')$ is the current density. If the current distribution flows in a loop with line element $d\mathbf{l}'$ and current I , this becomes:

$$\mathbf{m} = \frac{I}{2} \oint \mathbf{x}' \times d\mathbf{l}' \quad [2.1.35]$$

For a plane loop [2.1.35] reduces to the simple form:

$$\mathbf{m} = I\mathbf{\Omega} \quad [2.1.36]$$

Where $\mathbf{\Omega}$ is the vector area of the plane loop ($\mathbf{\Omega} = A \hat{\mathbf{n}}$ where $\hat{\mathbf{n}}$ is the normal vector to the plane of the loop). From [2.1.33] the \mathbf{B} field for a dipole at the origin aligned in the z -direction is given (in spherical co-ordinates) by:

$$\mathbf{B} = \frac{\mu_0}{4\pi} \frac{m}{r^3} \left[(2 \cos \theta) \hat{\mathbf{r}} + (\sin \theta) \hat{\boldsymbol{\theta}} \right] \quad [2.1.37]$$

Which is identical in form to the \mathbf{E} field of an electric dipole. Microscopically a magnetic dipole will look very different from an electrical dipole, as a magnetic dipole is created by a current loop whereas an electric dipole is created by separated charges. However, it is possible to consider magnetic dipoles and, in fact, macroscopic magnetic objects in terms of magnetic charges. This will be discussed in a section 2.2.1 on the magnetic charge model.

2.1.5 Forces and Torques on Magnetic Dipoles

For steady electric currents the lowest order contribution to the magnetic force on an infinitesimal magnetic dipole with permanent moment \mathbf{m} in an external field \mathbf{B} is given by:

$$\mathbf{F} = \nabla(\mathbf{m} \cdot \mathbf{B}) \quad [2.1.38]$$

The above holds for time-varying fields. We can expand [2.1.38] by using the identity:

$$\nabla(\mathbf{m} \cdot \mathbf{B}) = \mathbf{m} \times (\nabla \times \mathbf{B}) + \mathbf{B} \times (\nabla \times \mathbf{m}) + (\mathbf{m} \cdot \nabla)\mathbf{B} + (\mathbf{B} \cdot \nabla)\mathbf{m} \quad [2.1.39]$$

If there is no free or bound current at the location of the dipole the curl of \mathbf{B} is zero²⁵ and as the magnetic moment \mathbf{m} is constant the force is then given by:

$$\mathbf{F} = (\mathbf{m} \cdot \nabla)\mathbf{B} \quad [2.1.40]$$

This shows that the force on a dipole depends on *spatially varying* \mathbf{B} . The torque, \mathbf{N} , on an infinitesimal dipole with moment \mathbf{m} is given by:

$$\mathbf{N} = \mathbf{m} \times \mathbf{B} \quad [2.1.41]$$

such that dipoles attempt to align parallel to an external applied magnetic field.

2.1.6 Magnetisation

Magnetization Vector \mathbf{M}

The net magnetic dipole moment per unit volume, \mathbf{M} , is given by:

$$\mathbf{M} = \lim_{\Delta V \rightarrow 0} \frac{\sum \mathbf{m}_i}{\Delta V} \quad [2.1.42]$$

where $\sum \mathbf{m}_i$ is the vector sum of the dipole moments in the infinitesimal volume ΔV .

The magnetic field, \mathbf{H} , is defined by:

$$\mathbf{H} \equiv \frac{1}{\mu_0} \mathbf{B} - \mathbf{M} \quad [2.1.43]$$

We adopt the Kennelly convection, which defines the *magnetic polarisation* \mathbf{J} as:

$$\mathbf{B} = \mu_0 \mathbf{H} + \mathbf{J} \quad [2.1.44]$$

where:

$$\mathbf{J} = \mu_0 \mathbf{M} \quad [2.1.45]$$

Linear Approximations to \mathbf{M}

One way of relating \mathbf{M} and \mathbf{H} is using a simple linear approximation; in linear, homogeneous and isotropic media the magnetisation \mathbf{M} can be assumed to be proportional to the applied field \mathbf{H} :

$$\mathbf{M} = \chi_m \mathbf{H} \quad [2.1.46]$$

Where χ_m is called the *magnetic susceptibility*. Substituting [2.1.46] into [2.1.4] we find that for linear, homogeneous and isotropic materials \mathbf{B} is given by:

$$\mathbf{B} = \mu_0 \mu_r \mathbf{H} \quad [2.1.47]$$

Where μ_r , the *relative permeability*, is given by:

$$\mu_r = 1 + \chi_m \quad [2.1.48]$$

Equations [2.1.46] and [2.1.47] must be modified for nonlinear materials. Equation [2.1.47] becomes:

$$\mathbf{B} = \mu_0 \mu_r(\mathbf{H}) \mathbf{H} \quad [2.1.49]$$

Where the relative permeability depends on \mathbf{H} . The linear approximation is useful in materials which have no intrinsic magnetization, such as diamagnetic materials, or to approximate the initial magnetization curve of ferromagnetic materials (point A to B in figure 1.3), and is the simplest to deal with.




Figure 2.1 – A linear approximation to a B-H hysteresis curve.

Another way of relating \mathbf{M} and \mathbf{H} is by splitting the magnetization term into two parts:

$$\mathbf{M} = \mathbf{M}_i + \mathbf{M}_r \quad [2.1.50]$$

Where \mathbf{M}_i can be thought of as the *induced* magnetization in the material by the applied field \mathbf{H} and \mathbf{M}_r can be thought of as the intrinsic (*remnant*) magnetization of the material (which is given by a constant magnetization vector). This is effectively using a linear approximation for the demagnetization curve of a ferromagnetic material (points A, B and C in figure 1.3). The induced magnetization is related to \mathbf{H} by the simple linear relationship above:

$$\mathbf{M}_i = \chi_m \mathbf{H} \quad [2.1.51]$$

Which gives:

$$\mathbf{B} = \mu_0 (\mu_r \mathbf{H} + \mathbf{M}_r) \quad [2.1.52]$$




Figure 2.2 – An approximation to a B-H hysteresis curve for hard materials.

This approximation is of particular use for hard magnetic materials. In the Kennelly convection we define the induced polarization \mathbf{J}_i and the remnant polarization \mathbf{J}_r as:

$$\begin{aligned} \mathbf{J}_i &= \mu_0 \mathbf{M}_i \\ \mathbf{J}_r &= \mu_0 \mathbf{M}_r \end{aligned} \quad [2.1.53]$$

2.2 Magnetostatic Field Calculations

2.2.1 Charge Model

The charge model is especially useful for permanent magnets; in general it allows us to split the magnetization within a material into equivalent magnetic volume and surface charges². If the regions under consideration are electric current-free, [2.1.2] becomes:

$$\nabla \times \mathbf{H} = 0 \quad [2.2.1]$$

We can then use the fact that any irrotational vector field can be written as the gradient of some function (since an identity of vector calculus gives the curl of the gradient of any function as zero), i.e.:

$$\mathbf{H} = -\nabla \varphi_m \quad [2.2.2]$$

Where φ_m is the magnetostatic scalar potential. This procedure corresponds to the formulation of potential for the electrostatic case (the scalar potential here is give the subscript m to differentiate it from the electrostatic case). From [2.1.4]:

$$\mathbf{B} = \mu_0 (\mathbf{H} + \mathbf{M}) \quad [2.2.3]$$

Taking the divergence of both sides and using [2.2.2], gives:

$$\nabla \cdot \mathbf{B} = \mu_0 (-\nabla^2 \varphi_m + \nabla \cdot \mathbf{M}) \quad [2.2.4]$$

\mathbf{B} is divergenceless [2.1.1], so this reduces to:

$$\nabla^2 \varphi_m = \nabla \cdot \mathbf{M} \quad [2.2.5]$$

Which is a magnetostatic Poisson equation. If the boundary surfaces are at infinity, we can solve [2.2.5] using the Green's function [2.1.22], which gives:

$$\varphi_m(\mathbf{x}) = -\frac{1}{4\pi} \int_{\tau} \frac{\nabla' \cdot \mathbf{M}(\mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|} d\tau' \quad [2.2.6]$$

Where τ is the source volume. If the magnetization is localised within some volume V and *abruptly* falls to zero outside the volume, we can modify [2.2.6] to include a surface potential caused by the discontinuity at the interface:

$$\varphi_m(\mathbf{x}) = \frac{1}{4\pi} \int_V \frac{\rho_m'}{|\mathbf{x} - \mathbf{x}'|} dV' + \frac{1}{4\pi} \oint_S \frac{\sigma_m'}{|\mathbf{x} - \mathbf{x}'|} dS' \quad [2.2.7]$$

Where the magnetic volume charge density ρ_m is given by:

$$\rho_m = -\nabla \cdot \mathbf{M} \quad [2.2.8]$$

And the magnetic surface charge density is given by:

$$\sigma_m = \mathbf{M} \cdot \hat{\mathbf{n}} \quad [2.2.9]$$

Where $\hat{\mathbf{n}}$ is the unit outward normal from the surface. If the magnetization is uniform the magnetic volume charge density, ρ_m , is zero causing the first term to disappear; only the surface term contributes, which is particularly useful for calculating the magnetic field of structures containing only hard permanent magnets.

It can be seen that in terms of the surface charge, the normal component of \mathbf{H} is discontinuous by an amount proportional to the surface charge:

$$H_{>} - H_{<} = \sigma_m \quad [2.2.10]$$

Where $H_{>}$ and $H_{<}$ represent the radial components of \mathbf{H} above and below the surface, respectively.

Approximation for Hard Materials

Typical modern hard magnetic materials have very low values of magnetic susceptibility, often around $\chi_m \approx 10^{-2}$. If we then make the approximation that $\chi_m = 0$, $\mu_r = 1$ from [2.1.48] this gives [2.1.52] as:

$$\mathbf{B} = \mu_0 (\mathbf{H} + \mathbf{M}_r) \quad [2.2.11]$$

Physically, a low susceptibility value for a material means that the intrinsic magnetization vector of a material is very resistant to change by applied fields. By setting the susceptibility to zero we are assuming the material is unaffected by *any* applied field sources, and is therefore completely magnetically transparent. If the

material has a fixed, uniform, magnetization the charge volume density [2.2.8] is zero and the surface charge density [2.2.9] becomes:

$$\sigma_m = \mathbf{M}_r \cdot \hat{\mathbf{n}} \quad [2.2.12]$$

and the magnetic potential can be expressed in terms of *only* the constant remnant magnetization vector for the material; from [2.2.7]:

$$\varphi_m(x) = \frac{1}{4\pi} \oint_S \frac{\sigma_m}{|\mathbf{x} - \mathbf{x}'|} dS' \quad [2.2.13]$$

2.2.2 Current Model

The current model makes use of the fact that magnetization within a material can be reduced to equivalent surface and volume current densities, in a similar way to the charge model reducing the magnetization to equivalent surface and volume charge distributions. Using [2.1.9]:

$$\mathbf{B} = \nabla \times \mathbf{A} \quad [2.1.9]$$

and from [2.1.4]:

$$\mathbf{B} = \mu_0(\mathbf{H} + \mathbf{M}) \quad [2.2.14]$$

gives:

$$\nabla \times \mathbf{A} = \mu_0(\mathbf{H} + \mathbf{M}) \quad [2.2.15]$$

Taking the curl of both sides gives:

$$\nabla(\nabla \cdot \mathbf{A}) - \nabla^2 \mathbf{A} = \mu_0(\mathbf{j}_f + \nabla \times \mathbf{M}) \quad [2.2.16]$$

Using the Coulomb gauge [2.1.14] and the definition for the bound current [2.1.6] gives:

$$\nabla^2 \mathbf{A} = -\mu_0(\mathbf{j}_f + \mathbf{j}_b) \quad [2.2.17]$$

This is a second magnetic Poisson equation (compared to [2.2.5]) involving bound and free currents. If there is no free current and we set our domain boundaries at infinity, we can express [2.2.17] in integral form using the free-space Green's function:

$$\mathbf{A}(\mathbf{x}) = \frac{\mu_0}{4\pi} \int_V \frac{\mathbf{j}_b(\mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|} dV' \quad [2.2.18]$$

If the magnetization is confined to some volume V and falls abruptly to zero outside V , [2.2.18] can be expressed as²⁵:

$$\mathbf{A}(\mathbf{x}) = \frac{\mu_0}{4\pi} \int_V \frac{\mathbf{j}_b(\mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|} dV' + \frac{\mu_0}{4\pi} \oint_S \frac{\mathbf{j}_s(\mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|} dS' \quad [2.2.19]$$

With volume current density \mathbf{j}_b :

$$\mathbf{j}_b = \nabla \times \mathbf{M} \quad [2.2.20]$$

And surface current density \mathbf{j}_s :

$$\mathbf{j}_s = \mathbf{M} \times \hat{\mathbf{n}} \quad [2.2.21]$$

Where $\hat{\mathbf{n}}$ is the unit outward normal from the surface. As with the magnetic charge model, if the magnetization is fixed and uniform only the surface second term in [2.1.19] contributes as the volume current density is zero.

2.2.3 Green Function Expansion

The problem can sometimes be re-cast into an easier form by using expansions of the Green's function, for example in cylindrical magnetic designs. It is often convenient to use spherical harmonic expansions for the Green's function; in spherical polar coordinates, with boundary surfaces at infinity, the expansion of the Green's function is²⁵:

$$\frac{1}{|\mathbf{x} - \mathbf{x}'|} = 4\pi \sum_{l=0}^{\infty} \sum_{m=-l}^l \frac{1}{2l+1} \frac{r_{<}^l}{r_{>}^{l+1}} Y_{lm}^*(\theta', \phi') Y_{lm}(\theta, \phi) \quad [2.2.22]$$

Where $r_{<}$ is the smaller value of r and r' and $r_{>}$ is the greater value of r and r' . $Y_{lm}(\theta, \phi)$ are *spherical harmonics*. The expression [2.2.22] is useful as the Greens function can be expressed as the product of spherical harmonics separately associated with both the point and the source (at the price of a double sum). Any arbitrary (spherical) function can be expanded in terms of spherical harmonics as they form a complete set of orthogonal functions over the unit sphere:

$$f(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l g_{lm} Y_{lm}(\theta, \phi) \quad [2.2.23]$$

Where g_{lm} are expansion coefficients given by:

$$g_{lm} = \int_S Y_{lm}^*(\theta, \phi) f(\theta, \phi) dS \quad [2.2.24]$$

[2.2.22] is of particular use in the magnetic charge model; [2.2.7] with uniform magnetization (that is, with no volume charge density) gives:

$$\varphi(r, \theta, \phi) = \frac{1}{4\pi} \int_S \frac{\sigma(r', \theta', \phi')}{|\mathbf{x} - \mathbf{x}'|} dS' \quad [2.2.25]$$

Where σ is the magnetic surface charge density and φ the magnetostatic potential (we have dropped the 'm' subscripts to avoid confusion with the sum of m later in the expansion). This can now be expressed as:

$$\varphi(r, \theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l \frac{1}{2l+1} Y_{lm}(\theta, \phi) \int_{r'_>}^{r'_<} \sigma(r', \theta', \phi') Y_{lm}^*(\theta', \phi') d^3 x' \quad [2.2.26]$$

Which may be easier to evaluate in the case of certain geometries. If we consider the potential for points where the distance to the observation point r is less than the distance to the source point r' , this can be expressed as:

$$\varphi_{<}(r, \theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l q_{lm} r^l Y_{lm}(\theta, \phi) \quad [2.2.27]$$

And for points where the observation point r is greater than the source point r' :

$$\varphi_{>}(r, \theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l p_{lm} \frac{1}{r^{l+1}} Y_{lm}(\theta, \phi) \quad [2.2.28]$$

Where we have defined the expansion coefficients q_{lm} and p_{lm} as:

$$q_{lm} = \frac{1}{2l+1} \int_S \frac{1}{r'^{l+1}} \sigma(r', \theta', \phi') Y_{lm}^*(\theta', \phi') dS' \quad [2.2.29]$$

$$p_{lm} = \frac{1}{2l+1} \int_S r'^l \sigma(r', \theta', \phi') Y_{lm}^*(\theta', \phi') dS' \quad [2.2.30]$$

It can be noted that this expression for the expansion coefficients is expressed entirely in terms of source coordinates. The potential is continuous at $r = r'$, so by comparing [2.2.27] and [2.2.28] we can see that the expansion coefficients must obey:

$$p_{lm} = r^{2l+1} q_{lm} \quad [2.2.31]$$

A further condition is given by the discontinuity of \mathbf{H} across a boundary (see section 2.1.3); the normal component of \mathbf{H} across any boundary is discontinuous by an amount proportional to the surface charge on the boundary. In terms of the potential this gives (using [2.2.10]):

$$\nabla \varphi_{<}(r, \theta, \phi) - \nabla \varphi_{>}(r, \theta, \phi) = \sigma(r, \theta, \phi) \quad [2.2.32]$$

If the boundary is a spherical surface at a radius r_p , this becomes:

$$\frac{\partial \varphi_{<}(r, \theta, \phi)}{\partial r} - \frac{\partial \varphi_{>}(r, \theta, \phi)}{\partial r} = \sigma(r, \theta, \phi), r = r_p \quad [2.2.33]$$

as the normal is in the radial outward direction. Using this with [2.2.27] and [2.2.28] gives:

$$\sum_{l=0}^{\infty} \sum_{m=-l}^l \left[q_{lm} l r_p^l + p_{lm} \frac{(l+1)}{r_p^{l+2}} \right] Y_{lm}(\theta, \phi) = \sigma(r_p, \theta, \phi) \quad [2.2.34]$$

Using [2.2.31] then gives us:

$$\sum_{l=0}^{\infty} \sum_{m=-l}^l q_{lm} (2l+1) r_p^{l-1} Y_{lm}(\theta, \phi) = \sigma(r_p, \theta, \phi) \quad [2.2.35]$$

As our surface charge is defined over a spherical surface with $r = r_p$, we can expand it using [2.2.23]:

$$\sigma(r_p, \theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l g_{lm} Y_{lm}(\theta, \phi) \quad [2.2.36]$$

We can compare this directly to [2.2.27] to ascertain the expansion coefficients g_{lm} (there is no need to use [2.2.24]); the comparison shows that:

$$q_{lm} = \frac{1}{2l+1} \frac{1}{r_p^{l-1}} g_{lm} \quad [2.2.37]$$

Of course, the distribution of surface charge will rarely be spherical, but we can use Dirac delta functions to define a point charge on a spherical surface²⁵:

$$\sigma(r_p, \theta', \phi') = \frac{Q}{r_p^2} \delta(\phi' - \phi_p) \delta(\cos(\theta') - \cos(\theta_p)) \quad [2.2.38]$$

The delta functions simply represent a point charge Q at the spherical coordinate angles θ' and ϕ' at a radius r_p ; this can be thought of as a charge distribution over a sphere of radius r_p with zero charge everywhere except in an infinitely small region at the angles θ' and ϕ' . The properties of the delta function simplify our expression for the expansion coefficient q_{lm} to:

$$q_{lm} = \frac{Q}{2l+1} \frac{Y_{lm}^*(\theta_p, \phi_p)}{r_p^{l+1}} \quad [2.2.39]$$

With a similar expansion for p_{lm} . This gives the coefficients in the expansion of the potential g_{lm} as:

$$g_{lm} = \frac{Q}{r_p^2} Y_{lm}^*(\theta_p, \phi_p) \quad [2.2.40]$$

We can now use this by treating each point on a surface carrying magnetic surface charge as a point charge Q , which contributes $d\varphi$ to the potential. Integrating over all surfaces and summing these point charges will then give us the total potential φ , from which we can find the magnetic field \mathbf{H} .

2.3 Magnetic Structures

2.3.1 One-Sided Flux

In 1973 Mallinson discovered a 'magnetic curiosity' in which all the flux from a planar magnetic structure magnetized according to a simple rule would be confined to one side of the structure²⁰. This can be intuitively visualised using Mallinson's

original diagram (figure 2.3; this uses the *negative* y-component unlike the diagram in Mallinson's paper, given below in [2.3.2], with $x = 0$ at the left edge and $x = 2\pi$ at the right edge):




Figure 2.3 The x component (top left), y component (top right) and resulting total magnetization of Mallinson's flux sheet (bottom). The flux 'cancels' below the sheet and 'reinforces' above.

The key point is that the flux will cancel below the plane and reinforce itself above the plane. Mallinson shows that *any* magnetization pattern where the components of magnetization are $\pi/2$ out of phase with each other will result in a one-sided flux. If the x-component of the magnetization is a real function of x it can be broken into Fourier components; shifting each of these components by $\pi/2$ will give the y component of the magnetization required for a one-sided flux. The mathematical transform that changes the phase of all components is the Hilbert transform²¹:

$$\hat{f}(x) = \frac{1}{\pi} \int_{-\infty}^{\infty} \frac{f(x')}{(x-x')} dx' \quad [2.3.1]$$

A magnetization pattern based on any Hilbert transform pair (subject to physical constraints) will therefore give a one-sided flux.

The field from a flux sheet can be determined by considering the magnetic scalar potential above, inside and below an idealised infinite flux sheet. The magnetization, \mathbf{M} , of the sheet is given by:

$$\mathbf{M} = M_r (\sin(kx)\hat{\mathbf{i}} - \cos(kx)\hat{\mathbf{j}}) \quad [2.3.2]$$

where k is a chosen wavenumber. Using [2.2.5] the potentials satisfy:

$$\begin{aligned} \nabla^2 \varphi_{above} &= 0 \\ \nabla^2 \varphi_{inside} &= M_r k \cos(kx) \\ \nabla^2 \varphi_{below} &= 0 \end{aligned} \quad [2.3.3]$$

That is, a Laplace equation above and below the sheet and a Poisson equation inside the sheet. A solution to the Poisson equation inside the sheet is given by:

$$\varphi_{inside} = -M_r \frac{\cos(kx)}{k} \quad [2.3.4]$$

For the Laplace equations above and below the sheet we choose general solutions of the form:

$$\varphi = (K_1 e^{-ky} + K_2 e^{ky}) \cos(kx) \quad [2.3.5]$$

The general solutions are:

$$\begin{aligned} \varphi_{above} &= A e^{-ky} \cos(kx) \\ \varphi_{inside} &= \left(B e^{-ky} + C e^{ky} - \frac{M_r}{k} \right) \cos(kx) \\ \varphi_{below} &= D e^{ky} \cos(kx) \end{aligned} \quad [2.3.6]$$

where we have used the requirement that the potential must approach zero as $y \rightarrow \pm\infty$. Using the conditions for the potential between interfaces given in section 2.1.3 and taking the magnetization in the y -direction given by [2.3.2] above, we have:

$$A = \frac{M_r}{k} (e^{-kd} - 1), \quad B = \frac{M_r}{k} e^{-kd}, \quad C = 0, \quad D = 0 \quad [2.3.7]$$

Giving the potentials as:

$$\begin{aligned} \varphi_{above} &= \frac{M_r}{k} (e^{-kd} - 1) e^{-ky} \cos(kx) \\ \varphi_{inside} &= \frac{M_r}{k} (e^{-k(d+y)} - 1) \cos(kx) \\ \varphi_{below} &= 0 \end{aligned} \quad [2.3.8]$$

and the fields (using [2.2.2]) as:

$$\begin{aligned} \mathbf{H}_{above} &= M_r (e^{kd} - 1) e^{-ky} \{ \sin(kx) \hat{\mathbf{i}} + \cos(kx) \hat{\mathbf{j}} \} \\ \mathbf{H}_{inside} &= M_r \{ (e^{-k(d+y)} - 1) \sin(kx) \hat{\mathbf{i}} + e^{-k(d+y)} \cos(kx) \hat{\mathbf{j}} \} \\ \mathbf{H}_{below} &= 0 \end{aligned} \quad [2.3.9]$$

Showing that for a flux sheet with magnetization given by [2.3.2], the field is confined to the upper side of the sheet only.

Several Hilbert pairs are shown below for comparison, $x = 0$ is taken to be the centre of the block, and the length of the block ranges from $-\pi$ to $+\pi$:

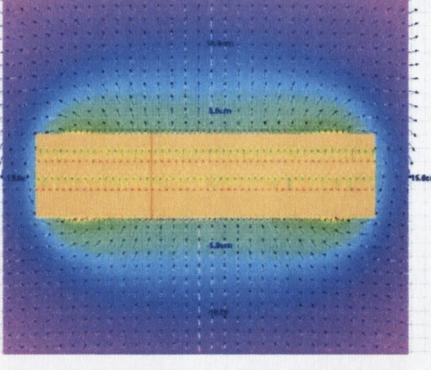
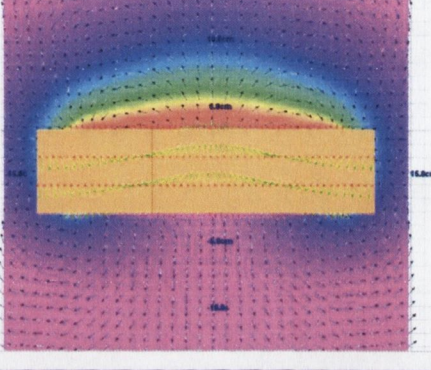
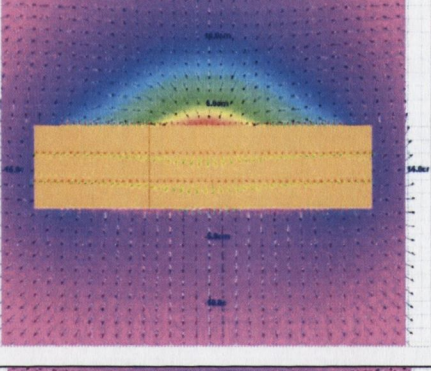
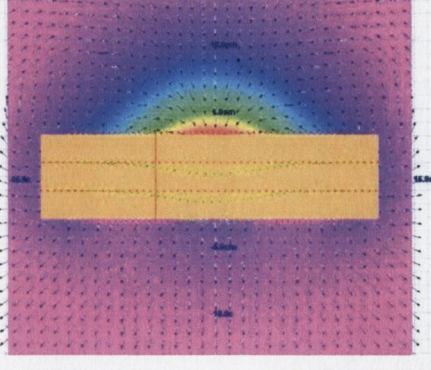
Hilbert Pair	Magnetization	Notes
$J_x = 0$ $J_y = 1$		<p><i>This is a uniformly magnetized block included for comparison purposes</i></p>
$J_x = \sin(x)$ $J_y = -\cos(x)$		<p><i>This is the magnetization pattern given by [2.3.2]</i></p>
$J_x = \frac{x}{1+x^2}$ $J_y = -\frac{1}{1+x^2}$		<p><i>This is a transform given by Mallinson's original paper. The magnetization $\rightarrow 0$ at the edges of the block</i></p>
$J_x = -x {}_1F_1\left\{1, \frac{3}{2}, -x^2\right\}$ $J_y = e^{-x^2}$		<p><i>${}_1F_1$ is a hypergeometric function of the first kind.</i></p>

Table 2.1 – A comparison of various Hilbert pair one-sided flux magnetization patterns.

Other Hilbert transform pairs exist, but many would correspond to non-physical distributions of magnetization, e.g. the Dirac delta function.

2.3.2 Infinite Length Halbach Cylinder

We derive a Green function expansion for a infinite (and later finite) length Halbach cylinder. A perfect Halbach cylinder is an infinite length cylinder (see figure 2.4 for the finite length case) with magnetization given in cylindrical co-ordinates by:

$$\mathbf{M}(\phi') = M_r \{ \sin(\phi') \hat{\rho} - \cos(\phi') \hat{\phi} \} \quad [2.3.10]$$




Figure 2.4 – A ferromagnetic cylinder of length L magnetized with a flux distribution given by [2.3.10] with $k = 1$ will give a magnetized cylinder known as a Halbach cylinder.

The magnetic volume charge density ρ_m is given by [2.2.8] (see section 2.2.1):

$$\rho_m = -\nabla \cdot \mathbf{M} \quad [2.3.11]$$

Which gives the magnetic volume charge density of the cylinder as:

$$\rho_m(\rho', \phi') = -2M_r \frac{\sin \phi'}{\rho'} \quad [2.3.12]$$

The primed co-ordinates represent source points. The potential for points (in *spherical* co-ordinates) where the observation point r is less than the source point r' is given by equation [2.2.27]:

$$\varphi_<(r, \theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l q_{lm} r^l Y_{lm}(\theta, \phi) \quad [2.3.13]$$

With coefficient q_{lm} given by equation [2.2.39]:

$$q_{lm} = \frac{Q}{2l+1} \frac{Y_{lm}^*(\theta_P, \phi_P)}{r_P^{l+1}} \quad [2.3.14]$$

Where P is the location of the magnetic charge Q . Here, Q is a infinitesimal volume charge ρ_m in cylindrical volume element dV at point ρ' , θ' . The potential for a ring of infinitesimal thickness $d\rho'$ at radius ρ' is then given by:

$$q_{lm} = \int_{-L}^L \int_0^{2\pi} \left\{ \left(-2M_r \frac{\sin \theta'}{r'} \right) \frac{1}{2l+1} \frac{Y_{lm}^*(\theta', \phi')}{r'^{l+1}} \right\} \rho' d\phi' dz' d\rho' \quad [2.3.15]$$

Where the cylinder extent is from $-L$ to L (we will let $L \rightarrow \infty$ at a later stage). To simplify this expression we use the relationship between spherical harmonics and the associated Legendre functions²⁵:

$$Y_{lm}(\theta, \phi) = (-1)^m \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} P_l^m(\cos\theta) e^{im\phi}, \quad m > 0 \quad [2.3.16]$$

The expression for spherical harmonics with negative m can be given by:

$$Y_{l-m}(\theta, \phi) = (-1)^m Y_{lm}^*(\theta, \phi), \quad m > 0 \quad [2.3.17]$$

We can then use [2.3.15] with the orthogonality relationship:

$$\int_0^{2\pi} \sin(\phi') e^{-im\phi} d\phi = \frac{\pi}{i} (\delta_{m,1} - \delta_{m,-1}) \quad [2.3.18]$$

where δ_{ij} is the Kronecker delta, to express the coefficient q_{lm} as:

$$q_{lm} = \int_{-L}^L \left\{ (-1)^m - \frac{2\pi}{i} \frac{M_r}{2l+1} \frac{1}{r'^{l+1}} \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} P_l^m(\cos\theta') (\delta_{m,1} - \delta_{m,-1}) \right\} dz' d\rho' \quad [2.3.19]$$

The expression for the potential then becomes:

$$\varphi_{<}(r, \theta, \phi) = \frac{2\pi M_r}{i} d\rho' \sum_{l=1}^{\infty} \left\{ \frac{1}{2l+1} \frac{r^l}{r'^{l+1}} \sqrt{\frac{2l+1}{4\pi} \frac{(l-1)!}{(l+1)!}} \int_{-L}^L [P_l^1(\cos\theta') Y_{l,1}(\theta, \phi) - P_l^{-1}(\cos\theta') Y_{l,-1}(\theta, \phi)] dz' \right\} \quad [2.3.20]$$

as all terms for m not equal to 1 or -1 drop out due to the Kronecker deltas (note that this also means l must be summed from 1 to ∞ as $l=0$ implies $m=0$). Using the relationship between spherical harmonics and associated Legendre functions [2.3.16], and the relationship²⁵:

$$P_l^{-m}(\cos\theta) = (-1)^m \frac{(l-m)!}{(l+m)!} P_l^m(\cos\theta) \quad [2.3.21]$$

[2.3.20] reduces to:

$$\varphi_{<}(r, \theta, \phi) = -M_r d\rho' \sum_{l=1}^{\infty} r^l \frac{(l-1)!}{(l+1)!} P_l^1(\cos\theta) \sin\phi \int_{-L}^L \frac{1}{r'^{l+1}} P_l^1(\cos\theta') dz' \quad [2.3.22]$$

This expression, however, uses both z and ρ from the cylindrical co-ordinate system and r , ϕ and θ from the spherical co-ordinate system. We can convert cylindrical co-ordinates to spherical with the following relationships:

$$\begin{aligned} \rho' &= r' \sin(\theta') \\ z' &= r' \cos(\theta') \end{aligned} \quad [2.3.23]$$

For an infinitesimally thin cylinder ρ' is constant, giving:

$$z' = \rho' \cot(\theta')$$

$$dz' = \frac{\rho'}{\sin^2(\theta')} d\theta' \quad [2.3.24]$$

If we only consider points inside the bore of the cylinder where the observation point r is *always* less than ρ' (i.e. the observation point is within a sphere, centred at the origin, with radius less than the inner bore radius) we can express r' as:

$$r' = \frac{\rho'}{\sin(\theta')} \quad [2.3.25]$$

and we need only consider our expression for the potential $\varphi_{<}$ as for all source points, r' , the observation point $r < r'$. This gives our expression for the potential as:

$$\varphi_{<}(r, \theta, \phi) = -M_r \sum_{l=1}^{\infty} \frac{(l-1)!}{(l+1)!} P_l^1(\cos \theta) \sin \phi \frac{r^l}{\rho'^l} d\rho' \int_{\theta_L}^{\pi-\theta_L} \sin^{l-1} \theta' P_l^1(\cos \theta') d\theta' \quad [2.3.26]$$

Where we have defined the limit θ_L as the value of θ when $z = L$ (as the cylinder is symmetric, the limit at $z = -L$ is simply $\pi - \theta_L$).

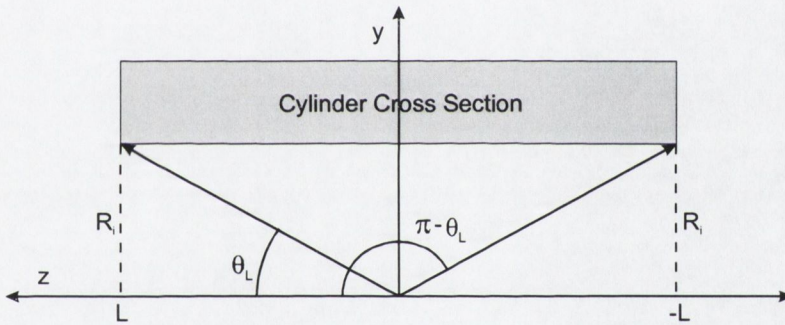


Figure 2.5 – The angle θ_L from expression [2.3.26] used in the spherical harmonic Halbach cylinder model.

The integral can be evaluated by noticing that for Legendre functions with $m = l$:

$$P_l^l(\cos \theta') = C_l \sin^l \theta', \quad C_l = (-1)^l (2l-1)!! \quad [2.3.27]$$

Where C_l is constant for some particular l (the double factorial is intentional). The integral can then be expressed as:

$$I_l(\theta_L) = \frac{1}{C_l} \int_{\theta_L}^{\pi-\theta_L} \frac{P_l^l(\cos \theta') P_l^1(\cos \theta')}{\sin^2 \theta'} [\sin \theta' d\theta'] \quad [2.3.28]$$

Where we have split up the $\sin^{l-1}(\theta')$. This can be expressed in terms of $\cos(\theta)$ as:

$$I_l(\theta_L) = \frac{1}{C_l} \int_{-\cos \theta_L}^{\cos \theta_L} \frac{P_l^l(\cos \theta') P_l^1(\cos \theta')}{1 - \cos^2 \theta'} d(\cos \theta') \quad [2.3.29]$$

This integral is easy to evaluate for the case of infinite length; as $L \rightarrow \infty$, $\theta_L \rightarrow 0$ and we can use the useful m -orthogonal property of associated Legendre functions:

$$\int_{-1}^1 \frac{P_l^m(\cos\theta') P_l^{m'}(\cos\theta')}{1 - \cos^2\theta'} d(\cos\theta') = \delta_{mm'} \frac{1}{m} \frac{(l+m)!}{(l-m)!} \quad [2.3.30]$$

To give:

$$I_l(0) = \frac{1}{C_l} \frac{(l+1)!}{(l-1)!} \delta_{1l} \quad [2.3.31]$$

This gives the expression for the potential as:

$$\varphi_{<}(r, \theta, \phi) = -M_r P_1^1(\cos\theta) \sin\phi \frac{r}{\rho'} d\rho' \quad [2.3.32]$$

due to the presence of the Kronecker delta δ_{1l} causing all terms $l \neq 1$ to drop out, and also since C_1 is simply 1, as $P_1^1(\cos\theta) = \sin\theta$. Recognising that $[(\sin\theta)(\sin\phi)r]$ is simply y in Cartesian coordinates, and integrating over infinitesimal thickness cylinders from inner cylinder radius R_i to outer radius R_o , we have:

$$\varphi_{<}(r, \theta, \phi) = -M_r y \int_{R_i}^{R_o} \frac{1}{\rho'} d\rho', \quad r < R_i \quad [2.3.33]$$

Giving:

$$\varphi_{<}(r, \theta, \phi) = -M_r y \ln\left(\frac{R_o}{R_i}\right), \quad r < R_i \quad [2.3.34]$$

And, taking the negative gradient of this expression gives the \mathbf{H} field as:

$$\mathbf{H}(r, \theta, \phi) = M_r \ln\left(\frac{R_o}{R_i}\right) \hat{\mathbf{y}}, \quad r < R_i \quad [2.3.35]$$

Which is a uniform field in the y -direction inside the bore of the cylinder, the magnitude of which depends on the logarithm of the ratio of the outer to inner radii, as expected.

2.3.3 Finite Length Halbach Cylinder

Our expression for the potential [2.3.26] must be evaluated numerically for the finite length case, but some further reduction can be achieved due to the integral within the expression being symmetric:

$$I_l(\theta_L) = \int_{\theta_i}^{\pi-\theta_i} \sin^{l-1}(\theta') P_l^1(\cos\theta') d\theta' \quad [2.3.36]$$

The associated Legendre functions can be given in terms of Legendre polynomials by the expression:

$$P_l^m(x) = (-1)^m (1-x^2)^{\frac{m}{2}} \frac{dP_l(x)}{dx^m} \quad [2.3.37]$$

Using this for $m=l$ the integral expression [2.2.36] becomes:

$$I_l(\theta_L) = - \int_{\theta_L}^{\pi-\theta_L} \sin \theta' \frac{dP_l(\cos \theta')}{d(\cos \theta')} d\theta' \quad [2.3.38]$$

The Legendre polynomials $P_l(x)$ are polynomials of degree l which contain *only* even powers of x if l is even and only odd powers of x if l is odd. This means that the derivative will give a polynomial with only odd powers if l is even and vice-versa.

For example:

$P_l(x)$	Polynomial	Derivative
$P_0(x)$	1	0
$P_1(x)$	x	1
$P_2(x)$	$\frac{1}{2}(3x^2 - 1)$	$\frac{1}{2}(6x)$
$P_3(x)$	$\frac{1}{2}(5x^3 - 3x)$	$\frac{1}{2}(15x^2 - 3)$
$P_4(x)$	$\frac{1}{8}(35x^4 - 30x^2 + 3)$	$\frac{1}{8}(140x^3 - 60x)$

Generally, the integral can be expressed as:

$$I_l(\theta_L) = \sum_i k_i \int_{\theta_L}^{\pi-\theta_L} \sin \theta' \cos^{n_i} \theta' d\theta' \quad [2.3.39]$$

Where k_i and n_i depend on the terms in the derivative of the polynomial and the sum is taken over the number of terms in the derivative of the polynomial P_l . The integral above gives:

$$I_l(\theta_L) = \sum_i \frac{k_i}{n+1} [1 - (-1)^{n+1}] \cos^{n+1} \theta_L \quad [2.3.40]$$

Which disappears for all n odd. This means that the only contribution to the integral comes from terms where the derivative of the polynomials has n even, i.e. the *odd* Legendre polynomials which occur for l odd. The result of this is that the integral in the potential [2.3.26] is zero for all l even, and only the odd l terms contribute; this is a direct result of the symmetry about the z -axis. The full expression for the potential of the finite length cylinder is then given by [2.3.26], where we have also integrated ρ' from R_i to R_o :

$$\varphi_{<}(r, \theta, \phi) = M_r \ln \left(\frac{R_o}{R_i} \right) \sum_{l_{\text{odd}}=1}^{\infty} \frac{(l-1)!}{(l+1)!} P_l^1(\cos \theta) \sin \phi \ r^l \int_{\theta_L}^{\pi-\theta_L} \sin \theta' \frac{dP_l(\cos \theta')}{d(\cos \theta')} d\theta' \quad [2.3.41]$$

and the summation is taken with l odd only. If we take $\theta = 0$ and $\phi = \pi/2$, r is the Cartesian y-axis (and r from now on is expressed as y), giving:

$$\varphi_{<}(r, \theta, \phi) = M_r \ln \left(\frac{R_o}{R_i} \right) \sum_{l_{\text{odd}}=1}^{\infty} \frac{(l-1)!}{(l+1)!} P_l^1(1) y^l \int_{\theta_L}^{\pi-\theta_L} \sin \theta' \frac{dP_l(\cos \theta')}{d(\cos \theta')} d\theta' \quad [2.3.42]$$

We define S_l , the sum terms (without the factor y^l), as:

$$S_l = \frac{(l-1)!}{(l+1)!} P_l^1(1) \int_{\theta_L}^{\pi-\theta_L} \sin \theta' \frac{dP_l(\cos \theta')}{d(\cos \theta')} d\theta' \quad [2.3.43]$$

Giving [2.3.42] as:

$$\varphi_{<}(r, \theta, \phi) = M_r \ln \left(\frac{R_o}{R_i} \right) \sum_{l_{\text{odd}}=1}^{\infty} y^l S_l \quad [2.3.44]$$

Using numerical evaluation, we can plot each of the S_l terms to show their relative contribution to the entire potential for the range $\theta_L = 0$ (representing an infinite length cylinder) $\rightarrow \pi/2$ (representing a cylinder of zero length):




Figure 2.6 – S_1, S_3, S_5 and S_7 from expression [2.3.43] plotted together (top), S_1, S_3 and S_5 plotted separately (bottom; left to right, respectively).

By inspecting the graph, we can see that the $l = 1$ term provides the only contribution to the potential at infinite length, where the entire summation comes to unity and agrees with our special case for the potential for infinite length derived above

[2.3.35]. The remainder of the l terms can therefore be thought of as a series of diminishing ‘corrections’ caused by the finite cylinder length. In cases where homogeneity is important, we seek to reduce all l terms *apart* from the $l = 1$ term as these are the terms that represent the departure from the ideal case and introduce inhomogeneities into the field. This is to try to make the field as close as possible to the ideal infinite length case.

3 Numerical Methods for Magnetic Field Calculations

3.1 Surface Charge Model

For magnet arrays composed from uniformly magnetized segments of hard, magnetically transparent, permanent magnetic material it is possible to calculate the field using the magnetic charge model (see section 2.2.1). Since the magnetization is fixed and uniform within each segment the calculation reduces to finding the field due to a scalar magnetic charge density on each surface of every segment.

The surfaces of each segment can be broken into *simplexes* in which the entire magnet array can be expressed. Our strategy for field calculation is to, firstly, find an analytical expression for the field created by an individual simplex then, secondly, to break the magnet array into simplexes and simply sum over all the simplexes to find the field at any required point.

3.1.1 2D Surface Charge Model

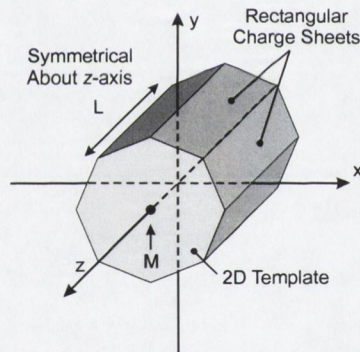


Figure 3.1 – A magnetic design subdivided into rectangular charge sheets.

Designs based on a 2D template translational invariant about the z-axis, composed of segments with uniform magnetization \mathbf{M} , such as a segmented cylinder, can be reduced into simplexes of finite length magnetically charged sheets (see figure 3.1). A general analytic expression can be derived for a charged sheet in the x - z plane. We can then use geometrical transformations to move the sheet to the ‘correct’ location within the design and find the field at any point by summing over all the charged sheets in the design.




Figure 3.2 A rectangular magnetic charge sheet 'simplex'.

We use a right handed co-ordinate system throughout the calculations. For derivation of the fields due to the sheet, please refer to the appendix A.1. The component parallel to the plane of the sheet, H_x , is given by:

$$H_x(P) = \frac{\sigma_m}{4\pi\mu_0} \left[\ln \left(\frac{(z+L) + \sqrt{(z+L)^2 + (x-a)^2 + y^2}}{(z-L) + \sqrt{(z-L)^2 + (x-a)^2 + y^2}} \right) + \ln \left(\frac{(z+L) + \sqrt{(z+L)^2 + (x+a)^2 + y^2}}{(z-L) + \sqrt{(z-L)^2 + (x+a)^2 + y^2}} \right) \right] \quad [3.1.1]$$

Where L is the length of the sheet in the z -direction and σ_m is the magnetic charge density. The expression for H_z is similar:

$$H_z(P) = \frac{\sigma_m}{4\pi\mu_0} \left[\ln \left(\frac{(x+a) + \sqrt{(x+a)^2 + (z-L)^2 + y^2}}{(x-a) + \sqrt{(x-a)^2 + (z-L)^2 + y^2}} \right) + \ln \left(\frac{(x+a) + \sqrt{(x+a)^2 + (z+L)^2 + y^2}}{(x-a) + \sqrt{(x-a)^2 + (z+L)^2 + y^2}} \right) \right] \quad [3.1.2]$$

The expression for the component normal to the sheet, H_y , is given by:

$$H_y(P) = \frac{\sigma_m}{4\pi\mu_0 y} \left[\tan^{-1} \left(\frac{(z-L)(x-a)}{y\sqrt{(x-a)^2 + y^2 + (z-L)^2}} \right) - \tan^{-1} \left(\frac{(z+L)(x-a)}{y\sqrt{(x-a)^2 + y^2 + (z+L)^2}} \right) - \tan^{-1} \left(\frac{(z-L)(x+a)}{y\sqrt{(x+a)^2 + y^2 + (z-L)^2}} \right) + \tan^{-1} \left(\frac{(z+L)(x+a)}{y\sqrt{(x+a)^2 + y^2 + (z+L)^2}} \right) \right] \quad [3.1.3]$$

The magnetic surface charge density is given by:

$$\sigma_m = \mathbf{M} \cdot \hat{\mathbf{n}} \quad [3.1.4]$$

Where \hat{n} is the unit outward normal from the surface.

3.1.2 2D Transformations

Having derived the analytic solution for a sheet of magnetic charge we can now find the field at any point P in space created by a magnetic design of the type shown in figure 3.1. This is carried out by breaking the design into charge sheets and computing the linear sum of fields from each of these charge sheet. However, we have derived expressions for H_x , H_y and H_z for a charge sheet lying flat in the x - z plane and centred at the origin. To sum over the effects of variously orientated and positioned sheets we need to use a global frame; we must first express the position of P in terms of axes translated and rotated relative to these the global (x, y) axes of this frame. This is fairly straightforward in 2D, we first express the position $P(x,y)$ in terms of coordinate axes x', y' related to x, y by a translation. Then we express $P(x',y')$ in terms of the rotated coordinate system x'', y'' , which is a rotation of the axes and not the point (figure 3.3 - the rotation shown here is $+\theta$, as the positive rotation sense is clockwise looking down the z -axis).




Figure 3.3 – An arbitrary charge sheet in a magnetic design (thick line, left), expressing P in the local un-rotated sheet frame (right).

The expression derived above for H_x , H_y and H_z can then be used to find the field at P due to the charge sheet with centre at (S_x, S_y) and parallel to the x'' axis. We can express the translation in homogeneous coordinates as:

$$\begin{pmatrix} P_{x'} \\ P_{y'} \\ 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & -S_x \\ 0 & 1 & -S_y \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} P_x \\ P_y \\ 1 \end{pmatrix} \quad [3.1.5]$$

Where x' and y' refer to the translated coordinate system and S_x and S_y refer to the midpoint of the 2D charge sheet. Homogeneous co-ordinates simply allow for translations to be easily included in the matrix expression; the last row of the co-

ordinate vectors containing unity can be ignored after calculation. The rotation matrix which transforms x' and y' into the frame with x'' parallel to the sheet and y'' normal to the sheet by a rotation through θ is given by:

$$\begin{pmatrix} P_{x''} \\ P_{y''} \\ 1 \end{pmatrix} = \begin{pmatrix} \cos(\theta) & -\sin(\theta) & 0 \\ \sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} P_{x'} \\ P_{y'} \\ 1 \end{pmatrix} \quad [3.1.6]$$

The full transformation is then given by the product of the two matrices [3.1.5] and [3.1.6]:

$$\begin{pmatrix} P_{x''} \\ P_{y''} \\ 1 \end{pmatrix} = \begin{pmatrix} \cos(\theta) & -\sin(\theta) & [-S_x \cos(\theta) + S_y \sin(\theta)] \\ \sin(\theta) & \cos(\theta) & [-S_x \sin(\theta) - S_y \cos(\theta)] \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} P_x \\ P_y \\ 1 \end{pmatrix} \quad [3.1.7]$$

We can use [3.1.7] to transform the coordinates of the point P in the global frame to coordinates in the sheet frame. Using the expressions for \mathbf{H} , we will calculate $H_{x''}$ and $H_{y''}$ in the sheet frame. The components $H_{x''}$ and $H_{y''}$ must be transformed back into the world frame by an inverse rotation, that is, [3.1.7] with $\theta = -\theta$.




Figure 3.4 - The inverse rotation of the components of the magnetic field vector H .

3.1.3 2D Computational Realisation

The first computational tool, called Magnetic Solver, created for this project summed the contributions from charge sheets of equal length L in the z -direction (these are shown as black lines in figure 3.5). Setting $z = 0$ in the expressions given above gives H_x as:

$$H_x(P) = \frac{\sigma_m}{4\pi\mu_0} \left[\ln \left(\frac{\sqrt{L^2 + (x-a)^2 + y^2} + L}{\sqrt{L^2 + (x-a)^2 + y^2} - L} \right) + \ln \left(\frac{\sqrt{L^2 + (x+a)^2 + y^2} + L}{\sqrt{L^2 + (x+a)^2 + y^2} - L} \right) \right] \quad [3.1.8]$$

And H_y as:

$$H_y(P) = \frac{\sigma_m}{2\pi\mu_0 y} \left[\tan^{-1} \left(\frac{L(x-a)}{y\sqrt{(x-a)^2 + y^2 + L^2}} \right) + \tan^{-1} \left(\frac{L(x+a)}{y\sqrt{(x+a)^2 + y^2 + L^2}} \right) \right] \quad [3.1.9]$$

To represent a physical magnet all surfaces enclosing each individual magnet segment must be considered. The surface charge on each of these segment faces was calculated (using [2.2.12]) and stored prior to any calculation. The coefficients of the transform matrix for each face were also stored prior to each calculation.

The user input was given in terms of primitives through a window user interface. Primitives included simple N-sided polygons blocks with user defined parameters being the fixed magnetization and length to more complex primitives such as Halbach cylinders with user-defined parameters including polarization, number of segments, inner and outer radii and length. Data could be obtained by directly writing values along a line through the design to a file. Visualization was also included; either a vector plot or a colour slice through the $z = 0$ plane could be rendered, this was carried out using Windows GDI routines.




Figure 3.5 – Designs of constant cross-section can be modelled with the first version of the Magnetic Solver tool.

The magnetic field is represented by the blue arrows, the magnetization vector in each block by the green arrows.

Figure 3.5 shows two designs with constant cross section and finite length perpendicular to the plane of the diagram: a magnetized bar with irregular cross section and a Halbach cylinder. Each of the black lines in the designs represents a charge sheet, the magnetic charge on each of these sheets is calculated by taking the

dot product of the normal vector (shown in red) with the magnetization vector (shown in green); the \mathbf{H} field is shown by the blue arrows.

This particular application was fairly simple and simply served as a step towards the next development, an application which could investigate fully three dimensional magnetostatic designs.

3.1.4 3D Surface Charge Model

The 2D charge model suffers from a number of restrictions. First, it only allows us to calculate the field in the $z = 0$ plane; we would like to examine the field due to the magnetic segments at any point in space. Secondly, it only allows designs with a constant cross-section based on some 2D template, as shown in figure 3.1. We would like to explore fully three dimensional designs in which the design is not restricted in these ways.




Figure 3.6- A magnetic design subdivided into right-angle triangle surface simplexes

These restrictions led to the development of the second version of the Magnetic Solver program, which can handle arbitrary sheets of magnetic charge in any orientation in 3D space. A simplex in 3D space is a triangle, and this can be broken further into two right-angled triangles (figure 3.6). Any design can then be represented in terms of simplexes of right-angled triangles.




Figure 3.7 – A right-angled triangular magnetic charge sheet 'simplex'.

The analytical expression for a right-angled triangle is much more complex than for that of a charge sheet; please see appendix A.2 for derivations. The calculations are carried out for a triangle lying in the x - z plane, with the y -axis normal to the surface of the triangle as shown in figure 3.7.

The component parallel to the sheet, H_x , is given by:

$$H_x(P) = \frac{\sigma_m}{2\pi\mu_0} \left[\ln \left(\frac{b + \sqrt{(a-x)^2 + y^2 + (b-z)^2} - z}{\sqrt{(a-x)^2 + y^2 + z^2} - z} \right) + \frac{b}{\sqrt{a^2 + b^2}} \ln \left(\frac{r - \frac{ax+bz}{\sqrt{a^2+b^2}}}{\frac{a(a-x)+b(b-z)}{\sqrt{a^2+b^2}} + \sqrt{(a-x)^2 + y^2 + (b-z)^2}} \right) \right] \quad [3.1.10]$$

There is a similar expression for H_z :

$$H_z(P) = \frac{\sigma_m}{2\pi\mu_0} \left[\ln \left(\frac{r-x}{a + \sqrt{(a-x)^2 + y^2 + z^2} - x} \right) + \frac{a}{\sqrt{a^2 + b^2}} \ln \left(\frac{\frac{a(a-x)+b(b-z)}{\sqrt{a^2+b^2}} + \sqrt{(a-x)^2 + y^2 + (b-z)^2}}{r - \frac{ax+bz}{\sqrt{a^2+b^2}}} \right) \right] \quad [3.1.11]$$

Where:

$$r = \sqrt{x^2 + y^2 + z^2} \quad [3.1.12]$$

The component normal to the sheet, H_y , is a sum of arctangents, in a similar way to the expression for the rectangular charge sheet in section 3.1.1:

$$H_y(P) = \frac{\sigma_m}{2\pi\mu_0} \left[\tan^{-1} \left(\frac{(x-a) + z + \sqrt{(x-a)^2 + y^2 + z^2}}{y} \right) - \tan^{-1} \left(\frac{x+z+r}{y} \right) + \frac{y}{\alpha\sqrt{B^2 - A^2 - C^2}} \left\{ \tan^{-1} \left(\frac{C + (A+B) \left(\sqrt{1 + \frac{(\alpha+\beta)^2}{\gamma^2}} + \frac{\alpha+\beta}{\gamma} \right)}{\sqrt{B^2 - A^2 - C^2}} \right) - \tan^{-1} \left(\frac{C + (A+B) \left(\sqrt{1 + \frac{\beta^2}{\gamma^2}} + \frac{\beta}{\gamma} \right)}{\sqrt{B^2 - A^2 - C^2}} \right) \right\} \right] \quad [3.1.13]$$

Where:

$$\alpha = \sqrt{1 + \frac{b^2}{a^2}}, \quad \beta = -\frac{x + \frac{bz}{a}}{1 + \frac{b^2}{a^2}}, \quad \gamma = \sqrt{-\beta^2 + \frac{r^2}{1 + \frac{b^2}{a^2}}} \quad [3.1.14]$$

And:

$$A = -\gamma \frac{b}{a}, \quad B = \gamma \sqrt{1 + \frac{b^2}{a^2}}, \quad C = z + \beta \frac{b}{a} \quad [3.1.15]$$

3.1.5 3D Transformations

The transformations between local charge sheet frame and the global frame in 3D are similar to the 2D transformations, apart from the addition of two extra rotation angles

and the fact that each polygon consists of two back-to-back right-angle triangle simplexes. The use of three rotation angles means that care has to be taken about the order of rotation; a rotation of a point around the x -axis followed by a rotation around the y -axis is not the same as a rotation about the y -axis followed by the rotation about the x -axis. We chose to rotate the frame and not the vectors, with the positive sense of rotation the *clockwise* direction looking down the positive axis at the origin.




Figure 3.8 - Positive rotations for the transformations in a right-handed co-ordinate system.

We first translate our polygon to the origin, so that the point along the longest edge (edge AC) where the right angled triangles meet (point D) lies at the origin. We then rotate the polygon about the y -axis (by an angle β) so the edge AC lies in the x - y plane, then rotate the polygon about the z -axis (by an angle α) so that edge AC lies parallel to the x -axis. Then we rotate about the final axis, the x -axis (by an angle γ), such that point B is rotated into the x - z plane. The polygon then lies flat in the x - z plane.

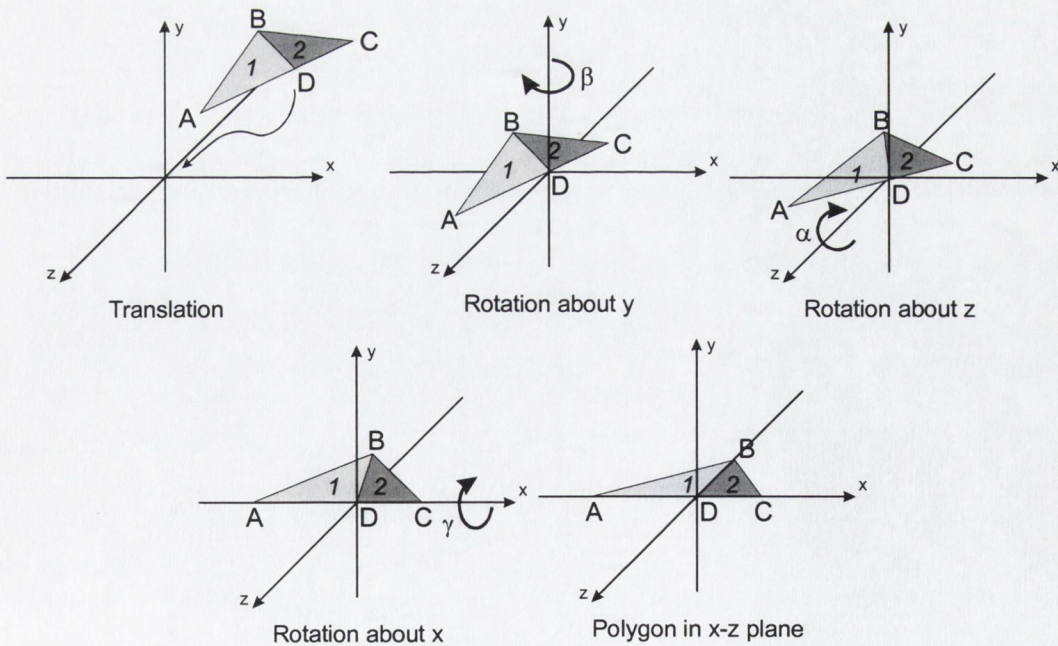


Figure 3.9 – The various transformations required for transforming an arbitrary polygon from the global frame to the charge sheet frame.

The polygon consists of two back-to-back triangles; this means triangle ‘2’ must be flipped as it is ‘upside-down’ with respect to our standard right-angled triangle simplex (figure 3.10). If the polygon must undergo the rotations α , β and γ , triangle ‘1’ must undergo the same rotations and the second ‘upside-down’ triangle ‘2’ must undergo the rotations $-\alpha$, $\pi+\beta$ and $\pi-\gamma$ to flip it to the same orientation as the standard triangle. The component triangles must also be *separately* translated after the rotation; triangle ‘1’ must be translated so that the vertex A lies at the origin and triangle ‘2’ must be translated so that the vertex C lies at the origin.




Figure 3.10 – The two triangles, triangle 1 (top) and triangle 2 (bottom), require different transformations.

The rotation matrix $\mathbf{R}_x(\gamma)$ for a rotation clockwise around the x -axis by γ is given by³²:

$$\begin{pmatrix} P_{x''} \\ P_{y''} \\ P_{z''} \\ 1 \end{pmatrix}_\gamma = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\gamma) & \sin(\gamma) & 0 \\ 0 & -\sin(\gamma) & \cos(\gamma) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} P_{x'} \\ P_{y'} \\ P_{z'} \\ 1 \end{pmatrix} \quad [3.1.16]$$

The rotation matrix $\mathbf{R}_y(\beta)$ for a rotation clockwise around the y -axis by β is given by:

$$\begin{pmatrix} P_{x''} \\ P_{y''} \\ P_{z''} \\ 1 \end{pmatrix}_\beta = \begin{pmatrix} \cos(\beta) & 0 & -\sin(\beta) & 0 \\ 0 & 1 & 0 & 0 \\ \sin(\beta) & 0 & \cos(\beta) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} P_{x'} \\ P_{y'} \\ P_{z'} \\ 1 \end{pmatrix} \quad [3.1.17]$$

The rotation matrix $\mathbf{R}_z(\alpha)$ for a rotation clockwise around the z -axis by α is given by:

$$\begin{pmatrix} P_{x''} \\ P_{y''} \\ P_{z''} \\ 1 \end{pmatrix}_\gamma = \begin{pmatrix} \cos(\alpha) & \sin(\alpha) & 0 & 0 \\ -\sin(\alpha) & \cos(\alpha) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} P_{x'} \\ P_{y'} \\ P_{z'} \\ 1 \end{pmatrix} \quad [3.1.18]$$

The rotation matrix \mathbf{R}_1 for triangle 1 is given by:

$$R_1(\alpha, \beta, \gamma) = R_x(\gamma)R_z(\alpha)R_y(\beta) \quad [3.1.19]$$

And the rotation matrix \mathbf{R}_2 for triangle 2 is given by:

$$R_2(\alpha, \beta, \gamma) = R_x(\pi - \gamma)R_z(-\alpha)R_y(\pi + \beta) \quad [3.1.20]$$

The translations are omitted here, but are carried out in the same way as the translations on 2D charge sheet. As with the 2D charge sheet model, the expressions for \mathbf{H} will be in the local sheet frame and the components H_x'' and H_y'' must be transformed back into the global frame by an inverse transformation. To transform to $\mathbf{H}(x, y)$ the inverse rotation must be used: The components of \mathbf{H} are rotated by an angle $-\gamma$ about the x -axis, by an angle $-\alpha$ about the z -axis and by an angle $-\beta$ about the y -axis. The inverse rotation matrix \mathbf{R}^{-1} is given by:

$$R_1^{-1}(\alpha, \beta, \gamma) = R_y(-\beta)R_z(-\alpha)R_x(-\gamma) \quad [3.1.21]$$

3.1.6 3D Computational Realisation

The second application developed, called Magnetic Solver 2, was much more complex than the earlier 2D version. This version included very limited scripting support which allowed for design parameters to be automatically varied and field data to be written to a file, which could be examined and compared in detail later.




Figure 3.11- Magnetic Solver version 2 can simulate magnetic fields from arbitrarily shaped (closed) blocks of hard magnetic material. The field is represented by the black arrows, the magnetization vector by the small green arrows.

The final version allowed primitive creation both through a window user interface and through the scripting interface. Primitive construction objects included quadrilaterals, in which all eight vertices could be defined, and cylinders, in which radius and length could be defined. The construction objects also included complex primitives such as Halbach cylinders, in which many extra parameters could be defined such as the inner bore radius and inner and outer cylinder lengths and Halbach spheres, in which the inner and outer radii could be defined. Primitive input was controlled by **Window.cpp** and **Script.cpp**.

All primitives are split into polygons, which were further split into right angled triangles. This is shown in figure 3.11; the grey lines define surface polygons, the magnetic charge on which is calculated from the dot product of the surface normal vector (red arrow) and the fixed magnetization vector for each block (green arrow). At construction-time all possible pre-calculations for the more complex field calculation stage are carried out. This includes finding the correct rotation angle and charge for each polygonal face. The calculation stage consists of finding, for a given point, the sum of the magnetic fields from all surface polygons in the design. During this stage each right angle triangle is rotated and translated to the correct position in 3D space, the field contribution calculated and the resulting field vector transformed to the global frame, this procedure is carried out in **Magnet.cpp**.

The resulting field can be rendered to a colour-coded or vector field slice in the x - y , x - z or y - z planes. The graphics are handled using OpenGL routines in **Graphics.cpp**. The field values can also be determined at any point or along any straight line through the design and written to a file, for examination at a later stage or for plotting by some external application.

3.2 Finite Element Method

The methods in the previous section calculated magnetic fields by summing over the individual contributions from field sources within a design. Designs were discretised into sources of magnetic surface charge and at the heart of the calculations were general analytic solutions to surface simplexes (a sheet in the 2D case and a triangle in the 3D case). However, the programs developed using the charge model were based on the transparency property of modern hard magnetic materials, and for designs

incorporating materials which do *not* possess this property a totally different computational modelling strategy had to be adopted.

The finite-element based program was therefore developed to allow the calculation of fields from magnetic designs incorporating *soft* magnetic materials. In general, the finite element model allows for complete flexibility in the choice of materials used in the design, whereas the charge model is restricted to only hard modern magnetic materials. The finite element method uses a totally different approach to the solution of the problem; in general, the method approximates the exact solution to the Poisson equation for the fields [2.2.18] as a series of independent functions and attempts to minimize the error or residual, which is the difference between the true solution and approximate solution³³.

3.2.1 Weighted Residual Methods

If we take the general differential equation:

$$Lu(x) = f(x) \quad [3.2.1]$$

where L is a linear operator and $f(x)$ is a known function, the solution of [3.2.1] can be shown to be equivalent to the solution of:

$$\langle Lu, v \rangle = \langle f, v \rangle \quad [3.2.2]$$

where $v(x)$ is any continuous function and the L^2 inner product is given by:

$$\langle p, q \rangle = \int_{\tau} pq d\tau \quad [3.2.3]$$

so:

$$\int_V [Lu(x) - f(x)]v(x)dV = 0 \quad [3.2.4]$$

Equation [3.2.4] is called the *weak form* of [3.2.1]. [3.2.4] will also be subject to the boundary conditions:

$$Mu(x) = g(x) \quad [3.2.5]$$

where M is a linear operator and $g(x)$ is a known function. This expression can also be given in the weak form, and a residual expression for the whole domain can be obtained by summing both the boundary expression and the volume expression³⁴:

$$\int_V [Lu(x) - f(x)]v(x)dV + \oint_S [Mu(x) - g(x)]w(x)dS = 0 \quad [3.2.6]$$

where $v(x)$ and $w(x)$ are two independent weighting functions. This is the form which we will attempt to solve, but for the time being we will only consider the first part of

the expression; the contribution of the boundary condition will be considered later. We introduce the approximation:

$$u'(x) = \sum_{i=1}^N a_i \phi_i(x) \quad [3.2.7]$$

or, in vector notation:

$$u'(x) = \mathbf{a}\boldsymbol{\phi}(x) = \boldsymbol{\phi}^T(x)\mathbf{a}^T \quad [3.2.8]$$

where:

$$\mathbf{a} = (a_1 \quad a_2 \quad \dots \quad a_N)$$

$$\boldsymbol{\phi}(x) = (\phi_1(x) \quad \phi_2(x) \quad \dots \quad \phi_N(x))^T \quad [3.2.9]$$

and $\phi_i(x)$ are a set of linearly independent functions and a_i the set of coefficients in [3.2.7]. Using the approximation u' , [3.2.4] becomes:

$$\int_V [\mathbf{L}u'(x) - f(x)]v(x)dV = 0 \quad [3.2.10]$$

But, in general:

$$[\mathbf{L}u'(x) - f(x)] - [\mathbf{L}u(x) - f(x)] \neq 0 \quad [3.2.11]$$

Leading to the definition of the *error* (or *residual*) $\varepsilon(x)$ as the difference between the real solution and our approximate function:

$$\varepsilon(x) = [\mathbf{L}u'(x) - f(x)] - [\mathbf{L}u(x) - f(x)] \quad [3.2.12]$$

From equation [3.2.1], this gives:

$$\varepsilon(x) = \mathbf{L}u'(x) - f(x) \quad [3.2.13]$$

Substituting our expression for u' , this becomes:

$$\varepsilon(x) = \mathbf{a}[\mathbf{L}\boldsymbol{\phi}(x)] - f(x) = [\mathbf{L}\boldsymbol{\phi}(x)]^T \mathbf{a}^T - f(x) \quad [3.2.14]$$

Equation [3.2.10] then becomes:

$$\int_V v(x)\varepsilon(x)dV = 0 \quad [3.2.15]$$

where $v(x)$ is still an arbitrary function. If we now represent $v(x)$ by another set of linearly independent functions $\psi_j(x)$ with coefficients b_j :

$$v(x) = \sum_{j=1}^N b_j \psi_j(x) \quad [3.2.16]$$

In vector notation:

$$v(x) = \mathbf{b}\boldsymbol{\psi}(x) \quad [3.2.17]$$

Where:

$$\mathbf{b} = (b_1 \quad b_2 \quad \dots \quad b_N) \quad [3.2.18]$$

$$\boldsymbol{\psi}(x) = (\psi_1(x) \quad \psi_2(x) \quad \dots \quad \psi_N(x))^T$$

and $\psi(x)$ is known as the weighting function. Substituting into equation [3.2.15] gives:

$$\mathbf{b} \int_V \boldsymbol{\psi}(x) \varepsilon(x) dV = 0 \quad [3.2.19]$$

This must hold for all \mathbf{b} , so:

$$\int_V \boldsymbol{\psi}(x) \varepsilon(x) dV = 0 \quad [3.2.20]$$

[3.2.20] is the basic equation for all weighted residual methods. Substituting the transposed form of our expression for $\varepsilon(x)$ [3.2.14] gives:

$$\left[\int_V \boldsymbol{\psi}(x) [\mathbf{L}\boldsymbol{\phi}(x)]^T dV \right] \mathbf{a}^T = \int_V \boldsymbol{\psi}(x) f(x) dV \quad [3.2.21]$$

This is a matrix equation of the form $\mathbf{A}\mathbf{x} = \mathbf{b}$:

$$\begin{pmatrix} \int_V \psi_1(x) \mathbf{L}\phi_1(x) & \int_V \psi_1(x) \mathbf{L}\phi_2(x) & \dots & \int_V \psi_1(x) \mathbf{L}\phi_N(x) \\ \int_V \psi_2(x) \mathbf{L}\phi_1(x) & \int_V \psi_2(x) \mathbf{L}\phi_2(x) & \dots & \int_V \psi_2(x) \mathbf{L}\phi_N(x) \\ \dots & \dots & \dots & \dots \\ \int_V \psi_N(x) \mathbf{L}\phi_1(x) & \int_V \psi_N(x) \mathbf{L}\phi_2(x) & \dots & \int_V \psi_N(x) \mathbf{L}\phi_N(x) \end{pmatrix} \begin{pmatrix} a_1 \\ a_2 \\ \dots \\ a_N \end{pmatrix} = \begin{pmatrix} \int_V \psi_1(x) f(x) \\ \int_V \psi_2(x) f(x) \\ \dots \\ \int_V \psi_N(x) f(x) \end{pmatrix} \quad [3.2.22]$$

This can then be solved for the unknown coefficients a .

Collocation Method

The simplest approach for minimizing the residual is to force it to be zero at N specified points x_i within the domain. This approach is called the collocation method and simply involves setting the weighting function $\psi(x)$ to a Dirac delta function³⁵.

$$\psi_i(x) = \delta(x - x_i), \quad i = \{1, \dots, N\} \quad [3.2.23]$$

For each component of [3.2.20] this gives:

$$\int_V \delta(x - x_i) \varepsilon(x) dV = 0 \quad [3.2.24]$$

So that:

$$\varepsilon(x_i) = 0 \quad [3.2.25]$$

Substituting [3.2.23] into [3.2.21] gives:

$$[\mathbf{L}\boldsymbol{\phi}(x_i)]^T \mathbf{a}^T = f(x_i) \quad [3.2.26]$$

Which are N simultaneous equations, one for each x_i . These can be solved by using the matrix system:

$$\begin{pmatrix} L\phi_1(x_1) & L\phi_2(x_1) & \dots & L\phi_N(x_1) \\ L\phi_1(x_2) & L\phi_2(x_2) & \dots & L\phi_N(x_2) \\ \dots & \dots & \dots & \dots \\ L\phi_1(x_N) & L\phi_2(x_N) & \dots & L\phi_N(x_N) \end{pmatrix} \begin{pmatrix} a_1 \\ a_2 \\ \dots \\ a_N \end{pmatrix} = \begin{pmatrix} f(x_1) \\ f(x_2) \\ \dots \\ f(x_N) \end{pmatrix} \quad [3.2.27]$$

Subdomain Collocation Method

The collocation method is simple, but a better method is to divide the domain into a set of subdomains and make the *average* of the residual over these subdomains zero. This more closely approximates the requirement that the residual must be zero over the entire domain. This is known as the subdomain collocation method. To make the average of the residual zero over a subdomain we require:

$$\int_{x_i}^{x_{i+1}} \mathcal{E}(x) dx = 0 \quad [3.2.28]$$

Where the interval $x_i \rightarrow x_{i+1}$ is the i^{th} subdomain. This means that the weighting function $\psi(x)$ must be of the form:

$$\psi_i(x) = \begin{cases} 1 & x_i \leq x \leq x_{i+1} \\ 0 & \text{otherwise} \end{cases} \quad [3.2.29]$$

Substitution of this expression for $\psi(x)$ into [3.2.21] gives the matrix form of this as:

$$\left[\int_{x_i}^{x_{i+1}} [L\phi(x)]^T dx \right] \mathbf{a}^T = \int_{x_i}^{x_{i+1}} f(x) dx \quad [3.2.30]$$

Which are N simultaneous equations, one for each subdomain, and can be solved for in the same way as the collocation method.

Least Squares Method

A more accurate method is to minimize the square of the error in a global sense. This is called the least squares method, used for line-fitting in this project. We require³⁵:

$$\text{Min} \left\{ \int_V \mathcal{E}(x)^2 dV \right\} \quad [3.2.31]$$

Which is:

$$\frac{\partial}{\partial a_i} \left\{ \int_V \mathcal{E}(x)^2 dV \right\} = 0 \quad [3.2.32]$$

Giving our residual expression as:

$$2 \int_V \varepsilon(x) \frac{\partial \varepsilon(x)}{\partial a_i} dV = 0 \quad [3.2.33]$$

Comparing to [3.2.20], this gives the trial function as:

$$\psi_i(x) = \frac{\partial \varepsilon(x)}{\partial a_i} \quad [3.2.34]$$

Differentiating [3.2.14] with respect to a_i gives this as:

$$\psi_i(x) = L\phi_i(x) \quad [3.2.35]$$

The matrix form [3.2.21] is then given by:

$$\left[\int_V [L\phi(x)][L\phi(x)]^T dV \right] \mathbf{a}^T = \int_V [L\phi(x)]f(x)dV \quad [3.2.36]$$

Galerkin's Method

In Galerkin's method the trial functions are the same as the weighing functions, i.e:

$$\psi_i(x) = \phi_i(x) \quad [3.2.37]$$

Substituting this into [3.2.21] gives the matrix form of Galerkin's method as:

$$\left[\int_V \phi(x)[L\phi(x)]^T dV \right] \mathbf{a}^T = \int_V \phi(x)f(x)dV \quad [3.2.38]$$

Expanding this gives:

$$\begin{pmatrix} \int_V \phi_1(x)[L\phi_1(x)] & \int_V \phi_1(x)[L\phi_2(x)] & \dots & \int_V \phi_1(x)[L\phi_N(x)] \\ \int_V \phi_2(x)[L\phi_1(x)] & \int_V \phi_2(x)[L\phi_2(x)] & \dots & \int_V \phi_2(x)[L\phi_N(x)] \\ \dots & \dots & \dots & \dots \\ \int_V \phi_N(x)[L\phi_1(x)] & \int_V \phi_N(x)[L\phi_2(x)] & \dots & \int_V \phi_N(x)[L\phi_N(x)] \end{pmatrix} \begin{pmatrix} a_1 \\ a_2 \\ \dots \\ a_N \end{pmatrix} = \begin{pmatrix} \int_V \phi_1(x)f(x) \\ \int_V \phi_2(x)f(x) \\ \dots \\ \int_V \phi_N(x)f(x) \end{pmatrix} \quad [3.2.39]$$

This is the weighting of choice for the finite element method.

3.2.2 Finite Element Method

In the finite element method, the domain of interest is divided into a number of simply-connected separate elements, or mesh³⁶. At the vertices of each element we would like a good approximation to the exact solution, then across the element we can interpolate the solution from the values at the vertices.




Figure 3.12 - The subdomain of interest is divided into simply-connected triangular elements, here in 2D.

The approximate solution to our linear differential equation, u' , could then be expressed as the ‘stitched together’ approximations from the separate elements within the mesh.

$$u'(x) = \sum_{i=1}^N u_i' N_i(x) \tag{3.2.40}$$

where there are N vertices in the mesh, u_i' are the values of u' at the vertices of the mesh elements and N_i is some type of interpolation function local to each vertex. If we only consider linear interpolation in two dimensions (for now) we can use the 2D pyramid functions $N_i(x,y)$. These have the property of being one at x_i , falling linearly to zero at the neighbouring nodes and being zero everywhere else in the domain³⁶.




Figure 3.13- A 2D pyramid function has the value one at a certain node and zero at all other nodes within the mesh.

Comparing [3.2.40] to [3.2.7] we can see that the unknown values of u_i' at the mesh vertices are the unknown constants a_i in the weighted residual method and $N_i(x,y)$ is the trial function $\phi_i(x,y)$. The Galerkin matrix therefore has the form:

$$\begin{pmatrix} \int_V N_1[LN_1] & \int_V N_1[LN_2] & \dots & \int_V N_1[LN_N] \\ \int_V N_2[LN_1] & \int_V N_2[LN_2] & \dots & \int_V N_2[LN_N] \\ \dots & \dots & \dots & \dots \\ \int_V N_N[LN_1] & \int_V N_N[LN_2] & \dots & \int_V N_N[LN_N] \end{pmatrix} \begin{pmatrix} u_1 \\ u_2 \\ \dots \\ u_N \end{pmatrix} = \begin{pmatrix} \int_V N_1 f(x,y) \\ \int_V N_2 f(x,y) \\ \dots \\ \int_V N_N f(x,y) \end{pmatrix} \tag{3.2.41}$$

This is a matrix equation of the form $\mathbf{Ax} = \mathbf{b}$; the matrix \mathbf{A} is known in finite element nomenclature as the *stiffness matrix*, the vector \mathbf{x} the *solution vector* and the vector \mathbf{b} as the *forcing vector*. Each of the integrals in [3.2.41] are taken over the entire domain. However, each of the pyramid functions is non-zero only over a limited

region within the domain. This means that *only* terms where the pyramid functions overlap contribute to the matrix [3.2.41]. Due to this, the matrix is a *sparse* matrix (which will be important later when we choose by which method it will be solved) with terms clustered around the diagonal. Conversely, this means that we can build our matrix up step-by-step from the overlapping pyramid functions. This is straightforward as we know that, for example in 2D, each triangular element has three overlapping pyramid functions (that is, one for each vertex of the triangular element) and we can therefore build up our matrix element-by-element.




Figure 3.14 – An individual pyramid function(left), and one component shape function (right).

The form of N_i within each element is a section of the pyramid function; in the 2D case this is a function defined over a triangle equal to zero at two vertices of the triangle and unity at the third. These special functions are known as linear shape functions, S , and depend solely on the geometry of the triangle. In 2D there are three shape functions, $S(x,y)$, for each element (one associated with each of the three vertices of the triangle); these allow us to express $u'(x,y)$ within the n^{th} element as:

$$u'_n(x,y) = u_i^1 S_n^1(x,y) + u_j^2 S_n^2(x,y) + u_k^3 S_n^3(x,y) \quad [3.2.42]$$

Some care has to be taken with the numbering system and definitions; in 2D the vertices have the global numbers i, j and k and the local numbers 1, 2 and 3. Superscripts refer to the local vertex number (of which, in 2D, there are 3) and the subscript refers to the global number (of which there are N in a mesh with N vertices); a subscript applied to the functions $u'(x,y)$ or $S(x,y)$ refers to the global element number whereas a subscript applied to the constants u refers to the global vertex number. Each element has three vertex values of u' and three local shape functions.

The number of shape functions that constitute the pyramid function is the number of elements that surround the pyramid function's central vertex. The exact forms of the pyramid functions therefore depend on the way the vertices are related to their surrounding elements within the mesh.




Figure 3.15 – A sample mesh showing shape functions for the pyramid function of vertex 1.

For example, the 1st vertex in figure 3.15 (which is considered a boundary vertex) has a pyramid function given by:

$$N_1 = S_2^1(x, y) + S_1^2(x, y) \quad [3.2.43]$$

That is, the 1st shape function from the 2nd element plus the 2nd shape function from the 1st element. The pyramid function for the 7th node will have six terms, each being a shape function from one of the six triangular elements surrounding it.

We can therefore build up the matrix from the shape functions within each element. As the pyramid function N_i is the sum of the shape functions incident on the i^{th} vertex, we can go through the mesh element by element adding the overlapping shape functions within each element to the global matrix. Each element will add the contribution:

$$\begin{pmatrix} \int_{\Omega} S_i[LS_i] & \dots & \int_{\Omega} S_i[LS_j] & \dots & \int_{\Omega} S_i[LS_k] \\ \dots & \dots & \dots & \dots & \dots \\ \int_{\Omega} S_j[LS_i] & \dots & \int_{\Omega} S_j[LS_j] & \dots & \int_{\Omega} S_j[LS_k] \\ \dots & \dots & \dots & \dots & \dots \\ \int_{\Omega} S_k[LS_i] & \dots & \int_{\Omega} S_k[LS_j] & \dots & \int_{\Omega} S_k[LS_k] \end{pmatrix} \begin{pmatrix} u_i \\ \dots \\ u_j \\ \dots \\ u_k \end{pmatrix} = \begin{pmatrix} \int_{\Omega} S_i f(x) \\ \dots \\ \int_{\Omega} S_j f(x) \\ \dots \\ \int_{\Omega} S_k f(x) \end{pmatrix} \quad [3.2.44]$$

to the global matrix [3.2.41], Where Ω is the element volume and S is the shape function associated with the vertices of the element. The global vertex numbers are given by i , j and k . The stiffness matrix is the sum of these contributions from the elements within the mesh. The global vector on the right-hand side (known as the forcing vector) is also built up in a similar way by summing the contributions from each element. The matrix [3.2.44] is known as the local element matrix.

If the shape functions are linear, [3.2.41] can take a particularly simple form if $S_i[LS_j]$ is constant within the element (this occurs for the Poisson equation as shown in a later section). The integrals then simply reduce to the area of each element and each of the terms in [3.2.44] is constant and *only* depends on the geometry of the mesh. Each of the terms for each element can then be pre-calculated and stored.

Shape Functions

In 2D the simplest elements are triangles; the linear 2D triangular shape functions have the form^{37 38}:

$$S^1(x, y) = \frac{1}{2\Delta} \begin{vmatrix} x & x^2 & x^3 \\ y & y^2 & y^3 \\ 1 & 1 & 1 \end{vmatrix} \tag{3.2.45}$$

$$S^2(x, y) = \frac{1}{2\Delta} \begin{vmatrix} x & x^3 & x^1 \\ y & y^3 & y^1 \\ 1 & 1 & 1 \end{vmatrix}, S^3(x, y) = \frac{1}{2\Delta} \begin{vmatrix} x & x^1 & x^2 \\ y & y^1 & y^2 \\ 1 & 1 & 1 \end{vmatrix}$$

Where the superscripts refer to the local vertex numbers and Δ is the area of the triangle, given by:

$$\Delta = \frac{1}{2} \begin{vmatrix} x^1 & x^2 & x^3 \\ y^1 & y^2 & y^3 \\ 1 & 1 & 1 \end{vmatrix} \tag{3.2.46}$$

S is defined as zero outside the element.

In 3D, the simplest elements are tetrahedra. The expressions above can be extended to cover the 3D case, and linear 3D tetrahedral shape functions have the form³⁸:

$$S^1(x, y, z) = \frac{1}{6\Delta} \begin{vmatrix} x & x^2 & x^3 & x^4 \\ y & y^2 & y^3 & y^4 \\ z & z^2 & z^3 & z^4 \\ 1 & 1 & 1 & 1 \end{vmatrix}, S^2(x, y, z) = \frac{1}{6\Delta} \begin{vmatrix} x & x^3 & x^4 & x^1 \\ y & y^3 & y^4 & y^1 \\ z & z^3 & z^4 & z^1 \\ 1 & 1 & 1 & 1 \end{vmatrix} \tag{3.2.47}$$

$$S^3(x, y, z) = \frac{1}{6\Delta} \begin{vmatrix} x & x^4 & x^1 & x^2 \\ y & y^4 & y^1 & y^2 \\ z & z^4 & z^1 & z^2 \\ 1 & 1 & 1 & 1 \end{vmatrix}, S^4(x, y, z) = \frac{1}{6\Delta} \begin{vmatrix} x & x^1 & x^2 & x^3 \\ y & y^1 & y^2 & y^3 \\ z & z^1 & z^2 & z^3 \\ 1 & 1 & 1 & 1 \end{vmatrix}$$

Where the superscripts refer to the local vertex numbers and Δ is the area of the triangle, given by:

$$\Delta = \frac{1}{6} \begin{vmatrix} x^1 & x^2 & x^3 & x^4 \\ y^1 & y^2 & y^3 & y^4 \\ z^1 & z^2 & z^3 & z^4 \\ 1 & 1 & 1 & 1 \end{vmatrix} \tag{3.2.48}$$

3.2.3 Application to Magnetostatic Systems

Poisson Equation

The linear Poisson equation is given by:

$$k\nabla^2 u + q = 0 \quad [3.2.49]$$

The boundary S is divided such that $S = S_D + S_N$, where on boundary S_D the Dirichlet boundary condition applies with u_S given:

$$u = u_S \quad [3.2.50]$$

On the boundary S_N the Neumann boundary condition applies with u_N given:

$$k \frac{\partial u}{\partial n} = u_N \quad [3.2.51]$$

Taking the residual expression [3.2.6] (which includes the boundary conditions) and using the approximation [3.2.7] gives:

$$\int_V [Lu' - f]v dV + \oint_S [Mu' - g]w dS = 0 \quad [3.2.52]$$

Using the Poisson operators for the volume and surface from [3.2.49] \rightarrow [3.2.51] and splitting the boundary contribution gives:

$$\int_V [k\nabla^2 u']v dV + \int_V qv dV + \oint_{S_N} \left[k \frac{\partial u'}{\partial n} - u_N \right] w dS + \oint_{S_D} [u' - u_S] w dS = 0 \quad [3.2.53]$$

If we use the finite element method we can immediately satisfy the Dirichlet condition [3.2.50] by specifying the values of u on the boundary vertices *directly*³⁴. This means that the second (Dirichlet) surface term in [3.2.53] will drop out, as the minimization of this residual over that part of the boundary is exactly satisfied. This gives:

$$\int_V [k\nabla^2 u']v dV + \int_V qv dV + \oint_{S_N} \left[k \frac{\partial u'}{\partial n} - u_N \right] w dS = 0 \quad [3.2.54]$$

We can convert the first term using the Green identity:

$$\int_V (\nabla v \cdot \nabla u) dV = \int_V (\nabla v \cdot \nabla u) dV + \oint_S v \frac{\partial u}{\partial n} dS \quad [3.2.55]$$

Where \hat{n} is the outward normal from the boundary surface. This gives:

$$k \int_V (\nabla v \cdot \nabla u') dV + \oint_S k \frac{\partial u'}{\partial n} v dS + \int_V qv dV + \oint_{S_N} k \frac{\partial u'}{\partial n} w dS - \oint_{S_N} u_N w dS = 0 \quad [3.2.56]$$

Since our weighting function w is arbitrary, we can choose it to be $-v$, which removes all the normal terms in the S_N part of the boundary:

$$k \int_V (\nabla v \cdot \nabla u') dV + \oint_{S_D} k \frac{\partial u'}{\partial n} v dS + \int_V q v dV + \oint_{S_N} u_N v dS = 0 \quad [3.2.57]$$

If we now restrict our choice of the weighting function such that v is zero at the boundary S_D , the remaining part of the normal term also disappears. This gives:

$$k \int_V (\nabla v \cdot \nabla u') dV + \int_V q v dV + \oint_{S_N} u_N v dS = 0 \quad [3.2.58]$$

If we now use:

$$u' = \sum_{i=1}^N a_i \phi_i(x) \quad [3.2.59]$$

And:

$$v = \sum_{j=1}^N b_j \psi_j(x) \quad [3.2.60]$$

This gives [3.2.57] as:

$$\left[k \int_V (\nabla \psi \cdot \nabla \phi^T) dV \right] \mathbf{a}^T = \int_V \psi q dV + \oint_{S_N} u_N \psi dS \quad [3.2.61]$$

Hence, if we set our Dirichlet boundary points explicitly and ensure that our weighting functions are zero at these points, the contribution from the Neumann part of the boundary takes the form of an extra source term on the side of the forcing vector. In the Galerkin formulation, with trial and weighting functions both equal to \mathbf{N} , this becomes:

$$\left[k \int_V (\nabla \mathbf{N} \cdot \nabla \mathbf{N}^T) dV \right] \mathbf{a}^T = \int_V \mathbf{N} q dV + \oint_{S_N} \mathbf{N} u_N dS \quad [3.2.62]$$

with N_i at the Dirichlet boundary points equal to zero. If we ignore the Neumann term, we are effectively setting u_N to zero over that part of the boundary. However, if we completely specify the boundary conditions using only Dirichlet conditions this term also goes to zero as $S = S_D + S_N$, and $S = S_D$ hence $S_N = 0$. If we specify $u = 0$ over the entire boundary this results in the following system (in component form):

$$\left[k \int_V \left(\frac{\partial N_i}{\partial x} \frac{\partial N_j}{\partial x} + \frac{\partial N_i}{\partial y} \frac{\partial N_j}{\partial y} + \frac{\partial N_i}{\partial z} \frac{\partial N_j}{\partial z} \right) dV \right] a_i = \int_V N_i q_i dV \quad [3.2.63]$$

Magnetostatics

Taking [2.2.16]:

$$\nabla(\nabla \cdot \mathbf{A}) - \nabla^2 \mathbf{A} = \mu_0 (\mathbf{j}_f + \nabla \times \mathbf{M}) \quad [3.2.64]$$

We apply the Coulomb gauge [2.1.14] giving^{39 40}:

$$\frac{1}{\mu_0} \nabla^2 \mathbf{A} + \nabla \times \mathbf{M} + \mathbf{j}_f = 0 \quad [3.2.65]$$

This is a vector Poisson equation with the source terms being the electric free current and the curl of the magnetization. By formulating the equation in this way we are effectively using the current model and reducing the magnetization to current distributions. The free current has been included here, in the final expression it takes the form of an extra term in the forcing vector which can be easily set to zero if required. If we define \mathbf{A} explicitly over all outer boundaries we can use [3.2.62] to give the finite element formulation of [3.2.65] as:

$$\left[\frac{1}{\mu_0} \int_V (\nabla N_i \cdot \nabla N_j) dV \right] A_{xi} = - \int_V N_i ([\nabla \times \mathbf{M}]_x + j_{xf}) dV \quad [3.2.66]$$

With a similar expression for A_y and A_z . The magnetization term in the forcing vector can be expressed as:

$$\int_V N_i [\nabla \times \mathbf{M}]_x dV = \int_V [\nabla \times (N_i \mathbf{M})]_x dV + \int_V [\mathbf{M} \times (\nabla N_i)]_x dV \quad [3.2.67]$$

And using the divergence theorem:

$$\int_V [\nabla \times (N_i \mathbf{M})]_x dV = \int_S [N_i (\hat{n} \times \mathbf{M})]_x dS \quad [3.2.68]$$

But we have defined the shape functions for Dirichlet conditions as zero on the boundary, and we have given Dirichlet boundary conditions over the entire boundary. This means that this term will disappear, giving:

$$\left[\frac{1}{\mu_0} \int_V (\nabla N_i \cdot \nabla N_j) dV \right] A_{xi} = - \int_V [\mathbf{M} \times (\nabla N_i)]_x - \int_V N_i j_{xf} dV \quad [3.2.69]$$

Which is the system we will solve. In all magnetostatic designs within this project we define \mathbf{A} to be zero over the entire outer boundary. In component form this gives:

$$\begin{aligned} & \left[\frac{1}{\mu_0} \int_V \left(\frac{\partial N_i}{\partial x} \frac{\partial N_j}{\partial x} + \frac{\partial N_i}{\partial y} \frac{\partial N_j}{\partial y} + \frac{\partial N_i}{\partial z} \frac{\partial N_j}{\partial z} \right) dV \right] A_{xi} \\ & = \int_V \left[M_y \frac{\partial N_i}{\partial z} - M_z \frac{\partial N_i}{\partial y} \right] dV - \int_V N_i j_{xf} dV \end{aligned} \quad [3.2.70]$$

Where \mathbf{M} , the magnetization, and \mathbf{j}_f , the free current, is taken to be fixed within each element. It should be noted that the matrix part of the system, \mathbf{S} (the stiffness matrix), *only depends on the geometry of the mesh* and will not change for A_y and A_z . The full expression for \mathbf{A} then consists of three matrix expressions:

$$\mathbf{S}A_{x_i} = \int_V \left[M_y \frac{\partial N_i}{\partial z} - M_z \frac{\partial N_i}{\partial y} \right] dV - \int_V N_i j_{x_f} dV \quad [3.2.71]$$

$$\mathbf{S}A_{y_i} = \int_V \left[M_x \frac{\partial N_i}{\partial z} - M_z \frac{\partial N_i}{\partial x} \right] dV - \int_V N_i j_{y_f} dV \quad [3.2.72]$$

$$\mathbf{S}A_{z_i} = \int_V \left[M_y \frac{\partial N_i}{\partial x} - M_x \frac{\partial N_i}{\partial y} \right] dV - \int_V N_i j_{z_f} dV \quad [3.2.73]$$

These three matrix expressions must be evaluated and \mathbf{B} can be determined on an element-wise basis by taking the curl of \mathbf{A} within the element. From [3.2.42] \mathbf{A} within the n^{th} element is given (in 3D, with global vertex numbers i, j, k and l) by:

$$\begin{aligned} A_{xn} &= A_{xi}^1 S_n^1 + A_{xj}^2 S_n^2 + A_{xk}^3 S_n^3 + A_{xl}^4 S_n^4 \\ A_{yn} &= A_{yi}^1 S_n^1 + A_{yj}^2 S_n^2 + A_{yk}^3 S_n^3 + A_{yl}^4 S_n^4 \\ A_{zn} &= A_{zi}^1 S_n^1 + A_{zj}^2 S_n^2 + A_{zk}^3 S_n^3 + A_{zl}^4 S_n^4 \end{aligned} \quad [3.2.74]$$

Where S_n^i is the shape functions for vertex i of element n . Then, for example, B_x is given by:

$$B_{xn} = \begin{pmatrix} A_{zi}^1 \frac{\partial S_n^1}{\partial y} + A_{zj}^2 \frac{\partial S_n^2}{\partial y} + A_{zk}^3 \frac{\partial S_n^3}{\partial y} + A_{zl}^4 \frac{\partial S_n^4}{\partial y} \\ A_{yi}^1 \frac{\partial S_n^1}{\partial z} + A_{yj}^2 \frac{\partial S_n^2}{\partial z} + A_{yk}^3 \frac{\partial S_n^3}{\partial z} + A_{yl}^4 \frac{\partial S_n^4}{\partial z} \end{pmatrix} - \quad [3.2.75]$$

With similar expressions for B_y and B_z .

Soft Materials in Magnetostatic Systems

The main advantage of the finite element method over the charge model is the ability to handle soft materials within the magnetic design. Unlike the charge simplexes used in the charge model calculation analytic closed-form solutions for designs incorporating soft material can only be derived for extremely simple shapes. The difficulty lies in the fact that each infinitesimal element of soft material in the design magnetizes and *becomes* magnetized by every other element of soft material within the design. However, if we assume that the magnetization of the soft material within a given small area is constant the finite element method can provide a natural solution to the problem by *iteratively* computing the effects of the soft iron parts.

We assume that the magnetization within each of the tetrahedral elements that compose the soft iron parts of our design is constant at every iteration. We then update these magnetization vectors at each iteration by using the linear **B-H** relationship [2.1.47]:

$$\mathbf{B} = \mu_0 \mu_r \mathbf{H} \quad [2.1.47]$$

Where the relative permeability μ_r is a constant and taken to be, for example, around 1000 in soft iron. Substituting into [2.1.4]:

$$\mathbf{B} = \mu_0 (\mathbf{H} + \mathbf{M}) \quad [2.1.4]$$

We find that the magnetization **M** can be expressed in terms of the magnetic field **B**:

$$\mathbf{M} = \frac{1}{\mu_0} \left(1 - \frac{1}{\mu_r} \right) \mathbf{B} \quad [3.2.76]$$

We can use this relationship to define our iterative scheme:

$$\mathbf{M}_{i+1} = \frac{1}{\mu_0} \left(1 - \frac{1}{\mu_r} \right) \mathbf{B}_i \quad [3.2.77]$$

Where \mathbf{M}_{i+1} is the new magnetization of the tetrahedral element of soft material and \mathbf{B}_i is the magnetic induction inside the tetrahedral element. This can be thought of in physical terms as solving over time; the applied field causes a resultant magnetization in each element of the soft material at each time step. This resultant magnetization causes a different resultant magnetization at the next time step, and so on until the magnetization of the entire system has stabilised. Another method which takes into account the *reluctivity* (the reciprocal of the permeability) of the material is given by Sykulski⁴¹. However, this method is much harder to implement due to the formulation requiring an integral to be performed for each entry in the stiffness matrix.

The iterative scheme is implemented in the project by repeatedly calculating the magnetization of the design and updating the magnetization of the elements of soft material until the maximum change in magnetization in any element is less than some (small) amount. At this point the system is assumed to have converged. However, simply updating the magnetization in each of the elements according to [3.2.77] at each step does not lead to the physically correct solution, as the magnetization within each segment is unbounded and can keep growing without limits. A very simple remedy to this is to *cap* the value of the magnetization, so that [3.2.77] becomes:

$$\mathbf{M}_{i+1} = -\mathbf{M}_s, \mathbf{M}_{i+1} < -\mathbf{M}_s \quad [3.2.78]$$

$$\mathbf{M}_{i+1} = \frac{1}{\mu_0} \left(1 - \frac{1}{\mu} \right) \mathbf{B}_{i'} - \mathbf{M}_s < \mathbf{M}_{i+1} < \mathbf{M}_s$$

$$\mathbf{M}_{i+1} = \mathbf{M}_s, \mathbf{M}_{i+1} > \mathbf{M}_s$$

Where \mathbf{M}_s is the saturation value of the material. In soft iron this is around $2.15/\mu_0$ A/m. This can be thought of as the simplest approximation possible to the initial magnetization curve (curve A to B in figure 1.3) of a soft material (see figure 3.16). This method is very easy to implement, if the new value of the magnetization is above or below the saturation value the magnetization is simply set to the saturation value. More complex approximations to the magnetization curve are possible, for example, by using a table of magnetization values and linearly interpolating between them⁴⁰.




Figure 3.16 - A 'capped' ferromagnetic hysteresis curve takes into account the saturation of the material.

Unfortunately, convergence of the above scheme is fairly slow. It is believed that the convergence speed could be increased using a minimization technique such as the Newton-Raphson method, although this was not attempted within the project⁴².

3.2.4 Solution of the Matrix System

The solution of large matrix systems is fundamental in computational science; using suitable discretization, many physical systems can be reduced to expressions of the following form:

$$\mathbf{Ax} = \mathbf{b} \quad [3.2.79]$$

where \mathbf{A} is the stiffness matrix, \mathbf{x} is the solution vector and \mathbf{b} the forcing vector. To solve for \mathbf{x} , one can either find \mathbf{x} using some iterative scheme, or directly invert \mathbf{A} .

If we split our matrix \mathbf{A} in some way (for example, into diagonal and off-diagonal elements as in the Jacobi method, see the section on stationary iterative methods below), we can generally express [3.2.79] as:

$$\mathbf{x} = \mathbf{B}\mathbf{x} + \mathbf{c} \quad [3.2.80]$$

where \mathbf{B} is some combination of the split parts of \mathbf{A} , and \mathbf{c} is some vector combination of \mathbf{b} and the split parts of \mathbf{A} . We can then form an iterative method⁴³:

$$\mathbf{x}_{i+1} = \mathbf{B}\mathbf{x}_i + \mathbf{c} \quad [3.2.81]$$

Iterative methods are divided into stationary and non-stationary methods. Stationary methods have constant \mathbf{B} and \mathbf{c} for each i . We can define an error \mathbf{e}_i as⁴⁴:

$$\mathbf{e}_i = \mathbf{x}_i - \mathbf{x} \quad [3.2.82]$$

which is a vector that points from our current iteration of \mathbf{x} , \mathbf{x}_i , to the true solution \mathbf{x} . We can then express the error in terms of \mathbf{B} :

$$\mathbf{e}_{i+1} = \mathbf{B}\mathbf{e}_i \quad [3.2.83]$$

If \mathbf{B} is symmetric, a set of n linearly independent eigenvectors exist for \mathbf{B} . Since this set of eigenvectors form a basis for \mathfrak{R}^n , the error can be expressed in terms of the eigenvectors of \mathbf{B} . If all the corresponding eigenvalues of \mathbf{B} are *all* less than one, all the components of \mathbf{e} will eventually decrease to zero (if \mathbf{B} is non-symmetric and defective, this rule still holds, but for *generalised* eigenvectors). If we define the spectral radius of \mathbf{B} , $\rho(\mathbf{B})$ as:

$$\rho(\mathbf{B}) = \max|\lambda_i| \quad [3.2.84]$$

where λ_i are the set of eigenvalues for \mathbf{B} , it can be seen that the iterative scheme will only converge if $\rho(\mathbf{B}) < 1$.

The stiffness matrices arising from our particular formulation of the finite element method (and indeed many forms of the method) give rise to a *symmetric positive definite* matrix \mathbf{A} . The definition of positive definiteness for some matrix \mathbf{A} is:

$$\mathbf{x}^T \mathbf{A}\mathbf{x} > 0 \quad [3.2.85]$$

for any non-zero vector \mathbf{x} . This means that the quadratic form of the matrix has some minimum \mathbf{x} that satisfies [3.2.79] where the quadratic form is given by⁴⁵:

$$f(x) = \frac{1}{2} \mathbf{x}^T \mathbf{A}\mathbf{x} - \mathbf{b}^T \mathbf{x} + c \quad [3.2.86]$$

where \mathbf{A} is a matrix, \mathbf{b} is a vector and c a scalar constant. The gradient of this (for symmetric \mathbf{A}) is:

$$f'(x) = \mathbf{A}\mathbf{x} - \mathbf{b} \quad [3.2.87]$$

Hence the minimum of the quadratic form is the solution to [3.2.79]. We use the positive-definiteness of our matrix to choose the best solution method for the matrix equation.

Direct Methods

Direct methods involve solving for \mathbf{x} directly; this can be done by inverting \mathbf{A} but a number of other much more computationally efficient methods exist. The simplest method is Gaussian elimination with backsubstitution⁴⁶. The Gaussian stage of this procedure results in a triangular matrix, and \mathbf{x} can be easily calculated from this by direct evaluation of each coefficient x_i (the backsubstitution stage).

$$\begin{pmatrix} a_{11} & a_{12} & \dots & a_{1N} \\ 0 & a_{22} & \dots & a_{2N} \\ 0 & 0 & \dots & \dots \\ 0 & 0 & 0 & a_{NN} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ \dots \\ x_N \end{pmatrix} = \begin{pmatrix} b_1 \\ b_2 \\ \dots \\ b_N \end{pmatrix} \quad [3.2.88]$$

A more efficient means of solving [3.2.79] is LU-decomposition (lower-upper decomposition). We can express \mathbf{A} in two parts such that $\mathbf{A} = \mathbf{LU}$ where \mathbf{L} and \mathbf{U} are upper and lower triangular matrices. This gives:

$$\mathbf{L}(\mathbf{U}\mathbf{x}) = \mathbf{b} \quad [3.2.89]$$

As the solution of a triangular matrix is straightforward (as we know from Gaussian elimination, above) we can first solve by putting $\mathbf{U}\mathbf{x} = \mathbf{y}$ (say), solving for $\mathbf{L}\mathbf{y} = \mathbf{b}$ and then solving again for $\mathbf{U}\mathbf{x} = \mathbf{y}$. LU decomposition is achieved using a very efficient algorithm called Crout's algorithm. According to Numerical Recipes⁴⁴, LU-decomposition is a factor of 3 faster than Gaussian elimination with backsubstitution which in turn is a factor of 1.5 times better than straight Gauss-Jordan elimination.

If the matrix is symmetric and positive definite, however, an even more efficient LU decomposition method exists called Cholesky factorization⁴⁷. This simply expresses the lower and upper parts as the transpose matrix of a single lower (or upper part), for example $\mathbf{A} = \mathbf{LL}^T$ where \mathbf{L} is the lower triangular matrix.

$$\mathbf{L}(\mathbf{L}^T \mathbf{x}) = \mathbf{b} \quad [3.2.90]$$

This can be solved more efficiently than standard LU decomposition due to the symmetry of the matrix; Cholesky factorization is about a factor of two faster than LU-decomposition.

We use Cholesky factorization in this project for two purposes; the first is for a direct solution of the finite element matrix as it can sometimes (depending on matrix size

and problem) be faster when the stiffness matrix has to be repeatedly re-applied to different forcing vectors (which happens when soft materials are included in the design). Secondly, *incomplete Cholesky factorization* is used as a preconditioner for the conjugate gradient method as the factorization process will, in general, assign non-zero values to \mathbf{L} during the factorization process when the corresponding element in \mathbf{A} is zero. These non-zero entries are known as the *fill entries* and the idea behind incomplete factorization is to drop some of these fill entries according to some strategy⁴⁸. The simplest strategy of *all* is to drop all fill entries of \mathbf{L} in which the corresponding entry of \mathbf{A} is zero; this saves considerable computation time (as we know which entries of \mathbf{A} we have and therefore do not need to calculate the entire \mathbf{L} -matrix, furthermore if we overwrite \mathbf{A} we do not need to assign extra space to store elements). Other drop strategies include keeping the number of row or column entries the same but re-assigning them so that they are the largest values of all the row or column values (including the newly-completed entries from the factorization), or keeping a set number of new entries and dropping the rest. Preconditioners will be discussed in a following section on the conjugate gradient method.

Stationary Iterative Methods

The main stationary methods are the Jacobi method, the Gauss-Seidel method and Successive-Over-Relaxation (SOR). The Jacobi method solves for each \mathbf{x}_i while every other \mathbf{x}_i remains fixed throughout the process and assumes that the system will converge to the correct \mathbf{x} after a certain number of iterations. It can also be thought of as splitting the matrix \mathbf{A} into two parts; the matrix \mathbf{D} , the diagonal elements of \mathbf{A} , and \mathbf{E} , all other off-diagonal elements of \mathbf{A} ⁴⁹. This gives:

$$\mathbf{B} = -\mathbf{D}^{-1}\mathbf{E}, \mathbf{c} = \mathbf{D}^{-1}\mathbf{b} \quad [3.2.91]$$

This method is simple to implement as finding the inverse of a diagonal matrix is straightforward. However, the convergence rate is slower than other stationary methods⁵⁰ (if the method converges at all).

The Gauss-Seidel method improves on the Jacobi method by solving for \mathbf{x}_i using the results from the values of \mathbf{x} up to i already computed by the algorithm⁴⁷. It can be thought of as splitting the matrix into three parts; the matrix of diagonal elements of \mathbf{A} , \mathbf{D} , the matrix of the strictly upper triangular part of \mathbf{A} , $-\mathbf{U}$, and the strictly lower triangular part of \mathbf{A} , $-\mathbf{L}$. This gives:

$$\mathbf{B} = (\mathbf{D} - \mathbf{L})^{-1} \mathbf{U}, \mathbf{c} = (\mathbf{D} - \mathbf{L})^{-1} \mathbf{b} \quad [3.2.92]$$

This process of calculating \mathbf{x}_{i+1} can be thought of as adding an increment to \mathbf{x}_i of $(\mathbf{x}_{i+1} - \mathbf{x}_i)$. If each increment is too small we can speed up the convergence of the Gauss-Seidel method by multiplying the increment by an extrapolation factor ω ; effectively a weighted average between the last and current iteration. This is known as the successive-over-relaxation method:

$$\mathbf{B} = (\mathbf{D} - \omega\mathbf{L})^{-1}[\omega\mathbf{U} + (1 - \omega)\mathbf{D}], \mathbf{c} = \omega(\mathbf{D} - \omega\mathbf{L})^{-1} \mathbf{b} \quad [3.2.93]$$

The value of ω must fall between $\{0,2\}$ and, in theory, the optimal ω can be calculated using the spectral radius of the matrix. However, determining the spectral radius requires the calculation of all eigenvalues for the matrix and therefore possibly more computation effort than solving the entire matrix system itself. Heuristic approximations are therefore often used⁴⁷.

Non-Stationary Iterative Methods

The main non-stationary method is the conjugate gradient method which can be introduced via a related method, the steepest descent method⁴⁵. The steepest descent method is very intuitive; if we are looking for an \mathbf{x} that minimizes the quadratic form of [3.2.79] and this form has a global minimum, we can just take a step in the direction in which by $f(\mathbf{x})$ decreases by the greatest amount (this being the direction of steepest descent). We repeatedly take steps until the error (which represents distance between our current position and the true solution) becomes as small as we require. Of course, the step size is not constant for each step and the method of computing the step size is to find the point at which $f(\mathbf{x})$ is minimized along our path.

If we define our residual as the direction of steepest descent:

$$\mathbf{r}_i = -f'(\mathbf{x}_i) \quad [3.2.94]$$

Note that this means the residual is also given by $-(\mathbf{Ax}-\mathbf{b})$ from [3.2.87]. Each step is then given by:

$$\mathbf{x}_{i+1} = \mathbf{x}_i + \alpha_i \mathbf{r}_i \quad [3.2.95]$$

The parameter α_i is minimised when the residual is orthogonal to the previous residual (this can be shown by minimising $f\{\mathbf{x}_{i+1}\}$ with respect to α_i). The parameter is given by⁴⁵:

$$\alpha_i = \frac{\mathbf{r}_i^T \mathbf{r}_i}{\mathbf{r}_i^T \mathbf{A} \mathbf{r}_i} \quad [3.2.96]$$

Where:

$$\mathbf{r}_i = \mathbf{b} - \mathbf{A}\mathbf{x}_i \quad [3.2.97]$$

A recurrence relation for the residual can be found by multiplying [3.2.95] by $-\mathbf{A}$ and adding \mathbf{b} . This gives:

$$\mathbf{r}_{i+1} = \mathbf{r}_i - \alpha_i \mathbf{A}\mathbf{r}_i \quad [3.2.98]$$

Unfortunately the steepest descent algorithm is usually inefficient due to the fact that each step is orthogonal to the previous step. In the worst case a bad starting point means the algorithm can very slowly converge as it zigzags in on the solution.




Figure 3.17 - In the method of steepest descent the starting position strongly affects the convergence time.

A much better solution would be if we only moved in one of these orthogonal directions *once*. If we picked exactly the right step size each time we would only then need N steps in N dimensional space. Unfortunately as, Shewchuk points out⁴⁵, if we use orthogonal directions we have to know the error first to work out the step size, and if we knew the error the problem would already be solved. The solution is to take an orthogonal step in \mathbf{A} -space. Two vectors are \mathbf{A} -orthogonal if:

$$\mathbf{v}_i^T \mathbf{A}\mathbf{v}_j = 0 \quad [3.2.99]$$

We can then calculate the step size α_i :

$$\alpha_i = \frac{\mathbf{d}_i^T \mathbf{r}_i}{\mathbf{d}_i^T \mathbf{A}\mathbf{d}_i} \quad [3.2.100]$$

Where \mathbf{d}_i are our \mathbf{A} -orthogonal search directions. This is the same as steepest descent if we use the residuals \mathbf{r}_i as the search directions. We can construct a set of \mathbf{A} -orthogonal search directions using a process called Gram-Schmidt orthogonalization⁴⁵. This takes a set of N -dimensional linearly independent trial vectors and creates a N -dimensional orthogonal basis from them by starting with the first and subsequently subtracting all components from the following vectors that have already been covered by the process. For example, it uses the first trial vector as the

first basis vector. It then removes all components in the direction of the first trial vector from second trial vector, which then becomes the second basis vector. It then removes all components of the first and second trial vectors from the third trial vector, which becomes the third basis vector and so on until no trial vectors are left. This procedure is useful in 3D for creating a local co-ordinate system on some surface defined by two vectors (the resulting basis vectors must be normalised here). For example it is used in this project in to work out the perpendicular height of a point above the plane defined by two vectors, and to calculate whether a point falls inside a 2D contour on a plane defined by 3D vectors.

We can express our basis as:

$$\mathbf{d}_i = \mathbf{v}_i + \sum_{k=0}^{i-1} \beta_{ik} \mathbf{d}_k \quad [3.2.101]$$

Where \mathbf{d} are the basis vectors and \mathbf{v} are our trial vectors. The constant β_{ik} is given by:

$$\beta_{ik} = -\frac{\mathbf{v}_i^T \mathbf{A} \mathbf{d}_j}{\mathbf{d}_j^T \mathbf{A} \mathbf{d}_j}, i > k \quad [3.2.102]$$

Unfortunately, all the old vectors must all be kept in memory and used to calculate the new vectors which makes calculating a full orthogonal basis set of vectors a very computationally expensive task. However, it can be shown that the residuals have the very useful property:

$$\mathbf{d}_i^T \mathbf{r}_j = 0, i < j \quad [3.2.103]$$

That is, each residual is orthogonal to all prior search directions. This means that if we use the residuals as the trial vectors (i.e. $\mathbf{v}_i = \mathbf{r}_i$) we no longer have to subtract all non-orthogonal parts from the residuals to form each basis vector; all non-orthogonal parts are *already* zero and the only non-zero part is that of the current residual. Herein lies the power of the conjugate gradient method; the fact that we no longer have to process the entire set of basis vectors reduces the algorithm from $O(N^2)$ complexity (for a full Gram-Schmidt process in N dimensions) to $O(m)$ complexity, where m is the number of non-zero elements of \mathbf{A} ⁴⁵ Conjugate gradients are therefore especially useful for positive-definite sparse matrices, such as stiffness matrices in the finite element method.

Taking the vector product of the residual and the recurrence relation for the residual [3.2.98] and re-arranging we have:

$$\mathbf{r}_i^T \mathbf{A} \mathbf{d}_j = \frac{1}{\alpha_j} (\mathbf{r}_i^T \mathbf{r}_j - \mathbf{r}_i^T \mathbf{r}_{j+1}) \quad [3.2.104]$$

Since our search direction are built from the residuals and each residual is orthogonal to all prior search directions, each residual is also orthogonal to each prior residual:

$$\mathbf{r}_i^T \mathbf{r}_j = 0, \quad i \neq j \quad [3.2.105]$$

This means only the two terms $i = j$ and $i = (j+1)$ contribute in the expression [3.2.104]. If we use [3.2.104] for β , we need only consider the $i = (j+1)$ term as β is defined *only* for $i > j$. This means that the only non-zero term of β is:

$$\beta_{ij} = -\frac{1}{\alpha_{i-1}} \frac{\mathbf{r}_i^T \mathbf{r}_i}{\mathbf{d}_{i-1}^T \mathbf{A} \mathbf{d}_{i-1}}, \quad i = j+1 \quad [3.2.106]$$

Using [3.2.100] for α_{i-1} , simplifies [3.2.106] to:

$$\beta_{ij} = \frac{\mathbf{r}_i^T \mathbf{r}_i}{\mathbf{d}_{i-1}^T \mathbf{r}_{i-1}} \quad [3.2.107]$$

Using the identity:

$$\mathbf{d}_i^T \mathbf{r}_i = \mathbf{v}_i^T \mathbf{r}_i \quad [3.2.108]$$

Where \mathbf{v} is a trial vector (This can be proved by taking the vector product of \mathbf{r}_i and [3.2.101]) and putting $\mathbf{v}_i = \mathbf{r}_i$ gives [3.2.107] as:

$$\beta_i = \frac{\mathbf{r}_i^T \mathbf{r}_i}{\mathbf{r}_{i-1}^T \mathbf{r}_{i-1}} \quad [3.2.109]$$

Where β_i is defined as $\beta_{i,i-1}$ and the entire method is then given by:

$$\alpha_i = \frac{\mathbf{r}_i^T \mathbf{r}_i}{\mathbf{d}_i^T \mathbf{A} \mathbf{d}_i} \quad [3.2.110]$$

$$\mathbf{x}_{i+1} = \mathbf{x}_i + \alpha_i \mathbf{d}_i \quad [3.2.111]$$

$$\mathbf{r}_{i+1} = \mathbf{r}_i - \alpha_i \mathbf{A} \mathbf{d}_i \quad [3.2.112]$$

$$\beta_{i+1} = \frac{\mathbf{r}_{i+1}^T \mathbf{r}_{i+1}}{\mathbf{r}_i^T \mathbf{r}_i} \quad [3.2.113]$$

$$\mathbf{d}_{i+1} = \mathbf{r}_{i+1} - \beta_{i+1} \mathbf{d}_i \quad [3.2.114]$$

The algorithm would terminate (in the ideal case) when the residual becomes zero after N iterations, however, due to floating point roundoff error accruing in the iterative formulation and the fact that N may be extremely large means that it is usually better to define a stopping criteria. This can be done by specifying some

maximum number of iterations and/or allowing the algorithm to terminate when the norm of the residual falls below a certain value or fraction of the initial residual.

The conjugate gradient method was initially used in the project, but the slow speed of convergence led to the use of the *preconditioned* conjugate gradient method.

Preconditioning

The convergence of the conjugate gradient method depends on the clustering of the eigenvalues of the matrix \mathbf{A} . A measure of the clustering of eigenvalues is given by the condition number of a matrix, κ :

$$\kappa = \frac{\lambda_{\max}}{\lambda_{\min}} \quad [3.2.115]$$

Where λ_{\max} and λ_{\min} are the largest and smallest eigenvalues of the matrix. Preconditioning attempts to make the iterative method converge faster by warping the quadratic form into a ‘better’ shape, namely a more spherical one in which the eigenvalues of \mathbf{A} are closer in value to each other, by solving the alternative system:

$$\mathbf{M}^{-1}\mathbf{A}\mathbf{x} = \mathbf{M}^{-1}\mathbf{b} \quad [3.2.116]$$

If κ for the alternate system is lower than κ for the original system, we can solve [3.2.116] faster than for the original problem. The problem is in choosing an \mathbf{M} that we can not only easily invert but which also produces better clustered eigenvalues. Full derivation of the preconditioned conjugate gradient method is given in⁴⁵, and the form is presented below.

$$\alpha_i = \frac{\mathbf{r}_i^T (\mathbf{M}^{-1}\mathbf{r}_i)}{\mathbf{d}_i^T \mathbf{A} \mathbf{d}_i} \quad [3.2.117]$$

$$\mathbf{x}_{i+1} = \mathbf{x}_i + \alpha_i \mathbf{d}_i \quad [3.2.118]$$

$$\mathbf{r}_{i+1} = \mathbf{r}_i - \alpha_i \mathbf{A} \mathbf{d}_i \quad [3.2.119]$$

$$\beta_{i+1} = \frac{\mathbf{r}_{i+1}^T (\mathbf{M}^{-1}\mathbf{r}_{i+1})}{\mathbf{r}_i^T (\mathbf{M}^{-1}\mathbf{r}_i)} \quad [3.2.120]$$

$$\mathbf{d}_{i+1} = (\mathbf{M}^{-1}\mathbf{r}_{i+1}) - \beta_{i+1} \mathbf{d}_i \quad [3.2.121]$$

Several choices of \mathbf{M} exist; the simplest is a diagonal or Jacobi preconditioner which is just the matrix of diagonal elements of \mathbf{A} . This preconditioner provides a reasonable improvement, is trivial to invert (the inverse of a diagonal matrix is the reciprocal of the diagonal elements) and unconditionally stable.

We have already mentioned the other most common type of preconditioner, the incomplete Cholesky preconditioner. It should be noted that the process of generating the incomplete conditioner is not, however, always stable. In our application the process reverts to the diagonal conditioner if the Cholesky factorization fails (which we found occurs very rarely). The matrix solution in the project used two drop-strategies; the first was a standard ‘no-fill’ strategy in which all the entries in \mathbf{L} which corresponded to a zero element in \mathbf{A} were dropped, which is the simplest drop strategy. The second was an ‘improved’ factorization proposed by Jones and Plassmann⁵¹ in which the number of non-zero elements per column of \mathbf{L} was the same as the corresponding column of \mathbf{A} and the values kept in each column was chosen from the highest values arising from the factorization. However, we saw no great improvement using the second method and, in fact, appeared to be more unstable than the first.




Figure 3.18 - A comparison of solution methods for matrix equation $Ax = b$ with different preconditioners.

Figure 3.18 shows a comparison of four solution methods for a magnetostatic system (meshed at different density levels) in which [3.2.79] was repeatedly solved for the same stiffness matrix \mathbf{A} , but different forcing vectors \mathbf{b} at each iteration (see section 3.2.3, on the inclusion of soft materials). The four methods were the preconditioned conjugate gradient with a diagonal preconditioner (PCG/D), preconditioned conjugate

gradient with an incomplete Cholesky preconditioner (PCG/C), a direct solution using Cholesky factorization (DC) and an un-preconditioned conjugate gradient method (CG). In systems such as this the cost of generating the preconditioner is very small in comparison to the large benefits in terms of solution time.

It can be seen that even the simple diagonal preconditioner (PCG/D) performs much better than the un-preconditioned conjugate gradient method (CG). A direct solution of the problem (DC) is almost as efficient time-wise as the incomplete Cholesky preconditioner (PCG/C), although at high N it can be seen that the incomplete Cholesky preconditioner becomes the more efficient method. We therefore chose the incomplete Cholesky preconditioner as our default.

3.3 Automated Mesh Generation

3.3.1 Overview

In order to use the finite element method, we need to break the problem domain into a large number of smaller elements, often the tetrahedron (being the simplex in 3D) or, for example, quadrilaterals or hexahedral elements; this subdivided domain structure is called the mesh. The purpose of meshing is to provide a description of the problem domain in terms of a set of basic blocks; this discretization of the domain can then be used as a basis for calculations involving partial differential equations (PDEs). Here, we apply it to the Finite Element Method (FEM), but it can also be applied to other methods such as the Boundary Element Method (BEM). Mesh generation for PDEs is typically more complicated than mesh generation for other types of applications (for example computer graphics) in that the mesh elements have to be ‘well shaped’. This is due to a number of factors; large angles in an element will cause errors in derivatives across the element as a consequence of interpolation, and small angles can also cause the matrix resulting from the discretization to be ill-conditioned⁵². Therefore, we would like the elements to be as ‘round’ as possible, that is (in the 2D and 3D simplex cases) to have the edges of any element as close in length to each other as possible.

Mesh generation is not straightforward. Approaches are broken into two distinct methods called structured and un-structured mesh generation⁵³. Structured mesh generation is, seemingly, the most straightforward way to generate a mesh; the basic

idea is to take some simple structured mesh pattern and map it using some transformation to fit the design we are trying to model. This simple pattern will have a given element connectivity and is therefore more computationally straightforward to store and process. Frequently, however, the design is too complicated to be represented by a simple mapping and must be broken up into simpler sections, each with a different mapping transformation. The mesh is then constructed by patching these sections together such that the sections share boundary vertices (called a multiblock scheme). Even this method is not always satisfactory, however, and decisions also have to be made as to where and how to join these blocks. Unstructured mesh generation, on the other hand, can handle arbitrarily complex designs and, furthermore, the density of the mesh can also be tailored to fit the problem, meaning regions of interest within the mesh can be refined to any degree desired. It is, however, much more difficult than structured mesh generation. Unstructured mesh generation a vast field and some approaches include:

- Octree methods, in which the domain is recursively subdivided into an octree grid (according to boundary curvature) and mesh vertices are either created or relaxed onto the grid to fit the design boundaries⁵⁴. The octree generation is fairly straightforward, but problems occur for certain geometries and for decisions on whether to insert a vertex or move a grid boundary vertex to fit the boundary. A neat modern approach closely related to this is red-green mesh generation⁵⁵.
- Advancing front methods, in which tetrahedra are crystallised inwards from the boundary⁵³. This approach gives excellent boundary surfaces, as the method can choose to pave the boundary with perfect or near-perfect simplexes. However, problems arise when boundary fronts meet and complex decisions have to be made about vertex insertion between opposing fronts.
- Incremental Delaunay point insertion, in which a Delaunay mesh is built up by inserting vertices to eliminate badly-shaped and/or large triangles. This is the method we will be considering here and forms the basis for Shewchuk's algorithm⁵⁶.

We decided to use adopt the incremental Delaunay approach over the octree method and advancing front method. The octree method proved to be difficult to implement in 2D (as a quadtree method) due to the difficulties involved in resolving boundary/tree interfaces and we expected these difficulties to be exacerbated in 3D (although the building of the tree itself was extremely straightforward). The advancing front method was more difficult than the quadtree method to implement in 2D due to the complex decisions to be made about vertex placement between colliding fronts and, similarly, we expected these difficulties to become worse in 3D. The incremental Delaunay approach, although difficult to implement in 2D, was attractive as many of the algorithms used could be extended to 3D relatively straightforwardly and of the three methods it has the least complex (and most elegant) set of rules for the placement of new vertices.

3.3.2 Delaunay Incremental Insertion Algorithms

A Delaunay triangulation of a point set is a triangulation in which every triangle has no vertex inside the triangle's circumcircle⁵². The circumcircle of a triangle is the unique circle which passes through each of the triangle's three vertices. The Delaunay triangulation is the 'best' triangulation to choose for mesh generation as it minimises the minimum containment circle in all dimensions. In two dimensions only, this triangulation maximises the minimum angle in the triangulation and minimises the largest circumcircle. The minimum containment circle is the smallest circle that contains a triangle, which is not always the triangle's circumcircle (for example, an isosceles triangle with a pair of angles approaching zero has a circumcircle that approaches infinite radius, but we can always find a circle that will contain it based on its maximum edge length).

We can build up a Delaunay triangulation incrementally in the following manner (see Figure 3.19): First we find the *circumball* around the point we wish to insert where the circumball is the set of triangles that would not have an empty circumcircle if the point were inserted (2). Then, we delete these triangles from the mesh creating a cavity (3). Lastly, we re-connect the insertion vertex to the cavity vertices (4).




Figure 3.19 – The stages of an incremental Delaunay triangulation, point location (1) cavity construction (2), cavity deletion (3) and cavity processing (4).

There are some points are worth noting here: The cavity will always be free from other vertices as the triangles are Delaunay and therefore have no vertices within them. Also, all of the newly created triangles are all Delaunay, which is somewhat counter-intuitive, but can be understood by noting that each of the newly created edges from the cavity vertices to the insertion vertex are Delaunay (an edge is classed as Delaunay if any empty circumsphere passing through the two edge vertices can be found). In this case the sphere tangential to the original sphere circumscribing an original vertex and the insertion vertex will always be Delaunay (see figure 3.20, top). If we then ‘grow’ this Delaunay circumsphere in either direction orthogonal to the edge, the first vertex it will come into contact with will be a vertex on the cavity. The triangle formed in this way will be Delaunay and, since this holds for each edge, every new triangle formed by this vertex insertion method will also be Delaunay⁵⁵.




Figure 3.20 - The triangles created by an incremental Delaunay insertion algorithm will all be Delaunay.

The above algorithm describes the Bower Watson^{57 58} algorithm and generalises to any dimension (an algorithm that produces the same result but is based on edge flips is the Lawson algorithm⁵⁹; however, this algorithm, although more robust in certain cases, only works in two dimensions and although there are 3D generalisations of it, we will use this more intuitive and somewhat simpler algorithm).

The advantage of this process is that we are fully in control of where to place the new vertices in the mesh. We can choose to eliminate a ‘bad’ triangle or tetrahedron by placing a new vertex at its circumcentre (The bad triangle or tetrahedron will then be deleted in the cavity creation phase). A bad triangle or tetrahedron can be either one with a sharp angle (so called skinny simplex) or one over a certain area or volume. One of the first 2D refinement algorithms was Chew’s (first) algorithm⁶⁰ which splits any triangle in the mesh with a circumradius greater than the smallest edge length in the mesh by placing a new vertex into the triangle’s circumcentre. This results in an extremely uniform mesh with well shaped tetrahedra, although unfortunately it does not guarantee that boundaries will survive the meshing process.

3.3.3 Geometric Primitives

There are two main geometric primitives we will use within this project, called the *orientation* test and the *incircle* test⁶¹; both are crucial in building up the mesh correctly when using an incremental algorithm.




Figure 3.21 The two basic geometrical tests.

The orientation test determines whether a point lies to the left, right or on a line and can be reduced to a test on a determinant (for the 2D case here):

$$D_O = \begin{vmatrix} a_x & a_y & 1 \\ b_x & b_y & 1 \\ p_x & p_y & 1 \end{vmatrix} \quad [3.3.1]$$

Where a positive value indicates the point lies on one side of the line, a negative value indicates the point lies on the other side of the line and a zero determinant indicates that the point lies exactly on the line. It can be noted that the geometrical meaning of [3.3.1] is related to the area of the triangle ABP .

The incircle test determines whether a point lies inside, outside or on the circumcircle of a triangle. This test can also be reduced to a determinant; if the triangle vertices are given by the points ABC , this (again, for the 2D case) is:

$$D_I = \begin{vmatrix} a_x & a_y & a_x^2 + a_y^2 & 1 \\ b_x & b_y & b_x^2 + b_y^2 & 1 \\ c_x & c_y & c_x^2 + c_y^2 & 1 \\ p_x & p_y & p_x^2 + p_y^2 & 1 \end{vmatrix} \quad [3.3.2]$$

A positive value indicates that the point lies outside the circle, a negative value indicates that the point lies within the circle and a value of zero indicates that the point lies exactly on the circle.

Unfortunately, we cannot implement a naïve evaluation using these determinants; if the point lies very close to the surface of the circle a simple evaluation of the determinants can lead to an incorrect result due to the limited precision of the floating point calculations within the computer. A simple evaluation of the determinants in order to incrementally build up a mesh using the Bower-Watson algorithm can therefore cause the algorithm to fail catastrophically due to the exclusion (or inclusion) of triangles that should (or should not) be within the cavity; this failure occurs much more frequently than would be assumed.

Therefore, we have to use *exact* methods of evaluating the determinant which is one of the largest drawbacks to using the Bower-Watson algorithm. Shewchuk approaches this problem by writing a full floating point library which evaluates the determinant to increasing precision until the certainty of the sign of the determinant is established⁶². In this project, however, we use a much more modest implementation by clamping all values within the mesh to integers and using the Clarkson algorithm⁶³ to evaluate the exact sign on the determinant.

The Clarkson algorithm works by using the fact that if we need to find the determinant of the matrix \mathbf{A} , $|\mathbf{A}|$, we can decompose our matrix \mathbf{A} into \mathbf{CR} where \mathbf{R} is the unit upper triangular matrix. The decomposition $\mathbf{A} = \mathbf{CR}$ is simply a Gram-Schmidt decomposition, where our matrix \mathbf{C} is composed of orthogonal basis vectors constructed from the vectors in \mathbf{A} . Then $|\mathbf{A}| = |\mathbf{C}|$, so if we can find \mathbf{C} we can find the determinant of our original matrix. The use of floating-point calculations, however, gives us a matrix \mathbf{B} where $\mathbf{B} \approx \mathbf{C}$, and our Gram-Schmidt decomposition may fail if a vector of \mathbf{B} is very small, resulting in a near division by zero. The algorithm works by building up \mathbf{B} making sure that each orthogonal component is 'large enough' to ensure correct determinant evaluation. If an orthogonal component is *not* 'large enough' it is multiplied by some scalar s and undergoes a second Gram-Schmidt orthogonalization to magnify the component in the orthogonal direction (the multiplication by a scalar does not affect the *sign* of the determinant). If this is done subject to certain bounds, we can be certain that the matrix \mathbf{B} resulting from the process will give the same determinant as \mathbf{A} . It should be noted that the Gram-Schmidt orthogonalizations should be 'modified' Gram-Schmidt orthogonalizations which are numerically stable as they use linear combinations of the input vectors ($\mathbf{b}_i = \mathbf{a}_i + \text{Sum}[\mathbf{a}_i]$) and *not* a linear combination of the trial vectors ($\mathbf{b}_i = \mathbf{a}_i + \text{Sum}[\mathbf{b}_i]$); the repeated summation over calculated vectors can cause build-up of floating point errors (although Brönnimann and Yvinec⁶⁴ claim that only the second orthogonalization needs to be of the 'modified' variety and the first can be the classical orthogonalization). The details are complicated and presented in Clarkson's original paper and undergo an even more comprehensive analysis by Brönnimann and Yvinec⁵³.

3.3.4 Fast Point Location

In order to proceed with Delauney triangulation we need to determine into which element in the mesh a point falls (called the query point); the most straightforward way to do this would be to simply go through all the elements and test whether the query point lies inside the element. This is obviously not optimal, at most it would take $O(M)$ calculations for a mesh with M elements. Furthermore, the test to see whether a point lies within an element is (relatively) computationally expensive as it requires 3 orientation tests in 2D (which requires evaluation of 3 determinants) or 4 orientation tests in 3D. The determinant evaluations do not, however, have to be exact; if a point lies on the boundary between two elements it does not matter which element gets picked as the point will fall into the circumcircles of both elements and both elements will be removed anyway.

A more optimal way to locate the element is to use a *jump and walk* algorithm. In this algorithm, a number (M') of random elements are picked and the nearest element to the query point is chosen as the start element (this is the jump). The algorithm then walks through the mesh in a straight line towards the query point and terminates when it finds the element that contains the query point. This method is presented by Mücke, Saias and Zhu⁶⁵.




Figure 3.22 – In the standard jump and walk algorithm the distances from random M' elements are measured (left), and the algorithm proceeds to element containing the query point in a straight line(right).

This method, however, requires M' distance tests before it begins the walk stage. Furthermore, the walking part of the algorithm requires calculation of intercepts between the walk line and the mesh elements. An alternative method is presented as follows: We start at some element (this can be any element, although we can keep the jump part of the standard jump and walk algorithm) and take the dot products of the vector from the centre of the element to the query point (the search vector) and the normal vectors of the surface of the element. We then walk to the element through the surface that gives the *highest* value from these dot products. This means that we will *always* move towards the element that contains the query point (as element surfaces

pointing away from the query point will always give us negative dot products). The algorithm terminates when *all* the dot products are less than zero, meaning we have to be inside the element that contains the point.




Figure 3.23 – In our normal walking algorithm the search vector is calculated (left), and the algorithm walks through the surface that gives the largest dot product of search vector and element surface normal (right).

The algorithm has to make more steps than the standard jump and walk algorithm, but each step is much quicker than the standard algorithm. This is due to the fact that determination of the search vector is easy and each of the normals are already stored for each element (calculated during the meshing process). This algorithm is implemented in the project.

3.3.5 Ruppert Algorithm

The problem with incremental insertion algorithms lies in the processing of boundaries in the design (in contrast to advancing front and octree methods in which the mesh is built around the boundary). This is due to the fact that the Delaunay triangulation of a set of vertices may or may not include boundary edges (figure 3.24).




Figure 3.24 - User-defined boundaries may not appear in a Delaunay triangulation.

Ruppert was the first⁶⁶ (slightly before Chew introduced his second algorithm⁶⁷) to suggest a 2D algorithm which handles boundary segments. Furthermore, Ruppert's algorithm automatically grades the mesh according to feature sizes of the boundary whilst guaranteeing an angle bound of the triangles in the final mesh. The algorithm works by using the concept of encroached edges (see figure 3.25). An edge is defined as encroached if a candidate vertex falls within its diametral circle (1) (the circle that passes through the two endpoints of the edge with diameter equal to the length of the edge). If any boundary edge is encroached the candidate vertex is rejected and the

edge broken by insertion of a new vertex at its midpoint (2). The cavity processing then goes ahead using this newly inserted vertex (3,4).

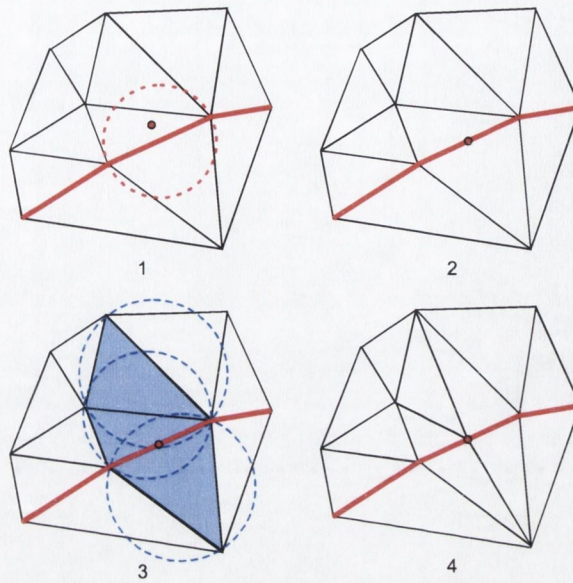


Figure 3.25 – Stages for boundary recovery in the Ruppert algorithm.

The mesh is incrementally constructed by taking a skinny triangle and attempting to insert a vertex in its circumcentre. If the candidate vertex encroaches any edges, these are split. The algorithm then takes a skinny triangle and repeats until there are no skinny triangles or encroached edges in the mesh. Breaking an encroached edge will sometimes remove the skinny triangle; if not, the edge splitting will continue until either the circumcentre no longer encroaches any edges, at which point a vertex can be inserted into the circumcentre, or the triangle is removed by an edge break. Pseudocode for this algorithm is as follows:

```

while (Bad Triangle List Has Members) {
    Find the Circumcentre of the Current Bad Triangle, p
    Find Boundary Edges p Encroaches
    Add Encroached Boundary Edges to Encroachment List

    if (Encroachment Edge List Has Members) {
        while (Encroachment List Has Members) {
            Find Midpoint of Current Encroached Edge m
            Insert Vertex at m
            Process Cavity Around m
        }
    }
    else {
        Insert Vertex at p
        Process Cavity Around p
    }
}

```





Figure 3.26- 2D Meshes generated by a quadtree method (top), Chew's algorithm (middle) and Ruppert's algorithm (bottom).
Figure taken from "Lecture Notes on Delaunay Triangulation" by Jonathan Shewchuk, p.55⁵⁵.

3.3.6 Shewchuk Algorithm

Mesh generation in 3D is much more complicated than 2D, though most of the algorithms covered previously can be extended to 3D with minor complications. The advantage of using the Bower-Watson algorithm becomes apparent here as this immediately extends to three dimensions with no difficulty (of course, a circumcircle becomes a circumsphere and tests now involve the candidate point being inside or outside a sphere rather than a circle). The main difficulty lies in the fact that we now have to worry about two types of boundary, boundary edges and boundary facets which both have to be recovered. The Shewchuk algorithm⁵⁵ is a 3D extension to the 2D Ruppert meshing algorithm which introduces the concept of an encroached subfacet. A subfacet is a triangular sub-face of a boundary facet.




Figure 3.27- A sample mesh showing boundary facet and boundary sub-facets.

A subfacet is encroached if a vertex falls within its equatorial sphere. The equatorial sphere of a triangle is the smallest sphere that passes through each of the triangle's vertices. Again, the diametral sphere is defined as the smallest sphere that passes through the endpoints of a boundary edge and the circumsphere of a tetrahedron is defined as the unique sphere that passes through the four points of the tetrahedron (figure 3.28).




Figure 3.28 – The three types of sphere used in the meshing process.

The algorithm maintains three lists: A list of tetrahedra to be split (these will be either large or badly shaped tetrahedra), a list of triangles to be split (these will be encroached boundary subfacet triangles) and a list of edges to be split (these will be

encroached boundary edges). Edges have priority over triangles and tetrahedra, triangles have priority over tetrahedra.

The algorithm progresses by taking a large or skinny tetrahedron from the list and trying to insert a vertex at its circumcentre. If the candidate vertex encroaches any edges, the vertex is rejected and the edges are added to the edge list to be split (which may or may not eliminate the skinny triangle; if not, the tetrahedron will remain in the list of bad triangles and will be split at a later stage). If the candidate vertex does not encroach any edges, boundary subfacets are checked next; if the candidate vertex encroaches any boundary subfacets the vertex is rejected and all the boundary facet triangles it encroaches are added to the triangle list for splitting. If the candidate vertex does not encroach any edges or facets, the vertex insertion goes ahead, eliminating the large or badly shaped tetrahedron. The next tetrahedron in the list is taken and the process repeats.

However, the process is complicated by the priority relationship. If any encroached subfacets are in the list, this list must be emptied before the processing of any further bad tetrahedra goes ahead. The process is illustrated by the following pseudocode:

```
while (Bad Tetrahedra List Has Members) {
    Find the Circumcentre of the Current Bad Tetrahedron, q
    Find Boundary Edges and Boundary SubFacet q Encroaches
    Add Encroached Edges to Edge List
    Add Encroached SubFacet to SubFacet List

    if (Edge List Has Members OR SubFacet List Has Members) {
        if (Edge List Has Members) {
            while (Edge List Has Members) {
                Find Midpoint of Current Encroached Edge m
                Find Boundary Edges m Encroaches
                Add Encroached Edges to Edge List
                Insert Vertex at m
                Process Cavity Around m
            }
        }
        if (SubFacet List Has Members) {
            while (SubFacet List Has Members) {
                Find the Circumcentre of the Current Facet, p
                Find Boundary Edges and Boundary SubFacet p Encroaches
                Add Encroached Edges to Edge List
                Add Encroached SubFacet to Facet List

                if (Edge List Has Members) {
                    while (Edge List Has Members) {
```



```
Find Midpoint of Current Encroached Edge
m
Find Boundary Edges m Encroaches
Add Encroached Edges to Edge List
Insert Vertex at m
Process Cavity Around m
}
}
else {
Insert SubFacet Circumcentre Vertex at p
Process Cavity Around p
}
}
}
else {
Insert Tetrahedron Circumcentre Vertex at q
Process Cavity Around q
}
}
```

As Shewchuk points out, this algorithm generates the best positions for the vertices on the boundary facets rather than, as with some Delaunay algorithms, triangulation of the boundaries as a prior stage to the 3D triangulation. As with the Ruppert algorithm, this process also guarantees both a quality and size bound on the tetrahedra in the final mesh. Furthermore, the mesh element size automatically adjusts to the feature size (see figure 3.29).




Figure 3.29 - The size of the mesh elements automatically grade depending on the local feature size in Shewchuk's algorithm.

3.3.7 Boundary Recovery

Unfortunately, boundary subfacets (and even boundary facets) may not appear in the final triangulation and have to be recovered. This boundary recovery process is complicated as it is difficult to quickly ascertain whether or not a facet is completely represented in the mesh by a set of subfacets. Many meshing routines⁶⁸ (including Shewchuk's own process for boundary recovery⁵⁵) use the fact that a set of vertices coplanar and within a boundary facet form a 2D Delaunay triangulation over the facet. By forming the 2D triangulation of all vertices on the facet and then searching 3D mesh for the triangles within the 2D triangulation, the missing facets can be ascertained and recovered by insertion of a vertex at their circumcentre. Shewchuk suggests using a 2D Delaunay triangulation for each boundary facet, and matching these 2D triangulations with the 3D mesh at each step to catch missing facets as they happen. This takes a large amount of processing, and is subject to co-circular degeneracies such that the 2D triangulations do not match the 3D mesh even for complete triangulations over the facet. Furthermore, Shewchuk does not cover thin boundary cases where two boundaries in close proximity interfere with each other. This is similar to the above case but can occur recursively; the two close boundaries can repeatedly remove subfacets in each other as they are processed. This can be seen in figure 3.30 below, which shows the effect of interfering boundary facets on a cuboid of shrinking length.




Figure 3.30– As the length of the cuboid diminishes, the boundary faces increasingly interfere with each other (the back of boundary faces is coloured differently to highlight the effect).

We introduce a very straightforward way to recover boundary facets which relies on the cavity creation stage and the maintenance of a 'ghost' subfacet list. During the cavity processing stage we delete the interior elements of the cavity, but we check whether we are about to remove a boundary subfacet (as these are marked as discussed later). If the boundary subfacet is about to be deleted by a vertex *not* on its

own boundary, it is moved to a ghost subfacet list (if the algorithm is implemented with a list of pointers to faces, this simply requires the addition of a pointer to the ghost list and the removal of a pointer from the face list). At this stage the ghost subfacet is also added to the list of subfacets to be broken (if it isn't included in this list already). Ghost subfacets and normal subfacets are then treated in the same manner in the subfacet break list; this means that a subfacet not physically present in the mesh can recover itself at a later stage (see figure 3.31).




Figure 3.31 - The interference of the boundary facets can be greatly reduced by using a 'ghost face' list.

Ghost subfacets are not checked for encroachment; this means that subfacets present in the mesh and the ghost subfacets can overlap (if, for example, the addition of another vertex to the boundary by another ghost subfacet on the same boundary has partially patched the boundary already). Because of this, we test whether vertex insertions for edges and facets on a boundary lie within a certain proximity of all the ghost subfacet Delaunay centres on the same boundary. If the processing of a ghost subfacet would introduce a vertex too close to another vertex on the boundary, the ghost subfacet is deleted from the list. The only remaining problem then is to choose a 'good' proximity' - if the proximity value is too small, clustering occurs. If it is too large, removing the ghost subfacet may cause a portion of the mesh not to be recovered.

3.3.8 Implementation

We can deduce whether an edge is a boundary edge, or if a face is a boundary subfacet using a simple vertex labelling procedure. Each vertex contains a list of the IDs of the boundaries to which it is attached; a vertex with no items in the list will be a 'free' vertex, a vertex with one item in the list will be a vertex on a boundary and a vertex with more than one items in its list will lie on a boundary edge. These vertex ID lists can be built incrementally as the only time we introduce a boundary vertex is

when we break a boundary subfacet; the new vertex can then be marked with the boundary's ID value. Similarly when an edge is broken the new vertex introduced at its midpoint is given the common ID values of the two vertices on either end on the edge. Keeping to these simple rules as the mesh is built incrementally ensures that every vertex is correctly labelled.

Any face in the mesh can then be checked as to whether it is a subfacet by finding the common ID value over the three vertices of the face; if the three vertices have a common ID, the face is a subfacet of the boundary with the common ID value. If this is done as the face is created, each face can be immediately marked as to whether it lies on a boundary and, if so, the ID value of the boundary it lies on. If any face with a boundary ID value is deleted, we can then decide whether or not to move it to the ghost subfacet list.

3.3.9 Computational Realisation

The entire meshing, matrix assembly and matrix solution process was packaged into a single application named 'Manifest'. In an advance on the previous versions of magnetic software (the Magnetic Solver packages), a fully programmable C syntax scripting interface was used with no predetermined shapes (unlike Magnetic Solver where, for example, a Halbach cylinder was a primitive construction object). Magnetic designs, including cylinders and flux sheets, can easily be created using scripts. Scripts also allow complex instruction flow which could work, for example, in a fully automated optimization algorithm. The script language also contains full file IO, which allows large batches of designs to be processed and compared efficiently later.

The current version only allows two primitives (although this could be extended very straightforwardly); a quad object, which is a six sided object with quadrilateral faces and user-defined vertices, and a cylinder object of which the radius and length are user-defined. The primitives are moved within the global bounding box by means of a local transformation, followed by a global translation, followed by a global rotation. This allows rotationally symmetric designs to be built up very easily by simply looping over the transformation and creation of the primitive whilst incrementing the angular rotation at each step. The scripting interface is handled by **Script.cpp**.

Once the design is constructed (and environmental variables such as the global box size and maximum mesh element volume have been set), the design is automatically

meshed using Shewchuk’s algorithm as outlined in the previous section; this is by far the most complex part of the application and uses **Mesh.cpp**. The mesh algorithm loops until no large or badly shaped elements are left in the design. Then, it further breaks elements which the user has requested be refined further; this is of particular use in observation volumes where a higher mesh density is required or in parts of the design containing soft materials, where high resolution is required for the flux. The mesh is then ‘painted’ using an algorithm that loops through every element and decides which material to assign it using the initial design building data. The elements, faces and vertices are stored in a doubly-linked list for efficiency.

Once the mesh is constructed, the matrix assembly begins and the application hands control to the **FEMengine.cpp**, which runs in a separate thread to the program shell. The matrix assembly is straightforward, although Manifest uses ‘compacted’ vectors (the class is contained in **Matrix.cpp**). Compacted vectors are fairly common. Here they are implemented by having a list of vector indexes and a list of vector values, with zero values in the vector not stored (see figure 3.32); this saves large amounts of processing when dealing with sparse matrices, as in this project. Arithmetic operations between compacted vectors is, however, fairly tricky; a number of routines were tried and those included in the project were found to be most efficient out of all those tested.




Figure 3.32 – A full vector(left) and the resulting ‘compacted’ vector.

After matrix assembly, the preconditioning matrix is constructed (or, the matrix is directly inverted using Cholesky factorization if desired). The conjugate gradient or direct solution is then carried out using compacted vector routines in **Matrix.cpp**.

Once the solution is complete data can be extracted from the system by routines that give the field at either any point or along any line through the design. If required, the

field can be displayed as an arrow plot or a 2D slice; the graphics routines are carried out by using OpenGL routines in **Graphics.cpp**. The mean field and standard deviation in any observation volumes can also be determined and used by the scripting engine, for example, to be written to a file for plotting by an external application.

4 Improvement of Homogeneity in Halbach Cylinders

4.1 Introduction

4.1.1 Effects of Cylinder Length on Field Homogeneity




Figure 4.1 – A standard Halbach cylinder of length L and inner and outer radius R_1 and R_2 respectively.

An ideal, infinite length, Halbach cylinder produces a field which is constant and uniform within the bore (see section 1.1.3). Using [1.1.12] with [2.1.3] and the definition of magnetic polarization [2.1.53] gives the magnetic induction, \mathbf{B} , in the bore of a Halbach cylinder as:

$$\mathbf{B} = J_r \ln\left(\frac{R_2}{R_1}\right) \hat{y} \quad [1.1.4]$$

where J_r is the magnetic remanence of the material, R_o and R_i are the inner and outer cylinder radii and \mathbf{B} is in the \hat{y} direction. For a finite length Halbach cylinder stray fields are produced at the ends of the cylinder; these *end effects* strongly affect the uniformity of the field within the cylinder as well as introducing small field components in the x and z directions. We seek to diminish these end effects whilst keeping the mean field as high as possible.




Figure 4.2 - B vector field and field magnitude for 2cm long Halbach cylinder




Figure 4.3 - B vector field and field magnitude for 4cm long Halbach cylinder




Figure 4.4 - B vector field and field magnitude for 6cm long Halbach cylinder

In practice a Halbach cylinder is *segmented*, as shown in figure 4.1, as opposed to having a circular ring-shaped cross section in the ideal case. The magnetic flux \mathbf{B} in the bore of an infinite length segmented cylinder is given by³:

$$\mathbf{B} = J_r \left(\frac{\sin(\frac{2\pi}{N})}{\frac{2\pi}{N}} \right) \ln \left(\frac{R_2}{R_1} \right) \hat{\mathbf{y}} \quad [4.1.1]$$

Where N is the number of segments. Segmentation decreases and also affects the uniformity of the field within the bore, but both of these effects can be reduced by increasing the number of segments. The effect of segmentation on field uniformity is, however, minor in comparison to the end effects.

Figures 4.2 – 4.4, show the \mathbf{B} field within the bore of an eight segment Halbach cylinder, calculated using the charge model, for three cylinders of increasing length with dimensions $R_1 = 3\text{cm}$, $R_2 = 6\text{cm}$, $J_r = 1.2\text{T}$ (the remanence of NdFeB). The \mathbf{B} field is shown on a slice through the central x - y plane of the design, and also on the central x - z plane of the design. z is the direction of the cylinder axis and the central plane is $z = 0$. The \mathbf{B} field direction is indicated by the black arrows and the \mathbf{B} field magnitude by a colour scale. It can be seen immediately that the majority of the flux produced by the design is confined to the bore of the cylinder

The end effects can be seen as a ‘bowing’ of the field towards the centre from the inner surface of the cylinder in the x - y plane and a bowing outwards from the centre of the cylinder in the y - z plane. The flux confinement is also affected by the length of the cylinder; the stray field outside the cylinder can be seen to vary inversely with the length of the cylinder.

We measured the homogeneity by measuring the field values within a thin sampling volume; we choose this to be a cylinder of very short length, S_l , and a radius, S_r , centred on $z = 0$ and slightly smaller than the inner radius of the Halbach cylinder under consideration, as shown in Figure 4.1. We sample the field on a regular grid of points within the cylinder when using the charge model, and at each tetrahedral element within the cylinder when using the finite element model. Taking these values we calculated the standard deviation about the mean field:

$$\Delta x = \sqrt{\frac{1}{N} \sum_{i=1}^N (x_i - \bar{x})^2} \quad [4.1.2]$$

Where \bar{x} is the mean value of the sample set. Figure 4.5, below, shows a plot of the mean field component B_y and its standard deviation ΔB_y for Halbach cylinders of

increasing length, calculated using the charge and finite element computer programs. The dimensions of the cylinder are given above, $R_1 = 3\text{cm}$, $R_2 = 6\text{cm}$ and remanence $J = 1.2\text{T}$. The sampling volume is set to be a cylinder of radius $S_r = 2\text{cm}$ and length $S_l = 0.1\text{cm}$ (centred on $z = 0$) for the finite element calculation and $S_l = 0\text{cm}$ (sampled in a plane at $z = 0$) for the charge calculation.




Figure 4.5 – Mean field, B_y , and standard deviation ΔB_y for Halbach cylinders of varying length

The finite element calculations used large numbers of mesh elements resulting in a matrix system with large N (around 40,000) and calculation time for any one design was typically be several orders of magnitude higher than the calculation for the same design carried out using the charge model.

The results from both programs can be seen to be in close agreement. It can be seen that the finite element model gives higher values for ΔB_y for each cylinder due to the slight difference in the sampling method. The maxima are, however, in the same location. It can be seen from figure 4.5 that ΔB_y for these particular dimensions has a maximum at a cylinder length of around 2cm.

4.1.2 Increasing Field Homogeneity

Field homogeneity can be improved using *shimming* methods. Shimming is the placement of small magnets (or current carrying coils) inside or outside the design (known as *active* shimming) or the placement of small pieces of soft magnetic material inside the design (known as *passive* shimming). Active shims directly

produce a field that compensates for the field inhomogeneities, whereas passive shims become magnetically polarized and the careful placement of these can then compensate for inhomogeneities. Typically shims are designed after fabrication of a design to smooth out any small inhomogeneities in the field. Apart from inhomogeneities caused by finite length and those due to the segmentation of the cylinder, inhomogeneities can also arise from several other sources. The magnetic segments are assumed to have their magnetization aligned precisely in a given direction, but often small deviations from this true alignment can arise during manufacture. This is also true of the geometrical construction of the design; it is assumed that the magnetic blocks are perfectly joined, which may not always be the case after fabrication. The blocks are also assumed to be magnetically transparent; however, in reality, the blocks have a small susceptibility and furthermore are also slightly anisotropic.

Here, we investigate two methods of increasing field homogeneity in short cylinders, which do not involve shimming. These are purely geometrical methods and do not affect any of the manufacturing considerations mentioned above. The first is to slice the cylinder into two parts and to separate the parts by a small distance S in such a way that the $l=3$ term of the spherical harmonic expansion for a Halbach cylinder is exactly cancelled⁶⁹ (see section 2.3.3). The second, which is a novel design, involves increasing the inner bore radius linearly from either end such that it reaches some maximum inner radius R_3 in the centre of the design. The main idea behind this design is to compensate for the ‘bowing’ of the field in the central plane caused by the end effects by focusing the stray fields into the bore of the cylinder. We also investigate a method of increasing field homogeneity by shimming a convectional Halbach cylinder by inserting a thin cylinder of soft iron into the bore.

4.2 Split Halbach Cylinder

4.2.1 Design Overview

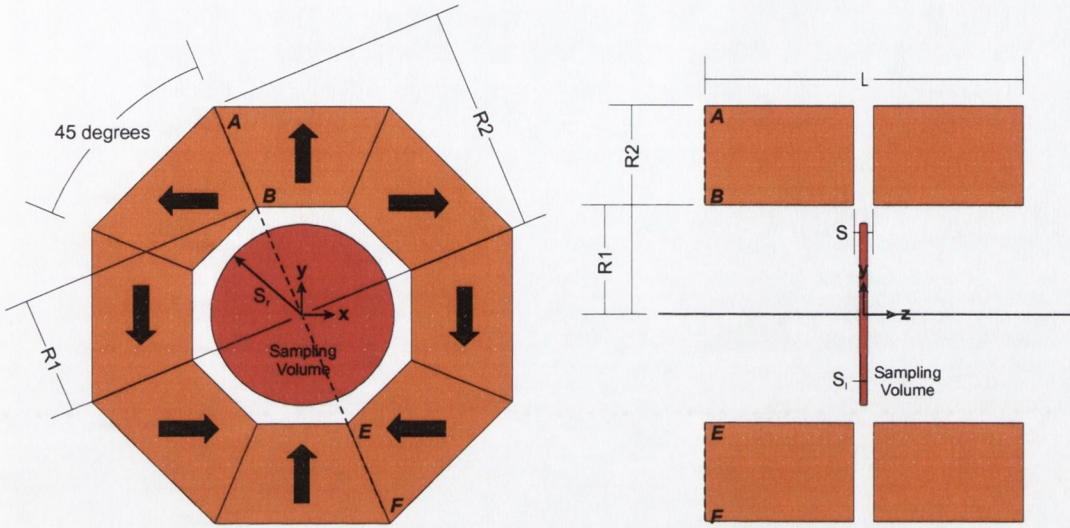


Figure 4.6 – A split Halbach cylinder with split with S .

In the spherical harmonic expansion expression for a Halbach cylinder [2.3.45] the $l = 1$ term gives a uniform field in the bore of the cylinder, and the higher odd l terms add progressively decreasing distortion terms to this field (there are no even l terms due to the symmetry, see section 2.3.3). The $l = 3$ term represents the greatest distortion term and an ingenious way to remove this⁷⁰ comes from considering the curve of S_3 (defined in [2.3.42]) against θ_L :




Figure 4.7 – The S_3 term in the spherical harmonic expansion for a Halbach cylinder.

The value of S_3 can be seen to be the same at two values of the angle θ_L . If the inner radius of the cylinder is fixed, these two angles represent two different points along the cylinder axis where S_3 takes the same value. If cylinders of these two lengths and opposite magnetization are superimposed on top of one another, the contributions from their respective S_3 will cancel as they will be of equal value but opposite in sign, and the dominant distortion term of the superposition of both cylinders will be zero.

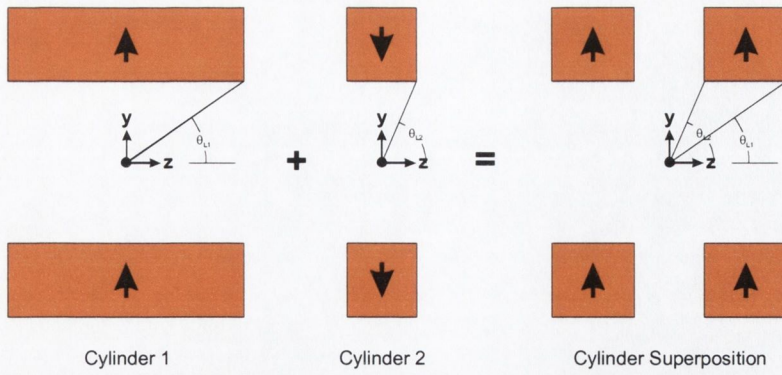


Figure 4.8 – Linear supposition of magnetic cylinders, shown in cross-section. The combination of two magnets of opposite polarization results in a area of zero magnetization, or a void, in the design.

As magnetization is linear, the superposition on a block of material on top of a block of material magnetized in the opposite sense gives a material of zero magnetization; this corresponds to a split in the cylinder, the separation of the two parts of which can be calculated from [2.3.42]. Further distortion terms can be cancelled in the same way, although the relative contribution to the distortion becomes very small as l increases.

Figures 4.9 - 4.11, below, show the \mathbf{B} field within the bore of an eight segment Halbach cylinder, calculated using the charge model, for three cylinders with different split widths S , of dimensions $R_1 = 3\text{cm}$, $R_2 = 6\text{cm}$, $J_r = 1.2\text{T}$ (the remanence of NdFeB) and length 5cm.




Figure 4.9 - B vector field and field magnitude for 0cm split (a regular 5cm long cylinder).




Figure 4.10 - B vector field and field magnitude for 1cm split.





Figure 4.11 - B vector field and field magnitude for 2cm split.



4.2.2 Effect of Split Width on Homogeneity

We measured the effect of increasing the split distance, S , whilst keeping the total length of the cylinder, L , and inner and outer radii, R_1 and R_2 , constant.




Figure 4.12 -Field mean B_y and standard deviation ΔB_y for a split Halbach cylinder with varying split width.

Figure 4.12 shows a plot of the mean field B_y and its standard deviation ΔB_y for a split Halbach cylinder with increasing separation, calculated using both the charge and finite element models. The dimensions of each of the cylinder parts are chosen to be $L = 5\text{cm}$, $R_1 = 3\text{cm}$, $R_2 = 6\text{cm}$, and the remanence $J = 1.2\text{T}$. The sampling volume is again a cylinder of radius $S_r = 2\text{cm}$ and length $S_l = 0.1\text{cm}$ for the finite element calculation and $S_l = 0\text{cm}$ for the charge calculation. It can be seen from figure 4.11 that ΔB_y for these particular dimensions has a minimum at a separation of around 8mm.

Also shown on figure 4.12 as a dashed vertical line is the separation needed to remove the $l = 3$ term from the harmonic expansion calculated using [2.3.42]. For this calculation we chose the *inradius* of the octagon (the distance from the origin to the centre of each inner face) and not the circumradius. For a cylinder with inner circumradius of 3cm this gives an inradius of $3\cos(\pi/8) \approx 2.77\text{ cm}$. θ_L is then given by $\tan^{-1}(2R_i/L) = \tan^{-1}(2.77/5)$ so $\theta_L \approx 0.84$ radians. The second value of $\theta_L = 1.34$ radians, as can be seen from inspecting figure 4.7. The second length (which is the

separation distance, S) is then given by $L_2 = 2 \times 2.77 / \tan(1.34)$, which gives $L_2 \approx 1.3\text{cm}$.

The minimum found differs from the minimum for the ideal case for two main reasons. The first is the effect of the remaining l terms which are not cancelled by the cut; the separation distance found by the simulation represents an average minimum of *all* the harmonic distortion terms. The second reason is the effect of segmentation. Figure 4.13 shows the effect of increasing N on the ΔB_y curve for the 5cm long cylinders (note the ΔB_y axis has changed from the right to the left vertical axis in the plot). As N becomes larger, the minimum moves closer to our ideal calculated cut width.




Figure 4.13 – Effect of increasing the number of component segments in the split Halbach design.

Figure 4.14 shows a family of these graphs, to determine the trend of the minima with respect to the total length of the cylinder:




Figure 4.14 – Standard deviation as a percentage of the mean for varying cylinder length of split Halbach design $N=8$

It can be seen from figure 4.14 that each cylinder length has a different optimum split width, as expected. Shorter cylinders can be seen to have a slightly higher ΔB_y at optimum separation distance than longer cylinders.

4.3 Halbach Cylinder With Non-Uniform Inner Radius

4.3.1 Design Overview

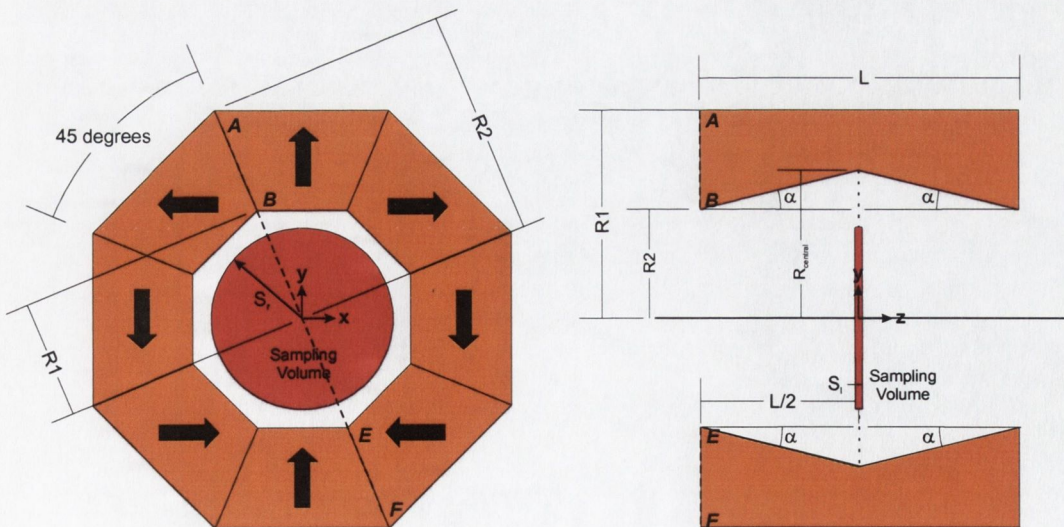


Figure 4.15 - A Halbach cylinder with a non-uniform inner bore radius

We introduce a novel design in which the bore radius of the Halbach cylinder increases linearly from both ends to a maximum central radius. This design was

introduced to see whether we could reduce the ‘bowing’ effect of the field caused by the end effects by focusing the flux inwards by varying the inner face surfaces. We expect the homogeneity to be affected by the angle of slope, α , and in particular, we expect there to be a optimum value of this slope angle which gives us the best homogeneity over the sample volume for any chosen length of cylinder. Figures 4.16 - 4.18, below, show the effect on the field within the bore of the cylinder for three different values of central radii (corresponding to three different slope angle values):




Figure 4.16- B vector field and field magnitude for inner radius of 3cm (a regular 5cm long cylinder).




Figure 4.17 - B vector field and field magnitude for inner radius of 4.5cm.





Figure 4.18 - B vector field and field magnitude for inner radius of 5.5cm.



4.3.2 Effect of Increasing Central Bore Radius

We vary R_{central} (which changes the slope angle α), while keeping the length the cylinder, L , and the radii R_1 and R_2 fixed (see figure 4.15):




Figure 4.19 - B_y field mean and standard deviation for a Halbach cylinder with varying central radius.

Figure 4.19 shows a plot of the mean field B_y and its standard deviation ΔB_y for a Halbach cylinder with increasing central inner radius, R_{central} , calculated using both the charge and finite element models. The dimensions of the cylinder is chosen to be $L = 5\text{cm}$, $R_1 = 3\text{cm}$, $R_2 = 6\text{cm}$, and the remanence $J = 1.2\text{T}$. The sampling volume is once again a cylinder of radius $S_r = 2\text{cm}$ and cylinder $S_l = 0.1\text{cm}$ for the finite element calculation and $S_l = 0\text{cm}$ for the charge model calculation. In figure 4.20 we plot a family of such curves to ascertain the trend of the minima with respect to the length of the cylinder:




Figure 4.20 - Standard deviation as a percentage of the mean for cylinder with non-uniform bore and varying length.

It can be seen that longer cylinders give an improved standard deviation as a percentage of the mean in comparison to shorter cylinders.

4.4 Halbach Cylinder With Soft Iron Insert

4.4.1 Design Overview




Figure 4.21 - A Halbach cylinder shimmed with central soft iron cylinder of thickness t .

This design is composed of a Halbach cylinder shimmed with a second soft iron cylinder inserted inside its bore. The outer radius of the iron cylinder is the same as the inner radius of the Halbach cylinder, and both cylinders are the same length (see

figure 4.21). Initially it would appear that this design would give no great improvement in homogeneity and the addition of the soft iron cylinder would simply reduce the strength of field in the central bore. This is due to the fact that soft iron is a flux conduit, and therefore the iron cylinder would simply channel all the flux around the central region, reducing the field. However, Brown⁵⁸ showed experimentally that although the expected drop in the mean field occurs, the homogeneity is improved inside designs incorporating thin soft iron cylinders (of thickness of around 1mm inside a Halbach cylinder of $L = 5\text{cm}$, $R_1 = 3\text{cm}$, $R_2 = 6\text{cm}$). As the Magnetic Solver program can only carry out calculations for designs composed of hard materials, calculations for this design were carried out using Manifest (our finite element program).

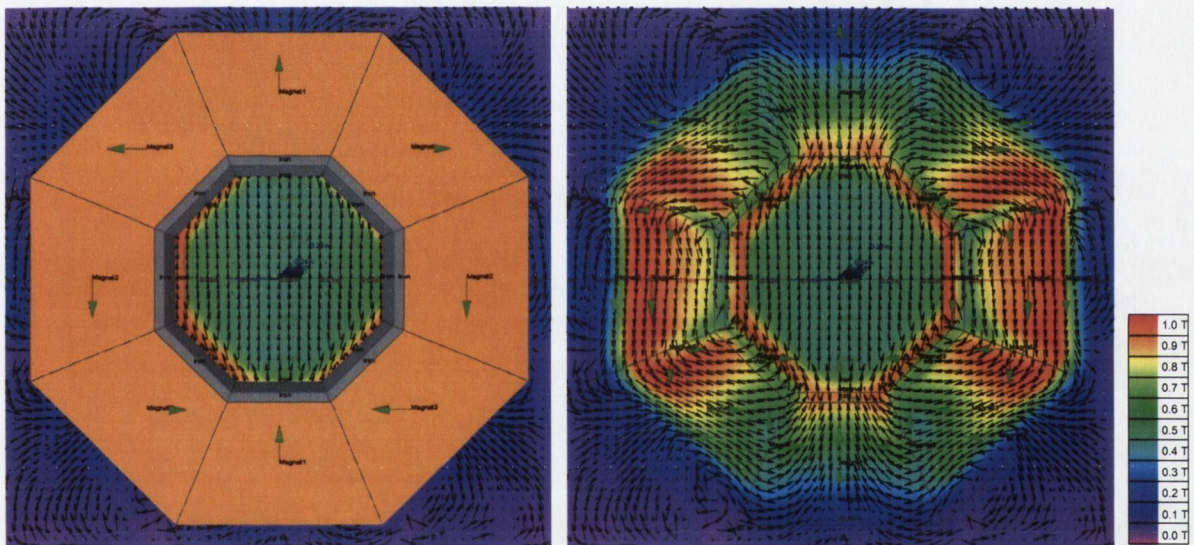


Figure 4.22 - B vector field and field magnitude for shim thickness of 2.5mm, calculated using Manifest: the magnetic field in the bore and external to the magnet (left), and the field inside and outside the design (right).

4.4.2 Effect of Insert Thickness on Homogeneity

We investigated the effect of increasing the thickness t of the soft iron cylinder, whilst keeping the length of the cylinders and inner and outer Halbach radii, R_1 and R_2 , constant.




Figure 4.23 - B_y field mean and standard deviation for a Halbach cylinder with soft iron yoke of varying thickness.

Figure 4.23 shows a plot of the mean field B_y and its standard deviation ΔB_y for a Halbach cylinder shimmed with a soft iron cylinder. The dimensions of the cylinder are $L = 5\text{cm}$, $R_l = 3\text{cm}$, $R_r = 6\text{cm}$, and the remanence $J = 1.2\text{T}$. The relative permeability of the soft iron was chosen to be $\mu_r = 1000$ and the saturation value of the iron, M_s , was chosen to be $2.15/\mu_0 \text{ A/m}$ (see section 3.2.3). The sampling volume is a cylinder of radius $S_r = 1.2\text{cm}$ and cylinder $S_l = 0.5\text{mm}$. The minimum thickness of iron cylinder was 1mm due to difficulty in meshing extremely thin boundaries. However, without using the techniques in boundary recovery described in section 3.3.7 this minimum thickness would have been approximately an order of magnitude higher.

The three points marked in green on figure 4.23 are the average values of B_y measured by Brown⁵⁸. The reason that these points lie slightly above the calculated values could be due to the soft iron cylinders in his experiment having some residual magnetization; the field values calculated by Manifest use the method outlined in section 3.2.3, which uses the initial magnetization curve of the soft material in the simulations (curve A to B in figure 1.3). It can be seen from figure 4.23 that the standard deviation curve has a minimum at around 1.75mm. Brown quotes the thickness of iron cylinder giving the best homogeneity in a Halbach cylinder with

dimensions $L = 5\text{cm}$, $R_1 = 3\text{cm}$, $R_2 = 6\text{cm}$ as just over 1mm^{58} , although he takes the sampling area over a square of side 1.4cm in the central plane of the magnet.

4.5 Comparison of the Methods of Improving Homogeneity

To compare the designs we plot the mean B_y against its standard deviation as a percentage of the mean. A regular eight sided Halbach cylinder with $R_1 = 3\text{cm}$, $R_2 = 6\text{cm}$ and remanence $J = 1.2\text{T}$ is included for reference.




Figure 4.24 – Comparison of the three designs from section 4.1, 4.2 and 4.3.

Figure 4.24 Shows the standard Halbach cylinder (section 4.1, blue line), the split cylinder design (section 4.2, green lines) and the Halbach cylinder design with the non-uniform inner radius (section 4.3, red lines). Each point on the blue line represents a regular cylinder of different length; moving along the line from the lower right corner each point represents an increase in cylinder length by 0.5cm . The green lines represent five split cylinders with varying split width. The point at which the green and blue curves intersect represents a split distance of zero; each point on the green line moving away from the blue line then represents an increase of 1mm of the split width, up to a maximum split width of 2cm . The red curves represent five cylinders with varying inner central bore radius. The red and blue curves intersect where the inner radius for each cylinder is constant; each point along the red line then

represents an increase in the inner radius at the centre by 1.5mm (keeping the inner radii at either end fixed), up to a maximum central bore radius of 5.85cm.

The designs that produce the highest field with the best homogeneity are the designs with minima as close as possible to the bottom right of the graph: it can be seen that the best possible design is a very long cylinder. A cylinder with $N = 8$, $J_r = 1.2$ T and $R_o/R_i = 2$ and length approaching infinity has a value of B_y in the bore of 0.749 T, from [4.1.1]. The standard deviation of the field ΔB_y of the *ideal* infinite length cylinder will approach zero, but segmentation will cause ΔB_y to approach a value slightly above zero; this can be seen in figure 4.24.

The sample radius S_r for the design comparison in figure 4.24 is 2cm. To compare the split design and the Halbach cylinder design with the non-uniform inner radius with the shimmed cylinder design the sample radius must be decreased (due to the fact that the shim would otherwise overlap the sampling volume). Figure 4.25, below, shows the shimmed Halbach design in comparison the other three designs for a sampling radius S_r of 1.2cm and a total cylinder length of 5cm:




Figure 4.25 – Comparison of the designs from section 4.1, 4.2, 4.3 and 4.4.

Figure 4.25 Shows the standard Halbach cylinder (section 4.1, blue line), the split cylinder design (section 4.2, green line), the Halbach cylinder design with the non-uniform inner radius (section 4.3, red line) and the Halbach cylinder with the soft iron shim (section 4.4, yellow line). The points on the red, green and blue curves are

generated using the charge model in the same way as for figure 4.24. Each point on the yellow curve moving to the left of its intersection with the blue curve represents an increase in the shim cylinder thickness of 0.25mm. There is a small difference between the point of interception of the red and green curves with the blue curve and the intercept of the yellow curve with the blue curve. This is due to the sample being taken over a sampling volume with $S_l = 0$ for the first three of these designs and over a sampling volume with $S_l = 0.5\text{mm}$ for the shimmed design (which was calculated using the finite element method). It can be seen that the shimmed design provides a good increase in homogeneity over the standard Halbach cylinder and provides a much higher field at its minimum than either the split design or Halbach design with the non-uniform inner radius.

The minimum of ΔB_y as a percentage of the mean for the split design (section 4.2) and the design with the non-uniform inner radius (section 4.3), sampled over a volume with $S_r = 2\text{cm}$ and $S_l = 0$ is given in table 4.1:

L (cm)	Stdev%Mean 4.2 Min.	Stdev%Mean 4.3 Min.	4.2 Mean (T)	4.3 Mean (T)
3	1.35	-	0.16	-
4	1.04	-	0.31	-
5	0.76	0.29	0.42	0.25
6	0.57	0.20	0.51	0.34
7	0.41	0.14	0.57	0.43
8	0.31	0.12	0.62	0.50
9	0.25	0.11	0.65	0.54
10	0.25	0.10	0.67	0.59

Table 4.1 – Minima for split and non-uniform bore designs.

Of these two designs, the design producing the highest mean field with the greatest homogeneity for any particular cylinder length is the split design. However, the design with the non-uniform inner radius produces better homogeneity than the split design, although this is achieved for lower fields. Also, as the length of a standard cylinder increases, the standard deviation as a percentage of the mean converges to approximately 0.15, with a mean field of approximately 0.74 T. It can be seen that the design with the non-uniform inner radius can approach this field homogeneity for cylinders of considerably shorter length. This design, therefore, can be put to use in applications in which low field strength with high field homogeneity is required.

The Halbach design shimmed with the soft iron cylinder produces a higher mean field at the minimum of the standard deviation than either of these designs, although the homogeneity achieved is not as good. The minimum of ΔB_y as a percentage of the mean for the split design (section 4.2), the design with the non-uniform inner radius (section 4.3) and the shimmed design (section 4.4) sampled over a volume with $S_r = 1.2\text{cm}$ is given in table 4.2:

L (cm)	Stdev%Mean 4.2 Min.	Stdev%Mean 4.3 Min.	Stdev%Mean 4.4 Min.	4.2 Mean (T)	4.3 Mean (T)	4.4 Mean (T)
5	0.094	0.035	0.779	0.36	0.21	0.47

Table 4.2 – Minima for split and cut designs.

It should be noted that the width of the split in the design of section 4.2 and the value of the central radius in the design of section 4.3 which gave the minimum value of ΔB_y as a percentage of the mean is *different* in the case of a sample volume of radius 2cm from that of a smaller sample volume of radius 1.2cm. This accounts for the difference between the values of the minimum mean field B_y for a cylinder of 5cm length in table 4.1 and table 4.2.

It is also important to take into account the volume of hard material used in the design; rare earth alloys are typically very expensive and we would ideally like to use the minimum possible volume of material. Figure 4.26 gives a comparison of the volume of the two designs from section 4.2 and 4.3 against the mean field. Soft materials are typically much cheaper than hard materials, and the shimmed design from section 4.4 is not considered here.




Figure 4.26 – Comparison of the volume of magnetic material used for the three designs.

The lines converge where the split width is 0cm in the split design and the central bore radius $R_{inner} = R_2$ in the design with non-uniform inner bore radius. Each point to the left of this represents an increase in the split width by 1mm for the split cylinders (green lines) and an increase in the central radius by 1.5mm for the Halbach design with non-uniform inner radius (red lines). It can be seen that the split design achieves any given mean field strength with the lowest volume of material. Figure 4.27 shows ΔB_y as a function of the volume.




Figure 4.27– Standard deviation for different volumes of material for each of the three designs.

It can be seen that for a minimum ΔB_y the designs of section 4.3 require the least volume of hard magnetic material when compared to the split designs of the same length.

In summary, there is no 'best' design; each has advantages and disadvantages and in practice the design must be picked to match the requirements of the particular application. If a high field is required and homogeneity is not important, a standard cylinder should be used as it provides the highest field strengths for any particular cylinder length. If homogeneity and field strength are both important, either the split cylinder or the shimmed cylinder should be chosen. The shimmed design provides high field with a good increase in homogeneity, whereas the split cylinder provides high homogeneity with lower mean field (see table 4.2). If field homogeneity is important and the fields required are low, the Halbach design with non-uniform inner radius should be chosen as it provides the best homogeneity for short cylinders out of the three designs considered.

5 Designs Producing a Constant Field Gradient

5.1 Introduction

We have investigated methods of producing a uniform field gradient using cylindrical designs. Ideally, we required designs which produced the highest and most uniform field gradient over the largest possible working volume. We considered three designs; firstly, a design in which the bore of a Halbach cylinder increases down the length of the cylinder giving a field gradient orthogonal to the field (figures 5.1, 5.2). Secondly, a design in which the bore of a Halbach cylinder is offset giving a field gradient parallel to the field (figure 5.8). The final design considered was an ‘open’ design in which a sheet of one-sided flux is partially rolled, giving a field gradient parallel to the field (figure 5.12). The remnant magnetic polarisation J_r is fixed at 1.2 Tesla throughout all these simulations.

5.2 Halbach Cylinder With Linearly Varying Inner Radius

5.2.1 Design Overview




Figure 5.1 – A Halbach design with linearly varying inner radius; front and back faces.




Figure 5.2 - Halbach design with linearly varying inner radius; cross section.

In this design the inner radius of the bore of a Halbach cylinder increases linearly along the z -axis, as shown in figures 5.1 and 5.2. The idea behind this design is that the central field in a Halbach cylinder depends on the ratio of inner and outer radii and if we consider the extremely simplified model of two semi-infinite Halbach cylinders placed end-to-end, shown in figure 5.3, we can see that the higher field at point A must change smoothly (though not necessarily linearly) to the lower field at point B . This is due to the fact that the field within the bore is given by a magnetostatic Laplace equation, and as there are no magnetic source terms between point A and B the field must be smooth along this path.




Figure 5.3 - A schematic view of two semi-infinite Halbach cylinders placed end-to-end.

The question arises as to whether we can approximate a linear transition by some choice of R_{i1} and R_{i2} . The same reasoning can be applied to a design incorporating any number of component cylinders with different radii, and in the design shown in figures 5.1 and 5.2 we take the limit in which the length of each of these component cylinders becomes zero, producing a smoothly sloping inner surface. The

question becomes whether we can obtain an approximately linear gradient along the z axis by some choice of θ .

We found that a design with parameters $R_1 = 1\text{cm}$, $R_2 = 3.3\text{cm}$, $R_3 = 5\text{cm}$ and $L = 10\text{cm}$ achieves a constant uniform field gradient within a region along the z -axis. The vector field from a design with these parameters is shown in figure 5.4:




Figure 5.4 – The B vector field and field magnitude for Halbach cylinder with linearly varying inner bore.

The main field component is in the y direction and can be seen to decrease along the z -axis. The B_y field component and gradient dB_y/dz (taken from positive to negative z) is plotted for these parameters in figure 5.5:




Figure 5.5 – B_y and dB_y/dz for standard design for Halbach with linearly varying inner bore; field is taken from -7cm to $+7\text{cm}$ along the z -axis (black line in figure 5.4).

The region of approximately constant gradient is marked on figure 5.5 by the two vertical dashed lines and the ends of the cylinder by the unbroken vertical lines; we have defined this to be a 3cm long region from -1cm to +2cm along the z-axis in this design. The magnitude of dB for this region for the parameters above is around 11 T/m. This region of approximately constant gradient will be referred to as the *working region* of the design.

In an ideal Halbach cylinder the magnetization direction changes continuously in the plane perpendicular to its axis, according to [1.1.12]. Using the charge model (section 2.2.1) we can derive an analytic expression for the field component B_y on the axis of such an un-segmented cylinder with inner radius linearly varying along the axis and finite length L (see appendix, A.1.3):

$$B_y(0,0,z) = \frac{J_r}{2} \left[\ln \left| \frac{\sqrt{R_3^2 + z^2} - z}{\sqrt{R_3^2 + (L-z)^2} + L - z} \right| + \frac{1}{\sqrt{1 + \tan^2 \theta}} \ln \left| \frac{\sqrt{(L + \alpha)^2 + \gamma^2} + L + \alpha}{\sqrt{\alpha^2 + \gamma^2} + \alpha} \right| \right] \quad [5.2.1]$$

where:

$$\alpha = \frac{R_1 \tan \theta - z}{1 + \tan^2 \theta} \quad [5.2.2]$$

$$\gamma = \sqrt{\frac{R_1 - z \tan \theta}{1 + \tan^2 \theta}} \quad [5.2.3]$$

This field is shown in figure 5.6 for a design with parameters $R_1 = 1\text{cm}$, $R_2 = 3.3\text{cm}$, $R_3 = 5\text{cm}$ and $L = 10\text{cm}$. Note that in the analytic expression $z = 0$ is the small end of the cylinder and not the centre.




Figure 5.6 - The analytic expression for a cylinder giving a constant field gradient.

The region of approximately constant gradient corresponding to figure 5.5 is marked by the two vertical dashed lines and the ends of the cylinder by the unbroken vertical lines. When compared to figure 5.5 it can be seen that the expression gives a B_y profile that starts at a higher value at the smaller end of the cylinder and the maximum is slightly lower than the segmented design. The gradient is, however, of similar magnitude within the working area in the two cases, although the analytic expression predicts a slight increase in the gradient along the axis. The segmentation therefore has little effect on the maximum achievable value of the gradient.

5.2.2 Design Summary

We chose cylinders with dimensions $R_1 = 1\text{cm}$, $R_3 = 5\text{cm}$, $L = 10\text{cm}$ and $\tan(\theta) = 0.23$, and examined the effects of varying the slope angle θ , the length L , and the outer radius R_3 one at a time. Graphs of the field B_y and field gradient $\text{d}B_y/\text{d}z$ as a function of z for different $\tan(\theta)$, L and R_3 are shown in appendix A.2.1. It should be noted that B_x and B_z are negligible in comparison to B_y (the mean value of B_z is approximately less 10^{-3} the mean value of B_y and the mean value of B_x is approximately less than 10^{-4} the mean value of B_y). We also calculated the standard deviation of $\text{d}B_y/\text{d}z$ within the working area to assess its uniformity. We varied the slope angle θ and found that the most constant field gradient is achieved for a value of $\tan(\theta) = 0.23$ corresponding to a gradient of magnitude approximately 11 T/m and a minimum standard deviation of 0.063 T/m (figure A.2.3). Intuitively, we would

expect the largest change of field along the z -axis to occur when $R_2 \approx R_3$ as the field contribution from the wider end should then be close to zero. However, this does not give a constant field gradient (as can be seen from figure A.2.2).

We found that a larger region of constant gradient cannot be achieved by simply changing the cylinder length (figure A.2.4). We also found that changing the outer radius of the cylinder does not increase the magnitude of the gradient within the working region.

This design is the only one developed in this project which produces a field gradient orthogonal to the field. The standard design with parameters $R_1 = 1\text{cm}$, $R_2 = 3.3\text{cm}$ (corresponding to $\tan(\theta) = 0.23$), $R_3 = 5\text{cm}$ and $L = 10\text{cm}$ produces a constant field gradient of approximately 11 T/m within the working region, from -1cm to +2cm along the z -axis. The uniform gradient cannot be tuned by a simple variation of a single parameter, for example, the outer radius or slope angle. However, using *scaling* of the design as a whole we can control the field gradient.

The length of the region of constant gradient varies as $1/s$ where s , the scaling factor, is a constant multiplying the parameters R_1 , R_2 , R_3 and L in the design. Scaling the entire design down *increases* the gradient which, although somewhat counter-intuitive, has a straightforward explanation which can be visualised by scaling down the system shown in figure 5.3. The fields at points A and B depend only on the ratio of inner to outer radii for both cylinders respectively, so this ratio and therefore the field at both of these points does *not* change as the system is scaled. However, the distance between the points on the z -axis does change, and so the gradient dB_y/dz along the z -axis is inversely proportional to the scaling factor. Extremely high gradients are therefore possible in small cylinders, but the working region decreases correspondingly.




Figure 5.7 – The value of the field gradient in the linear region of the design scales as $1/s$.

Once the gradient for a particular application is chosen, figure 5.7 can be used to pick the appropriate scaling factor for the design.

5.3 Halbach Cylinder With Offset Bore

5.3.1 Design Overview




Figure 5.8 – A Halbach design with offset bore, the bore offset is given by O_y .

This design is a variation of the standard eight segment Halbach cylinder in which the central bore is offset by an amount O_y with respect to the outer cylinder axis as shown in figure 5.8. A constant field gradient dB_y/dy is produced in the bore, depending on the inner radius, R_1 , the outer radius, R_3 , and the bore offset, O_y . The idea behind this design is that the higher volume of magnetic material in the lower part of the design

and the lower volume of magnetic material in the upper part of the design will produce a gradient along the y-axis.




Figure 5.9 - The B vector field in Halbach cylinder with an offset bore.

The \mathbf{B} field from a design with $R_1 = 2\text{cm}$, $R_3 = 6\text{cm}$, $O_y = 3\text{cm}$ and $L = 10\text{cm}$ is shown in figure 5.9. It can be seen that the design has good field confinement with little stray field outside the cylinder. The field gradient for this cylinder is approximately 25 T/m.

5.3.2 Design Summary

We chose a cylinder with dimensions $L = 10\text{cm}$, $R_3 = 6\text{cm}$, $R_1 = 1\text{cm}$ and examined the effects of varying the bore offset O_y and the length L individually. Graphs of the field gradient as a function of z for different O_y and L are shown in appendix A.2.2. We found that the magnitude of the gradient is proportional to the bore offset and the gradient has good uniformity over most of the bore (figure 5.10). Defining a working area of length 0.6cm from 0.45cm to 0.15cm along the y-axis, we found that the value of bore offset giving the lowest standard deviation is $O_y = 2.5\text{cm}$, which produces a field gradient of 44 T/m with a standard deviation of 0.087 T/m (figure A.2.7). We investigated the effect of varying the length of the cylinder and found that the gradient is most uniform for long cylinders (figure A.2.8). This suggests that end effects caused by short cylinders strongly affect the region over which the gradient is constant.

This design is therefore slightly different to the other designs considered as the end effects have a significant influence on the constancy of the field gradient (figure

A.2.8). The cylinders must be fairly long in comparison to their radii to overcome this effect.




Figure 5.10 – The field gradient in the linear region is directly proportional to the bore offset.

Figure 5.10 shows the value of the gradient at the centre of the bore (which is the midpoint of the region of uniform gradient) calculated for a cylinder with $R_1 = 1\text{cm}$, $R_2 = 15\text{cm}$ and $L = 50\text{ cm}$. If the entire design is scaled up or down by multiplying the parameters R_1 , R_3 , O_y and L by the scaling factor s , the magnitude of the gradient within the working area scales as $1/s$:




Figure 5.11 – The effect of scaling the Halbach design with offset bore.

In summary, this design would suit applications in which a low uniform field gradient parallel to the field direction is required over a small working area.

5.4 Rolled Flux Sheet

5.4.1 Design Overview




Figure 5.12- A design composed of a semi-rolled sheet of one-sided flux of length L .

This design is a novel ‘open’ cylinder design composed of a partially rolled one-sided flux sheet. Our motivation for this design came from the offset bore design covered in section 5.3. It can be seen that the gradient of the field for the design with the offset bore is proportional to the size of the offset amount. If the offset for this design is increased the bore will, at some point, touch the outer surface of the cylinder. A design with the inner bore touching the outer surface of the cylinder would waste large amounts of magnetic material in the lower part of the cylinder and would also be difficult to manufacture due to the thin slivers of material in the top part of the cylinder. We therefore investigated whether a rolled flux sheet would give the same effect, but would produce the field more efficiently, that is, with a lower volume of magnetic material for the same gradient. All designs we considered had twelve segments (as shown in figure 5.12).




Figure 5.13 – A one-sided flux sheet being rolled. The flux is almost entirely confined within the fold of the rolled sheet.

Figure 5.13 shows how the flux distribution changes as a one-sided flux sheet is rolled into a cylinder. As can be seen, the flux is concentrated within the fold and always confined to one side of the sheet. However, we found that an open cylinder with a semi-circular cross-section could *not* give a constant field gradient. After trying a number of different cross-sectional surfaces (such as parabolic and exponential) we found that a region of constant gradient could be achieved using a cylinder with an *elliptical* cross section.

The flux sheet was rolled into an open cylinder with elliptical inner and outer surfaces with different semi-axes. To make it easier to compare the rolled design with a standard circular cylinder, with inner and outer radii ir and or , the semi-minor and semi-major axes for the inner and outer elliptical surfaces were given as multiples of ir and or .

The inner ellipse has semi-major axis given by $ai \times ir$ and semi-minor axis given by $bi \times ir$. The outer ellipse has semi-major axis given by $ao \times or$ and semi-minor axis given by $bo \times or$, shown in figure 5.12. These dimensionless constants ai , ao , bi , bo can be thought of as elliptical distortions to a regular cylinder with inner radius ir and outer radius or , and have values around unity. The *angular range* of the ellipse was defined as the angle from the position of the centre of the first block to that of the last block, shown in figure 5.12.

We investigated the effect of varying the inner and outer ellipse parameters, as well as the range. The pattern of magnetic polarization was set so that the *outer* surface of the cylinder is the x -axis in an un-rolled frame of reference:

$$J_x = \sin(x), J_y = -\cos(x) \quad [5.1.1]$$

Several magnetization patterns were investigated, and the highest field strength with lowest flux leakage was found to be the distribution [5.1.1] (see section 2.3.1). Shown in figure 5.14 is a design producing a region of constant uniform field gradient along the y-axis, with an angular range of 1.3π radians, $b_i = 0.78$, $b_o = 0.7$, $a_i = 1.0$, $a_o = 1.0$, $L = 10\text{cm}$ and the inner and outer ellipse radial parameters $i_r = 4$ and $o_r = 8$ respectively.




Figure 5.14 – The cross section of the B vector field for the semi-rolled flux sheet design.

The flux can be seen to have good confinement, with low flux leakage from the open upper part of the design. The main field component is in the y direction and can be seen to decrease along the y-axis. The field component B_y and gradient dB_y/dy (taken from positive to negative y) are plotted for this design in figure 5.23:




Figure 5.15 - B_y and dB_y/dz for standard flux sheet design; field is taken from -3.5cm to $+3.5\text{cm}$ along the y -axis (black line in figure 5.14, above).

The region of linear field gradient is marked by the vertical dashed lines; it is 2cm long and is lies between -2cm and 0cm on the y -axis. The region of uniform gradient shown in figure 5.15 has a magnitude of approximately 17 T/m .

5.4.2 Design Summary

We chose a flux sheet with dimensions $b_i = 0.7$, $b_o = 0.7$ radial parameters $i_r = 4$, $o_r = 8$, length $L = 10\text{cm}$, and angular range of 1.3π radians, and examined the effects of individually varying the parameters b_i , b_o , angular range and the length L . Graphs of the field gradient as a function of z for these parameters are shown in appendix A.2.3. We also calculated the standard deviation of dB_y/dy within the working area to asses its uniformity. To measure the effects of changing the semi-minor and semi-major axes of the inner and outer ellipses we fixed the semimajor parameters a_i and a_o at 1 and varied only the semi-minor parameters b_i and b_o .

We investigated the effect of the inner ellipse on the field by varying b_i (figure A.2.9). We found that minimum value of the standard deviation within the working area is at $b_i = 0.78$, with a value of 0.062 T/m (figure A.2.10), corresponding to a uniform gradient close to 17.2 T/m over a length of approximately 2 cm . We also investigated the effect of the outer ellipse on the field gradient by varying b_o (figure A.2.11). and found, somewhat surprisingly, there is a second configuration of the inner and outer ellipse parameters which gives rise to a uniform gradient within the design with a minimum value of standard deviation at $b_o = 0.58$ (figure A.2.12). However, the field

gradient for the configuration $b_o = 0.58$, $b_i = 0.7$ is approximately 17 T/m with a standard deviation of 0.125 T/m, whereas field gradient obtained with the configuration $b_o = 0.7$ and $b_i = 0.78$ is higher, 17.2 T/m, and has a lower standard deviation of 0.062 T/m. Consequently, the following investigations were based on the configuration with $b_o = 0.7$ and $b_i = 0.78$.

We then investigated the effect of changing the angular range for a flux sheet with $b_o = 0.7$ and $b_i = 0.78$ and found that the angular range of the ellipse strongly affects the field within the cylinder and a uniform field gradient only exists for a range of approximately 1.3π radians (figure A.2.13). Finally, we investigated the effect of changing the length using parameters $b_i = 0.7$, $b_i = 0.78$ and found that short flux sheets destroy the uniformity of the field gradient (figure A.2.14). It should also be noted that the effect of changing the length is similar to the effect produced in the offset bore design of section 5.3: short lengths cause the region of uniform field gradient to decrease, suggesting a similar mechanism distorts the field in both designs. The rolled flux sheet with parameters $b_i = 0.78$, $b_o = 0.7$, $a_i = 1.0$, $a_o = 1.0$, $L = 10\text{cm}$ and angular range 1.3π radians, provides an uniform field gradient of around 17 T/m with a standard deviation of 0.062 T/m over a working area of 2cm. As with the Halbach design with the linearly sloping inner bore (section 5.2), this design has a region of constant field gradient only for these parameters. However, this 'standard' design can be scaled to give the field gradient required by multiplying the radial parameters i_r , o_r and the length L by a constant scaling factor s . The ellipticity parameters a_i , a_o , b_i , b_o and the angular range are *not* scaled. The field gradient in this design then scales as $1/s$, as shown in figure 5.16:




Figure 5.16 - The value of the field gradient in the linear region of the design scales as $1/s$.

Extremely high field gradients are therefore achievable for flux sheets with small dimensions.

Figure 5.17 shows the field gradient dB_y/dy on the plane $z = 0$. The field gradient can be seen to be approximately uniform (between 17 and 17.2 T/m) over a region measuring 1 cm on the x-axis by 2 cm on the y-axis.

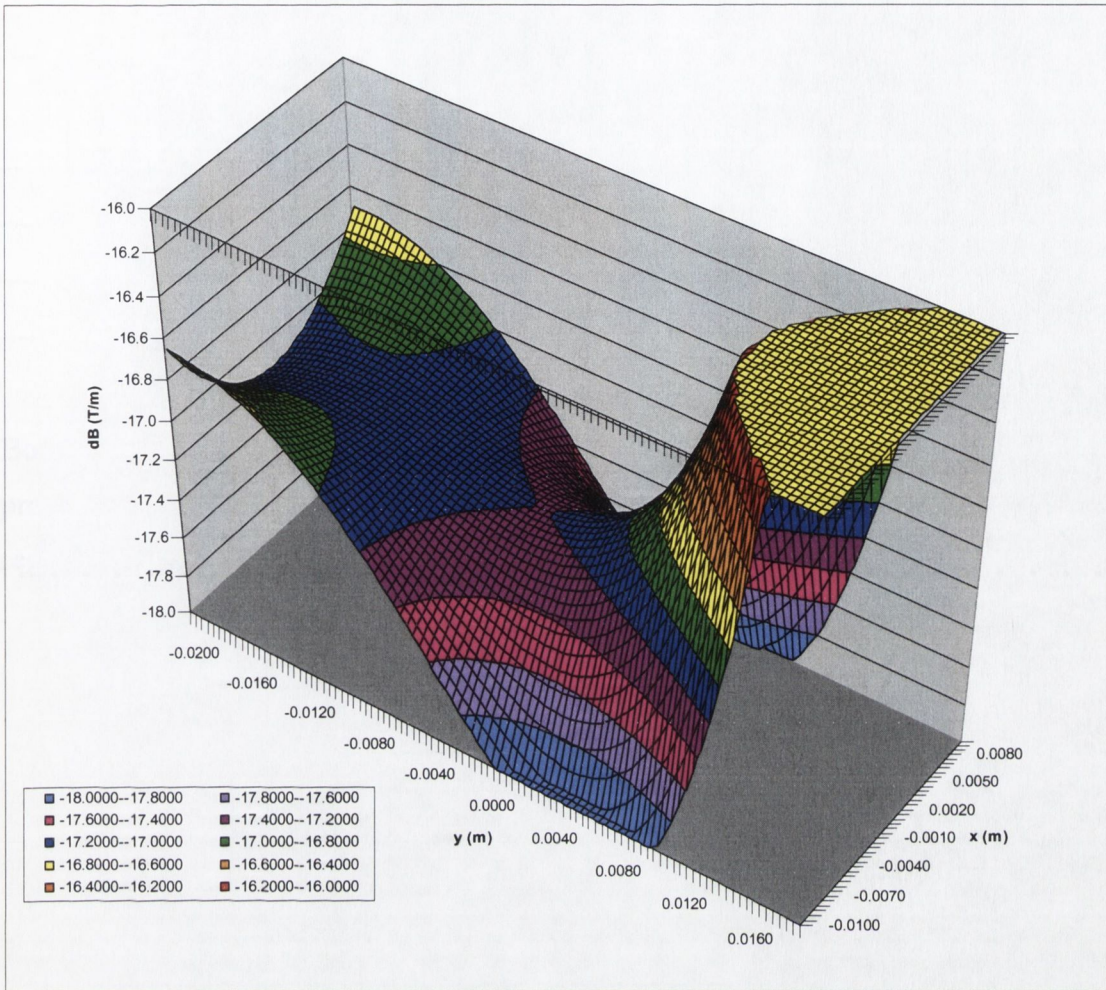


Figure 5.17 - The region of approximately linear field gradient covers a wide area in the central plane of the semi-rolled flux sheet design

5.5 Comparison of Designs Producing a Constant Field Gradient

The design of section 5.2 achieves a constant uniform field gradient orthogonal to the field, while those of sections 5.3 and 5.4 can achieve constant uniform field gradients parallel to the field. We have shown that the field gradient can be tuned by scaling the design, and that extremely high gradients are achievable if the working area is very small. Table 5.1, gives a comparison of the designs considered in section 5.2, 5.3 and 5.4. The magnitude of the gradient achieved within a working region of 1cm is given, as well as the standard deviation within the region, using the optimum values of relevant parameters:

Design	Gradient	dB over 1cm	Stdev	Notes
5.2	dB_y/dz	32	0.063	Segmentation has little effect
5.3	dB_y/dy	26	0.088	Strongly affected by short lengths
5.4	dB_y/dy	34	0.062	Open Design

Table 5.1 – Design summary.

Both the Halbach design of section 5.2 and the flux sheet design of section 5.4 produce field gradients of similar magnitude and similar constancy within the working region.

6 Designs Producing a Constant Product of Field and Field Gradient

6.1 Introduction

We investigated methods of creating a regions within a permanent magnetic design where the product of the field (B) with the field gradient (dB) is uniform. Ideally, we require as large a value of the product BdB as possible over the largest possible volume. To achieve a uniform BdB , the field within the region has to fall off with the square of the distance (for example, z), so that the field magnitude must be of the form:

$$B = \sqrt{a + bz} \quad [6.1.1]$$

Where a and b are constants; Taking the derivative of [6.1] shows that $B(dB/dz)$ is constant:

$$B \frac{dB}{dz} = \frac{b}{2} \quad [6.1.2]$$

Slight variations in the designs of section 4.2 and 4.3 which produce a constant field gradients transform them into magnets which generate regions of constant BdB . We considered three designs; The first is a cylinder with a linearly sloping inner bore, uniformly magnetized parallel to the cylinder axis (figures 6.1, 6.2), giving a field gradient parallel to the field. The second is a Halbach cylinder with a linearly increasing bore, similar to the design of section 4.2, where the field gradient is at right angles to the field (figure 6.8, 6.9). The final design, similar to that of section 4.3, is a partially rolled flux sheet design (figure 6.24). This produces a field gradient parallel to the field.

6.2 Uniformly Magnetized Cylinder With Linearly Varying Inner Radius

6.2.1 Design Overview




Figure 6.1; A uniformly magnetized cylinder with linearly varying inner radius; front and back faces.




Figure 6.2; A uniformly magnetized cylinder with linearly varying inner radius; cross section.

This design, unlike the other designs, can be manufactured as an un-segmented cylinder and subsequently given a uniform magnetization. This fabrication process is much easier than manufacturing separate segments which must be individually magnetized and then bonded together. For this design we use an analytic expression derived from the current model (see appendix A.1.3). However, it will also be modelled in a segmented version, using Magnetic Solver, for comparison.

The cylinder is uniformly magnetized in the z -direction, and an analytic expression for the ideal, non-segmented cylinder can be derived by considering a truncated cone

magnetised along the symmetry axis, with a surface current distribution \mathbf{j} , shown in Figure 6.3.




Figure 6.3 – The current distribution of a truncated cone magnetised parallel to cone axis.

The current model is easier to use here than the charge model, as the charge model replaces the cone by three charged surfaces, whereas the current model only needs a current on one surface, as shown in figure 6.3. The design can then be modelled as two such cones with opposite magnetisation superimposed on top of each other (The outer surface has constant radius, so [6.2.1] can be applied with $\theta = 0$). Along the symmetry axis, the field in the z direction, B_z , is given by:

$$B_z(z) = \frac{J_r \cos(\theta)}{2} \left\{ \frac{1}{\sec^2 \theta} \left(\frac{R \tan \theta + z}{\sqrt{R^2 + z^2}} - \frac{R \tan \theta + L(\tan^2 \theta - 1) + z}{\gamma} \right) + \frac{\tan^2 \theta}{\sec^3 \theta} \left(\ln \left[\gamma + \frac{R \tan \theta + L(\sec^2 \theta) - z}{\sec \theta} \right] - \ln \left[\frac{R \tan \theta - z}{\sec \theta} + \sqrt{R^2 + z^2} \right] \right) \right\} \quad [6.2.1]$$

where R is the smaller radius of the cone, L is the length, θ is the angle of slope of the cone (see figure 6.3) and γ is given by:

$$\gamma = \sqrt{(R + L \tan \theta)^2 + (L - z)^2} \quad [6.2.2]$$

For $\theta = 0$, this reduces to:

$$B_z(0,0,z) = \frac{J_r \cos(\theta)}{2} \left\{ \frac{z}{\sqrt{R^2 + z^2}} - \frac{z - L}{\sqrt{R^2 + (z - L)^2}} \right\} \quad [6.2.3]$$

A similar design constructed using a large number of segments ($N = 128$) instead of a continuous cylinder was modelled using magnetic solver. The field on a cross-section of this design is shown in figure 6.4 in the case of a magnet with dimensions $L = 6\text{cm}$, $R_1 = 1\text{cm}$, $\tan(\theta) = 0.3$ (defining $R_2 = 2.8\text{cm}$) and $R_3 = 6\text{cm}$.




Figure 6.4 – The cross section of a uniformly magnetized cylinder with linearly varying inner radius.

As can be seen, the flux leakage around the cylinder is considerable. The field within the bore is *anti-parallel* to the magnetization, and the gradient falls off from the narrow end of the cylinder to the wide end. Figure 6.5 shows the z component of the field and the product $B_z(dB_z/dz)$ (with the gradient taken from negative to positive z) on the z -axis for this design:




Figure 6.5 - B_z and $B_z(dB_z/dz)$ for uniformly magnetized cylinder with linearly varying inner radius on the z -axis.

The region of uniform BdB is marked by the vertical dashed lines at 2.5cm and 4cm and has magnitude of approximately $4T^2/m$. The ends of the cylinder are marked by the solid vertical lines.

6.2.2 Effect of Segmentation

The field calculated from the analytic expression [6.2.1] matches closely the field produced by the *segmented* version of the design. This is because all the surfaces on which two segments touch are parallel to the magnetisation vector, and so the magnetic charge on them is zero.




Figure 6.6 – Uniformly magnetized cylinder with linearly varying inner radius;
 B_z on z -axis for various numbers of segments, N .

As can be seen from figure 6.6, segmentation has little effect on the field along the bore of this design; the field from a cylinder with eight segments is very similar to a cylinder with 128 segments.

6.2.3 Design Summary

For a magnet with dimensions $R_1 = 1\text{cm}$, $\tan(\theta) = 0.3$ (defining $R_2 = 2.8\text{cm}$), $R_3 = 6\text{cm}$ and $L = 6\text{cm}$, we investigated the effects of varying the slope angle θ , the length L , and the outer radius R_3 one at a time. Graphs of the field B_y and field and field gradient product $B_z(dB_z/dz)$ as a function of z for different $\tan(\theta)$, L and R_3 are shown in appendix A.3.1. We also calculated the standard deviation of $B_z(dB_z/dz)$ within the working area to assess its uniformity. We varied the angle θ (figure A.3.1, A.3.2) and found the minimum value of the standard deviation of the product BdB over the

working region falls at $\tan(\theta) = 0.3$, with a standard deviation of $0.015 \text{ T}^2/\text{m}$ (figure A.3.3) and magnitude of BdB of $4\text{T}^2/\text{m}$. We also found that the magnitude of BdB increases as the length, L , of the cylinder becomes shorter (figure A.3.4), although the only length giving constant BdB within the working region is $L = 6\text{cm}$. Finally, we investigated the effect of changing the outer radius, R_3 (figure A.3.5), and found the product $B_z(dB_z/dz)$ is only constant for $R_3 = 6\text{cm}$ within the working region.

For constant BdB with field and field gradient in the same direction, the design with linear sloping inner bore and uniform magnetization is the simplest design to manufacture, due to the fact that the magnet can be fabricated and magnetized in one piece. This design scales as $1/s$ (in the same way as the other gradient designs considered in section 5) where s , the scaling factor, is a constant multiplying R_1 , R_2 , R_3 and L , shown in figure 6.7:




Figure 6.7– The effect of scaling on uniformly magnetized cylinder with linear varying inner bore.

6.3 Halbach Cylinder With Linearly Varying Inner Radius

6.3.1 Design Overview




Figure 6.8: A Halbach design with linear varying inner radius; front and back faces.




Figure 6.9: Halbach design with linearly varying inner radius; cross section.

The design of section 4.2 successfully produced a region of uniform field gradient for a particular choice of angle θ . We investigated the possibility of producing a region of uniform BdB by varying the design parameters R_2 , R_3 , θ and L and found that a region of constant BdB can be achieved in a design with parameters $\tan(\theta) = 0.23$ (corresponding to $R_2 = 2.6\text{cm}$), $R_3 = 5\text{cm}$, $R_1 = 1\text{cm}$ and $L = 7\text{cm}$, which we will call the standard design. Figure 6.10 shows the field B_y and the product $B_y(dB_y/dz)$ (with the gradient taken from positive to negative z) on the z -axis for a magnet with these parameters:




Figure 6.10 - B_y and $B_y(dB_y/dz)$ for Halbach cylinder with linearly varying inner radius on the z -axis.

The working region in which BdB is approximately uniform is marked above by the two vertical dashed lines at 0cm and 2cm; The magnitude of BdB is approximately 13 T^2/m in this region. This design was modelled using Magnetic Solver (based on the charge model).

The analytic expression for the field on the axis of an un-segmented Halbach cylinder with a linearly increasing inner bore radius is given in section 5.2.1, equation [5.2.1]. Figure 6.11 shows a plot of this expression as a function of z for the parameters values $\tan(\theta) = 0.23$ (corresponding to $R_2 = 2.6\text{cm}$), $R_3 = 5\text{cm}$, $R_1 = 1\text{cm}$ and $L = 7\text{cm}$:




Figure 6.11 - Halbach cylinder with linearly varying inner radius; closed form expression for standard design for B_y and $B_y dB_y/dz$ on z axis.

Comparing figure 6.11 to the field plot of the segmented cylinder, figure 6.10, it can be seen that the segmentation has a beneficial effect on the constancy of BdB in this design. The segmentation appears to increase the length of the region in which BdB is constant.

6.3.2 Design Summary

For a standard cylinder with dimensions $R_1 = 1\text{cm}$, $\tan(\theta) = 0.23$ (corresponding to $R_2 = 2.6\text{cm}$), $R_3 = 5\text{cm}$ and $L = 7\text{cm}$, we examined the effects of varying the slope angle θ , the length L , and the outer radius R_3 individually. Graphs of the field B_y and field and field gradient product $B_y(dB_y/dz)$ as a function of z for various values of these parameters are shown in appendix A.3.2. We also calculated the standard deviation of $B_y(dB_y/dz)$ within the working area to assess its uniformity.

We found that the minimum value of standard deviation of BdB within the working area occurs for a length of $L \approx 7\text{cm}$ and has a value of $0.084 \text{ T}^2/\text{m}$ (figure A.3.7) with magnitude of BdB approximately $13 \text{ T}^2/\text{m}$. We also found that only the value $\tan(\theta) = 0.23$ gives an approximately constant BdB within the working region. Finally, we found that a region of constant gradient can only be achieved for an outer radius $R_3 = 5\text{cm}$ although the *maximum* value of BdB increases for larger outer radii. This can be explained by the fact that although the gradient does not increase for larger outer radii (as found in section 5.3, see figure A.2.5) the field B *does* increase as the outer radius increases

From our results we found that as $\tan(\theta)$ or the length, L , increases, BdB *decreases* along the axis, but as the outer radius increases BdB *increases* along the axis. We investigated whether, by appropriate choice of parameters, we could balance these effects by scaling all of these three variables linearly. This turns out to be the case, scaling the parameters according to [6.3.1] increases the working region, although it has only a small effect on the value of BdB within this region.

$$\begin{aligned} R_3, R_4 &= s'+4 \\ \frac{1}{2}L &= 0.5s'+3 \\ R_2 &= 0.2737s'+2.3363 \end{aligned} \quad [6.3.1]$$

Here s' is a dimensionless scaling parameter and $R_1 = 1\text{cm}$ is unchanged. BdB is shown in figure 6.12 for various values of this scaling parameter. The constants in [6.3.1] are chosen such that $s' = 1$ gives the standard design of $R_3 = 5\text{cm}$, $L = 7\text{cm}$ and

$R_2 = 2.61\text{cm}$ (giving $\tan(\theta)$ as 0.23), with a region of constant BdB . The constants given in [6.3.1] were found by trial and error by examining the BdB curves for a large number of parameter sets.




Figure 6.12 - Halbach cylinder with linearly varying inner radius;
 $B_y(dB_y/dz)$ on z axis for a range of scaling parameters.

The line corresponding to $s' = 1$ in figure 6.12 corresponds to the standard design. As the scaling parameter s' is increased the region of constant BdB increases in length and the magnitude of BdB within the region also increases very slightly. Scaling the design using the expression [6.3.1] therefore provides a way of tuning the region of uniform BdB . Figure 6.13 shows the standard deviation over a working region from 0cm to 2cm for the scaling parameters corresponding to figure 6.12:




Figure 6.13 - Halbach cylinder with linearly varying inner radius;
Standard deviation over working region for range of scaling parameters.

It can be seen that a minimum occurs for a scaling parameter $s' = 3.5$, which is a cylinder with $R_3 = 7.5\text{cm}$, $L = 9.5\text{cm}$ and $R_2 = 3.29\text{cm}$. These parameters greatly improve the uniformity of BdB within the working region with a standard deviation of $0.013 \text{ T}^2/\text{m}$ and magnitude $14 \text{ T}^2/\text{m}$ for $s' = 3.5$, in comparison to $0.087 \text{ T}^2/\text{m}$ and magnitude $13 \text{ T}^2/\text{m}$ for the $s' = 1$ standard design.

The standard design produces a constant BdB of approximately $13 \text{ T}^2/\text{m}$, with the field gradient orthogonal to the field, within the working region. We can also scale the entire design by multiplying the parameters R_1 , R_2 , R_3 and L of the standard design by some constant scaling factor s (which is a different scaling parameter to s') to tune the magnitude of BdB . The effect of scaling the design in this way is shown in figure 6.14:




Figure 6.14 – The effect of scaling the Halbach cylinder with a linearly varying inner radius.

The scaling behaves in the same way as the constant gradient design (section 5.2): scaling the entire design (including the length) affects the magnitude of dB but not the magnitude of B . A high BdB is achievable if the design is scaled down, but this reduces the working region.

6.3.3 Performance of Magnets Built According to Design

We commissioned two cylinders to be manufactured according to the design of section 6.3.1, one (*design 1*) with the standard dimensions $R_1 = 1\text{cm}$, $R_2 = 2.6\text{cm}$, $R_3 = 5\text{cm}$ and length $L = 7\text{cm}$ (shown in figures 6.15 – 6.16) and the second (*design 2*) scaled to double the standard dimensions, $R_1 = 2\text{cm}$, $R_2 = 5.2\text{cm}$, $R_3 = 10\text{cm}$ and length $L = 14\text{cm}$ (shown in figures 6.17 – 6.18). The magnets were manufactured according to our specifications by the China National Electronics Import and Export Ningbo Company, Ningbo, People’s Republic of China.

Photographs and plans for the two ends of design 1 are shown in figures 6.15 - 6.16:




Figure 6.15 – A comparison of the manufactured and planned design 1; front, 7cm length

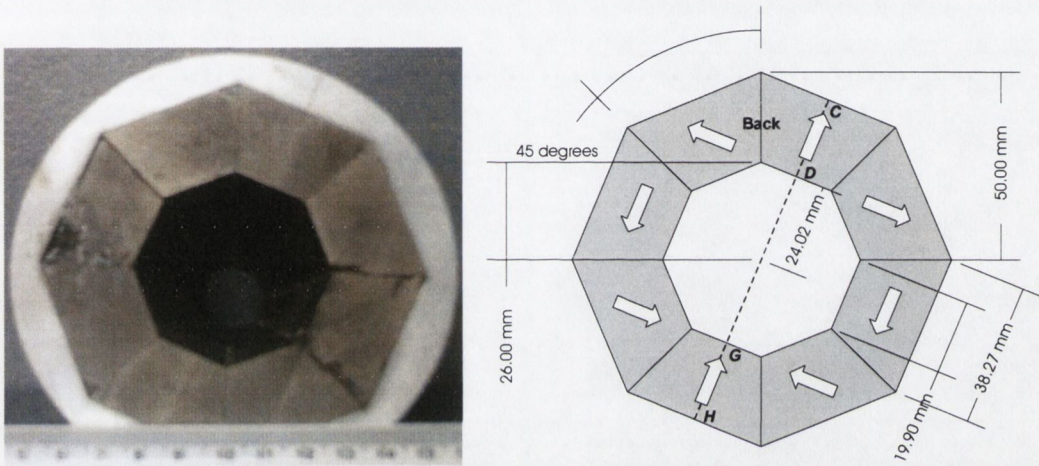


Figure 6.16 – A comparison of the manufactured and planned design 1; front, 7cm length.

Slight imperfections arising during the manufacturing process can be seen in the photographs; there are several chips around where the blocks meet. These manufacturing imperfections do not seem to affect the field, however. The entire magnet was cased within an aluminium cylinder.

Design 2 has double the dimensions of the first design. Photographs and plans for the two ends of this design are shown in figures 6.17 - 6.18:




Figure 6.17- A comparison of the manufactured and planned design 2; front, 14cm length.




Figure 6.18 – A comparison of the manufactured and planned design 2; front, 14cm length.

Design 2 was constructed from three section of equal length which did not fit together to give a completely smooth inner surface to the cylinder, as can be seen in figure 6.18 (left). Although we were concerned this would distort the uniformity of BdB , the actual effect of this ‘stepping’ seems to be negligible. The magnet has a thick aluminium shell which completely encases the design.

Measurements

The fields of both designs were measured using a mounted Hall probe, shown in figure 6.19. The probe was mounted on the vertical (z) axis and the field measured at a range of points along the y -axis. Steps were taken in the z -direction using the winding mechanism shown, and the field recorded at each interval.




Figure 6.19 – The experimental set up for measuring the field from the sloping bore Halbach design.

The measurements on design 1 and 2 are compared with the predictions of Magnetic Solver in the following graph. The error in the Hall probe measurement was 2% of the measured field value.

Figures 6.20 and 6.21 show the measured and predicted field B_y and BdB product for design 1:




Figure 6.20 – The comparison of predicted and measured field B_y for design 1.




Figure 6.21 – The comparison of predicted and measured $B_y(dB_y/dz)$ for design 1.

Figures 6.22 and 6.23 show the measured and predicted field B_y and BdB product for design 2:




Figure 6.22 – The comparison of predicted and measured field B_y for design 2.




Figure 6.23 -The comparison of predicted and measured $B_y(dB_y/dy)$ for design 2.

Discussion

It can be seen from figures 6.20 and 6.23 that the measurements on the manufactured design are in excellent agreement with the field and BdB product predicted using the charge model (BdB was calculated from the experimental data using forward differencing). The tail of the field, at high positive z values, appears to be slightly higher than predicted and the reason for this is unclear, although it could be due to the uncertainty in the exact position of the sampling tip of the Hall probe. A value of 4mm from the end of the probe housing to a mark indicating the position of the Hall sensor was measured. This mark was drawn onto the probe at a later stage, so the true position of the Hall sensor could be a slightly smaller distance from the end of the probe housing, meaning that the measured values could refer to points slightly shifted along the z -axis.

6.4 Rolled Flux Sheet

6.4.1 Design Overview




Figure 6.24 - A design composed of a semi-rolled sheet of one-sided flux of length L .

This design is similar to that in section 5.4, which produced a constant field gradient. Parameters ai , ao , bi , bo , ir , or and the angular range are defined as before (see figure 6.24). We investigated whether we could produce a region of constant BdB within the design by altering the extent to which it is rolled, and the eccentricity of the elliptical cross-sections. Figure 6.25 shows the field produced by a design with parameters $bi = 0.52$, $bo=0.5$, $ir =4$, $or =8$, angular range 1.2π radians and $L = 10\text{cm}$:




Figure 6.25 – The cross section of the standard flux sheet design producing uniform BdB .

The field on the plane $z = 0$ is shown, and was calculated using Magnetic Solver. As can be seen, most of the flux is confined within the fold of the rolled flux sheet. Within the bore the field is mainly in the y -direction, and falls off from the centre towards the opening. The field B_y and the product $B_y(dB_y/dy)$ (with the gradient taken from positive to negative z) in the region from $+1.5$ to -3.5 cm (black line in figure 6.25) is shown in figure 6.26:




Figure 6.26 – B_y and $B_y(dB_y/dz)$ for standard flux sheet design.

The working region, where BdB is approximately constant, is marked by the black dashed vertical lines at -3cm and -1.5cm and has a length of 1.5cm. For the design with parameters given above the magnitude of BdB in the working region is approximately $12 \text{ T}^2/\text{m}$. This design was modelled using Magnetic Solver, based on the charge model.

6.4.2 Design Summary

For a flux sheet with parameters $bi = 0.52$, $bo=0.5$, $ir =4$, $or =8$, angular range 1.2π radians and $L = 10\text{cm}$ we investigated the effect of changing each of these parameters one at a time. Varying the inner ellipse parameter bi (figure A.3.11), we found that the minimum of the standard deviation within the working region occurs at $bi = 0.52$ with a value of $0.070 \text{ T}^2/\text{m}$ and magnitude of BdB of approximately $12\text{T}^2/\text{m}$ (figure A.3.12). We also found that, as for the rolled flux sheet design which produced a constant field gradient (section 5.4), there is a second combination of inner and outer ellipticity that gives rise to a region of constant BdB (figure A.3.13) with $bo = 0.48$. However, the minima in standard deviation occur at $0.116 \text{ T}^2/\text{m}$ for the design with $bi = 0.5$, $bo = 0.48$ and at $0.070 \text{ T}^2/\text{m}$ for the design with $bi = 0.52$, $bo = 0.5$. Also, the magnitude of the product $B_y(dB_y/dy)$ at the minimum standard deviation is $11.5 \text{ T}^2/\text{m}$ for the first case and $12 \text{ T}^2/\text{m}$ in the second. We therefore used the parameters $bi = 0.52$, $bo = 0.5$ in our further investigations.

We found that a larger angular range reduces the magnitude of BdB and also causes it to increase slightly over the working region from within the bore to the opening, whereas a smaller angular range increases the magnitude of BdB , but causes it to decrease from within the bore to the opening (figure A.3.15). It is likely that the reason for the reduced magnitude of BdB in the flux sheets with larger angular range is due to the fact that the flux sheets are almost closed, providing a high field but low field gradient (a fully closed flux sheet is a cylinder which has zero gradient). As the main contribution to BdB comes from the gradient, reducing the gradient will significantly reduce the magnitude of BdB .

Lastly, we found that as the length, L , of the design increases the curve of variation of BdB with y approaches a limiting curve; this curve has a slight linear increase in BdB in the working region from inside the bore to the opening (figure A.3.16). However, $B_y(dB_y/dy)$ is only constant with the working region for $L \approx 10\text{cm}$.

This rolled flux sheet design provides another design which produces a field in the same direction as the field gradient. A uniform BdB of approximately $12 \text{ T}^2/\text{m}$ is obtained for the parameters $a_i = 1.0$, $a_o = 1.0$, $b_i = 0.52$, $b_o = 0.5$, $i_r = 4$, $o_r = 8$, angular range 1.2π radians and $L = 10\text{cm}$ over a working region of 1.5 cm with a standard deviation within this region of 0.07 . The magnitude of BdB can be changed by scaling the entire design, which scales the same way as the rolled flux sheet designed to produce a constant field gradient (section 5.4) as B is not affected by scaling. As with the design of section 5.4, scaling is carried out by multiplying the radial parameters i_r , o_r and the length L by a constant scaling factor s and the ellipticity parameters a_i , a_o , b_i , b_o and the angular range are *not* scaled.




Figure 6.27 – The effect of scaling on rolled flux sheet producing constant BdB .

Figure 6.28 shows the field and field gradient product, $B_y(dB_y/dy)$ on the plane $z = 0$:




Figure 6.28 - The region of uniform *BdB* within the central plane of the rolled flux sheet design.

It can be seen that a region of uniform *BdB* extends from $x = -0.5\text{cm}$ to $+0.5\text{cm}$ on the x -axis and $y = -1\text{cm}$ to -2.5cm on the y -axis, and product *BdB* has a magnitude of approximately $12 \text{ T}^2/\text{m}$ over this region.

6.5 Comparison of Designs Producing a Constant Product of Field and Field Gradient

We have examined three designs capable of producing a uniform product *BdB* over a region of space. Table 6.1, below, gives a comparison of the three designs. The magnitude of *BdB* which can be achieved within a working region of length 1cm is given, as well as the standard deviation of *BdB* for the optimum parameter values.

Design	B	Gradient	BdB for 1cm	Stdev	Notes
6.2	B_z	dB_z/dz	5.8	0.015	Low flux confinement
6.3	B_y	dB_y/dz	26.0	0.084	Standard deviation can be improved using scaling relationship [6.3.1]
6.4	B_y	dB_y/dy	17.9	0.070	Open design

Table 6.1 – Design summary.

The standard deviation of $0.07 \text{ T}^2/\text{m}$ within the working area of the flux sheet design (section 6.4) is an improvement on the Halbach design of (section 6.3), which has a minimum standard deviation of 0.084. The rolled flux sheet design therefore gives a more uniform BdB , although over a slightly smaller working region (1.5cm for this design as opposed to 2cm for the Halbach design.). The magnitude of BdB is also lower within the working region of the flux sheet design compared to the Halbach design.

7 Conclusions

Modern permanent magnetic materials provide opportunities for completely new design strategies, as their magnetic transparency property can be used to greatly simplify the calculation of the fields in devices incorporating them. Much work has been done on magnetic designs based on the geometrical arrangement of magnetic segments in order to produce a uniform field inside some cavity (such as the Halbach cylinder⁹ or designs based on rules given by Abele¹⁸), but most of these designs are based on the 2D ideal infinite length case. As this ideal infinite case could never be realised in practice, it is important to consider these designs in the case of finite length, which typically introduces inhomogeneities into the uniform field within the cavity. Little work has been done on compensating for these inhomogeneities purely through geometrical modification of the design. Furthermore little, if any, work has been done on modifying these designs to produce a uniform field gradient, or a uniform field times field gradient product, within the cavity.

In this project we looked at new designs tailored to produce:

- Improved homogeneity of the field within a cavity.
- A constant and uniform field gradient within a cavity.
- A constant and uniform field times field gradient product within a cavity.

We developed two computational tools allowing us to model the fields from magnetic designs involving both hard and soft magnetic materials. The first of these tools allowed us to simulate the fields from magnetic designs involving only permanent magnets in which the material susceptibility was assumed to be zero. The second of these tools allowed much more freedom in the design, as soft materials could be incorporated. This second program, Manifest, also had a number of advantages over the only commercial magnetic finite element program available to us, MagNet v6, as it allowed for the creation of any shape in 3D space. The MagNet application only allows designs based on a 2D template and extruded linearly or around a circular axis whereas Manifest allows the creation of any closed surface and uses true Delauney meshing. Manifest also uses a method developed by us (see section 3.3.7) for handling thin boundaries in meshes. This allowed us, for example, to investigate the properties of thin cylinders of soft material inside Halbach cylinders. We used these tools to investigate a number of designs composed from modern permanent magnets

using computational techniques, and new designs were presented for the three design classes given above.

For the first of these, the generation of a homogeneous uniform magnetic field within a cavity, we provided a new class of design which gave greatly improved homogeneity over a known design, the split Halbach cylinder, for short length cylinders. This design involved shaping the inner bore of a Halbach cylinder to better confine the flux within the bore. This new design is compared to existing designs in tables 4.1 and 4.2.

For the second design category, involving magnetostatic designs producing a uniform field gradient, we presented three new classes of design. These were the Halbach design with the linearly sloping inner bore, the Halbach design with the offset bore and the semi-rolled flux sheet design. To our knowledge, these designs are all novel; existing classes of designs producing a field gradient usually involve pairs of permanent magnets with opposing faces bridging the working region. In a set-up of this type the magnetic field becomes zero in the centre of the working region. However, in all the designs presented by us this zero value of the field is avoided within the working region. The field gradient magnitude and standard deviation of the gradient over the working region for each of the three new designs are compared to each other in table 5.1.

For the final category of design, involving magnetostatic designs producing a uniform field and field gradient product (BdB), three classes of design were presented. These were the Halbach with a linearly sloping inner bore, the uniformly magnetized cylinder with a linearly sloping inner bore and the semi-rolled flux sheet design. Again, to our knowledge, all designs are novel. Regions of constant BdB are typically produced using electromagnets with shaped pole pieces, and our designs provide a clear improvement over these in terms of power consumption, size and weight. Two scalings of the Halbach design with the sloping inner bore were constructed and it was shown that our predictions closely matched the measured fields from the designs. The magnitude of the product of the field and field gradient, BdB , and standard deviation of BdB over the working region for each of these three designs are compared in table 6.1.

The designs are summarised in the table below. The offset bore design and the uniformly magnetized cylinder have both been omitted from the table. These designs

both have disadvantages over other designs presented which produce the same field patterns; the offset bore design requires a cylinder of long length to produce a high constant gradient, and the uniformly magnetized cylinder has a low magnitude BdB with a large stray field.

Application Requirement	Design	Section
High Field Homogeneity	Halbach Cylinder With Non-Uniform Inner Radius	4.3
Field Gradient Orthogonal to Field	Halbach Cylinder With Linearly Sloping Inner Bore	5.2
Field Gradient Parallel to Field	Gradient Flux Sheet	5.4
BdB , Field Gradient Orthogonal to Field	Halbach Cylinder With Linearly Sloping Inner Bore	6.3
BdB , Field Gradient Parallel to Field	BdB Flux Sheet	6.4

Table 7.1 – Summary of all designs.

This work provides a reference point for anyone constructing an application requiring any of the field configurations covered above. The computational tools developed are also intended to be used for future work in magnetic design and can be developed further, for example, we intend include a full non-linear curve for soft materials in the Manifest program. Furthermore, the mesh-generation and matrix manipulation routines used in Manifest are stand-alone C++ modules and can be used in any computational application requiring the use of triangulation or linear algebra routines. We also intend to investigate various geometrical designs further and attempt to develop a mechanical *variable* field gradient design based on our design using the semi-rolled flux sheet.

Appendix

A.1 Derivations

A.1.1 Rectangular Charge Sheet

The expression for a charged sheet in the x-z plane is derived:




Figure A.1.1 - The rectangular charge sheet used in the analytic model.

The general form of the potential is given by:

$$\varphi_m(x) = \frac{1}{4\pi\mu_0} \oint_S \frac{\sigma_m}{|x - x'|} dS' \quad [A.1.1]$$

With:

$$\sigma_m = \mathbf{M} \cdot \hat{\mathbf{n}} \quad [A.1.2]$$

Where \mathbf{n} is the unit outward normal from the surface. For the charged sheet described above, [A.1.1] becomes:

$$\varphi_m(P) = \frac{\sigma_m}{4\pi\mu_0} \int_{-a-L}^a \int_{-L}^L \frac{1}{\sqrt{(x-x')^2 + y^2 + (z-z')^2}} dz' dx' \quad [A.1.3]$$

Where a is the half-width of the sheet and L the half-length of the sheet (see figure 2.1). We can make the integral in [A.1.3] easier by operating with the del operator, since from [2.1.27]:

$$H = -\nabla \varphi_m \quad [2.1.27]$$

Giving:

$$H_x(P) = \frac{\sigma_m}{4\pi\mu_0} \int_{-a-L}^a \int_{-L}^L \frac{(x-x')}{[(x-x')^2 + y^2 + (z-z')^2]^{\frac{3}{2}}} dz' dx' \quad [A.1.4]$$

$$H_y(P) = \frac{\sigma_m}{4\pi\mu_0} \int_{-a-L}^a \int_{-L}^L \frac{y}{[(x-x')^2 + y^2 + (z-z')^2]^{\frac{3}{2}}} dz' dx' \quad [A.1.5]$$

$$H_z(P) = \frac{\sigma_m}{4\pi\mu_0} \int_{-a-L}^a \int_{-L}^L \frac{(z-z')}{[(x-x')^2 + y^2 + (z-z')^2]^{\frac{3}{2}}} dz' dx' \quad [A.1.6]$$

These give:

$$H_x(P) = \frac{\sigma_m}{4\pi\mu_0} \int_{-L}^L \frac{1}{\sqrt{(x-a)^2 + y^2 + (z-z')^2}} - \frac{1}{\sqrt{(x+a)^2 + y^2 + (z-z')^2}} dz' \quad [A.1.7]$$

We use:

$$\int \frac{1}{\sqrt{a^2 + u^2}} du = \ln(u + \sqrt{u^2 + a^2}) + C \quad [A.1.8]$$

Which gives:

$$H_x(P) = \frac{\sigma_m}{4\pi\mu_0} \left[\ln \left(\frac{(z+L) + \sqrt{(z+L)^2 + (x-a)^2 + y^2}}{(z-L) + \sqrt{(z-L)^2 + (x-a)^2 + y^2}} \right) + \ln \left(\frac{(z+L) + \sqrt{(z+L)^2 + (x+a)^2 + y^2}}{(z-L) + \sqrt{(z-L)^2 + (x+a)^2 + y^2}} \right) \right] \quad [A.1.9]$$

With a similar expression for H_z :

$$H_z(P) = \frac{\sigma_m}{4\pi\mu_0} \left[\ln \left(\frac{(x+a) + \sqrt{(x+a)^2 + (z-L)^2 + y^2}}{(x-a) + \sqrt{(x-a)^2 + (z-L)^2 + y^2}} \right) + \ln \left(\frac{(x+a) + \sqrt{(x+a)^2 + (z+L)^2 + y^2}}{(x-a) + \sqrt{(x-a)^2 + (z+L)^2 + y^2}} \right) \right] \quad [A.1.10]$$

The expression for H_y is slightly more difficult to evaluate:

$$H_y(P) = \frac{\sigma_m}{4\pi\mu_0} \left[\int_{-L}^L \frac{(x-a)}{(y^2 + (z-z')^2)\sqrt{(x-a)^2 + y^2 + (z-z')^2}} dz' + \int_{-L}^L \frac{(x+a)}{(y^2 + (z-z')^2)\sqrt{(x+a)^2 + y^2 + (z-z')^2}} dz' \right] \quad [A.1.11]$$

We use:

$$\int \frac{1}{(a+u^2)\sqrt{b+u^2}} du = \frac{1}{\sqrt{a(b-a)}} \tan^{-1} \left(\frac{u\sqrt{b-a}}{\sqrt{a(b+u^2)}} \right) + C \quad [A.1.12]$$

Which gives:

$$H_y(P) = \frac{\sigma_m}{4\pi\mu_0 y} \left[\begin{aligned} & \tan^{-1} \left(\frac{(z+L)(x-a)}{y\sqrt{(x-a)^2 + y^2 + (z+L)^2}} \right) - \\ & \tan^{-1} \left(\frac{(z-L)(x-a)}{y\sqrt{(x-a)^2 + y^2 + (z-L)^2}} \right) + \\ & \tan^{-1} \left(\frac{(z+L)(x+a)}{y\sqrt{(x+a)^2 + y^2 + (z+L)^2}} \right) - \\ & \tan^{-1} \left(\frac{(z-L)(x+a)}{y\sqrt{(x+a)^2 + y^2 + (z-L)^2}} \right) \end{aligned} \right] \quad [A.1.13]$$

A.1.2 Triangular Charge Sheet




Figure A.1.2 - The right-angled triangular charge sheet used in the analytic model.

The general form of the potential is given by:

$$\varphi_m(x) = \frac{1}{4\pi\mu_0} \int_S \frac{\sigma_m}{|x-x'|} dS' \quad [A.2.1]$$

With:

$$\sigma_m = \mathbf{M} \cdot \hat{\mathbf{n}} \quad [A.2.2]$$

Where \mathbf{n} is the unit outward normal from the surface. For the triangular sheet described above, [A.2.1] becomes:

$$\varphi_m(P) = \frac{\sigma_m}{4\pi\mu_0} \int_0^a dx' \int_0^{\frac{b}{a}x'} dz' \frac{1}{\sqrt{(x-x')^2 + y^2 + (z-z')^2}} \quad [A.2.3]$$

Where a and b are the triangular dimensions, shown in figure A.2 above.

H_x and H_z Field Components

The field components for H_x and H_z can be found by using:

$$\mathbf{H} = -\nabla\varphi_m \quad [A.2.4]$$

on [A.2.3] before evaluating the integral. For H_z , this gives:

$$H_z(P) = \frac{\sigma_m}{4\pi\mu_0} \int_0^a dx' \int_0^{\frac{b}{a}x'} dz' \frac{(z-z')}{\left[(x-x')^2 + y^2 + (z-z')^2 \right]^{\frac{3}{2}}} \tag{A.2.5}$$

The first integral over z' can be evaluated straightforwardly to give:

$$H_z(P) = \frac{\sigma_m}{4\pi\mu_0} \int_0^a dx' \left[\frac{1}{\sqrt{(x-x')^2 + y^2 + z^2}} - \frac{1}{\sqrt{(x-x')^2 + y^2 + (z-\frac{b}{a}x)^2}} \right] \tag{A.2.6}$$

The integral in this expression was evaluated using *Mathematica*, giving:

$$H_z(P) = \frac{\sigma_m}{2\pi\mu_0} \left[\ln \left(\frac{r-x}{a + \sqrt{(a-x)^2 + y^2 + z^2 - x}} \right) + \frac{a}{\sqrt{a^2 + b^2}} \ln \left(\frac{\frac{a(a-x)+b(b-z)}{\sqrt{a^2+b^2}} + \sqrt{(a-x)^2 + y^2 + (b-z)^2}}{r - \frac{ax+bz}{\sqrt{a^2+b^2}}} \right) \right] \tag{A.2.7}$$

Where r is given by:

$$r = \sqrt{x^2 + y^2 + z^2} \tag{A.2.8}$$

The expression for H_x can be evaluated in a similar way by changing the order of integration in [A.2.3]:

$$\varphi_m(P) = -\frac{\sigma_m}{4\pi\mu_0} \int_0^b dz' \int_a^{\frac{a}{b}z'} dx' \frac{1}{\sqrt{(x-x')^2 + y^2 + (z-z')^2}} \tag{A.2.9}$$

Using [A.2.4] and taking the integral over x' gives:

$$H_x(P) = \frac{\sigma_m}{4\pi\mu_0} \int_0^b dz' \left[\frac{1}{\sqrt{(a-x)^2 + y^2 + (z-z')^2}} - \frac{1}{\sqrt{(x-\frac{a}{b}z)^2 + y^2 + (z-z')^2}} \right] \tag{A.2.10}$$

The integral in this expression was also evaluated using *Mathematica*, giving:

$$H_x(P) = \frac{\sigma_m}{2\pi\mu_0} \left[\ln \left(\frac{b + \sqrt{(a-x)^2 + y^2 + (b-z)^2} - z}{\sqrt{(a-x)^2 + y^2 + z^2} - z} \right) + \frac{b}{\sqrt{a^2 + b^2}} \ln \left(\frac{r - \frac{ax+bz}{\sqrt{a^2+b^2}}}{\frac{a(a-x)+b(b-z)}{\sqrt{a^2+b^2}} + \sqrt{(a-x)^2 + y^2 + (b-z)^2}} \right) \right] \tag{A.2.11}$$

H_y Field Component

It is more difficult to determine the H_y field component as *Mathematica* does not give a simple expression for the integral of the y -component of the gradient of [A.2.3]. In this case, we must evaluate the integral over z' in [A.2.3] *first*, then take the gradient of the result. The integral over z' in [A.2.3] is:

$$\varphi_m(P) = \frac{\sigma_m}{4\pi\mu_0} \int_0^a dx' \left[\begin{aligned} &\ln\left(z + \sqrt{(x-x')^2 + y^2 + z^2}\right) - \\ &\ln\left(z - \frac{b}{a}x' + \sqrt{(x-x')^2 + y^2 + \left(z - \frac{b}{a}x'\right)^2}\right) \end{aligned} \right] \quad [A.2.12]$$

The y -component of the gradient of this expression gives H_y as:

$$H_y(P) = \frac{\sigma_m}{4\pi\mu_0} y \int_0^a dx' \left[\begin{aligned} &\frac{1}{\left(z + \sqrt{(x-x')^2 + y^2 + z^2}\right)\sqrt{(x-x')^2 + y^2 + z^2}} + \\ &\frac{1}{\left(z - \frac{b}{a}x' + \sqrt{(x-x')^2 + y^2 + \left(z - \frac{b}{a}x'\right)^2}\right)\sqrt{(x-x')^2 + y^2 + \left(z - \frac{b}{a}x'\right)^2}} \end{aligned} \right] \quad [A.2.13]$$

The first of these integrals can be evaluated by using the substitutions:

$$(x-x') = p \sinh \theta, \quad p^2 = y^2 + z^2 \quad [A.2.14]$$

The integral then becomes:

$$\int dx' \frac{1}{\left(z + \sqrt{(x-x')^2 + y^2 + z^2}\right)\sqrt{(x-x')^2 + y^2 + z^2}} = -\int d\theta \frac{1}{z + p \cosh \theta} \quad [A.2.15]$$

Using:

$$\int d\theta \frac{1}{z + p \cosh \theta} = \frac{2}{\sqrt{p^2 + z^2}} \tan^{-1} \left(\frac{pe^\theta + z}{\sqrt{p^2 + z^2}} \right) \quad [A.2.16]$$

Gives:

$$\begin{aligned} &\int dx' \frac{1}{\left(z + \sqrt{(x-x')^2 + y^2 + z^2}\right)\sqrt{(x-x')^2 + y^2 + z^2}} \\ &= \frac{2}{y} \tan^{-1} \left\{ \frac{(x-x') + \sqrt{(x-x')^2 + y^2 + z^2} + z}{y} \right\} \end{aligned} \quad [A.2.17]$$

Putting in the limits $a, 0$ gives:

$$\begin{aligned} &\int_0^a \frac{1}{\left(z + \sqrt{(x-x')^2 + y^2 + z^2}\right)\sqrt{(x-x')^2 + y^2 + z^2}} \\ &= \frac{2}{y} \left\{ \tan^{-1} \left(\frac{(x-a) + \sqrt{(x-a)^2 + y^2 + z^2} + z}{y} \right) - \tan^{-1} \left(\frac{x+r+z}{y} \right) \right\} \end{aligned} \quad [A.2.18]$$

Where r is defined in [A.2.8]. The second integral in [A.2.13] is more difficult to evaluate, first we look for a form:

$$(x-x')^2 + y^2 + \left(z - \frac{b}{a}x'\right)^2 = \alpha^2(x'-\beta)^2 + \alpha^2\gamma^2 \quad [A.2.19]$$

Which gives α, β , and γ as:

$$\alpha = \sqrt{1 + \frac{b^2}{a^2}}, \beta = -\frac{x + \frac{b}{a}z}{1 + \frac{b^2}{a^2}}, \gamma = \sqrt{-\beta^2 + \frac{r^2}{1 + \frac{b^2}{a^2}}} \quad [A.2.20]$$

The second integral in [A.2.13] then becomes:

$$\int dx' \frac{1}{\left(z - \frac{b}{a}x' + \sqrt{(x-x')^2 + y^2} + \left(z - \frac{b}{a}x'\right)^2\right) \sqrt{(x-x')^2 + y^2 + \left(z - \frac{b}{a}x'\right)^2}} \quad [A.2.21]$$

$$= \int dx' \frac{1}{\left(z - \frac{b}{a}x' + \alpha \sqrt{(x'-\beta)^2 + \gamma^2}\right) \alpha \sqrt{(x'-\beta)^2 + \gamma^2}}$$

We can then use the same form of substitution [A.2.14] as before:

$$x' - \beta = \gamma \sinh \theta \quad [A.2.22]$$

This gives:

$$\int dx' \frac{1}{\left(z - \frac{b}{a}x' + \alpha \sqrt{(x'-\beta)^2 + \gamma^2}\right) \alpha \sqrt{(x'-\beta)^2 + \gamma^2}} \quad [A.2.23]$$

$$= \frac{1}{\alpha} \int \frac{1}{A \sinh \theta + B \cosh \theta + C}$$

where:

$$A = -\gamma \frac{b}{a}, B = \alpha \gamma, C = z + \beta \frac{b}{a} \quad [A.2.24]$$

The using hyperbolic functions in the denominator of the integral can be expressed as:

$$\int d\theta \frac{1}{A \sinh \theta + B \cosh \theta + C} = \int d\theta \frac{1}{\frac{1}{2}(B-A)e^{-\theta} + \frac{1}{2}(A+B)e^{\theta} + C} \quad [A.2.25]$$

Which can be evaluated using *Mathematica* (although it is fairly easy to evaluate by multiplying top and bottom by e^{θ} , using the substitution $u = e^{\theta}$ and taking the integral over u):

$$\int d\theta \frac{1}{\frac{1}{2}(B-A)e^{-\theta} + \frac{1}{2}(A+B)e^{\theta} + C} \quad [A.2.26]$$

$$= \frac{2}{\sqrt{B^2 - A^2 - C^2}} \tan^{-1} \left(\frac{C + (A+B)e^{\theta}}{\sqrt{B^2 - A^2 - C^2}} \right)$$

So:

$$\int dx' \frac{1}{\left(z - \frac{b}{a}x' + \sqrt{(x-x')^2 + y^2} + \left(z - \frac{b}{a}x'\right)^2\right) \sqrt{(x-x')^2 + y^2 + \left(z - \frac{b}{a}x'\right)^2}} \quad [A.2.27]$$

$$= \frac{2}{\alpha \sqrt{B^2 - A^2 - C^2}} \tan^{-1} \left(\frac{C + (A+B) \left[\frac{x'+\beta}{\gamma} + \sqrt{\left(\frac{x'+\beta}{\gamma}\right)^2 + 1} \right]}{\sqrt{B^2 - A^2 - C^2}} \right)$$

Putting in the limits $a, 0$ into this expression and putting all the parts [A.2.6] together gives the full expression for H_y as:

$$\begin{aligned}
 H_y(P) = \frac{\sigma_m}{2\pi\mu_0} \left[\right. & \tan^{-1} \left(\frac{(x-a) + z + \sqrt{(x-a)^2 + y^2 + z^2}}{y} \right) - \tan^{-1} \left(\frac{x+z+r}{y} \right) + \\
 & \frac{y}{\alpha\sqrt{B^2 - A^2 - C^2}} \left\{ \tan^{-1} \left(\frac{C + (A+B) \left(\sqrt{1 + \frac{(\alpha+\beta)^2}{\gamma^2}} + \frac{\alpha+\beta}{\gamma} \right)}{\sqrt{B^2 - A^2 - C^2}} \right) - \right. \\
 & \left. \left. \tan^{-1} \left(\frac{C + (A+B) \left(\sqrt{1 + \frac{\beta^2}{\gamma^2}} + \frac{\beta}{\gamma} \right)}{\sqrt{B^2 - A^2 - C^2}} \right) \right\} \right] \quad [A.2.28]
 \end{aligned}$$

A.1.3 Cylindrical Halbach With Linear Inner Bore




Figure A.1.3 – A Halbach cylinder with linearly varying inner bore.

The magnetization of an ideal Halbach cylinder, in cylindrical co-ordinates, is given by:

$$\mathbf{M}(\phi') = M_r \{ \sin(\phi') \hat{\rho} - \cos(\phi') \hat{\phi} \} \quad [A.3.1]$$

where the magnetic volume charge density ρ_m is given by [2.2.8] (see section 2.2.1):

$$\rho_m = -\nabla \cdot \mathbf{M} \quad [A.3.2]$$

The magnetic volume charge density is therefore:

$$\rho_m(\rho', \phi') = -2M_r \frac{\sin \phi'}{\rho'} \quad [A.3.3]$$

where the primed co-ordinates represent source points. Using the expression for magnetic potential from a volume charge distribution (see section 2.2.1):

$$\varphi_m(\mathbf{x}) = \frac{1}{4\pi} \int \frac{\rho_m'}{|\mathbf{x} - \mathbf{x}'|} d\tau' \quad [A.3.4]$$

gives:

$$\varphi_m(\mathbf{r}) = -\frac{J_r}{2\pi\mu_0} \int dz' \int dx' \int dy' \frac{\sin \phi'}{\rho'} \frac{1}{|\mathbf{r} - \mathbf{r}'|} \quad [A.3.5]$$

Where we have used [2.1.52b] for J_r , the magnetic remnant polarization of the material, in Tesla. We are interested in the magnetic field in the y -direction within the cylinder bore, H_y . Using:

$$\mathbf{H} = -\nabla \varphi_m \quad [A.3.6]$$

gives:

$$H_y(x, y, z) = \frac{J_r}{2\pi\mu_0} \int dz' \int dx' \int dy' \frac{\sin \phi'}{\rho'} \frac{\partial}{\partial y} \frac{1}{\sqrt{(x-x')^2 + (y-y')^2 + (z-z')^2}} \quad [A.3.7]$$

which becomes:

$$H_y(x, y, z) = \frac{J_r}{2\pi\mu_0} \int dz' \int dx' \int dy' \frac{\sin \phi'}{\rho'} \frac{(y-y')}{\left[(x-x')^2 + (y-y')^2 + (z-z')^2 \right]^{\frac{3}{2}}} \quad [A.3.8]$$

In cylindrical co-ordinates:

$$x' = \rho' \sin \phi', \quad y' = \rho' \cos \phi' \quad [A.3.9]$$

So:

$$(x-x')^2 + (y-y')^2 + (z-z')^2 = \rho^2 + \rho'^2 - 2\rho\rho' \cos(\phi - \phi') + (z-z')^2 \quad [A.3.10]$$

which gives:

$$H_y(x, y, z) = \frac{J_r}{2\pi\mu_0} \int dz' \int dx' \int dy' \frac{\sin \phi'}{\rho'} \frac{(y - \rho' \sin \phi')}{\left[\rho^2 + \rho'^2 - 2\rho\rho' \cos(\phi - \phi') + (z-z')^2 \right]^{\frac{3}{2}}} \quad [A.3.11]$$

On the z -axis, $\rho = 0$ (and $y = 0$). Therefore the expression [A.3.11] for the field H_y on the z -axis is:

$$H_y(0,0,z) = \frac{J_r}{2\pi\mu_0} \int dz' \int dx' \int dy' \frac{\sin^2 \phi'}{\left[\rho'^2 + (z-z')^2 \right]^{\frac{3}{2}}} \quad [A.3.12]$$

Converting to cylindrical co-ordinates gives:

$$H_y(0,0,z) = \frac{J_r}{2\pi\mu_0} \int dz' \int d\rho' \frac{\rho'}{\left[\rho'^2 + (z-z')^2 \right]^{\frac{3}{2}}} \int d\phi' \sin^2 \phi' \quad [A.3.13]$$

For a Halbach cylinder with a variable inner bore, this is:

$$H_y(0,0,z) = \frac{J_r}{2\mu_0} \int_0^L dz' \int_{R(z')}^{R_3} d\rho' \frac{\rho'}{\left[\rho'^2 + (z-z')^2 \right]^{\frac{3}{2}}} \quad [A.3.14]$$

Taking the integral over ρ gives:

$$H_y(0,0,z) = \frac{J_r}{2\mu_0} \int_0^L dz' \left[\frac{1}{\sqrt{R(z')^2 + (z-z')^2}} + \frac{1}{\sqrt{R_3^2 + (z-z')^2}} \right] \quad [A.3.15]$$

The second integral can be solved straightforwardly, giving:

$$H_y(0,0,z) = \frac{J_r}{2\mu_0} \left[\ln \left| \frac{\sqrt{R_3^2 + z^2} - z}{\sqrt{R_3^2 + (L-z)^2} + L-z} \right| + \int_0^L dz' \frac{1}{\sqrt{R(z')^2 + (z-z')^2}} \right] \quad [A.3.16]$$

For a linear increase in the inner radius:

$$R(z') = R_1 + z' \tan \theta \quad [A.3.17]$$

The second integral in [A.3.16] can then be expressed as:

$$\int_0^L dz' \frac{1}{\sqrt{R(z')^2 + (z-z')^2}} = \frac{1}{\sqrt{1 + \tan^2 \theta}} \int_0^L dz' \frac{1}{\sqrt{(z'+\alpha)^2 + \gamma^2}} \quad [A.3.18]$$

where:

$$\alpha = \frac{R_1 \tan \theta - z}{1 + \tan^2 \theta} \quad [A.3.19]$$

and:

$$\gamma = \sqrt{\frac{R_1 - z \tan \theta}{1 + \tan^2 \theta}} \quad [A.3.20]$$

The integral in the expression [A.3.18] can be solved to give:

$$\int_0^L dz' \frac{1}{\sqrt{(z'+\alpha)^2 + \gamma^2}} = \ln \left| \frac{\sqrt{(L+\alpha)^2 + \gamma^2} + L + \alpha}{\sqrt{\alpha^2 + \gamma^2} + \alpha} \right| \quad [A.3.21]$$

This gives:

$$H_y(0,0,z) = \frac{J_r}{2\mu_0} \left[\ln \left| \frac{\sqrt{R_3^2 + z^2} - z}{\sqrt{R_3^2 + (L-z)^2} + L-z} \right| + \frac{1}{\sqrt{1 + \tan^2 \theta}} \ln \left| \frac{\sqrt{(L+\alpha)^2 + \gamma^2} + L + \alpha}{\sqrt{\alpha^2 + \gamma^2} + \alpha} \right| \right] \quad [A.3.22]$$

A.1.4 Field From a Conical Magnet




Figure A.1.4 – A truncated conical magnet with uniform magnetization.

The current model gives the field B from a uniformly magnetized permanent magnet as:

$$B(\mathbf{x}) = \frac{\mu_0}{4\pi} \oint_S \mathbf{j}_m(\mathbf{x}') \times \frac{(\mathbf{x} - \mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|^3} dS' \quad [A.4.1]$$

Where the surface current density \mathbf{j}_m is given by:

$$\mathbf{j}_m(\mathbf{x}') = \mathbf{M} \times \hat{\mathbf{n}} \quad [A.4.2]$$

Where $\hat{\mathbf{n}}$ is the unit outward normal from the surface and \mathbf{M} the fixed magnetization. For the conical block shown in figure A.1, the surface current, in cylindrical coordinates, is then given by:

$$\mathbf{j}_m(\mathbf{x}') = \frac{J_r}{\mu_0} \cos(\theta) \hat{\phi} \quad [A.4.3]$$

Where we have used [2.1.52b] for J_r , the magnetic remnant polarization of the material, in Tesla. The vectors \mathbf{x} and \mathbf{x}' are given by:

$$\begin{aligned} \mathbf{x} &= z\hat{\mathbf{z}} \\ \mathbf{x}' &= (R + z' \tan \theta)\hat{\mathbf{r}} + \phi'\hat{\phi} + z'\hat{\mathbf{z}} \end{aligned} \quad [A.4.4]$$

Giving:

$$|\mathbf{x} - \mathbf{x}'| = \sqrt{(R + z' \tan \theta)^2 + (z - z')^2} \quad [A.4.5]$$

Substituting into A.x gives:

$$B(\mathbf{x}) = \frac{1}{4\pi} \int_0^L \int_0^{2\pi} J_r \cos(\theta) \hat{\phi} \times \frac{(R + z' \tan \theta)\hat{\mathbf{r}} + \phi'\hat{\phi} + (z' - z)\hat{\mathbf{z}}}{[(R + z' \tan \theta)^2 + (z - z')^2]^{\frac{3}{2}}} (R + z' \tan \theta) d\phi' dz' \quad [A.4.6]$$

Where the extra term $(R + z' \tan \theta)$ comes from the extra radial term when using cylindrical co-ordinates. The cross product gives:

$$B(\mathbf{x}) = \frac{1}{4\pi} \int_0^L \int_0^{2\pi} \begin{vmatrix} \hat{\mathbf{r}} & \hat{\boldsymbol{\phi}} & \hat{\mathbf{z}} \\ 0 & J_r \cos(\theta) & 0 \\ (R + z' \tan \theta) & \phi' & (z' - z) \end{vmatrix} \frac{(R + z' \tan \theta)}{\left[(R + z' \tan \theta)^2 + (z - z')^2 \right]^{\frac{3}{2}}} d\phi' dz' \quad [A.4.7]$$

We expect the field along the symmetry axis to be mainly confined to the z-direction (especially for our application of the expression which involves cavities inside cylinders). Therefore, we can concentrate on the z-component of the above expression, which is:

$$B_z(\mathbf{x}) = \frac{J_r \cos(\theta)}{2} \int_0^L \frac{(R + z' \tan \theta)^2}{\left[(R + z' \tan \theta)^2 + (z - z')^2 \right]^{\frac{3}{2}}} dz' \quad [A.4.8]$$

Where we have also evaluated the integral over ϕ . The above expression can be solved to give:

$$B_z(z) = \frac{J_r \cos(\theta)}{2} \left\{ \frac{1}{1 + \tan^2 \theta} \left(\frac{R \tan \theta + z}{\sqrt{R^2 + z^2}} - \frac{R \tan \theta + L(\tan^2 \theta - 1) + z}{\gamma} \right) + \frac{\tan^2 \theta}{(1 + \tan^2 \theta)^{\frac{3}{2}}} \left(\ln \left[\gamma + \frac{R \tan \theta + L(1 + \tan^2 \theta) - z}{\sqrt{1 + \tan^2 \theta}} \right] - \ln \left[\frac{R \tan \theta - z}{\sqrt{1 + \tan^2 \theta}} + \sqrt{R^2 + z^2} \right] \right) \right\} \quad [A.4.9]$$

Where:

$$\gamma = \sqrt{(R + L \tan \theta)^2 + (L - z)^2} \quad [A.4.10]$$

A. 2 Effect of Varying Parameters For Designs In Section 5

A.2.1 Design of Section 5.2

We chose a cylinder of dimensions $L = 10\text{cm}$, $R_3 = 5\text{cm}$, $R_I = 1\text{cm}$ and the angle θ , defining R_2 (see figure 5.2), was varied. Figure A.2.1 shows B_y on the cylinder axis for a range of $\tan(\theta)$. The solid lines at $+5\text{cm}$ and -5cm show the position of the ends of the cylinder, and the working region of approximately constant gradient is corresponding marked by the two vertical dashed lines.




Figure A.2.1 – Halbach design with linear inner bore of varying slope; B_y on z axis within bore.

Figure A.2.2 shows $\frac{dB_y}{dz}$ on the z -axis for a section from $z = -1.5\text{ cm}$ to $z = 2.5\text{ cm}$ within the cylinder bore(taken from positive to negative z):




Figure A.2.2 - Halbach design with inner bore of linearly varying slope; dB/dz on the z axis within bore.

The standard deviation of the field gradient with the working region is plotted as a function of $\tan(\theta)$ in figure A.2.3:




Figure A.2.3 - Halbach design with inner bore of linearly varying slope; standard deviation of dB/dz over working region.

We examined how changing the length of the design affected the field gradient over the working region, taking $L = 10\text{cm}$, $R_3 = 5\text{cm}$, $R_1 = 1\text{cm}$, $\tan(\theta) = 0.23$ and varying the length:




Figure A.2.4 - Halbach design with linear sloping inner bore of varying length; $\frac{dB_y}{dz}$ on the z axis within bore.

It can be seen that a long region of uniform gradient only exists for a particular length ($L = 10\text{cm}$). We then varied the outer radius of the cylinder R_3 , taking $\tan(\theta) = 0.23$, $L = 10\text{cm}$ and $R_1 = 1\text{cm}$:




Figure A.2.5 - Halbach design with linear sloping inner bore of varying R_3 ; $\frac{dB_y}{dz}$ on the z axis within bore.

Figure A.2.5 shows the gradient of the B_y component of the field for R_3 ranging from 4 to 8.5cm. With the parameters $L = 10\text{cm}$, $R_1 = 1\text{cm}$ and $\tan(\theta) = 0.23$ a uniform field gradient is achieved only for $R_3 = 5\text{cm}$.

A.2.2 Design of Section 5.3

For a cylinder with $L = 10\text{cm}$, $R_3 = 6\text{cm}$, $R_1 = 1\text{cm}$ the bore offset, O_y was varied and figure A.2.6 shows how the field gradient dB_y/dy varies in the y direction within the bore:




Figure A.2.6 - Halbach design with offset bore; dB_y/dy on the y axis within bore for various bore offsets.

Defining a working area of length 0.6cm from -0.45cm to 0.15cm within the bore (marked above in figure A.2.6 as the dashed vertical lines), we assessed the constancy of the gradient within this working region by taking the standard deviation of the gradient over it:




Figure A.2.7 – Halbach design with offset bore; standard deviation within working region for various bore offsets.

The bore offset which gives the lowest standard deviation is $O_y = 2.5\text{cm}$. For a cylinder with $O_y = 2.5\text{cm}$, $R_3 = 6\text{cm}$, $R_1 = 1\text{cm}$ the length was varied and figure A.2.8 shows how the field gradient dB_y/dy varies in the y direction within the bore:




Figure A.2.8 - Halbach design with offset bore; dB_y/dy on the y axis within bore for various cylinder lengths.

A bore offset of 3cm could also have been chosen, as this gives a higher gradient of around 53 T/m with only slightly larger standard deviation of 0.1 T/m, as can be seen from figure A.2.7.

A.2.3 Design of Section 5.4

We investigated the effect of the inner ellipse on the field, taking a range of 1.3π radians, with $bo = 0.7$, $L = 10\text{cm}$, inner and outer ellipse radii as $ir = 4$, $or = 8$ respectively and varying bi . The field gradient is shown as a function of z in figure A.2.9:

The minimum
effect of bi
is $0.7, L$




Figure A.2.9 - Rolled flux sheet design with varying inner ellipse parameter b_i ; dB_y/dy field gradient on y axis.

Figure A.2.10 shows the standard deviation of the gradient over the working region against the ellipticity parameter b_i :




Figure A.2.10 - Standard deviation over working area for flux sheet design for various values of b_i .

The minimum value of the standard deviation is at $b_i = 0.78$. We investigated the effect of the outer ellipse parameter b_o on the field gradient, with range 1.3π radians, $b_i = 0.7$, $L = 10\text{cm}$, inner and outer ellipse multipliers $i_r = 4$ and $o_r = 8$ respectively:




Figure A.2.11 - Rolled flux sheet design with varying outer ellipse parameter bo ; dBy/dy field gradient on y axis.

Figure A.2.12 shows the standard deviation of the gradient over the working region against the ellipticity parameter bo :




Figure A.2.12 - Standard deviation over working area for flux sheet design for various values of bo .

The minimum value of standard deviation is at $bo = 0.58$. We investigated the effect of changing the angular range on a flux sheet with parameters $bi = 0.78$, $bo = 0.7$, $L = 10\text{cm}$, $ir = 4$ and $or = 8$:




Figure A.2.13 - Rolled flux sheet design with varying range parameter; dB/dy field gradient on y axis.

The range is given in units of π . We then investigated the effect of changing the length using parameters $bi = 0.7$, $bi = 0.78$, $ir = 4$ and $or = 8$, with angular range 1.3π radians:




Figure A.2.14 - Rolled flux sheet design of varying length; dB/dy field gradient on y axis.

It can be seen that short lengths destroy the uniformity of the field gradient.

A. 3 Effect of Varying Parameters For Designs In Section 6

A.3.1 Design of Section 6.2

We chose parameters $L = 6\text{cm}$, $R_1 = 1\text{cm}$, $R_3 = 6\text{cm}$ and allowed $\tan(\theta)$ (defining R_2) to vary. Figure A.3.1 shows B_z for a range of $\tan(\theta)$; the dotted lines at 0cm and 6cm show the ends of the cylinder:




Figure A.3.1- Uniformly magnetized cylinder with linearly varying inner radius; B_z on z axis for a range of $\tan(\theta)$.

The product of the field B_z and field gradient dB_z/dz is shown in figure A.3.2 below:




Figure A.3.2 – Uniformly magnetized cylinder with linearly varying inner radius; $B_z(dB_z/dz)$ on z axis for a range of $\tan(\theta)$.

Taking the standard deviation of the product $B_y(dB_y/dz)$ within the working region from 2.5cm to 4cm on the z-axis, we can assess the constancy of BdB :




Figure A.3.3 - Uniformly magnetized cylinder with linearly varying inner radius; Standard deviation of $B_z(dB_z/dz)$ over working area as a function of $\tan(\theta)$.

The minimum value falls at $\tan(\theta) = 0.3$, with a standard deviation of 0.015.

We took $R_1 = 1\text{cm}$, $R_3 = 6\text{cm}$ and allowed L to vary. As $\tan(\theta)$ is fixed at 0.3, the length defines R_2 :




Figure A.3.4 - Uniformly magnetized cylinder with linearly varying inner radius; $B_z(dB_z/dz)$ on z axis for various lengths.

Taking a cylinder with $L = 6\text{cm}$, $R_1 = 1\text{cm}$, $\tan(\theta) = 0.3$ (defining $R_2 = 2.8\text{cm}$) we allowed R_3 to vary. Figure A.3.5 shows BdB along the z-axis for a range of outer radii:




Figure A.3.5 – Uniformly magnetized cylinder with linearly varying inner radius; $B_z(dB_z/dz)$ on z axis for various outer radii.

Changing the outer radius causes the magnitude of BdB to decrease with increasing z for small outer radii, and increase with increasing z for large outer radii.

A.3.2 Design of Section 6.3

Using parameters $\tan(\theta) = 0.23$, with $R_3 = 5\text{cm}$ and $R_1 = 1\text{cm}$, we investigated the effect of changing the cylinder length, L :




Figure A.3.6 - Halbach cylinder with linearly varying inner radius; $B_y(dB_y/dz)$ on z axis for a range of lengths.

Figure A.3.7 shows the standard deviation within the working region (defined in the previous section as the region of the z -axis from -2cm to 0cm) for a range of cylinder lengths:




Figure A.3.7 - Halbach cylinder with linearly varying inner radius; Stdev over working region for a range of lengths.

For a cylinder with parameters $L = 7\text{cm}$, $R_3 = 5\text{cm}$ and $R_1 = 1\text{cm}$, we allowed $\tan(\theta)$ to vary. Figure A.3.8 shows B_y for a range of $\tan(\theta)$; the black vertical lines at $+3.5\text{cm}$ and -3.5cm show the positions of the ends of the cylinder:




Figure A.3.8 - Halbach cylinder with linearly varying inner radius; B_y on z axis for a range of $\tan(\theta)$.

The product of the field and field gradient, $B_y(dB_y/dz)$ is shown in figure A.3.9:




Figure A.3.9 - Halbach cylinder with linearly varying inner radius; $B_y(dB_y/dz)$ on z axis for a range of $\tan(\theta)$.

It can be seen that BdB decreases along the z -axis if $\tan(\theta)$ is less than 0.23 and increases with z for values of $\tan(\theta)$ greater than 0.23 (figures A.3.8, A.3.9).

For a cylinder with parameters $L = 7\text{cm}$, $R_1 = 1\text{cm}$, $\tan(\theta) = 0.23$ (defining R_2) we allowed the outer radius (R_3) to vary. Figure A.3.10 shows B_y for a range of outer radii:




Figure A.3.10 - Halbach cylinder with linearly varying inner radius; $B_y(dB_y/dz)$ on z axis for a range of outer radii R_3 .

Changing the outer radius, R_3 , destroys the region of constant BdB .

A.3.3 Design of Section 6.4

Using flux sheet parameters $bo = 0.5$, $ir = 4$, $or = 8$, $L = 10\text{cm}$, range of 1.2π radians, we varied bi . As with the flux sheet design of section 5.4, we fixed ai and ao at 1.0 and changed the ellipticity by varying bi and bo only. Figure A.3.11 shows $B_y(dB_y/dy)$ against y for a range of bi :




Figure A.3.11 - Rolled flux sheet design with varying inner ellipse parameter bi ; $B_y(dB_y/dy)$ field on y axis.

Figure A.3.12 shows the standard deviation of $B_y(dB_y/dy)$ within the working area:




Figure A.3.12 - Rolled flux sheet design with varying inner ellipse parameter; standard deviation over working region.

We took a flux sheet with $ir = 4$, $or = 8$, length = 10cm, range = 1.2π radians with $bi = 0.5$ and varied bo . Figure A.3.13 shows $B_y(dB_y/dy)$ against y for a range of outer ellipses parameters bo :




Figure A.3.13 - Rolled flux sheet design with varying outer ellipse parameter; $B_y(dB_y/dy)$ on y axis.

We examined the standard deviation within the working region to assess the uniformity of BdB :




Figure A.3.14 - Rolled flux sheet design with varying outer ellipse parameter; standard deviation over working region.

The value of the standard deviation at this minimum is slightly higher than that found when varying the inner ellipse (figure A.3.12); the minima in standard deviation occur at $0.116 \text{ T}^2/\text{m}$ for the design with $bi = 0.5$, $bo = 0.48$ (varying the inner ellipse) and at $0.070 \text{ T}^2/\text{m}$ for the design with $bi = 0.52$, $bo = 0.5$ (varying the outer ellipse).

For a flux sheet with parameters $bi = 0.52$, $bo = 0.5$, $ir = 4$, $or = 8$, length = 10cm, figure A.3.15 shows $B_y(dB_y/dy)$ as a function of y for varying angular ranges (in π radians):




Figure A.3.15 - Rolled flux sheet design with varying range; $B_y(dB_y/dy)$ on y axis.

We allowed the length to vary in a flux sheet with parameters $b_i = 0.52$, $b_o = 0.5$, $i_r = 4$, $o_r = 8$, and range 1.2π radians.

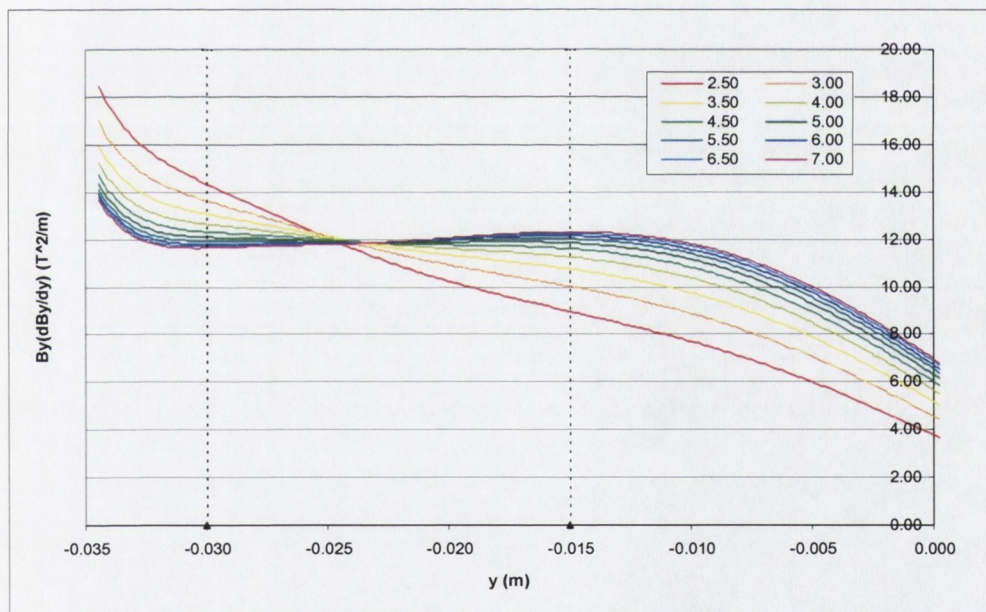


Figure A.3.16 - Rolled flux sheet design with varying length; $B_y(dB_y/dy)$ field on y axis.

Figure A.3.16 shows $B_y(dB_y/dy)$ as a function of y for varying design lengths.

References

- 1 M.V. Berry and A. K. Geim, "Of Flying Frogs and Levitrons", *Eur. J. Phys*, **18**, p307-313 (1997)
- 2 E. P. Furlani, *Permanent Magnet and Electromechanical Devices*, Academic Press (2001)
- 3 J. M. D. Coey and T.R. Ní Mhíocháin, "Permanent Magnets", *High Magnetic Fields: Science and Technology, Volume 1*, ed. F. Herlach and N. Miura, World Scientific Publishing, p25 - 47 (2003)
- 4 M. Mansuripur, "The Faraday Effect", *Optics & Photonics News*, November 1999, <http://www.mmresearch.com/articles/article3/> (1999)
- 5 J. P. Hornak, *The Basics of MRI*, <http://www.cis.rit.edu/htbooks/mri/> (2005)
- 6 A. Yariv, *An Introduction to Theory and Applications of Quantum Mechanics*, Wiley, (1982)
- 7 C. M. Kozierok, *Giant Magnetoresistive (GMR) Heads*, <http://www.pcguide.com/ref/hdd/op/heads/techGMR-c.html> (2001)
- 8 R. Butner, *Computing Unplugged*, http://www.research.ibm.com/thinkresearch/pages/2001/20010202_mram.shtml (2001)
- 9 Q. Peng, S. M. McMurry and J.M.D. Coey, "Cylindrical Permanent-Magnet Structures Using Images in an Iron Shield", *IEEE Transactions on Magnetics*, **39**, p1983 - 1989 (2003)
- 10 K. Halbach, *Nuclear Instruments and Methods*, **169**, 1, (1980)
- 11 E. Potenziani II, H. A. Leupold, and D. J. Basarab , "A Novel Self-Shielding Permanent-Magnet Rotor Assembly", *Journal of Applied Physics*, **64**, 10, p5986-5987 (1988)
- 12 T.R. Ní Mhíocháin, D. Weaire, S. M. McMurry and J. M. D. Coey, "Analysis of Torque in Nested Magnetic Cylinders", *Journal of Applied Physics*, **86**, 11, p6412-6424 (1999)
- 13 O. Cugat and F. Bloch, "4-Tesla Permanent Magnetic Flux Source", *Proc. 15th International Workshop on Rare Earth Magnets and Their Applications*, published by MAT INFO, Dresden 1998, p807 (1998)

- ¹⁴ O. Cugat and F. Bloch, *The 5 Tesla Sphere*, http://mag-net.ee.umist.ac.uk/reports/P14/p14_2.html (1999)
- ¹⁵ M. Kumada, T. Fujisawa, and Y. Hirao, “Development of 4 Tesla Permanent Magnet”, *Proceedings of the 2001 Particle Accelerator Conference, Chicago* (2001)
- ¹⁶ M. Kumada, Y. Iwashita, M. Aoki and E. Sugiyama, “The Strongest Permanent Dipole Magnet”, *Proceedings of the 2003 Particle Accelerator Conference* (2003)
- ¹⁷ J. H. Jensen and M. G. Abele, “Generation of highly uniform magnetic fields with magnetized wedges”, *IEEE Transactions on Magnetics*, **34**, p2316 - 2323 (1998)
- ¹⁸ M. G. Abele, *Structures of Permanent Magnets*, Wiley (1993)
- ¹⁹ O. Cugat, P. Hansson and J.M.D. Coey, “Permanent magnet variable flux sources”, *IEEE Transactions on Magnetics*, **30**, p4602 – 4604 (1994)
- ²⁰ J. C. Mallinson, “One Sided Fluxes – A Magnetic Curiosity?”, *IEEE Transactions on Magnetics*, **9**, p678 – 682 (1973)
- ²¹ E. W. Weisstein. “Hilbert Transform”, MathWorld, <http://mathworld.wolfram.com/HilbertTransform.html> (1999)
- ²² H. A. Shute, J. C. Mallinson, D. T. Wilton and D. J. Mapps, “One – Sided Fluxes in Planar, Cylindrical and Spherical Magnetized Structures”, *IEEE Transactions on Magnetics*, **36**, p440 - 451 (2000)
- ²³ J. A. Clarke, “The Magnetic Design of a High Field Permanent Magnet Multipole Wiggler for the SRS”, *17th IEEE Particle Accelerator Conference (PAC 97)*, DL Preprints, DL-P-97-002 (1997)
- ²⁴ CPSMA, *Free Electron Lasers and Other Advanced Sources of Light: Scientific Research Opportunities*, <http://www.nap.edu/books/NI000099/html/> (1994)
- ²⁵ J. D. Jackson, *Classical Electrodynamics*, Wiley (1962)
- ²⁶ W. Drenckhan, F. Elias, S. Hutzler, D. Weaire, E. Janiaud and J. C. Bacri, “Bubble Size Control and Measurement in the Generation of Ferrofluid Foams”, *Journal of Applied Physics*, **93**, p10078 – 10083 (2003)

- 27 J. C. P. Klaasse, *The Faraday Balance*,
<http://www.science.uva.nl/research/mmm/fdb.html> (1999)
- 28 HFML, *The Frog That Learned to Fly* ,
<http://www.hfml.sci.kun.nl/froglev.html> (2000)
- 29 S. Earnshaw, “On the nature of the molecular forces which regulate the constitution of the luminiferous ether”, *Transactions of the Cambridge Philosophical Society*, **7**, p97 - 112. (1842)
- 30 M. D. Simon, L. O. Heflinger, A. K. Geim, “Diamagnetically Stabilized Magnet Levitation”, *American Journal of Physics*, **69**, p702 - 713 (2001)
- 31 D. J. Griffiths, *Introduction to Electrodynamics*, Prentice Hall (1981)
- 32 E. W. Weisstein. “Rotation Matrix”, *MathWorld*,
<http://mathworld.wolfram.com/RotationMatrix.html> (1999)
- 33 R. K. Livesley, *Finite Elements: An Introduction for Engineers*, Cambridge University Press (1983)
- 34 O. C. Zienkiewicz and K. Morgan, *Finite Elements and Approximation*, Wiley (1983)
- 35 T. Lepikult, *Weighted Residual Methods for Solving ODE and PDE*,
http://www.cs.ut.ee/~toomas_1/linalg/dv/node7.html (1997)
- 36 G. James, *Advanced Modern Engineering Mathematics*, p679 - 686, Addison Wesley (1993)
- 37 R. Wait and A. R. Mitchell, *Finite Element Analysis and Applications*, Wiley (1985)
- 38 Joseph E. Flaherty , *Finite Element Analysis Course Notes, Lecture 4*,
<http://www.cs.rpi.edu/~flaherje/pdf/fea4.pdf> (2005)
- 39 Y. Kanai, T.Abe, M. Sengoku, T. Iijima, M. Iizuka and K. Mukasa, “New Formulation of Finite-Element Method With Gauge Condition for Three Dimensional Magnetic Field Analysis”, *IEEE Transactions on Magnetics*, **24**, p3123 - 3125 (1988)
- 40 Y. Kanai, T.Abe, M. Sengoku, T. Iijima, M. Iizuka and K. Mukasa, “Further Discussion on Magnetic Vector Potential Finite-Element Formulation for Three Dimensional Magnetostatic Field Analysis”, *IEEE Transactions on Magnetics*, **26**, p411 - 414 (1990)

- 41 J. K. Sykulski, *Computational Magnetism*, Chapman & Hall (1995)
- 42 L. Jänicke and A. Kost, “Convergence Properties of the Newton-Raphson
Method for Nonlinear Problems”, *IEEE Transactions on Magnetics*, **34**, p2505
- 2508 (1998)
- 43 J. Dongarra, *Iterative Methods*,
http://netlib2.cs.utk.edu/linalg/html_templates/node9.html (1995)
- 44 J. Gilbert, *Analysis of Convergence*,
<http://www.maths.lancs.ac.uk/~gilbert/m306b/node16.html> (1999)
- 45 J. R. Shewchuk, *An Introduction to the Conjugate Gradient Method Without
the Agonizing Pain*, [http://www.cs.cmu.edu/~quake-papers/painless-
conjugate-gradient.pdf](http://www.cs.cmu.edu/~quake-papers/painless-conjugate-gradient.pdf) (1994)
- 46 W. H. Press, B. P. Flannery, S. A. Teukolsky and W. T. Vetterling, *Numerical
Recipes in C : The Art of Scientific Computing*, Cambridge University Press
(1992)
- 47 R. K. Bock, “Cholesky Decomposition” *The Data Analysis BriefBook*,
<http://rkb.home.cern.ch/rkb/titleA.html> (1999)
- 48 M. Benzi, C. D. Meyer and M. Tüma, “A Sparse Approximate Inverse
Preconditioner for the Conjugate Gradient Method”, *SIAM Journal on
Scientific Computing*, **17**, p1135 – 1149 (1996)
- 49 R. Barrett *et al.*, *Templates for the Solution of Linear Systems: Building Blocks
for Iterative Methods, 2nd Edition*, SIAM,
http://www.netlib.org/linalg/html_templates/Templates.html (1994)
- 50 R. J. Hartman-Baker, *Introduction to Iterative Methods*,
<http://www.cse.uiuc.edu/~rjhartma/cs257/gaussseidel.pdf> (2002)
- 51 M. T. Jones and P. E. Plassmann, “An Improved Incomplete Cholesky
Factorization”, *ACM Transactions on Mathematical Software*, **21**, p5 – 17
(1995)
- 52 J. R. Shewchuk, *Lecture notes on Delaunay Mesh Generation*, [http://www-
2.cs.cmu.edu/~jrs/jrspapers.html](http://www-2.cs.cmu.edu/~jrs/jrspapers.html) (1999)
- 53 P.J. Frey and P. L. George, *Mesh Generation*, Hermes Science Publishing
(2000)

- 54 M. S. Shephard and M. K. Georges, “Automatic Three-Dimensional Mesh Generation by the Finite Octree Technique”, *International Journal for Numerical Methods in Engineering*, **32** p709 - 749 (1991)
- 55 N. Molino, R. Bridson, J. Teran, and R. Fedkiw, “A Crystalline, Red Green Strategy for Meshing Highly Deformable Objects with Tetrahedra”, *12th Int. Meshing Roundtable*, p103 - 114, (2003)
- 56 J. R. Shewchuk, “Tetrahedral Mesh Generation by Delaunay Refinement”, *Proceedings of the Fourteenth Annual Symposium on Computational Geometry*, p86 - 95 (1998)
- 57 A. Bower, “Computing Dirichlet Tessellations”, *The Computer Journal*, **24**, p162-166 (1981)
- 58 D. F. Watson, “Computing the n-Dimensional Delaunay Tessellation with Application to Voronoi Polytopes”, *The Computer Journal*, **24**, p167-172 (1981)
- 59 C. L. Lawson, “Software for C surface Interpolation”, *Mathematical Software III*, p161-194 (1977)
- 60 L. P. Chew, “Guaranteed-Quality Triangular Meshes”, *Technical Report TR-89-983*, Dep. of Computer Science, Cornell University (1989)
- 61 Steven S. Skiena, “Robust Geometric Primitives”, *The Algorithm Design Manual*,
<http://www2.toki.or.id/book/AlgDesignManual/BOOK/BOOK4/NODE184.HTM#geometricprimitives> (1997)
- 62 J. R. Shewchuk, *Adaptive Precision Floating-Point Arithmetic and Fast Robust Predicates for Computational Geometry*, <http://www-2.cs.cmu.edu/~quake/robust.html> (1996)
- 63 K. L. Clarkson, “Safe and effective determinant evaluation”, *Proc. 31st IEEE Symposium on Foundations of Computer Science*, p387 - 395 (1992)
- 64 H. Brönnimann and M. Yvinec , “A complete analysis of Clarkson's algorithm for safe determinant evaluation”, *Research Report 3051 INRIA* (1996)
- 65 E. P. Mücke, I. Saias and B. Zhu, “Fast Randomized Point Location Without Pre-processing in Two and Three Dimensional Delaunay Triangulations”,

Proceedings of the twelfth annual symposium on Computational geometry, p274 – 283 (1996)

- ⁶⁶ J. Ruppert, “A Delaunay Refinement Algorithm for Quality 2-Dimensional Mesh Generation”, *Journal of Algorithms*, **18**, p.548-585 (1995)
- ⁶⁷ L. P. Chew, “Guaranteed-Quality Mesh Generation for Curved Surfaces”, *Proceedings of the Ninth Annual Symposium on Computational Geometry*, p274 - 280 (1993)
- ⁶⁸ P.R. Calvalcanti and U. T. Mello, “Three-dimensional Constrained Delaunay Triangulation: a Minimalist Approach” , *Proceedings of the 8th International Meshing Roundtable, Lake Tahoe, California*, p119 - 129 (1999)
- ⁶⁹ P. Brown, *Characterization and Shimming of Magnetic Field Profiles Produced By Halbach Cylinders*, M.Sc. Thesis, Trinity College Dublin (1996)
- ⁷⁰ M. G. Abele and H. Rusinek, “Optimum Design of Yokeless Permanent Magnets”, *Journal of Applied Physics*, **67**, p4644 - 4646 (1990)