

MRI Magnet Design Using Genetic Algorithms

Nicholas R Shaw
Corpus Christi College
Cambridge
UK

June 2003

A dissertation submitted to the University of Cambridge for
the degree of Doctor of Philosophy

MRI Magnet Design using Genetic Algorithms

Nicholas R Shaw

This thesis describes the design of superconducting magnetic resonance imaging (MRI) magnets using genetic algorithms. Recent advances in the availability of high performance computers including Beowulf clusters and supercomputers have allowed the application of stochastic optimization techniques to highly complex real-world engineering problems. For this project, a real-coded genetic algorithm (GA) has been written to take advantage of such parallel architectures and is applied to the design of axisymmetric MRI magnets.

A substantial part of this thesis is dedicated to a general framework for the design and optimization of axisymmetric MRI magnets. This framework includes a summary of methods for calculating magnetic fields, forces and stresses. Two novel analytical methods are also included in this framework. A rapid method for the evaluation of the fringe field, using spherical harmonic decomposition, is presented. Furthermore, an analytical approach which measures design sensitivity to engineering tolerances is proposed.

Two particular MRI design configurations are investigated. Firstly, an ultra-short whole-body MRI magnet design is considered. Conventional MRI magnets have bore lengths between 1.5 and 2.5 m. This causes problems due to patient claustrophobia and also impairs access to critically ill patients. Using the GA, an ultra-short whole-body MRI design is optimized. The resulting design has a bore length of 1.0 m and inner bore diameter of 0.94 m with a standard deviation of the inhomogeneity over a 40 cm diameter spherical volume (DSV) of 1 ppm and operates at 1.0 T. It is found that short bore designs are limited by hoop stress constraints and that the cost of such a design is substantially higher than for conventional designs.

Secondly, a split-coil MRI magnet suitable for simultaneous positron emission tomography (PET) and MRI is considered. Dual-modality PET and MRI is an exciting combination of complimentary imaging techniques, offering detailed *in vivo* functional and anatomical information. A split-coil design is optimized and compared to a commercially produced design. The final design boasts a 20 cm inter-coil gap, with a 16 cm DSV giving a four-fold increase in the homogeneous imaging volume. This design is currently under construction and will be installed in the Cavendish laboratory during 2003.

Preface

The work presented in this dissertation was performed at the Cavendish Laboratory in the Low Temperature Physics group. Access to facilities at the Wolfson Brain Imaging Centre and the Herchel Smith Laboratory for Medicinal Chemistry is also acknowledged.

This dissertation is my own work and contains nothing which is the outcome of work done in collaboration with others, except as specified in the text and acknowledgements. The number of words in this thesis (including footnotes, tables, bibliography and appendices) is less than 60,000.

Nicholas Shaw
Cambridge

1 June 2003

Contents

Preface	i
Acknowledgements	xv
1 Introduction	1
2 Theory Part I: Magnetic Resonance Imaging	5
2.1 Introduction	5
2.2 The Emergence of MRI	5
2.3 Physics of Nuclear Spin Systems	7
2.4 Acquiring a Signal	11
2.4.1 RF excitation pulse	11
2.4.2 Free induction decay	12
2.4.3 Signal detection	15
2.5 The MRI System Overview	16
2.5.1 Main magnet	16
2.5.2 Gradient coils	17
2.5.3 RF coils	17
2.5.4 Computer system	18
2.6 Signal Localization	18
2.6.1 Slice selection	18
2.6.2 Frequency and phase encoding	19
2.7 Pulse Sequences	19
2.7.1 Spin echo imaging	20
2.7.2 Gradient echo imaging	22
2.7.3 Echo planar imaging	22

2.8	Effects of Field Inhomogeneity	22
3	Theory Part II: Genetic Algorithms	27
3.1	Introduction	27
3.2	The Simple Genetic Algorithm	28
3.2.1	Coding the genetic information	28
3.2.2	Initialization	31
3.2.3	Fitness evaluation	31
3.2.4	Selection	32
3.2.5	Recombination	33
3.2.6	Mutation	35
3.3	Constraint Handling	35
3.4	Advanced GA Methods	39
3.5	The Schema Theorem	40
3.6	Parallel Implementation	42
4	Design Framework for Axisymmetric Magnet Optimization	45
4.1	Introduction	45
4.2	Magnetic Field Evaluation and Central Field Homogeneity	46
4.2.1	Direct Biot-Savart integration	46
4.2.2	Hybrid integral methods	48
4.2.3	Harmonic decomposition methods	49
4.2.4	Other methods	51
4.2.5	Homogeneity - a single figure of merit	51
4.3	Novel Fringe Field Evaluation Method	52
4.3.1	Fringe field expansion of loops, thin solenoids and thick solenoids in spherical harmonics	52
4.3.2	Radius of convergence	61
4.3.3	Using the weighting constants	62
4.4	Design Sensitivity	63
4.4.1	Ensemble averaging	63
4.4.2	Novel analytical approach to design sensitivity	64
4.5	Forces, Stresses and Strains	71

4.5.1	Forces and stresses in solenoids	72
4.5.2	Material properties	74
4.5.3	Finite element analysis with ABAQUS	75
4.6	Quenching, Stored Energy and Inductance	76
4.6.1	Introduction to quenching	76
4.6.2	Quench modelling	79
4.6.3	Quench prevention and protection	80
4.6.4	Training	82
4.6.5	Stored energy and inductance	82
4.7	Physical Limitations	83
5	Genetic Algorithms for Optimizing Short Whole-Body MRI Magnets	85
5.1	Introduction	85
5.2	Specifications and Requirements	86
5.3	Methods I: Implementation of Genetic Algorithm	88
5.3.1	Encoding the parameters	88
5.3.2	Details of the GA	88
5.3.3	Parallel computation and MPI	91
5.3.4	Parallel GA	91
5.4	Testing the GA	92
5.4.1	Helmholtz pair	92
5.4.2	Coil configurations	93
5.5	Methods II: Applying the GA to the Short Bore Magnet	95
5.5.1	Main coil optimization	95
5.5.2	Superconducting shim coil optimization	99
5.5.3	Quantization of parameters	101
5.6	Results and Analysis	101
5.6.1	Overall configuration	101
5.6.2	Homogeneity	102
5.6.3	Fringe fields	105
5.6.4	Peak fields	105
5.6.5	Stress and Forces	106

5.6.6	Sensitivity	107
5.7	Discussion	109
5.8	Conclusions	112
6	Split-coil MRI Magnet for Combined PET/MRI	115
6.1	Introduction	115
6.2	Combined PET and MRI	116
6.2.1	What is PET?	116
6.2.2	Challenges of combining PET and MRI	117
6.2.3	Previous attempts	118
6.2.4	A novel approach	119
6.3	Specifications and Requirements	120
6.4	Methods I: Implementation of Genetic Algorithm	121
6.4.1	Specifying the homogeneity	122
6.4.2	Fringe field evaluation	123
6.4.3	Constraint handling	124
6.4.4	Summary of fitness function	125
6.5	Methods II: Applying the GA to the Split-Coil Magnet	126
6.5.1	Initial investigations of the search space	126
6.5.2	Preliminary results	128
6.5.3	Comparison with manufacturer’s design	128
6.5.4	Including the sensitivity constraint	130
6.6	Results and Analysis	132
6.6.1	Overall configuration	133
6.6.2	Homogeneity	135
6.6.3	Fringe fields	136
6.6.4	Peak fields	136
6.6.5	Stress and Forces	136
6.6.6	Sensitivity	137
6.6.7	Comparison with manufacturer’s design	137
6.7	Discussion	139
6.8	Conclusions	144

7	Conclusions	145
A	Magnetic field evaluation in the near-field using spherical harmonics	151
B	Partial differentials of near-field weighting constants	155
C	Post-processing software	159
	Bibliography	163

List of Figures

2.1	Nuclear magnetic moments precessing around B_0	9
2.2	Motion of the bulk magnetization vector in the presence of a radio-frequency excitation pulse.	13
2.3	Schematic of a z gradient coil and the effect on the static main field strength.	17
2.4	Spin echo pulse sequence showing the timing of the RF pulses, the linear gradients and the signal acquisition.	21
2.5	Example of (a) the acquired \mathbf{k} -space data (magnitude) and (b) the corresponding real-space image.	21
2.6	Example of the frequency spectrum for a sample in an inhomogeneous magnetic field.	24
3.1	A flowchart showing the different stages of a simple genetic algorithm.	29
3.2	Illustration of the formation of the chromosome structure in (a) binary encoded GAs and (b) real coded GAs.	30
3.3	Schematic representation of a single point crossover in a binary encoded GA.	34
3.4	A real coded crossover operator, BLX- α	34
3.5	Illustration of mutations using different encodings.	36
3.6	Heuristic constraint handling method.	38
3.7	Parallel genetic algorithms.	44
4.1	Comparison of Cartesian coordinates with spherical coordinates.	47
4.2	Axisymmetric loop defined by (r', θ') , carrying current I and observation point \mathcal{P}	54

List of Figures

4.3	A thick solenoid with current density J''	58
4.4	Correctly combining the weighting constants, $h_n(s)$	60
4.5	Pattern of forces within a solitary axisymmetric coil.	73
4.6	Finite element analysis of ultra-short whole-body MRI magnet design.	77
4.7	Von Mises stress for a laminated coil.	77
4.8	Critical current density vs B characteristics of NbTi and Nb ₃ Sn multifilamentary superconductors.	81
5.1	Illustration of (a) the definitions of the parameters used in axisymmetric coil optimization and (b) the encoding of the parameters to form a chromosome.	89
5.2	Convergence of thick Helmholtz coil test problem.	94
5.3	Comparison of different combinations of positive and negative coils.	94
5.4	Sketch of the points at which the magnetic field was evaluated within the ROI.	96
5.5	Contour plot of inhomogeneities for GA optimized short bore magnet.	98
5.6	Cross-section of the final optimized short, whole-body MRI magnet.	103
5.7	Close up of one quadrant of short, whole-body MRI magnet.	103
5.8	Homogeneity of short, whole-body MRI magnet over the ROI.	104
5.9	Fringe fields of short, whole-body MRI magnet.	105
5.10	Cutaway view of the short, whole-body MRI magnet showing hoop stress, σ_{33} (MPa)	107
5.11	Simulation of the homogeneity degradation due to 0.5 mm coil positioning errors.	108
6.1	Two geometries for simultaneous PET and MRI.	120
6.2	The optimized diameter of the ROI, given as a function of volume of superconducting wire (cost) and inter-coil separation.	127
6.3	Comparison between (a) the preliminary design from the GA and (b) the commissioned design from Magnex Scientific Ltd.	130

6.4	Contour plot of the homogeneity of (a) the preliminary design from the GA and (b) the commissioned design.	131
6.5	2D poloidal cross-section of final split-coil magnet design.	133
6.6	Homogeneity of the final split-coil design.	134
6.7	Magnetic far field region of the split-coil MRI magnet.	135
6.8	Hoop stress analysis of the split-coil magnet using Abaqus.	138
6.9	Shear stress analysis of the split-coil magnet using Abaqus.	138
6.10	Sensitivity of split-coil MRI magnet design to positioning errors.	139
6.11	Cut-away view of the split-coil MRI magnet with the PET system installed.	142
6.12	Figure showing simultaneously acquired PET and MR images.	143
7.1	Example showing the interaction of constraints on the number and classification of minima in a simplified optimization problem.	148
C.1	Summary sheet from post-processing software	162

List of Tables

5.1	Positions and current density of four main superconducting coils optimized using the genetic algorithm, in the short whole-body MRI magnet design.	98
5.2	Resulting short bore MRI magnet design with the addition of five pairs of superconducting shim coils.	100
5.3	Positions and current density of the coils in final short, whole-body MRI magnet design.	102
5.4	Peak field, hoop stress and axial force for final short, whole-body MRI magnet design.	106
6.1	Comparison of the preliminary design with the manufacturer's proposed design.	129
6.2	Positions and current density of the coils in the final split-coil MRI magnet design.	132
6.3	Peak fields, axial forces and approximate hoop and shear stresses for the three pairs of coils in the split-coil MRI magnet design.	136
6.4	Comparison of split-coil magnet designs.	140
C.1	Key-phrases, descriptions and default values used in the MRI magnet design evaluation software.	160

Acknowledgements

A dissertation such as this cannot and should not be written without the involvement, help and support of many people. In my case, I feel especially fortunate to have been surrounded by so many colleagues, friends and family who have supported me throughout the past three and a half years.

Much credit must go to my supervisor, Dr. Richard Ansorge, who was the first to get me interested in this project and has always been available for advice, guidance and support. As part of a small academic group, his presence has been absolutely vital and I am grateful for his willingness to get involved in the details from day one.

During my first year, I was fortunate enough to work closely with Dr. Moray Grieve. His methodical approach to every aspect of work was inspiring and complimented my desire to have everything completed yesterday. Having moved on to fresh pastures, I am particularly grateful for the time he spent proof-reading the first draft of this thesis. Many of his suggestions have found their way into the final document.

I have also been helped by Dr. Peter Hanley who introduced me to a lot of the black-art surrounding superconducting magnet design and Dr. Michael Sutcliffe who helped arrange access to the Engineering department for the stress analysis of the final split-coil magnet design.

Continuing on an academic theme, I would also like to thank Dr. Adrian Carpenter, Prof. Laurie Hall, Dr. Guy Williams and Dr. Da Xing for many useful conversations both concerning this project and MRI in general. Also, Prof. Mark Warner and the Rev. Mark Pryce must be thanked for their help during my years as an undergraduate.

Thanks are also directed at the people with whom I have shared an office

Acknowledgements

for the past three and a half years: Dr. Dominic Forsythe, Nicolas Doiron-Leyraud, Alix McCollam, Peter Mann and especially Patricia Alireza. Never a day has passed without some light-hearted joke or kind word spoken. Such an atmosphere engenders more work rather than less. I hope that they all know how much I have appreciated this.

The Low Temperature Physics group has been a vibrant unit led by Dr. Stephen Julian. I have enjoyed my own interactions with strongly correlated electron systems through weekly seminars and tea-time conversations, especially with Dr. Prabhat Agarwal, Dr. Karl Sandeman, Dr. May Chiao, Dr. Christoph Bergemann, Dr. Fabien Beckers, Dr. Sibel Özcan, Philipp Niklowitz, Ben Morgan, Samira Barakat, Pon Satangput, Ramzy Daou, Tim Hosey and Tim Moran.

The encouragement and support from Michèle Lévêque throughout these long years of research has been phenomenal, especially in the final months before completing this thesis. She has remained interested and inquisitive in my work despite the fact that progress was sometimes slow. My parents must also take a lot of credit. The love that they have shown can only be matched by my admiration for them.

I acknowledge the financial support of the Engineering and Physical Sciences Research Council, Oxford Instruments Plc and Corpus Christi College. And finally, I can only hope that some of the interest and excitement which I have felt for my research will be transmitted through these pages.

Chapter 1

Introduction

Since its invention in the late 1970s, magnetic resonance imaging (MRI) has developed rapidly to become one of the most important imaging modalities in modern medicine. The ability of magnetic resonance to image arbitrarily orientated cross-sections of the human body at high resolution has led to its wide acceptance by radiologists and clinicians. Offering unprecedented soft-tissue contrast, MRI is currently the method of choice for disease diagnosis in many parts of the body, especially the brain, spine and joints. The lack of harmful ionizing radiation or radioactive sources also makes MRI suitable for repetitive studies. Ultimately, though, the overriding advantage of MRI is its flexibility. By altering the excitation pulse sequence, a wide range of images exhibiting different contrasts may be acquired. Examples include proton density, T_1 or T_2 weighted, volumetric imaging, flow imaging, diffusion, perfusion and functional imaging.

During the past twenty years, major improvements have been made to all parts of the MRI system in order to keep up with the requirements of the new imaging sequences. The main-field magnet, which is by far the largest component of the hardware necessary for MRI, is no exception. Advances have included higher field strengths, shorter bore lengths and the introduction of active shielding to reduce the fringe fields. However, the considerable physical size of these magnets still causes problems. MRI magnets are typically 1.5-2.0 m in length and affect patient compliance due to claustrophobia in

a significant proportion of studies. Furthermore, access to patients in need of close supervision, due to acute trauma or during intervention, remains marginal. Various geometries have been investigated in the past to increase access to the patient during imaging, including ‘C’-shaped magnets. These magnets often operate at lower field strengths than conventional solenoidal magnets and suffer from a lower signal to noise ratio, requiring longer scan times. Therefore, much scope for the improvement of the design of MRI main-field magnets remains.

This thesis describes my own investigation into the optimization of MRI main-field magnets using genetic algorithms as a stochastic optimization method. Traditionally, magnet design has been performed analytically and much of the expertise has remained hidden from the general scientific community due to commercial secrecy. However, a significant increase recently in the availability of computing resources, specifically the advent of high-performance parallel Beowulf clusters, has enabled the consideration of much more complex real-world optimization problems. For this project, a genetic algorithm (GA), capable of exploiting these parallel computer architectures, has been written¹ and applied to the design of axisymmetric MRI magnets. In particular, two types of novel magnet design have been considered: an ultra short bore whole-body MRI magnet and an animal scale, split-coil magnet suitable for combined MRI and positron emission tomography (PET).

The immense advantages offered by a simultaneous MRI and PET imaging system are clear. Positron emission tomography is in many ways a complementary imaging modality to MRI. Whereas MRI offers detailed information on anatomical structure, PET is fundamentally used for functional imaging. By labelling various chemicals with positron emitters, PET can assess chemical and physiological changes related to the metabolism. Reconstruction of simultaneously acquired images would be much simplified and motion blurring artifacts on the PET images could theoretically be deconvolved with rapid acquisition MRI images.

As well as computational optimization of novel magnet designs, this dissertation contains two pieces of original theoretical work. One deals with a

¹The GA was originally written in collaboration with Dr. Moray Grieve.

method of calculating the magnetic field outside the magnet using a spherical harmonic expansion technique and the other deals with an analytical approximation of the buildability (or sensitivity to build-errors). These methods are presented within a framework which also includes a review of the standard methods of calculating the magnetic field from thick solenoidal coils and a discussion of the effects of forces and stresses.

Outline of thesis

Chapter 2 contains a basic description of the theory of MRI, starting with the physical basis for nuclear magnetic resonance (NMR). This is followed by an explanation of the need for magnetic gradient fields for spatially localizing the NMR signal and an investigation into the effects of an inhomogeneous static magnetic field over the region of interest (ROI).

Chapter 3 introduces the concept of optimization using a genetic algorithm (GA). The basic operations involved in a simple GA are presented, followed by some of the more advanced GA techniques. A taste of GA theory is given by reference to the Schema theorem and finally, the parallel implementation of GAs is discussed.

Chapter 4 presents a general framework for the design of axisymmetric MRI magnets. Most aspects of MRI magnet design are covered, including homogeneity requirements, fringe field evaluation, measurements of sensitivity to manufacturing errors and evaluation of stresses. Details of two novel theoretical approaches are considered and relevant proofs are included.

Chapter 5, the first of the two main results chapters, looks at the design of ultra short bore MRI magnets to reduce the effects of claustrophobia on patients and improve access for surgeons. Issues about the performance of the GA are investigated and methods of magnet shimming are examined.

Chapter 6 deals with the application of the GA to animal scale, split-coil MRI magnets suitable for combined PET and MRI. Improvements to the original GA are presented and comparisons are made with results from traditional design techniques.

In Chapter 7, conclusions are drawn and opportunities for future work

are suggested.

Appendix [A](#) revisits important work on the evaluation of magnetic fields from axially symmetric coils. Formulae are remodelled to take account of SI units and the weighting constants are expressed in Tesla.

Appendix [B](#) contains some of the proofs required for the derivation of the analytical sensitivity formulae presented in [Chapter 4](#).

Finally, Appendix [C](#) summarizes a few of the programs that were written as part of this project to post-process axisymmetric magnet designs.

Chapter 2

Theory Part I: Magnetic Resonance Imaging

2.1 Introduction

Details of the basic theory of magnetic resonance imaging (MRI) are presented in this chapter. After a brief history of MRI, the behaviour of nuclear spins in a magnetic field is described. Next, the application of a radio-frequency (RF) excitation pulse to the nuclear spin system is considered, as well as free induction decay (FID) and signal detection. The main magnet, gradient coils, RF coils and computer system - all essential components of an MRI scanner - are then introduced, followed by a description of the magnetic gradients and their role in spatial localization. The application of magnetic gradients leads to a discussion of three common imaging sequences: spin echo, gradient echo and echo planar imaging. Finally, the effects of magnetic field inhomogeneity on the acquired signal and the resulting image are examined.

2.2 The Emergence of MRI

Magnetic resonance imaging has its roots in a series of experiments by Stern and Gerlach between 1921 and 1933. The most famous of these, in 1922,

showed that collimated silver atoms passing through an inhomogeneous magnetic field separated into two beams of atoms [1] (later explained as spin-up and spin-down atoms by Pauli in 1924). This experiment indicated the existence of electronic intrinsic angular momentum and in 1933, an improved experiment on hydrogen atoms demonstrated a similar effect due to an intrinsic nuclear magnetic moment. Twelve years later in 1945 two groups, one led by Purcell at MIT¹ and the other led by Bloch at Stanford University, almost simultaneously measured nuclear resonance effects [2, 3]. Nuclear magnetic resonance (NMR) quickly became a key tool of fundamental physics research, allowing the properties of atoms and their nuclei to be measured accurately. As the systems under study became more complicated, atomic shielding was encountered along with the resultant chemical shift artifact, creating a basis for the enormously powerful NMR spectroscopy techniques currently used in analytical chemistry.

The major step towards *spatially* resolved nuclear magnetic resonance came in the early 1970's when Lauterbur proposed that magnetic field inhomogeneities could be used to localize the signals from NMR spectra [4]. Originally termed 'Zeugmatography', Lauterbur described the implementation of a back-projection reconstruction technique. Very soon Kumar, Welti and Ernst [5] realized that switched magnetic field gradients in the time domain could also be used to create an image and they published their method under the title 'NMR Fourier Zeugmatography' which forms the basis for all modern magnetic resonance imaging. During the late 1970's, faster scanning techniques were proposed including echo planar imaging by Sir Peter Mansfield [6] in 1977. The first human scale NMR images were also taken in 1977 using a machine called Indomitable and in 1980, the first clinically viable MRI scanners became available.

Although intrinsically insensitive, nuclear magnetic resonance imaging, or MRI, has grown to become a major rival of the older imaging modalities such as computed tomography (CT) or positron emission tomography (PET), whilst avoiding the use of potentially harmful high-frequency ionizing radiation. Indeed MRI is in some senses a form of *emission* tomography,

¹Massachusetts Institute of Technology

picking up signals created within the imaged object. Other types include *transmission* tomography such as X-ray or CT, and *diffraction* tomography such as ultrasound. Another incredible aspect of MR imaging, especially to the physicist, is its ability to image objects well below the characteristic length-scale of the working frequency. For example, routine scans at 1.5 Tesla (T) may provide sub-millimetre voxel resolution, using a radio-frequency of typically 65MHz with 2.5 m wavelength.

Since 1980, advances in magnetic field homogeneity, gradient set performance and computational power have all led to a remarkable increase in the numbers of routine scans. MRI offers superb soft tissue contrast for almost all parts of the body, all of which may be imaged using a wide range of pulse sequences. Only in the bones, lungs and possibly the gastrointestinal tract, does CT outperform MRI, due to lack of signal and the high contrast in CT of calcification. This chapter reviews the basic physics principles behind MR imaging. The reader can find further details in any one of several good references on the subject including Liang and Lauterbur [7] and Suetens [8].

2.3 Physics of Nuclear Spin Systems

A fundamental result of quantum mechanics is that nuclei with odd numbers of protons or neutrons possess an intrinsic angular momentum \mathbf{J} , often called *spin*. Although truly characterized by the equations of quantum mechanics, this property can in many ways be classically modelled by a spinning top. Associated with this spinning is a microscopic nuclear magnetic dipole moment, $\vec{\mu}$, which is related to the angular momentum by the equation

$$\vec{\mu} = \gamma \mathbf{J} \quad (2.1)$$

where γ is known as the *gyromagnetic ratio* and is dependent on the specific nucleus. For example $\bar{\gamma} = \frac{\gamma}{2\pi} = 42.6$ MHz/T for hydrogen (^1H) nuclei. The magnitude of $\vec{\mu}$ is then given, based on the theories of quantum mechanics, by

$$|\vec{\mu}| = \gamma \hbar \sqrt{I(I+1)} \quad (2.2)$$

where h ($= 2\pi\hbar$) is Planck's constant and I is the nuclear spin quantum number which takes the allowable values $I = \{0, \frac{1}{2}, 1, \frac{3}{2}, 2, \dots\}$. Nuclei with $I = 0$ therefore have no magnetic moment and cannot be probed using NMR.

In the absence of any external magnetic fields, the magnitude of the nuclear magnetic dipole is fixed, but its direction is completely random. Therefore at thermal equilibrium no net nuclear magnetic field exists at the macroscopic level. When an external magnetic field (B_0), acting along the z axis, is applied to the nuclear spin system², the z -component of $\vec{\mu}$ is constrained to take on quantized values given by

$$\mu_z = \gamma m_I \hbar \quad (2.3)$$

where m_I is called the quantum number and for any nucleus with $I \neq 0$, m_I takes the following set of $2I + 1$ values

$$m_I = \{-I, -I + 1, -I + 2, \dots, I\} \quad (2.4)$$

Whilst the longitudinal component of the nuclear magnetic dipole is quantized, the transverse component in the x - y plane is randomly orientated and has a magnitude given by $|\vec{\mu}_{xy}| = \gamma\hbar\sqrt{I(I+1) - m_I^2}$. In proton MRI, the nuclear spin quantum number $I = \frac{1}{2}$, so the components of $\vec{\mu}$ can be written as

$$|\vec{\mu}_z| = \frac{\gamma\hbar}{2} \quad (2.5)$$

$$|\vec{\mu}_{xy}| = \frac{\gamma\hbar}{\sqrt{2}} \quad (2.6)$$

with each microscopic magnetic vector occupying one of two possible orientations: parallel (spin-up) or anti-parallel (spin-down). As well as causing a quantization of μ_z , the magnetic field also causes the nuclear magnetic dipole to precess. Figure 2.1 shows a classical representation of nuclear magnetic dipoles, some spin-up and some spin-down, precessing around the z axis. Ig-

²Conventionally, the direction of the external static magnetic field is taken to be parallel to the z axis.

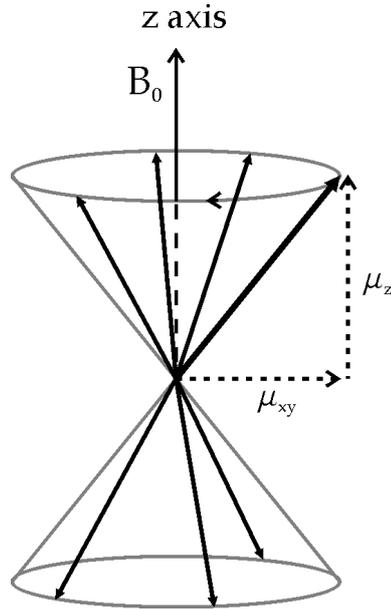


Figure 2.1: Schematic representation of spin-up and spin-down nuclear magnetic moments precessing around the direction of the applied external magnetic field (B_0).

noring interactions between dipoles, and according to classical mechanics, the nuclear magnetic dipole experiences a torque and so moves in the following manner

$$\frac{d\vec{\mu}}{dt} = \gamma\vec{\mu} \times B_0\mathbf{k} \quad (2.7)$$

From (2.7) we can identify γB_0 as the frequency of precession and the direction of precession as clockwise when viewed in the opposite direction to the magnetic field vector. Therefore, the angular velocity vector, $\vec{\omega}_0$ is defined as

$$\vec{\omega}_0 = -\gamma B_0\mathbf{k} \quad (2.8)$$

where $\omega_0 = |\vec{\omega}_0|$ is called the Larmor frequency. The Larmor frequency is an intrinsic property of every nucleus and arises from the gyromagnetic ratio γ . A list of magnetically active nuclei and their properties may be found in several references (e.g. [7, p.59] and [9, p.7]). For example, the most common nucleus used for NMR is the hydrogen nucleus (^1H), which has a Larmor frequency of 42.6 MHz in a 1 T magnetic field. As well as precession,

quantum theory also predicts that each state will have an associated energy of

$$E = -\gamma\hbar m_I B_0 \quad (2.9)$$

For spin systems with $I = \frac{1}{2}$, this leads to an energy splitting, commonly termed the Zeeman splitting

$$\Delta E = \gamma\hbar B_0 \quad (2.10)$$

Returning to the classical description of the nuclear magnetic dipole moments, which is particularly appropriate on the macroscopic level, we can define the bulk magnetization vector \mathbf{M} of a sample by

$$\mathbf{M} = \sum_{n=1}^{N_s} \vec{\mu}_n \quad (2.11)$$

where N_s is the total number of microscopic spins in the system. In thermal equilibrium, there is a slight excess of spins in the lower energy state and Boltzmann statistics show that

$$N_{\uparrow} - N_{\downarrow} \approx N_s \frac{\gamma\hbar B_0}{2k_B T_s} \quad (2.12)$$

where N_{\uparrow} and N_{\downarrow} are the number of spins in the spin-up and spin-down states, k_B is Boltzmann's constant and T_s is the sample absolute temperature. Although the difference is small, (2.12) leads to a net bulk magnetization vector \mathbf{M} which points along the z axis, given by³

$$M_z^0 = |\mathbf{M}| = \frac{\gamma^2 \hbar^2 B_0 N_s}{4k_B T_s} \quad (2.13)$$

For example, a 1 cm³ sample of H₂O in a 1 T field at room temperature gives a net magnetization of 3.2×10^{-9} Am². This is comparable to the magnetization from a 1 cm diameter conductive loop carrying a current of 0.1 mA.

³For general I , $M_z^0 = |\mathbf{M}| = \frac{\gamma^2 \hbar^2 B_0 N_s I(I+1)}{3k_B T_s}$

2.4 Acquiring a Signal

Until now, only two components of an NMR experiment have been discussed: the nuclear magnetic moment and the externally applied magnetic field. We now consider the excitation of the bulk magnetization vector from its equilibrium position. This is followed by an introduction to the relevant relaxation processes and finally signal detection is discussed.

2.4.1 RF excitation pulse

In equilibrium, the net bulk magnetization vector, \mathbf{M} , points directly along the direction of the magnetic field B_0 due to the random phases of the microscopic magnetization vectors. When a radio-frequency electromagnetic wave is applied to a nuclear spin system in equilibrium, the randomized transverse components of the microscopic magnetization vector, $\vec{\mu}_{xy}$, resonate and establish phase coherence. Planck's law states that the energy of electromagnetic radiation is given by

$$E_{\text{rf}} = \hbar\omega_{\text{rf}} \quad (2.14)$$

For coherent transition of spins, the RF photon energy must equal the Zeeman energy gap (2.10). Thus for resonance, $E_{\text{rf}} = \Delta E = \gamma\hbar B_0$, and

$$\omega_{\text{rf}} = \omega_0 \quad (2.15)$$

The RF excitation may be thought of as an oscillating magnetic field denoted by $\mathbf{B}_1(t)$. In a frame rotating at the Larmor frequency, as long as (2.15) holds, $\mathbf{B}_1(t)$ will be stationary and the bulk magnetization vector \mathbf{M} will precess around this new effective magnetic field. For example, an RF excitation pulse, polarized along the x axis could be expressed as

$$\mathbf{B}_1(t) = 2B_1^e(t) \cos(\omega_{\text{rf}}t + \Psi)\mathbf{i} \quad (2.16)$$

where $B_1^e(t)$ is the envelope function of $\mathbf{B}_1(t)$ and the magnetization vector in the rotating frame would move as follows

$$M_{x'}(t) = 0 \quad (2.17)$$

$$M_{y'}(t) = M_z^0 \sin\left[\int_0^t \gamma B_1^e(\hat{t}) d\hat{t}\right] \quad \text{for } 0 \leq t \leq \tau_P \quad (2.18)$$

$$M_{z'}(t) = M_z^0 \cos\left[\int_0^t \gamma B_1^e(\hat{t}) d\hat{t}\right] \quad (2.19)$$

where τ_P is the length of the RF excitation pulse and $M_{x'}$, $M_{y'}$, $M_{z'}$ are the projected components of \mathbf{M} in the rotating frame. Usually the envelope function of the RF pulse is either rectangular or a sinc function. In the case of a rectangular pulse of length τ_P and strength B_1 , the angle by which the bulk magnetization vector rotates, also known as the flip angle (α), is given by

$$\alpha = \gamma B_1 \tau_P \quad (2.20)$$

Figure 2.2 shows the motion of the bulk magnetization vector in both the rotating frame and the stationary frame. There are complications to this simple formula for multiple RF excitations, however the basic description remains the same. See Stark [10] for further details.

2.4.2 Free induction decay

At the end of the RF pulse, the spins in the system enter the epoch of relaxation and free induction decay. The bulk magnetization vector continues to precess around the static magnetic field and relaxes to its equilibrium position along the z axis. The exact mechanisms by which this relaxation occurs are very diverse and complex, but can be ascribed to the interaction of \mathbf{M} with the time-dependent microscopic fields, due to random thermal motion. This effect is usually separated into two processes, spin-spin relaxation and spin-lattice relaxation, which are described below.

- **Spin-spin relaxation:** The transverse component of the magnetization relaxes due to the time-dependent microscopic fields from neigh-

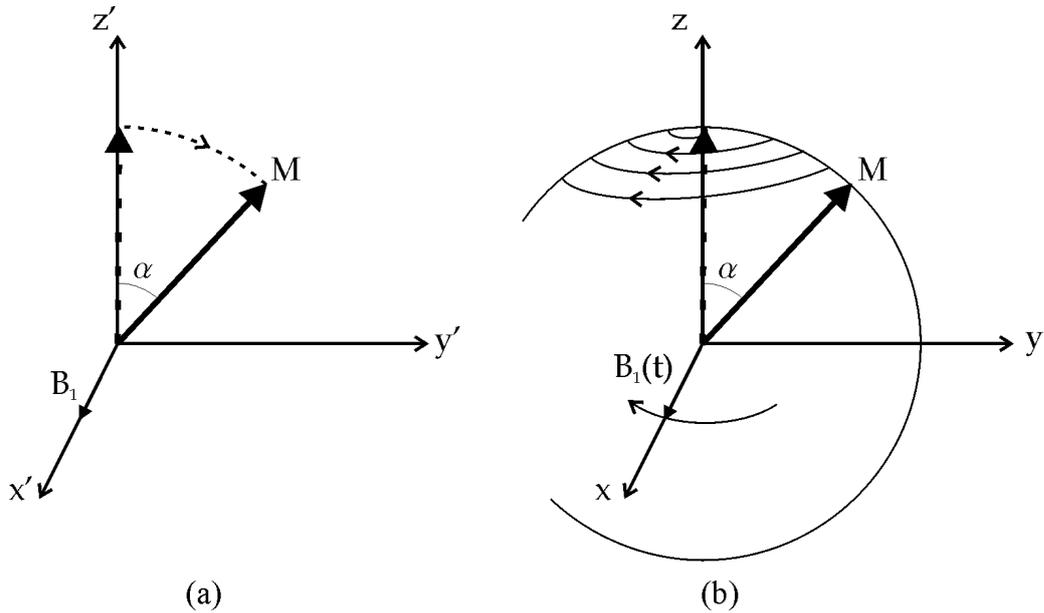


Figure 2.2: Motion of the bulk magnetization vector in the presence of a radio-frequency excitation pulse as observed in (a) the rotating frame, and (b) the laboratory frame [7, p.79].

bouring atoms generated by thermal motion. Effectively, spins experience slightly different magnetic fields and dephase, leading to net cancellation of the transverse component of magnetization. Spin-spin interactions can be considered as an entropy driven process, so the disorder of the system increases without any change in the occupation of the energy levels. The relaxation of the transverse component may be modelled by a single exponential decay with a half-life given by the time constant T_2 .

- **Spin-lattice relaxation:** Concurrently with the relaxation of the transverse component, the longitudinal component of magnetization regains its original value due to spin-phonon interactions. This process involves energy transfer and consequently occurs over a longer timescale. Energy transferred to the lattice is absorbed primarily by vibrations which dissipate as heat⁴. The spins thus return to their

⁴The heat produced in this process is inconsequential compared to the heating from RF absorption.

preferred lower energy state. Again, the relaxation of the longitudinal component of the bulk magnetization vector may be modelled using a single exponential. In this case, the half-life is given by the time constant T_1 .

Phenomenologically, the spin-spin and spin-lattice processes may be described, together with precession by the Bloch equation [7, p.77]

$$\frac{d\mathbf{M}}{dt} = \gamma\mathbf{M} \times \mathbf{B} - \frac{M_x\mathbf{i} + M_y\mathbf{j}}{T_2} - \frac{(M_z - M_z^0)\mathbf{k}}{T_1} \quad (2.21)$$

or in the rotating frame by the first-order differential equations

$$\frac{dM_{z'}}{dt} = -\frac{M_{z'} - M_z^0}{T_1} \quad (2.22)$$

$$\frac{dM_{x'y'}}{dt} = -\frac{M_{x'y'}}{T_2} \quad (2.23)$$

Equations (2.22) and (2.23) may be solved to give the exponential decay of the transverse component and the return to equilibrium of the longitudinal magnetization vector. Note that this single-exponential model, especially in the transverse direction, applies only to weak spin-spin interacting spin systems such as those found in liquids. Spin-spin relaxation in solids is much faster and the mechanisms are much more complicated. Finally, if $T_2 \ll T_1$, then the magnitude of the magnetization vector disappears due to dephasing effects well before the longitudinal magnetization is regained and hence the magnitude of \mathbf{M} is clearly a function of time.

In practice, the spin-spin interactions will be stronger than expected due, for example, to slight inhomogeneities in the static magnetic field. When field inhomogeneity effects are incorporated, the decay of the transverse component of the bulk magnetization vector is characterized by a new time constant known as T_2^* . In general T_2^* is much shorter than T_2 and is given by [7, p.111]

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \gamma\Delta B_0 \quad (2.24)$$

Typical values for T_1 and T_2 vary depending on tissue composition and on the field strength. In a 1.0 T field, water, CSF, grey matter, white matter and fat have T_1 values of 3000, 2000, 800, 650 and 250 ms respectively. The same components have T_2 values of 3000, 150, 100, 90 and 80 ms respectively [11, p.222].

2.4.3 Signal detection

Faraday's law of induction states that a time-varying magnetic flux may be detected as a voltage in a suitably orientated conductive loop. Specifically

$$\text{e.m.f} = -\frac{\partial\phi_B}{\partial t} = -\frac{\partial}{\partial t} \int_S \mathbf{B} \cdot d\mathbf{s} \quad (2.25)$$

where ϕ_B is the magnetic flux and the integral is performed over a surface bounded by the conductive loop. Any conductive loop which resonates at the Larmor frequency may be used, usually consisting of a simple LC circuit with a variable capacitor for tuning. The detection sensitivity of the receiver coil is determined by the principle of reciprocity, so we write the basic formula of all NMR signal detection

$$V(t) = -\frac{\partial}{\partial t} \int_{\text{sample}} \mathbf{B}(\mathbf{r}) \cdot \mathbf{M}(\mathbf{r}, t) d\mathbf{r} \quad (2.26)$$

where $V(t)$ is the measured voltage and \mathbf{r} is a general point within the sample. In general $V(t)$ is a fairly high frequency signal so phase sensitive detection (PSD) is used to move the signal to a lower frequency range. Unfortunately, PSD removes information about the direction of rotation so a second PSD system is usually included with a $\frac{\pi}{2}$ phase shift. Signal detection using two phase sensitive detectors is known as quadrature detection and is commonly used in most modern MRI scanners. The two signals are identified as the real and imaginary parts of the complex signal, $S(t)$, and assuming a

homogeneous coverage of the sample by the reception coil⁵ we have

$$S(t) \propto \int_{\text{object}} M_{xy}(\mathbf{r}, 0) e^{-i\Delta\omega(\mathbf{r})t} d\mathbf{r} \quad (2.27)$$

See Liang [7, p.99] for more details.

2.5 The MRI System Overview

So far, we have dealt with a basic NMR signal acquired from the entire sample. However, it is possible, by altering the magnetic field strength across the sample (i.e. creating a magnetic field gradient), to spatially localize the acquired signal. By scanning across the sample, images may be created. This is the basis of magnetic resonance imaging. To understand the concepts involved, we first examine the main components of an MRI system.

2.5.1 Main magnet

The acquisition of any NMR signal requires the presence of a static magnetic field. Main magnets for MRI typically operate at between 0.5 and 3.0 Tesla (T). For low field strengths (< 0.7 T), permanent or resistive magnets may be used, whilst superconducting magnets are employed for higher field strengths (> 0.7 T). In general, higher field strengths are preferred due to the better signal to noise ratio observed and narrower spectral widths⁶. However, problems with RF deposition and penetration limit, to a certain extent, the use of very high field strengths (> 4.0 T).

At the centre of the main magnet is an area of high field homogeneity called the region of interest (ROI). Typically, the ROI is spherical and exhibits an inhomogeneity in the magnetic field strength of less than 1-10 ppm over a 40-50 cm diameter spherical volume for human scale whole-body imaging. The term region of interest is often used interchangeably with diameter

⁵Non-homogeneous RF receiver coil reception is the key to the parallel imaging techniques such as SMASH [12] and SENSE [13]

⁶A notable exception to this is work by Agic *et al.* [14] at ultra-low field strength.

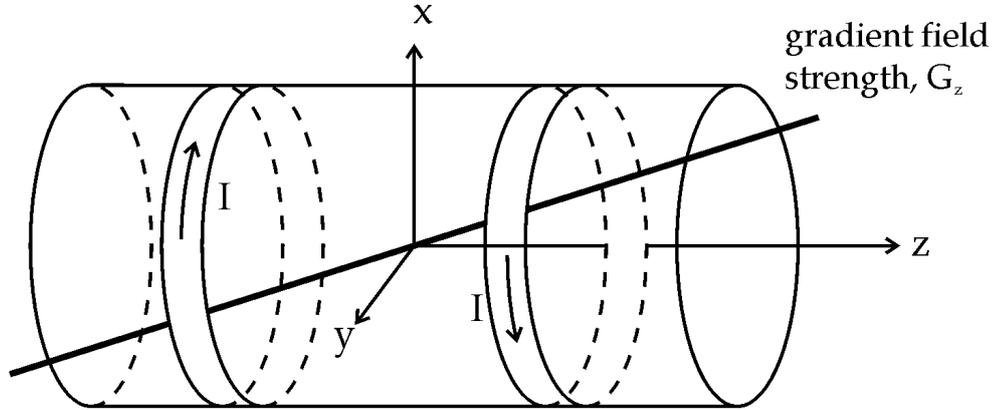


Figure 2.3: Schematic of a z gradient coil and the effect on the static main field strength. Gradient strength is given by $G_z = \frac{\partial B_z}{\partial z}$.

spherical volume (DSV) as both represent the area of high homogeneity where imaging occurs.

2.5.2 Gradient coils

The gradient coil system is the workhorse of an MRI scanner. It usually consists of three orthogonal coils which provide spatial resolution of the NMR signal by altering the magnetic field strength in a nonuniform manner. Neglecting any fields in the transverse direction, the coils create gradients in the magnetic field strength given by $\frac{\partial B_z}{\partial x}$, $\frac{\partial B_z}{\partial y}$ and $\frac{\partial B_z}{\partial z}$. An example of a z gradient set is shown in Figure 2.3. Typical gradient field strengths are 10-30 mT/m with rise times of the order 0.5 ms.

2.5.3 RF coils

The radio-frequency coils perform two important tasks. The transmit coil produces an RF excitation field, known as the B_1 field and the receive coil

acquires an electrical signal from the precessing bulk magnetization vector. In certain cases, the transmit and receive coils are the same. In solenoidal magnets, the transmit and receive coils are saddle shaped, although an alternative design, the *birdcage* coil is becoming more widespread. Usually, the receive coils are specially constructed for the area of body being imaged due to noise considerations and improvements in signal strength.

2.5.4 Computer system

A computer system is required for the overall control of the operation of an MRI system. Computers are involved in the control of gradient pulse sequences and the shaping and timing of RF pulses as well as data collection, manipulation and storage. Large amounts of data are likely to be created on a well used MRI system in a clinical situation and auto-archiving systems are now commonplace.

2.6 Signal Localization

The major difference between MRI and conventional NMR studies is the application of time-varying magnetic gradient fields to the sample. By imposing a linear variation on the z component of the magnetic field strength in the x , y and z directions, spatial information about the distribution of the nuclear spins may be obtained.

2.6.1 Slice selection

Many pulse sequences in MRI begin with a slice (or slab) selection gradient. A linear field gradient, G_z , is applied in the z direction resulting in a z dependent distribution of resonant frequencies. By altering the excitation frequency, it is therefore possible to acquire signals from different slices of the sample without any movement. The slice thickness, Δz , is given by

$$\Delta z = \frac{\Delta\omega}{\gamma G_z} \quad (2.28)$$

where $\Delta\omega$ is the RF pulse bandwidth. Due to limitations on gradient field strength, limits to the power of the RF pulses and their duration, the minimum slice thickness is about 2 mm in 1.5 T and 3 mm in a 1.0 T MRI system.

2.6.2 Frequency and phase encoding

Whilst the slice selection gradient gives resolution in the z direction, to obtain spatial resolution in the remaining two dimensions (x and y), it is necessary to introduce the mathematical concept of \mathbf{k} -space or reciprocal space. When a magnetic gradient \mathbf{G} is applied to a sample, the angular frequency of rotation in the rotating frame is

$$\omega(\mathbf{r}) = \gamma \mathbf{G} \cdot \mathbf{r} \quad (2.29)$$

where \mathbf{r} is a general three dimensional position vector. The signal received therefore becomes [9, p.348]

$$S(t) = \int_{\text{sample}} \rho(\mathbf{r}) e^{-i\gamma \int_0^t \mathbf{G}(\tau) \cdot \mathbf{r}(\tau) d\tau} d\mathbf{r} \quad (2.30)$$

which is exactly the Fourier transform of the spatial distribution of spin density if we define the \mathbf{k} vector as [7, p.160]

$$\mathbf{k}(t) = \frac{\gamma}{2\pi} \int_0^t \mathbf{G}(\tau) d\tau \quad (2.31)$$

In order to produce a complete two dimensional image of an excited slice, it is necessary to sample the whole of \mathbf{k} -space. Different sampling methods require different pulse sequences.

2.7 Pulse Sequences

There are a wide range of methods available to sample \mathbf{k} -space. This is one of the advantages of MRI as each pulse sequence potentially offers a different contrast mechanism. In this section three of the most common imaging

sequences are briefly described. See Brown [15] for a good overview of the range of available pulse sequences.

2.7.1 Spin echo imaging

In 2DFT⁷ spin echo imaging, a single line of \mathbf{k} -space is acquired for every RF excitation pulse. In the simplest case, the primary RF pulse rotates the magnetization vector by 90° , removing any longitudinal component. Field inhomogeneities and spin-spin interactions dephase the transverse magnetization vector on the time scale given by T_2^* (see Section 2.4.2), so to refocus these spins, a 180° pulse is applied at a time $TE/2$ which leads to an echo at a time TE after the initial RF pulse. Spatial encoding is achieved by applying gradients to manoeuvre to the appropriate location in \mathbf{k} -space using (2.31). Specifically, there are two dimensions which remain to be resolved. The first is resolved by frequency encoding. During signal acquisition, a gradient (G_x) is applied causing the spins to rotate at different frequencies across the sample in the x -direction. The second is resolved by phase encoding. A single application of gradient G_y straight after the slice selection RF pulse affects the phase of the spins across the sample in the y direction. Figure 2.4 shows the relative timings of the three gradients, the RF pulse and the signal acquisition. In the figure, each gradient is named using a term which describes its use in the 2DFT spin echo sequence. Thus, the terms slice selection (G_z), phase encoding (G_y) and frequency encoding (G_x) are used. As the data is acquired in \mathbf{k} -space, it must be Fourier transformed to create the final image. Figure 2.5 shows an example of some raw data acquired in \mathbf{k} -space and the resulting real-space image.

The major advantage of spin echo imaging is that, by refocusing the spins with a 180° RF pulse, it is possible to overcome the decay in signal strength due to field inhomogeneities and thus the signal decay is governed by the time constant T_2 rather T_2^* . However, spin echo imaging suffers from a relatively long acquisition time.

⁷2-dimensional Fourier transform

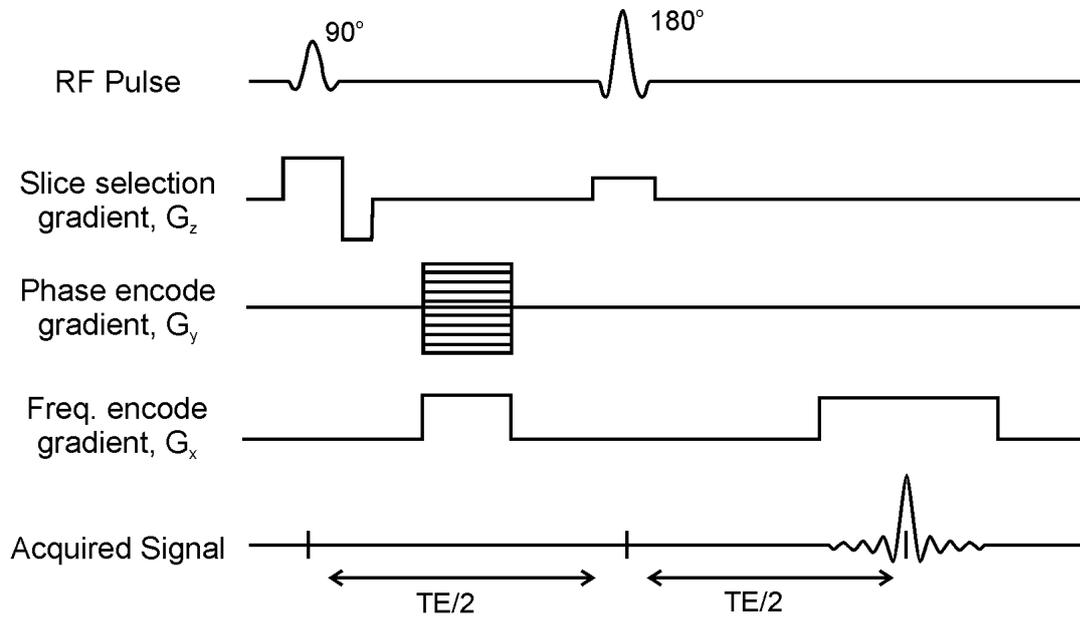


Figure 2.4: Spin echo pulse sequence showing the timing of the RF pulses, the linear gradients and the signal acquisition.

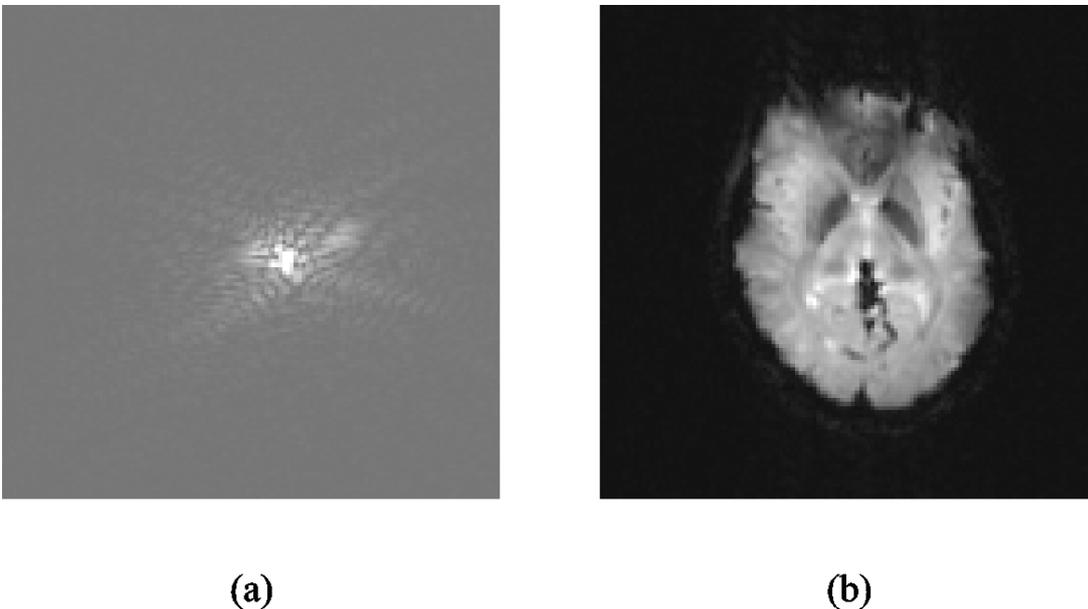


Figure 2.5: Example of (a) the acquired k -space data (magnitude) and (b) the corresponding real-space image. Each image contains 128×128 pixels. Courtesy of Pon Satangput.

2.7.2 Gradient echo imaging

It is also possible to refocus the spins using gradients. The echo occurs when the positive and negative contributions of a gradient, usually the frequency encoding gradient, cancel [15, p.60]. Although magnetization is recovered, gradient echoes cannot reduce the effects of inhomogeneity of the static magnetic field. Once again, it is necessary to sample the whole of \mathbf{k} -space. This is done by varying the magnitude of the phase encoding gradient after every RF pulse.

Two factors make gradient echo imaging faster than spin echo imaging. A small flip angle is generally used and no refocusing 180° RF pulse is required. A further advantage is that the RF energy deposition is reduced. Gradient echo sequences usually result in heavily T_1 weighted images.

2.7.3 Echo planar imaging

Fastest of all imaging sequences is echo planar imaging (EPI), in which the whole of \mathbf{k} -space is sampled within the T_2^* decay envelope from a single RF pulse. The image, which may be acquired in under 100 ms, is created by scanning \mathbf{k} -space in a raster fashion, using continuous oscillation of the frequency encoding gradient to create multiple echoes. In between each echo, the phase encoding gradient is pulsed, moving the location in \mathbf{k} -space vector along in the \mathbf{k}_y direction [15, p.68].

Echo-planar imaging is characterized by low resolution, low signal to noise ratio and images which are highly susceptible to field inhomogeneities or chemical shift effects. However, the power of such a rapid imaging technique and the continuous improvement in hardware, especially gradient hardware have made EPI an incredibly useful imaging sequence.

2.8 Effects of Field Inhomogeneity

NMR imaging is a technique with very high specificity due to an intrinsic property of the nucleus: the gyromagnetic ratio. The hydrogen (^1H) nucleus, as already discussed, has a gyromagnetic ratio of 42.6 MHz/T and

other nuclei will have different resonant frequencies such as ^{31}P which has a gyromagnetic ratio of 17.3 MHz/T [15, p.3]. Due to the natural abundance of ^1H nuclei in clinical situations, almost all MRI is performed using protons. When a group of nuclei are simultaneously excited at exactly the same frequency, they are collectively known as an isochromat.

In practice, a sample will not respond as a single isochromat exactly at the Larmor frequency. Instead, the existence of magnetic field inhomogeneities and the chemical shift effect will lead to a spreading of the resonant frequency. The chemical shift effect is due to the chemically heterogeneous environments of the individual spins. Each nucleus is ‘shielded’ to a certain extent by orbiting electrons and any interaction with neighboring atoms through chemical bonds will increase or decrease this shielding. Thus, each nucleus experiences a slightly different magnetic field

$$\hat{B}_0 = B_0(1 - \delta) \quad (2.32)$$

$$\text{and } \hat{\omega}_0 = \omega_0(1 - \delta) \quad (2.33)$$

where δ is usually of the order a few ppms. For example ‘fat’ protons (CH_2) display a 3.5 ppm shift from ‘water’ protons. If the maximum chemical shift is represented by ω_M , then the chemical shift bandwidth of the spin system is given by

$$|\omega - \omega_0| = \frac{\omega_M}{2} \quad (2.34)$$

The effect of inhomogeneities in the magnetic field cause a decrease in the time constant for the signal decay, see (2.24). In severe cases, where the transverse decay is effectively instantaneous, complete signal loss from regions of the image can be expected. Less severe consequences include image warping, image blurring and poor slice selection. These artifacts can be dealt with using correction algorithms if they are sufficiently small, however, large scale signal loss cannot be recovered by any image post-processing. Fortunately, unlike the dephasing effect of the time-variable microscopic spin-spin interactions, the dephasing due to the static magnetic inhomogeneities may be reversed by applying a spin echo RF pulse (Section 8). Thus a spin echo

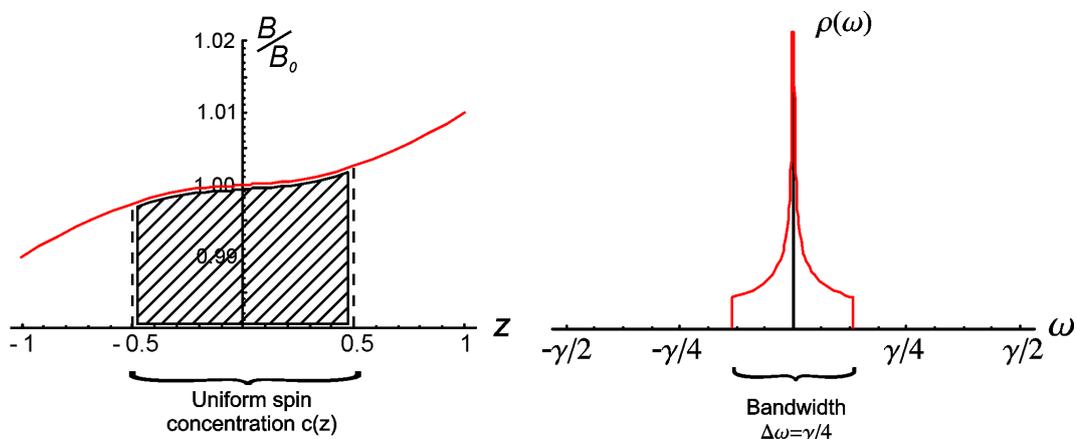


Figure 2.6: Example of the frequency spectrum for a sample in an inhomogeneous magnetic field $B(z) = B_0 + z^3$. The calculated spectral function is in practice smoothed by sensitivity and accuracy limits.

sequence can sometimes be used when no image is possible with a gradient echo or echo planar sequence.

Another effect of field inhomogeneity is to alter the frequency spectrum of the received signal, such as the signal from proton NMR spectroscopy. To examine this, let us define the spin spectral density function $\rho(\omega)$ such that

$$M = \int_{-\infty}^{\infty} \rho(\omega) d\omega \quad \text{and} \quad (2.35)$$

$$S(t) = \int_{-\infty}^{\infty} \rho(\omega) e^{-t/T_2(\omega)} e^{-i\omega t} d\omega \quad (2.36)$$

Now if, for example, a uniform sample with spin concentration $c(z) = \rho_0 \Pi(z)$ is placed in an inhomogeneous field $B(z) = B_0 + z^3$, we can pick an isochromat at an arbitrary position z and write its resonance frequency as

$$\omega(z) = \begin{cases} \gamma(B_0 + z^3) & |z| \leq 0.5 \\ 0 & \text{otherwise} \end{cases} \quad (2.37)$$

Substituting this into (2.36) gives

$$S(t) = \int_{-1/2}^{1/2} \rho_0 e^{-t/T_2} e^{-i\gamma(B_0 + z^3)t} dz \quad (2.38)$$

and using a change of variable, $\omega = \gamma(B_0 + z^3)$, we recover an expression for the spin spectral density function

$$\rho(\omega) = \begin{cases} \frac{\rho_0}{3\gamma^{1/3}(|\omega - \omega_0|)^{2/3}} & |\omega - \omega_0| < \gamma/8 \\ 0 & \text{otherwise} \end{cases} \quad (2.39)$$

The extent of the sample and the inhomogeneous field are shown in Figure 2.6 along with the resulting frequency spectrum of the received signal. In general⁸, asymmetric lineshapes result from misadjusted even-powered Z shims, whilst symmetrically broadened lines are usually the effect of misadjusted odd-powered Z shims such as Z^3 or Z^5 .

⁸In NMR experiments where the samples are spinning, more exotic side bands may appear at shifted frequencies due to X, Y, ZX and ZY inhomogeneities. Second order bands are a consequence of $X^2 - Y^2$ and XY inhomogeneities. As the samples are not spun in MRI, this effect is not seen.

Chapter 3

Theory Part II: Genetic Algorithms

3.1 Introduction

This chapter presents an introduction to the concepts and some of the basic theory behind genetic algorithms (GAs). Firstly a simple GA is described in detail and each stage of the optimization process is considered. GAs are presented as a series of operators acting on possible solutions to the optimization problem. The idea of constrained optimization is then examined and various methods of applying the constraints are presented. Next, some advanced GA techniques are introduced including elitism, sharing and crowding. A brief description of the theory behind GAs is presented which introduces the concept of building blocks and an important cornerstone of GAs: the Schema Theorem. Finally, some of the advantages of implementing GAs on parallel computer architectures are considered.

Genetic algorithms [16] are optimization techniques based on simulating the natural adaptation of biological systems to their environment. They are closely linked to the ideas of natural selection and survival of the fittest, first proposed by Darwin. Acting as a direct search method, GAs require only an evaluation of the objective (or fitness) function, whereas other optimization methods often involve the determination of the function gradients. Although

gradient based searches tend to locate an optimum faster for smoothly varying, single modal functions, they fail to perform well on discontinuous or non-differential problems [17]. This makes genetic algorithms important for real-world engineering design applications, where little information is known *a priori* about the fitness landscape. Another key advantage of genetic algorithms is their ability to deal with discrete values. Both of these reasons make genetic algorithms a logical choice as the optimization method in the design of novel magnetic resonance imaging magnets. Further details of comparisons between genetic algorithms and other stochastic optimization techniques can be found in Fisher, 1996 [18] and Williams, 2001 [19].

3.2 The Simple Genetic Algorithm

The execution of a simple genetic algorithm consists of five stages (or operators): initialization, evaluation, selection, recombination and mutation. Of these five stages, the last four are iterated, once every generation, until convergence criteria are met or a time limit expires. Figure 3.1 illustrates the basic sequence of events in a genetic algorithm. There are a myriad of possible implementations of the basic genetic algorithm, most of which are beyond the scope of this thesis. However, by taking each of the five main stages of the genetic algorithms in turn, key differences in approaches will be highlighted.

3.2.1 Coding the genetic information

Before exploring the details of the flowchart in Figure 3.1, the encoding of the optimized parameters must be examined. Any optimization problem may be cast in the following form

$$\text{Find } \vec{x} \text{ which optimizes } f(\vec{x}) \tag{3.1}$$

subject to

$$g_i(\vec{x}) \leq 0, \quad i = 1, \dots, n \tag{3.2}$$

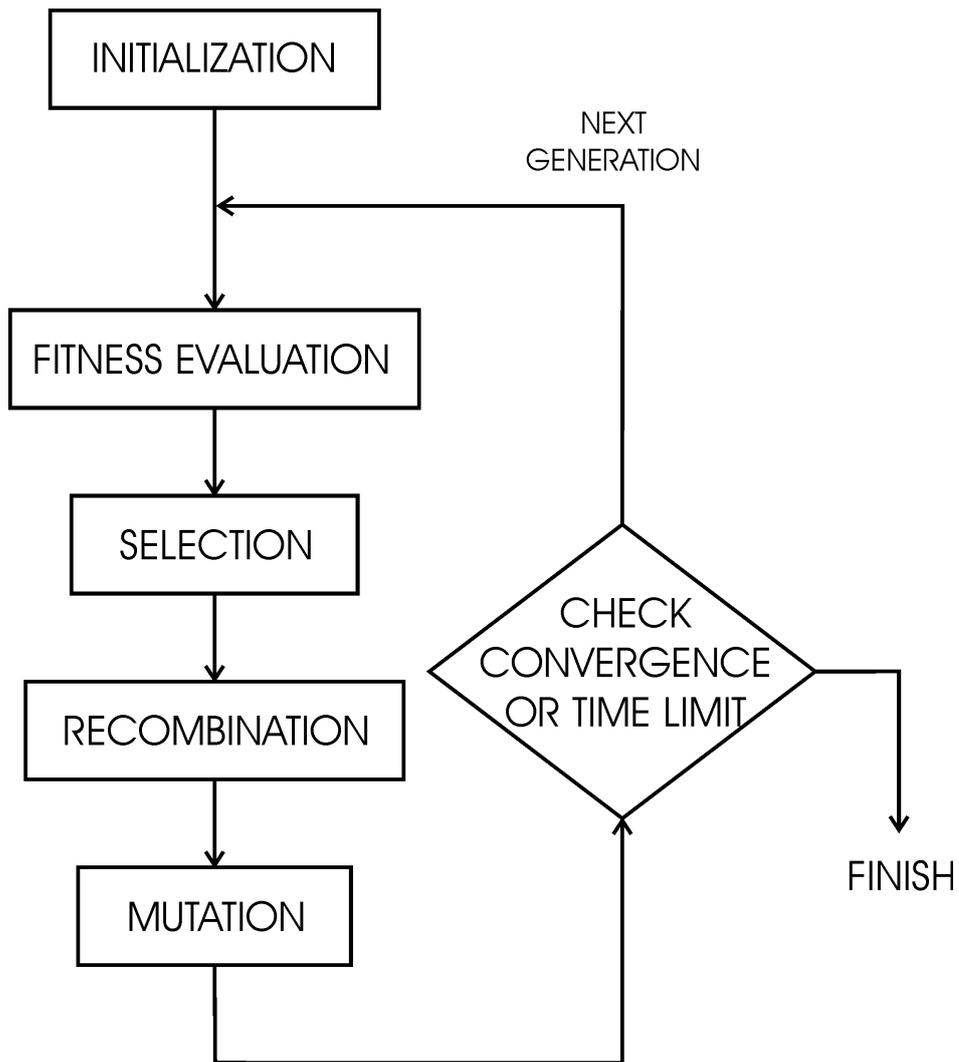


Figure 3.1: A flowchart showing the different stages of a simple genetic algorithm.

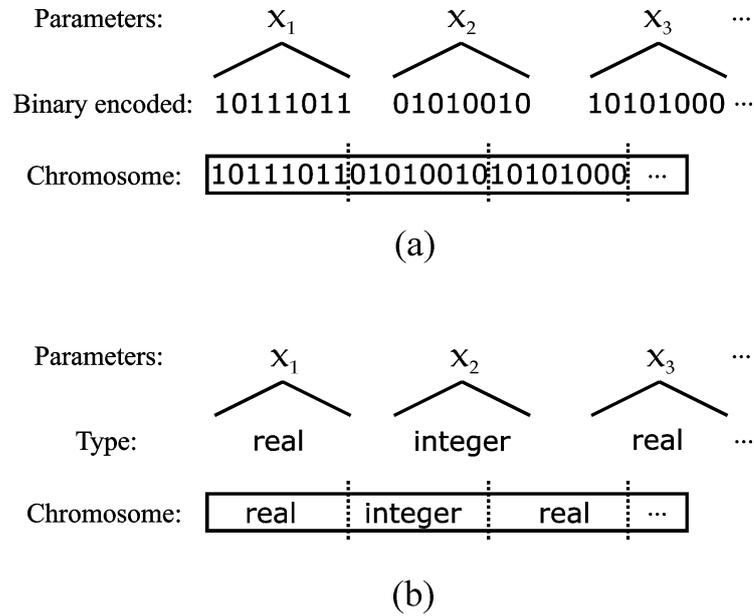


Figure 3.2: Illustration of the formation of the chromosome structure in (a) binary encoded GAs and (b) real coded GAs.

$$h_j(\vec{x}) = 0, \quad j = 1, \dots, p \quad (3.3)$$

where $\vec{x} = [x_1, x_2, \dots, x_r]^T$ is a vector consisting of the parameters to be optimized, $f(\vec{x})$ is the function to be optimized, n is the number of inequality constraints and p is the number of equality constraints. If the entire search space is given by \mathcal{S} and the feasible region of the search space by \mathcal{F} , then $\mathcal{F} \subseteq \mathcal{S}$. The functions $f(\vec{x})$, $g_i(\vec{x})$ and $h_j(\vec{x})$ may be linear or non-linear combinations of the r parameters x_1, \dots, x_r .

These parameters (x_1, \dots, x_r) are encoded into a string, known as a chromosome, which acts as the primary data structure. Encoding schemes vary between genetic algorithms and the parameters may be stored in binary or real-number format. A binary encoded genetic algorithm converts each parameter into a binary string whose length is determined by the required accuracy and range. The binary strings are then concatenated with the other parameters to form a single string which represents the vector \vec{x} (see Figure 3.2a). Binary encoding is usually associated with the simplest form of the genetic algorithm and leads to the formation of the building block schema

theorem proposed by Holland [20, 21]. On the other hand, real-coded genetic algorithms (RCGA) have become popular recently for engineering optimizations [22]. RCGAs create a chromosome by concatenating the real number parameters into a single data structure. The data structure may include integers as well, depending on the parameters encoded (see Figure 3.2b).

3.2.2 Initialization

There are two methods of initializing a GA population. The first is by far the most common method and allocates members of the population entirely randomly throughout the search space. The second method is to seed the GA with solutions which require prior knowledge about the problem. For example, a previously calculated solution may be used as a template for the entire population. Then, small random mutations are applied to each of the parameters to create a cluster of solutions around the initial solution. The other important factor in the initialization stage is the definition of the population size. Too large a population will unnecessarily prolong the search, whereas too small a population is likely to lead to premature convergence in the vicinity of a local minimum.

3.2.3 Fitness evaluation

The fitness evaluation step determines the shape of the search space and is often very tricky to construct in real-world optimization problems. Whereas optimization of simple mathematical functions may be straightforward, in real-world optimizations, the objective function is usually a combination of several competing factors. This leads to a messy and often indecipherable mix of objectives, all combined into a single overall solution fitness. If weights are associated with each of the m competing functions, $f_k(\vec{x})$, to be minimized we write

$$\begin{aligned} &\text{Optimize } f(\vec{x}) \\ &\text{where } f(\vec{x}) = \sum_{k=1}^m w_k f_k(\vec{x}) \end{aligned} \quad (3.4)$$

The weights, w_k , may be part of the original problem description or may have to be set by repeated runs of the genetic algorithm until a suitable balance of the required objectives is reached.

3.2.4 Selection

This stage is inspired directly from Darwin's theory of survival of the fittest. Members of the population are selected depending on their fitness and pass into a so-called parent pool, which undergoes the remaining operations of crossover and mutation, before being replaced in the general population. Several methods of selection are represented in the GA literature [16]

1. *Tournament selection*: two or more members of the population are chosen and their fitnesses are compared. The one which performs better is selected to have a copy placed in the parent pool. Both members are replaced in the population so multiple copies of a single member may be placed in the parent pool. This method has the property that the worst member of the population is never chosen. When two members are chosen from the population for comparison, this method is known as binary tournament selection.
2. *Roulette-wheel selection*: the probability of choosing a member of the population is directly proportional to its relative fitness compared to the average fitness of the population. Fitter members of the population therefore have a natural tendency to be selected.
3. *Ranking selection*: the probability of selection is proportional to the rank of the member within the population. This is similar to roulette-wheel selection except that the calculated fitness values are not important. This leads to a more consistent selection process.
4. *Deterministic sampling*: a selected percentage of the total population is chosen from predetermined ranks of the ordered population. This method is simple and straightforward to implement. For example the top half of the population may be placed automatically into the parent pool.

3.2.5 Recombination

Recombination is the process which now operates on solutions which have been passed into the parent pool. Two solutions are chosen at random from the parent pool and a random number $\chi \in \{0, 1\}$ is created. If $\chi \leq p_c$, where p_c is the probability of recombination, then crossover occurs, otherwise, if $\chi > p_c$ then the two chromosomes are placed, unaltered, into the next generation.

If crossover occurs in a binary encoded GA, a random site n is chosen along the chromosome string of length N , such that $n \in \{0, 1, \dots, N - 1\}$. The chromosome contents are copied directly to the offspring for all sites until the crossover point, after which, the genetic material is swapped between the two solutions (see Figure 3.3). The two offspring chromosomes are then placed into the next generation. Such a crossover is known as a *single-point* crossover operator. A *two-point* crossover operator may also be used whereby the genetic material between the two crossover points is swapped between the chromosomes. Recombination of the genetic material is very similar to choosing the best parts of two designs and seeing what happens if they are put together. Often the result will not be as fit, but occasionally, the result is better than the original at which point the overall fitness of the population will improve.

In real-number representation, recombination acts slightly differently. When a crossing point is chosen, there are two possibilities

1. In the first case, the crossing point lies at the boundary between two parameters. In this case, recombination is very similar to the binary method. The first offspring receives the first part of parent 1, and the second part of parent 2. Whilst the second offspring receives the remaining genetic information.
2. In the second case, the crossing point may lie within a single parameter. The most common way of dealing with this situation is using a crossover operator called BLX- α [22]. If the parameters of the two parents are given by x_i and x_j with, for example, $x_i < x_j$, then the offspring parameter is randomly assigned a value in the range $\{x_i - \alpha d, x_j + \alpha d\}$,

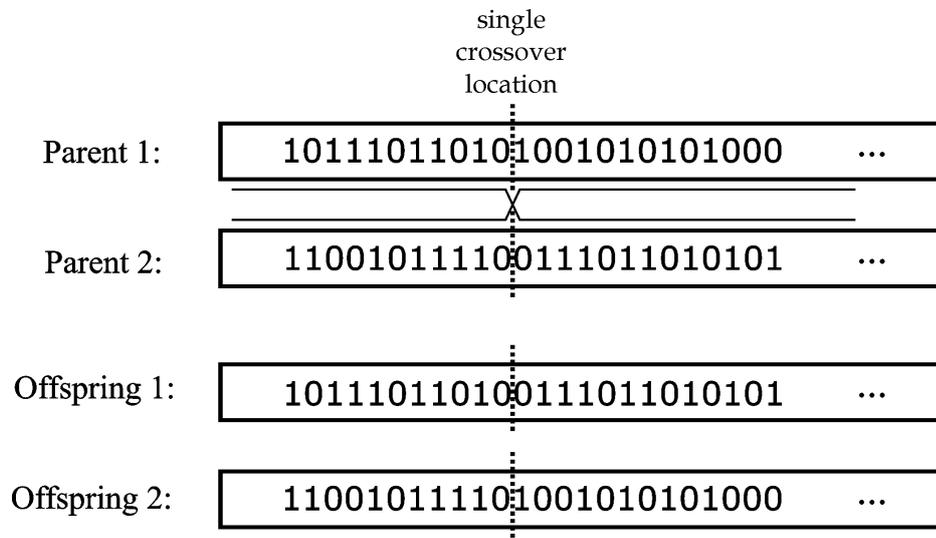


Figure 3.3: Schematic representation of a single point crossover in a binary encoded GA.

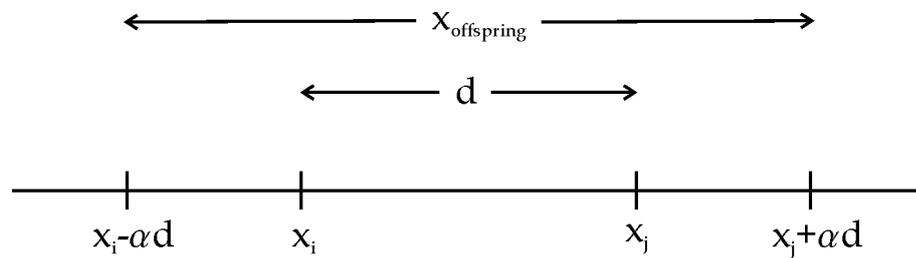


Figure 3.4: A real coded crossover operator, BLX- α . The offspring parameter is chosen with uniform distribution over the range $\{x_i - \alpha d, x_j + \alpha d\}$.

where $d = (x_j - x_i)$ (see Figure 3.4). If the offspring is assigned a value in the range $\{x_i, x_j\}$ then exploitation is said to have occurred. If the offspring is assigned a value between $\{x_i - \alpha d, x_i\}$ or $\{x_j, x_j + \alpha d\}$, then exploration is said to have occurred. This balance between exploration and exploitation is crucial and can be controlled by altering the parameter α . When $\alpha = 0.5$, there is equal probability of exploration and exploitation.

3.2.6 Mutation

After crossover, the chromosomes undergo mutation, thereby introducing new genetic material into the population. In binary encoded GAs, the mutation operator is very simple to understand. Each prospective member of the next generation must pass through a random process which flips a single bit with a given probability, say p_m . The spread of where the mutated chromosome bit string will map to in real space, compared to its pre-mutated location, depends on the type of encoding.

With binary encoding, one quickly recognizes the appearance of the Hamming Cliff problem [16]. For example, the binary string ‘10000’ (representing 16) must mutate all 5 bits to decrement its represented value by 1. Clearly, as p_m is in general a small number, the chances of altering all 5 bits are very slim. Figure 3.5a illustrates this problem and the dark region which runs diagonally across the diagram represents mutations which are particularly unlikely.

A way of avoiding this is to implement Gray binary encoding. In this encoding, any two adjacent numbers may be reached by the flipping of a single bit. Thus, the Hamming Cliff problem is avoided and, as illustrated in Figure 3.5b, the dark regions representing unlikely mutations are more widely dispersed, especially away from the diagonals.

In real-number encoded GAs, the action of the mutation operator is simulated by the application of a Gaussian distributed random variable. The spread of this mutation may be controlled by the standard deviation of the distribution, σ_m and, for comparison, the effect of a Gaussian distributed mutation operator is shown in Figure 3.5c.

3.3 Constraint Handling

Although it may be possible to pose simple problems in an unconstrained manner, many practical optimizations require several constraints or limits to be applied. A solution which violates one or more of these constraints is said to be unfeasible. Many researchers in evolutionary computing have examined

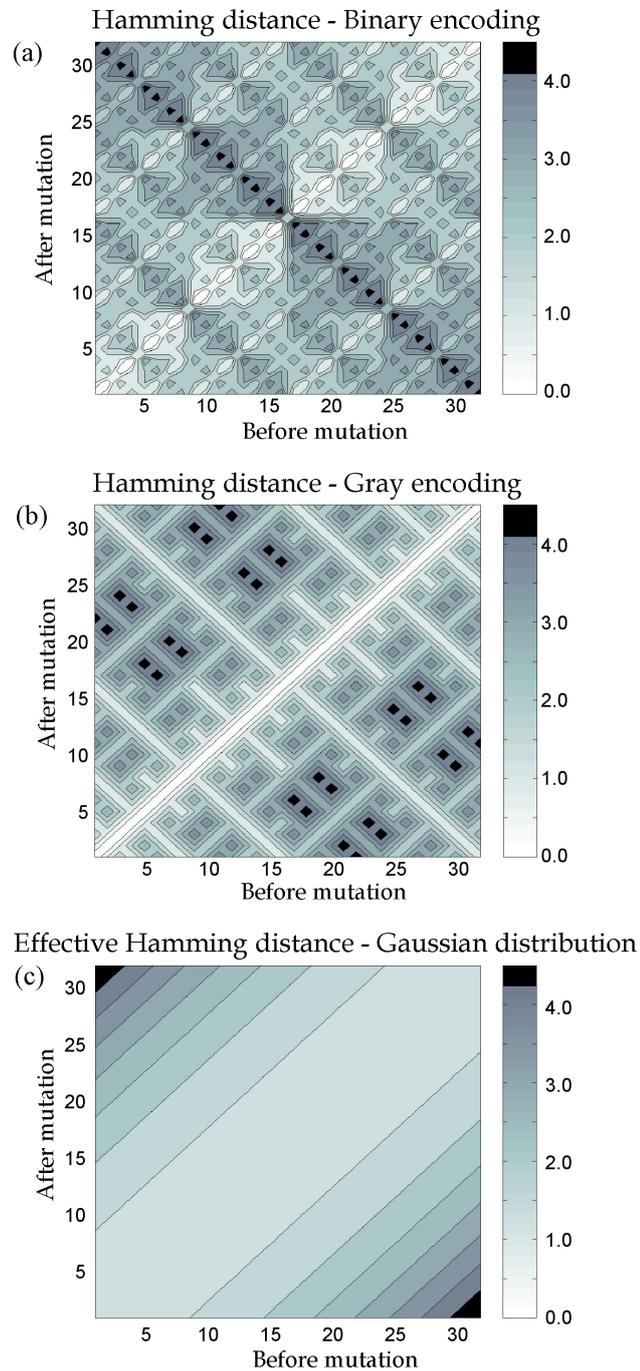


Figure 3.5: Illustration of mutations using different encodings. The value before mutation is plotted on the x axis and the value after mutation on the y axis. The probability of a particular mutation is given by the shading: white represents the most likely and black represents the least likely. Three possible encoding methods are shown, (a) binary encoded (b) Gray encoded (c) real number encoded.

ways of dealing with constrained optimization and detailed surveys exist on the various techniques currently employed [23, 24]. In this section, only a few of the most common constraint handling methods are introduced.

By far the most used constraint handling method in GAs is the *penalty method* [16]. A penalty function is added to solutions which violate one or more of the constraints, thereby degrading their solution fitness [16, 18, 19]. Thus

$$\bar{f}(\vec{x}) = \begin{cases} f(\vec{x}) + \sum_{i=1}^n w_i \Phi_i(f(\vec{x})) & \text{for unfeasible solutions} \\ f(\vec{x}) & \text{for feasible solutions} \end{cases} \quad (3.5)$$

In some cases, the penalty function, $\Phi(r(\vec{x}))$, may be a constant and in other cases it may be a measure of the unfeasibility of the solution. For each constraint, w_i is a measure of its importance. The main advantage of this method is its simplicity. However, many extra parameters are introduced which can be difficult to set and usually require some prior knowledge of the level of constraint violation. Furthermore, the inclusion of penalty functions warps the search space, which may be an undesirable side-effect.

Another very commonly used method is the rejection of unfeasible solutions, known as the *death penalty method* [23]. Every new solution created by the recombination and mutation operators is checked for constraint compliance. If the new solution violates one or more constraints then it is rejected and the operator is repeated iteratively until all constraint conditions are met. In optimizations which include a high proportion of unfeasible regions, this method can be very inefficient. However, implementation is very straightforward.

Thirdly, a novel method from Schoenauer and Michalewicz [25] may be used. The idea is to only search regions which lie close to the boundary between feasible and unfeasible solutions. Many constrained problems have global optima close to or directly on the boundary of at least one of the constraints. By focusing on these boundaries, the effectiveness of the search may be dramatically improved. Clearly, this type of approach is only possible where the global minimum is known (or at least suspected) to lie on the

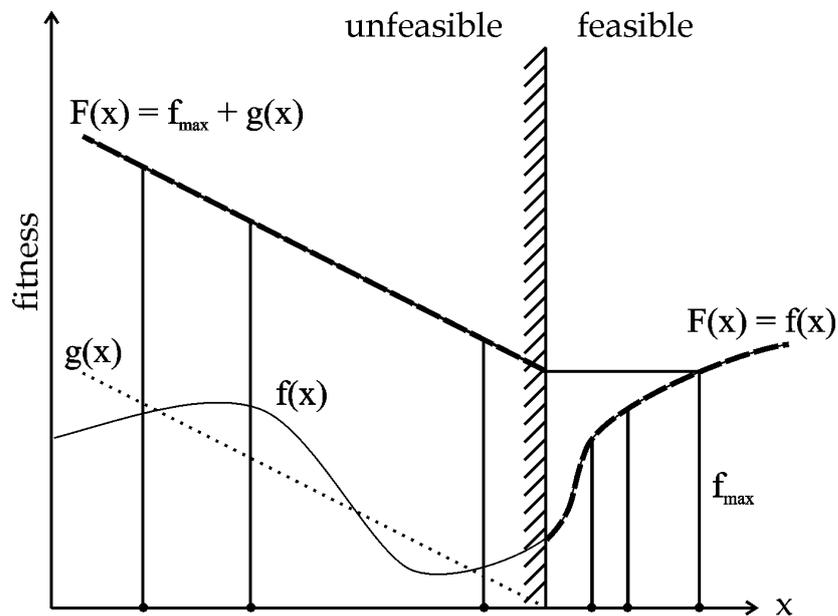


Figure 3.6: Heuristic constraint handling method. Six solid circles represent solutions in a GA population. $f(x)$ is the unconstrained fitness. The dotted line, $g(x)$, is the applied penalty function and the modified fitness of each solution is given by the heavily dashed line, $F(x)$. Reproduced from [17].

boundary of the feasible region. Also, this method provides no possibility to pass between unconnected feasible regions. However, in applicable situations, this method is efficient and produces good results [23].

Finally, heuristic rules may be applied to the evaluation of constraints. One such method, formulated by Deb [17], involves the following set of rules

- Any feasible solution is preferred to any unfeasible solution.
- Among two feasible solutions, the one having the better objective function value is preferred.
- Among two unfeasible solutions, the one having the smaller constraint violation is preferred.

The effect of this method is to converge rapidly on feasible regions of the search space whilst generally avoiding unfeasible regions during the optimization. An example of the operation of this constraint handling method is shown in Figure 3.6.

3.4 Advanced GA Methods

The basic operators used in a simple genetic algorithm have been described in the previous sections. In an attempt to improve the performance, both in terms of efficiency and flexibility, numerous modifications have been proposed [16, 26]. This section briefly describes some of the more widespread advanced GA methods.

- **Elitism** is commonly applied to a simple GA to prevent good solutions being lost. The elitist operator automatically includes the best individual from one generation in the next generation, bypassing the usual recombination and mutation operators.
- **Sharing** is the degradation of a particular individual's fitness due to the proximity of other individuals [16, 27]. In other words, a solution which is in the same neighbourhood as several other solutions will not be as fit as a similar solution without any near neighbours. One can think of sheep grazing on a hill side. If a single sheep finds a lush area, its fitness increases. However, if many sheep decide to also graze the same lush region, the amount of grass per sheep will be less and each sheep will have a lower fitness. When applied, this method tends to force the optimization away from premature convergence. Mathematically, the shared fitness $f_s(x_i)$ is given by

$$f_s(x_i) = \frac{f(x_i)}{\sum_{j=1}^n s(d(x_i, x_j))} \quad (3.6)$$

where $d(x_i, x_j)$ is a measure of the separation of the n solutions and s is the sharing factor. Although sharing is useful for avoiding premature convergence, it can be lengthy to calculate and the fitness of a specific location depends not only on the function being optimized, but also on the number of near neighbours in the population. This can lead to inconsistent fitness evaluations and poor exploitation of a particular optimum.

- **Crowding** is a similar method to sharing, but involves less calculation and is considered a ‘softer’ version of sharing [27, p.70]. The method was introduced by DeJong [28] in 1975 and uses the idea that, by forcing individuals to take the place of similar individuals in the next generation, higher levels of population diversity may be maintained. Typically, for every offspring, 2 or 3 individuals are chosen at random, where the number chosen is called the *crowding factor*. The Hamming distances, $d(x_i, x_j)$, are calculated and the most similar individual is then replaced by the offspring. Although on average, the effect of a crowding scheme is to allow niches to form, the effect is less dramatic when compared to sharing, as offspring may well replace individuals from other niches [29, p. 5].

Closely linked to crowding is the technique of *mating restriction* [30], where recombination is only permitted with a similar individual, calculated using the Hamming distance.

Many other advanced GA techniques exist including diploid or polyploid GAs, adaptive GAs, hybrid GAs and multi-objective GAs. Details of these techniques may be found in [16, 27, 31, 32].

3.5 The Schema Theorem

The discussion of schema theory is necessarily brief and offers merely a taste of the possible power of genetic algorithms. Originally formulated by Holland, the theory presented here is tailored from Goldberg’s book [16] and concentrates on the binary encoded simple genetic algorithm.

Consider first an extended binary alphabet $\{0, 1, *\}$ where the $*$ symbol represents the ‘don’t care’ symbol. Strings made up from this extended alphabet are called schemata and can be matched to actual chromosomes by comparing the individual bits, implicitly accepting either a 0 or a 1 when there is a $*$ in the schemata. Certain schemata will encode clear similarities between chromosomes and these schemata will naturally be passed on more readily to the next generation.

If we consider for example two schemata of length 5, $1***0$ and $**11*$. The first of these is much more likely to be disrupted by recombination or crossover than the second, which will tend to pass through to the next generation unmodified. Mutations usually have a much smaller effect and will tend to affect both schemata only slightly. From this it may be shown that highly-fit, short-length schemata are propagated most effectively whereas low-fitness or long-length schemata tend to disappear quickly. Highly-fit, short-length schemata are called *building blocks* and the optimal solution is usually a combination of these building blocks. This is called the *building block hypothesis*. It turns out that for a population of size n , requiring n fitness evaluations per generation, about n^3 schemata are usefully processed per generation. This implicit parallelism is one of the unique advantages of genetic algorithms [21, 33].

In schema theory, two measures play an important role, the *order* of individual schema and their *defining length*. The order of a schema (\mathcal{H}) is the number of fixed elements in the schema, written as $o(\mathcal{H})$ and is used to consider the effects of mutations. The defining length, $\delta(\mathcal{H})$, is the distance between the first and last fixed elements in the schema and is a measure of the probability of a particular schema being disrupted during recombination. For example, the schema $\mathcal{H}_1 = 10**0$ has $o(\mathcal{H}_1) = 3$ and $\delta(\mathcal{H}_1) = 4$ whereas the schema $\mathcal{H}_2 = **10*$ has $o(\mathcal{H}_2) = 2$ and $\delta(\mathcal{H}_2) = 1$. The probability of a schema being disrupted during recombination at every generation is

$$p_c \frac{\delta(\mathcal{H})}{l-1} \quad (3.7)$$

where l is the length of \mathcal{H} . Similarly, the probability of the schema being disrupted during mutation for a single generation may be calculated

$$o(\mathcal{H})p_m \quad (3.8)$$

At a certain time, t , if the schema \mathcal{H} occurs m times in the population, then we can form a lower bound on the expected occurrence of \mathcal{H} in the next

generation [16, p.33]

$$m(\mathcal{H}, t + 1) \geq m(\mathcal{H}, t) \cdot \frac{f(\mathcal{H})}{\bar{f}} \left[1 - p_c \frac{\delta(\mathcal{H})}{l - 1} - o(\mathcal{H})p_m \right] \quad (3.9)$$

where \bar{f} is the average fitness of the entire population and $f(\mathcal{H})$ is the average fitness of the strings representing schema \mathcal{H} at time t . The conclusion to be drawn from (3.9) is that short, low-order, above-average schema are by far the most likely to survive and will therefore form the building blocks of any solution. This statement is known as the *Schema Theorem* or Fundamental Theory of Genetic Algorithms.

A genetic algorithm's success depends upon the recombination of building blocks to seek better function values. If the building blocks are misleading due to the use of an unsuitable coding scheme, then the efficiency of the genetic algorithm may become reduced and the algorithm will tend towards a limited random search.

3.6 Parallel Implementation

One of the most appealing qualities of genetic algorithms is their intrinsic parallel nature. To take advantage of this, the fitness evaluation routine may be spread over a parallel computing architecture. Although there must still remain some serial processing associated with the selection, recombination and mutation operators, for real-world optimization problems with complex fitness evaluation routines, these operators are often a negligible part of the overall run-time.

There are many ways of categorizing parallel genetic algorithms (pGAs) and often a division is drawn between coarse-grained and fine-grained models [34, 35]. Coarse-grained pGAs come in two flavours. The first, consists of multiple subpopulations, where each subpopulation evolves in an isolated or semi-isolated manner. This method is also known as the island model GA and each subpopulation is called a *deme*. A second coarse-grained method is the master-slave configuration, where the recombination and evaluation steps are shared out to individual processors, whilst the master processor

maintains control over the selection operations. Figure 3.7 illustrates these two coarse-grained pGA models.

Fine-grained pGAs divide the population into much smaller subpopulations, often containing only 1 or 2 individuals, and implement a network of communications in the vicinity of each subpopulation for selection and recombination, [29, p. 139] and [36, p. 67].

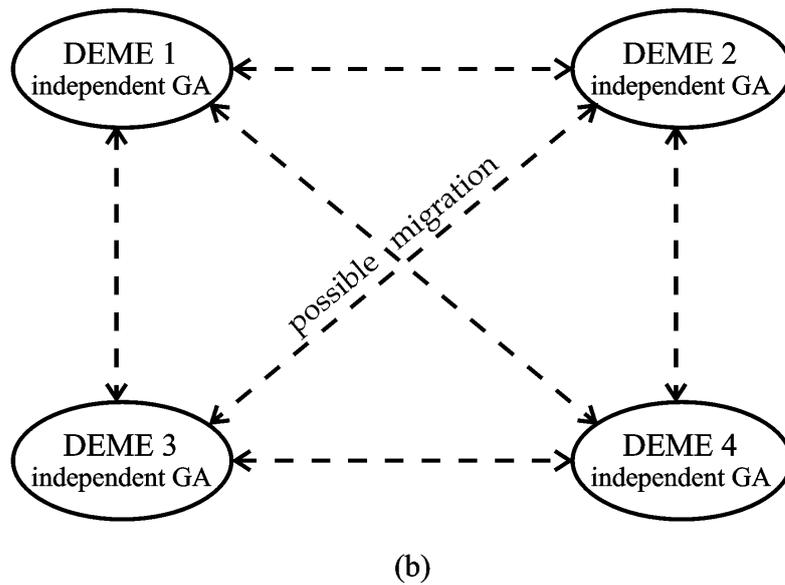
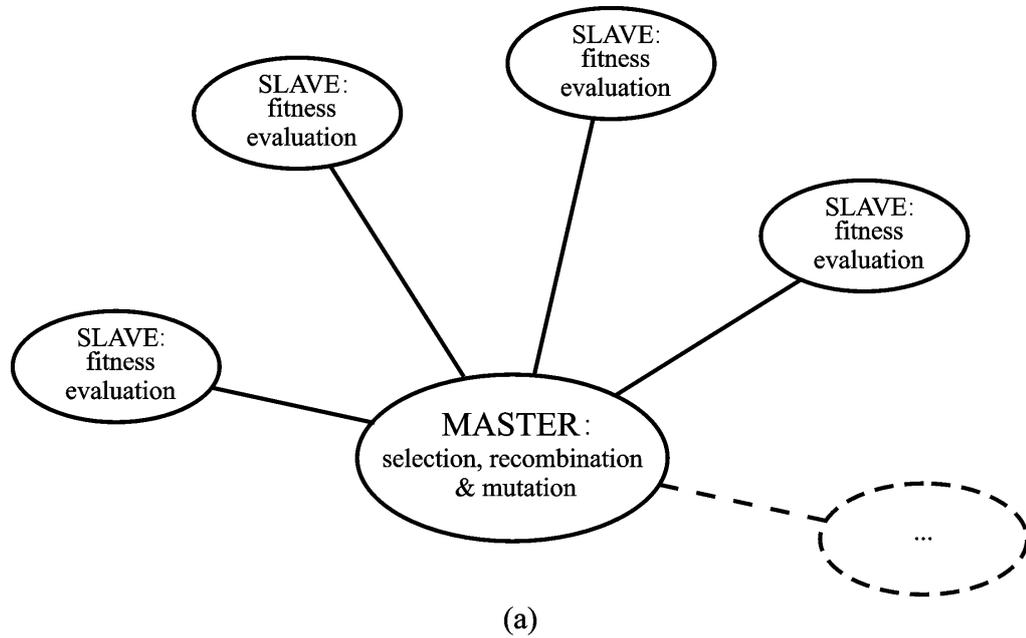


Figure 3.7: Two methods are shown for parallelizing GAs. In (a), the master-slave paradigm is used, where the fitness evaluation operation is shared in parallel. In the island model (b), GAs are operated in parallel with possible migration between demes at predetermined intervals.

Chapter 4

Design Framework for Axisymmetric Magnet Optimization

4.1 Introduction

In order to create and compare axisymmetric magnet designs, it is necessary to define criteria by which potential designs can be judged. In this chapter, a robust and detailed framework for considering the problem of axisymmetric magnet design is presented. Whilst most of what is presented can be applied to general axisymmetric magnet design, emphasis is particularly placed on the task of MRI main-field magnet design. The framework covers:

- Magnetic field evaluation, especially related to the homogeneous central region.
- Fringe-field evaluation methods, including a novel theoretical approach to rapid fringe-field calculation.
- Design sensitivity to build errors, including a Monte Carlo method and a novel analytical approach.
- Force and stress calculations, both approximations and finite element methods.

- Stored energy, inductance calculation and quench modelling.
- Physical constraints such as magnet size and wire volume.

The main aims of this chapter are to draw together previous work on axisymmetric MRI magnet design into a coherent framework and to present two pieces of original and as yet unpublished theoretical work contained within Sections 4.3 and 4.4.

4.2 Magnetic Field Evaluation and Central Field Homogeneity

Accurate methods to evaluate the magnetic field produced by axisymmetric coils are important in all magnet design work. For MRI magnet design, accuracy is particularly important as the intrinsic inhomogeneity of the imaging region should be equal to or less than the susceptibility inhomogeneities induced by the patient [37]. Typically, in an MRI magnet with a 1.5 T central field ($B_0 = 1.5 T$), the magnitude of the inhomogeneity is in the range 1.5 to 30 mgauss (or 0.1 parts per million (ppm) to 2 ppm) over the imaging region and therefore, calculations should be accurate to at least 1 part in 10^8 . In general, the homogeneous region in the centre of an MRI magnet is known as the region of interest (ROI) or the diameter spherical volume (DSV). Usually, spherical ROIs are used for optimization, but recent work has shown the feasibility of non-spherical geometries [37].

The remainder of this section introduces several methods currently in use for magnetic field evaluation [38, 39, 40].

4.2.1 Direct Biot-Savart integration

To evaluate the field produced by a thick solenoid, the simplest method uses the Biot-Savart law directly. The evaluation of the Biot-Savart law for any configuration of wires is a fairly straightforward exercise. Coordinate systems, though, used for the evaluation of the magnetic fields can be rather confusing. In this chapter, two systems are employed - Cartesian coordinates

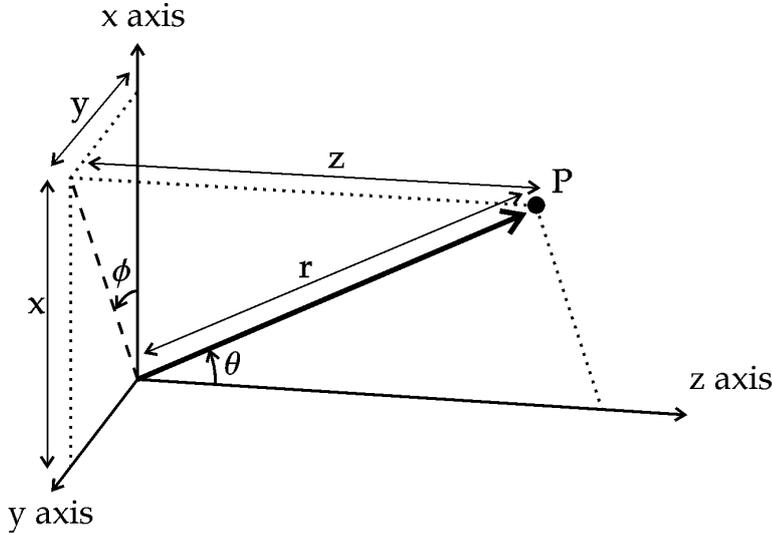


Figure 4.1: Comparison of Cartesian coordinates with spherical coordinates.

and spherical coordinates. Figure 4.1 shows the relationship between the two systems. As an example, the expression for the magnetic field from a loop carrying a current, I , is [41]

$$\mathbf{B} = \frac{\mu_0 I}{4\pi} \int_0^{2\pi} \frac{d\mathbf{s} \times (\mathbf{r}' - \mathbf{r})}{|\mathbf{r}' - \mathbf{r}|^3} \quad (4.1)$$

where \mathbf{r}' is the vector to the current source, \mathbf{r} is the position vector of the observation point, $d\mathbf{s}$ is a vector around the loop and μ_0 is the permeability of free space. Numerous references are available in which formulae are derived for other geometries such as arcs and straight-line segments [18, 41].

In some optimization processes such as when designing gradient coils, it is sufficient to model the coils as a collection of infinitely thin loops, arcs and segments [18, 19]. However, for the main field magnets, the coils must be approximated as thick windings, carrying an evenly spread current density. Thus, the current in a single loop, I , is replaced by the effective current, $J'' dz' dx'$. For the case of an axisymmetric thick coil around the z axis (see Figures 4.1 and 4.3) the azimuthal component of the magnetic field vanishes

and, in the plane $y = 0$, B_z and B_x are given by:

$$B_x = \frac{\mu_0 J''}{4\pi} \int_{x'} \int_{z'} \int_{\phi'} \frac{r' \sin \theta' \cos \phi' (r \cos \theta - r' \cos \theta') d\phi' dz' dx'}{d^3} \quad (4.2)$$

$$B_z = \frac{\mu_0 J''}{4\pi} \int_{x'} \int_{z'} \int_{\phi'} \frac{r' \sin \theta' (r' \sin \theta' - r \cos \theta \sin \phi') d\phi' dz' dx'}{d^3} \quad (4.3)$$

$$\text{where } d = (r^2 + r'^2 - 2rr' \cos \theta \cos \theta' - 2rr' \sin \theta \sin \theta' \cos \phi')^{1/2} \quad (4.4)$$

Depending on the exact method of winding employed, the coil may also have a slight spiral aspect. This, however, does not significantly affect the magnetic field, as the layer above is wound in the opposite direction and cancels any non-azimuthal effects.

Although the above integrals (4.2)-(4.4) may be performed using Gaussian quadrature, for a large number of coils this method tends to be painstakingly slow, due to integrating over three dimensions (x' , z' and ϕ') and is only worthwhile as a cross-check when more advanced techniques are tested.

4.2.2 Hybrid integral methods

Rather than performing all three of the integrals of (4.2)-(4.4) numerically, it is possible to analytically evaluate the z' integral, leaving the following expressions [42] involving the complete elliptic integral $cel(k_c, \rho, \alpha, b)$

$$B_z = 0.1 J'' \left[\int_{a_1}^{a_2} \frac{4(z-w)t}{\beta \sqrt{(z_b + \beta)}} cel\left(\gamma, \frac{\alpha}{\beta}, t-x, t+x\right) dt \right]_{w=b_2}^{b_1} \quad (4.5)$$

$$B_x = 0.1 J'' \left[\int_{a_1}^{a_2} \frac{4t}{\sqrt{(z_b + \beta)}} cel(\gamma, 1, 1, -1) dt \right]_{w=b_1}^{b_2} \quad (4.6)$$

where

$$z_b = (z-w)^2 \quad (4.7)$$

$$\alpha = (x-t) \quad (4.8)$$

$$\beta = (x-t) \quad (4.9)$$

$$\gamma = \sqrt{\frac{z_b + \alpha}{z_b + \beta}} \quad (4.10)$$

and

$$\text{cel}(k_c, \rho, \alpha, b) = \int_0^\infty \frac{\alpha + b\zeta^2}{(1 + \rho\zeta^2)\sqrt{(1 + \zeta^2)(1 + k_c^2\zeta^2)}} d\zeta \quad (4.11)$$

The dimensions of the coil are given by : $b_1 = z_{\min}$, $b_2 = z_{\max}$, $a_1 = x_{\min}$, $a_2 = x_{\max}$ with the current density, J'' , expressed in A/m². The observation point at which the field is calculated is defined by $(x, 0, z)$, with no loss of generality. The above method is referred to in this thesis as the CEL method due to the central role played by the evaluation of the complete elliptic integral. Fortunately, efficient methods exist for the numerical evaluation of the complete elliptic integral [43] and this method is consistently much faster than the direct integration of the Biot-Savart equation.

A similar method is one recently proposed by Forbes *et al.* [40]. In their derivation, two integrals (x' and z') are solved analytically, resulting in the expressions

$$B_z = \frac{\mu_0 J''}{2\pi r} [[\mathbb{L}(R_0 + \beta, \alpha; r, z)]_{\alpha=-a}^a]_{\beta=-b}^b \quad (4.12)$$

$$B_z = -\frac{\mu_0 J'' x}{2\pi r} [[\mathbb{N}(R_0 + \beta, \alpha; r, z)]_{\alpha=-a}^a]_{\beta=-b}^b \quad (4.13)$$

where the functions $\mathbb{L}(r', z'; r, z)$ and $\mathbb{N}(r', z'; r, z)$ are defined in [40] and involve a single integral of the angle ϕ' over the range $0 \rightarrow \pi$. This method is very useful for calculating the magnetic fields at points within current carrying coils. Care must be taken though near coil edges as accuracy can be lost more rapidly during integration when compared to the CEL method [44].

4.2.3 Harmonic decomposition methods

The harmonic decomposition methods reduce the magnetic field to a weighted series of complete orthogonal functions. The functions employed depend on the coordinate system: in cylindrical coordinates, the modified Bessel functions [45] are used; whereas in spherical coordinates, the Legendre polynomials are employed. For historical reasons, spherical coordinates are more

widely used, even in coaxial magnet design [38, 39]. The decomposition method is similar to the conventional description of atomic orbitals, where the orbitals s, p, d, f, . . . have increasing angular dependence. In magnet design, each of these terms is known as an *order*, so for example, an infinitely long solenoid contains only a *zeroth order* magnetic component. Likewise, a perfect z -gradient would only contain a *first order* component. In practice, all of the orders are present in a magnetic field to a certain degree due to the finite dimensions of coils and manufacturing errors.

Garrett (1951) [38] and Romeo and Hoult (1984) [39] both showed that the weightings of the orthogonal harmonic functions can be calculated uniquely from the configuration of the current sources (circular loops, cylinders or solenoids). Using these weights (q_n), the magnetic field may be calculated by

$$B_z = \sum_{n=0}^{\infty} q_n \left(\frac{r}{r_0}\right)^n P_n(\cos \theta) \quad (4.14)$$

$$B_x = \sum_{n=0}^{\infty} \frac{1}{(n+1)} q_n \left(\frac{r}{r_0}\right)^n P_n^1(\cos \theta) \quad (4.15)$$

where $P_n(\cos \theta)$ and $P_n^1(\cos \theta)$ are the Legendre and associated Legendre polynomials respectively (see Appendix A for details). The evaluation of the magnetic field using the weightings of orthogonal harmonic functions is even more efficient than the hybrid integral methods.

The accuracy of this method is determined by a combination of two quantities: the distance from the origin and the number of terms taken in the infinite expansion. The radius of convergence, r_c , is the maximum radius of a sphere centred on the origin which does not contain any current sources. If an attempt is made to use the central field expansion beyond this radius, then the infinite series will not converge and other methods must be employed.

Appendix A revisits Garrett's analysis of the spherical harmonic decomposition method, reformulating the results where necessary to take account of modern SI units¹.

¹Siystème International d'Unités

4.2.4 Other methods

Both finite element methods (FEM) and boundary element methods (BEM) may be used to calculate the magnetic fields of arbitrarily shaped magnet designs. These powerful techniques however are complex and require significant computational resources due to the creation of a mesh and the simultaneous solution of partial differential equations. Therefore these methods are not used in conjunction with any stochastic optimization techniques. Instead finite element programs such as FEMM [46] or OPERA [47] are better suited to checking the performance of existing designs and, for example, evaluating the effects of additional bulk iron.

Lastly, it is possible to create a scaled, two-dimensional lookup table for the calculated field from a solid, semi-infinitely long solenoid. This method was promoted by Brown and Flax [48], where they suggest the use of the dimensionless parameters $\alpha = \frac{\text{outer diameter}}{\text{inner diameter}}$ and $\beta = \frac{\text{length}}{\text{inner diameter}}$. Clearly, any thick solenoid can subsequently be created from the superposition of four such semi-infinite solenoids. However, for high accuracy, the size of the lookup table rapidly diverges and therefore this method is included for completeness only.

4.2.5 Homogeneity - a single figure of merit

Homogeneity, as a single figure of merit, may be determined in a combination of ways. There is a lack of consistency in the literature making direct comparison difficult, if not impossible, unless the exact form of the field or the original coil dimensions are available [49, 50, 51]. Methods for describing field homogeneity fall into two basic categories,

- The maximum to minimum deviation of $|\mathbf{B}|$ or B_z over the ROI expressed in parts per million (ppm) where B_0 is the average central field strength

$$\Delta B = \frac{|B_{\max} - B_{\min}|}{B_0} \times 10^6 \quad (4.16)$$

- The standard deviation of $|\mathbf{B}|$ or B_z over the ROI, also expressed in ppm [52, p. 495].

In both cases, the deviation of the field may be taken from the average field strength or from the field strength at the iso-centre of the magnet system. Another possible method to calculate the field distribution is to sample the field strength at a grid of points contained within the ROI. The maximum to minimum deviation and the standard deviation are then found for this subset of points. As the number of sampled points increases, so the calculated value approaches the results from the continuous methods.

4.3 Novel Fringe Field Evaluation Method

Evaluating the fringe field is important in the design of magnets and especially MRI magnets due to the potential safety hazards associated with stray fields. The magnetic field in the region outside the magnet is usually calculated using one of the direct hybrid methods discussed in Section 4.2.2. In this section however, a novel decomposition method is presented for the fringe fields which reduces calculation times significantly. Spherical coordinates are used in this derivation to maintain consistency with previous work in this field and a general bias in industry [38]. It is understood that expansions of the magnetic field have been previously performed in cylindrical polar coordinates, although these are unpublished [53].

4.3.1 Fringe field expansion of loops, thin solenoids and thick solenoids in spherical harmonics

Infinitely Thin Loops

Let us consider a single, infinitely thin loop of wire carrying a current I . The loop, also known as a filament (f), has rotational symmetry about the z axis and its position can be determined entirely in standard spherical coordinates by the primed variables $(r', \theta', 0)$ where a point on this loop is given by the vector (r', θ', ϕ') . We also create an observation point \mathcal{P} , which is denoted by the unprimed variables (r, θ, ϕ) . Due to the symmetry of the problem about the z axis, we can constrain the observation point \mathcal{P} to lie in the xz plane,

denoted by $y = 0$, without any loss of generality.

Now the magnetic vector potential, \mathbf{A} , which is related to the magnetic field (\mathbf{B}), by $\mathbf{B} = \nabla \times \mathbf{A}$, is given by the generalized Biot-Savart law [41]:

$$\mathbf{A}(\mathbf{r}) = \frac{\mu_0}{4\pi} \int \int \int_{\text{volume}} \frac{\mathbf{J}(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r}'^3 \quad (4.17)$$

where the source coordinates are all given by $\mathbf{r}' = (r', \theta', \phi')$, the observation point \mathcal{P} by $\mathbf{r} = (r, \theta, \phi)$, μ_0 is the permeability of free space and $\mathbf{J}(\mathbf{r}')$ is the current density at point \mathbf{r}' .

An infinitely thin loop is illustrated in Figure 4.2, and hence the following quantities are trivially expressed in Cartesian coordinates:

$$\mathbf{ds} = \begin{pmatrix} -r' \sin \theta' \sin \phi' \\ r' \sin \theta' \cos \phi' \\ 0 \end{pmatrix} d\phi' \quad (4.18)$$

$$\mathbf{r} = \begin{pmatrix} r \sin \theta \\ 0 \\ r \cos \theta \end{pmatrix} \quad (4.19)$$

$$\mathbf{r}' = \begin{pmatrix} r' \sin \theta' \cos \phi' \\ r' \sin \theta' \sin \phi' \\ r' \cos \theta' \end{pmatrix} \quad (4.20)$$

Substituting these into (4.17) we find that:

$$\mathbf{A} = \frac{\mu_0 I a}{4\pi} \int_0^{2\pi} \frac{(-\sin \phi', \cos \phi', 0) d\phi'}{|\mathbf{r} - \mathbf{r}'|} \quad (4.21)$$

where a is the radius of the loop and

$$|\mathbf{r} - \mathbf{r}'| = \sqrt{r^2 + r'^2 - 2rr'(\cos \theta \cos \theta' + \sin \theta \sin \theta' \cos \phi')} \quad (4.22)$$

It is clear that, as \mathcal{P} lies in the xz plane, A_x vanishes due to symmetry,

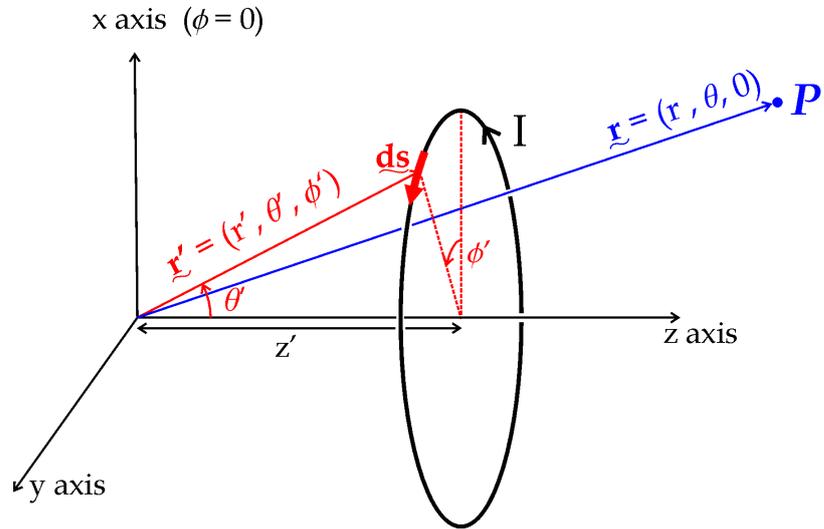


Figure 4.2: Axisymmetric loop defined by (r', θ') , carrying current I and observation point \mathcal{P} .

leaving only A_y , or equivalently, A_ϕ , when expressed in spherical coordinates

$$A_\phi = \frac{\mu_0 I a}{4\pi} \int_0^{2\pi} \frac{\cos \phi' d\phi'}{|\mathbf{r} - \mathbf{r}'|} \quad (4.23)$$

We now use the Green's function expansion in spherical coordinates [41],

$$\frac{1}{|\mathbf{r} - \mathbf{r}'|} = 4\pi \sum_{l=0}^{\infty} \sum_{m=-l}^{m=l} \frac{1}{2l+1} \frac{(r')^l}{r^{l+1}} Y_{lm}^*(\theta', \phi') Y_{lm}(\theta, \phi) \quad (4.24)$$

and the standard definitions for spherical polynomials², to rewrite (4.23) as

$$A_\phi = \frac{\mu_0 I a}{4\pi} \Re \left\{ \sum_{l=0}^{\infty} \sum_{m=-l}^{m=l} \int_0^{2\pi} \frac{(r')^l (l-m)!}{r^{l+1} (l+m)!} P_l^m(\cos \theta') P_l^m(\cos \theta) e^{-i(m-1)\phi'} d\phi' \right\} \quad (4.25)$$

Where $Y_{lm}(\theta, \phi)$ are the normalized spherical harmonics and $P_l^m(\cos \theta)$ are the associated Legendre polynomials. The next step is key to this whole argument and hinges on the fact that the integral over ϕ' averages to zero for all m except the special case where $m = 1$. This allows the sum over m to be removed from the equation, leaving

$$A_\phi = \sum_{l=1}^{\infty} \underbrace{\frac{\mu_0 I \sin \theta' (r')^{l+1}}{2l(l+1)} P_l^1(\cos \theta')}_{\text{source}} \underbrace{\frac{1}{r^{(l+1)}} P_l^1(\cos \theta)}_{\text{observation}} \quad (4.26)$$

Note that the limits of the sum in (4.26) have altered because the term with $l = 0$ now vanishes. Also, it is clear that (4.26) can be separated into two terms, one determined entirely by the positioning of the current sources, and the other dependent only on location of the observation point (\mathcal{P}). Thus, when evaluating the fringe field from a specific design, the source dependent part of A_ϕ may be calculated as an array of constants, allowing much faster evaluation of the magnetic field subsequently at any location in the far-field regime.

We are now in a position to find an expression for the magnetic field vector \mathbf{B} from a filament using $\mathbf{B} = \nabla \times \mathbf{A}$.

$$\begin{aligned} B_r &= \frac{1}{r \sin \theta} \frac{\partial}{\partial \theta} (\sin \theta A_\phi) \\ &= \frac{1}{r \sin \theta} \sum_{l=1}^{\infty} \frac{\mu_0 I \sin \theta'}{2l(l+1)} \frac{(r')^{l+1}}{r^{l+1}} P_l^1(\cos \theta') \frac{\partial}{\partial \theta} \left\{ \sqrt{1 - \cos^2 \theta} P_l^1(\cos \theta) \right\} \end{aligned}$$

²To be able to fully understand the following arguments, the reader should be familiar with many of the standard results for spherical harmonics and Legendre polynomials. References are widely available on this subject including Abramowitz and Stegun [54] and [55]

$$\begin{aligned}
&= \sum_{l=1}^{\infty} -\frac{\mu_0 I \sin \theta'}{2l(l+1)} \frac{(r')^{l+1}}{r^{l+2}} P_l^1(\cos \theta') \frac{\partial}{\partial \cos \theta} \left\{ \sqrt{1 - \cos^2 \theta} P_l^1(\cos \theta) \right\} \\
B_r &= \sum_{l=1}^{\infty} -\frac{\mu_0 I \sin \theta'}{2} \frac{(r')^{l+1}}{r^{l+2}} P_l^1(\cos \theta') P_l(\cos \theta) \tag{4.27}
\end{aligned}$$

and similarly,

$$\begin{aligned}
B_\phi &= -\frac{1}{r} \frac{\partial}{\partial r} (r A_\phi) \\
B_\phi &= \sum_{l=1}^{\infty} \frac{\mu_0 I \sin \theta'}{2(l+1)} \frac{(r')^{l+1}}{r^{l+2}} P_l^1(\cos \theta') P_l^1(\cos \theta) \tag{4.28}
\end{aligned}$$

Examining (4.26), (4.27) and (4.28), it is useful to define the far field spherical harmonic weighting constants³ $h_n(f)$ as:

$$h_n(f) = -\frac{\mu_0 I}{2} \sin \theta' \frac{(r')^{n+1}}{r_0^{n+2}} P_n^1(\cos \theta') \tag{4.29}$$

so

$$A_\phi(f) = -\sum_{n=1}^{\infty} \frac{1}{n(n+1)} h_n(f) \frac{r_0^{n+2}}{r^{n+1}} P_n^1(\cos \theta), \tag{4.30}$$

$$B_r(f) = \sum_{n=1}^{\infty} h_n(f) \frac{r_0^{n+2}}{r^{n+2}} P_n(\cos \theta) \text{ and} \tag{4.31}$$

$$B_\phi(f) = -\sum_{n=1}^{\infty} \frac{1}{n+1} h_n(f) \frac{r_0^{n+2}}{r^{n+2}} P_n^1(\cos \theta) \tag{4.32}$$

where the ‘f’ refers to the fact that the constants are for a filament and an arbitrary radius, r_0 , has been included to convert the units of $h_n(f)$ to Tesla. It is conventional to take $r_0 > r'_{\max}$ so that the weighting constants tend to zero for large n .

The principle of superposition means that the far-field source constants from many filaments may be added together to create an overall set of constants h_n^{total} which entirely describe the form of the magnetic fringe field. In the following two sections, expressions are developed for $h_n(c)$ and $h_n(s)$, the

³These are also known as the far-field source constants in a manner similar to [38] (see Appendix A).

weighting constants for infinitely thin cylinders and thick solenoids respectively.

Thin cylinders

This mathematical derivation is extended to an infinitely thin cylinder by integrating the result for $A_\phi(f)$ along the z' axis and replacing the current, I , in (4.29) by $J'(A/m)$, the current per unit length.

$$A_\phi(\text{thin cylinder}) = \int_{z'=\text{lower}}^{\text{upper}} A_\phi(\text{filament}) dz' \quad (4.33)$$

Taking the expression for $h_n(f)$ (4.29) and [56]

$$(u^2 - 1) \frac{d}{du} P_n^1(u) = nuP_n^1(u) - (n+1)P_{n-1}^1(u)$$

it can be shown that

$$\frac{\partial h_n(f)}{\partial z'} = \frac{n+1}{r_0} h_{n-1}(f) \quad (4.34)$$

which can be directly employed to yield an expression for $A_\phi(c)$

$$A_\phi(c) = - \sum_{n=1}^{\infty} \frac{r_0}{n(n+1)(n+2)} [h_{n+1}(f)]_{\text{limits}} \frac{r_0^{n+2}}{r^{n+1}} P_n^1(\cos \theta) \quad (4.35)$$

Notice that we only need to evaluate the weighting constants of the two filaments located at the extreme ends of the cylinder and combine them with the appropriate signs. Following the same procedure as before (4.27) and (4.28), we calculate the magnetic field components

$$B_r(c) = \sum_{n=1}^{\infty} [h_n(c)]_{\text{limits}} \frac{r_0^{n+2}}{r^{n+2}} P_n(\cos \theta) \quad (4.36)$$

$$B_\phi(c) = - \sum_{n=1}^{\infty} \frac{1}{n+1} [h_n(c)]_{\text{limits}} \frac{r_0^{n+2}}{r^{n+2}} P_n^1(\cos \theta) \quad (4.37)$$

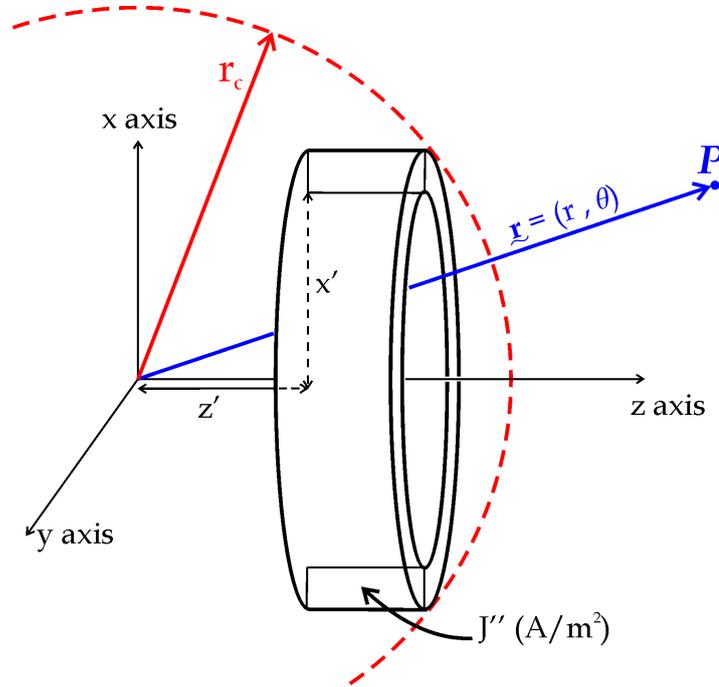


Figure 4.3: A thick solenoid with current density J'' . The observation point (\mathcal{P}) is constrained to lie outside the radius of convergence (r_c).

where the weighting constants for infinitely thin cylinders, $h_n(c)$, have been defined by

$$\begin{aligned} h_n(c) &= \frac{r_0}{(n+2)} h_{n+1}(f) \\ &= -\frac{\mu_0 J'}{2(n+2)} \sin \theta' \frac{(r')^{n+2}}{r_0^{n+2}} P_{n+1}^{-1}(\cos \theta') \end{aligned} \quad (4.38)$$

Thick solenoids

The final stage in this exposition is the integration of the magnetic vector potential, \mathbf{A} , over x' to create a thick solenoid as illustrated in Figure 4.3.

Thus

$$A_\phi(\text{thick solenoid}) = \int_{x'=\text{inner radius}}^{\text{outer radius}} A_\phi(\text{thin cylinder}) dx' \quad (4.39)$$

Examining (4.29), we replace I by the current density $J''(\text{A/m}^2)$ and write

$$A_\phi(s) = - \sum_{n=1}^{\infty} \frac{1}{n(n+1)} \frac{r_0^{n+2}}{r^{n+1}} P_n^1(\cos \theta) \int_{\text{limits}} h_n(c) dx' \quad (4.40)$$

It can be shown, following a similar argument to (4.34), that

$$\frac{\partial h_n(f)}{\partial x'} = \frac{n+1}{x'} h_n(f) - \frac{n+1}{r_0} \frac{z'}{x'} h_{n-1}(f) \quad (4.41)$$

which can be rearranged into the more useful form

$$h_{n+1}(f) = \frac{x'}{n+2} \frac{\partial h_{n+1}(f)}{\partial x'} + \frac{z'}{r_0} h_n(f) \quad (4.42)$$

At this point, the weighting constants for a thick solenoid, $h_n(s)$ are introduced and defined by:

$$h_n(s) = \int_{\text{limits}} h_n(c) dx' \quad (4.43)$$

Combining (4.38), (4.42) and (4.43), we integrate by parts to yield a recurrence relation of the form

$$h_n(s) = \left[\frac{r_0 x' h_{n+1}(f)}{(n+2)(n+3)} \right]_{\text{limits}} + \frac{(n+1)z'}{(n+3)r_0} h_{n-1}(s) \quad (4.44)$$

with a stopping condition given by

$$h_0(s) = \left[-\frac{\mu_0 J'' x'^3}{12 r_0^2} \right]_{\text{limits}} \quad (4.45)$$

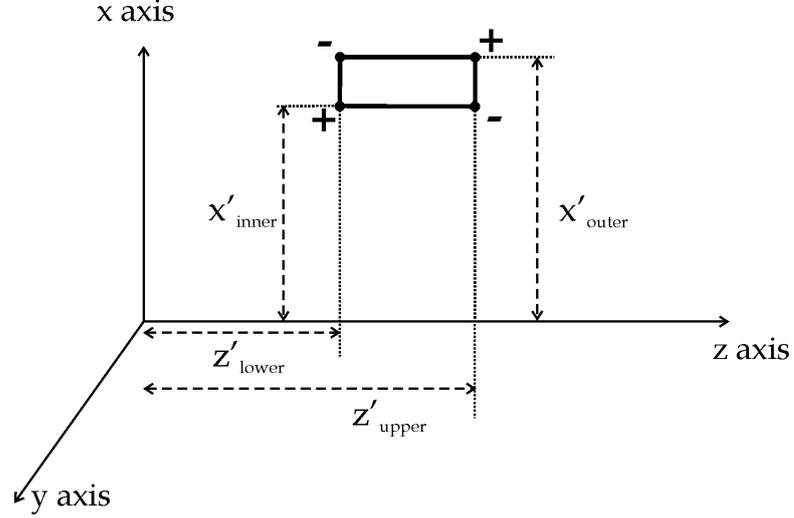


Figure 4.4: The weighting constants, $h_n(s)$, must be evaluated for each corner of a thick solenoid and combined using the correct parity. Using the abbreviations lower(l), upper(u), inner(i) and outer(o), $[h_n(s)]_{\text{limits}} = h_n(s)\{x'_i, z'_l\} - h_n(s)\{x'_o, z'_l\} + h_n(s)\{x'_o, z'_u\} - h_n(s)\{x'_i, z'_u\}$.

Finally, we take the curl of $\mathbf{A}(r, \theta)$ and recover two expressions for the magnetic field components, B_r and B_θ

$$B_r(r, \theta) = \sum_{n=1}^{\infty} [h_n(s)]_{\text{limits}} \frac{r_0^{n+2}}{r^{n+2}} P_n(\cos \theta) \quad (4.46)$$

$$B_\theta(r, \theta) = - \sum_{n=1}^{\infty} \frac{1}{n+1} [h_n(s)]_{\text{limits}} \frac{r_0^{n+2}}{r^{n+2}} P_n^1(\cos \theta) \quad (4.47)$$

where $[h_n(s)]_{\text{limits}}$ must be evaluated at the four corners of the thick solenoid and combined with the correct parity (+, -, +, -) as shown in Figure 4.4.

In a similar manner to the central field expansion, the weights of the spherical harmonics can be determined explicitly for a certain coil configuration and then used to evaluate the magnetic field at any point in the far-field region with very little additional effort. Note also that although an expression for $h_0(s)$ is given in (4.45), it is not directly used to calculate the spherical harmonic weighting, but is used instead along with the recurrence

relation to calculate the term $h_1(s)$.

The formula for A_ϕ (thick solenoid) is very satisfactory as it displays a remarkable symmetry with the central field expansion [38], where the spherical harmonic weights from thick solenoids are expressed as a sum (see Appendix A). In fact, (4.44) and (4.45) can be remodelled⁴ to give a single, final expression for $h_n(s)$

$$h_n(s) = \frac{\mu_0 J'' (x')^3 (z')^n}{2(n+2)(n+3)r_0^{n+2}} + \sum_{m=1}^n \frac{x' (z')^{n-m}}{(n+2)(n+3)r_0^{(n-m-1)}} h_{m+1}(f) \quad (4.48)$$

4.3.2 Radius of convergence

The series expansion of the magnetic field must be tested for convergence. This is achieved by first investigating the convergence of (4.31) and (4.32). Ignoring constant factors and the finite factor, $P_n(\cos \theta)$, we can employ the integral test [57]

$$\begin{aligned} \int^{\infty} a_n dn &\sim \int^{\infty} (r')^{n+1} r^{-(n+2)} dn \\ &\sim \lim_{n \rightarrow \infty} \frac{(r')^{n+1}}{r^{(n+2)} \ln \left| \frac{r'}{r} \right|} \end{aligned} \quad (4.49)$$

and we find that

$$\text{if } \begin{cases} r > r' & \frac{1}{r^{n+2}} \text{ dominates leading to convergence} \\ r = r' & \ln \left| \frac{r'}{r} \right| = 0 \text{ causing divergence} \\ r < r' & (r')^{n+1} \text{ dominates leading to divergence} \end{cases} \quad (4.50)$$

This result, applicable for infinitely thin filaments, can be generalized to thin cylinders and thick solenoids using (4.38) and (4.48). It is found that the same result applies, but now we only get convergence if r completely encompasses all cylinders and solenoids in the design. We therefore define r_c

⁴Compare with (A.10)-(A.13)

as

$$r_c = \max\{r'\} \quad (4.51)$$

and apply the far-field expansion formulae only in the region $r > r_c$.

4.3.3 Using the weighting constants

When the spherical harmonic weighting constants $h_n(f, c, s)$ were first introduced (4.29), an arbitrary constant r_0 was included in the definition. It is convenient to constrain r_0 to satisfy the inequality $r_0 > r_c$ to ensure that the higher order weighting constants tend to zero. Furthermore, it is reasonable to take r_0 as a distance at which the fringe fields are important. In the case of whole-body MRI magnets, smaller fringe fields allow easier siting and reduce health hazards. A standard specification could well be to contain the 5 gauss line within 5 m for a high field (1.0 - 3.0 T) magnet. So, setting r_0 as 5 m, it can then easily be shown that the fringe field at this distance (r_0) along the z axis is given by

$$\text{At } (z = r_0, x = 0), \text{ fringe field} = \sum_{n=1}^{\infty} h_n(s) \quad (4.52)$$

and, if the design has symmetry about the plane $z = 0$, then the fringe field at a distance r_0 along the x axis is given by

$$\text{At } (z = 0, x = r_0), \text{ fringe field} = - \sum_{n=1, \text{ odd}}^{\infty} \frac{1}{n+1} h_n(s) P_n^1(0) \quad (4.53)$$

Of course, the precise shape of the fringe field cannot be described in any detail by only these two numbers, but they are still useful and allow an understanding of the relative sizes of the fringe fields on both axes. For all other points, (4.46) and (4.47) must be used to evaluate the fringe field. In the far field region, both components of \mathbf{B} must be found and then combined to give an absolute value for the field strength, $|\mathbf{B}| = \sqrt{B_r^2 + B_\theta^2}$.

A final fact to be mentioned in this section is that some authors and mathematical software packages differ on their definition of the associated Legen-

dre polynomials. The problem involves an extra factor of $(-1)^m$ (see A.7) and allowance must be made for this in any specific implementation of the above formulae.

4.4 Design Sensitivity

When a magnet is designed, the ROI will have a theoretical uniformity which can be calculated by the usual methods. However, in practice, the coils which make up this design will not be positioned exactly according to the specifications and these positioning errors may well spoil the homogeneity of the central magnetic field and diminish the usefulness of the magnet. The sensitivity to these positioning errors (manufacturing errors) depends on the exact size and location of the coils and is termed ‘coil sensitivity’.

A theoretically sub-optimal design, which is hardly affected by manufacturing errors will generally be preferred to a better theoretical design which is overly susceptible to such positioning errors. Two techniques which approximate the sensitivity of a design are presented in this section. The first method creates an ensemble of designs, each slightly different and evaluates the homogeneity of each design within the ensemble. Unfortunately this method is very computationally intensive. The second method approaches the problem analytically and results in several novel expressions for the design sensitivity.

4.4.1 Ensemble averaging

A robust method of measuring the actual sensitivity of a magnet design is the ensemble method [58]. An ensemble of \mathcal{N} designs is formed where \mathcal{N} is large and each design is a slight variation on the original design. The size of the difference between the altered designs and the original is determined by the expected manufacturing tolerances (or margins). Both coil positioning and number of turns may be considered as having tolerances for these purposes. Each of the \mathcal{N} designs is then evaluated for homogeneity and an average of the ensemble is taken. The advantages of this method are that:

- a wide variety of tolerances can be investigated including winding the wrong number of turns and adding a layer.
- measurements other than homogeneity (e.g. fringe fields, hoop stresses) can be investigated as a function of manufacturing tolerance.

However, the ensemble method suffers from being relatively slow. To get a good estimate of the sensitivity for a fairly complex design, several thousand homogeneity evaluations are required. Therefore, this method is most suitable as a means of post-processing designs to determine their susceptibility to manufacturing errors.

By taking an analytical approach to the sensitivity of coils, a significant reduction in the time taken for evaluation may be obtained. The following section presents just such an approach.

4.4.2 Novel analytical approach to design sensitivity

The following assumptions are made:

- The number of turns in each coil is fixed along with the width and depth of each coil.
- To first order, the sensitivity can be approximated by the sensitivity to z' movements and x' movements independently.
- The homogeneity of the ROI is expressed in terms of a spherical harmonic expansion q_n (as given in Appendix A).

It is known [38], and has already been stated, that the magnetic field in the ROI, \mathbf{B} , can be expressed as

$$B_z = \sum_{n=0}^{\infty} q_n \left(\frac{r}{r_0}\right)^n P_n(\cos \theta) \quad (4.54)$$

$$B_x = \sum_{n=0}^{\infty} \frac{1}{(n+1)} q_n \left(\frac{r}{r_0}\right)^n P_n^1(\cos \theta) \quad (4.55)$$

where the constants q_n are determined by the set of equations (A.10)-(A.13). In systems where the goal is a homogeneous central field, the x component of the magnetic field can often be ignored and the magnetic field strength is approximated by $|\mathbf{B}| \approx B_z$. Using this approximation, the average field strength over the ROI can be calculated by

$$\text{av}(B_z) = \frac{1}{V} \int B_z dV = q_0 \quad (4.56)$$

where use has been made of the following relationship

$$\int_{\theta=0}^{\pi} \sin \theta P_n(\cos \theta) d\theta = \begin{cases} 2 & \text{if } n = 0 \\ 0 & \text{if } n \neq 0 \end{cases} \quad (4.57)$$

Furthermore, a measure of the homogeneity over the ROI is given by the variance of B_z

$$\text{var}(B_z) = \langle B_z^2 \rangle - \langle B_z \rangle^2 \quad (4.58)$$

where

$$\langle B_z^2 \rangle = \frac{1}{V} \int \sum_{n=0}^{\infty} \sum_{m=0}^{\infty} q_n q_m \left(\frac{r}{r_0} \right)^{n+m} P_n(\cos \theta) P_m(\cos \theta) dV \quad (4.59)$$

Equation (4.59) may be expressed in spherical coordinates and integrated using the property

$$\int_{\theta=0}^{\pi} \sin \theta P_n(\cos \theta) P_m(\cos \theta) d\theta = \begin{cases} \frac{2}{2n+1} & \text{if } n = m \\ 0 & \text{if } n \neq m \end{cases} \quad (4.60)$$

to give

$$\langle B_z^2 \rangle = \frac{1}{V} \sum_{n=0}^{\infty} \int_{r=0}^R \frac{4\pi}{2n+1} q_n^2 \left(\frac{r^{2n+2}}{r_0^{2n}} \right) dr \quad (4.61)$$

$$= \sum_{n=0}^{\infty} \frac{3 q_n^2}{(2n+1)(2n+3)} \left(\frac{R}{r_0} \right)^{2n} \quad (4.62)$$

where we have integrated over the whole of the DSV (diameter = $2R$) such that $V = \frac{4}{3}\pi R^3$. Substituting into (4.58) removes the first term from the sum and we recover a simple formula for the variance of B_z over the ROI

$$\text{var}(B_z) = \sum_{n=1}^{\infty} \frac{3 q_n^2}{(2n+1)(2n+3)} \left(\frac{R}{r_0}\right)^{2n} \quad (4.63)$$

where q_n can represent the harmonic weighting constants from filaments, thin cylinders or thick solenoids without changing the above formulae.

Sensitivity of a single filament

To find the changes in the homogeneity of the central field when a single filament is displaced by a small amount either in the z' or the x' direction, we consider the general case of displacing a single filament by a distance $\Delta \mathbf{s}$ and thus

$$\Delta B_z \approx \frac{\partial B_z}{\partial \mathbf{s}} \Delta \mathbf{s} = \sum_{n=0}^{\infty} \frac{\partial q_n}{\partial \mathbf{s}} \left(\frac{r}{r_0}\right)^n P_n(\cos \theta) \Delta \mathbf{s} \quad (4.64)$$

We calculate the average and variance of ΔB_z using the same methods as (4.56) and (4.63):

$$\text{av}(\Delta B_z) = \left(\frac{\partial q_0}{\partial \mathbf{s}}\right) \Delta \mathbf{s} \quad (4.65)$$

$$\text{var}(\Delta B_z) = \sum_{n=1}^{\infty} \frac{3}{(2n+1)(2n+3)} \left(\frac{R}{r_0}\right)^{2n} \left(\frac{\partial q_n}{\partial \mathbf{s}}\right)^2 (\Delta \mathbf{s})^2 \quad (4.66)$$

where R is the radius of the ROI.

The above formulae need to be carefully interpreted to understand their relevance to the sensitivity of a single filament. Firstly, the average ΔB_z is the change in the central field strength that occurs as the filament moves. For example if a filament moves further away from the origin, we expect the central field strength to decrease. This component of the sensitivity is termed the ‘*absolute sensitivity*’ because it alters the overall field strength without significantly changing the homogeneity of the ROI⁵. More important in the

⁵Of course for large movements, the homogeneity of the field would be affected. How-

understanding of the intrinsic coil sensitivity is the variance of ΔB_z (4.66), which measures the change in homogeneity of the central field and is termed the ‘*relative sensitivity*’. To allow comparison between different coils and designs, the following quantities are created with units of ppm per mm.

$$\text{‘absolute sensitivity’ (ppm/mm)} = \frac{\text{av}(\Delta B_z)}{\text{av}(B_z) \Delta \mathbf{s}} \times 10^3 \quad (4.67)$$

$$\text{‘relative sensitivity’ (ppm/mm)} = \frac{\sqrt{\text{var}(\Delta B_z)}}{\text{av}(B_z) \Delta \mathbf{s}} \times 10^3 \quad (4.68)$$

Sensitivity of several filaments

The above formulae (4.67) and (4.68) deal with a design containing a single filament. However, most designs tend to contain many filaments and the question of how to measure their overall sensitivity is addressed here.

The method proposed is to evaluate the sensitivity of each filament individually and then to add them in quadrature. Formalizing this, we take a design containing M filaments and denote a particular filament by m , such that $1 \leq m \leq M$. Some filaments in the design will be more susceptible to positioning than others, so we calculate for each filament (note the superfix)

$$\text{abs. sens.}^{(m)} = \frac{1}{\text{av}(B_z)} \left(\frac{\partial q_0^{(m)}}{\partial \mathbf{s}} \right) \times 10^3 \quad (4.69)$$

$$\text{rel. sens.}^{(m)} = \frac{1}{\text{av}(B_z)} \sqrt{\sum_{n=1}^{\infty} \frac{3}{(2n+1)(2n+3)} \left(\frac{\partial q_n^{(m)}}{\partial \mathbf{s}} \right)^2 \left(\frac{R}{r_0} \right)^{2n}} \times 10^3 \quad (4.70)$$

where the source constants from an individual filament, m , are denoted by $q_n^{(m)}$ and the overall magnetic field is determined by the superposition of the

ever, in this section we are talking only about small coil movements due to manufacturing errors.

filaments. So

$$q_n^{(\text{overall design})} = \sum_{m=1}^M q_n^{(m)} \quad (4.71)$$

As mentioned above, in most applications, we can ignore the absolute sensitivity as being an insignificant percentage of the overall field strength and concentrate only on the relative sensitivity which determines the change in homogeneity of the magnetic field. We combine the relative sensitivities of the M coils in quadrature to give a single figure measure of robustness to movements in the direction represented by the vector $\Delta \mathbf{s}$

$$\text{relative sensitivity along } \mathbf{s} = \sqrt{\sum_{m=1}^M [\text{rel. sens.}^{(m)}]^2} \quad (4.72)$$

Up until this point the sensitivity of filaments with respect to movements in a general direction, $\Delta \mathbf{s}$, has been discussed. Exactly the same treatment can be applied to find the sensitivity of the filaments in the z' and x' directions, by replacing $\Delta \mathbf{s}$ and $\partial q_n / \partial \mathbf{s}$ by, for example, $\Delta z'$ and $\partial q_n / \partial z'$.

The advantage of considering movement along the z' and x' axes in particular is two-fold. Firstly, the manufacturing process determines that placement in these axes is controlled by different parameters, z' being controlled by the accuracy of the slots cut into the magnet former, whereas accuracy in x' is determined by the ability to lathe accurately at a constant radius. Secondly, the partial differentials of q_n can be found relatively easily in these directions.

For source constants from infinitely thin filaments, $q_n(f)$, defined as in (A.6), we can determine that

$$\frac{\partial q_n(f)}{\partial z'} = -\frac{n+1}{r_0} q_{n+1}(f) \quad (4.73)$$

and with some tricky manipulation of the Legendre polynomials⁶, we find

⁶See Appendix B for details.

the differential with respect to x'

$$\frac{\partial q_n(f)}{\partial x'} = -\frac{(n+1)(n+2)}{(2n+5)} \left[\frac{(r')^2}{x'r_0^2} q_{n+2}(f) - \frac{1}{x'} q_n(f) \right] \quad (4.74)$$

Equations (4.73) and (4.74) can be directly substituted into (4.69) and (4.70) to give a measure of the relative sensitivity of the design to movements in the z' and x' directions separately. Notice that since the derivatives of $q_n(f)$ can be evaluated as multiples of $q_n(f)$, this method is extremely fast, especially if the code is written in such a way as to store the weighting coefficients as an array in memory to begin with.

We then combine the contributions from all M filaments in both the z' and x' directions to find an overall single figure of merit:

$$\text{Absolute relative sensitivity} = \sqrt{\sum_{m=1}^M \left(\left[\text{abs. sens.}_{z'}^{(m)} \right]^2 + \left[\text{abs. sens.}_{x'}^{(m)} \right]^2 \right)} \quad (4.75)$$

$$\text{Overall relative sensitivity} = \sqrt{\sum_{m=1}^M \left(\left[\text{rel. sens.}_{z'}^{(m)} \right]^2 + \left[\text{rel. sens.}_{x'}^{(m)} \right]^2 \right)} \quad (4.76)$$

Equation (4.76) can now be used as a single figure measure of the robustness of a design and is an extremely powerful tool of comparing different designs objectively.

Although the derivation has concentrated on designs containing infinitely thin filaments, it is clear from the mathematics that it is a simple case of substituting $q_n(c)$ or $q_n(s)$ in place of $q_n(f)$ to achieve a similar result for designs containing thin cylinders and thick solenoids. In the next section, results for the differentials of these source constants are presented.

Sensitivity of thin cylinders and thick solenoids

To determine the absolute and relative sensitivities of thin cylinders and thick solenoids, it is necessary to determine the differentials of $q_n(c)$ and $q_n(s)$ with respect to both z' and x' . Examining (A.8-A.9) and (A.10-A.12) we can develop the following relationships

$$\frac{\partial q_n}{\partial z'}(c) = -\frac{(n+1)}{r_0} q_{n+1}(c) \quad (4.77)$$

$$\frac{\partial q_n}{\partial x'}(c) = \frac{(n+1)}{(2n+3)} \left[\frac{(n+2)(r')^2}{x' r_0^2} q_{n+2}(c) - \frac{n}{x'} q_n(c) \right] \quad (4.78)$$

evaluated at both ends of cylinder $[z = z_{\text{lower}}] - [z = z_{\text{upper}}]$

$$\frac{\partial q_n}{\partial z'}(s) = -\frac{(n+1)}{r_0} q_{n+1}(s) \quad (4.79)$$

$$\frac{\partial q_n}{\partial x'}(s) = \begin{cases} \frac{\mu_0 J''}{2} \cos \theta' & \text{for } n = 0 \\ \frac{\mu_0 J''}{2n} \left(\frac{r_0}{r'}\right)^n \sin^2 \theta' P_n^1(\cos \theta') & \text{for } n \geq 1 \end{cases} \quad (4.80)$$

evaluated at all four corners (see Figure 4.4)

Notice the symmetry which exists when differentiating $q_n(f, c, s)$ with respect to z' . Remarkably, we recover exactly the same formulae for filaments, thin cylinders and thick solenoids, which makes the calculation of the sensitivity much faster than would otherwise be possible. As usual, (4.77)-(4.78) need to be evaluated at the ends of each thin cylinder and (4.79)-(4.80) need to be evaluated at the four corners of each thick solenoid before combining the results with the appropriate signs to calculate an overall figure for $\frac{\partial q_n(c)}{\partial z'}$, $\frac{\partial q_n(c)}{\partial x'}$, $\frac{\partial q_n(s)}{\partial z'}$ and $\frac{\partial q_n(s)}{\partial x'}$.

Finally, we substitute (4.77)-(4.80) into the formulae for absolute and relative sensitivity and then combine the contributions from many individual cylinders or solenoids to obtain the ‘**overall relative sensitivity**’ (4.76).

Nth order sensitivity

Up to this point, we have been attempting to find a single ‘figure of merit’ for the sensitivity. In fact, the inhomogeneity induced over the central field can be described in much more detail as a sum of spherical harmonic terms. This leads us to consider the fact that the inhomogeneities created during manufacture will eventually be shimmed with active or passive shimming methods which in general correct for each of the terms in the spherical harmonic expansion. Therefore, rather than combining the terms in (4.70), we maintain them as individual harmonics and find the relative sensitivity of the n^{th} order. These figures can then be compared to the various strengths of the shimming coils and a decision made about acceptable margins during fabrication. Thus

$$\text{‘}n^{\text{th}} \text{ order relative sensitivity’ (ppms/mm)} \approx \frac{1}{\text{av}(B_z)} \sqrt{\frac{3(n+1)^2}{(2n+1)(2n+3)}} (q_{n+1}) \frac{R^n}{r_0^{n+1}} \times 10^3 \quad (4.81)$$

These values are useful for detailed analysis of the magnet’s response to build errors, however, ‘**overall relative sensitivity**’ (4.76) provides a useful single figure of merit, which allows straightforward comparison between different designs. All of the presented formulae are straightforward and rapid to implement.

4.5 Forces, Stresses and Strains

Whenever a current flows in a magnetic field, the conductor experiences a force, perpendicular to both the direction of the current (\mathbf{I}) and the direction of the magnetic field (\mathbf{B}), known as the Lorentz force. The size and direction of the Lorentz force (\mathbf{F}) is given, per unit length, by [41]

$$\mathbf{F} = \mathbf{B} \times \mathbf{I} \quad (4.82)$$

where the magnetic field is the combination of self-field and the background field from the surrounding coils.

In many different types of magnet design, these forces, exerted on the conductors can be very large. Unless care is taken with the design and construction of such magnets, the forces can damage the superconductor and the insulation, induce premature quenching, or even lead to catastrophic failure of the entire structure [59]. The following sections describe methods of calculating the forces and resulting stresses.

4.5.1 Forces and stresses in solenoids

Taking a slice perpendicular to the axis of symmetry, the magnetic field lines can be considered as a gas which exert a pressure of $B^2/2\mu_0$ against the cylindrical coil [60]. Indeed this is exactly how Faraday visualized the intrinsic repulsion of magnetic flux lines. The pressure exerted by this ‘gas’ creates a radial force on the current carrying coils, but there are also axial forces acting on the windings due to the curvature of the magnetic field lines at the extremities of the windings.

A cross-section through a solitary axisymmetric coil is shown in Figure 4.5 with the pattern of forces shown on the right and lines of magnetic flux shown on the left. Notice how the axial forces act towards the middle of the coil, thereby compressing the windings and consolidating the structure. Furthermore, near the outer edge of the coil, the direction of the field reverses and the forces act to partially contain the winding with an inward radial force.

Axial forces in solenoidal magnets tend to pose few problems due to their symmetry and consolidating nature. However, in unusual geometries such as split-coil magnets or ‘C’ shaped vertical axis magnets, these forces can reach several tonnes, leading to extreme difficulties in maintaining the correct separation of the pole pieces. More often though, it is the radial forces which act as the limiting factor, leading to very high tensions being exerted on the superconducting matrix, known as the *hoop stress*. It is therefore essential that the hoop stresses are considered during the design of magnets, especially for large bore MRI designs.

To evaluate the hoop stress, a reasonable approximation is to assume

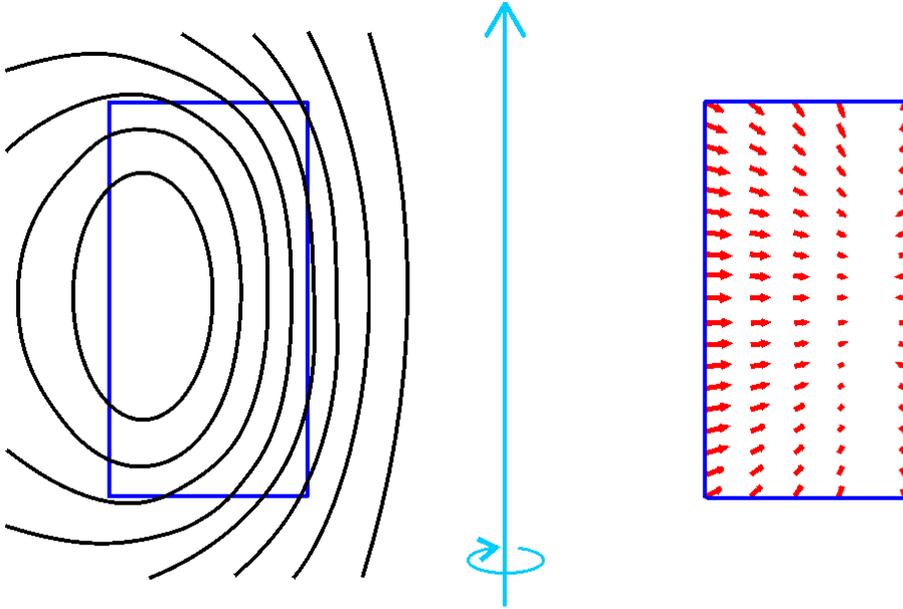


Figure 4.5: Pattern of forces and magnetic flux lines within a solitary axisymmetric coil. Notice how the net force over the cross-section is in the positive radial direction and that the axial forces are acting to consolidate the coil.

that each turn acts independently of its neighbours [60, p.42]. Thus, a single turn, with radius r , develops a tension given by $|T| = B_z I r$. Dividing by the cross-section of the turn, we obtain a value for the unsupported uniaxial hoop stress, σ_{hoop} .

$$\sigma_{\text{hoop}} = B_z J'' r \quad (4.83)$$

where J'' is the current density in A/m^2 . The above formula, whilst correct for a coil composed entirely of superconductor, must be altered slightly in practice, where the superconducting wire is surrounded by insulation and embedded in epoxy. These materials are much softer than the copper/superconducting wire matrix and consequently do not contribute to the stiffness of the coil. To account for this factor, an engineering current density J_E is referred to [52], such that

$$J_E = \frac{1}{1 - \epsilon} J'' \quad (4.84)$$

where ϵ is the fraction of cross-sectional area containing epoxy and insulation. Thus (4.83) becomes

$$\sigma_{\text{hoop}} = B_z J_E r \quad (4.85)$$

Shear stresses and strains are harder to approximate, but are usually less of a constraint compared to the uniaxial hoop stresses. However it must be born in mind that, in general, materials are much more susceptible to shear stresses than uniaxial stresses, especially when the uniaxial stress is compressive.

4.5.2 Material properties

To obtain a better indication of the stresses and strains within a particular magnet design, it is necessary to take into account the actual material properties of the coil windings and eventually, the coil formers. All materials are affected by both uniaxial and shear stresses which produce elastic deformation, followed, near the yield stress by eventual fracture. Stresses can also alter the superconducting properties of the materials leading to premature quenching. In order to study them more accurately, a finite element model must be created which should include the superconducting wire, stabilizing matrix, insulation and epoxy. Inclusion of the formers in the finite element model is also necessary for a final evaluation of the total deformation of the magnet structure, however, this is beyond the scope of this thesis.

Deformation is governed by up to 21 independent constants for general materials. If we can assume isotropic conditions, then just two constants, Young's modulus of elasticity and Poisson's ratio, are needed. Young's modulus is given by [61]

$$Y = \frac{\sigma_{xx}}{\epsilon_{xx}} \quad (4.86)$$

and Poisson's ratio by [61]

$$\nu = -\frac{\epsilon_{yy}}{\epsilon_{xx}} \quad (4.87)$$

where σ and ϵ are the 2nd rank stress and strain tensors such that

$$\sigma_{ij} = \frac{F_i}{\mathbf{A}_j} \quad (4.88)$$

and

$$\epsilon_{ij} = \frac{\partial u_i}{\partial x_j} \quad (4.89)$$

Often a volume-weighted average can be taken of the individual parts of the winding to find an approximate value of Young's modulus and Poisson's ratio, which can be used in a finite element analysis. For even more accuracy, a finite element model must be created, which includes the superconducting wire, stabilizing matrix, insulation material and epoxy individually. From the finite element model, all necessary stresses and strains may be found, including the hoop stress, σ_{33} and particularly, the von Mises stress which is the universally accepted measure of stress in industry given by [62]

$$\sigma_\nu = \sqrt{\frac{(\sigma_{11} - \sigma_{22})^2 + (\sigma_{22} - \sigma_{33})^2 + (\sigma_{11} - \sigma_{33})^2}{2}} \quad (4.90)$$

4.5.3 Finite element analysis with ABAQUS

A detailed finite element analysis of the stresses, shears and strains incurred in an MRI magnet design is beyond the scope of this thesis. Indeed, even commercial MRI manufacturers often subcontract the finite element analysis to a specialist company. However, using a state of the art finite element package, ABAQUS(c) [63], a two-stage procedure was created to perform preliminary finite element analysis of candidate designs after optimization.

1. The entire design is modelled, approximating each coil by a solid, thick winding with volume-averaged Young's modulus and Poisson's ratio. The coils are restricted in the axial direction on the relevant side by a frictionless, infinitely strong former. This allows radial expansion/contraction to be visualized (see Figure 4.6). Shear stresses are

also evaluated, but may not be as accurate due to the interaction with an infinitely strong former.

2. Individual coils which are found to have a high hoop stress are subsequently modelled by taking each layer of winding as a laminated combination of a copper/superconductor matrix and epoxy. The mesh used in this model is much finer and each coil is examined individually (see Figure 4.7). As in the previous analysis, the coils are only contained axially by frictionless, infinitely strong formers.

By accurately modelling the stresses present in a particular magnet design, informed decisions can be made about which designs are feasible and which are unfeasible.

4.6 Quenching, Stored Energy and Inductance

This section gives an overview rather than specific details of the problem of magnet underperformance due to quenching. The emphasis is placed on the physical indicators that can be controlled during the design process and their impact on the phenomenon of quenching. This is an important part of magnet design and factors such as the stored energy and the inductance of a magnet must be carefully considered especially with larger scale superconducting magnets like those used in whole-body MRI.

4.6.1 Introduction to quenching

In normal operation, the entire volume of a magnet coil is maintained at a temperature below the critical temperature, $T_c(B)$, such that the superconducting properties of the wire can be exploited. However, if due to some instability⁷, part of the wire heats up, this region can be driven into its normal state and the resultant Joule heating may well lead to runaway and a

⁷See Seeber [52, p. 528] for a brief explanation of causes.

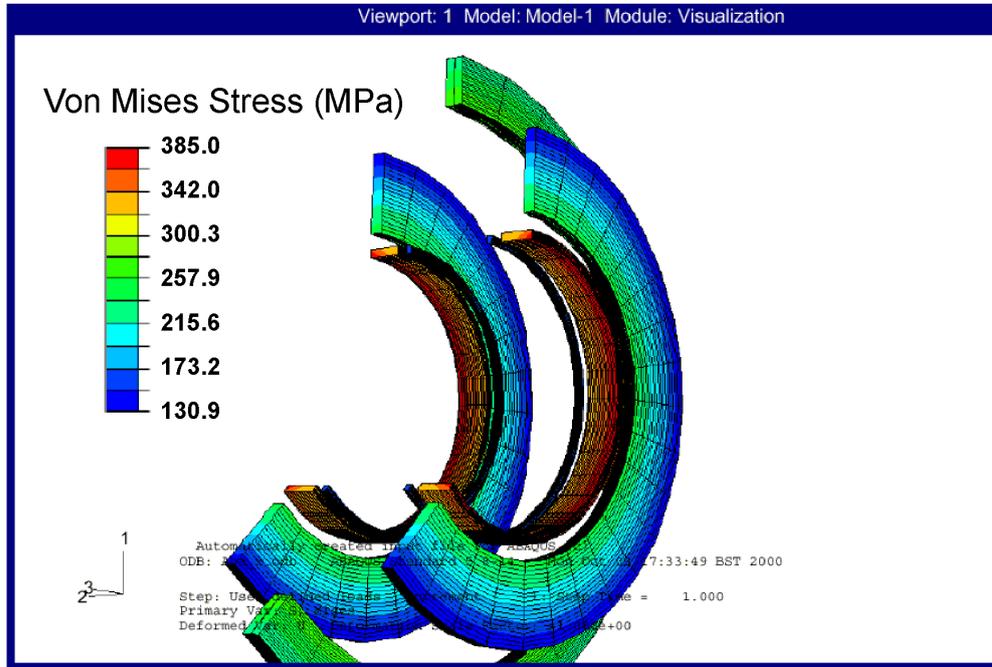


Figure 4.6: This figure shows the finite element calculated von Mises stress for a short, whole-body MRI magnet design. Maximum stress is roughly 385 MPa.

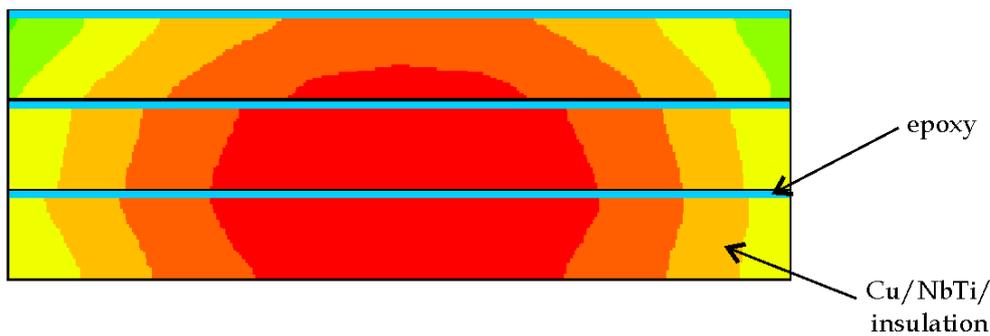


Figure 4.7: Detailed finite element analysis of a laminated coil with three 1 mm layers of epoxy and copper. The maximum von Mises stress can be seen near the inner edge shown in red.

quench. Only if the heat conduction away from the initiation point is greater than the Joule heating can a quench be prevented. Under balanced conditions we can match the heat conduction to the Joule heating in the following formula

$$\frac{2kA(T_c - T_0)}{l} = J^2 \rho A l \quad (4.91)$$

where k is the thermal conductivity, l is the length of the normal zone and $\frac{T_c - T_0}{l}$ is the approximate temperature gradient between the normal and the superconducting zones. A is cross-sectional area of the wire, J is the current density and ρ is the normal state resistivity. This rough calculation shows that a normal zone greater than approximately $0.5 \mu\text{m}$ will provoke a quench [60, p.74]. Once a quench has been initiated, heat spreads by a combination of thermal conduction and Joule heating.

Historically, the causes of quenching have been very hard to pin down. Flux jumping certainly played a role in the early days, however with modern multicore wire, this problem has been virtually eliminated. Probably the most important factor in magnet degradation⁸ is due to mechanical disturbances. As the magnet is energized, the stress endured by the coil increases and can cause the wire to shift, releasing energy. The amount of energy needed to drive a section of wire into its normal state is very small, typically of the order of 10^{-9} J compared to a total energy in the region of 10^6 J. Evidence for this comes from the recording of the sounds emitted as the magnet ramps up to its operational current. The frequency of cracking increases as does the intensity until the point where quenching occurs. In general, coils which are subjected to large forces between themselves and the formers tend to suffer greater degradation.

Once the normal zone region has started to grow there is no alternative but to allow the coil to quench completely. In a few seconds, all of the energy stored within the magnetic field must be dissipated, without permanent damage to the superconducting windings. For a magnet with 10 MJ of internal energy which quenches in 5 seconds, this is equivalent to coping with the en-

⁸The effect that large magnets do not obtain the same critical currents as would be expected from experiments on short samples of superconducting wire is known as degradation.

ergy output from 2000 single bar heaters! Three factors need consideration, each of which can lead to permanent magnet damage.

Firstly, Joule heating from the now resistive conductor leads to the formation of a hot spot, usually at the point of initiation. The quench must be controlled in such a way that the peak temperature does not exceed a certain critical limit, T_{max} , and that differential thermal expansion does not have a significant impact on the structural stability. Secondly, although the voltage across the coil terminals may only be a few volts, turn-to-turn voltages can reach several thousand volts (kV) resulting in shorting of the insulation and a reduction in the magnet performance. Thirdly, as parts of the magnet tend to quench before others, asymmetric forces exceeding the design limits can be produced. Obviously the magnet designers must take all of these effects into account before building such a magnet to avoid permanent damage. From an economic standpoint as well, quenches are very expensive both in the time taken to re-cool the magnet and in the consumption of liquid Helium.

4.6.2 Quench modelling

The techniques developed to model normal zone propagation come in two basic flavours. Some attempt to solve the problem analytically, often including many approximations, whereas others use a finite element technique and solve the resulting coupled differential equations numerically. The main aims of the modelling are to find the peak turn-to-turn voltage⁹, the peak temperature¹⁰ and the velocity of quench propagation¹¹. The details of these methods are not discussed in any depth here as many good references exist on the subject; see [52] and [60] for the analytical approaches and [64, 65, 66] for numerical techniques.

⁹In general, the turn-to-turn voltage is much greater than the voltage measured across the magnet terminals.

¹⁰Peak temperature is usually assumed to occur at the point of quench initiation.

¹¹The propagation velocity depends on the thermal conductivity of the wire. In general, the longitudinal conductivity is much greater than the transverse conductivity due to the intervening layer of insulation, creating a quasi-1D model.

4.6.3 Quench prevention and protection

Multi-filament, superconducting wire in a copper matrix is now almost universally used in large magnet projects. The fine filaments reduce AC losses and prevent flux-jumping which can release enough energy to initiate a quench on its own. Better thermal contact between the filaments and the copper matrix also mean that heat transfer to the entire cold mass is improved.

Combined with the type of wire used, the method of winding is also crucial. For high performance magnets, it is necessary to dry wind the coils and follow this by a carefully studied process of vacuum epoxy impregnation. The exact epoxy properties are very important and must include an ability to flow rather than crack at low temperatures combined with inherent strength. Glass fibre reinforcement is also used to improve insulation along with a more surprising quality of allowing long term slippage. Too much insulation and reinforcement however reduces the maximum current density that can be obtained in a coil.

Apart from the manufacturing considerations for quench prevention, there are design criteria that can be examined. These are known as the safety margins, which are expressed either as a current margin or as a temperature margin¹². The maximum current density is always limited by the peak field strength within windings. This is substantially greater than the central field strength in wide-bore magnets such as those used in whole-body MRI. Figure 4.8 shows some of the typical current density limits for NbTi and Nb₃Sn at 4.2 K and 2 K. However, as well as current fluctuations, there may also be thermal fluctuations. Safety margins for thermal fluctuations can be expressed either as a temperature difference, ΔT , or as an enthalpy margin.

If all of the preventive measures fail and a quench occurs, it is crucial that a magnet is protected to cope with the ensuing conditions. Specifically there are two main threats. One is the peak temperature throughout the

¹²Copper, the main component of the wire has a specific heat of $0.1 \text{ J kg}^{-1} \text{ K}^{-1}$ at 4.2 K and $0.5 \text{ J kg}^{-1} \text{ K}^{-1}$ at 8 K.

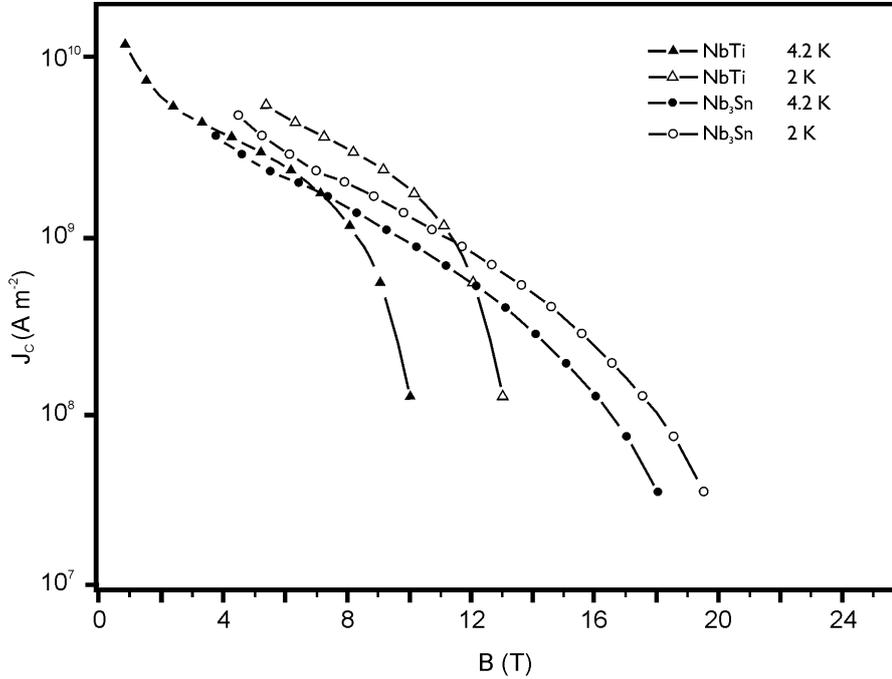


Figure 4.8: Critical current density vs B characteristics of NbTi and Nb₃Sn multifilamentary superconductors (reproduced from Seeber [52, p. 400]).

magnet and the other is the peak voltage. The copper matrix is the primary means of reducing the Joule heating effects and increasing the thermal mass of the magnet system thereby stabilizing normal zone fluctuations. Both Nb₃Sn and NbTi have normal phase resistances that are far too high for safe dissipation of the currents so instead, the current bypasses into the copper matrix ($\rho_{\text{copper}} \approx 3 \times 10^{-10} \Omega\text{m}$, $\rho_{\text{NbTi}} \approx 6.5 \times 10^{-7} \Omega\text{m}$ at 4.2 K [52]). External resistors may then be used to extract part of the stored energy, whilst the conserved energy will cause the magnet to heat up. During the design process, a limit is placed on the peak temperature. Sometimes this is taken as room temperature, but if differential thermal contraction is an issue then 100 K is a more appropriate target [52, 60]. Similarly, the insulation surrounding each turn must be capable of surviving several kV. The supply voltage is usually only a few volts, but this must be turned off as soon as

possible after a quench is detected.

Devices to protect superconducting magnets can be either passive or active. Passive methods include external dump resistors or coil subdivision, whilst active systems include heaters to spread the quench more evenly across the magnet and the activation of a cold diode [67].

4.6.4 Training

If a magnet quenches well below its design current, there is a chance that by repeated quenching, the maximum current in the coils can be increased significantly. This process is known as training. Training allows stressed regions to relax to a lower energy configuration thus removing possible locations for quench initiation and consequently improving magnet performance. However, training is an unreliable and expensive procedure. Some magnets will improve dramatically during the first few training quenches, whilst others show little improvement. In general, all magnets reach a plateau where further quenches do not increase performance. Further problems can be encountered if the magnet is thermally cycled to room temperature. This can have a detrimental effect on the operational current limit and require a subsequent retraining period.

4.6.5 Stored energy and inductance

The self-inductance of a solenoid, C_i , is given by [41]

$$L_i = \frac{\mu_0}{4\pi I_i^2} \int_{C_i} d^3r_i \int_{C_i} d^3r'_i \frac{\mathbf{J}(\mathbf{r}_i) \cdot \mathbf{J}(\mathbf{r}'_i)}{|\mathbf{r}_i - \mathbf{r}'_i|} \quad (4.92)$$

and the mutual inductance of solenoids, C_i and C_j , by [41]

$$M_{ij} = \frac{\mu_0}{4\pi I_i I_j} \int_{C_i} d^3r_i \int_{C_j} d^3r'_j \frac{\mathbf{J}(\mathbf{r}_i) \cdot \mathbf{J}(\mathbf{r}'_j)}{|\mathbf{r}_i - \mathbf{r}'_j|} \quad (4.93)$$

For general coaxial thick solenoids, complications arise as there is only partial flux transfer between the two coils. As Grover says [68]:

For the most part, calculations are not simple and in some cases the formulae are very complicated.

For this reason, many approximations are usually made and common techniques for calculating the inductances involve tables and approximations such as those found in [68, 69, 70]. However, self and mutual inductances may also be evaluated by numerical methods using a knowledge of the field strength within the magnet design and the flux linkage [71].

Once the self and mutual inductances have been found, the stored energy of a magnet can be calculated by an extension of the formula for two coils carrying currents I_1 and I_2

$$E = \frac{1}{2}L_1I_1^2 + M_{12}I_1I_2 + \frac{1}{2}L_2I_2^2 \quad (4.94)$$

or by an integral of the magnetic field strength over all space

$$E = \int \frac{B^2}{2\mu_0} dV. \quad (4.95)$$

4.7 Physical Limitations

Finally in this chapter, the overall size of the magnet is considered as well as limitations such as the maximum allowable current. In most cases, the overall size of the magnet is determined at the start of the design procedure, so such limitations can be dealt with fairly simply. Due to the complexity of the engineering involved, the increased forces and the increased volume of materials required, the cost of a project increases rapidly for large designs.

The following quantities form part of the design specification: magnet length, outer diameter, inner bore diameter, mid-plane separation (if necessary) and maximum current supply. Two supplementary quantities may well be used to compare designs.

The volume of superconductor V_{sc} is given by

$$V_{sc} = (1 - \epsilon) \sum_{\text{coils}} \pi w d (d + 2r_{in}) \quad (4.96)$$

where ϵ is the fraction of cross-sectional area containing insulation and epoxy, w is the width of a coil, d is the depth of the coil, r_{in} is the inner radius of the coil and the sum is over all coils in the design.

Finally, the number of amp-turns for magnet designs may be calculated as follows

$$\text{Amp.turns} = \sum_{\text{coils}} \frac{w d I}{A} \quad (4.97)$$

where A is the wire cross-sectional area and I is the current flowing in a single strand of wire.

Chapter 5

Genetic Algorithms for Optimizing Short Whole-Body MRI Magnets

5.1 Introduction

Magnetic resonance imaging (MRI) has heralded a revolution in the *in vivo* understanding of soft tissue structure and is becoming one of the most important modalities within a clinician's toolkit. Throughout the past twenty-five years, the hardware for MRI has continually improved, leading to faster imaging at higher field for less cost. However the goal to create an ultra-short main field magnet, comparable in length to the currently available computed tomography (CT) scanners, remains elusive. The need for ultra-short MRI magnets though is well documented [37] and is driven by two major factors. Firstly, many patients feel uncomfortably confined inside the magnet bore especially during some of the longer scans up to an hour and a half. This claustrophobia seriously affects¹ about 14% of adult patients [72] and it is generally accepted that this percentage increases in children, especially as noisy gradient sets can accentuate the problem. Images of patients affected

¹14.3% of adults required some form of sedation: oral, intra-venous or general anesthesia.

by claustrophobia often include motion artifacts which can in turn impact on the diagnostic quality of the images [73]. In many such cases, the patients are obliged to undergo subsequent examinations with other imaging modalities which are less suitable for their condition and potentially risky. A second reason for ultra-short main field magnets is that clinicians are increasingly performing operations under MRI guidance. The advantages of such procedures include very good spatial resolution, real-time imaging and non-ionizing radiation, making it relatively safe for both surgeon and patient. However, conventional length MRI scanners mean that the patient must be moved into the bore of the magnet for each scan during the operation and then removed before the operation can continue². Another solution is to split the magnet into two halves. Such magnets include the GE Signa SP/i (double doughnut) system³ [75], which operates at 0.5 T and is used for adult and paediatric interventional work⁴.

This chapter details the optimization of a suitable ultra-short MRI main field magnet using GAs.

5.2 Specifications and Requirements

Since the early 1980s, MRI magnets have undergone a revolution in their design, becoming smaller, cheaper and easier to site. However, there are still problems with claustrophobic patients and accessibility to critically ill patients. One method to address these problems is to further reduce the length of the main field magnet. Traditional design techniques, however, have been poor at reducing whole-body magnet lengths much below 1.5 m, so novel design methodologies need to be investigated. Short magnet design is known to be challenging [37] and especially so when the aim is to provide the shorter magnet for a similar cost to existing systems [76, 77]. Currently, magnets have lengths ranging from 1.5-2.5 m and room temperature bore

²More recently, a system has been developed where the patient remains stationary at all times and the *magnet* is moved into position for imaging [74].

³<http://www.gemedicalsystems.com/rad/mri/products/spi/index.html>

⁴For example, the Interventional Magnetic Resonance unit based in St. Mary's Hospital, London.

diameters in the region of 80-90 cm.

Preliminary work [78, 79] to optimize short bore MRI magnets using a GA illustrated some of the problems faced when reducing the overall length. The problems can be categorized into 3 key areas.

- *Stress limits:* As the magnet length is shortened, the coils need to produce magnetic field more efficiently, leading to strong fields in coils with high currents. The Lorentz force equation, $\mathbf{F} = \mathbf{B} \times \mathbf{I}$, shows that the force depends on the magnetic field strength and the current. Hence, the hoop stress is proportional to the square of the current.
- *Critical current and critical field limits:* All high field MRI magnets rely on the operation of superconducting NbTi or Nb₃Sn wire, which loses its superconducting properties when exposed to a combination of strong magnetic fields and high current densities. As the magnet length is reduced, the magnetic flux density is further concentrated exacerbating the problems.
- *Region of Interest (ROI) homogeneity limits:* Homogeneity of the magnetic field over the ROI becomes increasingly difficult to maintain as the magnet length decreases. A corollary is that the size of the useful imaging region usually decreases.

With the above limitations in mind, the following specifications were decided upon:

- 1.0 m long MRI magnet (coils only) with 1.0 m wide bore access.
- 1.0 T field strength⁵.
- Outer diameter less than 2.5 m (allowing space for shielding coils).
- Homogeneous region of at least a 30 cm diameter spherical volume (DSV).

⁵A 1.5 T field was avoided due to problems with hoop stresses at this higher field strength.

- Inhomogeneity over 30 cm DSV of less than 1 ppm max. to min. or roughly 1 ppm standard deviation over a 40 cm DSV.
- Necessary critical current and critical field constraints satisfied.
- Design robust enough to still achieve good homogeneity when subjected to manufacturing tolerances.

5.3 Methods I: Implementation of Genetic Algorithm

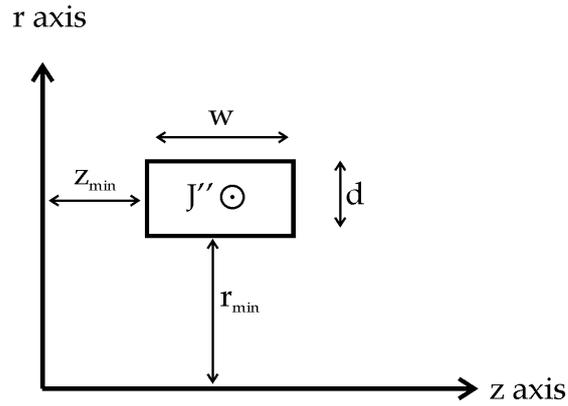
5.3.1 Encoding the parameters

A real-coded genetic algorithm (see Section 3.2.1) was written in C to optimize the positions of pairs of coils within an axisymmetric MRI superconducting magnet design. Each coil was uniquely specified by five parameters: z_{\min} , r_{\min} , w (width), d (depth) and J'' , the current density (see Figure 5.1). Width and depth were quantized throughout the optimization process to the wire size, which was generally taken as $1 \text{ mm} \times 1 \text{ mm}$. These parameters also specified a symmetrical partner, reflected in the plane $z = 0$, where necessary and formed a basic building block for the GA optimization. Details of the parameters and the formation of a chromosome are illustrated in Figure 5.1. Using this representation, a single chromosome contains all of the necessary parameters that define a single magnet design.

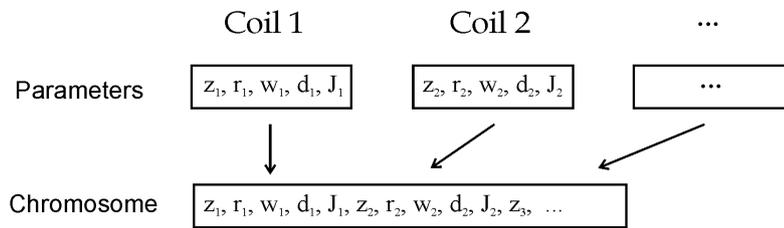
5.3.2 Details of the GA

The GA was configured so that each member of the population represented a single potential MRI magnet design. A population size of 128 was chosen after testing, since larger population sizes were found to slow down the GA, whilst smaller population sizes led to rapid convergence and stagnation of the GA.

To initialize the population, designs were created by assigning a random value to each parameter. Each initialized design was checked and rejected



(a)



(b)

Figure 5.1: Illustration of (a) the definitions of the parameters used in axisymmetric coil optimization (using cylindrical coordinates) and (b) the encoding of the parameters to form a chromosome.

if the design failed to comply with all of the constraints. Specifically, the constraints ensured that all of the coils were located within the boundaries of the maximum MRI magnet size and that none of the coils overlapped. To check for overlapping coils, the following condition was used:

$$\begin{aligned} z_{\min}(1) &< z_{\max}(2) \text{ and} \\ z_{\min}(2) &< z_{\max}(1) \text{ and} \\ r_{\min}(1) &< r_{\max}(2) \text{ and} \\ r_{\min}(2) &< r_{\max}(1) \end{aligned} \tag{5.1}$$

where $z_{\max} = z_{\min} + w$, $r_{\max} = r_{\min} + d$ and the numbers in brackets relate to the two coils being compared.

Selection was performed using a binary tournament method (see Section 3.2.4) and selected individuals were copied into a group, known as the parent pool. Members were then chosen randomly from the parent pool and recombination was performed with a probability (p_m) of 0.9 to create a single offspring. Crossover was permitted at a single-point between parameters and parents were replaced in the parent pool after recombination.

Random mutations were applied to the offspring by the addition of a Gaussian-distributed random variable to the parameters. Three possible mutations were performed, a *spread* mutation, a *jump* mutation and a *kick* mutation. For a *spread* mutation, the standard deviation of the Gaussian-distributed random variable was fairly small (of the order 1 mm). *Jump* mutations had a larger standard deviation (1-2 cm) and the *kick* mutations re-initialized the parameter to a random value over the entire allowable range. Each type of mutation was controlled by two variables, the probability that an individual would be subjected to this type of mutation, and the probability that once chosen, a particular parameter of that individual would be selected for mutation. Designs were checked after both recombination and after mutation to see if they remained within the boundaries and no coils overlapped. If the constraints were broken, then recombination or mutation were performed again until a feasible solution was created.

Finally, elitism (see Section 3.4) was implemented to ensure that the best

genetic information was not lost between generations.

5.3.3 Parallel computation and MPI

A *de facto* industry standard parallel communication protocol, Local Area Multicomputer (LAM) / Message Passing Interface (MPI) [80, 81, 82] was used for communication between processors (or nodes) on both cluster computers and supercomputers. MPI offers a portable and easily-implemented interface in both homogeneous and heterogeneous parallel environments, whilst maintaining efficiency. For simplicity, a single instruction, multiple data (SIMD) format was implemented. This involves the simultaneous execution of the same program on all processors and relies on function calls to a MPI library written in C.

The parallel GA code was tested on a small scale Linux cluster and then ported to a Hitachi SR2201 supercomputer, part of the High Performance Computing Facility at the University of Cambridge [83], for the production runs. Details of the two parallel computer architectures are given below:

- *Four-node Linux cluster*: Pentium-II 333 MHz processors with 256 MB of RAM linked by 100 BaseT ethernet switch.
- *Hitachi SR2201 supercomputer*: 256×150 MHz processors with 2 floating point pipelines, 256 MB of RAM and 3D crossbar network operating at 300 MB/s. Three partitions were available for production runs and contained 64 processors each.

An almost-linear speedup was found to occur due to the very low rate of communication between nodes. For example, a production run of 4 parallel demes for 60,000 generations took 2-3 days on the Linux cluster. However, using a 64-node partition on the SR2201 supercomputer, 32 parallel demes could be processed in a standard 8 hour run.

5.3.4 Parallel GA

A coarse-grained parallel GA topology, known as the island model, was used in this implementation (see Figure 3.7b). Periodic migration was permitted

between the demes, which otherwise were allowed to evolve independently as if the populations were confined to islands. The actual separation of the calculation across the nodes depended on the number of available nodes and the requested number of demes. To minimize delays from communication, demes were equally divided among the nodes whenever possible. In the cases, though, where the number of demes was less than the number of nodes, the fitness evaluation stage was performed by sharing members of the population amongst the available nodes. Once all nodes had completed the fitness evaluation, they communicated the resulting fitness to the controlling node, which would perform the serial operations of selection, recombination and mutation.

Migration was performed between demes at periodic intervals and involved the following heuristic method

The member from the source deme which undergoes migration should be the fittest member such that, when migration is completed, it will not become the fittest member of the destination deme.

Therefore, a particular deme would send a copy of one of its fittest individuals, as long as it did not become the elite member of the destination deme. This paradigm avoided premature convergence between the demes and instead, the genetic material of the migrants was slowly absorbed into the population of the destination deme, possibly improving the existing designs.

5.4 Testing the GA

5.4.1 Helmholtz pair

In order to test the GA optimization code, a simple design task was created. The task consisted of optimizing a pair of thick Helmholtz coils. Usually Helmholtz coils are assumed to be infinitely thin and carry identical currents. If the separation is given by s , the diameter by d and the following

relationship [84, p.172] holds

$$s = \frac{1}{2} d \quad (5.2)$$

then the first three derivatives of the magnetic field cancel at the centre of the system, creating a fairly homogeneous magnetic field. In this case, the problem was altered slightly to optimize the placement of thick (10 cm wide and 10 cm deep) Helmholtz coils to minimize the inhomogeneity over an 8 cm DSV. Constraints were imposed on the optimization to ensure that a minimum field strength of 1 T was achieved using a fixed current density of 100 A/mm².

16 demes were optimized in parallel without migration, each with a sub-population size of 32 which was initialized randomly and then left to evolve for 500 generations. The results (see Figure 5.2) clearly show that all of the runs converged within 350 generations to a single optimal design given by

$$\begin{aligned} z_{\text{mid}} &\approx 44.96 \text{ cm} \\ r_{\text{mid}} &\approx 89.86 \text{ cm} \end{aligned}$$

where $z_{\text{mid}} = z_{\text{min}} + 0.5w$ and $r_{\text{mid}} = r_{\text{min}} + 0.5d$ define the location of the centre of the Helmholtz coils. Comparing the results to (5.2), we find $\frac{44.96}{89.86} \approx 0.5003$ (4 s.f.). The small deviation from exactly 0.5 is due to the finite size of the coils used in the design⁶.

5.4.2 Coil configurations

Having shown that the GA is capable of optimizing a single pair of coils, the next task was to investigate how the GA dealt with multiple coils. Therefore, a series of experiments was performed with varying numbers of coils. In particular, the coils making up a design were divided into positively wound coils, which contributed to the main field strength, and negatively wound coils, which reduced the fringe fields and improved the homogeneity. Figure 5.3 shows the different convergence rates using a wide range of possible

⁶See Kaminishi and Nawata [85] for an in depth description of the thick Helmholtz coils problem.

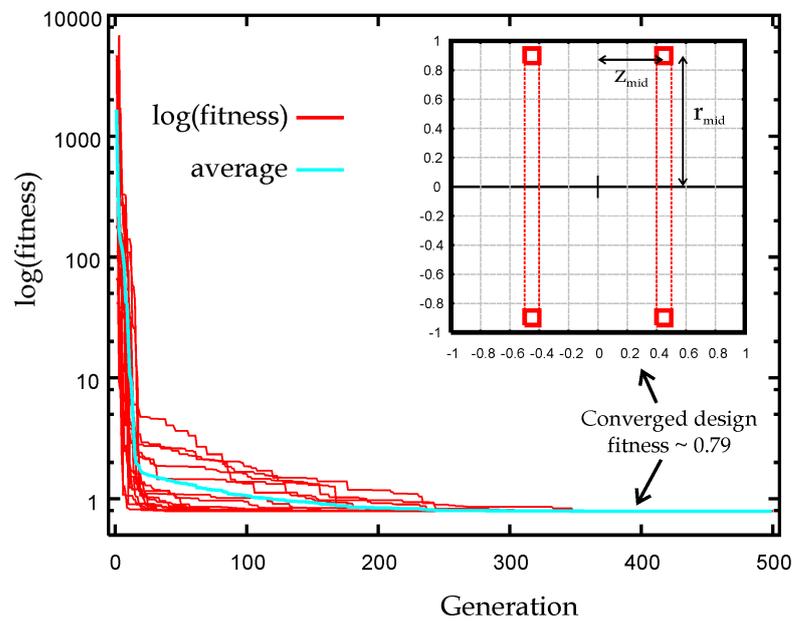


Figure 5.2: Convergence of thick Helmholtz coil test problem. The fitness is the standard deviation of the magnetic field over an 8 cm diameter sphere. The optimal design is shown in the inset.

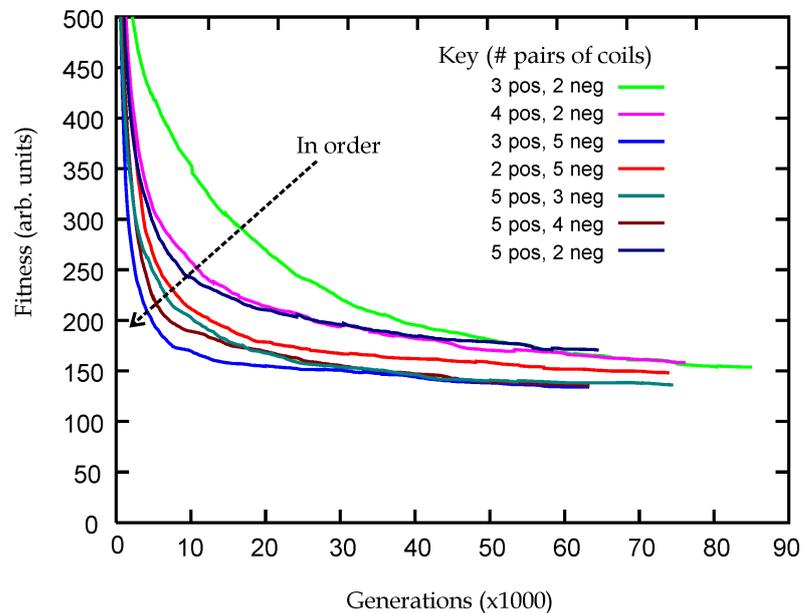


Figure 5.3: Comparison of the preliminary genetic algorithm optimization on various combinations of positive and negatively wound coils. The results showed little difference between the combinations, although the rate of convergence is slightly higher with more pairs of coils.

configurations. Surprisingly, adding more coils did not significantly improve the homogeneity of the final designs. It was found instead that roughly the same homogeneity (or fitness) could be obtained using very different numbers of coils, although designs containing more pairs of coils tended to converge to their final fitness slightly faster.

Using this information and considering the complexity of the designs, a decision was made to optimize a design containing four pairs of coils; two positively wound and two negatively wound.

5.5 Methods II: Applying the GA to the Short Bore Magnet

5.5.1 Main coil optimization

The GA described in the previous sections was used to optimize a short bore MRI magnet consisting of four pairs of coils. The coils were configured during initialization so that two pairs were positively wound and two pairs were negatively wound. The maximum length of the magnet bore was set to 1.0 m with a 1.0 m bore diameter and a maximum outer diameter of 2.5 m.

A fitness function was used which calculated the standard deviation of the z component of the magnetic field at a few sampled points within the ROI. Therefore the minimized function was

$$\text{Fitness} = \sqrt{\frac{1}{N} \sum_{i=1}^N (B_{z,i})^2 - \left(\frac{1}{N} \sum_{i=1}^N B_{z,i} \right)^2} \quad (5.3)$$

where N points in total were sampled across the ROI of radius R . In this case, a 6×6 grid of points was formed and the 28 points located within $1.05R$ of the origin were used to calculate the fitness (see Figure 5.4).

Various constraints were placed on the optimization and these were implemented by adding a penalty value to the basic fitness (see Section 3.3). The penalty values were determined by a process of iterative testing, resulting in the following formulae

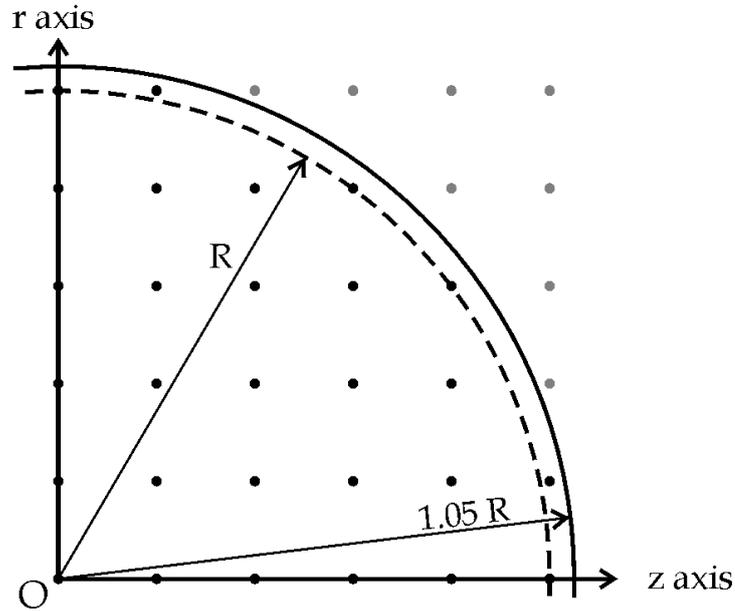


Figure 5.4: Sketch of the points at which the magnetic field was evaluated within the ROI.

- *Central field strength:* if the field strength at the centre of the ROI, $B_z(0,0)$, was less than the required field strength, $B_{\text{req.}}$, then the fitness value was given by

If $B_z(0,0) < B_{\text{req.}}$ then

$$\text{Fitness} = \text{Fitness} \times [1.0 + 10(B_{\text{req.}} - B_z(0,0))] \times 100$$

- *Fringe fields:* the fringe fields were constrained by examining the field strength at two points along the z axis. In this case, a fringe field of less than 5 gauss was required at 5 m. Therefore the field strength was calculated at 5 m and 6 m along the z axis. A penalty was added to the fitness value if the field strength at either point exceeded the required 5 gauss. Evaluating the field at two points along the z axis prevented the fringe field from simply inverting at $z = 5$ m.

If $B_z(z = 5\text{m}) > B_{\text{req.}}(\text{fringe})$ then

$$\text{Fitness} = \text{Fitness} + [B_z(z = 5) - B_{\text{req.}}(\text{fringe})] \times 10^6$$

If $B_z(z = 6\text{m}) > B_{\text{req.}}(\text{fringe})$ then

$$\text{Fitness} = \text{Fitness} + [B_z(z = 6) - B_{\text{req.}}(\text{fringe})] \times 5 \times 10^6$$

- *Hoop stress*: the final constraint was a limit imposed on the hoop stress. This was calculated by finding the z component of the field inside each coil and multiplying by the current density and the radius. Unfortunately, evaluation of the field within a coil was very slow, so only a single point could be chosen in each coil. This point was located at $(z_{\text{min}} + \frac{w}{2}, r_{\text{min}} + 0.5\text{mm})$, just inside the edge of each coil, where the field strength is highest for a solitary coil (see Figure 4.5). The fitness was penalized if the approximated hoop stress, S_{approx} , exceeded the nominal maximum stress, S_{max} , of 200 MPa.

If $S_{\text{approx}} > S_{\text{max}}$ then

$$\text{Fitness} = \text{Fitness} + [S_{\text{approx}} - S_{\text{max}}] \times 100$$

Production runs were performed on the Hitachi SR2201 supercomputer using 64 demes for 60,000 generations and a deme sub-population size of 128. Unfortunately, it was impossible to optimize the homogeneity of such a magnet to the required levels of 1 ppm max. to min. over a 40 cm DSV. In fact, typical inhomogeneities of roughly 50 ppm max. to min. over a 40 cm DSV were the best that could be achieved by the GA.

Taking into consideration fitness and the layout of coils, the best design was chosen with the following properties; 1.0 T field with homogeneity of 9.9 ppm and 48.4 ppm max. to min. over a 30 cm and 40 cm DSV respectively. Fringe fields were within the required distance and hoop stresses were properly constrained. The field in the ROI for this design is shown in Figure 5.5 and details of the positions and current densities of the four pairs of coils are given in Table 5.1.

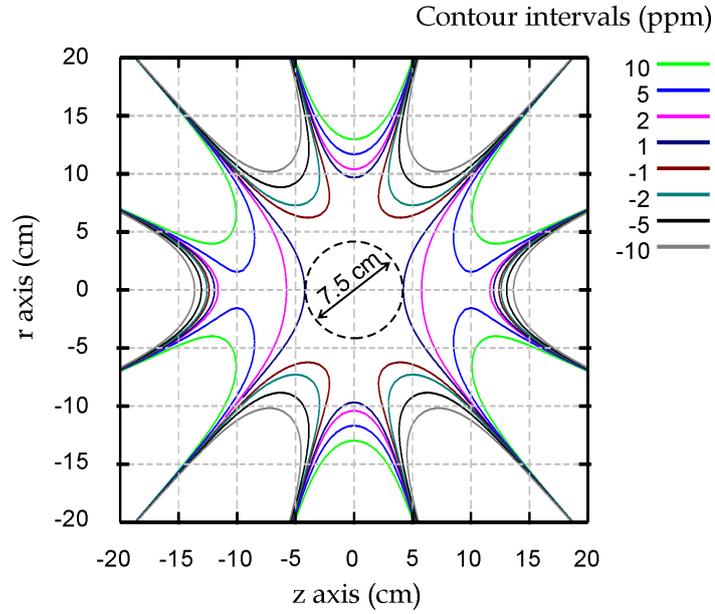


Figure 5.5: Contour plot of inhomogeneities over a 40×40 cm cylinder for four main superconducting coils. The highlighted region has 1 ppm max. to min. deviation over a 7.5 cm diameter sphere.

Coil	z_{\min} (cm)	r_{\min} (cm)	width (cm)	depth (cm)	J'' (A/mm ²)
m1	10.80	54.08	2.70	12.90	77.91
m2	28.71	57.86	4.70	12.90	-58.07
m3	40.72	63.00	9.20	24.70	66.43
m4	20.97	112.88	7.70	11.80	-65.61

Table 5.1: Positions and current density of four main superconducting coils optimized using the genetic algorithm, in the short whole-body MRI magnet design.

5.5.2 Superconducting shim coil optimization

To improve the homogeneity of this design, a second optimization stage was performed, involving the addition of a set of superconducting shims. It was found that by cancelling all non-zero orders up to and including 8th order⁷, that good homogeneity could be created over a 40 cm DSV. This was achieved with a superconducting shim set consisting of five pairs of coils. To set the currents in the shim set, a matrix inversion technique was used.

Firstly the spherical harmonic expansion weighting coefficients, q_n , were calculated for the four coils of the existing magnet design. Due to symmetry about the plane $z = 0$, these coefficients were zero for all values of odd n . The first five non-zero weighting coefficients q_0 , q_2 , q_4 , q_6 and q_8 were used to form a vector \mathbf{Q} . Secondly, the spherical harmonic weighting coefficients were calculated for each of the shim coils with unitary current density ($J'' = 1 \text{ A/mm}^2$) and these values were used to create a 5×5 matrix, \mathbf{M} . Finally, to find the required currents in the shim coils, the following matrix equation was solved

$$\mathbf{M}\mathbf{J} = \mathbf{Q} \quad (5.4)$$

where \mathbf{J} is a vector containing the required current densities in the superconducting shims. To maintain the field strength, the first component of the vector \mathbf{Q} was set to zero. Equation (5.4) may be rearranged to give

$$\mathbf{J} = \mathbf{M}^{-1}\mathbf{Q} \quad (5.5)$$

providing \mathbf{M}^{-1} , the inverse of \mathbf{M} , exists.

There remains the problem of where exactly the superconducting shim coils should be placed. To answer this, a second stage GA was used to optimize the positioning of the superconducting shims coils with the objective of minimizing the size of the remaining weighting coefficients in the spherical harmonic expansion.

⁷See Section 4.2.3 for an explanation of the meaning of ‘order’.

Coil	z_{\min} (cm)	r_{\min} (cm)	width (cm)	depth (cm)	J'' (A/mm ²)
m1	10.80	54.08	2.70	12.90	77.91
m2	28.71	57.86	4.70	12.90	-58.07
m3	40.72	63.00	9.20	24.70	66.43
m4	20.97	112.88	7.70	11.80	-65.61
s1	9.42	50.00	4.30	0.90	133.94
s2	14.81	50.02	4.90	3.40	-124.80
s3	20.91	50.00	8.60	3.50	114.68
s4	32.34	51.00	9.90	5.00	-104.90
s5	43.30	54.49	6.70	5.50	99.99

Table 5.2: Resulting short bore MRI magnet design with the addition of five pairs of superconducting shim coils.

The fitness function was expressed as

$$\text{Fitness} = \sum_{n=10, \text{even}}^{20} |q_n| \quad (5.6)$$

By iterative optimization of this fitness function, it was found that the hoop stresses in the superconducting shim coils were too high to consider construction. In an attempt to control these stresses, limits were placed on the allowed current densities in the five superconducting shims. Starting from the outer shim, the limits were set to $J_{\max} = 100, 105, 140, 140$ and 140 A/mm². These limits were enforced by adding a penalty function to the fitness if the current in the i^{th} coil, J_i , exceeded the limit, $J_{\max(i)}$, thus

$$\begin{aligned} &\text{If } J_i > J_{\max(i)} \text{ then} \\ \text{Fitness} &= \text{Fitness} + [J_i - J_{\max(i)}] \times 10^6 \end{aligned} \quad (5.7)$$

The results of this second optimization stage are shown in Table 5.2. The inhomogeneity of this design showed a standard deviations of 0.07 ppm over a 30 cm DSV and 1.0 ppm over a 40 cm DSV. Maximum to minimum deviations over the same DSVs were 1.0 ppm and 14 ppm respectively.

5.5.3 Quantization of parameters

The design, detailed in Table 5.2, would in practice be very difficult to maintain due to the fact that each coil operates at a different current density. Therefore, a final stage of the optimization process was implemented whereby the primary coils, m1-m4, were configured to operate at 68 A/mm² and the superconducting shim coils, s1-s5, were configured to operate at 110 A/mm². This alteration was performed using the first order relationship

$$w \times h \times J'' = \text{constant} \quad (5.8)$$

Thus, the change in cross-sectional area of each coil was inversely proportional to the change in current density. Once these alterations had been applied to the design however, a further process of re-shimming was required. To achieve this, five pairs of room temperature shim coils were added to the design and the currents were set using the matrix inversion technique described in the previous section. As well as removing much of the inhomogeneity from the design, the room temperature shim coils also allowed for the compensation of manufacturing errors. For simplicity, the shim coils were all located at an inner radius of 47 cm and had a cross-section of 2 × 2 cm.

5.6 Results and Analysis

The final magnet design is detailed in Table 5.3 and analyzed in the following sections.

5.6.1 Overall configuration

The overall configuration of the design is shown in Figure 5.6 and a detailed view of one quadrant is shown in Figure 5.7. The coils were numbered using the convention ‘m1-m4’ for the main coils carrying 68 A/mm², ‘s1-s5’ for the superconducting shims carrying 110 A/mm² and ‘r1-r5’ for the room temperature shims used to correct for positioning errors. The main coils and the superconducting shims were slightly flared from an inner radius of 50 cm

Coil	z_{\min} (cm)	r_{\min} (cm)	width (cm)	depth (cm)	J'' (A/mm ²)
m1	10.71	54.08	2.90	13.50	68.00
m2	28.89	57.86	4.30	11.90	-68.00
m3	40.77	63.00	9.10	24.40	68.00
m4	21.04	112.88	7.60	11.60	-68.00
s1	9.20	50.00	4.70	1.00	110.00
s2	14.65	50.02	5.20	3.70	-110.00
s3	20.82	50.00	8.90	3.60	110.00
s4	32.46	51.00	9.70	4.80	-110.00
s5	43.45	54.49	6.30	5.10	110.00
r1	8.00	47.00	2.00	2.00	4.2177
r2	15.00	47.00	2.00	2.00	-0.9457
r3	21.00	47.00	2.00	2.00	-1.9129
r4	32.00	47.00	2.00	2.00	-6.2943
r5	48.00	47.00	2.00	2.00	5.1421

Table 5.3: Positions and current density of the coils in final short, whole-body MRI magnet design. Main superconducting coils are labelled m1-m4, superconducting shims by s1-s5 and room temperature shim coils by r1-r5.

at the centre to 54.5 cm at the edge of the magnet. The central field strength was 1.0 T.

5.6.2 Homogeneity

The magnitude of the z component of the magnetic field over the ROI is shown in Figure 5.8 with the deviation of the field from the central value expressed in parts per million (ppm). The homogeneity has a standard deviation of 0.06 ppm over a 30 cm ROI and 1.06 ppm over a 40 cm ROI. Maximum to minimum deviations are 0.87 ppm and 14.65 ppms over DSVs of 30 cm and 40 cm respectively. The size of this DSV is perfectly suitable for surgical interventions and in particular brain surgery where the typical size of the region examined is of the order 25 cm.

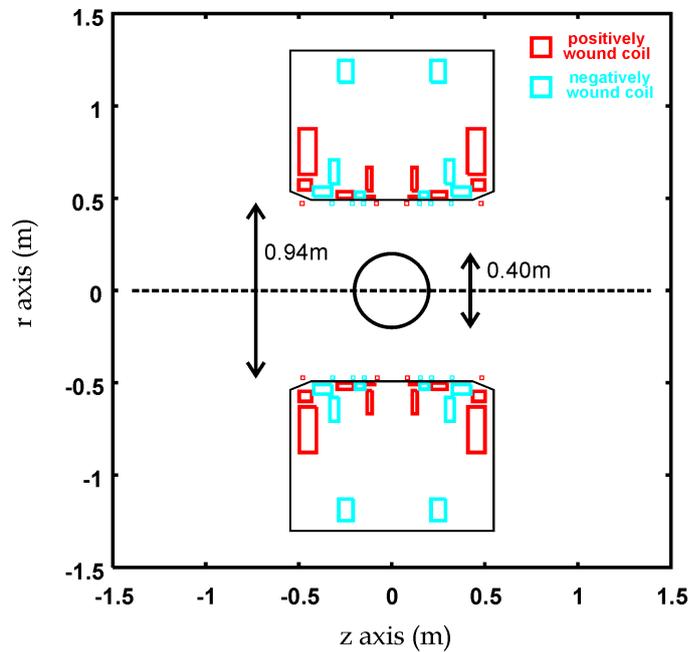


Figure 5.6: Cross-section of the final optimized short whole-body magnet. The positively wound coils are shown in red, whereas the negatively wound coils are shown in blue.

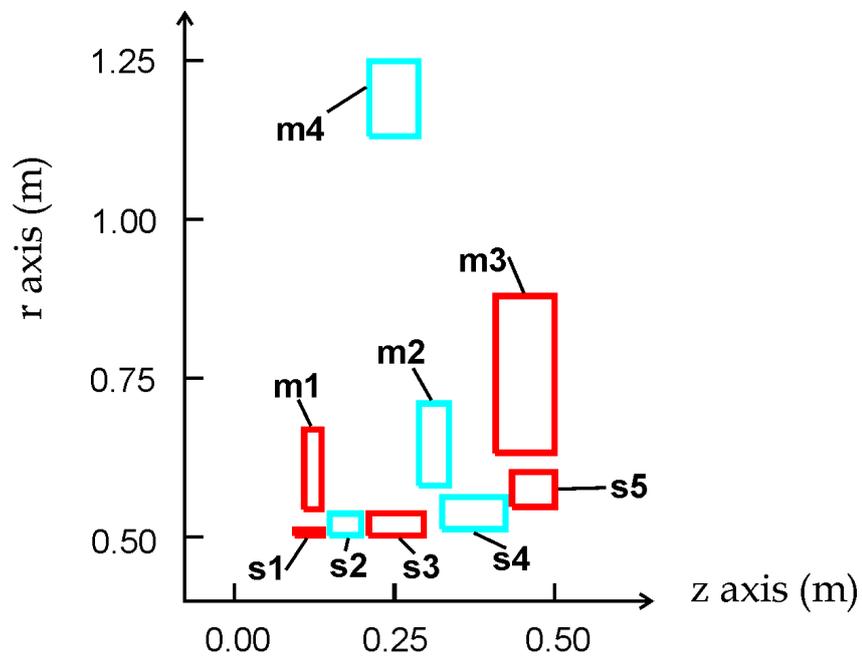


Figure 5.7: Close up of one quadrant showing the 4 main coils and the 5 superconducting shim coils. For clarity the room temperature shim coils are not included.

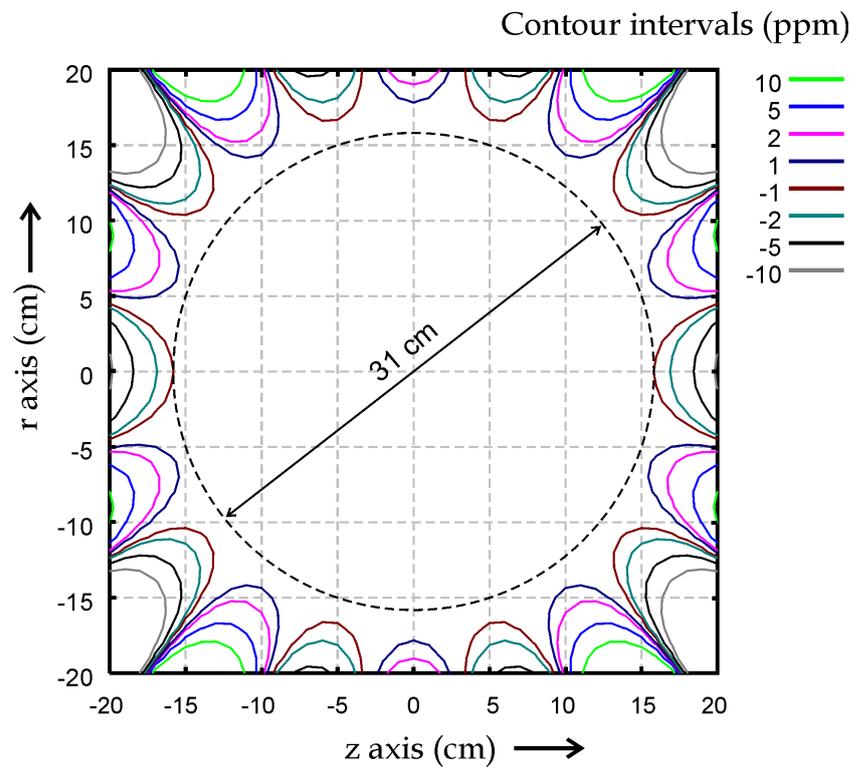


Figure 5.8: Homogeneity of short, whole-body magnet over the ROI. The uniform region is highlighted and is sufficiently large for routine brain imaging.

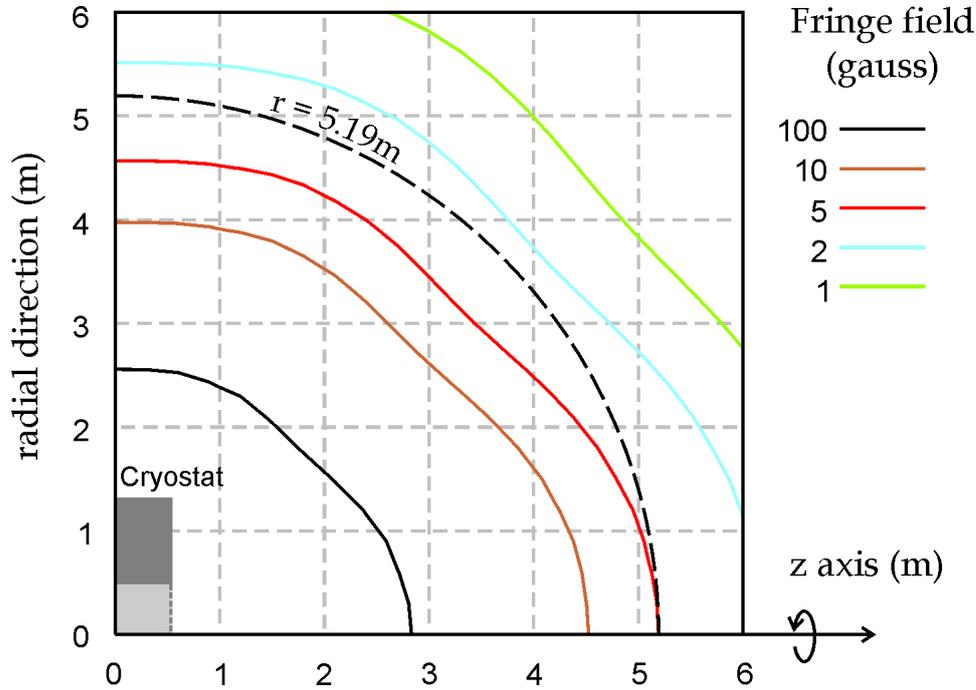


Figure 5.9: Fringe fields of short, whole-body MRI magnet. The 5 gauss line can be seen to lie entirely within a sphere of radius 5.19 m.

5.6.3 Fringe fields

The active shielding in the short MRI magnet reduces the magnetic footprint considerably. On axis, the magnetic field strength decays to 5 gauss at 5.19 m, whereas radially, the distance to the 5 gauss line is 4.56 m. The overall region in which the fringe field exceeds 5 gauss is contained entirely within a sphere of radius 5.19 m, covering a floorspace of approximately 54 m^2 . See Figure 5.9 for details. These dimensions compare well to other similar strength magnets especially considering the open aspect of the 1.0 T magnet.

5.6.4 Peak fields

The maximum field within the conductors, 5.1 T, occurs in coil s4 of the superconducting shim set. From Table 5.4 it can also be seen that s5 and m3 endure fairly high fields of 5.0 T and 4.2 T respectively. However, the main

Coil	Peak Field (T)	Hoop Stress (MPa)	Axial Force (MN)
m1	2.18	75.25	0.08
m2	3.09	-101.37	-1.17
m3	4.20	-179.63	2.30
m4	2.85	214.31	-1.03
s1	1.66	89.43	-0.07
s2	2.46	-135.51	-0.15
s3	2.48	134.65	-0.15
s4	5.10	-234.66	1.50
s5	5.00	257.44	1.00

Table 5.4: Peak field, hoop stress and axial force for final short, whole-body MRI magnet design. Hoop stresses are calculated using the independent hoop approximation (Section 4.5.1). The room temperature shim coils are omitted due to their low current densities.

coils (m1-m4) only operate at 68 A/mm² compared to 110 A/mm² for the superconducting shim set. This means that our attention should be placed, on the two outermost superconducting shim coils (s4 and s5) to check that they lie within the critical field boundary. Operating at 110 A/mm² and 4.2 K, the critical field for NbTi superconductor is roughly 10 T⁸. Often an operational temperature margin of 0.5 K is used to ensure that small temperature fluctuations do not initiate a quench (as discussed in Section 4.6.3). Even at 4.7 K, the critical field remains well above the operating field at about 9 T. This ensures that critical current limits do not restrict this design, rather the design is limited by the hoop stresses.

5.6.5 Stress and Forces

Stresses (including axial stress, shear stress and von Mises stress) were calculated for this short bore design using the finite element modelling program ABAQUS⁹. Figure 5.10 shows the calculated axial hoop stresses with a maximum tensile stress of 197 MPa and maximum compressive stress of 94 MPa. The remaining coils in the design had substantially lower stresses, as can

⁸See Figure 4.8.

⁹See Section 4.5.3.

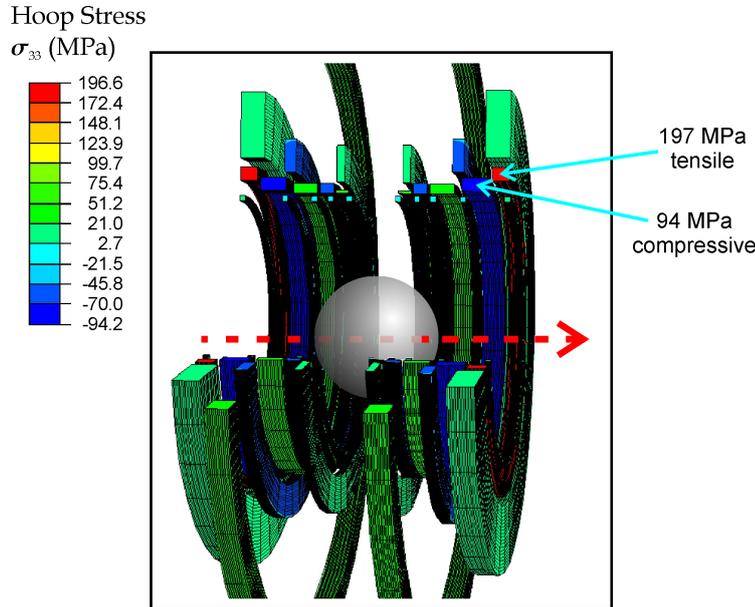


Figure 5.10: Cutaway view of the short, whole-body MRI magnet showing hoop stress, σ_{33} (MPa)

be seen from the figure. The outer two superconducting shims (s4 and s5) have fairly high internal stresses and would require some structural reinforcement, for example with steel reinforced superconductor. Such cable is formed by winding multiple strands of standard NbTi/Cu superconductor around a steel core. Axial forces are detailed in Table 5.4.

5.6.6 Sensitivity

Sensitivity of the main superconducting coils to manufacturing errors was measured using a Monte Carlo simulation. Movements in the axial direction were considered by shifting each coil by 0.5 mm and imperfections in the milling radius of the coil formers were also considered by expanding or contracting the inner and outer radii by 0.5 mm. Coils on both sides of the plane where $z = 0$ were subjected to different movements and the homogeneity of the ROI was recalculated. The results are shown in the main part of Figure 5.11, where the new homogeneity of the ROI is plotted against the

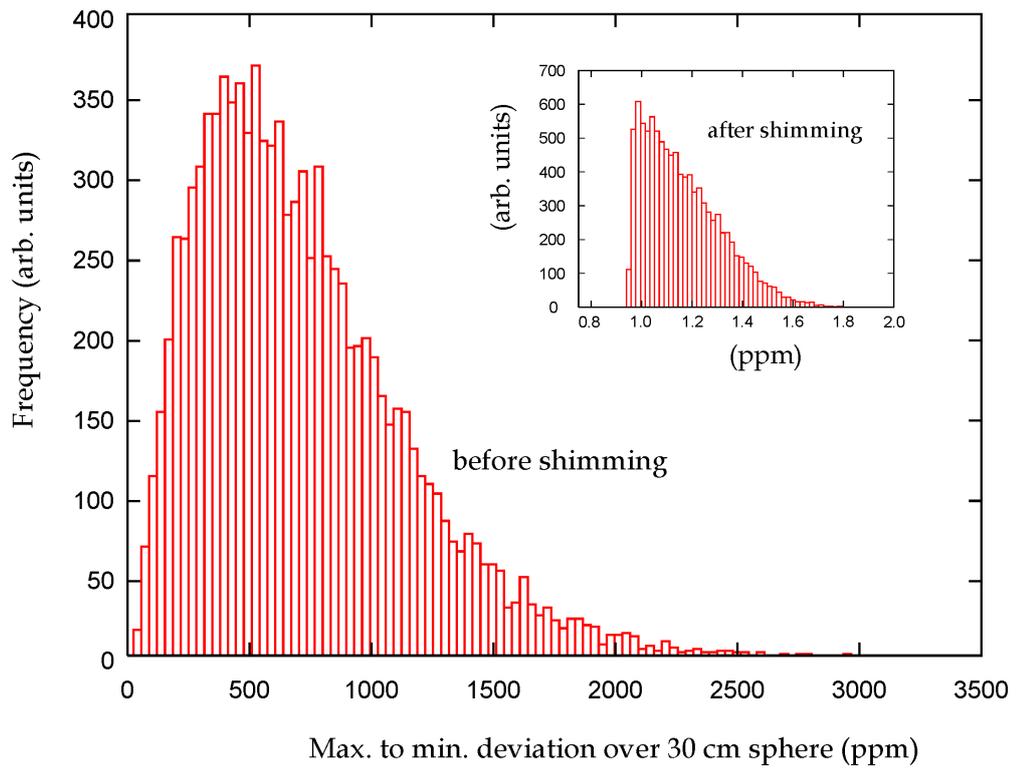


Figure 5.11: Simulation of 0.5 mm manufacturing errors showing frequency distribution of maximum to minimum homogeneity over a 30 cm sphere. Average degradation of the homogeneity was approximately 710 ppm. The figure shown in the inset details the analogous distribution after reshimming with the room temperature shim coils. The average homogeneity was reduced to 1.16 ppm.

frequency of the outcome. The average maximum to minimum deviation for 0.5 mm movements was 710 ppm with the tail of the graph reaching almost 3000 ppm.

In order to reshim the magnet a matrix inversion technique similar to the method in Section 5.5.2 was used. However, in this situation, there was no symmetry about the plane $z = 0$ resulting in the need to set the currents in all 10 room temperature shim coils independently. The inset of Figure 5.11 shows the corresponding graph when the room temperature shim coils are set to cancel the first 10 spherical harmonic terms. It can be seen that the homogeneity is restored to roughly its original value with an average inhomogeneity of 1.16 ppm. The maximum currents used in the room temperature shim coils to restore the homogeneity did not exceed 16 A.

It should also be noted, that general susceptibility effects of the patient could also be cancelled using the room temperature shims, however, more localized effects around the sinuses and ears would still have to be dealt with using specialized techniques [86].

5.7 Discussion

On several counts, the ultra short bore MRI magnet design proposed in this chapter compares well with conventional MRI magnets. The bore length of 1.0 m is significantly shorter than existing magnet designs, typically 1.5-2.5 m, and the ratio of bore length to bore diameter is very nearly unity, reducing the effects of claustrophobia and increasing access for interventional surgery. Homogeneity over the central ROI is acceptable with 1.06 ppm standard deviation over a 40 cm DSV or 0.87 ppm max. to min. over a 30 cm DSV, whilst the fringe fields are reduced to less than 5 gauss within a sphere of radius 5.19 m. Using a GA, multiple constraints have been simultaneously applied, limiting the fringe-fields, the hoop stresses and the maintaining the central field strength. By including all of the relevant constraints in the fitness function, it has been shown that an *ab initio* approach to MRI magnet design using stochastic optimization is feasible.

Recent work by Zhao *et al.* [87] has added to the theoretical understand-

ing of short bore MRI magnet design. Zhao investigated the properties of MRI magnet designs found by the target field¹⁰ method as a function of the overall magnet length. It was found that for magnets with a length to diameter ratio less than 1.4, the solutions to the target field method exhibit regions with negative current density. These regions are interpreted as ‘bucking’ coils wound in the opposite direction to the main field producing coils. This leads to two effects seen in ultra-short, whole-body MRI magnets. Firstly, the overall efficiency of the magnet is dramatically reduced, with the counter-wound coils working against each other to cancel field strength. This is not too much of a problem as the 1.0-3.0 T field strengths used in MRI are relatively low field strengths for superconducting magnets. However, any increase in the required volume of superconductor can have a large impact on the cost of the system. Secondly, the bucking coils create regions within the magnet where the field strength is much higher than usually encountered due to the compaction of the flux lines between positive and negatively wound coils. This leads to problems with strong forces and stresses as well as potentially limiting critical currents.

It is interesting to note that the original design from the genetic algorithm (see Table 5.1) used a single negatively wound coil ‘m2’ in precisely the expected position for this bucking coil. Also, attempts to include more coils in the optimization did not significantly improve the results (see Figure 5.3), possibly due to the coils clumping together to form localized regions of current density.

Several other techniques have been used in the literature for the design of ultra short bore MRI magnets. For example, Crozier *et al.* [89] use the target field method as part of a hybrid optimization process combined with simulated annealing. The two methods are required as the target field method produces a continuous current distribution on the inner bore surface, which must subsequently be approximated by quantized, thick solenoidal coils. As a last step in the process, simulated annealing is used to reoptimize the posi-

¹⁰The target field method, popularized by Turner [88] for the design of gradient coils, is still widely used in magnet and gradient coil design. The technique is very efficient, but cannot deal with integer parameters as easily as genetic algorithms or other stochastic optimization methods.

tions. Alternatively, the designs may be finally optimized using a non-linear least-squares scheme [87]. Other methods that have been examined include the use of matrix subset selection [90] and genetic algorithms [91]. However, the literature has so far tended to ignore fundamental constraints which make MRI magnets feasible. An example is a symmetric magnet design in [87] with a bore length of roughly 1.0 m, a bore diameter of roughly 0.9 m and a central field strength of 1.0 T. Whilst the homogeneity of the 50 cm DSV is very impressive (standard deviation of 3.4 ppm) and the peak fields within the superconductor are reasonable at 6.5 T, an approximation (4.83) of the hoop stress in the one of the coils gives over 800 MPa, which clearly rules out this magnet as these stresses cannot be supported.

In terms of commercial designs currently on the market, the shortest magnet design is the Philips *Infinion* [37]. This 1.5 T magnet has a bore length of 1.4 m and has an elliptical ROI, 50 cm across in the anterior-posterior direction and 40 cm across in the axial direction. Ignoring the trimming and cryostat, the coil dimensions do not exceed 1.27 m in length and a flared whole-body gradient gives the impression of an even shorter bore. Compared to the design optimized with the GA, the *Infinion* compromises its length (25% longer) for a larger imaging volume and a higher operating field strength. Unfortunately, the methods of optimization used in industry are still shrouded in secrecy, so further comparisons are very difficult.

At this point, the feasibility of the short bore magnet proposed in this chapter should be considered. Although many important constraints have been enforced, the sheer complexity of the design makes construction a difficult prospect. The magnet consists of 18 superconducting coils and 10 room temperature shim coils, many more than conventional MRI magnets which typically contain about 10 coils. Furthermore, the high stresses on the two outermost superconducting shim coils may well cause problems either through premature quenching or simply in the construction of sufficiently strong formers. In cases where hoop stresses are slightly higher than usual, it is possible to reinforce the superconductor with stainless steel. However, these factors would probably lead to a higher cost when compared to existing systems. Cost is also likely to be significantly higher due to the use of

0.515 m³ of superconductor, which is 3-4 times the quantity used in conventional MRI designs.

Cost optimization has been studied by only a handful of authors including Xu *et al.* [76] and Prestmon *et al.* [77]. Xu arrived at the conclusion that magnets with a length (L) shorter than

$$L_{\text{critical}} = 1.2DSV + 0.77D, \quad (5.9)$$

become very expensive to build¹¹. The methods used to obtain this result remain unpublished, but the evidence from the ultra-short magnet proposed in this chapter and the design of the *Infinion* magnet appear to support his conclusions.

It may well be that although ultra short bore magnet design is possible in principle, building these magnets at the same cost as conventional designs remains unfeasible.

5.8 Conclusions

In conclusion, it has been shown that ultra short bore magnet design using genetic algorithms is possible and merits further investigation. An *ab initio* approach to the design problem, including all of the relevant constraints on design has produced a design which is 30-40% shorter than current conventional designs and is therefore suitable for interventional use and in cases where the patient is claustrophobic. A real-coded genetic algorithm was used in a two-stage optimization procedure, firstly to position the main coils and then to position the superconducting shim coils. A third stage involved the inclusion of a set of room temperature shims. The resulting magnet design compares favourably with the existing literature on short bore designs.

It was found that the inclusion of a set of room temperature shim coils allows for the correction of manufacturing errors as large as 0.5 mm with currents below 16 A and that the shim coils would also be suitable for correcting general inhomogeneities caused by susceptibility of patients in the scanner.

¹¹Where DSV is the diameter spherical volume and D is the bore diameter.

Almost all of the constraints placed on the design were met. Surprisingly, the most difficult constraint placed on the MRI magnet design was found to be the hoop stress limit. By confining the hoop stress to below 200 MPa, compromises had to be made on the homogeneity of the central field strength.

Unfortunately, the complexity of the design has contributed to an elevated cost projection for this magnet, which would make the design uncompetitive compared to existing designs¹². Indeed, keeping the costs down is one of the most difficult problems to overcome as the length of MRI magnets decreases. This is due to the need for bucking coils near the ROI and the need to have increasingly high alternate current densities along the inner magnet bore. Ways of relaxing the apparent constraints remain an area of research in which genetic algorithms may well play a role.

Finally, in the search for interventional magnets, other possible magnet configurations may be considered, for example, ‘C’ shaped magnets [92]. Another option is to create a split-coil magnet system with radial access for the surgeon [93]. Designs with radial access to the central homogeneous ROI, although for a different purpose, are the subject of the next chapter.

¹²The design presented in this chapter has been quoted at 5,000,000 US\$ for one-off manufacture.

Chapter 6

Split-coil MRI Magnet for Combined PET/MRI

6.1 Introduction

A wide range of biomedical imaging techniques are now available to physicians including magnetic resonance imaging (MRI), computed tomography (CT), positron emission tomography (PET) and ultrasound. Each technique or modality provides slightly different information with some, such as ultrasound, providing better anatomical information and others, such as PET, providing primarily functional information. Recently, interest has been sparked in combined imaging systems following very encouraging results from a commercial dual-imaging PET/CT system [94, 95]. Combining PET and CT, two complimentary imaging modalities, into a single system allows near-simultaneous images to be obtained and results in several important advantages. Firstly, images taken from the two modalities may be more accurately matched (co-registered) as the patient does not have to be repositioned between scans [95]. Secondly, photon attenuation maps for PET reconstruction may be generated from the CT images [96]. Thirdly, the combined system produces better overall sensitivity and specificity for diagnosis of primary tumours [95], therefore improving patient management and avoiding unnecessary treatment.

Another exciting candidate for dual modality imaging is combined PET and MRI [97]. A truly simultaneous combined PET/MRI scanner would be remarkable due to the good sensitivity of MRI to soft tissues along with the quantitative functional activation data available from PET. As in combined PET/CT, the simultaneous acquisition of PET and MRI data would offer more diagnostic information than the sum of the individual modalities. Combined PET/MRI would allow almost perfect spatial registration using the known characteristics of each imaging modality and would permit the correction of partial volume effects. Furthermore, temporal correlation between the modalities would allow the removal of motion blurring from the PET data. Photon attenuation maps could be reconstructed from the MRI data and perhaps most importantly of all, targeted molecular imaging would be feasible with nano to picomolar sensitivity [97] from PET combined with anatomical information available from MRI.

In this chapter, a novel way of combining PET and MRI is considered using a split-coil magnet. The genetic algorithm (GA) used in the previous chapter for the design of the ultra short bore MRI magnet is adapted and improved. Results are compared to a commercial design from a leading manufacturer.

6.2 Combined PET and MRI

6.2.1 What is PET?

Positron Emission Tomography (PET) is a form of emission tomography that provides unique functional information *in vivo* about the viability of tissue and organs. The technique is used to measure the spatial distribution of positron-emitting radionuclides. Such radionuclides include proton-rich isotopes of carbon, nitrogen, oxygen and fluorine. The most common example used in PET is ^{18}F , an unstable isotope of fluorine which has a half-life of 112 minutes [11, p. 122] and decays via the following nuclear process



where p =proton, n =neutron, ν =neutrino and β^+ =positron. The emitted positron (β^+) will, *in vivo*, travel only a short distance (≈ 1 mm) before coming to rest and annihilating with an electron in the surroundings. In this annihilation, both energy and momentum are conserved and two 511 keV gamma rays are emitted with a 180° angular separation. Detectors are used to localize the source of the electromagnetic radiation. They rely on coincidence detection over a short timing window (2-10 ns) and the fact that the annihilation photons are produced back-to-back.

A PET scanner usually contains several rings of scintillation crystals¹ which produce a flash of light for each incident photon. Light from several scintillation crystals is amplified using a position sensitive photomultiplier tube (PMT) which creates an electrical pulse for analysis. PET images are acquired in either 2D mode or 3D mode. In 2D (multislice) mode, lead or tungsten shields are placed between detectors to absorb scattered radiation and restrict the angular range of each crystal. In 3D mode, no shields are used and each scintillation crystal is sensitive to photons from a much wider area. Using 3D mode allows lower doses of radiation to be administered, but more sophisticated scatter correction algorithms must be implemented. For further details about PET many good text books are available, e.g. [98].

6.2.2 Challenges of combining PET and MRI

Although the advantages of combined PET and MRI are enormous, there are many technical challenges that must be overcome before simultaneous PET/MRI imaging becomes a reality. These challenges may be divided into three categories

1. *Effects of MRI on PET*: Photomultiplier tubes (PMTs), a key component of any PET system are very sensitive to even very small magnetic fields [99, 100]. State-of-the-art position sensitive PMTs can operate in field strengths up to 1 gauss, although with soft-iron shielding, they

¹Scintillation crystals are made from many materials including Bismuth Germanate Oxide (BGO), Lutetium Oxyorthosilicate (LSO) or Thallium-doped Sodium Iodide (NaI-Tl).

may be used in fields strengths up to 500 gauss.

2. *Effects of PET on MRI*: Additional inhomogeneity over the region of interest (ROI) may be induced by magnetically active materials from the PET system placed within the magnetic footprint. Whilst it has been shown that scintillation crystals do not cause any problems [97], shielding associated with the PMTs may well affect the flux lines and thus the homogeneity of the MRI system. Any conducting materials placed near the central ROI may also cause problems due to eddy currents [101, 102].
3. *Electromagnetic interference (EMI)*: Interference between the imaging modalities is a serious cause for concern. For example, radio frequency interference from the PET electronics may well lead to banding artifacts in the magnetic resonance images. Similarly, strongly switching gradients and RF pulses from MRI could well lead to spurious activations in the PMTs and related electronics [97].

Ideally, the aim for a combined PET/MRI scanner must be to achieve a performance which is as good as the state-of-the-art performance attainable from the individual modalities.

6.2.3 Previous attempts

In recent years, Cherry, Marsden *et al.* have performed ground-breaking preliminary work on simultaneous PET and MRI [101, 97, 102]. Their approach has been to develop a miniaturized PET detector capable of operating in the environment of a conventional clinical MRI scanner. Two prototypes have been built so far using roughly the same geometry. The original prototype consisted of a ring of 48 LSO crystals ($2 \times 2 \times 10$ mm) in a ring with diameter 48 mm [101]. Each scintillation crystal was linked via 4 m of optical fibre to one of three position sensitive PMTs located outside the region of strong magnetic field. Figure 6.1a illustrates the experimental setup. More recently, a prototype with 72 LSO scintillation crystals ($2 \times 2 \times 5$ mm) has been constructed with an inner bore diameter of 56 mm [97].

Remarkable results have been achieved with these systems. In 1997, simultaneous MRI and PET images were successfully acquired within a 0.2 T open bore magnet [101]. With the second prototype, a study was made of the artifacts induced by the combination of imaging modalities and showed no significant interaction even with echo-planar imaging [97]. More recently, simultaneous PET and nuclear magnetic resonance (NMR) has been performed on isolated, perfused rat hearts [103] using the same system.

However, the geometry of this prototype leads to several problems. Firstly, the orientation of the scintillation crystals leads to poor sensitivity to the gamma rays. A better orientation, with the long direction parallel to the direction of the electromagnetic radiation, gives a much higher interaction rate. Secondly, there is significant attenuation in the long (4 m) optical fibres which are required to site the PMTs in a region of low field strength. Thirdly, with this system, a very small imaging region is available and access to the imaging volume is restricted by the location of the LSO crystals and the optical fibres. Finally, the small ring prototype may only obtain 2D images. This 2D data must be compared to usually 3D images from MRI, possibly introducing an additional partial volume effect.

6.2.4 A novel approach

In this chapter, a novel geometry for combined PET and MRI is proposed to overcome some of the problems encountered by previous attempts. The main feature of this proposal requires the construction of a split-coil MRI superconducting magnet. By splitting the magnet into two halves and creating a room temperature gap, the PET components may be positioned radially around the imaging volume, which is located at the centre of the magnet. Within the gap, scintillation crystals may be located, offering 360° coverage. A number of PMTs may then be placed radially away from the centre of the magnet in a region where the magnetic field is sufficiently reduced. As the magnetic field tends to drop away faster radially than axially, relatively short optical fibres may be used to link the scintillation crystals and the PMTs (see Figure 6.1b).

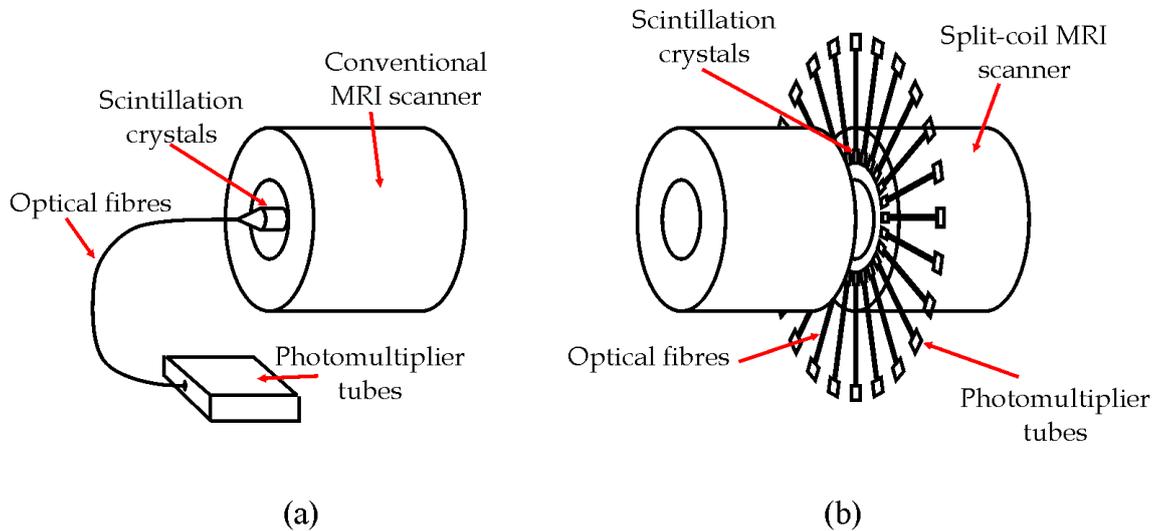


Figure 6.1: Two geometries for simultaneous PET and MRI. In (a) a small ring of scintillation crystals is inserted into the bore of a conventional MRI scanner. In (b) the PET detection system is located in a gap between the two halves of a split-coil MRI magnet.

Several advantages may be noted as a consequence of this novel geometry. Firstly, the orientation of the scintillation crystals is optimized to allow the maximum length of crystal for interaction with the gamma rays. Secondly, the attenuation caused by the fibre optics is reduced by shortening their length. Finally, depending on the size of the gap, it is likely that full 3D PET could be implemented. This is a significant advantage as the images obtained from MRI are often full 3D volume data sets.

Given these advantages, a decision was taken to use the GA developed in the previous chapter to optimize a split-coil main field MRI magnet.

6.3 Specifications and Requirements

As the design was novel, exact specifications were not available and could only roughly be envisaged. A leading manufacturer of magnetic resonance magnets in the UK, Magnex Scientific Limited [104], was approached and discussions were held about the feasibility of such a split-coil design. Very

soon, it became clear that our efforts should be focussed on designing an animal scale system due to the cost constraints and also the importance of the potential applications in drug testing. After consultation, the following flexible requirements were agreed upon

Field strength	1.0 T
Homogeneity	as good as possible
Length	< 1.5 m
Outer diameter	< 1.5 m
Inner bore diameter	20-30 cm
Split-coil gap	10-40 cm (as wide as possible)
Fringe fields	5 gauss within < 5 m
Volume of superconductor	minimal
Number of coils	3 pairs (for simplicity)

As usual, limits on the critical currents and hoop stresses had to be enforced. At this point, the optimization proceeded along two separate avenues. Magnex Scientific Ltd were commissioned to perform a design study of the split-coil magnet and produce their own optimal design. Simultaneously, the GA which had been used for the design of the ultra short bore MRI magnet in the previous chapter was implemented to optimize our split-coil magnet design. The two designs are subsequently compared in Section 6.5.3.

6.4 Methods I: Implementation of Genetic Algorithm

Work on the ultra short bore, whole-body MRI scanner showed some of the promise of stochastic optimization techniques [78, 79] on real-world engineering designs. For the optimization of the split-coil MRI magnet, the same GA (see Section 5.3) was used, but several improvements were made to the efficiency of the code which allowed more of the search space to be examined in less time. For completeness, the GA parameters are summarized below:

Population size	128 per deme
Number of demes	up to 128
Initialization	random
Selection	binary tournament
Prob. of crossover	0.9
Crossover	single point
Mutation type	Gaussian real-number
Mutation rates	$\left\{ \begin{array}{l} \text{spread} : p_{\text{indiv.}} = 0.5 \\ \phantom{\text{spread}} : p_{\text{comp.}} = 0.1 \\ \text{jump} : p_{\text{indiv.}} = 0.1 \\ \phantom{\text{jump}} : p_{\text{comp.}} = 0.05 \\ \text{kick} : p_{\text{indiv.}} = 0.05 \\ \phantom{\text{kick}} : p_{\text{comp.}} = 0.05 \end{array} \right.$
Elitism	yes

6.4.1 Specifying the homogeneity

In the previous chapter, the homogeneity over the ROI was found by taking the standard deviation of the magnetic field evaluated at 28 points located on a grid within the ROI. The actual magnetic field was calculated using the well known spherical harmonic expansion method (see Section 4.2.1).

In this chapter, a slight variation of this technique was applied. Rather than using a grid of points, the standard deviation was derived analytically from the spherical harmonic weighting constants (q_n). Thus, if the z component of the central field (4.54) was given by

$$B_z = \sum_{n=0}^{\infty} q_n \left(\frac{r}{r_0} \right)^n P_n(\cos \theta) \quad (6.2)$$

then the standard deviation of the inhomogeneity of the magnetic field over

a ROI with radius R could be expressed (4.63) as

$$\sigma(B_z) = \sqrt{\sum_{n=1}^{\infty} \frac{3 q_n^2}{(2n+1)(2n+3)} \left(\frac{R}{r_0}\right)^{2n}} \quad (6.3)$$

where r_0 is an arbitrary radius. The main advantage of using this method was that the homogeneity depended only on the radius of the ROI (R) and not on the choice of grid points.

6.4.2 Fringe field evaluation

One of the most important constraints placed on the GA was the limit applied to the magnetic fringe field. Various methods were tested in the previous chapter to evaluate and constrain the fringe field, whilst maintaining the speed of the fitness function. In the end, a compromise was reached whereby, only two points along the z axis were evaluated. However, this method did not always completely constrain the fringe field.

In order to constrain the fringe fields properly, the magnetic field should be evaluated at several points around the magnet (typically 12 points). However, fringe field evaluation is generally much slower than the central field evaluation as a method involving numerical integration has to be used to calculate both components of the magnetic field (see Section 4.2.2). For example, finding the fringe field using numerical evaluation at 12 points took up over 50% of the time spent on the fitness function.

To improve the speed of the fitness function, a rapid fringe field evaluation method was invented and implemented (see Section 4.3 for full details). The method provides a way of directly calculating the far-field spherical harmonic weighting constants, h_n , from thick solenoidal coils (4.48). These weighting constants can then be used to evaluate the magnetic fringe field at arbitrary points outside the magnet by applying

$$B_r = \sum_{n=1}^{\infty} h_n \left(\frac{r_0}{r}\right)^{n+2} P_n(\cos \theta) \quad \text{and} \quad (6.4)$$

$$B_\phi = - \sum_{n=1}^{\infty} \frac{1}{n+1} h_n \left(\frac{r_0}{r}\right)^{n+2} P_n^{-1}(\cos \theta) \quad (6.5)$$

This novel technique was found to be 216 times faster than previous hybrid methods [105]. Consequently, the percentage of time spent evaluating the fringe field constraint diminished to less than 5% of the overall fitness function.

6.4.3 Constraint handling

A new way of handling constraint violation was also implemented in this chapter. Whereas previously, a large penalty constant was added to the fitness function for constraint violation [18, 19], this new method allows better comparison of widely different constraints and avoids the need for large, arbitrary penalty constants. Suggested by Deb [17], the constraint handling technique has the following properties:

- Any feasible solution is preferred to any unfeasible solution
- Among two feasible solutions, the one having the better objective function value is preferred
- Among two unfeasible solutions, the one having smaller constraint violation is preferred.

When a constraint is violated then a normalized value in the range [0-1] is added to the overall penalty. Mathematically, the constraint penalty can be expressed as a sum of contributions from each of the N constraints. If $x_i(c)$ is the limiting value of x_i and the scaling parameter is σ_i , then

$$Penalty = \sum_{i=1}^N \begin{cases} \tanh\left(\frac{x_i - x_i(c)}{2\sigma_i}\right) & \text{if } x_i > x_i(c) \\ 0 & \text{if } x_i \leq x_i(c) \end{cases} \quad (6.6)$$

The scaling parameters σ_i are expressed in units of the constraint x_i and give an approximation to the value of constraint violation when the penalty reaches a value of 1/2.

6.4.4 Summary of fitness function

A summary of the fitness function used in the optimization of the split-coil MRI magnet is given below with the associated constraints

$$Fitness = \sigma(B_z) \quad (6.7)$$

where $\sigma(B_z)$ is given in 6.3. Under the new constraint handling technique, penalties were calculated separately from the fitness. The following constraints were applied to the optimization

- *Central field constraint:* if the central field strength (B_0) was less than the desired central field strength ($B_{req.}$), a penalty term was added using $\sigma_B = 0.5$ T

$$Penalty = \tanh\left(\frac{B_{req.} - B_0}{2\sigma_B}\right) \quad (6.8)$$

- *Fringe field constraint:* field magnitude was calculated for 12 points at a fixed distance (r_{fringe}) from the iso-centre. If the maximum field strength at any of the 12 points (B_{fringe}) exceeded the fringe field constraint ($B_{constraint}$), then a penalty term was enforced where $\sigma_{fringe} = 200$ gauss

$$Penalty = Penalty + \tanh\left(\frac{B_{fringe} - B_{constraint}}{2\sigma_{fringe}}\right) \quad (6.9)$$

- *Hoop stress limit:* the hoop stress was found by calculating the z component of the magnetic field at a single point within each coil, located at $(z_{min} + \frac{w}{2}, r_{min} + 0.5\text{mm})$ (c.f. Section 5.5.1). If the maximum approximate stress found (S_{approx}) in any coil was greater than the maximum allowable stress (S_{max}), then the penalty term was set using $\sigma_{stress} = 50$ MPa

$$Penalty = Penalty + \tanh\left(\frac{S_{\approx} - S_{max}}{2\sigma_{stress}}\right) \quad (6.10)$$

A final attempt was made to streamline the fitness evaluation code by comparing the level of required accuracy to the level of obtained accuracy. It was found that all of the evaluation routines contained much more accuracy than was necessary. By reducing the accuracy, for example, of the numerical integration for stress evaluation, the typical evaluation time of a 6 coil design was reduced from 4.1 ms to 2.6 ms. In other words, 60% more designs could be examined in the same period of time.

Finally two new computer systems, which had not been available previously were used for this optimization

1. *16-node Linux cluster*: Forming a major part of the Wolfson Brain Imaging Centre computing facilities [106], this cluster contains 16 dual Pentium-III 866 MHz processors with 1 GB RAM each connected by 100 MB/s Fast ethernet.
2. *IBM SP-II supercomputer*: Part of the HPCF [83], the IBM SP-II contains 10 nodes each with 16 Power3-II 375 MHz processors offering 4 floating point pipelines. Each node contains 12 GB memory and is linked to 500 MB/s communication channels. Partitions of 64 processors are readily available giving 96 Gflops peak performance.

6.5 Methods II: Applying the GA to the Split-Coil Magnet

6.5.1 Initial investigations of the search space

An initial investigation of the search space for designs consisting of three pairs of superconducting coils was performed. In particular, the effect of two important parameters, the inter-coil gap and the volume of superconductor, was considered. Multiple runs of the GA were executed with different constraints on these parameters and the size of the homogeneous region (max. to min. deviation < 1 ppm) was found.

Every run consisted of 32 parallel demes with a population size of 128, optimized for 100,000 generations. If a particular deme did not improve for

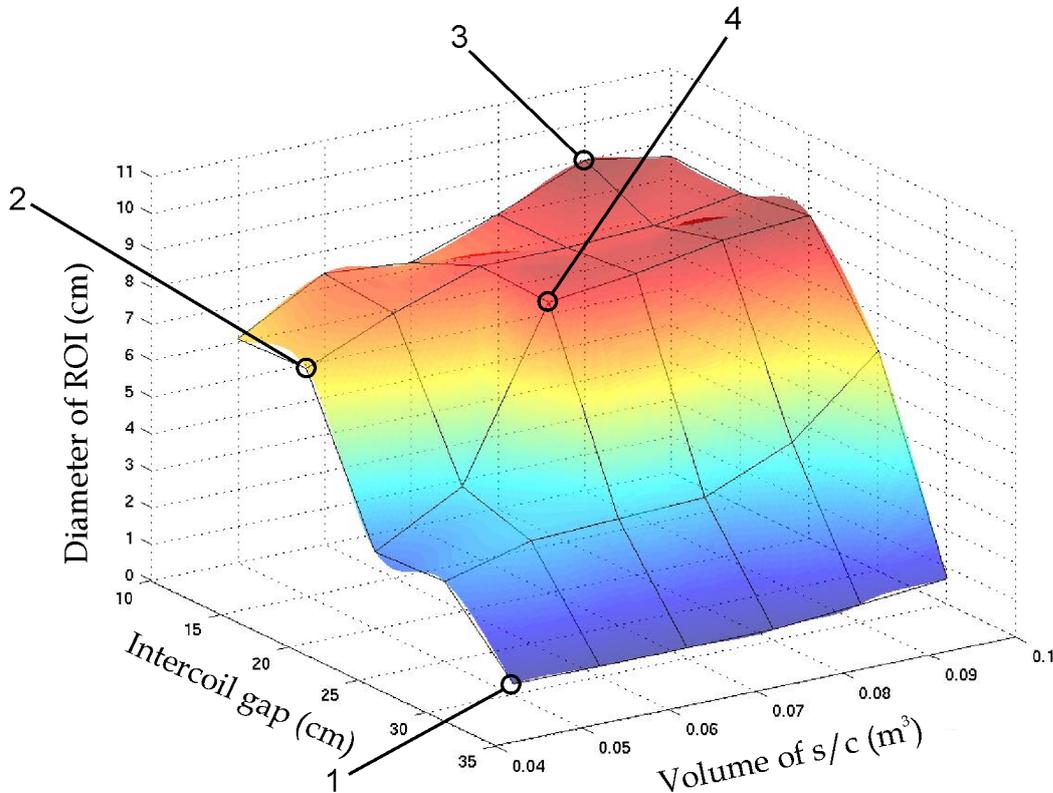


Figure 6.2: The optimized diameter of the ROI, given as a function of volume of superconducting wire (cost) and inter-coil separation.

more than 2000 generations, it was said to have stagnated and the deme was randomly re-initialized. On average, about 150 different designs were found during the course of each run. This spread indicated that the GA was becoming entrapped in a myriad of local minima, rather than converging to a global minimum. However, fitnesses of the best designs were usually very similar.

The size of the homogeneous region of the best design is plotted in Figure 6.2 against the two parameters, inter-coil gap and volume of superconductor. There are four points of special interest on this plot:

1. This point represents the smallest volume of superconductor (0.05 m^3) and most accessible (30 cm gap) design considered. It suffers though

from a very small homogeneous ROI with a diameter of only 1 cm.

2. Maintaining minimal cost (0.05 m^3), this point represents a good compromise between size of ROI (7 cm) and inter-coil separation (15 cm). It can be seen that reducing the inter-coil separation to 10 cm appears to allow little improvement in the size of the ROI.
3. This point represents the design with the largest ROI, almost 10 cm in diameter. The inter-coil gap, however, is rather small (10 cm) and the volume of superconductor used in the design (0.09 m^3) indicates higher than desirable cost.
4. This point appears to be a good compromise with a decent size for the ROI of 8 cm whilst reducing the volume of superconductor to 0.07 m^3 and allowing an inter-coil gap of 20 cm.

This initial investigation shows not only which of the compromises are important for such a design procedure, but also provides quantitative data about the actual trade-offs involved.

6.5.2 Preliminary results

Using the understanding of the search space gained, further computational runs were performed with an inter-coil gap of 20 cm and a maximum volume of 0.07 m^3 superconductor. The resulting design exhibited a standard deviation of 0.5 ppm and maximum deviation 2.1 ppm over a 10 cm diameter spherical volume (DSV). The design had a peak field within the conductor of 1.9 T and a very conservative maximum hoop stress of 16 MPa. The fringe field constraint was satisfied, with the 5 gauss line lying entirely within a sphere of radius 3.1 m. Figure 6.3a shows a cross-section of the preliminary magnet design and Figure 6.4a shows the homogeneity over the ROI.

6.5.3 Comparison with manufacturer's design

At this point, the commissioned design study was nearing completion and it was possible to compare our preliminary results against a design from indus-

	Preliminary Results	Manufacturer Design
Field strength	1.0 T	1.0 T
Homogeneity (10 cm DSV)	2.1 ppm	1.2 ppm
Vol. of s/c (m ³)	0.07	0.058
Sensitivity (ppm/mm) over 10 cm DSV } }	463	244

Table 6.1: Comparison of the preliminary design with the manufacturer’s proposed design.

try. The configuration of the commissioned design is shown in Figure 6.3b and the homogeneity is shown in Figure 6.4. Both designs had a 1.0 T central field strength with a 20 cm inter-coil gap, an overall length of 1.5 m and outer bore diameter of 1.0 m. Remarkably, the two designs also displayed very similar homogeneity over a 10 cm DSV (see Table 6.1). The single major difference, though, between the designs was the operating current density. Whereas our design operated at only 33.3 A/mm², the design from industry was rated at 100 A/mm². This led to clear differences in the hoop stresses and in the volume of superconductor used. Interestingly, our design was based upon two negatively wound coils, one to actively shield the magnet and the other to improve the homogeneity, whereas the manufacturer’s design only had a single negatively wound coil.

During discussions about the various advantages of the two designs, it was suggested by the manufacturers that our design would be too sensitive to coil positioning and was therefore unbuildable. This objection was quite disturbing, as although sensitivity to build-errors had been considered previously, a Monte Carlo type calculation of design sensitivity during optimization was unrealistic because of the complexity of such a calculation. In order to overcome this issue, a rapid method of analytically approximating the coil sensitivity was created (see Section 4.4.2 for details).

Using the spherical harmonic weighting constants for the central field expansion, the novel analytical method obtains a value for the *overall relative sensitivity* of each design with minimal extra calculation. The speed of this technique is such that when tested, the evaluation of the sensitivity took

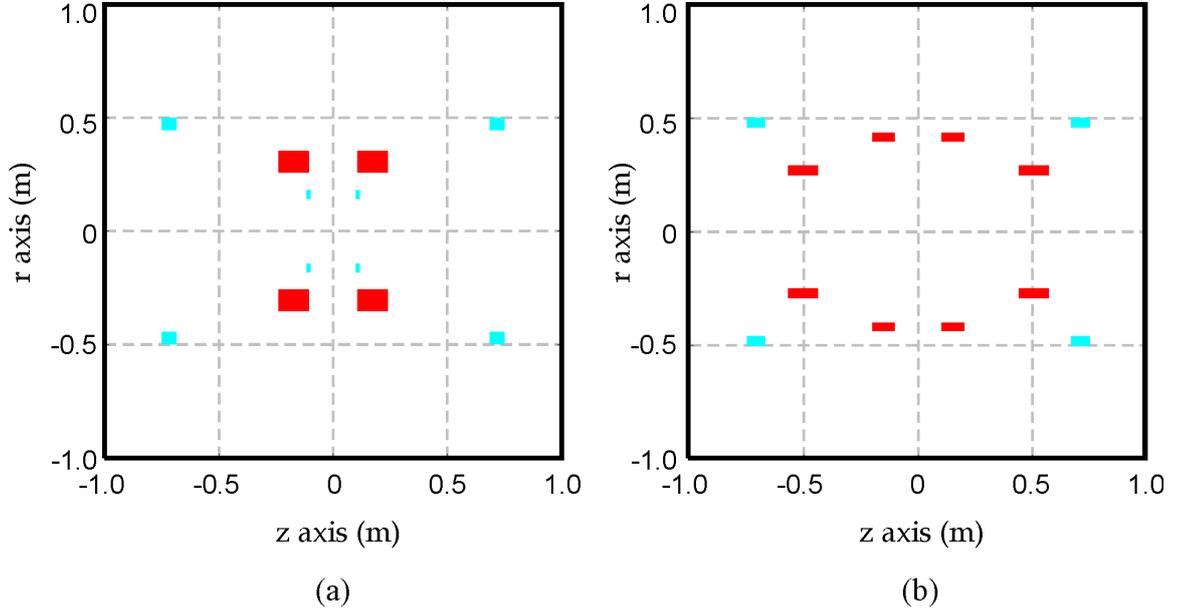


Figure 6.3: Comparison between (a) the preliminary design from the GA and (b) the commissioned design from Magnex Scientific Ltd.

up less than 0.5% of the overall fitness evaluation time. This technique was applied to both of the designs and the sensitivity over a 10 cm DSV was found to be 463 ppm/mm for our preliminary design and 244 ppm/mm for the manufacturer's design. The designs are briefly summarized and compared in Table 6.1. In the following section, the new rapid technique to evaluate coil sensitivity is added to the GA.

6.5.4 Including the sensitivity constraint

The measure of coil sensitivity was included in the GA by adding another constraint to the penalty function. In order to try to match the design from Magnex, the sensitivity limit was set to 80 ppm/mm per coil over a 10 cm DSV, thus

$$Penalty = Penalty + \tanh\left(\frac{\mathcal{S}_{\text{coil}} - \mathcal{S}_{\text{max}}}{2\sigma_S}\right) \quad (6.11)$$

where $\mathcal{S}_{\text{coil}}$ was the maximum sensitivity of all of the coils, \mathcal{S}_{max} was the maximum allowable sensitivity and σ_S was set to 10 ppm/mm. For consistency,

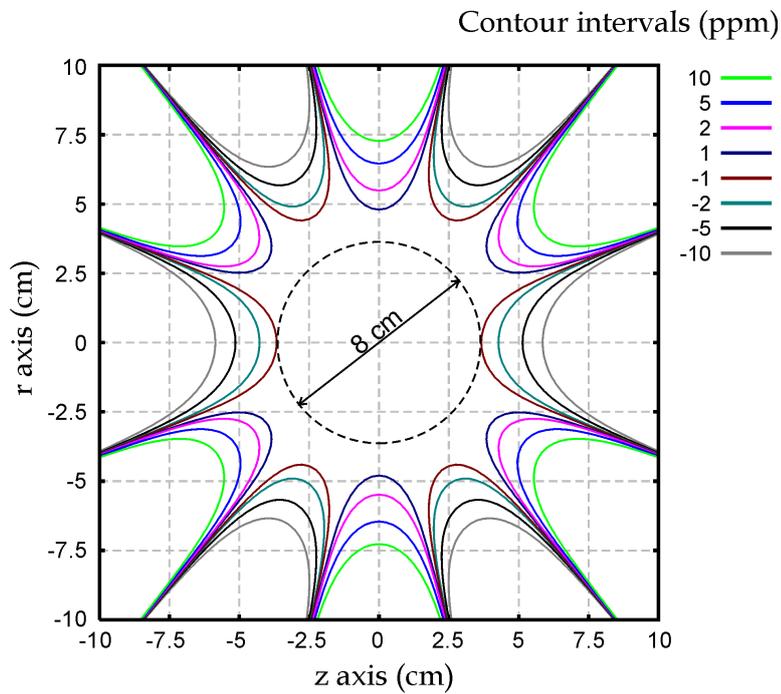
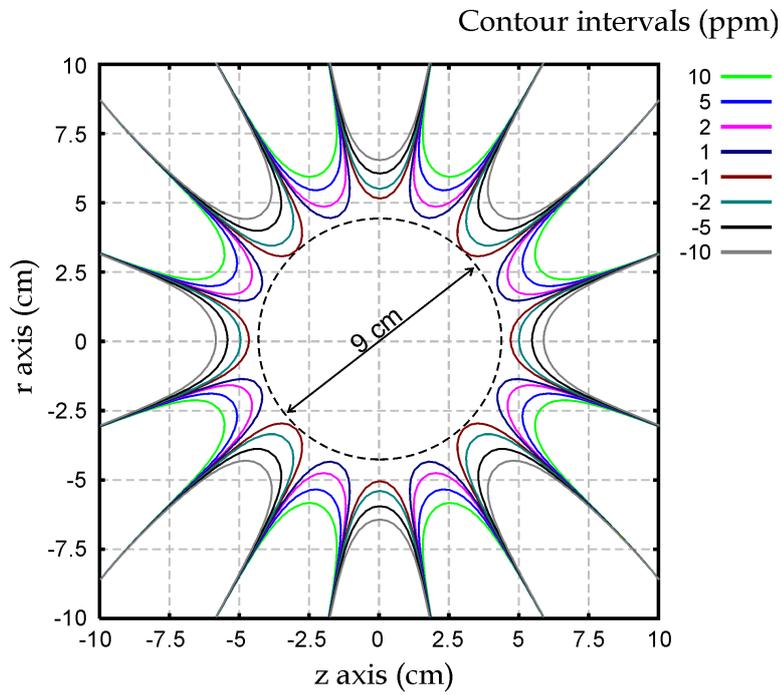


Figure 6.4: Contour plot of the homogeneity of (a) the preliminary design from the GA and (b) the commissioned design.

the operating current density was increased to 100 A/mm² and following discussions with the manufacturers, a few final design criteria were enforced

- Wire size was set to 1.5×1.5 mm.
- Number of layers was forced to be even.
- Coil depth was limited to 45 mm.
- Fringe field less than 5 gauss at 3.5 m.
- Inner bore diameter set to 50 cm.

The GA was run again for 10,000 generations using 128 demes each with a population of 128. Using the 16-node Linux cluster from the Wolfson Brain Imaging Centre, the computation time was approximately 3 hours. At the end of the run, the results were examined and the best design is presented in the next section.

6.6 Results and Analysis

With the inclusion of the sensitivity constraint, the GA produced a very promising final design which is presented and analyzed in this section. See Table 6.2 for the positions and current density of the coils in the final split-coil MRI magnet design.

Coil	z_{\min} (cm)	r_{\min} (cm)	width (cm)	depth (cm)	J'' (A/mm ²)
c1	10.48	42.06	8.55	4.50	100.00
c2	30.40	25.01	8.85	2.40	100.00
c3	65.36	50.84	8.10	3.30	-100.00

Table 6.2: Positions and current density of the coils in the final split-coil MRI magnet design.

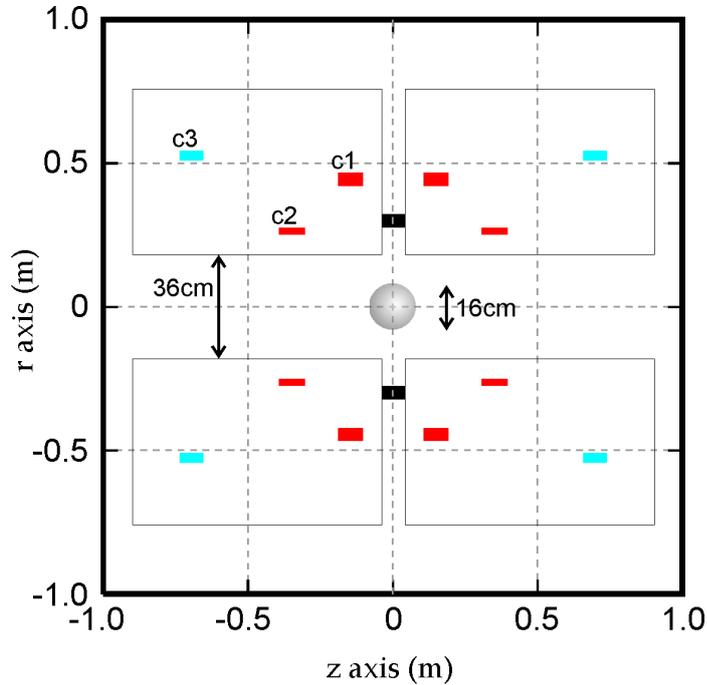


Figure 6.5: 2D poloidal cross-section of final split-coil magnet design. Positively wound coils are represented in red and negatively wound coils in blue. The ROI shown has a maximum inhomogeneity of < 1.15 ppm. The location of the cryostat and connecting struts are shown schematically (see Section 6.7 for details).

6.6.1 Overall configuration

The general configuration of this magnet is that of a split-coil MRI magnet with a gap between the two symmetrical halves of 20 cm. Containing only three pairs of superconducting coils, the magnet is a relatively simple design. Two pairs of coils are positively wound and produce a strong magnetic field parallel to the z axis, whilst a single pair of coils is negatively wound and actively shields the magnet design. A 2D poloidal cross-section of the magnet is shown in Figure 6.5. The superconducting coils are labelled c1-c3 and the operating magnetic field strength is 1.0 T.

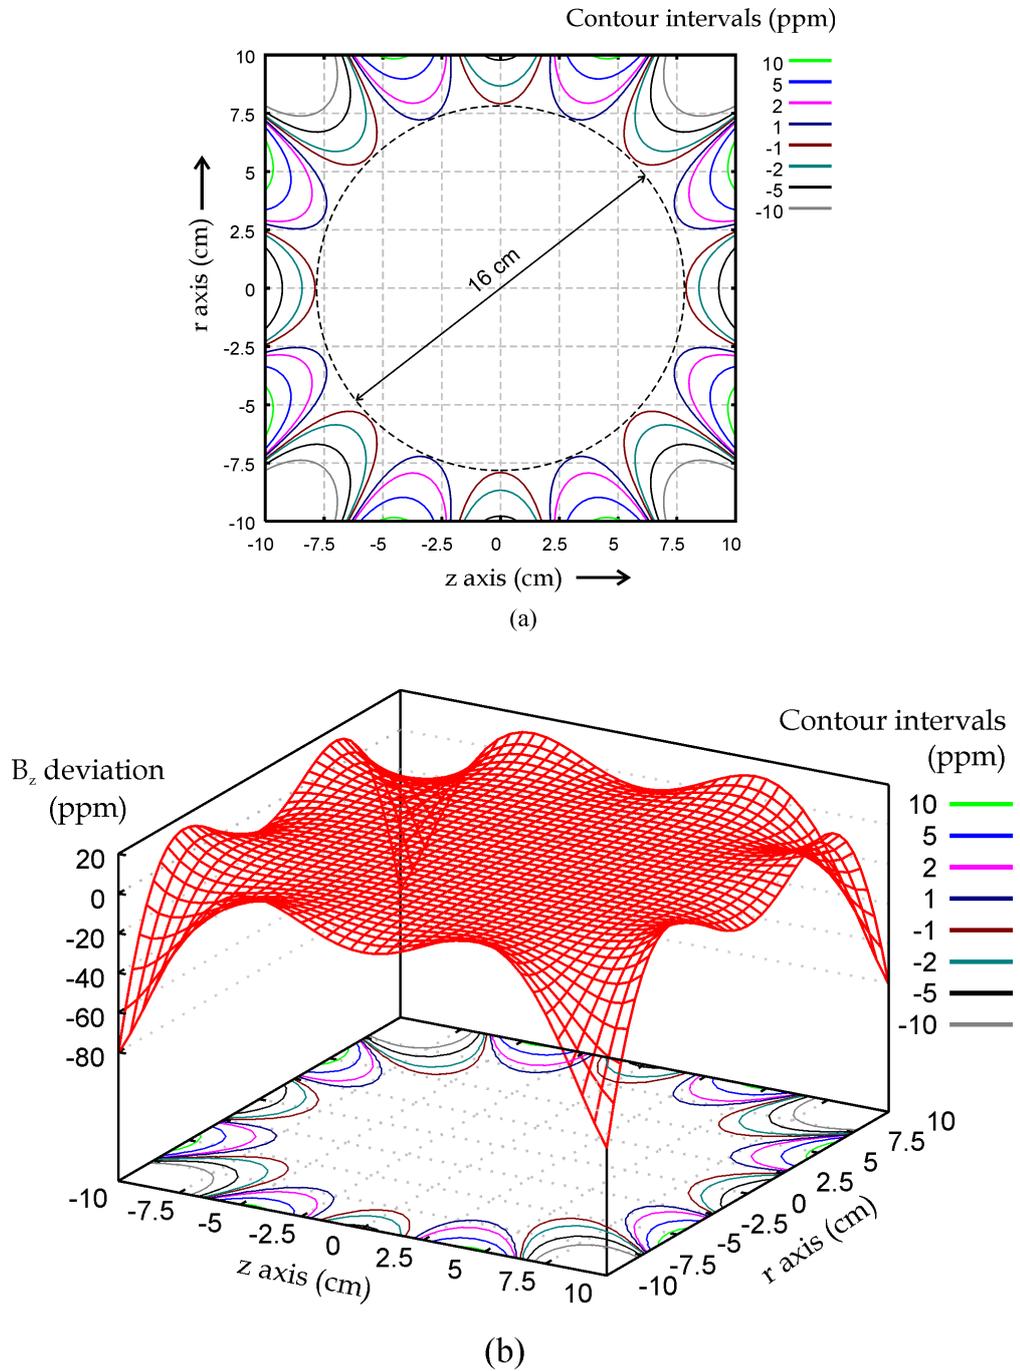


Figure 6.6: Homogeneity of the final split-coil design. (a) Contour plot and (b) surface plot of the z component of the magnetic field over the ROI showing the deviation from uniformity.

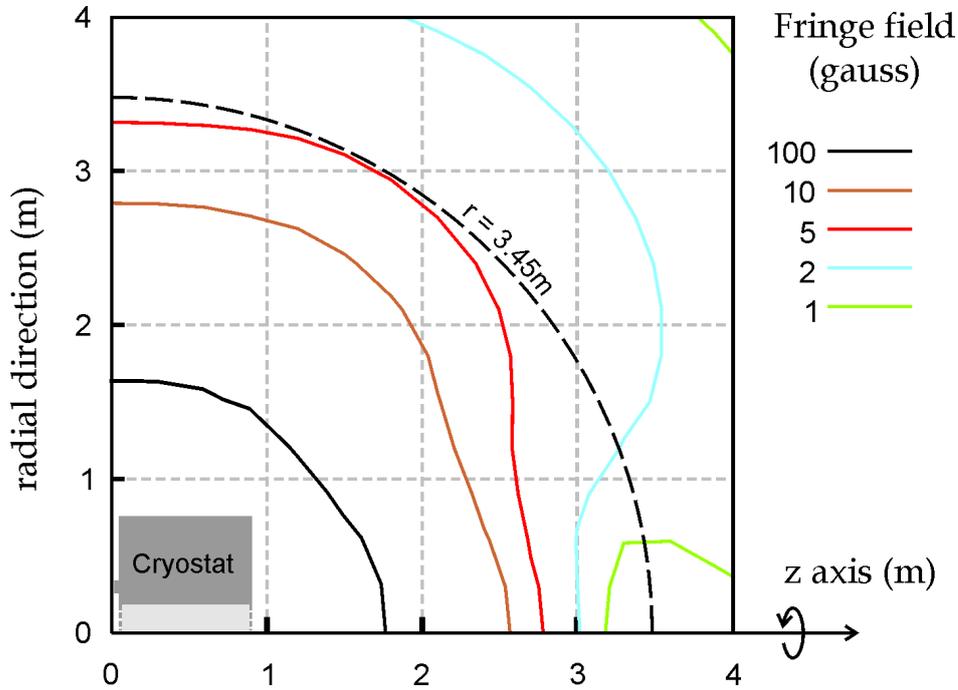


Figure 6.7: Magnetic far field region of the split-coil MRI magnet showing the 5 gauss contour line lying entirely within a sphere of radius 3.45 m.

6.6.2 Homogeneity

The deviation of the z component of the magnetic field from the central value, expressed in parts per million, is shown in Figure 6.6. Over a 10 cm DSV, the standard deviation of the magnetic field is 0.01 ppm. Extending the DSV to 16 cm and 20 cm, the standard deviation becomes 0.30 ppm and 1.67 ppm respectively. The maximum deviation from the central field strength is 0.05 ppm over a 10 cm DSV, 1.15 ppm over a 16 cm DSV and 9.92 ppm over a 20 cm DSV. It is very clear that the diameter of the ROI is about one and a half times as large as our preliminary design and also the manufacturer's design. This translates into a four-fold increase in homogeneous imaging volume.

6.6.3 Fringe fields

The split-coil MRI magnet design is actively shielded using a single pair of negatively wound coils to reduce the magnetic footprint. The 5 gauss line lies entirely within a sphere of radius 3.45 m. Axially, the field drops to less than 5 gauss at 2.76 m from the iso-centre and radially at 3.31 m. The 500 gauss line, which is important for the siting of the PMTs is located at 1.27 m along the z axis and 1.06 m radially. Details of the fringes are shown in Figure 6.7.

The compact nature of the fringe fields allows for easier siting of the magnet and reduces the possibility of interaction with surrounding magnetically active materials. The overall 5 gauss footprint of the magnet covers an area of approximately 31.5 m².

6.6.4 Peak fields

The peak fields were calculated for each of the three pairs of coils within the magnet design. The results in Table 6.3 show that the maximum field within the superconducting region is 2.4 T in coil c1, well within the intrinsic short-sample limits of standard NbTi superconductor at 4.2 K and 100 A/mm² (see Figure 4.8).

Coil	Peak Field (T)	Axial Force (kN)	Hoop Stress approx (MPa)	Shear Stress (MPa)
c1	2.4	191 (attractive)	101.7	0.802
c2	1.9	108 (attractive)	47.4	0.434
c3	1.8	101 (repulsive)	86.7	1.440

Table 6.3: Peak fields, axial forces and approximate hoop and shear stresses for the three pairs of coils in the split-coil MRI magnet design.

6.6.5 Stress and Forces

A finite element analysis of the stresses in the split-coil MRI magnet design was performed using Abaqus [63] (see Section 4.5.3). The percentage of epoxy and insulation was estimated as 10% of the overall wire cross-sectional

area ($\epsilon = 0.1$). It was found that coil c1 undergoes the largest hoop stress of 93 MPa, whilst the maximum shear stress of 1.4 MPa is located within coil c2. See Figures 6.8 and 6.9 for details. The hoop stresses quoted in Table 6.3 are approximated from the self-supporting hoop stress analysis (4.83) which also accounts for the slight difference. Both values are well within the intrinsic limits for a Cu/NbTi superconducting matrix.

Axial forces were found by numerical integration of \mathbf{F}_z over each individual annular coil. The results are shown in Table 6.3. Summing the forces shows that a net attractive force of 198 kN or almost 20 tons must be supported by the interconnecting rods between the two coil halves.

6.6.6 Sensitivity

The sensitivity of the design to build errors was evaluated using the two available methods. Firstly, a Monte Carlo simulation was performed using 0.5 mm positioning errors in both the axial location of the coils and the inner bore radius. Results are shown in Figure 6.10. The standard deviation of the inhomogeneity over a 10 cm DSV was 131 ppm on average, equivalent to 262 ppm/mm. Using the novel analytical approximation of the sensitivity, a value of 229 ppm/mm was found for the *overall relative sensitivity* for a 10 cm DSV.

6.6.7 Comparison with manufacturer's design

A comparison of the latest split-coil MRI magnet design against the manufacturer's suggested design is shown in Table 6.4. It can be seen that for almost all measures of performance, the stochastically optimized magnet design has outperformed the manufacturer's proposed design. Especially encouraging is the four-fold increase in the imaging volume with homogeneity < 1 ppm combined with a reduction of 20% in the volume of superconductor used in the design.

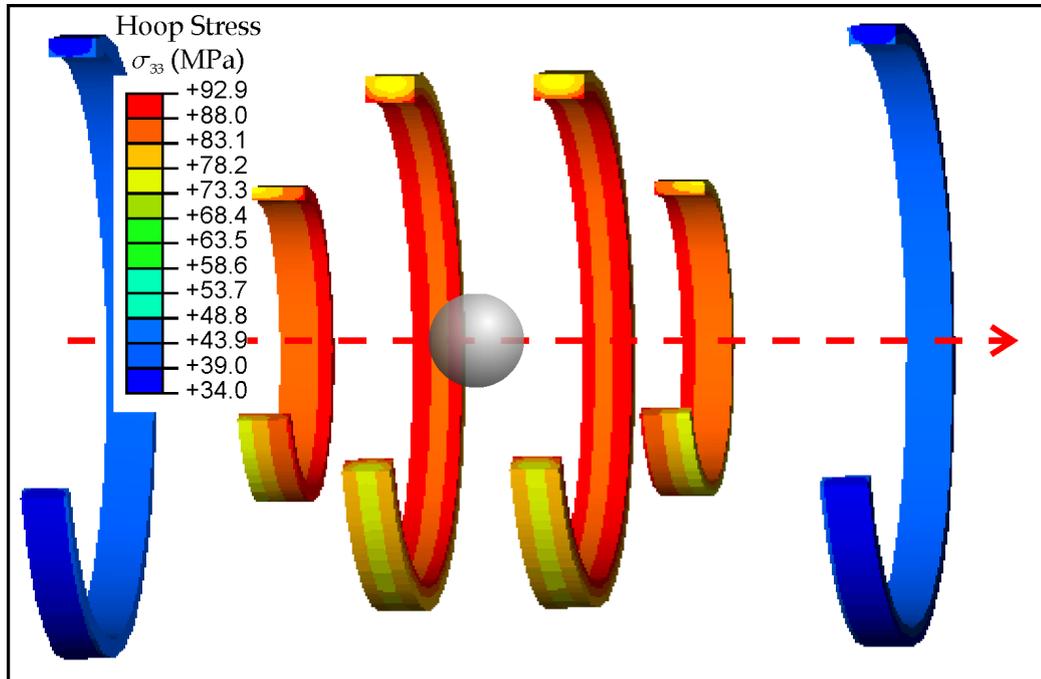


Figure 6.8: Analysis of the split-coil magnet for hoop stresses using Abaqus.

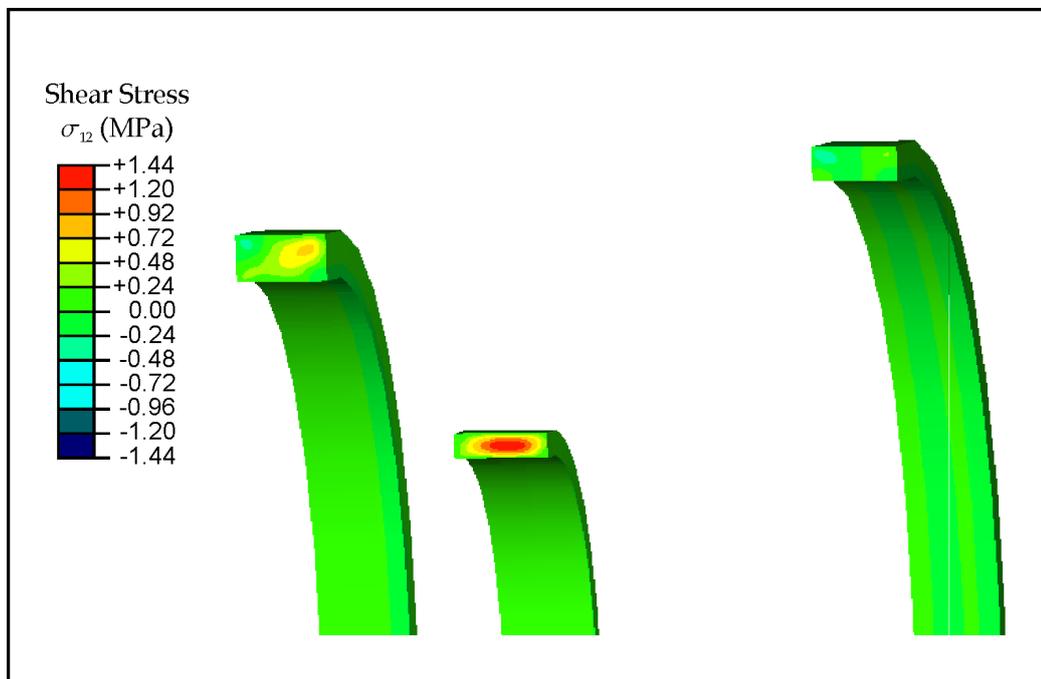


Figure 6.9: Analysis of the split-coil magnet for shear stresses using Abaqus.

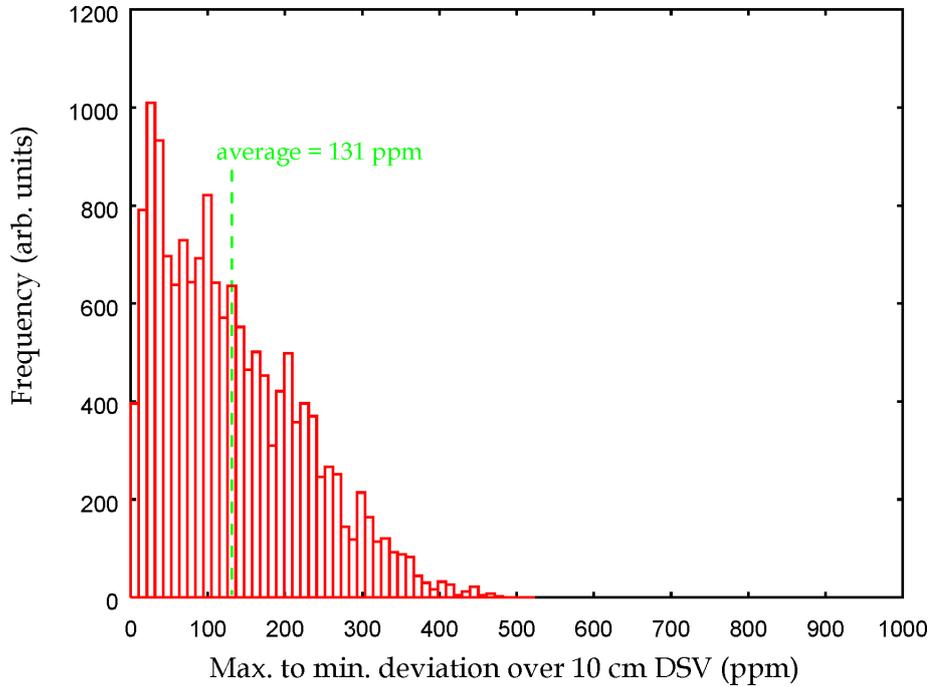


Figure 6.10: Sensitivity of split-coil MRI magnet design to positioning errors. In this case, 0.5 mm errors were simulated in the z position and the inner radius of each coil.

6.7 Discussion

Using an improved GA optimization, a split-coil MRI magnet design has been created which significantly outperforms a commercially produced design. It is a shame that the optimization methods used in industry are a closely guarded secret and therefore direct comparison of our method is impossible.

However, there are a few research groups that have published stochastic optimization techniques for MRI magnet design recently. As discussed in Chapter 5, Crozier *et al.* [89] have used a hybrid method combined with simulated annealing, but their work has focussed almost exclusively on short bore magnets. Cavaliere *et al.* [58, 91, 107] on the other hand have applied their stochastic optimization technique to human scale split-coil magnets. An example of this is given in [107], where Cavaliere presents a design for a human scale split-coil magnet with 1 ppm deviation over a 10 cm DSV,

Parameter	Manufacturer's design	GA design	Improved?
Field strength	1.0 T	1.0 T	
Operating current	225 A	225 A	
Wire size	1.5×1.5 mm	1.5×1.5 mm	
Inner bore ϕ	50 cm	50 cm	
Outer diameter ϕ	100 cm	100 cm	
Length	150 cm	150 cm	
Inhomogeneity	1.22 ppm (10 cm)	0.04 ppm (10 cm)	+
	21.0 ppm (16 cm)	1.15 ppm (16 cm)	+
Fringe fields	5 gauss at 3.70 m	5 gauss at 3.45 m	+
Sensitivity	249 ppm/mm	229 ppm/mm	+
Volume of s/c	0.058 m ³	0.046 m ³	+
Max hoop stress	98 MPa	102 MPa	
Peak field	2.7 T	2.4 T	+
Axial force	177 kN	214 kN	-

Table 6.4: Comparison of the design optimized by GAs and the design available from the manufacturers as part of a commissioned design study. The advantages and disadvantages are highlighted.

created using a real-coded GA. As part of the GA, a single fitness function value is made up from a measure of the homogeneity, the volume of superconductor and a factor called the ‘compactness’ of the design. Although the design offers a substantial gap between the coils of 24 cm and a homogeneous ROI of 10 cm, it is clearly unfeasible to build as it suffers from very large hoop stresses exceeding 900 MPa. Indeed one of the major advantages of the method presented in this chapter is that its ability to optimize designs with all of the necessary constraints, ranging from homogeneity and fringe fields to hoop stresses and even buildability.

Turning now to the absolute performance of the GA in this chapter, the inclusion of a new measure of the homogeneity and a new fringe field evaluation method improved the efficiency of the fitness function evaluation routine. This reduced the time spent on the evaluation of each design and allowed more of the search space to be covered and therefore better designs to be produced at the end of each run. Optimization in parallel with multiple demes was also particularly useful as it offered multiple solutions to the

problem and could be used to counteract the natural tendency of the GA to converge prematurely to local minima.

Altering the constraint handling method was also beneficial. In the previous chapter, large arbitrary parameters were used as penalties and added to the fitness of a design in violation of the relevant constraint. This created a delicate balance between the magnitude of the constraint penalties and the expected fitness value. Typically, several iterations were required to set the penalty factors to reasonable levels. In this chapter, the new constraint handling mechanism [17, 108] required little or no iteration to initialize as each penalty was measured using a normalizing function (see Equation 6.6). Setting the length-scale, σ_i is much easier as it can be taken as a rough measure of the violation of the constraint when the penalty grows to a value of 0.5.

In the introduction to this chapter and Section 6.2, the design of a split-coil magnet was motivated as an essential component of a combined PET and MRI dual-imaging system. Now that a theoretical design has been optimized, various additional components need to be considered in order to create the final combined PET/MRI system.

Firstly, the superconducting coils have to be enclosed in a cryostat. The dimensions of the cryostat were determined by the manufacturers who take into consideration the size of the coils and the requirements for good thermal shielding. The resulting structure reduces the intercoil separation from 20 cm to a room temperature gap of 8 cm. Secondly, there is a large force of about 20 tons pulling the two halves of the magnet together. In order to support this force, four struts (with 5 cm diameter) have been located between the two halves of the magnet at a radius of 30 cm. To achieve complete coverage by the PET detection system, the scintillation crystals will be placed close to the ROI at a radius of 20 cm. Using flexible optical fibres, the light produced will then be channelled around the struts to the PMTs which will be located at a radius of roughly 1.1 m, outside the 500 gauss field lines. Figure 6.11 illustrates the composition of the final combined PET/MRI system and includes the split gradient coils and RF coils for MRI.

In the literature, simultaneous PET/MRI has been achieved in both phan-

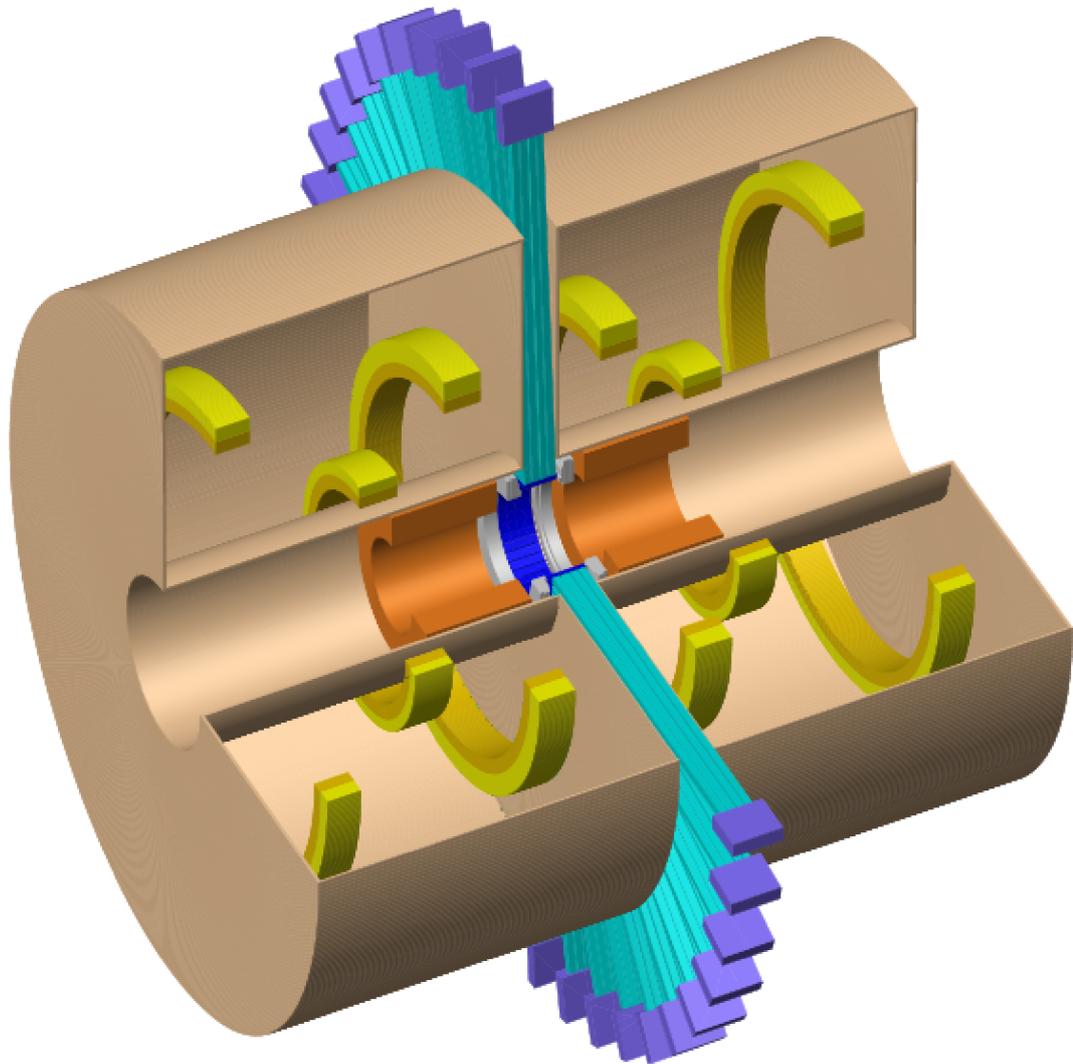


Figure 6.11: Cut-away view of the split-coil MRI magnet with the PET system installed. The scintillation crystals are linked via optical fibres to the PMTs which are located outside the 500 gauss magnetic field line. Courtesy of O. Barret.

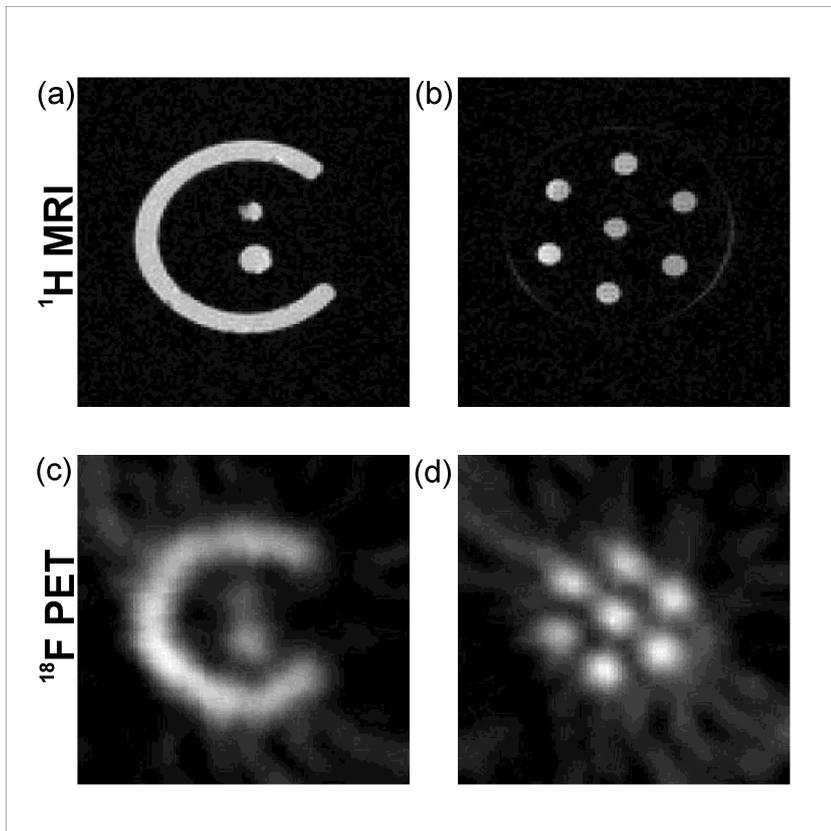


Figure 6.12: Simultaneous PET and MR images taken in 15 minutes. (a) and (b) show spin echo ($TE = 30$ ms, $TR = 2000$ ms) MR images of the two phantoms acquired on a 4.7 T, 30 cm bore system. Corresponding PET images are shown in (c) and (d). These images were reconstructed from 300,000 counts with 10 MBq activity in the phantoms. The PET resolution is roughly 2 mm FWHM [102].

toms [102] and isolated, perfused rat hearts [103]. An example of the remarkable images obtained is shown in Figure 6.12.

Using our novel geometry, it is hoped that even more exciting combined PET and MRI images may be obtained. Examples include molecular imaging of targeted receptor sites, the direct comparison of ^{18}F -Fluorodeoxyglucose-PET with functional MRI and deconvoluting PET data to reduce motion blurring and partial volume effects. Furthermore, by taking advantage of the improved PET system location, fully 3D PET data sets will become available for comparison to 3D MRI data.

6.8 Conclusions

Improvements and additions made to the GA optimization method have been described in this chapter, leading to a robust, homogeneous and feasible split-coil MRI magnet design. An exploratory evaluation of the search space, using simplified variables of cost and inter-coil gap was found to be useful in understanding the potential trade-offs. Using this information, a preliminary design was optimized and compared to a commercial design. Comparison showed encouraging results, but also that our design was highly sensitive to build errors. To avoid this problem, a novel analytical method was incorporated into the fitness evaluation function which calculated the design sensitivity during optimization. Using this novel sensitivity evaluation, a final design was optimized *ab initio* using the GA.

The final design presented in this chapter consists of a 1.0 T split-coil magnet with homogeneity over a 16 cm DSV of approximately 1 ppm (max. to min. deviation). As well as improving on the homogeneity of the imaging volume, the optimized design also satisfies all of the additional constraints placed on the magnet dimensions, the fringe fields, the stress and the sensitivity to build errors. The room temperature gap between the magnet halves is sufficiently wide to allow radial access to the ROI for an MRI compatible PET detector. The magnet system is due for delivery to the Cavendish Laboratory in 2003 and testing of simultaneous combined PET and MRI will begin shortly afterwards.

Chapter 7

Conclusions

A robust, simple genetic algorithm (GA) has been written. Whilst it can be applied to general optimization problems, this dissertation has focussed on its application to the design and optimization of axisymmetric MRI main-field magnets. Taking advantage of the recent development of parallel computer architectures, especially cluster computing, the GA has been written using Message Passing Interface (MPI), an industry standard protocol for communicating efficiently on homogeneous or heterogeneous networks. Two methods of parallelization have been considered. The first uses a single node of the cluster for the serial genetic processes such as crossover and mutations, whilst spreading the function evaluation step evenly across the available nodes. In the second method, entire demes are evaluated on separate nodes in parallel and occasional migration between demes is allowed (see Figure 3.7).

Separate to the GA optimization technique, a theoretical framework for the design of axisymmetric MRI magnets has also been presented. A significant part of this framework is the exposition of two new theoretical approaches. The first offers a rapid fringe-field evaluation method using a spherical harmonic expansion technique. The second is an analytical, rather than numerical method to determine the sensitivity of a design to build-errors. Both of these methods reduce the evaluation time significantly when compared to previous methods and therefore allow the associated constraints to be included in stochastic optimizations such as the GA. In passing, Gar-

rett's work from 1951 [38] on the central field expansion has been revisited and the formulae have been updated to SI units (see Appendix A).

The GA software has been applied to the design and optimization of novel MRI main-field magnets. In particular, two classes of magnet were examined. Firstly, in Chapter 5, the design of ultra short bore MRI magnets was investigated. Shortening the bore length in such magnets reduces claustrophobia for patients and improves access for clinicians during scanning which can be especially important when close supervision is required (e.g. critical head trauma imaging). Via a three-stage process, a particular ultra short bore design was proposed and fully evaluated. The design consists of a 1.0 T magnet with good homogeneity of less than 1 ppm max. to min. deviation over a 30 cm diameter spherical volume (DSV). The bore length of 1.0 m is significantly shorter than that in any other current design (> 1.25 m) and the aspect ratio of this magnet is close to 1:1. Critical currents and fringe fields are within the desired limits and hoop stresses have been considered using finite element analysis.

In Chapter 6, the design of an animal scale, split-coil MRI magnet was investigated. A novel constraint handling technique was included in the GA and faster evaluation of the fringe field allowed more of the search space to be considered. Preliminary results were examined, but it was found that the split-coil magnet suffered from high sensitivity to build errors. By introducing a direct measure of this sensitivity as an additional constraint in the fitness function, it was possible to find a less sensitive design with improved performance. This 1.0 T magnet has a coil inner-bore diameter of 50 cm and a length of 1.46 m. The design features a 20 cm gap between the coils so that the magnet can be used for simultaneous positron emission tomography (PET) and MRI. With the addition of a vacuum jacket and final trimmings, this gap will be reduced to 8 cm. Comparisons were made between this magnet and one resulting from a commercial design study. Our design was found to outperform the manufacturer's design in almost all aspects, especially the volume of imaging. Whereas the manufacturer's design exhibited an homogeneity¹ of 1 ppm over a 10 cm DSV, our magnet quadrupled

¹Standard deviation of z component of magnetic field.

the imaging volume with a 16 cm DSV maintaining the same homogeneity. This shows the power of a stochastic optimization method, especially when applied to problems where the overall optimum cannot be predetermined.

Although the results of the GA outperformed MRI magnet designs currently available, lack of convergence to a global optimum for these designs was slightly disappointing. If certain problems are addressed, it may well be that optimization using GAs holds even more promise for the future. One such issue is the encoding of the parameters which should be re-examined. As Goldberg says [16]

The user should select a coding so that short, low-order schemata are relevant to the underlying problem and relatively unrelated to schemata over other positions.

This point is crucial, but difficult to implement for the optimization of MRI main field magnets. The crux of the problem lies in the fact that if one of the parameters of a coil is altered, other parameters naturally need to compensate. Without this compensation, the optimization may quickly become entrapped in a local minimum whereby the alteration of any parameter will result in a design which does not perform as well. However, if two or three parameters were altered simultaneously then movement between regions of high fitness might become more probable.

Another issue, the combination of several objective functions into a single fitness, was found to be rather tricky especially as the weights (see Equation 3.4) of each objective could not be calculated in advance. Instead, in Chapters 5 and 6 a single objective function was chosen (homogeneity) whilst the other objectives were encoded as constraints. Recently, novel approaches to problems requiring the optimization of multiple objectives have been published [31, 32]. These approaches, called *multi-objective evolutionary algorithms* also act on a population of designs. However, rather than combining competing objectives into a single fitness, a term called non-dominance is introduced, which describes a solution for which no other solution exists which is better than it for all objectives. It may be that such an approach would be worthwhile attempting for MRI magnet design.

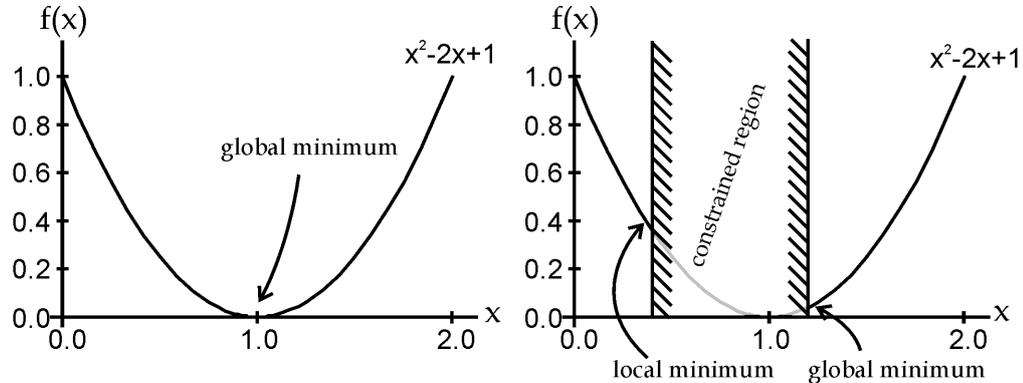


Figure 7.1: Example showing the interaction of constraints on the number and classification of minima in a simplified optimization problem.

Linked to the problem of not finding the global minimum, is the suggestion that a significant proportion of local minima were also not sampled. Evidence for this comes from the fact that stagnation within populations of solutions typically appeared towards the end of each optimization run and that the stagnated designs rarely matched when different starting configurations were used. This behaviour, indicative of a highly complex search space, is thought to be due to the large number of constraints that are imposed during optimization.

Taking a simple example, minimizing $f(x) = x^2 - 2x + 1$ in the range $x \rightarrow \{0, 2\}$, we see that only a single global minimum occurs at the point $x = 1$ where $f(x) = 0$. However, if we constrain x such that $|x - 0.8| \geq 0.4$, then the solutions alter significantly. A local minimum appears on one boundary of the constraint with $f(x) = 0.36$ at $x = 0.4$ and the global minimum moves to the point $x = 1.2$ where $f(x) = 0.04$ (see Figure 7.1). In the case of the main-field MRI magnet designs discussed in this thesis (Chapters 5 and 6), the many boundary conditions and constraints imposed on the optimization have rapidly produced a very complicated search space, even for relatively small numbers of parameters. It may well be that focussing the search on the boundary regions between feasible and unfeasible solutions [25] would yield better performance.

An issue not explicitly considered in this thesis is the inclusion of iron in MRI magnet designs to improve the homogeneity and passively shield high field systems. A new rapid method for the calculation of the magnetic field induced by iron has been recently presented by Zhao [109] and should allow the investigation of such designs. Another technique which has not been considered is that of grading the coils with different current densities depending on the magnetic field strength [60, p.23]. This not only reduces the possibility of exceeding the critical current limits, but also reduces the hoop stresses exerted on the magnet design, which is an important limiting factor. Inclusion of both these factors is strongly recommended for any future work in this field.

Whilst the ultra short bore magnet presented in Chapter 5 has many advantages over the current generation of short bore magnets, the cost of construction has so far prevented it from leaving the drawing board. However, the split-coil combined MRI and PET magnet, using the design in Chapter 6, is currently under construction and will be installed in the Cavendish laboratory during the second half of 2003. The radial access to the ROI offered by this magnet and its dedication to combined modality imaging research mean that the future is very exciting. Combination of PET and MRI may well lead to a better understanding of the fMRI results that are currently obtained. Certainly, whilst the physics behind PET response is well understood, the physics and physiological responses associated with fMRI remain the subject of intensive research. Furthermore, the removal of motion blurring of PET images with high temporal resolution MRI and almost perfect image co-registration should lead to imaging throughout the body at a molecular level.

As the power of MRI continues to advance, so the hardware necessary to support the clinicians must improve. In this thesis, the power of GAs applied to highly constrained real-world problems has been shown. The results have been very encouraging and the future is sure to bring even greater advances.

Appendix A

Magnetic field evaluation in the near-field using spherical harmonics

It is commonly known, and quoted, that the magnetic flux density field (\mathbf{B}) can be described near the origin of a system as an expansion of weighted spherical harmonic functions. For the general case this expansion contains both tesseral and zonal harmonics. However, for problems which are constrained to axisymmetrical geometries, the tesseral harmonics can be ignored and the magnetic field expressed as

$$B_z = \sum_{n=0}^{\infty} q_n \left(\frac{r}{r_0}\right)^n P_n(\cos \theta) \quad (\text{A.1})$$

$$B_x = \sum_{n=0}^{\infty} \frac{1}{(n+1)} q_n \left(\frac{r}{r_0}\right)^n P_n^1(\cos \theta) \quad (\text{A.2})$$

where the magnetic field is to be evaluated at the point (r, θ) , r_0 is an arbitrary radius (usually the radius of the region of interest) and q_n are the spherical harmonic weighting constants. There is an unfortunate mixing of coordinate systems in (A.1) and (A.2) with locations described in spherical coordinates and the magnetic field vector described in Cartesian coordinates. However, this is inevitable as it is common practice to concentrate on the

z component of the magnetic field in the region of interest rather than the radial component. Indeed for most homogeneous regions of interest, the approximation $|\mathbf{B}| \approx B_z$ is perfectly valid.

This appendix follows the mathematical treatment given by Garrett [38] to calculate the ‘source’ constants (q_n) for a given geometry of thick solenoids. However, the units have been updated to S.I. units and the source constants, which in [38] are dimensionless, are given here with dimensions of magnetic flux density (Telsa). Both representations require the inclusion of an arbitrary radius, denoted by r_0 , but by forcing the source constants to have dimensions of magnetic flux density, we can directly interpret them as a measure of the field inhomogeneity. Thus, for a design with symmetry in the plane $z = 0$,

$$\text{Central field strength} = q_0 \quad (\text{A.3})$$

$$\text{Deviation (in T) at } (r_0, 0) = \sum_{n=1}^{\infty} q_n \quad (\text{A.4})$$

It is therefore clear that taking r_0 as the radius of the region of interest is sensible. Furthermore, for a system with symmetry in the plane $z = 0$, the off-axis deviation at the edge of the ROI is given by

$$\begin{aligned} \text{Deviation (in T) at } (r_0, \pi/2) &= \sum_{n=2, \text{ even}}^{\infty} q_n P_n(0) \\ &= -\frac{1}{2}q_2 + \frac{3}{8}q_4 - \frac{5}{16}q_6 + \dots \end{aligned} \quad (\text{A.5})$$

Notice the close links between this and the far-field spherical harmonic expansion as calculated in Section 4.3.1. The remainder of this appendix gives the remodelled source constants for an infinitely thin loop, an infinitely thin cylinder and a thick solenoid.

For an infinitely thin loop, the source constants, $q_n(\text{f})$, are given by:

$$q_n(\text{filament}) = \frac{\mu_0 I}{2} \sin^2 \theta' \frac{r_0^n}{(r')^{n+1}} P_{n+1}'(\cos \theta') \quad (\text{A.6})$$

where the position of the filament is defined by the primed variables (r', θ') and the current is given by I . $P_m'(u)$ is the first derivative of $P_m(u)$ w.r.t. u and can be related to the associated Legendre polynomials using¹

$$P_m^1(u) = -\sqrt{1-u^2}P_m'(u) \quad (\text{A.7})$$

Replacing I by $J' dz'$ and integrating along z' , we find the source constants for thin cylindrical systems

$$q_0(\text{thin cylinder}) = \frac{\mu_0 J'}{2} \cos \theta' \quad (\text{A.8})$$

$$(n \geq 1), \quad q_n(\text{thin cylinder}) = -\frac{\mu_0 J'}{2n} \left(\frac{r_0}{r'}\right)^n \sin^2 \theta' P_n'(\cos \theta') \quad (\text{A.9})$$

where (A.8) and (A.9) must be evaluated at the limits of the cylinder. Finally by integrating $q_n(\text{thin cylinder})$ radially, the source constants are found for thick solenoidal coil systems [38]:

$$q_0(\text{solenoid}) = -\frac{\mu_0 J_0''}{2} z' [1 - \ln(r' + x')] \quad (\text{A.10})$$

$$q_1(\text{solenoid}) = \frac{\mu_0 J_0''}{2} r_0 [\sin \theta' - \ln(r' + x')] \quad (\text{A.11})$$

$$(n \geq 2), \quad q_n(\text{solenoid}) = \frac{\mu_0 J_0''}{2n(n-1)} \frac{r_0^n}{(z')^{n-1}} s_{n-1}(\cos \theta') \quad (\text{A.12})$$

where

$$s_{n-1}(u) = 1 - \sin^3 \theta' \sum_{m=1}^{n-1} u^{m-1} P_m'(u) \quad (\text{A.13})$$

Again, the limits of the integration require that (A.10)-(A.13) be evaluated at each of the four corners of the thick solenoid and the results combined with signs as shown in Figure 4.4. An observant reader who compares the above equations to [38] will notice that there are several slight alterations to the original set of equations. Specifically, an extra factor of $-\frac{\mu_0 J_0'' z'}{2}$ appears

¹The sign in (A.7) can be either positive or negative depending on the precise definition of $P_l^m(u)$. This appendix follows the convention used by Mathematica [110] rather than Belousov [111]. For more details see Abramowitz and Stegun [54]

in (A.10) and the logarithmic arguments have been altered from $\ln(\frac{1+\sin\theta'}{\cos\theta'})$ to $\ln(r' + x')$. These are cosmetic changes and alter only the arbitrary constants of integration, which arise from $\int q_n(\text{thin cylinder}) dx'$. To confirm the validity of (A.10)-(A.13), we can check that they all satisfy the much more rigorous condition

$$\frac{\partial q_n(\text{solenoid})}{\partial x'} = \frac{J''}{J'} q_n(c) \quad (\text{A.14})$$

The advantage of altering the form of these equations is to allow the weighting coefficients to be used as a basis for finding the relative coil sensitivities as proposed in Section 15.

Appendix B

Partial differentials of near-field weighting constants

To find the partial differentials of the near-field weighting constants, several standard results [54, 110, 111] involving the associated Legendre polynomials are required. The following are used in this appendix

$$(2n + 3)\mu P_{n+1}^1(\mu) = (n + 2)P_n^1(\mu) + (n + 1)P_{n+2}^1(\mu) \quad (\text{B.1})$$

$$(\mu^2 - 1)\frac{dP_{n+1}^1(\mu)}{d\mu} = (n + 2)P_n^1(\mu) - (n + 1)\mu P_{n+1}^1(\mu) \quad (\text{B.2})$$

$$(1 - \mu^2)^{1/2}\frac{dP_{n+1}^1(\mu)}{d\theta} = (n + 1)\mu P_{n+1}^1(\mu) - (n + 2)P_n^1(\mu) \quad (\text{B.3})$$

and

$$(l - m)P_l^m(\mu) = (2l - 1)\mu P_{l-1}^m(\mu) - (l + m - 1)P_{l-2}^m(\mu)$$

For example

$$(n + 1)P_{n+2}^1(\mu) = (2n + 3)\mu P_{n+1}^1(\mu) - (n + 2)P_n^1(\mu) \quad (\text{B.4})$$

$$(n + 2)P_{n+3}^1(\mu) = (2n + 5)\mu P_{n+2}^1(\mu) - (n + 3)P_{n+1}^1(\mu) \quad (\text{B.5})$$

Taking the formula for $q_n(f)$ (A.6) and dropping the primed notation for simplicity, we write

$$q_n(f) = -\frac{\mu_0 I r_0^n \sin \theta}{2} \frac{1}{r^{n+1}} P_{n+1}^1(\mu) \quad (\text{B.6})$$

where $\mu = \cos(\theta)$ and the Cartesian and spherical coordinates are defined in the usual manner¹. The partial differential of $q_n(f)$ w.r.t. z may be calculated using the chain rule

$$\frac{\partial q_n(f)}{\partial z} = -\frac{\mu_0 I r_0^n}{2} \frac{\partial}{\partial z} \left\{ \frac{x}{(x^2 + z^2)^{\frac{n+2}{2}}} P_{n+1}^1(\mu) \right\} \quad (\text{B.7})$$

$$= -\frac{\mu_0 I r_0^n}{2} \left[-\frac{(n+2)xz}{(x^2 + z^2)^{\frac{n+4}{2}}} P_{n+1}^1(\mu) \right. \quad (\text{B.8})$$

$$\left. + \frac{x}{(x^2 + z^2)^{\frac{n+2}{2}}} \cdot \frac{x^2}{(x^2 + z^2)^{\frac{3}{2}}} \cdot \frac{dP_{n+1}^1(\mu)}{d\mu} \right] \quad (\text{B.9})$$

since

$$\left. \frac{\partial \mu}{\partial z} \right|_x = \frac{x^2}{(x^2 + z^2)^{3/2}} \quad (\text{B.10})$$

Applying (B.2) and simplifying with (B.1) we obtain

$$\frac{\partial q_n(f)}{\partial z} = -\frac{\mu_0 I r_0^n}{2} \left[-\frac{(2n+3)x}{r^{n+3}} \mu P_{n+1}^1(\mu) + \frac{(n+2)x}{r^{n+3}} P_n^1(\mu) \right] \quad (\text{B.11})$$

$$= \frac{\mu_0 I r_0^n}{2} \cdot \frac{x}{r^{n+3}} \cdot (n+1) P_{n+2}^1(\mu) \quad (\text{B.12})$$

Thus

$$\frac{\partial q_n(f)}{\partial z'} = -\frac{(n+1)}{r_0} q_{n+1}(f) \quad (\text{B.13})$$

where the primed variables, representing locations of the current sources, have been reintroduced.

Next, we need to find the differential of $P_{n+1}^1(\mu)$ w.r.t. x . Using the chain rule again, we note

$$\frac{\partial P_{n+1}^1(\mu)}{\partial x} = \frac{dP_{n+1}^1(\mu)}{d\theta} \cdot \frac{\partial \theta}{\partial x} = \frac{z}{x^2 + z^2} \cdot \frac{dP_{n+1}^1(\mu)}{d\theta} \quad (\text{B.14})$$

into which (B.3) may be substituted.

¹See Figure 4.2 for details.

Furthermore,

$$\frac{\partial q_n(f)}{\partial x} = -\frac{\mu_0 I r_0^n}{2} \left\{ \frac{\sin \theta}{r^{n+1}} \cdot \frac{z}{r^2} \left[\frac{(n+1)\mu P_{n+1}^1(\mu)}{\sin \theta} - \frac{(n+2)P_n^1(\mu)}{\sin \theta} \right] + \frac{z^2 - (n+1)x^2}{r^{n+4}} \cdot P_{n+1}^1(\mu) \right\} \quad (\text{B.15})$$

Substituting $x^2 = r^2 - z^2$ and $\mu = z/r$ we recover

$$-\frac{\mu_0 I r_0^n}{2} \left\{ \frac{(2n+3)\mu^2}{r^{n+2}} P_{n+1}^1(\mu) - \frac{(n+2)\mu}{r^{n+2}} P_n^1(\mu) - \frac{(n+1)}{r^{n+2}} P_{n+1}^1(\mu) \right\} \quad (\text{B.16})$$

The first two terms in (B.16) can now be combined with (B.4) to give

$$-\frac{\mu_0 I r_0^n}{2} \left\{ \frac{(n+1)\mu P_{n+2}^1(\mu) - (n+1)P_{n+1}^1(\mu)}{r^{n+2}} \right\} \quad (\text{B.17})$$

before substituting (B.5) to remove any dependence on $P_{n+2}^1(\mu)$

$$-\frac{\mu_0 I r_0^n}{2} \cdot \frac{(n+1)}{r^{n+2}} \left\{ \frac{(n+2)}{(2n+5)} P_{n+3}^1(\mu) + \left[\frac{(n+3)}{(2n+5)} - 1 \right] P_{n+1}^1(\mu) \right\} \quad (\text{B.18})$$

which simplifies to

$$-\frac{\mu_0 I r_0^n}{2} \cdot \frac{(n+1)(n+2)}{(2n+5)} \left\{ \frac{P_{n+3}^1(\mu) - P_{n+1}^1(\mu)}{r^{n+2}} \right\} \quad (\text{B.19})$$

This equation (B.19) may finally be rewritten, as required, in terms of $q_{n+2}(f)$ and $q_n(f)$,

$$\frac{\partial q_n(f)}{\partial x'} = -\frac{(n+1)(n+2)}{(2n+5)} \left[\frac{(r')^2}{x' r_0^2} q_{n+2}(f) - \frac{1}{x'} q_n(f) \right] \quad (\text{B.20})$$

where the primed notation has been reintroduced as before.

Appendix C

Post-processing software

This appendix contains details of supplementary software which was written to analyse the magnet designs created by the genetic algorithm code used in this thesis. The software is divided into various components, each of which performs a single task including automated evaluation of axisymmetric magnet designs, shimming of designs, analysis of the sensitivity to build errors and conversions to other formats (e.g. for finite element analysis). These four programs are discussed below.

C.1 Detailed axisymmetric design evaluation

This program processes axisymmetric MRI magnet designs directly from the output of the GA software. It calculates central field strength, homogeneity and fringe field magnitude, as well as the spherical harmonic expansion coefficients (both near and far-field). Furthermore, it approximates the peak field strength within the coils, the peak hoop stresses and the sensitivity of the design to build-errors in the z direction. Physical factors such as the volume of superconductor and wire length are also evaluated. The input file is text only and begins with a header section detailing important information about the overall design. The program extracts these variables using key-phrases as markers. A list of key-phrases is given in Table C.1 along with the default values which are used if the key-phrase is missing. The design itself is encoded using one line for each coil (or pair of coils). Coils are represented

Appendix C

Key-phrase	Description	Default (units)
MRI.length <double>	Overall length of magnet	1.5 m
MRI.height <double>	Overall height of magnet	1.5 m
ROI.radius <double>	Radius of region of interest	0.2 m
R.nought (for expansion) <double>	Typical radius r_0	0.1 m
Required central field <double>	Expected field strength	1.0 T
Fringe field limit <double>	Expressed in Teslas	0.0005 T
Fringe field distance <double>	Radius of the fringe field sphere	5.0 m
Number of coils <int>	<i>Pairs</i> of coils if symmetric	1
Mid-plane symmetry <int>	Boolean expression, 1=true, 0=false	1
Wire width <double>		0.001 m
Wire height <double>		0.001 m
Maximum allowable wire volume <double>	Expressed in m^3	0.1 m^3
Maximum relative sensitivity <double>	Only w.r.t. z movements	100.0 ppm/mm

Table C.1: Key-phrases, descriptions and default values used in the MRI magnet design evaluation software.

by 5 parameters: left edge of coil (m), inner radius of coil (m), width (m), depth (m) and current in a single strand of wire (A).

The software calculates the magnetic field for the complete design and useful figures are automatically created (see Figure C.1). These include a 2D cross-section of the magnet configuration, a surface plot of the field inhomogeneities over the region of interest, plots of the field strength along the axes and a contour plot of the field strength in the far-field region.

C.2 Autoshim

Autoshim calculates the shim currents required to cancel the first $(n - 1)$ non-zero spherical harmonic terms, using n pairs of shim coils. This program requires a file containing the magnet design parameters and a separate file detailing the locations of the shim coils. A matrix inversion method is implemented, as discussed in Section 5.5.2, which maintains the central field strength and calculates the required currents for the shim coils.

C.3 Sensitivity to builderrors

A Monte Carlo analysis is performed to simulate the effects of manufacturing errors by moving the positions of coils in an optimized design. Typically two error modes are examined: accuracy along the z axis and accuracy of the

coil inner radius. Shim coils are subsequently used in an attempt to restore the homogeneity and graphs are automatically created which summarize the results (see Figure 5.11).

C.4 Conversion routines

Two conversion routines have been created. One routine converts the parameters of an MRI magnet design to a suitable format for the finite element program FEMM [46]. The other produces a model of the magnet design suitable for ABAQUS [63] as well as a database of forces which is used in the finite element stress analysis. To apply forces to the model in ABAQUS, a user subroutine, ‘*dload.f*’, was developed in Fortran.

Bibliography

- [1] W. Gerlach and O. Stern, *Ann. Phys.* **74**, 673 (1924). [6](#)
- [2] E. M. Purcell, H. C. Torrey, and R. V. Pound, *Physical Review* **98**, 1787 (1946). [6](#)
- [3] F. Bloch, W. W. Hansen, and M. Packard, *Physical Review* **69**, 127 (1946). [6](#)
- [4] P. Lauterbur, *Nature* **242**, 190 (1973). [6](#)
- [5] A. Kumar, D. Welte, and R. R. Ernst, *Journal of Magnetic Resonance* **18**, 69 (1975). [6](#)
- [6] P. Mansfield, *J. Phys. C: Solid State Phys.* **10**, L55 (1977). [6](#)
- [7] Z.-P. Liang, *Principles of Magnetic Resonance Imaging: a Signal Processing Perspective* (Bellingham, Wash.: SPIE Optical Engineering Press, 2000). [7](#), [9](#), [13](#), [14](#), [16](#), [19](#)
- [8] P. Suetens, *Fundamentals of Medical Imaging*, 1st ed. (Cambridge University Press, 2002). [7](#)
- [9] B. Cowan, *Nuclear Magnetic Resonance and Relaxation* (Cambridge University Press, Cambridge, UK, 1997). [9](#), [19](#)
- [10] D. D. Stark and W. G. J. Bradley, editors, *Magnetic Resonance Imaging*, 3 ed. (Mosby, St. Louis, 1999). [12](#)
- [11] R. F. Farr and P. J. Allisy-Roberts, *Physics for Medical Imaging* (WB Saunders, republished 2000). [15](#), [116](#)

- [12] D. Sodickson and W. Manning, *Magnetic Resonance in Medicine* **38**, 591 (1997). [16](#)
- [13] K. Pruessmann, M. Weiger, M. Scheidegger, and P. Boesiger, *Magnetic Resonance in Medicine* **42**, 952 (1999). [16](#)
- [14] M. Agic *et al.*, Fatal imaging and spectroscopy at zero tesla, in *Proceedings of the International Society for Magnetic Resonance in Medicine* Vol. 2, p. 1387, Denver, Colorado, 2000. [16](#)
- [15] M. A. Brown and R. C. Semelka, *MRI: Basic Principles and Applications*, 2nd ed. (Wiley-Liss, 1999). [20](#), [22](#), [23](#)
- [16] D. E. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning* (Addison-Wesley, 1989). [27](#), [32](#), [35](#), [37](#), [39](#), [40](#), [42](#), [147](#)
- [17] K. Deb, *Sadhana* **24**, 293 (1999). [28](#), [38](#), [124](#), [141](#)
- [18] B. Fisher, *Gradient Design by Genetic Algorithm for Magnetic Resonance Imaging*, PhD thesis, Herchel Smith Laboratory for Medicinal Chemistry, University of Cambridge, 1996. [28](#), [37](#), [47](#), [124](#)
- [19] G. Williams, *Novel Geometry Gradient Coils for MRI Designed by Genetic Algorithms*, PhD thesis, Herchel Smith Laboratory for Medicinal Chemistry, University of Cambridge, 2001. [28](#), [37](#), [47](#), [124](#)
- [20] J. H. Holland, *Hierarchical Descriptions of Universal Spaces and Adaptive Systems* (Ann Arbor, University of Michigan, 1968). [31](#)
- [21] J. H. Holland, *Adaptation in Natural and Artificial Systems* (Ann Arbor, University of Michigan, 1975). [31](#), [41](#)
- [22] L. J. Eshelman and J. D. Schaffer, *Foundations of Genetic Algorithms* Vol. 2 (Morgan Kaufman, San Mateo, 1993). [31](#), [33](#)
- [23] C. Coello, *Computer Methods in Applied Mechanics and Engineering* **191**, 1245 (2002). [37](#), [38](#)

-
- [24] Z. Michalewicz, A survey of constraint handling techniques in evolutionary computation methods, in *Proceedings of 4th Annual Conference on Evolutionary Programming*, pp. 135–155, MIT Press, Cambridge, MA, 1995. 37
- [25] M. Schoenauer and Z. Michalewicz, *Lecture Notes in Computer Science* **1141**, 245 (1996). 37, 148
- [26] V. Kureichik, *Journal of Computer and Systems Sciences International* **38**, 137 (1999). 39
- [27] L. D. Chambers, editor, *Practical Handbook of Genetic Algorithms: Applications, Volume 1* (CRC Press, Inc., 1995). 39, 40
- [28] K. DeJong, *An Analysis of the Behaviour of a Class of Genetic Adaptive Systems*, PhD thesis, University of Michigan, 1975. 40
- [29] L. D. Chambers, editor, *Practical Handbook of Genetic Algorithms: New Frontiers, Volume 2* (CRC Press, Inc., 1995). 40, 43
- [30] K. Deb and D. Goldberg, An investigation of niche and species formation in genetic function optimization, in *Proceedings of the Third International Conference on Genetic Algorithms*, edited by J. D. Schaffer *et al.*, pp. 42–50, San Mateo, CA, 1989, Morgan Kaufmann Publishers. 40
- [31] E. Zitzler, K. Deb, and L. Thiele, *Evolutionary Computation* **8**, 173 (2000). 40, 147
- [32] K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan, A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II, in *Proceedings of the Parallel Problem Solving from Nature VI Conference*, pp. 849–858, Paris, France, 2000, <http://www.iitk.ac.in/kangal/pub.htm>. 40, 147
- [33] J. H. Holland, *Scientific American* **267**, 66 (1992). 41
- [34] E. Cantu-Paz, IlliGAL Report No. 97003, 1997 (unpublished). 42

Bibliography

- [35] B. H. V. Topping and J. P. B. Leite, *Engineering Optimization* **31**, 65 (1998). [42](#)
- [36] L. D. Chambers, editor, *Practical Handbook of Genetic Algorithms: Complex Coding Systems, Volume 3* (CRC Press, Inc., 1999). [43](#)
- [37] G. DeMeester *et al.*, *Magnetic Resonance Materials in Physics, Biology and Medicine* **13**, 193 (2002). [46](#), [85](#), [86](#), [111](#)
- [38] M. W. Garrett, *Journal of Applied Physics* **22**, 1091 (1951). [46](#), [50](#), [52](#), [56](#), [61](#), [64](#), [146](#), [152](#), [153](#)
- [39] F. Romeo and D. Hoult, *Magnetic Resonance in Medicine* **1**, 44 (1984). [46](#), [50](#)
- [40] L. K. Forbes, S. Crozier, and D. M. Doddrell, *IEEE Transactions on Magnetics* **33**, 4405 (1997). [46](#), [49](#)
- [41] J. D. Jackson, *Classical Electrodynamics*, 3rd ed. (J. Wiley & Sons, 1999). [47](#), [53](#), [54](#), [71](#), [82](#)
- [42] P. Hanley, Personal communication, 1999. [48](#)
- [43] R. Bulirsch, *Numerische Mathematik* **13**, 305 (1969). [49](#)
- [44] N. R. Shaw and R. E. Ansorge, *IEEE Transactions on Applied Superconductivity* **12**, 733 (2002). [49](#)
- [45] T. Tominaka, K. Hatanaka, and T. Katayama, *Proceedings of the 1997 Particle Accelerator Conference* **3**, 3434 (1998). [49](#)
- [46] D. Meeker, Finite element method magnetics, <http://femm.berlios.de/>. [51](#), [161](#)
- [47] Vector Fields Ltd: OPERA, <http://www.vectorfields.co.uk/>. [51](#)
- [48] G. V. Brown and L. Flax, *Journal of Applied Physics* **35**, 1764 (1964). [51](#)

-
- [49] W. H.-G. Müller, *IEEE Transactions on Applied Superconductivity* **10**, 752 (2000). [51](#)
- [50] F. Davies, G. DeMeester, D. Melotte, M. Morich, and A. Thomas, Short magnets for whole-body magnetic resonance imaging, in *Proceedings of the 9th ISMRM International Conference*, p. 1149, 2001. [51](#)
- [51] J. W. Burgoyne, P. D. Daniels, K. W. Timms, and S. H. Vale, *IEEE Transactions on Applied Superconductivity* **10**, 703 (2000). [51](#)
- [52] B. Seeber, editor, *Handbook of Applied Superconductivity* Vol. 1&2 (Institute of Physics Publishing, Bristol, Philadelphia, Pa, 1998). [51](#), [73](#), [76](#), [79](#), [81](#)
- [53] R. Turner, Personal communication, 2002. [52](#)
- [54] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions* U.S. National Bureau of Standards (Dover, New York, 1964). [55](#), [153](#), [155](#)
- [55] Legendre polynomials, <http://mathworld.wolfram.com/LegendrePolynomial.html>. [55](#)
- [56] C. Snow, *Hypergeometric and Legendre Functions with Applications to Integral Equations of Potential Theory*, 2nd ed. (National Bureau of Standards, Washington, D.C., 1952). [57](#)
- [57] M. L. Boas, *Mathematical Methods in the Physical Sciences*, 2nd ed. (J. Wiley & Sons, DePaul University, 1983). [61](#)
- [58] V. Cavaliere, M. Cioffi, A. Formisano, and R. Martone, *IEEE Transactions on Applied Superconductivity* **12**, 1413 (2002). [63](#), [139](#)
- [59] J. R. Sims *et al.*, *IEEE Transactions on Applied Superconductivity* **12**, 480 (2002). [72](#)

Bibliography

- [60] M. C. Wilson, *Superconducting Magnets*, 1st ed. (Oxford University Press, 1983). 72, 73, 78, 79, 81, 149
- [61] S. R. Elliott, *Physics and Chemistry of Solids* (John Wiley and Sons, 1998). 74
- [62] Engineer's Edge: Strength of Materials, http://www.engineersedge.com/strength_of_materials.htm. 75
- [63] Abaqus: Finite element analysis, <http://www.hks.com>. 75, 136, 161
- [64] D. Hagedorn and W. Nagele, CERN Report No. LHC Note 148, 1991 (unpublished). 79
- [65] T. Tominaka, K. Mori, and N. Maki, IEEE Transactions on Magnetics **28**, 727 (1992). 79
- [66] S. Pissanetzky and D. Latypov, Cryogenics **34**, 795 (1994). 79
- [67] J. H. Schultz, IEEE Transactions on Applied Superconductivity **12**, 1390 (2002). 82
- [68] F. W. Grover, *Inductance Calculations: working formulas and tables*, Special ed. (Instrument Society of America, 1981). 82, 83
- [69] W. R. Smythe, *Static and Dynamic Electricity*, 2nd ed. (McGraw-Hill, New York, 1950). 83
- [70] M. N. El-Derini, Journal of Applied Physics D **17**, 503 (1984). 83
- [71] M. T. Thompson, Power Control and Intelligent Motion **25**, 40 (1999). 83
- [72] K. J. Murphy and J. A. Brunberg, Magnetic Resonance Imaging **15**, 51 (1997). 85
- [73] K. Dantendorfer *et al.*, Magnetic Resonance Imaging **15**, 301 (1997). 86

-
- [74] D. I. Hoult *et al.*, *Journal of Magnetic Resonance Imaging* **13**, 78 (2001). 86
- [75] E. T. Laskaris *et al.*, *IEEE Transactions on Applied Superconductivity* **5**, 163 (1995). 86
- [76] H. Xu, S. Conolly, and B. Hu, Short bore-length asymmetric electromagnets for Magnetic Resonance Imaging, U.S. Patent #6 064 290, 2000. 86, 112
- [77] S. Prestmon, J. Miller, and Y. Eyssa, *IEEE Transactions on Applied Superconductivity* **10**, 1356 (2000). 86, 112
- [78] N. R. Shaw, Design of an ultra-short whole-body MRI magnet using genetic algorithms, MSci thesis, University of Cambridge, 1999. 87, 121
- [79] R. E. Ansorge, T. A. Carpenter, L. D. Hall, N. R. Shaw, and G. B. Williams, *IEEE Transactions on Applied Superconductivity* **10**, 1368 (1999). 87, 121
- [80] J. Bruck, D. Dolev, C. T. Ho, M. C. Rosu, and R. Strong, *Journal of Parallel and Distributed Computing* **40**, 19 (1997). 91
- [81] The Message Passing Interface (MPI) Standard, <http://www.lam-mpi.org/>. 91
- [82] The Message Passing Interface (MPI) Standard, <http://www-unix.mcs.anl.gov/mpi/>. 91
- [83] University of Cambridge: High Performance Computing Facility, <http://www.hpcf.cam.ac.uk/>. 91, 126
- [84] W. J. Duffin, *Electricity and Magnetism*, 4th ed. (McGraw-Hill, 1990). 93
- [85] K. Kaminishi and S. Nawata, *Review of Scientific Instruments* **52**, 447 (1981). 93

Bibliography

- [86] J. L. Wilson *et al.*, *NeuroImage* **17**, 967 (2002). [109](#)
- [87] H. Zhao, S. Crozier, and D. M. Doddrell, *Magnetic Resonance in Medicine* **45**, 331 (2001). [109](#), [111](#)
- [88] R. Turner, *Journal of Physics D* **19**, L147 (1986). [110](#)
- [89] S. Crozier and D. M. Doddrell, *Journal of Magnetic Resonance* **127**, 233 (1997). [110](#), [139](#)
- [90] P. N. Morgan, S. M. Conolly, and A. Macovski, *Magnetic Resonance in Medicine* **41**, 1221 (1999). [111](#)
- [91] V. Cavaliere, A. Formisano, R. Martone, and M. Primizia, *IEEE Transactions on Applied Superconductivity* **10**, 1376 (2000). [111](#), [139](#)
- [92] G. Danby, J. Linardos, J. Damadian, J. Dworkin, and R. V. Damadian, *A Proposed New Magnet Design for MR-Guided Surgery at 0.6 Tesla* (Mosby, St. Louis, Missouri, 1999), chap. 4, *Interventional MRI*, ed. R.B. Lufkin, pp. 29–34. [113](#)
- [93] S. G. Hushek, *Design of a Mid-Field Intraoperative MR System at 0.5 Tesla* (Mosby, St. Louis, Missouri, 1999), chap. 3, *Interventional MRI*, ed. R.B. Lufkin, pp. 15–28. [113](#)
- [94] T. Beyer *et al.*, *Journal of Nuclear Medicine* **41**, 1369 (2000). [115](#)
- [95] P. G. Kluetz *et al.*, *Clinical Positron Imaging* **3**, 223 (2000). [115](#)
- [96] B. Hasegawa *et al.*, *Nuclear Instruments and Methods in Physics Research A* , 140 (2001). [115](#)
- [97] R. B. Slaters *et al.*, *Physics in Medicine and Biology* **44**, 2015 (1999). [116](#), [118](#), [119](#)
- [98] M. E. Phelps, J. C. Mazziotta, and H. R. Schelbert, editors, *Positron Emission Tomography and Autoradiography* (Raven Press, New York, 1986). [117](#)

-
- [99] B. Hammer and N. Latham, NMR-PET Scanner Apparatus, U.S. Patent #4 939 464, 1990. [117](#)
- [100] N. Christensen, B. Hammer, B. Heil, and K. Fetterly, Physics in Medicine and Biology **40**, 691 (1995). [117](#)
- [101] Y. Shao *et al.*, Physics in Medicine and Biology **42**, 1965 (1997). [118](#), [119](#)
- [102] P. Marsden, D. Strul, S. Keevil, S. Williams, and D. Cash, The British Journal of Radiology **75**, S53 (2002), Special Issue. [118](#), [143](#)
- [103] P. B. Garlick, The British Journal of Radiology **75**, S60 (2002), Special Issue. [119](#), [143](#)
- [104] Magnex Scientific Ltd, <http://www.magnex.co.uk/>. [120](#)
- [105] N. R. Shaw and R. E. Ansorge, Far field spherical harmonic expansion for MRI main magnet design, in *Proceedings of the International Society for Magnetic Resonance in Medicine*, Toronto, Canada, 2003. [124](#)
- [106] Wolfson Brain Imaging Centre, <http://www.wbic.cam.ac.uk/>. [126](#)
- [107] V. Cavaliere *et al.*, IEEE Transactions on Applied Superconductivity **10**, 759 (2000). [139](#)
- [108] K. Deb, Computer Methods for Applied Mechanics and Engineering **186**, 311 (2000). [141](#)
- [109] H. Zhao and S. Crozier, Measurement Science and Technology **13**, 198 (2002). [149](#)
- [110] S. Wolfram, Mathematica, <http://www.wolfram.com/products/mathematica/>. [153](#), [155](#)
- [111] S. L. Belousov, *Tables of Normalized Associated Legendre Polynomials* (Permagon Press, Oxford, 1962). [153](#), [155](#)