Parallel Transmission Methods for Arterial Spin Labelling Magnetic Resonance Imaging

Aaron Oliver-Taylor

The Advanced Magnetic Resonance Imaging Group

Department of Medical Physics and Bioengineering

Ph.D Thesis

Submitted for the Degree of Doctor of Philosophy

University College London

February 2012

I, Aaron Oliver-Taylor confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

Vessel selective arterial spin labelling (ASL) is a magnetic resonance imaging technique which permits the visualisation and assessment of the perfusion territory of a specific set of feeding arteries. It is of clinical importance in both acute and chronic cerebrovascular disease, and the mapping of blood supplied to tumours. Continuous ASL is capable of providing the highest signal-to-noise (SNR) ratio of the various ASL methods. However on clinical systems it suffers from high hardware demands, and the control of systematic errors decreases perfusion sensitivity. A separate labelling coil avoids these problems, enabling high labelling efficiency and subsequent high SNR, and vessel specificity can be localised to one carotid artery. However this relies on the careful and accurate positioning of the labelling coil over the common carotid arteries in the neck.

It is proposed to combine parallel transmission (multiple transmit coils, each transmitting with different amplitudes and phases) to spatially tailor the labelling field, removing the reliance on coil location for optimal labelling efficiency, and enabling robust vessel selective labelling with a high degree of specificity. Presented is the application of parallel transmission methods to continuous ASL, requiring the development of an ASL labelling coil array, and a two channel transmitter system. Coil safety testing was performed using a novel MRI temperature mapping technique to accurately measure small temperature changes on the order of $0.1°C$. A perfusion phantom with distinct vascular territories was constructed for sequence testing and development. Phantom and in-vivo testing of parallel transmit CASL using a 3D-GRASE acquisition showed an improvement of up to 35% in vessel specificity when compared with using a single labelling coil, whilst retaining the high labelling efficiency and associated SNR of separate coil CASL methods.

The work presented in this thesis has featured in the following publications:

Parallel Transmit Vessel Selective Arterial Spin Labelling: Phantom and In-Vivo Results Oliver-Taylor, A., Ordidge, R., Randell, C. and Thomas, D. L. 1995. Proc. Intl. Soc. Mag. Reson. 20, 2012

Improving SNR in Small Temperature Change MR Thermometry to Acquire SAR Maps of a Pair of ASL Labelling Coils Oliver-Taylor, A., Ordidge, R., Randell, C. and Thomas, D. L. 2665. Proc. Intl. Soc. Mag. Reson. 20, 2012

A Perfusion Phantom with Distinct Vascular Territories Oliver-Taylor, A., Ordidge, R. and Thomas, D. L. Proc. Brit. Chap. ISMRM 2011

Parallel Transmit Vessel Selective Arterial Spin Labelling: A Proof of Concept Simulation Oliver-Taylor, A., Ordidge, R. and Thomas, D. L. 2086. Proc. Intl. Soc. Mag. Reson. 19, 2011

Contents

List of Figures

List of Tables

Acknowledgements

Throughout the course of this thesis I have received help, support and advice from the kind people around me, without which it would not have been possible for me finish. First and foremost I would like express my sincere thanks and gratitude to my primary supervisor Dr David Thomas for his support and guidance, for sharing his knowledge and skills with me, and for giving me the opportunity to work on this project. While sometimes it may have felt like there was little light at the end of the tunnel, his enthusiasm and positivity helped me to focus on my goals. To Professor Roger Ordidge I am ever grateful for his good advice, wealth of knowledge and experience, stories about the early days of MRI, and for giving me the opportunity to earn some money for attending conferences. My thanks to Chris Randell for his support in this project, particularly in giving me the opportunity to gain some hands on experience building coils at PulseTeq. To the rest of the staff at PulseTeq, particularly Nav Chhina, Colin Sunderland, David Gilderdale and Mark Cooper I am extremely grateful for the time, patience and expertise they expended in teaching me the black arts of RF coils; and to Richard Brisland who designed and manufactured the mechanical structure of the ASL Labelling Coils.

I would like to thank my colleagues in the AMRIG lab: Dr David Carmichael, Dr Enrico DeVita, Dr Harry Parkes, Dr Jack Wells, and Magda Sokolska. They have been a pleasure to work with, have contributed immensely to both my personal and professional time here at UCL. Also to the members of the WTCN, particularly Oliver Josephs to extremely useful discussions and advice, Dr Nikolaus Weiskopf and Dr Antoine Lutti for pulse sequence code, and Peter Aston for permitting me to use their workshop facilities, without which I wouldn't have been able to build many of the things featured in this thesis. Also my acknowledgements and thanks to Professor Xavier Golay for his advice and support.

I wish to thank my friends and family for putting up with me, particularly over the past 18 months. To my parents and brother I am ever grateful for their unequivocal love and encouragement. Particularly to my dad for financially supporting me this past year. To my wonderful girlfriend Isabel I am ever indebted for her love, patience, wisdom and personal support, I only hope I can do the same for her when she writes her thesis. Also to her parents for their support, advice and roast lamb dinners after a long day in the lab.

Finally, I would like to acknowledge the Medical Research Council for funding my project.

Chapter 1

Introduction

Arterial Spin Labelling (ASL) is a magnetic resonance imaging (MRI) technique which enables a quantitative, non-invasive assessment of cerebral perfusion [53, 237], the delivery of nutritive, arterial blood to the brain. Inflowing blood is used as an endogenous contrast agent, permitting visualisation of the spatial distribution of perfusion, which is of clinical importance in acute stroke [37], chronic cerebrovascular disease [52], epilepsy, neurodegenerative disorders and neuropsychiatric diseases [238]. Unlike exogenous contrast agent techniques such as dynamic susceptibility contrast (DSC) MRI [190], positron and single photon emission tomography (PET and SPECT) [15, 234], xenon-enhanced computed tomography [253], ASL can be performed repeatedly, enabling its use in functional studies [156]. Vessel selective ASL restricts the labelling to specific arteries, permitting the visualisation and assessment of their perfusion territory. This is of clinical significance in chronic and acute cerebrovascular disease [99, 221, 220, 141].

Perfusion is a microscopic process, where arterial blood pumped by the heart is carried via arteries and arterioles to the capillary bed. Oxygen and nutrients are exchanged with the tissue, waste products returned to the bloodstream, and water is also exchanged between the blood and the extracellular and intracellular space. Insufficient perfusion is known as ischaemia and can occur in both acute and chronic disease states. A shortage of glucose and oxygen, essential for cellular metabolism can impair cell function resulting in irreversible damage. Acute stroke, occurring either when blood vessels become blocked (ischaemic stroke) or from a rupture of blood vessels (haemorrhagic stroke), is the consequence of sustained neuronal ischaemia. It results in a region of irreversible damage; an infarction, surrounded by a hypoperfused ischaemic penumbra, which if reperfused within a certain time can regain its function [17]. The ability to visualise the infarction and ischaemic penumbra is an extremely important clinical tool, aiding treatment planning and facilitating the monitoring of therapy. ASL is capable of providing these quantitative perfusion measurements, providing important information about whether an ischaemic organ is viable [180].

This thesis presents a novel vessel selective arterial spin labelling method, which aims to improve

arterial specificity whilst retaining high signal to noise and sensitivity in perfusion measurements. It combines the merits of separate coil continuous arterial spin labelling [54, 259] with the concepts of parallel transmission [116, 155, 249, 228] to spatially tailor the transmitted radio frequency (RF) magnetic field, B_1^+, used when labelling. Continuous ASL (CASL) utilises long (lasting up to several seconds) RF and gradient pulses to invert inflowing arterial blood. Standard clinical MRI hardware utilising volume transmitter coils [91, p.850] cannot always produce long RF pulses, and also introduces magnetisation transfer contrast (MTC) [240, 47, 101] to the imaging volume, decreasing perfusion sensitivity, and requiring this same contrast in a control acquisition [10] otherwise systematic errors are introduced when quantifying perfusion.

An alternative method is to use a localised labelling coil, positioned directly over the common carotid arteries. MTC is eliminated because no labelling field exists within the imaging volume, and power requirements are reduced, such that continuous wave capable, low power RF amplifiers can be used. Separate labelling coil CASL is capable of high labelling efficiencies, therefore high signal to noise perfusion images, and left–right selectivity is possible by using a small labelling coil located over one of the carotid arteries. However it relies upon precise, anatomy dependent positioning of the coil. Previous implementations [259, 176] require repositioning of the coil to label each carotid artery, and a different coil to label both.

In this work it is proposed to spatially tailor the labelling field with parallel transmission; multiple simultaneous RF transmitter coils, each transmitting with a different amplitude and phase, which sum to create a total B_1^+ field that is their interference pattern. Using anatomical scans of the arterial structure and acquired B_1^+ maps, the ASL acquisition can be planned in a similar manner to normal scanning, and the appropriate sets of amplitudes and phases computed automatically. The resulting field will spatially vary such that it is of a suitable amplitude for efficient labelling over the specified arteries, yet of a minimal amplitude in other arteries to eliminate their contribution to the perfusion image. By doing so the dependence on the precise location of the coils is removed, and it permits both vessel selective and global ASL scans to be acquired without repositioning. Continuous ASL has been shown theoretically and experimentally to have the largest available perfusion signal [241], however it is mired by the aforementioned imperfections in implementation. A separate labelling coil circumvents these imperfections, permitting highly efficient labelling and a large perfusion signal. Therefore parallel transmit CASL has the potential to provide improved flexibility in vessel specificity, reduced contralateral labelling, and improved reproducibility whilst retaining the large perfusion signal that is associated with separate coil CASL.

This work has focused on implementing a parallel transmission enabled set of labelling coils on a clinical MR system, and their use in a CASL sequence. Chapter 3 presents a computational simulation of the continuous ASL labelling procedure. Parallel transmission was utilised to improve vessel specificity when labelling the right-hand side neck arteries by forming a null in the labelling B_1^+ field over the left-hand side arteries, reducing contralateral labelling and their contribution to

the perfusion signal. Simulations were performed for two and four coils, and results compared with those from one coil (positioned at the right-hand side of the neck).

A pair of labelling coils were designed, constructed and characterised in chapter 4. Coils were optimised for good left–right specificity when used independently; require no per-subject tune and match adjustment; and include safety mechanisms to mitigate the risk of patient burns. Chapter 5 describes a low power two channel transmitter system that was developed to drive the RF coils. Based around a recycled spectrometer, it has full independent amplitude and phase control on each channel, facilitating parallel transmission.

The human body, consisting of tissues which have some degree of conductivity, interacts with radio frequency electromagnetic fields, resulting in the deposition of power into the body. As RF magnetic fields are generated in MRI it is important to assess and if necessary limit any power deposition to ensure that scanning remains safe. In chapter 6 the RF power deposition from these coils is assessed using an MRI technique to measure temperature [118, 185, 41, 168]. A novel method was developed which permits the accurate measurement of very small temperature changes, which would otherwise be smaller than the noise level of any two repeated measurements.

In order to perform parallel transmission it is necessary to obtain a map of the transmitted B_1^+ field for each coil element. Furthermore, for small surface coils, where the transmitted field is highly inhomogeneous over the field of view, regions of low excitation have a very low signal to noise ratio, introducing large errors into the acquired B_1^+ maps. An optimised method for acquiring accurate B_1^+ maps that are not affected by this low excitation, low SNR limit is presented in chapter 7. Also presented in this chapter is an advanced, parallel transmit enabled continuous ASL sequence which utilises a 3D-GRASE acquisition module [173, 69] developed from the code base of a FAIR pulsed ASL [130, 90] 3D-GRASE sequence. The sequence includes interleaved background suppression [56, 150, 250, 44] pulses to suppress signal contribution from static background tissue.

To provide a controlled environment for the development and testing of parallel transmit CASL, a novel perfusion phantom with distinct vascular territories was constructed. It utilised computer aided design and 3D printing (a rapid prototyping technique) to generate a complex network of inflow pipes to simulate the macroscopic flow within the carotid and cerebral arteries; and small glass beads to simulate incoherent microscopic flow at the capillary bed.

Chapter 8 presents the culmination of this work: an initial phantom and in-vivo study using parallel transmit CASL. Vessel selective CASL scans were run with the objective of obtaining left–right selectivity of the common carotid arteries, using a single labelling coil and two coil parallel transmission.

Finally, chapter 9 presents a critical review and discussion of the work in this thesis, with particular emphasis on the technique as a whole. Moreover, the potential for this technique under certain idealistic circumstances is hypothesised and discussed, and directions for immediate further work are suggested.

Chapter 2

Background

2.1 Magnetic Resonance Imaging

Magnetic Resonance Imaging (MRI) is a non-ionising imaging technique that makes use of the phenomenon known as Nuclear Magnetic Resonance (NMR) to visualise biological tissue. NMR occurs when nuclei with non-integer spin, for example ^{1}H, ^{13}C, ^{19}F, ^{23}Na, are placed within a magnetic field. Normally the energy levels for spin up and spin down are degenerate, but in the presence of a magnetic field the state which is aligned with the magnetic field has a lower energy, and consequently this becomes the equilibrium state. This energy difference is much smaller than the average thermal energy at room temperature, and as a result only a very small excess of nuclei are in the spin up state, producing a small net magnetisation. Furthermore, the nuclei can efficiently absorb and emit electromagnetic radiation at a frequency equal to the energy difference between the spin up and spin down states. This frequency is known as the Larmor frequency. The small amount of excess magnetisation can be manipulated out of equilibrium by irradiating a sample with a magnetic field at the Larmor frequency, and as the magnetisation returns back to its equilibrium state it produces a signal which can be detected, from which information about the sample can be deduced.

Using this ability to interact with non-integer spin nuclei, MRI can produce detailed, quantitative images of the human body. Water, consisting of two hydrogen atoms and one oxygen atom, makes up approximately 60% of the human body, providing a large source of protons for ^{1}H MRI. Consequently it is the most common form of MRI; the reduced abundance of other nuclei results in a smaller signal. Furthermore, many different contrast mechanisms which can be exploited, producing images which are sensitised to anatomy, pathology, physiology and function.

2.1.1 Nuclear Magnetic Resonance

2.1.1.1 Spin and Precession

A fundamental property of all subatomic particles is their intrinsic angular momentum, otherwise known as spin. In a classical sense, this is analogous to considering subatomic particles to be spheres which rotate. Those which carry an electric charge, such as a proton, can be considered to be loops of moving charge; loops of electrical current, which generate a magnetic field. Thus, subatomic particles with non-zero spin and charge also have a magnetic moment.

The classical interaction between the magnetic moment of a proton, $\boldsymbol{\mu}$ and an external magnetic field, $\mathbf{B} = B_0\hat{\mathbf{z}}$, describes that the magnetic moment experiences a torque (turning force), forcing it to align with the direction of the field, as shown in figure 2.1.

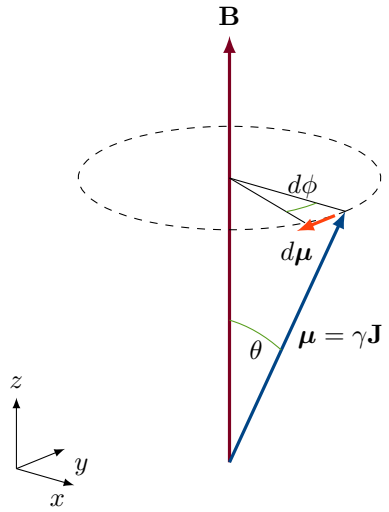

Figure 2.1: Classical model of precession

Because nuclear spin is quantised (see below) the magnetic moment cannot fully align with $\mathbf{B}$, and consequently precesses about it, in same way that a gyroscope precesses about an axis defined by gravity. The rate of change of the magnetic moment is given by the cross product of the external magnetic field and the magnetic moment:

$$\frac{d\boldsymbol{\mu}}{dt} = \boldsymbol{\mu} \times \mathbf{B} \tag{2.1}$$

The magnitude of this change is:

$$|d\boldsymbol{\mu}| = \gamma\mu B_0 \sin\theta dt \tag{2.2}$$

From figure 2.1, the differential change in the magnetic moment is also

$$|d\boldsymbol{\mu}| = \mu \sin\theta d\phi \tag{2.3}$$

Equating equations 2.2 and 2.3 results in $\frac{d\phi}{dt} = \gamma B_0$, and by substituting $\frac{d\phi}{dt} = \omega_0$ this becomes the familiar form of the Larmor equation (given in equation 2.4). This states that the precession occurs at a frequency which is proportional to the strength of the external magnetic field by a factor known as gyromagnetic ratio, γ, which for protons is $2.7 \times 10^8 \text{rad}\, s^{-1} T^{-1}$.

$$\omega_0 = \gamma B_0 \tag{2.4}$$

In quantum mechanics all observable quantities, including angular momentum are quantised into discrete values. These discrete values are characterised by their quantum number, and for spin, the intrinsic angular momentum of a particle, its quantum number s, can take on values which are positive integer multiples of $1/2$. Therefore the allowed values of spin, S, are:

$$S = \hbar\sqrt{s\,(s+1)} = \frac{\hbar}{2}\sqrt{n\,(n+2)} \tag{2.5}$$

Where n is a non-negative integer and $\hbar$ is Planck's constant divided by 2π. Thus S is quantised in positive integer multiples of $\frac{\hbar}{2}$. The direction of the spin is also quantised, however at the quantum level this is not as simple as the classical example of a sphere spinning either clockwise or anticlockwise. Certain measurable quantum mechanical quantities have an uncertainty relationship (from Heisenberg's Uncertainty Principle); that is to say that the more accurately one is measured, the greater the uncertainty in the other. Orthogonal components of the total angular momentum, which consists of spin and orbital angular momentum (only valid for electrons), have an uncertainty relationship, which implies that all three components of the spin direction (i.e. S_x, S_y, and S_z) cannot have simultaneous definite values. Therefore only the component of the spin along one direction can be measured, and this is conventionally chosen to be the z-axis. For every spin quantum number, there are $2s + 1$ values of S_z, hence for a proton, where $s = 1/2$, two states are possible: where the proton's spin points along the $+z$ direction, referred to as "spin up", and where it points along the $-z$ direction, referred to as "spin down". Therefore, the magnetic moment of a proton is discretised:

$$\boldsymbol{\mu}_p = \gamma m_s \hbar \tag{2.6}$$

Where $m_s = 1/2$ for spin up, and $m_s = -1/2$ for spin down. It has energy levels:

$$E = -\boldsymbol{\mu} \cdot \mathbf{B} = -\mu_z B_z = -\gamma m_s \hbar B_z \tag{2.7}$$

For a constant external magnetic field $\mathbf{B_0} = B_0\hat{\mathbf{z}}$, where $\hat{\mathbf{z}}$ is a unit vector in the $+z$ direction, the difference between the two energy levels is:

$$\Delta E = \gamma \hbar B_0 \tag{2.8}$$

Wave-particle duality states that a photon with this energy is equivalent to an electromagnetic wave with frequency equal to the rate of precession given in the Larmor equation (equation 2.4).

2.1.1.2 Equilibrium Magnetisation

In the previous section the concept of spin up and spin down states was introduced, and the energy difference between the two states was derived. By considering how this energy difference compares with the thermal energy of the sample, this section will show how the equilibrium macroscopic magnetisation arises, and also how further MRI phenomenon can be explained using classical physics.

The energy difference between the spin up and spin down states varies linearly with magnetic field strength and the gyromagnetic ratio. For protons in a 3T magnet this has a value of $5.3 \times 10^{-7} eV$. The thermal energy of a system of protons is on the order of kT[1], where k is Boltzmann's constant. At body temperature ($37°C$) $kT = 0.027eV$, hence $\gamma \hbar B_0 \ll kT$. For a very large number of protons, and in the absence of a magnetic field, there is no energy difference between spin up and spin down states, so at any given moment there would an equal number of protons in the spin up and spin down states, and no net magnetisation because all of the magnetic moments sum to zero. With a magnetic field present the degenerate spin up and spin down states split, with the spin up state lower in energy. The average thermal energy of the protons is much larger than this energy difference, so protons are in a continuous flux between the two states. However, because it is lower in energy there is a small, but non-negligible tendency for the spin up state to be more populated than the spin down state. Therefore the magnetic moments do not sum to zero, but instead have a small net magnetisation aligned with the external magnetic field.

The Boltzmann distribution describes the probability of a particle having a given energy ϵ, when in thermal contact with a system at a defined temperature T:

$$P(\epsilon) = \frac{e^{-\frac{\epsilon}{kT}}}{\sum_{\epsilon} e^{-\frac{\epsilon}{kT}}} \tag{2.9}$$

The denominator is a normalisation function, which is summed over all energies. For a large number of particles the probability is also equal to the ratio of the number of particles with energy ϵ, and the total number of particles. The probability of a proton being in the spin state m_s, and the subsequent population ratio is:

$$P(\epsilon(m_s)) = \frac{N_{m_s}}{N} = \frac{e^{-\frac{m_s \gamma B_0}{kT}}}{\sum_{\epsilon} e^{-\frac{\epsilon}{kT}}} \tag{2.10}$$

[1] A distribution of energies exists, given by the Boltzmann distribution.

The ratio of spin up to spin down protons is therefore:

$$\frac{N_u}{N_d} = e^{\left(\frac{\Delta E}{kT}\right)} = e^{\left(\frac{\gamma\hbar B_0}{kT}\right)} \tag{2.11}$$

This can be expanded as a Taylor series, which as an approximation becomes:

$$\frac{N_u}{N_d} = 1 + \frac{\gamma\hbar B_0}{kT}$$

At body temperature this is $1 + 6.6 \times 10^{-6} B_0$ [91, p.86] , or in other words for each tesla of magnetic flux density, if there are a million protons in the spin down state, one million and six will be in the spin up state. In 9g of water there are 6.02×10^{23} protons[2], so even for modest field strengths this sums to a small, but measurable macroscopic magnetisation. For a sample with proton density ρ, the equilibrium macroscopic magnetisation per unit volume is:

$$M_0 = \frac{\rho\gamma^2\hbar^2 B_0}{4kT} \tag{2.12}$$

While the macroscopic magnetisation is formed from many quantum objects, each of which obey the rules of quantum mechanics, it itself is a macroscopic entity, representing the vector sum of all the magnetic moment from all of the protons within a sample. Therefore, its behaviour is described by classical physics.

2.1.1.3 The Rotating Frame

For simplicity, it is convenient to consider the behaviour of the macroscopic magnetisation in a rotating reference frame which rotates about the z-axis at the Larmor frequency. Consider a rotating frame with rotational angular velocity vector $\mathbf{\Omega}$ (the axis of rotation is general, i.e. not necessarily the z-axis). The orthogonal unit vectors within the rotating frame are $\hat{\mathbf{x}}'$, $\hat{\mathbf{y}}'$ and $\hat{\mathbf{z}}'$. In the absence of relaxation (section 2.1.1.6) the equation of motion of the macroscopic magnetisation vector $\mathbf{M}$ within the static laboratory frame is

$$\frac{d\mathbf{M}}{dt} = \gamma\mathbf{M} \times \mathbf{B_0} \tag{2.13}$$

In the rotating frame this becomes

$$\left(\frac{d\mathbf{M}}{dt}\right)' = \gamma\mathbf{M} \times \left(\mathbf{B_0} + \frac{\mathbf{\Omega}}{\gamma}\right) \tag{2.14}$$

The rotating frame act as an additional source of magnetic field, therefore the magnetisation vector behaves according to an effective magnetic field, $\mathbf{B_{eff}}$ which is the superposition of the

[2]Water, H_2O has a molecular weight of 18, therefore one mole weighs 18g. As there are two hydrogen atoms per water molecule, one mole of hydrogen atoms is present in every 9g of water.

external magnetic field and the fictitious magnetic field resulting from the rotating frame:

$$\mathbf{B_{eff}} = \mathbf{B_0} + \frac{\mathbf{\Omega}}{\gamma} \tag{2.15}$$

If the rotating frame is chosen such that it rotates with the precession of the magnetisation, and the subsequent motion of the magnetisation within the field is zero, then the fictitious field $\frac{\mathbf{\Omega}}{\gamma}$ exactly opposes $\mathbf{B}_0$, i.e $\mathbf{\Omega} = -\gamma\mathbf{B_0}$.

2.1.1.4 Excitation

The macroscopic magnetisation is very small in comparison to the main magnetic field (by a factor on the order of 10^6), making it difficult to measure when aligned with $\mathbf{B_0}$. However, if it is rotated into the plane orthogonal to $\mathbf{B}_0$ (the x-y plane), then the magnetisation can be detected independently of the main magnetic field by the voltage generated in an appropriate receiver coil through Faraday Induction. Rotation of the magnetisation is achieved by applying a second magnetic field, $\mathbf{B_1}$, perpendicularly to $\mathbf{B_0}$. In the rotating frame where $\mathbf{\Omega} = -\gamma\mathbf{B_0}$, the equation of motion becomes:

$$\left(\frac{d\mathbf{M}}{dt}\right)' = \gamma\mathbf{M} \times \mathbf{B_1} \tag{2.16}$$

$\mathbf{B_1}$ is most effective at rotating the magnetisation when it is static with respect to the rotating frame. This requires that $\mathbf{B_1}$ oscillates at the Larmor frequency, or in other words is on-resonance. As typical clinical and experimental values for B_0 range from 0.5 - 9.4T, the frequency of the applied B_1 field ranges from approximately 32 - 400MHz, which is within the radio frequency (RF) spectrum. The polarisation of $\mathbf{B_1}$ is also important. Consider a $\mathbf{B_1}$ field which is linearly polarised along +x-axis within the laboratory frame:

$$\mathbf{B_1^{lin}} = B_1 \cos(\omega t)\,\hat{\mathbf{x}} \tag{2.17}$$

This can also be expressed in terms of the basis unit vectors within the rotating frame:

$$\mathbf{B_1^{lin}} = \frac{1}{2}B_1\left(\hat{\mathbf{x}}'(1 + \cos 2\omega t) + \hat{\mathbf{y}}' \sin 2\omega t\right) \tag{2.18}$$

Which has a constant term plus two oscillating terms, both of which average to zero over long durations, subsequently only half of the amplitude B_1 actually contributes towards nutation of the magnetisation. A more efficient $\mathbf{B_1}$ can be produced by using a circularly polarised field, which is constant within the rotating frame. This is constructed from two linearly polarised fields in the laboratory frame, with a 90° phase shift between the two, as shown in equation 2.19, and is hence termed a "quadrature" field. Because it is constant within the rotating frame all of the amplitude

contributes towards nutating the magnetisation, therefore the same nutation angle can be achieved using half the B_1 amplitude ($\sqrt{2}$ of the power).

$$\mathbf{B_1^{cir}} = B_1 \left(\hat{\mathbf{x}} \cos \omega t + \hat{\mathbf{y}} \sin \omega t\right) = B_1 \hat{\mathbf{x}}' \tag{2.19}$$

The positive circularly polarised RF field which is used to rotate the magnetisation is also known as the $\mathbf{B_1^+}$. The act of rotating the magnetisation into the transverse plane is also known as "excitation" (because the spins are excited), "flipping", or "tipping", referring to the resultant angle which the magnetisation is rotated to. For on-resonance pulses, the flip angle, α is proportional to the time integral of the amplitude envelope of the applied B_1 field:

$$\alpha = \gamma \int_0^\tau B_1(t) dt \tag{2.20}$$

2.1.1.5 Detection

Once tipped into the x-y plane, the transverse component of the magnetisation precesses about the main magnetic field, creating a small magnetic field which is on the order of μT. This field is known as the B_1^-, and it is negatively circularly polarised. A (well engineered, see section 2.2) conducting coil positioned such that it the flux of this transverse magnetic field passes through the surface contained within the coil's conductors has an electromotive force (emf) induced according to Faraday's law of induction:

$$emf = \frac{d}{dt} \int_S \mathbf{B_1^-} \cdot d\mathbf{S} \tag{2.21}$$

Once detected, the signal is demodulated by multiplication with a corresponding reference sinusoidal signal that has an angular frequency equal or close to ω_0. This produces two signals[3], one at a low frequency, and the other at a much higher frequency than ω_0. Low pass filtering removes the high frequency signal, and retains what is the equivalent of the motion of the magnetisation within the rotating frame. This procedure is performed with cosine and sine reference signals to obtain real and imaginary channels, respectively. By doing so the phase of the magnetisation can be deduced, this is called phase sensitive detection. It is convenient to use complex notation to represent the measured MR signal, $s(t)$, from which the magnitude and phase become apparent:

$$s(t) = \mathrm{Re}(s) + i\mathrm{Im}(s) = |s| \, e^{i\phi}$$

Where the signal the signal phase, ϕ is the argument of the complex signal, and given by

$$\phi = \tan^{-1}\left(\frac{\mathrm{Im}(s)}{\mathrm{Re}(s)}\right)$$

[3]From the trigonometric identity: $\sin a \sin b = \frac{1}{2}\left(\cos(a-b) - \cos(a+b)\right)$

Typically the real and imaginary channels of the signal are sampled by an analogue to digital converter (ADC), and stored within computer memory, ready for processing and reconstruction. Analogous to creating a quadrature B_1^+ field, magnetisation detection can also be performed in quadrature. By independently detecting the x' and y' components of B_1^-, then combining with a 90° phase shift the signal is increased in amplitude compared to linear detection.

As with any measurement, noise is also present when detecting the MR signal, its source being thermal noise within the detector coil. Therefore, the signal-to-noise ratio, SNR is given by [110, 114, 91, p.852]:

$$\text{SNR} \propto \frac{M_{xy}\mathcal{B}_\perp}{\sqrt{4kT_{coil}\Delta f\mathcal{R}}} \propto \frac{\omega_0^{\frac{7}{4}}\mathcal{B}_\perp}{\sqrt{T_{coil}\Delta f V_{sens}}}$$

Where $\mathcal{B}_\perp$ is the receive coil sensitivity, defined as the magnetic field generated by the receive coil per unit current; k is the Boltzmann constant; T_{coil} is the temperature of the coil; Δf is the experiment bandwidth; V_{sens} is the volume over which the receive coil is sensitive; $\mathcal{R}$ is the effective coil resistance, which has contributions from the sample, the resistivity of the coil's conductors and components, and noise introduced by the electronics such as the pre-amplifier. The factor $\omega_0^{\frac{7}{4}}$ arises from radio frequency losses within the sample which make the noise resistance proportional to $\sqrt{\omega_0}$; without these losses the signal would be proportional to ω_0^2 [62]. Hence, higher magnetic field strengths result in better SNR.

2.1.1.6 Relaxation Mechanisms

Longitudinal Relaxation The equilibrium state for a system of protons within a magnetic field is for all of the magnetic moments to align with the field. If perturbed from this minimum energy equilibrium state, for example by the application of a RF B_1^+ pulse, the system returns back to equilibrium by giving up the energy gained in producing the initial perturbation. This energy is transferred to the protons' surrounding environment, or lattice, hence this process is termed "spin-lattice" relaxation. As excited protons return to the lower energy level and their magnetic moments realign with B_0, the longitudinal component of the macroscopic magnetisation (M_z) returns back to its equilibrium value M_0, at a rate proportional to $M_0 - M_z$:

$$\frac{dM_z}{dt} = \frac{M_0 - M_z}{T_1} \tag{2.22}$$

The longitudinal relaxation time, T_1 shown in equation 2.22, is the time constant for the longitudinal magnetization to return to equilibrium. If the magnetisation is tipped fully into the transverse plane, i.e. $\alpha = 90°$, $M_z(t = 0) = 0$, then equation 2.22 has solution:

$$M_z(t) = M_0\left(1 - e^{-\frac{t}{T_1}}\right) \tag{2.23}$$

The mechanism behind spin-lattice relaxation is as follows. Once promoted to the higher spin down energy state, and when the B_1^+ field has ceased, a proton is in a stable energy state and will not return to the lower state spontaneously, but rather requires an external stimulating field. Neighbouring protons, and other nuclei and molecules that have magnetic moments can provide this field, primarily through the movement of one proton in a water molecule with respect to the other. These motions cause a fluctuation in the magnetic field, and for these fluctuations to have an effect and stimulate a transition of spin states they need to occur at the Larmor frequency. The physical and chemical composition of a material affects the T_1 relaxation time, as does the field strength.

Transverse Relaxation Following excitation the transverse component of the magnetisation, M_{xy} is a manifestation of the phase coherence of an ensemble of microscopic magnetic moments. So long as the precession frequency of the ensemble of magnetic moments remains the same, they will continue to have the same phase, and M_{xy} will only decay at a rate equal to the longitudinal relaxation. However local magnetic field inhomogeneities and subsequent differences in precession frequency result in a "fanning out" of the magnetic moments and loss in phase coherence.

This decay is characterised by the "spin-spin", or transverse relaxation time, T_2. In the rotating frame this has the form:

$$\left(\frac{d\mathbf{M_{xy}}}{dt}\right)' = -\left(\frac{M_x}{T_2}\hat{\mathbf{x}}' + \frac{M_y}{T_2}\hat{\mathbf{y}}'\right) \tag{2.24}$$

Following a 90° excitation pulse this has solution

$$\mathbf{M_{xy}}(t) = (M_{x,0}\hat{\mathbf{x}} + M_{y,0}\hat{\mathbf{y}})\, e^{-\frac{t}{T_2}}$$

Where $M_{x,0}$ and $M_{y,0}$ are the magnitudes of the x and y components of the magnetisation at t=0 (immediately after the 90° pulse). Two separate sources are responsible for the spin-spin interaction: intrinsic time dependent B_0 inhomogeneities, and external static B_0 inhomogeneities. These can be characterised by two independent relaxation times, T_2 for the intrinsic inhomogeneities, and T_2' for the external field inhomogeneities, which are combined according to equation 2.25 to make the total spin-spin relaxation time, T_2^*.

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'} \tag{2.25}$$

Because the inhomogeneities responsible for T_2' do not change over the duration of the decay, the dephasing that they cause can be reversed by applying an additional 180° pulse to the excited magnetisation, leading to the recovery of the original phase relationship and the creation of a "spin echo" (see section 2.1.1.9).

Like T_1, the intrinsic T_2 relaxation is affected by molecular motion, however unlike T_1 the motion does not have to occur close to the Larmor frequency. Where there is fast molecular tumbling, for example within a liquid, the field fluctuations average out to a homogeneous field, resulting in long T_2 relaxation times. On the other hand slow tumbling, for example bound protons within large molecules see inhomogeneities which do not average out, and consequently more dephasing occurs, hence shorter T_2 times. T_2' is affected by the static inhomogeneities in the external magnetic field, which are caused by differences in magnetic susceptibility between materials.

2.1.1.7 The Bloch Equation

The motion of the magnetisation is described by the Bloch equation. This combines the effect of external magnetic fields, longitudinal and transverse relaxation:

$$\frac{d\mathbf{M}}{dt} = \gamma\mathbf{M} \times \mathbf{B} + \frac{M_0 - M_z}{T_1}\hat{\mathbf{z}} - \frac{\mathbf{M_{xy}}}{T_2}$$

For a static field $\mathbf{B} = B_0\hat{\mathbf{z}}$, this has solutions:

$$M_x(t) = (M_{x,0}\cos\omega_0 t + M_{y,0}\sin\omega_0 t)\, e^{-\frac{t}{T_2^*}}$$

$$M_y(t) = (M_{y,0}\cos\omega_0 t - M_{x,0}\sin\omega_0 t)\, e^{-\frac{t}{T_2^*}}$$

$$M_z(t) = M_{z,0}e^{-\frac{t}{T_1}} + M_0(1 - e^{-\frac{t}{T_1}})$$

Where $M_{z,0}$ is the value of M_z at $t = 0$.

2.1.1.8 Free Induction Decay

The Free Induction Decay (FID) is the simplest MRI experiment. It involves exciting the magnetisation of a sample with a 90° pulse, then immediately detecting the transverse magnetisation as it decays. This is illustrated in figure 2.2, where the FID signal has been demodulated to produce both real (blue) and imaginary (orange) signals, both of which have an amplitude envelope governed by T_2^* decay.

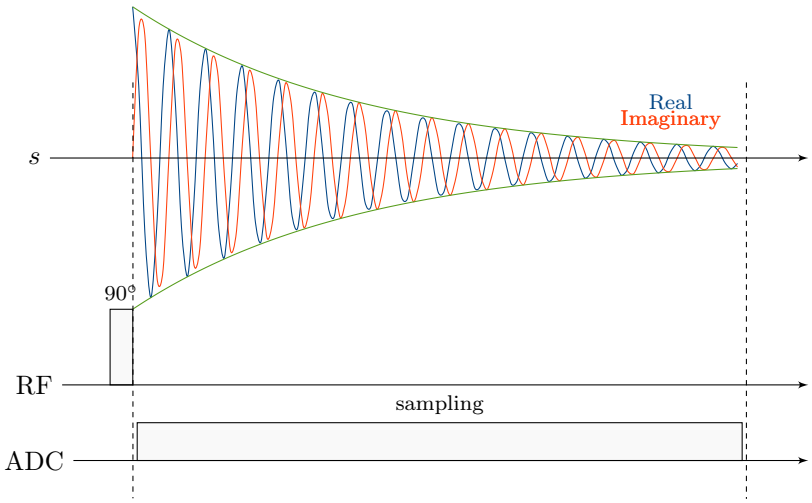

Figure 2.2: Free Induction Decay Pulse Sequence and Signal

2.1.1.9 Spin Echo

The spin echo sequence adds an additional pulse after the initial excitation pulse which refocuses the dephasing caused by static inhomogeneities. This is shown schematically in figure 2.3. After excitation, the magnetisation (and available signal) decays according to e^{-t/T_2^*}, as with the free induction decay, shown by the orange amplitude envelope. A pulse with a flip angle of 180°, termed the refocusing pulse, is applied at a time $^{\mathrm{TE}}/_2$ after the initial excitation. It has the effect of reversing the sign of the phase accrued in the time between the excitation pulse and the refocusing pulse. Following the pulse, phase continues to be accrued, but because the sign of the phase has been reversed the net accrual reverses, and at a time $^{\mathrm{TE}}/_2$ after the refocusing pulse the original phase is recovered and a spin echo is formed. At its peak the spin echo (shown by the yellow amplitude envelope in figure 2.3) has an amplitude which is proportional to $e^{-\mathrm{TE}/T_2}$, the green amplitude envelope in figure 2.3. Multiple echoes can be acquired by repeating the refocusing pulse every TE. The signal is still subject to T_2 decay, therefore repeated echoes have a lower amplitude. This technique is utilised in the RARE (Rapid Acquisition Relaxation Enhanced) [100], fast spin echo [121], and 3D-GRASE [173] acquisition techniques.

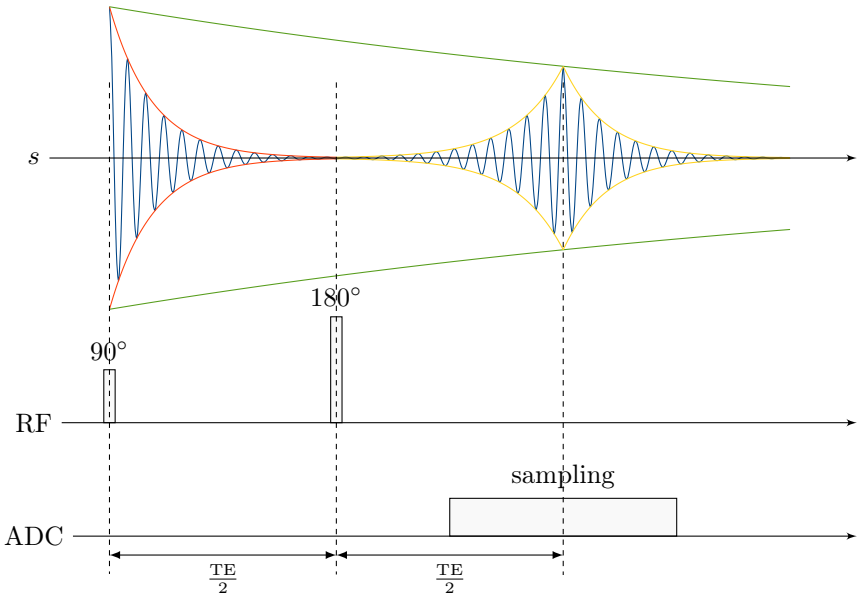

Figure 2.3: Basic Spin Echo Pulse Sequence and Signal

2.1.2 Image Formation

Magnetic Resonance Imaging, uses the principles of NMR to produce images. The preceding description of NMR has only considered the magnetisation of a large sample, immersed in a uniform magnetic field such that the Larmor frequency, ω_0 is the same at all points in space. In this section, the principles of image formation, and some basic imaging sequences are discussed, from which almost all subsequent imaging techniques have their origins.

2.1.2.1 Frequency Encoding

The aims of image formation are to impose a spatial dependence onto some property of the magnetisation, which upon detection can be interpreted and an image produced. One such way is to spatially modulate the amplitude of the static magnetic field, such that different locations correspond to different precession frequencies. This is termed frequency encoding, and is achieved using a linear magnetic field gradient. This is an additional magnetic field which is in the direction of $\mathbf{B_0}$, and has constant gradient, i.e

$$G_x = \frac{\partial B_z}{\partial x}, G_y = \frac{\partial B_z}{\partial y}, G_z = \frac{\partial B_z}{\partial z} \tag{2.26}$$

Resulting in a total field

$$\mathbf{B} = \mathbf{B_0} + \mathbf{G} \cdot \mathbf{r}\hat{\mathbf{z}} = \mathbf{B_0} + (G_x x + G_y y + G_z z)\hat{\mathbf{z}} \tag{2.27}$$

Using a linear combination of gradients in the x, y and z direction, a gradient can be constructed in any arbitrary direction. A linear gradient creates a mapping between precession frequency and position. By acquiring a MR signal whilst a gradient is imposed, and performing a 1D Fourier

transform, the spatial signal variation can be resolved in one dimension. During the application of a gradient of duration t the accumulated signal phase is given by:

$$\phi(\mathbf{r}, t) = -\gamma \int_0^t \mathbf{G} \cdot \mathbf{r} dt' \tag{2.28}$$

Consider frequency encoding along the x-axis in the absence of relaxation. Neglecting the initial phase, the demodulated signal at time t is modulated by the gradient induced phase accrual:

$$s(t) = \int \rho(x) e^{-i\gamma x \int G(t')dt'} dx \tag{2.29}$$

Where ρ is a function representing the spatial proton density. By defining a spatial frequency (see section 2.1.2.4) k_x as

$$k_x = \frac{\gamma}{2\pi} \int_0^t G_x(t')dt$$

Equation 2.29 can be rewritten in terms of the spatial frequency

$$s(k_x) = \int \rho(x) e^{-2\pi i k_x x} dx \tag{2.30}$$

Which shows that for linear gradients, the signal is the Fourier transform of the proton density. By taking the inverse Fourier transform, the spatial proton density can be computed:

$$\rho(x) = \int s(k_x) e^{2\pi i k_x x} dk_x \tag{2.31}$$

2.1.2.2 Phase Encoding

Spatial encoding in a second dimension, orthogonal to the frequency encoding dimension is achieved using phase encoding. A 2D spin density function, i.e. $\rho(x, y)$, requires a 2D "signal" in the spatial frequency domain. Frequency encoding allows the acquisition of lines of spatial frequencies along the frequency encoding axis, but for a 2D object the spatial frequencies along the other axis must also be measured. To do so, a pulsed approach to frequency encoding is taken. Consider a gradient pulse of amplitude G_y and duration t_p applied in the $+y$ direction prior to the frequency encoding step described in the preceding section. The fixed duration of the phase encoding gradient prescribes a linear phase ramp along the y-axis of the magnetisation:

$$\phi(y) = -\gamma G_y t_p y \tag{2.32}$$

Which corresponds to a spatial frequency

$$k_y = \frac{\gamma}{2\pi} G_y t_p \tag{2.33}$$

The signal during the frequency encoding gradient is

$$s(t) = s(k_x, k_y) = \int\int \rho(x, y) e^{-2\pi i(k_x x + k_y y)} dx\, dy$$

As k_y is constant during the measurement, only one "line" of the spatial frequency domain can be sampled in each measurement. The experiment must be repeated to sample the entire spatial domain in the k_y direction (it is assumed that the frequency encoding step sufficiently samples in the k_x direction). Once all spatial frequencies have been sampled, the 2D proton density function is obtained via inverse Fourier transform:

$$\rho(x, y) = \int\int s(k_x, k_y) e^{2\pi i(k_x x + k_y y)} dk_x\, dk_y$$

Phase encoding can also be performed along the z-axis, resulting in 3D acquisitions. Lines must then be acquired from the entire spatial domain in both the k_y and k_z directions, and a 3D inverse Fourier transform performed to obtain the 3D proton density. In this case the generalised signal equation becomes:

$$s(\mathbf{k}) = \int\int \rho(\mathbf{r}) e^{-2\pi i \mathbf{k}\cdot\mathbf{r}} d\mathbf{r} \tag{2.34}$$

with

$$\mathbf{k} = \frac{\gamma}{2\pi} \int_0^t \mathbf{G}(t') dt' \tag{2.35}$$

2.1.2.3 Slice Selection

Typically it is desired to obtain images which are slices through a sample. While phase encoding can be performed in 3D, the entire object must be scanned, which requires sampling of the entire 3D spatial frequency domain, and is not always convenient or necessary. The MR signal can be localised by only exciting thin slices of magnetisation, on which 2D Fourier transform imaging (frequency and phase encoding to acquire a 2D spatial frequency domain data set) is performed. This is achieved by using the linear frequency mapping creating by a linear gradient. If a linear gradient is applied during excitation, and the frequency content restricted of the excitation RF B_1^+ field, only those protons which have a Larmor frequency within the bandwidth of the pulse will be significantly affected, and hence only a narrow slice of magnetisation is excited. For a gradient of amplitude G_z applied in the z direction, the flip angle varies as a function of the Fourier transform of the temporal RF pulse waveform, $B_1(t)$:

$$\alpha(z) = \gamma \int_0^t B_1(t') e^{i\gamma G_z z t'} dt' = \gamma \int_0^t B_1(t) e^{2\pi i \Delta f t'} dt'$$

The slice thickness is given by the bandwidth of the RF pulse, and the slice location is given by its central frequency, the frequency of the RF carrier wave which $B_1(t)$ modulates. Δf is the

frequency offset between the Larmor frequency, f_0 and the RF carrier wave frequency, f_1. An ideal slice profile is that of a rectangle function:

$$\text{rect}(z) = \begin{cases} 1 & z_0 - \frac{\delta z}{2} < z < z_0 + \frac{\delta z}{2} \\ 0 & \text{elsewhere} \end{cases}$$

The spatial frequency content of this slice profile, given by its Fourier transform is that of a sinc function :

$$\text{FT(rect (z))} = \delta z \frac{\sin(\pi f \delta z)}{\pi f} \tag{2.36}$$

As discussed, the application of a gradient creates a direct mapping between spatial frequency and temporal frequency, therefore, by modulating the amplitude of the RF B_1^+ field by a sinc function, and setting the RF carrier frequency accordingly a slice can be excited at an arbitrary position along z. The slice position is given by:

$$z_{slice} = \frac{2\pi\Delta f}{\gamma G_z}$$

And the thickness, TH of the excited slice is related to the bandwidth of the RF pulse, BW_{RF} by:

$$\text{TH} = \frac{2\pi \text{BW}_{RF}}{\gamma G_z}$$

2.1.2.4 k-Space

k-space is a mathematical construct that formalises the concept of acquiring data within the spatial frequency domain. Image space, and k-space are a Fourier pair, that is to say that one can be reconstructed from the other using the appropriate Fourier transform. Each point in k-space refers to a specific spatial frequency, and it is the objective of an MR experiment to obtain a full set of k-space for the excited magnetisation (which may be the whole object or a slab on which 3D phase encoding is performed, or a thin slice on which 2D phase encoding is performed). Gradients are used to traverse k-space, and for a general gradient $\mathbf{G}$, the k-space position is given by its time integral, which is proportional to the gradient area, also known as the gradient moment:

$$\mathbf{k} = \frac{\gamma}{2\pi} \int_0^t \mathbf{G}(t')dt \tag{2.37}$$

A MRI scanner samples the signal digitally, therefore in order to accurately reproduce the k-space of the object and reconstruct an image which is artefact free, data must be sampled according to the Nyquist sampling theorem [151]. This states that in order to accurately reproduce an analog signal which has frequency content up to B Hz, the sampling must occur at a frequency of $2B$,

otherwise aliasing will occur. As k-space is the Fourier transform of the image space, the sampling criteria places implications on the properties of the k-space points that must be sampled. Figure 2.4 shows the relationship between k-space and image space. The field of view (FOV) of the image is inversely proportional to the k-space resolution, Δk, and the resolution of the image is inversely proportional to the amount of k-space that is sampled (or k-space field of view, FOV_k). For the frequency encoding direction this has implications for the rate that the signal is sampled (i.e. the ADC clock speed), and for the phase encoding direction it places criteria for the number of repetitions of the experiment that must be made.

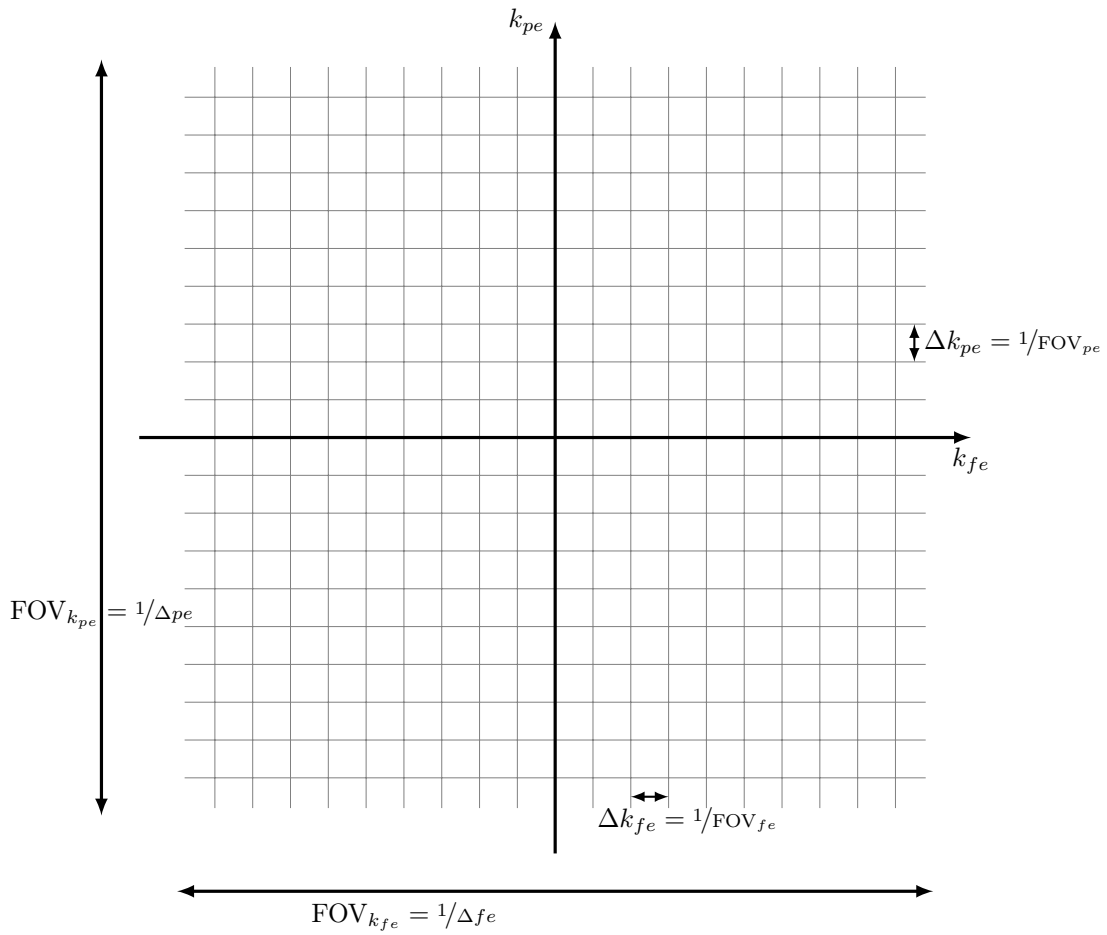

Figure 2.4: 2D k-space

2.1.2.5 Gradient Echo

The gradient echo is used to obtain both positive and negative k values during frequency encoding. The FID experiment with a frequency encoding (also known as the readout) gradient applied immediately after excitation provides a means for spatial localisation, however depending on the polarity of the frequency encoding gradient, only half of the k-space line can be acquired (for positive gradients only the positive half of k-space, and vice-versa). A gradient with polarity opposite to the readout gradient can be used to traverse the opposite direction in k-space, such that during the readout both positive and negative spatial frequencies are sampled. Consider a

rectangular readout gradient waveform of amplitude G_{re} and duration τ. By applying a gradient with moment equal to half of the readout moment, and opposite polarity, for example amplitude $-2G_{re}$ and duration $\tau/4$ before the readout gradient, k-space is traversed "back" by half the amount that is covered in the readout gradient. This is known as a read dephase gradient. When the readout gradient is then applied, at the point at which the gradient moment (the time integral of the gradient) of the readout gradient so is equal in magnitude but opposite in polarity to the moment of the dephase gradient an echo is formed, which corresponds to the centre of k-space.

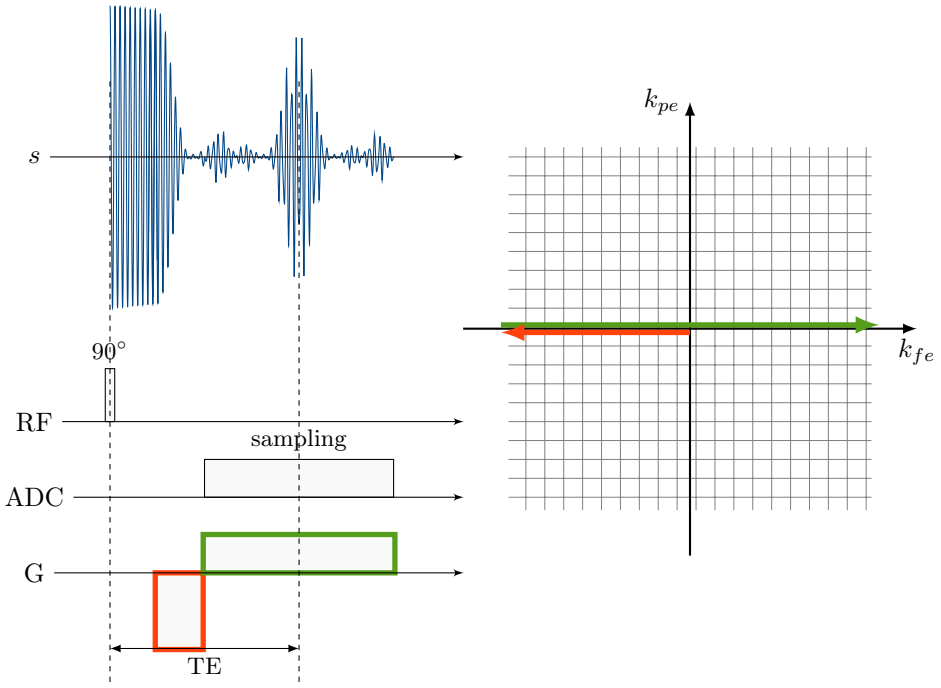

Figure 2.5: 1D Gradient Echo

Figure 2.5 shows a basic 1D gradient echo sequence. Following the 90° excitation a dephase gradient (orange outline) is applied, traversing k-space in the $-k_{fe}$ direction. The transverse magnetisation, and subsequent signal is dephased. A positive readout gradient is then applied, k-space is traversed in the $+k_{fe}$ direction, and the signal is sampled. When the moment of the readout gradient is equal in magnitude to the moment of the dephase gradient, a gradient echo is formed. The time between the excitation pulse and this point is defined as the echo time, or TE, and corresponds to $k = 0$. Following the echo time the remaining half of the k-space line is sampled. A gradient echo sequence is subject to T_2^* decay, and because this is much shorter than T_2 the sequence is typically run with the shortest echo time possible to maximise the available signal.

A 2D gradient echo sequence, which includes a slice selective excitation pulse, phase encoding and frequency encoding is shown in figure 2.6. G_r is the readout gradient channel, G_p is the phase encoding gradient channel and G_s is the slice select gradient channel. The slice select sinc pulse and gradient start at $k = 0$, and move away from the centre of k-space as the magnetisation is tipped. To refocus the magnetisation, a gradient with half the moment and the opposite polarity to the

slice selective gradient is performed afterwards. The slice select refocus, phase encoding and read dephase gradients can all be applied simultaneously because they act in orthogonal, independent directions. To acquire the full set of k-space values, the gradient echo sequence must be repeated with a different phase encoding gradient, acquiring a new k-space line each time, as shown in the k-space diagram in figure 2.6. The combined phase encoding and read dephase gradients move to a $-k_{fe}$ location, as shown in orange, and then with the application of the readout gradient a positive line is acquired, as shown in green. A new pulse sequence parameter, the repeat time, TR is introduced. This is the time between successive excitation pulses, and in this sequence it must be set to be several times longer than the longest T_2 within the sample to ensure there is no transverse magnetisation upon the next excitation pulse. If TR is made to be shorter than several times the longest T_1 within the sample then the magnetisation will not fully relax back to equilibrium, resulting in less available longitudinal magnetisation upon the next excitation pulse, a reduction of signal and an introduction of T_1 contrast to the image (see section 2.1.2.7).

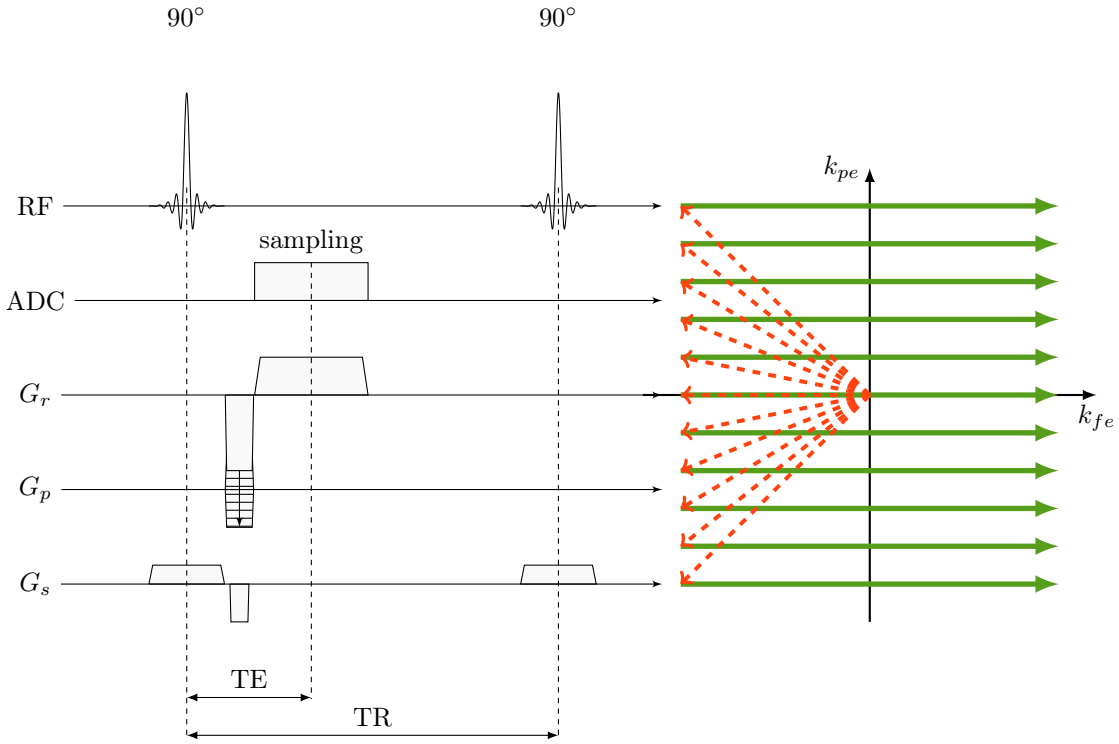

Figure 2.6: 2D Gradient Echo Sequence

2.1.2.6 Spin Echo Imaging Sequence

With the addition of a slice select pulse, and phase encoding and frequency encoding gradients, the global spin echo sequence described in section 2.1.1.9 can also be used to generate an image. A dephase gradient is still used to form a gradient echo, and this is timed appropriately such that it occurs at the same time as the spin echo formed by the 180° RF pulse. The application of the refocusing RF pulse reverses the sign of the phase accrued by the magnetisation, and in k-space

this is equivalent to reversing the polarity of the k vector. A slice selective refocusing pulse is also not necessary for the 180° pulse, because the pulse is inherently self-focusing. The phase encoding and read dephase gradients can occur either before or after the refocusing RF pulse. If they occur before, then their effects are reversed in sign by the refocusing pulse, therefore the gradient amplitudes must be opposite in polarity to those shown in the gradient echo sequence. If placed afterwards then they are not subject to the refocusing pulse, and no change is necessary to their polarity. Figure 2.7 shows a 2D spin echo sequence, with the phase encode and read dephase gradients placed before the refocusing RF pulse. The k-space diagram shows the movement of the phase encode and read dephase gradients with an orange dashed line, the polarity reversal from the refocusing pulse with a blue dashed line, and the readout with a solid green line.

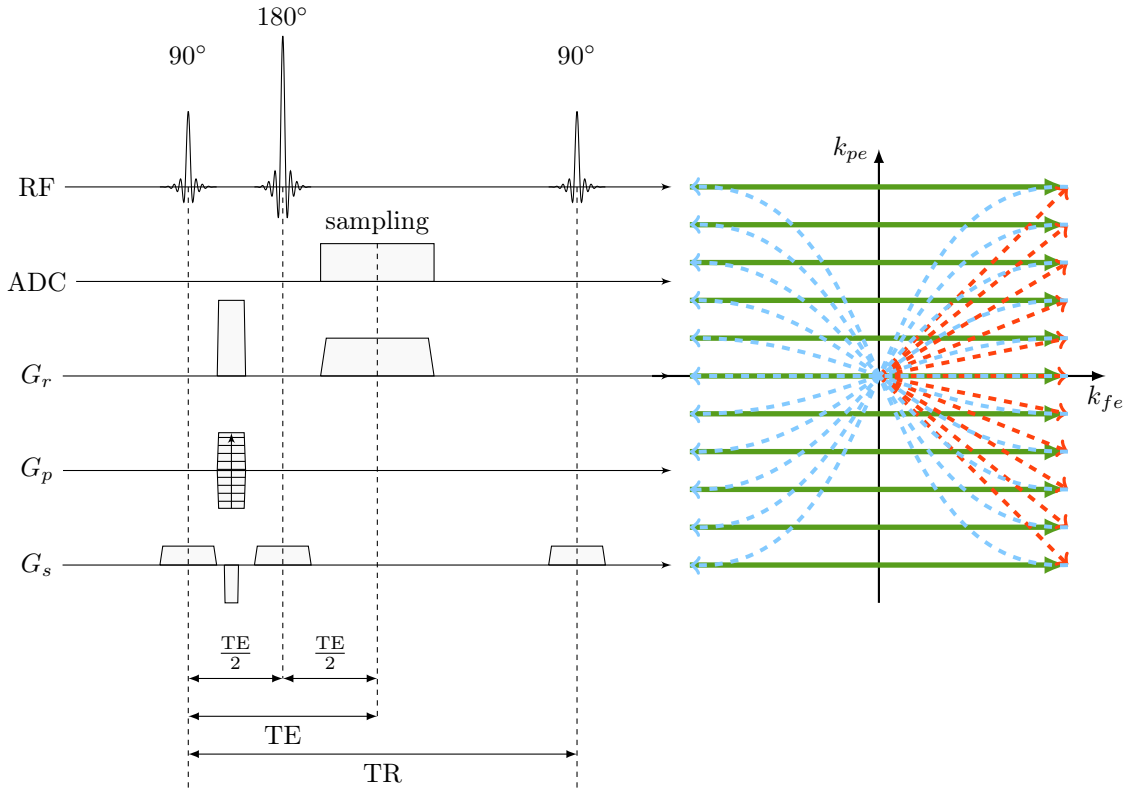

Figure 2.7: 2D Spin Echo Sequence

2.1.2.7 Contrast

MR images can be sensitised to show contrast from differences in physical parameters which affect the MR signal. The primary contrast mechanisms are proton density, T_1 relaxation, and T_2 (including T_2^*) relaxation. Proton density images do not contain much contrast between different tissues, because the proton density is similar for most tissue type. On the other hand, both T_1, T_2^* and T_2 vary with tissue type and certain pathologies.

The signal equation for a spin echo sequence is:

$$s_{se} = M_0 \frac{(1 - e^{-\frac{\mathrm{TR}}{T_1}})}{\left(1 - \cos\alpha e^{-\frac{\mathrm{TR}}{T_1}}\right)} \sin\alpha e^{-\frac{\mathrm{TE}}{T_2}} \tag{2.38}$$

Where M_0 is the equilibrium magnetisation, TR is the repeat time, TE is the echo time, α is the excitation flip angle, T_1 is the longitudinal relaxation time and T_2 transverse relaxation time. The signal equation for a gradient echo sequence is:

$$s_{ge} = M_0 \frac{(1 - e^{-\frac{\mathrm{TR}}{T_1}})}{\left(1 - \cos\alpha e^{-\frac{\mathrm{TR}}{T_1}}\right)} \sin\alpha e^{-\frac{\mathrm{TE}}{T_2^*}} \tag{2.39}$$

T_2^* is the transverse relaxation time which includes dephasing due to external inhomogeneities. Consider a sample consisting of two different tissue types placed within a MR scanner. For simplicity tissue one has longer T_1, T_2 and T_2^* than tissue two. T_2 contrast for a spin echo sequence, and T_2^* contrast for a gradient echo sequence can be obtained by choosing TE such that the transverse magnetisation from tissue two has decayed significantly more than tissue one, as shown by the dashed line in figure 2.8.a. For such an image, the signal intensity from tissue one will be brighter. T_1 contrast can be introduced by choosing a TR value such that tissue one and tissue two relax to different longitudinal magnetisations, as shown by the dashed line in figure 2.8.b. Because the amount of available magnetisation differs, once tipped into the transverse plane the corresponding signals differ, resulting in contrast between the two tissues within the image. Proton density contrast is obtained by using a short echo time, and long TR such that no T_1 or T_2 contrast is imparted.

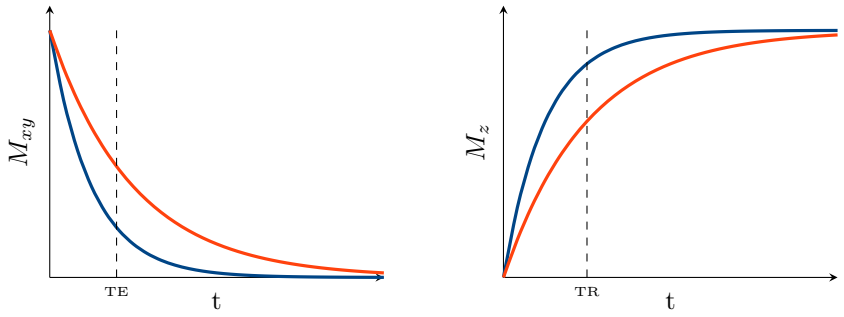

a. T_2-weighting b. T_1-weighting

Tissue 1: Orange
Tissue 2: Blue

Figure 2.8: T_1 and T_2 Contrast

2.1.2.8 Spoiling

To facilitate faster imaging, where it is not possible to wait for the transverse magnetisation to fully decay, or when using magnetisation preparation pulses which affect the longitudinal magnetisation such as an inversion recovery sequence [91, p.364], it is desirable to destroy, or "spoil" any coherent transverse magnetisation. If this is not done, both the longitudinal and transverse components of

the magnetisation will be subject to the excitation pulses, and multiple echo pathways will be formed, resulting in the wrong image contrast and potential image artefacts. Analogous to the manner in which a 90° and 180° pulse pair produces a spin echo, any combination of RF pulses can form a partial echo, known as a Hahn echo [153, p.239]. In general, an RF pulse will flip longitudinal magnetisation into the transverse plane generating a FID, flip transverse magnetisation longitudinally which later pulses can refocus, and refocus existing transverse magnetisation into an echo.

Gradient spoiling is the application of a gradient to purposefully dephase the magnetisation within a voxel by forcing the phase to vary by at least 2π across the voxel, resulting in a transverse magnetisation vector sum of zero. With a sufficiently large gradient, varied in a linearly incremental or pseudorandom manner from one excitation to another, only magnetisation components with zero phase prior to each RF pulse will add coherently and contribute to the detected signal [42, 264, 140, p.297].

Gradient spoiling is used to effect in the spoiled gradient echo sequence, also known as the FLASH (Fast Low Angle SHot) [92] sequence, where in conjunction with a low flip angle, it permits TR $< T_2$ resulting in an the acquisition of an entire 2D image in a few seconds. Over the course of multiple sequence repetitions the longitudinal magnetisation reaches a steady state value that is less than the equilibrium magnetisation. Therefore, for correct contrast it is important that enough repetitions of the sequence are run prior to data acquisition. A 3D FLASH sequence is shown in figure 2.9. The slice selective sinc pulses are set to excite a large volume, rather than thin slices in a 2D sequence. Phase encoding now occurs on both the phase and slice gradient axes (G_p and G_s respectively), with the G_s phase encoding gradient coinciding with the slice select refocusing gradient. Following the readout and acquisition of a k-space line, spoiling gradients are applied in both the phase and slice directions. In addition, gradients equal but opposite in sign to the phase encoding gradients are shown, this is representative that the spoiling gradient must take into account the effects of phase encoding. In practise this can either occur by combining individual gradient waveforms, or from calculating the total gradient moment required to perform both these operations and designing a suitable gradient.

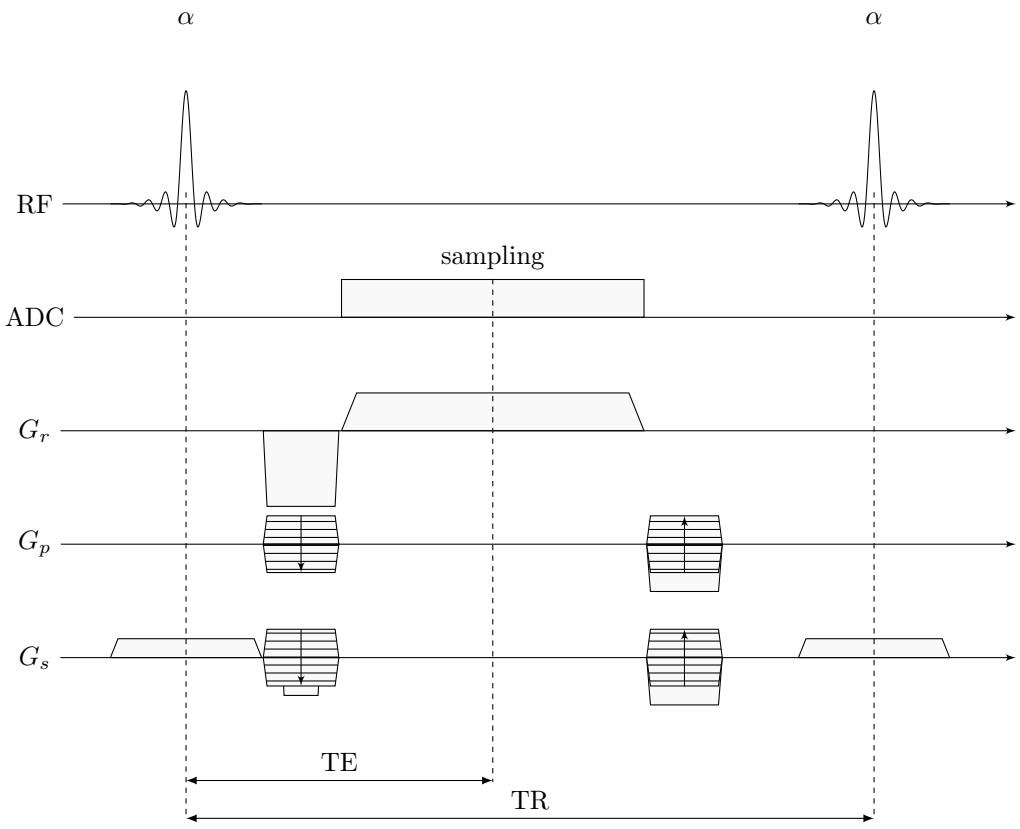

Figure 2.9: FLASH Sequence

The phase shift introduced by the spoiler gradient can also be achieved by incrementing the RF transmitter phase, ϕ by a linearly increasing amount, resulting in quadratically increasing phase. This technique is called RF spoiling. For the k^{th} pulse, where the base phase increment is $\Delta\phi$ the transmitter phase is given in equation 2.40. In [264], adequate RF spoiling for all values of T_1/TR and T_2/TR encountered clinically was found for $\Delta\phi$ of 117° or 123°.

$$\phi_k = \phi_{k-1} + k\Delta\phi \tag{2.40}$$

2.1.2.9 Echo Planar Imaging

Echo Planar Imaging (EPI) [171] is an MR technique which uses rapidly switching gradients to acquire multiple gradient echoes after a single excitation pulse, making it possible to acquire an entire image in under 100ms. Because it offers such speed, it has facilitated methods which require high temporal resolution, such as diffusion imaging [18], perfusion imaging (see section 2.3), BOLD functional MRI (fMRI) [136, 164, 16], and capturing motion, for example within the heart.

As shown in section 2.1.2.5, a gradient echo is formed from the dephasing and then rephasing of the transverse magnetisation. Following the readout gradient, if another readout gradient of opposite polarity is applied, a second echo can be formed. However as there is now a negative gradient the k-space trajectory will be in the $-k_{fe}$ direction, and the signal from the second echo

is lower in intensity due to T_2^* decay. By applying a phase encoding gradient which moves up by one Δk_{pe} between these two readout gradients, two successive lines in k-space can be acquired from a single excitation. Repeating this process forms an echo train, making it possible to acquire a full set of 2D k-space lines from a single excitation, this is known as single-shot EPI. Figure 2.10 shows a gradient echo blipped EPI (GE-EPI) pulse sequence and its accompanying k-space trajectory. Prior to the start of the echo train, the read dephase and phase encoding gradients applied move to the negative extent of the k-space that is to be acquired. Then during the echo train the readout gradients move in the positive and negative k_{fe} direction (shown as green lines), with small phase encoding gradients known as blips occuring between each readout gradient, incrementing in the $+k_{pe}$ direction (shown as orange lines). The gradient echo forms at the centre of k-space, which corresponds to the centre of the echo train. Other acquisition trajectories are possible, for example constant phase-encode EPI has a constant phase encoding gradient applied during the echo train, resulting in a zigzag k-space trajectory. This samples in a non-cartesian manner, consequently requiring re-gridding of the data before an inverse Fourier transform can be applied. Spin echo EPI (SE-EPI) is also possible by positioning a 180° refocusing pulse half way between the excitation pulse and centre of the echo train, and 3D EPI acquisitions can be made by phase encoding on the slice direction and acquiring an entire partition with the echo train instead of a single slice.

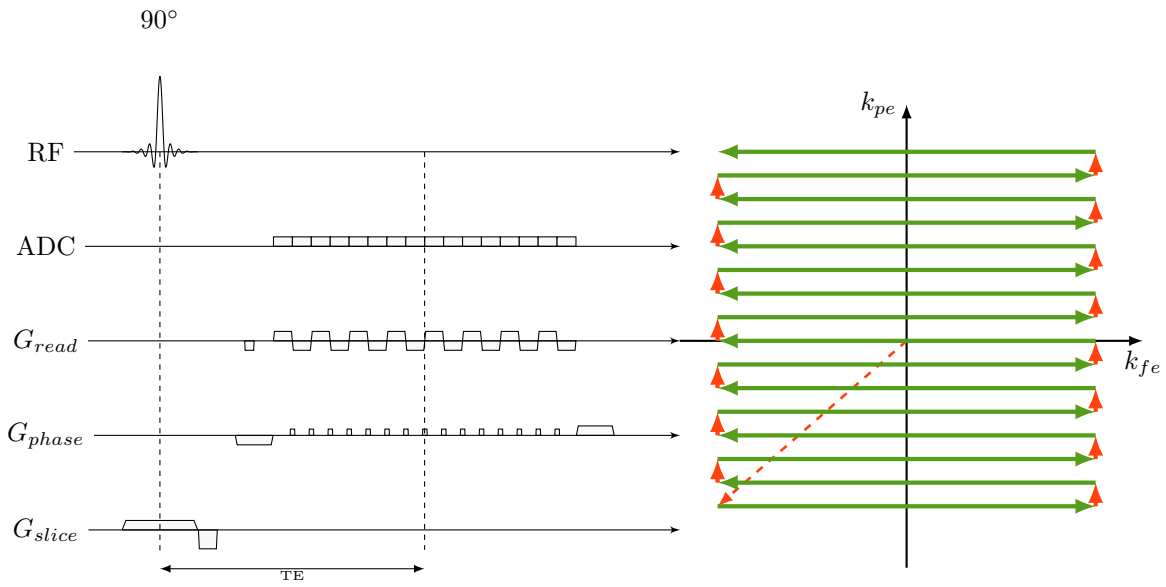

Figure 2.10: Gradient Echo EPI

Due to the long echo time used in GE-EPI (typically 30–60ms), T_2^* decay can cause signal dropout in regions with susceptibility gradients, such as tissue–air interfaces and also from gradient switching induced eddy currents. Because of its refocusing pulse SE-EPI does not suffer from signal dropout, however it generally requires a longer echo time. Susceptibility gradients acting during the

readout can be misinterpreted as being due to phase encoding blips, resulting in image distortion. EPI can also be multi-shot, or segmented, where a portion of k-space is acquired for each excitation pulse, resulting in a shorter echo time and less T_2^* decay for GE-EPI.

High performance hardware, particularly gradients, is required for EPI. Gradient switching must be very fast (see section 2.1.3.2), and while limits are placed due to peripheral nerve stimulation [71], this can result in the induction of eddy currents within large conductors within the magnet. The large electrical currents flowing through the gradient coils dissipate heat, potentially leading to instabilities in passive shims (see section 2.1.3.1). Very precise sequence timing is necessary to ensure that the positive and negative echoes both pass through the centre of k-space at the same time; a mismatch produces a periodic shift between positive and negative samples in k-space and results in a "$N/2$" ghost [153, p.328]. Furthermore, the EPI acquisition uses a very large signal bandwidth in the frequency encoding direction, resulting in high sensitivity to RF noise (also known as spikes), and a very small signal bandwidth in the phase encoding direction, resulting in a physical shift from the fat signal, requiring mandatory fat suppression in EPI, and image distortions from small magnetic field inhomogeneities.

2.1.3 MRI Hardware

MRI requires three different types of magnetic field: a homogeneous static magnetic field (B_0), switched spatially varying magnetic fields (the gradients), and a radio frequency magnetic field (B_1). It also requires a detector coil, a means to demodulate and sample the acquired signal, and a computer to reconstruct the images and control the acquisition. This section describes the equipment that makes up a typical MR system.

2.1.3.1 Main Magnet

The main magnet is designed to produce a strong, homogeneous magnetic field, which the sample is placed in during an MR experiment. It must be large enough for the sample to be placed wholly inside and provide a region which is highly homogeneous at the centre of the magnet (known as the isocentre). Typically a solenoid magnet is used, because it provides a homogeneous magnetic field inside its bore, and its shape allows a sample, or person to be easily inserted inside. Magnets can be made from permanent magnets, resistive (i.e. copper wire) electromagnets, or superconducting electromagnets. Because of the large magnetic field strengths (normally $>$0.5T, with standard clinical field strengths of 1.5 and 3T) required for good signal-to-noise ratio (SNR), superconducting magnets are widely used as they do not lose energy once energised (unlike a resistive magnet, in which the current flowing around the windings dissipates power as heat). Superconductivity is a property that certain materials have when cooled close to absolute zero ($-273°C$): below their critical temperature, critical current and critical magnetic flux density, the

resistivity of the material becomes zero, making them capable of carrying large currents without dissipating any energy. MR magnets are typically wound with Niobium Titanium wire[4], cooled to 4.2K by immersion in liquid helium.

Though well engineered, the homogeneity of the magnet alone does not meet the requirements of MRI, therefore shims are used to improve homogeneity. There are two types of shims in a typical MR system: fixed shims, which are set when the scanner is installed/during maintenance, and dynamic shims, which are adjusted on a per-patient basis to correct for the inhomogeneities and distortions introduced by different body shapes and sizes. The fixed shims can either be passive, or superconducting. Passive shims are pieces of ferromagnetic material (such as iron) inserted into the inside of the magnet bore. Superconducting shims are electromagnets wound from the same superconducting wire and are also bathed in liquid helium. Passive shims are not normally cooled and are prone to heating up under heavy gradient usage, resulting in unstable shim conditions. Dynamic shims are typically resistive electromagnets which are driven with direct current (DC), shimming the magnetic field. Typically the shims are arranged as such that they produce spherical harmonics [160], which are then used to improve the shim over the imaging volume.

The magnet is sited within a room that is electromagnetically shielded from the outside world. This is achieved by lining the walls with copper sheeting. In addition the magnet room is surrounded by thick iron or steel plates to reduce the extent of the fringe fields from the magnet, which would otherwise interact with ferrous objects and equipment in neighbouring rooms. All electrical signals must pass through shielded (and sometimes filtered) connectors in a panel known as the filter plate.

2.1.3.2 Gradients

Magnetic field gradients are created with gradient coils; sets of resistive coils which are designed produce a linear spatial variation in three orthogonal directions. For a solenoid magnet, where the main magnetic field, and hence the z-axis runs along the axis of the bore (i.e. $\mathbf{B_0} = B_0\hat{\mathbf{z}}$) the gradient coils are arranged to produce gradients in the x, y and z directions. Optimal gradient designs minimise the resistance and inductance of the gradient coils, permitting fast switching. Gradients switching occurs in the audio-frequency (AF, kHz) range, and requires high power amplifiers with fast slew rates, particularly for sequences such as echo planar imaging.

2.1.3.3 RF Coils

Radio frequency coils are used to create the B_1^+ field (known as transmission), and to detect the weak B_1^- field generated by the precessing magnetisation (known as reception). To maximise efficiency and signal-to-noise, the coils used for transmission and reception are tuned circuits, designed and built to resonate at the specific Larmor frequency of the MRI scanner and imaging nucleus. A

[4]Normally filaments of NbTi are held within a copper or aluminium wire matrix.

tuned circuit is achieved using a combination of the intrinsic inductance of the coil, and added capacitance such that the inductive and capacitive reactance exactly cancel at a specified frequency. This minimises the losses within the coil, making it extremely efficient at carrying current; at the resonant frequency a large RF current can efficiently flow to produce an intense B_1^+field, or a weak B_1^- will induce a emf and subsequent RF current with good sensitivity. To maximise power transfer to and from the coil, it must also be matched to present the same impedance as the connecting cables, and the RF power amplifier or RF pre-amplifier, 50Ω being standard[5]. This is achieved by using a matching network, which transforms the resistance of the coil (on the order of 1Ω when loaded by a sample) to the system characteristic impedance specifically for the resonant frequency of the coil. Both tuning and matching are affected by the sample, therefore coils are constructed such that either the tune and match can be adjusted, for example using variable capacitors; or if the sample loading is likely to be similar as is the case when a coil is used only for a specific use, they are tuned and matched to fixed, optimal values during construction. Minimal loss is essential for efficient transmission, or good SNR detection, and RF coils must not contain any magnetic materials, which can create image distortions, therefore specialist low loss, non-magnetic RF components are used.

While the principle of reciprocity [110, 109] shows that a good transmitter coil makes a good receiver coil, and initially a single coil was used for both, over the past two decades there have been many developments that has led to the divergence of transmit and receive coil technology (for example the NMR Phased Array [194]) and as a result modern MR systems have dedicated coils for each. RF Coils are described in more detail in section 2.2.

Transmit Coils The primary goal of a transmit coil is to produce a uniform B_1^+ of suitable strength over the entire imaging volume, resulting in homogeneous excitation. This is achieved using a volume coil, such as Birdcage [97] or Transverse-Electro-Magnetic (TEM) [223] resonators. Both of these coil designs permit quadrature excitation, and produce a homogeneous B_1^+ field within the centre of the coil, with minimal electric fields (only present at the periphery of the coil). RF homogeneity is affected by the sample, and gets increasingly worse at higher frequencies (corresponding to higher field strengths), for human imaging this occurs at field strengths of 3T and above due to increased power deposition[6] and wavelength associated effects[7]. For ultra-high field

[5]This impedance is known as the characteristic impedance, and it is the ratio of voltage to current on a transmission line (a special type of cable designed to carry alternating currents which have a high enough frequency/short enough wavelength in relation to the length of the cable that their wave nature must be taken into account; the RF signal in MRI is within this regime even for short cables). At interfaces between equipment and cables, if their is a mismatch in impedances reflections occur, resulting in non-optimal power transfer. By ensuring that all devices have the same output or input impedance as the characteristic impedance of the cables that connect them, reflections are minimised and power transfer maximised. Transmission lines and characteristic impedance are explained in more detail in section 2.2.1.6.

[6]The incident RF field generates eddy currents within the sample, which themselves produce a RF magnetic field that opposes the incident field. Consequently the field is attenuated with depth.

[7]Within biological tissues the electric permittivity is much higher than in free space or air, resulting in shorter wavelengths. If the sample or anatomy being scanned is of similar dimensions to a multiple of $\lambda/4$ (where λ is the wavelength within the material), a resonance condition can occur (known as dielectric resonance), focusing of the RF field and resulting in highly inhomogeneous excitation and increased localised power deposition.

(7T and above) systems, the RF inhomogeneity is so poor that new transmitter designs have been necessary, that use multiple, localised transmitter coil elements (rather than a single resonator), which can be individually shimmed by adjusting their relative amplitudes and phases [115] (see section 2.1.4). A spatially confined transmit field is sometimes desirable, and this can be achieved using a surface coil (a single turn solenoid, tuned to resonate at the Larmor frequency).

When a sample is placed within a RF coil, it interacts with the generated field, causing a perturbation to the coil's operation in comparison to an unloaded coil. Capacitive (electric field) interactions cause a shift in the resonant frequency of the coil, whilst the inductive interaction with a lossy conductor, as most biological materials are, results in induced eddy currents within the sample, which dissipate power into the sample. This increases the resistive losses in the coil, and is known as "loading". The power required to generate a prescribed B_1^+strength is affected by the coil loading. Typical volume transmitters used in whole body human MR systems require many kilowatts of peak power, resulting in large voltages across the capacitors, and large currents through the conductors. Consequently high power RF components are necessary to cope with these conditions. The power deposited within the sample is also of concern as it can cause an increase in temperature, which for humans or animals would result in burns or a dangerous increase in the core body temperature. Power deposition is quantified by the specific absorption ratio, SAR, which is defined as the power absorbed per unit mass. Regulatory bodies such as the IEC and FDA place limits on the maximum allowed SAR, which all human MR systems must adhere to. More information about SAR is provided in chapter 6.

Receive Coils Receive coils are used to detect the weak B_1^- field. An emf is induced within a coil by the changing magnetic flux due to the precessing transverse magnetisation, this signal must then be transferred to the receiver instrumentation for amplification, demodulation and sampling. The goal of a well-designed receive coil is to detect the MR signal from the sample over the volume of interest, with minimal introduction of noise. Thermal noise from the coil, noise from the interfacing electronics and noise from the sample (due to coil loading and physiological noise) all contribute to the noise in an MR image. For human imaging, where the coils are large in comparison to the RF wavelength, sample losses dominate, with noise being generated from all of the sample that is within the coil's field of view (FOV). A volume coil, such as a birdcage resonator can be used as a receive coil, providing good homogeneity. However the detected signal is also subject to noise from all of the sample volume that is inside the coil. On the other hand, a surface coil provides high SNR at a depth equivalent to the diameter of the coil, however the FOV is spatially localised to the size of the surface coil.

A third option is to combine multiple surface coils into a phased array [194, 35], retaining the high SNR available by using surface coils, but extending the FOV. Each coil element is interfaced to an independent receive channel, and the individual images from each are then combined in

an optimum manner, providing the coverage of a volume coil with the SNR of a surface coil. In order to function correctly and yield high SNR images, the noise between each coil element must be uncorrelated, therefore the coil elements must be electromagnetically decoupled. Neighbouring coils, which exhibit the strongest coupling, are commonly decoupled by overlapping them such that their mutual inductances cancel and RF currents on one coil do not induce RF currents on the other. Other coil elements within the array are decoupled from one-another using a low input impedance preamplifier to resonate the coil's matching network, resulting in a large blocking impedance at the coil's terminals, minimal RF current flow and hence no noise or MR signal coupled between elements. Receive arrays are also suitable for parallel imaging: a method that utilises multiple RF receiver coils with spatially different receive sensitivities to perform some of the spatial encoding, such that a reduced set of k-space lines can be acquired. Theoretically it is capable of reducing scan times by a factor equal to the number of independent receive coils, however, SNR is reduced from system imperfections, and in practise acceleration factors of 2–4 are typically used. Parallel imaging can be performed both in the image domain and in k-space. SENSE (SENSitivity Encoding) [186], is the major image domain method, and utilises sensitivity maps of the receive coils to unfold aliased images acquired when the acquisition is undersampled. Conversely, GRAPPA (GeneRalised Auto calibrating Partially Parallel Acquisitions) [88], a k-space method, does not require the acquisition of full receive sensitivities, instead it acquires some additional k-space lines to reconstruct skipped k-space lines.

2.1.3.4 Transmitter System

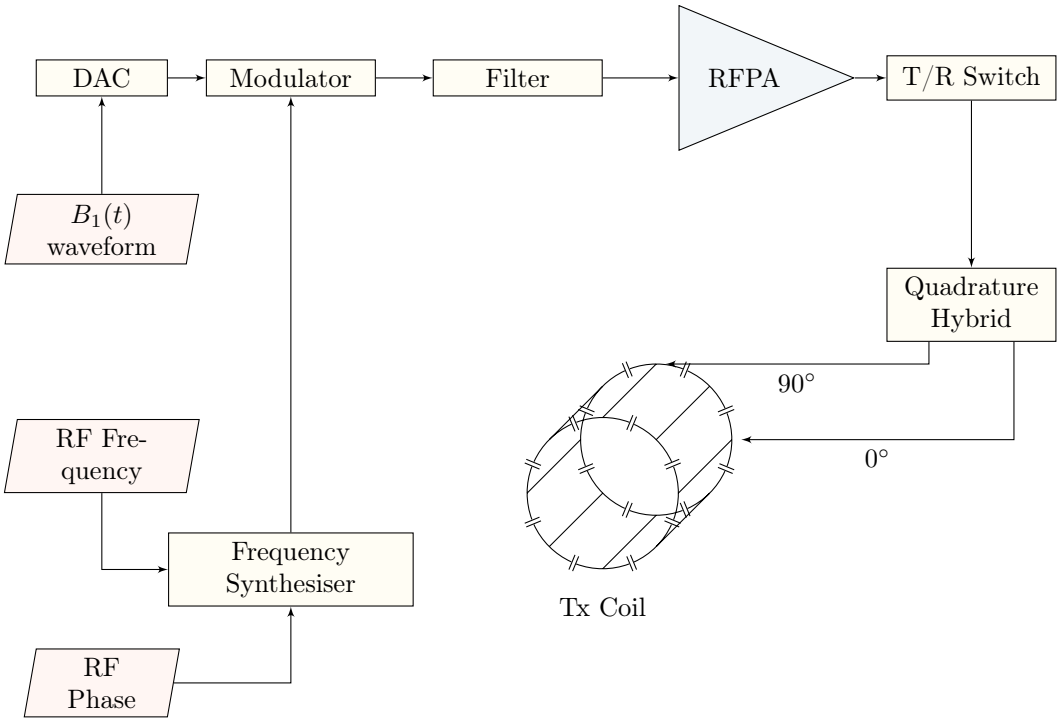

Figure 2.11: Transmitter System

The role of a transmitter system is to generate the amplitude and phase modulated RF signal, then drive the transmitter coil with an amplified version of this signal. Figure 2.11 shows a typical transmitter system that is used on many current MR systems[8]. The RF amplitude, frequency and phase waveforms begin as digital representations residing in the memory of the control computer. When an RF pulse is run, the correct RF carrier wave frequency and phase is generated by a frequency synthesiser, and then modulated in amplitude by the amplitude waveform. Unwanted frequencies generated by the digital to analog converter's system clock are filtered before the signal is amplified by the RF power amplifier (RFPA). Following the RFPA the amplified signal is gated by a T/R switch to prevent noise from the RF amplifier (a result of thermal noise at the input being amplified) reaching the transmitter coil when it is not in use, and if the transmitter coil is also used as a receive coil it provides a way to switch between the transmitter system and receiver system without damaging the sensitive pre-amplifiers during transmit. The high power RF is then split by a quadrature hybrid into two signals with half the power of the original, but with a phase shift of 90° between the two. On some systems the quadrature hybrid and T/R switches are reversed in order (requiring two T/R switches, and a separate quadrature hybrid for the transmitter and receiver). These two signals are then used to drive a quadrature transmitter coil, such as the birdcage coil shown in figure 2.11.

[8]As technology has progressed, the analog signal processing that was used predominantly in the late 1970's has shifted to the digital domain, providing improved performance, scalability, processing power and flexibility. For transmitter systems this has meant that the frequency synthesis, and the amplitude and phase modulation has moved to the digital domain, with DACs directly generating the RF waveform.

2.1.3.5 Receiver System

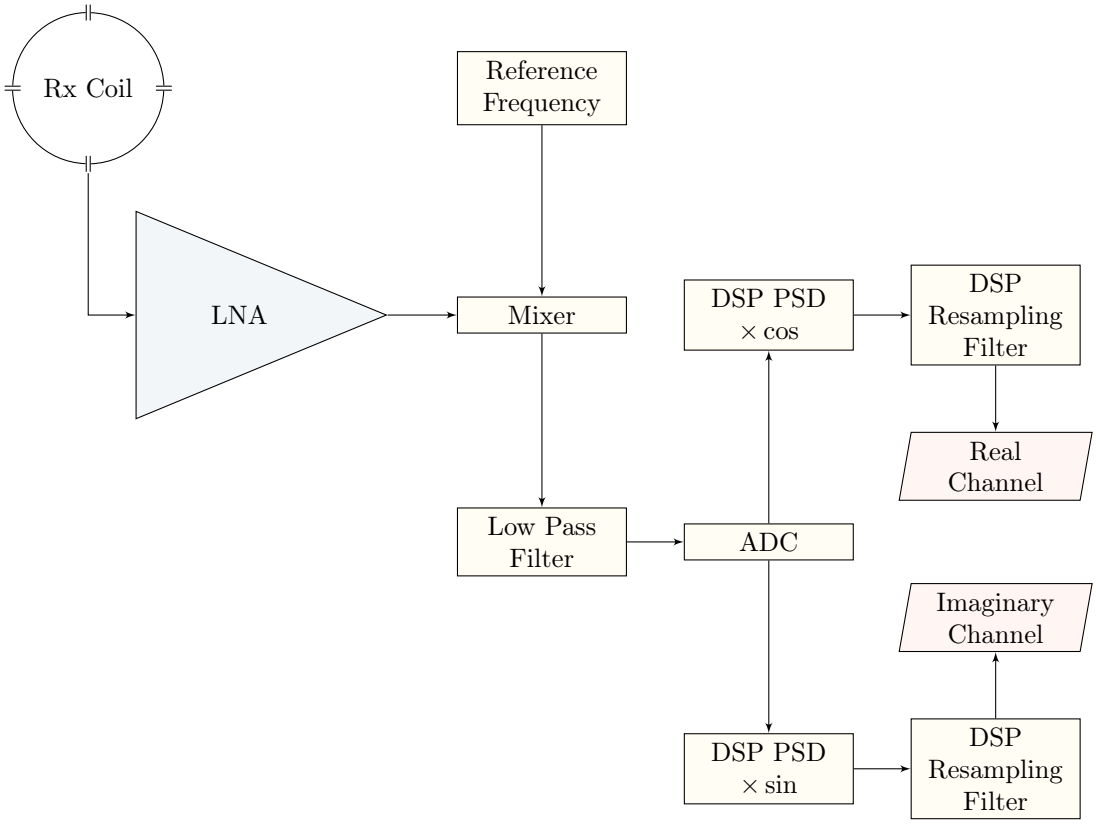

Figure 2.12: Receiver System

The receiver system amplifies, demodulates, filters and samples the signal from the receive coil(s), so that it is ready to be reconstructed into an image. Figure 2.12 shows the main features of a typical receive channel. Modern MR systems have multiple channels, facilitating the use of phased arrays and parallel imaging. Signal from the receive coils is very weak and must be amplified to a useful level so that further processing does not add large amount of noise. Extremely low noise figure ($\approx$ 0.5dB) pre-amplifiers (LNA) are crucial for the first gain stage, and tend to be placed close to, if not at the coils to boost the signal and reduce losses from long lengths of cable. The signal is then demodulated by multiplying with a reference RF signal. Historically this was set to be the Larmor frequency; demodulating the MR signal to an audio frequency signal centred around DC, however because of improvements in digital technology, modern systems use a reference frequency that is lower than the Larmor frequency, resulting in a modulated "intermediate frequency" (IF) signal, which can range anywhere from approximately 125kHz [107] to many MHz. The IF is then sampled, typically with an ADC running at four times the IF such that phase sensitive detection (PSD) can be implemented as a digital filter[9] using digital signal processing (DSP). The real and imaginary signals are then resampled to the sampling rate required by the imaging parameters (specifically the frequency encoding direction's FOV) using a second DSP digital filter.

[9]With a digital clock at 4× the IF, phase sensitive detection is implemented digitally multiplying by a sine and cosine at the IF, which corresponds to multiplication with +1, 0, -1, 0... for the real channel and 0, +1, 0, -1... on the imaginary channel.

2.1.4 Parallel Transmission

Parallel transmission utilises multiple RF transmitter coils/elements to facilitate some degree of spatial control over the total transmitted field. It was developed to overcome the RF inhomogeneities observed at high field strengths due to increased RF absorption and dielectric resonances [111]. It requires a multi-element transmit array [219], and multiple transmitter channels to drive each element independently.

Initially all coil elements were driven with the same the same temporal B_1 waveform, but different amplitude weightings and phase shifts, a method which has been termed "RF shimming" [116]. For a set of R transmitter coils, each being driven by a current waveform $I(t)$ with amplitude weighting a_r, phase shift ϕ_r and transmit sensitivity (B_1^+ per unit current for the coil element) $\mathcal{B}_r^+$, the total B_1^+ field is the vector sum of these the transmit field from each coil:

$$B_1^+ = \sum_r^R a_r e^{i\phi_r} I(t) \mathcal{B}_r^+ \tag{2.41}$$

i is the imaginary unit ($\sqrt{-1}$). RF shimming is usually performed with the goal of homogenising the B_1^+ field. If $\mathcal{B}_r^+$ is spatially discretised into M samples, equation 2.41, with the time dependent term $I(t)$ removed can be formulated as a matrix equation:

$$\mathbf{b_1} = \mathbf{Sc}$$

Where $\mathbf{b_1}$ is a M element column vector containing spatial samples of the total B_1^+ per unit current, $\mathbf{S}$ is a $M \times R$ element matrix formed by the concatenation of vectors of the spatial samples of each $\mathcal{B}_r^+$ (see section 3.2.1 and figure 3.12 for further details on this operation), and $\mathbf{c}$ is a R element column vector with each element $c_r = a_r e^{i\phi_r}$, representing the complex driving coefficient for the r^{th} coil. By specifying a desired $\mathbf{b_1}$, $\mathbf{c}$ can be obtained by solving the inverse problem [14]. For RF homogeneity, $\mathbf{b_1}$ is set to 1, and the problem is typically solved in the least squares sense:

$$\mathbf{c} = \arg\min \left\{ \|\mathbf{S} \cdot \mathbf{c} - 1\|^2 \right\} \tag{2.42}$$

A more generalised version of parallel transmission, "Transmit SENSE" [124, 125, 257, 87], utilises different complex driving waveforms for each transmit coil, which combined with an excitation k-space trajectory permits acceleration of multi-dimensional RF pulses, analogous as the name suggests to the acceleration of acquisition when using parallel imaging methods such as SENSE [186]. Transmit SENSE is capable of producing excitation pattens that are not limited by Maxwell's equations, whereas RF shimming uses the interference between the B_1^+fields created by each coil. The complex waveforms for each transmitter channel are typically calculated in an iterative manner, and can also include terms which account for B_0 inhomogeneities or that minimise SAR [87]

Parallel transmit requires extensive additional hardware, most noticeably multiple transmitter channels [228]. For RF shimming where a_r and ϕ_r are time independent, implementations have included using a single RF amplifier, fixed power splitter and adjustable phase shifters [155]; splitting the low power RF output from an MR console, then using vector modulators for amplitude and phase control before amplifying with individual RF amplifiers [249]. Transmit SENSE typically requires an individual transmitter channel for each element, enabling full waveform control of the amplitude and phase of the RF. Furthermore, precise gradient and RF coordination is warranted, also meaning that it is desirable for these transmitter channels to be a part of the main MR system and not an additional module.

Both methods require all of the RF coils to be sufficiently decoupled from each other. Multiple circuits which are resonant at the same frequency can strongly couple energy to each other. Not only does this perturb the properties of the resonant circuits (such as resonant frequency and loss), for MRI transmit coils it means that the B_1^+ field created by one coil can induce an RF current within another coil, interfering with the driving current such that the desired current amplitude and phase is not realised. Decoupling is necessary to ensure that the B_1^+ of each coil can be precisely controlled. Methods to decouple transmit coils include: overlapping coils to minimise their mutual inductance[194], capacitive [183] or inductive [245] networks, cartesian feedback [112, 113], digital vector decoupling [127, 199, 225], current source RF amplifiers [135], and ultra-low output impedance RF amplifiers [40].

With regards to the work in this thesis, the term "parallel transmission" or "parallel transmit" specifically refers to RF shimming methods (modifying the static amplitudes and phases of each RF channel to spatially tailor the B_1^+field).

2.2 Radio Frequency NMR Coils

In NMR, Radio Frequency (RF) coils[10] are employed to excite the nuclear spin system, by applying an oscillating transverse magnetic field at the Larmor frequency, manipulating the direction of the macroscopic magnetisation from its thermal equilibrium pointing along B_0, and also to detect the transverse component of the macroscopic magnetisation. For transmission a current is required to create a transverse magnetic field, as given by Ampère's Circuital Law:

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J} + \mu_0 \varepsilon_0 \frac{\partial \mathbf{E}}{\partial t} \tag{2.43}$$

in differential form (where $\mathbf{J}$ is the current density), and

[10]Coil refers to the solenoid coils used in the early days of NMR, however for practical reasons they have been superseded by other forms of NMR probe, most of which are not coil shaped at all. Nevertheless the word coil has remained.

$$\oint_{\partial S} \mathbf{B}\cdot d\mathbf{l} = \mu_0 I_S + \mu_0 \varepsilon_0 \frac{\partial \Phi_{E,S}}{\partial t} \tag{2.44}$$

in integral form (where ∂S is the boundary of surface S and $\Phi_{E,S}$ is the electric field flux passing through it). The second term in both equations is Maxwell's famous displacement current term, from which self-propagating electromagnetic waves were predicted, however for the case of transmission, where the aim is to solely produce a magnetic field they can be ignored. Reception of the signal is achieved because a time-varying magnetic field generates an electromotive force (EMF) in a closed circuit equal to the rate of change of the magnetic field flux which passes through the circuit; Faraday's Law of Induction:

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \tag{2.45}$$

in differential form, and in integral form:

$$\oint_{\partial S} = \mathbf{E}\cdot d\mathbf{l} = -\frac{\partial \Phi_{B,S}}{\partial t} \tag{2.46}$$

Where again ∂S is the boundary of surface S, and $\Phi_{B,S}$ is the magnetic field flux passing through it. More information about the derivation and application of Maxwell's Equations can be found in an electromagnetic theory text such as Electricity and Magnetism by E.M. Purcell [187]. Through the application of the principle of reciprocity in an NMR context [110, 109], which states that the signal induced in a coil by a precessing magnetic moment is proportional to the strength of a hypothetical field created by the same coil carrying unit, the qualities that make a coil a good transmitter, also make it a good receiver and vice versa.

NMR coils are designed to transmit and receive radio frequency magnetic fields. Apart from ultra high field (7T and above) human MR systems, this occurs solely in the near field regime [108] as coil and sample sizes are smaller than the wavelength of the magnetic fields. Coils are built to be electrically resonant at the Larmor frequency for a specific nuclear spin species at a specific field strength. At resonance the impedance in a coil is minimised, enabling efficient driving of large currents during transmission, and efficient conversion of the EMF induced by the precessing transverse magnetisation into a measurable RF signal during reception.

2.2.1 Electricity at Radio Frequencies

A water pipe analogy is often used to describe the behaviour of electricity. Electric current being analogous to the flow of water in a pipe, voltage the backing pressure, and resistance something which obstructs the flow (such as a narrowing of a pipe). While useful for DC and low frequencies, this concept breaks down as frequency increases to the RF domain, where the size of conductors and components become comparable to the wavelength, hence the wave nature of the electromagnetic

signal must be taken into account. The electrical length of a conductor is defined as the ratio of the conductor length to the wavelength of the signal frequency:

$$\ell = \frac{L}{\lambda} \tag{2.47}$$

In general, any device whose electrical length is greater than approximately $1/20$ can be considered to be "electrically long", therefore the full wave nature of the signal must be taken account [198, p.8]. Devices with electrical lengths less than $1/20$ are "electrically short", and basic circuit theory can be used to describe their behaviour. At low frequencies the current and voltage distributions along a conductor can be considered to be uniform. However, if a conductor is electrically long, the signal must be considered as an electromagnetic wave propagating within, therefore its intensity varies sinusoidally with position, and can be reflected at discontinuities.

Two types of electromagnetic fields need to be considered: reactive (also termed storage or near) fields and radiative fields. Reactive fields are created in the vicinity of a source (with intensity spatially decreasing at a rate quicker than $1/r^2$), do not persist when the source is turned off, stores energy in the field, and can be exclusively electric, magnetic or a combination of the two. For example, a wire carrying a DC current creates a magnetic field around the wire, which is a reactive field. Energy is stored within the field when the current is initiated, and when the current is stopped the energy returns to the wire. In contrast a radiative electromagnetic field propagates away from a source, with its intensity spatially decreasing with rate $1/r^2$, persists even when the source has ceased radiating, and has both electric and magnetic components. A further distinction is that the external absorption of energy from a reactive field has an effect on the source, whereas absorbing energy from a radiated field does not. A consequence of this is that measurements on reactive fields perturb their source, consequently changing the field itself.

Reactive electric fields are created when two spatially separated objects are at different electric potentials, i.e. there is a non-zero voltage between the two. A reactive magnetic field is produced by a current flowing along a conductor. A radiative field is created by an alternating current flowing on a conductor, or antenna. For electrically small conductors the radiated power at a given frequency is proportional to the square of the product of the current and the antenna's electrical length, $P_{rad} \propto (I\ell)^2$. For antennas of length near or above a wavelength this relationship becomes a slowly increasing, oscillating function. An optimal (good radiative power, yet not too large) length for an antenna is $\lambda/2$, which is the length typically used for a dipole antenna[11] [93, p.263-264].

When working with radio frequency circuits, any conductors which are electrically long may radiate unintendedly, and it is important to ensure that this does not occur. Components such as capacitors and resistors can no longer be considered to be ideal, but must have their "parasitic"

[11] A half wave dipole, is an efficient antenna made from two conductors, each $\lambda/4$ long, arranged such that they emanate out from a centre point in opposite directions. They are ideally driven in a balanced fashion at the centre, with equal but opposite currents and voltages on each half. A monopole antenna is just one $\lambda/4$ section positioned above a ground plane, creating an "image" monopole.

properties considered as shown in section 2.2.1.5. Interconnections between devices must be made with "transmission lines", a special type of cable which maintains a consistent ratio between current and voltage, further described in section 2.2.1.6.

2.2.1.1 Impedance and Reactance

Impedance is the ratio of voltage to current, incorporating both the ratio and the phase difference between the two. It is mathematically convenient to represent impedance as a complex number:

$$\mathbf{Z} = \frac{\mathbf{V}}{\mathbf{I}} = |\mathbf{Z}|\, e^{j\phi}$$

Where $\mathbf{V}$ and $\mathbf{I}$ are complex voltage and current phasors, $|\mathbf{Z}|$ is the ratio of the voltage difference to current amplitude, ϕ is the phase difference between the voltage and current, and j is the imaginary unit given by $j = \sqrt{-1}$[12]. This can also be written in cartesian form

$$\mathbf{Z} = R + jX \tag{2.48}$$

R is the resistance, and X is the reactance. Inductors and capacitors have a frequency dependent reactance given by:

$$X_L = \omega L \tag{2.49}$$

$$X_c = -\frac{1}{\omega C} \tag{2.50}$$

2.2.1.2 Skin Depth

The skin effect is a phenomenon exhibited by conductors that cause RF currents to only flow on their outer surface. An incident oscillating electromagnetic field on a conductor will induce mobile charges (usually electrons) to oscillate at the same frequency as the field, resulting in an alternating electric current, which generates its own magnetic field that opposes the incident field. Thus, the current density decreases exponentially with depth, allowing the skin depth δ to be defined as the distance over which the current density falls to $1/e$ of its original value. It is a function of frequency f, resistivity ρ, the magnetic permeability of the conductor μ_r, and is given by the formula.

$$\delta = \sqrt{\frac{\rho}{\pi f \mu_0 \mu_r}} \tag{2.51}$$

[12] In electrical engineering j is commonly used to represent the imaginary unit because i represents instantaneous current.

2.2.1.3 Tuned Circuits

An electrical circuit is resonant when the reactance from the circuit inductance and circuit capacitance are of equal magnitude, but opposite in sign and so cancel. The frequency at which this occurs is called the resonant frequency ω_0, given by equation 2.52. At resonance the tuned circuit does not have any reactive component to its impedance, only the the real resistive value. Energy oscillates between being stored as an electric field in the capacitor, and a magnetic field in the inductor.

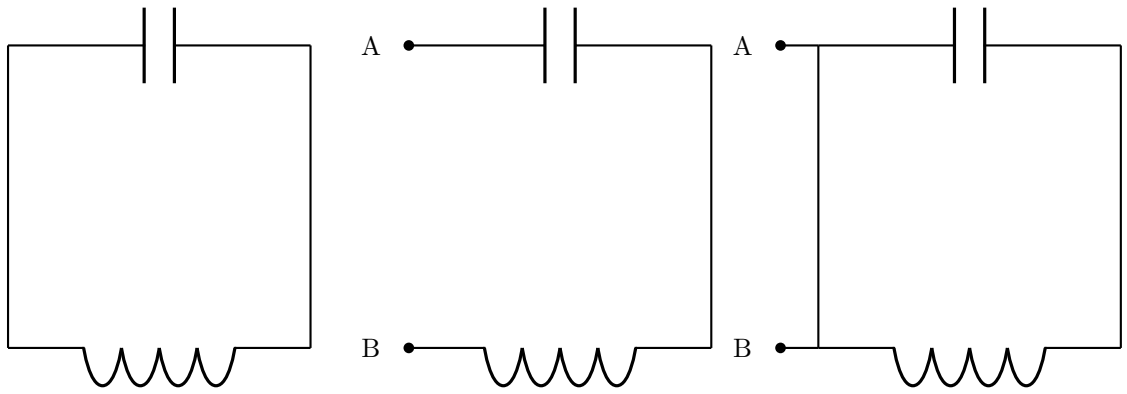

a. Isolated LC circuit b. Series LC circuit c. Parallel LC circuit

Figure 2.13: Simple LC circuits

The resonant frequency of an LC circuit is given by:

$$\omega_0 = \sqrt{\frac{1}{LC}} \tag{2.52}$$

An isolated LC circuit is shown in figure 2.13.a. Usually this circuit is interfaced to other components, which can either be done in series or parallel. At resonance, due to the cancellation of the reactances in a series LC circuit (figure 2.13.b), the impedance measured between points A and B is zero, yielding maximum current flow. In a resonating parallel LC circuit the impedance is infinite, blocking current flow between points A and B. Real-world components also have a resistance, which introduces losses such that in a series LCR circuit the minimum impedance is equal to the sum of all the resistances in the circuit, and in a parallel LCR the blocking impedance is reduced to a finite value. Generally, for a given frequency a larger blocking impedance is obtained by using a larger value of inductance.

2.2.1.4 Quality Factor

The quality factor Q of a resonating circuit is the ratio of energy stored per cycle over energy dissipated:

$$Q = 2\pi \frac{E_{stored}}{E_{dissipated}} \tag{2.53}$$

For a coil this is the ratio of the energy stored in the electric and magnetic fields, over the energy dissipated by the coil resistance. It can be calculated as:

$$Q = \frac{f_0}{\Delta f} = \frac{\omega_0}{\Delta \omega} \tag{2.54}$$

2.2.1.5 The Effect of Parasitic circuit Elements

At radio frequencies the assumption that components are "ideal" is no longer valid. It is important to consider the parasitic inductance and parasitic capacitance of components in the design of RF circuits. In addition all conductors have a non-zero resistance, which can have a significant effect on the performance of a circuit. Described below are non-ideal circuit models of a resistor, capacitor and inductor.

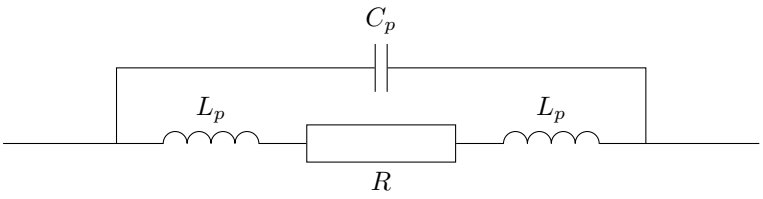

Figure 2.14: Real Resistor Schematic

Resistor A real resistor has additional series inductance from its leads, and parallel capacitance due to the voltage drop (and hence electric field) across the component. As a result, the resistor will only behave as a pure resistance at low frequencies where the inductive and capacitive reactance can be ignored, and at the frequency the capacitive and inductive reactances cancel, the resistor will resonate as a parallel LCR circuit, resulting in a high blocking impedance. Above the self resonant frequency the reactance from the capacitor dominates, making it easier for high frequencies to pass.

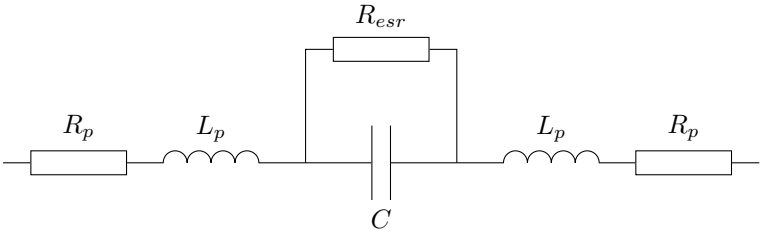

Figure 2.15: Real Capacitor Schematic

Capacitor A real capacitor has parasitic inductance and resistance from its leads, and in addition has an equivalent series resistance (ESR, R_{esr}) which is typically frequency dependent. ESR arises due to DC dielectric leakage and AC dielectric loss [198, p.142], and has implications for the amount of AC current that the capacitor can pass. At RF frequencies parasitic inductance can be

problematic, as the capacitor will be self-resonant at the frequency where the reactances from L_p and C cancel, above which the impedance of the capacitor will increase with frequency. Capacitors also have a dielectric breakdown voltage, above which arcing can occur, damaging the capacitor and posing a fire risk.

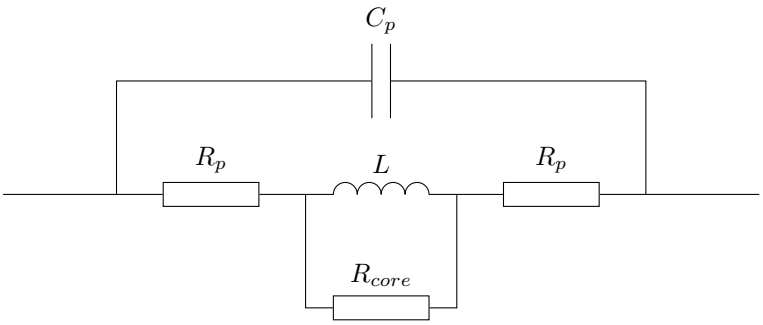

Figure 2.16: Real Inductor Schematic

Inductor A real inductor has parasitic capacitance between the leads and loops of wire, resistance in the leads and the wire of the inductor itself, and an equivalent resistance due to core losses (for ferromagnetic cores). Again a resonant circuit is formed from the combination of the inductance and parasitic capacitance, and above this self-resonant frequency the impedance of the inductor decreases with frequency.

To produce circuits which behave in a predictable manner it is important to use components which are designed for operation at radio frequencies. RF resistors and capacitors are generally of the chip type, no leads minimises their inductance and raises their self-resonant frequency, extending their working frequency range. Low ESR capacitors are also favourable, these are typically made from a ceramic material.

2.2.1.6 Transmission Lines

Electrical signals are carried on pairs of conductors; one that carries the signal, and the other which makes up the return path. For electrically small conductors, where the voltages and currents can be assumed to be spatially uniform, these conductors can theoretically assume any geometry[13]. This idealisation does not hold for electrically large conductors, where the conductors themselves must be considered as circuit components: a transmission line. While any signal–return pair of conductors can be considered as a transmission line, the term usually refers to a specific pair of conductors with uniform geometry (consistent width, height, spacing, material etc.), designed to carry high frequency signals with minimal power loss.

[13]The source impedance of a signal, and the electromagnetic environment makes some conductor geometries more susceptible to noise or electromagnetic interference, and of course, even at low frequencies a large enough inductance will have an effect. It is good practise even at low frequencies to assume what approaches a transmission line geometry.

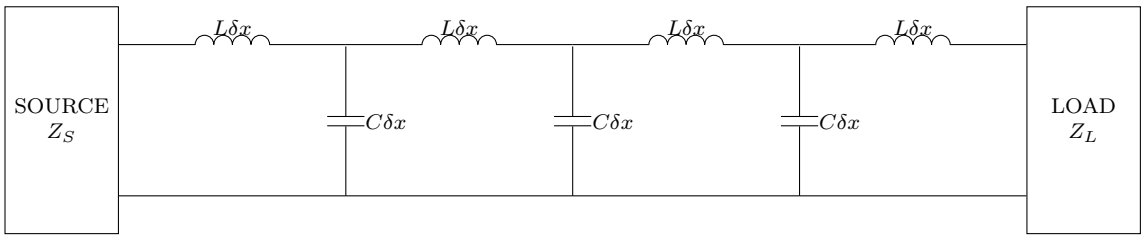

a. Ideal Transmission Line

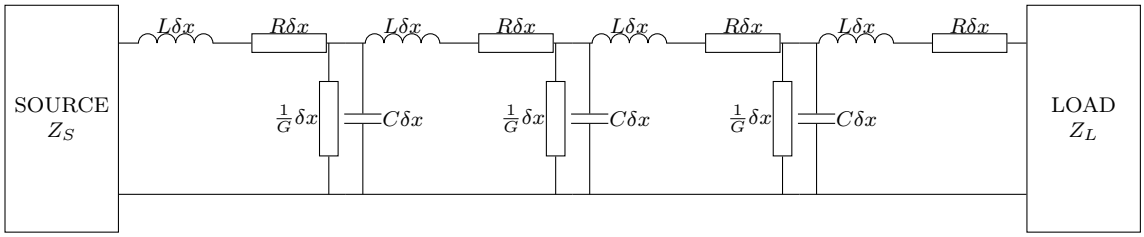

b. Lossy Transmission Line

Figure 2.17: Ideal and Lossy Transmission Lines

An ideal, lumped element transmission line model is shown in figure 2.17.a. A set of signal-return conductors with uniform geometry has a capacitance C per unit length due to electric field created between the two conductors, and an inductance L per unit length due to the magnetic field created by the current. In the transmission line model, this is represented as an infinitesimal circuit element consisting of a series inductance $L\delta x$ and a parallel (or shunt) capacitance $C\delta x$. A transmission line is then built out of a series of these infinitesimal circuit elements. At low frequencies, the capacitance and inductance does not have much of an effect on the signal. However, at radio frequencies they cannot be ignored. Transmission lines have a characteristic impedance, Z_0, which is given by:

$$Z_0 = \sqrt{\frac{L}{C}} \tag{2.55}$$

Where L and C are the inductance and capacitance per unit length, respectively. The characteristic impedance of a transmission line is a function of its geometry and material composition (for example the dielectric between the two conductors), but does not depend on length. Changing the length of an ideal transmission line introduces a phase shift to the current and voltage waveforms, but otherwise does not affect them.

Real-world conductors also have a finite series resistance per unit length, and dielectric materials exhibit a small amount of leakage, which is a small shunt conductance per unit length. Therefore, a lossy transmission line model, as shown in figure 2.17.b also has contributions from a series resistance $R\delta x$, and shunt conductance, which has resistance $1/G\delta x$. The characteristic impedance of a lossy transmission line is given by equation 2.56. A well constructed transmission line minimises these loss factors, however there will be a degree of attenuation per unit length, which can be of significance for very weak or high power signals.

$$Z_0 = \sqrt{\frac{R + j\omega L}{G + j\omega C}} \tag{2.56}$$

Another way to consider a transmission line is that it is a guide for an electromagnetic wave. The wave propagates along the transmission line, with a voltage–current ratio given by Z_0. At interfaces where Z_0 differs (for example two transmission lines with different characteristic impedances connected together), or if the load Z_L or source impedance Z_S differ to the transmission line's Z_0, reflections occur, resulting both a forward and reverse wave. For optimal power transfer it is important to correctly match Z_S, Z_L and Z_0, particularly because certain devices such as power amplifiers cannot tolerate reflected power. To correctly match a source or load impedance, a matching network [93, Ch.2] is used, which consists of inductors and capacitors with values chosen to transform the load or source impedance to another value. When a forward and reverse wave are present on a transmission line, their summation produces an interference pattern known as a standing wave[14]. The ratio of the maximum to minimum voltage magnitudes is called the voltage standing wave ratio, or VSWR. With zero reflected power the VSWR is 1, corresponding to all of the transmitted power being transferred to the load. If all of the power is reflected the VSWR is infinite.

Another property of transmission lines is that they can be used to transform impedances. A transmission line that is always terminated with its characteristic impedance will have an input impedance of Z_0. However, if it is terminated with another impedance, $Z \neq Z_0$, the input impedance of the transmission line will be a transformed version of the terminating impedance, Z':

$$Z' = \frac{Z + jZ_0 \tan(\frac{2\pi l}{\lambda})}{Z_0 + jZ \tan(\frac{2\pi l}{\lambda})} \tag{2.57}$$

Where l is the length of the transmission line and λ is the wavelength of the electromagnetic wave within the transmission line (the permittivity and permeability of the transmission line material must be taken into account). This has several implications. A transmission line whose length is an integer multiple of $\lambda/2$ will not transform a load impedance ($Z' = Z$). If a transmission line's length is an odd multiple of $\lambda/4$, it inverts the terminating impedance, for example transforming an open circuit to a short circuit, short to an open circuit, an inductor to a capacitor and vice versa. For lengths less than $\lambda/4$, an open circuit is transformed into a capacitance, a short circuit is transformed to an inductance. For lengths between $\lambda/4$ and $\lambda/2$ the inverse occurs: an open circuit is transformed into an inductance, and a short circuit is transformed into a capacitance. This behaviour repeats for every half-wavelength period.

The most common form of transmission line used at radio frequencies is the coaxial cable. This

[14] A standing wave is a wave which does not propagate. Consider a transmission line with equal forward and reverse propagating voltage waves: $V(x,t) = V_0 \sin(kx - \omega t) + V_0 \sin(kx + \omega t)$. Using the appropriate trigonometric sum-to-product identity this is $V(x,t) = 2V_0 \cos(\omega t) \sin(kx)$. If the forward and reverse waves are not equal in amplitude then the net result is a standing wave of amplitude equal to the weakest of the two waves, and a travelling wave of amplitude equal to the remainder of the stronger wave.

consists of an inner conductor, surrounded by a dielectric (usually a low loss plastic dielectric such as polytetrafluoroethylene, PTFE, also known by its trademark Teflon), which is then surrounded by an outer shield, which acts as the return current path. This outer shield can either be a metallic braid (such as woven copper filaments), foil, or solid metal. The radii of the inner conductor and outer shield are carefully controlled such that the characteristic impedance is uniform along its length. The electric and magnetic fields are contained within the outer shield, providing good screening from noise, and minimal coupling of the transmitted signal to the surroundings. The standard characteristic impedance for most RF systems is 50Ω, chosen because for coaxial cables it is a good compromise between maximum power handling and maximum power efficiency, which occur at $Z_0 = 30\Omega$ and $Z_0 = 77\Omega$, respectively [198, p.175]. Further information on transmission lines can be found in [93, Ch.10] and [198, Ch.8].

2.2.2 RF Coil Principles

In section 2.1.3.3, a brief overview of transmit and receive coils was given. In this section the basic principles and considerations required are detailed for the practical design and construction of basic RF coils for use in MRI. A RF coil is a tuned circuit, sensitive to magnetic fields at its resonant frequency. It has an inductance due to the geometry of the conductors, to which an appropriate capacitance is added such that the coil resonates at a chosen frequency, and it also has a resistance due to losses in the conductors and components. It is possible to represent a coil as a tuned circuit with lumped elements representing each of these, similar to those shown in figure 2.13 in section 2.2.1.3. While this aggregate approach has its uses when analysing the gross properties of the tuned circuit, it does not provide insight into the practical aspects of coil design or construction. It is important to note that standard RF coils used in MRI are designed to produce and detect a reactive (near) magnetic field. If the coil produces, or is sensitive to electric or radiative fields, efficiency and SNR is lost, even if the coil is resonating and has a high Q. Therefore, in order to design an efficient coil for MR it is important to understand how to minimise these electric and radiative fields.

Two categories of RF coils exist: surface coils and volume coils. A surface coil is constructed from conductors arranged in a loop, for example a circle, square, or figure-eight shape. They have a transmit and receive sensitivity that falls off with distance from the coil, and consequently have poor homogeneity. However their SNR is very good close to the coil, therefore surface coils are typically used for detecting the NMR signal. Surface coils can be combined into arrays to increase their field of view whilst retaining the high sensitivity and SNR. Volume coils have very homogeneous transmit and receive sensitivities, and permit the signal to be driven/detected in quadrature. Modern clinical MR systems tend to use surface coil arrays for reception, and a large volume coil for transmitting the imaging pulses.

A simple surface coil can be constructed from a loop of wire and a single capacitor to resonate the inductance of the loop of wire. A matching network can be constructed using a capacitor geometrically opposite the tuning capacitor, with its value selected to transform the coil's resistance to 50Ω (see section 2.2.2.2), allowing the coil to be connected to a MR RF transmit or receive channel. While it will be possible to use this coil for excitation and reception, performance will be sub optimal for several reasons:

First, the single tuning capacitor concentrates the electric field in one location on the coil. If the coil is used as a transmitter, localised power deposition close to this capacitor will be very high. Second, the coil will capacitively and inductively interact with a biological (hence conductive) sample, and other conducting objects within the magnet bore. The stray capacitive interaction is typically on the same order of magnitude as the tuning capacitance, shifting the resonant frequency such that the coil requires re-tuning. The inductive interaction causes RF eddy currents to be induced in the sample, which being lossy results in an increased resistance in the coil (this is known as loading). Third, the coil is a balanced (i.e. equal but opposite currents on each side of the coil) device, but the coaxial cable is unbalanced (the inner conductor carries the signal and the outer is at ground potential, 0V). Unless correctly converted from balanced to unbalanced signals, a loss in power transfer can occur (resulting in reduced transmit efficiency and decreased SNR in reception) because the signal does not travel purely as a differential mode along the coaxial cable. Fourth, if the coil is used in conjunction with other transmitter/receiver coils within the magnet bore, it will couple strongly to those coils, resulting in image distortions and potentially increased power deposition.

2.2.2.1 Distributed Capacitance

To reduce the electric field interaction between a coil and the sample, the tuning capacitance is spatially distributed around the coil. Ideally, an RF coil should only generate an RF magnetic field, however a conservative electric field is present at each capacitor. If this electric field penetrates a conductive sample, loss is introduced (known as dielectric loss) and the resonant frequency of the coil is shifted. Furthermore, if the conductors used to construct the coil are electrically large they behave as a transmission line and phase shifts along these lines cause destructive interference that decreases the peak of the generated magnetic field (radiation loss). If the length approaches $\lambda/4$ then the conductor behaves as an antenna and a significant proportion of the power is radiated as an electromagnetic wave. Distributed capacitance is employed to counter dielectric and radiation losses.

Distributing the tuning capacitance around a coil will, through the application of Kirchhoff's Voltage Law reduce the large voltage across one capacitor into smaller voltages across more capacitors. Providing the capacitor values at each point are the same, so too are the voltages. In addition, symmetrically arranging the capacitors around the loop geometrically balances the electric fields

on the coil, and it reduces the length of any one conductor, reducing radiation losses.

Capacitance adds in series according to $\frac{1}{C_T} = \sum \frac{1}{C_n}$. When splitting a single tuning capacitance C_T into N distributed capacitors, each will have a capacitance of NC_T. Coil-sample capacitive coupling, and associated dielectric losses can be minimised by choosing the number of distributed capacitors so that the value of the capacitance at each point is larger than the stray coil-sample capacitance. It is hard to predetermine what the value of this will be, for 123.2MHz a general rule of thumb is that it is desirable to have at least 100pF at each position so that this capacitance dominates over the stray capacitance. In doing so the coil will not detune significantly when loaded by a sample.

There are disadvantages to distributing the tuning capacitance. Each capacitor has an equivalent series resistance (ESR), and this sums linearly as more capacitors are added in series. Hence it is important to strike a balance as adding too many capacitors in series will increase the losses in the coil, decreasing the Q, too few will result in increased dielectric losses and even B_1 inhomogeneities from current standing waves [96]. The total ESR at each capacitor gap can be reduced by using more than one capacitor in parallel.

2.2.2.2 Impedance Matching

Standard RF preamplifiers and power amplifiers are designed to operate with 50Ω coaxial cable transmission lines. To ensure maximum transfer of power the impedance of the coil must be transformed to this. Parallel LCR circuits are used as the basis of a coil because the resistance in the coil can be easily transformed to any desired resistance. This is achieved by placing another capacitor in the coil and connecting across it (as shown in figure 2.18). More complicated "L", "T" and "π" [93, Ch.2] networks can also be used, however they are typically required when transforming a complex impedance, or when the bandwidth of the match needs to be controlled, hence not necessary for the narrow band, purely resistive impedance of a resonant RF coil. Matching can also be achieved using transformers [93, p.11].

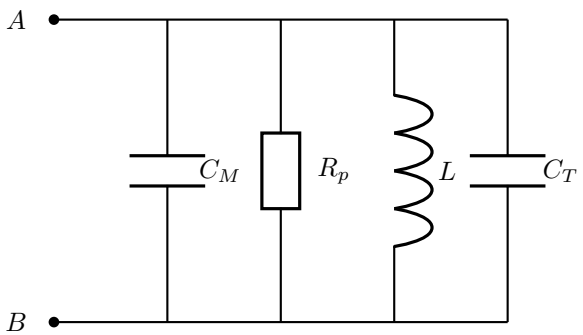

Figure 2.18: Match Capacitor Configuration

At resonance and without the match capacitor the coil has a purely resistive impedance, R_p, the equivalent parallel resistance. The origin of this resistance is a combination of the resistance

in the conductors within the coil, component losses and sample losses. When constructing a coil, this resistance can be calculated from measurements of the quality factor when loaded, Q_L:

$$R_p = \frac{X_L}{Q_{Loaded}} \tag{2.58}$$

The required reactance to transform the coil impedance R_p to 50Ω is then calculated by[38]

$$X_{match} = \sqrt{R_T R_p} \tag{2.59}$$

Where R_T is the value to transform the coil impedance to, in this case it is set to 50Ω. The value of C_M is then computed with equation 2.50.

2.2.2.3 Detuning

Detuning (also referred to as Q-spoiling) is necessary when multiple coils resonating at the same frequency are used (for example transmit with a volume coil and receive with a localised surface coil). If during transmit the receive coil is not detuned it will couple strongly to the intense B_1 field created by the volume coil, inducing large RF currents in the receive coil. These could damage the preamplifier, and also focus the B_1 field, increasing local power deposition.

Detuning is the achieved using PIN (Positive-Intrinsic-Negative) diodes [57]. A PIN diode has a lightly doped intrinsic region in between heavily doped p-type and n-type semiconductor regions. The wide intrinsic region makes the diode suitable as an RF attenuator or switch as it has a low resistance ($< 1\Omega$) when forward biased. At zero bias a PIN diode has a small capacitance (less than 3-6pF), which decreases when a reverse bias voltage is applied. A PIN diode does not pass RF so long as the peak to peak RF voltage does not surpass the reverse bias voltage, and so PIN diodes used to switch high power RF (such as in a body transmit coil) are reverse biased by several hundred volts.

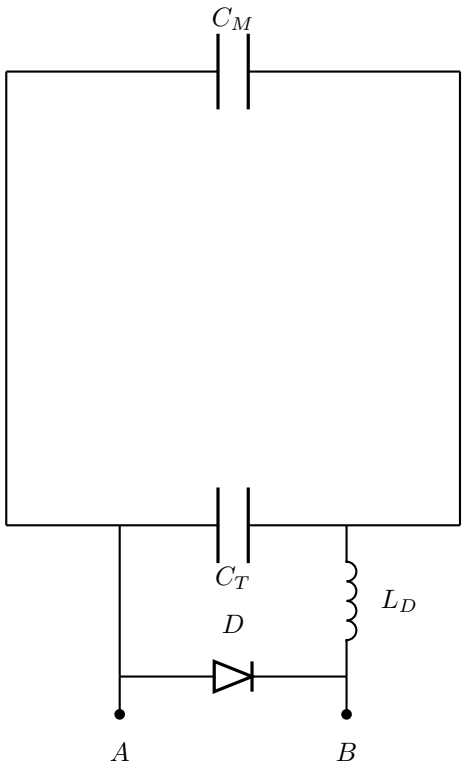

Figure 2.19: PIN Diode Detuning

PIN diode detuning was first proposed by Mellor and Checkley [154]. It is achieved by constructing a circuit such that when the PIN diode is biased the match capacitor is resonated with an inductor L_D in parallel, as shown in figure 2.19. The value of the inductor is chosen so that the parallel resonant circuit resonates at the operating frequency of the coil. This places a large series impedance in the coil, blocking any currents from flowing, detuning it. PIN diodes can also be placed directly in series within a coil to act like a switch, so that they must be forward biased for the coil to operate. Their non-zero resistance when forward biased introduces noise, making this topology unsuitable for receive coils.

2.2.2.4 Balancing

Ideally RF signals are transmitted along a coaxial cable in what is known as differential mode: equal but opposite currents on the inner conductor and shield with zero net current flow. The skin effect causes this differential mode current to travel on the inside of shield, therefore it does not couple to any external influences. However due to imperfections there can also exist common mode currents. An alternate path to the return conductor (in the case of the coaxial cable this is the cable shield), such as a parasitic capacitance results in a non-zero net current on the cable. Common mode currents can also occur due to improper balancing between the coil (which is a balanced device) and the coaxial cable (an unbalanced device). Two methods are used to prevent common mode currents. A balun (BAlanced to UNbalanced), a device which converts an unbalanced signal to a balanced signal, and vice versa, can be placed in between the coil and coaxial cable to

prevent common mode currents from flowing on the cable [38, p.138]. Common mode currents are problematic because they represent a loss in efficiency, can make a coil's performance very with what interacts with the cable, and if large common mode currents are induced they can cause patient burns.

2.2.3 Network Analyser

The vector network analyser (shortened to network analyser) is a powerful tool for testing RF circuitry. It applies a signal to a device, and simultaneously measures its response. Typically they consist of:

- A signal generator. This produces an oscillating signal which can be sent to the device under test (DUT), usually this will be a swept frequency range around a centre frequency, but the output power can also be swept at a single frequency.
- A test set. This routes the signal generator output both to the output ports (and to DUT), and to the receivers for a phase reference. A directional coupler or power divider will also be present which allows the splitting of forward and reverse travelling signals.
- Receiver. This makes measurements of the DUT connected to its test ports.

Two types of measurements can be made with a network analyser: reflection and transmission. A typical network analyser makes measurements of the scattering parameters (S-parameters) [93, Ch.28], which are measurements of the forward and reflected power between the ports on a network analyser and related to the VSWR (see section 2.2.1.6). A two port network analyser (hardware permitting) can make measurements of the following:

S_{11} Voltage reflection at port 1. $S_{11} = \frac{V_1^-}{V_1^+}$

S_{21} Voltage gain from port 1 to port 2 (also known as forward voltage gain). $S_{21} = \frac{V_2^-}{V_1^+}$

when port 2 is terminated with a load identical to the characteristic impedance (50Ω, i.e no reflections at port 2), and

S_{12} Voltage gain from port 2 to port 1 (also known as reverse voltage gain). $S_{12} = \frac{V_1^-}{V_2^+}$

S_{22} Voltage reflected at port 2. $S_{11} = \frac{V_2^-}{V_2^+}$

when port 1 is terminated with a load identical to the characteristic impedance (50Ω, i.e no reflections at port 1). $V_n^+$is the forward voltage at port n, and V_n^- is the received voltage at port n. When building and testing MRI coils usually only the S_{11} reflection and S_{21} transmission coefficients are used. Typically the S-parameters are displayed on a logarithmic scale (decibels) against the frequency. Using the S-parameters the network analyser is able to also make impedance measurements. More information on the use of a network analyser can be found in the

Agilent Application Note 1287-1, "Understanding the Fundamental Principles of Vector Network Analysis"[6].

2.3 Arterial Spin Labelling

Arterial Spin Labelling (ASL) is a non-invasive MR technique used for measuring cerebral perfusion, the delivery of nutritive arterial blood to the tissue's capillary bed. Using a combination of RF and gradient pulses, the magnetisation of protons within arterial blood water is inverted, resulting in an endogenous contrast agent. This contrast agent is then delivered to the tissue bed at a rate proportional to the perfusion rate, decreasing the available magnetisation when a MR image is formed. Two images are normally acquired, one sensitised to perfusion, a "tag" or "label" image, and the other without perfusion sensitisation, a "control" image, which are subtracted to obtain a perfusion weighted image. The decrease in signal is very small, approximately 0.5–1.5% of the total signal, therefore multiple averages are required to obtain sufficient SNR. The technique benefits from the increased SNR of high field MR systems (3T and above), and fast imaging techniques such as EPI or 3D-GRASE [173] are preferable for high SNR per unit time.

ASL was first proposed by Detre et al. in 1992 [53] to generate perfusion maps in the rat brain. Arterial blood was saturated with slice selective RF pulses applied to the neck, which upon reaching the capillary bed exchange with the bulk water in the brain to create a steady state between inflow and T_1 decay. After a delay to allow for inflow time, a spin echo acquisition was made. The off resonance pulses used to saturate the water also saturate macro-molecules within the imaging slice, imparting magnetisation transfer (MT) [240, 101] contrast into the image, so control images with the saturation slice positioned distally to the imaging slice (either by reversing the slice selective gradient polarity or the frequency offset of the RF pulse) were also acquired and subtracted from the perfusion weighted image, cancelling the MT effect (see section 2.3.1.1 for more details).

That same year Williams et al. [237] improved on this method by inverting the inflowing spins to double the perfusion weighted signal. Blood was labelled continuously using the principles of adiabatic fast passage by applying continuous-wave radio frequency power in the presence of a magnetic field gradient in the direction of arterial flow (see section 2.3.2.1). This method became known as continuous arterial spin labelling (CASL). CASL places large demands on MR equipment as it requires a continuous wave (CW) RF pulse. Consequently CASL has not been employed routinely in clinical practise.

In 1994 Edelman et al. [60] demonstrated labelling of a 10-15cm slab proximal to the imaging slice with a single 180° pulse, with a control image acquired using an inversion slab placed distally outside the head to induce identical MT effects. This was the first pulsed ASL sequence (PASL), known as "Echo-Planar MR Imaging and Signal Targeting with Alternating Radio frequency", or EPISTAR. Shortly afterwards a symmetric, MT compensating method known as "Flow Alternating

Inversion Recovery" (FAIR) was proposed by Kim [130] which uses a non-selective inversion pulse to sensitise the magnetisation to perfusion, and a slice selective inversion pulse for the control. PASL can be implemented on standard hardware and so has become a popular choice for perfusion measurements despite it theoretically having a lower signal difference compared to CASL [241].

The most recent ASL development has been dubbed pseudo-Continuous ASL (pCASL) [246, 45]. It is a pulsed approximation to the CASL sequence, and uses a train of concomitant slice selective gradients with Hanning windowed RF pulses to reduce sidebands, followed by refocusing gradients. Blood is continuously labelled by this train of pulses with similar efficiency to CASL, however no additional hardware is required and the sequence can be run on standard clinical systems, making this a very attractive technique.

The ability to noninvasively measure perfusion makes ASL a very attractive technique. Unlike dynamic susceptibility contrast perfusion MRI [190, 153, p.336], in which a a bolus of contrast agent (typically Gadolinium based) is injected into the blood stream and circulates around the vasculature, causing a reduction in T_2 and T_2^*, ASL does not require any external contrast agents and can be performed repeatedly without any adverse side-effects, even on patients with kidney disorders. Many MR system vendors now include a product ASL sequence, which can be used clinically. However the intrinsic low SNR of ASL is a limiting factor, and has only become clinically viable with the introduction of 3T systems. Multiple averaging makes the sequence very susceptible to motion artefacts, which can be misinterpreted as a perfusion signal, therefore subject collaboration is essential. While modern ASL sequences can acquire perfusion maps with good SNR in approximately five minutes, this is still a long time for some patients to keep still; therefore it is critical to use SNR efficient methods so that the acquisition time can be minimised.

2.3.1 Perfusion Quantification

A quantitative measure of perfusion, the cerebral blood flow (CBF) is obtained by fitting the perfusion weighted signal to a suitable model. In the original ASL experiments [53, 237], a CBF estimation method was presented that adapted the nitrous oxide tracer clearance theory proposed by Kety and Schmidt [128], resulting in a modified Bloch equation:

$$\frac{dM_t(t)}{dt} = \frac{M_{t,0} - M_t(t)}{T_{1,t}} + f\left(M_a(t) - \frac{M_t(t)}{\lambda}\right)$$

Where M_t is the longitudinal tissue magnetisation, $M_{t,0}$ is the equilibrium tissue magnetisation, $T_{1,t}$ is the longitudinal relaxation time of the tissue, M_a is the arterial magnetisation, f is the cerebral blood flow, and λ is the blood-brain partition coefficient. This model assumed uniform plug flow and equal relaxation rates for both tissue and arterial blood, however in subsequent work it was improved to include differences in tissue and arterial relaxation rates, the bolus length and the transit time [180].

In a different approach, Buxton et al. [33] proposed the "General Kinetic Model", which considered the magnetisation difference between control and tag measurements, ΔM to be a deliverable quantity, which arterial blood carries to a voxel of tissue. The amount of ΔM within the tissue depends on the history of its delivery by the arterial flow, clearance by venous flow and longitudinal relaxation. Three functions are defined: $c(t)$, the normalised concentration of ΔM arriving to a voxel at time t (the delivery function); $r(t, t')$, the fraction of tagged water molecules that arrived at time t' that still reside within the voxel at time t (the residue function); and $m(t, t')$, the fraction of the original longitudinal magnetisation that was carried by water molecules to the voxel at time t', that still remains at time t (the magnetisation relaxation function).

Therefore, $\Delta M(t)$ is constructed as a sum over the history of delivery, clearance and relaxation as a convolution integral:

$$\begin{aligned} \Delta M(t) &= 2M_{0,b}f \int_0^t c(t')r(t-t')m(t-t')\,dt' \\ &= 2M_{0,b}f\,[c(t) \otimes (r(t)m(t))] \end{aligned} \tag{2.60}$$

Where $M_{0,b}$ is the equilibrium magnetisation of arterial blood and $\otimes$ denotes a convolution. Using three key assumptions, which form the standard ASL kinetic model, the forms of $c(t)$, $r(t)$, and $m(t)$ can be deduced. The standard ASL kinetic model applies to both CASL and PASL, and assumes that both of these methods produce a temporally defined labelled bolus of duration τ and inversion efficiency α, which travels to the tissue within the imaging plane in a time Δt. The assumptions are:

1. Labelled blood is assumed to arrive as uniform plug flow, such that before an initial transit delay Δt no labelled blood arrives, then uniformly labelled blood arrives between $t = \Delta t$ and $t = \Delta t + \tau$, followed by no labelled blood for $t > \Delta t + \tau$.

2. Single compartment kinetics describe the mechanisms of exchange between tissue and blood, i.e. the ratio of the tissue concentration to the venous concentration is constant and equal to the brain-blood partition coefficient, λ.

3. The labelled blood initially decays with the longitudinal relaxation time of blood $T_{1,b}$, then upon reaching the tissue the magnetisation decays with the relaxation time of the tissue, T_1. Here the assumption is that all of the labelled blood is immediately extracted into the voxel.

These assumptions are summarised in equation 2.61.

$$c(t) = \begin{cases} 0 & 0 < t < \Delta t \\ \alpha e^{-\frac{t}{T_{1,b}}} \text{ (PASL)} & \Delta t < t < \Delta t + \tau \\ \alpha e^{-\frac{\Delta t}{T_{1,b}}} \text{ (CASL)} & \\ 0 & t > \Delta t + \tau \end{cases} \tag{2.61}$$

$$r(t) = e^{-\frac{ft}{\lambda}} \tag{2.62}$$

$$m(t) = e^{-\frac{t}{T_1}} \tag{2.63}$$

Using these assumptions the PASL signal is:

$$\Delta M(t) = \begin{cases} 0 & 0 < t < \Delta t \\ 2M_{0,b} f (t - \Delta t) \alpha e^{-\frac{t}{T_{1,b}}} q_p(t) & \Delta t < t < \Delta t + \tau \\ 2M_{0,b} f \tau \alpha e^{-\frac{t}{T_{1,b}}} q_p(t) & t > \Delta t + \tau \end{cases} \tag{2.64}$$

Where

$$q_p(t) = \begin{cases} \frac{e^{kt}(e^{-k\Delta t} - e^{-kt})}{k(t - \Delta t)} & \Delta t < t < \Delta t + \tau \\ \frac{e^{kt}(e^{-k\Delta t} - e^{-k(\tau + \Delta t)})}{k\tau} & t > \Delta t + \tau \end{cases} \tag{2.65}$$

$$k = \frac{1}{T_{1,b}} - \frac{1}{T_1'} \tag{2.66}$$

$$\frac{1}{T_1'} = \frac{1}{T_1} + \frac{f}{\lambda} \tag{2.67}$$

The CASL signal is:

$$\Delta M(t) = \begin{cases} 0 & 0 < t < \Delta t \\ 2M_{0,b} f T_1' \alpha e^{-\frac{\Delta t}{T_{1,b}}} q_c(t) & \Delta t < t < \Delta t + \tau \\ 2M_{0,b} f T_1' \alpha e^{-\frac{\Delta t}{T_{1,b}}} e^{-\frac{t - \tau - \Delta t}{T_1'}} q_c(t) & t > \Delta t + \tau \end{cases} \tag{2.68}$$

Where

$$q_c(t) = \begin{cases} 1 - e^{-\frac{t - \Delta t}{T_1'}} & \Delta t < t < \Delta t + \tau \\ 1 - e^{-\frac{\tau}{T_1'}} & t > \Delta t + \tau \end{cases} \tag{2.69}$$

CBF can be quantified from either a single measurement, or from multiple measurements with acquisitions made at different times after labelling. For the former the suitable form of either equation 2.64 for PASL or 2.68 for CASL is solved for the CBF, f, using the perfusion weighted

image and the appropriate parameters. For the multi-acquisition method, the ΔM data is fit to the appropriate PASL or CASL $\Delta M(t)$ equation, resulting in the CBF and transit time, Δt. The various parameters that feature in the equations must either be estimated or measured, for example a T_1 map can be acquired with a multiple inversion recovery sequence [91, p.649], λ is assumed to be 0.9ml/g, and $T_{1,b}$ can be taken from literature values.

2.3.1.1 Magnetisation Transfer

Magnetisation transfer (MT) occurs when saturated bound protons from the bound water pool (for example bound within macromolecules such as proteins and lipids) transfer their longitudinal magnetisation to protons in the free water pool. These bound protons have a broad resonance and so can be saturated by an off-resonance RF pulse. MT contrast provides a way of viewing the macromolecular content of a tissue, as the bound protons would not normally be visible due to their short T_2 relaxation times. An off-resonance pulse is used to saturate the macromolecules, which then transfer their magnetisation to the free protons, reducing the available magnetisation in an imaging sequence (in a similar manner to how reducing TR imparts T_1 contrast). In ASL, the off-resonance labelling RF pulses have this effect. Perfusion weighted contrast from an ASL sequence is generated by the reduction in available magnetisation due to the inverted magnetisation of a bolus of labelled blood, and so unless mitigated, magnetisation transfer contrast (MTC) confounds the ASL measurement, resulting in an overestimation of CBF. Magnetisation transfer contrast also decreases the available signal by approximately 2.5-fold [261], reducing the sensitivity of the perfusion measurement.

A second consequence of MT is that for long pulses, such as those performed in CASL it reduces the apparent T_1 of tissue, and to a lesser extent blood [101, 47]. This has two consequences: blood that has already arrived at the imaging volume and exchanged into the tissue will decay at a faster rate than normal, due to the decreased tissue T_1; the labelled bolus of blood will decay at a faster rate, resulting in reduced inversion efficiency because of the decreased arterial blood relaxation time. Measurements by Hernandez-Garcia et al. [101] at 7T of blood within a phantom found its $T_1$to reduce from 1810ms to 1357ms due to the MT effects of off-resonance irradiation at 5000Hz applied between acquisitions (achieving an equilibrium magnetisation transfer) in a multi-TR T_1 measurement.

2.3.2 Continuous ASL

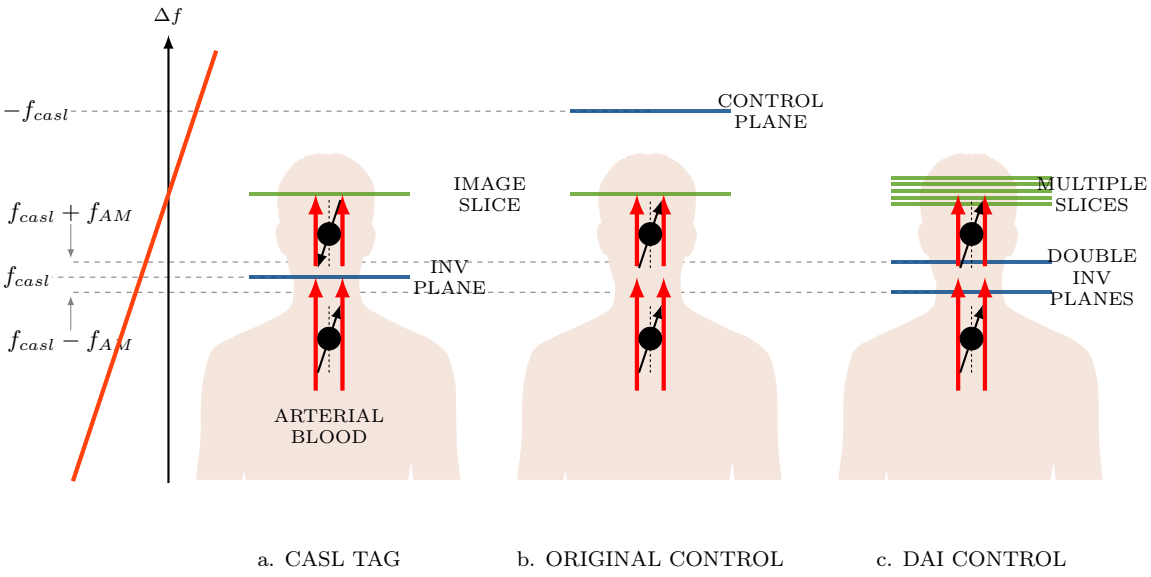

Figure 2.20: Continuous ASL Labelling Schemes

Continuous arterial spin labelling uses a long RF pulse (lasting up to several seconds) and a gradient applied in the direction of arterial flow to invert blood continuously using a velocity driven adiabatic fast passage. As the duration of the RF pulse is long compared to the evolution of the magnetisation, its behaviour can be described by the steady state equations of the Bloch equation; precession about the effective magnetic field, B_{eff} formed by the vector summation of the static and radio frequency magnetic fields. With the applied gradient this defines a plane (offset proximally by Δz with respect to the imaging slice), usually located proximally to the Circle of Willis, or within the neck to label the common carotid arteries, in which the Larmor frequency of the protons is equal to the frequency of the RF pulse. Flowing protons experience a frequency sweep, which causes the effective magnetic field to rotate from aligned with$+z$ to aligned with $-z$. Under the correct conditions (see section 2.3.2.1 for these), the magnetisation of the protons can be made to tightly precess about B_{eff} as they pass through the inversion plane, resulting in a flow driven adiabatic inversion. Typically, labelling is performed for 1–3s, after which there is a delay known as the post-labelling-delay [9] to allow the tagged blood to reach the imaging volume and perfuse into the tissue. Without this delay the ΔM signal becomes highly sensitive to the tissue transit time.

Figure 2.20 is a diagram of this process, with 2.20.a depicting the labelling of arterial blood at the inversion plane, which then flows to the imaging slice. Static magnetisation at the inversion plane does not experience this frequency sweep, and is consequently only saturated, and elsewhere the free water protons are unaffected. However, the long RF pulse also saturates the bound protons within macromolecules such as proteins and lipids, which then impart their magnetisation onto protons in the free water pool, resulting in a MT induced reduction in longitudinal magnetisation

in the imaging slice. As the ASL signal also comes from a reduction in longitudinal magnetisation, the MT contrast must also be present on the control image, otherwise it will be misinterpreted as part of the ASL signal and lead to a CBF overestimation.

The degree of macromolecular saturation is specific to the RF carrier frequency, therefore control acquisitions must keep this constant whilst not labelling any of the inflowing magnetisation. The original continuous ASL method by Williams et al. [237] used a control plane, positioned at $-\Delta z$ to the imaging slice, resulting in the same MT effect, but no labelling of the inflowing blood, as shown in figure 2.20.b. This can be achieved by either reversing the gradient polarity, or the RF frequency offset (with respect to the imaging slice frequency). Because Δz is specific to each slice, a different control acquisition is required each time, and multi-slice acquisitions cannot be accomplished for a single labelling pulse.

To permit multislice acquisition, the double adiabatic inversion (DAI) was proposed for the control acquisition by Alsop et al. [10]. DAI is also referred to as amplitude modulated control (AMC). By modulating the labelling RF pulse by a sinusoidal envelope at frequency f_{AM}, two inversion planes are created, spaced equidistantly about the central frequency of the RF pulse. Typically a low frequency ($f_{AM} < 500Hz$) sinusoidal envelope is used, creating two inversion planes separated by $2 \times f_{AM}$ around the RF carrier frequency (typically offset by 10kHz from the Larmor frequency). Flowing blood is inverted at the first labelling plane, then re-inverted at the second, theoretically resulting in unperturbed magnetisation. Because both tag and control acquisitions use the same labelling gradient and RF carrier frequency, both exhibit the same degree of MTC, hence multislice acquisitions are possible. Due to imperfections in the double inversion, AMC results in a reduction in inversion efficiency, which can be dependent on the velocity of inflowing blood, and this must be properly taken into account for correct quantification [166].

Not all RF power amplifiers, particularly those in clinical MR systems are able to produce a continuous wave (CW) RF pulse at the power required to successfully invert inflowing blood. Even if the capability is there, SAR limits are often exceeded in attempting to create sufficient B_1. Therefore in many implementations of CASL the CW pulse is broken down into a series of long RF pulses with short gaps between the pulses (duty cycle constrained). Simulation and experimental results have shown the dependence of the inversion efficiency on the labelling duty cycle to be highly non-linear [79, 218], hence performance is varied between different MR systems, and even on a subject-by-subject basis due to SAR constraints.

An alternative way of performing CASL that circumvents both RF duty cycle constraints, and MTC is to use a separate, localised labelling coil, positioned directly on the neck over the carotid arteries. Because the RF field is spatially localised to the neck, there is no saturation of macromolecules and subsequent MTC, and substantially less power is required because the B_1^+ field from the labelling coil only has to penetrate a few centimetres to the carotid arteries. Furthermore, vessel specific labelling can be performed by using a small surface coil which confines the labelling

field to only the left or right common carotid artery and produces perfusion maps specific to the perfusion territories of each region.

Separate coil CASL was first implemented by Detre et al. [54] in 1994, with more details on the implementation given by Zhang et al. [260] and Silva et al. [204]. A small 0.5cm diameter butterfly-shaped surface coil was used to label inflowing blood within the neck of rats. The labelling coil was decoupled from the imaging saddle volume coil using PIN diodes (see section 2.2.2.3). A high inversion efficiency of approximately 82% was seen with the labelling coil, 10% higher than that from an equivalent 7cm volume coil. Because of the spatially localised labelling field, no RF power was applied during the control acquisition.

The first implementation on a clinical scaner was by Zaharchuk et al. [259] in 1999. Two PIN didoe detunable labelling coil designs were used; a circular 5cm diameter surface coil for positioning over a single carotid artery, and a $5 \times 20cm$ flexible rectangular coil for labelling both carotid arteries. An independent RF transmitter system was constructed, consisting of a frequency synthesiser, gated by mixers and amplified by a 1W wideband RF amplifier module, then further gated by a T/R switch. Power was also monitored by an RF power meter to ensure that excessive power was not deposited into the subject. Results showed that both coils completely eliminated MTC within the brain, and single vessel labelling also produced perfusion territories of the right carotid artery with good specificity. Maximal labelling was found to be achieved with a power of 0.3W.

More recent implementations have used similar hardware [103, 157, 156, 213], generally consisting of either a small single artery coil or larger coil for labelling both arteries, and a low power transmitter channel with basic pulse gating. A heavy reliance is placed on the labelling geometry, which can very from subject to subject, requiring accurate positioning of the coil for optimal inversion efficiency. Additionally, the labelling plane is typically located further away from the brain than in regular CASL, resulting in increased transmit time and more T_1 decay of the labelled blood. Despite the combination of high labelling efficiency, no MTC, and low power deposition, the additional custom hardware requirements have meant that this technique is not suitable in the clinical setting, and developments in the clinically viable pulsed and pseudo-continuous ASL methods have meant that there has been no reason for MRI vendors to build separate coil capabilities into their systems. Consequently, separate coil CASL has remained the speciality of only a few research groups.

Because CASL continuously inverts fresh inflowing blood, it has been shown to theoretically and experimentally produce a larger ΔM signal than PASL [241]. However, the difficulties in perfect implementation on unmodified hardware mean that it can be difficult to realise these benefits, and subsequently other methods are more favoured.

2.3.2.1 Velocity Driven Adiabatic Fast Passage

In continuous arterial spin labelling (CASL) flowing blood is inverted by what is known as a Velocity Driven Adiabatic Fast Passage [55] (AFP). Blood is subjected to a continuous radio frequency transverse field $\mathbf{B_1^+}$, and a labelling gradient $\mathbf{G} = G_z\hat{\mathbf{z}}$, where the direction of arterial flow is assumed to be in along z. This defines a labelling plane, which in the rotating frame at the isocentre of the magnet the position is given by

$$z_{label} = \frac{f_{casl}}{\gamma G_z} \tag{2.70}$$

Where γ is the proton gyromagnetic ratio, G_z is the amplitude of the labelling gradient, and f_{casl} is the frequency offset of the applied B_1 field (typically 5-20kHz) from the B_0 isocentre frequency, f_0.

A proton in the blood water flowing with velocity v will experience an effective magnetic field $\mathbf{B_{eff}}$, which is the vector summation of the RF B_1 field and the magnetic field gradient:

$$\mathbf{B_{eff}} = \mathbf{B_1^+} + (\mathbf{G}\cdot\mathbf{z})\hat{\mathbf{z}} \tag{2.71}$$

In the rotating frame of the applied RF field, $\mathbf{B_{eff}}$ will sweep from pointing along $+z$, to along the x axis at the labelling plane (where $\mathbf{G}.\mathbf{z} = 0$) and then down to $-z$. Adiabatic inversion is achieved if this sweep is slow enough so that the magnetisation of the proton remains locked to and precesses about $\mathbf{B_{eff}}$, but must be fast enough so that T_1 and T_2 relaxation effects are negligible. This is written mathematically as the adiabatic condition:

$$^1/_{T_1}, ^1/_{T_2} \ll \frac{Gv\left|\mathbf{B_1^+}\right|}{\left|\mathbf{B_{eff}}\right|^2} \ll \gamma\left|\mathbf{B_{eff}}\right| \tag{2.72}$$

By considering the value of $\mathbf{B_{eff}}$ at the labelling plane, where $\mathbf{B_{eff}} = \mathbf{B_1^+}$, a lower limit for B_1 can be defined from the right hand side of the inequality:

$$\beta \gg 1$$

Where β is the adiabaticity factor [193], defined as

$$\beta = \frac{\gamma\left|\mathbf{B_1^+}\right|^2}{Gv} \tag{2.73}$$

2.3.2.2 Pseudo-Continuous ASL

Pseudo-Continuous ASL (pCASL) was developed to address the duty cycle constraints of CASL without using additional hardware. A train of rectangular RF pulses, as used in duty cycle constrained CASL, produces multiple aliased inversion planes at locations governed by the pulse

spacing, resulting in an inefficient inversion. pCASL, first proposed by Dai et al. [45] and Wu et al. [246] modifies the RF and gradient waveforms that make up the labelling pulse train such that the aliased labelling planes are eliminated. This is achieved by using a Hanning[15] pulse as the RF envelope, which has a smaller frequency content than a rectangular pulse; and an increased gradient amplitude during the Hanning pulse (as opposed to a constant gradient). The repeated RF pulses create a pulsed steady state, with the time-averaged RF amplitude comparable to the constant RF amplitude present in CASL, and an unbalanced refocusing gradient lobe in the space between pulses such that the net gradient is also comparable to that found in CASL. The phase of each RF pulse must also be incremented by the phase accrued by the magnetisation at the labelling plane over each pulse cycle (which is proportional to the average gradient over the pulse cycle).

The control acquisition can be accomplished in one of two ways. Balanced pCASL [246] uses the same gradient waveforms, hence average gradient and phase increment for both control and tag, and alternates the sign of the RF pulse between pulses for the control scheme. Unbalanced pCASL [45] has different gradient waveforms for control and tag: the tag is not fully refocused to produce an average gradient comparable to that in CASL (as described above); the control experiment has fully balanced gradient waveforms (therefore average gradient is zero, with similarities to a balanced steady state free precession sequence [91, p.472]), and the RF phase is alternated in sign. In both cases, during the control experiment there is an effective 180° shift between successive pulses at all positions. The flowing blood's magnetisation is unperturbed because the time averaged B_1^+ is zero, yet because the RF frequency is the same, MTC contrast is identical to the tag measurement.

pCASL is capable of retaining up to 96% of the inversion efficiency of CASL, whilst only having a RF duty cycle of 33% [45], resulting in SAR levels well within safety limits. It can be easily implemented on standard scanner hardware, and the larger gradient amplitude results in reduced macromolecule saturation in the imaging volume, increasing perfusion sensitivity. However, pCASL is sensitive to B_0 inhomogeneities or gradient imperfections which cause a mismatch between the per pulse phase accrual and the RF phase increment, reducing inversion efficiency because the RF does not successfully track the phase evolution [122]. Shimming at the label plane can aid this, however this can introduce inhomogeneities into the imaging volume if the MR hardware does not permit rapid switching between two shim conditions.

[15] A Hanning pulse is defined as $H(t) = \frac{1+cos(2\pi t)}{2}$.

2.3.3 Pulsed ASL

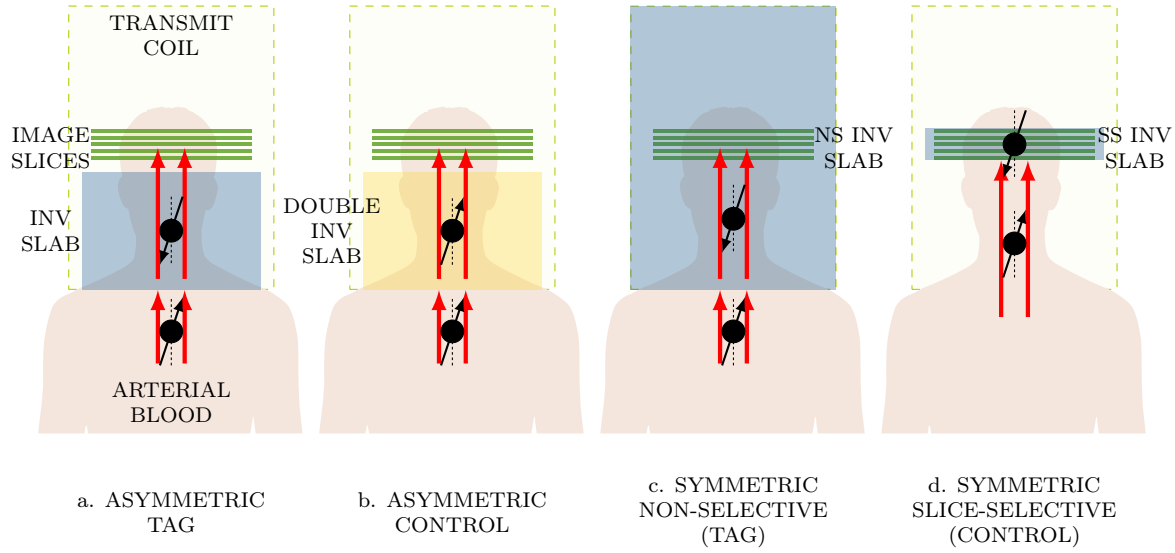

Figure 2.21: Asymmetric and Symmetric Pulsed Labelling Schemes

Pulsed Arterial Spin Labelling refers to the creation of an inverted slab of blood using a short (on the order of ms) RF pulse. Because they can be implemented on standard clinical MR hardware, they have been met with increasing success and are starting be seen as product sequences provided by the MR system vendors. The two PASL sequences discussed above, EPISTAR, with its inversion slab placed asymmetrically to the imaging volume, and FAIR which uses symmetrically positioned non-selective and slice selective inversions, represent the bases from which most PASL sequences are derived, subsequently termed asymmetric and symmetric PASL sequences, respectively. Figure 2.21 illustrates the differences in the inversion slab positions of the two methods. In general, an inversion pulse is applied to invert inflowing blood, then after a delay TI, corresponding to the inversion time, an image is acquired.

In the original EPISTAR implementation by Edelman et al. [60] the control acquisition was made using an inversion slab placed distally to the imaging slice, analogous to the basic CASL technique. The slab of inverted blood flows to the imaging volume, reducing the available longitudinal magnetisation when an image is acquired. This only permitted a single slice acquisition, because MTC is specific to the offset between the inversion slab and the imaging slice. To facilitate multislice acquisitions, several asymmetric sequences have been developed. PICORE (Proximal Inversion with Control for Off Resonance Effects) [72], uses a non-selective pulse in the control acquisition. EPISTAR was adapted by Edelman et al. for multislice acquisitions [61] by using an inversion pulse with an amplitude significantly higher than required by the adiabatic condition [82], and two 180° pulses in the control phase, which cycle the magnetisation back to equilibrium. Both pulses have the same MTC contributions, and these cancel when tag and control acquisitions are subtracted. Using a similar concept, the TILT (Transfer Insensitive Labelling Technique) utilises

a pair of 90° pulses to invert inflowing blood, and for the control another pair of 90° pulses, with a 180° phase shift applied to the second pulse such that the magnetisation is "flipped back" into the transverse plane. This was further refined in the PULSAR [80] and QUASAR sequences [181], where adiabatic pulses were used to improve robustness to B_1^+ inhomogeneities at high field. For the control acquisition two 180° adiabatic fast passage pulses were used with half the RF power of the labelling pulse to produce a net nutation of 0°. Figure 2.21a. and b. shows a schematic summary of how these methods label blood.

Conversely, the FAIR sequence is inherently symmetric about the imaging slices/volume, automatically compensating for MTC. It is based on the subtraction of a non-selective inversion acquisition from a slice-selective inversion acquisition. The non-selective inversion inverts blood in all vessels within the transmitter coil, as shown in figure 2.21.c, which during TI flows into the imaging slice. The slice-selective inversion only inverts magnetisation within the imaging volume, therefore non-labelled, fully relaxed blood flows into the imaging volume and no perfusion weighting is present in the control acquisition. Symmetric PASL sequences label blood both proximally and distally to the imaging volume, which can mean that venous flow is also be labelled and contribute to the perfusion signal; this is not observed in double inversion asymmetric PASL (excluding the original EPISTAR because of its distal control slab) because only proximal arterial blood is labelled.

In both symmetric and asymmetric PASL sequences the size of the inverted bolus of blood is limited by the size of the transmitter coil. A typical body transmitter coil has a length of approximately 50cm, which with the imaging slices positioned at the magnet isocentre (hence centre of the transmitter coil) means that a slab of length less than 25cm can be labelled. This finite bolus size means that the ΔM signal, hence SNR for a PASL experiment is less than for an equivalent CASL experiment. Another problem is that the bolus of blood is defined in space, and not in time, as is required in the General Kinetic Model (section 2.3.1), making the ΔM signal sensitive to the transit time as well as the flow rate. Three techniques have been proposed which solve this issue, QUIPSS, QUIPSS II, and Q2TIPS, which use saturation pulses to truncate the labelling bolus after a specified duration. More details on these techniques are given in section 2.3.3.1. Despite the reduced theoretical sensitivity in comparison to CASL (which should provide greater SNR by a factor of e [241]), in practise the difference between the two methods is less because: transmit time is increased in CASL, so more T_1 decay of the bolus; less MTC is present in PASL than in single coil CASL because a larger gradient is used; the inversion efficiency is higher in PASL because frequency-modulated adiabatic pulses are used; and the optimal TR in PASL is shorter than in CASL (where long labelling durations are favourable), therefore the SNR per unit time is higher (i.e. more averages in less time). This, combined with their ease of implementation on clinical systems and overall flexibility have meant that PASL sequences have become the prevailing ASL methodology.

2.3.3.1 Defining the Bolus Duration

In the initial implementations of pulsed ASL, the dependence of the ΔM signal on the transit time meant that quantitative CBF measurements were not possible with a single subtraction; i.e. one set of tag and control measurements at a single TI. Different regions within the brain have different tissue transit times, which in humans is on the order of the T_1 of arterial blood, resulting in variable amounts of T_1 decay between the inversion and acquisition of an image, which unless accounted for becomes a source of systematic error in the CBF quantification. It was only possible to produce quantitative results by acquiring multiple tag and control images for a series of TI times, from which the time course ΔM signal at each voxel can be fitted to a kinetic model, such as the one given in equation 2.64 section 2.3.1, to obtain both the CBF and the transit time. Whilst useful, this requires a long scan time which is not always practical for some patients.

Modifications to the labelling schemes have been suggested to reduce sensitivity to transit times and permit a quantitative measurement in a single subtraction. The first two modifications, proposed by Wong et al. were QUIPSS (QUantitative Imaging of Perfusion using a Single Subtraction) and QUIPSS II [243]. Both apply a saturation pulse[16] at a delay of TI_1 after the initial inversion pulse, which has inversion time TI_2 with respect to the acquisition, and the image is then acquired after a further delay of $\Delta\mathrm{TI} = \mathrm{TI}_2 - \mathrm{TI}_1$. Figure 2.22 shows this schematically.

In QUIPSS the imaging slice is saturated, destroying the contribution of any labelled blood which has arrived into the imaging slice during the time TI_1. Providing TI_1 is greater than the tissue transit time, Δt (from equation 2.64), and TI_2 is shorter than the tissue transit time plus the actual duration of the bolus of blood, τ (this differs with subject and is dependent on the size of the RF transmitter coil), then the inverted blood entering the slice has a well defined duration of $\Delta\mathrm{TI}$, andΔM is given by equation 2.74,which does not contain the terms Δt or τ :

$$\Delta M(\mathrm{TI}_2) = 2M_{0,b} f \Delta\mathrm{TI} e^{-\mathrm{TI}_2/T_{1,b}} q_p \tag{2.74}$$

QUIPSS II saturates the inversion slab at a time TI_1, such that any blood flowing into the slice after this pulse does not contribute to the ΔM signal. Providing TI_1 is shorter than the tagged bolus length, and $\Delta\mathrm{TI}$ is greater than the tissue transit time, the inverted blood has a well defined duration of TI_1, removing the dependence ofΔM on τ and Δt, as shown in equation 2.75. QUIPSS II can be implemented in multislice acquisitions, whereas QUIPSS cannot.

$$\Delta M(\mathrm{TI}_2) = 2M_{0,b} f \mathrm{TI}_1 e^{-\mathrm{TI}_2/T_{1,b}} q_p \tag{2.75}$$

A further modification was made to QUIPSS II by Luh et al. [145], reducing residual errors from incomplete saturation and spatial mismatch of the saturation and inversion slice profiles.

[16] The magnetisation is saturated and then spoiled with a strong gradient to eliminate any transverse coherence.

Incomplete saturation arises due to B_1 inhomogeneities, resulting in a CBF overestimation, and the adiabatic inversion pulses used in PASL methods have a much sharper slice profile than that produced by the sinc saturation pulses. To minimise these errors the QUIPSS II saturation pulse was replaced by a periodic train of thin slice saturation pulses at the distal end of the tagged region. Inflowing blood travelling through this region will be saturated, and the slice profile from a thin sinc pulse has much sharper edges than a thick slice.

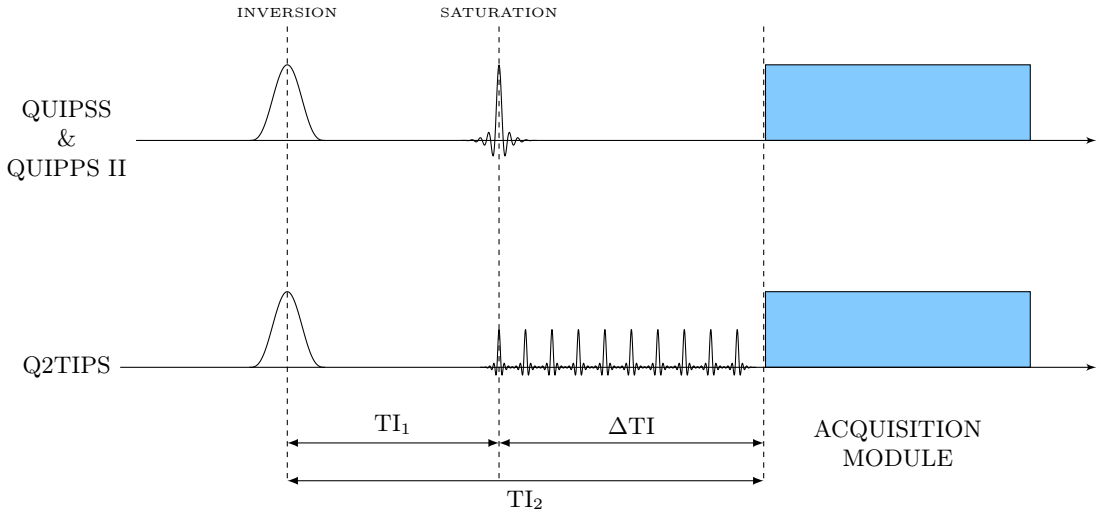

Figure 2.22: QUIPSS, QUIPSS II and Q2TIPS Bolus Truncating Schemes

2.3.4 Vessel Selective ASL

Vessel selective ASL (VS-ASL, also termed regional ASL, selective ASL, territorial ASL, perfusion territory imaging) is a technique which permits the visualisation of the perfusion territory from a feeding set of arteries. Techniques exist for all three labelling schemes (PASL, CASL and pCASL), spatially localising the labelling only to certain arteries, such that when a perfusion image is acquired only they contribute to the ΔM signal. The ability to visualise individual perfusion territories is of clinical interest in many applications.

In healthy population studies it was found that there is a large interindividual variability of perfusion territories, predominantly due to variability in the structure of the circle of Willis [220, 221].

In the diagnosis and treatment planning of cerebrovascular disease VS-ASL can be used to identify regions at risk from ischaemia, plan treatment, and assess recovery progress. Normally the carotid arteries are the main source of blood supply to the brain, however in cerebrovascular disease flow through these arteries is compromised, for example due to the formation of a stenosis (abnormal narrowing), or increased vulnerability to changes in blood pressure, resulting in the formation of a blood clot or rupture of the blood vessels. Ultimately, cerebrovascular disease

can lead to a stroke, resulting in an acute reduction in perfusion to a region of the brain and subsequent ischaemia. VS-ASL can be used to diagnose cerebrovascular disease, such as occlusions, view the extent of ischaemic regions, and determine the adaptation of the vasculature to provide a prognostic assessment of collateral circulation for aiding surgery in patients requiring vascular interventions [99]. In acute stroke VS-ASL can be used to identify regions susceptible to ischaemia and infarction[17], demonstrate the source of collateral blood supply tissue surrounding the ischaemic core [141], provide guidance for treatment planning, and a means of monitoring its progress [220].

In Moyamoya disease, a chronic cerebrovascular disease characterised by progressive stenosis of the supraclinoid (a segment of the internal carotid artery) circulation, patients show extensive development of collateral vessels by the time of diagnosis. Presence of collaterals is associated with better recovery after thrombolytics (drugs used to dissolve blood clots), therefore this disease is of interest in the study of acute stroke. Digital subtraction angiography (DSA)[18] is commonly used for judging the presence of collateral vessels, however it is time consuming, invasive, and costly. Studies in which Moyamoya patients were scanned with regional [39] and global [258] ASL sequences and compared with routine DSA for identifying collateral vessels; the diagnostic power of the ASL methods was found to be equivalent to DSA, and in the global ASL study it was suggested that perfusion territory imaging may further aid diagnosis.

Vessel selective ASL also has potential use in identifying hyperperfused regions of tumours [239, 221], providing functional information about their vascular territory and supply [197].

Many different methods exist to selectively label blood flowing to the brain. A small surface coil, positioned directly over one of the carotid arteries to localise the labelling B_1^+ of a continuous ASL (CASL) labelling scheme, provides good left/right delineation, and complete elimination of magnetisation transfer (MT) effects as the imaging volume is not exposed to the long, off-resonance RF pulse [259]. Similarly, a labelling coil with coverage of both carotid arteries, combined with an oblique (rotated by 60°) labelling gradient, ensures that the resonance condition, and hence adiabatic inversion only occurs within one of the common carotid arteries [176, 54]. Both of these methods require specialist hardware, and the correct placement of the labelling coil on the neck.

When a separate labelling coil system is unavailable, for example in the clinical setting, selective inversion is possible using only a volume transmitter in combination with gradients. The oblique labelling gradient method by Werner et al. has been implemented using a head only transmit-receive coil to perform the labelling, utilising the fact that outside a volume transmitter the B_1^+ field drops off to low levels [233]. In another CASL based method by Werner et al., an oscillating labelling plane is centred about the artery of interest to provide good labelling at this point, yet causes dephasing of the magnetisation of flowing blood in other arteries [232]. However the labelling

[17] Necrosis due to a localised lack of oxygen.

[18] A fluoroscopic imaging method used to image blood vessels. X-ray images are continuously acquired (typically using an x-ray image intensifier and CCD camera) whilst a bolus of a contrast agent such as iodine is injected into a vessel, and a secondary image, acquired without any contrast agent is subtracted, removing the distracting structures.

mechanism can contaminate the perfusion weighted images with dark bands as the static tissue is also dephased; and with a volume transmitter coil MTC must be accounted for, for example by using amplitude modulated control to facilitate multi-slice acquisitions.

Several pulsed ASL (PASL) based methods have used slab selective inversion pulses, oriented to label specific cerebral arteries. [214, 99, 80, 221, 220, 59]. Taoka et al. [214] used the subtraction of two inversion slabs positioned equidistantly either side of oblique imaging planes to provide perfusion weighting and account for MT. Hendrikse et al. [99] introduced Regional Perfusion Imaging (RPI), which used two consecutive 90° pulses to create inversion slabs, in the same manner as a transfer insensitive labelling technique (TILT) [81]. A control acquisition was achieved by inverting the phase of the second 90° pulse. RPI was used to selectively label the left internal carotid artery (ICA) , right ICA and the posterior circulation (basilar and vertebral arteries). A more robust method, PULSAR (PULsed Star labelling of Arterial Regions) was developed by Golay et al. [80], combining adiabatic inversion pulses to counter high field flip angle variation within the brain, and a water suppression enhanced through T_1 effects (WET) pre-saturation technique. More recent developments have also introduced time-efficient methods to produce the perfusion territories of multiple regions, by acquiring combinations of tagged regions and to separate them in post-processing [89, 263, 242].

Two-dimensional RF pulses have been used to create a region of inversion along the length of arteries [46]. A temporal RF waveform was computed based on a spiral gradient trajectory, generating a Gaussian-weighted "pencil" inversion profile for selecting individual arteries. A labelling bolus was built up from several "pencil beam" inversions. While the technique showed promising results, it suffered from contralateral labelling, in part due to B_0 inhomogeneities, and the perfusion signal was significantly lower than in a comparatory PICORE experiment.

Pseudo-continuous labelling has been made selective by adding additional transverse gradients which create phase shifts between vessels of interest, and modulating the phase of the labelling RF pulses to place vessels in tag and control conditions [242]. This was combined with a Hadamard [177] encoding scheme to obtain perfusion territories from a set of vessels in a SNR efficient manner.

2.4 Summary

To summarise, this chapter has described the key principles of magnetic resonance imaging, radio frequency coils for MRI, and arterial spin labelling. The quantum origin and classical description of the NMR signal has been introduced, along with the principles of excitation, signal detection, and relaxation mechanisms. Concepts regarding the underlying principles of image formation which apply to most imaging techniques; the gradient echo and spin echo sequences; T_1, T_2 and proton density contrast; gradient and RF spoiling; and FLASH and EPI sequences have been detailed. The hardware components which form an MRI scanner have been introduced, and an overview of

parallel transmission methods has been given. A theoretical and practical overview of the design of RF coils for MRI was presented. The three main ASL schemes have been introduced: continuous ASL, pseudo-continuous ASL and pulsed ASL, along with how they are implemented, challenges that have been overcome, and quantification methods. Finally, an overview of vessel selective ASL methods was given.

The work presented in the following chapters combines the concepts of parallel transmission and vessel selective arterial spin labelling, with the objective of improving vessel specificity whilst retaining the high labelling efficiency and SNR that is associated with separate-coil continuous arterial spin labelling.

Chapter 3

Computer Simulation of ASL

This chapter describes the development and use of an *in-silico* simulation of blood flowing within the arteries in the neck, and the effect of the gradients and radio frequency magnetic fields imposed on the bulk magnetisation of the inflowing blood. Primarily the aim was to use the simulation as a development platform for applying and optimising parallel transmission methods to ASL. Additionally the simulation framework is highly configurable and has potential to be used for sequence optimisation of a variety of ASL scenarios.

Bloch equation simulators have been previously employed to optimise the labelling parameters of an ASL experiment. Maccotta et al. [147] developed theoretical and numerical methods to calculate the labelling efficiency over a wide range of experimental and physiological parameters which were compared with in-vivo results. It was found that the RF amplitude can be reduced significantly whilst maintaining labelling efficiency by decreasing the labelling gradient strength, and that there is an optimum RF amplitude for a given gradient strength. When experimental parameters are appropriately chosen labelling efficiency is largely independent of flow velocity. A higher labelling efficiency was found to be achieved than was previously assumed.

Utting et al. [218] established optimum conditions for velocity driven adiabatic fast passage (AFP) in both a human and animal model, and determined the effect of pulsatile flow due to the cardiac cycle and RF duty cycle on the degree of inversion. Variations in the blood velocity over the course of a cardiac cycle were found to affect the instantaneous labelling efficiency, however the mean degree of labelling was found to be similar to that of blood moving at a velocity equal to the time-averaged mean. Substituting the continuous RF for a train of RF pulses was found to result in a highly non-linear dependence of the degree of inversion on the RF duty cycle.

Gach et al. [79] made MR measurements of the inversion efficiency in a flow phantom capable of laminar and turbulent flow to study the effect of flow dynamics, labelling field amplitude and duty cycle on the inversion efficiency of the AFP. Again a non-linear relationship was observed between labelling efficiency and duty cycle, and with the exception of high labelling field amplitudes

good agreement was shown with models based on Zhernovoi's theory [262]. Numerical Bloch equation simulations of AFP with a train of RF pulses were found to reproduce the non-linear behaviour of inversion efficiency with increasing B_1 amplitude, but did not accurately reproduce the experimental findings of duty cycle due to the simulation not incorporating spin dephasing and field inhomogeneities.

Gach and Dai [78] developed a simple analytic model for double adiabatic inversion (DAI) amplitude modulated control CASL (see section 2.3.2), which was compared with experimental and numerical Bloch simulation data. DAI efficiency can be significantly less than 100%, and so it is important to quantify this in an examination for accurate CBF quantification. They found the Bloch equation simulations to grossly mimic the experimental results when many velocities are averaged together, but not useful in estimating DAI or single adiabatic inversion (SAI) efficiencies for real flows due to the high sensitivity of the non-linear equations to experimental conditions.

3.1 Simulation Development

3.1.1 Anatomic and Physiological Model

A simplified model of blood flow in the neck arteries is required to efficiently simulate the labelling process (by limiting the computational complexity), yet still retaining accuracy to ensure the usefulness of results. Labelling only concerns the neck region, specifically the carotid and vertebral arteries. Gray's Anatomy [209] shows that over this region they run approximately parallel with the neck, which can be considered cylindrical in shape. The heart pumps blood in a pulsatile fashion, and its velocity waveform can be characterised using only a few parameters as shown by Holdsworth et al. [106]. As the arteries are smooth walled, elliptical in shape (but close to circular), and a labelling plane is usually defined in the middle of the neck (in the head-foot direction) some distance away from any bifurcation, the flow profile within one of the neck arteries can be considered to be laminar [235]. Therefore it is a reasonable approximation to construct a model where the neck is a cylinder along the z-axis, the two carotid and two vertebral arteries are smaller cylinders running parallel to the axis of the neck, and that the flow of blood within the neck is laminar and has a parabolic velocity profile.

Thus a simplified model can be constructed with the centre of the neck at the origin and all features of the neck parametrised into the following:

- The radius of the neck R_{neck}.
- Positions of each artery within the neck $\mathbf{r_{artery}}$.
- Arteries have circular cross sections with radius R_{artery}.

- The blood velocity within each artery is given by the function

$$\sigma_{artery}(r,t) = v_{artery}(t)P_{artery}(r) \tag{3.1}$$

where $v_{artery}(t)$ is a temporal velocity waveform (given by a model such as [106] or from measurements); and $P_{artery}(r)$ is a normalised radially parabolic velocity profile within the artery in the x-y plane with roots ($P_{artery} = 0$) at the edges of the artery, R_{artery}, given by

$$P_{artery}(r) = -\left(\frac{r}{R_{artery}}\right)^2 + 1 \tag{3.2}$$

where r is the radial position of the blood from the centre of the artery.

Method

The model was verified with anatomical data collected from a MRI scan of a 24-year-old healthy male volunteer using the following protocol:

1. Gradient Echo Sagittal and Coronal slices through the head and neck, used for planning the angiograms in 2. and 3. TE = $5.25ms$, TR = $15ms$, $\alpha = 40°$, 10mm slice thickness, $matrix\ size = 256 \times 256$ FOV $= 350 \times 350mm$, $BW = 180Hz/px$.

2. 3D time of flight (TOF) (p704-724 [91]) angiogram of the carotid and vertebral arteries, TE = $3.69ms$, TR = $20ms$, $\alpha = 24°$, $matrix\ size = 384 \times 384 \times 218$ FOV $= 384 \times 384 \times 200mm$, $BW = 250Hz/px$.

3. A transverse, single slice, cardiac gated, phase contrast angiogram (p725-733 in [91]), located half-way between the branching of the right subclavian and right carotid arteries from the brachiocephalic artery, and the bifurcation of the carotid artery into the internal and external carotids. 30 frames were taken over the cardiac cycle, each having TE = $5.88ms$, TR = $31.1ms$, $\alpha = 10°$, 1mm slice thickness, $matrix\ size = 384 \times 384$ FOV $= 200 \times 200mm$, $BW = 250Hz/px$, $v_{enc} = 1.3ms^{-1}$.

Images were acquired on a Siemens 3T Tim Trio (Erlangen, Germany) using a body coil for transmission and combined twelve channel head matrix coil and four channel neck matrix coil for reception. Analysis was performed in Matlab (The Mathworks Inc., Natick, MA) using a semi-automated script. Elliptical regions of interest (ROIs) were drawn on the first transverse slice of the 3D TOF dataset to identify each of the arteries, then the centre of each artery was found using the assumption that the brightest pixel is at the centre. The artery width and height were found by stepping along the pixel index in the $\pm$x and $\pm y$ directions until the pixel brightness dropped below a threshold value (half the maximum was empirically found to be a suitable threshold). An elliptical ROI was drawn on every 10^{th} slice to manually segment the neck. Pixel indices were

converted to dimensions using the dicom meta-data of the slice position, image size and voxel dimensions.

A rectangular ROI was drawn on a sagittal slice through the centre of the head and neck to indicate the extent of the neck in the head-foot direction and also around the brain representing the imaging volume in an ASL experiment. Over the region of the neck as given by the ROI the average neck radius, artery position and artery radius were calculated along with associated standard deviations.

Phase contrast data was converted from a 12-bit unsigned integer to a velocity map[1] and then ROIs were drawn around each artery on the first image in the temporal dataset. A temporal velocity waveform was computed from the maximum velocity in each ROI and a normalised velocity profile in the x and y directions for comparison against the parabolic flow profile model from equation 3.2.

Results

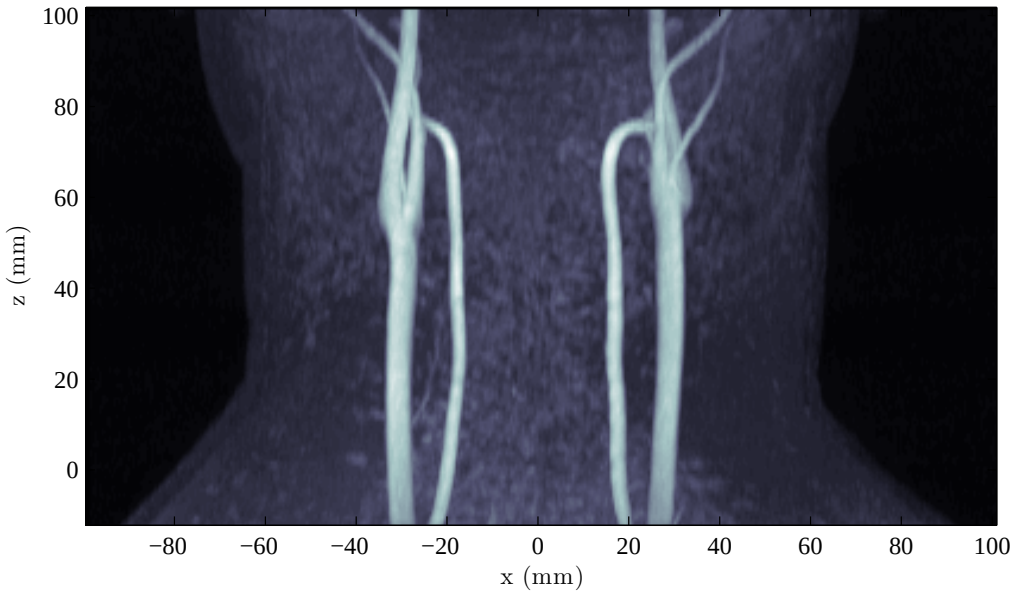

Figure 3.1: Coronal Maximum Intensity Projection (MIP) from the 3D TOF dataset

[1] A 12 bit unsigned integer has a value between 0 and 4095 (2^{12} values). In the case of the phase contrast dicom data a value of 0 represents flow at $-v$ with respect to the direction of the flow encoding gradients, and a value of 4095 represents flow at $+v_{enc}$ with respect to the direction of the flow encoding gradients. All other values are scaled linearly between these extents. The velocity v_i at each i^{th} pixel, p_i is calculated as

$$v_i = v_{enc}\left(\frac{2p_i}{4095} - 1\right)$$

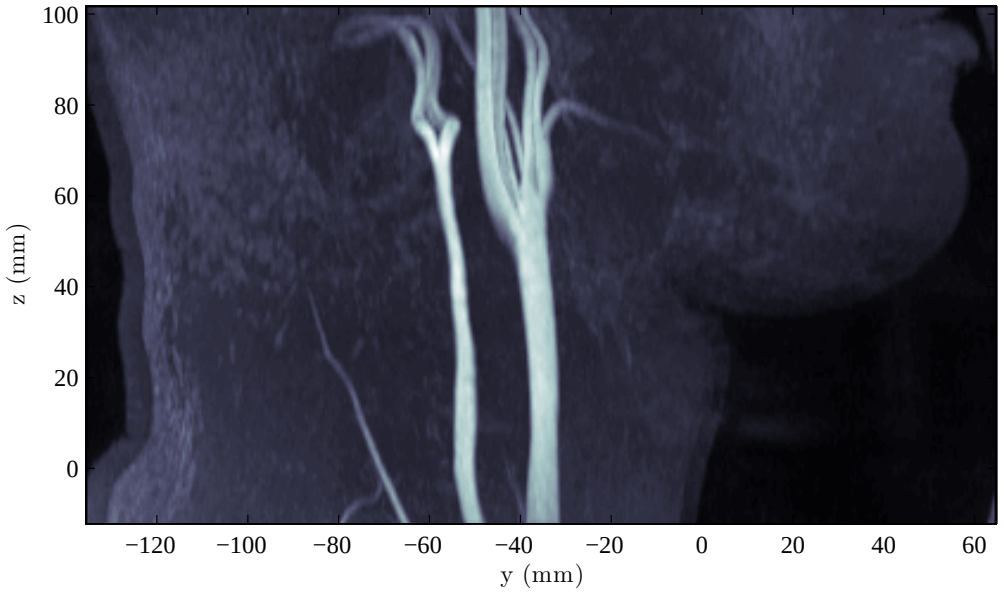

Figure 3.2: Sagittal MIP from the 3D TOF dataset

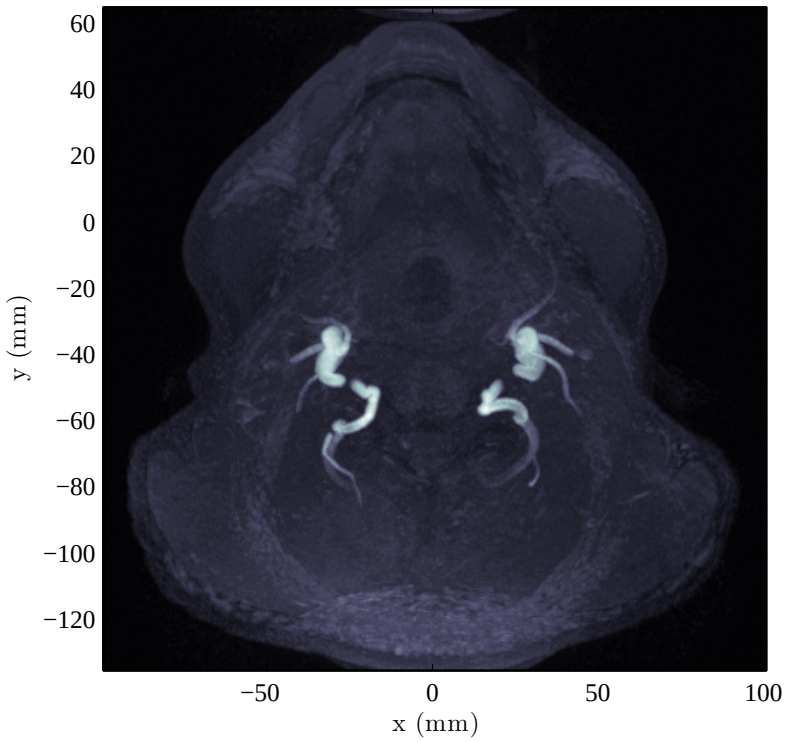

Figure 3.3: Transverse MIP from the 3D TOF dataset

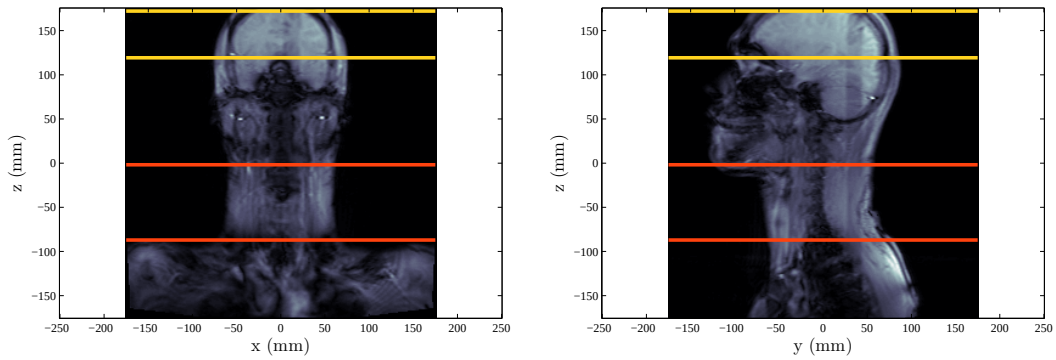

Figure 3.4: Coronal and Sagittal localiser images
Red lines represent the extent of the neck, defined to be between the shoulders to the point the chin juts out. Yellow lines represent a typical imaging volume in an ASL acquisition, for example twelve 5mm thick slices covering the most of the frontal and parietal lobes. These extents were used to define a new co-ordinate system with the centre of the brain at $z' = 0$.

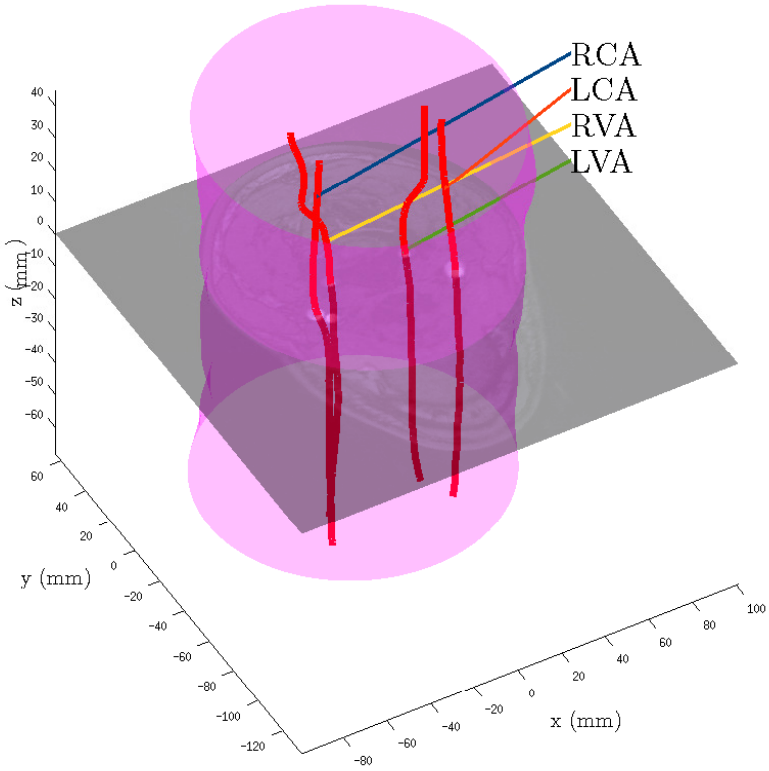

Figure 3.5: Schematic plot of the neck and artery information extracted from the TOF dataset
The 3D pink surface is interpolated between the 10 elliptical ROIs drawn around the neck at their respective slice positions. The red lines are 3D plots of the coordinates of the central point of each artery for each slice. RCA - Right Carotid Artery, LCA - Left Carotid Artery, RVA - Right Vertebral Artery, LVA - Left Vertebral Artery.

x_{neck}, width (mm)	125.0±3.5
y_{neck}, depth (mm)	130.7±4.1
R_{neck}(mm)	63.9±1.3
$z_{neck\ start}$(mm)	-232.8
$z_{neck\ length}$(mm)	85.2
$z_{brain\ length}$(mm)	52.9

Table 3.1: Anatomical dimensions established from the MRI scans
It should be noted that all of these measurements were made by manually drawing ROIs on MR images. The R_{neck}measurements are an average of the x and y dimensions of the ROIs drawn around the neck, over the range defined by $z_{neck\ start}$ and $z_{neck\ end}$ (between z = -232.8 and z = -147.6 mm).

Artery Locations As can be seen from figure 3.5 and table 3.1, the neck's radius does not vary significantly in size and position. The difference between the width and depth is approximately 1mm larger than the standard deviation in width and depth, thus the neck can be considered cylindrical.

	Position		**Size**		
	x(mm)	y(mm)	x(mm)	y(mm)	R_{artery}(mm)
RCA	28.8±3.5	27.5±4.4	7.0±0.6	7.2±0.7	3.5±0.4
LCA	-29.6±4.2	27.9±4.2	6.8±0.4	6.9±0.3	3.4±0.3
RVA	17.4±3.8	11.2±4.3	4.3±0.3	4.3±0.3	2.2±0.2
LCA	-18.6±3.6	10.5±4.2	4.7±0.2	4.6±0.3	2.3±0.2

Table 3.3: Arterial positions and dimensions

Table 3.3 shows the average positions and sizes of the arteries over the extent of the neck and their standard deviations. The average variation of the centre of each artery is 3.8mm and 4.3mm in the x and y directions respectively, and when compared to the size of R_{neck} this deviation is minimal. The x and y dimensions of each artery are almost the same, varying by less than 10%. Thus the assumption that the arteries run parallel to the z axis and have a circular cross section is fair.

Flow Velocity A total of 44 phase contrast images (figure 3.6 shows one of these images) were acquired over a cardiac cycle lasting 910ms, with the time interval between images alternating between 20ms and 22.5ms (i.e. interval between 1st and 2nd image is 20ms, 2nd and 3rd is 22.5ms, 3rd and 4th is 20ms). Figure 3.7 is a plot of the maximum velocity within each artery over the time series.

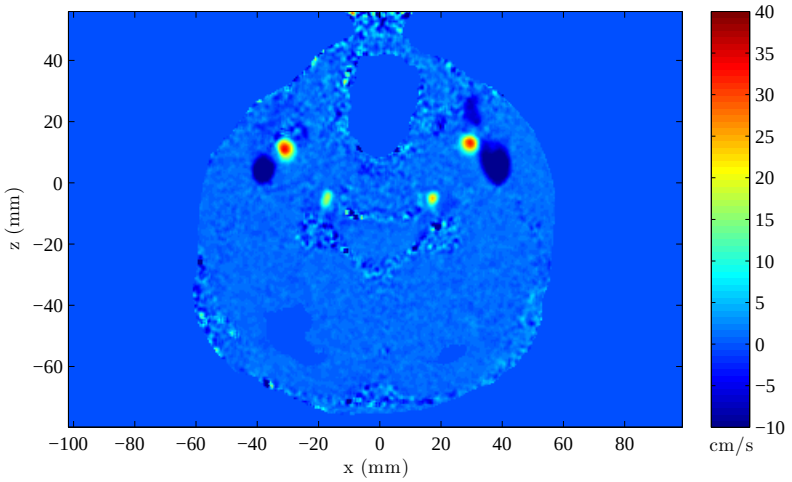

Figure 3.6: Phase contrast image of the neck at the peak velocity
Positive values correspond to flow towards the head (anterior direction). The carotid and vertebral arteries can be clearly seen as the four regions of high anterior flow towards the centre of the neck and next to each carotid artery the jugular veins can also be seen as regions of posterior flow.

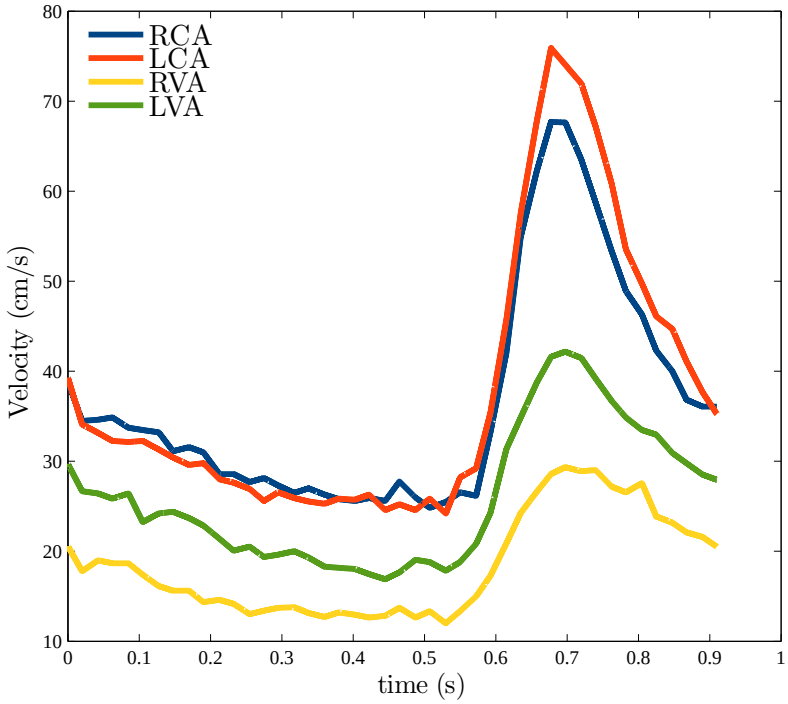

Figure 3.7: Arterial Velocity Waveforms

Artery	v_{max}(cm/s)	v_{min}(cm/s)	$\langle v \rangle$(cm/s)	σ_v(cm/s)
RCA	67.7	24.8	36.5	12.5
LCA	75.9	24.2	37.7	15.3
RVA	29.3	12.0	18.4	5.7
LVA	42.2	16.9	26.0	7.6

Table 3.5: Arterial velocity data

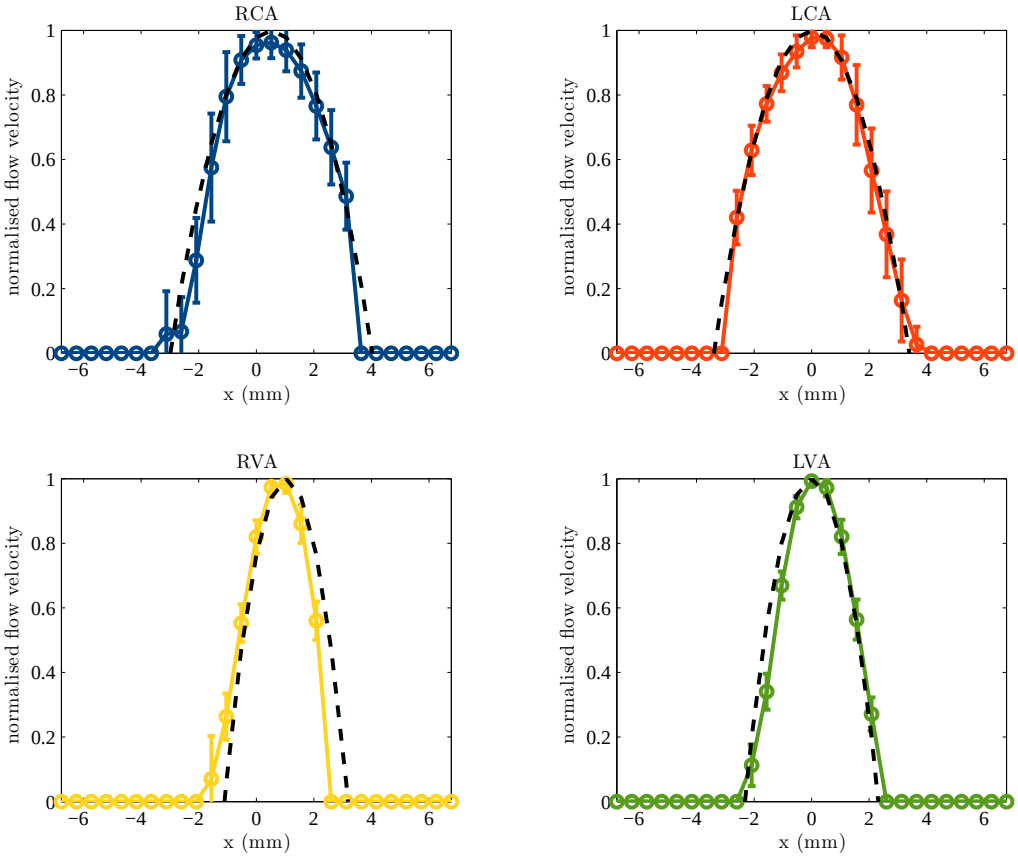

Figure 3.8: Normalised arterial flow profile in the x direction

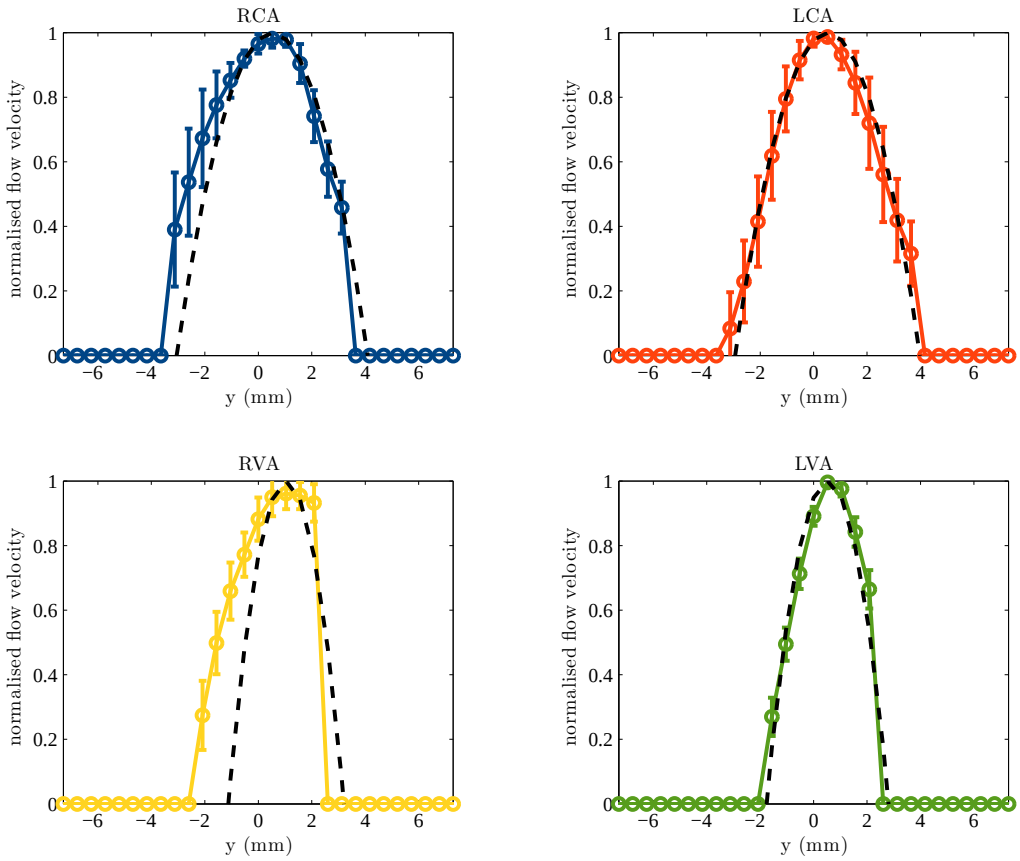

Figure 3.9: Normalised arterial flow profile in the y direction

Figures 3.8 and 3.9 are normalised arterial flow profiles in the x and y directions through the centre of each artery, along with a plot of equation 3.2 for the corresponding arterial width (black dotted line). Each data-point was normalised to the peak velocity within the artery at that particular time-point and then an average normalised velocity for every time point was calculated. Error bars of length one standard deviation above and below indicate the deviation of the normalised profile from the average over the time-series. For the most part the plot of equation 3.2 sits within these error bars.

Given that only a few pixels span each artery (image resolution 0.5802mm/pixel in-plane and the arterial width is approximately 4.5 - 7mm) and associated noise and imaging artefacts that may confound any flow measurements, the normalised velocity profile plots in figures 3.8 and 3.9 show a good agreement with the modelled arterial flow, therefore not only is it suitable to use equation 3.2 in designing the cross-sectional flow profile, but it also validates the hypothesis of laminar flow at this particular point within the arteries.

Discussion and Conclusion

Presented is a simplified model of the anatomy involved in an ASL experiment for use in an *in-silico* simulation of the RF tagging procedure. The model's assumptions were:

1. The neck can be represented by a cylinder running parallel to the z-axis.
2. The two carotid and two vertebral arteries run within the neck also parallel to the z-axis.
3. Each of these arteries is cylinder (and so has a circular cross section).
4. The velocity of the blood flowing in each artery has a spatial and temporal dependence given by equation 3.1.
5. The spatial dependence is based on the assumption that the flow in each artery is laminar and hence has a parabolic flow profile through the centre of each artery.

MRI images were acquired and an analysis performed to test this hypothesis. Results show that for this particular volunteer all five assumptions are valid. The output of this analysis is the neck radius R_{neck}, positions of the neck along the z axis, the length along the z-axis of the brain/imaging volume, radii for each artery R_{artery}, position vectors of these arteries with respect to the centre of the neck $\mathbf{r_{neck}}$, and a velocity waveform for each artery. From these a model which approximates the arterial flow of a subject and is appropriately scaled can be constructed. While the analysis has proved to be a successful preliminary study for this particular subject more volunteers would need to be scanned with the same protocol and the same procedure before the parametrised model of the neck can be completely validated. Some bias may have been introduced into the analysis script because it was developed using the acquired data as test data.

3.1.2 Computational Model

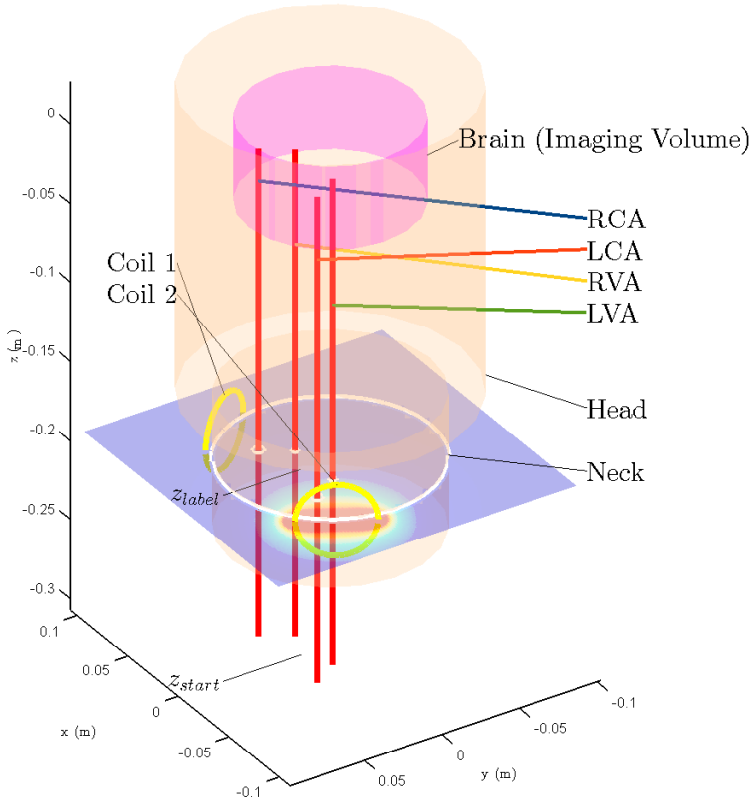

Figure 3.10: Simulation Model
Dimensions obtained in section 3.1.1 were used to create this 3D model. The "head" cylinder is for illustrative purposes, and the field plot is from two surface coils positioned at the labelling plane, operating in a "Helmholtz" (in-phase) configuration.

Blood was represented as a uniform density flow of particle objects, each of these an ensemble of spins with a macroscopic magnetisation $\mathbf{M}$, and position $\mathbf{r}(x, y, z)$, referred to as a spin object. Upon initialisation each spin object was randomly assigned a position in the x-y plane within the confines of the artery's radius. Blood is of uniform density within the arteries (as blood mostly consists of water which is incompressible at the pressures found within the human body), so there is equal probability of the spin object being assigned any transverse position within a given artery. Using the randomly generated transverse position the parabolic velocity weighting function $P_{artery}(r)$ was computed as given in equation 3.2. Spin objects were initially placed uniformly along z, from $z = 0$ to $z = z_{start}$ (so that the most recently created spin was initially at $z = z_{start}$, see figure 3.10), with the spacing between spin objects determined by ρ_{spins}, the spin density. During execution of the simulation new spin objects were regularly created and initialised at $z = z_{start}$ to maintain the spin object density. When reaching the location $z = 0$, the tissue bed, the velocity of a spin object was set to zero. Spin objects at $z = 0$ with longitudinal magnetisation

(M_z) greater than 0.999 were removed from memory as they had negligible contribution to the perfusion signal.

Over the simulated space the labelling gradient G_z was assumed to be constant and the frame of reference was that of the rotating frame at the isocentre of the magnet ($z = 0$); $\mathbf{G} \cdot \mathbf{r} = 0$ at this position. For an in-vivo ASL experiment the shims are typically optimised over the imaging volume (particularly for B_0 sensitive imaging sequences such as EPI or 3D-GRASE which are commonly used in perfusion studies for their high temporal resolution and SNR), resulting in what can be large inhomogeneities over the labelling region. An option could be selected to include these inhomogeneities as an additional component to B_z, such as a 3D raster of noise, or field maps obtained from in-vivo scans.

The radio frequency, transverse, positively polarised magnetic field, B_1^+, could be modelled for two types of coil configuration: a body coil or an array of surface coils positioned around the neck. On a typical whole-body scanner the birdcage body coil is approximately 50cm long and at 3T and for head and neck geometry it can be assumed that the transmit field is homogeneous within the coil, therefore in this case the value of B_1^+ is the same at all points in space.

Surface coils were modelled as circular current loops placed on the neck at distance R_{neck} from the z axis and angle $\theta_{coil}$from the x axis. A quasi-static approximation to the B_1^+ field produced by a surface coil was calculated by numerically integrating the Biot-Savart law around each loop using elliptical integrals (p266, [207]), yielding a 3D raster of spatial field samples covering $[x, y, z]$ coordinates of $[-1.5R_{neck}, -1.5R_{neck}, z_{start}]$ to $[1.5R_{neck}, 1.5R_{neck}, 0]$, each sample having B_x, B_y and B_z(calculated for completeness) components. It is the positive circularly polarised field (B_1^+) that interacts with the nuclear magnetisation, and this is calculated in complex form with $\mathbf{B_1^+} = \frac{B_x + iB_y}{2}$[109]. The complex B_1^+ field from each coil is then summed according to the following equation:

$$\mathbf{B^+_{1,total}} = \sum_{r=1}^{R} a_r e^{i\phi_r} \mathbf{B^+_{1,r}} \tag{3.3}$$

a_r and ϕ_r are the modulation of the amplitude and phase of the driving current in coil r respectively, see section 3.2 for the calculation of these.

Coil-sample interactions present at 3T are complex and the Biot-Savart Law is only a quasi-static approximation. For improved accuracy the field from each coil could be calculated with a full-wave analysis of Maxwell's equations [244], or obtained from experimentally acquired B_1^+ maps (see section 7.2).

A small time-step ($dt = 10\mu s$) was used to ensure that the rotation of the B_1^+ field was not aliased. At each time point every spin object was moved along by $\Delta z = \sigma_{artery}(r, t)dt$ (where $\sigma_{artery}(r, t)$ is the arterial velocity waveform weighted by the position of the spin object in the artery, as given by equation3.1). Every spin object was then operated on by the Bloch equation

kernel (see section 3.1.3).

3.1.3 Bloch Equation Kernel

For each time-step the evolution of the magnetisation vector of every spin was be determined by a rotation matrix formalism of the Bloch equation [21]. The magnetisation of a spin isochromat is a vector of unit length, aligned with the z-axis at equilibrium. The Bloch equation shows that the magnetisation will experience a torque when subjected to a magnetic field, and will subsequently be rotated. By determining the magnitude and direction of the applied magnetic fields the degree of rotation that the magnetisation experiences due to this torque can be calculated and the magnetisation vector operated on accordingly.

$\mathbf{B_1^+}$ is a vector in the x-y plane of length

$$\left|\mathbf{B_1^+}\right| = \sqrt{B_{1,x}^2 + B_{1,y}^2} \tag{3.4}$$

at azimuthal angle

$$\phi = \phi_{B_1} + \phi_{\Delta f} \tag{3.5}$$

to the x axis, where $\phi_{B_1} = \arctan\left(\frac{B_{1,y}}{B_{1,x}}\right)$, $\phi_{\Delta f} = 2\pi\Delta f t$ ($\mathbf{B_1^+}$ rotates about the z axis with frequency Δf).

$\mathbf{B_z}$ is a vector aligned with the z axis of length

$$|\mathbf{B_z}| = G_z z\hat{\mathbf{k}} + \Delta\mathbf{B_0} \tag{3.6}$$

The magnetisation of the spins is subjected to the vector summation of $\mathbf{B_z}$ and $\mathbf{B_1^+}$; the effective magnetic field, $\mathbf{B_{eff}}$ which has length

$$|\mathbf{B_{eff}}| = \sqrt{|\mathbf{B_z}|^2 + |\mathbf{B_1^+}|^2} \tag{3.7}$$

and makes zenith angle

$$\theta = \arctan\left(\frac{\left|\mathbf{B_1^+}\right|}{|\mathbf{B_z}|}\right) \tag{3.8}$$

to the z axis.

During a time interval dt the magnetic moment of a spin precesses about $\mathbf{B_{eff}}$ by the angle

$$\varphi_{precess} = 2\pi\gamma\left|\mathbf{B_{eff}}\right| dt \tag{3.9}$$

This simple set of equations calculates three angles - the angles $\mathbf{B_{eff}}$ makes with the z (θ)

and x (ϕ) axis, and the precession angle ($\varphi_{precess}$). The Bloch equation kernel determines the magnitude and direction of the applied magnetic field (a summation of both the transmit $\mathbf{B_1^+}$ and the gradients), and the amount that the spin's magnetisation vector will precess about the magnetic field. Using these it numerically integrates the Bloch equation to first order approximation by rotating the magnetisation accordingly:

$$\mathbf{M_{n+1}} = \mathbf{R}\mathbf{M_n} \tag{3.10}$$

The rotation matrix $\mathbf{R}$ is a concatenation of rotations; rotate the co-ordinate system from that of the rotating frame at isocentre to that of the applied $\mathbf{B_1}$ field, rotate the co-ordinate system to that of the $\mathbf{B_{eff}}$ field so that the magnetisation vector can precess about this field. The co-ordinate system must then be returned back to the rotating frame at isocentre by applying inverse rotations:

$$\mathbf{R} = \mathbf{R_z}(\phi)\mathbf{R_y}(\theta)\mathbf{R_z}(\varphi)\mathbf{R_y}(-\theta)\mathbf{R_z}(-\phi) \tag{3.11}$$

Following this the effects of T_1 and T_2 relaxation are calculated and applied to the magnetisation vector using the following formula:

$$\mathbf{M} = \begin{pmatrix} 1-\frac{dt}{T_2} & 0 & 0 \\ 0 & 1-\frac{dt}{T_2} & 0 \\ 0 & 0 & 1 \end{pmatrix} \mathbf{M} + \begin{pmatrix} 0 \\ 0 \\ 1-\frac{M_z dt}{T_1} \end{pmatrix} \tag{3.12}$$

Figure 3.11 shows the results of a simulation of four spins travelling at different, constant velocities. Spins were subjected to a labelling gradient of $3mT/m$ and B_1 of $3.5\mu T$. Inversion efficiency is a function of velocity, if the velocity is too low, as with the 10cm/s case, the magnetisation vector remains tightly locked with $\mathbf{B_{eff}}$ (M_z does not have any oscillations), but the frequency sweep does not occur fast enough for the effects of longitudinal relaxation to be disregarded. At 35cm/s (approximately the average carotid arterial velocity in a healthy adult) the inversion is near optimal, as one would expect from the typical labelling gradient and B_1 strengths. With faster flow the magnetisation vector does not remain as tightly locked with $\mathbf{B_{eff}}$ and the inversion is imperfect. Crucially it is the velocity of the blood as it passes through the labelling plane that has the largest effect on the inversion efficiency. For the pulsatile flow present in the arteries the peak velocity can be two to three times the average velocity, and so the inversion efficiency at this point will be reduced. The plot appears to show different rates of relaxation for different velocities - this is misleading as the x-axis is the position of the spin along z. A faster flow velocity means the blood will move quicker along z, and will relax less as it covers the distance between the labelling plane and the imaging volume (at $z = 0$).

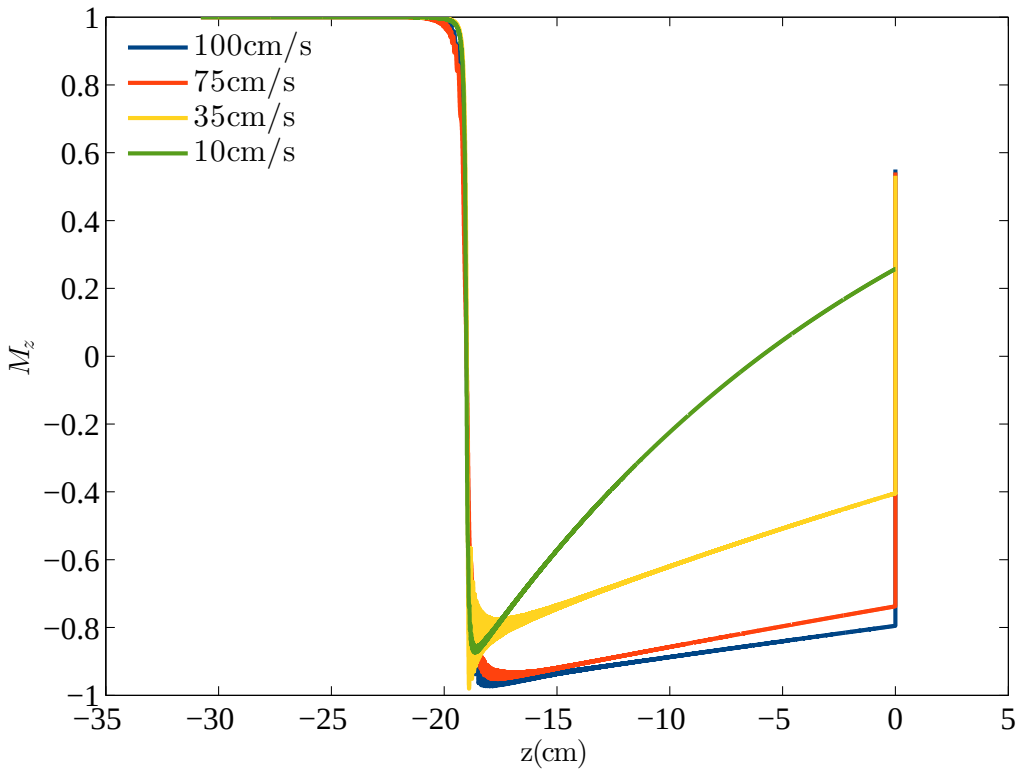

Figure 3.11: $M_z$of a single spin as it travels from z_{start} to $z = 0$ for a range of velocities

3.2 Spatial Localisation of the B_1^+ field

Vessel-selective labelling can be achieved with just one surface coil placed over the arteries to be labelled. Consider the case where a surface coil is placed over the right hand side of the neck with the aim of labelling both the right carotid and vertebral arteries (RCA and RVA), whilst leaving the left carotid and vertebral arteries (LCA and LVA) unlabelled. This arrangement is rarely ideal because:

1. Knowledge about the underlying vasculature must be known to accurately position the surface coil over the RCA and RVA. Correlating imaging data with placing a coil over the neck is challenging and results are not always repeatable.

2. The B_1^+ field along the axis of a surface coil falls off with distance by an inverse-square law, and so will extend across to the contralateral arteries. In order to obtain a sufficiently high B_1^+ for efficiently labelling the RVA and RCA the field over the LCA and LVA will inevitably be large enough to cause some partial labelling.

3. Finer selectivity (such as the RCA only) might be possible in certain instances, such as using a very small surface coil (approximately 2cm diameter), however this will be difficult to position (see point 1. above).

A spatially tailored B_1^+ field is possible using the superposition of multiple B_1^+ fields from independent transmit array elements. Modulating the amplitude and phase of the driving current in

each coil causes constructive and deconstructive interference in the total B_1^+ field. This technique is employed to homogenise the B_1^+in high field (7T+) human systems and is known as B_1^+ Shimming [116]. It has also been implemented for spatial localisation [169]. The degree of control over the transmitted field is dependent on the number of independent transmit coils; typical parallel transmit systems consist of 8 to 16 transmitter channels, meaning that cost and complexity is high. Calculation of the complex weights applied to the driving currents is usually calculated by specifying a desired transmit field pattern and solving the appropriate equations, as explained in the next section.

3.2.1 Forward Problem

Consider a system of R transmit coils, each coil carrying a current I_r and producing a time harmonic magnetic field. The positively polarised field is defined as:

$$\mathbf{B_1^+}(\mathbf{r})=\frac{B_x(\mathbf{r})+iB_y(\mathbf{r})}{2} \tag{3.13}$$

If the driving current in each coil is modulated in amplitude and phase, $I_r = Ia_re^{i\phi_r}$, the superposition of the fields from all coils is:

$$\mathbf{B^+_{1,T}}(\mathbf{r})=\sum_{r=1}^{R}\mathbf{c_r}\mathbf{B^+_{1,r}}(\mathbf{r}) \tag{3.14}$$

Where

$$\mathbf{c_r} = a_re^{i\phi_r} \tag{3.15}$$

Commonly $\mathbf{B^+_{1,r}}$ and $\mathbf{B^+_{1,T}}$ exist as two or three-dimensional arrays consisting of M spatial samples of the B_1 field. It is more convenient to convert these into column vectors (see figure 3.12) of length M by "stacking" each column, doing so makes it possible to operate on them using matrices. Matlab stores all arrays (regardless of how many dimensions) as single dimension vectors in memory. Accessing these in this manner is straightforward; a linear index coexists with standard subscript indexing.

a	b	c
d	e	f
g	h	i

becomes

a
d
g
b
e
h
c
f
i

Figure 3.12: Converting a 2D array into a 1D vector
The 3×3 array has a total of 9 elements, and can be converted into a 9 element column vector by stacking the columns.

The right hand side of 3.14 can be written as a vector-matrix multiplication:

$$\mathbf{b_1} = \mathbf{S} \cdot \mathbf{c} \tag{3.16}$$

where $\mathbf{b_1}$ is a M element column vector of spatial samples of the total field from all the transmit coils, $\mathbf{S}$ is a M by R matrix comprising of a concatenation in the second dimension of the column vectors representing the field produced by each coil, and $\mathbf{c}$ is a R element column vector of complex weights for each coil.

By specifying a desired target field $\mathbf{T}$ and substituting for $\mathbf{b_1}$, equation 3.16 can be solved for $\mathbf{c}$:

$$\mathbf{S^{-1}} \cdot \mathbf{T} = \mathbf{c} \tag{3.17}$$

3.2.2 Finding a Robust Solver Method

To achieve successful vessel-selective labelling, B_1^+ is required to have a target magnitude at the arteries to be labelled to fulfil the adiabatic condition [192], whilst as low as possible over any other arteries. Outside of these areas it does not really matter what the transmit field is, providing it remains within SAR limits. These constraints are more relaxed than those of B_1^+ Shimming (where the aim is a uniform field over the entire imaging region) or the arbitrary field localisation described by [169], and so it is plausible to suggest that a suitable field might be achieved using fewer transmit coils.

In this section the development of a robust solver method for producing suitable fields for vessel-selective ASL is presented, with the aim of being able to efficiently label in one region whilst ensuring that no labelling occurs in another, for as few transmit elements as possible. In this context this refers not only to the mathematical method of calculating the complex weights, but how the target field should be defined and any other constraints that need to be imposed, and as will become clear this was not as straightforward to arrive at as equation 3.17 makes it seem.

A typical vessel-selective ASL labelling scenario was devised - successful labelling of the right carotid artery (RCA) and right vertebral artery (RVA), and no labelling of the left carotid artery (LCA) and left vertebral artery (LVA). Anatomical measurements obtained from measurements detailed in section 3.1.1 of a 24-year-old healthy male were used to construct the geometric model also described in that section. Two planar circular loop transmit coils of radius 22.5mm were positioned on the neck such that the vectors normal to the surface and passing through the centres of the coils subtended angles of 30° and 150° with the x-axis. Two coils were chosen because this is the bare minimum for parallel transmission, and if a solver method is able to produce a satisfactory field with only two coils this is a good measure of robustness. A label region and null region was selected by manually drawing elliptical regions of interest (ROIs) on a schematic of the positions

of the arteries within the neck (see figure 3.13). A target field value for within the label region was chosen to be $3.5\mu T$, a typical B_1^+ strength usually employed in human CASL studies. To simplify computation only the B_1^+ field profiles at the labelling plane were used in calculating the complex weights.

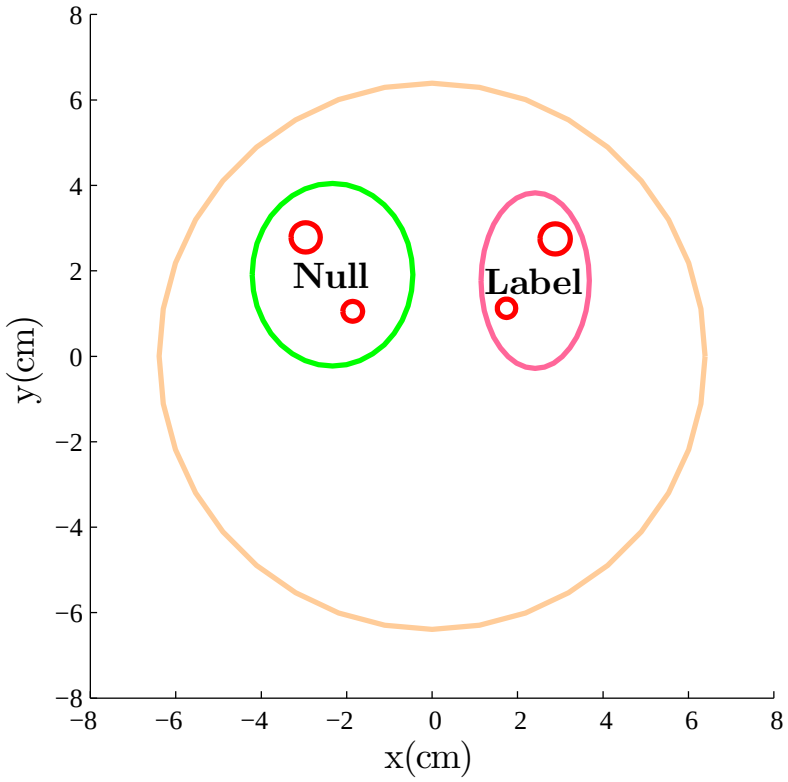

Figure 3.13: Selected label and null Regions
This figure is used during the "planning" of a simulation (echoing the patient planning of a real MRI scan). It calls for two elliptical ROIs to be drawn around a labelling region (shown in pink) and a null region (shown in green). In this instance the labelling region is the RCA and RVA, and the null region is the LCA and LVA. These regions were used throughout this section.

3.2.3 Solution to $\mathbf{S} \cdot \mathbf{c} = \mathbf{T}$

Equation 3.17 calls for a target field to be specified so that the equation can be solved for $\mathbf{c}$. Three different normalised target fields were specified (which were then multiplied by the target field amplitude). Plots of these fields are shown in figure 3.14.

1. Zero within the null region ellipse, unity elsewhere.
2. Unity within the label region ellipse, zero elsewhere.
3. A 2D elliptical Gaussian centred at the label region, with full-width-half-maximum widths those of the label region ellipse, zero within the null region ellipse.

Matrix $\mathbf{S}$ is not square and so does not have an inverse, so for each of the three target fields a solution to equation 3.17 was found using two methods:

1. Iteratively by solving the overdetermined[2] least squares problem:

$$\mathbf{c} = \arg\min \left\{ \|\mathbf{S} \cdot \mathbf{c} - \mathbf{T}\|^2 \right\} \tag{3.18}$$

2. Finding the Moore-Penrose pseudoinverse of $\mathbf{S}$, $\mathbf{S}^+$and solving equation 3.17, with $\mathbf{S}^+$ substituted for $\mathbf{S^{-1}}$.

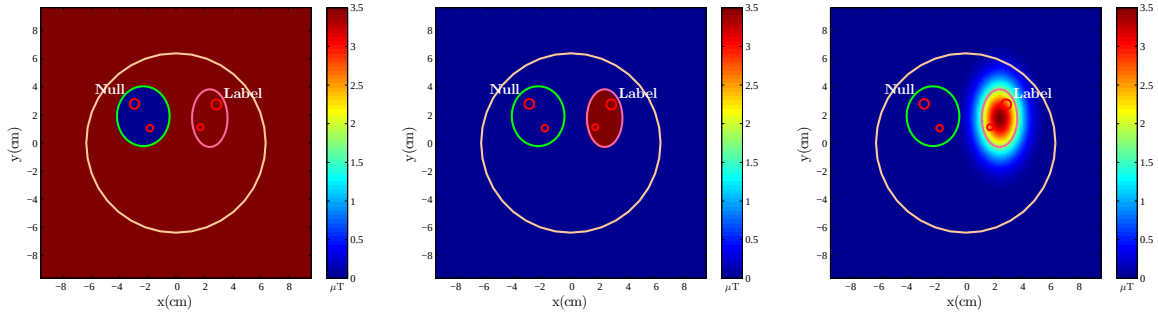

Target Field 1 Target Field 2 Target Field 3

Figure 3.14: Target field plots

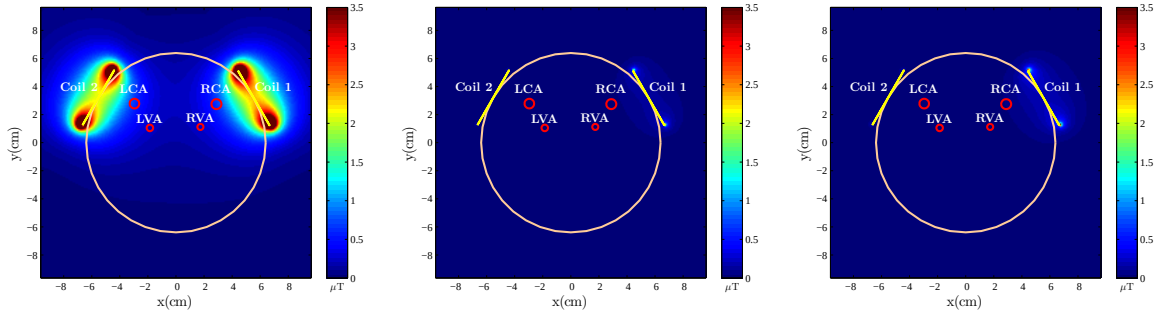

a. b. c.

Figure 3.15: Computed fields from the least squares solution

Figure 3.15.a-c shows the computed fields from the least squares solution with target fields 1, 2 and 3 as inputs respectively. The results obtained using the pseudoinverse matrix were identical to those found with the least squares algorithm. As matrix $\mathbf{S}$ is of full rank ($rank(\mathbf{S}) = 2$), the overdetermined least squares problem in equation 3.18 has a unique solution, and so the solution obtained with the pseudoinverse of $\mathbf{S}$ is identical to the least squares solution. All three target fields do not yield a suitable solution, the reason for this is that specifying a field value at each point results in a lot more equations than unknowns (as there are two coils there are only two

[2] A system of linear equations is considered to be overdetermined if more equations exist than there are unknowns. Each unknown represents a degree of freedom, and each equation a restraint on that degree of freedom. If a system is overdetermined, then multiple solutions exist that satisfy some of the equations, but no solution that satisfies all. There is one case where an overdetermined systems has a unique solution, when the overdetermined set of equations consists of linearly dependent equations.

unknowns). Only in the case where the equations are a linearly dependent set will there exist a unique solution, as there is not the method of least squares gives only an approximate solution. A linearly dependent set of equations arises when the target field is a linear combination of the B_1^+ field from each coil, so in order to obtain a suitable solution it is important that the target field is physically realisable.

Without prior knowledge of an optimal solution which the coil array is capable of producing, a specified target field may not physically be realisable (i.e. not permitted by Maxwell's Equations). For example target field #2, where the field is defined to be $3.5\mu T$ within an elliptical ROI and zero elsewhere. Using the calculated B_1^+ maps it is impossible to obtain this target field because the B_1^+ map for each coil is strongest at the coil and diminishes with distance, so no deconstructive interference can occur would result in zero field at the positions of the coils and a strong field confined to just one region. In high field systems where travelling wave phenomenon occurs and the B_1^+ amplitude oscillates in space it may be possible to construct a field which approaches this, however it would also require more than two transmitter channels to obtain a well-defined region.

3.2.4 Reducing the number of constraints

For a system of linear equations with two unknowns the optimal number of equations for finding a unique solution is two. The matrix $\mathbf{S}$ could consist only of the values of the B_1^+ field at the centre of the label and null regions respectively, and a target field consisting of only two values: $3.5\mu T$ at the centre of the label region, and $0\mu T$ at the centre of the null region. A unique solution will be physically realisable, however it might not be within hardware or safety limits, or be the most optimal solution for vessel selective labelling. Ideally, the B_1^+ field within the null region should be as low as is possible, the B_1^+ field within the label region at a target magnitude, and the magnitude of each complex weight cannot exceed hardware or safety constraints. These requirements can be written as a constrained minimisation problem:

$$minimise\ \left\|\mathbf{S_i}\cdot\mathbf{c}\right\|^2,\ i\epsilon Null\ Region \tag{3.19}$$

subject to

$$\mathbf{S_i}\cdot\mathbf{c} = B_{1,target},\ i = Label\ Centre \tag{3.20}$$

$$|\mathbf{c}| \leq 1 \tag{3.21}$$

Qualitatively the problem becomes "Minimise the B_1^+ field within the null region, whilst ensuring that the B_1^+ at the centre of the label region is equal to a target value, and that the absolute values of the complex weights cannot exceed unity".

These constraints were formulated in Matlab using CVX [85, 84], a package for specifying and solving convex programs. CVX simplifies the formulation and solution of least squares, linear programming and quadratic programming problems, by using an approach to convex optimisation called disciplined convex programming[86]. Syntax is similar to writing the problem mathematically (such as in equation 3.19), allowing the user to focus on the application without prerequisite expertise in numerical optimisation.

Using the label and null regions specified in figure 3.13 a solution was found for two coils:

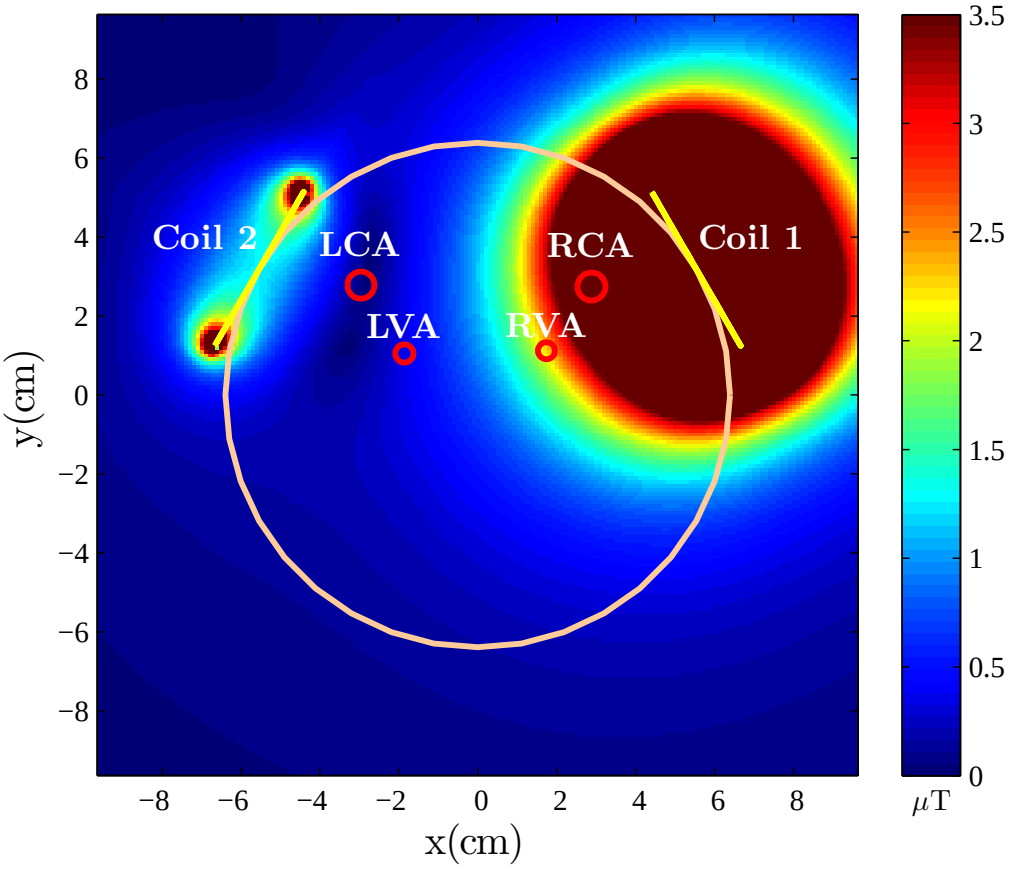

Solved

Figure 3.16: Computed B_1^+ field

	$B_1^+(\mu T)$ at centre of artery				Coil 1		Coil 2	
Iterations	RCA	RVA	LCA	LVA	a	$\phi(^\circ)$	a	$\phi(^\circ)$
9	6.00	2.03	0.11	0.33	0.92	-5.30	0.056	-169.8

Table 3.7: Summary of the B_1 magnitude at the centre of each artery and coil amplitude and phase for the solution

As can be seen in figure 3.16 and table 3.7 the solution produces a total field that is in good agreement with the labelling criteria. The RCA receives a B_1^+ field well in excess of the $3.5\mu T$ target, the RVA less so, however, generally the flow velocity is lower in the vertebral arteries so a smaller B_1^+ is required to fulfil the adiabatic condition (see equation 2.72 in section 2.3.2.1). A well defined null was created directly at the location of the LCA, resulting in a B_1^+ over an order

of magnitude lower than the target labelling amplitude, and while its shape does not optimally cover the LVA, the field is partially nulled and its intensity reduced. Examination of the B_1^+ field profile in the $-y$ direction from the centre of coil 1 gives an indication of what the unperturbed field without the nulling effect of coil 2, stronger field strengths persists for a given radial distance from the centre, suggesting that contralateral labelling would be greater if only coil 1 was used.

As one might intuitively suspect, the coils are driven almost out of phase with each other, with most of the power coming from the coil closest to the labelling region.

3.3 Computer Simulation of Vessel Selective Continuous Arterial Spin Labelling using Parallel Transmission

Vessel Selective ASL provides visualisation and quantification of the perfusion territory from a feeding set of arteries. Methods to selectively label blood include using pseudo-Continuous ASL with additional transverse gradients and phase cycling [242], CASL with obliquely defined labelling planes by rotating the field of view [232] and the use of surface coils to spatially confine the B_1 field to one side of the neck [156]. The latter is an attractive option as the confined field does not produce any magnetisation transfer (MT) [101] effects within the brain, however the spatial profile of the surface coil extends across the neck, partially labelling the contralateral arteries.

In this next section the concept of spatially localising the B_1 field from a set of transmit coils is applied to an *in-silico* Vessel Selective Continuous Arterial Spin Labelling (CASL) experiment. As shown in section 3.2 modulation of the amplitudes and phases of multiple coils can be used to create a region where the field is minimised. This will be used to produce a minimum in the B_1 field at the Left Carotid (LCA) and Vertebral (LVA) arteries, whilst retaining the B_1 field at a target strength or higher for labelling the Right Carotid (RCA) and Vertebral (RVA) arteries. These fields will then be used in the ASL simulation to evaluate their labelling efficacy. Specifically, the specificity of labelling the RCA and RVA without contralateral labelling of the LCA and LVA will be compared for two and four coils with respect to one coil.

3.3.1 Method

The computer simulation described in section 3.1 was set up with the following parameters:

- Anatomical data from the 24 year old male volunteer was used to construct the simulation model (see tables 3.1, 3.3 and 3.7).
- Each virtual coil's field was computed over a 3D raster covering $[x, y, z]$ coordinates of $[-95.9, -95.9, -307.8]$mm to $[95.9, 95.9, 0.0]$mm, with 1mm isotropic spatial resolution. The current in each coil was calculated with the Biot-Savart law for the on axis magnetic field of a current loop, $B_a = \frac{\mu_o}{4\pi} \frac{2\pi R_2 I}{(a^2+R^2)^{3/2}}$ by specifying $B_z = 3.5\mu T$ at a target depth of $a = 5cm$.

- Uniform labelling gradient of $G_z = 3mT/m$.
- Label and null regions as given by figure 3.13.
- Labelling plane positioned half way along the neck at $z = -190mm$.
- Complex weights for two and four coils were calculated by minimising equation 3.19 according to the constraints in equations 3.20 and 3.21.
- Arterial blood relaxation times corresponding to $B_0 = 3T$: $T_1 = 2000ms$ and $T_2 = 300ms$ [211].
- Simulation time step of $10\mu s$.
- Spins initially positioned every 1mm: 307 spins per artery.
- Labelling gradients and RF for a duration of 2 seconds.

Coils one to four were placed around the neck at angles 30°, 150°, -30° and -150° respectively to the x-axis. Coils one and two are approximately over the left and right carotid arteries. Simulations were performed with only coil one (as per a single coil experiment), coils one and two, and all four coils.

Results were collected in two ways. Firstly for each artery the longitudinal magnetisation, M_z of the spin starting from $z = z_{start}$ at $t = 0$, and z position was recorded at each time step as it travelled through its artery. These spins were positioned in the centre of each artery so their velocity was only dictated by the velocity waveform $v_{artery}(t)$, and are hereafter referred to as the "tracked spins". Secondly the value of M_z and z of each spin in each artery was recorded at the end of the simulation ($t = 2s$).

Using these results the inversion efficiency of the tracked spins was calculated as:

$$\eta_{inv} = \frac{1 - M_{z,min}}{2} \times 100\% \tag{3.22}$$

Where $M_{z,min}$ is the minimum value of M_z for each spin. The percent change in inversion efficiency, $\Delta\eta_{inv}$ for each parallel transmit configuration ($\eta_{inv,pt}$) with respect to the single coil configuration was calculated as:

$$\Delta\eta_{inv,pt} = \frac{\eta_{inv,pt} - \eta_{inv,sc}}{\eta_{inv,sc}} \times 100\%$$

The mean longitudinal magnetisation of all the spins within the imaging volume (between $z = -26.45mm$ and $z = 0$) at the end of the simulation, $\langle M_z \rangle$, was computed, and the respective mean inversion efficiency and mean percent change were also calculated:

$$\langle \eta_{inv} \rangle = \frac{1 - \langle M_z \rangle}{2} \times 100\%$$

and

$$\Delta \langle \eta_{inv,pt} \rangle = \frac{\langle \eta_{inv,pt} \rangle - \langle \eta_{inv,sc} \rangle}{\langle \eta_{inv,sc} \rangle} \times 100\%$$

3.3.2 Results

	a	$\phi(^\circ)$	Artery	$B_1^+(\mu t)$ at centre	$M_{z,min}$	η_{inv}	$\langle M_z \rangle$	$\langle \eta_{inv} \rangle$	$\Delta\eta_{inv}$ (%)	$\Delta \langle \eta_{inv} \rangle$ (%)
One Coil										
Coil 1	1.000	0.000	**RCA**	6.44	-0.93	96.50%	0.088	45.60		
			RVA	2.25	-0.97	98.35%	0.198	40.11		
			LCA	0.39	0.87	5.70%	0.952	2.41		
			LVA	0.52	0.73	13.35%	0.889	5.56		
Two Coils										
Coil 1	0.915	-5.300	**RCA**	6.00	-0.93	96.70%	0.113	44.37	0.21	-2.71
Coil 2	0.056	-169.790	**RVA**	2.03	-0.97	98.60%	0.256	37.19	0.25	-7.28
			LCA	0.11	0.99	0.40%	0.995	0.26	-92.98	-89.42
			LVA	0.33	0.89	5.55%	0.954	2.33	-58.43	-58.18
Four Coils										
Coil 1	0.978	11.700	**RCA**	6.20	-0.93	96.55%	0.169	41.54	0.05	-8.91
Coil 2	0.010	-140.997	**RVA**	1.86	-0.93	96.60%	0.243	37.86	-1.78	-5.62
Coil 3	0.998	-129.105	**LCA**	0.034	0.999	0.05%	0.998	0.08	-99.12	-96.68
Coil 4	0.105	130.125	**LVA**	0.044	0.997	0.15%	0.998	0.08	-98.88	-98.47

Table 3.9: Quantitative summary of the B_1 field in the labelling plane and longitudinal magnetisation measurements

Figure 3.17 contains field plots of the magnitude of the total B_1^+. In comparison to using one single surface coil both two and four coils clearly reduce the magnitude of the field in the vicinity of the LCA and LVA, whilst retaining the field strength and distribution around the RCA and RVA. Two coils introduce a null which is positioned mostly around the LCA, whilst four coils improve the depth and shape so that both the LCA and LVA are in a region where the field is weak.

The effect of the vector summation of the B_1^+ fields along the z-axis is shown in figure 3.18. For one and two coils the field experienced by the spins within the RCA and RVA are similar, but for four coils the shape of each B_1^+ profile is narrower, particularly the field experienced in the RVA. Minima in the B_1^+ field can be seen in the LCA in figure 3.18.b and both the LCA and LVA in figure 3.18.c. In 3.18.b these minima do not occur in the labelling plane, but approximately one centimetre either side.

Figure 3.19 shows the labelling effect of these fields on spins moving at the maximum velocity through each artery and the columns $M_{z,min}$, and η_{inv} in table 3.9 are the corresponding values of M_z and the inversion efficiency for this figure. In comparison to the single coil case contralateral labelling in both the LCA and LVA is reduced by at least 58% for two coils and two orders of magnitude for four coils (to the extent where it is almost non-existent). Labelling of the RCA and RVA is slightly improved (by 0.2%) for two coils, but reduced by nearly 2% in the RVA when using four coils.

The effect of the labelling fields on a bolus of blood (under pulsatile flow and with a range

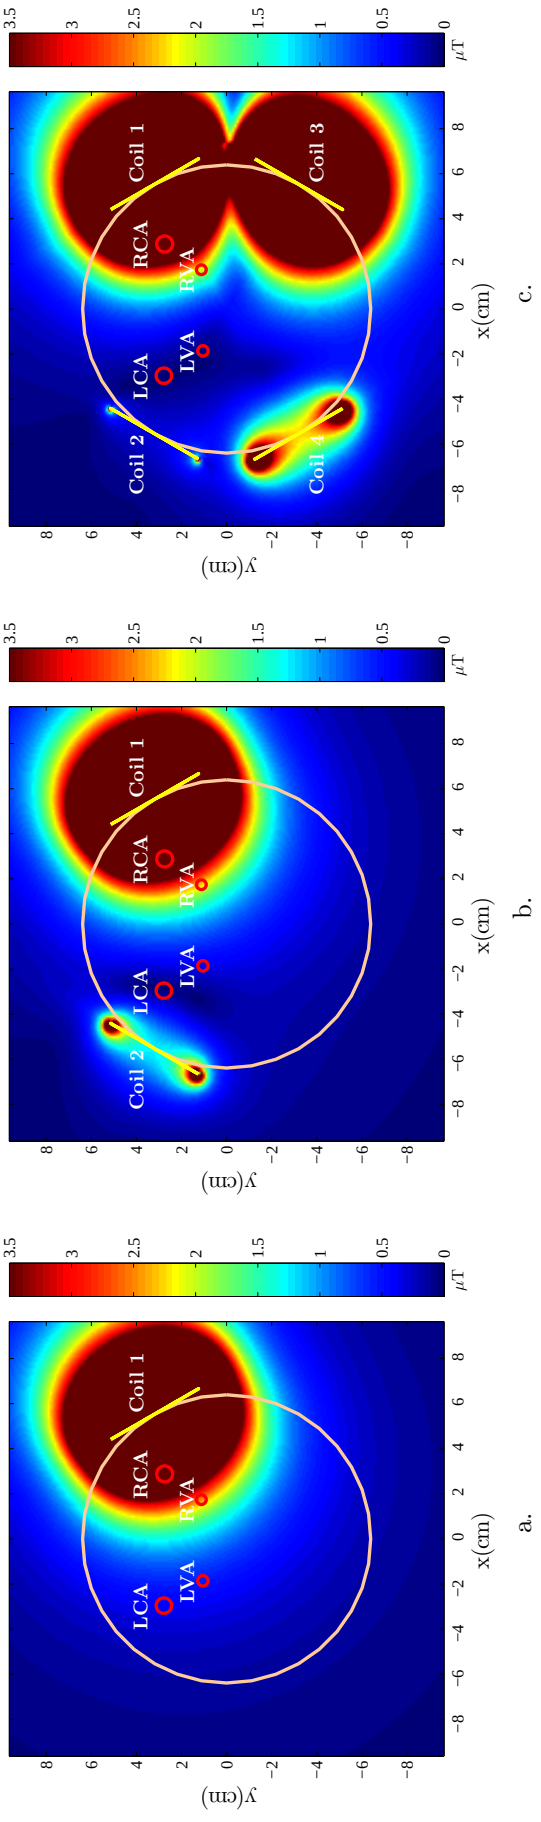

Figure 3.17: Field plots in the labelling plane from a. 1, b. 2 and c. 4 coils

Colours are scaled between 0 (dark blue) and 3.5μT (dark red). Interestingly in the case of four coils there is hardly any contribution from coil two (especially compared to b.), and coil four seems to be responsible for producing the null effect. This suggests there are optimal coil positions for being able to prescribe the most amount of specificity in the B_1^+ field.

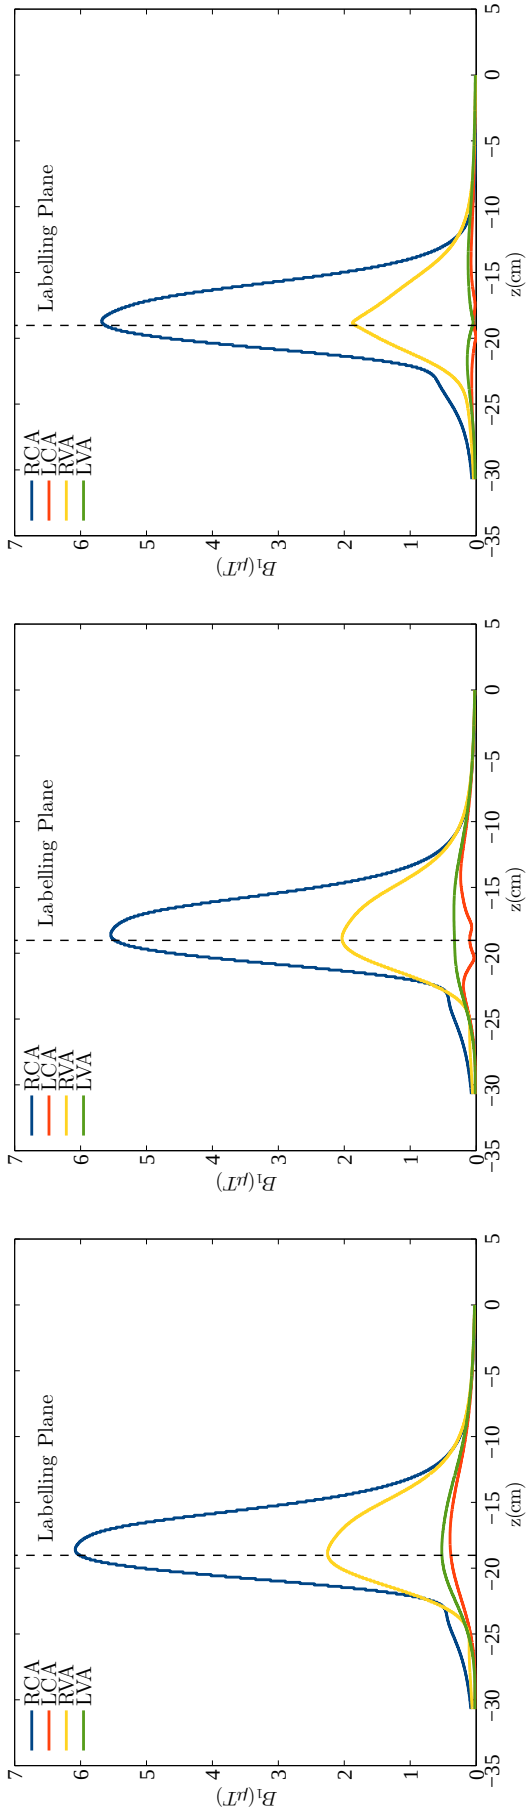

Figure 3.18: B_1^+ profile parallel along z through the centre of each artery for a. one coil, b. two coils and c. four coils

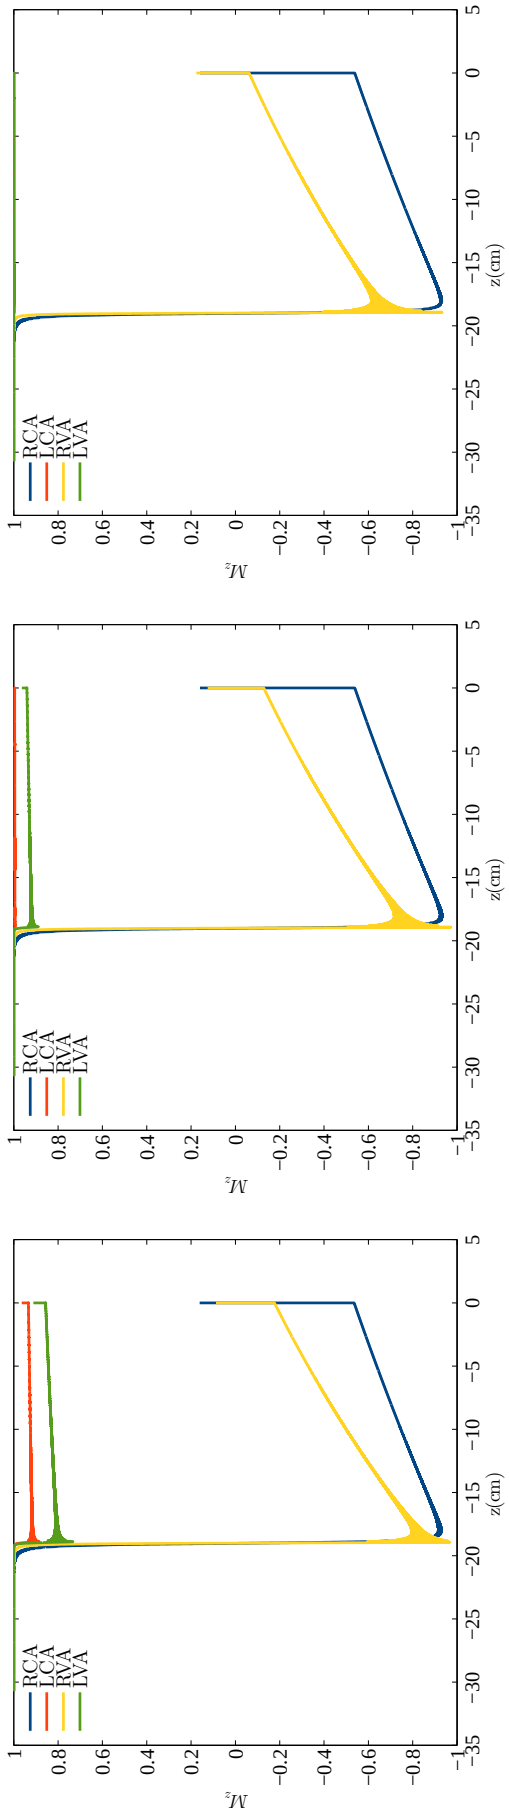

Figure 3.19: $M_z$of a single spin at each z position for a. one coil, b. two coils and c. four coils

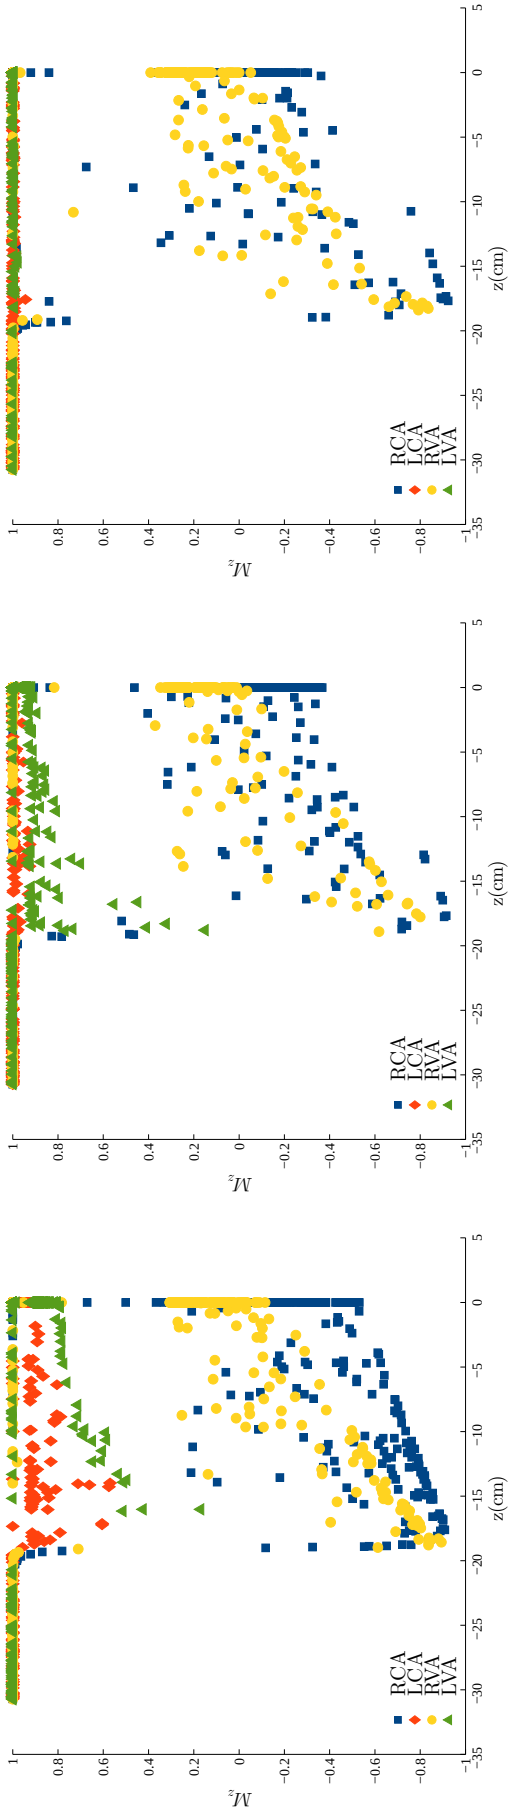

Figure 3.20: Final longitudinal magnetisation of a bolus of blood after labelling for a. one coil, b. two coils and c. four coils

of flow velocities due to the laminar flow) is shown in figure 3.20, and a mean value of the final longitudinal magnetisation of all the spin that have a z displacement greater than $z = -26.45mm$ is shown in the $\langle M_z \rangle$ column of table 3.9. Once again a reduction in contralateral labelling can be seen for the two coil case, and with four coils nearly all the spins are unperturbed in the LCA and LVA. However fewer spins are labelled with high efficiency for two and four coils compared to the single coil case, as can be seen on figure 3.20b and c and from the values of $\langle M_z \rangle$ (lower values are seen for the single coil case, indicating a higher degree of labelling). A reduction of as much as 9% was seen in comparison with the single coil case. Implications of this are that the labelled bolus of blood will provide less perfusion weighted contrast than expected and this needs to be accounted for when quantifying cerebral blood flow.

3.3.3 Discussion and Conclusion

Results show that selective labelling of the RCA and RVA with minimal labelling of the LCA and LVA is possible using parallel transmit methods to spatially localise the B_1field. As the number of coils increases there is an improvement in the ability to tailor the shape and depth of the null region. Using the calculated fields in a simulation of the labelling procedure shows that the labelling efficiency on the RCA and RVA can be maintained to within 10% when compared to the single coil case, whilst contralateral labelling is reduced by up to two orders of magnitude.

In the case of two and four coils the reduction in magnitude of the B_1 field at the RCA and RVA may be attributed to the reduced amplitude, a of coil 1 compared to the single coil case (where $a = 1$). Normalising the amplitudes such that $a_1 = 1$ in all cases would increase the magnitude of B_1at the RCA and RVA and would perhaps have been a better measure of comparison. However one of the benefits of using parallel transmission is that the scanner can be calibrated to the patient and the application, in this case less power was required in coil one to achieve the prescribed B_1 field at the centre of the label ROI. An extension would be to use a better constraint for the labelling region, such as defining the average B_1 intensity within the label region to be above a target magnitude, rather than a single point.

Spatially tailoring the B_1 field through modulation of the amplitudes and phases of independent transmit elements is purely due to the superposition of the magnetic fields produced by each array element, and so can be used for spatial localisation in any pulse sequence. In comparison Transmit SENSE requires a multi-dimensional k-space trajectory, and only works for short RF pulses because of T_2^* decay.

The calculation of the complex weights does not take into account power deposition and ensure Specific Absorption Rate (SAR) limits are not exceeded. Transmit SENSE methods limit power through Tikhonov regularisation of the temporal B_1 waveform [87, 125], however for continuous wave power this cannot be done (as the temporal B_1waveform has only one value, minimising this

would result in very small complex weights). Providing the coil array used for parallel transmit is designed and characterised such that safe operating conditions are known and possible, constraints can be placed on the complex weights such that they do not exceed the safe operating conditions of the transmitter system in a similar manner to the constraint in equation 3.21.

3.4 Summary and Conclusions

In this chapter a computer simulation of an ASL experiment has been developed and used as a test bed for applying parallel transmission methods to vessel selective ASL. The simulation consists of an *in-silico* arterial phantom that models the flow of blood through the carotid and vertebral arteries, with anatomic and physiological parameters based on MRI measurements of a volunteer. A rotation matrix Bloch equation kernel then operates on the bulk magnetisation of the blood as it is subjected to a labelling gradient and RF fields.

The computational model was designed to be flexible in its operation, allowing a variety of labelling scenarios to be implemented. For example it can be used to look at the evolution of just one spin isochromat, or a continuous flow of blood in both carotid and vertebral arteries. Only a small number of changes are required to adapt the simulation to the anatomy of another species, for example a mouse or rat. Likewise, pulsed RF and gradient waveforms could be incorporated to simulate pseudo-Continuous ASL. By incorporating a 3D B_0 field map, the effect of B_0 field inhomogeneities on labelling efficiency can also be investigated and techniques developed for ensuring robustness to them.

A mathematical method for spatially tailoring the transmitted B_1 field such that label and null regions are defined was developed and tested using simulated surface coils. This was then applied to Vessel Selective ASL as a method for selectively labelling the right carotid and vertebral arteries, whilst leaving the left carotid and vertebral arteries unperturbed. The performance of two and four coils were compared with one coil positioned on the right-hand-side of the neck. Spatial selectivity was shown to be achievable using just two coils, and four coils reduces the magnitude of the field over the left-hand-side arteries by two orders of magnitude in comparison to the single coil case. Calculated fields were utilised in the ASL simulation to label a bolus of blood under typical experimental and physiological parameters. Contralateral labelling was markedly reduced for two coils and almost negligible with four. A reduction of up to approximately 10% was seen in the average inversion efficiency in the labelling arteries, however this reduction was less severe (2%) when only considering a spin travelling at the maximum velocity within each artery. In general however the labelling efficiency within the right-hand side vessels was very high, so even with this reduction a well inverted bolus of blood was created.

To conclude, results from this simulation show that it is possible to specify the labelling of arteries using parallel transmit methods, and that an improvement over the single coil case can be

seen using as few as two coils.

Chapter 4

ASL Labelling Coil Array

Multi-element coil arrays have been employed for signal reception for over 20 years [194], providing improved SNR to FOV ratio in comparison to volume coils, and facilitating parallel imaging [186]. In recent years the RF inhomogeneity challenges faced in moving to ultra high field (7T and above) human systems has prompted the development of coil arrays for transmission. By transmitting independently on each coil the dielectric field focusing effect experienced as the wavelength of the B_1 field becomes comparable to (or smaller than) the anatomy being excited can be compensated for. Two distinct methodologies exist, B_1 shimming [31, 116, 169], where the current on each element is modulated in amplitude and phase (but all share the same pulse envelope) to spatially tailor the B_1 field, and Transmit SENSE [124, 257, 87], where each element transmits a unique pulse waveform simultaneously with gradients tracing out a trajectory through k-space (such as an EPI or spiral trajectory), resulting in a multidimensional excitation pattern. Transmitting from multiple coil elements introduces new engineering challenges in decoupling that cannot be addressed in the same way as with receive coils[1]. Neighbouring coils can be decoupled by overlapping them slightly [194] or with capacitive networks [183] but decoupling other elements is more challenging and requires more exotic approaches. Kurpad et al. [135] achieved >15dB of isolation by employing a power MOSFET as a voltage controlled current source for each element. Hoult et al. [112] custom built a spectrometer which utilised cartesian feedback (a stable form of negative feedback developed for RF, originally in the communications industry), which was then shown to provide decoupling in a transceive array [113].

In chapter 3 it was shown that in simulations of Continuous Arterial Spin Labelling (CASL) parallel transmission methods for spatially localising the B_1 improves specificity of arteries when performing vessel selective labelling. This chapter details the development of a suitable set of RF coils for implementing parallel transmit ASL on humans in conjunction with a Siemens Tim Trio 3T MRI scanner and Surrey Medical Imaging Systems (SMIS) spectrometer.

[1] Low input impedance pre-amplifiers are employed to effectively turn each coil element into an open circuit at the Larmour frequency, consequently no current can flow in each coil and no coupling can exist between elements.

4.1 Design

4.1.1 Specification

Characteristic	Target Requirement
Operating Frequency	Proton frequency for a 3T Siemens Tim Trio - nominally 123.2MHz.
Transmit Performance	Produce a B_1^+ field capable of selectively labelling without repositioning of the coils: 1. All the arteries within the neck 2. Selectively label either the right or left sets of carotid and vertebral arteries.
Receive Performance	No standard imaging will be performed with the coil array.
RF Power	Typical use will be continuous wave sufficient to produce a B_1^+ field of up to $10\mu T$ at a depth of 5cm. Previous studies have used approximately 0.3-2W [103, 213, 259]. All components must be capable of withstanding this power without any undue stress, and a suitable RF power amplifier is required.
Interfacing	Coil array will be interfaced to an independent spectrometer (SMIS MR5000), which has a maximum of two independent transmitter channels. Can be controlled via a TTL signal. Imaging will be performed by a 3T Siemens Tim Trio, which has an optical out for external triggering.
Tune/Match Adjustability	No tuning/matching facilities are available within the magnet room, and so the coils should be tuned and matched at a predetermined point to ensure maximum transmit efficiency based on average load conditions for a human neck.

Characteristic	Target Requirement
Decoupling	Each coil shall be decoupled from: 1. The other labelling coil 2. Trio body coil 3. Trio receive coil Decoupling must be greater than 15dB.
Matching	Matched to 50Ω for connecting with coaxial cable to standard RF power amplifiers.
Safety	Coils must remain decoupled from the Trio body coil even if there is a malfunction with the coil or PIN diode driver circuitry. Any points which could potentially have a high voltage across them must be well insulated so that the patient is not exposed to them. Cable shield currents, caused by the body transmit coil inducing currents on the outer shield of the interfacing coaxial cables, potentially burning the patient must be mitigated on all cables within the magnet bore.
Mechanics	The coils need to be fixed in position around the neck in a comfortable manner. The coils will be fixed to the Trio patient couch.
Materials	All components and materials must be non-magnetic and not produce a detectable ^1H NMR signal.

4.1.2 RF Design

In this section each aspect of the design is detailed and justified. When building a coil component values are determined through a combination of calculations from empirical measurements and trial and error. While the theory behind the operation of RF coils is well understood, there are a number of hidden parameters in every component and material that make it difficult or complex to accurately model; consequently it is easier to determine component values empirically, and this merges the design and manufacturing process.

4.1.2.1 Number of Coil Elements

As there are only two transmit channels available on the SMIS spectrometer, two coil elements were built. This is the bare minimum to perform parallel transmit spatial localisation. Additional transmit channels could be added by splitting the spectrometer output into separate signals, each of which passes through a vector modulator for amplitude and phase modulation. However complexity

increases. Two coils permits both parallel transmit ASL and vessel selective labelling using only one coil at a time. Direct comparisons between the two methods of spatial localisation can be drawn as the same sets of coils are used in all experiments.

4.1.2.2 Array Geometry

The positions of the coils with respect to one another will affect the spatial extent of the B_1^+ field, coil coupling, and the ability of the array to be able to spatially tailor the B_1^+ field when used as a transmit phased array. As the array consisted of only two coils, in order to comply with the transmit performance specification of labelling all arteries and selectively labelling either the right or left sets of arteries, without adjusting the positions of the coils the most sensible solution was to position each coil over each set of arteries. All arteries can then be labelled by transmitting with the same amplitude and phase on both coils, and selective labelling achieved either by parallel transmit phased array operation, or only using one coil at a time.

4.1.2.3 Coil Geometry

The coil shape and size will directly influence the spatial extent of its B_1 field. Circular surface coils are ideally suited as they produce a spatially confined B_1 field which penetrates by approximately the same amount as the diameter of the surface coil. Parameters to consider are:

- Diameter - Dictated by the desired B_1^+ penetration depth and the overall size of the coil. Neck sizes are variable, and the exact positioning of the coils on the neck is very subject specific. Coupling between the two coils will be reduced for smaller coil diameters, and so it is desirable to have the smallest diameter possible for a specified B_1^+ penetration depth.
- Track width - a wide track is preferable as it provides a larger area for current to flow (minimising losses, increasing the Q of the coil) and permits more discrete capacitors to be added in parallel with one another. However it also increases the total diameter of the coil, contradictory to the previous point.

The desired B_1^+ penetration depth is 5cm, and as the RF current distribution always takes the shorted path, the inside diameter should be approximately the same. For reasons given in 4.1.2.4 it is desirable to be able to fit three ATC B-Series or similar (approximate dimensions of 2.79 x 2.79 x 2.54mm) capacitors at each tuning capacitance gap, and so, allowing for dimension tolerances, a track width of at least 9mm is required.

With the above considerations taken into account an inner diameter of 45mm and outer diameter of 64mm was chosen.

4.1.2.4 Tuning and Matching

Each coil must be tuned to resonate at 123.2MHz with discrete capacitors, distributed around the loop to reduce electric field coupling with the sample (see section 2.2.2.1 and [8]). The number of distributed capacitors was determined with a pilot build (see below).

Matching to 50Ω can be achieved either inductively or capacitively. Inductive matching consists of a series resonant loop positioned to overlap the coil. It is intrinsically balanced, however it relies on precise positioning of the coil and feed loop to ensure the correct impedance transformation. Capacitive matching is easier to implement, as a direct connection is made across a capacitor in the coil, the value of which is chosen to transform the impedance of the coil to 50Ω (see equation 2.59).

Pilot build to determine the capacitor topology To determine a suitable number of distributed capacitors a 5cm inner diameter octagonal coil was constructed from 12mm wide copper foil on a piece of 0.4mm thick FR-4 (glass-fibre epoxy board). Fixed value low loss ceramic chip capacitors were used to resonate the loop to approximately 123MHz. The unloaded and loaded (with the coil directly against the neck of a subject) Q was measured. Then, using the nominal value of capacitance and the loaded Q the inductance, loaded resistance and match capacitance required to transform the coil resistance to 50Ω at 123.2MHz were calculated using the equations described in sections 2.2.1.1 to 2.2.2.2.

	Unloaded		Loaded					
Nominal C_{tune}	f_0	Q_{UL}	f_0	Q_L	L	R_{Loaded}	X_{match}	C_{match}
$15 + 3.9 = 18.9$pF	123.0MHz	500	122.0MHz	30	88.6nH	1.96Ω	9.89Ω	131pF

Table 4.3: Coil Measurements and Calculations

Results are given in table 4.3. Measurements of the frequency and Q were made with a double loop probe (see section 4.2.1) and HP 8712C Vector Network Analyser set to measure S_{21}. Unloaded measurements were made with the coil in the centre of a plastic table 1.5m away from any metal objects. Successfully distributed capacitance will overcome coil-sample capacitive coupling, and balance the electric fields on the coil. As mentioned in section 2.2.2.1 this is achieved by ensuring that a) at all positions there is at least 100pF of capacitance, and b) all values of capacitance are equal (or at least symmetrical about the feed point - the match capacitor). The calculated value of C_{match} sets a lower limit on the amount of capacitance at each gap, and to achieve this there needs to be at least seven gaps.

This calculation does not take into account the parasitic inductance of the capacitors, which causes each capacitor to have a self-resonant frequency and consequently the reactance of the capacitor is frequency dependent. The nominal value of C_{match} to obtain X_{match} at 123.2MHz will be lower than the value calculated here.

Following these calculations the correct match capacitance was added and the tuning capacitance distributed to the remaining seven gaps.

	f_0	Q
Unloaded	124.0MHz	270
Loaded	123.5MHz	27

Table 4.5: Final characterisation of the Pilot Build

Distributing the capacitance increased the losses in the coil, nearly halving the Q from 500 to 270. However the loaded Q was only affected by 10%, indicating that sample losses dominate over coil losses, and the measured drop in unloaded Q does not degrade coil performance significantly. When loaded the resonant frequency shifted by 0.5MHz, half of what was previously experienced. By choosing component values so that the coils tune and match at 123.2MHz for an average neck then the reduced dielectric losses from distributing the capacitors and subsequent frequency shifts should be small for most neck shapes/sizes/tissue combinations (as some necks contain more muscle or fat than others).

To conclude the pilot coil build confirmed that seven distributed capacitances and an eighth capacitor to transform the coil resistance to 50Ω was sufficient to reduce dielectric losses so that tuning is stable over a range of loading conditions. As such this configuration was adopted in each coil element of the ASL Labelling Coil Array.

4.1.2.5 Detuning

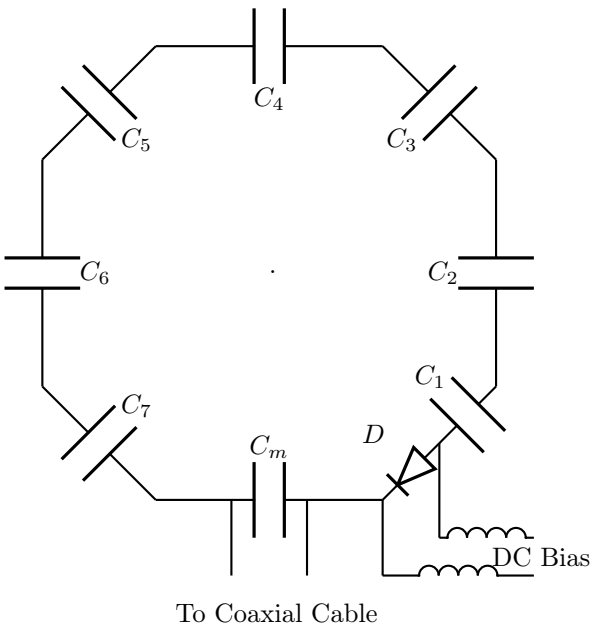

Figure 4.1: PIN Diode Detuning Circuit
PIN diodes are biased with direct current (DC). RF Chokes (large value inductors, typically $1\mu H$) are used to block RF from leaking down the DC line.

When not transmitting each coil needs to be detuned. Failing to do so will cause:

- Coupling to the body coil, focusing its field onto the neck, exceeding SAR limits and potentially causing burns.
- Coupling to the receive coils, which could introduce inhomogeneities into the receive B_1^- field and increase image noise.
- Coupling to the other coil in the ASL array. This would confound any attempt at labelling only one side of the neck with a single coil as the contralateral coil would focus the field onto the arteries close to it.

Placing a PIN diode in series with the coil loop as shown in figure 4.1 is not normally recommended practise for receive coils, as the resistance of the forward biased diode introduces noise. However for a small transmit coil where coil-sample losses dominate this will have very little effect. In this configuration the coil must be "actively" tuned; when unbiased the PIN diode will not pass an RF current and the coil will remain in a detuned state. In the event of a PIN diode driver malfunction, the coil would not function. For other PIN diode circuit topologies (see section 2.2.2.3) where the coil is detuned when the diode is forward biased, the coil would remain tuned and will couple to other coils.

To conclude it was decided that each array element will be detuned using a PIN diode in series with the loop.

4.1.2.6 Balancing

A coil is a balanced device, with equal but opposite currents flowing on opposite sides of the loop. It is however fed by an unbalanced coaxial cable (the shield of the cable is grounded). If a direct connection is made with the coaxial cable across the match capacitor, performance will be sub optimal. A balun (which stands for BAlanced to UNbalanced or BALancing UNit) is a device used to convert signals that are balanced about ground to those which are unbalanced and vice-versa. It presents a high impedance to common mode currents, but leaves differential mode currents unaffected.

A lattice balun (see figure 4.2) was used to transform the unbalanced signal from the coaxial cable to the balanced signal for driving the coil. It also reduced common mode currents induced in the coil by the body coil. The advantages of using a lattice balun are that it can be built directly on to the coil's circuit board, takes up a minimal

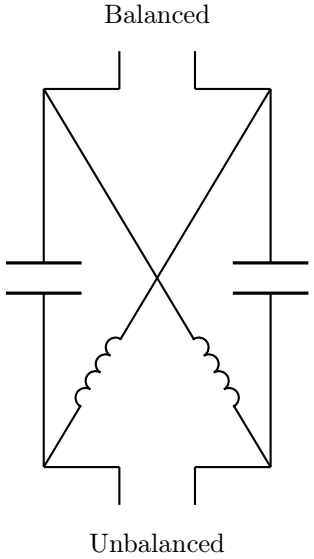

Figure 4.2: Lattice Balun

A lattice balun consists of two parallel resonant circuits. C and L are chosen such that $X_c = X_L = 50\Omega$ at the frequency of operation. For 123.2MHz this corresponds to $L = 64.46nH$ and $C = 25.8pF$.

amount of space and has high common mode rejection ratio (CMRR).

4.1.2.7 Mitigating Cable Currents

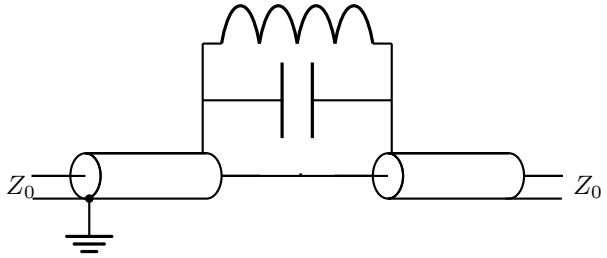

Figure 4.3: Concept of a Cable Trap
A parallel resonant circuit (resonant at the frequency of interest) is introduced to the shield of a coaxial cable, without breaking the shield so that the characteristic impedance between the inner conductor and shield remains constant - typically 50Ω.

Currents flowing on the outer shields of the coaxial cables are problematic as they can reduce the efficiency of a coil, couple noise into the magnet bore, and potentially cause patient burns. The ASL Labelling Coils were interfaced to a separate spectrometer console, located outside the magnet room. Therefore cabling ran out of the magnet bore to bulkhead BNC connectors on the filter plate, and then on the other side of that plate more cables are needed to run to the amplifiers, totalling at least 10m of cable. The vacuum wavelength of electromagnetic radiation at 123.2MHz is 2.43m, making the outer screen of this length of cable electrically large and hence good antenna. RF currents can be induced on it due to a. the body coil during transmission, inducing the most powerful currents that can cause patient burns, b. improper balancing of the coil, resulting in common mode currents, reducing transmit efficiency, and c. noise can be injected into the scanner bore via the low impedance outer shield.

Fortunately there are ways to reduce these cable currents. Outside of a magnetic field it is standard practise to clamp a ferrite bead around a cable, imposing a large impedance to high frequencies on the cable shield, the energy of which is absorbed by the ferrite. These beads are common on computer, video and switched mode power supply cables (for example a USB cable). As ferrite is magnetic it cannot be used inside the MR environment, so a non-magnetic equivalent must be used: a cable trap[2].

All cable traps use the principle of creating a parallel resonant circuit in series with the cable shield (this is sometimes called a tank, or trap circuit, hence the name cable trap). Parallel resonant circuits have a high impedance at their resonant frequency, so by tuning these circuits to the ^1H frequency the cable currents can be reduced. Figure 4.3 is a schematic of the concept of operation of a cable trap.

[2]A cable trap is in fact a Balun. The only difference being the application - a cable trap positioned at the feed point of a coil will balance the coil, whereas a cable trap positioned further along the cable will have less of a balancing effect but will reduce cable shield currents induced onto the cable.

Cable traps need to be positioned in specific places along a cable to have optimal effect. Common mode currents running along the cable shield will reflect from any impedance mismatch (such as at the coil, or the end of the cable), so there will be a superposition of RF currents running in both directions, forming a standing wave. Placing a trap at a current node will only have a minimal effect on the suppression of shield currents.

Multiple traps are also necessary to ensure that shield currents do not propagate anywhere along the cable as even a short length can radiate (and through the principle of reciprocity receive). A simple yet effective antennae is the half-wave dipole, which consists of two quarter wavelength sections. Transmission or reception of electromagnetic radiation is achieved from standing waves on each of the sections of the dipole. To ensure that a cable does not become radiative its length must be kept below a quarter-wavelength. Indeed, Siemens guidelines state that any third party coil that is to be used with their 3T systems must have an attenuation of 15dB or greater must be placed every 25cm along all cables within the magnet bore (this corresponds to less than $\lambda/8$ in free space).

As discussed in section 4.1.3 two types of coaxial cable are used - semi-rigid coaxial cable within the coil array housing, and RG223 cable for connecting between the coil array and the filter plate in the magnet room. Different types of cable traps are required to meet the needs of these two cable types:

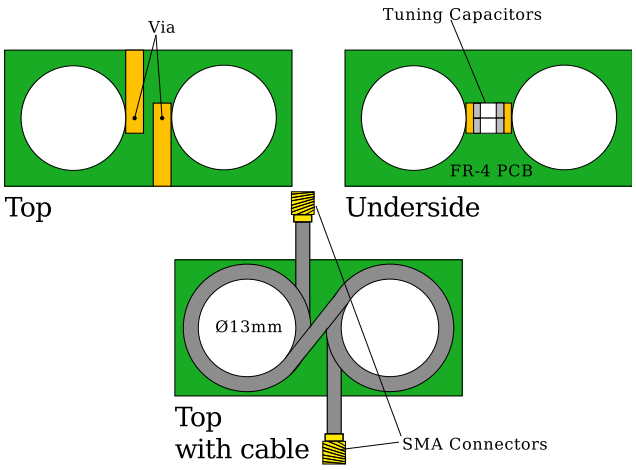

Figure 4.4: Figure-Eight Cable Trap PCB Design
Non-magnetic semi-rigid coaxial cable (Micro-Coax UT85C-TP, Pottstown, PA, USA), coiled into a figure eight is soldered to the top of a double sided PCB. On the reverse side there are pads for ceramic chip capacitors to be soldered for tuning the trap. Fine tuning is achieved by adjusting the spacing between the inductor coils, as the outer shield is made from solid copper they will retain their shape.

Figure-Eight Trap The figure eight trap forms an inductor on the outer shield of some coaxial cable by coiling it into a "figure eight" shape. The trap circuit is completed by bridging this inductor with a capacitor. B_1 and B_0 distortions are minimised as the figure eight shape cancels the magnetic field produced by the windings. The effectiveness of a cable trap is governed by the size of the inductor in the trap. Coiling the cable in this way creates a large inductance and

typically these traps provide >40dB of attenuation. However due to their high-Q nature, this degree of attenuation will only be available over a narrow bandwidth and so it is important that they are fine tuned to 123.2MHz.

Two of these traps were required, one for each coil, positioned within the base section (see 4.1.3) after the first 30cm of cabling. Figure 4.4 is a schematic of the design of these traps.

Floating Cable Trap The second type of trap was a variant of the $\lambda/4$ balun [38, p.138], following the design of Seeber et al. [200] to not require a direct electrical connection to the cable. These traps were placed on the long (5m+) lengths of cable between the coil array and the filter plate in magnet room.

A transmission line of length $\lambda/4$ is resonant, and when connected to the shield of a cable presents a high impedance to common mode currents travelling along the cable shield. This can be implemented as a second shield of length $\lambda/4$ over the coaxial cable, electrically connected to the cable at one end, as shown in figure 4.5.a. Inconveniently $\lambda/4$ is approximately 60cm for 123.2MHz, so to reduce the length of the balun capacitors are connected between the balun and cable at the opposite end to tune the 'bazooka' balun to resonate accordingly (figure 4.5.b). The lower inductance from the reduced length results in a lower Q and blocking impedance than the $\lambda/4$ balun. The physical connection to the cable can be completely removed by splitting the trap into two halves, each a resonant circuit, as shown in figure 4.5.c. Cable shield currents induce opposing circulating currents in the two halves, imposing a high impedance to the common mode currents. Adjusting the separation between each half fine tunes the trap. A shield (not shown) can be placed around this trap to confine the high voltages that form across the capacitors, and to protect the trap from external interactions which will affect its resonant frequency.

This style of trap has several advantages:

- It does not require an electrical connection to the cable and so can be easily installed.
- Each trap can be fitted to more than one cable.
- Unlike traps consisting of coaxial cable wound into an inductor (such as the figure-eight trap) no extra cable length is required.

A modified version of the trap featured in [200] was designed, incorporating a shield fully surrounding the two trap halves to protect patients from the high voltages developed across the capacitors, and also prevent any capacitive coupling between the trap and the patient or RF shield inside the magnet bore (which would shift the frequency and reduce the Q). Computer Aided Design (CAD) software (Inventor, Autodesk Inc.) was used to design the sections of the trap, as shown in figure 4.6. Each trap was manufactured from nylon using selective laser sintering (SLS), see section 4.1.3 for more details.

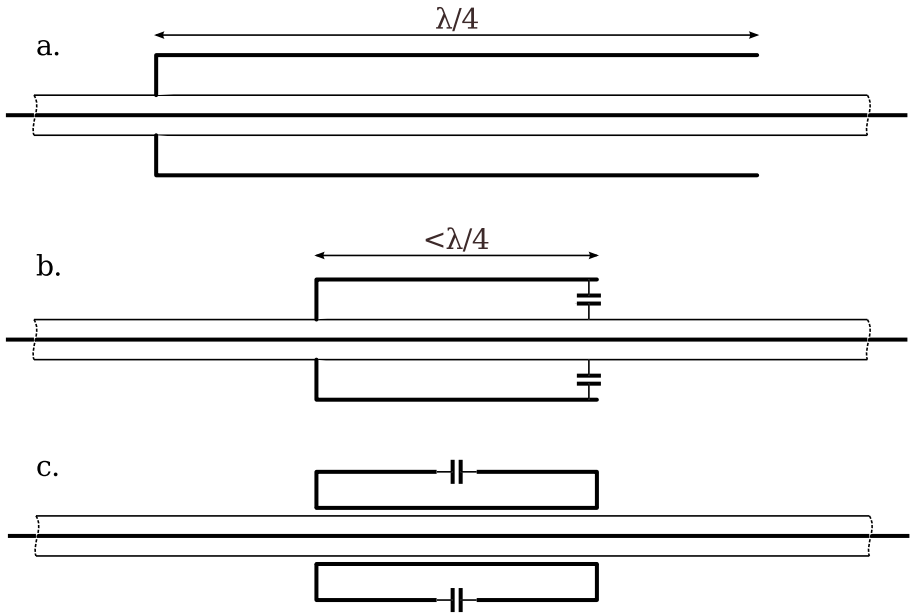

Figure 4.5: Schematic diagrams of a. $\lambda/4$ balun, b. bazooka balun, and c. floating cable trap

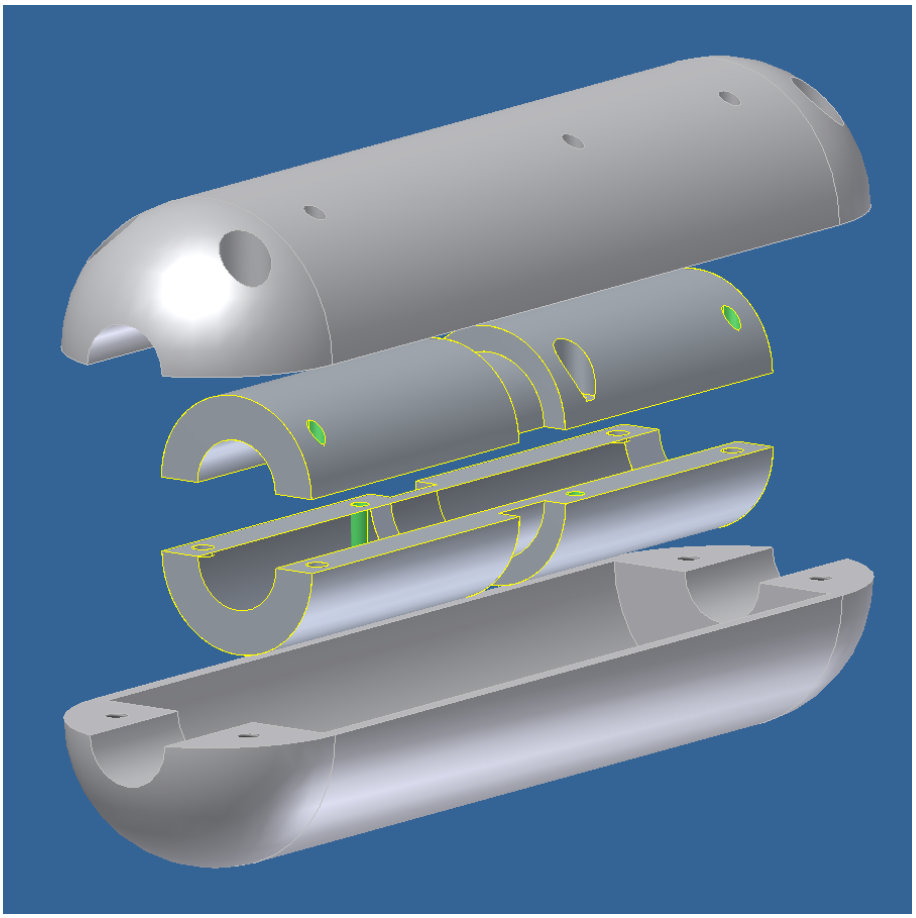

Figure 4.6: 3D CAD Model of the Floating Cable Trap Assembly
From bottom to top: bottom shield cover, bottom trap half, top trap half and top shield cover. When clamped together the inner diameter of the trap is 12.5mm, and can be clamped around up to four RG223 cables. As there are only two cables, two 10cm lengths of RG223 cable were placed inside to improve the grip on the cables, and also improve the coupling between the cables and trap.

4.1.2.8 Final Coil Design

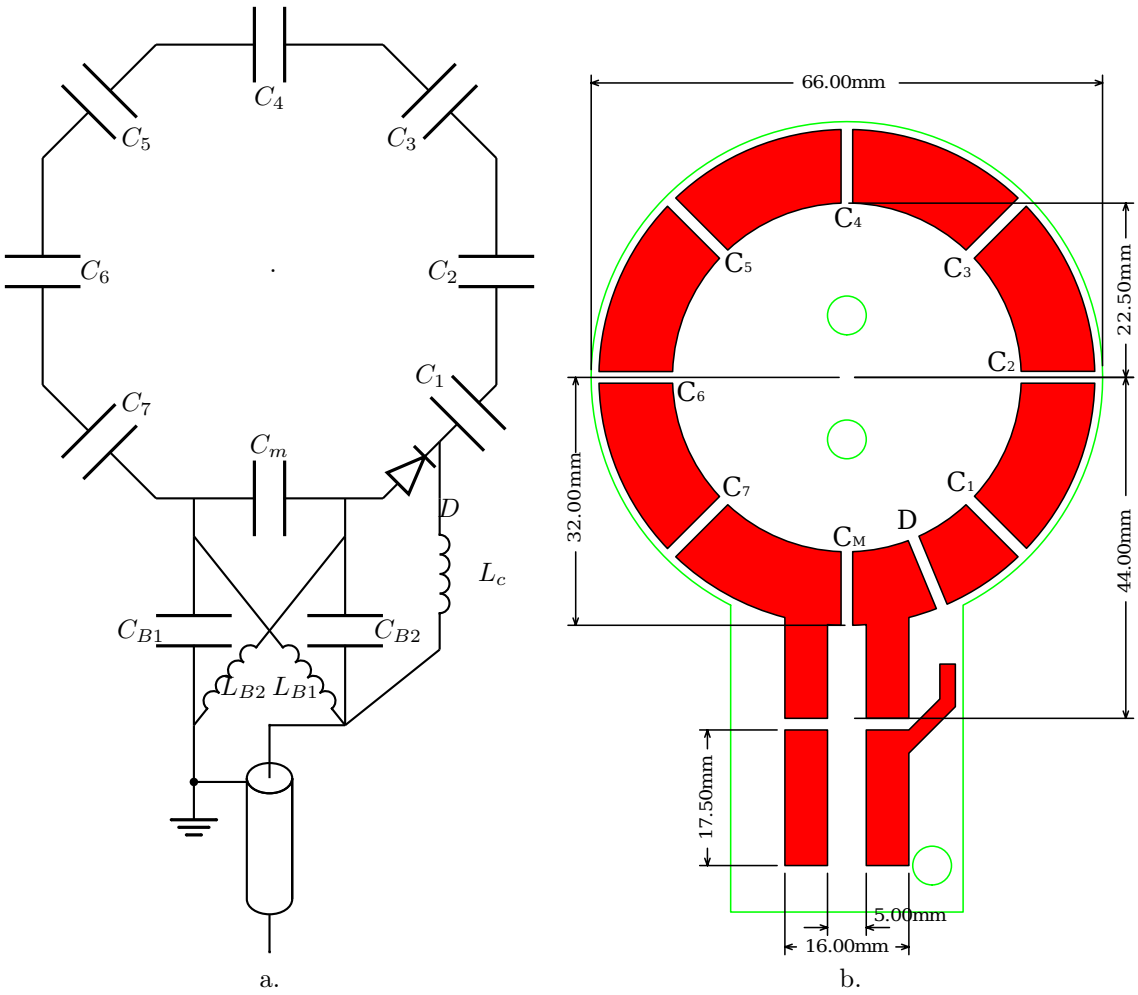

Figure 4.7: Final Circuit Diagram and PCB design

Combining the coil, PIN diode detuning circuit and balun (as shown on figure 4.7.a) simplifies the DC path so that only one choke is required. DC bias is provided by the same coaxial cable as the RF current; the bias current provided by the PIN diode driver in section 5.4 (the output MOSFET is sinking current) travels from ground to the anode of the diode via L_{B2} and then to the inner conductor of the coaxial cable via the RF choke L_c. This was realised as a printed circuit board (PCB) on copper clad FR-4, as shown in figure 4.7.b.

4.1.3 Mechanical Design

Housing was designed and constructed in collaboration with PulseTeq such that that:

1. Coils can be positioned anywhere on the subject's neck.
2. A separation between the coil and sample of approximately $h \sim {}^{d}/_{10} = 6.4mm$ ([158, p.441]) was maintained to optimise the coil loading, parasitic capacitance and B_1 penetration.
3. Nothing was placed across the windpipe as this could cause discomfort to the subject.

4. The coils can be easily used in conjunction with the Siemens Trio (including coils and the patient bed).

Figure 4.8 is a diagram of the coil housing. One of these was made for each coil out of nylon using selective laser sintering (SLS), a rapid prototyping technique[49][3]. A goose-neck made from 1/4 inch (6mm) bore coolant hose (Jeton R/D & MFG. INC., Taichung, Taiwan) connected each coil to a base designed to fit into a slot on the Trio patient couch which would normally hold a neck matrix coil (see figure 4.9). The base was made from clear sheets of acrylic, laser cut into pieces that assemble three dimensionally. Cable from each coil ran through the goose-neck and the bottom of the base, which also contained a figure-eight cable trap for each coil, to a BNC socket.

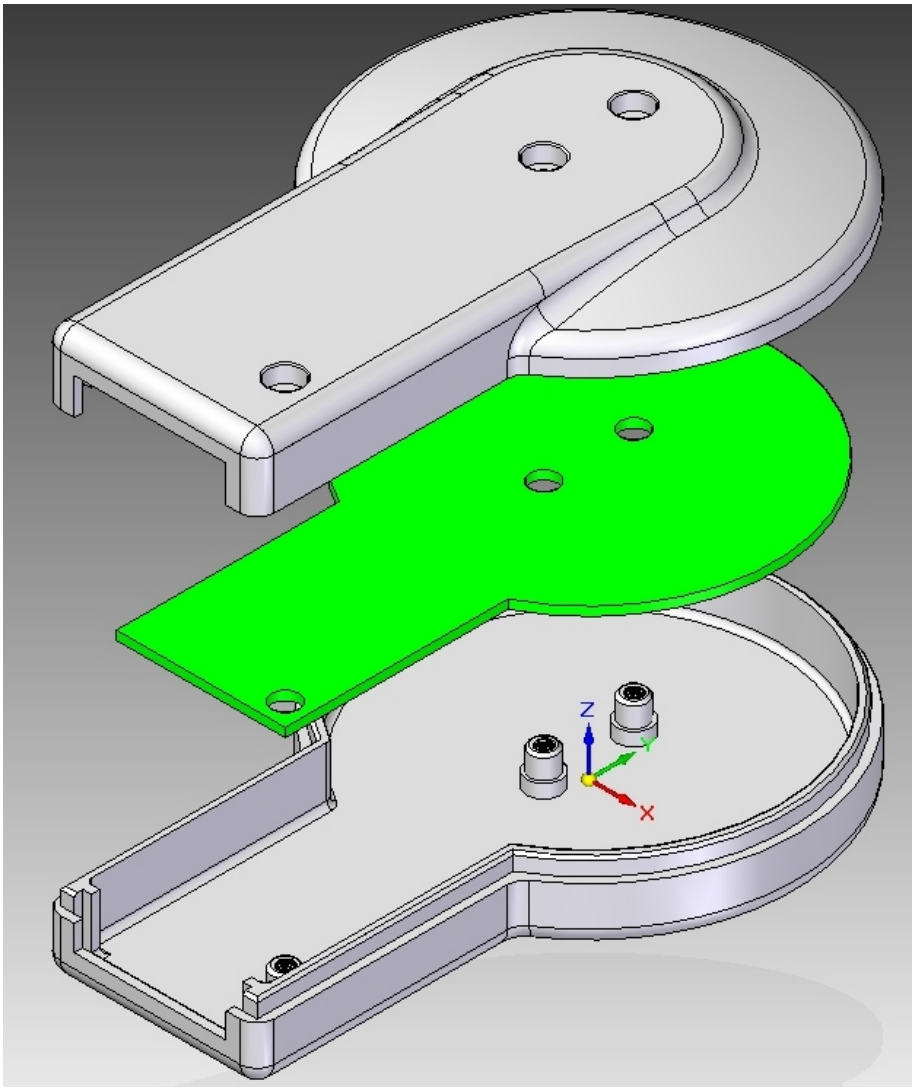

Figure 4.8: 3D CAD Model of the Coil Housing and Coil PCB (shown in green)

[3]Selective laser sintering is a rapid prototyping technique which fuses together powdered material with a laser to quickly build up a high quality object. Fine layers of powder are rolled out and sintered with a laser. Non-fused powder forms a support matrix for newer layers. The level of detail possible is limited only by the thickness of the layers and the resolution of the layer. Many different materials can be used, including nylon, glass reinforced nylon, aluminium and aluminium-plastic hybrids.

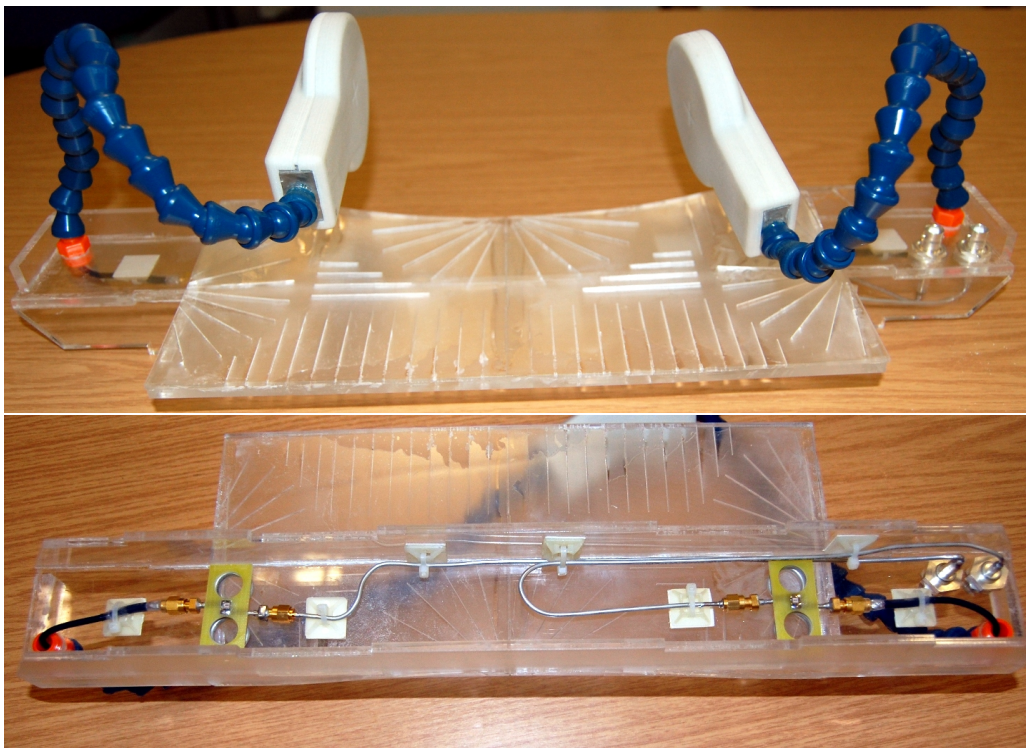

Figure 4.9: Complete Mechanical Assembly
The complete mechanical assembly consisting of the coil housing, goose-necks and base. Coil 1 is on the right and coil 2 on the left. RG316 cable runs from each coil's PCB through the goose-neck and into a cavity that runs the length of the base to connect to a figure-eight trap. UT85C-TP semi-rigid coaxial cable was then used for the remainder of the cabling up to bulkhead BNC connectors for its rigidity and low loss. Cable length is kept the same so that both coils are in phase with each other.

4.2 Implementation and Characterisation

4.2.1 RF Laboratory

An RF Laboratory was set up for the construction and testing of the coils. This consisted of:

- Clean environment free from magnetic/ferrous materials. Even the smallest particle of iron/steel will cause large field inhomogeneities and so coils must be quarantined from any materials/tools which could introduce ferrous particles.
- Non-Magnetic components segregated from magnetic components.
- Temperature controlled soldering iron for soldering of RF components at a relatively low temperature (240°C) so that no damage occurs to them. If heated for too long chip capacitors show increased loss (increased ESR, decreased Q), and PIN diodes are damaged.
- Plastic table for performing coil measurements. Any conductive object (metal, wood, people) in close proximity (<1m) to a coil perturbs measurements.
- Cable Current Test Bed: a controlled environment for testing cable traps and baluns. It consisted of a large conductive plane constructed from a sheet of plastic covered in aluminium

foil and secured with aluminium tape with wooden blocks fixed at each end for securing a cable to. When used a cable is strung taut across the conductive plane and supported with more wooden blocks to maintain a constant height of 45mm above the ground plane. Ferrite beads are placed at the ends of the cable, and any free cable outside of the test bed is wrapped toroidally around ferrite rings to confine the shield currents to the test bed. If this is done insufficiently measurements will become very sensitive to external influences, particularly body movements (making it impossible to take consistent measurements). Common mode currents on the cable couple strongly to the conductive plane, and very minimally to other objects in the room, facilitating the controlled testing of cable traps and baluns. See figure 4.10.

- Test Equipment: HP 8712C Vector Network Analyser (Agilent, Inc. Santa Clara, CA, USA). Tektronix TDS410 150MHz Digitising Oscilloscope (Tektronix, Inc. Beaverton OR, USA), TTi TG1010 Programmable 10MHz Function Generator (Thurlby Thandar Instruments LTD, Huntingdon, Cambridgeshire,UK).

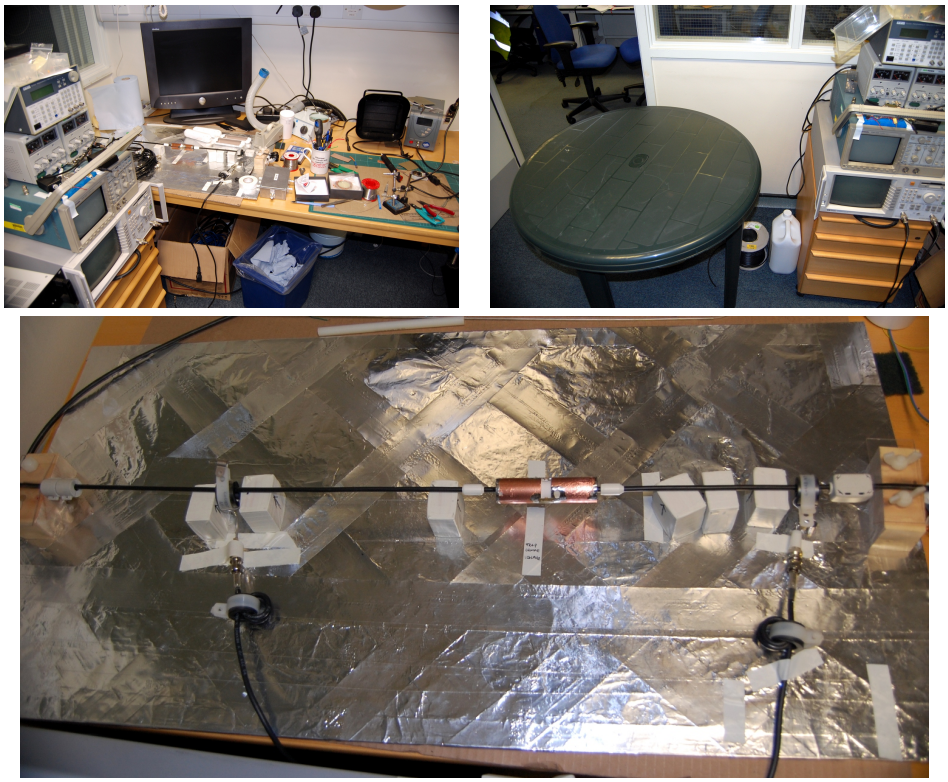

Figure 4.10: RF Laboratory Photos
Top Left: Work area.
Top Right: Plastic table for coil testing.
Bottom: Cable Current Test Bed (with cable, cable trap, current probes and ferrite beads in place).

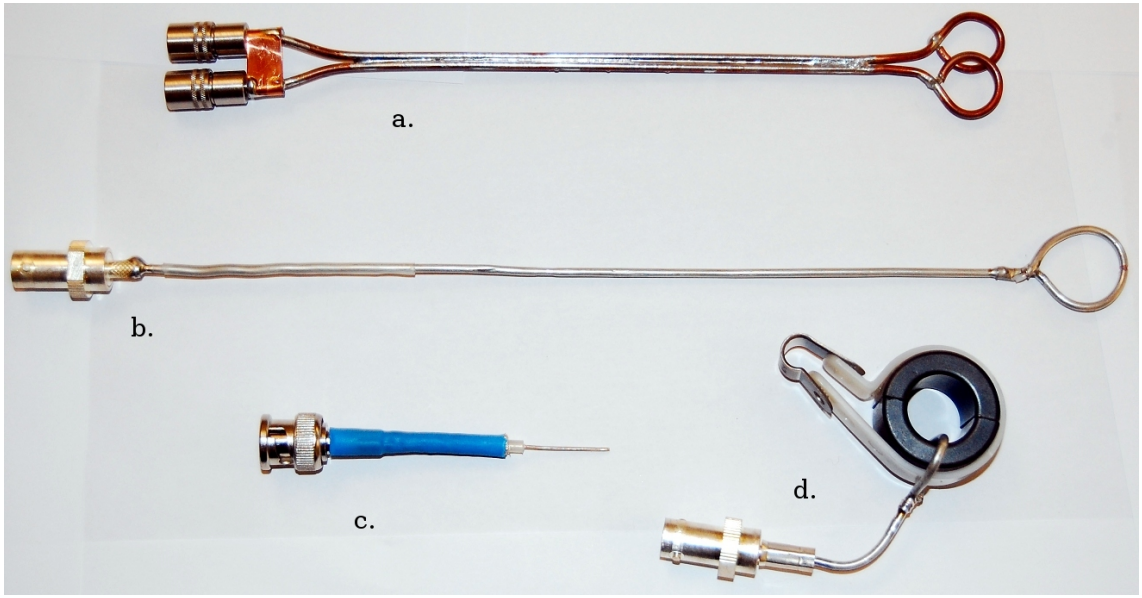

Figure 4.11: RF Probes
a. Shielded Double Loop Probe, b. Shielded Single Loop Probe, c. Monopole Antenna, d. Current Probe

Custom built RF probes were also made for testing of the coils, as shown in sections 4.2.1.1 to 4.2.1.4.

4.2.1.1 Shielded Single Loop Probe

A single turn loop of wire is an ideal magnetic field probe as the magnetic field will induce a detectable EMF in the loop (through the action of Faraday's Law). Enclosing the loop with a shield renders it insensitive to electric fields, and this can be easily realised by forming the loop from semi-rigid coaxial cable. This effectively forms an electrically small antennae which can be used to detect magnetic fields [236, 256, 63].

Figure 4.11b shows a shielded single loop probe, constructed from semi-rigid coaxial cable.

4.2.1.2 Shielded Double Loop Probe

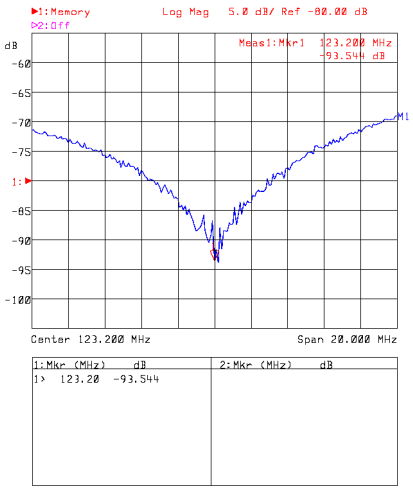

Figure 4.12: Nulling the Shielded Double Loop Probe

Two of these shielded loops can be used to detect resonant structures: one loop transmits and the second receives. A resonant structure will couple strongly to the magnetic field transmitted by the first loop if it is at the resonant frequency, focusing the field so that the second loop will detect it with more intensity. The transmit and receive loops can be assembled together and overlapped so that their mutual inductance cancels (and no signal is broadcast between the two).

Figure 4.11a shows the shielded double loop probe constructed for the ASL Labelling Coils, constructed from RG405/U semi-rigid coaxial cable. Two individual shielded loops were first built (the inner conductor is soldered to the cable shield at the end of the loop), and then soldered together. Loop circumference must be electrically small.

The probe is used with the network analyser in S_{21} (transmission) mode, with the loops connected to ports 1 and 2. Nulling of the probe is achieved by adjusting the separation of the two loops until there is a null at the desired frequency, as shown by figure 4.12.

4.2.1.3 Monopole Antennae

An electrically small monopole antenna is used to produce and detect electric fields, particularly useful in testing how well balanced a coil is. One was built (figure 4.11c) from a section of RG223 coaxial cable with 20mm of the inner conductor exposed.

4.2.1.4 Current Probe

A current probe is useful for measuring the common mode cable shield currents and can be built as a transformer around a ferrite ring. The cable runs through the ring and forms the primary winding, and a shielded single loop is formed toroidally around the ferrite ring as the secondary winding, as shown in figure 4.11d. The split ferrite ring allows the probe to be easily clamped over a cable. Two of these current probes were constructed and used for injecting and detecting common mode current on the shields of coaxial cables.

4.2.2 Construction

4.2.2.1 Phantoms

Saline Phantom Once the match capacitance of the first coil was determined (stage 3 of section 4.2.2.2) a saline phantom was constructed for repeatable loading conditions. Inside its housing, the coil was connected to port 1 of the network analyser whilst being loaded by a volunteer's neck. The network analyser trace for a 20MHz span around 123.2MHz was saved to memory. A 12cm diameter by 15cm cylindrical acrylic container was filled with distilled water and the coil fixed to the side with adhesive tape. Sodium chloride was added to the water until the loading conditions were such that the coil was matched to 50Ω.

Agarose Phantom Due to its conductivity and fluidity large circulating currents are induced in the saline phantom which shift the resonant frequency of a coil placed in close proximity. So, once the coils were completed a second phantom consisting of agarose gel was constructed in a 12cm diameter by 20cm cylindrical container. A network analyser trace of a S_{11} measurement of one coil when loaded by the saline phantom was saved to memory (20MHz span around 123.2MHz). The coil was then fixed to the side of the new phantom container with adhesive tape and sodium chloride was added to the water until the network analyser trace matched the one in memory. This water was then used to mix 2L of agarose, which was poured back into the phantom and allowed to set overnight.

In addition to not shifting the resonant frequency of the coil, an agarose phantom can be used for calorimetry and MR temperature mapping [195, 64, 118, 170, 203, 251, 41] to determine any localised hotspots as there is no convection.

4.2.2.2 Coils

All capacitors used were high voltage, high Q, P90 dielectric ceramic chip capacitors (1111P series, Passive Plus Inc. Huntingdon Station, NY, USA). PCBs were manufactured by milling copper-clad FR-4 laminate. PIN diodes were non-magnetic chip type (MA4P7464F-1072T, M/A-COM Technology Solutions Inc., Lowell, MA, USA).

Each coil was built to the following procedure:

1. All but one of the gaps in the PCB were bridged with a small piece of copper. A capacitor was soldered across the remaining gap (corresponding to C_4 in figure 4.7) and the resonant frequency of the coil measured using the shielded double loop probe[4]. Using the known value of capacitance the inductance of the coil was calculated using the equation for resonance (equation 2.52). Then the capacitance required to tune the coil (using the recently calculated value of inductance) was calculated and added to the coil. This was repeated iteratively until the coil was tuned to resonate at 123.2MHz.

2. Using the shielded double loop probe the unloaded and loaded Q of the coil was measured with the coil in its housing and positioned on the neck of a healthy 24-year-old male volunteer of average build. The match capacitance was then calculated and added to the coil in the match capacitor position. Then the value at C_4 was adjusted so that $1/C_T = 1/C_m + 1/C_4$, where C_T, the total tuning capacitance was calculated using the resonance equation. The match of the coil when loaded was checked by soldering a coaxial cable across the match capacitor and connecting to port 1 of the network analyser. Measurements were made in S_{11} (reflection) mode of the impedance at 123.2MHz. A correct match will result in $Z = 50\Omega$. Then the cable was de-soldered and the loaded resonant frequency of the coil measured.

[4]All measurements made using the shielded double loop probe are done with the coil in complete isolation - no RF connection is made across the match capacitor either directly with a coaxial cable, or to the lattice balun.

3. This process was repeated iteratively until the coil was correctly tuned and matched at 123.2MHz on the volunteer's neck. A saline phantom was then constructed according to the steps in 4.2.2.1. The associated frequency shift of the saline phantom was recorded and all further measurements involving the saline phantom took this shift into account.

4. The PIN diode was then added to the coil and RF chokes soldered either side for a DC bias. The tune and match capacitance was adjusted to account for the increased resistance ($\sim 0.8\Omega$) and inductance of the PIN diode.

5. With the correct tune and match capacitance values known, the tuning capacitance was then distributed to all seven capacitor positions such that there were three ceramic chip capacitors in parallel at each gap. Again this is an iterative process. Shimming of the capacitors so that not all positions have the same capacitance was required, and these changes were made symmetrically with respect to the axis running through C_m and C_4.

6. With the correct distributed tuning capacitance determined the capacitors at position C_4 were replaced with a 1-23pF high voltage non-magnetic variable capacitor (NMAF25HV, Voltronics Corporation, Denville, NJ, USA) and two chip capacitors. For coil 1, 56+56+56pF was replaced with the trimmer set half way (approximately 11.5pF) and 2x 82pF capacitors. For coil 2 this was with the trimmer set half way and 100+33pF.

Final values of the capacitors are shown in table 4.7

Coil 1	C (pF)				Coil 2	C (pF)			
	1	2	3	Total		1	2	3	Total
C_1	56	56	56	168	C_1	56	56	56	168
C_2	56	56	47	159	C_2	56	56	47	159
C_3	56	56	56	168	C_3	56	56	56	168
C_4	82	82	1-23	165-187	C_4	100	33	1-23	134-156
C_5	56	56	56	168	C_5	56	56	56	168
C_6	56	56	47	159	C_6	56	56	47	159
C_7	56	56	56	168	C_7	56	56	56	168
C_M	47	39	33	119	C_M	47	47	39	133
Total	19.71pF				Total	19.61pF			

Table 4.7: Coil Capacitor Values

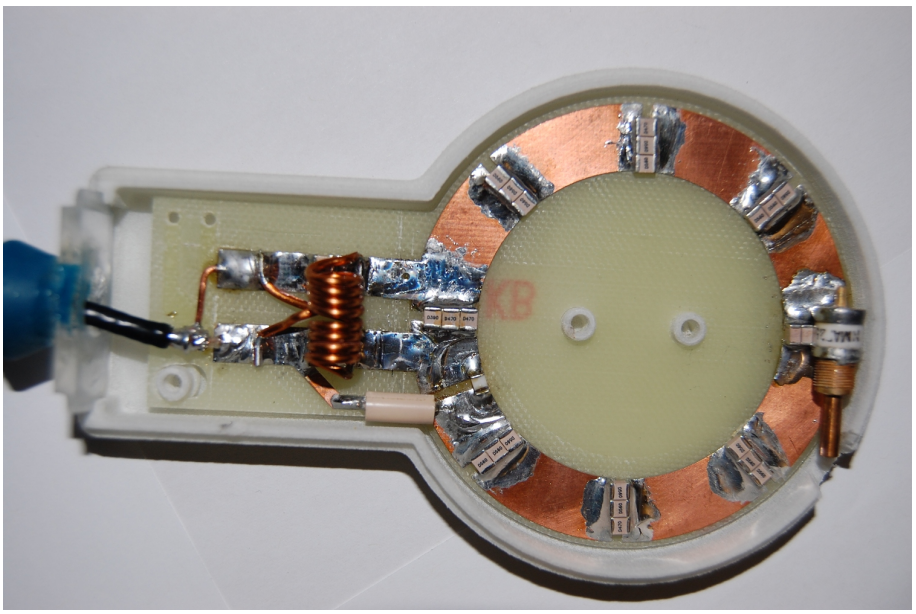

Figure 4.13: Populated Coil PCB (Coil 2)
At the very left of the image is the feeding coaxial cable, soldered to the feed-point of the lattice balun, which is identified from the two copper coils of wire. Current is sourced through the RF choke (beige cylinder below the balun) when the PIN diode (small white chip to the right of the RF choke) is forward biased, and +5V is applied during reverse bias to prevent the diode from switching on. The match capacitors are located at the three o'clock position on the coil loop, and tuning capacitors are located at all other gaps. There is a small slot cut into the housing to allow for fine-tuning of the trimmer capacitor once the housing has been fully assembled.

4.2.2.3 Balun

A lattice balun as shown in figure 4.2 were constructed on each coil PCB. Two 12pF capacitors were used in parallel on each half of the balun, and each inductor consisted of 1.25mm enamelled copper wire wound four times around a 5mm mandrel. The two inductors were counter wound with respect to each other.

4.2.2.4 Cable Traps

Figure-Eight Trap Two figure eight cable traps were constructed per the design in figure 4.4. A 12pF capacitor was required for each trap to tune to 123.2MHz.

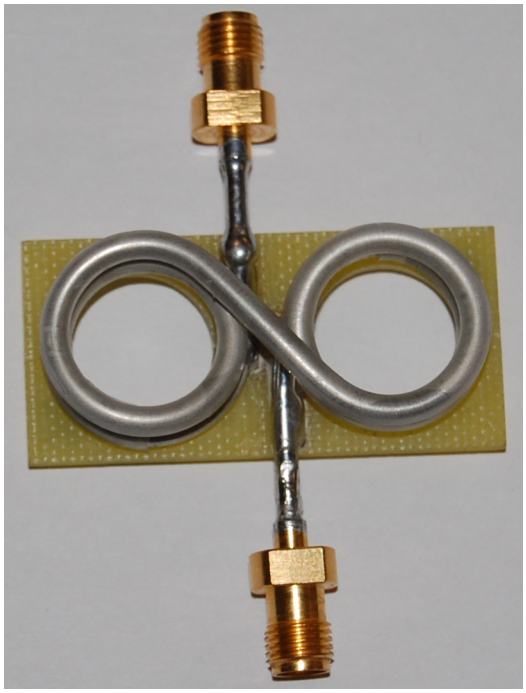

Figure 4.14: Finished Figure-Eight Cable Trap

Floating Cable Trap The floating cable trap CAD models (figure 4.6) were manufactured from nylon using the SLS 3D printing process. Copper tape was applied to the exterior surfaces of the inner trap halves, and aluminium tape to the inside of the shield covers followed by an insulating layer of PTFE tape. Ceramic chip capacitors were soldered across the gap at the top of each half to resonate them. The inner trap halves were held together around the cable with nylon screws and their separation adjusted for fine tuning using four grub screws.

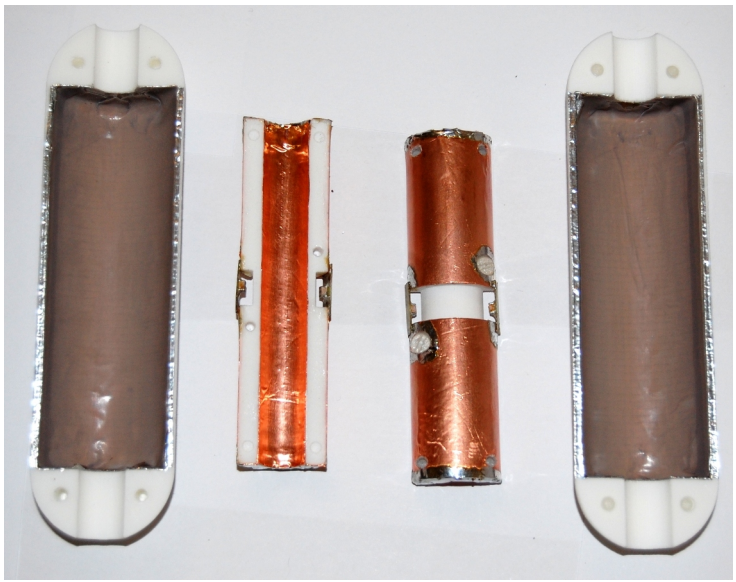

Figure 4.15: Finished Clamp-On Cable Trap

4.2.3 Characterisation

4.2.3.1 Coil Q

Measurements of the quality factor (Q) of each coil were performed using the double loop probe connected to ports 1 and 2 of the network analyser, which was set to measure the S_{21} transmission coefficient between the two loops in the probe. Each coil was disconnected from the balun to make it an isolated resonant circuit. RF Chokes were used to bring DC bias to the PIN diode. The probe was positioned in an aluminium vice on the edge of the plastic test table and held in position. Interactions between the probe and coil resulted in a shift of the resonant frequency and reduction in the Q of the coil, so it was desirable for the probe to be as far away from the coil as is possible to minimise perturbations. For all measurements the coil was slowly brought close to the probe until the received signal became strong enough for the network analyser to make a bandwidth measurement[5].

With each coil unloaded, and loaded by the agarose phantom the centre frequency and Q was measured. For all measurements the PIN diodes were forward biased with 100mA of current. Using these measurements and the nominal value of tuning capacitance on each coil, the coil inductance, loaded parallel resistance, series resistance[6] and power required to produce a B_1 of $3.5\mu T$ at a depth of 2.5cm and 5cm[7] were calculated.

Results of the measurements are shown in table 4.9. Figure 4.16 is a plot of the required power to produce a B_1 field of $3.5\mu T$ at a given depth based on the measured series resistance of the coils. The relationship between power and B_1 penetration depth is approximately a sixth order polynomial, and as can be seen in the graph there are rapidly diminishing returns for depths greater than 4cm. At the depth of the carotid arteries, approximately 2.5cm the power required is very low, at 4cm approximately 0.8W is required, and at 5cm 2.4W. Fortunately the vertebral arteries are not normally any deeper than 4-5cm, and their flow velocity is lower than the carotids, so successful labelling can be achieved with a lower B_1 amplitude.

Figure 4.17 shows network analyser traces of the measurements made by the double loop probe on each coil. The effect of loading each coil can clearly be seen as the resonant peak reduces in height and increases in width.

	Unloaded		Loaded							
	f_0(MHz)	Q	f_0 (MHz)	Q	$C_T(pF)$	$L(nH)$	$R_p(\Omega)$	$R_s(\Omega)$	$P_{2.5cm}(W)$	$P_{5cm}(W)$
Coil 1	123.319	119.1	123.195	44.6	19.71	84.7	1.36	0.74	0.13	2.4
Coil 2	123.3	110.81	123.198	43.0	19.61	85.1	1.36	0.74	0.13	2.4

Table 4.9: Unloaded and Loaded Q of Each Coil

[5]This is automatically done by the network analyser using three markers. It first searches the trace for a maxima, then places two markers -3dB either side of the maximum. The quality factor is then calculated as $Q = \frac{f_0}{\Delta f_{-3dB}}$.

[6]The inverse of the parallel resistance: $R_s = {}^1/_{R_p}$

[7]Calculated with $B_z = \frac{\mu_o}{4\pi} \frac{2\pi R^2 I}{(z^2+R^2)^{3/2}}$ with the inner radius of the coil as R; 1.814A of current is required to produce the prescribed field.

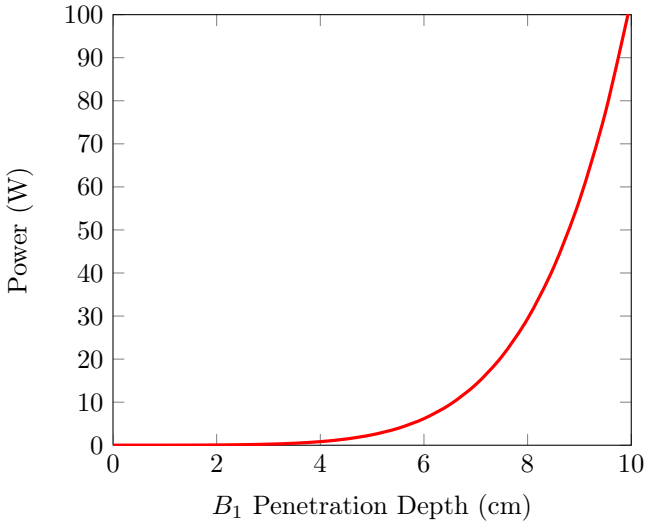

Figure 4.16: Power required to achieve $3.5\mu T$ at a given depth

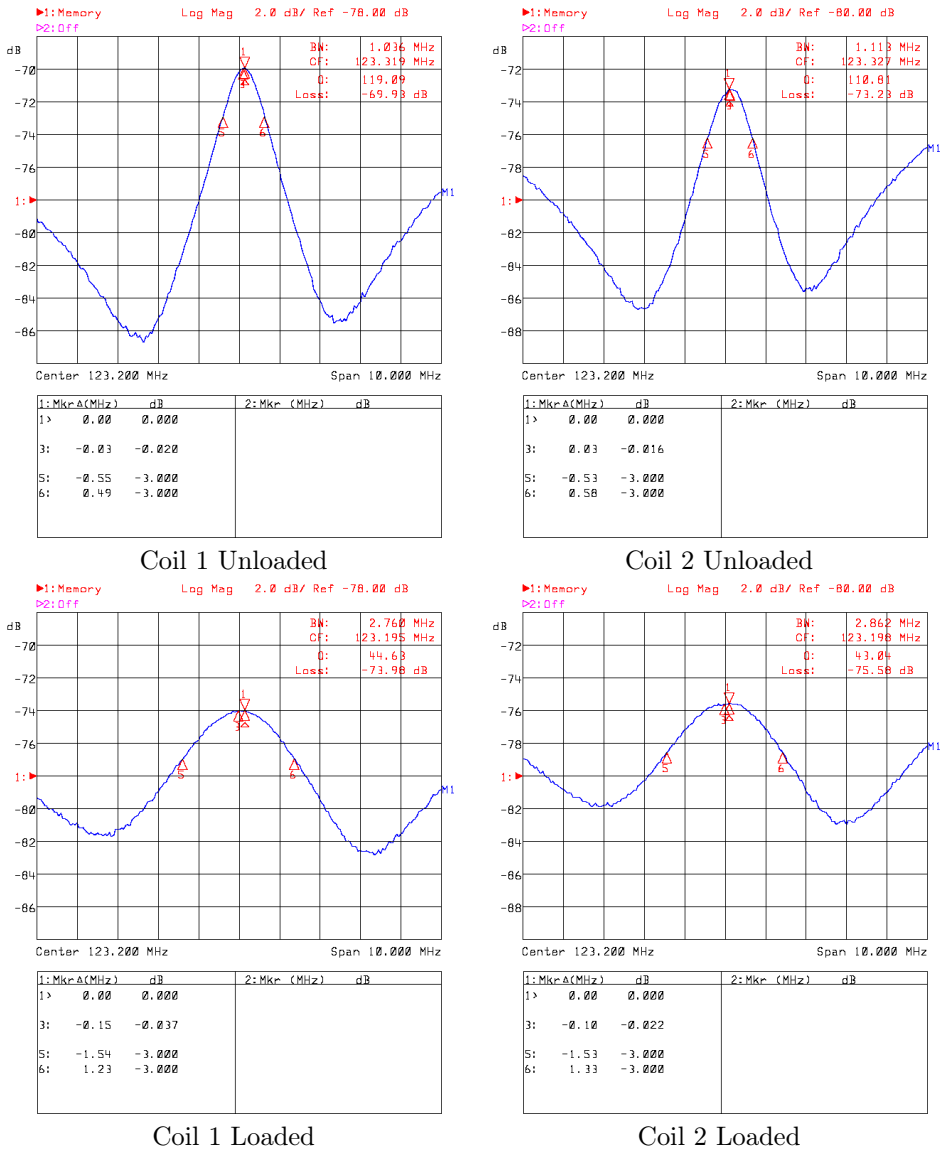

Figure 4.17: Network Analyser Traces of the Measured Q

4.2.3.2 Coil Match

The coil match was characterised by making S_{11} (reflection) measurements with the coil under test plugged directly into port 1 of the network analyser via a bias tee (see section . The network analyser was calibrated to the end of the cable that connected to the bulkhead BNC connector on the base section of the housing. Each coil was loaded by fixing to the side of the agar phantom and the value of the S_{11} coefficient and the impedance, Z, at 123.2MHz was measured.

	at 123.2MHz	
	S_{11} (dB)	Z (Ω)
Coil 1	-51.41	50.23 - 3.0i
Coil 2	-45.92	50.2 - 3.56i

Table 4.11: Match Parameters for Each Coil

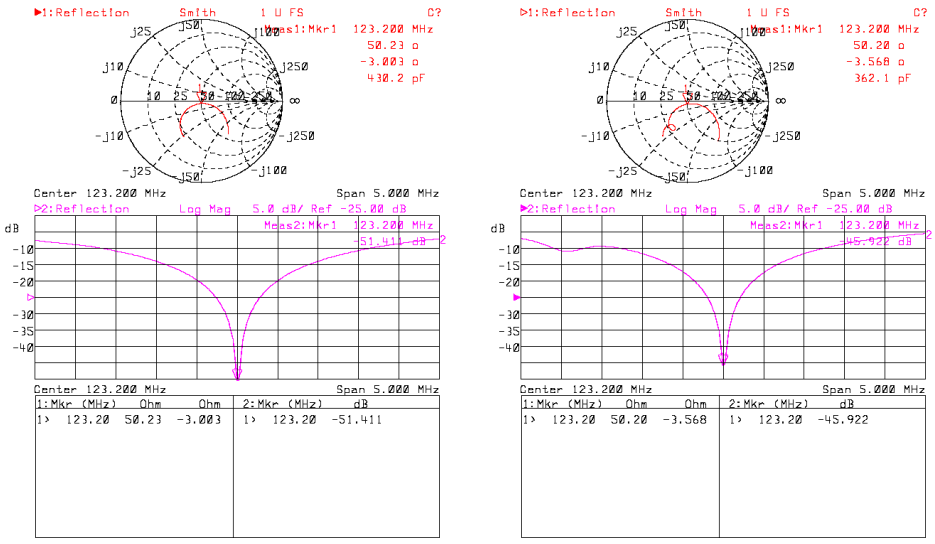

Figure 4.18: Network Analyser Traces of the Measured Q
Each trace shows both a logarithmic scale (decibels) magnitude plot of the reflected signal (bottom), and the measured impedance at each frequency plotted onto the complex plane, overlaid on a smith chart (top). The smith chart is set out such that 50Ω is at the centre, and so is useful when measuring the match of the coil. A marker is set to 123.2MHz, for the log-mag plot this is at the bottom of the null, and for the smith chart this is at the centre. Both are indicative of a well matched coil.

Table 4.11 shows the measured reflection coefficient and impedance of the coils at 123.2MHz. A reflection coefficient of -20dB or less is considered a good match, so in both cases the coils are very well matched. As the impedance measured is very close to a pure 50Ω, with only a slight capacitive component, most of the power delivered by the amplifiers will be used to generate a B_1 field.

4.2.3.3 PIN Diode Detuning

PIN diode detuning was tested in two ways: the result of detuning the coil on the Q on an unloaded S_{21} measurement of the coil Q similar to those made in section 4.2.3.1; and secondly when the

coil is plugged into port 1 of the network analyser and an S_{11} measurement is made as shown in section 4.2.3.2.

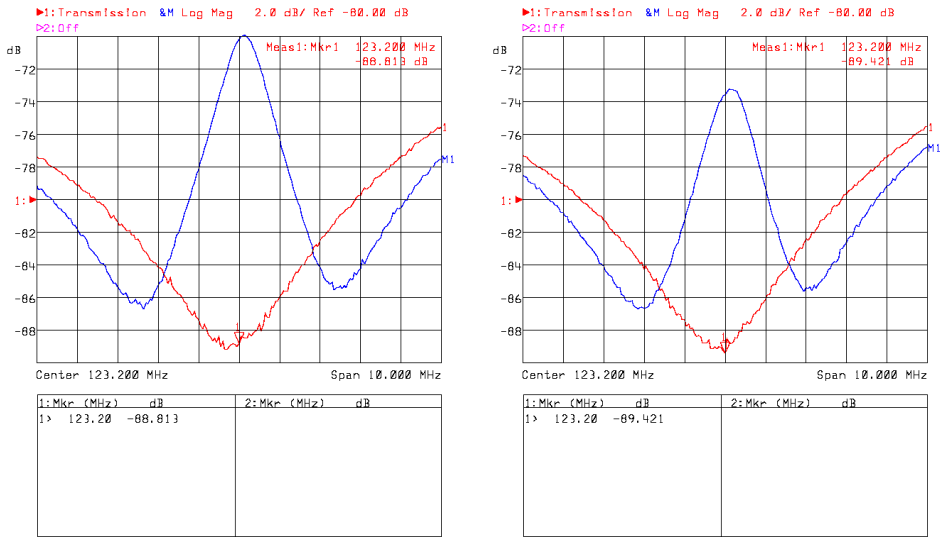

Coil 1 Coil 2

Figure 4.19: Q Spoiling By Reverse Biasing the PIN Diode
Blue trace – PIN diode forward biased (coil tuned); red trace – PIN diode reverse biased (coil detuned).

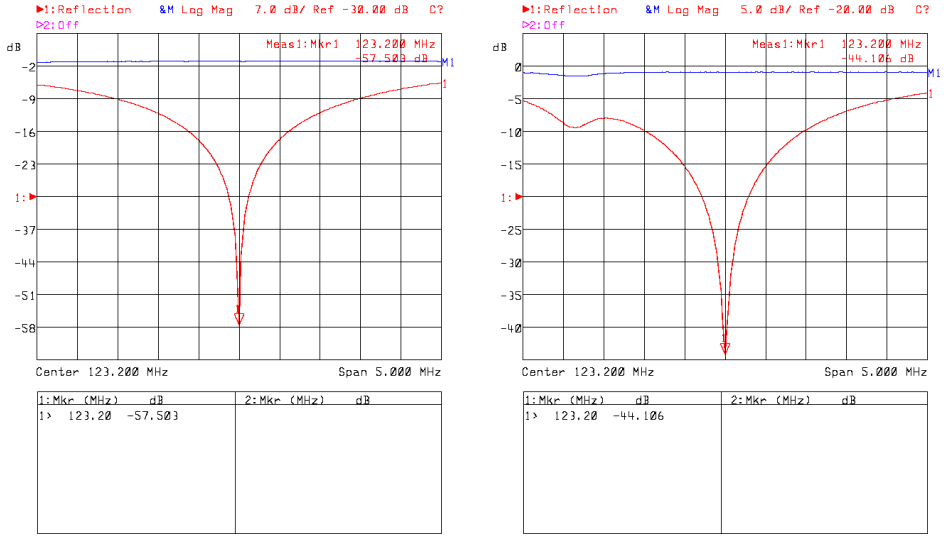

Coil 1 Coil 2

Figure 4.20: Match of the Coil When Reverse Biasing the PIN Diode
Blue trace – PIN diode forward biased (coil tuned); red trace – PIN diode reverse biased (coil detuned). When detuned, the value of S_{11} at 123.2MHz was measured to be -0.976dB and -0.932dB for coils 1 and 2 respectively.

Figure 4.19 shows the the measured trace from the S_{21} measurements. When forward biased the resonant peak at 123.2MHz can be clearly seen on both traces from both coils. With the coil detuned all that can be seen is the null from the double loop probes, and the measured value is

at the noise floor. Figure 4.20 is the measured trace from the S_{11} measurements. When forward biased both coils are well matched at 123.2MHz. When the diode is reverse biased the trace becomes very flat and almost at 0dB, which means that nearly all of the signal is being reflected back to the network analyser.

Both results from figures 4.19 and 4.20 indicate that the PIN diodes are very effective at detuning each coil.

4.2.3.4 Coil Coupling

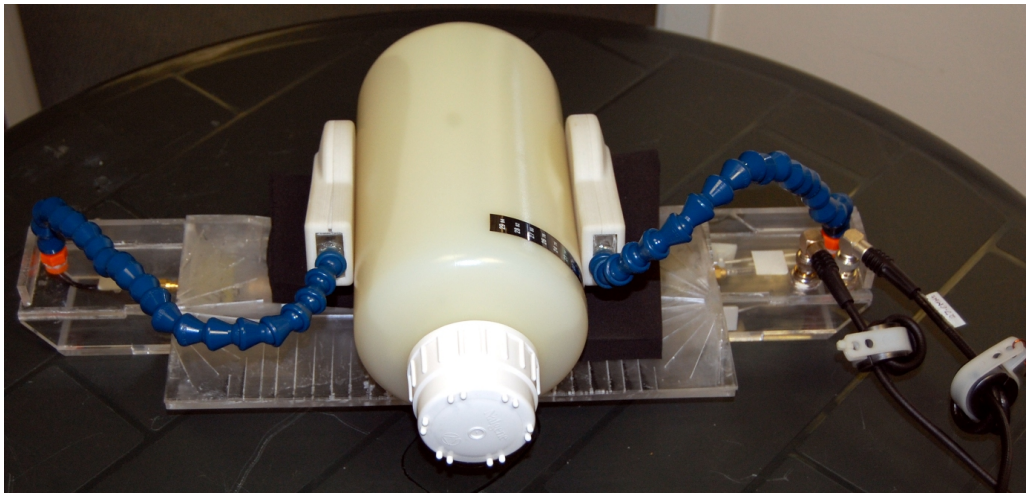

Figure 4.21: Testing the Coupling Between Each Coil

The coupling between the two coils was measured by transmitting on one coil and receiving with the second whilst they are both loaded either side of the agarose phantom, as shown in figure 4.21. This arrangement is similar to how the coils would be used in-vivo. Measurements were taken with the PIN diodes on the coils forward and reverse biased (100mA and -5V respectively), and with each coil as both a transmitter and receiver.

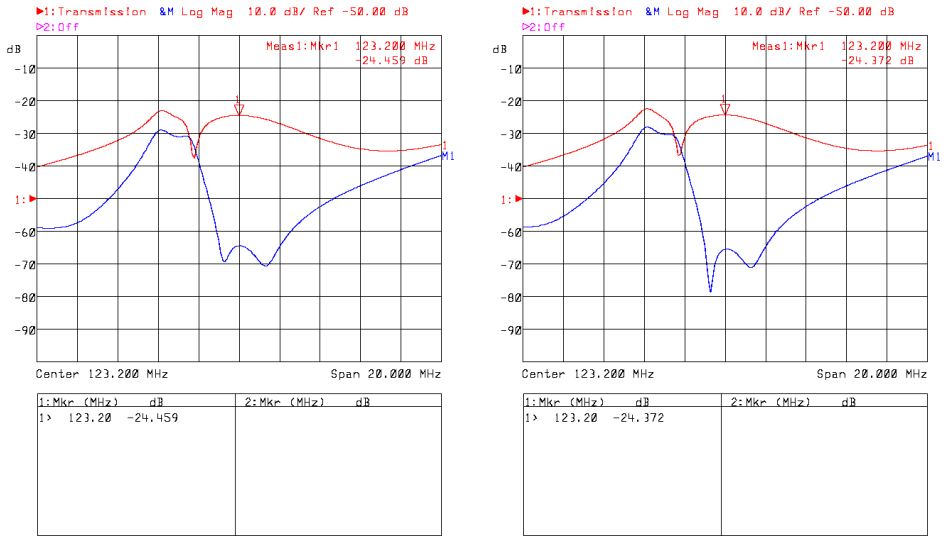

Tx: Coil 1, Rx: Coil 2 Tx: Coil 1, Rx: Coil 2

Figure 4.22: Network Analyser Trace of the Coupling Between the Two Coils
Blue trace – PIN Diode Forward Biased (coil tuned); red trace – PIN Diode Reverse Biased (coil detuned)

		Receive (dB)	
		Coil 1	Coil 2
Transmit (dB)	Coil 1		-24.48
	Coil 2	-24.37	

Table 4.13: S_{21} Transmission Coefficient Between the Coils

Coupling between the coils when the PIN diodes are reverse biased is minimal ($S_{21} \sim -65dB$). When the PIN diodes are forward biased the isolation between each coil is approximately 24.4dB. This is over 9dB greater than the isolation reported by Kurpad et al.[135] using voltage controlled current sources, and so it stands to reason that no further measures need to be taken to decouple the coils from each other.

4.2.3.5 Balun Common Mode Rejection

Each balun was tested by inducing common mode currents in the coil with the monopole antennae and detecting these with a current probe. At the resonant frequency of the balun a null in the common mode shield currents was seen. Fine tuning was accomplished by adjusting the separation of the inductor coils. The set up of this is shown in figure 4.23.

Measurements were made of the transmission coefficient (S_{21}) at 123.2MHz between the monopole antenna and current probe both with the balun active, and with one capacitor short circuited using a pair of conductive tweezers so that the balun no longer functions. These two measurements were subtracted to obtain the common mode rejection ratio (CMRR). Results are shown in table 4.15 and the network analyser traces in figure 4.24.

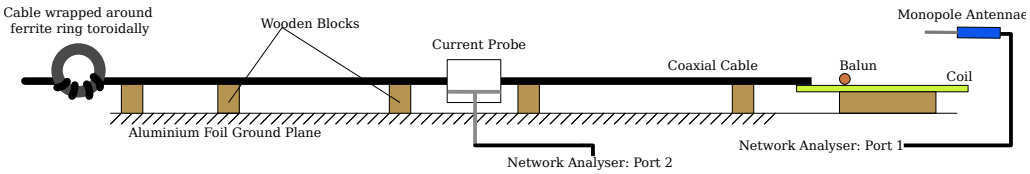

Figure 4.23: Lattice Balun Test Set Up

	CMRR @ 123.2MHz (dB)
Coil 1	26.1
Coil 2	35.3

Table 4.15: Balun CMRR

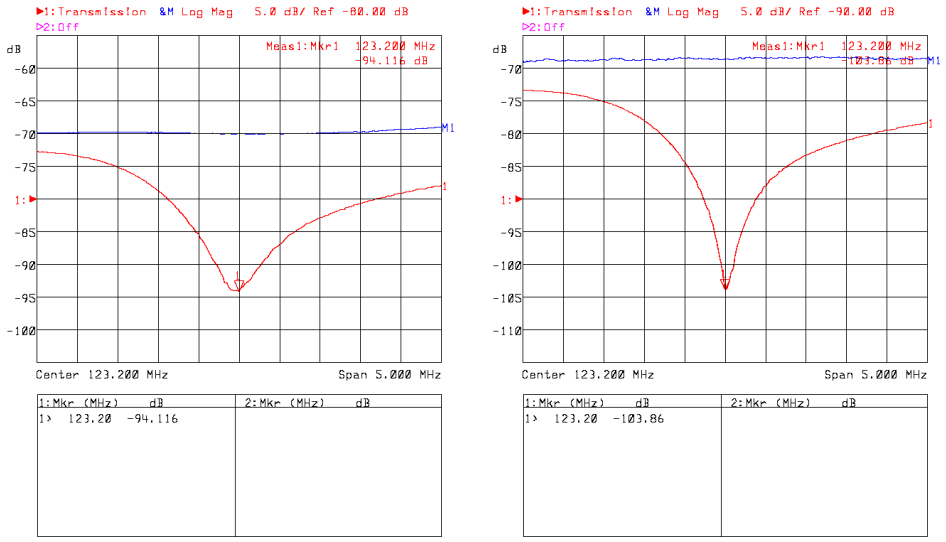

Figure 4.24: Balun Common Mode Attenuation
Blue trace – balun bypassed; red trace – Balun active

This measurement is highly sensitive to external interactions due to the radiative nature of the monopole antenna. Consequently the performance of the balun is difficult to quantify absolutely as the 9dB difference between coil 1 and coil 2 illustrates. Both however achieve a satisfactory CMRR.

4.2.3.6 Cable Traps

For both types of trap, tuning and characterisation was performed on the cable current test bed by injecting current onto the shield of a length of RG223 coaxial cable with one current probe and detecting the current on the shield with the second. Placing the trap between the two probes introduces a blocking impedance at the resonant frequency of the trap, attenuating the current detected by the receive probe. A schematic of this set up is shown in figure 4.25.

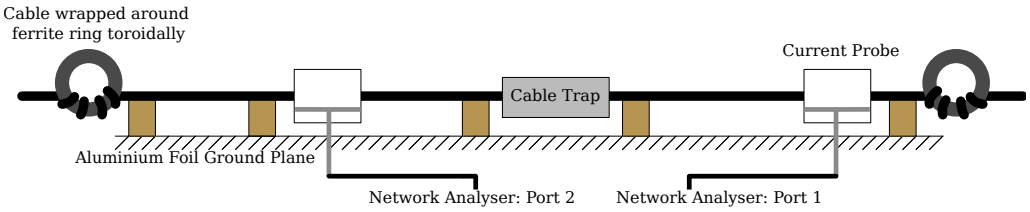

Figure 4.25: Cable Trap Test Schematic

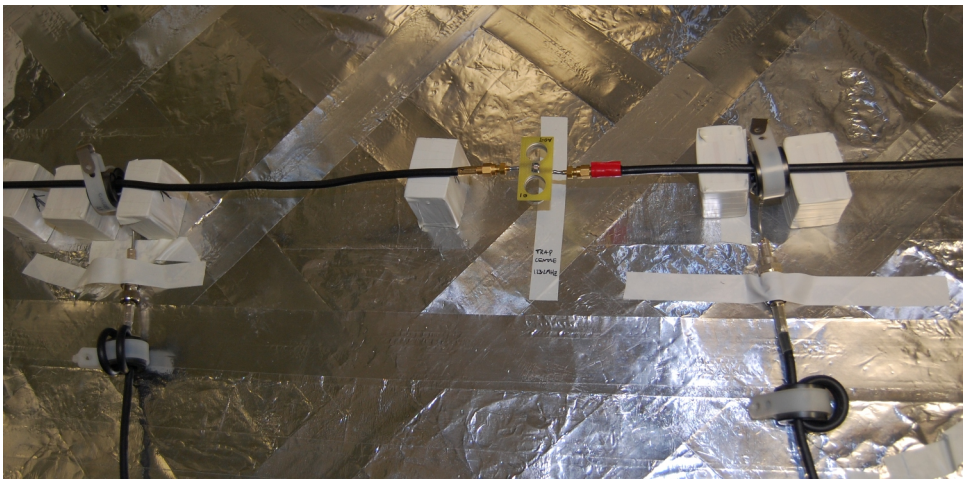

Figure 4.26: Figure-Eight Cable Trap Testing
Ferrites were also positioned on the current probe cables to prevent the common more currents on the test cable from coupling and perturbing measurements.

Figure-Eight Trap Measurements were made of the S_{21} transmission coefficient between the send and receive current probes with the trap active, and with the tuning capacitor shorted using a pair of metal tweezers, which were then subtracted to give the Shield Current Attenuation Ratio (SCAR) for each trap. Results are shown in table 4.17 and the network analyser traces in figure 4.27. Trap 1 is connected to coil 1, and likewise trap 2 to coil 2.

	SCAR @ 123.2MHz (dB)
Figure-Eight Trap 1	46.9
Figure-Eight Trap 2	47.5

Table 4.17: Figure-Eight Cable Trap SCAR

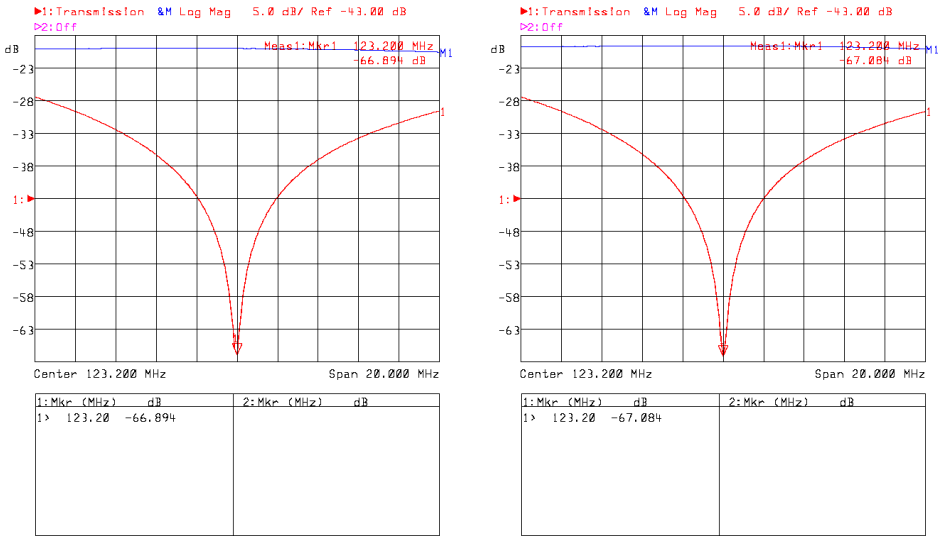

Figure-Eight Trap 1 Figure-Eight Trap 2

Figure 4.27: Figure-Eight Cable Trap Shield Current Attenuation
Blue trace – trap bypassed; red trace – trap active

Both figure eight traps provide a high degree of shield current attenuation - over three times the amount specified by Siemens' regulations. Inside the MRI scanner the coils are located in the middle of the body transmit coil, and radially off-centre, and consequently be subjected to strong magnetic and electric fields (the electric field from a birdcage coil increases at the periphery of the volume inside the coil). This high degree of attenuation will be advantageous in suppressing coupled RF currents due to the intense electric and magnetic fields.

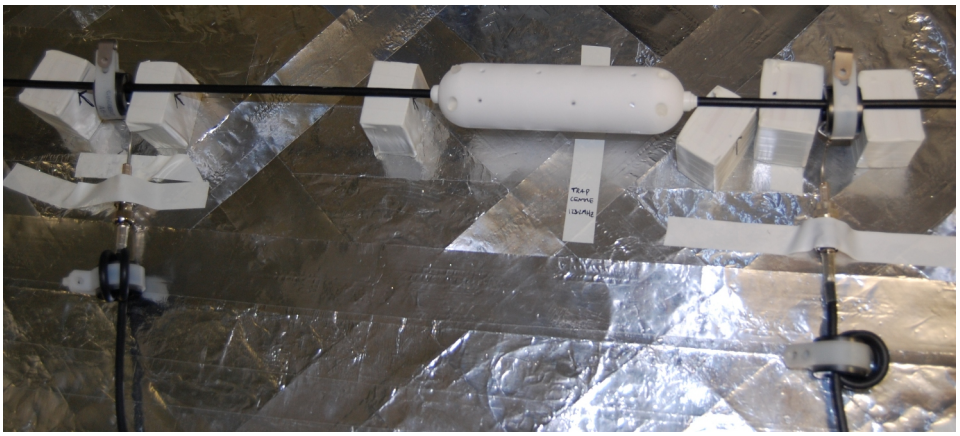

Figure 4.28: Floating Cable Trap Testing

Floating cable trap Testing and characterisation of the floating cable traps was performed on the cable trap test bed as per the set up in figures 4.25 and 4.28. A measurement of the S_{21} transmission coefficient was made without a trap on the cable as a baseline, and then each trap was fitted to the cable. The separation between each half of the inner trap was adjusted using the grub screws until a null was seen at 123.2MHz. Subtracting the value of S_{21} at 123.2MHz with the trap fitted from the baseline measurement at 123.2MHz provides the SCAR for each trap.

	SCAR @ 123.2MHz (dB)
Floating cable trap 1	24.8
Floating cable trap 2	24.9
Floating cable trap 3	24.0

Table 4.19: Floating Cable Trap SCAR

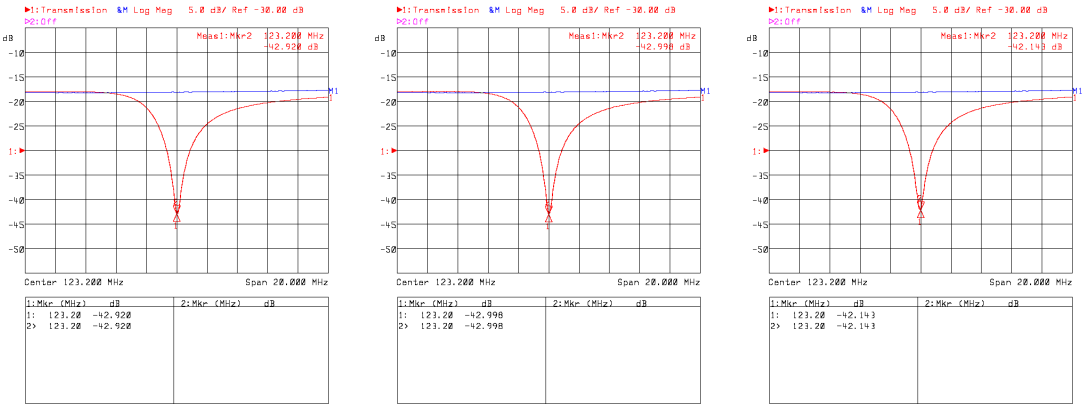

Floating cable trap 1 Floating cable trap 2 Floating cable trap 3

Figure 4.29: Floating Cable Trap Shield Current Attenuation
Blue trace – baseline measurement (cable with no trap); red trace – trap in place

All three traps exceed the 15dB Siemens specification by at least 9dB. After characterisation the traps were positioned on the two lengths of RG223 cable connecting the coil array to the magnet room filter plate. The high degree of common mode attenuation provided by both the figure eight and floating cable traps meant that trap spacing could be greater than 25cm. Trap one was positioned 20cm from the coil end of the cable, trap two 30cm from trap one, and trap three 30cm from trap two. Each trap was tuned to 123.2MHz before the next trap was fixed to the cable. When the coil array and cabling is in position within the magnet the cable traps extend outside the magnet bore. Not only does this minimise shield currents induced by the body coil during transmit as the traps extend well beyond the region where a significant radio frequency magnetic or electric field is produced, it also reduces any noise at 123.2MHz coupled into the scanner by the cable.

4.3 Conclusion

This chapter has detailed the design, construction and characterisation of two-element coil array for arterial spin labelling. The design of the coil array allows each element to be placed directly over the right and left side of the neck, optimising power requirements as only a few centimetres of B_1 penetration is required. Unlike other separate-coil CASL systems, the coils are supported by a goose-neck structure and hence not fixed to the neck; which may otherwise cause the subject discomfort, especially when resting on the windpipe. Each coil element is completely independent, and can used as an individual coil (for vessel-selective labelling), or a together as a coil array.

Intrinsic coupling between each coil is 10dB lower than that measured in parallel transmit coil arrays with specific decoupling technology (current source amplifiers, [135]) therefore decoupling mechanisms were not deemed necessary. Each coil was optimised to tune and match for the average loading conditions of the neck, removing the need to tune and match each coil on a subject-by-subject basis. To ensure subject safety, a bias current is required to tune each coil (in the event of a malfunction coils will remain detuned), and cable shield currents have been minimised through the use of a lattice balun and cable traps on the feeding coaxial cables.

Chapter 5

Transmitter System

Conventionally, clinical MRI scanners have only one proton transmit channel, possibly combined with a second broadband transmit channel for multi-nuclear capability. RF transmission is performed by a volume coil, which is designed to produce a homogeneous B_1^+ field over a region approximately the size of a beach ball at the isocentre of the scanner. High power transmitters capable of many kilowatts of pulsed power are necessary to achieve the high peak B_1^+ required to nutate the spin system by flip angles of 90° and 180° in a few ms. Increased RF absorption and wavelength associated field focusing effects [111] were not problematic until field strengths reached 3 Tesla, and so the power split and 90° phase shift employed to drive a birdcage coil in quadrature mode was achieved using a passive quadrature hybrid. At ultra high field strengths (7T+) RF inhomogeneity is so severe that new methods of driving transmitters have been required, and there has been a shift to using multi-channel transmitters, capable of RF shimming [116, 115, 228] and transmit SENSE [124, 87]. In the case of RF shimming multi-port/element volume coils transmit with different amplitudes and phases on each port/element to homogenise the RF field, and for transmit SENSE two or three dimensional excitation pulses are accelerated by transmitting independent temporal B_1 waveforms on each channel in a manner analogous to parallel imaging [186].

Separate-coil continuous arterial spin labelling (CASL) systems do not need to have the same requirements as an imaging transmitter system. Power specifications and pulse fidelity are much lower than that of imaging transmitters, so typically separate-coil CASL transmitter systems have consisted of a gated frequency source, a low power RF amplifier capable of continuous wave output, and associated PIN diode driver electronics[259, 213, 103]. Recently a low-cost programmable RF transmitter was devised by Xu et al.[247], incorporating a direct digital synthesiser (DDS) chip, amplitude modulation and a 400MHz DAC on a PCI bus mountable board, allowing all labelling parameters to be controlled remotely, rather than from the MR equipment room. Significant costs were saved by using the DDS chip, as opposed to a frequency generator such as a PTS 300

(Programmed Test Sources, Littleton, MA) and modulating RF signals using an analog mixer.

A clinical Siemens Trio (Erlangen, Germany) MRI scanner was available as the base system for this project. Hardware modifications were not possible, so the additional transmitter channels were implemented as a complete system which was synchronised to the Trio.

Parallel Transmission CASL has the following hardware requirements:

1. Independent amplitude and phase modulation of the driving current to each coil on a pulse-by-pulse basis.
2. Synchronisation between the Trio and transmitter channels for triggering pulses
3. Power amplifiers, capable of producing enough pulsed power for a $60° - 90°$ flip angle over the neck for B_1 mapping, and capable of continuous wave power for CASL.
4. Capable of RF spoiling.
5. Power monitoring.
6. Provide bias currents/voltages for PIN diode tuning/detuning of the coils.

These requirements are more relaxed than that of an imaging parallel transmit system, for which the power requirements are far greater and pulses must be perfectly synchronised. For example incorrect timing between a slice select sinc pulse and gradient could result in an imperfect slice profile, or improperly refocused slices. Incorrect RF pulse timing in gradient echo sequences will only result in an echo time which differs from the nominal, however for spin echo sequences this would result in an acquired spin echo offset from the centre of k-space.

This chapter details the equipment used to implement a two-channel parallel transmit system synchronised to the Trio. It is by no means the best way to construct a transmitter: most of the equipment was either surplus, sourced on a budget or constructed in-house with limited manufacturing facilities.

5.1 System overview

The transmitter system is based around a Surrey Medical Imaging Systems[1] (SMIS) spectrometer, which was previously used as the console for a 4.7T whole body MRI scanner [216, 215], and kept after the system was decommissioned in 2006. It is capable of up to four independent transmitter channels, each requiring its own PTS frequency synthesiser and proprietary SMIS modules. During the course of this project only two of the transmitter channels were fully operational, however if the appropriate SMIS modules can be sourced the system could be expanded to make use of all four. Each channel has full amplitude, phase and frequency control, and can use shaped waveforms, and

[1] SMIS was a UK company that manufactured PC based MR Scanners. It was bought by Marconi in 1999, which was subsequently bought by Philips in 2001.

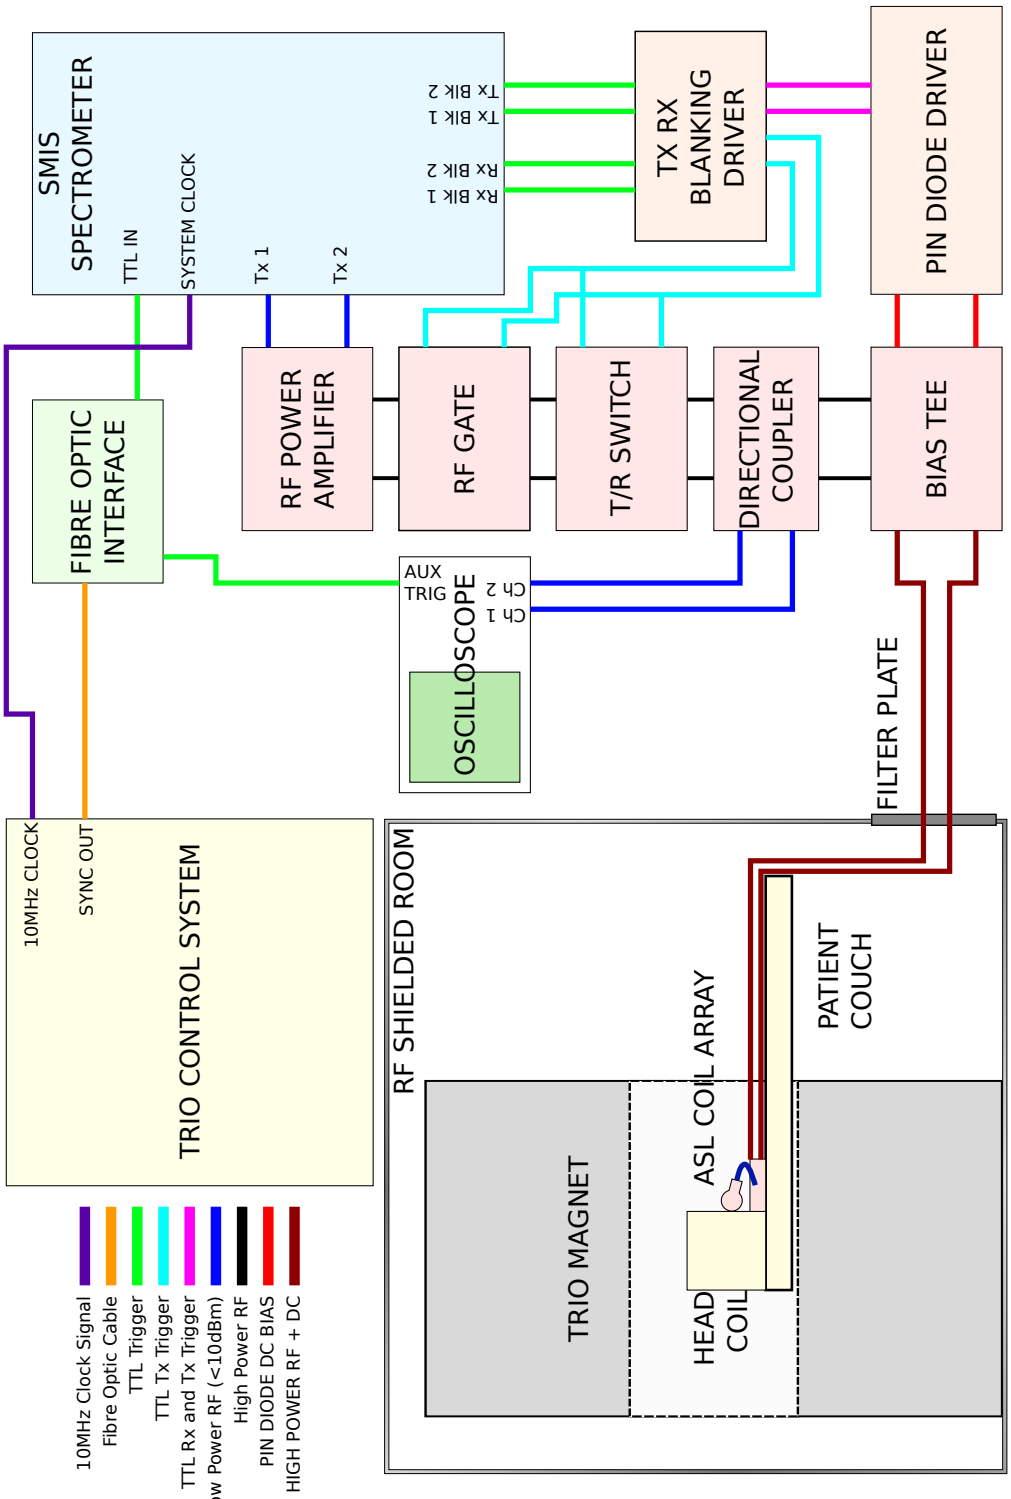

Figure 5.1: Transmitter system overview

Block colour: yellow–Siemens Trio system; green–interface between the Trio and SMIS systems; blue–SMIS spectrometer; red–RF signal chain; orange–non-RF control hardware.

in addition the spectrometer can be used to control additional hardware via Transistor-Transistor-Logic [83] (TTL) level digital inputs and outputs.

Figure 5.1 is a block diagram of the system components that make up the parallel transmit system, and its integration into the Trio scanner. Synchronisation with the Trio is achieved from a trigger pulse delivered by fibre optic cable, and a 10MHz clock signal to synchronise the system clock of the SMIS spectrometer and derive the RF signal used to transmit. The optical trigger pulse is converted to a TTL electrical signal via the fibre optic interface (section 5.3.1), which is connected to a digital input on the SMIS spectrometer, and also is used to trigger an oscilloscope (TDS410 150MHz Digitising Oscilloscope, Tektronix, Inc. Beaverton OR, USA). During operation a pulse sequence is executed by the SMIS spectrometer which waits for a trigger input, and when one is detected an RF pulse is played out on each channel, after which the pulse sequence returns to waiting for a trigger input. When a pulse is produced the spectrometer outputs an RF waveform and TTL blanking signal for each transmitting channel.

The low power (<10dBm; 10mW) RF waveform is fed to the input of RF power amplifiers (ZHL-3A, Mini-Circuits Inc. Brooklyn, NY, USA) where it is amplified to approximately 30dBm (1W). The RF power amplifiers do not have any noise blanking, so they are gated by a custom-built RF gate and transmit/receive (T/R) switch (PulseTeq Ltd, Wotton-Under-Edge, Gloucestershire, UK). Following that, the RF passes through directional couplers with -40dB of coupling, to allow the RF power to be monitored using the oscilloscope without significantly perturbing the output power. DC bias for the PIN diodes on the coils is added to the RF line via a custom built bias-tee (section 5.4.1), then the RF is brought to the magnet room filter plate where there are bulkhead BNC sockets.

Blanking signals are used to control the RF gates, T/R switches, and PIN diode driver. The Tx Rx Blanking Driver serves two purposes: to buffer the output of the blanking signal from the SMIS spectrometer, ensuring 0V low and 5V high signals when several devices are connected, and to control the RF gates and T/R switches only during transmit, but the PIN diode driver both during transmit and receive. While the ASL Labelling Coil Array has not been used as a receiver, this capability may come useful in the future.

5.2 SMIS Spectrometer

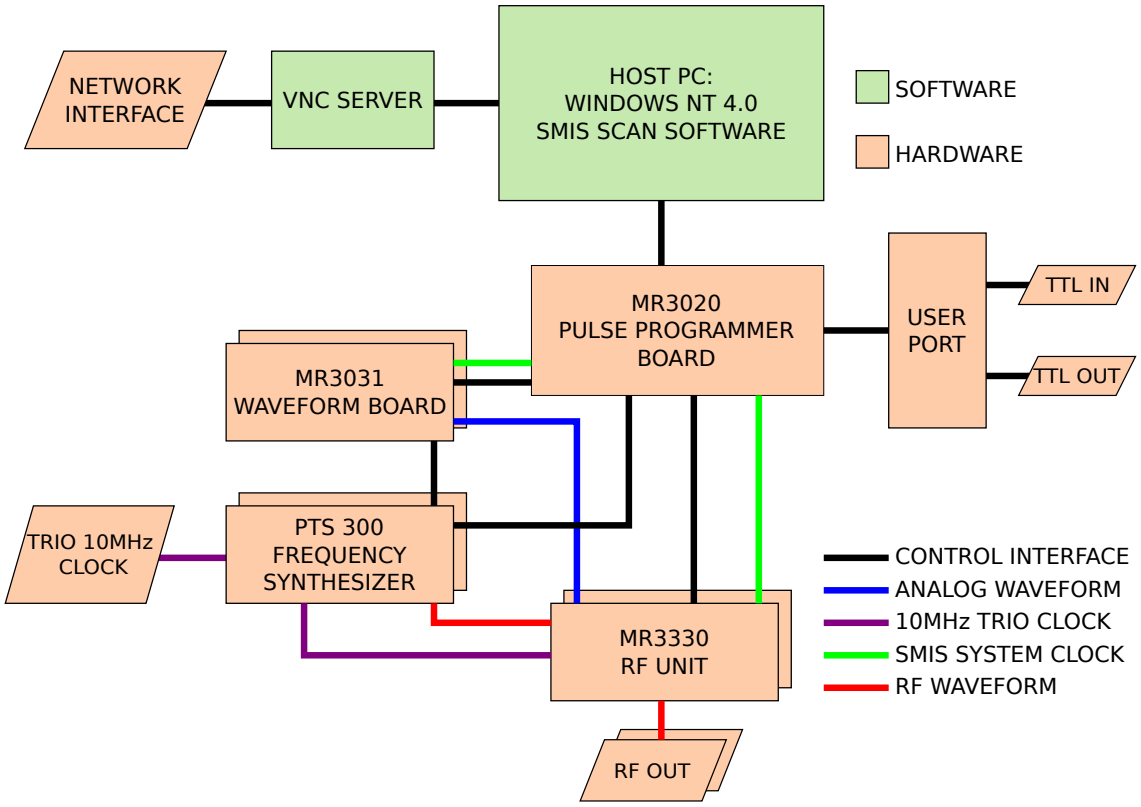

Figure 5.2: SMIS spectrometer system diagram

As detailed in section 5.1 the SMIS spectrometer was once used as the control system for a 4.7T whole body MRI scanner, providing control of RF, gradients, shims, and data acquisition. A single computer running Windows NT 4.0 (Microsoft Corporation, Redmond, WA, USA) runs the SMIS scan software and is mounted within a 19" rack mounted chassis, which also contains the MR3020 pulse programmer board, MR3031 waveform generator board, as well as other boards used for gradient control and data reconstruction. A compiled sequence is loaded onto the MR3020 board, which then controls all other hardware accordingly.

Generation of the radio frequency signal is performed by a dual channel PTS 300 (Programmed Test Sources, Littleton, MA, USA), producing an independent and stable frequency locked to a 10MHz source signal. As a standalone system the PTS300 has a built-in 10MHz crystal oscillator, or it can accept an external 10MHz sinusoidal signal, as was used to synchronise with the Trio. Frequency and phase control is performed by sending binary-coded-decimal [2] (BCD) data over a parallel interface. Frequencies can be specified to 0.1Hz and phase control to 0.225°. The MR3020 is used to set the base frequency and phase of the PTS on a pulse-by-pulse basis (for example exciting multiple slices in a single experiment, each will require a different frequency during both transmit

[2]BCD encodes each decimal in a number as a four-bit binary sequence, each of these sequences is called a nibble. For example to switch the frequency of the PTS to 123.2MHz the number 123200000 is sent as eight nibbles corresponding to:

Decimal	1	2	3	2	0	0	0	0	0
Binary	0001	0010	0011	0010	0000	0000	0000	0000	0000

and acquisition). The MR3031 can rapidly modulate the base frequency and phase provided by the MR3020, for example during adiabatic pulses.

The same 10MHz clock signal used by the PTS (external or from the internal oscillator) is also used by the MR3330 RF unit to generate a TTL level square wave system clock, from which all hardware is synchronised and timings calculated. Amplitude modulation of the RF signal is performed by the MR3330, using either an analog waveform generated by the MR3031, or digital control signal to modulate a PIN diode attenuator and analogue multiplier.

For each transmitter channel a MR3031 waveform generator, PTS frequency synthesiser and MR3330 RF unit is required. All are interfaced to the MR3020 pulse programmer board, which executes the pulse sequence, controls all other hardware and and communicates with the host PC. The user port provides eight TTL level outputs and four TTL level inputs, which were used for synchronisation and additional hardware control.

5.2.1 Pulse Sequence

A pulse sequence was developed to produce RF pulses on each channel and allow independent modulation of their amplitude and phase. SMIS pulse sequences are written in a language similar to C, then compiled to a format that can be executed by the MR3020 board. Pulse sequence parameters can be specified to be placed in the dual port RAM of the MR3020, which means that the host scanning software also has access to these variables during the execution of a sequence. This is used to effect in the "interactive setup" mode of the scanning software, where a pulse sequence is repeated indefinitely to allow for fine-tuning of the parameters. Typically this would have been used for power calibration of the RF pulses. The pulse sequence was written so that it could be used in "setup" mode and parameters adjusted on-the-fly without stopping and re-starting the sequence.

Two types of sequences were considered when the pulse sequence was devised: first a FLASH [92] style imaging sequence, used, for example, to acquire B_1 maps (see section 7.2), and consisting of short hard pulses repeated as often as every 20ms; and second a CASL sequence (see section 7.3) consisting of continuous pulses lasting several seconds, repeated every 5-10s. Figure 5.3 is a flow chart of the operation of the pulse sequence. Upon initialisation of the pulse sequence the scan parameters are obtained from the host computer and the hardware is initialised. The program then ensures that the input port is low before proceeding, this is because the user port's input will go to high if nothing is connected (there are internal pull-up resistors) resulting in false triggering of the RF. For a CASL sequence this would be dangerous as the sequence is set to transmit for as long as the input is high, so by halting the program until the input goes low this can never happen.

Once the user input is low, the main part of the program commences. The phase of each channel has two components - "static" phase, which is set by the user and does not change on a

pulse-by-pulse basis, and "dynamic" phase, which is used for RF spoiling and does change on a pulse-by-pulse basis. Two types of RF spoiling are available, basic RF spoiling where the spoil angle increments by the same amount each time, and "alternating" RF spoiling, where the spoil increment alternates between two spoil angles, a requirement for the formation of a steady state in an Actual Flip angle Imaging B_1 mapping sequence (section 7.2.1 and [162]). If the sequence is set to continuous wave (CW) mode one of the the user output ports is set to high to provide a CW enable signal required for some RF power amplifiers. The sequence then waits for a trigger pulse that lasts for longer than a specified amount of time ($10\mu s$ was found to provide stable results), and once triggered plays out a blanking pulse and RF waveform on each channel that has its power level greater than 0 (the power level can be set to 0, 1, 2 or 3, with 0 being off and 3 the most power). If CW mode is selected the pulses continue until the user input returns to low, otherwise the pulses last a specified duration. The pulse sequence then repeats, recalculating new parameters and preparing the hardware for another pulse once triggered.

It is convenient and time saving to run the pulse sequence in "interactive setup" mode and make changes to parameters as the sequence executes: the procedure of aborting the pulse sequence, changing parameters and running the sequence again takes at least thirty seconds and would unacceptably amount to many minutes of dead time in a typical scan session. In order to achieve this the dynamic component of the phase of each transmitter channel needs to be re-set accordingly when a sequence has ended, otherwise the transmitter will not be synchronised with the receiver of the Trio and images will become distorted. Additional conditions have been implemented in the period where the sequence waits for a trigger pulse that if met will reinitialise the sequence, resetting the dynamic component of the phase and the spoil angle to zero. This occurs if the time waiting for a trigger pulse is longer than a specified time (for example ten seconds), and if a reset required flag variable is true (set to be true if a pulse has run, and false once a reset has occurred).

The time taken between the generation of a 5V trigger pulse and RF appearing at the output of each channel is approximately $200\mu s$, and there is jitter on the order of 5ns in the starting time of the pulse. For a gradient echo sequence this results in an echo time that differs slightly from the nominal value, however this can be corrected by triggering $200\mu s$ in advance.

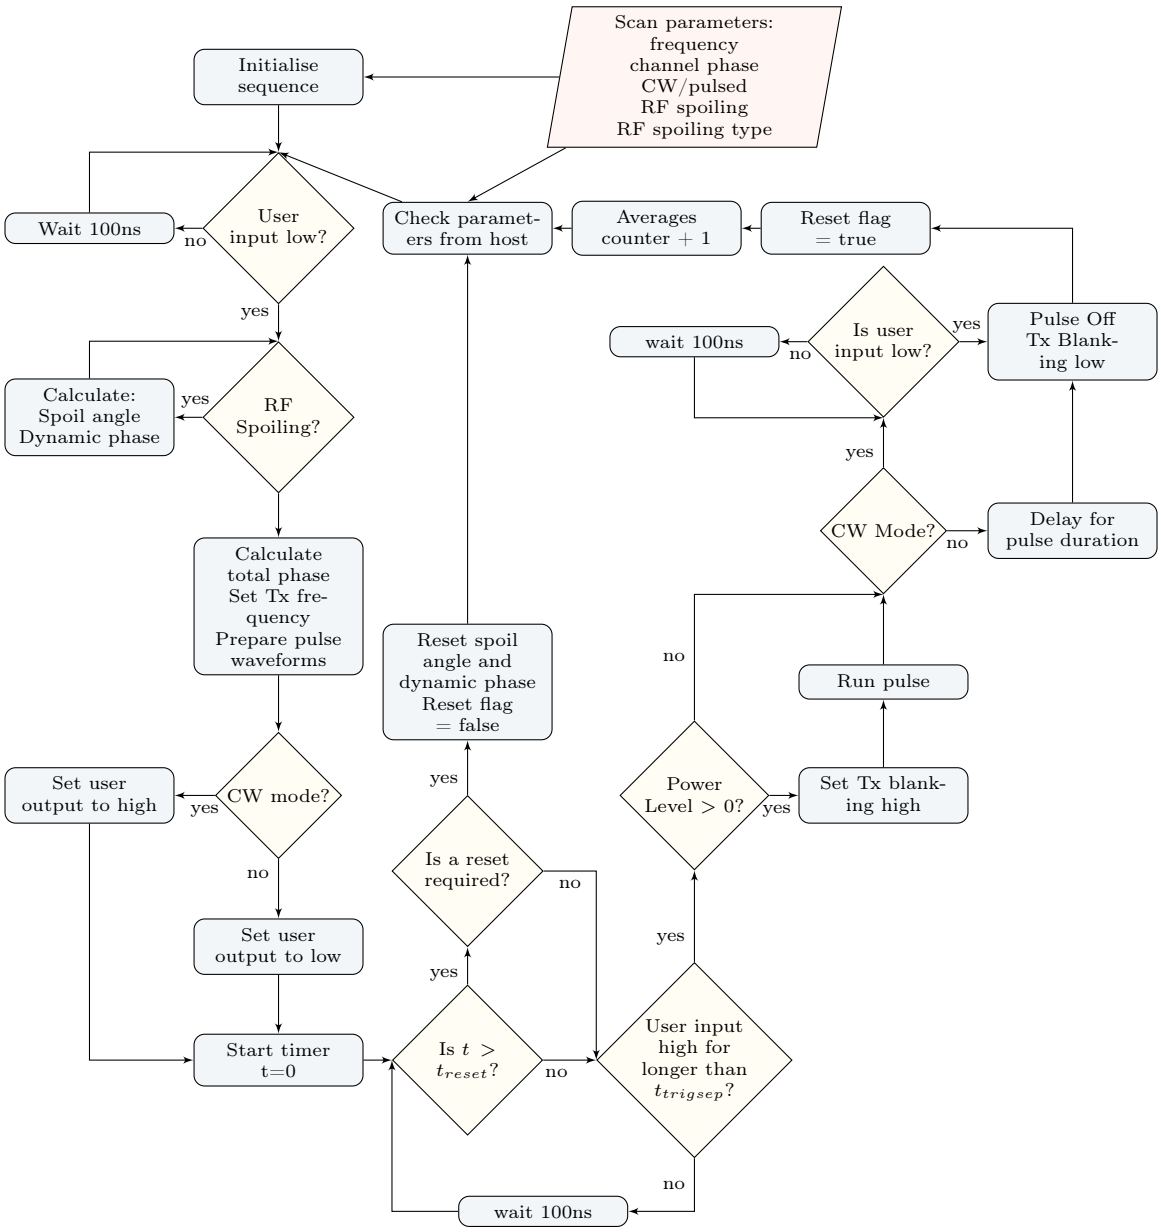

Figure 5.3: SMIS control pulse sequence

5.3 Trio Synchronisation

5.3.1 Fibre Optic Interface

The optical synchronisation signal provided by the Trio was converted into a TTL level signal by the fibre optic interface. Essentially this consisted of a photo-sensor which is switched with the light from the Trio's signal, however to ensure there was no false triggering due to ground loops the TTL output was electrically isolated from the power supply to the fibre optic interface by using an opto-coupler (KB817, Kingbright, Taipei, Taiwan) and isolated DC-DC power supply (NMA0515S, Murata Power Solutions Inc., Mansfield, MA, USA).

The SMIS spectrometer's user port has four TTL level inputs. Each has a pull-up resistor

connected between the input and the 5V power rail. In order to produce a low state a low impedance path from the input to ground must be created, and this is achieved when the MOSFET Q1 (IRLML2803, International Rectifier, El Segundo, CA, USA) is switched on. Three LED's are present to indicate power (D1), optical data (D2) and a TTL data output (D3).

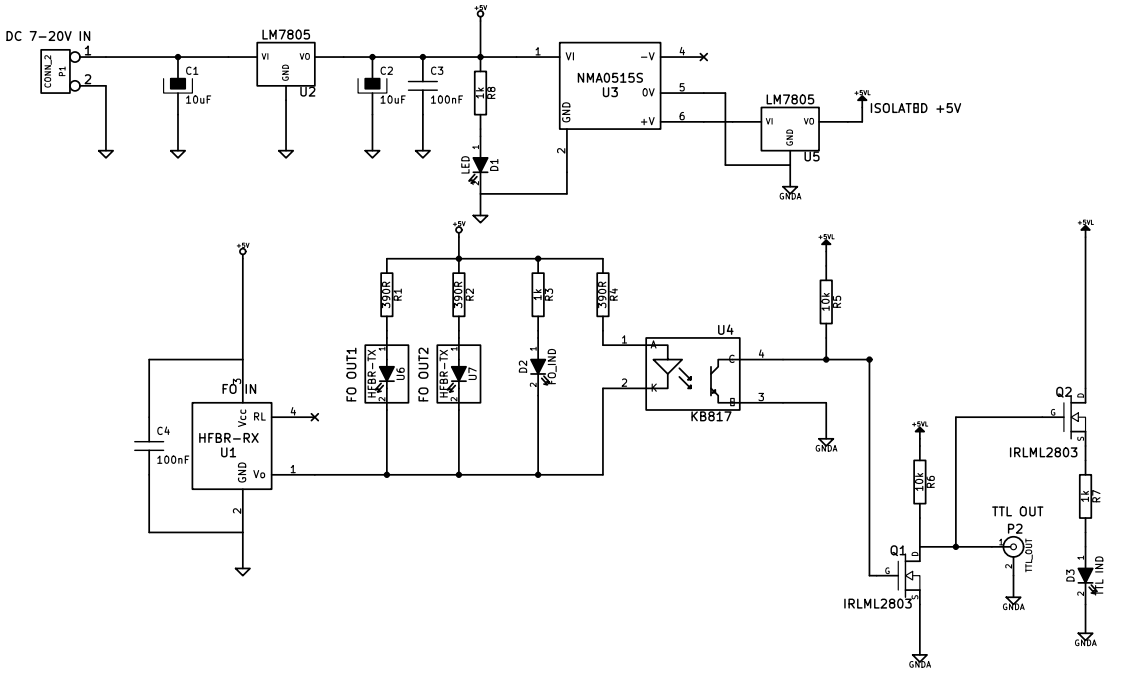

Figure 5.4: Fibre optic interface schematic
The fibre optic photo-diode receiver U1 has an open emitter transistor output which switches on four LEDs - two in corresponding fibre optic transmitters TX1 and TX2, a data indicator LED D2 and the LED in opto-coupler U4. U3 is an isolated DC to DC power supply and U5, a voltage regulator provides an isolated 5V and ground for the TTL output circuit. As such, ground loops between the Fibre Optic Interface and SMIS spectrometer were completely avoided.

5.3.2 System clock synchronisation

RF frequency synchronisation was found to be essential to ensure the frequency generated by the SMIS console exactly matches the Trio, and was achieved by using the 10MHz clock output of the Trio as the 10MHz reference signal that the SMIS derived its master clock and RF frequencies from. When transmitting in a FLASH style imaging sequence, if the transmitter and receiver frequencies are not identical, consecutive lines of k-space acquired an interval $t = TR$ apart exhibit a phase shift $\Delta\phi = \frac{2\pi\Delta f}{TR}$, amounting to a linear phase ramp in the phase encoding direction(s) of k-space, and a shift in image space. When the Trio and SMIS were not synchronised, but set to the same frequency a shift in images was seen, and due to different temporal variations in the two 10MHz clocks this shift was different for every acquisition. Synchronising the SMIS to the Trio alleviated this problem.

5.4 PIN Diode Driver

A custom PIN diode driver was built to could detune each coil independently. Each PIN diode requires:

- 100mA of current when forward biased so that the diode has a low resistance.
- 5V reverse bias to prevent the electro-motive-force (EMF) induced when the body coil is transmitting from forward biasing the diode and tuning the ASL coil. This would intensify the body coil's B_1 field in the vicinity of the ASL coil, increasing locally deposited power.

Figure 5.5 is a schematic of the PIN diode driver. A TTL level input signal is used to trigger the driver, which is configured so that the PIN diode is reverse biased when the input is low, and forward biased when the input is high, making it compatible with the amplifier gating signal generated for each channel by the SMIS console. Buffered status LEDs indicate when there is a TTL signal and the bias state (forward or reverse) of each coil.

A PIN diode driver module was required for each transmitter channel, so a PCB was designed to aid reproducibility. Both modules shared the same dual rail ±5V power supply. The PCBs were mounted in a 2U 19" chassis as shown by figure 5.6.

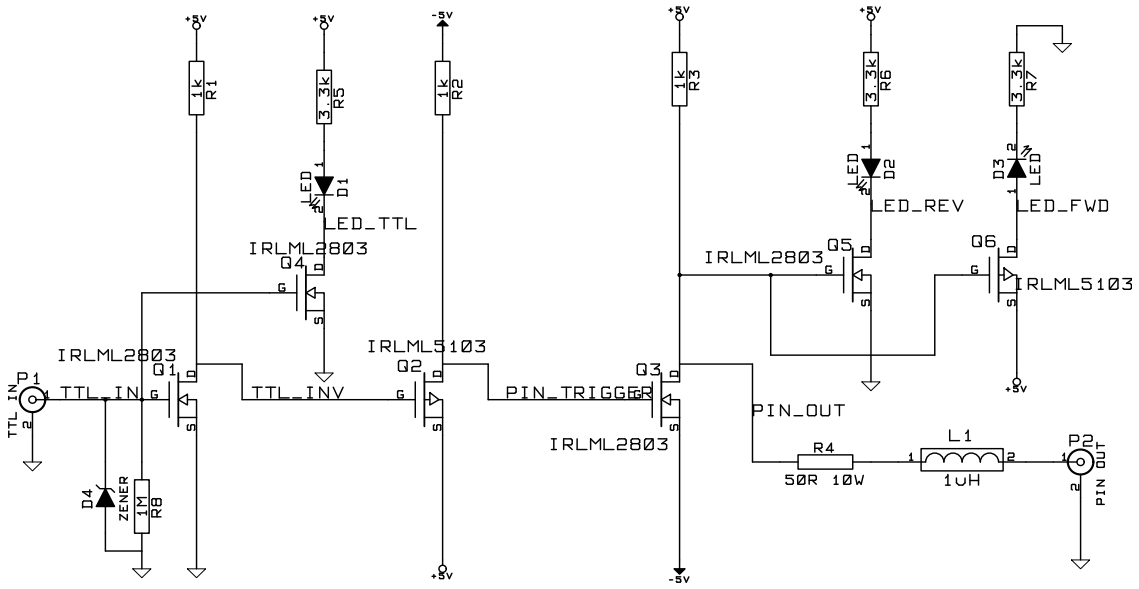

Figure 5.5: PIN diode driver schematic
When the TTL input is low Q3 is off and the output at P2 (PIN_OUT) is +5V. When the input is high Q3 switches on and sinks current limited by R4. As 100mA is required R4 is chosen to be 50Ω. The PIN diode must be connected such that the anode is connected to ground and the cathode to the output of the PIN diode driver.

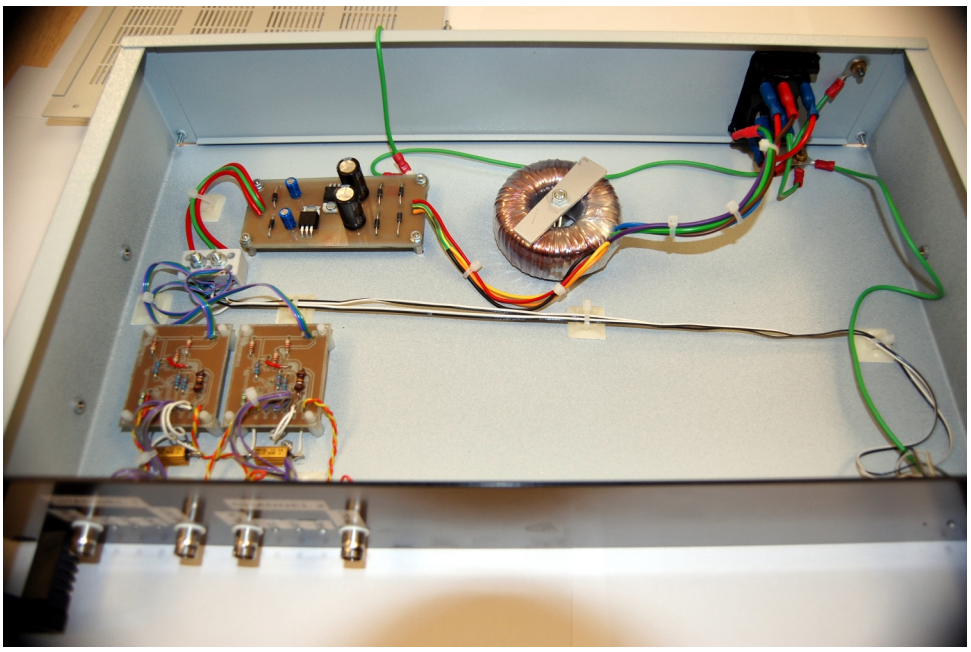

Figure 5.6: Complete PIN diode driver

5.4.1 Bias Tee

Bias currents/voltages can be fed to the PIN diodes on separate cables, however it is more practical for DC power to run down the coaxial cable simultaneously with the RF power. A bias tee is used to separate/combine RF and DC signals. In its simplest form it consists of a series capacitor sufficiently large to have minimal impedance at the frequency of interest and an inductor large with a high impedance at the frequency of interest. RF is blocked from travelling down the DC path by the inductor, and DC is blocked from travelling down the RF path by the capacitor. Values were chosen so that at 123.2MHz the capacitor's reactance is less than 1Ω, and the inductor's reactance is greater than $1k\Omega$:

$$C = 2nF \quad X_c = 0.65\Omega$$

$$L = 3\mu H \quad X_L = 2.32k\Omega$$

A real inductor has an additional capacitive component to its reactance, resonating with the inductance at a certain frequency, known as the self resonant frequency (SRF). When selecting an inductor as an RF choke (to block RF) it is essential that the SRF is higher than the operating frequency, otherwise the inductor will behave as a capacitor and no longer block RF. By selecting an inductor with a SRF just above the operating frequency (for example a 150MHz SRF inductor for use at 123MHz) the total reactance will be much higher than from a pure inductance alone. Figure 5.7 shows the impedance of a real $3\mu H$ inductor and ideal $3\mu H$ inductor as a function of frequency. The SRF of the real inductor was 150MHz, DC resistance 0.5Ω, and quality factor 5.

At 123.2MHz the blocking impedance of the real inductor is nearly three times that of the real inductor.

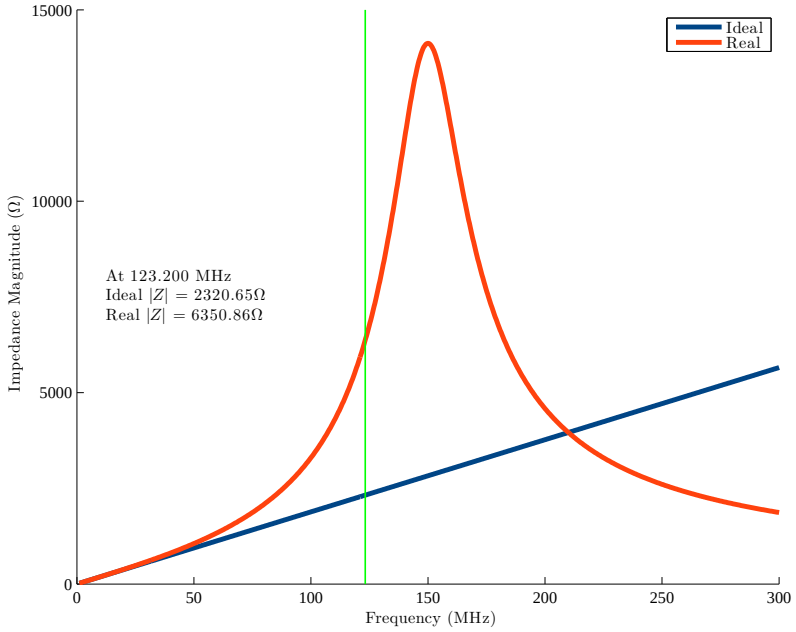

Figure 5.7: Impedance Magnitude Plot of an RF Choke

The coplanar waveguide PCB design ensures that 50Ω characteristic impedance is maintained for the entire RF path so the signal is not reflected, reducing transmit efficiency and potentially damaging the RF Power Amplifiers.

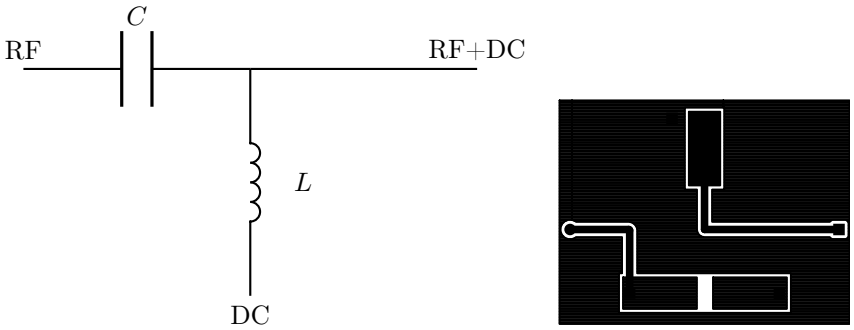

Figure 5.8: Basic bias tee circuit and PCB

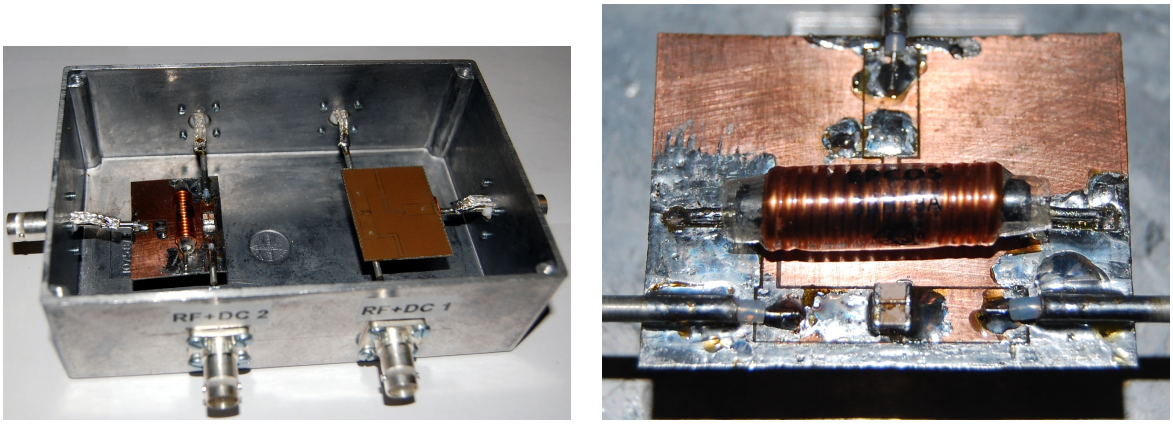

Figure 5.9: Images of the bias tee

It is important to maintain a characteristic impedance of 50Ω in the RF path, so the bias tee was built using a coplanar waveguide[205] PCB. Two bias tee's were built and mounted in an aluminium enclosure (see figure 5.9). Table 5.2 shows the performance of the bias tee. There is a low insertion loss of 0.06dB and 0.07dB between the RF and RF+DC ports for bias tee's 1 and 2 respectively, which shows that there is minimal effect on the RF current from the capacitors or coplanar waveguide. Between the DC port and the RF or RF+DC port there is an attenuation of 23.5dB and 26.1dB for each bias tee. While this is not ideal (commercial bias tee's can provide at much as 90dB of attenuation), it is sufficient for use here as the RF power will be low and the PIN diode driver has RF chokes on its output to further block RF. The reason for the poor blocking impedance could be due to the groundplane of the coplanar waveguide PCB surrounding the DC pad and coupling RF onto the DC line. A solution to this would be to re-design the PCB so that the ground plane only surrounds the RF path. Both bias tee's are well isolated from each other (approximately 70dB), as indicated by the RF1 – RF2, RF2 – DC1, DC1 – DC2, DC2 – RF+DC1 and RF+DC1 – RF+DC2 measurements.

	RF1	RF2	DC1	DC2	RF+DC1	RF+DC2
RF1	X	-71.09	-23.29	-70.49	-0.06	-69.39
RF2		X	-73.59	-25.82	-69.37	-0.07
DC1			X	-71.59	-23.69	-73.79
DC2				X	-70.49	-26.29
RF+DC1					X	-71.59
RF+DC2						X

Table 5.2: Bias Tee Characterisation
Through loss as measured by the network analyser. Ports 1 and 2 of the network analyser were connected to the corresponding ports on the bias tee and the value of S_{21} measured. A baseline measurement of the two coaxial cables connected with a BNC coupler was also taken (-0.607dB) and subtracted from all the results.

5.5 RF Signal Path

5.5.1 RF power amplifiers

The SMIS console delivers up to 10dBm (10mW) of RF power for each channel, which must be amplified further to produce the required B_1 field. Low power instrumentation/communication RF amplifiers (ZHL-3A, Mini-Circuits Inc. Brooklyn, NY, USA) were employed, providing 30dB of gain (1000× amplification factor), and a rated power output of 30dBm (1W) at 1dB compression, an IP3 point[3][29] of 38dBm and an operational bandwidth between 0.4 and 150MHz. Each amplifier was powered by a 24V, 40W DC power supply (ECM40US24, XP Power Limited, Singapore).

5.5.2 Reducing power amplifier noise

Thermal noise at the input of a RF power amplifier will be amplified to levels which are problematic in MRI. The weak nature of the NMR signal demands an extremely sensitive receiver system, which can be easily saturated from external signals at the frequency of interest, resulting in "spikes" in an image. It is for this reason that the MRI scanner is placed within a shielded room. Typically RF power amplifiers used in MRI have some form of blanking that gates the amplifier output to only when a pulse is being transmitted, and if employed a T/R switch provides additional isolation. No blanking circuitry was present on the ZHL-3A amplifiers, so to reduce the noise injected into the scan room custom RF gates were built and used in conjunction with T/R switches.

5.5.2.1 T/R Switches

Each T/R switch came as a complete pre-built module designed to operate at 123.2MHz. A 2V, 100mA bias voltage was supplied by TTL controlled n-channel MOSFET driver circuit. Table 5.3 shows the results of network analyser measurements of the RF power transmitted from the transmitter port to the coil port of each T/R switch at 123.2MHz. Measurements were made with the receiver port of the T/R switches terminated with a 50Ω load.

[3]The third order intercept point, or IP3 point is a measure of non-linearity in weakly non-linear devices, such as linear RF amplifiers. A polynomial is used to model the non-linearity, i.e. $V_{out} = a_1 V_{in} + a_2 V^2 + a_3 V^3 + ...$ The output power that results in the magnitude of the first order and third order components being equal is said to the be the IP3 point.

5.5.2.2 Transmitter Gate

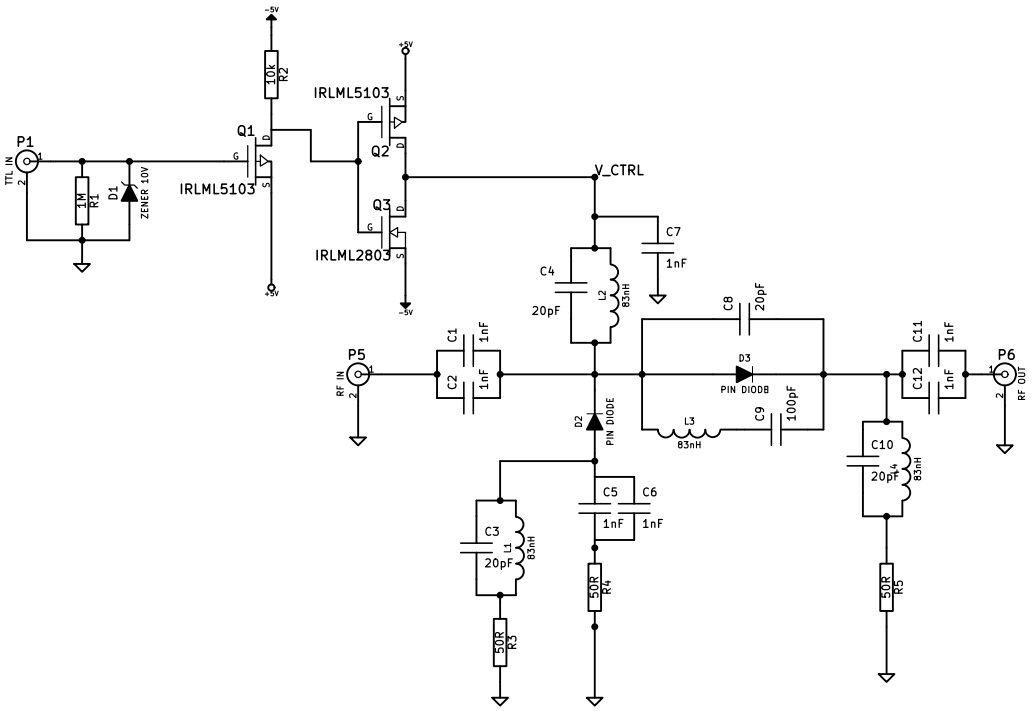

Figure 5.10: Transmitter gate schematic diagram

Figure 5.10 is a circuit diagram of the transmitter gate. It is based around two PIN diodes (PPD1200, Passive Plus Inc. Huntingdon Station, NY, USA), in a series-shunt arrangement such that when one is forward biased the other is reverse biased. The PPD1200 has a resistance of 0.3Ω when forward biased with 10mA of current, and a capacitance of 3pF when reverse biased by 50V. As the RF power that the gates carry is low the reverse bias voltage did not need to be this high, and a ±5V power supply was used. Forward bias current was decided to be 100mA, resulting in a measured resistance of 0.06Ω at 123.259MHz, and the reverse bias capacitance at 5V measured to be 3.7pF.

	S_{21}(dB)	
	Blanked	Transmit
T/R Switch 1	-34.5	-1.1
T/R Switch 2	-37.6	-1.2

Table 5.3: Measured transmission between transmitter and coil ports of each T/R switch at 123.2MHz

In transmit mode D3 is forward biased, providing a low impedance path between the RF IN and RF OUT ports, and D2 is reverse biased, isolating the RF from the ground by the reactance of 3.7pF at 123.2MHz, approximately 349Ω. During blanking D2 is forward biased, shunting the RF signal to a 50Ω load R4, and D3 is reverse biased, presenting a similar 349Ω blocking reactance. To further improve isolation a parallel resonant circuit was included around D3. The reverse bias capacitance of D3 and C8 are resonated by L3, creating a high blocking impedance. C9 is a DC blocking capacitor as the inductor would otherwise allow DC current to flow, bypassing D3.

Capacitor pairs C1 and C2, C5 and C6, and C11 and C12 are DC blocking capacitors to prevent

the DC bias being fed to the connecting coaxial cables, and through the 50Ω RF load. The parallel resonant circuits formed by C3 and L1, C4 and L2, and C10 and L4 block the RF signal and provide DC bias for the PIN diodes, during construction of the gates each was found to provide at least $13k\Omega$ of blocking impedance.

Control of the gate is with a TTL signal (0V during blanking, 5V in transmit), which is level shifted by Q1 so that it will correctly switch Q2 and Q3 to provide the correct DC bias of -5V during blanking and +5V during transmit. Bias currents are regulated by the 50Ω resistors, R3 and R5.

A double sided PCB was designed where the RF traces were implemented as 3.1mm microstrip transmission lines to ensure 50Ω characteristic impedance. This was manufactured from 1.6mm thick double sided copper clad FR-4, copper plated to $70\mu m$ thick to improve conductivity, then CNC milled. All inductors were hand-wound from 1.25mm diameter enamelled copper wire, capacitors were 500V rated, high Q, P90 dielectric ceramic chip capacitors (1111P series, Passive Plus Inc. Huntingdon Station, NY, USA), and R4 a 40W power RF chip resistor (CHF2525CNT500LW, Bourns Inc., Riverside, CA, USA). Each PCB was housed in a die-cast aluminium box with flanged BNC connectors for the TTL and RF signals, and feed-through filters for DC power. Figure 5.11 shows an assembled PCB within its aluminium box.

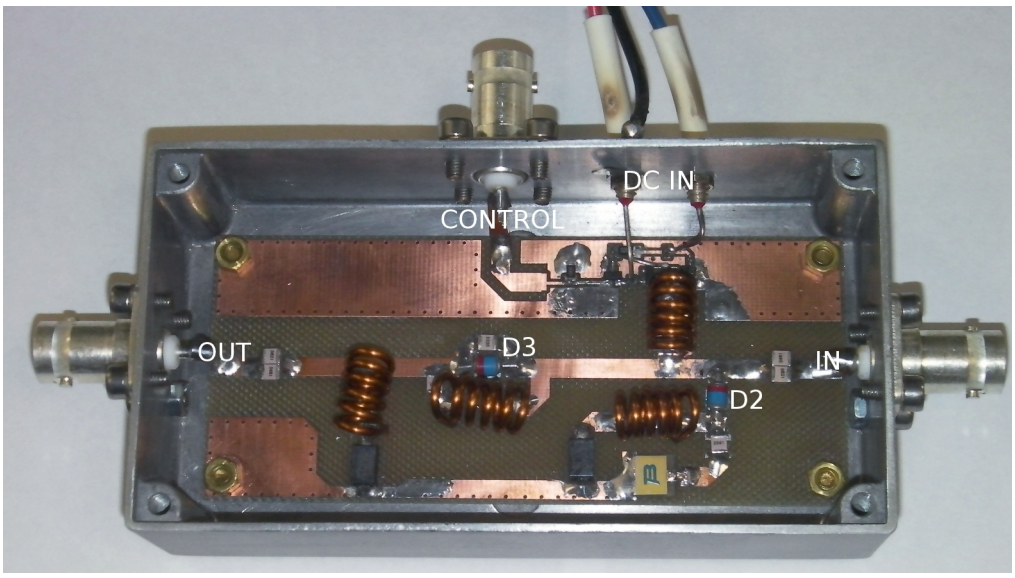

Figure 5.11: Complete Transmitter Gate

		S_{11} (dB)	Z (Ω)	S_{21} (dB)
Gate 1	Blank	-21.2	46.5 - j7.6	-37.4
	Transmit	-36.6	49.6 - j0.2	-0.8
Gate 2	Blank	-21.9	45.7 - j6.4	-37.3
	Transmit	-40.9	49.2 - j0.4	-0.8

Table 5.4: Reflection and Transmission measurements of the Transmitter Gates at 123.259MHz

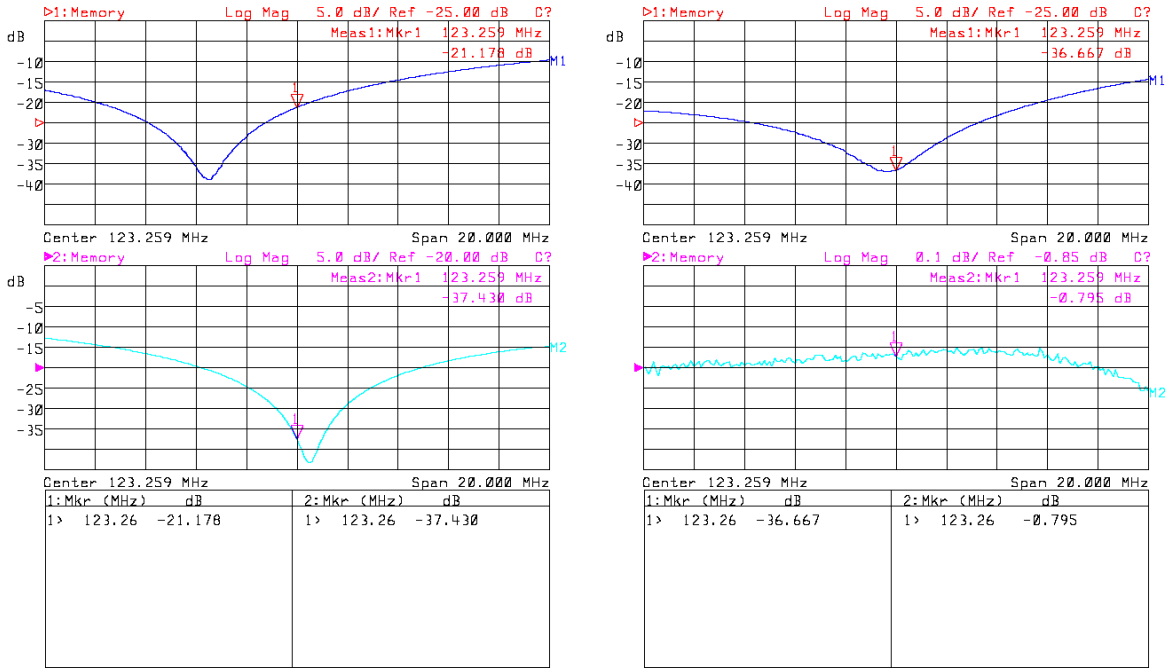

a. Gate 1: Blank b. Gate 1: Transmit

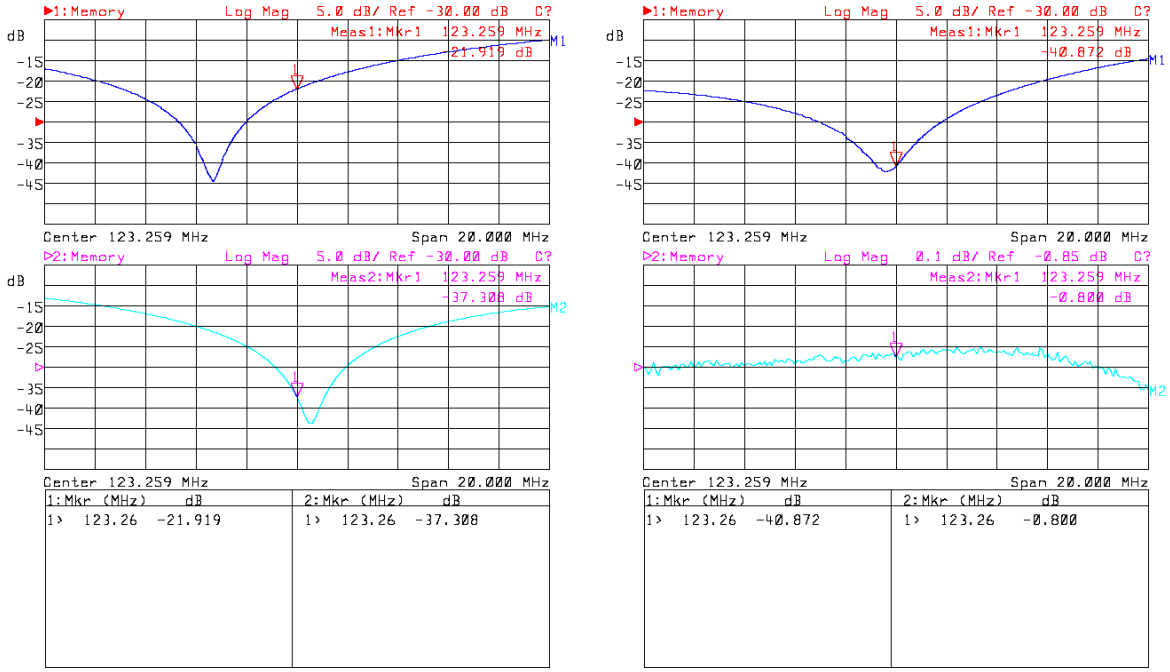

c. Gate 2: Blank d. Gate 2: Transmit

Figure 5.12: Network Analyser Plots of the Transmitter Gates
Blue Trace - Reflection (S_{11})
Turquoise Trace - Transmission (S_{21})

Figure 5.12 shows the network analyser traces for each of the two transmitter gates and table 5.4 shows the values measured at 123.259MHz. Traces are over a 20MHz sweep centred at 123.259MHz, with RF power set to 15dBm. In blanking mode both gates show approximately 37dB of isolation between the input and output ports, and the input port is reasonably close to 50Ω. A matching network prior to R4 would improve this , however during blanking no RF power is applied at the input of the amplifiers, so any reflected signal (consisting only of amplified noise) will be small and will not cause any damage to the amplifiers. Figure 5.12.c and d both show that the frequency

where the isolation is greatest is slightly higher in frequency than 123.259MHz. This is due to the additional reactance of the coaxial cables shifting the frequency of the parallel resonant trap circuit around D3, during construction each trap was tuned in isolation to ensure it had a pure resistive impedance that peaked at 123.259MHz, in all cases exceeding $9k\Omega$. In transmit mode both gates have low insertion loss and are closely matched to 50Ω.

5.5.2.3 Complete Amplifier unit

Amplifiers, power supplies, T/R switches and switching electronics to provide 2V bias for the T/R switches were housed within a 2U 19" rack chassis, as shown in figure 5.13. Due to insufficient space within the chassis transmitter gates were mounted on the top.

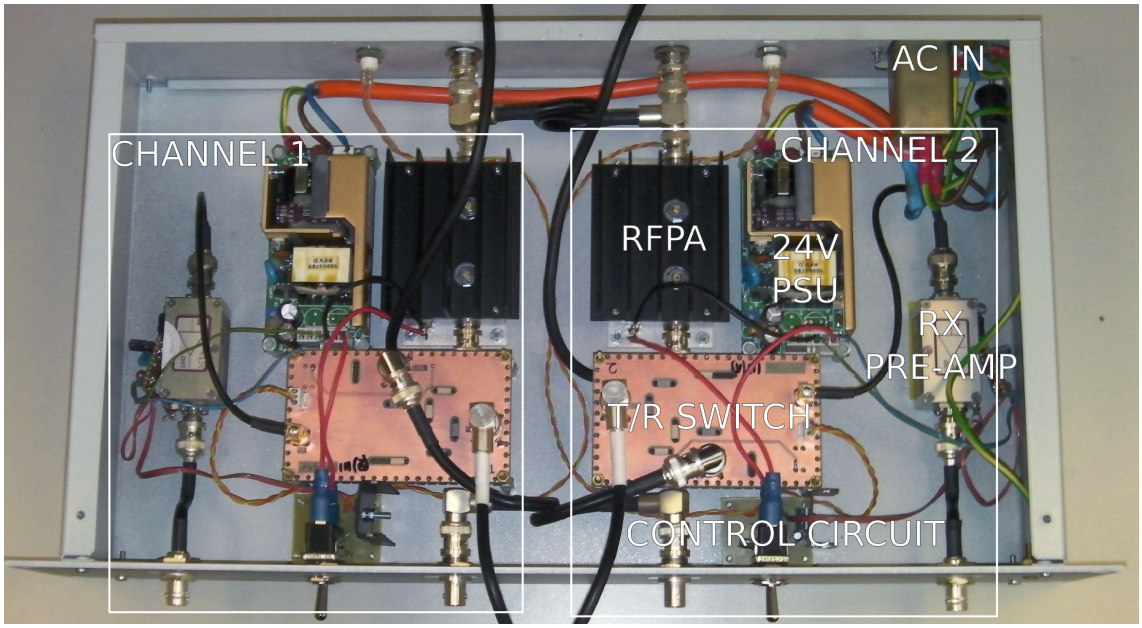

Figure 5.13: Inside the Amplifier Unit

5.5.3 Power monitoring

Monitoring the forward and reverse power delivered to the coils is a simple way initially characterise the system, and to ensure that it is operating to specification. A dual directional coupler (RF Power Components DDC-2112-R5-40) with nominally 40dB of coupling was used to sample the forward and reverse power of each channel, and for measurement the respective port was connected via a RG223 cable with BNC connectors to an oscilloscope (TDS410 150MHz Digitising Oscilloscope, Tektronix, Inc. Beaverton OR, USA), which is set to DC coupling, 50Ω input impedance and to automatically measure the peak to peak voltage, V_{pk}. Root mean square (RMS) power (P_{RMS}) through the directional coupler was then calculated with the following equation:

$$P_{RMS} = \frac{V_{pk}^2}{8R_{sc}} 10^{C/10} \tag{5.1}$$

Where C is the coupling factor of the directional coupler in dB (for example 40dB), R_{sc} is the

input impedance of the oscilloscope. A schematic diagram of the dual directional coupler is shown in figure 5.14. It is inserted in series with an RF line such that the RF power travels from the in port to the out port. Forward power is coupled to the forward port, and reverse power is coupled to the reverse port, the ratio of the through power to the coupled power is given by the coupling factor[4].

The forward and reverse coupling factor the directional coupler was measured at 123.259MHz to improve the accuracy of power measurements. An absolute power measurement was made with the network analyser of the coupled power to the forward and reverse ports. Output power was set to 10dBm and the loss of the cables was measured by measuring the transmitted power through only the connecting cables. To measure the forward power, power was applied to the in port of the directional coupler, the out port and reverse ports were terminated with 50Ω loads and the power at the forward port was measured. For a reverse power measurement power was applied at the out port (simulating reflected power), the in and forward ports were terminated with a 50Ω load and the coupled power was measured from the reverse port of the directional coupler. The coupling factor was then calculated by subtracting the coupled power from the transmitted power through the connecting cables. Table 5.5 has the results of these measurements.

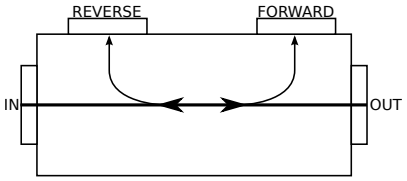

Figure 5.14: Dual Directional Coupler Schematic

	C_{fwd}(dB)	C_{rev}(dB)
Channel 1	41.4	41.1
Channel 2	41.9	42.3

Table 5.5: Forward and Reverse Coupling Factors of the Directional Couplers.

Forward and reverse power of each transmitter channel through a 50Ω load was measured using the described setup to characterise the output power with respect to RF channel amplitude in the SMIS scan software. Between the directional coupler and the load the bias tee see (section 5.4.1) was connected. RF amplitude was swept from 10% to 100% in 10% increments, and the peak to peak voltage at the oscilloscope was recorded. Using equation 5.1 and the measured coupling coefficients in table 5.5 the forward and reflected power was calculated. Graphs of the measured voltage and calculated power are shown in figure 5.15.

[4]An individual directional coupler has four ports: in, out, coupled and isolated. The isolated port must be terminated with the characteristic impedance of the line. The coupled and isolated ports can be swapped; their configuration determines the directivity of the coupler (forward or reverse). There exists a coupling factor between the coupled port and both the in port and out port. The ratio of incident power at port P_i to the coupled power at the coupled port, P_c is given by the coupling factor, expressed in dB: $C_{i,c} = -10log(\frac{P_c}{Pi})$.

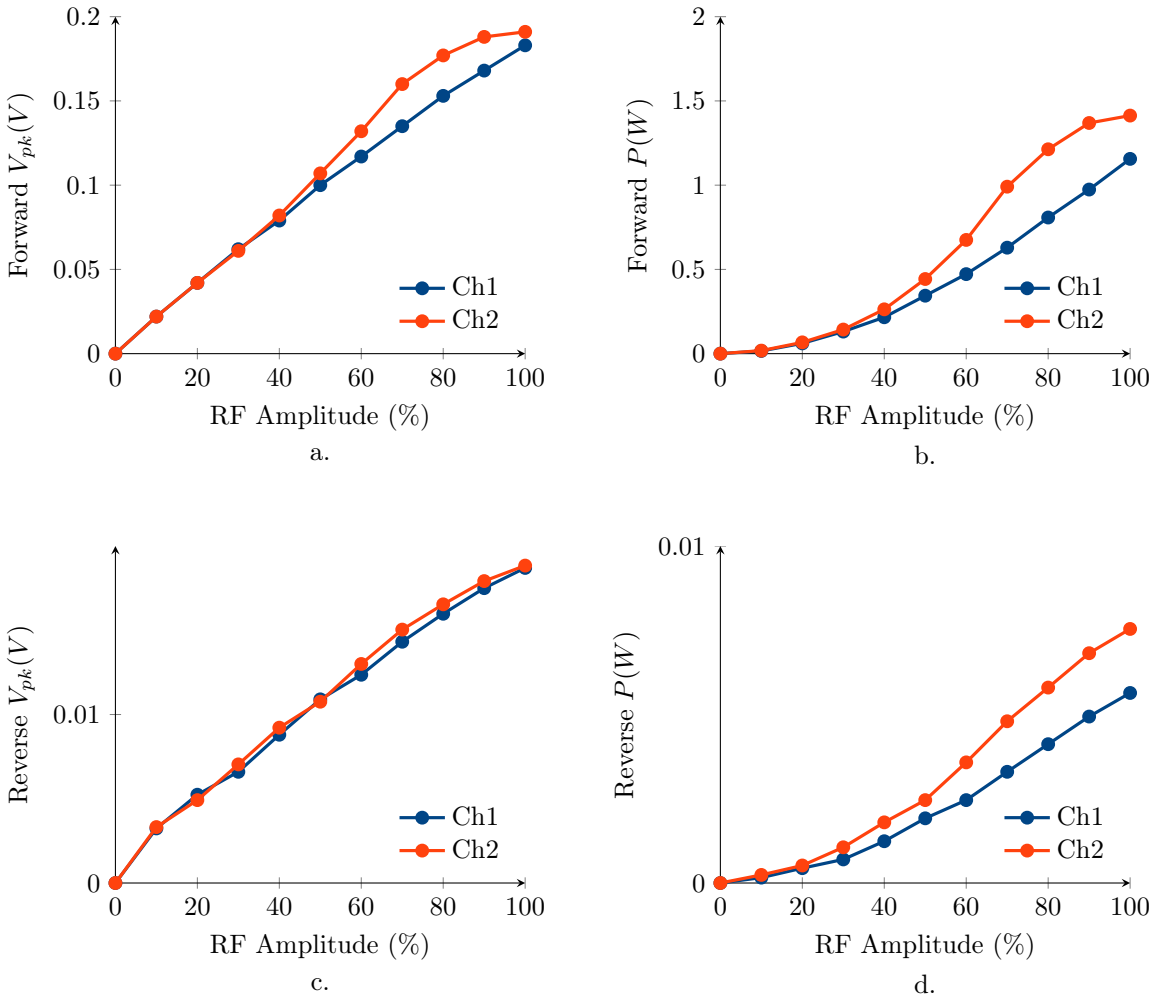

Figure 5.15: Measured voltage and power of each channel

Maximum power output was 1.16W (30.63dBm) and 1.41W (31.50dBm) for channels one and two. Channel two outputs more power than channel one, however as seen in the figure 5.15.a and c the output becomes increasingly compressed and non-linear, whilst channel one retains linearity better (particularly noticeable on figure 5.15.a, where the measured voltage data points are close to a straight line). Reflected power is low and the return loss remains constant at approximately 24dB, indicative that the connecting cables and bias tee are well matched to 50Ω.

The linearity of channel one is highly desirable for parallel transmission, as it relies on the linear superposition of the B_1 fields from each coil element. By operating in the linear region of an amplifier the calculated amplitudes map directly to the console RF amplitudes. For channel two, a correction curve could be used, however this adds additional complexity into calculations, requires the acquisition of an accurate calibration curve, which may differ for different load conditions introducing systematic errors. Most previous studies have required less than 2W of power for labelling, at this power level the non-linearities of channel two are reduced and a linear approximation may suffice.

When using the transmitter as the excitation source for imaging sequences non-linear behaviour

is problematic for shaped pulses (such as a slice selective sinc pulse) and without correction the resultant B_1 waveforms and excitation patterns will differ vastly from those predicted. As the transmitter system and coils will only be used for transmitting hard pulses this problem will not be encountered, at worst the resultant flip angle will differ from that prescribed.

5.6 Conclusion

Presented in this chapter was the hardware and software required to add an additional two transmit channels to a Siemens Trio using a SMIS spectrometer as a base system. Both transmit channels have independent amplitude and phase control as required for parallel transmission. The system is triggered by a single pulse from the Trio, and can be configured to output either pulsed or continuous wave RF power, making it suitable for both excitation during imaging sequences, and to provide the long labelling pulses in continuous ASL.

Chapter 6

Radio Frequency Safety

Biological interactions with radio frequency (RF) electromagnetic (EM) fields are a major concern in MRI. Two mechanisms exist whereby an oscillating electromagnetic field can deposit power within a biological sample, increasing its temperature: Joule (also known as ohmic) heating and dielectric heating. Joule heating is a consequence of induced RF eddy currents within a resistive medium; the passage of an electrical current through a conductor deposits energy due to the charges in the current (usually electrons) giving up energy each time they collide with particles within the conductor, increasing their thermal energy and the temperature of the conductor. Dielectric heating is a result of dipolar molecules rotating to align with an oscillating electromagnetic field, increasing the kinetic energy of the molecules, and consequently increasing the macroscopic thermal energy and hence temperature of the medium. Dielectric heating is dominant at high frequencies and in materials with poor conductivities, such that $\sigma \ll \omega\epsilon$. It is exploited in microwave ovens which use electric fields oscillating at 2.4GHz to rapidly heat food. More detail about these heating mechanisms in biological materials can be found in [222, Ch.4].

In MRI high power transmitters capable of many kilowatts peak power are employed to produce sufficient B_1^+ for nutating the macroscopic magnetisation. Transmitter coils such as birdcage coils are typically designed to efficiently produce an oscillating circularly polarised magnetic field, with minimal electric and electromagnetic (travelling wave/radiated) components inside the imaging volume. Therefore, the majority of the deposited power results from magnetically induced eddy-currents within the sample.

The amount of power deposited within the human body must be carefully controlled to prevent thermal injury: a local, or whole-body increase in temperature that is greater than the body can safely deal with. Therefore, in a MR scanning session it is important to know in advance how much RF energy will be absorbed by the human body and its spatial distribution, however this is not a trivial task. Because of the heterogeneity of the human body the distribution of deposited power cannot easily be predicted with analytical models. Without knowledge of the regions of

intense localised power deposition it is insufficient to only monitor/set limits on the total power delivered (for example by using directional couplers in the transmitter path), a convenient real-time measurement on an MRI scanner. The deposited power distribution must be either measured directly or simulated, from which limits on the globally delivered power can be set. As MRI moves to increasing field strengths, and consequently higher frequencies where significant B_1^+ distortions are present due to standing wave effects from the RF wavelength being similar in dimension to the anatomy of the human body, full-wave 3D solutions to Maxwell's equations have become the prevailing method for determining the deposited power distribution.

In contrast to the the high power transmitter system used for driving a body coil, the ASL Labelling Coils described in chapter 4 use low power (1W nominal RMS power) RF amplifiers to create a continuous wave (CW) B_1^+ field that labels blood in the neck. As close fitting surface coils are used, which produce their most intense electric and magnetic RF fields close to the conductors, there is the possibility of excessive local power deposition. Therefore it is crucial to assess what the peak power deposition and associated temperature rise is, and if necessary impose limits on the experimental parameters such as labelling amplitude, labelling duration and sequence repeat time.

This chapter presents an assessment of the risks from the ASL labelling coils and seeks to quantify them by determining the worst case temperature change and attributed deposited RF power. To do so, an optimised method using proton resonance frequency (PRF) MR thermometry [118] has been developed that allows the accurate measurement of small temperature changes using MRI. PRF thermometry takes advantage of the linear temperature dependency of the water proton chemical shift; approximately $-0.01 ppm/^{\circ}C$, encoding this into the phase of a gradient echo image. By comparing the phase between two gradient echo images, the change in temperature can be deduced on a voxel-by-voxel basis. However, the magnetisation phase has other time-varying dependencies and corrections must be made to account for this to ensure accuracy of the temperature measurements.

6.1 Risks

6.1.1 Radio Frequency Power Deposition

6.1.1.1 SAR

As discussed in the preceding section, the radio frequency electromagnetic field produced in MRI interacts with the human body due to the conductivity of tissue, resulting in the induction of eddy currents within the body which cause Joule heating in the tissue. The specific absorption ratio (SAR), defined as the power deposited per mass of tissue (units W/kg) is a metric of the rate at which energy from radio frequency electromagnetic fields is absorbed by the body. It is commonly

used for RF dosimetry in MRI and for mobile phones, and can be calculated from the electric field within the tissue:

$$\mathrm{SAR}(\mathbf{r}) = \frac{\sigma(\mathbf{r})\,|\mathbf{E}(\mathbf{r})|^2}{\rho(\mathbf{r})} \tag{6.1}$$

Where σ and ρ, the sample conductivity and density are spatially varying quantities and $\mathbf{E}$ is the RMS electric field. In human MRI, at field strengths below, but not including 7T (<300MHz) the electric field is predominantly due to Faraday induction within tissue from the RF magnetic field. MRI transmit RF coils do generate electric fields, particularly close to capacitors (which by their very definition store energy within an electric field), but these are minimised by using distributed capacitance (section (2.2.2.1)), and increasing the coil-sample distance. In human whole body MR systems at field strengths of 7T and above, travelling-wave waveguide modes [32] are permitted, potentially increasing the electric field component generated by the transmitter coil.

Regulatory limits within the European Union, defined by the International Electrotechnical Commission in document IEC 60601-2-33 [5] are shown in table 6.1. All SAR values are time averaged over a 6 minute period. Local SAR values are determined by averaging over a cubic volume containing 10g of tissue (a suitable post-processing averaging algorithm is described by Caputa et al. [34]).

	Average SAR (W/kg)			**Local SAR (W/kg)**		
Operating Mode	**Whole Body**	**Partial Body**	**Head**	**Head**	**Trunk**	**Extremities**
Normal	2	2 - 10	3.2	10 (c)	10	20
First Level Controlled	4	4 - 10	3.2	10 (c)	10	20
Second Level Controlled	>4	>(4 - 10)	>3.2	>10 (c)	>10	>20
Short term SAR	The SAR limit over any 10s period shall not exceed three times the stated values					

Table 6.1: IEC SAR Guidelines
Normal Operating Mode: Partial body SAR = 10 – (8 × exposed patient mass / patient mass) W/kg
First Level Controlled: Partial body SAR = 10 – (6 × exposed patient mass / patient mass) W/kg
For whole body exposure there is a limit of 0.08 W/kg averaged over the whole body.

The first incarnation of SAR guideline limits were based on ensuring the total power absorbed by the body should not exceed the basal metabolic rate, approximately 1W/kg for a healthy adult [27]. Local SAR limits are given to ensure that local SAR does not cause localised heating or burns, and is averaged over a mass (10g for the IEC, 1g for the FDA) to account for thermal dissipation effects.

6.1.1.2 SAR and Temperature Relationship

As SAR is defined as the deposited power per mass of tissue, it can also be obtained from the rate of change of temperature with respect to time. From the definition of heat capacity, $dQ = CdT$ (where Q is the internal energy, ΔT is the change in temperature and C is the specific heat capacity), the rate of change of temperature with respect to time can be written as

$$\frac{dT}{dt} = \frac{1}{C}\frac{dQ}{dt} \tag{6.2}$$

The quantity $\frac{dQ}{dt}$ is the rate of change of internal energy with respect to time. For living tissue exposed to RF energy this is

$$\frac{dQ}{dt} = \text{SAR} + P_m - P_c - P_p \tag{6.3}$$

Where P_m is the metabolic heating rate per unit mass, P_c is the rate of heat loss per unit mass due to conduction, and P_p is the rate of heat loss per unit mass due to blood flow (perfusion). If the tissue is at thermal equilibrium prior to RF exposure, i.e. $P_m = P_c + P_p$, and then the SAR can be directly related to the initial rate of change of temperature providing the effects of metabolism, conduction and perfusion are negligible in the time Δt:

$$\text{SAR} = C\frac{\Delta T}{\Delta t} \tag{6.4}$$

6.1.1.3 Factors which affect SAR

As has been eluded to, RF power absorption within biological materials is affected by many factors: the electric or magnetic nature of the incident RF field, its frequency (content), the electrical (e.g. conductivity, permittivity), and geometric properties of the biological material. In MRI, at frequencies up to those corresponding to a field strength of 3T (up to 128MHz) the main mechanism of interaction is the induction of RF eddy currents in conductive tissues, and the subsequent heating of this tissue through Joule heating. Therefore in order to determine the SAR, the nature of these eddy currents must be established.

The magnitude of the induced eddy currents within a conductive material is proportional to the induced electro-motive force (EMF), which, from Faraday's law of induction, is proportional to the rate of change of magnetic flux passing through a surface formed whose boundary is the eddy current loop (equation 2.45 in section 2.2). Material conductivity dictates the size of the induced eddy currents (through the application of Ohm's law), and deposited power, whilst the permittivity of a material affects an electromagnetic wave's speed of propagation and subsequently its wavelength. As a result of this, the wavelength of the induced EMF and eddy currents at the Larmor frequency can be of similar dimensions as parts of the anatomy, and wave-associated effects

such as reflections occur at material boundaries to produce standing waves. If the geometry of the medium is such that one or more of its dimensions are a quarter wavelength (or integer multiple of) long then a resonance condition can be met, accompanied by large peak eddy currents and power deposition. The induced eddy currents serve to oppose the incident B_1^+ field, reducing its intensity and as such, this phenomenon is responsible for B_1^+ inhomogeneity in 3T body imaging [73], and at 7T, where wavelengths within the body are less than 20cm it is a formidable challenge as these field distortions are present in most anatomy [224].

Another significant phenomenon is the skin effect. In section 2.2.1.2 the skin effect and skin depth was introduced for good conductors, and the skin depth was defined in equation 2.51. However, in lossy conductors such as tissue the electrical permittivity cannot be neglected and a more general expression is required:

$$\delta = \sqrt{\frac{2}{\sigma\omega\mu_0\mu_r}}\sqrt{\sqrt{1+(\frac{\omega\epsilon_0\epsilon_r}{\sigma})^2}+\frac{\omega\epsilon_0\epsilon_r}{\sigma}} \tag{6.5}$$

Eddy currents induced within a conductor produce a magnetic field which opposes the incident field, reducing its amplitude exponentially as it penetrates the conductor. With respect to power deposition this has the effect of forcing the induced electrical current and subsequent power deposition towards the surface of the conductor. The additional term in equation 6.5 is significant for the conductivities found in biological tissues and serves to lengthen the skin depth in comparison to the good conductor approximation. Figure 6.1.a illustrates this; skin depths for muscle, fat, grey matter, white matter and dry skin have been calculated for frequencies between 1 and 200MHz using the good conductor approximation (equation 2.51, dashed lines) and the full model (equation 6.5, solid lines). The conductivity of fat is an order of magnitude less than that of the other tissue, hence it does not attenuate EM fields as much and has a larger skin depth. The values of conductivity and relative permittivity used to calculate skin depth for each tissue type were obtained from [11], which is an internet resource that calculates the dielectric properties of body tissues based on a parametric model [74, 75, 76, 77]. Figure 6.1.b and c are plots of the relative permittivity and conductivity for each of these tissues between 1 and 200MHz, respectively. Values of skin depth, conductivity and relative permittivity at 123.2MHz are given in table 6.2.

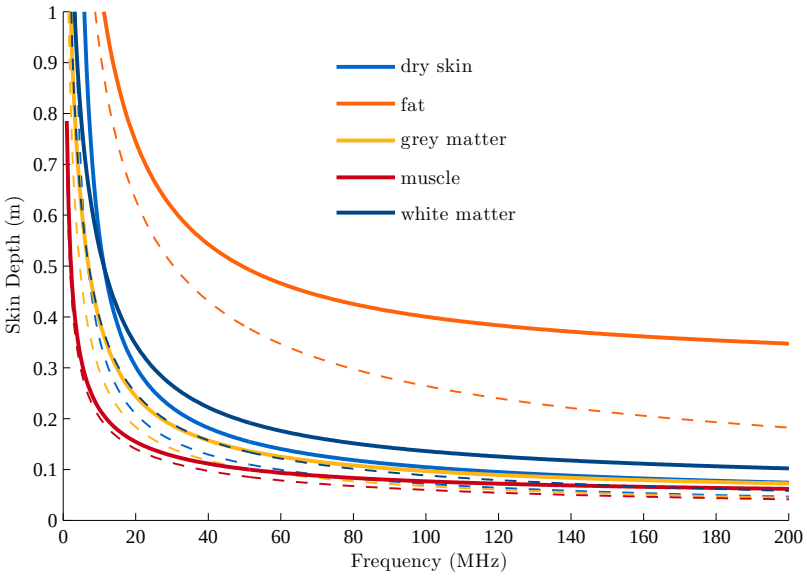

a. Skin Depth

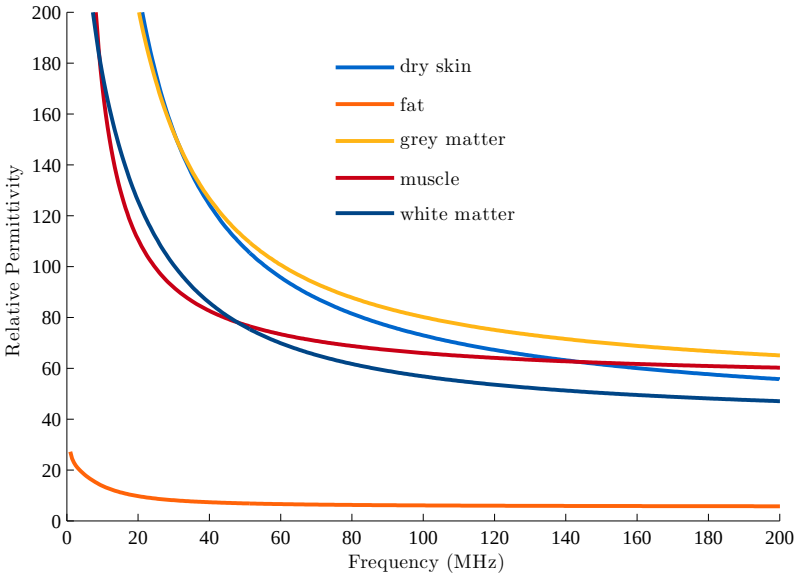

b. Relative Permittivity

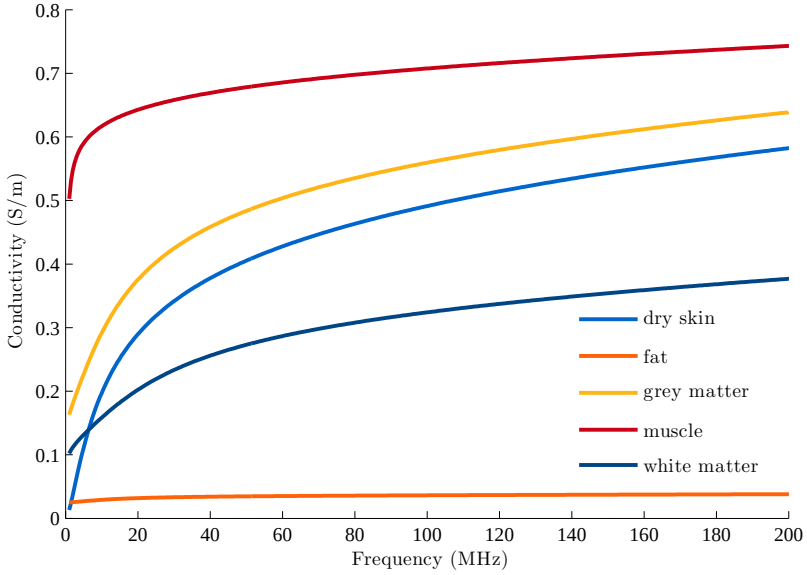

c. Conductivity

Figure 6.1: Electrical properties of tissue 1-200MHz

6.1.1.4 Determining the SAR distribution

Early analytical work to investigate the effect of the RF field used in MRI on biological materials [28, 27, 24, 23] modelled the human body as an infinitely long conductive cylinder, or the human head as a conductive sphere. For example, Bottomley and Andrew [23] investigated a planar RF magnetic field axially incident on a homogeneous 20cm radius tissue cylinder at frequencies between 1-100MHz. In this configuration, the skin depth of most biological tissues (a function of frequency, electrical permittivity, magnetic permeability and conductivity) was found to be 20cm or less at frequencies of 20MHz or greater, indicating significant attenuation (over one skin depth the RF field is attenuated by a factor of $1/e$, approximately 37%). Induced eddy currents and subsequent power absorption were found to be at their greatest at the surface of the cylinder where the flux linkage[1] is at its maximum. With increasing frequency, due to decreased magnetic field penetration the power absorption is reduced somewhat in comparison to what would be expected from an extrapolation of the low frequency results.

Tissue	δ (cm)	ϵ_r	σ(s/m)
Muscle	7.1	63.8	0.71
Fat	38.1	5.94	0.04
Grey Matter	8.8	74.4	0.58
White Matter	12.4	53.1	0.34
Dry Skin	9.4	66.5	0.52

Table 6.2: Dielectric properties at 123.2MHz

While homogeneous tissue models have shown that power deposition varies approximately as the square of the RF frequency and sample radius, with maximum local SAR and heating at the surface, they do not accurately predict the SAR distribution within the heterogeneous structure of the human body. Numerical methods that simulate the interaction between electromagnetic fields and body models have since become the prevailing approach. At the high frequencies used in modern MRI systems, full wave differential forms of Maxwell's equations must be solved, commonly using finite-difference time-domain (FDTD) [212] or finite element method (FEM) [120] algorithms.

A surface coil has a very non-uniform B_1^+ and hence SAR distribution, with peak amplitudes found closest to the coil. This was recognised in surface coil hydrogen (^1H) decoupling for carbon (^{13}C) spectroscopy, where a continuous wave B_1^+ field at the ^1H frequency is transmitted at the same time as a ^{13}C spectroscopy sequence to collapse multiplets into single resonances and increase sensitivity [26, 28]. The magnetic field produced by a figure-eight decouple coil was computed by numerical integration of the Biot-Savart equation, assuming that the sample and ^{13}C imaging coil do not perturb the field, then using the vector potential of this field, and Faraday's law the induced electric field within a semi-infinite plane of biological tissue of uniform conductivity and dielectric constant and SAR was determined. These calculations suggested that local power deposition levels could reach as high at 9-12W/kg per Watt of input power to the coil, however when averaged over

[1] the total magnetic flux passing through a surface, defined as $\Phi = \int_S \mathbf{B}.d\mathbf{S}$

1g of tissue this was reduced to less than 8W/kg, within FDA guidelines at the time. Approximately 90% of the power deposition from both a 6.5cm diameter surface coil, and a 8×13cm figure-eight coil, occurs within 2cm of the surface.

Like surface coil decoupling, separate coil CASL delivers continuous wave RF for several seconds at the ^1H frequency using a close fitting surface coil. As shown in [26, 28] there is cause for considerable concern regarding peak local power deposition during ^1H decoupling, and therefore these should also apply to separate coil CASL. The numerical calculations are a quasistatic approximation[2], and assume that there are no sample perturbations or high frequency effects such as the skin depth or dielectric resonance, which is not the case at 3T (123.2MHz). An additional factor to consider is that there are two coils, both of which may transmit simultaneously with different phase shifts applied, resulting in both constructive and deconstructive interference of the magnetic and electric field components generated by the coils. It is important to determine the SAR for both individual coils, and when transmitting simultaneously to diligently assess safety before subjects can be scanned.

6.1.2 Electromagnetic Compatibility Considerations

Multiple devices operating within the same electromagnetic environment can have an influence on each other by transmitting/intercepting stray electromagnetic fields. It is vital that interference effects between the electromagnetic fields they generate/receive are avoided. Electromagnetic compatibility (EMC) is a branch of electrical engineering that studies these effects with the aim of ensuring compatibility between devices. In general, electrical devices have to comply with certain regulations and laws, for example within the EU EMC tests must be run in order to comply with CE-labeling. A device may produce an unwanted electromagnetic field, which via a coupling mechanism (radiated, magnetic, electrostatic) is intercepted by another device, resulting in an unwanted effect. Neither of these devices have to be designed as a transmitter or receiver for them to behave as one, and many of the techniques employed in EMC are to ensure that only transmitters transmit, and receivers receive.

With regards to MRI, any device which is placed within the scanner bore will be subject to intense time-varying electromagnetic fields from both the gradients and RF transmitter. Any device which interacts with these fields presents a potential thermal safety hazard by either perturbing the SAR distribution and creating localised hotspots, or heating of the device itself from induced eddy currents. MRI gradients have a frequency content that extends to approximately 100kHz, and have been shown to deposit a negligible amount of power in biological tissues [24]. Eddy currents

[2]Using solutions obtained for DC currents for time varying fields. For alternating magnetic the concept of of "retarded time" is required due to the fact that an electromagnetic field cannot change instantaneously in response to changes at its source (the field changes occur at the speed of light). At low frequencies where the propagation time between the source and observation point is much smaller than the time period of field oscillations the quasistatic approximation holds. However at higher frequencies, including those used in MRI the propagation time should be taken into account. It is because of the non-negligible propagation time that there is a discrepancy between the B_1^+ and B_1^- fields above 1.5T .

induced within peripheral devices may cause susceptibility artefacts in MR images, or cause the device to malfunction, but are not large enough to cause significant heating of the device itself, and therefore from a safety perspective do not present a hazard. However the interaction between the RF electromagnetic fields generated within the MR scanner and a peripheral device can be a hazard, and as such the device must be well designed to minimise this. The ASL Labelling Coils present risks to subject safety from the transmitter coil interaction to the coil elements, and the connecting cables.

6.1.2.1 Coil Elements

Each coil is a tuned circuit, resonant at 123.2MHz and will strongly couple to the B_1^+ field transmitted by the body coil. Series PIN diodes have been used within each coil and are reverse biased by 5V to detune the coils. As the PIN diodes are in series, the coils are intrinsically detuned even when no bias voltage is applied. In the event that both of these fail-safe mechanisms fail and the coils become resonant during body coil transmission the B_1^+ from the body coil will be focused by the ASL coils, increasing the local SAR.

The PIN diode driver circuitry shown in section 5.4 includes status LED's to indicate when the PIN diodes are forward or reverse biased. If a PIN diode short circuits then the status LED's will show that the PIN diode is forward biased when no TTL trigger signal is present, providing a means of monitoring for errors. This was done through through visual inspection, however it would be possible to use a dedicated hardware and software solution using a microcontroller to feed back error information to the SMIS console using the user input port, with the pulse sequence programmed to stop in the event of an error.

6.1.2.2 Connecting Cables

Conductive wires/cables placed within the bore of a MR scanner which are electrically large in comparison to the wavelength of the EM fields at the frequency of interest are susceptible to strongly coupling to the RF field via two mechanisms: The first is induction of currents in loops of wire by the magnetic field, and the second is direct coupling of the electric field to the wire (the antenna effect). Both result in large voltages and currents on the wire, which heat up the wire itself through Joule heating, and also distort the SAR distribution, potentially amplifying it in close proximity to the wire. Both mechanisms have been observed to produce temperature increases of over $60°C$ in resonant lengths of wire [51]. Because of the interest in interventional MRI and combined EEG-fMRI, research into both the mechanisms and how to mitigate their effects is ongoing and active [12, 13, 64, 230, 231].

Magnetic induction on the interfacing coaxial cables used on the ASL coils was reduced by keeping both cables together to minimise their loop area. The antenna effect is strongest for resonant lengths of cable; integer multiples of $\lambda/4$, therefore cable traps (see section 4.2.2.4) were

used approximately every 30cm on all cables within the MR scanner bore, electrically shortening the lengths of cable such that they were less than an eighth of wavelength long, and also attenuating induced RF currents. High attenuation (>45dB, an attenuation factor of approximately 30,000) figure eight coaxial traps were used on the cabling within the base housing (onto which the two coils mounted to), and three floating cable traps were used on the long lengths of RG223 coaxial cable. Being resonant structures, both styles of cable traps have high voltages and circulating currents on them in order for them to correctly block RF currents on the cables, and the subject must be protected from these too. The figure eight cable traps were separated from the subject by approximately 3cm of air and perspex, and the floating cable traps were shielded with a layer of aluminium tape and nylon insulating material to contain the electric field within.

6.2 Methods to determine the SAR

The specific absorption rate of RF energy being deposited into biological tissue can be determined from either the electric field, using equation 6.1, or the temperature increase of the sample using equation 6.4. All methods are time consuming and as a consequence are not performed on a subject by subject basis.

6.2.1 Electric field measurements

The local electric field can be directly measured by mechanically moving a calibrated electric field probe, such as a monopole or dipole antenna, inside a tissue phantom with known conductivity and density, making point measurements over a regular mesh. This is a popular technique for mobile phone SAR compliance testing [22] and has been utilised for SAR mapping in MRI [123]. However it is time consuming, requires an appropriate mechanical assembly, can only be used to make measurements in free space or within liquid phantoms, and the electric field probes perturb the RF field. Subsequently, directly measuring the electric field has not become a widespread method in MRI SAR evaluation, although it can be used for validating simulation results.

6.2.2 EM Modelling

The prevailing method for determining the electric field and SAR distribution from MRI RF coils is to numerically compute the electromagnetic field by solving Maxwell's equations. The most widely used implementation is the finite difference time domain (FDTD) [115, 126, 155, 201] method.

FDTD uses the fact that the time derivative of the electric field is dependent on the curl of the magnetic field; from the Ampère's circuital law, equation 2.43, and the time derivative of the magnetic field is dependent on the curl of the electric field; from the Maxwell-Faraday law, equation 2.45. Both of these equations are discretised using central differencing[3], and evaluated

[3]the central difference of a function $f(t)$ with spacing Δt is given by $\delta f(t) = f(t + \Delta t/2) - f(t - \Delta t/2)$. It is more

over a geometric lattice. Each lattice cell, known as a Yee cell [252] is defined such that the electric field components run along the edges of the cell, and the magnetic field components run normal to the centre of the faces of the cell. A simulation is constructed by creating a lattice of Yee cells, each with a specified permeability, permittivity and conductivity to define a material. A source, for example a voltage or current is specified and simulation proceeds in the time domain by calculating the electric and magnetic field components on alternating time steps. By using a broadband pulse (such as a Gaussian) the frequency response of the system can be obtained via Fourier transform of the time domain fields.

For stability both the spatial and temporal step size must satisfy the "Courant-Friedrichs-Lewy" (CFL) stability criterion; the time step must be smaller than the time it takes for an electromagnetic wave to propagate across a Yee cell. As a result the FDTD method is computationally intensive. Non-uniform grids can be used in areas of low detail to reduce the number of cells, and absorbent boundary conditions to truncate the computational domain and place boundaries over what is simulated.

Whole body digital phantoms (voxel or boundary representation models with conductivities and permittivities set to mimic biological tissue) facilitate the simulation of all aspects of the complex interaction between the RF electromagnetic field and the human body. Some software packages also include thermal solvers, predicting the induced temperature rise by modelling the Pennes Bioheat equation [196, p.51], a modification of the general differential equation of heat conduction to include terms for metabolic heat production, and perfusion.

EM modelling was once computationally intensive, with run times on the order of days and weeks. As computing power increases exponentially and becomes more affordable, utilising advances such as multi-core processors and general purpose graphics processing unit (GPU) technology, these simulations can be conducted in mere hours on standard off-the-shelf equipment.

For the uninitiated however it can be difficult to find an entry point into electromagnetic simulation. A good grounding in electromagnetic theory, and computational electromagnetism is required to be able to set up simulations such that they will give sensible results. Commercial software packages (for example CST Microwave Studio, Remcom XFDTD, Speag SEMCAD, Ansoft HFSS) are expensive and have a steep learning curve, despite their graphical user interfaces and CAD-like modelling tools. It is relatively straightforward to write FDTD code in C++, however the code would then need to be validated against an existing software package, and a means of modelling the geometry and creating a grid devised. Open source software packages are also available such as [142, 174], however until recently they have lacked many of the features present in commercial software packages and have not been featured in MRI publications: OpenEMS [142] was recently used to simulate a novel 7T travelling wave antenna and human body model [66].

accurate than a forward or backward difference.

6.2.3 Calorimetry

Calorimetry is used to measure the heat of chemical or physical processes from the induced change in a thermodynamic property (such as temperature, pressure, volume). Its principles are used in USA test standards to determine the local [1] and global [2] power deposited from diagnostic MRI coils by measuring the temperature rise within a tissue mimicking phantom. Similar standards apply in the UK. These test procedures are outlined below.

Global SAR Global SAR is determined using a saline phantom (volume $<$ 25 litres) with conductivity such that it provides the same load as a 50–90kg human, thermally insulated such that it can maintain a $2°C$ elevation over the ambient temperature to within 5% for 1 hour. The phantom is allowed to equalise its temperature with the MR scanner bore and scan room. The temperature is recorded at the start, and then scanning is performed for long enough to produce a temperature rise at least 20 times greater than the error of the temperature measurement system. Then the temperature is immediately measured. Because the mass of the liquid within the phantom and its specific heat capacity is known, the power deposition, SAR of the phantom, and apparent SAR (the SAR of the human body which the phantom simulates) can be calculated. Global SAR is not a suitable metric for the safety compliance of surface coils because they produce an inhomogeneous RF field.

Local SAR Local SAR is determined using a gel phantom to inhibit convection, with sodium chloride and other additives to adjust the conductivity and relative permittivity of the phantom so that it is equivalent to muscle tissue at the frequency of interest. In addition the specific heat capacity and density should match muscle tissue. Thermal insulation is used to minimise heat loss. The phantom is allowed to reach thermal equilibrium with the scan room, and then the temperature is recorded continuously within the phantom whilst scanning with a high power fast spin echo sequence for approximately 20 minutes. The temperature probe is placed where SAR is expected to be highest. The SAR is then calculated from the initial temperature rise on the heating curve.

Both methods require an accurate thermometer. For the local SAR measurement this must be MRI compatible as the temperature is logged during scanning, ruling out all thermometer probes except for expensive ($>$£1000 per channel) fibre-optic thermometers. To correctly measure the maximum local SAR the thermometer probe has to be positioned at the "hot spots", which must be determined in advance, usually from numerical simulations.

6.2.4 MRI Temperature Mapping

The NMR signal can be sensitised to temperature changes, making it possible to noninvasively acquire quantitative maps at standard MRI image resolutions. Many NMR parameters display some

temperature dependence, but in most cases it is difficult to measure their values separately. Three parameters in which the temperature dependence can be measured separately are the longitudinal relaxation time (T_1), the diffusion coefficient (D) or the proton resonance frequency (PRF). Temperature difference maps showing the change in temperature before and after RF heating can be used to calculate the SAR, however as with the case with the calorimetric method for measuring local SAR, care must be taken to ensure that the temperature change is representative only of effects due to the SAR.

6.2.4.1 Longitudinal Relaxation Time Thermometry

T_1 based MRI temperature mapping was first reported in 1983 by Parker et al. [179], taking advantage of its linear temperature dependence over a small temperature range. Fast imaging techniques must be used to acquire T_1 measurements, as the conventional saturation or inversion recovery sequences do not provide sufficient temporal resolution. Typically either a short TR (therefore T_1 sensitive) RF spoiled gradient echo sequence, or single shot fast spin echo inversion recovery is used (only one inversion time is feasible) [152]. A reference image is acquired, and then subsequent images will show differences in intensity based on the temperature dependence of T_1, therefore only relative temperature measurements can be made. Temperature changes are on the order of $1\%/^{\circ}C$ [139]. For gradient echo images there is the possibility for interference with the PRF effect which causes intra-voxel dephasing of spins (see section 6.2.4.3) and subsequent signal reduction, therefore small voxels must be used and preference is given to the fast spin echo method because it avoids this problem altogether. The temperature dependency of T_1 to temperature becomes smaller with increasing field strength [25], decreasing sensitivity and negating the benefits of increased SNR from high field systems.

6.2.4.2 Diffusion Thermometry

There is an exponential relationship between the diffusion constant D and the Brownian motion of an ensemble of molecules, amounting to a change in D of approximately $2\%/^{\circ}C$ which is independent of field strength. MR images are sensitised to diffusion by using a strong bipolar gradient to dephase diffusing spins, resulting in signal attenuation proportional to the distribution of molecular displacement. Diffusion based thermometry shows good sensitivity and accuracy in phantom measurements, with an accuracy of $\pm 0.2^{\circ}C$ [50], however in-vivo the temperature dependence is non-linear due to the temperature dependent permeability of barriers such as cell structures, and physiological effects which can produce large changes in D. PRF induced signal changes are also present in gradient echo diffusion thermometry, with the effect more pronounced and problematic than in T_1 thermometry due to the long TE values required for diffusion sensitisation [189].

6.2.4.3 Proton Resonance Frequency Thermometry

The well known linear temperature dependence of the chemical shift was first investigated in 1966 by Hindman [104]. Stretching, bending and rupturing of the electronic bonds in water result in a shift of the water proton resonance frequency to lower frequencies with increased temperature by approximately $-0.01ppm/^{\circ}C$ for pure water. The proportionality constant for different types of in-vitro tissue is almost the same as that of pure water [133]. PRF thermometry has been demonstrated using chemical shift imaging [94], however because the change in the chemical shift is small accuracy is low and only large temperature differences can be observed qualitatively. 3D magnetic resonance spectroscopic imaging (3D-MRSI) can also be used [134], improving temperature resolution at the expense of poor spatial and temporal resolution.

High spatial and temporal resolution temperature mapping using the PRF is possible by using the magnetisation phase during a gradient echo sequence, a technique which has been named "Proton Resonance Frequency" (PRF) or "Proton Resonance Shift" (PRS) thermometry. The chemical shift is due to changes in the shielding constant of the protons, which can be represented as being caused by magnetic shielding effects; a B_0 inhomogeneity, which for a given echo time and field strength/Larmor frequency manifests as a phase shift in the magnetisation. There are other contributions to the magnetisation phase such as the phase of the RF transmitter and non-temperature dependent B_0 inhomogeneities which need to be removed in order to have a phase image which only depends on temperature. This is achieved by subtracting a phase map acquired at a reference temperature T_0, obtaining a phase difference map $\Delta\phi$ from which the temperature with respect to T_0 is calculated with:

$$\Delta T = \frac{\Delta\phi}{\alpha\gamma B_0 \mathrm{TE}} \tag{6.6}$$

Where α is the temperature dependent water proton chemical shift. Ishihara et al. [118] first proposed this method in 1995, demonstrating in-vitro temperature measurements of RF inductive heating in a saline phantom, agar, egg white, and chicken muscle. In addition the technique was performed in-vivo to monitor the cat brain temperature during whole body heating from hot air, and inductive heating of the brain. Temperature measurements were validated using copper-constantan thermocouples and the error was estimated to be less than $\pm 1^{\circ}C$ for temperature increases on the order of $10^{\circ}C$. Sources of error were found to be from magnetic field drifts due to system instabilities, image noise, variations in α between materials and temperature induced susceptibility changes.

The signal to noise ratio of PRF thermometry (SNR_T) is proportional to the product of the echo time (TE) and the T_2^* decay (equation 6.7), which has been shown to be at a maximum when $\mathrm{TE} = T_2^*$ [189]. Unlike thermometry methods based on T_1 and diffusion, temperature changes do not decrease the signal intensity and subsequent SNR of the images, unless a temperature gradient

across a voxel is strong enough to cause intra-voxel dephasing.

$$\mathrm{SNR}_T \propto \mathrm{TE} e^{-\mathrm{TE}/T_2^*} \tag{6.7}$$

Temporal system instabilities, if not accounted for attribute to large errors in PRF temperature measurements. The static magnetic field (B_0) drifts over time, particularly under heavy gradient duty cycle load, resulting in a temperature offset error the same for all voxels. Literature values [20, 70, 65] report that the typical drift of a MR system is specified to less than 0.1ppm/hr under no load, however this increases to 1Hz/min (around 0.5ppm/hr at 3T) during EPI or similar high gradient duty cycle scans, which would correspond to PRF temperature changes of $10°C$ and $50°C$ per hour, respectively. In addition, the spatial homogeneity of the B_0 field can change over time from the heating of the passive shims; iron plates positioned between the inner surface of the cryostat and the outer surface of the gradient coil that provide the first level of homogeneity correction. Heat from the gradient coils transfers to the shims, reducing their magnetisation [70] and causing thermal expansion, affecting the homogeneity of the B_0 field.

These instabilities can be corrected by estimating their contribution to the magnetisation's phase in each acquisition, then subtracting from the corresponding phase maps. The amount of correction required is dependent on the gradient loading of the scanner during the course of the temperature measurement; single slice gradient echo images acquired before and after a heating procedure are a minimal load, whereas continually acquired EPI images place a heavy load on the gradients. In the first instance it is sufficient to estimate/measure the drift in advance and apply this same correction factor to all voxels [203]. For higher gradient duty cycles it is necessary to measure the global drift (zeroth order) of the B_0 field, typically estimated from the phase within external reference phantoms placed around the object/anatomy scanned [184]. Long acquisitions and/or high gradient duty cycle sequences also introduce higher order perturbations to the B_0 field because of the markedly different shim conditions to those when the reference images were acquired. Spatial corrections must also be made, typically by using at least three external reference phantoms and linearly interpolating between to produce a 2D phase correction map, first proposed by De Poorter et al. in 1995 [185]. With this first order correction, temporally and spatially consistent temperature measurements with an accuracy of $0.2°C$ were made within an agar phantom by acquiring a gradient echo image every two minutes over the course of four hours.

PRF thermometry is of interest in interventional MR applications, for example real time monitoring of radio frequency [188], microwave [19, 131], ultrasound [208, 105, 217] or laser beam hyperthermia [229, 191, 19, 131]. Body motion and the desire to acquire more than a single slice at each time point has driven the development of fast PRF thermometry methods. Echo shifting combined with multiple gradient echoes [48] permit TE $>$ TR, providing fast acquisition speed (144ms per 64×60 image) yet retain the high SNR_T expected for TE values around T_2^*. EPI

sequences can also be used, however because each line of k-space is acquired at different TE's, the echo train length (ETL) must be kept short. Emphasis has been particularly placed on segmented and partial Fourier acquisitions, combined with parallel imaging to achieve a short ETL [229, 208, 36, 105, 217]. It has also been shown that single shot EPI is equivalent to the segmented acquisition [129], and that the gain in speed can be invested into a longer TE, resulting in improved PRF SNR and temperature accuracy.

6.3 SAR Mapping Using Proton Resonance Frequency Magnetic Resonance Thermal Imaging

Magnetic Resonance Thermal Imaging (MRTI) utilising the temperature dependence of the PRF can acquire 2D/3D temperature maps with good spatial and temporal resolution and a precision within phantoms of better than $0.5°C$ (in-vivo measurements are subject to physiological noise and motion, decreasing the precision and accuracy of measurements). In addition to monitoring thermal therapy procedures, PRF thermometry can be used to investigate the thermal effects of RF power deposition.

While MRTI methods have been used to investigate the SAR induced temperature rise within gel phantoms [64, 95, 202, 203, 201], these measurements do not necessarily equate to the SAR. To derive accurate SAR values, MRTI measurements need to satisfy the criteria given in 6.1.1.2: the system must be in thermal equilibrium before heating commences, and the temperature difference used to calculate SAR values must be over a period in which thermal losses through convection, conduction and radiation are negligible.

Cline et al. [41] utilised PRF MRTI to compare heating measurements with analytically and numerically derived SAR profiles. Agar spheres with different conductivities were heated using a high power, fast spin echo sequence for 300 seconds, with gradient echo images acquired before and after. During this heating period the thermal diffusion was estimated to be 0.9cm. A power of 10.6W was used at 1.5T, and 4.6W at 3T, resulting in temperature increases of $1.8 \pm 0.4°C$ and $0.32 \pm 0.2°C$, respectively. A 3-inch surface coil was also used for transmission, depositing power such that the temperature increase close to the coil was $6°C$. In all these cases the power was set to the maximum allowed by the MR scanner. Temperature maps were found to be in good agreement (RMS error of $0.4°C$) with the simulated temperature increase using the numerically and analytically computed SAR maps.

Oh et al. used PRF MRTI to assess MRI induced RF heating [168], explicitly using the measured temperature maps to compute SAR maps for a transmit receive head sized birdcage coil within a conductive agar sample (trapezoidal prism in shape). Preliminary experiments using fibre optic temperature sensors were conducted to determine the period in which the temperature

change was linear (therefore thermal diffusion effects were negligible), from which the SAR could be calculated. Subsequent SAR mapping performed in a 3T MR system consisted of two fast gradient echo sequences, separated by a high power, short TR spin echo sequence lasting the appropriately determined time, 120 seconds. Using the measured forward power and the loaded to unloaded quality factor ratio of the coil, the power dissipated in the sample was determined to be 5.25W. Experimentally measured SAR maps matched well with numerically computed SAR maps. This technique was also used in-vivo to measure the temperature change and SAR within a rat using a 7T pre-clinical MR system. Gradient echo images were acquired every four seconds, with a WALTZ-4 decoupling module placed between acquisitions, dissipating 11W of power into the rat for a total of 120 seconds. Time-course temperature measurements showed a linear fit, verifying the validity of SAR measurements in-vivo, however no numerical simulations SAR maps were produced for comparison.

Further work [167] introduced a high power (50W) solenoid coil continuous wave heating system to improve the quality of the temperature map by intensely heating a conductive agar sample. PRF MRTI measurements were validated against fibre-optic temperature sensors positioned within the phantom and found to be in agreement to within $0.12°C$, and an experimentally determined SAR map showed good agreement with one numerically simulated.

These three examples confirm the feasibility of using MRTI measurements to determine the SAR distribution of a MR transmitter coil. As shown by Oh et al. [168, 167] it is important to consider thermal diffusion and ensure its contributions are negligible over the duration of the temperature change used to calculate the SAR map. SAR Mapping using PRF MRTI measurements has been presented as a tool for validating numerical simulations, however its good agreement between measured and predicted SAR maps, and providing appropriate consideration is given to ensuring thermal diffusion is negligible, it is reasonable to propose that experimentally measured SAR maps derived from PRF MRTI could be used for safety compliance in their own right, without additional numerical simulations. Assuming experimental errors are minimised, experimentally determined SAR maps have some advantages over those numerically simulated. No assumptions have to be made about coil geometry, component and material losses, interactions with other objects, or the composition of the sample. The idealisations available when performing numerical simulations are useful for studying physical principles, or novel transmitter concepts, however for evaluating a specific coil design it can be difficult to capture the nuances/imperfections of the manufacturing process, and in doing so the simulation can become more complex (for example a larger or more finely graded FDTD mesh) and time consuming to run. Therefore PRF MRTI was selected as the method for determining the spatial SAR distribution of the ASL Labelling Coils, from which appropriate safety directives could be established.

6.3.1 Considerations

In order to produce SAR maps of the ASL Labelling Coils using PRF MRTI, certain factors had to be taken into consideration. High spatial resolution, three dimensional temperature maps were desired over a volume in which the majority of the power from each coil was dissipated. Thermal diffusion effects need to be minimised both during acquisition, calling for the use of fast imaging techniques. Previous studies have assumed that the only contribution to the temperature change within a voxel is from RF power deposition[168, 41, 201], valid only if thermal conduction is negligible between the two measurements, and if the object is at thermal equilibrium when heating commences. The low power amplifiers used to drive the ASL Labelling Coils means that only a small temperature rise occurs over the short period in which conduction is negligible, and consequently the measurements suffer from low SNR. The coils could be driven with a more powerful amplifier, however in doing so the equipment set-up would deviate from what would be used in CASL measurements, and the ASL Labelling Coils were not designed for high power use.

6.3.1.1 Power and induced temperature changes

Medium to high power sequences have been used in order to maximise the temperature change over the short period of time in which thermal diffusion is negligible and therefore the temperature increase is linear, Δt_{lin}. Each ASL Labelling Coil was driven by a nominal 1W RF power amplifier, which is around a tenth of the power reported in [41] and [168]. Consequently, the initial temperature change during Δt_{lin} will be smaller and have lower SNR. In order to increase SNR without compromising the accuracy of the SAR maps it is proposed to make multiple MRTI measurements for an extended duration (much longer than Δt_{lin}), and fit an appropriate function to the temperature time course on a voxel-by-voxel basis, from which the SAR can be calculated using the initial temperature change of the fitted data.

6.3.1.2 Sequence

A fast, high spatial resolution, multi-slice or 3D sequence is required to acquire gradient echo phase images at regular intervals. Fast gradient echo (fGRE, similar to a FLASH sequence) sequences have been used extensively for single slice PRF MRTI, acquiring an image in around four seconds, however this is not fast enough for high resolution multi-slice/3D coverage as significant thermal diffusion will occur. Gradient echo echo planar imaging (GE-EPI) has been shown to be suitable for acquiring PRF MRTI images with high temporal and spatial resolutions [229, 208, 36, 105, 217, 129], acquiring many slices in the same amount of time required for a single fGRE slice.

The high gradient duty cycles necessary for EPI will induce both zeroth and higher order B_0 drift, which must be corrected for accurate temperature measurements. In addition, the heat from the gradients will raise the air temperature within the magnet bore, introducing an extra

source of heat within the sample and confounding SAR measurements. To correct for this a state of dynamic thermal equilibrium can be established prior to scanning by running the same EPI acquisition sequence for several hours prior to commencing MRTI measurements.

6.3.1.3 Fitting Function

To increase the SNR and accuracy of the initial temperature change it is proposed to acquire PRF MRTI data for an extended duration such that the temperature change is many times the noise level, and then fit a suitable function which empirically describes the processes of heating and thermal diffusion. Within an agar gel phantom undergoing RF power deposition the temperature evolution is governed by the heat equation for a non-perfused material with an internal heat source:

$$\rho c \frac{dT}{dt} = k\nabla^2 T + \mathrm{SAR}\rho$$

Where k is the thermal conductivity, ρ is the material density, T is the temperature, t is time and c is the material specific heat capacity. The term $k\nabla^2 T$ is derived from Fourier's Law of Heat Conduction, which states that the time rate of heat transfer through a material is proportional to the negative local temperature gradient, and to the cross sectional area through which the heat is flowing. If RF power is deposited at a constant rate, the heat will conduct through the material until it is all at the same temperature. Solutions to the heat equation are non-trivial, however by considering the nature of the temperature change under different conditions, a suitable fitting function can be deduced which empirically describes the temperature change within a voxel.

First, consider a voxel in which RF power is being deposited, and is perfectly insulated from surrounding voxels. For constant RF power deposition, the rate of change of temperature is also constant, therefore the temperature increases linearly, as shown in figure 6.2.a.

Second, consider a voxel at an elevated temperature to its surroundings, where heat can flow to and from the voxel via conduction. Heat energy is lost to the surroundings at a rate proportional to the temperature difference, therefore the temperature decays exponentially until it is has equalised with the temperature of the surroundings, see figure 6.2.b.

Third, consider a voxel being heated by RF power deposition at a constant rate, in thermal contact with a thermal reservoir. Both voxel and reservoir are initially in thermal equilibrium. The temperature change within the voxel is at first linear, however as the heat is conducted to the surroundings at a rate proportional to the temperature difference, the rate of temperature increase will decrease until the heat generated by SAR is equal to the heat lost through conduction. At this point the temperature will no longer increase, but will remain constant so long as the conditions do not change, as shown in figure 6.2.c.

Fourth, consider a voxel surrounded by a heat reservoir whose temperature is increased linearly with time. Heat is conducted between the voxel and reservoir, however the rate of flow is slow

compared to the rate at which the reservoir's temperature increases, i.e. the process cannot be considered to be quasi-static. As the temperature difference between the voxel and the thermal reservoir increases, so does the rate at which heat flows from the reservoir to the voxel. The temperature of the voxel will initially increase exponentially, asymptotically approaching the linear increase in temperature of the thermal reservoir, as shown in figure 6.2.d.

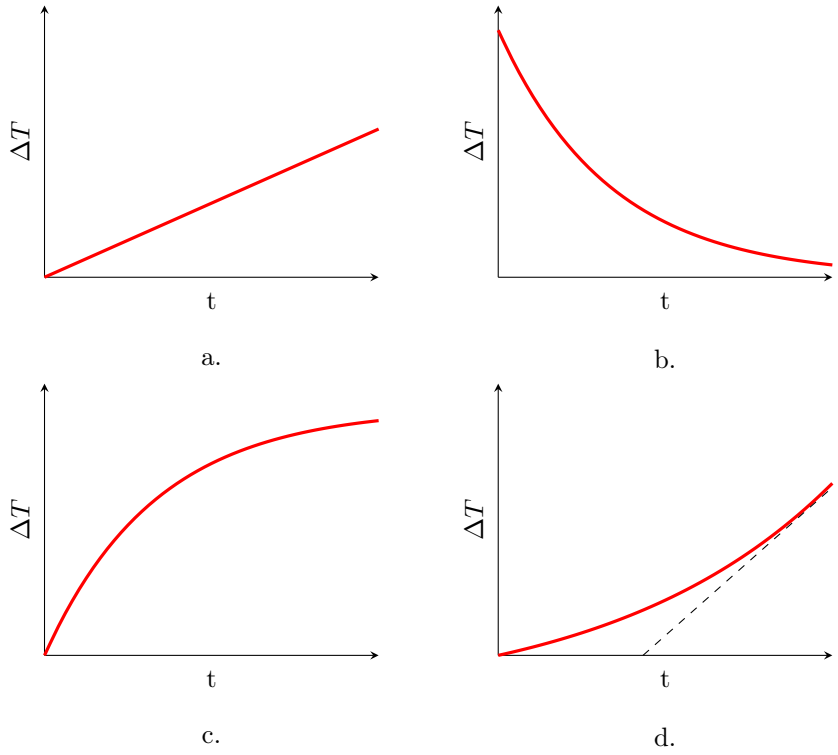

Figure 6.2: Temperature time evolution in a single voxel for a. heat generation with no loss, b. conductive loss to surroundings, c. conductive heating of the voxel from surroundings, d. combined heat generation, conductive loss, conductive heating, and internal heat generation

Any given voxel within the agar phantom will experience a combination of direct heating from SAR, cooling from conduction of heat to surrounding voxels that are at a lower temperature (for voxels at the surface of the phantom there will also be heat loss to the phantom container and surrounding air), and heating from conduction of heat from surrounding voxels that are at a higher temperature. By considering a function which can match any combination of the heating curves shown in figure 6.2a., c. and d., the temperature change in a voxel can be approximated. One suitable function is given in equation 6.8; the product of a linear and exponential function. For $c_2 \simeq 0$ the function is linear as the exponential evaluates to unity, $c_2 < 0$ the function resembles the exponential recovery in figure 6.2.c, and for $c_2 > 0$ and $t < \frac{1}{c_2}$ the function resembles the exponential growth in figure 6.2.d. For $t > \frac{1}{c_2}$, $T_{fit}(t) \to 0$, therefore the function is only valid in the domain $t < \frac{1}{c_2}$.

$$T_{fit}(t) = c_1 t e^{c_2 t} + c_3 \tag{6.8}$$

The SAR can then be computed from the value of T_{fit} at the first and second time points, t_1

and t_2 respectively:

$$\text{SAR} = \frac{c(T_{fit}(t_2) - T_{fit}(t_1)}{\Delta t} \tag{6.9}$$

6.3.1.4 Phase Stability

As discussed in section 6.2.4.3 phase stability is key for accurate PRF MRTI measurements. Ideally, the homogeneity of the static magnetic field, and the zeroth order offset should be the same between the reference gradient echo image and all subsequent gradient echo images, so that the phase difference only depends on the change in temperature. Unfortunately this is not the case; both the zeroth and higher order components of the magnetic field vary over time, with high gradient duty cycles exacerbating this effect.

B_0 inhomogeneities can be corrected for by using the phase information in regions where no heating has occurred to create extrapolated maps of the non-temperature phase change, which are subtracted from the object in which the temperature is being measured. This can be done with reference phantoms [185], or using the phase of regions inside the object that are not subject to heating [191]. The latter requires the temperature change to be spatially confined, for example laser heating, therefore it is unsuitable for RF heating. If reference phantoms are used then at least three are required, positioned around the object. Typically the reference phantoms consist of cylinders of water/agarose, positioned parallel to the in-plane image direction. The procedure for generating B_0 drift corrected temperature maps is as follows:

1. A time series of phase images are acquired, and the phase difference with respect to the starting image in the series is calculated.
2. Images are segmented into the object and reference regions.
3. The measured phase difference within the reference regions is extrapolated over the entire field of view to generate phase drift maps.
4. Phase drift maps are subtracted from the phase difference maps.
5. Temperature difference maps are calculated according to equation 6.6.

This method assumes that the temperature of the reference regions remains constant. If this is not the case then the phase of the reference regions must be corrected to account for temperature changes [185]. The temperature of the reference phantoms must be logged during scanning, and prior to spatial drift map generation the associated temperature dependent phase calculated and subtracted from the voxels in the corresponding reference phantom. As only a point temperature measurement is made in each reference phantom the temperature distribution needs to be as uniform as possible.

6.3.2 Methods

6.3.2.1 Hardware

Imaging was performed on a 3T (Tim Trio, Siemens Medical Systems, Erlangen, Germany) whole-body MR system, interfaced to a two-channel in-house built transmitter system based around a re-cycled spectrometer (Surrey Medical Imaging Systems, Guildford, UK) as described in chapter 5. Each channel utilised a 1W (nominal) RF power amplifier (ZHL-3A, Mini-Circuits Inc. Brooklyn, NY, USA) to drive one of the ASL Labelling Coils. A whole-body birdcage coil was used to transmit the imaging RF pulses, and a 12 channel receive-only head coil for reception.

Temperature measurements were made with a four channel fibre optic thermometer system (ProSens with PSP-62 Modules and OTP-A Sensors, OpSens, Quebec, Canada), sampling at 100Hz. An analog voltage corresponding to $10°C/V$ was output by each thermometer channel and logged using an analogue to digital interface (CED 1401, Cambridge Electronic Design Ltd, Cambridge, UK) and laptop running data acquisition software (Spike 2, Cambridge Electronic Design Ltd, Cambridge, UK).

The phantoms (see section 6.3.2.2) were placed within the receive coil, and the ASL coils positioned on opposite sides of the agar phantom, facing each other. A temperature probe was inserted into each of the reference phantoms and held in place with PVC adhesive tape, and a temperature probe was inserted into the centre of the agar phantom equidistant from the two ASL coils.

6.3.2.2 Phantom

A 12cm diameter, 20cm long cylindrical gel phantom was prepared using distilled water mixed with 2% agar, 1.6% $CuSO_4$ and 0.9% NaCl. The phantom diameter corresponds to the neck dimensions reported in section 3.1.1 in chapter 3. Three reference phantoms were constructed from 32mm diameter, 25cm sections of PVC pipe. Threaded ABS couplers were fixed to the ends of each pipe using solvent weld glue to allow corresponding screw on caps to be fit. A hole with M3 size thread was made in the centre of each phantom for thermometer probe access; a nylon M3 nut and O-ring provided a watertight seal at other times. Reference phantoms were filled with a solution of 1.6% $CuSO_4$ distilled water, sealed, then wrapped in four layers of 3.3mm thick expanded polyethylene foam for thermal insulation, with sections removed for access to the probe holes. Liquid reference phantoms were used to maximise thermal homogeneity through convection as only point temperature measurements were possible.

The three reference phantoms were positioned equilaterally around the agar phantom and supported using a frame constructed from blocks of expanded polyethylene foam held together with strips of adhesive backed hook and loop, as shown in figure 6.3. Access was maintained to the sides of the agar phantom for positioning of the RF coils, and the positions of the reference

phantoms minimised the RF power deposited within them. Markings were made on the agar phantom container to ensure identical positioning of the coils and alignment with the magnet isocentre.

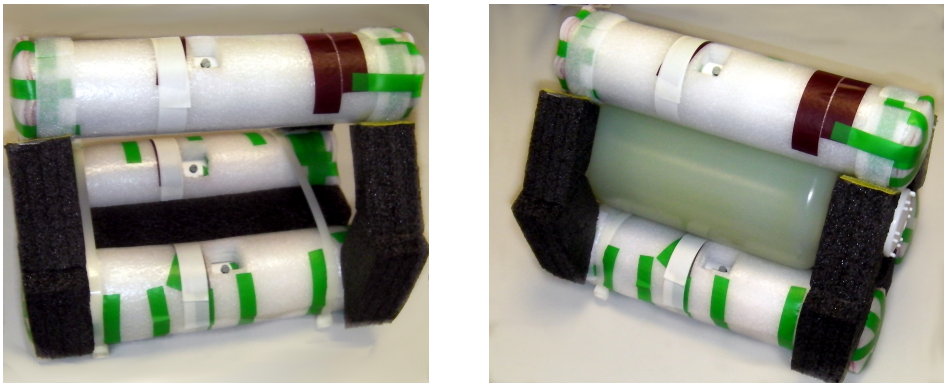

a. b.

Figure 6.3: Reference phantoms and foam support frame a. without and b. with the agar gel phantom in place

6.3.2.3 Sequence

A high resolution, 3D EPI sequence, provided by Dr. Antoine Lutti from the Wellcome Trust Centre for Cognitive Neuroimaging, University College London was modified to produce an optical trigger pulse for a specified duration prior to each EPI volume acquisition. The sequence had previously been optimised for fMRI use, which requires a high degree of stability, hence making it an ideal candidate for MRTI measurements.

EPI volumes of resolution 128×128×48 were acquired using 2× GRAPPA acceleration in 3840ms with the following parameters: FOV $= 210 \times 210 \times 120mm$, 2.5mm slice thickness, TE=32.86ms, TR per slice=80ms. Acquisitions were RF spoiled to suppress stimulated echoes. Continuous wave RF power was applied to the ASL Labelling Coils (100kHz higher than the imaging frequency to avoid saturation of spins) for 30s prior to each volume acquisition, resulting in a heating duty cycle of 88.7%. A single volume was acquired at the start of the sequence without any prior RF heating, followed by 99 sets of heating and an EPI volume acquisition.

To ensure a thermal steady state an equalisation EPI sequence was run prior to each set of MRTI measurements. This was identical to the MRTI measurement sequence, however no RF power was applied to the ASL coils. The temperature within the agar phantom and the reference phantoms was monitored during the control EPI sequence, and the sequence was run until they reached a steady state. The necessary duration of the equalisation scans varied with how the MR scanner had been used on the day the measurements were made[4], ranging between 250 and 500 repetitions (2 hours 20 minutes to 4 hours 40 minutes). For each coil configuration (see section 6.3.2.5) two sets of MRTI measurements were run; between the two sets of measurements the

[4]Some were made after the scanner had been left overnight, others after the scanner had been used all day scanning subjects, typically EPI fMRI with high gradient duty cycles.

equalisation EPI sequence was run for 250 measurements.

6.3.2.4 Post-Processing

All image processing was performed in Matlab (The Mathworks Inc., Natick, MA) using custom written scripts. Fibre optic thermometer data was smoothed using a 20 second window moving average, then interpolated to the acquisition times of the MRI data. Magnitude images from the first acquisition of each time series were used to construct a binary image mask by smoothing with a Wiener filter, setting all voxels below a dynamically determined threshold limit to 0, filling holes and then setting all non-zero pixels to have a value of 1. The mask was was then segmented into the agar phantom and each reference phantom.

For the n^{th} time point, phase difference maps, $\Delta\varphi_n$ were produced by complex dividing the phase map at each time point, by the first phase map in the time series:

$$\Delta\varphi_n = \arg\left(\frac{e^{i\varphi_n}}{e^{i\varphi_1}}\right) \tag{6.10}$$

Despite thermal insulation, a small temperature increase was detected in the reference phantoms due to non-negligible power deposition from the ASL coils and heating from the gradients. Therefore, at each time point the phase difference within each reference phantom was temperature corrected by complex dividing with the corresponding temperature dependent phase accrual for that phantom:

$$\Delta\bar{\varphi}_{ref,m,n} = \arg\left(\frac{e^{i\Delta\varphi_{ref,m,n}}}{e^{iT_{m,n}\alpha\gamma B_0 \mathrm{TE}}}\right) \tag{6.11}$$

Where m is the reference phantom index, $\Delta\bar{\varphi}_{ref,m,n}$ is the temperature corrected reference phantom phase difference, $\Delta\varphi_{ref,m,n}$ is the raw reference phantom phase difference, $T_{m,n}$ is the temperature measured by the thermometer probes. B_0 drift correction was then applied to each slice by generating a 2D linear fit from the temperature corrected phase difference within the reference phantoms, then complex dividing this correction phase map within the agar phantom by this correction map.

ΔT maps were calculated according to equation 6.6, with $\alpha = -0.0097ppm$, and the time series of each voxel within the agar phantom mask was fit to equation 6.8 using a non-linear least squares curve fitting algorithm. SAR maps were calculated with equation 6.9, where $c = 4200J/kg$, and divided by the RF duty cycle to obtain SAR maps at 100% duty cycle. In addition maps of the sum of squares residuals between the measured temperature data and the model fit were generated, providing spatial information about the confidence of the SAR maps.

Voxels which displayed changes due to susceptibility instabilities (for example air pockets forming between the agar and plastic phantom container over the duration of the scan) were identified with a region mask by calculating the absolute difference between magnitude images from the first

and last acquisitions in each time course, and setting a lower threshold of the mean plus 10 standard deviations of all voxels within the agar phantom volume. Connected regions with less than four voxels were discarded from this mask, and the region was dilated with a circular structuring element of radius 6 voxels. Voxels that were within this mask or that had a sum of squares error greater than 2 were excluded from further analysis. Histograms with 100 bins were then calculated for all qualifying positive value voxels of the measured SAR, and the temperature increase at the 100^{th} acquisition, providing measures of the temperature and SAR distribution due to heating from the coils.

For easy comparison SAR values at the temperature probe position calculated from the raw MRTI data, fitted MRTI data and thermometer data were normalised to the SAR value obtained from the fitted thermometer data (which was taken to be the gold standard): the Reference Normalised SAR defined as $\widehat{\mathrm{SAR}}$. From this a metric of accuracy was calculated, defined as the root-mean-square value of the $\widehat{\mathrm{SAR}}$ from all measurements:

$$\mathrm{RMS}\ \widehat{\mathrm{SAR}} = \sqrt{\frac{1}{N}\sum^{N}\left(\frac{\mathrm{SAR}}{\mathrm{SAR}_0}\right)^2} \tag{6.12}$$

Where N is the total number of measurements (10 – 5 coil configurations and 2 measurements per configuration) and SAR_0 is the reference SAR calculated from the fitted thermometer data. SAR values calculated with the raw data used the temperature change between the first and fourth measurement points, separated by approximately 2 minutes to maximise the temperature change yet remain within the linear region of the heating curve [168].

6.3.2.5 Coil Configurations

Two 100 acquisition MRTI sequences were run for each of the following coil configurations:

1. Right coil transmitting, left coil detuned. The right coil was connected to channel 1, and is hereafter referred to as coil 1.

2. Left coil transmitting, right coil detuned. The right coil was connected to channel 2, and is hereafter referred to as coil 2.

3. Transmit on both coils, set to 0° phase difference between channel 1 and 2 at the spectrometer, hereafter referred to as the "in-phase configuration".

4. Transmit on both coils, set to 90° phase difference channel 1 and 2 at the spectrometer, hereafter referred to as the "quadrature configuration".

5. Transmit on both coils, set to 180° phase difference channel 1 and 2 at the spectrometer, hereafter referred to as the "anti-phase configuration".

Transmit power was set to full on both channels. As the coils were geometrically identical, but placed opposite one another, in the absence of inter-coil coupling the in-phase and anti-phase configurations correspond to Maxwell and Helmholtz pair configurations. The three nominal phase shifts were chosen to cover a wide range of different driving conditions.

6.3.3 Results

6.3.3.1 Validation of PRF MRTI method

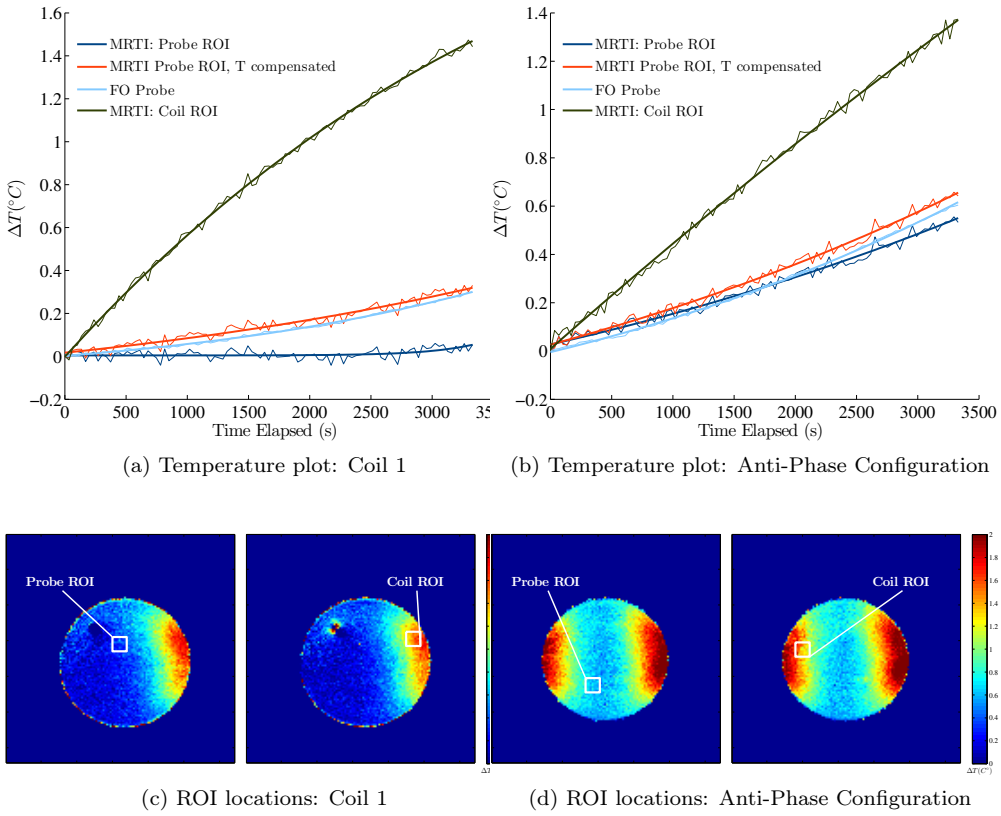

(a) Temperature plot: Coil 1

(b) Temperature plot: Anti-Phase Configuration

(c) ROI locations: Coil 1

(d) ROI locations: Anti-Phase Configuration

Figure 6.4: Plots of the thermometer and MRTI temperature data for a. the right coil transmitting, and b. both coils transmitting in the anti-phase configuration, and the ROIs used in the calculation on the relevant slices, c., d.

Figures 6.4.a and b show time-course plots of the temperature measured by the phantom thermometer probe (light blue), MRTI measurements averaged over a 9×9 voxel region of interest (ROI) around the voxel corresponding to the probe tip (termed Probe ROI in figure 6.4.c and d) with first order B_0 correction (dark blue) and temperature compensated first order B_0 correction (orange), and MRTI measurements averaged over a 9×9 voxel ROI close to the transmitting coil (termed Coil ROI in figure 6.4.c and d) with temperature compensated first order B_0 correction (dark grey-green). The locations of the corresponding ROI's are shown on ΔT maps at the 100th measurement in figures 6.4 c. and d. Fibre optic temperature measurements were logged at a sample rate of 100Hz, and subsequently have been smoothed with a 20 second window moving average and interpolated to the same time-base as the MRTI data. In both plots there is good agreement

between the temperature compensated MRTI data and the thermometer measurement. This was the case for all of the measured data. Most importantly the fitted temperature compensated MRTI data curve and the fitted thermometer data curves have very similar shapes, maintaining a similar offset throughout the entire measurement time course. Because the SAR is calculated from the subtraction of two temperatures at different times, this systematic offset cancels and makes no contribution.

The thermometer probe was generally positioned close to the centre of the phantom, where it did not directly receive a large SAR dose. The temperature rise at this position is largely attributed to the conduction of heat from regions close to the transmitting coil(s), characterised by a curve of increasing slope. Conversely, the temperature rise measured in the ROI positioned close to the transmitting coil is due to the incident RF field, where because of conduction to the surroundings it approaches a steady state temperature, characterised by a curve of decreasing slope. Fitted curves to the probe data, probe ROI and coil ROI's display the expected characteristics, confirming the suitability of equation 6.8 as a fitting function.

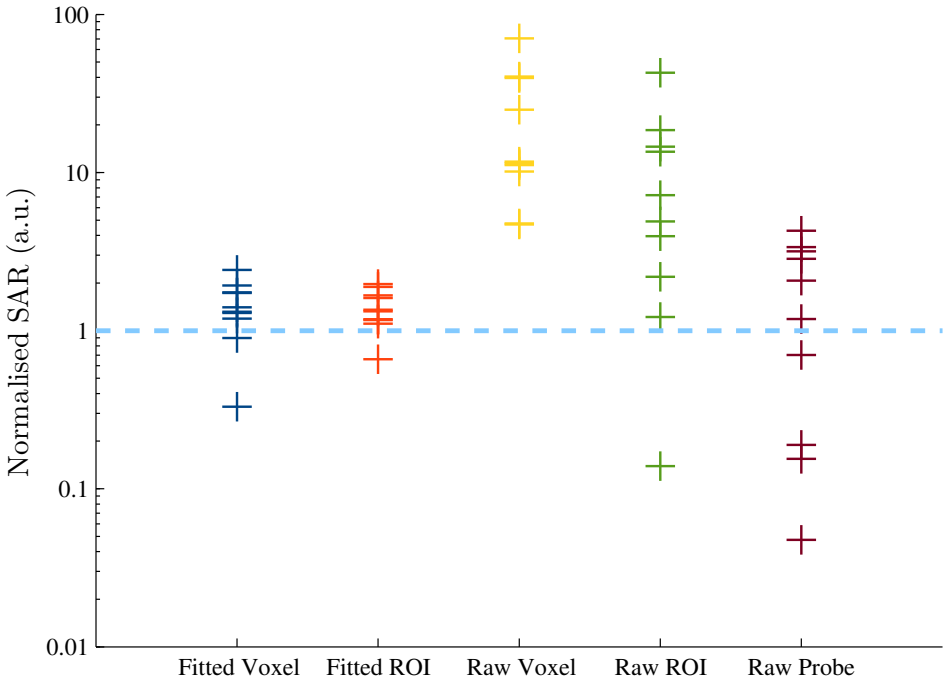

Figure 6.5: Reference Normalised SAR measured at the temperature probe position

Figure 6.5 is a plot of the calculated SAR's from each measurement, normalised to the SAR calculated using the corresponding fitted thermometer data. The y axis scale is logarithmic. A value close to unity indicates good agreement with the fitted thermometer data. In comparison to the raw (non-fitted) SAR values, including those from the raw thermometer data, the fitted MRTI SAR values show much better agreement with the fitted thermometer data. The RMS errors are given in table 6.3. Figure 6.6 shows single-slice SAR maps through the centre of the coils in the anti-phase configuration calculated from a. fitted MRTI data and b. raw MRTI data. The fitted

MRTI data provides a SAR map with good SNR and which clearly shows the SAR distribution of the coils, whereas the SNR of the raw MRTI data image is too poor to be of use.

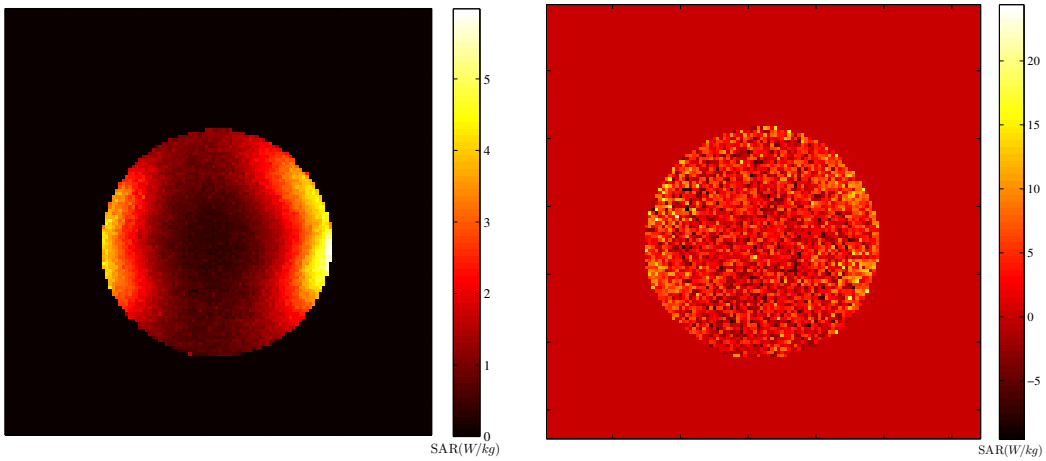

a. Fitted SAR map b. Raw SAR map

Figure 6.6: SAR calculated with fitted and raw MRTI data, anti-phase configuration

6.3.3.2 ΔT and SAR Results

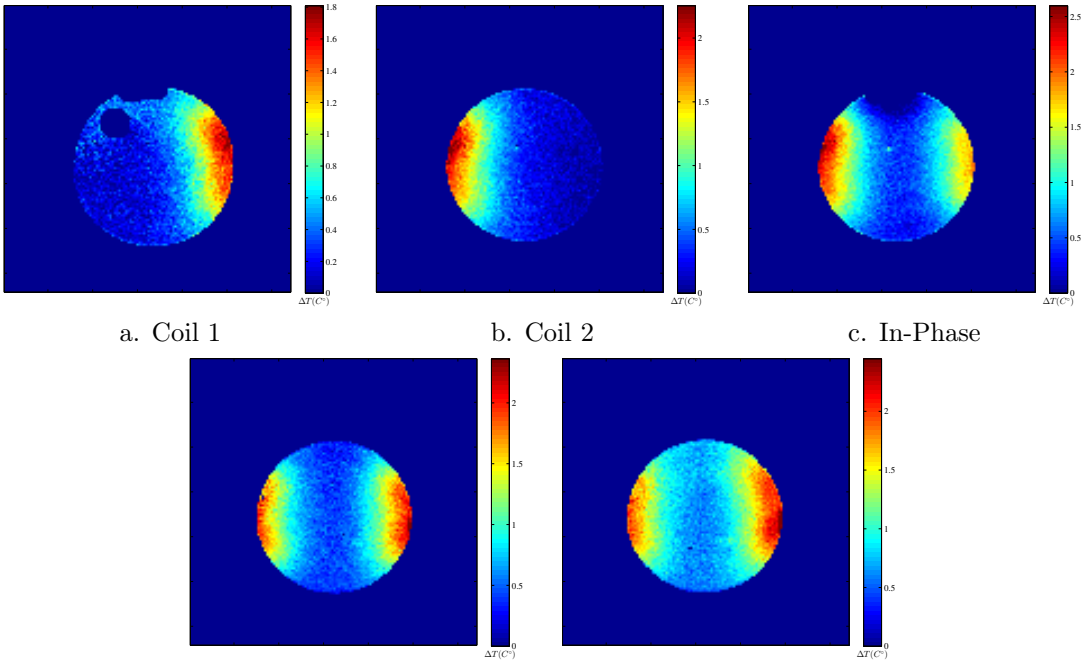

a. Coil 1 b. Coil 2 c. In-Phase

d. Quadrature e. Anti-Phase

Figure 6.7: Centre of coils slice ΔT Maps (colour scaled between 0 and the corresponding $\Delta T_{max,v}$ for each volume)

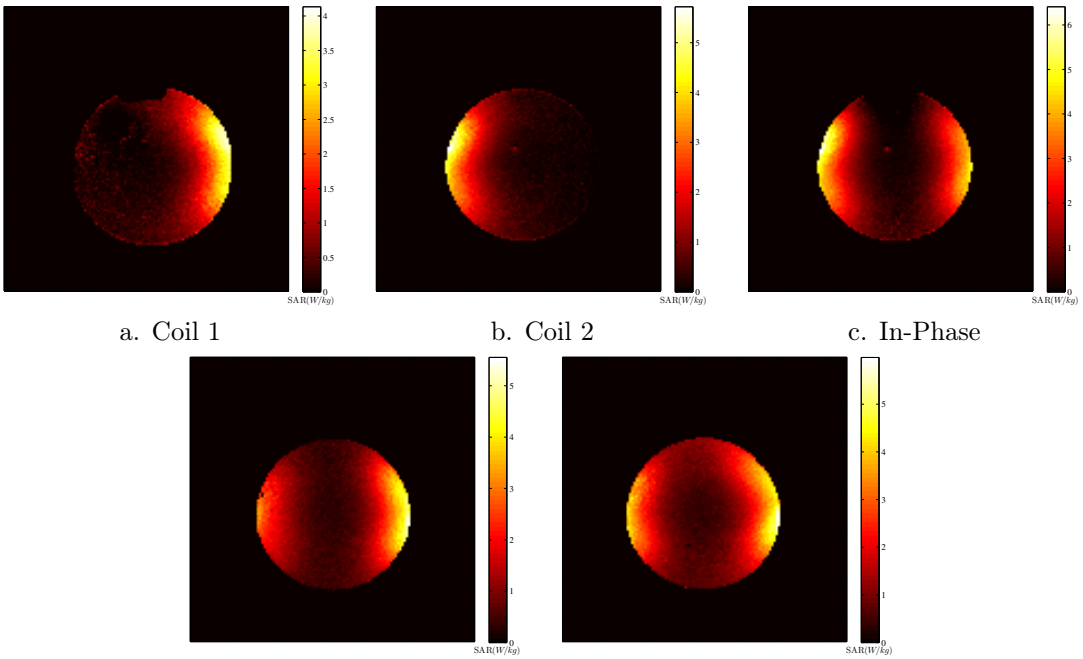

a. Coil 1 b. Coil 2 c. In-Phase

d. Quadrature e. Anti-Phase

Figure 6.8: Centre of coils slice SAR Maps (colour scaled between 0 and the corresponding $SAR_{max,v}$ for each volume)

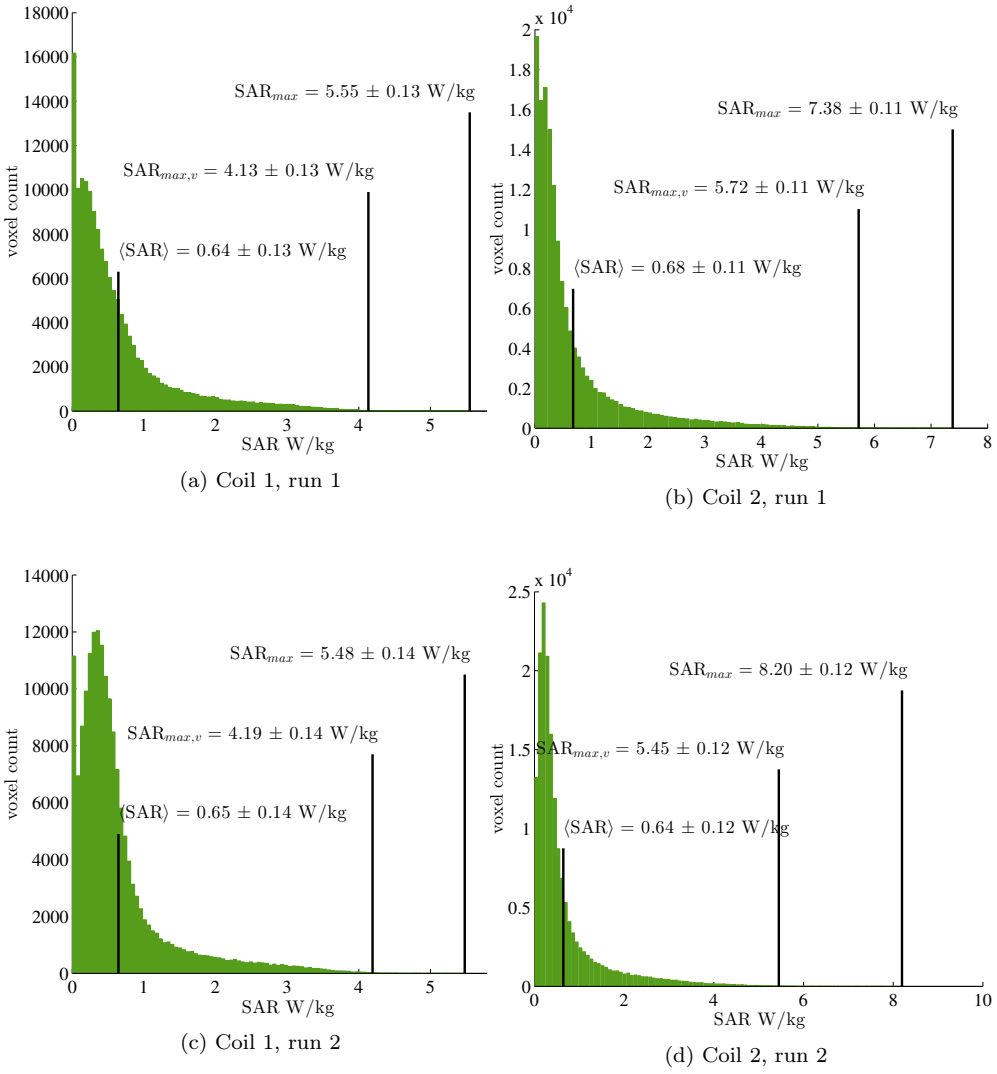

(a) Coil 1, run 1

(b) Coil 2, run 1

(c) Coil 1, run 2

(d) Coil 2, run 2

Figure 6.9: SAR Histograms of all suitable voxels for each coil transmitting individually

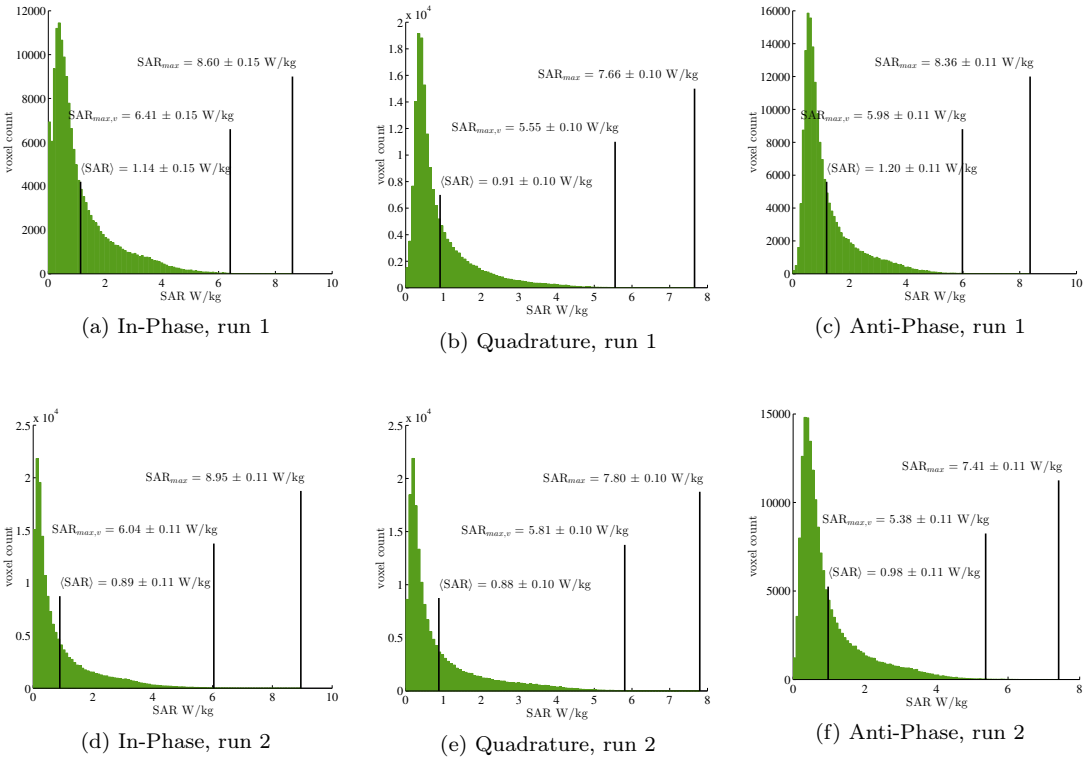

(a) In-Phase, run 1 (b) Quadrature, run 1 (c) Anti-Phase, run 1

(d) In-Phase, run 2 (e) Quadrature, run 2 (f) Anti-Phase, run 2

Figure 6.10: SAR Histograms of all suitable voxels for both coils transmitting simultaneously

			$(^\circ C)$			(W/kg)		
Coil Configuration	ϕ_1	ϕ_2	ΔT_{max}	$\Delta T_{max,v}$	$\langle \Delta T \rangle$	SAR_{max}	$\mathrm{SAR}_{max,v}$	$\langle \mathrm{SAR} \rangle$
Coil 1 (right)	0°	N/A	2.14 ± 0.23	1.76 ± 0.16	0.41 ± 0.13	5.51 ± 0.19	4.16 ± 0.18	0.64 ± 0.15
Coil 2 (left)	N/A	0°	2.98 ± 1.08	2.20 ± 0.11	0.42 ± 0.10	7.79 ± 0.69	5.59 ± 0.31	0.66 ± 0.14
In-Phase	0°	0°	3.09 ± 0.16	2.46 ± 0.28	0.64 ± 0.14	8.78 ± 0.38	6.23 ± 0.39	1.01 ± 0.30
Quadrature	0°	90°	2.84 ± 0.11	2.33 ± 0.13	0.60 ± 0.09	7.73 ± 0.20	5.68 ± 0.28	0.90 ± 0.12
Anti-Phase	0°	180°	3.00 ± 0.55	2.34 ± 0.24	0.73 ± 0.18	7.89 ± 0.78	5.68 ± 0.54	1.09 ± 0.27

Table 6.4: Measured temperature change and SAR

Results of the measured temperature change and computed SAR values are shown in table 6.4. Errors were estimated from the standard deviation of the two measurements for each coil configuration, and by subtracting smoothed SAR and temperature maps to obtain noise maps, from which the spatial standard deviation was measured to estimate the noise. ΔT and SAR maps of the slice at the centre of the coils are shown in figures 6.7 and 6.8 respectively. Histograms of the SAR distribution within the imaging volume are shown in figures 6.9 and 6.10. Full sets of ΔT maps, SAR maps, maps of the sum of squares error, and temperature histograms are shown in the appendix.

	RMS $\widehat{\mathrm{SAR}}$
Raw Voxel	30.60
Raw ROI	16.35
Raw Probe	2.33
Fitted Voxel	1.53
Fitted Probe	1.44

Table 6.3: RMS Reference Normalised SAR

ΔT_{max} and SAR_{max} are the value of the voxel with theΔT /SAR within all of the qualifying

voxels (i.e. all the voxels that were not excluded due to large errors) in each dataset. There may only be one voxel which has this value, and it could be a noisy measurement, therefore to improve the confidence in the measurement of the maximum ΔT and SAR a new metric was defined using the histograms: $\Delta T_{max,v}$ and $\text{SAR}_{max,v}$ are the values of the maximum ΔT/SAR histogram bin with a count of 25 voxels or greater. By excluding values attributed to individual voxels $\Delta T_{max,v}$ and $\text{SAR}_{max,v}$ provide a more realistic figure of what the peak ΔT and SAR within the phantom are.

ΔT maps show the greatest heating in close proximity to each coil element. Spatial and histogram temperature distributions are similar in both the individual coil configurations and the simultaneous transmission configurations. Additionally, the simultaneous transmission ΔT maps appear to be the superposition of each of the individual coil temperature maps, indicative that there no highly localised hotspots are formed from the addition of each coil's RF field. SAR maps also show similar agreement and features, with a more concentrated spatial distribution than the ΔT maps, indicating that by the 100^{th} acquisition the heat has been conducted towards the centre of the phantom. Additionally, the SAR maps show evidence of the skin effect as power deposition is attenuated in the centre of the phantom.

6.3.4 Discussion

6.3.4.1 Methodology

This methodology was optimised for improving the accuracy of the initial change in temperature when RF power is first applied, from which the SAR could be calculated. The combination of 3D EPI PRF MRTI measurements, B_0 drift correction using temperature compensated reference phantoms, and data fitting with equation 6.8 yielded temperature measurements accurate to less than $0.1°$ when compared to fibre optic thermometer measurements. The average noise level of the raw MRTI data, calculated as the standard deviation using the fitted function as a moving average is $0.06°C$ for the single voxel at the probe position, and $0.02°C$ for the 9×9 voxel ROI around the probe voxel, compared to $0.007°C$ for the thermometer data (the high temporal resolution of the thermometer data facilitates smoothing, substantially decreasing the noise; the manufacturer quoted accuracy of the thermometer is $\pm 1°C$). This is higher than the accuracy quoted in literature for PRF MRTI (in phantoms approximately $0.5°C$), and is only achievable because data is acquired for an extended period of time and the temperature change can be closely modelled with the fitting function. In-vivo measurements using this method would suffer from physiological noise, additional thermal processes such as perfusion might mean that the fitting function is no longer valid, and it may not be practical to run the equilibrium scans for the required period of time.

Because of the long measurement duration, phase stability was found to be crucial for accurate thermometry. The temperature compensated B_0 drift correction scheme applied was found to be

most effective if a long equilibrium scan had been run prior to MRTI measurements. A possible explanation is if the gradients and passive shims of the MR system are close to thermal equilibrium then any changes in shim condition over the measurement duration are spatially linear, which the 2D phase correction map can closely match. If the system is not at thermal equilibrium then changes in shim can be larger, introducing higher order harmonics which the linear phase correction map cannot correct for. A MR system which uses superconducting shims should provide better temporal stability.

The PRF MRTI SAR measurements made were found to over-estimate the SAR by approximately 50% compared with SAR measurements calculated from the fitted thermometer data, taken to be a "gold standard" SAR reference. The thermometer probe was placed at the centre of the agar phantom, a position which received a low SAR dose, therefore the temperature change and SNR of the MRTI signal was lower there than in regions of high SAR such as close to the coils. Averaging the MRTI data in a ROI had only a small effect on the magnitude of the over-estimation. Both voxel and ROI values for the Reference Normalised SAR are closer to unity than that computed for the raw thermometer data, however this is not of statistical significance (p-values of 0.48 and 0.43, respectively from a two-sample t-test). Assuming that this over-estimation factor is the same for all voxels (that have not been discounted because of detected errors), this systematic error could be corrected, although additional work would be necessary to determine its source and nature. Moving the thermometer probe closer to the coils so that it experiences more heating (therefore higher SNR MRTI measurements) would be one such way to achieve this; if the over-estimation reduces then SNR is a factor.

6.3.4.2 Susceptibility instabilities

Susceptibility gradient induced instabilities were found to be present in the coil 1 and in-phase configuration data, manifesting as an artefact in both the magnitude and phase image data which changed over the course of each acquisition, as shown in figure 6.11a-c. These artefacts were present on data acquired during both heating and equalisation scans. For the coil 1 data the artefact is present in a hole made by the fibre optic thermometer on a previous date. This hole extended through almost the entire length of the phantom, producing an inhomogeneity which could not be corrected by the shims. For the in-phase data the location of this artefact is towards the front of the phantom, close to the opening where it is exposed to the air. In both cases the agar gel may have experienced some "settling", or the water that was not trapped within the gel moved away from these regions, creating a small air pocket and subsequent changing susceptibility gradient causing a localised B_0 inhomogeneity which could not be compensated for using the first order B_0 drift correction described in section 6.3.1.4. In some voxels the time course of the phase data was of a similar shape to that from heating (an exponential curve), albeit the phase difference was much greater, and so the equation 6.8 could be fit with a sum of squares error less than 2 and

would result in a miscalculation of ΔT and the SAR. Fortuitously these regions do not occur in areas where there is peak SAR (i.e. close proximity to the coil conductors), therefore with these regions masked out the data is still valid and representative of the maximum delivered SAR in the corresponding configuration.

It is important to address why the phase changes in this region are attributed to susceptibility instabilities, and not a region of hyper-intense heating caused by the superposition of the electric fields of the two transmitting coils in this particular configuration. Evidence that shows the latter is not the case is:

1. These artefacts are present on data acquired when no RF was transmitted by the coils.

2. The artefacts are clearly present in the magnitude image data, and are shown to change over time, showing evidence of intra-voxel dephasing. With the imaging parameters used, a temperature gradient greater of $12.72°C$ across each voxel is required to cause dephasing; many times the maximum temperature change measured within the phantom.

3. Localised phase wrapping is seen in the phase image, which usually is a sign of a B_0 inhomogeneity.

4. The ΔT image shows regions of negative temperature change immediately next to large positive temperature changes. This temperature distribution is not physically possible after RF heating, and over the length of the measurement (one hour) the heat would have diffused throughout the phantom and its effect would not only be seen within this region. Beyond this region both the ΔT and SAR images are similar to those measured in the quadrature and anti-phase configuration.

5. If high localised electric fields were generated from the interaction of the two RF coils, this effect, perhaps not in the same position or intensity should also be seen in the quadrature and anti-phase configurations. However both the ΔT and SAR maps for these configurations only show heating close to the coils.

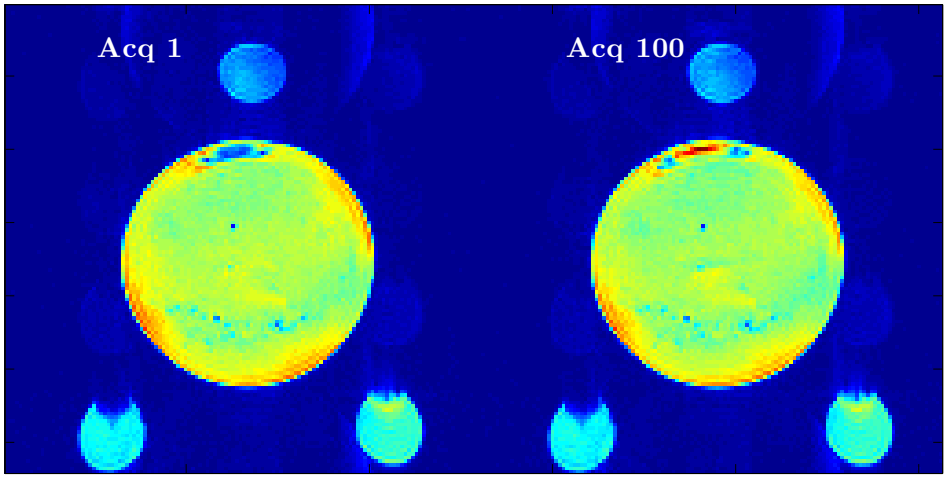

a. Magnitude

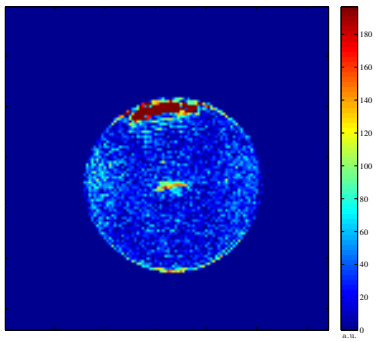

b. Absolute difference between the first and last magnitude images

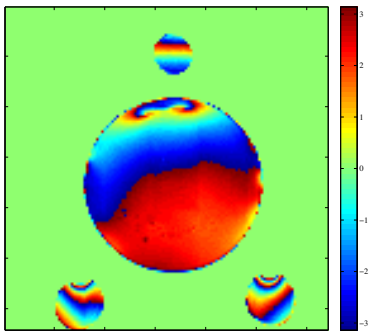

c. Phase

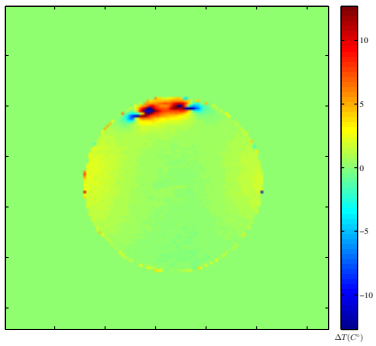

d. ΔT

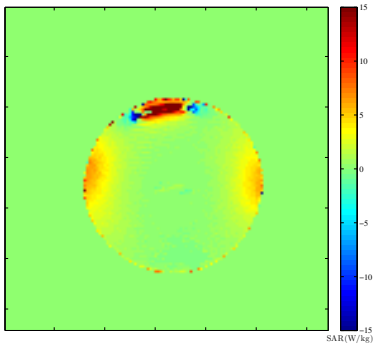

SAR

Figure 6.11: Errors in the In-Phase data, run 1, slice 12

6.3.4.3 SAR Results

Both the SAR and ΔT measured when coil 1 was transmitting independently (figure 6.7.a for ΔT and 6.8. for SAR maps) was found to be reduced in comparison to what was measured for coil 2 (figures 6.7.b and 6.8.b), and also the equivalent ΔT and SAR distributions with both coils transmitting in the quadrature and anti-phase configurations. This is also evident in table 6.4. A possible cause is that coil 1 was not positioned optimally on the agar phantom, decreasing the amount of absorbed power. The measured forward voltage was found to be the same in this configuration as in others, highlighting the need to ensure the coils are placed identically each time. One such, albeit impractical way would have been to fix the coils to the phantom and not remove them until measurements had been made for all configurations.

The power deposition varies with the phase difference between the two coils as a result of both the interference of the RF fields produced by the coils, and inter-coil coupling. The anti-phase configuration has an elevated SAR and temperature rise towards the centre of the phantom compared to the other two configurations, indicating constructive interference of the RF fields. Conversely the in-phase configuration has a reduced SAR in the centre of the phantom in comparison to both the anti-phase and quadrature configurations, indicating destructive interference of the RF fields. Inter-coil coupling appears to cause an imbalance in the power deposition between the two coils: the in-phase data shows a more intense SAR from the left coil compared to the right coil, whereas both the quadrature and anti-phase configuration data show the right coil to have a more intense SAR, however this may also be due to non-optimal coil positioning.

6.3.4.4 Safety Limits

In contrast to an agar phantom, the human neck has a heterogeneous structure with different tissue conductivities and permittivities which will affect the SAR distribution. Also it is highly perfused; both the carotid arteries and jugular veins transport a large quantity of blood through the neck, giving it the ability to dissipate heat well. Phantom conductivity was considerably higher than that of biological tissues due to the combination of salts used (0.9% NaCl and 1.5% $CuSO_4$), hence larger RF eddy currents are induced within the phantom than would be in a human neck. A subcutaneous layer of fat present on the neck, which has significantly lower conductivity than muscle (see figure 6.1 and figure 6.2), reduces the power deposition close to the coil conductors where the SAR would be at its most intense. All of these factors should mean that power deposition within the human neck, and any associated heating is less severe than in the agar phantom.

The maximum peak SAR, measured in the in-phase configuration was $8.78 \pm 0.38 W/kg$ over the entire phantom volume, and $6.23 \pm 0.39 W/kg$ when only accounting for histogram bins with more than 25 voxels. Both of these measures are within standard regulatory limits for local SAR, which are are defined as the maximum SAR within 1g (FDA) and 10g (IEC) cube shaped averaging

masses, typically calculated by spatially averaging the SAR maps using a spatially adaptive filter [34]. The standard method expands a cubic region around each voxel until a mass limit is reached, then averages all of the contributing voxels. The MRTI data was acquired with a voxel size of $1.6 \times 1.6 \times 2.5mm$. Assuming the agar phantom had a density of water, and the entire volume of the voxel was filled with agar gel, each voxel had a mass of 0.0061g, therefore the number of voxels to average over for the 1g and 10g masses are approximately 163 and 1638, respectively. 25 voxels have a mass of 0.153g. Therefore, were this averaging process applied the peak SAR values would be averaged over many voxels, of which it is likely that a large proportion will have a significantly lower SAR value, further reducing the peak SAR and associated heating.

Results show that at 100% duty cycle the ASL Labelling Coils do not present a thermal hazard as the SAR is well within regulatory limits. The 3D-GRASE CASL sequence described in section 7.3 is typically run with TR = $4350ms$ and a typical/maximum labelling durations of 2/4 seconds, resulting in a RF duty cycles of 46% and 91%. Using the maximum measured SAR over the entire phantom volume, 8.8W/kg, this results in peak SAR's of 4.0W/kg and 8.0W/kg. As this value is not from averaging it provides a conservative safety margin from which safe limits can be defined. Assuming that the non-labelling RF pulses used in the imaging sequence are homogeneous, the peak maximum peak SAR can be calculated by summing the labelling SAR and the SAR reported by the MRI scanner. For a background-suppressed 3D GRASE sequence this is typically 10%, 0.32W/kg; for typical imaging parameters both the 2 second and 4 second labelling pulses can be used without exceeding SAR limits. Likewise the maximum average SAR over the entire phantom from the labelling coils was calculated to be approximately 1.1W/kg, with the addition of the imaging RF pulses the SAR becomes approximately 1.4W/kg. As the head is being imaged, only the head, neck and shoulders receive a significant RF dose (as the rest of the body is outside the body coil), this is less than half the regulatory SAR limit of 3.2W/kg.

With the coils positioned opposite each other on the 12cm agar phantom and a phase shift such that their B_1^+ fields constructively interfere, there is an increase in power deposition in the gel between the two coils in comparison to a non-constructive phase shift or single coil. Because of the skin effect this is most noticeable at the top and bottom of the phantom. The coil diameter, and hence B_1^+ penetration of the coil is approximately equal to the radius of the phantom, therefore the increase in power deposition is small compared to the power deposition close to the coils, and has no effect on the peak SAR (although does affect the average SAR). On a smaller phantom/neck the coils would be positioned closer together, and regions of higher intensity B_1^+ will interfere, increasing the SAR. Additionally, for smaller radii the skin effect is reduced, therefore the intensity of the RF fields is greater at the centre, however the flux linkage through the phantom is less, reducing the intensity of the induced eddy currents.

To estimate the minimum permitted separation of the coils the SAR map from two coils placed on opposite sides of a cylinder was numerically calculated for a range of cylinder diameters. A

single slice SAR map was synthesized for each coil by assuming that it had the same spatial distribution as the magnetic field generated by a circular loop of wire of the same dimensions (3.25cm diameter), computed by numerically integrating the Biot-Savart Law; as discussed in the introduction to this chapter this is not valid because of the skin effect and wavelength effects, however for this purpose it serves as a suitable approximation. SAR maps for each coil within the cylinder were normalised to 8.78W/kg, and then summed to give the combined SAR. This was repeated for cylinder diameters ranging from 1cm to 12cm in 1cm increments. The maximum SAR for each coil separation is presented in figure 6.12. The minimum permitted separation that maintains a peak SAR of 10W/kg is 5.4cm, after which the peak SAR rapidly increases. Therefore, for neck widths which range between 6 and 12cm, which covers most adults, the local SAR is within the 10W/kg limit.

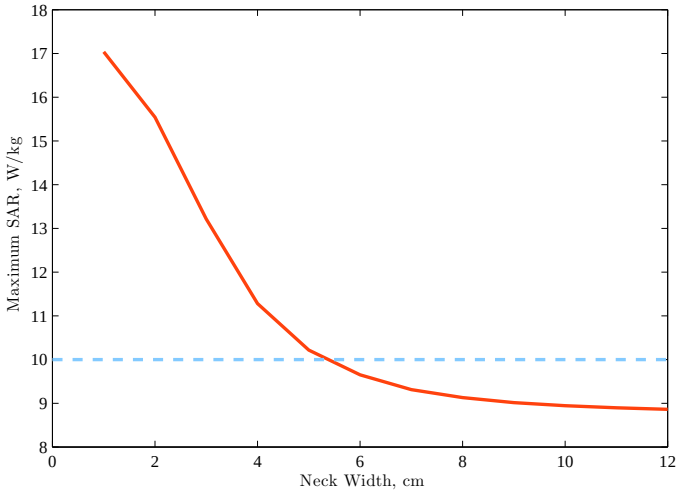

Figure 6.12: Maximum SAR as a function of coil separation

6.3.4.5 Further Work

At present this method is in its initial stages of development, and although it has proved to be suitable for the safety validation of the ASL Labelling Coils, there is much scope for improvement and optimisation:

1. An analysis of the errors in the MRTI measurements, particularly how the noise in the raw phase data, and the sum-of-squares residuals relate to the error in the SAR measurements.

2. Repeated measurements with the same coil configurations to estimate the repeatability of the method, and how factors such as equalisation scan time affect temperature and SAR measurements.

3. An investigation into factors affecting the temperature-compensated B_0 scheme, such as length of the equalisation scans, fitting higher order 2D polynomials to the reference phase. B_0 field maps could be acquired using a double echo EPI sequence to determine the nature of

the drift. Also determine how much data is required for an accurate fit; MRTI measurements might only be necessary for a fraction of this time to achieve comparable SAR accuracy, accelerating the acquisition time.

4. Place the agar phantom fibre optic thermometer probe in a region which receives a high SAR dose to further determine the accuracy and validity of the MRTI SAR measurements, particularly whether the over-estimation discussed in section 6.3.4.1 is due to low SNR in the centre of the phantom (where power deposition is minimal).

5. Compare SAR maps obtained from PRF MRTI measurements with those generated numerically by EM Modelling, both for the homogeneous agar phantom, and a tissue sample, for example a human or animal cadaver.

6.3.5 Conclusion

Presented in this chapter is a method for measuring the SAR of the ASL Labelling Coils using PRF MRTI. Because of the intrinsic low power of the coils, the observed temperature change is SNR limited, therefore data was acquired for an extended duration and a suitable function fit to the data to improve the accuracy of the calculated SAR. A 3D, high resolution (128×128 in plane resolution, 48 slices) EPI sequence was used to acquire gradient echo phase images of an agar gel phantom, interleaved with 30 seconds of continuous wave RF power delivered by the ASL Labelling Coils, resulting in time course temperature measurements. To correct for drifts in the B_0 field, the temperature compensated phase within surrounding reference phantoms was used to generate linear drift correction maps, and equalisation scans lasting up to 5 hours were run prior to MRTI measurements commencing to ensure thermal equilibrium of the MR system and phantoms. Temperature data was fit with a suitable function to remove noise, resulting in MRTI measurements accurate to within $0.1°C$ of fibre optic thermometer measurements. SAR maps were then calculated from the fitted MRTI data, and data consistency algorithms were used to automatically discard regions with unreliable SAR values.

As SAR is calculated from the initial temperature increase, this must be sufficiently noise free, however for low power coils the temperature increase is small and significantly affected by noise. Previous studies [41, 168, 167] have been able to use high power transmitter systems to deposit more power than in this study, subsequently heating the sample more, however in this instance it was not practical to do so. Using the method presented the high incident power requirement is removed, and the fast EPI sequence permits the acquisition of full volumetric data in a similar time as a single gradient echo slice. There is also scope for this technique to find use in measuring small temperature changes and the corresponding SAR from implantable devices, and communications equipment such as mobile telephones, which would normally have their SAR measured electric field probes within a tissue phantom.

The SAR maps obtained showed a maximum peak local SAR of $8.78 \pm 0.38 W/kg$. As this is the value within one single voxel, with respect to the IEC's limit of 10W/kg in 10g of tissue, the actual spatially averaged local SAR will be lower, providing a convenient conservative safety margin. As such, no restrictions need to be placed on the maximum RF duty cycle when using the ASL Labelling Coils, unless the combined ASL Labelling Coil SAR and the SAR produced by other pulses in the sequence exceeds the regulatory limit.

Chapter 7

Parallel Transmit CASL Implementation

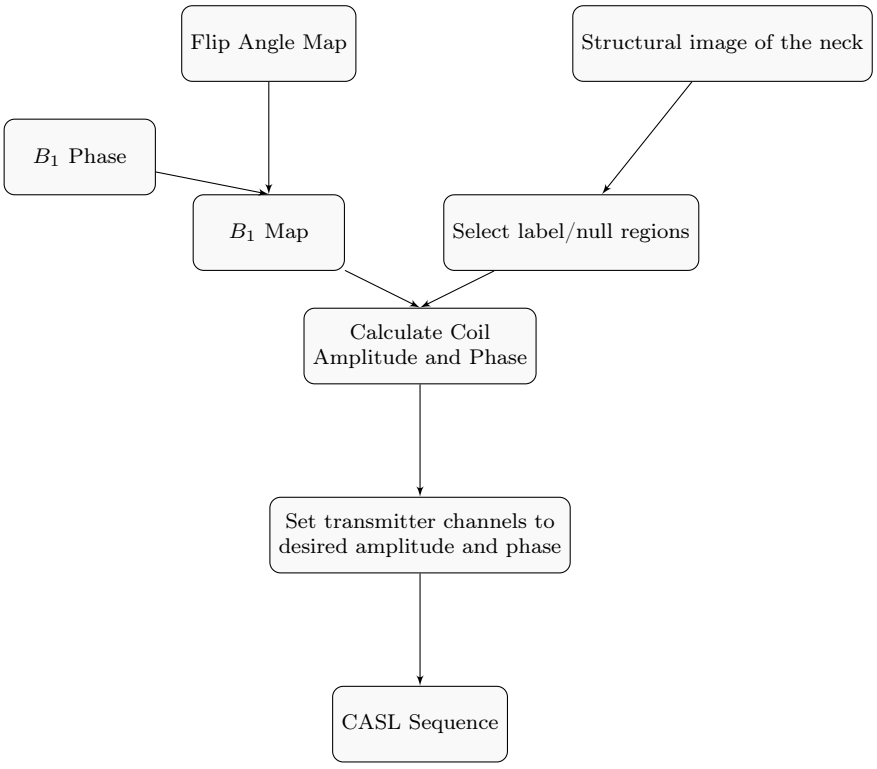

Figure 7.1: Parallel Transmit CASL Procedure

In this chapter the components which make up the implementation of parallel transmit CASL are detailed. Figure 7.1 illustrates the process required to achieve parallel transmit CASL. Knowledge is required of the vector B_1^+ maps for each coil, and the anatomical locations of the carotid and vertebral arteries within the neck. Both the B_1^+ maps and artery locations vary on a subject by subject basis and consequently must be acquired as part of a pre-scan routine. Appropriate label/null regions are selected, then the amplitude and phase of each channel determined using

the mathematical framework presented in section 3.2.4. Finally each transmitter channel is set accordingly for use in a continuous ASL sequence.

Section 7.2 details the sequence and processing required to obtain vector B_1^+ maps of both transmitter coils. Section 7.3 describes the continuous arterial spin labelling sequence used to acquire perfusion images with the ASL Labelling Coils. Section 7.4 presents a novel perfusion phantom with distinct vascular territories used for testing parallel transmit CASL, designed to simulate both the macroscopic flow within the arteries and incoherent microscopic flow at the capillary bed. Section 7.5 exhibits custom scan planning software which quickly and robustly performs the required image processing of the B_1^+maps and anatomical scans, and with which the scan operator can plan the parallel transmit CASL scans and obtain the appropriate driving amplitudes and phases.

7.1 Inter-Coil Coupling Issues

Multiple RF coils tuned to the same frequency placed in close proximity exhibit strong coupling, resulting in the efficient transfer of energy from one coil to another and the generation of coupled modes, splitting the resonant peaks of the coils. In other words, RF currents on one coil generate a magnetic field, which induces an EMF and subsequent RF currents on other coils. If the coils are used for reception, inter-coil coupling causes a loss in SNR due to detuning and mismatching of the coils from the resonance peak splitting, and couples coherent noise between multiple channels which does not cancel when signals from the coil elements are combined. Coupled transmitter coils suffer from the inability to produce a prescribed current distribution. Fortunately there are many decoupling methods available for both transmit and receive coils. Neighbouring elements can be overlapped to cancel their mutual inductance [194]; low input impedance preamplifiers prevent the induction of RF currents in receive only coils by effectively creating an open circuit at the coil's terminals [194]; off-coil 2N-port decoupling interfaces can be used to cancel mutual coupling [138]; cartesian feedback corrects for induced RF currents on coil elements by applying specially phased negative feedback [113]; current source amplifiers on transmit coils produce a specified current regardless of loading conditions, rather than produce a specified voltage in the manner that conventional power amplifiers operate[135]; vector decoupling measures the coupling between coil elements and applies a corresponding amplitude and phase shift in the digital domain to account for this [199, 127].

In chapter 4 the coupling between the two ASL Labelling Coils whilst being loaded by an agar phantom was measured to be approximately -24dB. This was assumed sufficient enough to not need any further decoupling mechanisms, however MR results indicate that coupling is present, as shown in the perfusion images of the perfusion phantom (section 7.4) in figure 7.2. Images were acquired according to the method described in section 7.4.3.2. With the left coil only transmitting

(PIN diode in the right side coil reverse biased, detuning the coil) there is a small amount of contralateral labelling due to the left coil's B_1^+ field extending across to the right inflow pipe (figure 7.2.a). Forward biasing the PIN diode to tune the right coil, yet still only applying power to the left coil results RF power is coupled to the right coil, which then generates its own B_1^+ field and increases the degree of labelling within the right inflow pipe, indicated by the increased intensity in the right perfusion territory in figure 7.2.b.

Most problematic is that prescribed amplitudes and phases are not correctly realised, affecting the ability to create specified linear combinations B_1^+ maps (section 7.2.3.2), and to calculate and set the transmitter phase in parallel transmit CASL measurements. Corrupt linear combinations, where the specified amplitudes and phases are not met, cannot be correctly inverse-transformed to single coil maps, because the inverse transformation matrix is not valid. Consequently the resultant single coil maps are also corrupt, affecting the calculated complex driving coefficients from which the amplitude and phase of each channel is set. Furthermore, these incorrect driving coefficients are then not realised by the transmitter channels.

The implication of this is that the automatic calculation of the appropriate amplitudes and phases for parallel transmit CASL cannot be performed. However, as shown in chapter 8, with only two transmitter coils it is possible to iteratively find suitable coil driving conditions to place a null in the labelling B_1^+ field over one of the carotid arteries. The rest of this chapter presents the implementation of parallel transmit CASL assuming that the RF coils used for labelling are sufficiently decoupled.

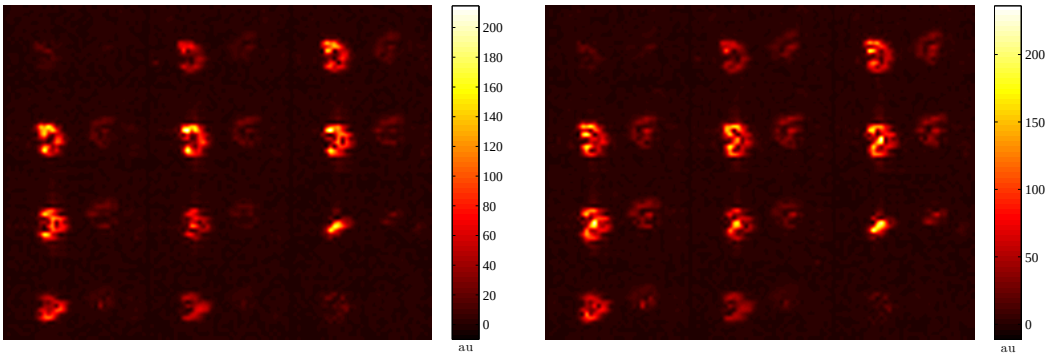

a. Left coil only b. Left coil full power, right coil tuned

Figure 7.2: Evidence of inter-coil coupling

7.2 B_1^+ Mapping

Knowledge of the spatial distribution of both the magnitude and phase of the transmit field from each coil is required to successfully utilise parallel transmission and spatially tailor the total B_1^+ field. B_1^+ mapping methods are used to acquire images which can be further processed to remove image contrast from all mechanisms (such as proton density, T_1, T_2, receiver sensitivity) except the flip angle.

Many methods exist to map the transmit field, for example the double angle method (DAM) and its variants [117, 43, 206], spin echo and stimulated echo image pairs [7], EPI readout single shot stimulated echo acquisition mode [98], steady state gradient echo with two interleaved TR's, also known as Actual Flip angle Imaging (AFI) [248, 162].

An AFI sequence had been previously implemented and tested on the Siemens Trio scanner by Dr Antoine Lutti at the Wellcome Trust Centre for Cognitive Neuroimaging, University College London[146]. The AFI sequence is a variant of a 3D spoiled gradient echo sequence, where two echoes with different TR's are acquired in an interleaved manner. It is well suited to mapping the B_1^+ of the ASL coils for the following reasons:

1. As it is a three dimensional (3D) acquisition, non-selective pulses may be used, which the SMIS spectrometer is capable of producing. Due to timing delays and inaccuracies, and power amplifier compression, as discussed in sections 5.2.1 and 5.5.3 respectively, a slice selective sinc pulse would not be correctly formed to produce well defined slices, ruling out multislice acquisition methods.

2. Signal formation is gradient echo, avoiding any issues associated with an accurate spin echo due to timing delays.

3. Only a single flip angle is required for both periods of the AFI sequence as opposed to the DAM, simplifying the operation of the SMIS spectrometer pulse sequence.

4. Acquisitions can be accelerated using partial Fourier and parallel imaging: an entire 3D volume can be acquired in less than five minutes.

7.2.1 3D Actual Flip angle Imaging Sequence

The 3D AFI sequence consists of two interleaved FLASH acquisitions with repetition times TR_1 and TR_2. A fast repetition rate is assumed such that $\mathrm{TR}_1 < \mathrm{TR}_2 < T_1$, establishing a pulsed steady state of magnetisation. Provided all coherent transverse magnetisation is destroyed at the end of each TR period the longitudinal magnetisation prior to each excitation pulse can be expressed as:

$$M_{z1} = M_0 \frac{1 - E_2 + (1 - E_1)E_2 \cos\alpha}{1 - E_1 E_2 \cos^2\alpha} \tag{7.1}$$

$$M_{z2} = M_0 \frac{1 - E_1 + (1 - E_2)E_1 \cos\alpha}{1 - E_1 E_2 \cos^2\alpha} \tag{7.2}$$

Where $E_{1,2} = e^{-\mathrm{TR}_{1,2}/T_1}$, M_0 is the equilibrium magnetisation, and α is the flip angle. Following an excitation pulse the the observed signals S_1 and S_2 at the echo time TE are given by:

$$S_{1,2} = M_{z1,2} e^{-\mathrm{TE}/T_2^*} \sin\alpha \tag{7.3}$$

The ratio of the two signals r, can be expressed as:

$$r = \frac{S_2}{S_1} = \frac{1 - E_1 + (1 - E_2)E_1 \cos\alpha}{1 - E_2 + (1 - E_1)E_2 \cos\alpha} \tag{7.4}$$

As $\mathrm{TR}_2 > \mathrm{TR}_1$, the signal intensity of a voxel in S_1 will always be greater than in S_2 and hence $r < 1$. For short TR_1 and TR_2 this ratio can be simplified to a first order approximation:

$$r \approx \frac{1 + n\cos\alpha}{n + \cos\alpha} \tag{7.5}$$

Where $n = {}^{\mathrm{TR}_2}/_{\mathrm{TR}_1}$. Hence, the flip angle can be calculated by:

$$\alpha \approx \cos^{-1}\left(\frac{rn - 1}{n - r}\right) \tag{7.6}$$

Practical implementation of this technique is detailed in [248, 162, 146]. Care must be taken to ensure that equation 7.6 is a reasonable approximation to the flip angle for the range of T_1 relaxation times observed in tissue. Simulations detailed in [248] show that there is minimal deviation between the first order approximation of the ratio and correct expression (equation 7.4) for a range of sequence parameters. The first order approximation to r is to a high degree T_1 insensitive, the exception being for high flip angles with short T_1 species (where $T_1 \sim \mathrm{TR}_2$). Sensitivity to the range of flip angles is determined by the ratio of TR periods n; an increase in n improves the dynamic range of r by increasing the sensitivity to flip angle variation of the intensity reduction of S_2 with respect to S_1, however larger values of n increase scan time.

As the AFI approach is based on the assumption that the transverse magnetisation is completely spoiled prior to each RF pulse it is important that adequate spoiling is implemented. In [162] the steady state properties of the AFI sequence are quantitatively and qualitatively investigated, resulting in a novel RF and gradient spoiling scheme. Standard RF spoiling [264], where the spoil angle increment is the same for each pulse repetition, is only applicable to equidistant pulse spacing. A new RF spoiling scheme was proposed by transforming the alternately spaced pulses into a series of virtual equidistant pulses, such that the two physical RF pulses are separated by $N_1 - 1$ and $N_2 - 1$ virtual RF pulses with zero flip angle and period $\mathrm{TR}' = {}^{\mathrm{TR}_1}/_{N_1}$, where the ratio ${}^{\mathrm{TR}_1}/_{\mathrm{TR}_2} = {}^{N_1}/_{N_2}$ is a rational number (both N_1 and N_2 integers). In order to establish a proper steady state the RF spoiling phase shift must be applied to both real and virtual pulses, cycling the spoil angle increment between two values.

In standard RF spoiling, each excitation pulse has a phase that has been incremented by a spoil angle which increments linearly with each repetition by the spoil increment $\Delta\phi$, i.e for the k^{th} pulse the transmitter phase ϕ_k will be:

$$\phi_k = \phi_{k-1} + k\Delta\phi \tag{7.7}$$

In AFI, the spoil increment alternates with each TR period. For pulses prior to the TR_1 period the spoil increment is $\Delta\phi_1 = \Delta\phi\frac{N_1}{N_2}$, and for pulses prior to the TR_2 period the spoil increment is given by $\Delta\phi_2 = n\Delta\phi\frac{N_1}{N_2} = \Delta\phi$.

Gradient spoiling was also analysed in the equidistant representation. Additional gradient spoiling was found to be required during the TR_2 cycle, with area $\Delta\mathbf{k}$ equal to:

$$\Delta\mathbf{k} = \frac{(TR_2 - TR_1)}{TR_1}\mathbf{k}_0 \tag{7.8}$$

Where $\mathbf{k}_0$ is the gradient area of the spoiling gradient during the TR_1 interval.

In [162] the effect of diffusion damping on the AFI steady state was also investigated, and found that spoiling improved with increased diffusion damping. Consequently, Lutti et al. [146] introduced flow-compensated diffusion weighted spoiler gradients (in addition to existing spoiler gradients), employing diffusion damping for the additional spoiling of transverse magnetisation.

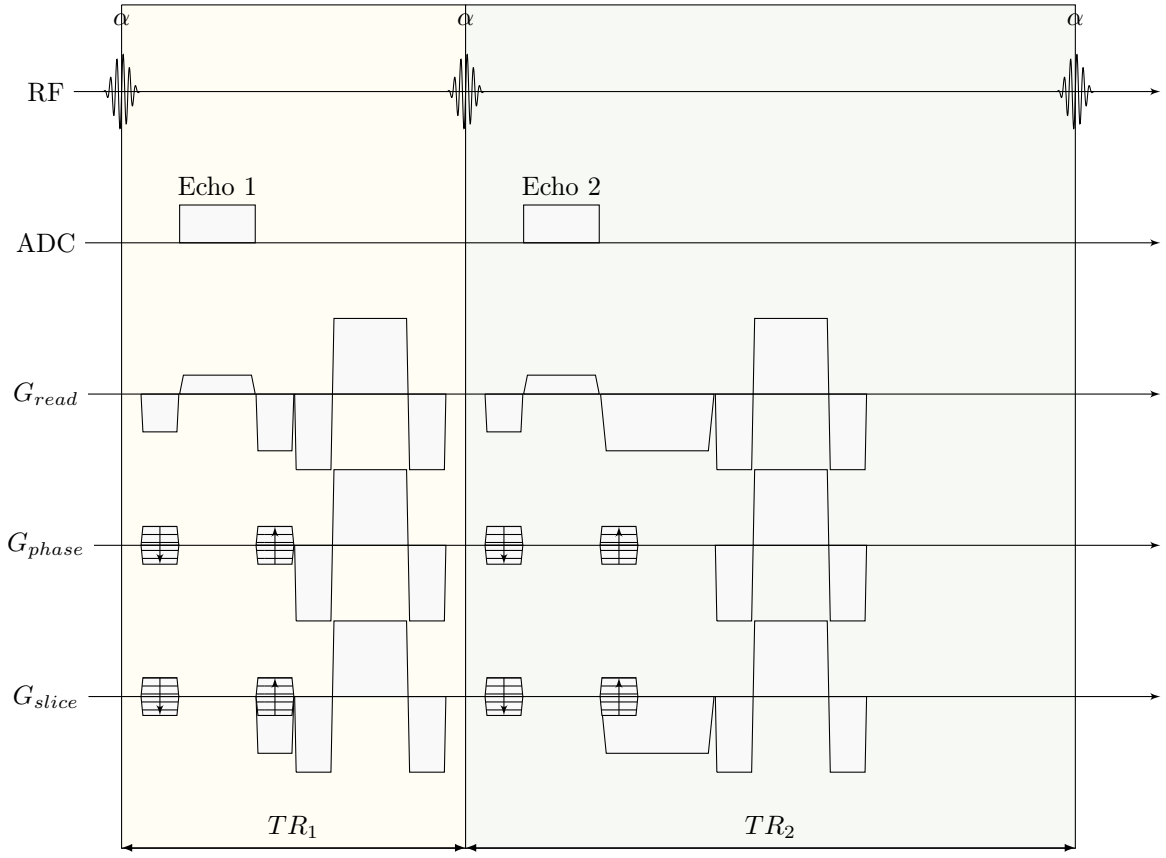

Figure 7.3: AFI Pulse Sequence

Figure 7.3 is a pulse sequence diagram showing the two cycles of the AFI sequence. As it is a 3D sequence, phase encoding was performed on both the phase encode and slice select logical

gradient channels. Cycle one, with TR = TR_1 is highlighted within the pale yellow box, and cycle two, with TR = TR_2 is highlighted within the pale green box. Cycle one and cycle two are each run once for each phase encoding step. Following each excitation pulse there are phase encoding gradients in the phase and slice directions, and a read dephase gradient. An echo time as short as possible is typically used as no relaxation based contrast is necessary, and to minimise T_2^* decay. After the readout, phase encode refocusing gradients occur on the two phase encoding channels at the same time as spoiler gradients in the read and slice directions; in cycle two the gradient moment is larger to account for the increased TR time, as given by equation 7.8. On the slice channel the resultant gradient waveform is the combination of the spoiler and phase encode refocusing pulses. Finally, tri-lobe diffusion weighted flow compensated gradient pulses with relative duration ratio 1:2:1 follow each spoiler gradient.

The AFI pulse sequence code was modified to produce an optical signal to trigger the SMIS spectrometer instead of playing out an RF pulse on the Trio. Sequence parameters similar to those used in [146] were used: TE = $3.05ms$, $26mT/m$ spoiler gradient amplitude, $30mT/m$ diffusion gradient amplitude, and $42ms$ total diffusion gradient duration. As a result TR_1 was set to $50ms$ and TR_2 set to 150ms, resulting in $n = 3$. The RF spoiling phase increment was $\Delta\phi = 129.6°$ [1], and spoiler gradient durations were $1ms$ in cycle one and $3ms$ in cycle two. These parameters permit RF excitation pulses up to $840\mu s$ long.

7.2.2 Vector B_1 Maps

Magnitude and phase information of the transmitted B_1^+ field is required for both of the ASL Labelling Coils. Typically AFI data is acquired whilst pulsing RF on each coil individually. Magnitude and phase information can then be calculated using the following methods:

Magnitude With knowledge of the excitation pulse shape and duration, the magnitude of the B_1^+ field can be computed from an AFI flip angle map, α. In the small tip angle regime the flip angle due to an RF pulse of duration τ with temporal waveform $B_1(t)$ is given by:

$$\alpha = \gamma \int_0^\tau B_1(t)dt \tag{7.9}$$

For a square pulse of amplitude B_1 this simplifies to

$$\alpha = \gamma B_1 \tau \tag{7.10}$$

Rearranging yields the B_1 magnitude:

[1] [162] and [146] both use 129.3°, however the SMIS spectrometer is limited to a phase resolution of 0.225° which sets a constraint on appropriate values for the spoil increments in cycles one and two, $\Delta\phi_1$ and $\Delta\phi_2$ as they must both be multiples of 0.225°.

$$B_1 = \frac{\alpha}{\gamma\tau} \tag{7.11}$$

Each AFI dataset consisted of two volumes. The voxel-wise ratio of the two volumes was calculated, and then the flip angle map computed using equation 7.6. The B_1^+ magnitude was then calculated by applying equation 7.11 to each voxel.

Phase A gradient echo image has phase φ which depends on the phase of the excitation B_1, φ_{B_1}, B_0 field inhomogeneities, $\varphi_{\Delta B_0}$ and other residual sources such as eddy currents and k-space centre offsets, φ_{res}:

$$\varphi = \varphi_{B_1} + \varphi_{\Delta B_0} + \varphi_{res} \tag{7.12}$$

$\varphi_{\Delta B_0}$ and φ_{res} are sequence dependent, and can be assumed to be the same for all images acquired with the same sequence and parameters. The relative phase between each RF channel can be obtained by subtracting a reference phase to remove the additional sources of phase accrual. Therefore the relative B_1^+phase, Φ_{B_1} is calculated from the cycle 1 AFI phase maps, according to:

$$\Phi_{B_1,n} = \frac{e^{i\varphi_n}}{e^{i\varphi_1}} \tag{7.13}$$

Where φ_1 is the phase from the cycle 1 AFI image of coil 1, and n is the coil index. Subsequently the relative phase map of coil 1 will be 0° everywhere.

With both the magnitude and phase of each coil calculated the vector field is given by:

$$\mathbf{B_{1,n}^+}(\mathbf{r}) = \frac{|\alpha_n(\mathbf{r})|}{\gamma\tau} e^{i\Phi_{B_{1,n}}(\mathbf{r})} \tag{7.14}$$

Where $\mathbf{r}$ is denotes spatial position.

7.2.3 Improving the SNR of B_1^+ maps

Parallel transmission CASL, described in section 3.3, uses the superposition of B_1^+ fields from multiple coils to spatially tailor the labelling B_1^+ field. Take the example of two coils placed either side of the neck, with the aim of labelling the right carotid artery only. The right hand side coil can be considered the labelling coil, as its power mostly contributes towards ensuring the B_1^+ field at the right carotid artery is at a target value. Its B_1^+ field extends across the neck to the left carotid artery, causing partial labelling, so the second contralateral coil is used to generate a field which at the position of the left carotid artery opposes the field generated by the labelling coil. As both coils are small the magnitude of the B_1^+ field due to the labelling field at the left carotid artery is low, perhaps a tenth of the magnitude at the right carotid artery. In order to correctly oppose the labelling field at the left carotid artery it is important to accurately know the magnitude of this

labelling field, however in an acquired B_1^+ map this corresponds to a region where excitation and hence SNR is low.

For a voxel n, the error in the ratio image (given by equation 7.4), σ_r is:

$$\sigma_{n,i} = \sqrt{\left(\frac{\sigma_1}{S_{1,n}}\right)^2 + \left(\frac{\sigma_2}{S_{2,n}}\right)^2}\frac{S_{2,n}}{S_{1,n}} \tag{7.15}$$

Where $S_{1/2,n}$ are the image intensity at voxel n, and $\sigma_{1/2}$ is the error in the AFI images acquired in cycles 1 and 2, respectively. As the SNR of the AFI images decreases due to reduced signal intensity the error in the ratio image increases. Furthermore, in regions where noise is a significant proportion of the image intensity, the ratio S_2/s_1 may become greater than unity, amplifying the associated error, and, for values greater than 1 the inverse cosine function returns an imaginary number, introducing additional errors in the phase of the B_1^+ map.

7.2.3.1 B_1^+ Map Simulator

To asses methods to improve the SNR of acquired B_1^+ maps, simulated data was generated in Matlab (The Mathworks Inc., Natick, MA) using the AFI signal equations (equation 7.3). Figure 7.4 is a flow chart of the B_1^+ map simulation procedure. A rasterised digital phantom was created based on a modified Shepp-Logan phantom. T_1 contrast was provided by segmenting each region of the phantom and assigning a T_1 value, and uniform proton density contrast by setting the equilibrium magnetisation to $M_0 = 1$ within the phantom and $M_0 = 0$ elsewhere.

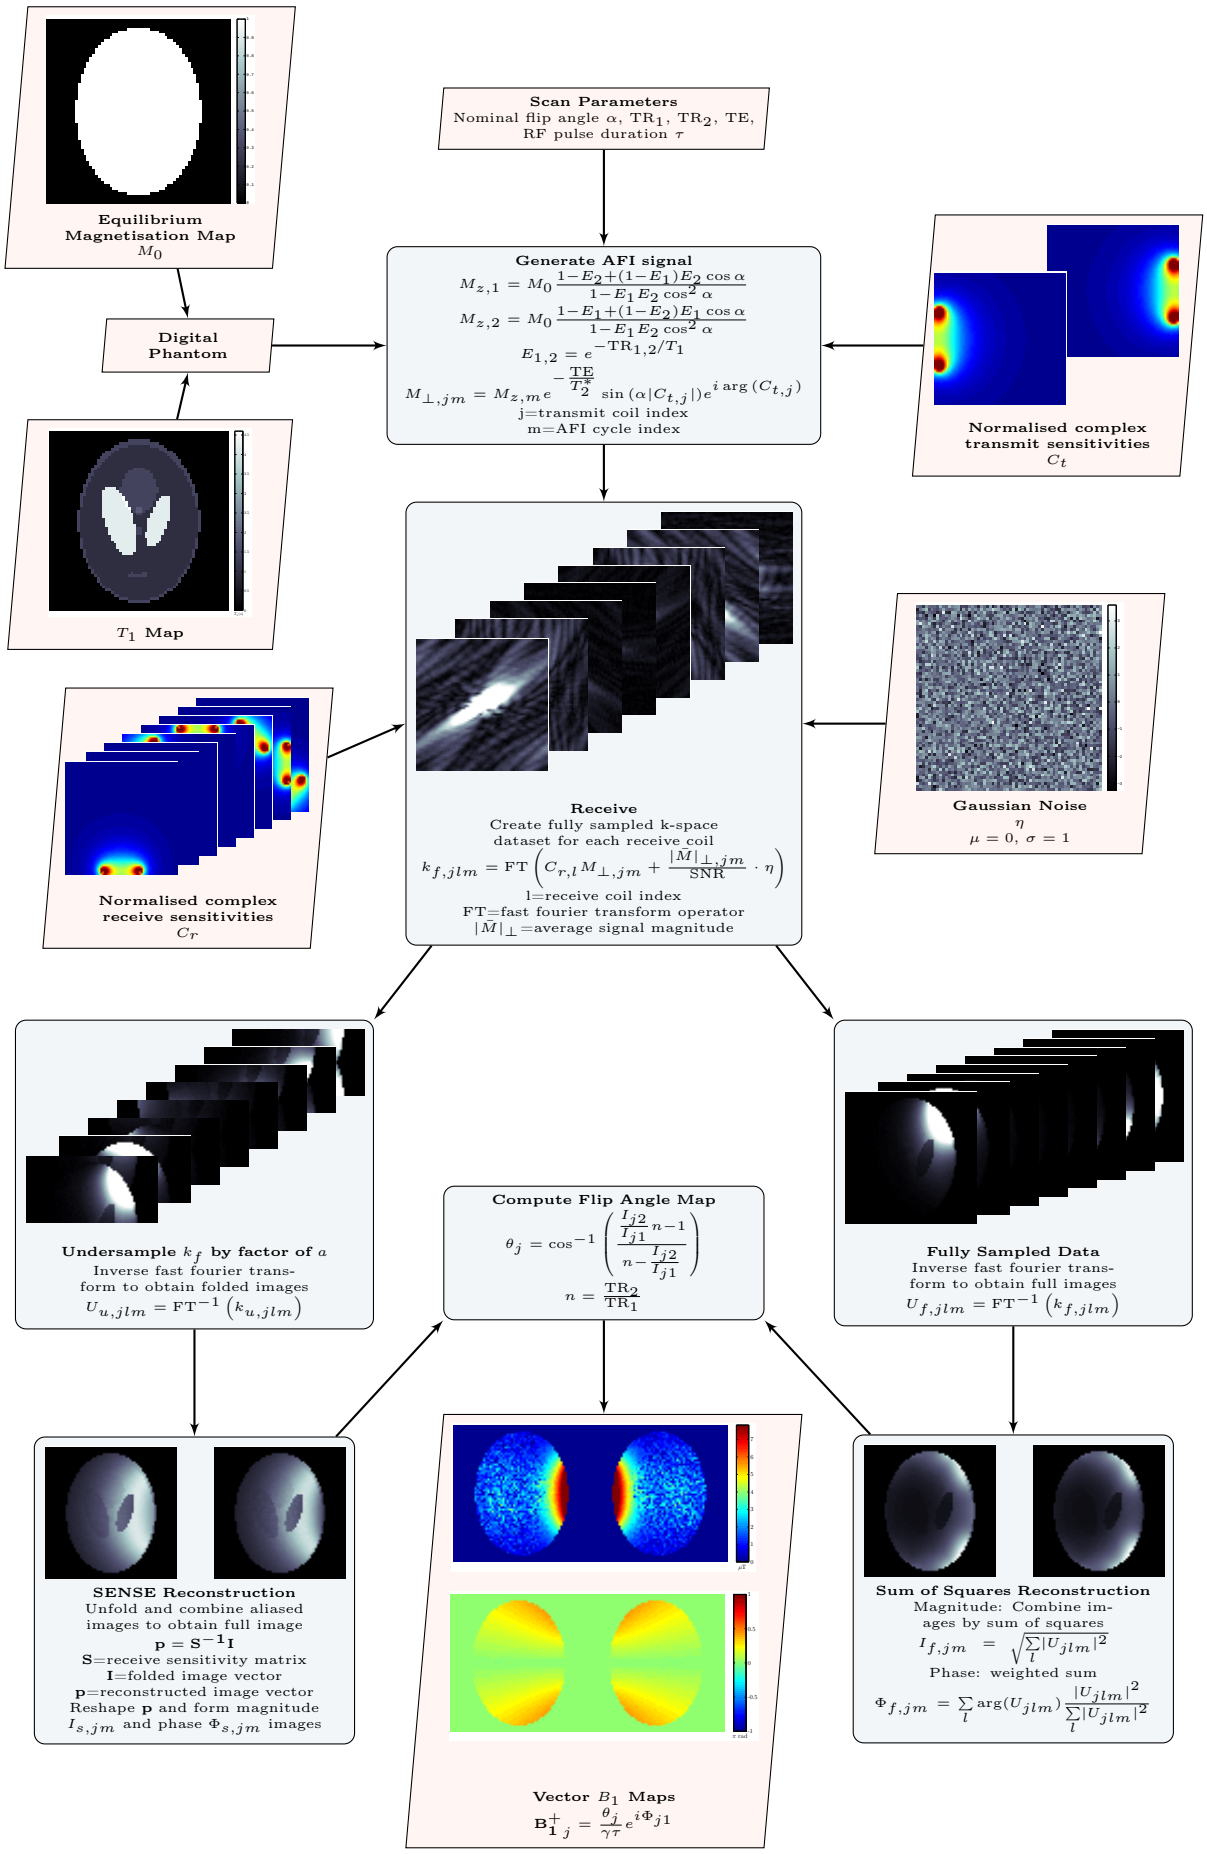

Figure 7.4: B_1^+ Map simulation procedure

Normalised complex transmit and receive coil spatial sensitivity distributions were synthesised for circular loop coil elements using an elliptical integral formulation of the Biot-Savart law (p266, [207]). All coils were positioned such that their axis was on the x-y plane and passed through the origin, and their B_1^+ was normalised to unity at a depth equal to a quarter of the coil radius. The transverse magnetisation was generated by calculating the longitudinal magnetisation at the start of each AFI cycle (see equation 7.3), rotating this vector into the transverse plane by an angle determined by the nominal flip angle and the B_1^+ magnitude and phase at that point in space, and then allowing for T_2^* relaxation to occur. K-space data was acquired for each receive coil by taking the two dimensional fast Fourier transform of the transverse magnetisation multiplied by the receive coil sensitivity, with complex Gaussian noise added (normally distributed magnitude, uniformly distributed phase).

Both fully sampled and undersampled (simulating parallel imaging accelerated) data was reconstructed using a sum of squares and simplified SENSE[186] (using only the pseudo-inverse of the coil sensitivity matrix) algorithms, yielding magnitude and phase images for each AFI cycle. Flip angle maps were generated from the magnitude images using equation 7.6, and the B_1^+ phase was set to be equal to the phase image acquired from the first AFI cycle. Vector B_1^+ maps were then formed using equations 7.11 and 7.14.

To illustrate the problem of how low SNR introduces large errors in the B_1^+ maps, simulated data was generated with the following parameters:

Geometry: FOV $= 150 \times 150mm$, image matrix $= 64 \times 64$,

Phantom: 138mm height, 103.5mm width, T_1 values chosen to similar to grey matter, white matter and CSF at 3T: 1400ms, 1000ms and 4200ms respectively. T_2 was set to $200ms$.

Transmit Coils: Two circular coils of radius 32.5mm, located at a radial distance of 71.25mm from the origin and at azimuthal angles 0 and 180° to the x axis.

Receive Coils: Eight circular coils of radius 25mm, located at a radial distance of 71.25mm from the origin and uniformly distributed azimuthally (spacing of 45°).

Imaging Parameters: Nominal flip angle $\alpha = 60°$,RF pulse duration $\tau = 500\mu s$ (corresponding square pulse amplitude of $7.8\mu T$), TE $= 5ms$, $\text{TR}_1 = 50ms$, $\text{TR}_2 = 150ms$, parallel imaging acceleration factor 2, signal to noise ratio (SNR) 100.

A metric of comparison was defined in accordance with those set out by Malik et al. [148] to quantify the error in the methods presented in this section: the root mean square (RMS) difference between the calculated B_1^+ maps and the reference B_1^+ maps used to generate the simulated data:

$$E_m = \sqrt{\frac{1}{N_c}\frac{1}{N}\sum_{j}^{N_c}\sum_{n}^{N}\left|\left|B_{1,jn}^+\right| - \left|B_{1,ref,jn}^+\right|\right|^2} \tag{7.16}$$

$$E_p = \sqrt{\frac{1}{N_c}\frac{1}{N}\sum_{j}^{N_c}\sum_{n}^{N}\left|\arg(B^+_{1,jn}) - \arg(B^+_{1,ref,jn})\right|^2} \tag{7.17}$$

Where N is the total number of voxels within each B_1 map, n is the voxel index, N_c is the total number of coils, and j is the coil index.

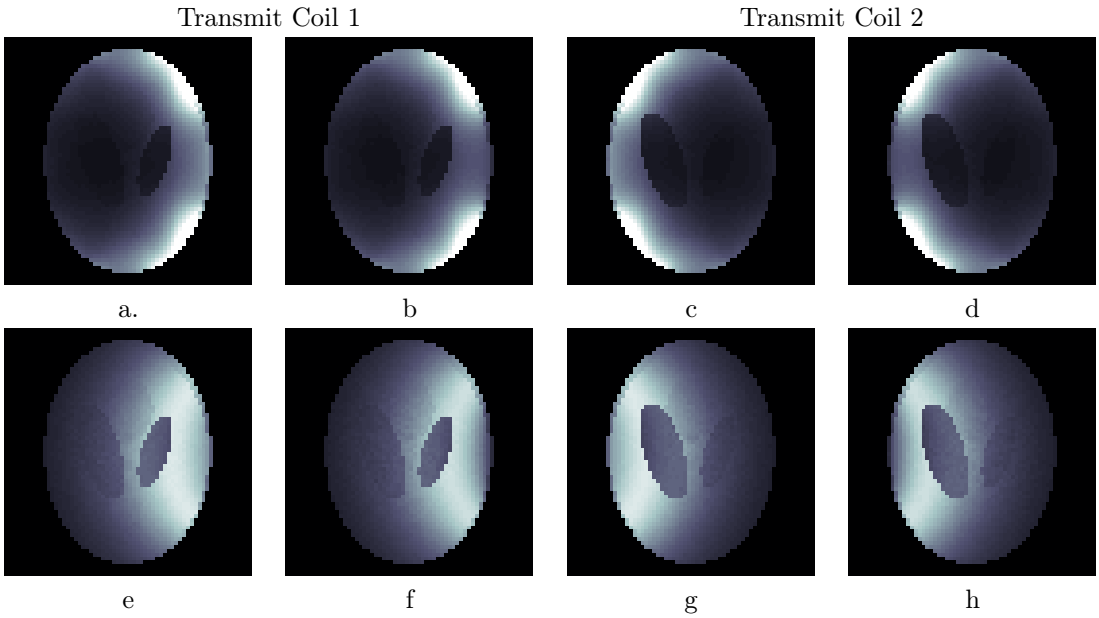

Figure 7.5: AFI images from cycle one (a,c,e,g) and cycle two (b,d,f,h), reconstructed using Sum of Squares (a-d) and SENSE (e-h) algorithms

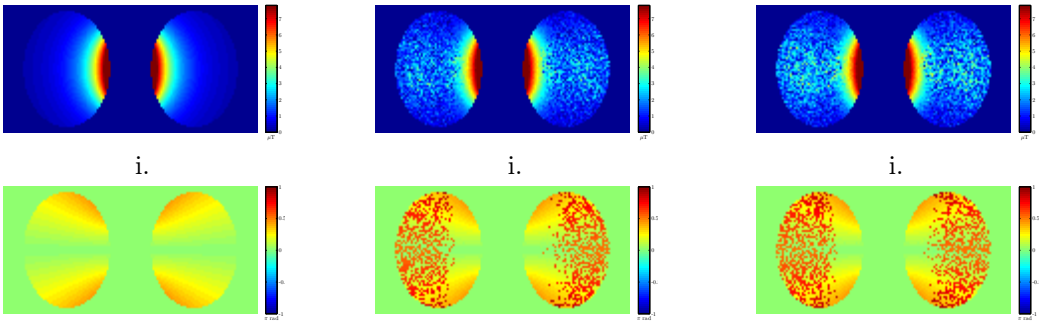

a. Reference b. Sum of Squares Reconstruction c. SENSE Reconstruction

Figure 7.6: Simulated i. Magnitude and ii. Phase B_1^+ Maps

Figure 7.5 shows the magnitude images reconstructed from simulated data for each AFI cycle. The basic sum of squares reconstruction algorithm does not use coil sensitivity data, and so the images show some weighting due to the receive coils' inhomogeneous sensitivities. Homogeneity is better in the SENSE reconstructed images, at the expense of notably decreased SNR in the centre of each image. Visually, the SNR of all the images is high.

Figure 7.6.a shows the actual B_1^+ magnitude and phase maps used to excite the digital phantom. Figures 7.6.b and 7.6.c show the computed B_1^+ magnitude and phase maps for the sum of squares and magnitude reconstructions. Both magnitude images show increasing amounts of noise as the

B_1^+ field strength decreases, such that noise dominates where the field strength is below approximately $3\mu T$ on the reference magnitude map . Greater errors are present in the SENSE images due to the intrinsic noise propagation of the reconstruction algorithm. In voxels where the ratio S_2/s_1 has exceeded unity the calculated flip angle from equation 7.6 becomes imaginary, introducing a phase shift of $\pi/2$, which is visible in the otherwise high SNR B_1^+ phase maps in figures 7.6.b.ii and 7.6.c.ii.

The $\pi/2$ phase shift error can be corrected by using only the magnitude of the flip angle map when calculating the total vector B_1^+ field, thus equation 7.14 becomes:

$$\mathbf{B_1^+}(\mathbf{r}) = \frac{|\alpha(\mathbf{r})|}{\gamma\tau} e^{i\varphi_{B_1,rel}(r)} \tag{7.18}$$

With this correction no phase shift occurs in regions of low SNR, as shown in figures 7.7.a and 7.7.b.

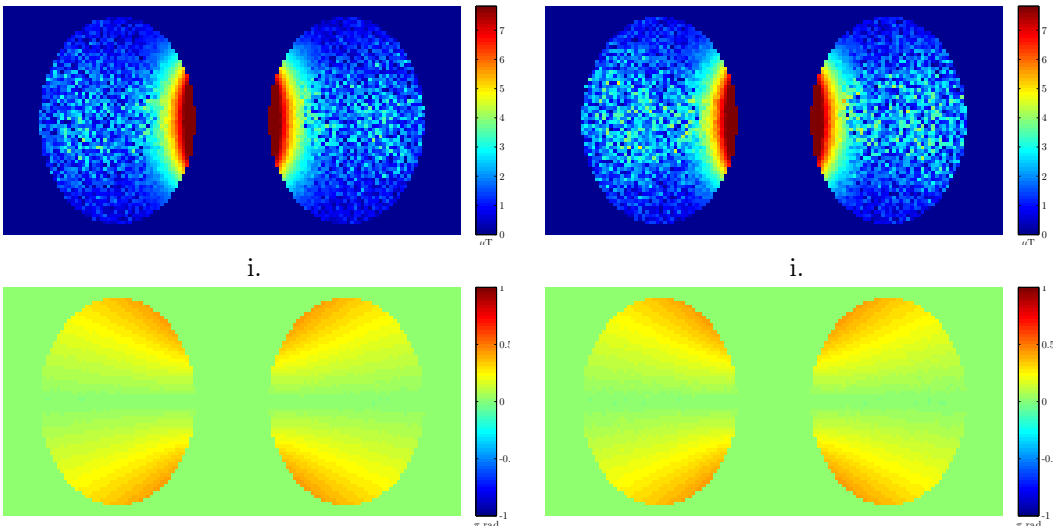

i. i.

ii. ii.

a. Sum of Squares Reconstruction b. SENSE Reconstruction

Figure 7.7: Corrected B_1 Maps: i. Magnitude and ii. Phase

The RMS errors for the sum of squares and SENSE reconstructions, with and without the phase correction of equation 7.18 applied are shown in table 7.1.

	E_m (μT)	E_p (rad)
Sum of Squares	0.78	0.58
Sum of Squares - phase corrected	0.61	0.006
SENSE	0.95	0.60
SENSE - phase corrected	0.77	0.008

Table 7.1: B_1^+ map RMS error

By using the magnitude of the flip angle when calculating the vector B_1 map the RMS error in the magnitude is reduced by approximately 20%, and the RMS error in the phase is reduced by two orders of magnitude.

7.2.3.2 B_1^+ Mapping By Interferometry

B_1^+ mapping methods are not inherently robust to a large dynamic range of flip angles. Large flip angles are penalised from limits in the maximum flip angle a B_1^+ mapping sequence is capable of measuring, and ultimately SAR constraints, and low flip angles constitute to low signal, decreasing SNR and the subsequent accuracy the measured B_1^+ field. A compromise between SNR, SAR and peak signal must be found, and for coils that are small in relation to the field of excitation this inevitably results in noisy B_1^+ maps.

A novel solution proposed by Nehrke [161], and Brunner and Pruessman in [31, 30] is to reduce the dynamic range of the B_1^+ field by measuring the perturbation of each given array element on a SAR-efficient baseline RF field distribution. In doing so the flip angle varies very little over the excitation volume and the power level can be appropriately chosen to maximise SNR and minimise SAR. To accomplish this linear combinations of coil elements were used for excitation rather than a single coil. For a set of N_c transmitter coils, each producing a RF field $B_{1,n}^+(\mathbf{r})$, the total field from linear combination (LC) configuration m, $\overline{B}_{1,}^+(\mathbf{r})$ is given by:

$$\overline{B}_{1,m}^+(\mathbf{r}) = \sum_{n=1}^{N_c} A_{mn} B_{1,n}^+(\mathbf{r}) \tag{7.19}$$

Where A_{mn} are individual elements of the transformation matrix $\mathbf{A}$, each a complex number representing the relative drive amplitude and phase of each element. By forming a set of M (where $M \geq N_c$) unique LC configurations, and applying the inverse of the transformation matrix $\mathbf{A}$, B_1^+ maps of each individual transmitter coil can be obtained:

$$B_{1,n}^+(\mathbf{r}) = \sum_{m=1}^{M} (\mathbf{A}^{-1})_{mn} \overline{B}_{1,m}^+(\mathbf{r}) \tag{7.20}$$

Matrix $\mathbf{A}$ represents the driving configuration of the transmitter array, and must be carefully chosen. As discussed in [31] there are three competing objectives; in all LC configurations a large B_1^+ throughout the sample is desired, high local electric fields are to be avoided, and $\mathbf{A}$ must be well conditioned. Previous implementations have used the circularly polarised (CP) mode of a multi-element transmit array[2] as the low dynamic range, low electric field baseline. Each LC configuration is constructed by perturbing the driving condition of one element:

$$A_{mn} = \begin{cases} 1 & n \neq m \\ ae^{i\phi} & n = m \end{cases} \tag{7.21}$$

Possible perturbations have included $a = 1$ and $\phi = \pi$ [31], $a = 0$ [161], and optimised values for a and ϕ specific to the imaged anatomy [148], taking into account non-uniformity due to sample

[2] A phase shift corresponding to the azimuthal position of each element is applied, so that the phase shift sums to 360° in total. For example a cylindrically shaped array consisting of eight equally spaced elements around the circumference will have successive phase shifts of 45° applied to each element.

loading and for the particular pulse sequence used to acquire the B_1^+ maps (both AFI and the double angle method were used). All these methods show a marked improvement in the SNR of the B_1^+ maps of the individual coil elements in a multi-element transmit array, independent of the choice of B_1^+ mapping method.

As the number of elements in the ASL transmit array is only two, a CP baseline is not possible. Therefore the choice of a and ϕ centres on the aim to compress the dynamic range of the total transmitted B_1^+ field. In [148] it was shown that the condition number of the transformation matrix $\mathbf{A}$ is equal to one when $a\cos\phi = -\left(\frac{N_c}{2} - 1\right)$, as all rows in $\mathbf{A}$ become orthogonal. For $N_c = 2$, $a\cos\phi = 0$, leaving two options: $a = 0$, which is equivalent to transmitting on one coil at a time, or $\cos\phi = 0$, resulting in $\phi = \pi/2$. Thus $\mathbf{A}$ becomes:

$$\mathbf{A} = \begin{bmatrix} i & 1 \\ 1 & i \end{bmatrix} \tag{7.22}$$

Which is non-singular and has inverse

$$\mathbf{A^{-1}} = \begin{bmatrix} -0.5i & 0.5 \\ 0.5 & -0.5i \end{bmatrix} \tag{7.23}$$

In each LC configuration both coils transmit with equal amplitude, and a phase shift of 90° is applied to one of the coils. To test this method, the B_1^+ map simulation described in figure 7.4 was modified to use LC coil sensitivities generated from the above transformation matrix for excitation rather than individual coils, resulting in LC B_1^+ maps from which individual coil B_1^+ maps were derived using the inverse transformation shown in equation 7.20. Simulation parameters were the same as those in section 7.2.3.1.

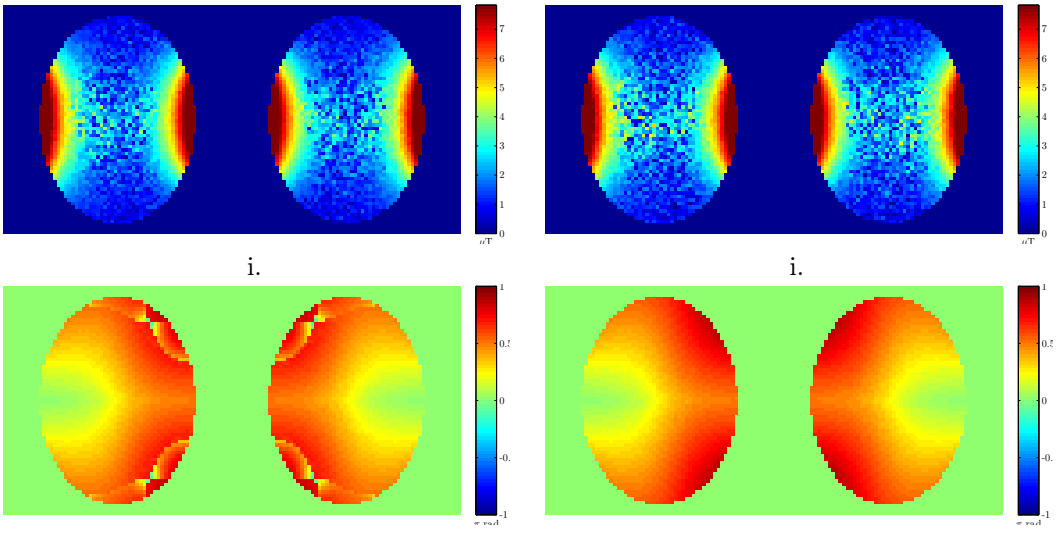

i. i.

ii. ii.

a. Sum of Squares Reconstruction b. SENSE Reconstruction

Figure 7.8: Simulated linear combination B_1^+ maps: i. Magnitude and ii. Phase

Figure 7.8 shows LC B_1^+ maps from each reconstruction method. SNR is reduced at the centre of the magnitude maps, where excitation is at a minimum. The discontinuities at the edges of the sum of squares reconstructed phase maps (figure 7.8.a) are artefacts due to the sum of squares reconstruction algorithm.

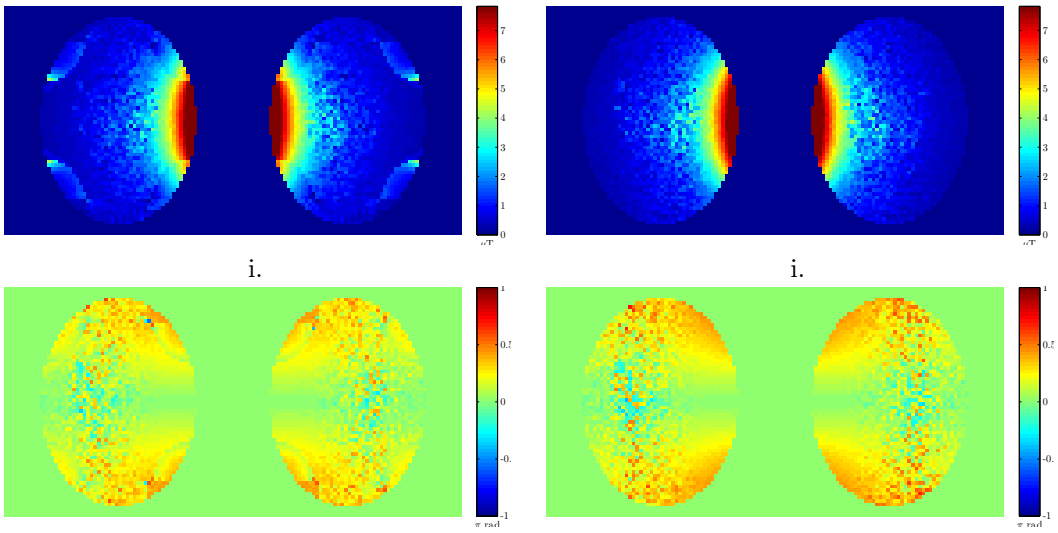

a. Sum of Squares Reconstruction b. SENSE Reconstruction

Figure 7.9: Individual coil B_1^+ maps derived from LC maps: i. Magnitude and ii. Phase

Figure 7.9 shows the individual coil B_1^+ maps derived from the simulated LC B_1^+ maps. The magnitude maps from both reconstruction algorithms show a large improvement in SNR, particularly far away from the coil. In the centre of the image noise is present due to the intrinsic low excitation of that region as shown in the LC maps in figure 7.8. There is however a substantial degradation in the SNR of the phase maps. Table 7.2 shows the RMS error in the phase and magnitude maps, with the phase corrected individual coil RMS errors for comparison. The RMS error of the LC derived SENSE reconstructed magnitude maps is less than half of the error of the individual coil map. The SOS reconstructed LC derived maps have an increased RMS error from localised artefacts due to the SOS reconstruction.

	E_m (μT)	E_p (rad)
LC derived, SOS	0.37	0.22
LC derived, SENSE	0.27	0.22
Individual Coil, SOS, phase corrected	0.61	0.006
Individual Coil, SENSE, phase corrected	0.77	0.008

Table 7.2: RMS error for LC derived B_1^+ maps with phase corrected individual coil

While previous methods have concentrated on utilising B_1^+ interferometry for mapping the transmit patterns of many-element arrays, results presented in this section show that the same technique can be used to provide reductions in B_1^+ mapping error for only two coils. Using linear combinations of coils for excitation produces flip angle maps (therefore B_1^+magnitude maps) with substantially higher SNR than from transmitting on only one coil at a time. Transmitting on

multiple coils reduces the dynamic range of the transmitted B_1^+, ensuring sufficient excitation over most of the FOV. The spatial SNR distribution of a resultant individual coil B_1^+ map is no longer dependent on the degree of excitation the coil can produce. However the inverse transformation procedure which extracts individual coil B_1^+ maps from the LC maps propagates noise from the magnitude maps to the phase maps.

7.2.3.3 Spatial Filtering

Up until this point no standard noise reduction techniques, such as low pass filtering have been used. A further gain in accuracy of the B_1^+ maps was made by applying a pixel-wise adaptive Wiener method [143] noise-removal filter to the magnitude LC B_1^+ maps prior to the inverse transformation. By only applying the filter to the magnitude component of the B_1^+ maps the high SNR and accuracy of the LC phase component is preserved.

The Wiener filter estimates the local image mean and standard deviation of neighbourhood of size m-by-m. If this neighbourhood is too small the filtering will be insufficient, conversely if the filter is too large detail will be lost in the filtered image. The optimum neighbourhood size was found by computing the magnitude and phase RMS error for a range of neighbourhood sizes. A filter size of m=1 corresponds to no filtering. Figure 7.10 is a plot of the RMS magnitude (shown in blue) and phase (shown in red) error, each normalised to the calculated RMS error with no filtering applied.

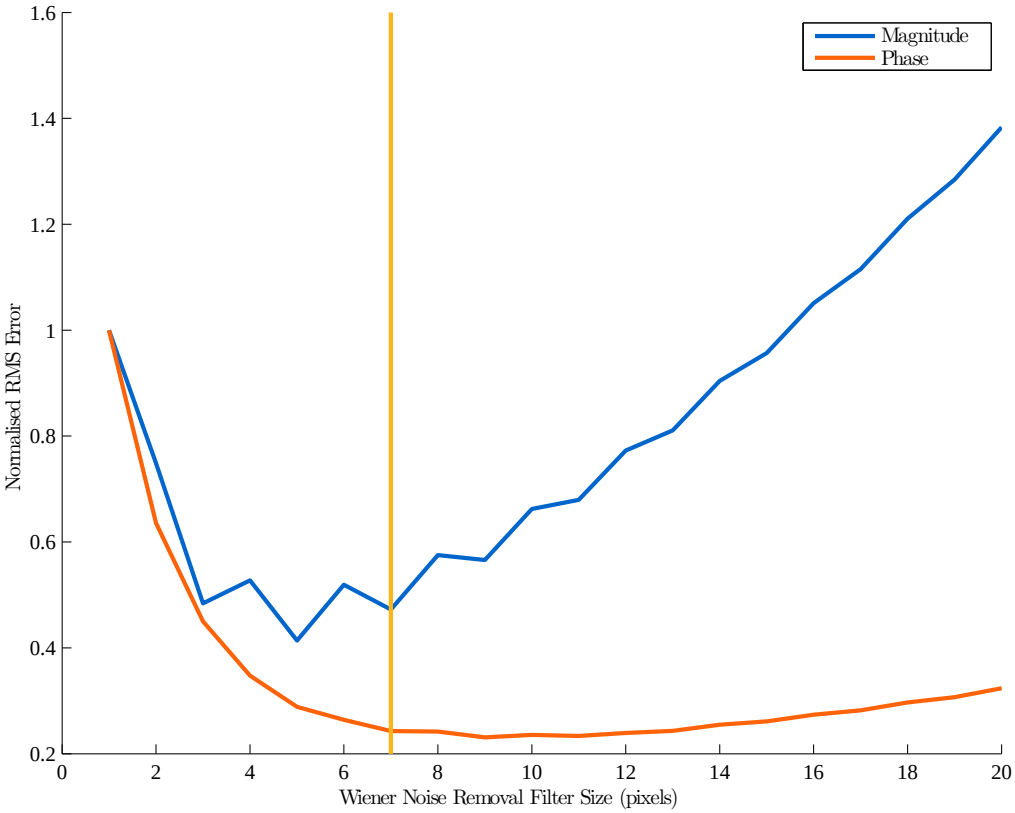

Figure 7.10: Normalised RMS error as a function of Wiener filter neighbourhood size

Optimal values for m were found to be 5 for the magnitude image, and 9 for the phase image, resulting in normalised RMS errors of 0.41 and 0.23 respectively. As they do not coincide the mean value ($m = 7$) was used, shown on figure 7.10 by the yellow line, yielding normalised RMS errors of 0.47 for the magnitude image and 0.24 for the phase image (RMS errors were $0.126\mu T$ and 0.054 rad). SENSE reconstructed B_1^+ magnitude and phase maps generated with this level of smoothing are shown in figure 7.11.

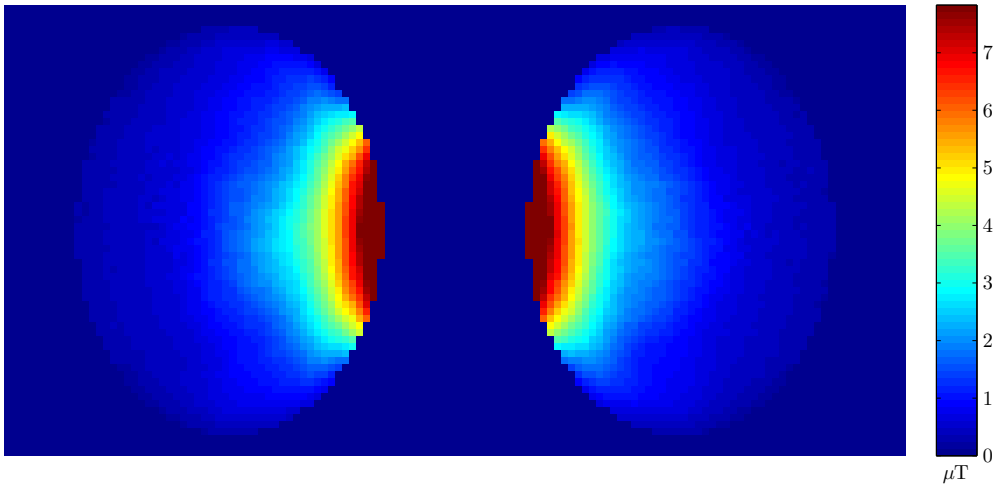

i.

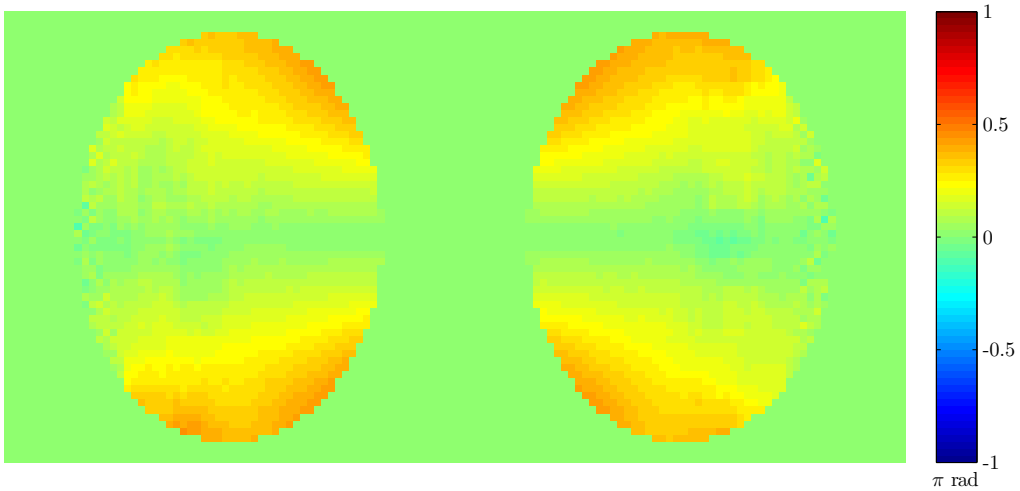

ii.

Figure 7.11: Wiener filtered B_1^+ maps

7.2.4 MRI measurements

B_1^+ maps were acquired and processed using the methods described above on an agar phantom and healthy volunteer.

7.2.4.1 Agar

The 12cm diameter cylindrical agar phantom (consisting of 2% agar, 0.9% NaCl and 1.6% $CuSO_4$) and three reference phantoms described in section 6.3.2.2 were imaged using the ASL Labelling

Coils for excitation and a 12 channel head array coil for reception. Labelling coils were positioned on opposite sides of the agar phantom. AFI datasets were acquired whilst transmitting at full power on (a) each coil individually, and (b) both simultaneously with phase shifts to create the linear combinations described in section 7.2.3.2 and equation 7.22.

AFI acquisition parameters for each scan were FOV = $220 \times 206mm$, 7.5mm slice thickness, 32 slices, 64×64 image matrix $TR_1 = 50ms$, $TR_2 = 150ms$, TE = $3.05ms$, pulse length $840\mu s$, 2×GRAPPA accelerated, $^6/_8$ partial Fourier in both phase encode dimensions, resulting in an acquisition time of 2m 48s.

Each run of the sequence produced a series of magnitude and phase images. Dicom images were converted to NIFTI in SPM 8 (Statistical Parametric Mapping, Wellcome Trust Centre For Cognitive Neuroimaging, Institute of Neurology, University College London) to separate the volumes of each AFI cycle, then the magnitude and phase NIFTI volumes of the for both coils/linear combinations and both AFI cycles (a total of eight images for each set of measurements) were loaded into Matlab. All further processing was performed with custom scripts. An object mask was generated by adaptively thresholding the sum of all four magnitude images on a slice-by-slice basis, the threshold value determined by estimating the background signal level. B_1^+ maps were calculated in accordance to the methods described in the preceding sections and in addition the linear combination maps were transformed to single coil maps as described in sections 7.2.3.2 and 7.2.3.3.

Computed B_1^+ maps are shown in figure 7.12. Transmitting with only one coil at a time yields B_1^+ maps with large errors on the contralateral side of the phantom, manifesting as voxels with value zero, see figure 7.12.a. In these voxels the signal difference between the two AFI images is smaller than the dynamic range of the images themselves (which is defined by the gain of the receiver pre-amplifiers, dynamic range of the analogue to digital converters, and also the 12 bit resolution available in dicom images), and as a result the voxel value is identical for both images. Consequently the ratio of the two voxels is unity, and because $\cos^{-1}(1) = 0$ the value within the B_1^+ maps at these voxels is zero. To illustrate the zero value voxel errors no spatial filtering has been performed on the individual coil B_1^+ maps. When the Wiener filter described in section 7.2.3.3 was applied the B_1^+ maps become smooth, but the zero value voxels propagate their errors into the surrounding voxels rendering the B_1^+ maps unsuitable for parallel transmission calculations despite their low SNR appearance.

B_1^+ maps acquired and processed from linear combinations of coils (figure 7.12.b) retain the SNR level on the contralateral side of the phantom, extending the dynamic range of the B_1^+ maps so that smaller flip angles can be resolved. Noise is now confined to where excitation is lowest at the centre of the phantom, and as artefacts from the linear combination inverse transformation at the sides of the phantom.

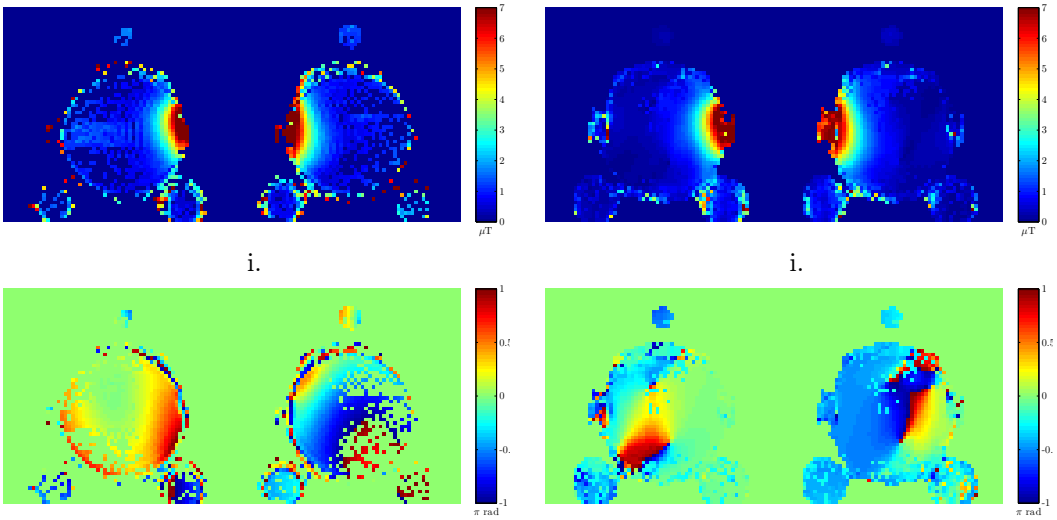

i. i.

ii. ii.

a. Individual Coil B_1^+ Maps b. Linear Combination Derived B_1^+ Maps

Figure 7.12: Phantom B_1^+ maps: i. Magnitude and ii. Phase

7.2.4.2 Human Subject

B_1^+ maps were acquired using the AFI sequence of a healthy, 25 year old male subject. The ASL Labelling Coils were positioned on opposite sides of the neck, and a 32 channel head array coil was used for reception. Linear combination datasets were acquired whilst transmitting at full power on each coil. AFI acquisition parameters were FOV $= 220 \times 185mm$, 7.5mm slice thickness, 32 slices, 64×54 image matrix, $\mathrm{TR}_1 = 50ms$, $\mathrm{TR}_2 = 150ms$, TE $= 3.05ms$, pulse length $840\mu s$, 2×GRAPPA accelerated, $^6/_8$ partial Fourier in both phase encode dimensions, resulting in an acquisition time of 2m 19s. Images were processed as described for the agar phantom in section 7.2.4.1.

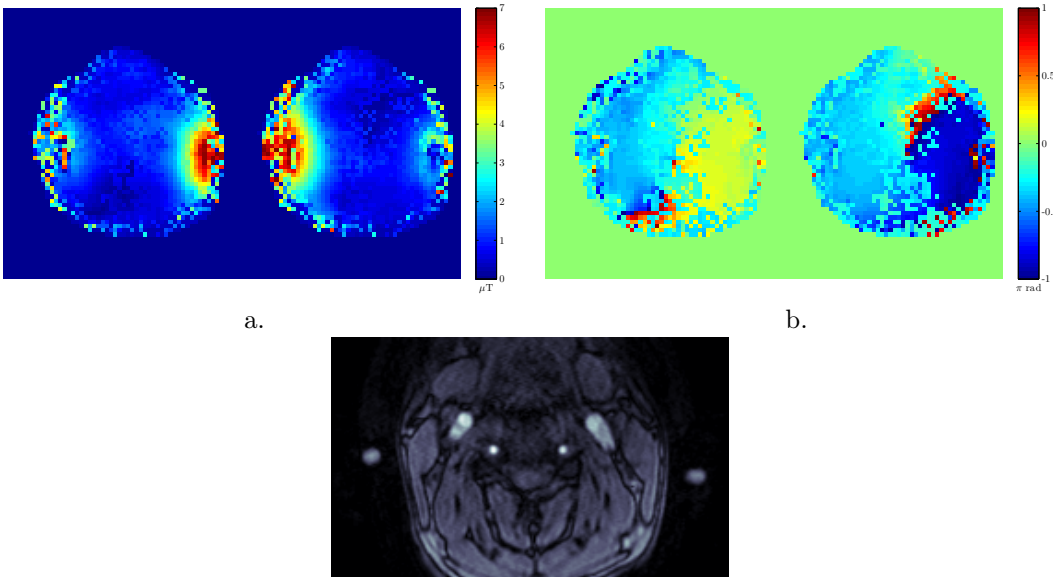

a. b.

c.

Figure 7.13: Human Neck Linear combination derived B_1^+ maps, a. magnitude, b. phase and c. corresponding TOF angiogram slice

Results are shown in figure 7.13. As with the phantom B_1^+maps SNR is good on the contralateral side. However, due to coil coupling the linear combination phase shifts were imperfect and as a result the inverse transformation required does not produce LC derived B_1^+ maps which solely consist of the transmit field from one coil, but have some contamination from the LC modes. This can be seen as noise close to the coils. Less coupling is present when using the agar phantom, most likely because its size and high salt concentration loads the coils more than a human neck, reducing the size of coupled RF currents and their subsequent perturbation on the driving conditions of each coil.

7.2.5 B_1^+ Mapping Conclusion

In this section methods have been presented to acquire high SNR B_1^+ maps of small surface coils. An Actual Flip angle Imaging (AFI) was chosen for its ease of implementation with a separate transmitter system, and for its ability to acquire 3D B_1^+ maps of the head and neck in under five minutes. Acquiring maps where the ASL Labelling Coils transmitted simultaneously, but with specific phase shifts to generate linear combinations reduces the dynamic range of the total excitation field, boosting the SNR significantly once these linear combination maps were transformed to single coil maps. Simulated and phantom results confirm this improvement, however due to inter-coil coupling (see section 7.2) it is not always possible to form perfect linear combinations, and consequently there is a degradation in the quality of the B_1^+ maps. Nevertheless, with the application of one of the well established decoupling methods for non-fixed position coils [199, 127, 135, 113], the methods described in this section should permit the acquisition of accurate B_1^+ maps.

7.3 Continuous Arterial Spin Labelling Pulse Sequence

The goal of an arterial spin labelling sequence is to acquire two datasets, a "tag" dataset where the magnetisation has been sensitised to the tissue perfusion rate, and a "control" dataset without perfusion sensitisation, but retaining all other effects of the pulse sequence such as magnetisation transfer. The tag and control datasets are subtracted to obtain a perfusion weighted image, also known as a ΔM image. All arterial spin labelling sequences consist of a labelling module followed by a readout module. During labelling, protons in the blood water are inverted to act as an endogenous contrast agent. As discussed in section 2.3 there are three main categories of labelling schemes - pulsed (PASL), continuous (CASL) and pseudo-continuous (pCASL), each of these with its own variants. To date ASL studies have been performed using a multitude of acquisition methods: 2D FT spin echo [53, 237, 204], gradient echo EPI [226, 246, 241, 130, 178, 46], spin echo EPI [102, 157, 156, 233, 232, 241, 247, 259], spiral EPI [242], fast spin echo [45, 159, 213]

and 3D GRadient and Spin Echo (3D-GRASE) [90, 69, 89, 68]. The inherent low SNR of the ASL signal (around 0.5-1-5% of the full signal) is a driving factor behind acquisition module choice, and has led to the adoption of methods which can rapidly acquire multislice images/3D volumes such as EPI and 3D-GRASE. Ultimately however, acquisition module choice is application specific and compromises must be made. For example gradient echo EPI can have a shorter echo time than spin echo EPI, resulting in a larger acquired signal. However gradient echo EPI images are highly susceptible to B_0 inhomogeneities and eddy-currents, resulting in distorted images with regions of signal dropout as well as distortion. Multi-slice EPI acquisitions are commonly employed in both a clinical and pre-clinical setting with CASL, pCASL and PASL labelling schemes. Multiple signal averages are required to obtain sufficient SNR for perfusion measurements, and as slices cannot be acquired simultaneously each slice has a different inflow time (TI).

An alternative, SNR efficient acquisition method, proposed by Günther et al. [90] is to acquire entire volumes in a single shot application of the 3D-GRASE [173] sequence. 3D-GRASE is a combination of a spin echo EPI and fast spin echo sequence, where each partition (a partition is a k_x–k_y plane in 3D phase encoding) is acquired with an EPI readout, and 180° pulses are used to refocus the magnetisation for phase encoding in the partition direction. SNR gains are made from the increased sensitivity of a larger excited volume, and all the voxels within the volume have the same inflow time. However, T_2 decay in the slice-select direction (each partition is acquired after a 180° refocusing pulse) causes blurring along this dimension, and at best the effective resolution in the slice-select direction is half of that nominally stated.

In [90], using a FAIR labelling scheme in conjunction with background suppression (see section 7.3.2) and saturation pulses to suppress signal from outside of the imaging volume, SNR gains of 2.8 were reported with 3D-GRASE compared to a multislice 2D gradient echo EPI sequence with the same nominal resolution.

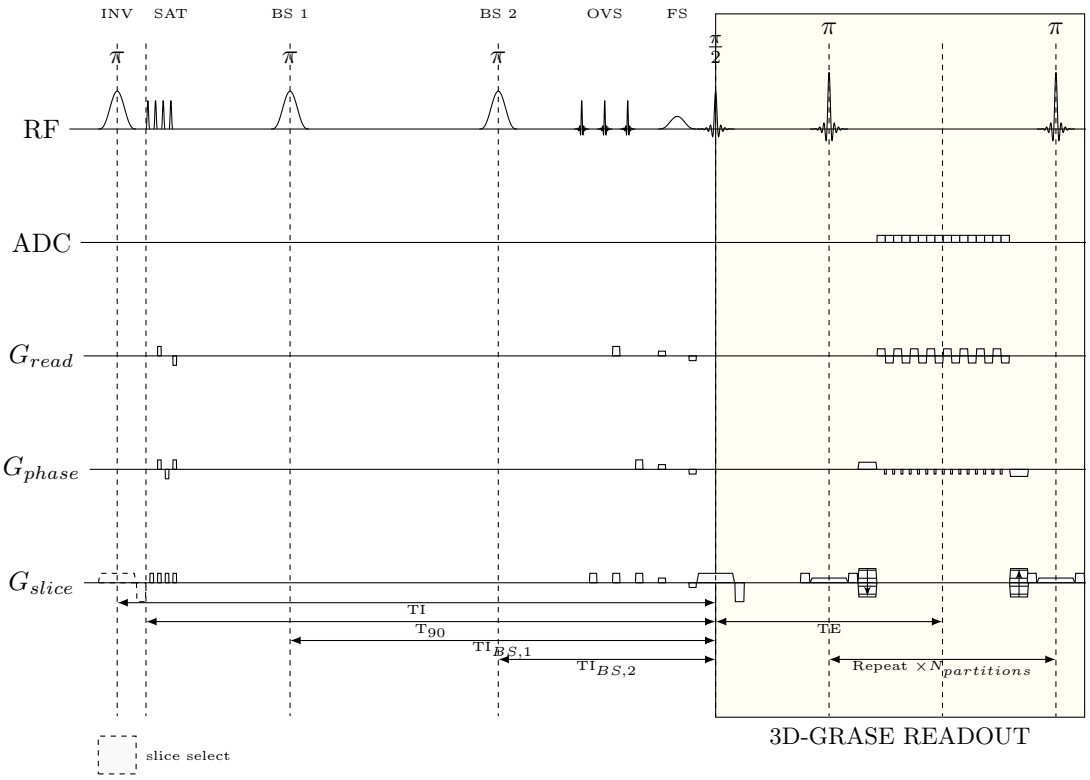

Figure 7.14: FAIR 3D-GRASE Pulse Sequence
INV–FOCI adiabatic inversion pulse; SAT–WET saturation pulses; BS 1/2–Background suppression FOCI adiabatic inversion pulses; OVS–Outer volume suppression modulated saturation pulses and Q2TIPS bolus truncation; FS–Fat saturation pulse

A derivative of the FAIR 3D-GRASE code developed by Matthias Günther had been previously implemented on the Siemens Trio. This code base was used as the starting point for a separate labelling coil CASL sequence. To produce a perfusion weighted dataset the FAIR 3D-GRASE sequence, as pictured in figure 7.14, acquires two inversion recovery images, one with a slice select inversion pulse in which relaxation within the inverted slice is sensitive to flow, and one with a global inversion where there is no flow sensitivity. Background suppression is employed to null signal from static tissue, reducing physiological noise. This is implemented by first saturating the magnetisation of the tissue within the imaging slab using an optimised four-pulse WET [165] presaturation scheme with optimised flip angles 88.9°, 98.7°, 82.5° and 159°, and then applying nonselective inversion pulses at specific times such that the longitudinal magnetisation of tissues with known T_1s is zero at the time of the imaging excitation pulse. Background suppression inversion pulses were applied with TIs computed according to the method in [90] for T_1 of 700ms and 1400ms. For the brain, where the static tissue consists of grey matter, white matter and CSF two inversion pulses are sufficient to nullify their signal. To ensure robustness to variations in B_1^+ the inversion and background suppression pulses are implemented as Frequency Offset Corrected

Inversion (FOCI) adiabatic pulses[172, 254].

Additional saturation pulses are then applied in a manner similar to Q2TIPS [145] to a) sharply define the temporal length of the tagged bolus, and b) suppress the signal from outside of the imaging volume. Q2TIPS applies a periodic train of thin slice saturation pulses proximally to the imaging slice, saturating the magnetisation of the inflowing bolus of tagged blood just before it enters the imaging volume, as described in section 2.3.3.1 in chapter 2. Inverted blood which enters the imaging volume before the application of the Q2TIPS pulses contributes to the perfusion weighted image, whereas blood which flows in during the application of the Q2TIPS pulses does not, hence the temporal length of the tagged bolus is defined. The sinc saturation pulses used in the FAIR 3D-GRASE are modulated with a sinusoid function so that a slice above and below the imaging volume is saturated. Three of these modulated saturation pulses are always run regardless of the bolus length, saturating the magnetisation of the tissue outside the imaging volume to provide a sharp slab profile.

Preceding the 3D-GRASE readout module a fat saturation pulse is applied to remove the fat signal. Then a 90° initiates the 3D-GRASE acquisition. At this point the effects of the background suppression, outer volume suppression and Q2TIPS modulated saturation pulses mean that (for ideal pulses) the longitudinal magnetisation is minimised for static tissue within the imaging volume. Magnetisation in slices surrounding the imaging volume, is also suppressed with the modulated outer volume suppression pulses. The blood which has flowed in during the TI and exchanged into the tissue will produce a ΔM signal. If Q2TIPS style pulses have been used to truncate the bolus of inflowing blood, any blood which flows into the imaging volume following the initiation of these pulses does not contribute to the ΔM signal. Acquisition is accelerated by acquiring a 6/8 partial Fourier dataset in the slice direction and using 2× accelerated GRAPPA parallel imaging.

7.3.1 CASL Implementation

The CASL sequence was implemented by adding an option within the FAIR 3D-GRASE protocol for CASL, which completely replaced the PASL labelling module, including the background suppression and modulated sats. In doing so, both labelling schemes can easily be compared as they share the same 3D-GRASE acquisition code. Additionally, as the Siemens pulse sequence programming environment is written in the object oriented programming language C++, many features of the PASL labelling scheme such as the WET presaturation sequence, FOCI inversion pulses and modified saturation pulses were implemented as methods of an ASL class, and consequently could simply be re-used within the CASL sequence code.

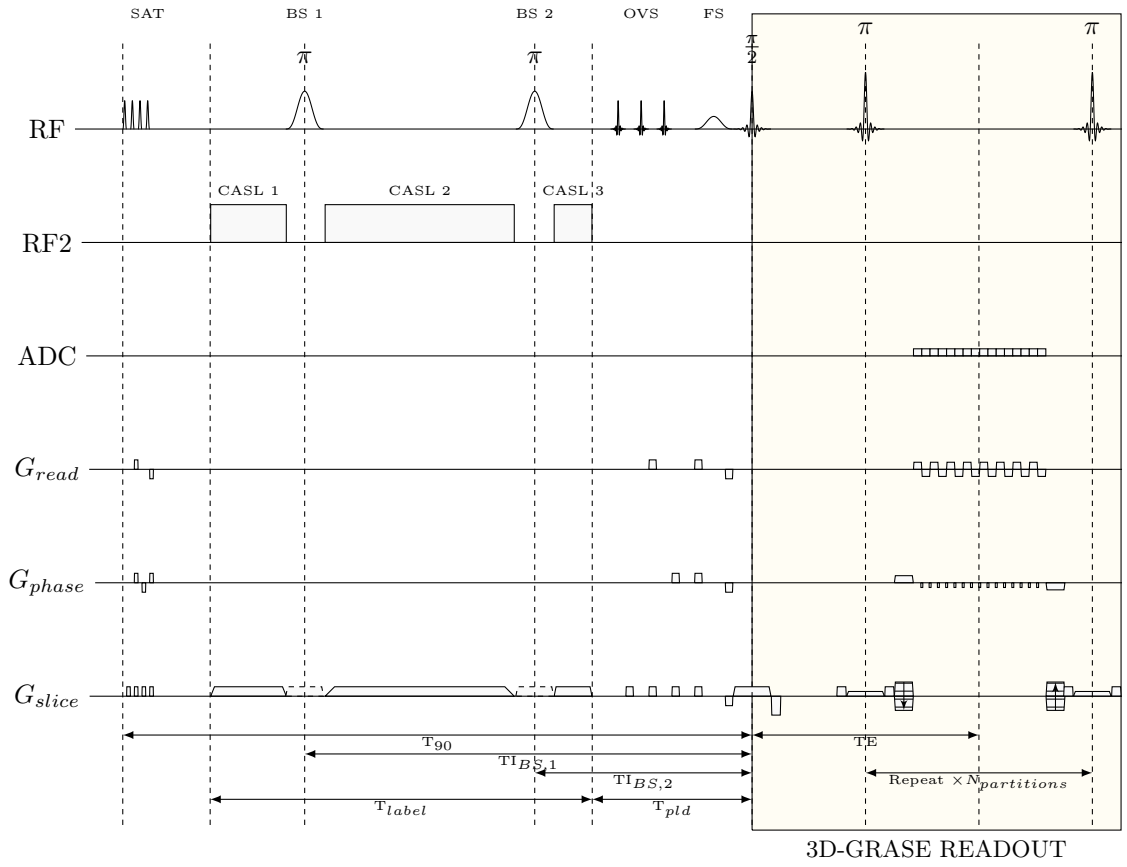

Figure 7.15: CASL 3D-GRASE Pulse Sequence
SAT–WET saturation pulses; BS 1/2–Background suppression FOCI adiabatic inversion pulses; OVS–Outer volume suppression modulated saturation pulses and Q2TIPS bolus saturation; FS–Fat saturation pulse

Figure 7.15 is a pulse sequence diagram of the CASL 3D-GRASE sequence. In contrast to the FAIR sequence, the CASL sequence commences with the WET presaturation pulses, thus the time between the end of the WET pulses and the start of the outer volume suppression pulses demarcates the time available for CASL labelling (T_{label}). The outer volume suppression and fat saturation pulses last for approximately 40ms and hence define the minimum post labelling delay (T_{pld}). Depending on the choice of T_{label} and T_{pld} the background suppression inversion pulses may occur inside the continuous labelling pulse, in this instance the labelling pulse is interrupted and the background suppression pulse run. More detail on the background suppression implementation is provided in section 7.3.2.

Synchronisation between the Siemens Trio and the SMIS spectrometer was achieved by using the trigger signal to gate continuous wave transmission (see section 5.2.1). Labelling gradient amplitude is controlled by the Trio sequence protocol, and the RF parameters such as amplitude and phase shift are controlled by the SMIS spectrometer.

7.3.2 Interleaved Background Suppression

Background suppression is a technique which uses multiple inversion recovery to suppress the signal from static background tissue. It was first proposed to remove the static background signal in angiograms [56, 150], allowing the resolution of smaller sized vessels, and has since become a feature in some arterial spin labelling sequences because of its ability to reduce physiological noise [68, 80, 90, 250, 44]. Ideally the signal from the static background tissue is identical in both control and tag images, however physiological noise and magnetic field fluctuations lead to instabilities in the background signal which differ between the tag and control images, contributing to the ΔM image when they are subtracted. These differences are more problematic for 3D datasets, and so background suppression techniques have been successfully employed to suppress noise due to inter acquisition magnetic field fluctuations and increase the sensitivity of 3D ASL methods.

The implementation of multiple inversion recovery based background suppression in the 3D-GRASE CASL sequence is shown in figure 7.16, along with the effect of the pulses on the longitudinal magnetisation of grey matter, white matter and cerebrospinal fluid (CSF). T_1 relaxation times were set to those exhibited at a B_0 strength of 3 Tesla. First the magnetisation of the tissue is saturated and gradient spoiled, destroying any transverse coherence. In CASL and pCASL ASL sequences this usually occurs before any labelling begins, and for PASL the saturation can occur a short time after the labelling inversion pulse. However if this delay is too long then labelled blood which has entered the imaging volume will also be saturated, destroying its contribution to the perfusion signal. Then inversion pulses are applied at specific times so that the longitudinal magnetisation of the tissue is zero at the time of the excitation pulse. By using multiple inversion pulses the signal from several different T_1 species can be nulled.

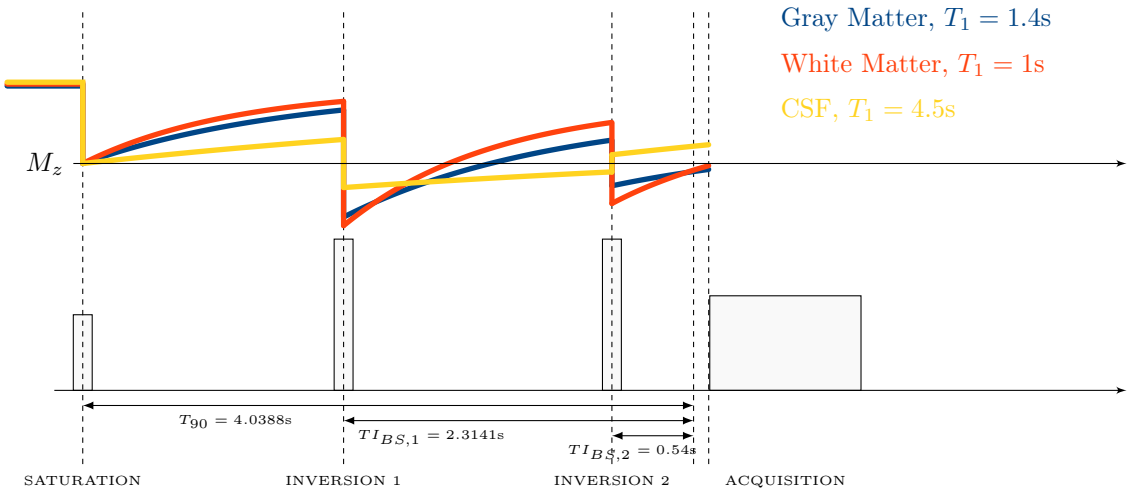

Figure 7.16: Multiple Inversion Background Suppression

TI times for the background suppression pulses were selected using the following procedure:

1. The time between the saturation pulses and the start of the acquisition module was selected to allow a good range of labelling durations and post labelling delays, without being so long that the scan time would be increased unnecessarily. Four seconds was chosen for the combined label and post label delay duration; with the additional duration of the WET presaturation sequence this became 4.038760s.

2. The number of inversion pulses was chosen. While it is theoretically better to have many inversion pulses as a smaller value of M_z can be reached for all T_1 species, the imperfect nature of the inversion pulses means that in practice fewer pulses are more effective. Two pulses were previously implemented and gave good results in the FAIR 3D-GRASE sequence, therefore it was decided to use two pulses in the CASL 3D-GRASE sequence.

3. The TI times of each inversion pulse, $\mathrm{TI}_{BS,1}$ and $\mathrm{TI}_{BS,2}$ were then calculated by solving at $t = T_{90}$ the signal equation for the background suppression sequence[150] for the T_1's of white matter (1s), grey matter (1.4s) and CSF (4.5s) at 3T:

$$M_z(T_1) = 1 + (-1)^{n+1} e^{-\frac{T_{90}}{T_1}} + 2\sum_{i=1}^{n} (-1)^i e^{-\frac{\mathrm{TI}_{BS,i}}{T_1}} \tag{7.24}$$

   Solutions to equation 7.24, with n=2 were determined using the function "fsolve" with the Levenberg-Marquardt algorithm in Matlab, yielding $TI_{BS,1} = 2.3141s$ and $TI_{BS,2} = 0.5437s$.

4. To ensure any residual longitudinal magnetisation is always positive at the point of excitation, an additional 100ms delay was introduced prior to the 90° excitation pulse

For any given post-labelling delay, it has been shown that the best signal can be achieved with a long labelling duration[44]. Therefore, it is desirable to interleave the background suppression with the CASL labelling pulses to make maximum use of the time available for labelling. However, in doing so the magnetisation of the inflowing blood is also inverted and will act to reduce the perfusion signal unless a special strategy is adopted to preserve the ASL signal. Once tagging has commenced, any background suppression pulses which occur during the labelling pulse must be slice selective, and the labelling pulse must also be switched. In other words, in the tag condition an interleaved background suppression pulse switches off the labelling RF until another interleaved background suppression pulse is run, and in the control condition an interleaved background suppression pulse switches on the labelling RF until another interleaved background suppression pulse is run, as shown in figure 7.17.

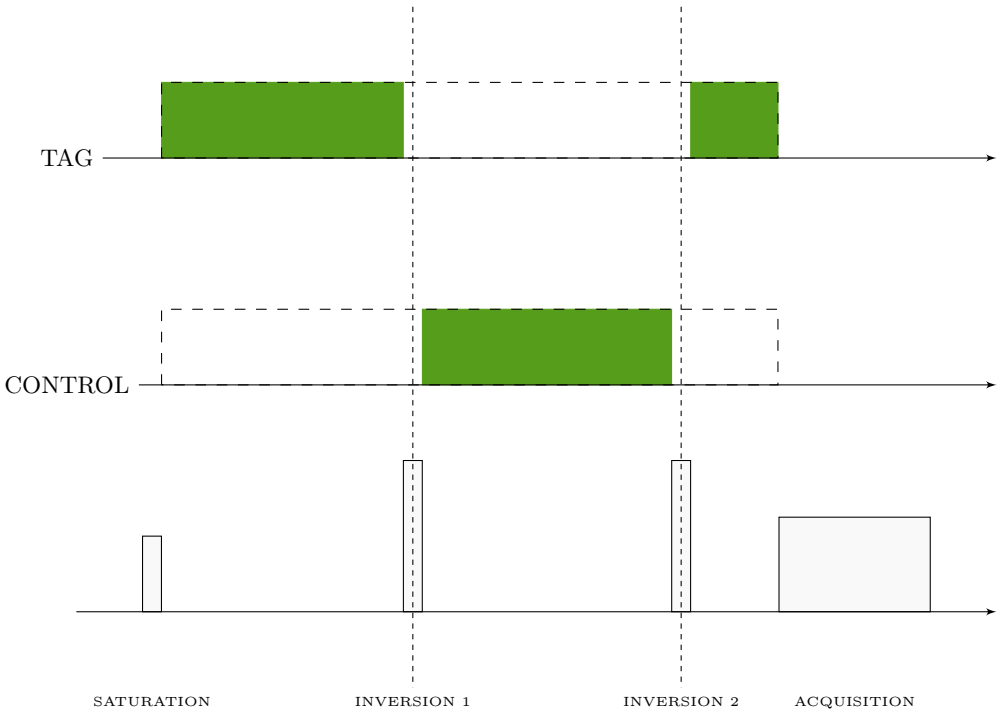

Figure 7.17: Interleaved Background Suppression

Labelling can be performed from the end of the saturation pulses until the start of the acquisition module (minus the fat saturation and outer volume suppression pulses, not shown for brevity), shown as the dotted outline in figure 7.17. Conventionally, labelling only occurs during the tag acquisition, however when background suppression pulses are interleaved this swaps between the labelling and control conditions, as shown by the green fill in figure 7.17. When using separate labelling coils the control condition is simply no RF, if other labelling schemes are used such as pCASL, or CASL with amplitude modulated control, then the applicable control condition is run.

7.4 Perfusion Phantom

To facilitate the development of parallel transmit continuous arterial spin labelling, a novel perfusion phantom was designed and constructed. Perfusion phantoms have been previously employed to provide a controlled environment for the optimisation of ASL sequences by simulating the macroscopic coherent flow encountered in large blood vessels, and the microscopic, incoherent flow synonymous with perfusion through capillaries to the tissue bed.

Flow based phantoms have been used to determine the inversion efficiency of continuous and pseudo-continuous labelling strategies [157, 79, 101, 119], typically consisting of tubes of flowing water embedded in a medium such as agar to provide background signal. For example, the flow phantom used by Werner et al.[232] for the development of continuous artery-selective spin labelling

(CASSL), a vessel selective ASL method consisted of two glass tubes embedded within agar, each glass tube having two regions: a 6mm inner diameter labelling section where the flow velocity could be adjusted to match physiological conditions; and a 12mm inner diameter section where the flow velocity was reduced and a measurement of the magnetisation could be made to determine the labelling efficiency.

Microvascular perfusion phantoms have been used for perfusion measurements as a direct replacement for in-vivo imaging. Phantom construction is more varied, but generally consists of a porous substrate which mimics the capillary bed, to which flowing water is delivered. In early work investigating the effects of perfusion on the MR signal Duewell et al. [58] utilised a dialysis filter containing 9000 parallel capillaries ($200\mu m$ diameter, $40\mu m$ wall thickness) for linear capillary flow, and an anion exchange resin consisting of $450\mu m$ average diameter beads for incoherent flow. A similar study by Lorenz et al. used an isolated canine kidney preserved in formalin [144] to simulate perfusion at the capillary bed.

Warmuth et al. used a living porcine kidney perfused with oxygenated autologous[3] blood to determine the accuracy of CBF values measured with a FAIR arterial spin labelling sequence [227]. Perfusion was conducted for three hours, during which the kidney consumed oxygen and produced urine. Using a living kidney enabled real perfusion to be measured, which includes exchange of blood water into the extra and intracellular space, rather than only the microscopic incoherent flow present in a synthetic perfusion phantom. However significant effort is required to keep the kidney alive, and the same kidney cannot be re-used.

Synthetic perfusion phantoms used in ASL studies have used a variety of materials as a perfusion medium: foam of different porosities [67], capillary tubes with plastic beads placed proximally to the tubes to divert flow [163], and a chamber filled with 3mm glass beads fed by an array of perforated pipes [137]. During MR imaging of the latter, the perfusion signal from within the chamber has components of both coherent macroscopic flow from the pipes, and incoherent flow around the glass beads, therefore this represents a good analogue for perfusion within the brain.

A perfusion phantom for parallel transmit continuous arterial spin labelling requires characteristics of both macroscopic and microscopic flow phantoms to allow full development of both labelling and acquisition techniques, and their comparison with other available sequences. Multiple, distinct perfusion territories are also necessary for evaluating vessel selective labelling methods. Finally the dimensions and geometry of the phantom should be similar to the anatomy of the neck and brain, and flow rates should be physiologically comparable. Therefore it was decided to design and construct a phantom which simulated a simplified model of perfusion to the left and right hemispheres of the brain from the carotid arteries, consisting of two sets of branching tubes (representing the carotid arteries and their branching into smaller arteries), each set providing flow for a perfusion chamber similar to that in [137]; perforated tubes inside a chamber filled with glass

[3]Blood from the donor; in this case blood taken from the same pig as the kidney.

beads. The phantom was designed using a Computer Aided Design (CAD) software package, and manufactured by 3D printing nylon, ensuring accuracy and detail greater than that achievable by hand.

7.4.1 Design

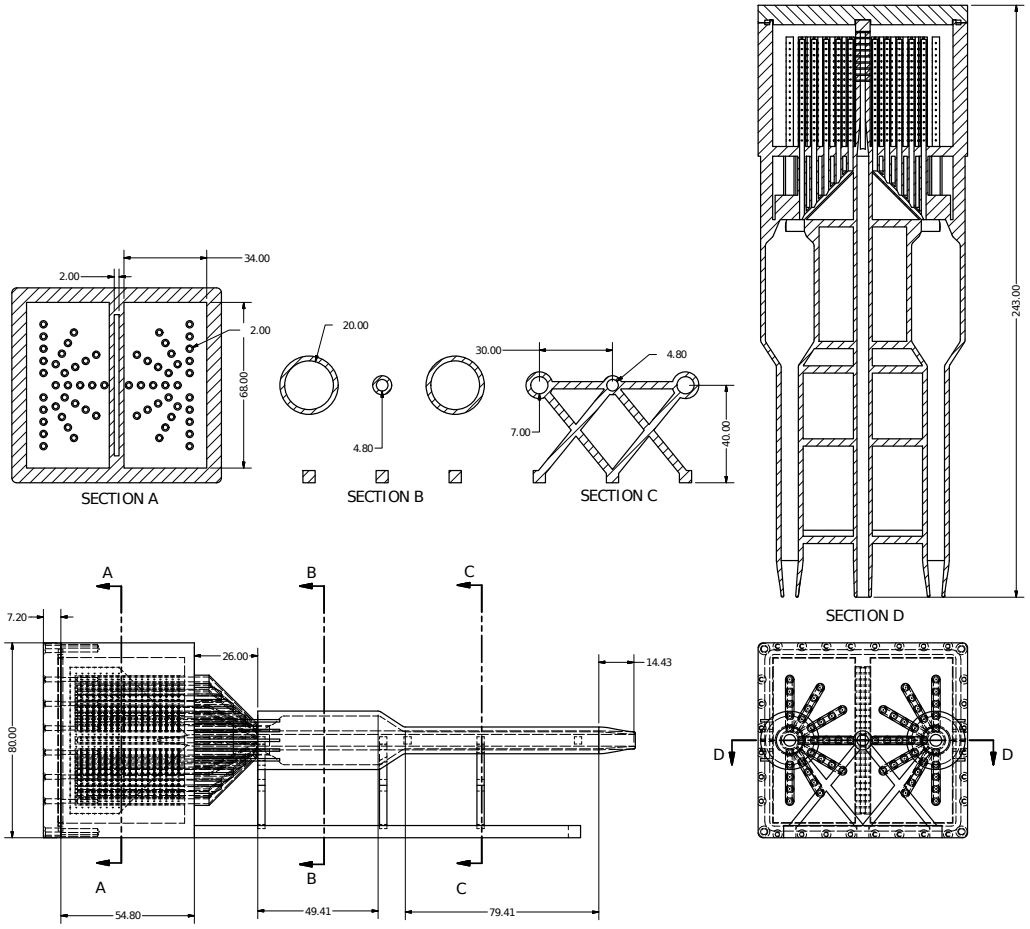

Figure 7.18: Phantom design schematic

Figure 7.18 shows a working drawing of the design of the phantom. Inflowing water travels through two pipes representing the carotid arteries, where the flowing protons can be labelled in a CASL imaging sequence as they are of similar dimensions to the carotid arteries in an adult. These pipes then expand in diameter, slowing the flow and enabling a MR measurement of the inversion efficiency following CASL labelling (in a similar manner to in [232]). Next, each pipe branches into thirty smaller pipes, representing the arterial tree branching out into smaller arteries. The small pipes are arranged radially, with the largest concentration of pipes (and hence, highest perfusion rate) at the centre of each group. The pipes then enter the perfusion chambers, which are filled with glass beads, and perforations along the length of the pipes allows water to pass. Water flows out of the perfusion phantom via small perforations in the dividing partition between the two perfusion chambers, which then lead to the outflow pipe.

Dimensions of the phantom were based on the measurements made in section 3.1.1, specifically:

- The diameter of each inflow pipe based on the diameter of the carotid arteries - 6.5mm.
- Separation between the two inflow pipes based on the separation between the carotid arteries - 60mm.
- Distance between the middle of the inflow pipes to the middle of the perfusion chamber - 168mm.

Basing the phantom's dimensions on real anatomy improves the utility of the phantom as a sequence development aid, as minimal adjustment to scan parameters will need to be made upon the transition to in-vivo scans.

7.4.2 Construction

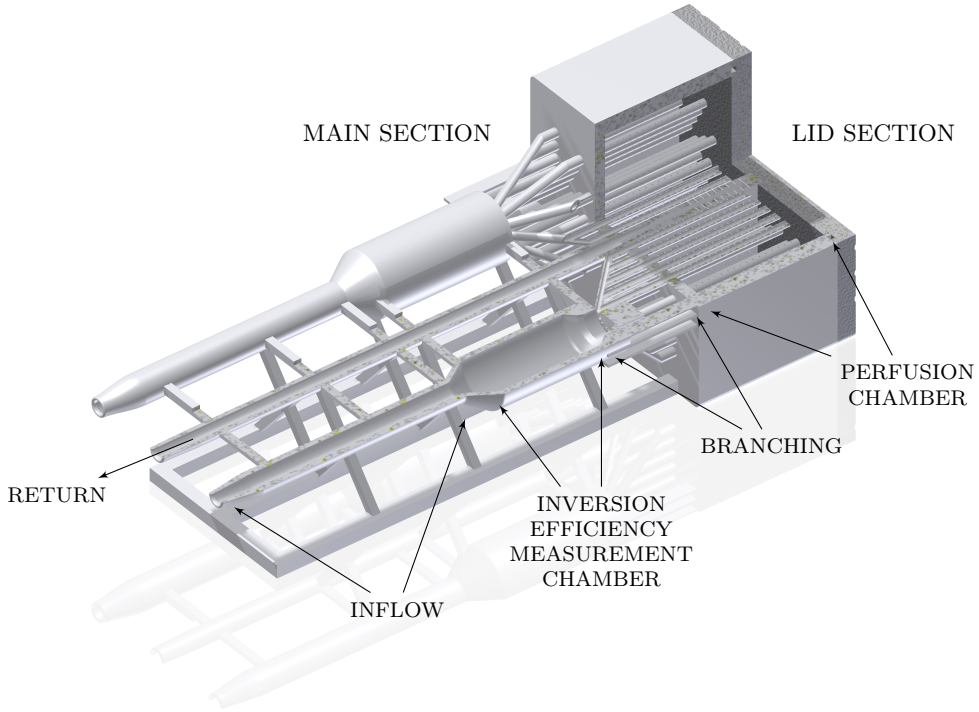

Figure 7.19: 3D render of the phantom design

CAD software (Inventor 2011, Autodesk Inc., San Rafael, CA, USA) was used to create a design that could be manufactured using selective laser sintering (SLS), a 3D printing process. The phantom was designed as two separate parts, the main section, containing all of the flow features, and a lid section to seal the perfusion chambers once they had been filled with glass beads. Once manufactured any unfused nylon powder was removed and screw holes were tapped. The perfusion chamber was then filled with 1mm diameter glass beads and the lid fixed in place with epoxy glue and nylon machine screws.

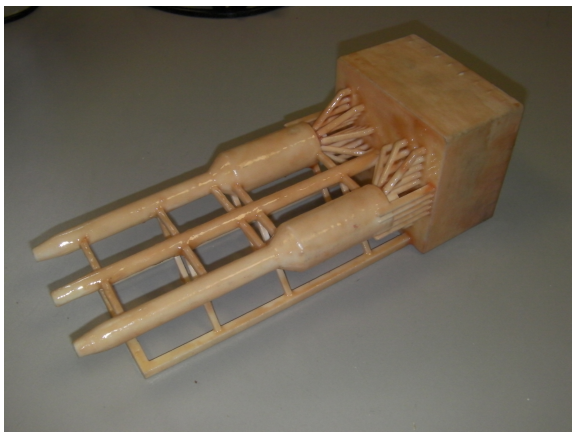

a. Main section, coated in epoxy sealant

b. Finished phantom, inside plastic container

Figure 7.20: Phantom Images

After initial tests showed the phantom to leak water at the interface between the main section and lid, three layers of glass fibre matting and epoxy resin were applied to the outer faces of the perfusion chamber to ensure a strong seal and no leaks from the perfusion chamber. Thin (less than approximately 4mm) regions of the nylon material fused by the SLS process can be slightly porous, so the exterior of the pipes were also coated with epoxy resin. The phantom was placed within a close-fitting watertight nylon plastic container as shown in figure 7.20.b, with holes drilled for 6.25mm inner diameter tubing (Tygon R3603, Saint Gobain Performance Plastics) to connect inflow and return pipes. Copper sulphate doped tap water was was used to fill the container and provide background contrast.

A 20m length of tubing was coiled inside a plastic cylindrical container and placed close to the phantom within the magnet bore, enabling the inflowing water to reach equilibrium magnetisation before entering the phantom, and providing a large volume of water to invert in pulsed ASL studies, as shown in figure 7.21. Pulsatile flow was provided by a peristaltic pump (502S, Watson Marlow, Cornwall, UK). Flow rate could be adjusted by varying the speed of the pump, and also by using a section of silicone hose with a wider diameter (10mm and 13mm have been used) in place of the regular 6.25mm tygon tubing.

The silicone hose, which was red in colour, was found to expand a small amount from the water

pressure at the peak of the pump cycle, changing the transmittivity/attenuation of light from a pulse oximeter (Siemens Medical Systems, Erlangen, Germany) placed around the tube, providing the MR scanner with a reliable signal from which to trigger gated scans. This phenomenon was not observed when using the transparent tygon tubing.

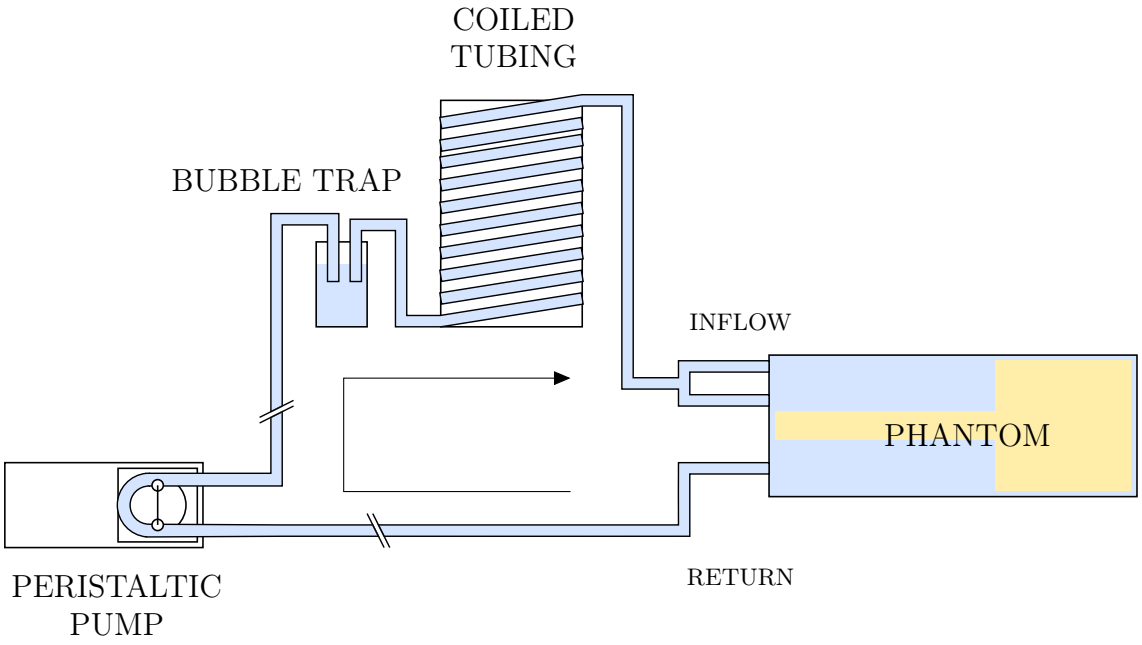

Figure 7.21: Phantom Operation Schematic

7.4.3 Characterisation

7.4.3.1 Flow

These tests were conducted prior to the addition of the fibre glass matting and epoxy resin around the perfusion chamber. A small amount of water was able to continuously leak out of the perfusion chamber into the phantom container, reducing the water pressure within the system, and permitting the use of the 13mm diameter silicone hose and much faster faster flow. Once the outside of the phantom had been sealed with epoxy the water pressure was so great that the tubing was regularly forced out of connectors, prohibiting the 10mm and 13mm silicone hoses from use.

Flow within the phantom was qualitatively assessed using a 3D time of flight (TOF) angiogram, and the flow velocity within the inflow and return pipes quantitatively measured using a single transverse slice gated phase contrast angiogram. Images were acquired using a 32 channel head coil array for reception and body coil for transmission. 3D TOF angiogram imaging parameters were FOV $= 180 \times 180 \times 206mm$, $384 \times 268 \times 206$ image matrix, TE $= 3.69ms$, TR $= 20ms$, $\alpha = 20°$. Phase contrast angiogram imaging parameters were FOV $= 180 \times 180mm$, 256×256 image matrix, 20mm slice thickness, 3 averages, $v_{enc} = 130cm/s$ in the through plane, anterior to posterior and right to left directions, TE $= 7.12ms$ and TR $= 132.9ms$, resulting in a time series of 11 images for each direction.

Acquired images were converted to velocity maps as described in section 3.1.1 in chapter 3. The inflow and return pipes were masked in the phase contrast images to remove the static background signal, and within each mask the velocity of the voxel at its centre of mass, and the mean velocity of all the voxels was computed.

Images from the TOF angiogram are shown in figure 7.22. Hyperintense voxels indicate flowing water. Flow within the perfusion chamber is most concentrated within the central tubes due to the reduced flow resistance. Figure 7.23.a is a transverse slice phase contrast angiogram at the peak of the velocity waveform. The two inflow pipes are clearly shown as regions with water flowing into the plane, and the return pipe as a region with flow out of the plane. The mean velocity over the pump cycle was found to be 10.2cm/s in the left inflow pipe, 10.5cm/s in the right inflow pipe and 34.6cm/s in the return pipe. The flow velocity is much lower than that encountered in the human carotid arteries (see section 3.1.1 and [106]).

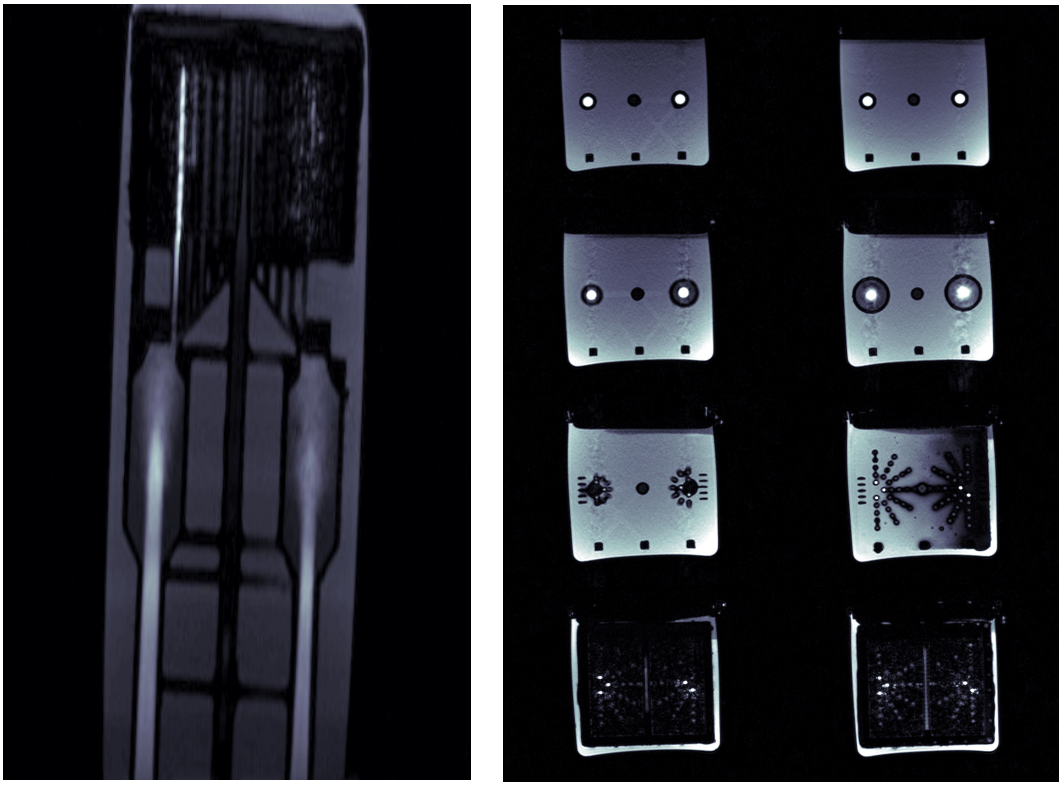

a. Coronal slice b. Transverse slices through the phantom

Figure 7.22: TOF Angiogram images

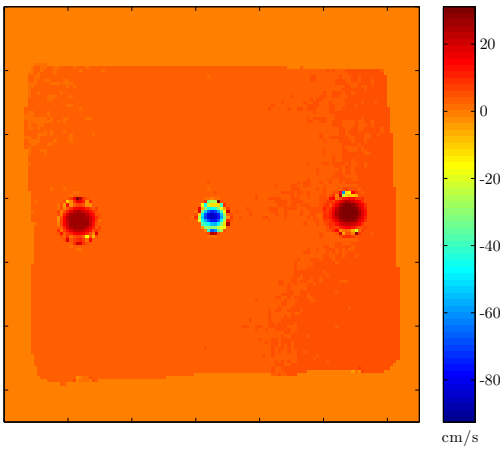

a. Velocity Map

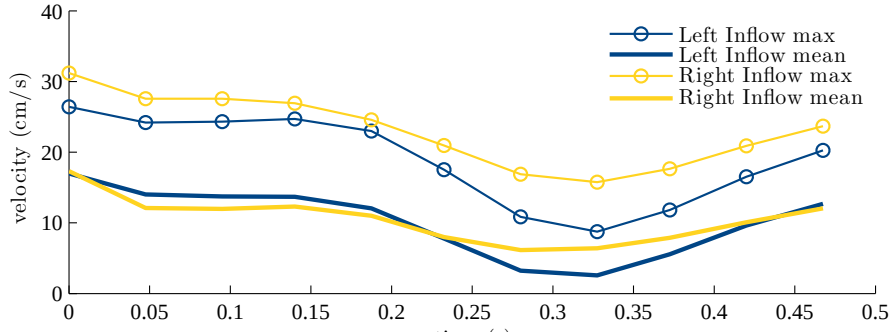

b. Peak and Average Velocity within the inflow pipes

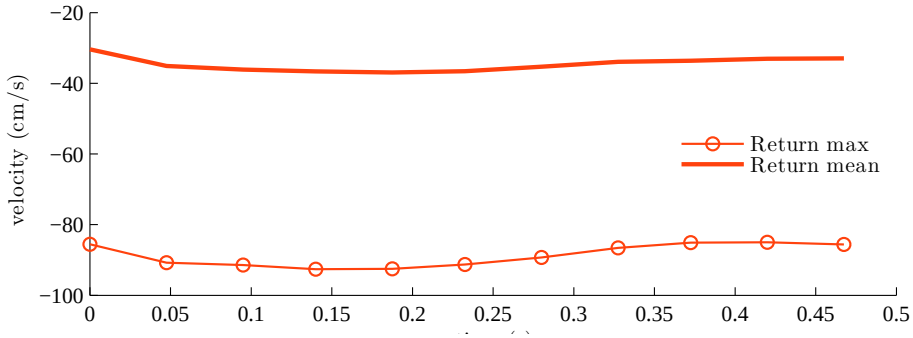

c. Peak and Average Velocity within the return pipe

Figure 7.23: Phase Contrast through plane results

7.4.3.2 Perfusion

Perfusion images were acquired using the PASL and CASL sequences described in section 7.3. Both sequences shared the same 3D-GRASE acquisition, hence imaging parameters were TE = $14.9ms$, TR = $4350ms$, FOV = $240 \times 145 \times 60mm$, $64 \times 36 \times 12$ image matrix, 30 averages, 6/8 partial Fourier and 2× GRAPPA accelerated. PASL images were acquired using a background-suppressed FAIR labelling scheme with an inversion time (TI) of 2500ms. CASL images were acquired with a labelling duration of 3500ms, post labelling delay of 500ms and labelling gradient strength of $3.5mT/m$. Two scanning sessions are presented, the first using the 13mm silicone hose within the pump (as in section 7.4.3.1) where the CASL sequence was run with both coils on full power; and the second using the tygon tubing within the pump (therefore a reduction in flow rate), where the

CASL sequence was run with both coils on full power, and with labelling from the left and right side coils individually. Images have been enlarged in Matlab by a factor of 3 using a Lanczos filter.

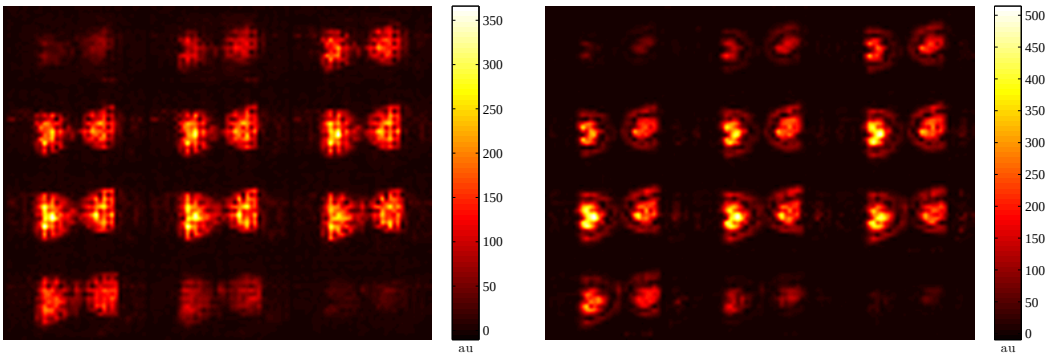

a. FAIR PASL b. CASL

Figure 7.24: High flow rate PASL and CASL ΔM perfusion weighted images

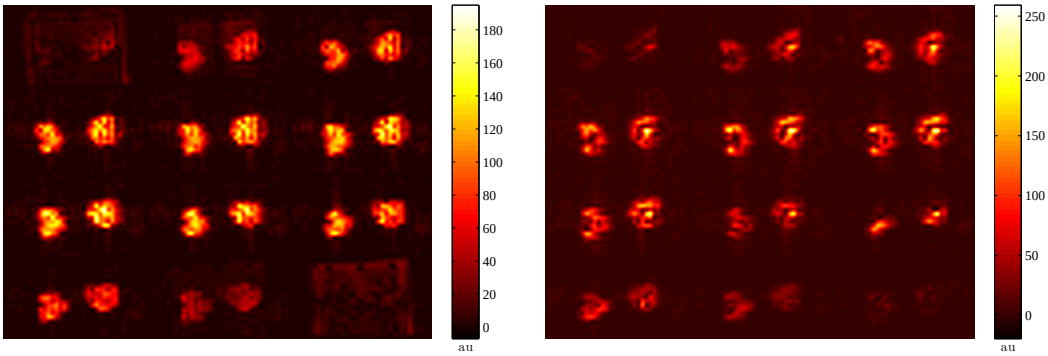

a. FAIR PASL b. CASL

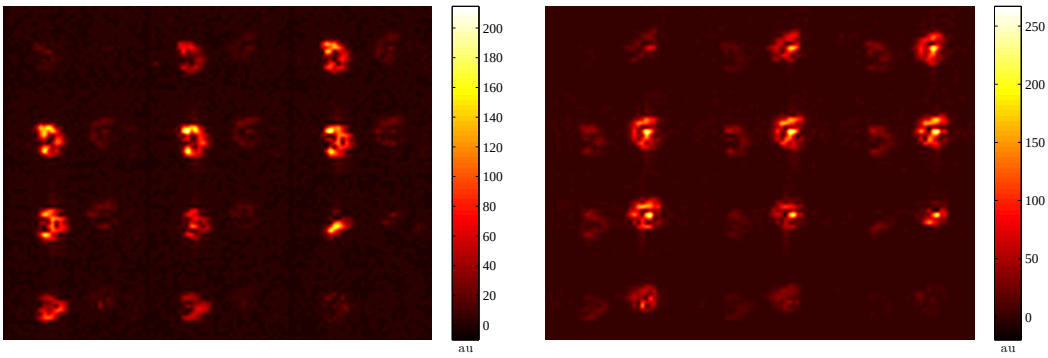

c. CASL: left coil only d. CASL: right coil only

Figure 7.25: Low flow rate PASL and CASL perfusion images

Figure 7.24 shows ΔM images obtained using the 13mm silicone hose within the peristaltic pump, resulting in a high flow rate. Perfusion can clearly be seen in both PASL and CASL images centred around the semicircular arrangement of pipes within each perfusion chamber. A dark band is visible on the CASL images, due to a sequence bug causing incorrect background suppression labelling from slice selective inversion pulses spatially offset from the imaging volume (explained in more detail in section 8.3.1.4). As with the 3D TOF images, flow and hence perfusion is most concentrated in the central pipes in each perfusion chamber.

Figure 7.25 shows ΔM images obtained at a lower flow rate as the the 6.5mm tygon tubing was

used within the peristaltic pump. Because of the lower flow rate not all labelled water enters the perfusion chamber, and consequently the ΔM images show smaller perfusion regions. The distinct perfusion territories of the phantom are exemplified in figures 7.25.c and d, where only the left or right coil was used for labelling. Contralateral partial labelling is also evident, as a result of a non-negligible B_1^+ field from the labelling coil at the opposite pipe.

7.4.4 Discussion and Conclusion

Within the inflow tubes the flow rate was not fast enough to simulate physiological conditions, primarily due to powder from the SLS manufacturing process remaining within the small branched pipes, and an insufficient return path for the water from the perfusion chambers. Once manufactured, the main section of the phantom was filled with nylon powder, which was manually removed using a combination of needles, copper wire and pipe cleaning brushes. However the positions of some of the branching pipes meant that it was not possible to clear away all of the powder within these pipes, obstructing flow. Additionally, water from within the perfusion chambers was supposed to flow via small 0.5mm holes to the return pipe, however the number and size of these was insufficient for the flow rate required.

Both these issues can easily be rectified by modifying the CAD design. To prevent obstructions from unremovable nylon powder the diagonal branching section could be replaced with a large conical chamber which allows easy access. The exit holes from the perfusion chambers could be widened, replaced with a SLS manufactured grill design, or a fine plastic mesh could be fitted over a hole in the partition between the perfusion chambers. Doing so will reduce the water pressure within the system, permitting the 10mm or 13mm silicone hoses without any risk of the tubing disconnecting at the connectors. Further improvements could be made to the design by increasing the size of the perfusion chamber to more realistically model the adult human brain; the current size is limited by the dimensions of the phantom container, and the inflow pipes could branch sooner into fewer, larger pipes, accompanied with more perfusion territories, for vessel selectivity more specific than simple right/left.

It is unfortunate that the entire phantom cannot be made as one single piece; a lid is required for removal of nylon powder remaining from the SLS manufacturing process, and this adds an additional manufacturing phase where the lid must be sealed to the main phantom body. The porous nature of SLS is also a disadvantage as tight seals cannot be made by clamping pieces together, and sections with thin walls must be sealed with epoxy sealant.

The perfusion phantom presented was designed to provide a controlled environment for testing ASL sequences, mimicking both the macro and micro vasculature of the neck and brain. It successfully meets these objectives, giving good ASL signal for both pulsed and continuous labelling schemes. A significant hindrance is the reduced flow rate, meaning that different parameters are

required in comparison with human imaging for optimal ΔM signal. Nonetheless, as shown in figure (7.25).c and d the distinct perfusion territories provide an ideal test-bed for developing vessel selective ASL sequences. By using CAD software and rapid prototyping the phantom was made inexpensively (total cost was less than £200), and the parameter driven design means that dimensions can easily be adapted to a specific anatomy, or even different species for pre-clinical work. Multiple, identical phantoms can also be cost effectively produced, for instance the phantom could be issued for quality assurance tests in multi-centre ASL studies.

7.5 Parallel Transmit CASL Scan Planning Software

An important factor which must not be overlooked is the time required to obtain the pre-requisite data for parallel transmit CASL: B_1^+ maps and anatomical images must be acquired, processed to allow the selection of the label and null regions, and then the amplitudes and phases for each channel calculated. To maximise the amount of useful data acquired in the scan session this needs to be as short as possible. Fast imaging techniques such as partial Fourier and parallel imaging make it feasible to acquire B_1^+ maps and anatomical scans in a matter of minutes, however it is also vital that there is minimal delay from the image processing and subsequent calculation of the complex weights. The proprietary scan planning software used to control an MR system automatically processes and loads images on-screen, robustly calculating appropriate scan parameters based on high-level user input (for example graphically positioning the field of view). Ideally the user input and calculations for parallel transmit CASL would be included as part of the MR system scan planning software, however this is technically challenging. Therefore, standalone parallel transmit CASL scan planning software with a graphical user interface (GUI) was developed, simplifying and automating the procedure.

7.5.1 Principle of Operation

The scan planning software was designed as a Matlab graphical user interface, executing custom Matlab scripts in response to user input. The software was run on a separate computer, accessed via secure socket shell (SSH) from a computer within the MR console room. Figure 7.26 shows the software in operation. Once acquired, AFI data sets for each linear combination and a 3D time of flight angiogram over the same field of view were transferred as dicom images from the MR scanner's console computer to a network share. Operation of the software is described below, each numbered paragraph refers to the boxes in figure 7.26, and is also the procedural order for using the software.

1. Subject details First, the path to the dicom directory with the subject's AFI and 3D TOF data, and a directory on the local computer in which to work from were entered into the appropriate

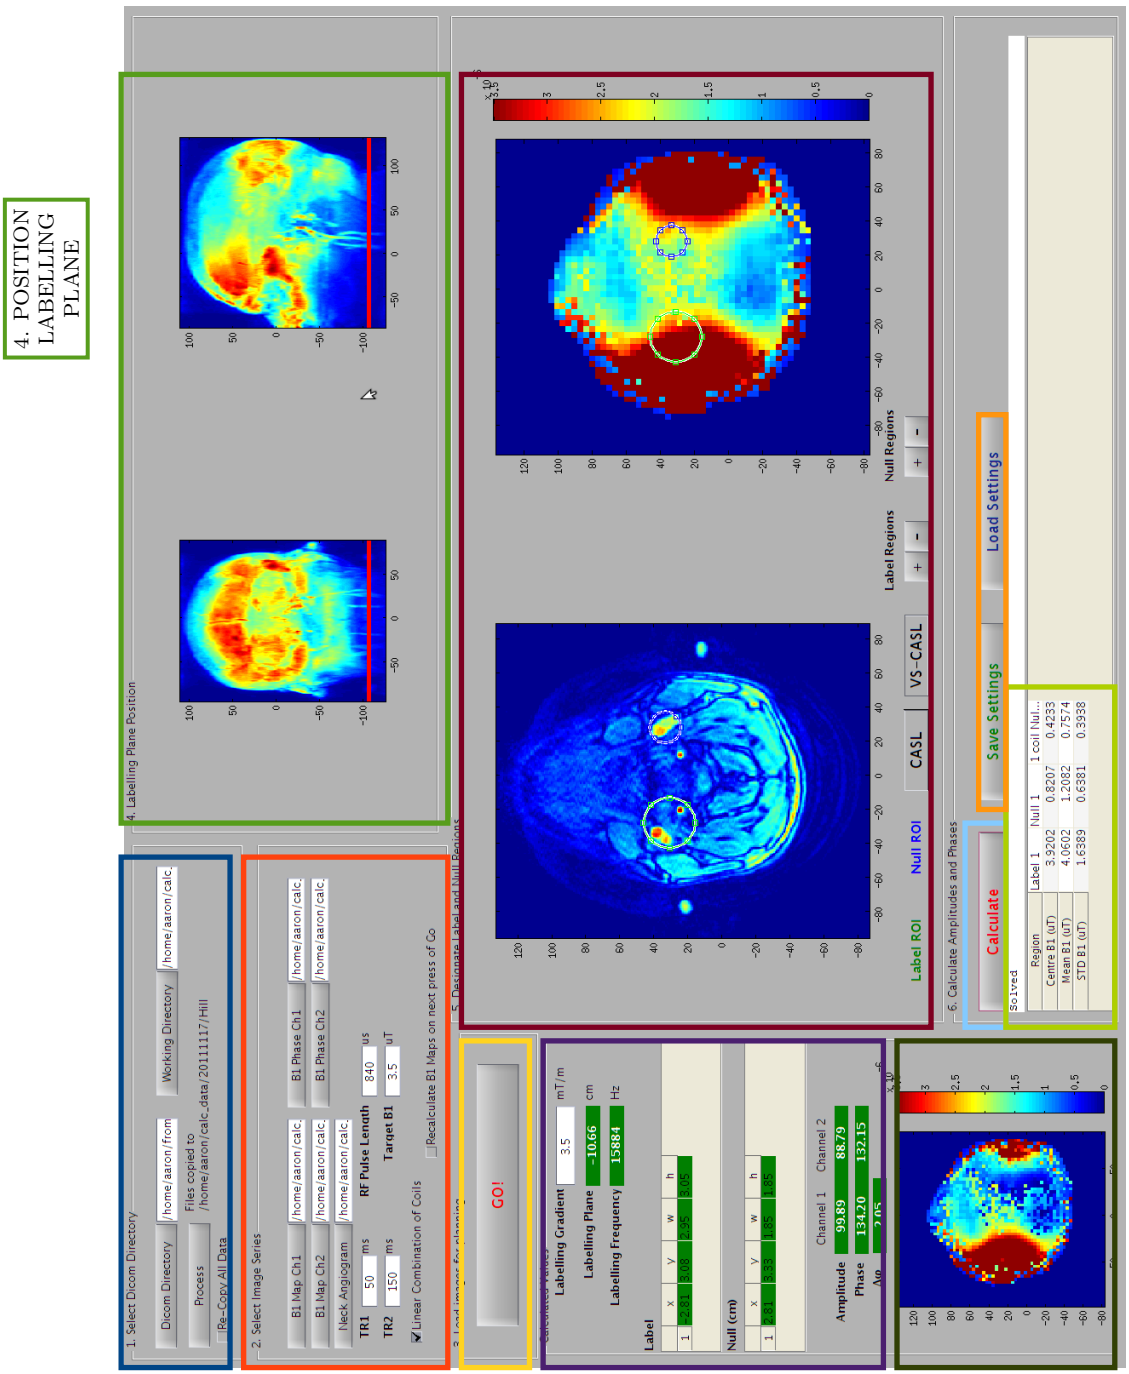

Figure 7.26: Parallel Transmit CASL Planning Software Overview

fields in box #1. Files were then copied to a generated folder, specific to the subject and study date within the working directory, renamed and sorted into numbered folders according to the dicom attribute "Series Number", amounting to being sorted by scan.

2. Image directories and parameters Once the files have been copied, renamed and sorted, the directory paths for the magnitude and phase AFI data, and the 3D TOF angiogram were loaded into the appropriate fields. Parameters required for calculating the B_1^+ maps were also provided: the two TR values from the AFI sequence (by default 50ms and 150ms), RF pulse duration in the AFI sequence (default $840\mu s$), a target B_1^+ strength within the label region (default $3.5\mu T$), a tick box to indicate linear combination B_1^+ maps, and a tick box to force a re-calculation of everything if step 3 is run again.

3. Load and process images A button, marked "Go" was pressed to initiate the image processing. AFI data was converted to NIFTI format using functions from SPM 8 (Statistical Parametric Mapping, Wellcome Trust Centre For Cognitive Neuroimaging, Institute of Neurology, University College London) before being loaded directly into the Matlab workspace. B_1^+ maps were calculated in accordance with the methods described in the section 7.2, and then inverse transformed to yield single coil maps as described in section 7.2.3.2. 3D TOF data, including coronal and sagittal maximum intensity projections (MIP) were also converted to NIFTI format and then loaded into the Matlab workspace. Header information was processed for both B_1^+ maps and angiograms for coordinate registration of the two sets of images. MIP images were then displayed in the position labelling plane box (box #4), and in the select regions box (box #5) transverse angiogram slices and a magnitude sum of the B_1^+ maps from each coil were displayed, the latter giving an indication of the maximum B_1^+ field that can be produced by both coils.

4. Position labelling plane Next, the labelling plane was positioned by the user on either the coronal and sagittal MIP images by dragging the red bar up and down. Oil capsules positioned at the centre of each coil provided a reference for the optimum position, these can be seen on the angiogram slice in box #5 as two hyperintense ellipses either side of the neck. Callback functions ensure that the red bar position is linked on both coronal and sagittal images, and also displays the corresponding angiogram and B_1^+ transverse slices in the select regions box. In the calculated values box (box #9), the labelling plane position with respect to the magnet isocentre was shown, and for a user provided labelling gradient (default $3.5mT/m$) the offset RF frequency for the CASL labelling pulse is calculated and displayed.

5. Select label and null regions Two modes were available: CASL or VS-CASL, toggled by buttons at the bottom of this box. In CASL mode only label regions were available, drawn around any arteries which were to be labelled. When calculating the the amplitudes and phases for

driving the coils, this mode aimed to optimise the average B_1^+ within the label regions to the target B_1^+ value. Vessel selective CASL (VS-CASL) mode allowed the user to select both label and null regions, which were used to calculate appropriate amplitudes and phases which minimised the B_1^+ field within the null region, whilst optimising the average B_1^+ within the label region to the target B_1^+ value. The label and null regions were defined by drawing ellipses on the transverse angiogram image, "+" and "-" buttons at the bottom of the box add or subtract ellipses (minimum of 1 label ellipse in CASL mode and 1 label and 1 null ellipse in VS-CASL mode), allowing the selection of multiple vessels, or the construction of non-elliptoid shaped regions by combining ellipses. In the calculated values box (box #9) the ellipse positions, widths and heights were shown, and could be modified manually by entering new values.

6. Calculate driving amplitudes and phases Once the label and null regions were selected the "Calculate" button was pressed. This determined the indices of all the voxels within each label and null ROI, and then invoked the corresponding solver method described in section 7.5.2.1 for CASL and 7.5.2.2 for VS-CASL to compute the optimised amplitudes and phases. If the calculation was successful (i.e. the optimisation converged) a message displaying "Solved" was shown in the status bar in the regional B_1^+ statistics box, the computed amplitudes and phases displayed in the calculated values box (box #9) and the text boxes in that box toggled from red to green. If the calculation failed, for example if the prescribed target B_1^+ is greater than that is achievable by the coils then a status message saying either "Inaccurate/Solved", "Infeasible", or "Failed", depending on the status of the solver method, the calculated amplitudes and phases will be displayed (however they may be "NaN" or "Inf"), but the text boxes will remain red.

7. Resultant B_1^+ field The resultant B_1^+ field from the coils' B_1^+ maps modulated by the computed amplitudes and phases was calculated and displayed within this box. This gave an indication of whether the amplitudes and phases produce an appropriate labelling field.

8. Regional B_1^+ statistics In this box the resultant B_1^+ fields within the label and null regions was determined and the B_1^+ amplitude at the centre of each region, the mean amplitude within each region, and the standard deviation of the B_1^+ within region was displayed in this box. For VS-CASL B_1^+ statistics were also calculated and displayed for a B_1^+ field where only the coil with the highest amplitude transmitted, providing a useful comparison to see whether parallel transmit VS-CASL provides more specificity than using just one coil.

9. Calculated values Providing the computed B_1^+ fields and statistics were satisfactory the values reported in this box could be used as parameters for a parallel transmit CASL imaging experiment.

10. Load and save settings A snapshot of loaded images, labelling plane position, region positions, parameters and calculated values could be saved as a standard Matlab file using the save settings button, and then recalled using the load settings button.

7.5.2 Amplitude and Phase Determination

Section 3.2 in chapter 3 describes a method for spatially localising the B_1^+ field based on the selection of a label and null region and computing complex driving coefficients which produce a total B_1^+ field that is minimised within the null region, subject to conditions where the B_1^+ must be a target value at the centre of the label region, and the magnitude of the complex driving coefficients must not exceed 1. This mathematical framework was used as a basis within the scan planning software, however modifications were made to improve the robustness of the resultant B_1^+ field, and provide an additional mode in which parallel transmission was used to optimise the labelling field in conventional separate labelling coil CASL.

In the formulation described by equation 3.19, the B_1^+ field is constrained to a target value in the voxel at the centre of the label region. As the carotid arteries typically have a cross-sectional area greater than one voxel, a better value to constrain which describes the B_1^+ field over the entire label region is the average B_1^+ inside the label region. This constraint helps to homogenise the field within the labelling region, and reduces the likelihood of large variations over the label region. In order to maximise the effectiveness of this constraint the label region should be drawn as close to the boundaries of the arteries as possible, using multiple ellipses to cover irregular shaped vessels.

7.5.2.1 CASL

Conventional separate labelling coil CASL typically uses an empirically determined power level, which is not calibrated on a per subject basis. Acquiring B_1^+ maps provides the exact labelling field strength at the carotid arteries, and so can be used to set the coil power appropriately. Furthermore, by using multiple coils, and minimising a cost function, the B_1^+ field within labelling region(s) can be optimised for multiple arteries, using the minimum required power.

Therefore, by specifying an arbitrarily shaped labelling region and providing B_1^+ maps for each coil, the constrained minimisation problem to determine the optimum complex coil driving coefficients is:

$$minimise \ \|\mathbf{c}\|^2 \tag{7.25}$$

subject to

$$\frac{1}{N_{label}} \sum_{i}^{N_{label}} \mathbf{S_i} \cdot \mathbf{c} = B_{1,target}, \ i = Label \ Region \tag{7.26}$$

$$|\mathbf{c}| \leq 1 \tag{7.27}$$

Qualitatively the problem is "Minimise the coil driving coefficients whilst ensuring that the mean B_1 within the label region is equal to a target value $B_{1,target}$, and that the absolute values of the coil driving coefficients do not exceed unity".

7.5.2.2 VS-CASL

Vessel Selective CASL uses the two transmitter coils to achieve a mean B_1^+ field equal to a target value within the label region, whilst minimising the B_1^+ field within the null region. The minimisation problem is:

$$minimise \ \left\|\mathbf{S_i} \cdot \mathbf{c}\right\|^2, \ i \epsilon Null \ Region \tag{7.28}$$

subject to

$$\frac{1}{N_{label}} \sum_{i}^{N_{label}} \mathbf{S_i} \cdot \mathbf{c} = B_{1,target}, \ i = Label \ Region \tag{7.29}$$

$$|\mathbf{c}| \leq 1 \tag{7.30}$$

Qualitatively the problem is "Minimise the B_1^+ field within the null region, whilst ensuring that the mean B_1^+ within the label region is equal to a target value $B_{1,target}$, and that the absolute values of the coil driving coefficients do not exceed unity".

7.6 Summary and Conclusion

In this chapter the methods required to implement parallel transmit CASL have been presented. Magnitude and phase B_1^+ mapping was implemented using an Actual Flip Angle Imaging sequence, and optimised image processing methods have been presented which increase the SNR and accuracy of the resultant B_1^+ maps. A background suppressed continuous arterial spin labelling sequence was implemented using a 3D-GRASE acquisition; a SNR efficient single shot measurement which is also used in a pulsed arterial spin labelling sequence, allowing for direct comparison between the two methods. For the development and testing of vessel selective ASL sequences a novel perfusion phantom with distinct perfusion territories was designed and constructed. Further development is required to accurately model physiological flow and perfusion conditions, however this prototype simulates both macroscopic arterial flow, and microscopic capillary flow/perfusion and provides consistent perfusion images for both continuous and pulsed ASL sequences. Finally scan planning software intended to automate the processing of pre-scan data required for parallel transmit CASL

was presented.

These methods are now ready to be brought together and used to acquire perfusion images on human subjects, which is the topic of chapter 8.

Chapter 8

Parallel Transmit Continuous Arterial Spin Labelling

8.1 Introduction

Arterial spin labelling permits the non-invasive measurement and assessment of cerebral perfusion, doing so by magnetically labelling the protons within blood water, making them an endogenous, diffusible contrast agent. Vessel selective ASL methods are designed to only label blood within specific arteries, permitting visualisation of their perfusion territory. Techniques to selectively label blood vessels have existed for almost as long as the ASL technique itself [54], and in recent years the ability to assess and visualise individual perfusion territories has found increasing clinical importance, for example in acute stroke and chronic cerebrovascular disease[221].

Separate labelling coils have two distinct advantages when used for vessel selective ASL: as they use the CASL labelling method a large bolus of blood can be labelled, hence a large ΔM signal; and they do not suffer from MTC within the imaging volume, therefore a control image can be acquired with no RF power applied. Vessel specificity is improved for smaller labelling coils, however this also affects the penetration of the coil, and means that the coil must be well placed to ensure optimum labelling. Increasing the size of the coil relaxes the need for exact placement, but as a result a non-negligible B_1^+ field will be present at the other arteries within the neck, causing partial labelling within these arteries.

To improve specificity it is proposed to use parallel transmission to null the B_1^+field in contralateral arteries, thus reducing partial labelling. By doing so, larger labelling coils can be used, and labelling within the selected vessels becomes more tolerant to non-optimal coil placement. In this work parallel transmission refers to transmitting independently on multiple RF coils, with the amplitudes and phases set such that the interference between their B_1^+ fields spatially tailors the total field.

In this chapter the ASL Labelling Coils presented are used in conjunction with the SMIS console based transmitter system to acquire CASL perfusion images on human subjects. Global (i.e., non-regional) perfusion maps are acquired with the ASL Labelling Coils and 3D-GRASE CASL sequence, and compared with perfusion images obtained with a FAIR 3D-GRASE pulsed ASL sequence. Parallel transmit vessel selective CASL is used to selectively label the right or left carotid arteries, and compared with using a single surface coil for specificity.

Inter-coil coupling means that it is not possible to rely only on the computed driving coefficients, as the acquired B_1^+ maps and the ability to control the amplitude and phase of the coils is compromised. However, this does not fully compromise the ability to perform parallel transmit vessel selective CASL. Appropriate driving coefficients can be manually found from FLASH images acquired whilst transmitting on the ASL Labelling Coils. An accelerated FLASH volume can be acquired in approximately 30s, and the average value within ROIs drawn around vessels within the labelling plane used to determine the relative changes in the transmit field in response to different driving coefficients. While this may not arrive at the optimal solution, it should however be possible to find a solution in which a null can be positioned over either the left or right carotid artery.

8.2 Method

All scans were performed on a 3T whole body scanner (Tim Trio, Siemens Medical Systems, Erlangen, Germany), using a 32 channel head coil for reception and birdcage body coil for transmitting the imaging RF pulses. Labelling was performed by the ASL Labelling Coils (described in chapter 4), interfaced to a low power, two channel transmitter system, based around a recycled spectrometer (MR5000, Surrey Medical Imaging Systems, Guildford, UK), as described in chapter 5. Triggering of the transmitter system was achieved by an optical pulse from the Trio, which was converted to a TTL level signal and detected by the SMIS pulse sequence. Control of the amplitudes and phases was achieved with three parameters: channel 1's amplitude, A_1, channel 2's amplitude, A_2, and the relative phase difference between the two channels, $\Delta\phi$. Amplitudes are given as percentages of the maximum RF output level from the spectrometer (10dBm).

8.2.1 Phantom Imaging Protocol

Vessel selective labelling was performed on the perfusion phantom presented in section 7.4. The phantom was positioned with the centres of the perfusion chambers at the isocentre of the MR scanner inside the receive coil. The ASL Labelling Coils were positioned either side of the phantom, such that the centre of the coils coincided with the middle of the inflow pipes. First, B_1^+ maps and a 3D TOF angiogram were acquired over the entire phantom volume. The data was then processed by the scan software (section 7.5), and set accordingly to provide the correct off-resonance labelling

frequency. Vessel selective scans were acquired for both sides with the labelling coil (the coil closest to the labelling target) amplitude set to 80% and 100%, using the procedure described in section 8.2.3. For each of these labelling amplitudes a global (both coils with the same amplitude and the phase difference, $\Delta\phi = 0°$) CASL scan was run. Scan parameters were as follows:

- B_1^+ maps using the AFI sequence described in section 7.2.1: FOV = $150\times140mm$, 6.3mm slice thickness, 40 slices, 64×60 image matrix, $TR_1 = 50ms$, $TR_2 = 150ms$, TE = $3.05ms$, pulse length $500\mu s$, 2× GRAPPA accelerated, 6/8 partial Fourier in both phase encode dimensions. Linear combination 1 was acquired with $\phi_1 = 90°$, $\phi_2 = 0°$; and linear combination 2 acquired with $\phi_1 = 0°$ and $\phi_2 = 90°$.

- 3D Time of Flight Angiogram over the same volume as the B_1^+ maps: FOV = $150 \times 140mm$, 3.75mm slice thickness, 64 slices, 256×240 image matrix, TE = $3.87ms$, TR = $20ms$, $\alpha = 17°$, 4× GRAPPA accelerated.

- CASL: TE = $14.9ms$, TR = $4350ms$, FOV = $240 \times 145 \times 60mm$, $64 \times 36 \times 12$ image matrix, 30 averages, 6/8 partial Fourier and 2× GRAPPA accelerated, 3.5s labelling duration, 0.5s post labelling delay, $3.5\mu T$ labelling gradient.

Images were processed as described in section 8.2.4.

8.2.2 Human Imaging Protocol

The ASL Labelling Coils were positioned on either side of the neck and held in place with pieces of foam, pushed against the head receive coil. On all subjects B_1^+maps, a 3D TOF Angiogram, global PASL and global CASL scans were run. Section 8.2.5 describes the differences in subsequent scans. For those subjects on which vessel selective labelling was performed, the procedure described in section 8.2.3 was performed for each labelling configuration. In all cases, the scan planning software was used to locate the labelling plane and calculate the off-resonance labelling frequency accordingly. Scan parameters were as follows:

- B_1^+ Maps using the AFI sequence: FOV = $185 \times 220mm$, 7.5mm slice thickness, 32 slices, 54×64 image matrix, $TR_1 = 50ms$, $TR_2 = 150ms$, TE = $3.05ms$, pulse length $840\mu s$, 2×GRAPPA accelerated, 6/8 partial Fourier in both phase encode dimensions. Linear combination 1 acquired with $\phi_1 = 90°$, $\phi_2 = 0°$, and linear combination 2 acquired with $\phi_1 = 0°$ and $\phi_2 = 90°$.

- 3D Time of Flight Angiogram: FOV = $185 \times 220mm$, 3.75mm slice thickness, 64 slices, 216×256 image matrix, TE = $3.87ms$, TR = $20ms$, $\alpha = 17°$, 4× GRAPPA accelerated, same volume as the B_1^+ maps

- 3D-GRASE acquisition parameters (shared by PASL, CASL, M_0 and the saturation sequences): 12 slices, 5mm slice thickness, TE = $14.9ms$, TR = $4350ms$, FOV = $240 \times 145 \times 60mm$, $64 \times 36 \times 12$ image matrix, 30 averages, $^6/_8$ partial Fourier and 2× GRAPPA accelerated in the PE direction.

- CASL: 2s labelling pulse, 1s post labelling delay, background suppression inversion pulses with TI's of 2.41s and 0.64s.

- PASL: 1.8s TI, 0.8s bolus length for scans with Q2TIPS pulses, background suppression inversion pulses with TI's of 1.15s and 0.38s.

- M_0 map: TR = $10000ms$,

- Saturation recovery: TR = $10000ms$, $T_{sat} = 250, 500, 750, 1000, 1500, 2500, 3500, 4500ms$.

Global CASL scans were acquired by pulsing on both coils at full power, $A_1 = A_2 = 100\%$, and$\Delta\phi = 0°$.

8.2.3 Parallel Transmit Vessel Selective CASL Procedure

Once acquired, the B_1^+ maps and 3D TOF angiogram data were processed using the scan planning software presented in section 7.5. For each labelling configuration (i.e. label the left side, nulling the right carotid artery, and label the right side, nulling the left carotid artery), label and null regions were designated within the scan software, and the corresponding amplitudes and phases computed for a target labelling amplitude of $3.5\mu T$. Because of inter-coil coupling these were then adjusted manually. For the right label condition, the right coil was set to the desired labelling amplitude (80% or 100% in the phantom measurements, 100% for humans), and then the left coil amplitude, A_2 and the relative phase difference between the two coils, $\Delta\phi$ was adjusted such that a null in the total transmitted B_1^+ field was positioned over the left carotid artery within the labelling plane. Conversely, in the left label condition the left coil was held at the labelling amplitude, and A_1 and $\Delta\phi$ was adjusted to position a null over the right carotid artery. In each case a solution was iteratively found through trial and error by checking in a FLASH acquisition acquired with coils transmitting to confirm placement of null. The FLASH acquisition had the same slice positions as the B_1^+ maps, double the in-plane resolution. During the iterative checking, 4× GRAPPA acceleration was used to acquire images in approximately 30 seconds, then once the location of the null was confirmed a 2× GRAPPA acceleration acquisition was made to produce an image with more SNR.

For each vessel selective configuration, a parallel transmit and single coil labelling CASL sequence was run. Scan sessions lasted 90-120 minutes, and as many vessel selective configurations were run as possible in this time. Because of the time-consuming nature of finding appropriate

amplitudes and phases, only subjects A2, A4 and B3 (see section 8.2.5 for subject codes) had both left and right labelling conditions performed.

8.2.4 Image Processing

Image processing performed in Matlab (The Mathworks Inc., Natick, MA). Masks were automatically generated using the global CASL perfusion weighted (ΔM) data. Two computationally determined thresholds were used to create a whole brain mask (excluding only the background) and a pseudo grey-matter mask (containing only the most intense voxels, roughly corresponding to the grey matter as it has the largest ΔM signal). Left and right ROIs were also drawn onto the central slice to demarcate the left and right hemisphere; these regions were then used to split the pseudo grey matter mask into left and right masks.

8.2.4.1 Cerebral Blood Flow Quantification

T_1 and M_0 maps were determined by fitting the saturation recovery data to $S = M_0\left(1 - e^{-t/T_1}\right)$ on a voxel-by-voxel basis using a non-linear least squares curve fitting algorithm. CBF was quantified using the general kinetic model [33]. For pulsed ASL where the time between the FAIR inversion pulse and the 90° excitation pulse commencing the 3D-GRASE acquisition, TI_2, is greater than the sum of the bolus length, TI_1, and the bolus arrival time, Δt, this was calculated on a voxel by voxel basis using:

$$\mathrm{CBF}_P = \frac{\lambda \Delta M}{2 M_0 \mathrm{TI}_1 e^{-\mathrm{TI}_2/T_{1a}} q_p} \tag{8.1}$$

$$q_p = \frac{e^{k\mathrm{TI}_2}\left(e^{-k\Delta t} - e^{-k(\mathrm{TI}_1+\Delta t)}\right)}{k\mathrm{TI}_1}$$

$$k = \frac{1}{T_{1a}} - \frac{1}{T_1'}$$

$$\frac{1}{T_1'} = \frac{1}{T_1} + \frac{f}{\lambda} \tag{8.2}$$

As the term $\frac{f}{\lambda}$ makes only a small change to the apparent longitudinal relaxation rate, T_1', it was omitted in equation 8.2 such that T_1' was set to equal to T_1, andΔt was set to reflect the mean transit time within the brain, 0.82s [182]. Where no Q2TIPS pulses were used to temporally define the PASL labelling bolus, the value for TI_1 was assumed to be 0.8s, a value previously determined by fellow members of the lab in work investigating the MT effects of Q2TIPS pulses [47].

CASL CBF was calculated on a voxel by voxel basis using:

$$\mathrm{CBF}_C = \frac{\lambda \Delta M}{2\eta_{inv} M_0 T_1 q_c} e^{\delta/T_{1a}} e^{(t_{pld}-\delta)/T_1'} \tag{8.3}$$

$$q_c = 1 - e^{-t_l/T_1'}$$

Where t_{pld} is the post labelling delay; η_{inv} is the CASL labelling efficiency, assumed to be 1; δ is the tissue transit time, set to 1.2s; λ is the partition coefficient for water, 0.9ml/g; T_{1a} is the longitudinal relaxation time of the arterial blood, 1.664s; For measurements made without the saturation recovery sequence to obtain a T_1 map, T_1 was set to that of grey matter, 1.3s.

8.2.4.2 Vessel Selective ASL Specificity

The relative inversion efficiency, ξ, was calculated as the ratio of the vessel selective CASL perfusion weighted image, ΔM_{vs} and global perfusion weighted image, ΔM_{global}, as given in equation 8.4.

$$\xi = \frac{\Delta M_{vs}}{\Delta M_{global}} \tag{8.4}$$

Left/right specificity was assessed using two metrics of the mean value within the left and right grey matter masks. The first metric was the arithmetic mean, and standard deviation of all the voxels within each mask. The second was determined from histograms calculated of ξ (500 bins spaced between 0 and 2) for global, left and right grey matter masks, and only for voxels within the data range of $0 < \xi \leq 2$. Global histogram bin counts were fitted to a dual gaussian function (equation 8.5) using a non-linear least squares curve fitting algorithm:

$$f(x) = c_1 e^{-\left(\frac{x-\mu_1}{2\sigma_1}\right)^2} + c_2 e^{-\left(\frac{x-\mu_2}{2\sigma_2}\right)^2} \tag{8.5}$$

Left and right histogram bin counts were fitted to a single gaussian function (equation 8.6) using a non-linear least squares curve fitting algorithm. Initial values for the gaussians were amplitude, $c_1 = c_2 = 10$, standard deviations $\sigma_1 = \sigma_2 = 0.1$, and mean values $\mu_1 = 0.1$ and $\mu_2 = 0.9$.

$$f(x) = ce^{-\left(\frac{x-\mu}{2\sigma}\right)^2} \tag{8.6}$$

Initial amplitude and standard deviations were the same as with the dual gaussians. To assist convergence of the fitting algorithm, the initial mean values were given *a priori*: for labelling left, left ROI histogram $\mu = 0.9$, right ROI histogram $\mu = 0.1$; labelling right, left ROI histogram $\mu = 0.1$, and right ROI histogram $\mu = 0.9$.

The percentage change in relative inversion efficiency between parallel transmit and single coil labelling was quantified for the mean value of ξ within each ROI as:

$$\Delta \langle \xi \rangle = \frac{\left(\langle \xi \rangle_{pt} - \langle \xi \rangle_{sc} \right)}{\langle \xi \rangle_{sc}} \times 100\%$$

And for the mean value of the fitted single gaussian function, μ, corresponding to parameter

c_2:

$$\Delta\mu = \frac{(\mu_{pt} - \mu_{sc})}{\mu_{sc}} \times 100\%$$

8.2.5 Subjects

A total of nine subjects were scanned in three separate groups, in accordance with the local research ethics protocol.

Group A Four subjects (2 female, mean age 25): PASL (without Q2TIPS pulses, TI=1.8s), global CASL, M_0 and vessel selective CASL sequences were run. All subjects were scanned with a label left condition, two (both female) were additionally scanned with a label right condition. A previously unknown error in the background suppression implementation (see section 8.3.1.4) was present in these scans. The error caused the background suppression slice selective inversion pulse to not spatially coincide with the imaging volume, but instead was shifted distally by half the slab thickness (30mm), resulting in half of the volume receiving correct background suppression, and the other half only seeing the saturation and non-selective inversion pulses. When subject 4 was scanned, the label left condition was performed with two different sets of amplitudes and phases, to assess how this affects the degree of nulling over the right carotid artery and subsequent labelling specificity.

Group B three subjects (2 female, mean age 23): PASL, global CASL, saturation recovery and vessel selective CASL sequences were run. All were scanned with a label left condition, one subject (female) was scanned with an additional label right condition. No background suppression was used in the CASL and vessel selective CASL scans.

Group C Two male subjects (mean age 32): PASL, global CASL and saturation recovery sequences were run. Background suppression was correctly implemented in these CASL scans by increasing the size of the saturation and inversion slab to 120mm; twice the imaging volume thickness, ensuring coverage over the entire imaging volume. Q2TIPS saturation pulses were used in the PASL scans to shorten the bolus length to 0.8s. The saturation recovery sequence used on subject 1 of this group was run with an error in the protocol, and consequently no M_0 or T_1 map could be produced. As an approximation the control images in a non-background suppressed CASL sequence were used as an M_0 map.

8.3 Results

8.3.1 Global Perfusion

8.3.1.1 Group A

Figures 8.1 to 8.4 show ΔM and calculated CBF maps for the PASL and global CASL scans for the subjects in group A. Table 8.1 presents the average CBF over the global pseudo grey-matter mask for the PASL and CASL scans. CASL ΔM images are reduced in intensity compared with the PASL images, also resulting in reduced CBF quantification. The average ratio of the average CASL and PASL perfusion signal within the global, left and right masks was calculated to be 0.54±0.10, 0.53±0.11 and 0.54±0.09, respectively; i.e. the CASL ΔM intensities were approximately half of that from the PASL images. The spatial distribution of the ASL signal in the CASL ΔM images is in good agreement with those acquired by the PASL sequence, particularly for the female subjects. Less perfusion signal is observed at the rear of the brain in the CASL images, this is more pronounced in the two male subjects. Quantification from the PASL scans show an overestimation of CBF, indicating that the bolus length was greater than 0.8s.

	$\langle CBF_P \rangle$	$\langle CBF_C \rangle$
A1	116.75	25.02
A2	185.46	48.65
A3	180.62	48.26
A4	151.11	28.93

Table 8.1: Group A PASL and CASL mean CBF values

8.3.1.2 Group B

Figures 8.5 to 8.7 show ΔM and calculated CBF maps for the PASL and global CASL scans for the subjects in group B. Table 8.2 presents the average CBF over the global pseudo grey-matter mask for the PASL and CASL scans. Both female subjects show CASL ΔM images with an intensity approaching that of the PASL ΔM images (average CASL/PASL ΔM ratios of 0.89 and 0.98 for subjects B1 and B3. Subject B2 does not fit with this trend and shows a reduced CASL perfusion signal (global ratio of 0.57), particularly in distal slices, suggesting labelling parameters were sub-optimal for the tissue transit time of this subject. As with the images from group A, the CASL perfusion weighted images show a reduction in the ASL signal at the rear of the brain.

	$\langle CBF_P \rangle$	$\langle CBF_C \rangle$
B1	144.02	55.75
B2	153.42	47.78
B3	175.22	75.39

Table 8.2: Group B PASL and CASL mean CBF values

8.3.1.3 Group C

Figures 8.8 and 8.9 show ΔM and calculated CBF maps for the PASL and global CASL scans for the subjects in group C. Table 8.3 presents the average CBF over the global pseudo grey-matter mask for the PASL and CASL scans. With Q2TIPS pulses correctly employed the CASL and PASL scans show similar values for CBF. The CASL ΔM images are significantly higher in intensity than the PASL ΔM images. The average CASL/PASL ratio 1.94 ± 0.15 within the global mask, 1.88±0.02 for the left mask and 1.97±0.26 for the right mask. Additionally,

	$\langle CBF_P \rangle$	$\langle CBF_C \rangle$
C1	55.16	52.72
C2	47.20	41.42

Table 8.3: Group C PASL and CASL mean CBF values

for subject C2 the spatial perfusion distribution in the CASL images is not as "full" as in the PASL images, indicating a tissue transit time longer than the post labelling delay.

8.3.1.4 Incorrect background suppression

The effects of incorrect background suppression are illustrated in figure 8.10. The pulsed ASL acquisition shown in 8.10.a was performed with a correct background suppression scheme. The slice selective (control) image at the top has a low signal magnitude (colour scales to approximately 300), and the resultant ΔM image has low levels of noise and other artefacts.

In 8.10.b the incorrect background suppression, as described in section results in the first half of the slab (the first six slices) receiving only the saturation and non-selective inversion pulses, and the last half (last six slices) receiving saturation, slice selective inversion and non-selective inversion pusles; subsequent image contrast in the top control images is different for the two halves of the volume. In this particular example the ΔM images have significantly more noise than would normally be present (the ΔM images from group A also have this incorrect background suppression but do not feature this degree of noise), possibly an artefact from motion (for instance eye motion or swallowing), but in general the incorrect background suppression reduces the perfusion signal.

No background suppression is used in figure 8.10.c, resulting in high magnitude control tag and control images as the static tissue signal is not attenuated at all. The ΔM images show greater intensity and good agreement with the spatial distribution of the perfusion regions in the PASL ΔM images, however there is a lot of noise and some residual static tissue signal.

Figure 8.10.d was acquired with correct background suppression pulses over the entire imaging volume, resulting in control images with an intensity similar to the PASL control images, and ΔM images which are reduced in noise compared to not using background suppression, and are approximately twice the intensity of the PASL ΔM images.

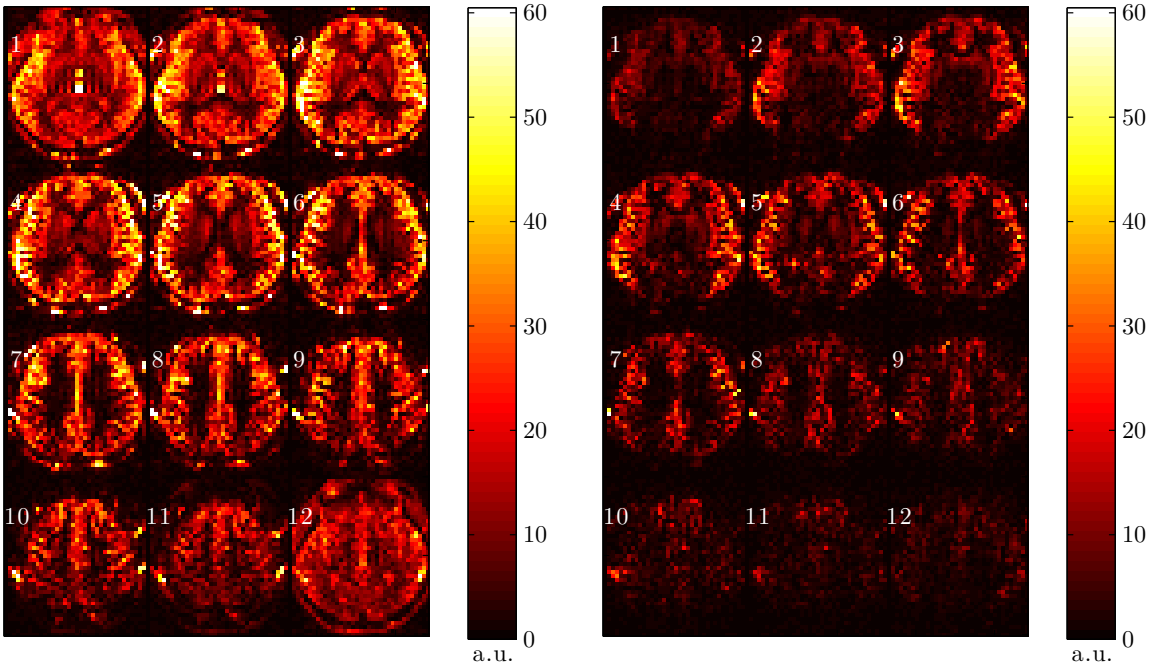

a. PASL ΔM

b. CASL ΔM

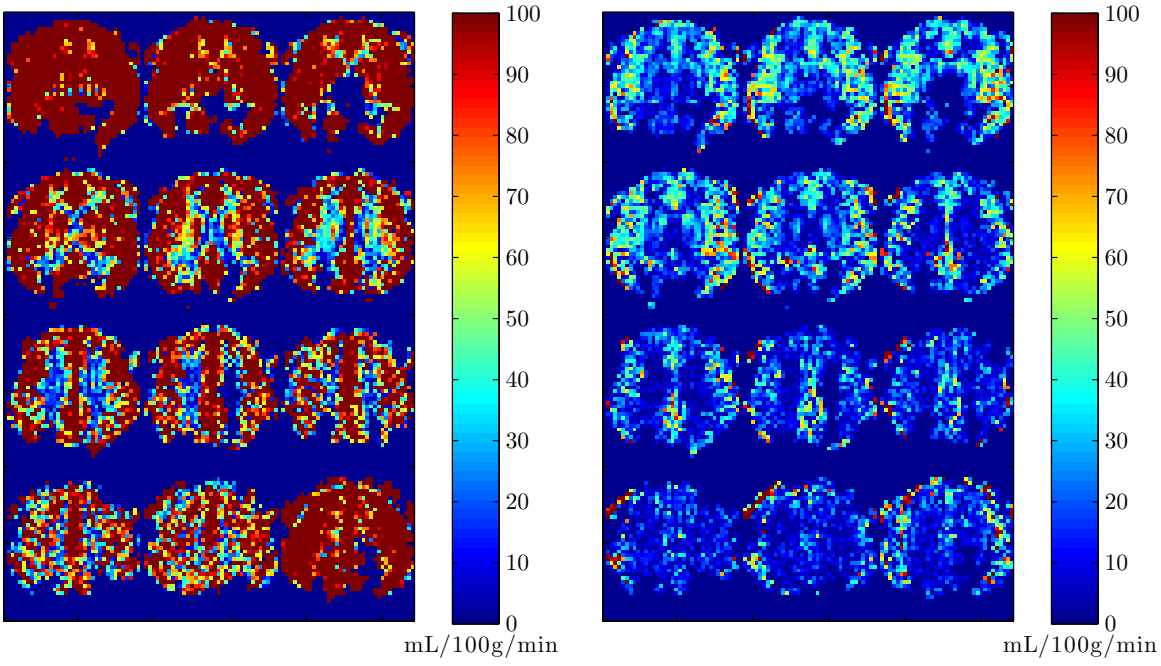

c. PASL CBF

d. CASL CBF

Figure 8.1: Group A, subject 1 (male, age 27) PASL and CASL ΔM and CBF images

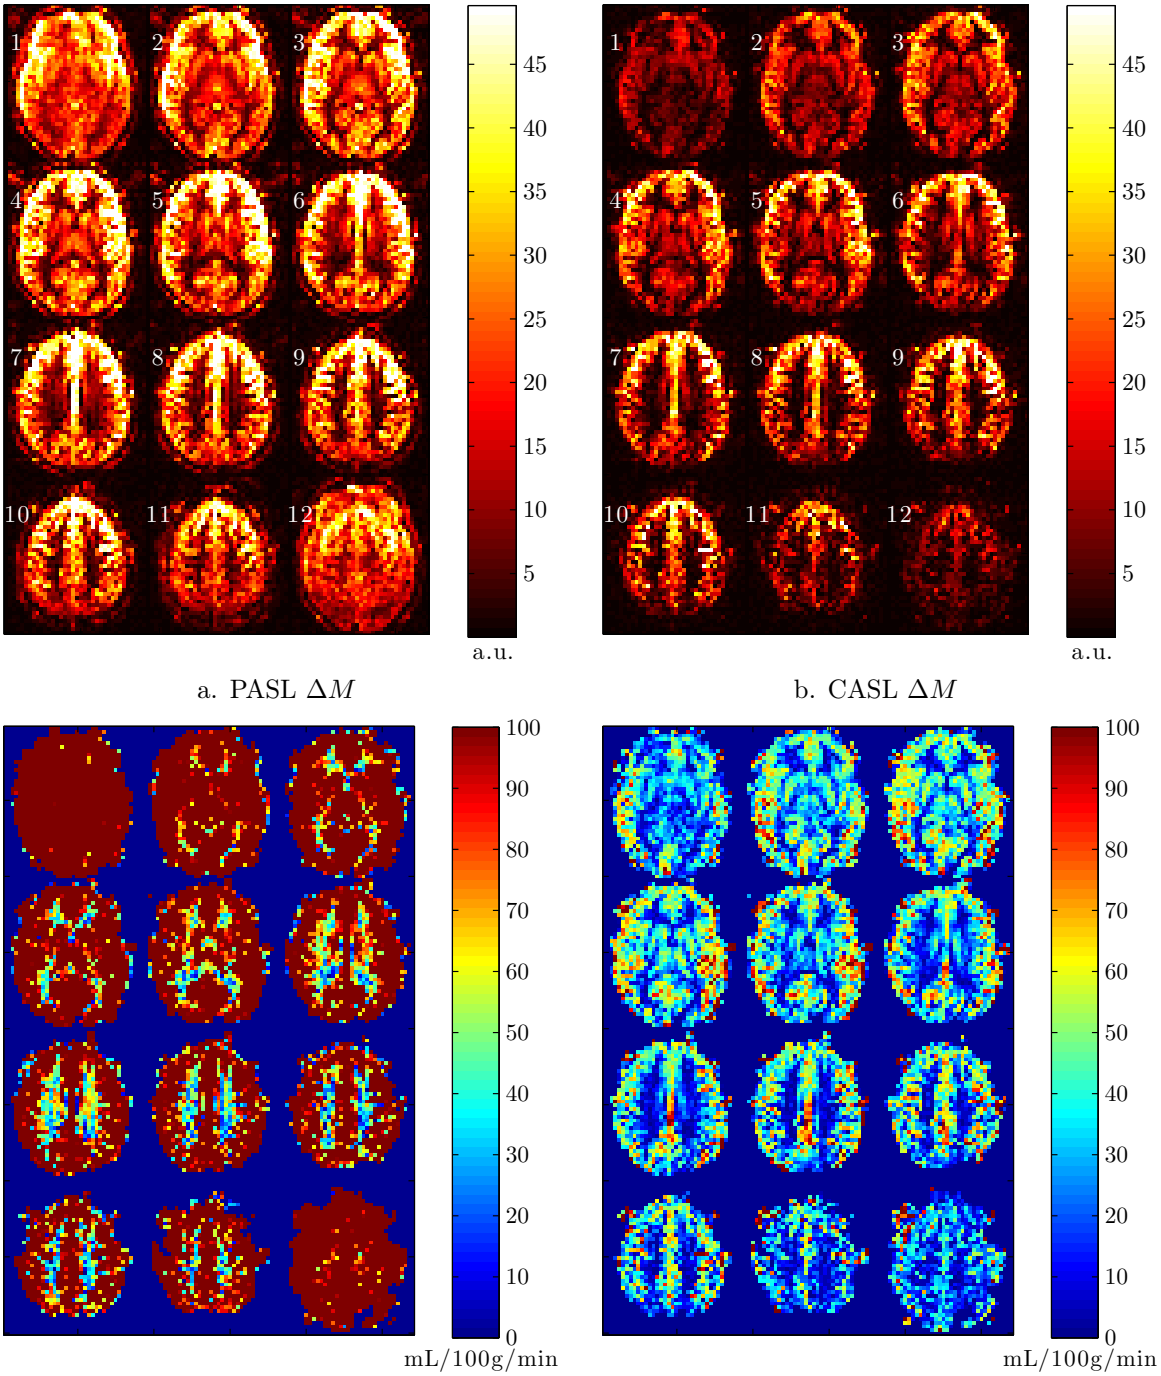

a. PASL ΔM

b. CASL ΔM

c. PASL CBF

d. CASL CBF

Figure 8.2: Group A, subject 2 (female, age 26) PASL and CASL ΔM and CBF images

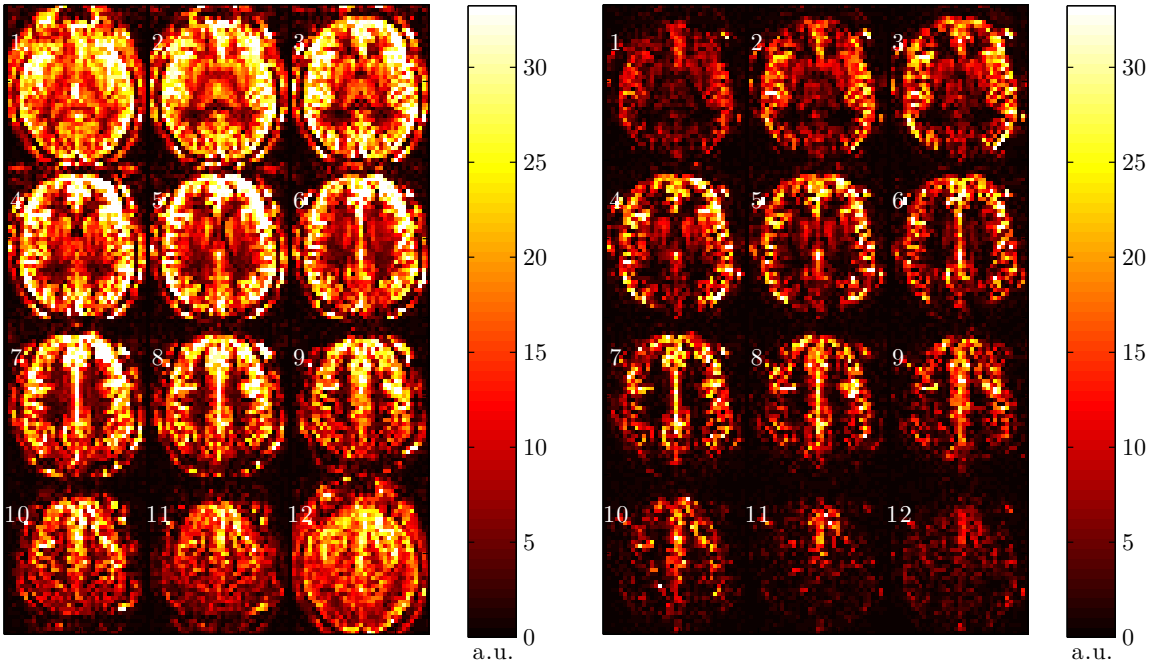

a. PASL ΔM

b. CASL ΔM

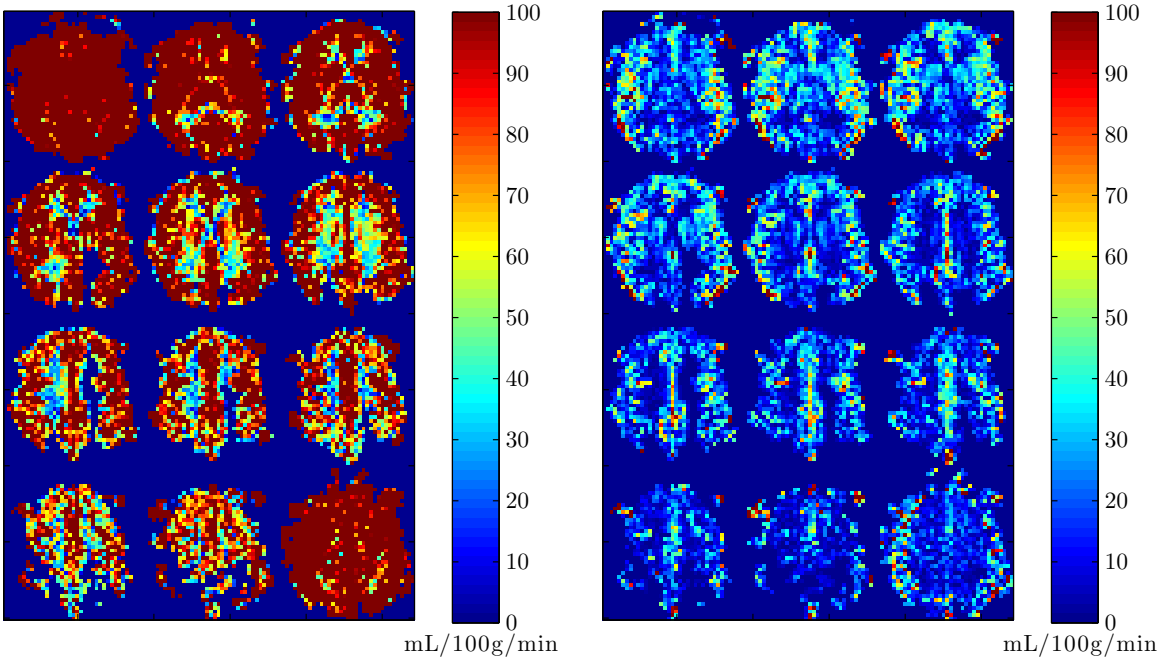

c. PASL CBF

d. CASL CBF

Figure 8.3: Group A, subject 3 (male, age 25) PASL and CASL ΔM and CBF images

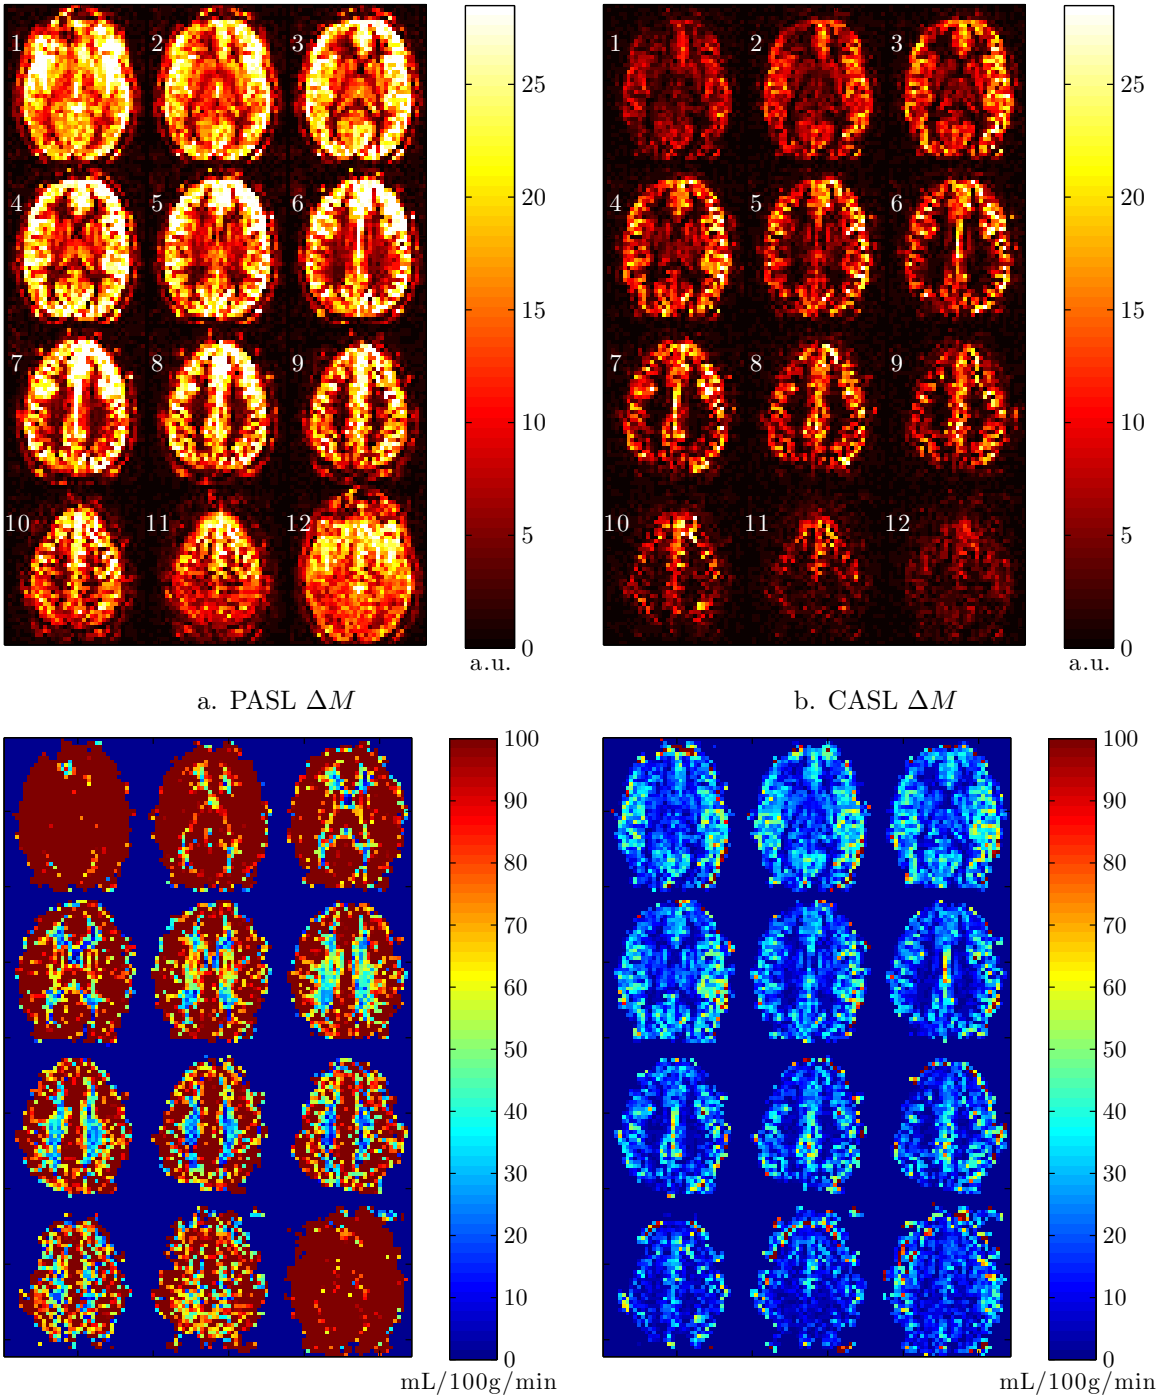

a. PASL ΔM

b. CASL ΔM

c. PASL CBF

d. CASL CBF

Figure 8.4: Group A, subject 4 (female, age 23) PASL and CASL ΔM and CBF images

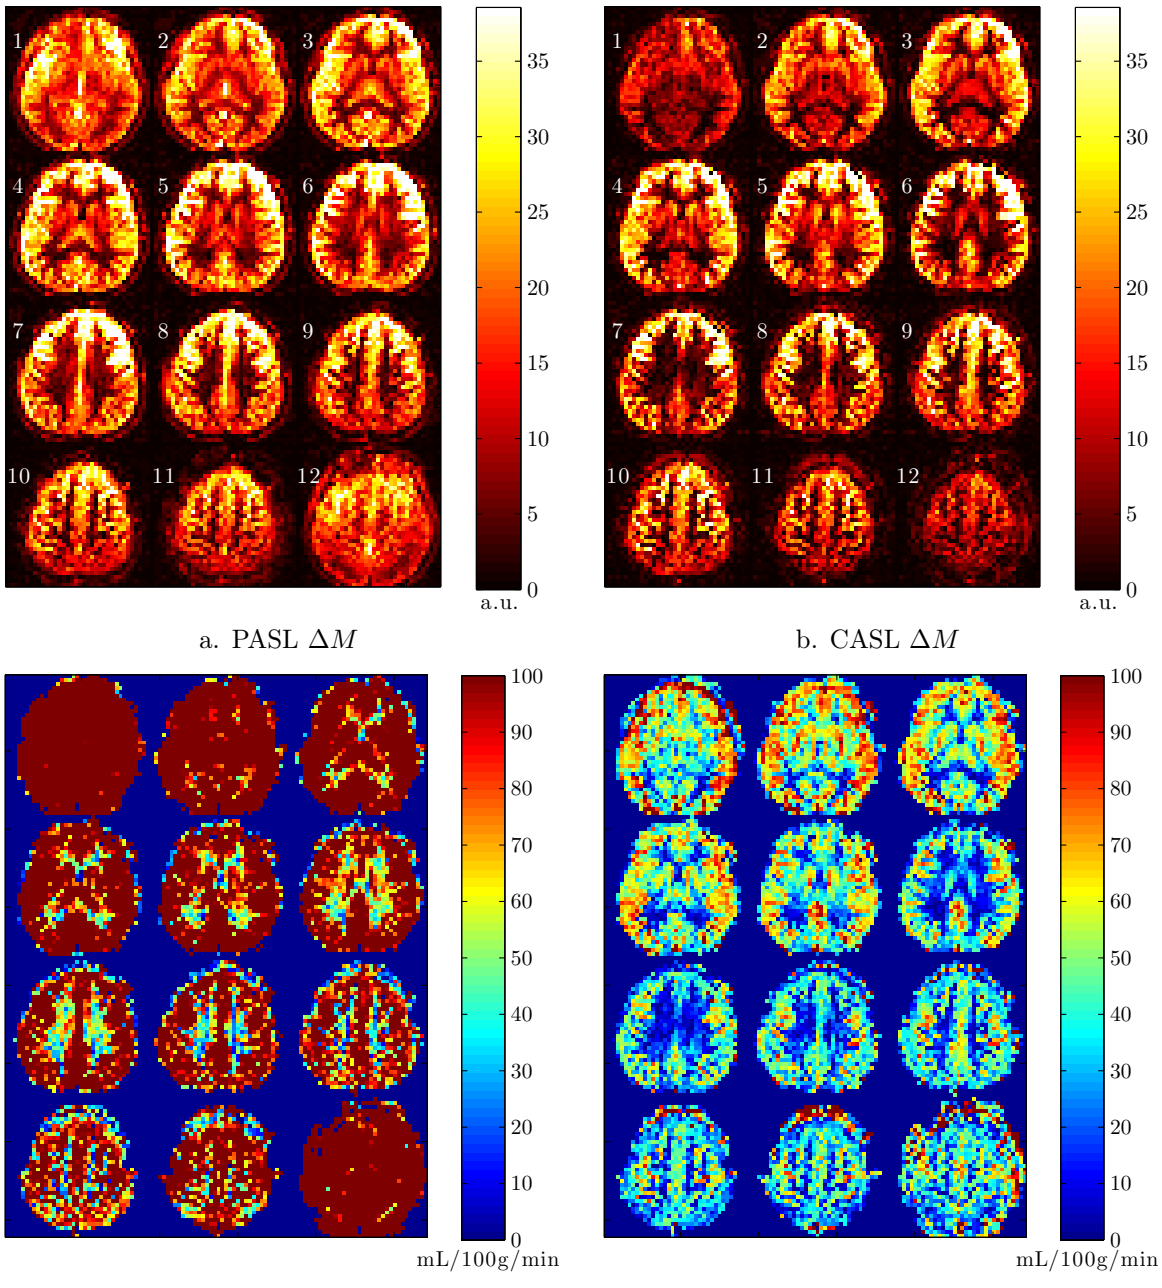

a. PASL ΔM

b. CASL ΔM

c. PASL CBF

d. CASL CBF

Figure 8.5: Group B, subject 1 (female, age 23) PASL and CASL ΔM and CBF images

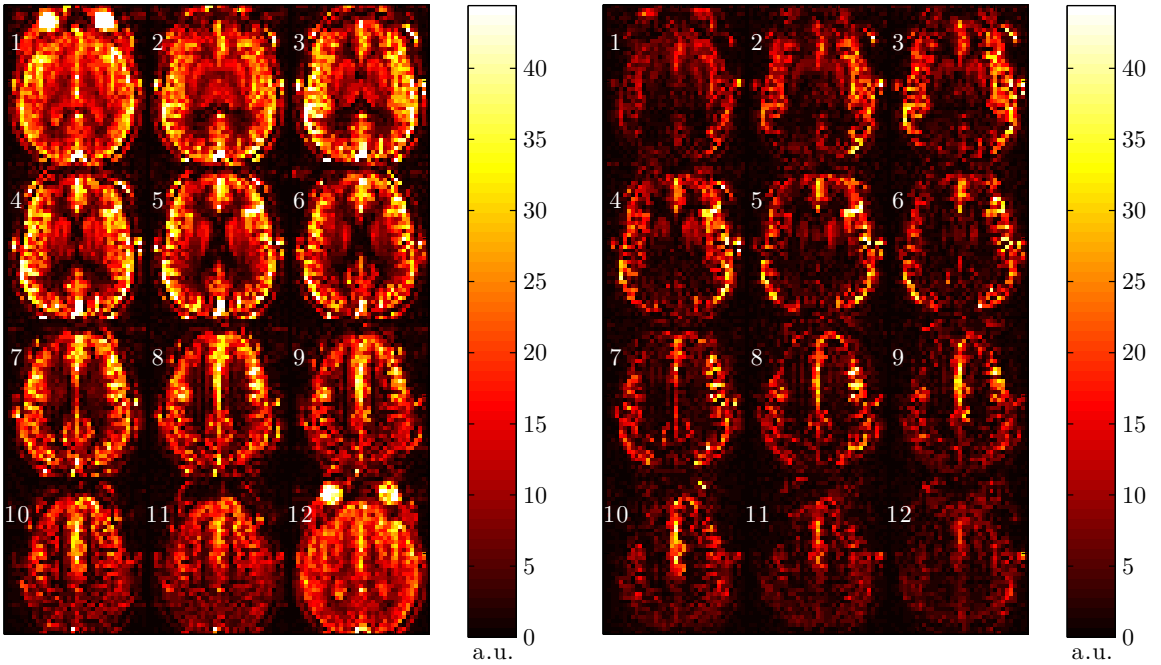

a. PASL ΔM

b. CASL ΔM

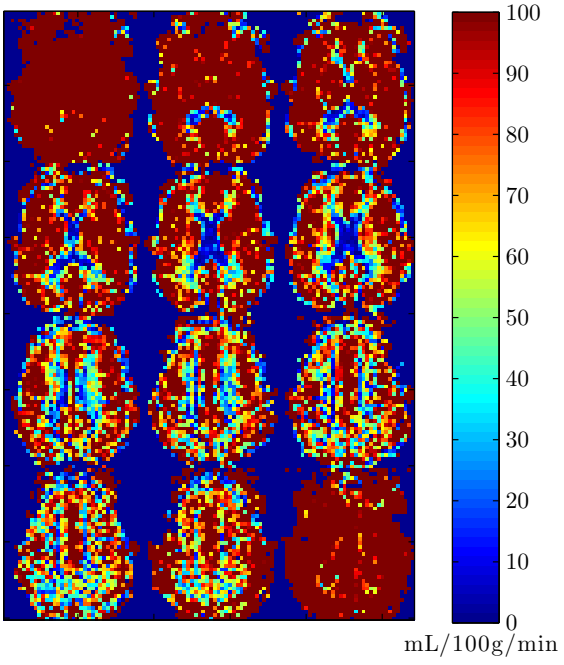

c. PASL CBF

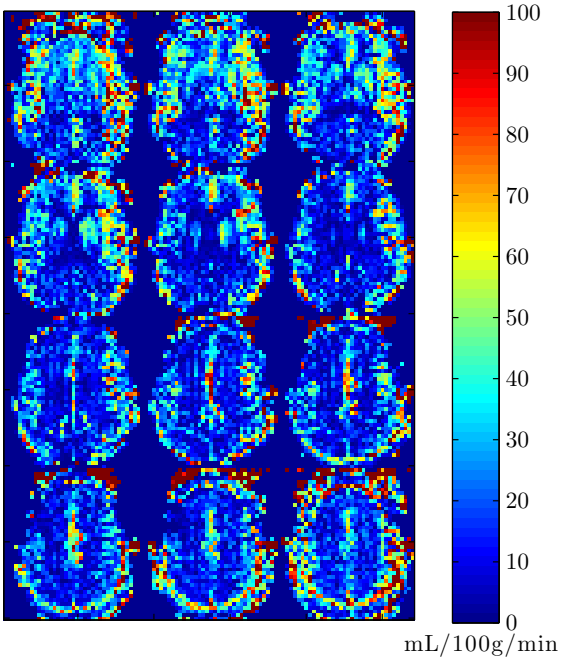

d. CASL CBF

Figure 8.6: Group B, subject 2 (male, age 23) PASL and CASL ΔM and CBF images

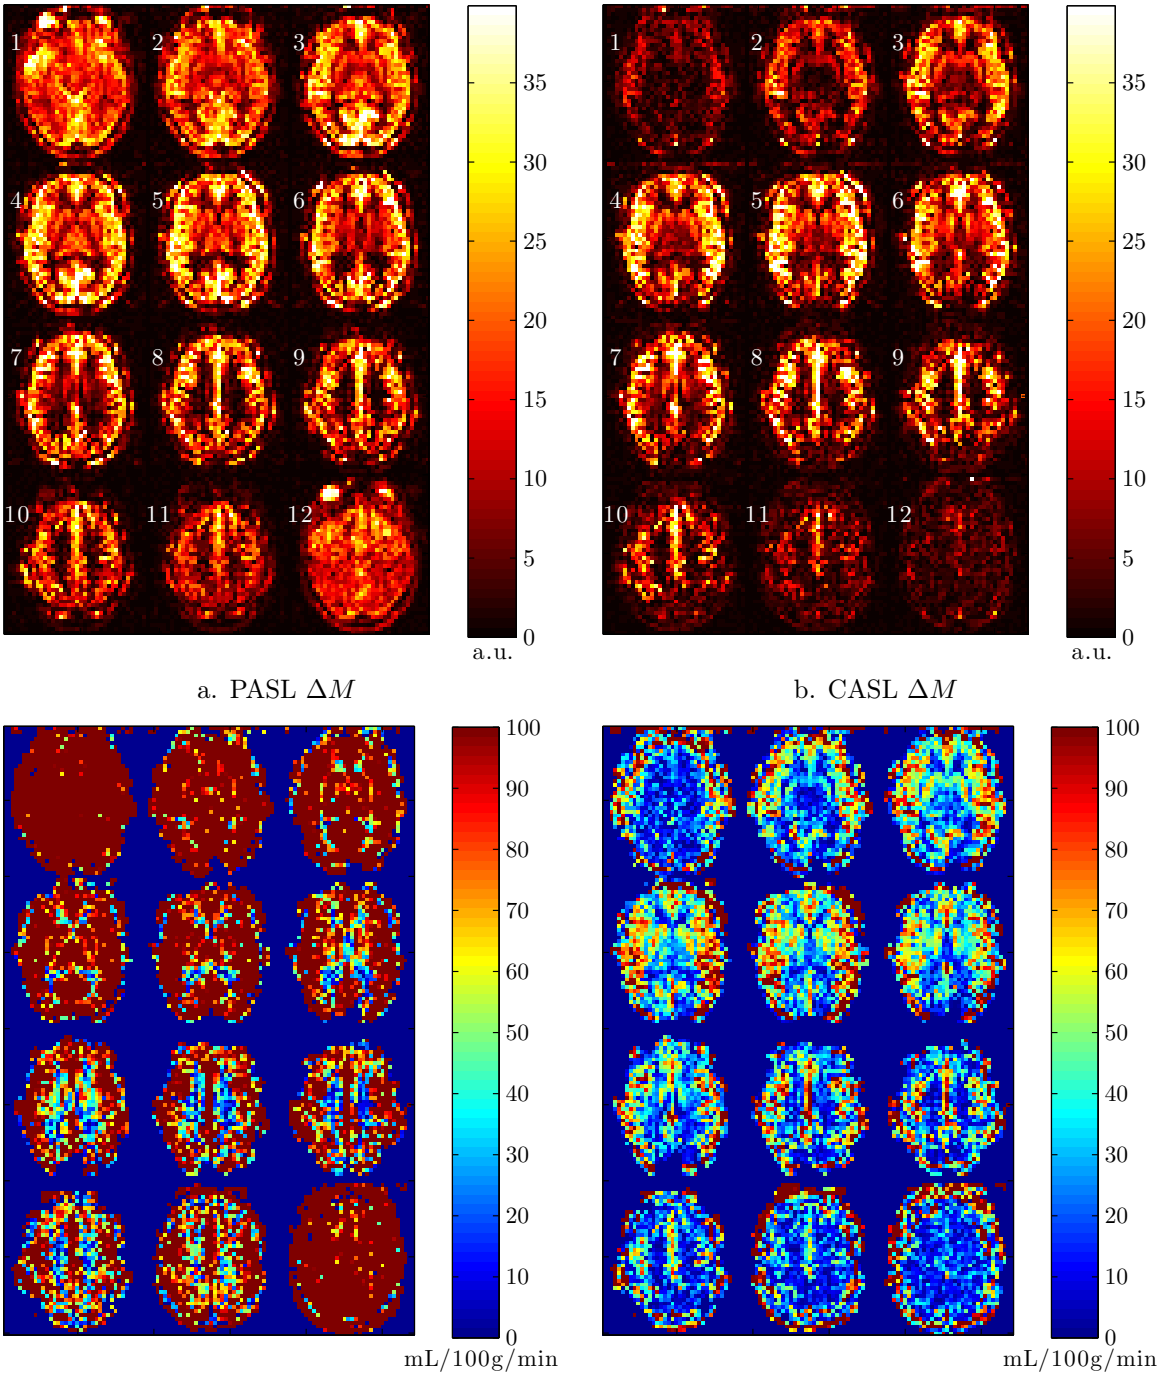

a. PASL ΔM

b. CASL ΔM

c. PASL CBF

d. CASL CBF

Figure 8.7: Group B, subject 3 (female, age 23) PASL and CASL ΔM images

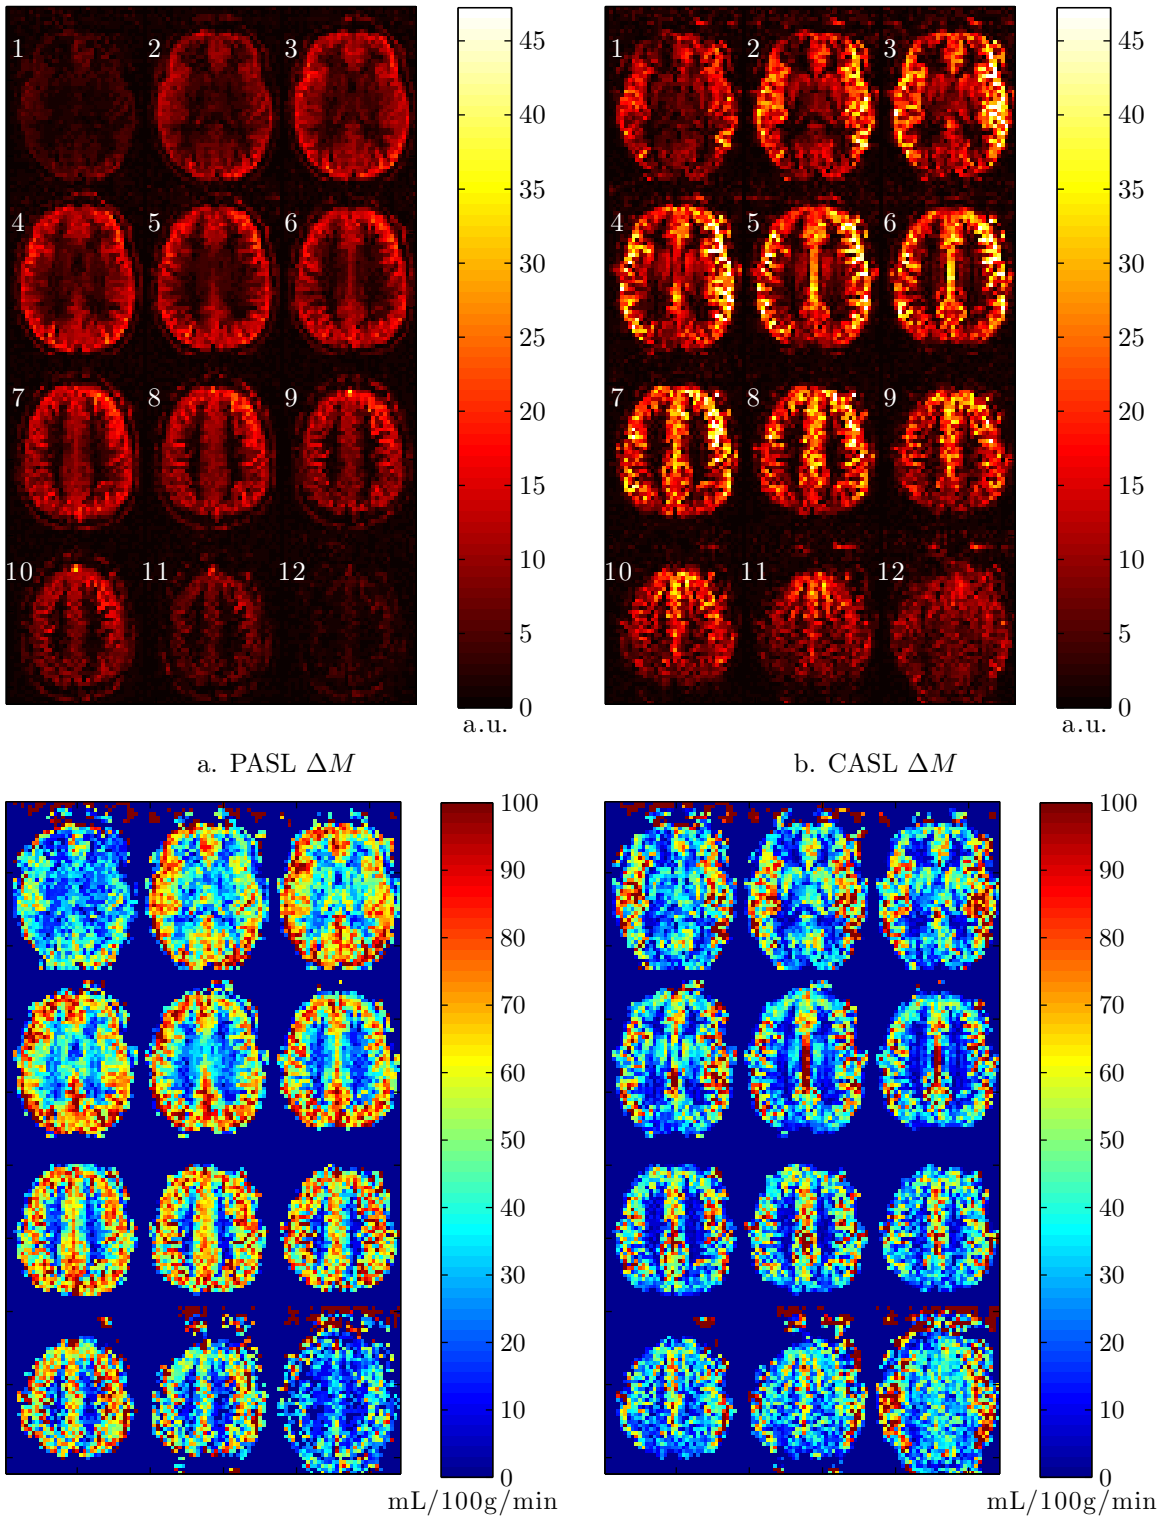

Figure 8.8: Group C, subject 1 (male, age 24) PASL and CASL ΔM and CBF images

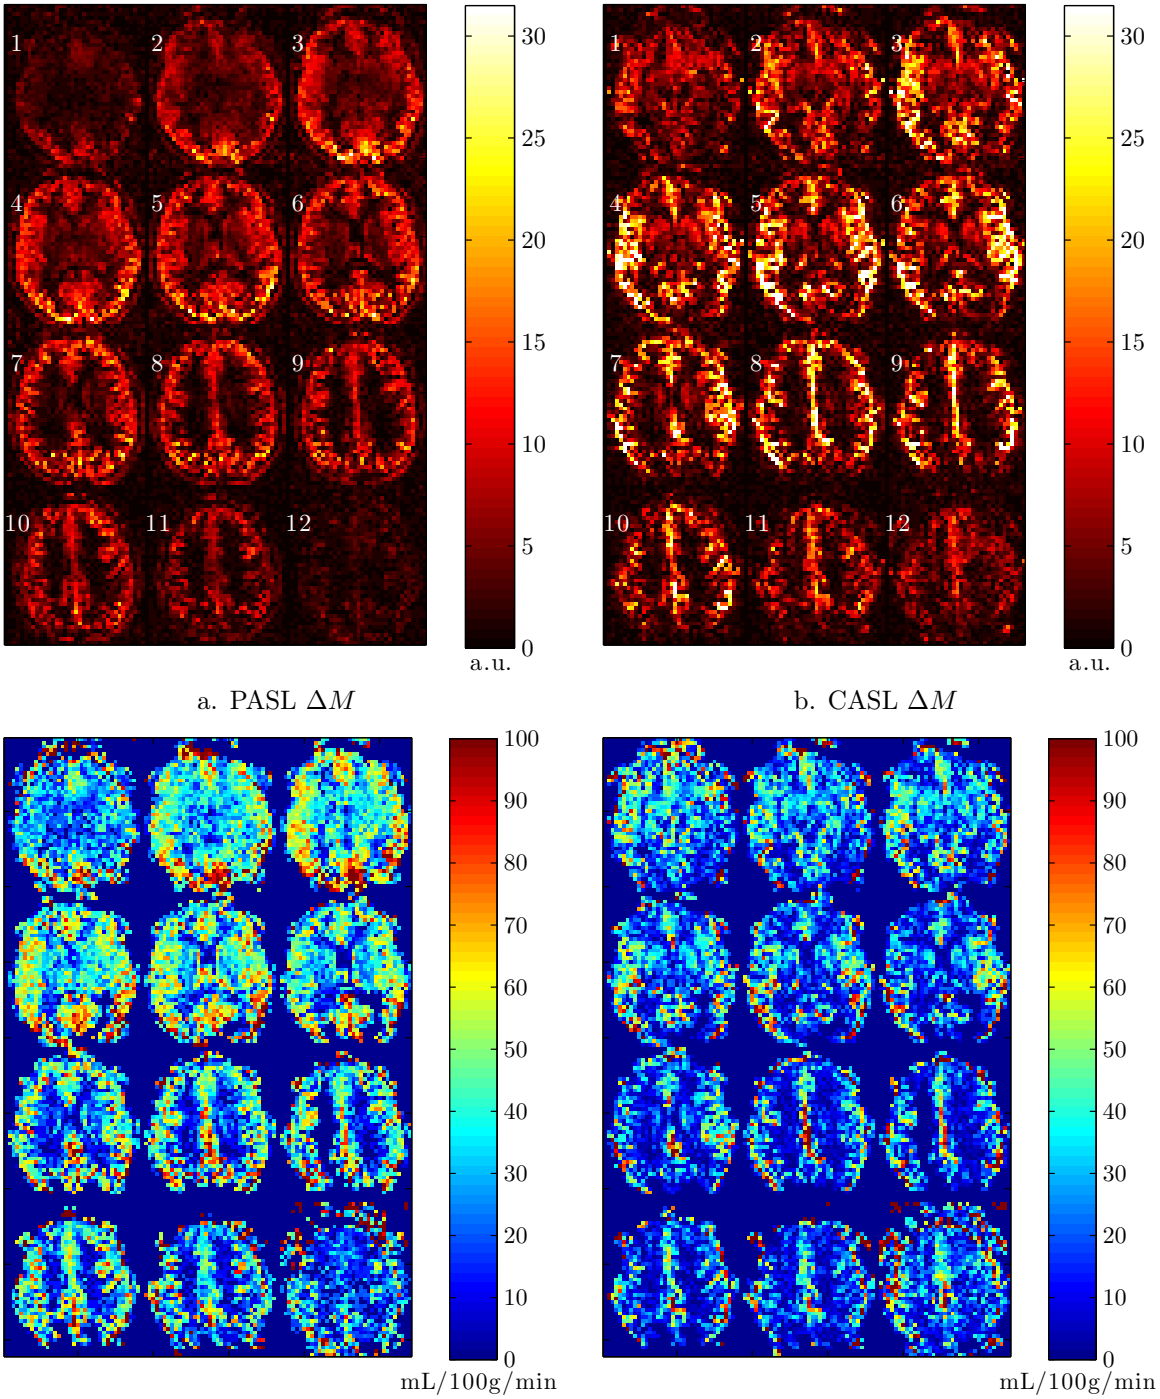

Figure 8.9: Group C, subject 2 (male, age 39) PASL and CASL ΔM and CBF images

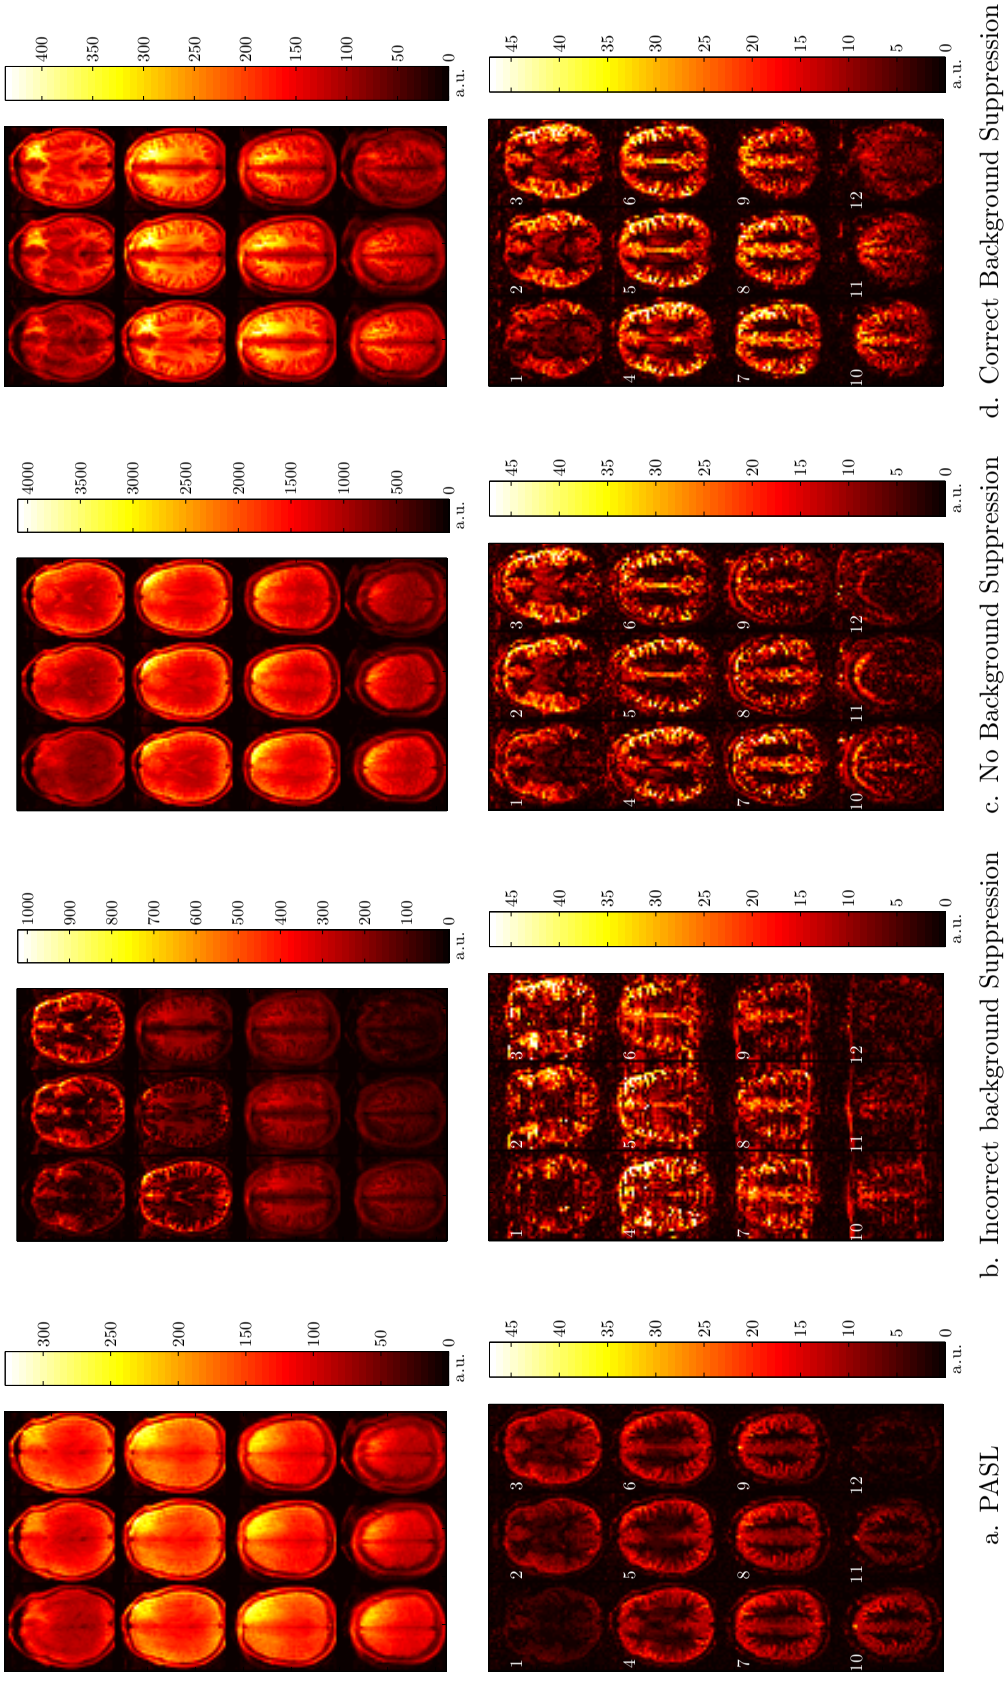

Figure 8.10: The effects of incorrect and correct background suppression, control (top) and ΔM (bottom) images

8.3.2 Regional Perfusion

8.3.2.1 Phantom

		Parallel Transmit				Single Coil			
		Left ROI		Right ROI		Left ROI		Right ROI	
Label Condition	Labelling Coil Amplitude	$\langle\xi\rangle$	σ_ξ	$\langle\xi\rangle$	σ_ξ	$\langle\xi\rangle$	σ_ξ	$\langle\xi\rangle$	σ_ξ
Label Left	80%	0.91	0.34	0.07	0.13	0.95	0.34	0.09	0.15
Label Left	100%	0.96	0.41	0.09	0.17	0.96	0.56	0.09	0.12
Label Right	80%	0.15	0.12	0.92	0.31	0.21	0.15	0.93	0.31
Label Right	100%	0.14	0.16	0.98	0.43	0.23	0.19	0.98	0.38

Table 8.4: Phantom relative inversion efficiency ROI statistics

		Parallel Transmit				Single Coil			
		Left ROI		Right ROI		Left ROI		Right ROI	
Label Condition	Labelling Coil Amplitude	μ	σ	μ	σ	μ	σ	μ	σ
Label Left	80%	0.91	0.10	0.03	0.03	0.94	0.09	0.05	0.02
Label Left	100%	0.94	0.09	-0.11	0.12	0.93	0.14	0.06	0.02
Label Right	80%	0.09	0.04	0.89	0.09	0.18	0.10	0.89	0.09
Label Right	100%	0.11	0.04	0.96	0.12	0.19	0.05	0.96	0.11

Table 8.5: Phantom relative inversion efficiency ROI gaussian fit parameters

Label Condition	Labelling Coil Amplitude	$\Delta\langle\xi\rangle_{left}$ (%)	$\Delta\mu_{left}$ (%)	$\Delta\langle\xi\rangle_{right}$ (%)	$\Delta\mu_{right}$ (%)
Label Left	80%	-4.33	-3.17	-26.51	-49.03
Label Left	100%	-0.50	0.65	-0.03	-275.38
Average		-2.42 ± 2.71	-1.26 ± 2.70	-13.27 ± 18.72	-162.21 ± 160.06
Label Right	80%	-31.86	-47.73	-0.98	-0.37
Label Right	100%	-39.17	-40.93	0.23	-0.19
Average		-35.52 ± 5.17	-44.33 ± 4.81	-0.37 ± 0.85	-0.28 ± 0.13

Table 8.6: Phantom parallel transmit inversion efficiency change

Figures 8.11 to 8.14 show perfusion weighted images and histograms for the left and right regions for the vessel selective labelling within the phantom. Each region within the histograms is well characterised by the fitted gaussian functions, and generally the dual gaussian function (the global fit) closely matches the individual left and right parameters. Figure 8.15 shows transverse slices at the labelling plane from the 3D FLASH sequence, obtained with the ASL Labelling Coils transmitting with amplitudes and phases to create nulls over the specified inflow pipes.

Table 8.4 summarises the average relative inversion efficiencies ($\langle\xi\rangle$) and standard deviations (σ_ξ) within the left and right regions for each labelling configuration. Table 8.5 presents the mean (μ) and standard deviation (σ) of the gaussian functions fitted to the histograms counts within the left and right regions, for each labelling configuration. The percent change in the mean value within each ROI, and the mean of the fitted gaussian functions are shown in table 8.6. Both the ROI mean value and the centre of the fitted gaussian functions provide metrics of left/right specificity which are in good agreement with each other.

The 100% amplitude label left configuration does not show any difference in the mean value of the right ROI for both single coil and parallel transmit scans, despite there being an obvious difference in the right side perfusion region on the ΔM images, and in the position of the centre of the right ROI gaussian. Both scans from the label right condition show a large reduction in contralateral labelling (at least 35%), whilst retaining almost identical labelling of the right perfusion territory.

Perfusion weighted images, histograms, the ROI statistics and the gaussian fit parameters all show that when using parallel transmission there is a reduction in contralateral labelling, whilst retaining good labelling efficiency within the labelling perfusion territory. There is a noticeable reduction in contralateral labelling on the perfusion weighted images, particularly for the left labelling condition. The 100% labelling amplitude measurements both show no loss in labelling efficiency within the labelling perfusion territory whatsoever. On average, within the labelling ROI the labelling efficiency of the parallel transmit vessel selective scans was approximately 93.5% of that achieved in the global CASL scans.

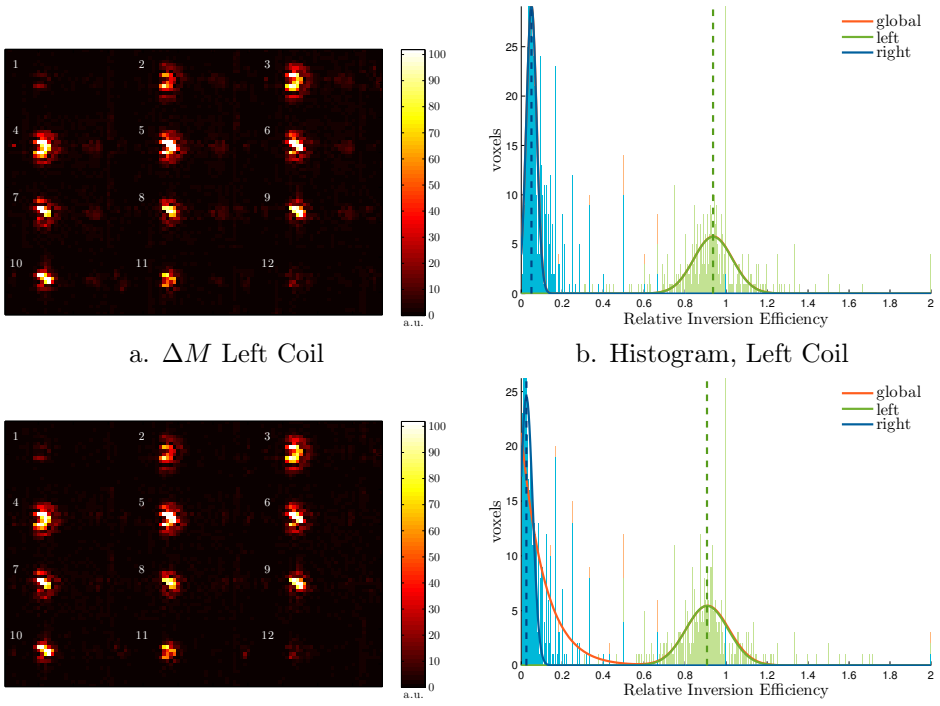

a. ΔM Left Coil

b. Histogram, Left Coil

c. ΔM Parallel Transmit

d. Histogram, Parallel Transmit

Figure 8.11: Phantom left labelling, 80% amplitude
$A_1 = 14.51\%$, $A_2 = 80.00\%$, $\Delta\phi = 180.00°$

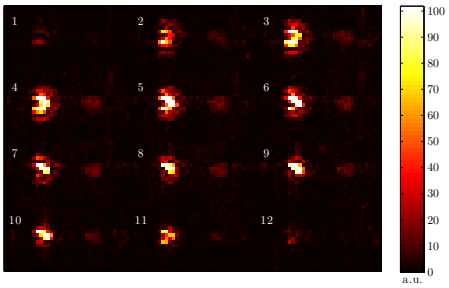

a. ΔM Left Coil

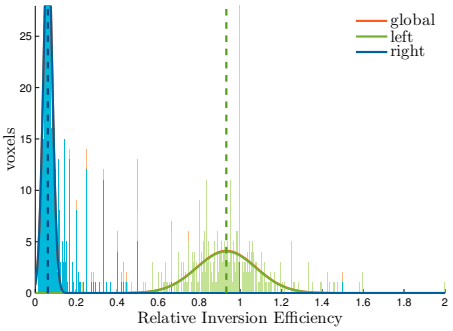

b. Histogram, Left Coil

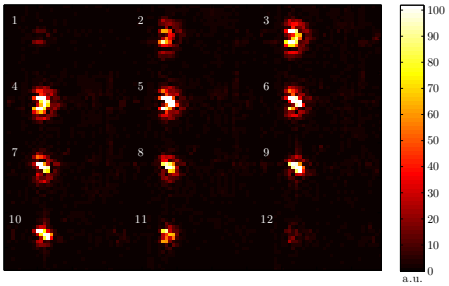

c. ΔM Parallel Transmit

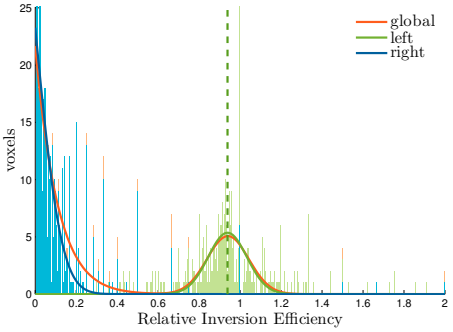

d. Histogram, Parallel Transmit

Figure 8.12: Phantom left labelling, 100% amplitude
$A_1 = 16.02\%$, $A_2 = 80.00\%$, $\Delta\phi = -173.25°$

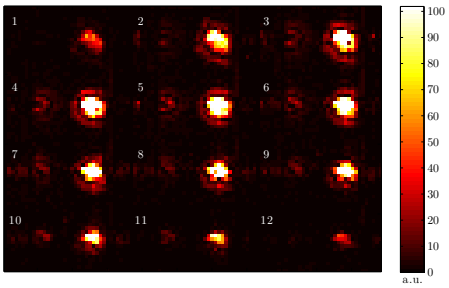

a. ΔM Right Coil

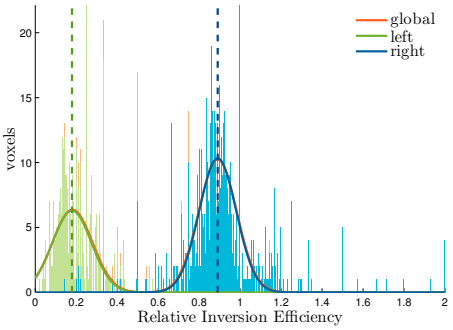

b. Histogram, Right Coil

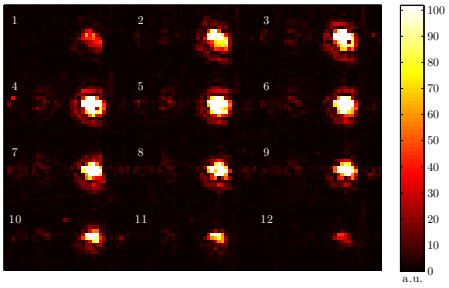

c. ΔM Parallel Transmit

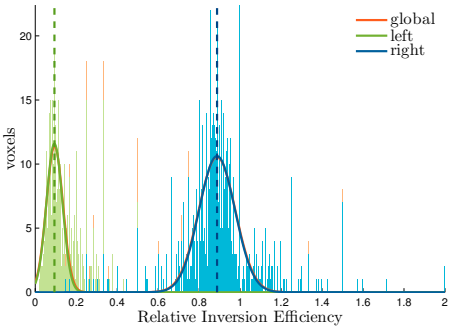

d. Histogram, Parallel Transmit

Figure 8.13: Phantom right labelling, 80% amplitude
$A_1 = 80.00\%$, $A_2 = 8.5\%$, $\Delta\phi = 86.625°$

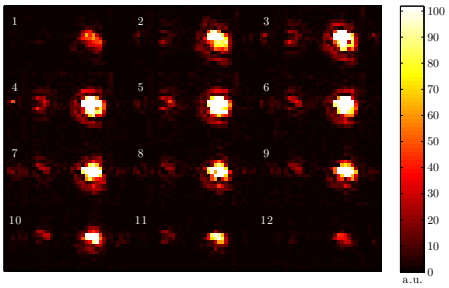

a. ΔM Right Coil

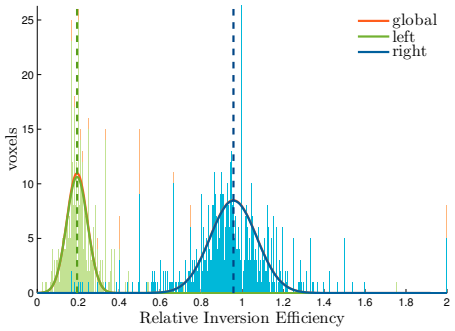

b. Histogram, Right Coil

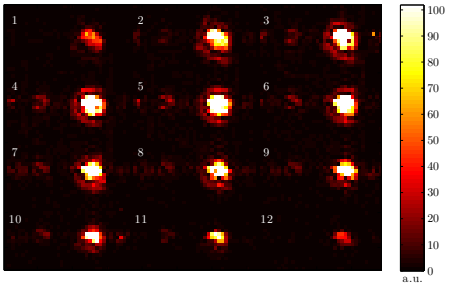

c. ΔM Parallel Transmit

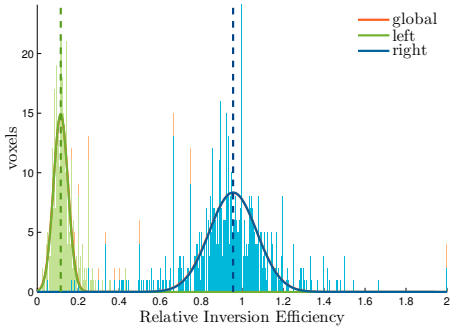

d. Histogram, Parallel Transmit

Figure 8.14: Phantom right labelling, 100% amplitude
$A_1 = 100\%$, $A_2 = 9.87\%$, $\Delta\phi = -81.00°$

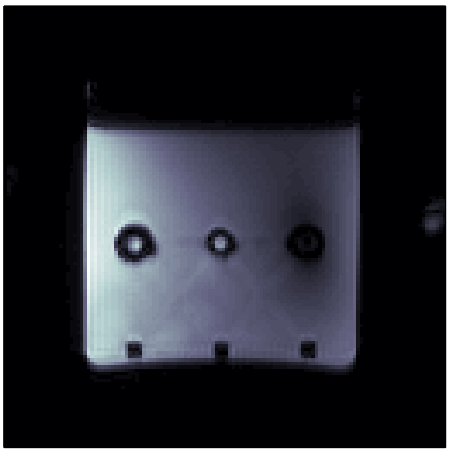

a. Label Left 80%

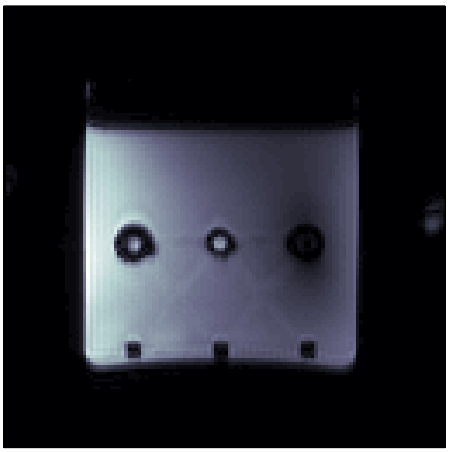

b. Label Left 100%

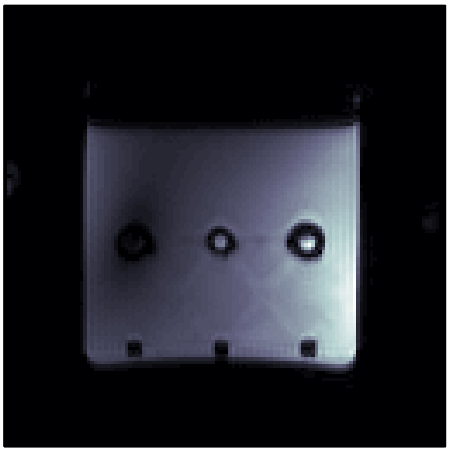

b. Label Right 80%

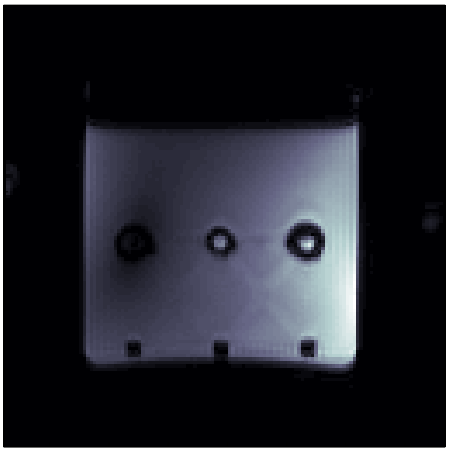

c. Label Right 100%

Figure 8.15: FLASH images at the labelling plane locating the B_1^+ null

8.3.2.2 Human Subjects

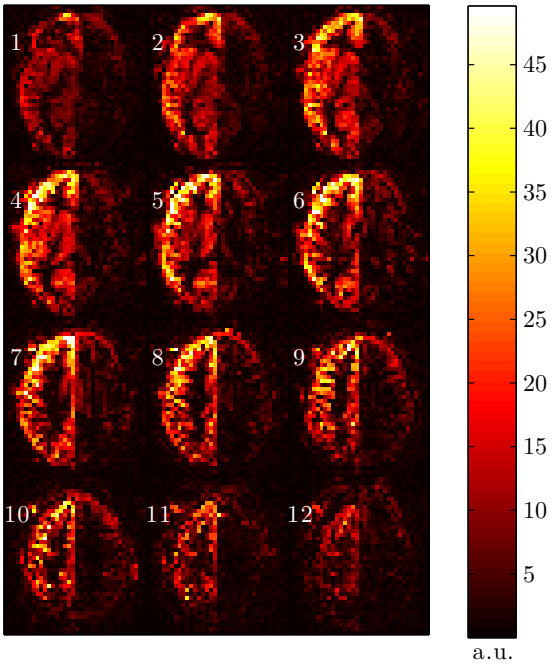

a. ΔM Left Coil

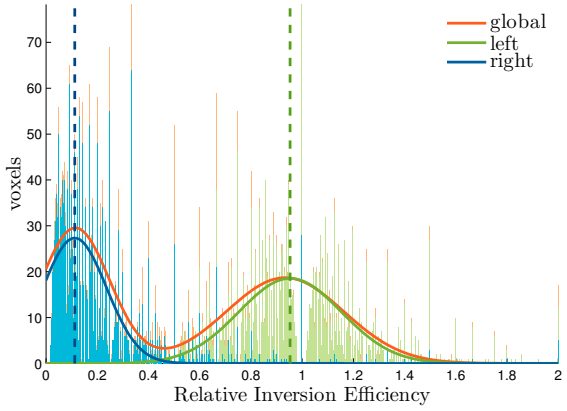

b. Histogram, Left Coil

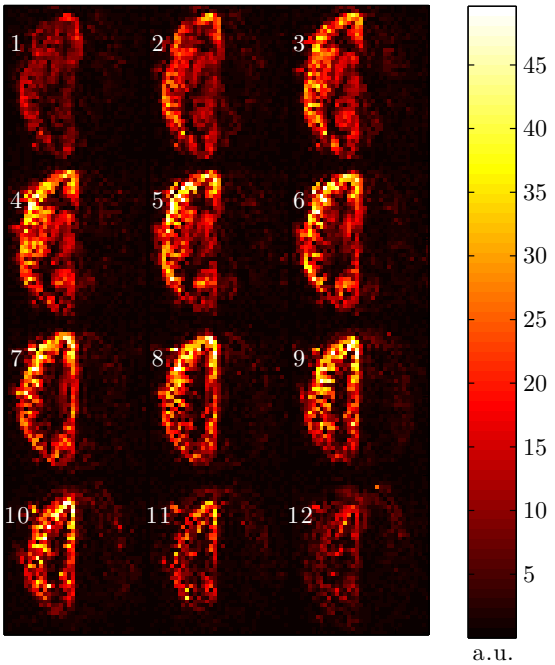

c. ΔM Parallel Transmit

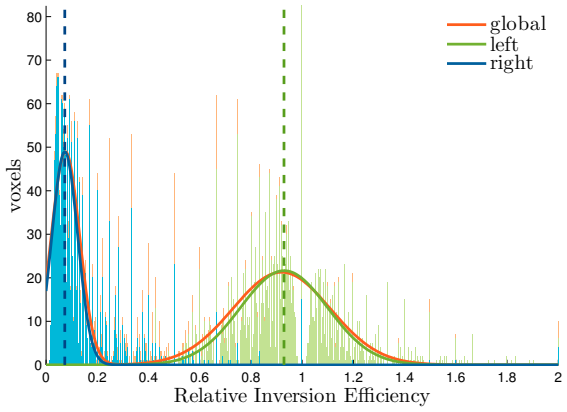

d. Histogram, Parallel Transmit

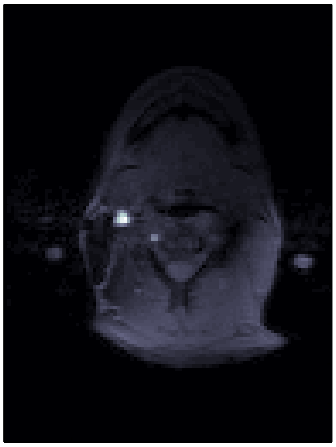

e. FLASH Image

Figure 8.16: Group A, Subject 2, Vessel Selective Label Left
$A_1 = 8.5\%$, $A_2 = 100\%$, $\Delta\phi = -51.75°$

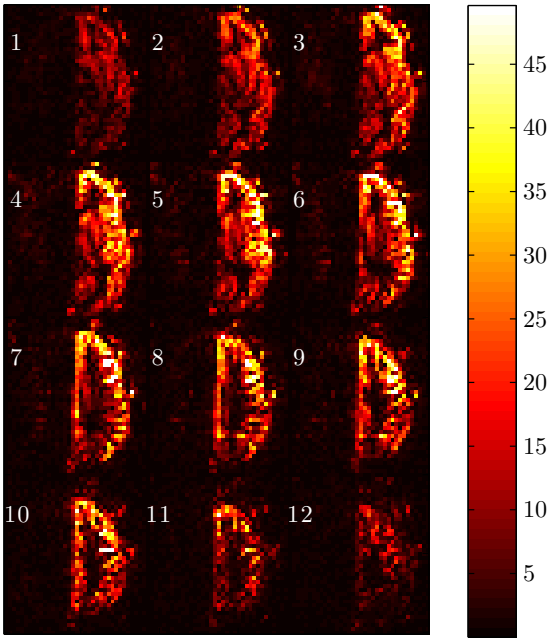

a. ΔM Right Coil

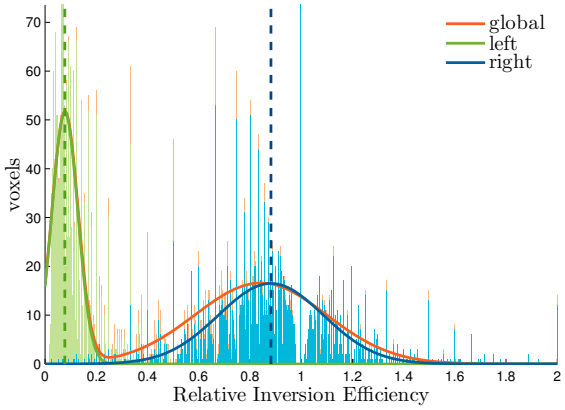

b. Histogram, Right Coil

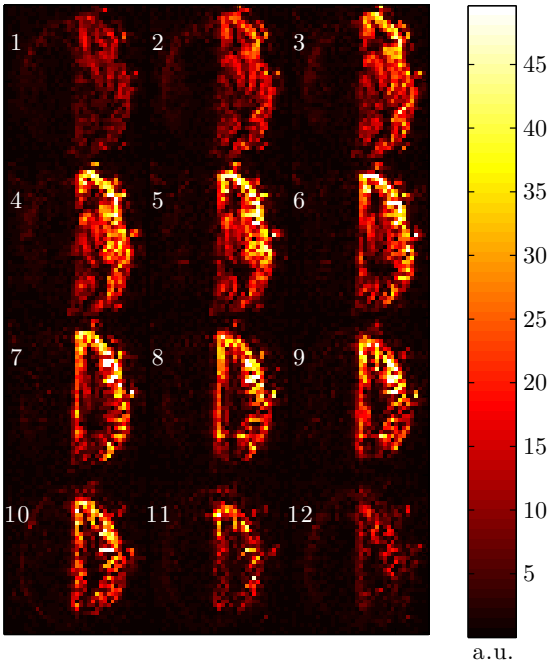

c. ΔM Parallel Transmit

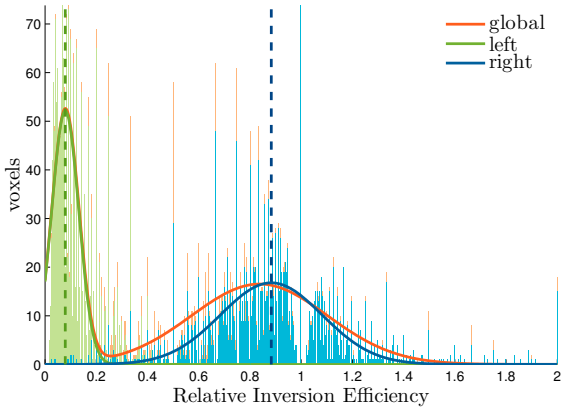

d. Histogram, Parallel Transmit

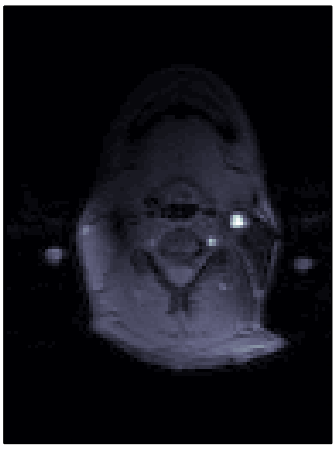

e. FLASH Image

Figure 8.17: Group A, Subject 2, Vessel Selective Label Right
$A_1 = 100\%$, $A_2 = 12.99\%$, $\Delta\phi = -140.40°$

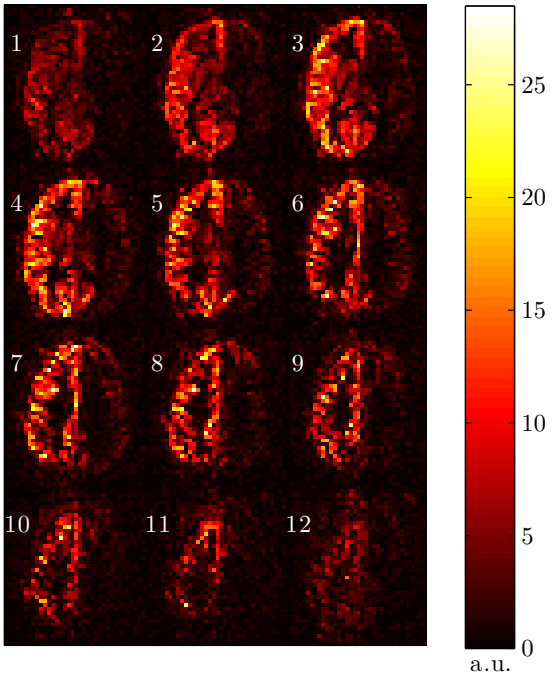

a. ΔM Left Coil

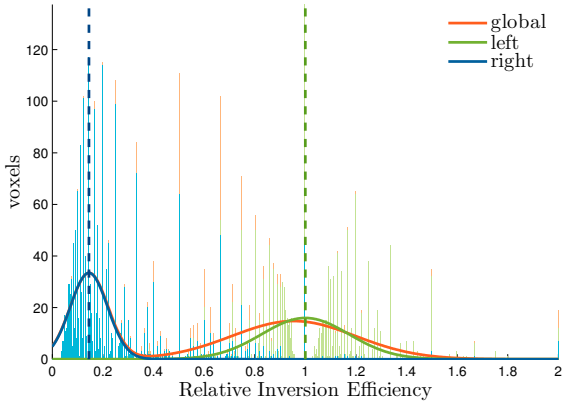

b. Histogram, Left Coil

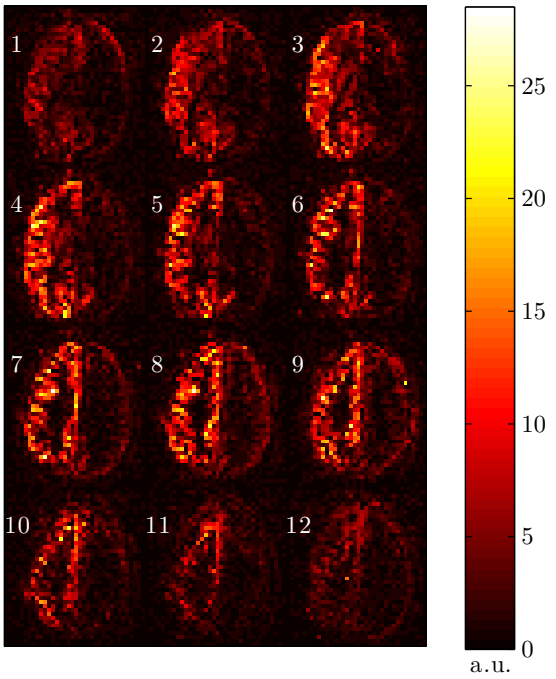

c. ΔM Parallel Transmit

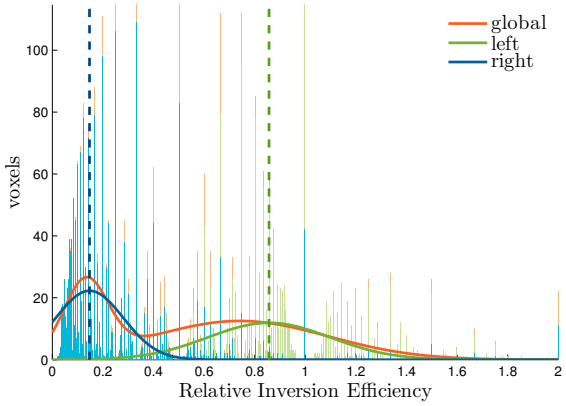

d. Histogram, Parallel Transmit

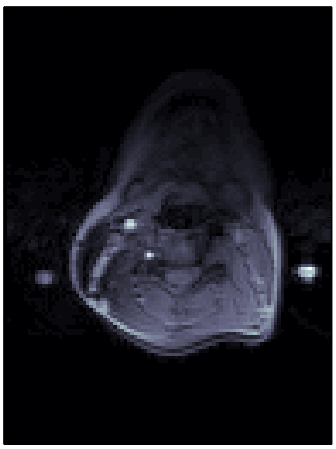

e. FLASH Image

Figure 8.18: Group A, Subject 4, Vessel Selective Label Left, Attempt 1
$A_1 = 12.02\%$, $A_2 = 100\%$, $\Delta\phi = 12.38°$

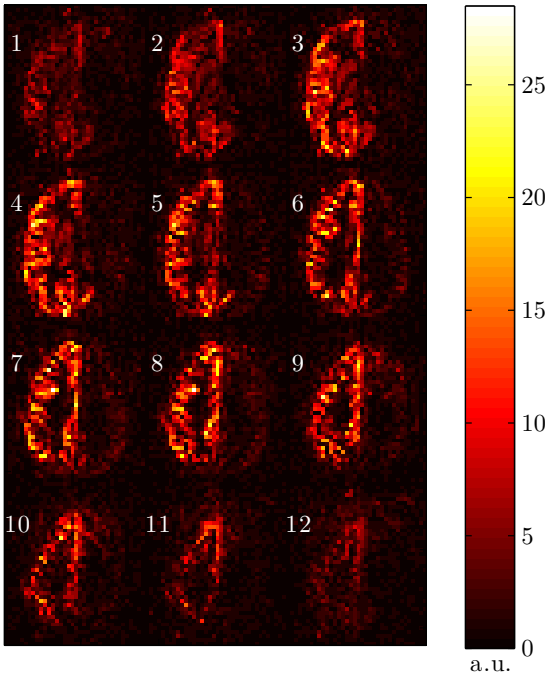

a. ΔM Left Coil

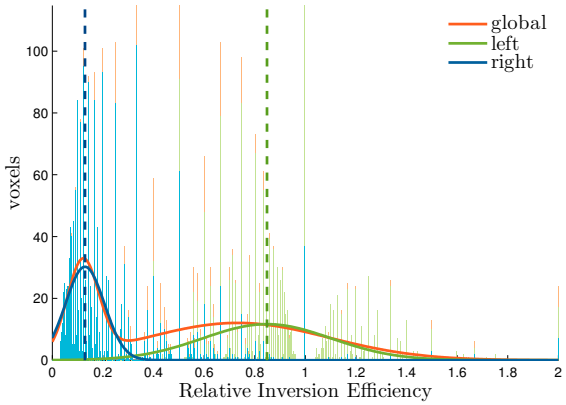

b. Histogram, Left Coil

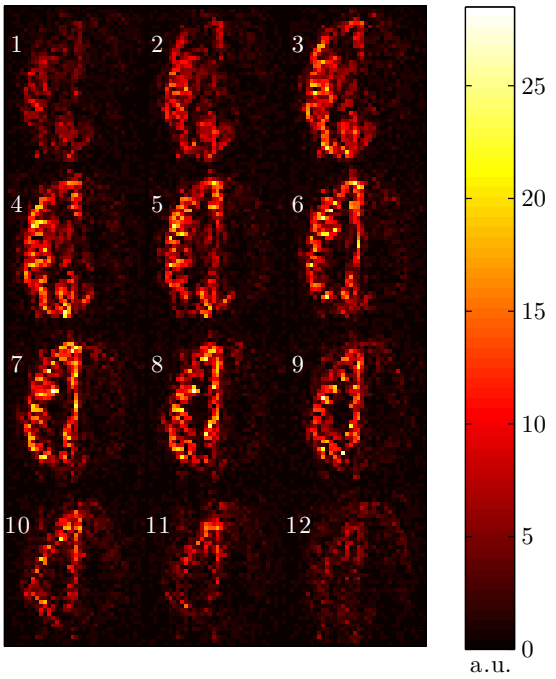

c. ΔM Parallel Transmit

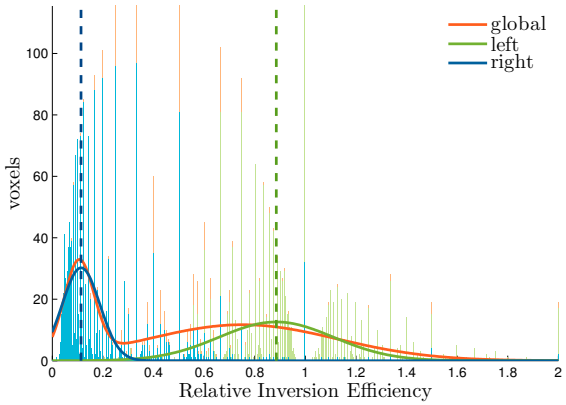

d. Histogram, Parallel Transmit

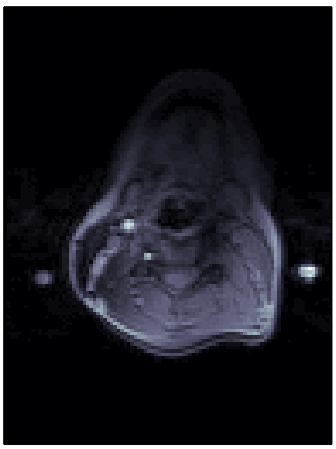

e. FLASH Image

Figure 8.19: Group A, Subject 4, Vessel Selective Label Left, Attempt 2
$A_1 = 15.73\%$, $A_2 = 100\%$, $\Delta\phi = 20.70°$

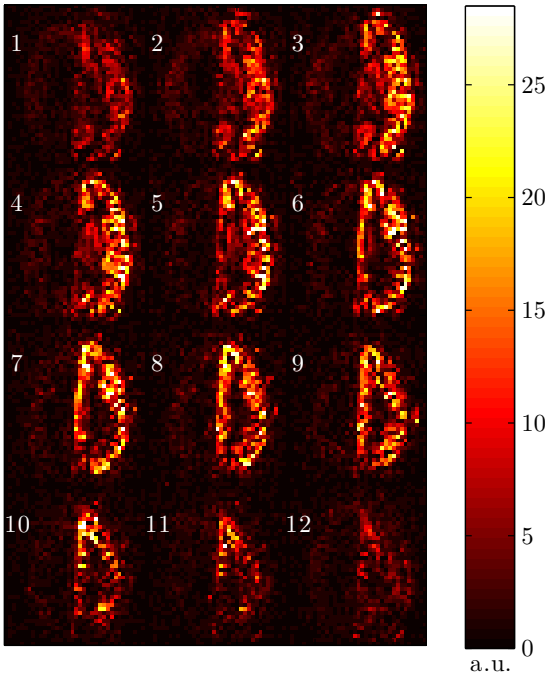

a. ΔM Right Coil

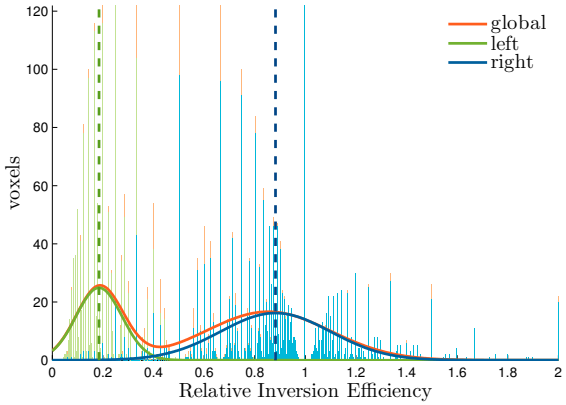

b. Histogram, Right Coil

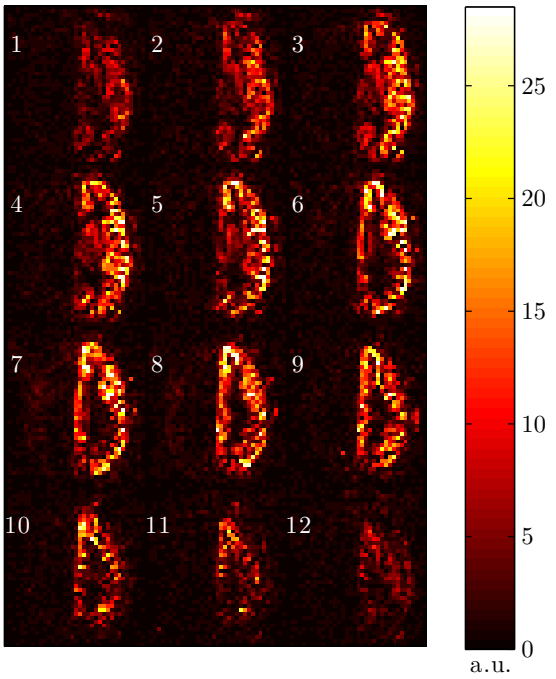

c. ΔM Parallel Transmit

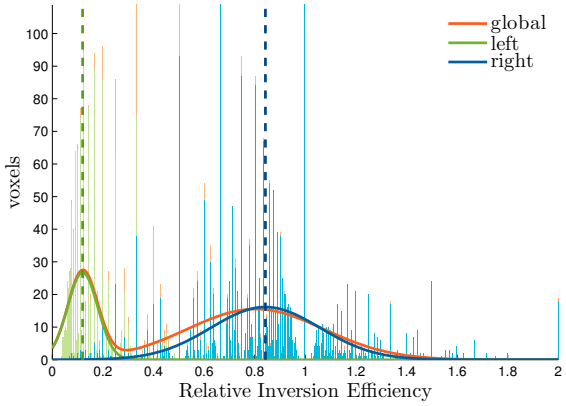

d. Histogram, Parallel Transmit

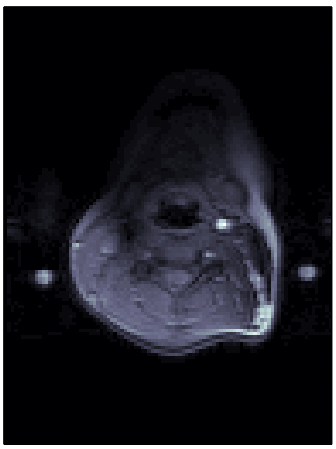

e. FLASH Image

Figure 8.20: Group A, Subject 4, Vessel Selective Label Right
$A_1 = 100\%$, $A_2 = 20.47\%$, $\Delta\phi = -67.95°$

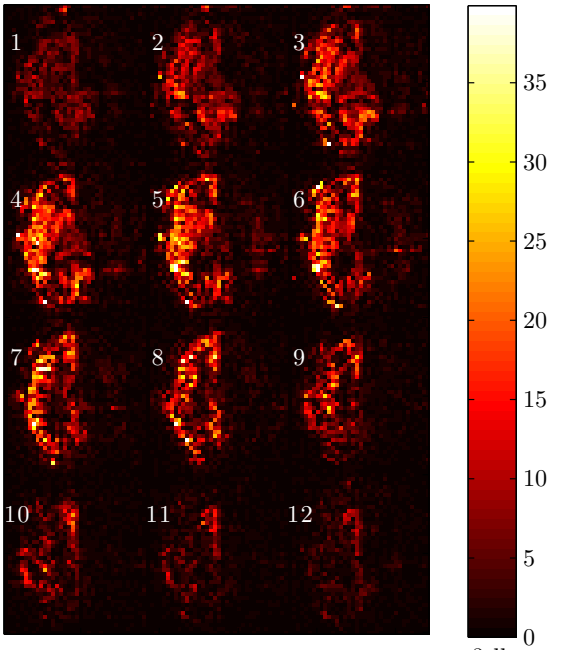

a. ΔM Left Coil

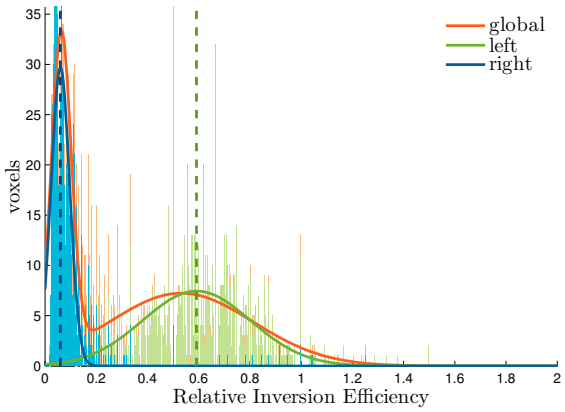

b. Histogram, Left Coil

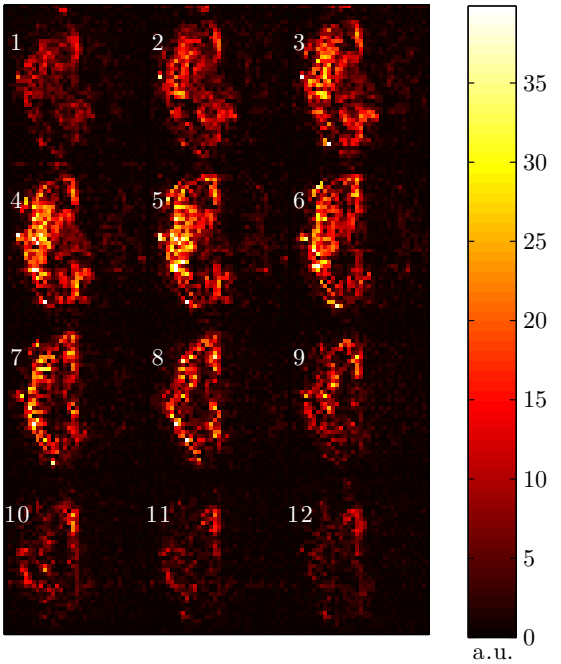

c. ΔM Parallel Transmit

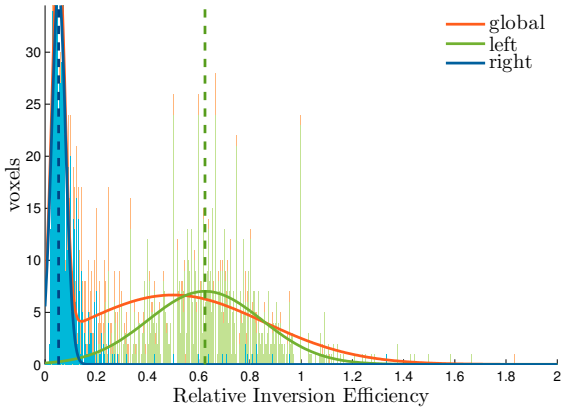

d. Histogram, Parallel Transmit

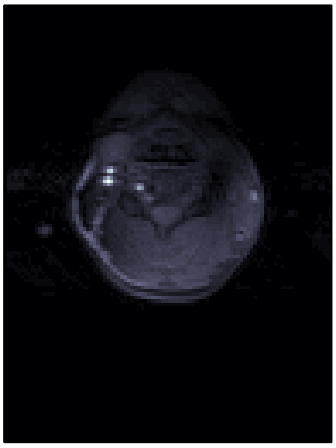

e. FLASH Image

Figure 8.21: Group B, Subject 3, Vessel Selective Label Left
$A_1 = 7.03\%$, $A_2 = 100\%$, $\Delta\phi = 93.15°$

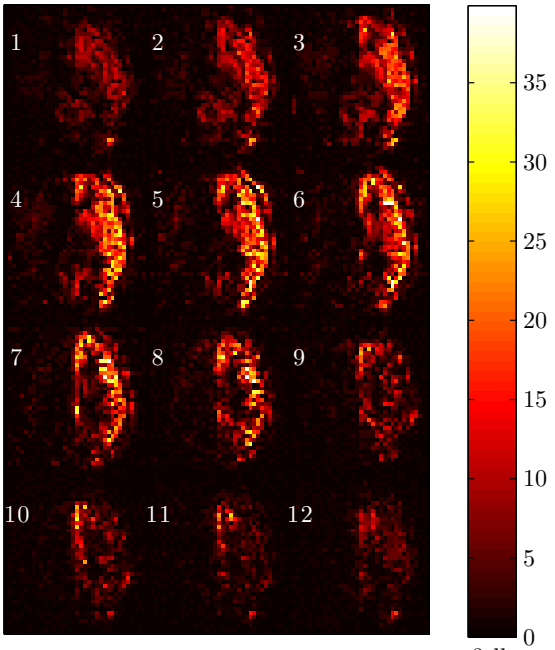

a. ΔM Right Coil

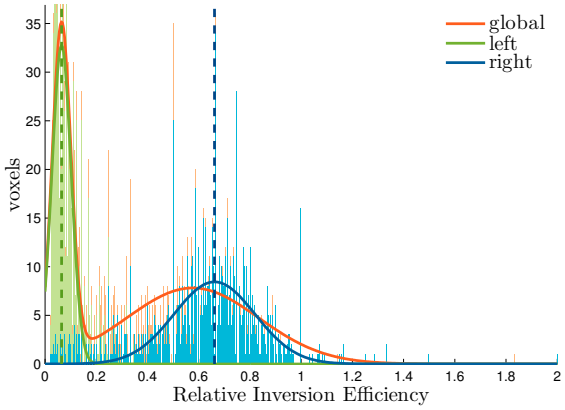

b. Histogram, Right Coil

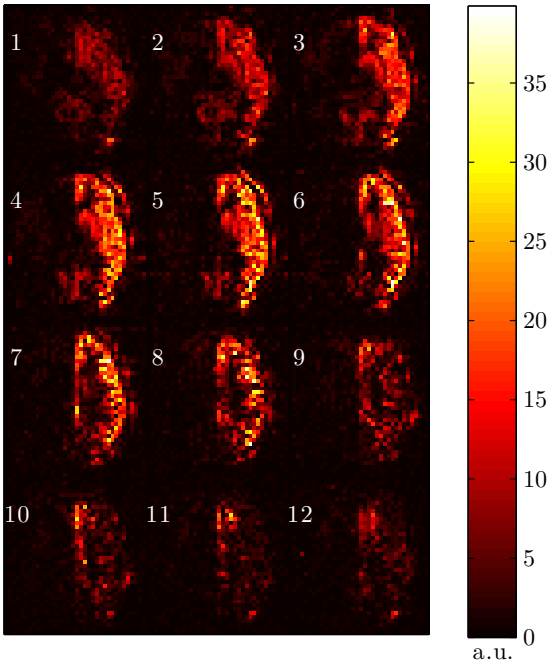

c. ΔM Parallel Transmit

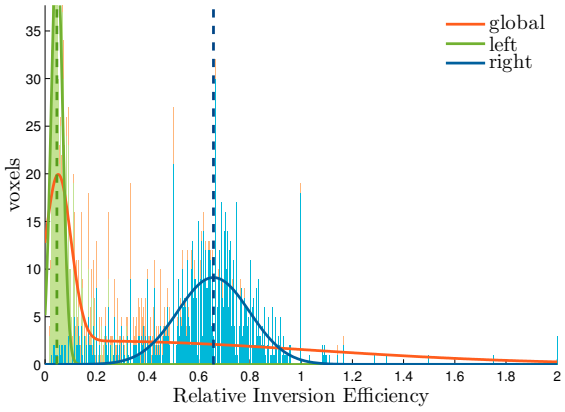

d. Histogram, Parallel Transmit

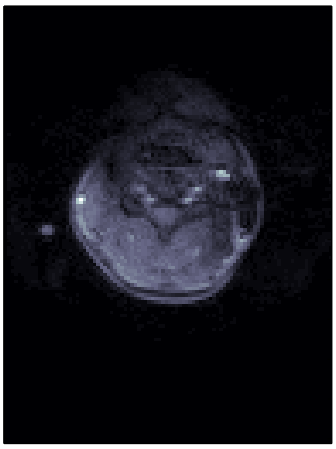

e. FLASH Image

Figure 8.22: Group B, Subject 3, Vessel Selective Label Right
$A_1 = 100\%$, $A_2 = 20.47\%$, $\Delta\phi = -67.95°$

Subject	$\Delta\langle\xi\rangle_{left}$ (%)	$\Delta\mu_{left}$ (%)	$\Delta\langle\xi\rangle_{right}$ (%)	$\Delta\mu_{right}$ (%)
A2	-0.33	-2.49	-65.95	-34.57
A4-1	-12.40	-14.33	10.80	1.82
A4-2	-3.79	4.33	-10.85	-11.92
B3	5.61	5.61	-16.03	-11.81

Table 8.7: Percentage fractional change in the mean inversion efficiency when using parallel transmission, label left

Subject	$\Delta \langle \xi \rangle_{left}$ (%)	$\Delta \mu_{left}$ (%)	$\Delta \langle \xi \rangle_{right}$ (%)	$\Delta \mu_{right}$ (%)
A2	26.29	2.01	0.00	0.15
A4	-37.97	-35.01	-3.85	-4.50
B3	-24.75	-28.44	-0.98	-0.59

Table 8.8: Percentage fractional change in the mean inversion efficiency when using parallel transmission, label right
have

Results are presented from the three subjects on who it was possible to perform both left and right labelling on within the allotted scan time. ΔM images, ROI histograms and a corresponding FLASH sequence slice at the position of the labelling plane are shown for subject A2 (figures 8.16 for left and 8.17 for right labelling), A4 (figures 8.18 and 8.19 for left and 8.20 for right labelling), and subject B3 (8.21 left, 8.22 right). The fractional change in the mean relative inversion efficiency for $\langle \xi \rangle$, and the mean value of the fitted gaussian in each ROI are shown in tables 8.7 and 8.8 for left and right labelling, respectively. Data from all other subjects, and tables with values for the ROI mean and standard deviations, and the fitted gaussian means and standard deviations are shown in appendix B.

For the left labelling condition, subject A2 shows the most effective use of parallel transmission CASL: an average reduction in contralateral labelling of at least 35% (ROI mean and gaussian mean), with only a minimal (approximately 2.5%) reduction in labelling on the left hemisphere of the brain. The ΔM images (figure 8.16.a and c) clearly show this difference, particularly around the edges of the left hemisphere on slices 8 and 9. A comparison of the histograms also show the standard deviation of the right ROI distribution to decrease significantly. The FLASH image in 8.16.e shows a hypointense region located at the right carotid artery, confirming the correct placement of the null. Conversely the right label condition for subject 2 shows no change in contralateral labelling, and in some slices (2 and 3) there appears to be an increase at the periphery of the brain.

Two different sets of amplitudes and phases were attempted in the left labelling condition for subject A4. In attempt 1, despite a null being present over the right carotid artery, as shown in figure 8.18, contralateral labelling was found to worsen when using parallel transmission. Attempt 2, in which parameters were found where A_1 was deliberately higher than previously, a better (from the FLASH images it appears to be larger in area) null was produced, resulting in a visible reduction in contralateral labelling in the ΔM images, a reduction in contralateral $\langle \xi \rangle$ of approximately 10%, and a small increase (approximately 4%) in $\langle \xi \rangle$ on the left side. Labelling on the right was effective in subject A4, with an average reduction in contralateral labelling of at least 35%, however at the expense of a small decrease (around 4%) in relative labelling efficiency in the right hemisphere.

In subject B3, parallel transmit vessel selective CASL resulted in a decrease in contralateral labelling in the right hemisphere by approximately 14%, and also shows a small increase (about 5%) in labelling in the right side. However ΔM images do not visibly show many signs of any

change, but a decrease in the standard deviation of the right ROI's gaussian function is observed. The right labelling condition of subject B3 has a noticeable decrease in contralateral labelling, evident in both the ΔM images and on the histograms. The change in the right hemisphere was less than 1%.

Subject averaged values for the left labelling condition were a relative inversion efficiency change in the right ROI of approximately $-9.5 \pm 28.4\%$ when using parallel transmission, with the left ROI changing by an average $-1.0 \pm 12.0\%$. In the right labelling condition this was $-16.3 \pm 39.3\%$ in the contralateral left ROI, and $2.4 \pm 3.2\%$ in the right ROI. On average, within the labelling ROI, the inversion efficiency of the parallel transmit vessel selective scans was approximately 80% of that obtained in the global CASL scans.

8.4 Discussion

8.4.1 Global CASL

In groups A and B the PASL and CASL sequences were not equivalent, due to incorrect background suppression in the CASL sequence used in group A, lack of background suppression in the CASL sequence in group B, and because Q2TIPS pulses were not used in the PASL sequence. Q2TIPS pulses shorten the labelling bolus, reducing the overall ΔM signal for the same TI, and also introduce an additional source of magnetisation transfer (MT) in both the control and tag images, further reducing the ΔM signal [47]. In group C, where both CASL background suppression and PASL Q2TIPS pulses are correctly employed, the ΔM images from the CASL sequence have almost twice the intensity of the PASL ΔM images, this confirms the theoretical prediction that the CASL signal is greater than the PASL signal [241]. Both the CASL and PASL CBF maps for group C show similar values.

In contrast with the PASL images a reduction in ΔM was observed in the occipital lobe, this being due to increased transit time; Petersen et al. [182] showed in a large multi-centre ASL reproducibility study that the occipital lobe has long transit times. ObservedΔM was generally higher in female subjects, again due to shorter transit time [182] and possibly from increased labelling efficiency: males tend to have larger necks, placing the carotid arteries further away from the labelling coils and potentially reducing the labelling B_1^+ field intensity and subsequent labelling efficiency.

Considering the different quantification models used for PASL and CASL data, and the assumptions made, in general the CASL data yields reasonably similar CBF values to that of the PASL sequence. As the focus of this study is vessel specificity, the CASL sequence and labelling parameters have not been adjusted for optimal CBF quantification, but as separate coil CASL in humans is an established technique [103, 157, 156, 213, 259] it should be possible to do so.

While this study has not investigated reproducibility, it is fair to say that between subjects there has been a wide variation in optimal labelling conditions. As previously discussed, what appears to be significant tissue transit time effects (characterised by inhomogeneous reduction in the ΔM signal, particularly in the upper slices and in the occipital lobe at the rear of the brain) were seen in four of the five male subjects, however these could also be in part caused by sub-optimal labelling, as a result of an insufficient labelling B_1^+field at the arteries. In section 7.5.2.1 in chapter 7, a method for optimising the B_1^+ field without specifying a region to null was proposed. Because of inter-coil coupling this method has not yet been tested, but as simulations and the vessel selective results in this chapter have shown that it is possible to spatially tailor the B_1^+ field to create label and null regions, it should also be possible to spatially tailor the B_1^+ field such that it optimally labels within all specified arteries, much in the same way that parallel transmission is used to homogenise B_1^+ fields in excitation. Using more local (coverage of the neck only) transmit coils will increase the ability to tailor the field, yet still spatially confine the B_1^+ field to the neck, resulting in zero magnetisation transfer effects and a large ΔM signal.

8.4.2 Vessel Selective CASL

In the controlled environment provided by using a phantom, a significant improvement in left right specificity was achieved using parallel transmit vessel selective CASL. The phantom provided consistent flow, zero motion from respiration during the labelling pulse, and zero motion between scans. Therefore the null was always positioned at the same location, and there were no changes to labelling parameters, apart from the labelling B_1^+field between the acquisition of the global CASL scans (which were required to calculate ξ), the parallel transmit scan and the single coil scan. In other words, the phantom minimised the sources of error that could contribute to a change in ξ. Another advantage of using a phantom was that there were no restrictions on the total scan session duration, therefore it was possible to ensure a well optimised null was created over the contralateral inflow pipe. All of these factors contributed to an improvement in labelling specificity of at least 35% in the scans performed.

Conversely, motion and flow fluctuations are present in humans, and these effects contribute to the results presented in section 8.3.2.2. Results showed that while it was possible to produce a null over the contralateral artery, as shown in the FLASH images this did not always result in an improvement in labelling specificity. For those scans in which there was an improvement this ranged from approximately 10 to 35%, approaching the improvement observed in the phantom measurements. However, even the phantom results are a long way from the theoretical 10× improvement in labelling specificity obtained in the simulations in chapter 3. While it is not fair to directly compare the theoretical and experimental results, as the simulations did not use realistic B_1^+ maps at 3T, nor take into account motion and physiological fluctuations, this difference does

suggest that there is further scope for improving specificity by using a more robust method for selecting the driving coefficients (A and $\Delta\phi$), i.e. the mathematical framework presented in section 7.5.2.2 in chapter 7. As discussed in section 7.1, inter-coil coupling prevents the use of this for the moment, but all else is in place for its use if a suitable decoupling system is implemented.

Sensitivity to motion is problematic for parallel transmit vessel selective CASL. As the technique relies upon the exact placement of a null in the B_1^+ field over an artery, motion will cause the null and artery to no-longer coincide, potentially increasing the B_1^+ field within the artery and subsequent labelling to greater than what would be experienced with only one coil. In addition to changing the position of the arteries with respect to the null location, subject motion can also change the loading conditions of the coils, which also determines the degree of coupling that exists. Variations in the amount of coupling could in turn change the actual B_1^+ field produced by each coil, potentially by enough to also move the null. A further complication in this work was that the coils were not fixed in position, as would normally be the case for a coil array. All of these factors could attribute in part to the less than expected improvement in specificity using parallel transmit over a single coil, and also explain why some scans did not show any significant improvement, despite confirmation of the null placement in the FLASH images. Due to their long duration, potential for motion in these scans was high. With sufficient decoupling such that optimal driving coefficients can be determined with the scan planning software, the scan time could be significantly reduced.

Compared to the global CASL scans, a decrease on average 5% and 20% was seen in the vessel selective scans (parallel transmit and single coil) for the phantom and the human scans, respectively. This suggests that there was a small, but significant contribution to the labelling field at the labelling artery/inflow pipe from the contralateral coil. Ideally, labelling parameters would be optimised such that labelling efficiency is insensitive to small variations in B_1^+. As there was only a small reduction in labelling efficiency in the vessel selective, 100% labelling coil amplitude phantom scans, this suggests that these scans were conducted in this regime. Optimal coil placement is necessary to ensure that a sufficient B_1^+ is possible at the carotid artery locations, and previous studies have located the labelling coils such that they are positioned directly on the carotid arteries [156, 213, 259]. However, in this study the coils were positioned on opposite sides of the neck, to keep conditions similar to those in which the SAR mapping was performed in chapter 6, minimise inter-coil coupling by maximising the distance between the two coils, and also take advantage of parallel transmission to spatially tailor the field rather than rely on coil placement. A consequence of the last objective was that even when running at 100% power, labelling was not always optimal, particularly in the male subjects. While this did have implications on the results in this study, in general this problem could be overcome by using a more suitable, optimised transmit coil array that ensures a sufficient labelling field can be created without driving the coils at the maximum permitted amplitude.

Chapter 9

Discussion, Conclusions and Further Work

This thesis has applied parallel transmit methods to arterial spin labelling. In this final chapter the entire technique will be discussed and important conclusions highlighted, along with directions for further work.

In chapter 3, the mathematical framework for the selection of the amplitudes and phases was presented and used in an anatomically representative Bloch equation simulation of the continuous labelling procedure. In comparison with a single surface coil positioned over the right carotid artery, parallel transmission, with the aim of creating a null over the left carotid and vertebral artery resulted in a reduction in contralateral labelling by one and two orders of magnitude when using two and four transmitter coils, respectively.

Chapter 4 described the design, implementation and characterisation of a pair of ASL Labelling Coils which could be used for parallel transmit vessel selective CASL. The coils consisted of two surface coils (outer diameter 62.5mm), tuned and matched to 123.2MHz for standard human neck loading conditions. A complementary transmitter system, designed to provide two independent low power transmit channels for use with a 3T human MR system was presented in chapter 5. Built around a recycled spectrometer, the system allows full control over amplitude and phase of each channel, and the spectrometer provides a full pulse programming environment enabling considerably more control than previous simple gated RF implementations [157, 156, 204, 213, 259] .

RF safety concerns were diligently considered and assessed in chapter 6. Intense localised SAR from the use of the coils was identified as a potential safety risk, and was measured using proton resonance frequency magnetic resonance temperature imaging (PRF-MRTI) to obtain SAR maps. A high resolution gradient echo EPI sequence acquired proton resonance frequency (PRF) temperature maps of an agar phantom being irradiated by the ASL Labelling Coils. An appropriate

function (equation 6.8) which empirically describes the SAR heating process was fitted to each voxel's time series, boosting the SNR of the initial temperature change with which the SAR is defined. Results showed that this method can provide a temperature accuracy of approximately $0.1°C$ in comparison to measurements made with a fibre optic thermometer. However this is only for the simple case of a homogeneous, non-convective gel phantom, in which the fitting function closely matches the temperature changes. Resultant SAR maps showed a maximum local SAR of $8.78 \pm 0.38 W/kg$, which was obtained without any mass-averaging, on an agar phantom with a higher conductivity, therefore higher power deposition than within muscle tissue. Hence, it was determined that the coils could be used on human subjects without any duty cycle constraints (providing SAR limits are not exceeded when the SAR from the sequence pulses are also taken into account), and in a typical CASL sequence this gives a large conservative safety margin of at least 2, without accounting for the fact that the peak SAR values are not mass-averaged, effectively increasing the safety margin.

Chapter 7 presents the components required to implement parallel transmit vessel selective CASL. A B_1^+ mapping sequence, and image processing techniques were described that can be used to acquire maps of each coil's transmit RF field with optimum SNR. An advanced ASL sequence which takes advantage of a SNR efficient single shot 3D-GRASE acquisition, and background suppression for static tissue nulling was implemented, capable of both PASL and CASL labelling for direct comparison of the two labelling schemes with the same acquisition parameters. To aid development of parallel transmit vessel selective arterial spin labelling a perfusion phantom with distinct vascular territories was developed. The phantom was manufactured by 3D printing, a rapid prototyping technique; and provides a large ASL signal with both continuous and pulsed labelling sequences. Scan planning software was also developed, which automates the processing of B_1^+ maps and anatomical scans, making it easy for the scan operator to select vessels for labelling and set the transmitter channels to the correct driving coefficients.

Chapter 8 presents initial in-vivo results using parallel transmit vessel selective CASL, with the aim of improving left–right specificity in comparison to using a single labelling coil. Because of inter-coil coupling it was not possible to fully utilise the automatic calculation of the driving coefficients for each coil. However by iteratively adjusting the combinations of amplitude and phase shift on the contralateral coil (i.e. for labelling the left carotid artery, the right side coil is the contralateral coil), and confirming the location of the B_1^+ null in a rapidly acquired FLASH sequence, it was possible to spatially locate a null over the contralateral carotid artery on all subjects. Human and phantom measurements showed an increase in specificity by up to 35% in comparison to using a single coil for labelling, with a reduction in labelling efficiency for the parallel transmit configurations of less than 5%. In comparison to using both coils for global perfusion measurements (full power on both coils), the labelling efficiency within the labelling hemisphere of the brain was around 80% for both parallel transmit and single coil measurements,

this was because the CASL labelling was not fully optimised to be insensitive to B_1 changes. With more optimised labelling parameters this reduction could be avoided.

9.1 Discussion of the work in this thesis

9.1.1 Separate Coil CASL Technique

To the author's knowledge, this is the first separate labelling coil system that permits labelling of both carotid arteries, or each individually, without repositioning of the coils. Furthermore the degree of vessel specificity attained from each coil transmitting independently was very good, with only a small amount of contralateral labelling. Despite non-optimal coil positioning (coils were placed on the sides of the neck and not directly over each carotid artery), most subjects showed a high degree of inversion efficiency and accompanying large ΔM signal. There is however still scope to optimise the labelling parameters to improve robustness to variations in B_1^+ and blood velocity.

9.1.1.1 Parallel Transmit CASL

Parallel transmit vessel selective CASL has been shown to provide an improvement in specificity in left/right labelling. Because of the good specificity and consequent low contralateral labelling obtained from using the labelling coils individually, this improvement was quite small (10–35%). It is anticipated that by using a pair of coils which are optimised for good global CASL labelling (more in plane coverage across the neck so more contralateral labelling when used individually) utilising parallel transmission would result in a larger improvement in specificity. Furthermore it is proposed that parallel transmission could also be used to ensure optimal global CASL labelling, ensuring sufficient labelling B_1^+ to account for flow variations, whilst controlling the degree of power deposition.

In its current state of development, the parallel transmit CASL implementation is cumbersome and time consuming to set up, therefore is not yet suitable for general use as a vessel selective labelling technique. The trial and error, iterative placement of the null over the contralateral carotid artery is a slow process, and also non-optimal, so an improvement over using a single labelling coil is not guaranteed. Ultimately this falls down to the lack of decoupling between the two coils. However, as decoupling methods are well established, it would be relatively straightforward to implement one of these, as detailed in section 9.3.1. With appropriate decoupling, correct linear combination B_1^+ maps can be acquired, from which accurate B_1^+ maps for each individual coil maps can be calculated. Using these, the parallel transmit scan planning software can determine the optimal combinations of amplitude and phase first time, and the coils can be driven exactly with these specified.

The technique is sensitive to any motion that occurs between the acquisition of the pre-scans

and the actual ASL measurement. The non-fixed nature of the labelling coils also exacerbates this effect. Generally the trial and error FLASH images were acquired immediately before the parallel transmit ASL sequence, however even in that short space of time some motion could occur.

9.1.1.2 Compared to other VS-ASL techniques

Parallel transmit CASL has several advantages over existing techniques. Being a CASL method it can produce a large bolus of labelled blood, resulting in a large ΔM signal and subsequent high SNR. Because it uses localised labelling coils, no MTC is imparted into the imaging volume therefore a control acquisition does not need RF power applied. In addition there is no MT induced shortening of the T_1 relaxation time of the blood or tissue which is present when CASL is performed with a volume transmitter coil, therefore the labelled bolus of blood has the highest possible inversion efficiency. Because the required labelling power is low, true continuous wave is easily accomplished, and the RF power deposition and associated SAR is well within regulatory limits.

Compared to performing CASL with a single small labelling coil, parallel transmit CASL has the advantage that it is less sensitive to placement of the coil, and it is possible to label either the right or left common carotid arteries without repositioning the coil, or globally label both arteries. Furthermore contralateral labelling effects can be minimised by nulling the labelling field over the opposite carotid artery.

In its present implementation however, parallel transmit CASL shares the same disadvantages as other separate labelling coil CASL methods: labelling can only be performed in the common carotid artery, resulting in increased tissue transit times (and larger T_1 decay of the labelled bolus); and additional hardware is required. However, recent developments in RF technology permit an RF channel to be implemented wholly in the digital domain [210, 247], significantly reducing component count and cost, which when combined with inexpensive low power RF amplifiers means that a turnkey off-the-shelf solution could be realised. Of course the design and manufacture of medical devices is not a trivial matter, as high quality standards (ISO 13485 [4]) must be granted and adhered to, but if a significant improvement in SNR (and associated scan time reduction) can be realised, the approach may offer a significant enough advantage to justify this effort, especially when applied to patients who cannot remain still for long periods.

As using parallel transmission to define labelling regions and nulls is purely an RF method, it is compatible with other labelling schemes such as pCASL, and could be used in conjunction with other vessel selective methods such as an oblique labelling plane. Furthermore it could be implemented with a Hadamard encoding scheme [89, 242] to accelerate the acquisition of multiple perfusion territories.

9.1.2 Pre-Scans

9.1.2.1 B_1^+ mapping

In section 7.2, it was theoretically shown to be possible to acquire high SNR B_1^+ maps of individual, small coils using an AFI sequence, with both coils pulsing in linear combinations (see section 7.2.3.2) from which the individual coil B_1^+ maps could be deduced via matrix inversion. However, the technique is corrupted by improperly formed linear combinations, which is the case when inter-coil coupling is present. With adequate decoupling, it should be possible to acquire highly accurate, high SNR in-vivo B_1^+ maps which are suitable for parallel transmission CASL.

In general, B_1^+ mapping methods are time-consuming, and parallel transmit ASL would benefit from further acceleration. At present the global excitation pulsed used by the coils mean that far more data is acquired than is necessary. Slab selective pulses could be used to limit excitation to the neck only, or potentially just a single slice at the labelling plane. The former is favoured, because there will be less slice profile effects to confound the B_1^+ map. Slice/slab selective pulses require more precise synchronisation between the Trio and SMIS console, and also additional data about the RF pulse would need to be passed to the SMIS console, with an increased chance of human error in entering these parameters correctly.

9.1.2.2 Anatomical scans

The time of flight angiogram acquired is of suitable resolution for defining the carotid arteries, and the acquisition time (approximately one minute) is acceptable. In addition it would be useful to acquire a gated phase contrast angiogram at the labelling plane, from which the blood velocity within the labelling planes could be measured. Using the peak and average blood velocity the labelling parameters could then be optimised on a per-subject basis. This would however increase the amount of post-processing, particularly because the carotid arteries would need to be identified. In chapter 3 a semi-automated method for determining the locations and size of the carotid and vertebral arteries was presented, this could be implemented in the scan planning software. It would also aid the selection of appropriate label and null regions.

9.1.3 ASL Labelling Coils

For global and single coil CASL, the labelling coils performed well. They produce a suitable B_1^+ field using low power RF amplifiers for efficient inversion of arterial blood even when not optimally positioned on the neck. By using fixed tune and match capacitor values, the time consuming tune and match procedure is eliminated. Series PIN diodes detuning ensured no coupling to the body transmit coil, even with no reverse bias voltage applied. The mechanical design permitted flexible positioning on the neck without excess weight or pressure being placed on the windpipe. SAR mapping confirmed that deposited power from the coils was within regulatory limits. However, as a

transmit array the lack of an appropriate decoupling scheme means that optimised drive coefficients cannot be calculated or applied correctly. A variety of decoupling methods are available, which are discussed in section 9.3.1.

Two coils were used because this was the number of available transmitter channels on the SMIS console, however this restricts the number of degrees of freedom to spatially tailor the labelling B_1^+ field and more coil elements would be advantageous. Being able to move the coils is advantageous for global and single coil CASL because they can be located in optimal positions. However for parallel transmit ASL this is not necessary. Movement of the coils during the scan affects their loading, the amount of inter-coil coupling, and reduces the validity of the acquired B_1^+ maps. Furthermore different coil positions would have an effect on power deposition. While the SAR measurements in chapter 6 suggest that there is flexibility in coil position without exceeding regulatory limits, it would be better to have fixed coil positions to reduce deviations from the known safe configuration. Fixed coils cannot be placed directly on the neck, making a more powerful RF amplifier necessary, however the increased distance between the subject's neck and the high power deposition close to the coil might mean that peak local power deposition is reduced. Moreover, fixed position coils also permit the use of coil overlapping and inductive or capacitive networks for decoupling, therefore additional transmit channel components might not be required as would be the case with alternative decoupling methods.

9.1.4 Transmitter System

A recycled spectrometer was used as the base of the transmitter system, which would have otherwise been condemned as junk. In addition to the merits of upcycling unused equipment, using a spectrometer provides a lot more flexibility compared to the gated frequency source normally seen in separate coil CASL implementation. Full amplitude and phase control was possible of each channel independently, and the ability to implement a pulse sequence enables more advanced techniques such as RF spoiling (which for the AFI sequence the increment alternated between two values), conditional triggering, and shaped pulses. Even for a single transmitter channel this configuration is has many benefits, and adds further flexibility for applications beyond ASL, such as multi-nuclear decoupling, or long off-resonance pulses without excessive SAR for magnetisation transfer contrast imaging [240].

Whilst fully functioning, the spectrometer technology is old compared to what is state of the art. Most noticeable is the scale and complexity of the entire system; it is composed of many proprietary modules from a manufacturer that no longer exists. It also takes up a large amount of space, approximately $2\times1\times1.5m$. The system is limited to a maximum of four transmitter channels (if additional modules can be sourced), uses a very old CPU (running at 500MHz), runs a long out of date operating system (Windows NT 4.0), and the pulse sequence language is proprietary,

hence it cannot have additional libraries incorporated as would be the case with a C/C++ based system. Triggering of the system achieved by polling (checking for a trigger signal within a loop) the input port, resulting in a not-insignificant delay of approximately $200\mu s$. For non-spatially selective pulses in gradient echo sequences this is not of consequence, but could be problematic for other pulse schemes which require more accurate synchronisation. Limited communication is available between the Trio and SMIS console. The Trio is capable of unidirectional (send only) communication via optical out, over which is it possible to send serial bytes, but additional hardware would be required to pass this data to the SMIS console.

The RF power amplifiers used are restricted in power, beneficially preventing excess power deposition, but also limits the maximum B_1^+ that can be achieved. Furthermore the amplifiers are not linear over their entire dynamic range, and they have generally been run at full power. A better solution would be to use amplifiers capable of outputting more power, for example 10W, so that the additional power is there if it is required, and ensuring that the amplifiers operate more in their linear region rather than at their maximum where the output is compressed. Furthermore, operating an amplifier continuously close to its maximum output level shortens its lifespan.

9.1.5 SAR Mapping Technique

A large conservative safety margin was present in the SAR measurements, due to possible over-estimation of the SAR (MRTI SAR measurements at voxels corresponding to the location of a fibre optic thermometer probe were found be be approximately 1.5× greater than measured by the thermometer), and because no mass averaging was performed. Because the objectives of the SAR measurements was to evaluate the safe use of the ASL Labelling Coils, this over-estimation worked in favour, however there are situations where accurate SAR measurements are required. Therefore a full analysis of the errors, and further measurements with more fibre optic thermometry measurements to validate against is essential. In addition the method could be validated against results from EM modelling (with thermal solvers) and MR signal modelling.

One obvious drawback for this method is the long duration required to make measurements. Several hours of system stabilisation whilst running the EPI sequence was mandatory for the B_0 drift compensation algorithm, and temperature measurements were made for an hour. However many MR systems are not used overnight or at weekends, and the measurement procedure could be automated.

9.2 Potential for the technique

This section describes the potential for the concept of using parallel transmission vessel selective arterial spin labelling. Whilst the work undertaken in this thesis has focused mainly on the specific labelling scenario of reducing contralateral labelling in comparison with using a single labelling coil,

parallel transmit CASL has the potential for further flexibility in vessel selection. Also, by using B_1^+ maps and anatomical scans to plan the labelling, rather than relying on the correct placement of the labelling coils, reproducibility should be improved.

As previously mentioned the main drawback is the need for additional hardware, however parallel transmit capabilities are becoming standard on 7T systems, and have found use at 3T [149], suggesting that in the future this capability might become a standard feature. Another drawback was the lack of automation between the Trio, SMIS console and the scan planning computer. With hardware control fully incorporated into a pulse sequence, the scan planning could be performed on the main MR console, such that all the pre-scan data is automatically processed, providing seamless experience for the MR operator.

9.2.1 Two Channel System

A two channel system, consisting of a small, low power, standalone transmitter system which can be automatically configured by a pulse sequence from any MR vendor, and a fixed position coil array would be optimal. Coil elements which are fixed in position with respect to each other eliminates errors due to motion, and facilitates decoupling through the use of passive capacitive or inductive networks. Geometry of the coil elements should be optimised to provide coverage over the entire neck for a range of neck sizes and artery locations, yet minimal coverage in the head-foot direction. This will minimise the amount of MTC induced in the imaging volume, permitting the coil and labelling plane to be positioned closer to the Circle of Willis (CoW), where transit times are shorter hence less T_1 decay of the labelled blood, and potentially facilitating labelling of arteries distal to the CoW. Because parallel transmission is used for specifying the labelling vessels, the geometry of the individual coil elements do not need to be optimised for good left–right differentiation. While it is important to maintain a fixed spatial relationship between the coils, the entire array could be positionable up or down the neck, or alternatively the array could consist of multiple pairs of coils arranged repeatedly along the z direction, which can be switched using PIN diodes (such that only two coils can be used simultaneously).

Figure 9.1 is a scale drawing of the preferred coil geometry. The new coils, with conductors shown in orange have dimensions $R_{ar} = 8cm$, $Z_{ar} = 2cm$ and $\phi_{ar} = 120°$ and are shown alongside the existing circular loop coils (shown in blue) which have radius 3.2cm, and are positioned directly on the neck (represented by the grey wire-frame cylinder of radius 5cm).

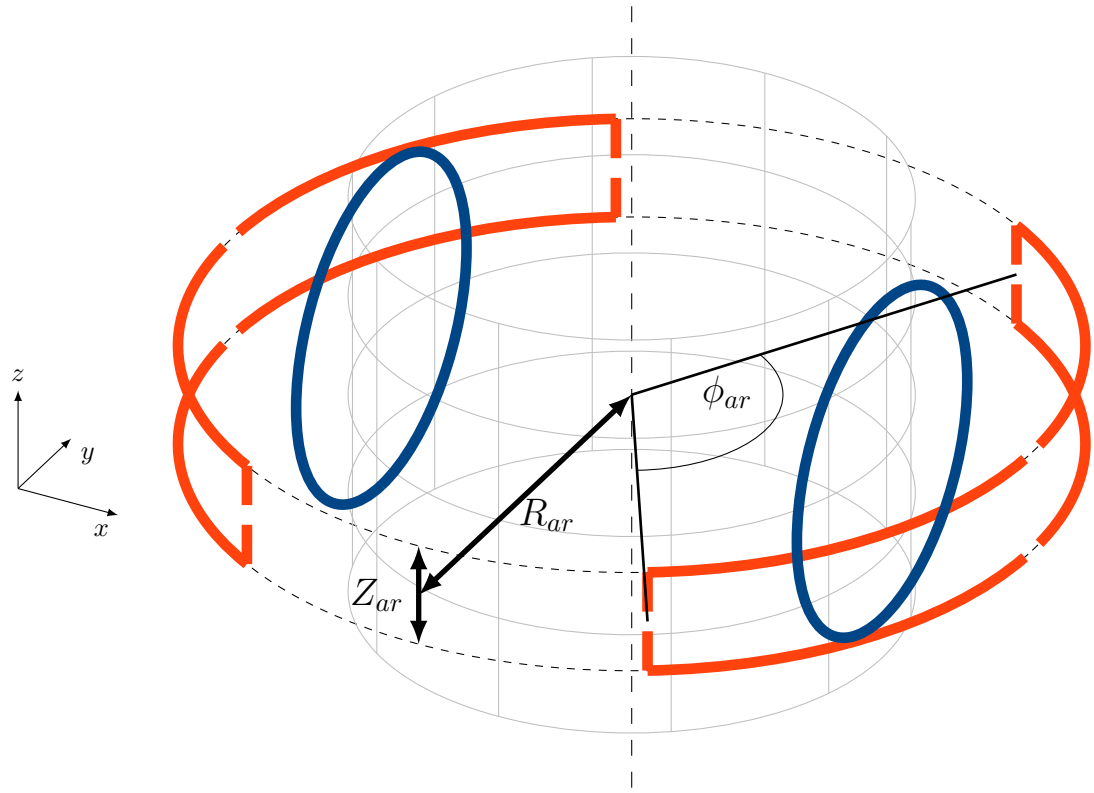

Figure 9.1: Possible Two Channel Coil Geometry

9.2.2 Multi Channel System

Multiple channels will provide more ability to spatially tailor the excitation field. By doing so the B_1^+ null can be optimised to have decreased sensitivity to motion. However, multiple coil elements will result in increased power deposition, therefore measures will need to be taken to control this.

9.2.2.1 General Parallel Transmit System

A typical parallel transmit system uses a multi-element head only or body transmit coil, for example a stripline array or degenerate birdcage. These could be used for parallel transmit vessel selective CASL, however because a volume transmitter is used, which also irradiates the imaging slice, similar issues to standard volume coil CASL will be introduces: MTC, a shortening of the T_1 of the blood and tissue and high power deposition. If there is control over the B_1^+ field in the z direction then parallel transmission could also be used to simultaneously null the B_1^+ field within the imaging volume to reduce MTC.

Alternatively, parallel transmission could be used with a pCASL labelling scheme, as suggested in [255], which is more favourable for volume transmitters because of the reduced power deposition, MTC and no need for continuous wave capable RF amplifiers.

9.2.2.2 Specialised Coil

A specialised coil could be designed to utilise the advantages of multi-channel transmit with the zero MT effects by enabling as much specificity in the x-y plane as possible, yet minimise coverage in the z direction. A possible labelling array is shown in figure 9.2. This has the same dimensions as those of figure 9.1, however there are now eight coil elements, each subtending 35°, resulting in a high degree of spatial specificity, potentially permitting the nulling of several arteries. An increased number of coils introduces new challenges, particularly in acquiring B_1^+ maps of all the coils in a fast enough time frame. Because of the decreased spatial coverage in the z-direction, it should be possible to acquire a B_1^+ map over a small slab (for example four to six slices) using the AFI sequence in less than one minute. Using a faster sequence, such as single-shot STEAM EPI [98], or through high undersampling and an iterative, non-linear reconstruction method [175], the acquisition of a full set of B_1^+ maps could be further accelerated.

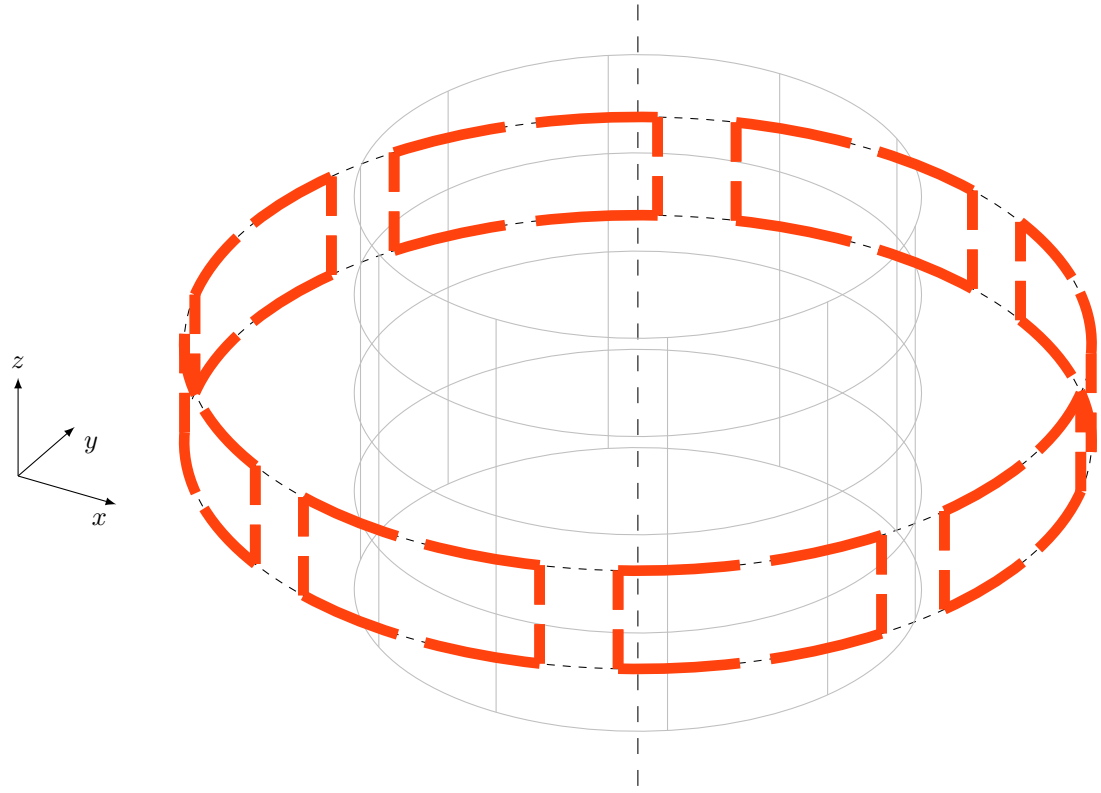

Figure 9.2: Possible Eight Channel Coil Geometry

9.3 Further work

The following areas have been identified as potential avenues for further work.

9.3.1 Decoupling

In order to correctly perform parallel transmit CASL, the ASL Labelling Coils need to be correctly decoupled from each other. For the current coil design, where the coils can be moved in space, three options are available: current source/ultra low-output impedance amplifiers, cartesian feedback, and digital decoupling.

A current source amplifier [135] can be created by placing a RF power MOSFET in the saturation region of its DC characteristics [132], such that it behaves as a voltage controlled current source. A series resonating stripline RF "rung" is then driven directly by this MOSFET, the high internal resistance of the current source meaning that it becomes an open circuit to induced current, suppressing any currents caused by inter-coil coupling. However the series resonant rung is severely mismatched to the MOSFET, degrading the maximum available output power.

A method analogous to using ultra-low input impedance preamplifiers has been developed which utilises ultra-low output impedance power amplifiers to isolate coil elements, removing constraints on transmit array geometry [40]. This also uses a MOSFET operating in saturation mode to create a voltage controlled current source, however to maximise the output power an output matching network is also used. This provides the correct output impedance for maximum power output from the MOSFET, and further transforms the output impedance to be extremely low, resonating the matching network at the coil and effectively creating a high blocking impedance to induced currents. A coaxial cable of length $^{n\lambda}/_{2}$, where n is an integer, is required to ensure that the low output impedance of the amplifier is not transformed to another value, therefore the amplifiers can be placed outside the magnet room, however loss within connecting coaxial cables affects the amount of decoupling that can be achieved.

Cartesian Feedback [112, 113] is an "active" decoupling method, in which the current within each array coil element is continuously monitored, compared with the desired amplitude and phase, and a correctly phased error signal is used as negative feedback to correct for deviations. Induced currents from inter-coil coupling are blocked, effectively adding a high resistance in series with the coil. This negative feedback is achieved using analogue electronics, which are able to continuously adapt to any changes, resulting in "absolute calibration" of the spectrometer.

Active decoupling can also be implemented in the digital domain, a technique called digital vector decoupling [199, 127, 225]. By transmitting on each coil element sequentially with a nominal amplitude and phase, and measuring the induced currents amplitude and phase in each coil, a matrix of the coil coupling terms can be determined. This coupling matrix relates desired RF currents to the voltages that must be applied to coupled RF coils to achieve those currents. Digital vector decoupling requires a calibrated means to measure the current in each coil, for example a small transformer or pickup coil in close proximity to the loop, and as induced currents are opposed by an inverse amplitude and phase from the RF amplifier and not blocked, induced power travels

from the coils to the amplifiers and must be re-routed using circulators to prevent power amplifier damage.

Out of these four methods, digital vector decoupling is the most suitable and easiest to implement with the existing hardware and system. Both the current source and ultra-low output impedance amplifier methods require significant modifications to the coils, and the advanced RF engineering methods are beyond the scope of this work. Similarly, the cartesian feedback system is extremely complex. However, these three methods are capable of automatically compensating for changes in coupling, hence robust to motion, whereas digital vector decoupling is only valid so long as conditions remain the same as when the coupling matrix was measured.

Digital vector decoupling could be implemented as a calibration phase prior to commencing scanning, much like the power calibration of the RF transmitter, or the acquisition of receive coil sensitivities. For two coils, only two measurements need to be made: transmit on coil one, and measure the induced current amplitude and phase on coil two, and vice-versa. For this a calibrated current transformer or pickup coil is necessary, these would be mounted directly on to each coil's PCB and calibrated with a S_{21} measurement on a network analyser to determine the amount amplitude and phase of the coupled signal. Alternatively a complex (amplitude and phase) power sensor can be used to sample the reverse power. The coupled signal must be sampled, and the amplitude and phase determined. This could be achieved by complex demodulation with the transmit frequency, much in the same way a MR receive channel operates, or by using a dedicated RF amplitude and phase detector (for example AD8302, Analog Devices, Norwood, MA, USA), and then sampling the magnitude and phase output with a precision ADC and microcontroller.

Once both coupling coefficients are determined, the driving coefficient correction factors need to be calculated, and applied to each coil in response to the desired amplitudes and phases. Ideally this would be done automatically on the SMIS spectrometer, such that an amplitude and phase could be specified and the correction factors automatically applied transparently to the user.

With correctly implemented decoupling it will be possible to utilise the B_1^+ maps acquired in the pre-scans to calibrate the RF amplitude of the labelling coils. Also, by acquiring of a phase contrast angiogram for carotid velocity measurements, the B_1^+ amplitude and labelling gradient can be appropriately chosen to provide a high inversion efficiency robust to variations in B_1^+ and flow velocity.

9.3.2 Updated Coil Design

A new, optimised labelling coil array could also be the subject of further work. As discussed in section 9.2.1, it is hypothesised that better labelling performance will be achieved using fixed position coils that have a large transmit sensitivity in the x-y plane, but minimal along the z-axis. With reduced z coverage the coil could be positioned closer to the CoW without causing MTC in

the imaging volume, reducing transit time and improving the inversion efficiency of the labelled bolus of blood. Furthermore, the larger extent in the x-y plane will mean that it is not necessary to carefully locate the coil close to the common carotid arteries. Figure 9.1 exemplifies this geometry, presenting a simple parametrised model suited for both simulation and manufacture.

For this new coil design, more precedence should be placed in ensuring robustness to variations in neck size and artery position, therefore anatomical measurements made in section 3.1.1 in chapter 3 should be made on a sample population with an emphasis on a wide range of body sizes (hence wide range of neck sizes). Optimised coil geometry can be found using numerical electromagnetic and Bloch equation simulations, with the aim of producing a suitable B_1^+ for efficient labelling over the a range of anatomical dimensions, and have a large degree of freedom in specifying a B_1^+ null, whilst conforming to the power deliverable by the RF amplifiers, and the SAR limits.

9.3.3 System Automation

A lot of user interaction is required to correctly set the parameters on the SMIS console. This contrasts to standard scanning where the console scan software takes care of all calculations in the background. As it is both time consuming, and increases the chances of an error, it would be very beneficial to automate setting the parameters in the SMIS console.

At a basic level, the frequency needs to be reliably synchronised between the two systems. At present this is done by manually entering the Trio frequency into the SMIS parameter editor, and an unknown error can sometimes result in there being a 1Hz mismatch between the two frequencies. Typically the SMIS console's centre frequency, f_s is set to the Trio's centre frequency f_t, however on some occasions a spatial shift is seen in the images, obtained when pulsing on the ASL Coils (for example in the AFI sequence) corresponding to the SMIS actually being 1Hz ahead, i.e. $f_s = f_t + 1Hz$. It is suspected that this is an error in the Trio user interface, where the frequency is not always displayed correctly. Therefore, as a calibration stage a low resolution scan (hence fast to acquire) is acquired to determine if there is a frequency difference, and if one occurs then the value of f_s is adjusted appropriately. If the Trio performs another frequency calibration, which is normally done when the imaging volume changes, then the SMIS frequency also must be changed and verified with a low resolution scan. A more reliable way to set the SMIS frequency would be for the Trio to automatically send its centre frequency to the SMIS console at the start of each scan.

Furthermore, it would be convenient for other parameters to be automatically sent to the SMIS console, for example the amplitudes and phases calculated in the parallel transmit CASL scan planning software. Doing so would reduce the amount of user interaction, and subsequent user error, and speed up the "dead time" between scans when parameters must be set.

A way to implement this automation would be to use a microcontroller to relay information

between the Trio console, the SMIS console and the scan planning software residing on a dedicated computer. A typical microcontroller has a UART (Universal Asynchronous Receiver Transmitter), facilitating serial communication, and can also be interfaced with an Ethernet controller for TCP/IP communication. The Trio has TCP/IP support within its pulse sequence environment, and its optical trigger output can also be used to send serial bytes (however speed is slow and this is a one-way communication channel). Within its pulse sequence environment the SMIS console only supports communication via eight TTL level outputs and four TTL level inputs, effectively eight and four bit parallel interfaces, respectively.

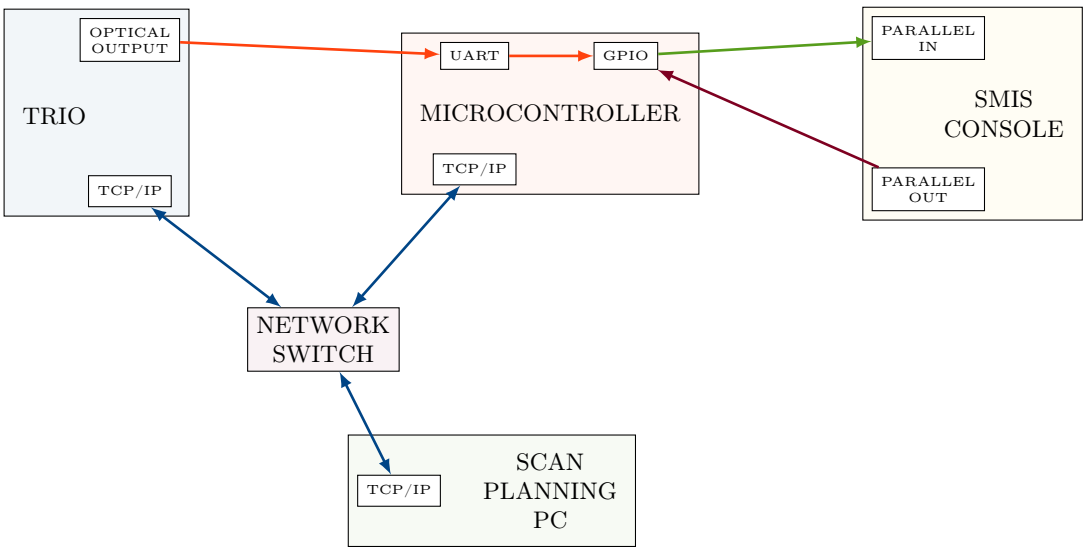

Figure 9.3: Using a microcontroller to aid communication between the Trio, SMIS and a standalone computer

Figure 9.3 illustrates how a microcontroller can aid communication between devices. Scan parameters such as the system frequency, offset frequency and pulse parameters can be relayed from the Trio to the microcontroller over a network connection, where, using its GPIO (General Purpose Input/Output; digital pins that can be software assigned for any purpose) parameters can be sent to the SMIS console, and confirmation of receipt sent back. Once all parameters have successfully been relayed, the Trio would indicate to the microcontroller that a scan was about to commence, and the microcontroller would then go into a mode whereby it passed on trigger pulses directly to the SMIS console.

9.3.4 Comprehensive Parallel Transmit CASL Study

With both decoupling and system automation implemented, and potentially a new labelling coil, parallel transmit CASL could be compared with other vessel selective ASL techniques in an in-vivo study on multiple volunteers, furthering the work of chapter 8. Of interest would be specificity, signal to noise and inter-session repeatability.

9.3.5 pCASL and Separate Coil CASL Comparison

There has not yet been a study which directly compares separate labelling coil CASL with pCASL. In the original pCASL implementation [45], a single slice volume coil CASL acquisition was used as a gold standard for comparing the inversion efficiency of the pCASL sequence. Their CASL implementation had a duty cycle of 53.33%, suggesting that it had an inversion efficiency that was reduced in comparison to using a dedicated labelling coil which is capable of 100% duty cycle. While the pCASL sequence showed that it was capable of comparable inversion efficiency to a volume coil CASL sequence, there may still be benefits from using a separate labelling coil, both in inversion efficiency, and in increased sensitivity because there are no MT effects.

It is proposed that the existing 3D-GRASE CASL sequence could be modified to include an option for a pCASL labelling module, resulting in a single sequence capable of directly comparable PASL, CASL and pCASL measurements. Additionally the pCASL pulses could be implemented on the SMIS console, which would allow for an assessment of the MT effects in a volume coil pCASL sequence, through comparison with the same labelling using localised transmitter coils.

9.3.6 Perfusion Phantom

A high flow resistance, and physiologically unrealistic flow velocities were found to be two areas that could be improved in the perfusion phantom. As suggested in section 7.4.4, a new phantom could be designed and manufactured which rectifies these problems. Computational fluid analysis (for example using the open source software package OpenFOAM [3]) of the CAD model can reveal regions responsible for restricting flow, and with a marked reduction in flow resistance, physiologically representative flow velocities should be possible.

Further improvements can also be made. The vessel structure could be made more representative of the cerebrovascular system, particularly the paths of the internal carotid arteries and the middle cerebral arteries. Perfusion territories which are representative of those found within the human brain would increase the utility of the phantom in the development of vessel selective ASL techniques. The size of the phantom was restricted by the dimensions of its container, however a purpose built vessel (for example from acrylic) could be manufactured to house the phantom, allowing for a larger perfusion chamber whose size is anatomically correct. Regions of doped water/gel with realistic T_1 decay times for grey matter, white matter and CSF would be of use when developing or testing background suppression schemes. Macromolecular content could also be introduced to these regions, enabling an assessment of MTC and other MT effects.

As shown in section 7.4, it was useful to obtain a gating signal for the pulsatile flow within the phantom. However, the means of acquiring this signal (pulse oximeter placed around a red silicone tube) was fortuitous, and only reliable at high flow rates, warranting a more robust method. A flow meter of some sorts could be included in the connecting tubing, and appropriate electronics

to condition and process the signal to obtain a reliable trigger pulse for sequences such as a phase contrast angiogram. Moreover this flow meter would provide an independent assessment of the global flow rate, providing a measure of the repeatability and reliability of the phantom.

It would also be useful if the phantom could have a way to inject a bolus of contrast agent, for example gadolinium for dynamic contrast enhanced MRI, or iodine for digital subtraction angiography, facilitating comparisons between the different methods and modalities.

Finally, rapid prototyping enables multiple phantoms to be produced quickly and inexpensively. It would be interesting to see whether it would be possible to obtain repeatable results across multiple phantoms and MR systems, enabling their use as a quality control measure in multi centre trials.

9.4 Conclusion

To summarise, the work in this thesis has focused on applying parallel transmission methods to vessel selective ASL. Initial computer simulations confirmed that an improvement in vessel specificity, whilst retaining labelling efficiency, was possible when using parallel transmission in comparison to a single labelling coil. A pair of radio frequency ASL Labelling Coils was designed, constructed and characterised. Each coil was tuned and matched to average neck loading conditions, thus removing the need for per-subject adjustment. Individual PIN diode detuning permits both individual and simultaneous use of the coils, and coils provide very good specificity and labelling efficiency without using parallel transmission.

A low power two-channel transmitter system for use in conjunction with a 3T clinical MR system was constructed, based around a recycled imaging spectrometer. Full amplitude and phase control was possible, and using a spectrometer provides much more flexibility in hardware control than the gated RF source that has been used in previous implementations. A novel technique for producing accurate temperature maps of very small temperature changes was developed and applied to measure the SAR of the ASL Labelling Coils. Other potential applications for this method include verification of electromagnetic simulation results and the testing of implanted devices.

A B_1^+ mapping technique was optimised to acquire accurate, high SNR B_1^+ maps of small transmitter coils, even in regions where the excitation flip angle is low. An advanced CASL sequence was developed that utilised a 3D-GRASE acquisition module and interleaved background suppression. This used the same acquisition module as a FAIR PASL sequence, facilitating direct comparison between the two labelling schemes. A perfusion phantom with distinct vascular territories was designed and constructed for the purpose of testing vessel selective ASL sequences. With some refinement, this phantom could be used as a universal perfusion phantom, with applications such as quality control measurements, and sequence development without the need for subjects.

Finally the parallel transmit ASL sequence was used on a group of healthy volunteers, with the aim of improving left–right specificity when labelling individual common carotid arteries. In comparison with using only one of the labelling coils an improvement of up to 35% in specificity was seen. With a coil array that is optimised for global coverage, and a high degree of freedom in specifying the labelling field, this improvement could be much larger, whilst at the same time improving repeatability by reducing the dependence on coil position.

References

[1] NEMA MS 10-2006: Determination of Local Specific Absorption Rate (SAR) in Diagnostic Magnetic Resonance Imaging (MRI).

[2] NEMA MS 8-2008: Characterization of the Specific Absorption Rate (SAR) for Magnetic Resonance Imaging Systems.

[3] OpenFOAM: an open source library designed for development of parallel or serial multi-dimensional modeling code.

[4] ISO 13485:2003 Medical devices – Quality management systems – Requirements for regulatory purposes, 2003.

[5] IEC 60601-2-33, Edition 3. Medical electrical equipment - Part 2-33: particular requirements for the safety of Magnetic Resonance equipment for medical diagnosis, 2010.

[6] Agilent Technologies. Understanding the fundamental principles of vector network analysis, agilent application note 1287-1. Retreived September 21 2010 from the World Wide Web, http://cp.literature.agilent.com/litweb/pdf/5965-7707E.pdf, 1997.

[7] S. Akoka, F. Franconi, F. Seguin, and A. Le Pape. Radiofrequency map of an nmr coil by imaging. *Magnetic Resonance Imaging*, 11(3):437 – 441, 1993.

[8] Donald W. Alderman and David M. Grant. An efficient decoupler coil design which reduces heating in conductive samples in superconducting spectrometers. *Journal of Magnetic Resonance (1969)*, 36(3):447 – 451, 1979.

[9] D C Alsop and J A Detre. Reduced transit-time sensitivity in noninvasive magnetic resonance imaging of human cerebral blood flow. *J Cereb Blood Flow Metab*, 16(6):1236–1249, November 1996.

[10] D C Alsop and J A Detre. Multisection cerebral blood flow mr imaging with continuous arterial spin labeling. *Radiology*, 208(2):410–416, 1998.

[11] D. Andreuccetti, R. Fossi, and C Petrucci. Dielectric properties of body tissues in the frequency range 10hz - 100ghz. Internet Resource: http://niremf.ifac.cnr.it/tissprop/ (November 2011), 1997.

[12] Cristina Armenean, Mircea Armenean, Emmanuel Perrin, Olivier Beuf, Frank Pilleul, and Hervé Saint-Jalmes. Rf heating comparison between conductive and resistive wires in interventional and endoluminal mri. In *Proceedings of the International Society for Magnetic Resonance in Medicine 11*, 2004.

[13] Cristina Armenean, Emmanuel Perrin, Mircea Armenean, Olivier Beuf, Frank Pilleul, and Hervé Saint-Jalmes. Rf-induced temperature elevation along metallic wires in clinical magnetic resonance imaging: Influence of diameter and length. *Magnetic Resonance in Medicine*, 52(5):1200–1206, 2004.

[14] R.C. Aster, C.H. Thurber, and B. Borchers. *Parameter estimation and inverse problems*. International geophysics series. Elsevier Academic Press, 2005.

[15] D.L. Bailey. *Positron emission tomography: basic sciences*. Springer, 2005.

[16] Peter A. Bandettini, Eric C. Wong, R. Scott Hinks, Ronald S. Tikofsky, and James S. Hyde. Time course epi of human brain function during task activation. *Magnetic Resonance in Medicine*, 25(2):390–397, 1992.

[17] J.-C. Baron. Perfusion thresholds in human cerebral ischemia: Historical perspective and therapeutic implications. *Cerebrovasc Dis*, 11(Suppl. 1):2–8, 2001.

[18] Peter J. Basser and Derek K. Jones. Diffusion-tensor mri: theory, experimental design and data analysis – a technical review. *NMR in Biomedicine*, 15(7-8):456–467, 2002.

[19] B. Behnia, M. Suthar, and A. G. Webb. Closed-loop feedback control of phased-array microwave heating using thermal measurements from magnetic resonance imaging. *Concepts in Magnetic Resonance*, 15(1):101–110, 2002.

[20] Thomas Benner, André J.W. van der Kouwe, John E. Kirsch, and A. Gregory Sorensen. Real-time rf pulse adjustment for b0 drift correction. *Magnetic Resonance in Medicine*, 56(1):204–209, 2006.

[21] J. Bittoun, J. Taquin, and M. Sauzade. A computer algorithm for the simulation of any nuclear magnetic resonance (nmr) imaging method. *Magnetic Resonance Imaging*, 2(2):113 – 120, 1984.

[22] J.-C. Bolomey. Overview of fast sar assessment techniques. In *Antenna Technology: Small Antennas and Novel Metamaterials, 2008. iWAT 2008. International Workshop on*, pages 47 –50, march 2008.

[23] P A Bottomley and E R Andrew. Rf magnetic field penetration, phase shift and power dissipation in biological tissue: implications for nmr imaging. *Physics in Medicine and Biology*, 23(4):630, 1978.

[24] Paul A. Bottomley and William A. Edelstein. Power deposition in whole-body nmr imaging. *Medical Physics*, 8(4):510–512, 1981.

[25] Paul A. Bottomley, Thomas H. Foster, Raymond E. Argersinger, and Leah M. Pfeifer. A review of normal tissue hydrogen nmr relaxation times and relaxation mechanisms from 1–100 mhz: Dependence on tissue type, nmr frequency, temperature, species, excision, and age. *Medical Physics*, 11(4):425–448, 1984.

[26] Paul A. Bottomley, Christopher J. Hardy, Peter B. Roemer, and Otward M. Mueller. Proton-decoupled, overhauser-enhanced, spatially localized carbon-13 spectroscopy in humans. *Magnetic Resonance in Medicine*, 12(3):348–363, 1989.

[27] Paul A. Bottomley, Rowland W. Redington, William A. Edelstein, and John F. Schenck. Estimating radiofrequency power deposition in body nmr imaging. *Magnetic Resonance in Medicine*, 2(4):336–349, 1985.

[28] Paul A. Bottomley and P. B. Roemer. Homogeneous tissue model estimates of rf power deposition in human nmr studies. *Annals of the New York Academy of Sciences*, 649(1):144–159, 1992.

[29] Christopher Bowick. *RF Circuit Design.* Newnes, February 1997.

[30] D. O. Brunner and K. P. Pruessmann. A matrix approach for mapping array transmit fields in under a minute. In *Proc. Intl. Soc. Mag. Reson. Med. 16*, 2008.

[31] D. O. Brunner and K. P. Pruessmann. $\mathbf{B}_1^+$ interferometry for the calibration of rf transmitter arrays. *Magn. Reson. Med.*, 61(6):1480–1488, 2009.

[32] David O. Brunner, Nicola De Zanche, Jurg Frohlich, Jan Paska, and Klaas P. Pruessmann. Travelling-wave nuclear magnetic resonance. *Nature*, 457(7232):994–998, February 2009.

[33] Richard B. Buxton, Lawrence R. Frank, Eric C. Wong, Bettina Siewert, Steven Warach, and Robert R. Edelman. A general kinetic model for quantitative perfusion imaging with arterial spin labeling. *Magnetic Resonance in Medicine*, 40(3):383–396, 1998.

[34] K. Caputa, M. Okoniewski, and M.A. Stuchly. An algorithm for computations of the power deposition in human tissue. *Antennas and Propagation Magazine, IEEE*, 41(4):102 –107, sep 1999.

[35] David W. Carmichael, David L. Thomas, Enrico De Vita, Maria A. Fernández-Seara, Navjeet Chhina, Mark Cooper, Colin Sunderland, Chris Randell, Robert Turner, and Roger J. Ordidge. Improving whole brain structural mri at 4.7 tesla using 4 irregularly shaped receiver coils. *NeuroImage*, 32(3):1176 – 1184, 2006.

[36] Alexandru Cernicanu, Matthieu Lepetit-Coiffe, Joerg Roland, Christoph D. Becker, and Sylvain Terraz. Validation of fast mr thermometry at 1.5 t with gradient-echo echo planar imaging sequences: phantom and clinical feasibility studies. *NMR in Biomedicine*, 21(8):849–858, 2008.

[37] Julio A. Chalela, David C. Alsop, Julio B. Gonzalez-Atavales, Joseph A. Maldjian, Scott E. Kasner, and John A. Detre. Magnetic resonance perfusion imaging in acute ischemic stroke using continuous arterial spin labeling. *Stroke*, 31(3):680–687, 2000.

[38] C.N. Chen and DI Hoult. *Biomedical magnetic resonance technology*. Adam Hilger Bristol, 1989.

[39] Soke Miang Chng, Esben Thade Petersen, Ivan Zimine, Yih-Yian Sitoh, C.C. Tchoyoson Lim, and Xavier Golay. Territorial arterial spin labeling in the assessment of collateral circulation. *Stroke*, 39(12):3248–3254, 2008.

[40] Xu Chu, Xing Yang, Yunfeng Liu, Juan Sabate, and Yudong Zhu. Ultra-low output impedance rf power amplifier for parallel excitation. *Magnetic Resonance in Medicine*, 61(4):952–961, 2009.

[41] Harvey Cline, Richard Mallozzi, Zhu Li, Graeme McKinnon, and William Barber. Radiofrequency power deposition utilizing thermal imaging. *Magnetic Resonance in Medicine*, 51(6):1129–1137, 2004.

[42] Adrian P. Crawley, Michael L. Wood, and R. Mark Henkelman. Elimination of transverse coherences in flash mri. *Magnetic Resonance in Medicine*, 8(3):248–260, 1988.

[43] Charles H. Cunningham, John M. Pauly, and Krishna S. Nayak. Saturated double-angle method for rapid b1+ mapping. *Magnetic Resonance in Medicine*, 55(6):1326–1333, 2006.

[44] W. Dai, P.M. Robson, A. Shankaranarayanan, and D. C. Alsop. Benefits of interleaved continuous labeling and background suppression. In *Proc. Intl. Soc. Mag. Reson. 18*, number 1748, 2010.

[45] Weiying Dai, Dairon Garcia, Cedric de Bazelaire, and David C. Alsop. Continuous flow-driven inversion for arterial spin labeling using pulsed radio frequency and gradient fields. *Magnetic Resonance in Medicine*, 60(6):1488–1497, 2008.

[46] N. P. Davies and P. Jezzard. Selective arterial spin labeling (sasl): Perfusion territory mapping of selected feeding arteries tagged using two-dimensional radiofrequency pulses. *Magn. Reson. Med.*, 49(6):1133–1142, 2003.

[47] Enrico De Vita, Matthias Günther, Xavier Golay, and David Thomas. Magnetisation transfer effects of q2tips pulses in asl. *Magnetic Resonance Materials in Physics, Biology and Medicine*, pages 1–14, 2011. 10.1007/s10334-011-0298-z.

[48] Jacco A. de Zwart, Peter van Gelderen, Dan J. Kelly, and Chrit T.W. Moonen. Fast magnetic-resonance temperature imaging. *Journal of Magnetic Resonance, Series B*, 112(1):86 – 90, 1996.

[49] CR Deckard and PF McClure. Selective laser sintering. Technical report, 1988.

[50] J. Delannoy, Ching-Nien Chen, R. Turner, R. L. Levin, and D. Le Bihan. Noninvasive temperature imaging using diffusion mri. *Magnetic Resonance in Medicine*, 19(2):333–339, 1991.

[51] Mary F. Dempsey, Barrie Condon, and Donald M. Hadley. Investigation of the factors responsible for burns during mri. *Journal of Magnetic Resonance Imaging*, 13(4):627–631, 2001.

[52] J. A. Detre, D. C. Alsop, L. R. Vives, L. Maccotta, J. W. Teener, and E. C. Raps. Noninvasive mri evaluation of cerebral blood flow in cerebrovascular disease. *Neurology*, 50(3):633–641, 1998.

[53] John A. Detre, John S. Leigh, Donald S. Williams, and Alan P. Koretsky. Perfusion imaging. *Magnetic Resonance in Medicine*, 23(1):37–45, 1992.

[54] John A. Detre, Weiguo Zhang, David A. Roberts, Afonso C. Silva, Donald S. Williams, Donald J. Grandis, Alan P. Koretsky, and John S. Leigh. Tissue specific perfusion imaging using arterial spin labeling. *NMR in Biomedicine*, 7(1-2):75–82, 1994.

[55] W. T. Dixon, L. N. Du, D. D. Faul, M. Gado, and S. Rossnick. Projection angiograms of blood labeled by adiabatic fast passage. *Magn. Reson. Med.*, 3(3):454–462, 1986.

[56] W. Thomas Dixon, Maziar Sardashti, Mauricio Castillo, and Guido P. Stomp. Multiple inversion recovery reduces static tissue signal in angiograms. *Magnetic Resonance in Medicine*, 18(2):257–268, 1991.

[57] WE Doherty Jr and R. Joos. *The Pin Diode Circuit Designers' Handbook*. 1998.

[58] S. Duewell, R. Wüthrich, A. Buck, C. von Weymarn, and G. K. von Schulthess. Mri signal loss due to microcirculation: Phantom studies. *Magnetic Resonance in Medicine*, 14(2):347–357, 1990.

[59] James D Eastwood, Chad A Holder, Patricia A Hudgins, and Allen W Song. Magnetic resonance imaging with lateralized arterial spin labeling. *Magn Reson Imaging*, 20(8):583–586, October 2002.

[60] R R Edelman, B Siewert, D G Darby, V Thangaraj, A C Nobre, M M Mesulam, and S Warach. Qualitative mapping of cerebral blood flow and functional localization with echo-planar mr imaging and signal targeting with alternating radio frequency. *Radiology*, 192(2):513–520, August 1994.

[61] Robert R. Edelman and Qun Chen. Epistar mri: Multislice mapping of cerebral blood flow. *Magnetic Resonance in Medicine*, 40(6):800–805, 1998.

[62] W. A. Edelstein, G. H. Glover, C. J. Hardy, and R. W. Redington. The intrinsic signal-to-noise ratio in nmr imaging. *Magnetic Resonance in Medicine*, 3(4):604–618, 1986.

[63] Roy Ediss. Probing the magnetic field probe. Retreived September 18, 2010 from the World Wide Web, http://www.compliance-club.com/archive/old_archive/030718.htm, March 2007.

[64] Philipp Ehses, Florian Fidler, Peter Nordbeck, Eberhard D. Pracht, Marcus Warmuth, Peter M. Jakob, and Wolfgang R. Bauer. Mri thermometry: Fast mapping of rf-induced heating along conductive wires. *Magnetic Resonance in Medicine*, 60(2):457–461, 2008.

[65] AbdEl El-Sharkawy, Michael Schär, Paul Bottomley, and Ergin Atalar. Monitoring and correcting spatio-temporal variations of the mr scanner's static magnetic field. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 19:223–236, 2006. 10.1007/s10334-006-0050-2.

[66] D. Erni, T. Liebig, A. Rennings, N.H.L. Koster, and J. Fröhlich. Highly adaptive rf excitation scheme based on conformal resonant crlh metamaterial ring antennas for 7-tesla traveling-wave magnetic resonance imaging. In *Proc. IEEE. EMBC. 33*, 2011.

[67] Emilio Esparza-Coss, Jarek Wosik, and Ponnada A. Narayana. Perfusion in rat brain at 7 t with arterial spin labeling using fair-truefisp and quipss. *Magn. Reson. Imag.*, 28(4):607 – 612, 2010.

[68] Maria A. Fernández-Seara, Jiongjiong Wang, Ze Wang, Marc Korczykowski, Matthias Guenther, David A. Feinberg, and John A. Detre. Imaging mesial temporal lobe activation during scene encoding: Comparison of fmri using bold and arterial spin labeling. *Hum. Brain Mapp.*, 28(12):1391–1400, 2007.

[69] María A. Fernández-Seara, Ze Wang, Jiongjiong Wang, Heng-Yi Rao, Matthias Guenther, David A. Feinberg, and John A. Detre. Continuous arterial spin labeling perfusion measurements using single shot 3d grase at 3 t. *Magnetic Resonance in Medicine*, 54(5):1241–1247, 2005.

[70] Bernd U. Foerster, Dardo Tomasi, and Elisabeth C. Caparelli. Magnetic field shift due to mechanical vibration in functional magnetic resonance imaging. *Magnetic Resonance in Medicine*, 54(5):1261–1267, 2005.

[71] Lawrence K Forbes and Stuart Crozier. On a possible mechanism for peripheral nerve stimulation during magnetic resonance imaging scans. *Physics in Medicine and Biology*, 46(2):591, 2001.

[72] Lawrence R. Frank, Eric C. Wong, and Richard B. Buxton. Slice profile effects in adiabatic inversion: Application to multislice perfusion imaging. *Magnetic Resonance in Medicine*, 38(4):558–564, 1997.

[73] Kendra M. Franklin, Brian M. Dale, and Elmar M. Merkle. Improvement in b1-inhomogeneity artifacts in the abdomen at 3t mr imaging using a radiofrequency cushion. *Journal of Magnetic Resonance Imaging*, 27(6):1443–1447, 2008.

[74] C. Gabriel. Compilation of the dielectric properties of body tissues at rf and microwave frequencies. Technical report, DTIC Document, 1996.

[75] C Gabriel, S Gabriel, and E Corthout. The dielectric properties of biological tissues: I. literature survey. *Physics in Medicine and Biology*, 41(11):2231, 1996.

[76] S Gabriel, R W Lau, and C Gabriel. The dielectric properties of biological tissues: Ii. measurements in the frequency range 10 hz to 20 ghz. *Physics in Medicine and Biology*, 41(11):2251, 1996.

[77] S Gabriel, R W Lau, and C Gabriel. The dielectric properties of biological tissues: Iii. parametric models for the dielectric spectrum of tissues. *Physics in Medicine and Biology*, 41(11):2271, 1996.

[78] H. M. Gach and W. Dai. Simple model of double adiabatic inversion (dai) efficiency. *Magn. Reson. Med.*, 52(4):941–946, 2004.

[79] H.M. Gach, A.W. Kam, E.D. Reid, and S.L. Talagala. Quantitative analysis of adiabatic fast passage for steady laminar and turbulent flows. *Magnetic Resonance in Medicine*, 47(4):709–719, 2002.

[80] Xavier Golay, Esben T. Petersen, and Francis Hui. Pulsed star labeling of arterial regions (pulsar): A robust regional perfusion technique for high field imaging. *Magnetic Resonance in Medicine*, 53(1):15–21, 2005.

[81] Xavier Golay, Matthias Stuber, Klaas P. Pruessmann, Dieter Meier, and Peter Boesiger. Transfer insensitive labeling technique (tilt): Application to multislice functional perfusion imaging. *Journal of Magnetic Resonance Imaging*, 9(3):454–461, 1999.

[82] Robin A. De Graaf and Klaas Nicolay. Adiabatic rf pulses: Applications to in vivo nmr. *Concepts in Magnetic Resonance*, 9(4):247–268, 1997.

[83] A. Grabel and J. Millman. *Microelectronics*. McGraw-Hill, 1987.

[84] M. Grant and S. Boyd. Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura, editors, *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pages 95–110. Springer-Verlag Limited, 2008. `http://stanford.edu/~boyd/graph_dcp.html`.

[85] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 1.21,http://cvxr.com/cvx, August 2010.

[86] M. Grant, S. Boyd, and Y. Ye. Disciplined convex programming. *Global Optimization*, pages 155–210, 2006.

[87] W. Grissom, C.-y. Yip, Z. Zhang, V. A. Stenger, J. A. Fessler, and D. C. Noll. Spatial domain method for the design of rf pulses in multicoil parallel excitation. *Magn. Reson. Med.*, 56(3):620–629, 2006.

[88] Mark A. Griswold, Peter M. Jakob, Robin M. Heidemann, Mathias Nittka, Vladimir Jellus, Jianmin Wang, Berthold Kiefer, and Axel Haase. Generalized autocalibrating partially parallel acquisitions (grappa). *Magnetic Resonance in Medicine*, 47(6):1202–1210, 2002.

[89] Matthias Günther. Efficient visualization of vascular territories in the human brain by cycled arterial spin labeling mri. *Magn. Reson. Med.*, 56(3):671–675, 2006.

[90] Matthias Günther, Koichi Oshio, and David A. Feinberg. Single-shot 3d imaging techniques improve arterial spin labeling perfusion measurements. *Magnetic Resonance in Medicine*, 54(2):491–498, 2005.

[91] Mark E. Haacke, Robert W. Brown, Michael R. Thompson, and Ramesh Venkatesan. *Magnetic Resonance Imaging: Physical Principles and Sequence Design*. Wiley-Liss, 1st edition, June 1999.

[92] A. Haase, J. Frahm, D. Matthaei, W. Hanicke, and K. D. Merboldt. Flash imaging. rapid nmr imaging using low flip-angle pulses. *Journal of Magnetic Resonance (1969)*, 67(2):258 – 266, 1986.

[93] J.B. Hagen. *Radio-frequency electronics: circuits and applications*. Cambridge University Press, 2009.

[94] LD Hall and SL Talagala. Mapping of ph and temperature distribution using chemical-shift-resolved tomography. *Journal of magnetic resonance*, 65(3):501–505, 1985.

[95] P. R. Harvey, H. van Hulten, and Z. Zhai. Quantification of the spatial distribution of sar at 3.0t using mr thermal imaging. In *Proc. Intl. Soc. Mag. Reson. 11 (2002)*, number 2649, 2003.

[96] A. Hassan, I. Elabyed, J. Mallow, T. Herrmann, J. Bernarding, and A. Omar. Optimal geometry and capacitors distribution of 7t mri surface coils. In *Microwave Conference (EuMC), 2010 European*, pages 1437 –1440, sept. 2010.

[97] Cecil E. Hayes. The development of the birdcage resonator: a historical perspective. *NMR in Biomedicine*, 22(9):908–918, 2009.

[98] Gunther Helms, Jürgen Finsterbusch, Nikolaus Weiskopf, and Peter Dechent. Rapid radiofrequency field mapping in vivo using single-shot steam mri. *Magnetic Resonance in Medicine*, 60(3):739–743, 2008.

[99] Jeroen Hendrikse, Jeroen van der Grond, Hanzhang Lu, Peter C.M. van Zijl, and Xavier Golay. Flow territory mapping of the cerebral arteries with regional perfusion mri. *Stroke*, 35(4):882–887, 2004.

[100] J. Hennig and H. Friedburg. Clinical applications and methodological developments of the rare technique. *Magnetic Resonance Imaging*, 6(4):391 – 395, 1988. <ce:title>Papers presented at the 1987 Topical Conference on Fast Magnetic Resonance Imaging Techniques</ce:title>.

[101] L. Hernandez-Garcia, D. P. Lewis, B. Moffat, and C. A. Branch. Magnetization transfer effects on the efficiency of flow-driven adiabatic fast passage inversion of arterial blood. *NMR Biomed.*, 20(8):733–742, 2007.

[102] Luis Hernandez-Garcia, Gregory R. Lee, Alberto L. Vazquez, and Douglas C. Noll. Fast, pseudo-continuous arterial spin labeling for functional imaging using a two-coil system. *Magnetic Resonance in Medicine*, 51(3):577–585, 2004.

[103] Stefan Hetzer, Toralf Mildner, Wolfgang Driesel, Manfred Weder, and Harald E. Möller. Shielded dual-loop resonator for arterial spin labeling at the neck. *Journal of Magnetic Resonance Imaging*, 29(6):1414–1424, 2009.

[104] J. C. Hindman. Proton resonance shift of water in the gas and liquid states. *The Journal of Chemical Physics*, 44(12):4582–4592, 1966.

[105] Andrew B. Holbrook, Juan M. Santos, Elena Kaye, Viola Rieke, and Kim Butts Pauly. Real-time mr thermometry for monitoring hifu ablations of the liver. *Magnetic Resonance in Medicine*, 63(2):365–373, 2010.

[106] D W Holdsworth, C J D Norley, R Frayne, D A Steinman, and B K Rutt. Characterization of common carotid artery blood-flow waveforms in normal human subjects. *Physiological Measurement*, (3):219, 1999.

[107] G. Neil Holland and James R. Macfall. An overview of digital spectrometers for mr imaging. *Journal of Magnetic Resonance Imaging*, 2(2):241–246, 1992.

[108] D. Hoult. The origins and present status of the radio wave controversy in nmr. *Concepts Magn. Reson.*, 34A(4):193–216, 2009.

[109] D. I. Hoult. The principle of reciprocity in signal strength calculations - a mathematical guide. *Concepts Magn. Reson.*, 12(4):173–187, 2000.

[110] D. I. Hoult and R. E. Richards. The signal-to-noise ratio of the nuclear magnetic resonance experiment. *Journal of Magnetic Resonance (1969)*, 24(1):71 – 85, 1976.

[111] David I. Hoult. Sensitivity and power deposition in a high-field imaging experiment. *Journal of Magnetic Resonance Imaging*, 12(1):46–67, 2000.

[112] D.I. Hoult, G. Kolansky, and D. Kripiakevich. A [']hi-fi' cartesian feedback spectrometer for precise quantitation and superior performance. *Journal of Magnetic Resonance*, 171(1):57 – 63, 2004.

[113] D.I. Hoult, G. Kolansky, D. Kripiakevich, and S.B. King. The nmr multi-transmit phased array: a cartesian feedback approach. *Journal of Magnetic Resonance*, 171(1):64 – 70, 2004.

[114] D.I Hoult and Paul C Lauterbur. The sensitivity of the zeugmatographic experiment involving human samples. *Journal of Magnetic Resonance (1969)*, 34(2):425 – 433, 1979.

[115] T S Ibrahim, R Lee, B A Baertlein, and P-M L Robitaille. B 1 field homogeneity and sar calculations for the birdcage coil. *Physics in Medicine and Biology*, 46(2):609, 2001.

[116] Tamer S. Ibrahim, Robert Lee, Brian A. Baertlein, Allahyar Kangarlu, and Pierre-Marie L. Robitaille. Application of finite difference time domain method for the design of birdcage rf head coils using multi-port excitations. *Magnetic Resonance Imaging*, 18(6):733 – 742, 2000.

[117] E. Insko. Mapping of the Radiofrequency Field. *Journal of Magnetic Resonance, Series A*, 103(1):82–85, June 1993.

[118] Yasutoshi Ishihara, Arturo Calderon, Hidehiro Watanabe, Kazuya Okamoto, Yoshinori Suzuki, Kagayaki Kuroda, and Yutaka Suzuki. A precise and fast temperature mapping using water proton chemical shift. *Magnetic Resonance in Medicine*, 34(6):814–823, 1995.

[119] Hesamoddin Jahanian, Douglas C. Noll, and Luis Hernandez-Garcia. B0 field inhomogeneity considerations in pseudo-continuous arterial spin labeling (pcasl): effects on tagging efficiency and correction strategy. *NMR in Biomedicine*, pages n/a–n/a, 2011.

[120] J.M. Jin. *The finite element method in electromagnetics*. A Wiley-Interscience publication. Wiley, 2002.

[121] K M Jones, R V Mulkern, R B Schwartz, K Oshio, P D Barnes, and F A Jolesz. Fast spin-echo mr imaging of the brain and spine: current concepts. *American Journal of Roentgenology*, 158(6):1313–20, 1992.

[122] Youngkyoo Jung, Eric C. Wong, and Thomas T. Liu. Multiphase pseudocontinuous arterial spin labeling (mp-pcasl) for robust quantification of cerebral blood flow. *Magnetic Resonance in Medicine*, 64(3):799–810, 2010.

[123] Alayar Kangarlu, Lin Tang, and Tamer S. Ibrahim. Electric field measurements and computational modeling at ultrahigh-field mri. *Magnetic Resonance Imaging*, 25(8):1222 – 1226, 2007.

[124] U. Katscher, P. Bornert, C. Leussler, and J. S. van den Brink. Transmit sense. *Magn. Reson. Med.*, 49(1):144–150, 2003.

[125] Ulrich Katscher and Peter Börnert. Parallel rf transmission in mri. *NMR in Biomedicine*, 19(3):393–400, 2006.

[126] B. Keong Li, F. Liu, and S. Crozier. Focused, eight-element transceive phased array coil for parallel magnetic resonance imaging of the chest - theoretical considerations. *Magn. Reson. Med.*, 53(6):1251–1257, 2005.

[127] A. B. Kerr, W. Grissom, P. Stang, G. C. Scott, and J. Pauly. B1 mapping and parallel excitation using vector decoupling. In *Proc. Intl. Soc. Mag. Reson. 17*, number 2618, 2009.

[128] Seymour S. Kety and Carl F. Schmidt. The nitrous oxide method for the quantitative determination of cerebral blood flow in man: Theory, procedure and normal values 1. *The Journal of Clinical Investigation*, 27(4):476–483, 7 1948.

[129] A. Kickhefel, Jörg Roland, Clifford Weiss, and Fritz Schick. Accuracy of real-time mr temperature mapping in the brain: A comparison of fast sequences. *Physica Medica*, 26(4):192 – 201, 2010.

[130] S.-G. Kim. Quantification of relative cerebral blood flow change by flow-sensitive alternating inversion recovery (fair) technique: Application to functional mapping. *Magn. Reson. Med.*, 34(3):293–301, 1995.

[131] M. Kowalski, B. Behnia, A.G. Webb, and Jian-Ming Jin. Optimization of electromagnetic phased-arrays for hyperthermia via magnetic resonance temperature estimation. *Biomedical Engineering, IEEE Transactions on*, 49(11):1229 –1241, nov. 2002.

[132] H.L. Krauss, C.W. Bostian, and F.H. Raab. *Solid state radio engineering*. Wiley, 1980.

[133] K. Kuroda, Y. Miki, N. Nakagawa, S. Tsutsumi, Y. Ishihara, Y. Suzuki, and K. Sato. Non-invasive temperature measurement by means of nmr parameters: use of proton chemical shift with spectral estimation technique. *Med Biol Eng Comput*, 29:902, 1991.

[134] Kagayaki Kuroda, Yutaka Suzuki, Yasutoshi Ishihara, Kazuya Okamoto, and Yoshinori Suzuki. Temperature mapping using water proton chemical shift obtained with 3d-mrsi: Feasibility in vivo. *Magnetic Resonance in Medicine*, 35(1):20–29, 1996.

[135] Krishna N. Kurpad, Steven M. Wright, and Eddy B. Boskamp. Rf current element design for independent control of current amplitude and phase in transmit phased arrays. *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering*, 29B(2):75–83, 2006.

[136] K K Kwong, J W Belliveau, D A Chesler, I E Goldberg, R M Weisskoff, B P Poncelet, D N Kennedy, B E Hoppel, M S Cohen, and R Turner. Dynamic magnetic resonance imaging of human brain activity during primary sensory stimulation. *Proceedings of the National Academy of Sciences*, 89(12):5675–5679, 1992.

[137] G. R. Lee, Hernandez-Garcia, and D. C. L., Noll. A phantom for quantitative spin tagging perfusion measurements. In *Proc. Intl. Soc. Mag. Reson. 10*, number 1057, 2002.

[138] Ray F. Lee, Randy O. Giaquinto, and Christopher J. Hardy. Coupling and decoupling theory and its application to the mri phased array. *Magnetic Resonance in Medicine*, 48(1):203–213, 2002.

[139] CJ Lewa and Z. Majewska. Temperature relationships of proton spin-lattice relaxation time t1 in biological tissues. *Bulletin du cancer*, 67(5):525, 1980.

[140] Z.P. Liang, P.C. Lauterbur, IEEE Engineering in Medicine, and Biology Society. *Principles of magnetic resonance imaging: a signal processing perspective*. IEEE Press series in biomedical engineering. SPIE Optical Engineering Press, 2000.

[141] David S. Liebeskind. Collaterals in acute stroke: Beyond the clot. *Neuroimaging Clin N Am*, 15(3):553–573, August 2005.

[142] Thorsten Liebig. openems - open electromagnetic field solver.

[143] Jae S. Lim. *Two-dimensional signal and image processing*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1990.

[144] Christine H. Lorenz, David R. Pickens, Donald B. Puffer, and Ronald R. Price. Magnetic resonance diffusion/perfusion phantom experiments. *Magnetic Resonance in Medicine*, 19(2):254–260, 1991.

[145] Wen-Ming Luh, Eric C. Wong, Peter A. Bandettini, and James S. Hyde. Quipss ii with thin-slice ti1 periodic saturation: A method for improving accuracy of quantitative perfusion imaging using pulsed arterial spin labeling. *Magnetic Resonance in Medicine*, 41(6):1246–1254, 1999.

[146] Antoine Lutti, Chloe Hutton, Jurgen Finsterbusch, Gunther Helms, and Nikolaus Weiskopf. Optimization and validation of methods for mapping of the radiofrequency transmit field at 3t. *Magn. Reson. Med.*, 64(1):229–238, 2010.

[147] L. Maccotta, J. A. Detre, and D. C. Alsop. The efficiency of adiabatic inversion for perfusion imaging by arterial spin labeling. *NMR Biomed.*, 10(4-5):216–221, 1997.

[148] Shaihan J. Malik, David J. Larkman, and Joseph V. Hajnal. Optimal linear combinations of array elements for b1 mapping. *Magnetic Resonance in Medicine*, 62(4):902–909, 2009.

[149] Shaihan J. Malik, David J. Larkman, Declan P. O'Regan, and Joseph V. Hajnal. Subject-specific water-selective imaging using parallel transmission. *Magnetic Resonance in Medicine*, 63(4):988–997, 2010.

[150] Sanjay Mani, John Pauly, Steven Conolly, Craig Meyer, and Dwight Nishimura. Background suppression with multiple inversion recovery nulling: Applications to projective angiography. *Magn. Reson. Med.*, 37(6):898–905, 1997.

[151] R.J. Marks. *Handbook of Fourier analysis & its applications*. Oxford University Press, 2009.

[152] Reiko Matsumoto, Robert V. Mulkern, Stephen G. Hushek, and Ferenc A. Jolesz. Tissue temperature monitoring for thermal interventional therapy: Comparison of t1-weighted mr sequences. *Journal of Magnetic Resonance Imaging*, 4(1):65–70, 1994.

[153] D.W. McRobbie, E.A. Moore, and M.J. Graves. *MRI from picture to proton*. Cambridge University Press, 2007.

[154] Paul Mellor and David Checkley. Active coil isolation in nmr imaging and spectroscopy using pin diodes and tuned transmission line: a practical approach. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 3:35–40, 1995. 10.1007/BF02426398.

[155] Gregory J. Metzger, Carl Snyder, Can Akgun, Tommy Vaughan, Kamil Ugurbil, and Pierre-Francois Van de Moortele. Local b1+ shimming for prostate imaging with transceiver arrays at 7t based on subject-dependent transmit phase measurements. *Magn. Reson. Med.*, 59(2):396–409, 2008.

[156] T. Mildner, R. Trampel, H. E. Moller, A. Schafer, C. J. Wiggins, and D. G. Norris. Functional perfusion imaging using continuous arterial spin labeling with separate labeling and imaging coils at 3 t. *Magn. Reson. Med.*, 49(5):791–795, 2003.

[157] Toralf Mildner, Harald E. Möller, Wolfgang Driesel, David G. Norris, and Robert Trampel. Continuous arterial spin labeling at the human common carotid artery: the influence of transit times. *NMR in Biomedicine*, 18(1):19–23, 2005.

[158] J. Mispelter, M. Lupu, and A. Briguet. *NMR probeheads for biophysical and biomedical experiments: theoretical principles & practical guidelines.* Imperial College Pr, 2006.

[159] Bradford A. Moffat, Thomas L. Chenevert, Daniel E. Hall, Alnawaz Rehemtulla, and Brian D. Ross. Continuous arterial spin labeling using a train of adiabatic inversion pulses. *Journal of Magnetic Resonance Imaging*, 21(3):290–296, 2005.

[160] C. Müller. *Spherical harmonics.* Lecture notes in mathematics. Springer-Verlag, 1966.

[161] K. Nehrke and P. Boernert. Improved b1-mapping for multi rf transmit systems. In *Proc. Intl. Soc. Mag. Reson. 16*, page 353, 2008.

[162] Kay Nehrke. On the steady-state properties of actual flip angle imaging (afi). *Magn. Reson. Med.*, 61(1):84–92, 2009.

[163] Tomoyuki Noguchi, Takashi Yoshiura, Akio Hiwatashi, Osamu Togao, Koji Yamashita, Kouji Kobayashi, Futoshi Mihara, and Hiroshi Honda. Quantitative perfusion imaging with pulsed arterial spin labeling: A phantom study. *Magn. Reson. in Med. Sci.*, 6(2):91–97, 2007.

[164] S Ogawa, D W Tank, R Menon, J M Ellermann, S G Kim, H Merkle, and K Ugurbil. Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. *Proceedings of the National Academy of Sciences*, 89(13):5951–5955, 1992.

[165] R.J. Ogg, R.B. Kingsley, and J.S. Taylor. Wet, a t1- and b1-insensitive water-suppression method for in vivo localized 1h nmr spectroscopy. *Journal of Magnetic Resonance, Series B*, 104(1):1 – 10, 1994.

[166] Ruth L. O'Gorman, Paul E. Summers, Fernando O. Zelaya, Steven C. R. Williams, David C. Alsop, and David J. Lythgoe. In vivo estimation of the flow-driven adiabatic inversion efficiency for continuous arterial spin labeling: A method using phase contrast magnetic resonance angiography. *Magnetic Resonance in Medicine*, 55(6):1291–1297, 2006.

[167] Sukhoon Oh, C.A. Roopnariane, M.-R. Tofighi, and C.M. Collins. Mri-based temperature and sar mapping with a new dual-coil solenoid/birdcage heating/measurement system. In *Radio and Wireless Symposium (RWS), 2010 IEEE*, pages 520 –522, jan. 2010.

[168] Sukhoon Oh, Andrew G. Webb, Thomas Neuberger, BuSik Park, and Christopher M. Collins. Experimental and numerical assessment of mri-induced temperature change and sar distributions in phantoms and in vivo. *Magnetic Resonance in Medicine*, 63(1):218–223, 2010.

[169] C. C. Olson, J. T. Vaughan, and A. Gopinath. Radio frequency field localization for magnetic resonance, Dec 2009.

[170] Johan Olsrud, Ronnie Wirestam, Sara Brockstedt, Annika M K Nilsson, Karl-Goran Tranberg, Freddy Stahlberg, and Bertil R R Persson. Mri thermometry in phantoms by use of the proton resonance frequency shift method: application to interstitial laser thermotherapy. *Physics in Medicine and Biology*, 43(9):2597, 1998.

[171] R. J. Ordidge, P. Mansfield, and R. E. Coupland. Rapid biomedical imaging by nmr. *British Journal of Radiology*, 54(646):850–855, 1981.

[172] Roger J. Ordidge, Marzena Wylezinska, James W. Hugg, Edward Butterworth, and Florence Franconi. Frequency offset corrected inversion (foci) pulses for use in localized spectroscopy. *Magnetic Resonance in Medicine*, 36(4):562–566, 1996.

[173] Koichi Oshio and David A. Feinberg. Grase (gradient-and spin-echo) imaging: A novel fast mri technique. *Magn. Reson. Med.*, 20(2):344–349, 1991.

[174] Ardavan F. Oskooi, David Roundy, Mihai Ibanescu, Peter Bermel, J. D. Joannopoulos, and Steven G. Johnson. MEEP: A flexible free-software package for electromagnetic simulations by the FDTD method. *Computer Physics Communications*, 181:687–702, January 2010.

[175] F. Padormo, S. J. Malik, and J. V. Hajnal. Direct reconstruction of b1 maps from undersampled acquisitions. In *Proc. Intl. Soc. Mag. Reson. 18*, 2010.

[176] Fernando F. Paiva, Alberto Tannús, S. Lalith Talagala, and Afonso C. Silva. Arterial spin labeling of cerebral perfusion territories using a separate labeling coil. *Journal of Magnetic Resonance Imaging*, 27(5):970–977, 2008.

[177] R. E. A. C. Paley. "on orthogonal matrices.". *J. Math. Phys.*, 12:311–320, 1933.

[178] Nikos G. Papadakis, Ian M. Devonshire, Jiabao He, and John E.W. Mayhew. Continuous arterial spin labeling with separate labeling and imaging coils: Implementation using a single rf channel and amplifier. *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering*, 29B(3):145–152, 2006.

[179] Dennis L. Parker, Vernon Smith, Philip Sheldon, Lawrence E. Crooks, and Lauren Fussell. Temperature distribution measurements in two-dimensional nmr imaging. *Medical Physics*, 10(3):321–325, 1983.

[180] E T Petersen, I Zimine, Y-C L Ho, and X Golay. Non-invasive measurement of perfusion: a critical review of arterial spin labelling techniques. *Br J Radiol*, 79(944):688–701, 2006.

[181] Esben Thade Petersen, Tchoyoson Lim, and Xavier Golay. Model-free arterial spin labeling quantification approach for perfusion mri. *Magnetic Resonance in Medicine*, 55(2):219–232, 2006.

[182] Esben Thade Petersen, Kim Mouridsen, and Xavier Golay. The quasar reproducibility study, part ii: Results from a multi-center arterial spin labeling test-retest study. *NeuroImage*, 49(1):104–113, 2010.

[183] Robert G. Pinkerton, Enzo A. Barberi, and Ravi S. Menon. Noise properties of a nmr transceiver coil array. *Journal of Magnetic Resonance*, 171(1):151 – 156, 2004.

[184] J. De Poorter, C. Dewagter, Y. Dedeene, C. Thomsen, F. Stahlberg, and E. Achten. The proton-resonance-frequency-shift method compared with molecular diffusion for quantitative measurement of two-dimensional time-dependent temperature distribution in a phantom. *Journal of Magnetic Resonance, Series B*, 103(3):234 – 241, 1994.

[185] John De Poorter, Carlos De Wagter, Yves De Deene, Carsten Thomsen, Freddy Ståhlberg, and Eric Achten. Noninvasive mri thermometry with the proton resonance frequency (prf) method: In vivo results in human muscle. *Magnetic Resonance in Medicine*, 33(1):74–81, 1995.

[186] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger. Sense: Sensitivity encoding for fast mri. *Magn. Reson. Med.*, 42(5):952–962, 1999.

[187] E.M. Purcell. *Electricity and magnetism*, volume 34. 1966.

[188] Bensheng Qiu, Abdel-Monem El-Sharkawy, Vaishali Paliwal, Parag Karmarkar, Fabao Gao, Ergin Atalar, and Xiaoming Yang. Simultaneous radiofrequency (rf) heating and magnetic resonance (mr) thermal mapping using an intravascular mr imaging/rf heating system. *Magnetic Resonance in Medicine*, 54(1):226–230, 2005.

[189] Bruno Quesson, Jacco A. de Zwart, and Chrit T.W. Moonen. Magnetic resonance temperature imaging for guidance of thermotherapy. *Journal of Magnetic Resonance Imaging*, 12(4):525–533, 2000.

[190] K A Rempp, G Brix, F Wenz, C R Becker, F Gückel, and W J Lorenz. Quantification of regional cerebral blood flow and volume with dynamic susceptibility contrast-enhanced mr imaging. *Radiology*, 193(3):637–641, 1994.

[191] Viola Rieke, Karl K. Vigen, Graham Sommer, Bruce L. Daniel, John M. Pauly, and Kim Butts. Referenceless prf shift thermometry. *Magnetic Resonance in Medicine*, 51(6):1223–1231, 2004.

[192] D. A. Roberts, L. Bolinger, J. A. Detre, E. K. Insko, P. Berge, and J. S. Leigh. Continuous inversion angiography. *Magn. Reson. Med.*, 29(5):631–636, 1993.

[193] D A Roberts, J A Detre, L Bolinger, E K Insko, and J S Leigh. Quantitative magnetic resonance imaging of human brain perfusion at 1.5 t using steady-state inversion of arterial water. *Proceedings of the National Academy of Sciences of the United States of America*, 91(1):33–37, 1994.

[194] P. B. Roemer, W. A. Edelstein, C. E. Hayes, S. P. Souza, and O. M. Mueller. The nmr phased array. *Magn. Reson. Med.*, 16(2):192–225, 1990.

[195] R. R. Ruan, K. Chang, P. L. Chen, and A. Ning. Simultaneous heat and moisture transfer in cheddar cheese during cooling. ii. mri temperature mapping. *Drying Technology: An International Journal*, 16(7):1459–1470, 1998.

[196] D.A. Sánchez-Hernández. *High frequency electromagnetic dosimetry*. Artech House electromagnetic analysis series. Artech House, 2009.

[197] A. Sasao, T. Hirai, S. Nishimura, H. Fukuoka, R. Murakami, M. Kitajima, T. Okuda, M. Akter, M. Morioka, S. Yano, H. Nakamura, K. Makino, J.-i. Kuratsu, K. Awai, and Y. Yamashita. Assessment of vascular supply of hypervascular extra-axial brain tumors with 3t mr regional perfusion imaging. *American Journal of Neuroradiology*, 31(3):554–558, 2010.

[198] Ron Schmitt. *Electromagnetics Explained: A Handbook for Wireless/RF, EMC, and High-Speed Electronics*. Butterworth-Heinemann, Newton, MA, USA, 2001.

[199] G. C. Scott, P. Stang, A. Kerr, and J. Pauly. General signal vector decoupling for transmit arrays. In *Proc. Intl. Soc. Mag. Reson. 16*, number 146, 2008.

[200] D. Seeber, J. Jevtic, and A. Menon. Floating shield current suppression trap. *Concepts Magn. Reson.*, 21B(1):26–31, 2004.

[201] Frank Seifert, Gerd Wübbeler, Sven Junge, Bernd Ittermann, and Herbert Rinneberg. Patient safety concept for multichannel transmit coils. *Journal of Magnetic Resonance Imaging*, 26(5):1315–1321, 2007.

[202] Erik Shapiro, Arijitt Borthakur, and Ravinder Reddy. Mr imaging of rf heating using a paramagnetic doped agarose phantom. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 10:114–121, 2000. 10.1007/BF02601846.

[203] Erik M. Shapiro, Arijitt Borthakur, Michael J. Shapiro, Ravinder Reddy, and John S. Leigh. Fast mri of rf heating via phase difference mapping. *Magnetic Resonance in Medicine*, 47(3):492–498, 2002.

[204] A. C. Silva, W. Zhang, D. S. Williams, and A. P. Koretsky. Multi-slice mri of rat brain perfusion during amphetamine stimulation using arterial spin labeling. *Magnetic Resonance Medicine*, 33(2):209–214, 1995.

[205] R. Simons. *Coplanar waveguide circuits, components, and systems*. Wiley series in microwave and optical engineering. Wiley-Interscience, 2001.

[206] John G. Sled and G. Bruce Pike. Correction for b1 and b0 variations in quantitative t2 measurements using mri. *Magnetic Resonance in Medicine*, 43(4):589–593, 2000.

[207] W.R. Smythe. *Static and Dynamic Electricity*. McGraw-Hill Inc, 1950.

[208] R. Jason Stafford, Roger E. Price DVM, Chris J. Diederich, Marko Kangasniemi, Lars E. Olsson, and John D. Hazle. Interleaved echo-planar imaging for fast multiplanar magnetic resonance temperature imaging of ultrasound thermal ablation therapy. *Journal of Magnetic Resonance Imaging*, 20(4):706–714, 2004.

[209] Susan. Standring and Henry Gray. *Gray's anatomy : the anatomical basis of clinical practice / editor-in-chief, Susan Standring*. Churchill Livingstone, Edinburgh :, 40th ed. edition, 2008.

[210] P. Stang, S. Conolly, and G. Scott. MEDUSA: A Scalable MR Console for Parallel Imaging. In *Proc. Intl. Soc. Mag. Reson. 15*, number 925, 2007.

[211] Greg J. Stanisz, Ewa E. Odrobina, Joseph Pun, Michael Escaravage, Simon J. Graham, Michael J. Bronskill, and R. Mark Henkelman. T1, t2 relaxation and magnetization transfer in tissue at 3t. *Magnetic Resonance in Medicine*, 54(3):507–512, 2005.

[212] A. Taflove and S.C. Hagness. *Computational electrodynamics: the finite-difference time-domain method*. Artech House antennas and propagation library. Artech House, 2005.

[213] S. Lalith Talagala, Frank Q. Ye, Patrick J. Ledden, and Scott Chesnick. Whole-brain 3d perfusion mri at 3.0 t using casl with a separate labeling coil. *Magnetic Resonance in Medicine*, 52(1):131–140, 2004.

[214] Toshiaki Taoka, Satoru Iwasaki, Hiroyuki Nakagawa, Akio Fukusumi, Shinji Hirohashi, Masahiko Sakamoto, Kimihiko Kichikawa, and Katsutoshi Murata. Distinguishing between anterior cerebral artery and middle cerebral artery perfusion by color-coded perfusion direction mapping with arterial spin labeling. *American Journal of Neuroradiology*, 25(2):248–251, 2004.

[215] David L. Thomas, Enrico De Vita, Ralf Deichmann, Robert Turner, and Roger J. Ordidge. 3d mdeft imaging of the human brain at 4.7 t with reduced sensitivity to radiofrequency inhomogeneity. *Magnetic Resonance in Medicine*, 53(6):1452–1458, 2005.

[216] David L. Thomas, Enrico De Vita, Steven Roberts, Robert Turner, Tarek A. Yousry, and Roger J. Ordidge. High-resolution fast spin echo imaging of the human brain at 4.7 t: Implementation and sequence characteristics. *Magnetic Resonance in Medicine*, 51(6):1254–1264, 2004.

[217] Nick Todd, Urvi Vyas, Josh de Bever, Allison Payne, and Dennis L. Parker. Reconstruction of fully three-dimensional high spatial and temporal resolution mr temperature maps for retrospective applications. *Magnetic Resonance in Medicine*, pages n/a–n/a, 2011.

[218] J. F. Utting, D. L. Thomas, D. G. Gadian, and R. J. Ordidge. Velocity-driven adiabatic fast passage for arterial spin labeling: Results from a computer model. *Magn. Reson. Med.*, 49(2):398–401, 2003.

[219] Pierre-François Van de Moortele, Can Akgun, Gregor Adriany, Steen Moeller, Johannes Ritter, Christopher M. Collins, Michael B. Smith, J. Thomas Vaughan, and Kamil Ugurbil. B1 destructive interferences and spatial phase patterns at 7 t with a head transceiver array coil. *Magnetic Resonance in Medicine*, 54(6):1503–1518, 2005.

[220] Peter Jan van Laar, Jeroen Hendrikse, Xavier Golay, Hanzhang Lu, Matthias J.P. van Osch, and Jeroen van der Grond. In vivo flow territory mapping of major brain feeding arteries. *NeuroImage*, 29(1):136 – 144, 2006.

[221] Peter Jan van Laar, Jeroen van der Grond, and Jeroen Hendrikse. Brain perfusion territory imaging: Methods and clinical applications of selective arterial spin-labeling mr imaging1. *Radiology*, 246(2):354–364, February 2008.

[222] André Vander Vorst, Arye Rosen, and Youji Kotsuka. *RF/Microwave Interaction with Biological Tissues*. John Wiley & Sons, Inc., 2006.

[223] J. Thomas Vaughan, Hoby P. Hetherington, Joe O. Otu, Jullie W. Pan, and Gerald M. Pohost. High frequency volume coils for clinical nmr imaging and spectroscopy. *Magnetic Resonance in Medicine*, 32(2):206–218, 1994.

[224] J.T. Vaughan, M. Garwood, C.M. Collins, W. Liu, L. DelaBarre, G. Adriany, P. Andersen, H. Merkle, R. Goebel, M.B. Smith, and K. Ugurbil. 7t vs. 4t: Rf power, homogeneity, and signal-to-noise comparison in head images. *Magnetic Resonance in Medicine*, 46(1):24–30, 2001.

[225] P. Vernickel and I. Graesslin. Active decoupling of transmitters in mri, January 5 2010. US Patent 7,642,782.

[226] Ze Wang, Jiongjiong Wang, Thomas J. Connick, Gabriel S. Wetmore, and John A. Detre. Continuous asl (casl) perfusion mri with an array coil and parallel imaging at 3t. *Magnetic Resonance in Medicine*, 54(3):732–737, 2005.

[227] Carsten Warmuth, Stefan Nagel, Oliver Hegemann, Waldemar Wlodarczyk, and Lutz Lüdemann. Accuracy of blood flow values determined by arterial spin labeling: A validation study in isolated porcine kidneys. *Journal of Magnetic Resonance Imaging*, 26(2):353–358, 2007.

[228] Andrew G. Webb and Christopher M. Collins. Parallel transmit and receive technology in high-field magnetic resonance neuroimaging. *International Journal of Imaging Systems and Technology*, 20(1):2–13, 2010.

[229] Claudia Weidensteiner, Bruno Quesson, Bénédicte Caire-Gana, Noureddine Kerioui, Anne Rullier, Hervé Trillaud, and Chrit T.W. Moonen. Real-time mr temperature mapping of rabbit liver in vivo during thermal ablation. *Magnetic Resonance in Medicine*, 50(2):322–330, 2003.

[230] Steffen Weiss, Bernd David, Oliver Lips, Jan Hendrik Wuelbern, and Sascha Krueger. Comparison of rf heating in cables equipped with different types of current limitations. In *Proc. Intl. Soc. Mag. Reson. 19*, number 3773, 2011.

[231] Steffen Weiss, Peter Vernickel, Tobias Schaeffter, Volkmar Schulz, and Bernhard Gleich. Transmission line for improved rf safety of interventional devices. *Magnetic Resonance in Medicine*, 54(1):182–189, 2005.

[232] R. Werner, D. G. Norris, K. Alfke, H. M. Mehdorn, and O. Jansen. Continuous artery-selective spin labeling (cassl). *Magn. Reson. Med.*, 53(5):1006–1012, 2005.

[233] Richard Werner, Karsten Alfke, Tobias Schaeffter, Arya Nabavi, H. Maximilian Mehdorn, and Olav Jansen. Brain perfusion territory imaging applying oblique-plane arterial spin labeling with a standard send/receive head coil. *Magnetic Resonance in Medicine*, 52(6):1443–1447, 2004.

[234] M.N. Wernick and J.N. Aarsvold. *Emission tomography: the fundamentals of PET and SPECT*. The Fundamentals of Pet and Spect Series. Elsevier Academic Press, 2004.

[235] F.M. White. *Fluid Mechanics, WCB*. McGraw-Hill, New York, 1999.

[236] H. Whiteside and R. King. The loop antenna as a probe. *Antennas and Propagation, IEEE Transactions on*, 12(3):291 – 297, may. 1964.

[237] D S Williams, J A Detre, J S Leigh, and A P Koretsky. Magnetic resonance imaging of perfusion using spin inversion of arterial water. *Proceedings of the National Academy of Sciences of the United States of America*, 89(1):212–216, 1992.

[238] Ronald L. Wolf and John A. Detre. Clinical neuroimaging using arterial spin-labeled perfusion magnetic resonance imaging. *Neurotherapeutics*, 4(3):346 – 359, 2007. <ce:title>Advances in Neuroimaging/Neuroethics</ce:title>.

[239] Ronald L. Wolf, Jiongjiong Wang, Sumei Wang, Elias R. Melhem, Donald M. O'Rourke, Kevin D. Judy, and John A. Detre. Grading of cns neoplasms using continuous arterial spin labeled perfusion mr imaging at 3 tesla. *Journal of Magnetic Resonance Imaging*, 22(4):475–482, 2005.

[240] Steven D. Wolff and Robert S. Balaban. Magnetization transfer contrast (mtc) and tissue water proton relaxation in vivo. *Magnetic Resonance in Medicine*, 10(1):135–144, 1989.

[241] E. C. Wong, R. B. Buxton, and L. R. Frank. A theoretical and experimental comparison of continuous and pulsed arterial spin labeling techniques for quantitative perfusion imaging. *Magn. Reson. Med.*, 40(3):348–355, 1998.

[242] Eric C. Wong. Vessel-encoded arterial spin-labeling using pseudocontinuous tagging. *Magnetic Resonance in Medicine*, 58(6):1086–1091, 2007.

[243] Eric C. Wong, Richard B. Buxton, and Lawrence R. Frank. Quantitative imaging of perfusion using a single subtraction (quipss and quipss ii). *Magnetic Resonance in Medicine*, 39(5):702–708, 1998.

[244] S. M. Wright. Full-wave analysis of planar radiofrequency coils and coil arrays with assumed current distribution. *Concepts in Magnetic Resonance*, 15(1):2–14, 2002.

[245] Bing Wu, Chunsheng Wang, R. Krug, D.A. Kelley, Duan Xu, Yong Pang, S. Banerjee, D.B. Vigneron, S.J. Nelson, S. Majumdar, and Xiaoliang Zhang. 7t human spine imaging arrays with adjustable inductive decoupling. *Biomedical Engineering, IEEE Transactions on*, 57(2):397 –403, feb. 2010.

[246] Wen-Chau Wu, María Fernández-Seara, John A. Detre, Felix W. Wehrli, and Jiongjiong Wang. A theoretical and experimental investigation of the tagging efficiency of pseudocontinuous arterial spin labeling. *Magnetic Resonance in Medicine*, 58(5):1020–1027, 2007.

[247] Qin Xu, Christopher Glielmi, Lei Zhou, Kisueng Choi, and Xiaoping Hu. An inexpensive and programmable rf transmitter setup for two-coil casl. *Concepts in Magnetic Resonance Part B: Magnetic Resonance Engineering*, 33B(4):228–235, 2008.

[248] Vasily L. Yarnykh. Actual flip-angle imaging in the pulsed steady state: A method for rapid three-dimensional mapping of the transmitted radiofrequency field. *Magn. Reson. Med.*, 57(1):192–200, 2007.

[249] P. Yazdanbakhsh and K. Solbach. Active phased array techniques for high-field mri. In *Antennas and Propagation, 2009. EuCAP 2009. 3rd European Conference on*, pages 1367 –1371, march 2009.

[250] Frank Q. Ye, Joseph A. Frank, Daniel R. Weinberger, and Alan C. McLaughlin. Noise reduction in 3d perfusion imaging by attenuating the static signal in arterial spin tagging (assist). *Magn. Reson. Med.*, 44(1):92–100, 2000.

[251] Xiaofei Ye, Roger Ruan, Paul Chen, Kehua Chang, Ke Ning, Irwin A. Taub, and Chris Doona. Accurate and fast temperature mapping during ohmic heating using proton resonance frequency shift mri thermometry. *Journal of Food Engineering*, 59(2-3):143 – 150, 2003.

[252] Kane Yee. Numerical solution of initial boundary value problems involving maxwell's equations in isotropic media. *Antennas and Propagation, IEEE Transactions on*, 14(3):302 –307, may 1966.

[253] H. Yonas. *Cerebral blood flow measurement with stable xenon-enhanced computed tomography.* Raven Press, 1992.

[254] Martin N. Yongbi, Craig A. Branch, and Joseph A. Helpern. Perfusion imaging using foci rf pulses. *Magnetic Resonance in Medicine*, 40(6):938–943, 1998.

[255] D. Yoon, H. Jananian, D.C. Noll, and L. Hernandez-Garcia. Spatially selective pcasl with parallel excitation. In *Proc. Intl. Soc. Mag. Reson. 19*, 2011.

[256] B. Yordanov, K. Doughty, and R. Yordanov. Near-field probes for EMC applications. *EMC Test and Design*, pages 18–26, 1994.

[257] Chun yu Yip, Jeffrey A. Fessler, and Douglas C. Noll. Iterative rf pulse design for multidimensional, small-tip-angle selective excitation. *Magnetic Resonance in Medicine*, 54(4):908–917, 2005.

[258] Greg Zaharchuk, Huy M. Do, Michael P. Marks, Jarrett Rosenberg, Michael E. Moseley, and Gary K. Steinberg. Arterial spin-labeling mri can identify the presence and intensity of collateral perfusion in patients with moyamoya disease. *Stroke*, 42(9):2485–2491, 2011.

[259] Greg Zaharchuk, Patrick J. Ledden, Kenneth K. Kwong, Timothy G. Reese, Bruce R. Rosen, and Lawrence L. Wald. Multislice perfusion and perfusion territory imaging in humans with separate label and image coils. *Magnetic Resonance in Medicine*, 41(6):1093–1098, 1999.

[260] Weiguo Zhang, Afonso C. Silva, Donald S. Williams, and Alan P. Koretsky. Nmr measurement of perfusion using arterial spin labeling without saturation of macromolecular spins. *Magnetic Resonance in Medicine*, 33(3):370–376, 1995.

[261] Weiguo Zhang, Donald S. Williams, John A. Detre, and Alan P. Koretsky. Measurement of brain perfusion by volume-localized nmr spectroscopy using inversion of arterial water spins: Accounting for transit time and cross-relaxation. *Magnetic Resonance in Medicine*, 25(2):362–371, 1992.

[262] AI. Zhernovoi. Fast adiabatic passage in nuclear magnetic resonance. *Soviet Physics-Solid State*, 9:523–524, 1967.

[263] Ivan Zimine, Esben Thade Petersen, and Xavier Golay. Dual vessel arterial spin labeling scheme for regional perfusion imaging. *Magnetic Resonance in Medicine*, 56(5):1140–1144, 2006.

[264] Y. Zur, M. L. Wood, and L. J. Neuringer. Spoiling of transverse magnetization in steady-state sequences. *Magnetic Resonance in Medicine*, 21(2):251–263, 1991.

Appendix A

RF Safety Results

A.1 SAR Maps

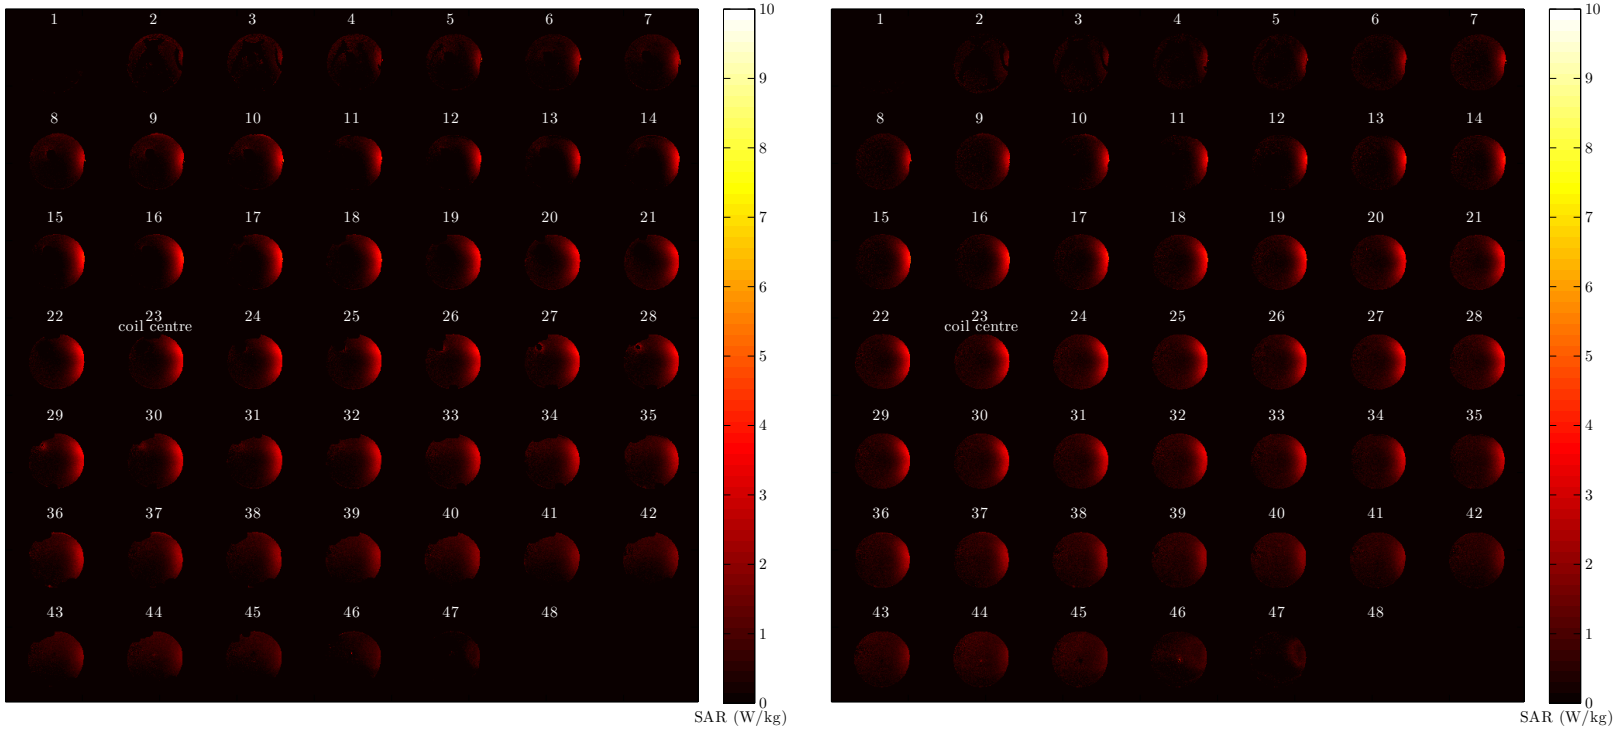

a. run 1 b. run 2

Figure A.1: SAR Maps for coil 1. Colour scaled 0-10 W/kg

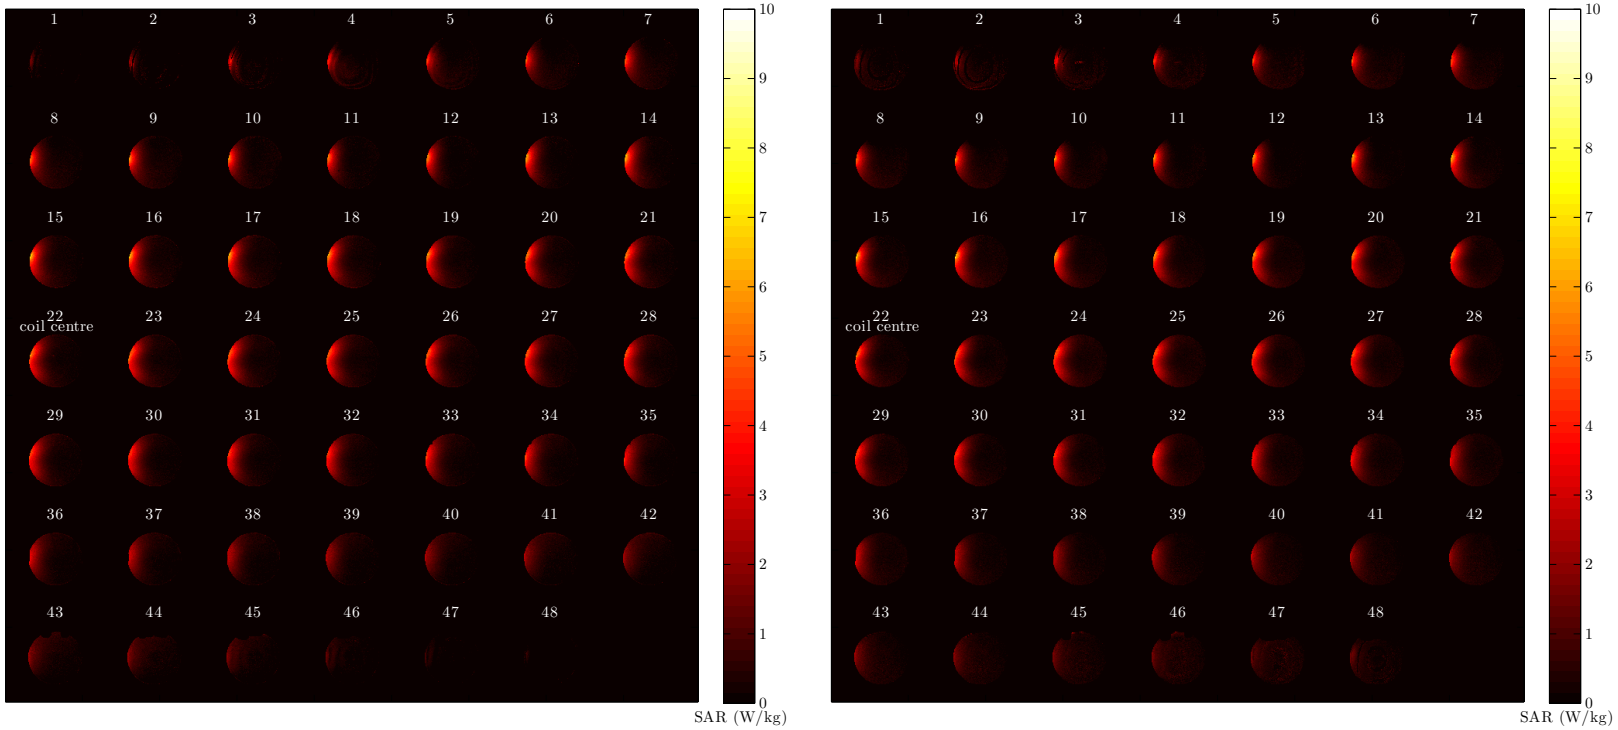

a. run 1

b. run 2

Figure A.2: SAR Maps for coil 2. Colour scaled 0-10 W/kg

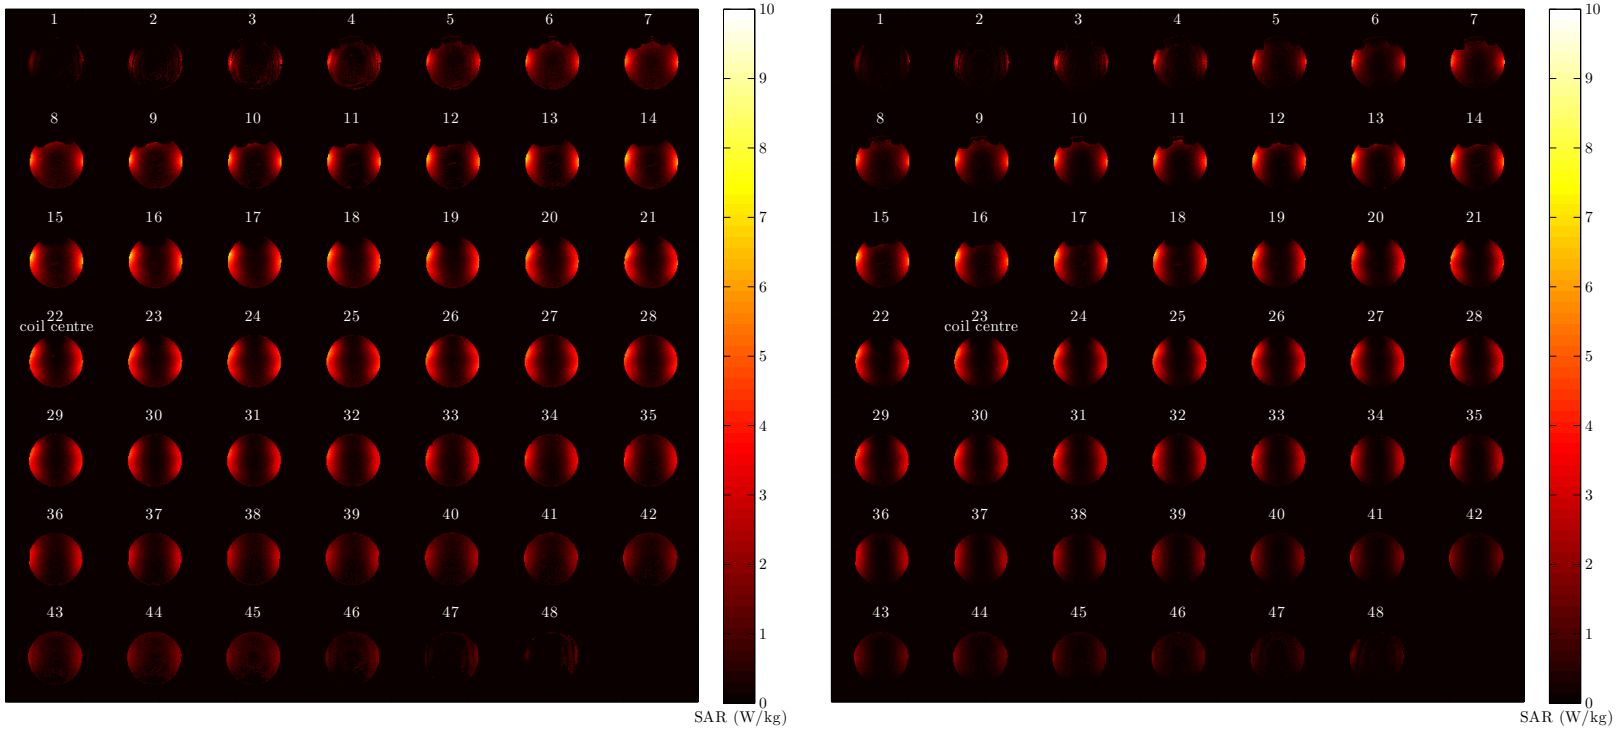

a. run 1

b. run 2

Figure A.3: SAR Maps for Anti-Phase configuration. Colour scaled 0-10 W/kg

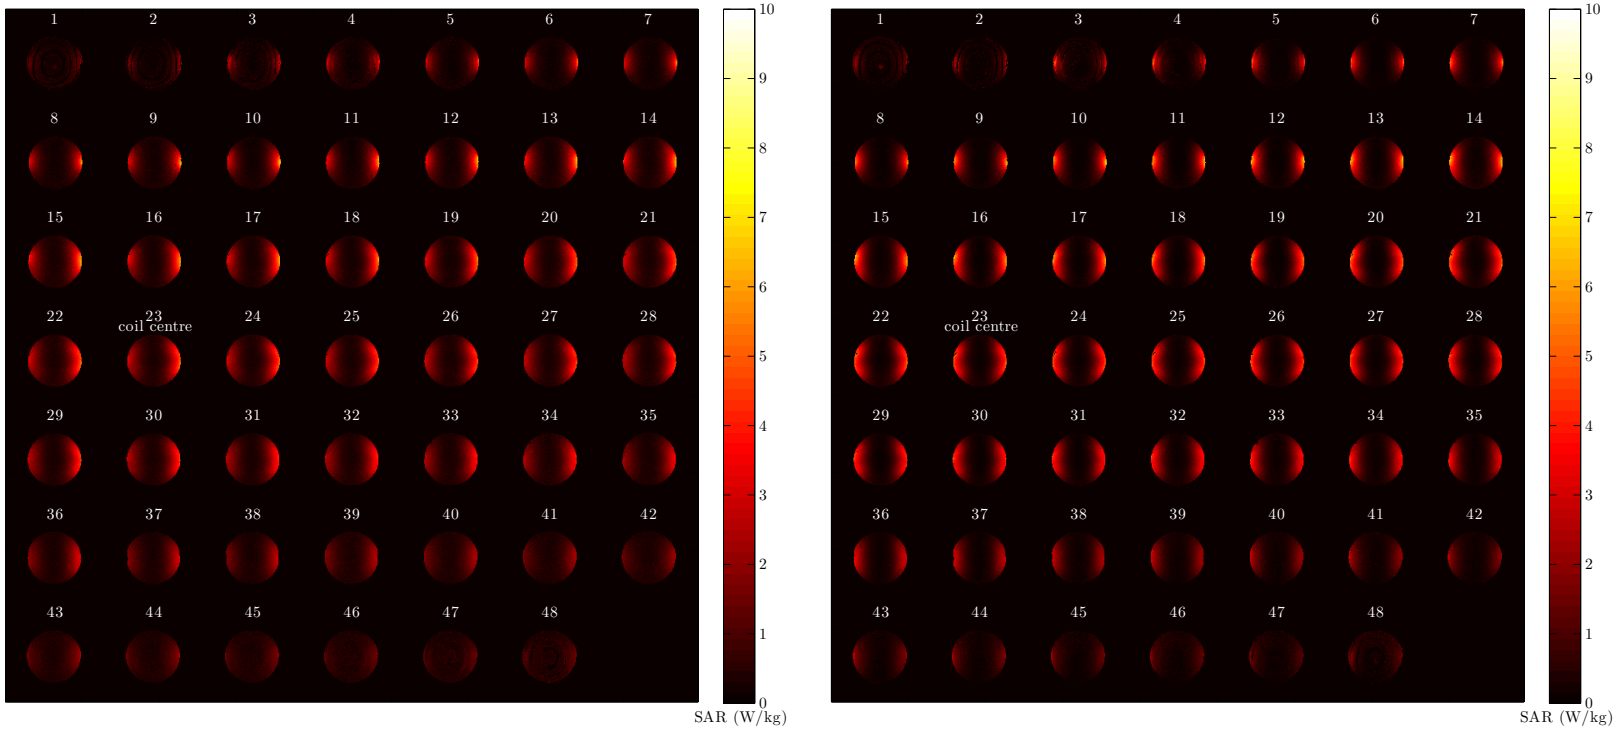

a. run 1

b. run 2

Figure A.4: SAR Maps for quadrature configuration. Colour scaled 0-10 W/kg

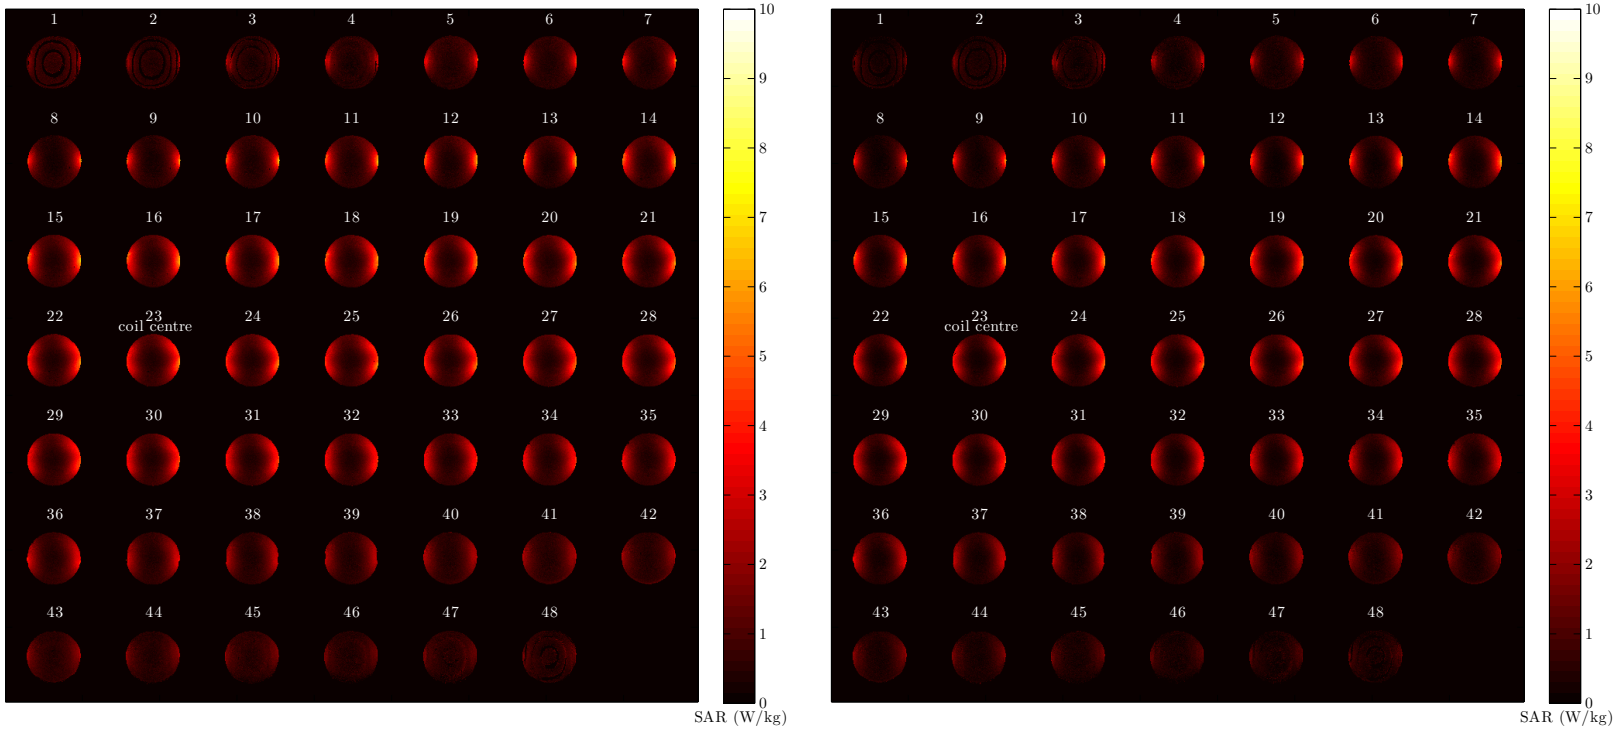

a. run 1

b. run 2

Figure A.5: SAR Maps for In-Phase configuration. Colour scaled 0-10 W/kg

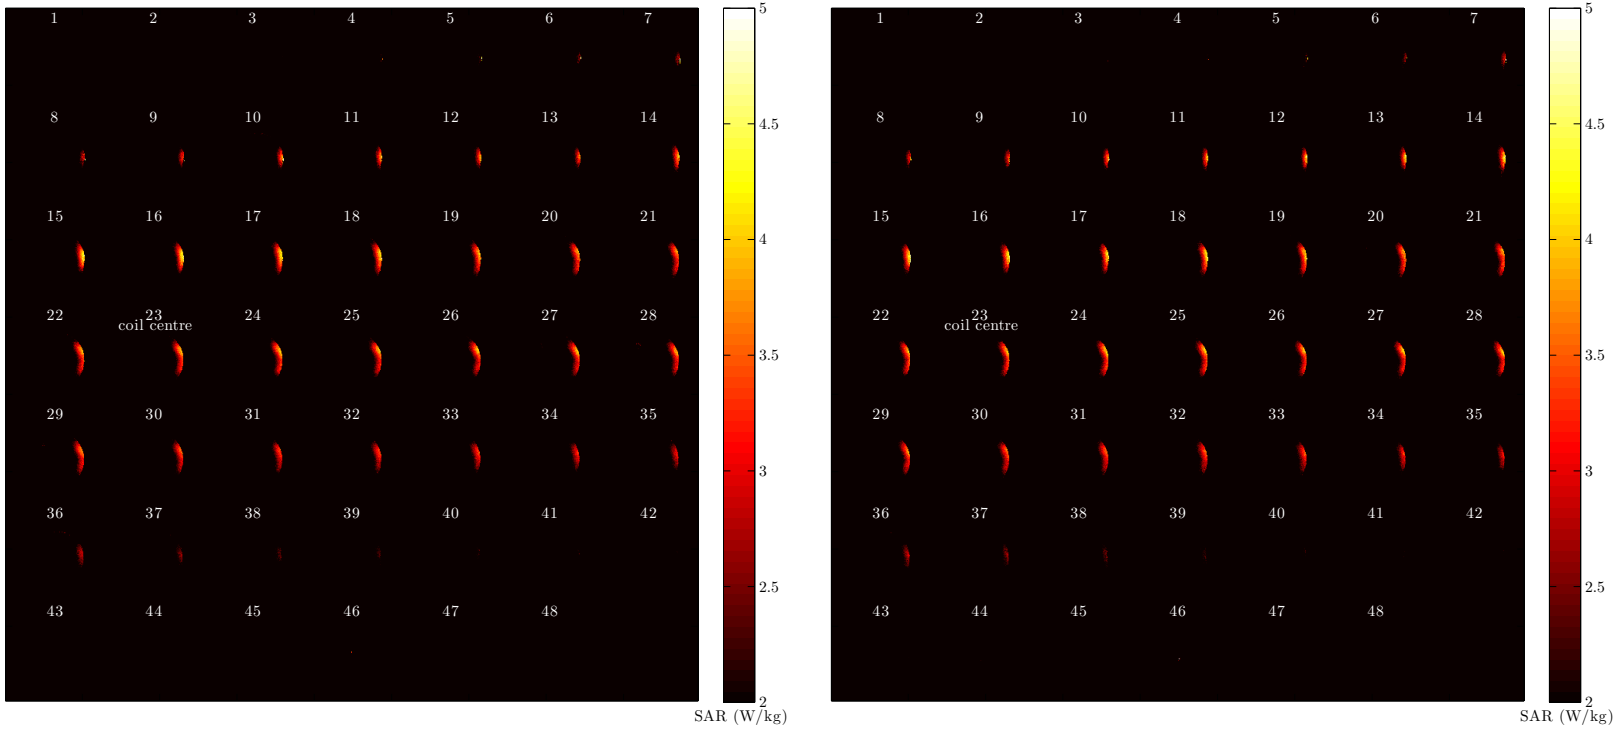

a. run 1

b. run 2

Figure A.6: SAR Maps for coil 1. Colour scaled 2-5 W/kg

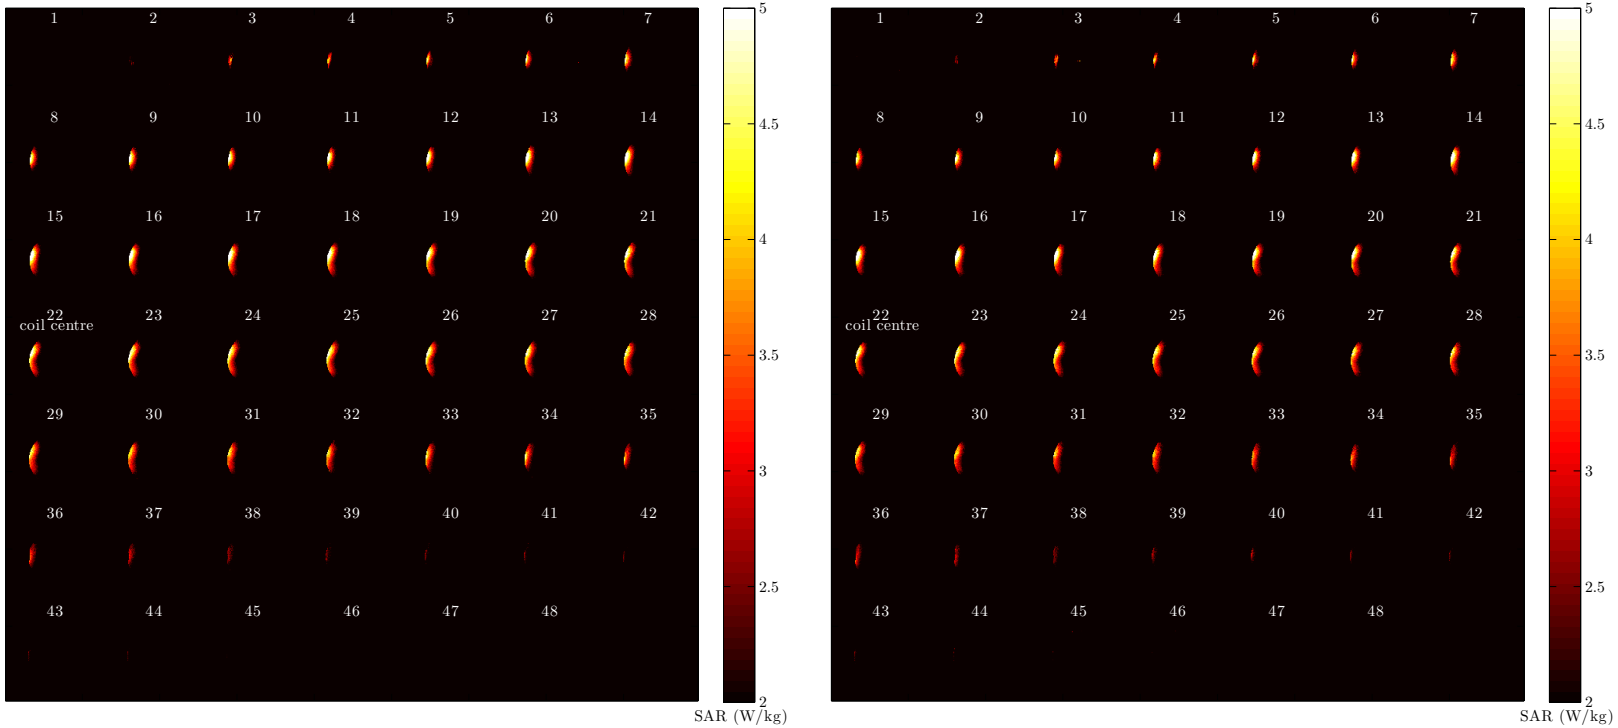

a. run 1

b. run 2

Figure A.7: SAR Maps for coil 2. Colour scaled 2-5 W/kg

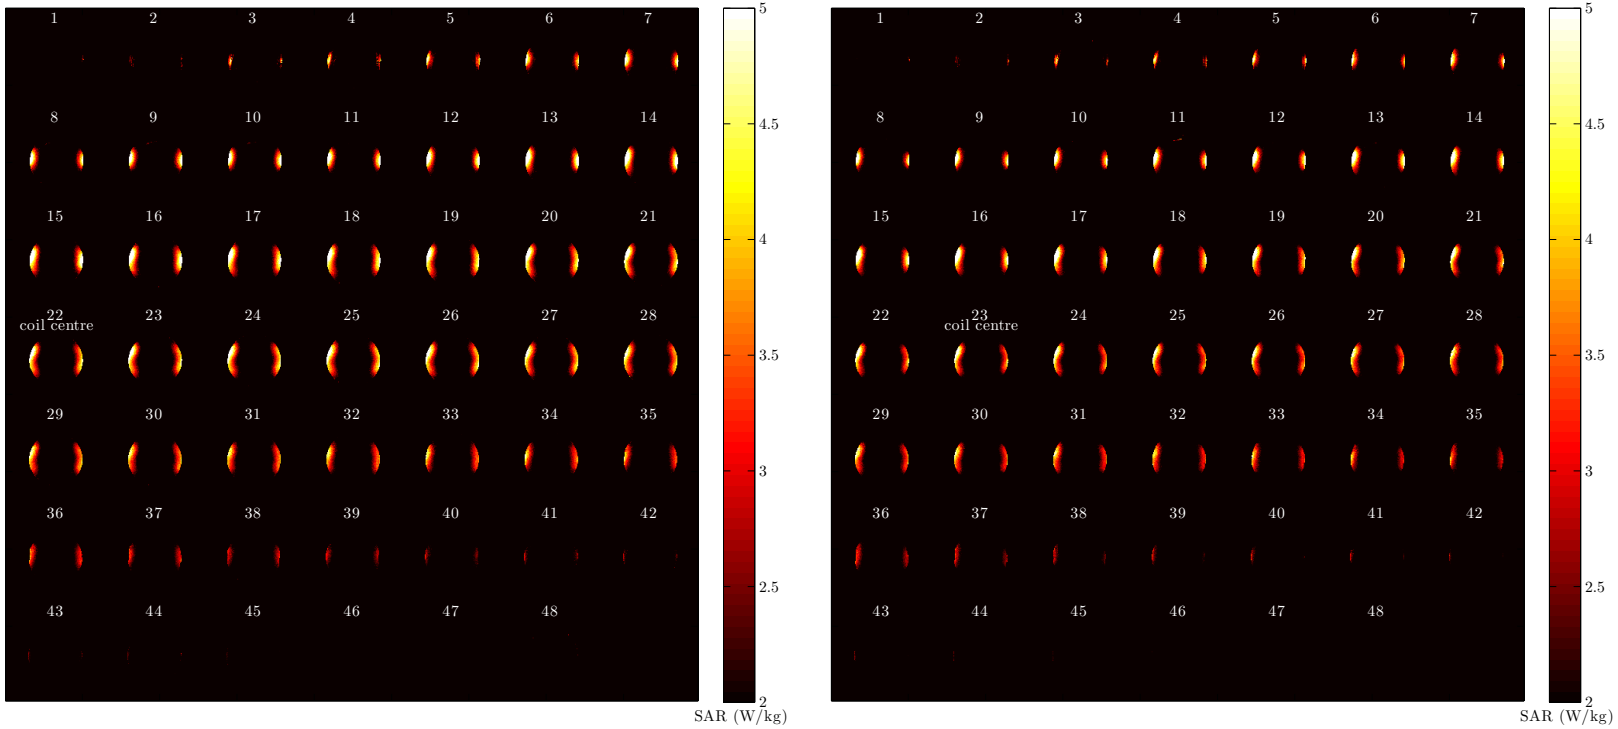

a. run 1

b. run 2

Figure A.8: SAR Maps for Anti-Phase configuration. Colour scaled 2-5 W/kg

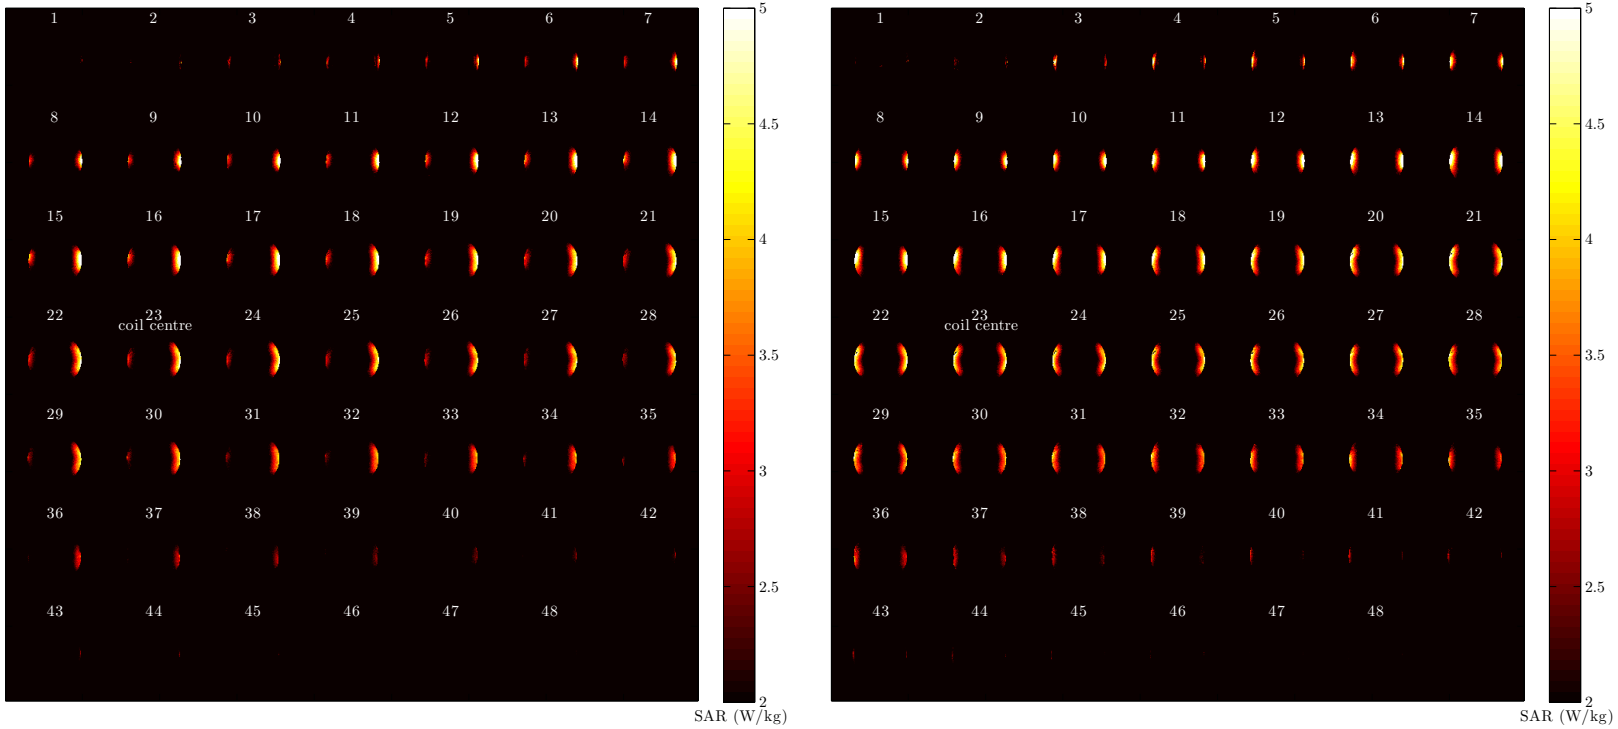

a. run 1

b. run 2

Figure A.9: SAR Maps for quadrature configuration. Colour scaled 2-5 W/kg

a. run 1

b. run 2

Figure A.10: SAR Maps for In-Phase configuration. Colour scaled 2-5 W/kg

A.2 ΔT Maps

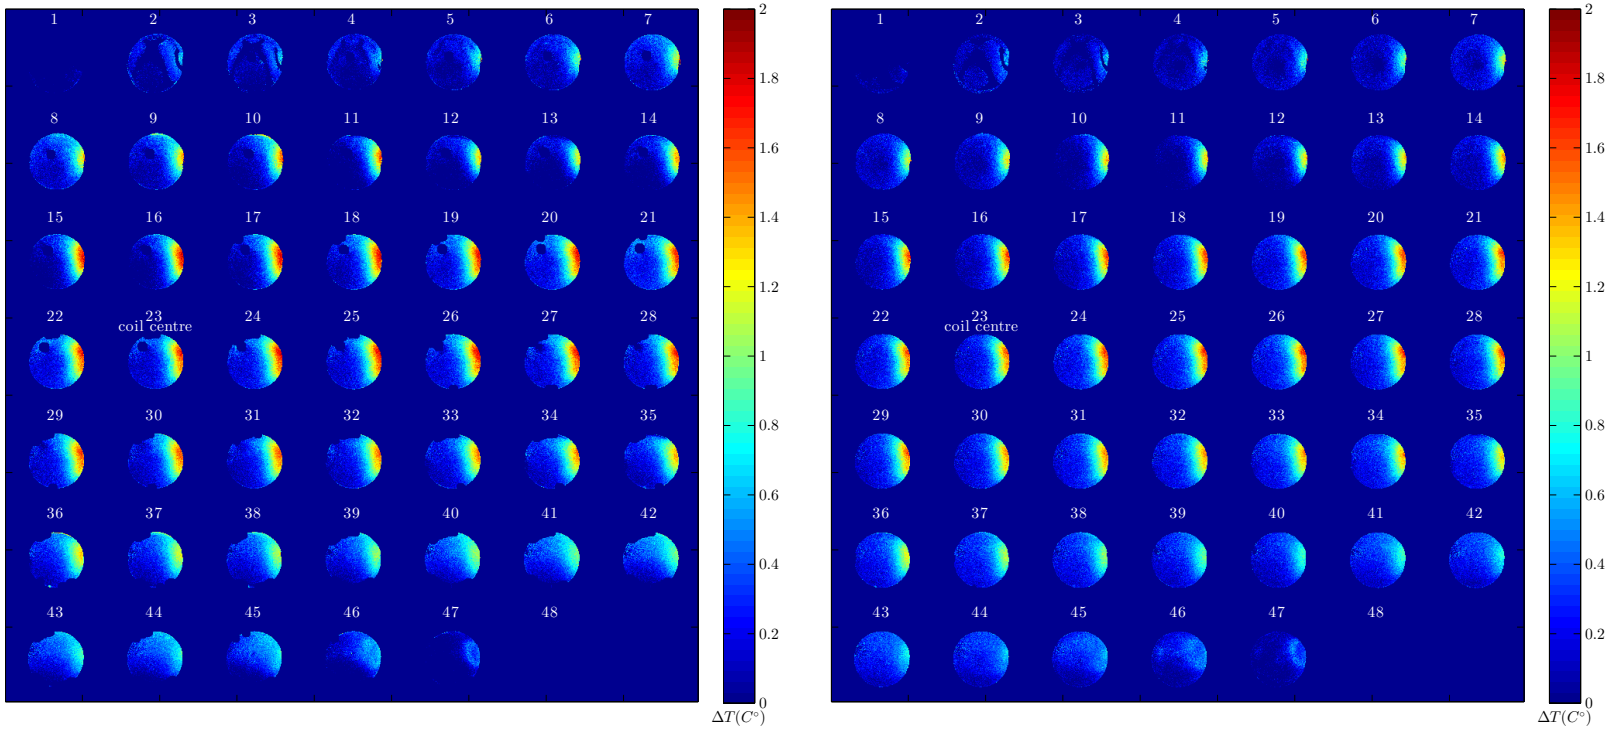

a. run 1

b. run 2

Figure A.11: ΔT maps for coil 1. Colour scaled $0 - 2°C$

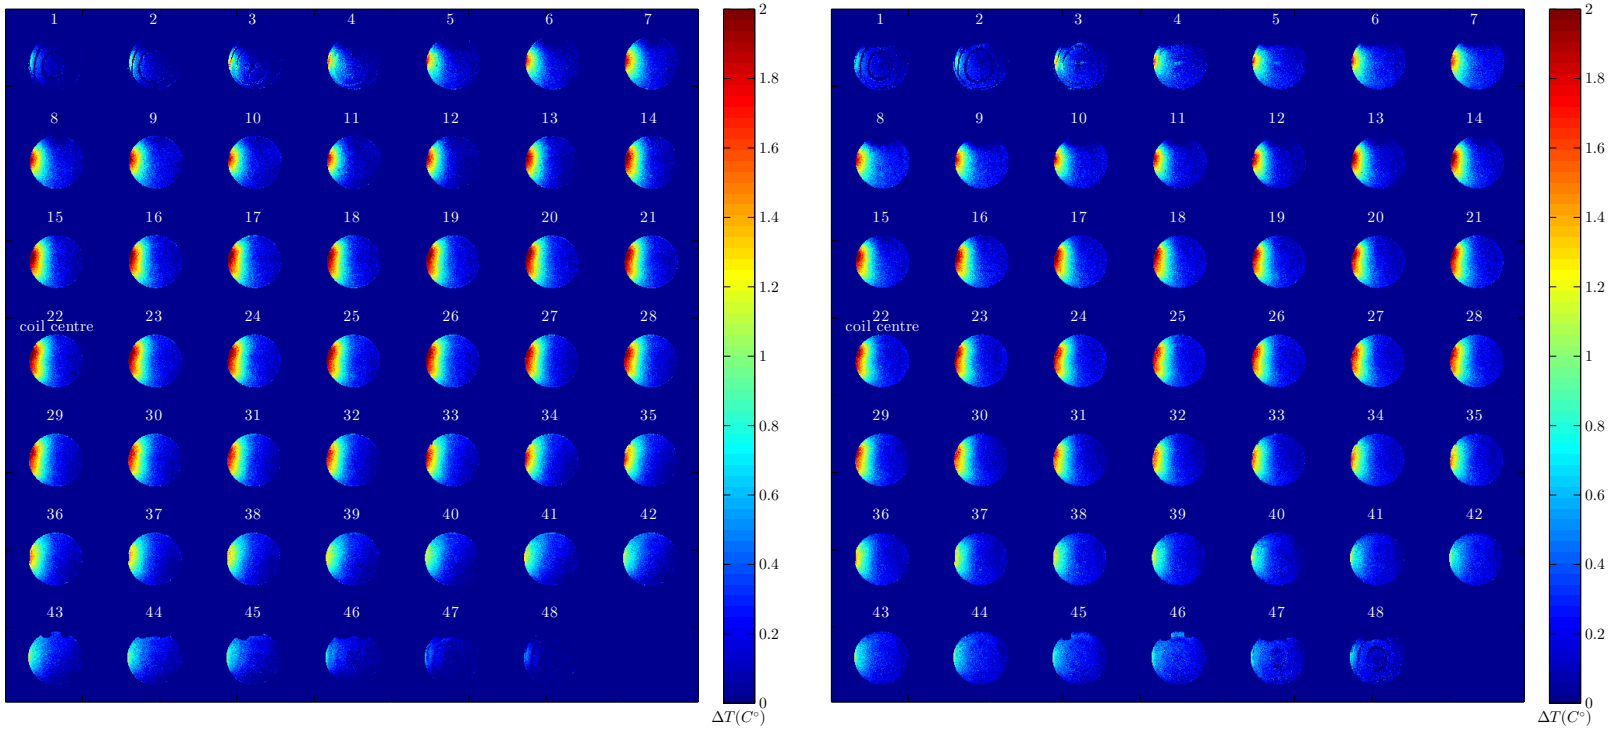

a. run 1

b. run 2

Figure A.12: ΔT maps for coil 2. Colour scaled $0 - 2°C$

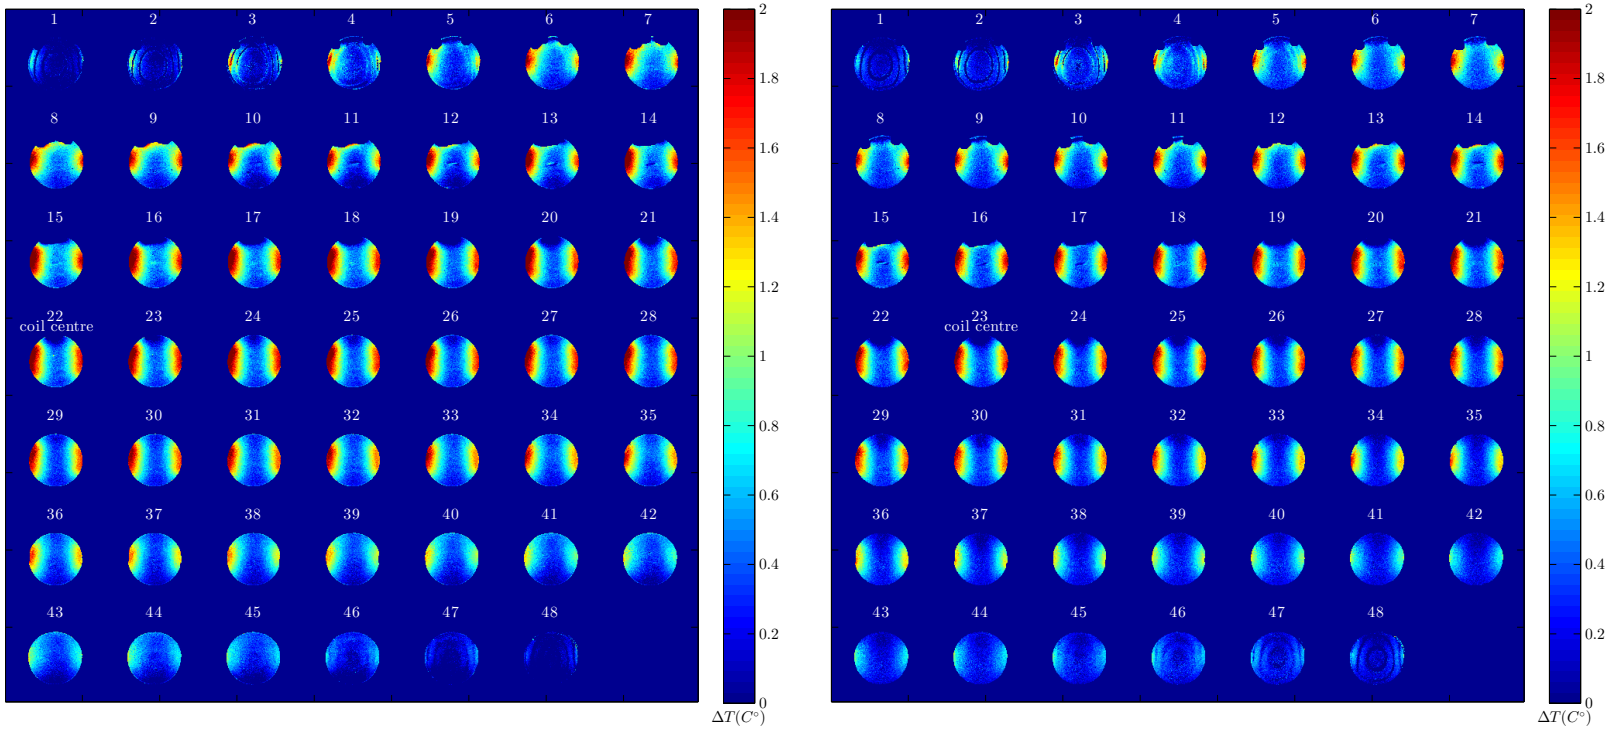

a. run 1 b. run 2

Figure A.13: ΔT maps for Anti-Phase configuration. Colour scaled $0 - 2°C$

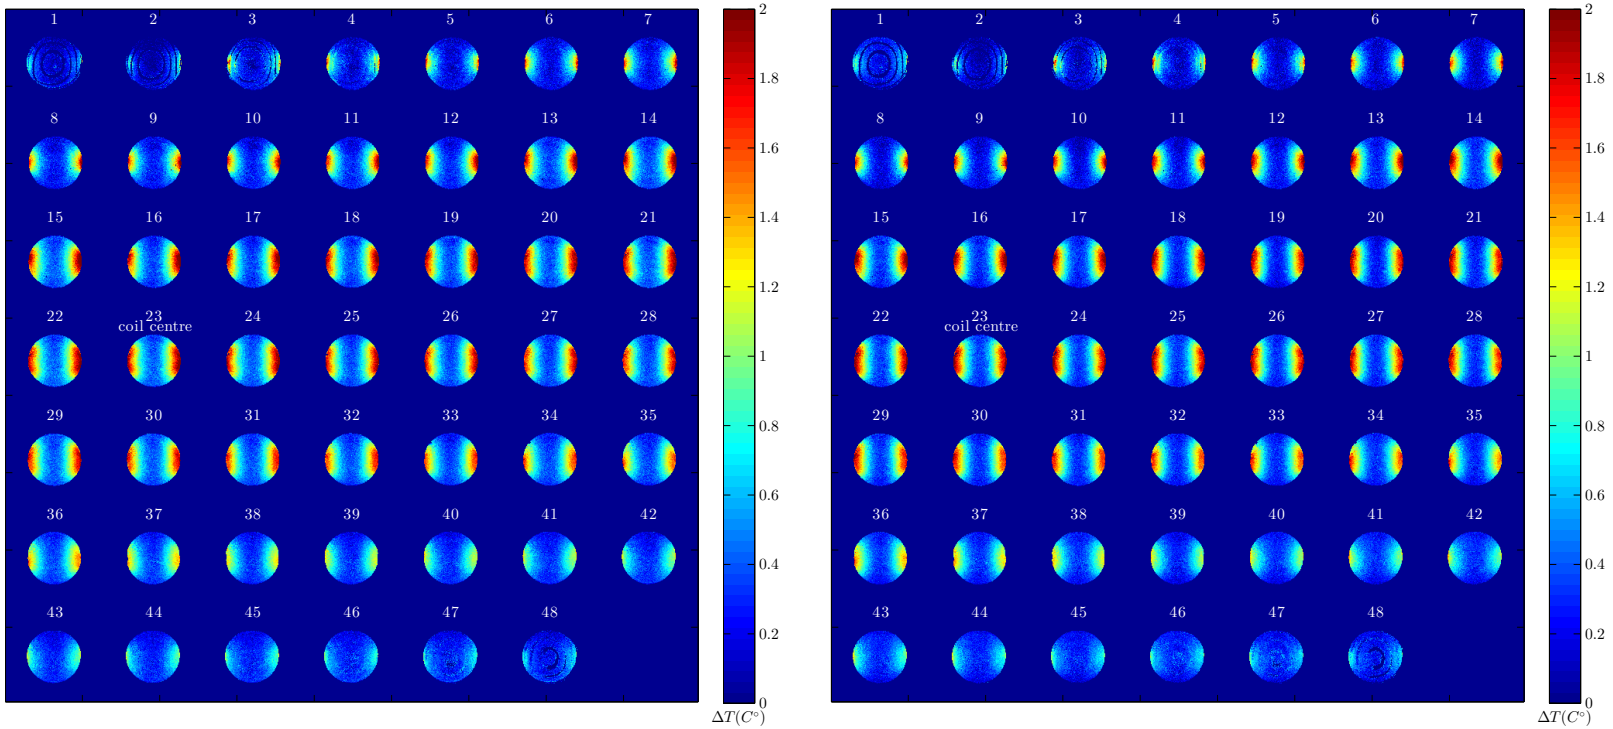

a. run 1 b. run 2

Figure A.14: ΔT maps for quadrature configuration. Colour scaled $0 - 2°C$

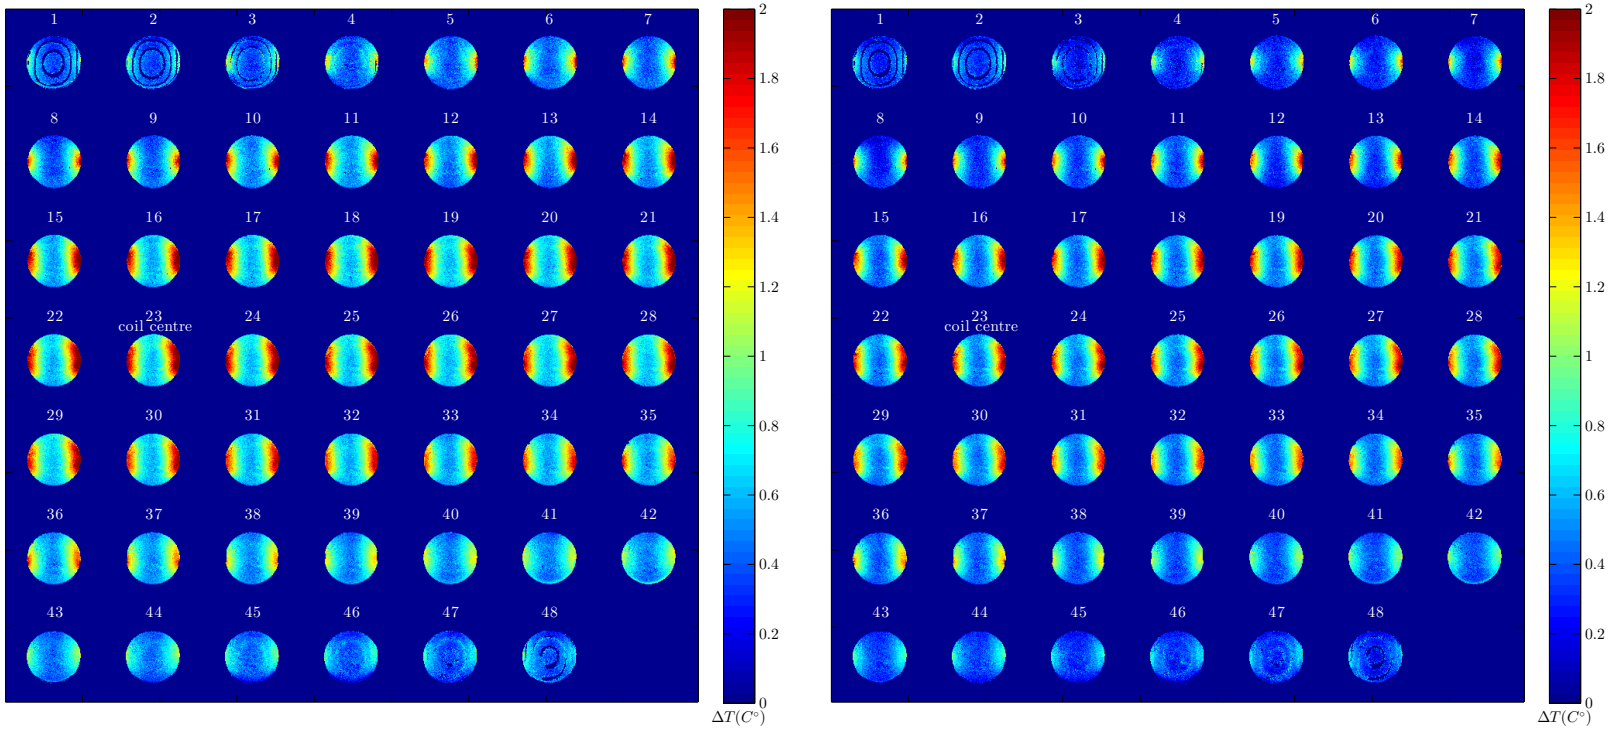

a. run 1

b. run 2

Figure A.15: ΔT maps for In-Phase configuration. Colour scaled $0 - 2°C$

A.3 Sum of Squares Error Maps

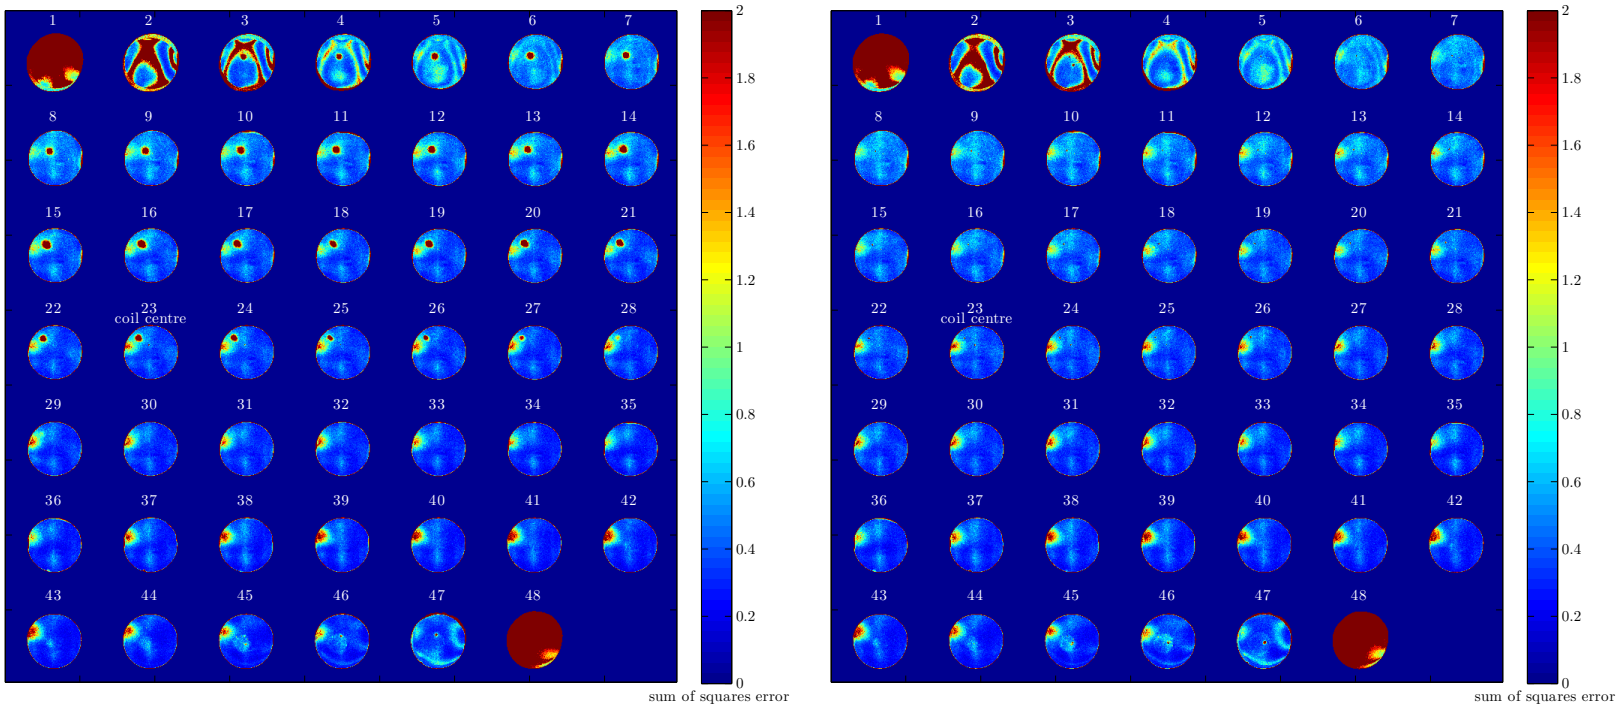

a. run 1

b. run 2

Figure A.16: Sum of squares error maps for coil 1. Colour scaled 0-2

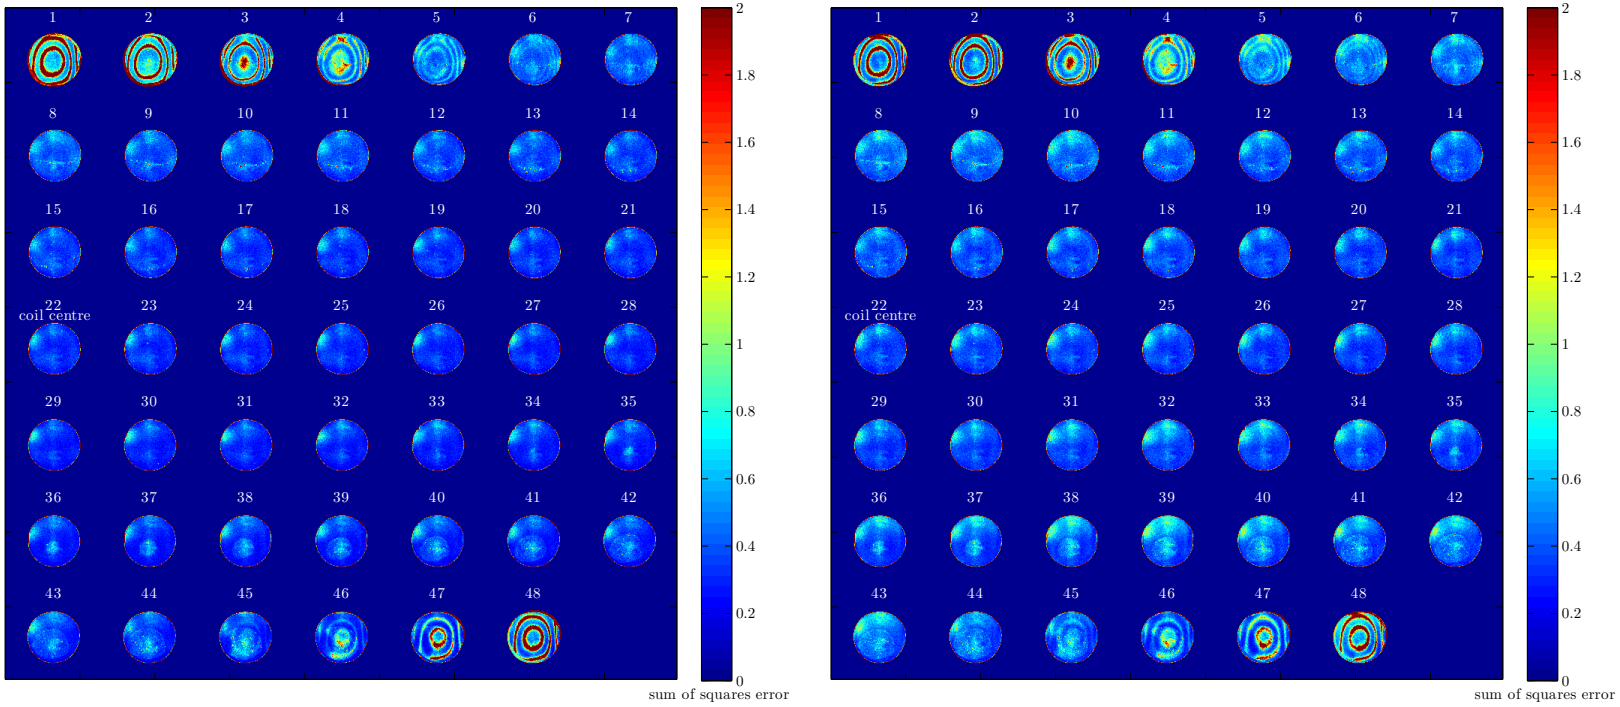

a. run 1 b. run 2

Figure A.17: Sum of squares error maps for coil 2. Colour scaled 0-2

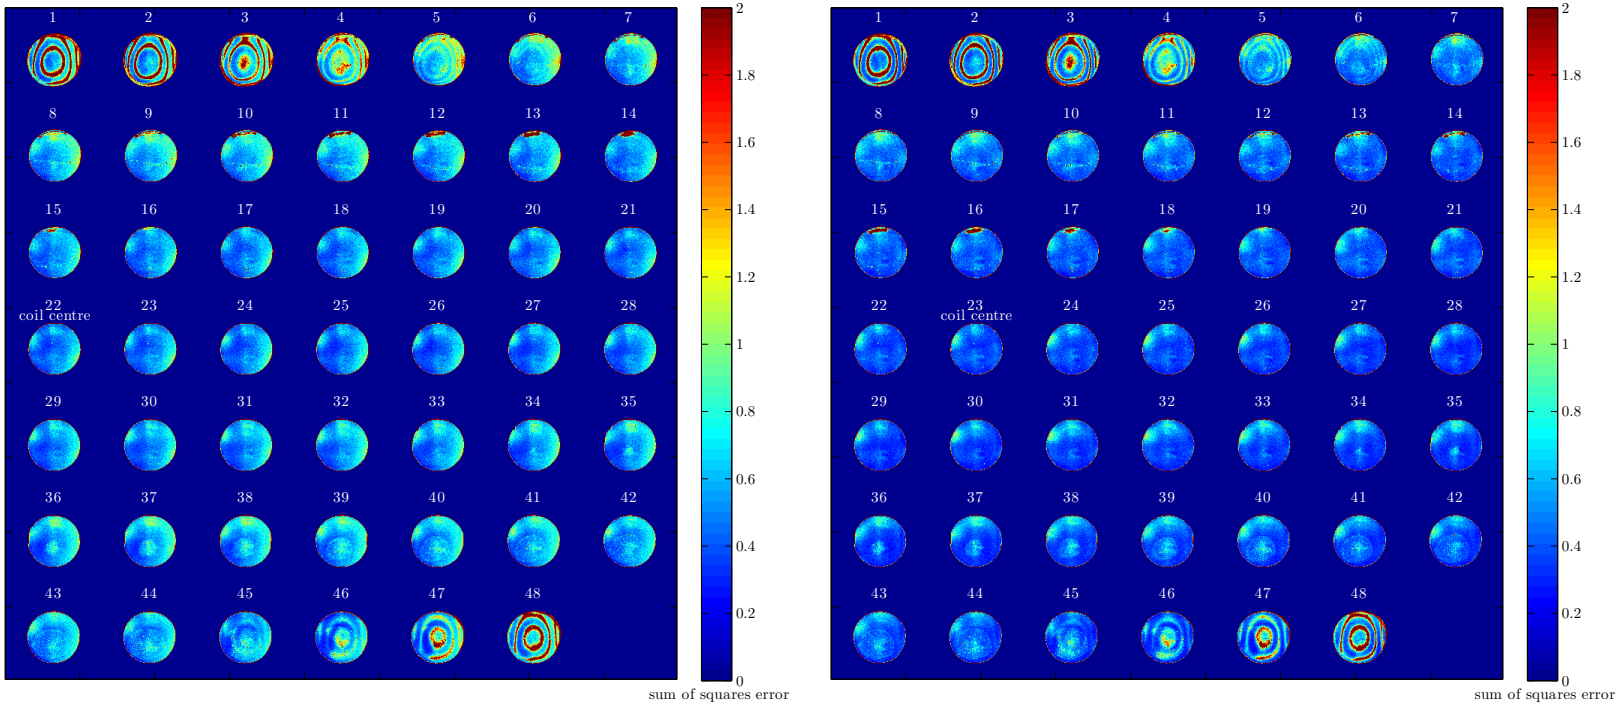

a. run 1

b. run 2

Figure A.18: Sum of squares error maps for Anti-Phase configuration. Colour scaled 0-2

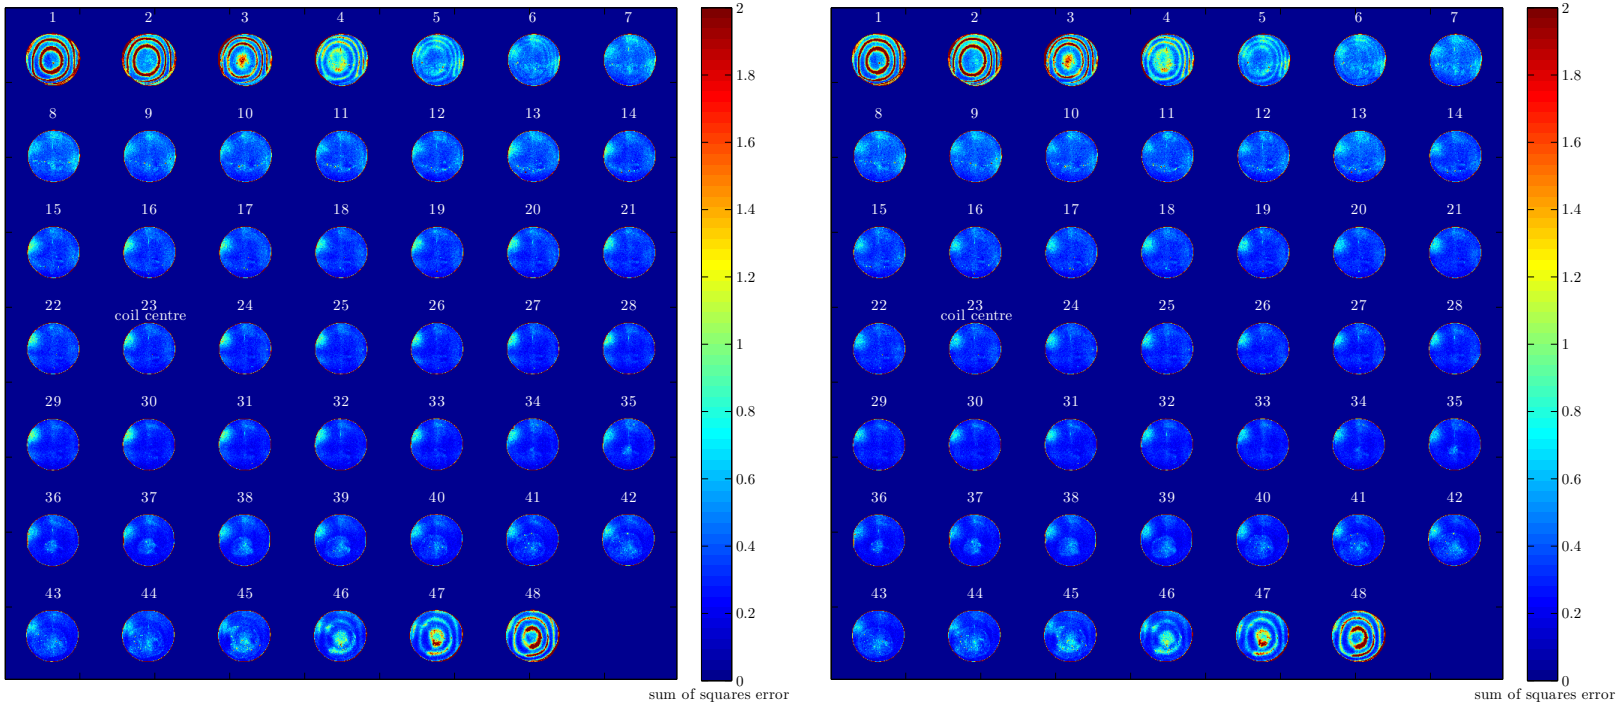

a. run 1

b. run 2

Figure A.19: Sum of squares error maps for quadrature configuration. Colour scaled 0-2

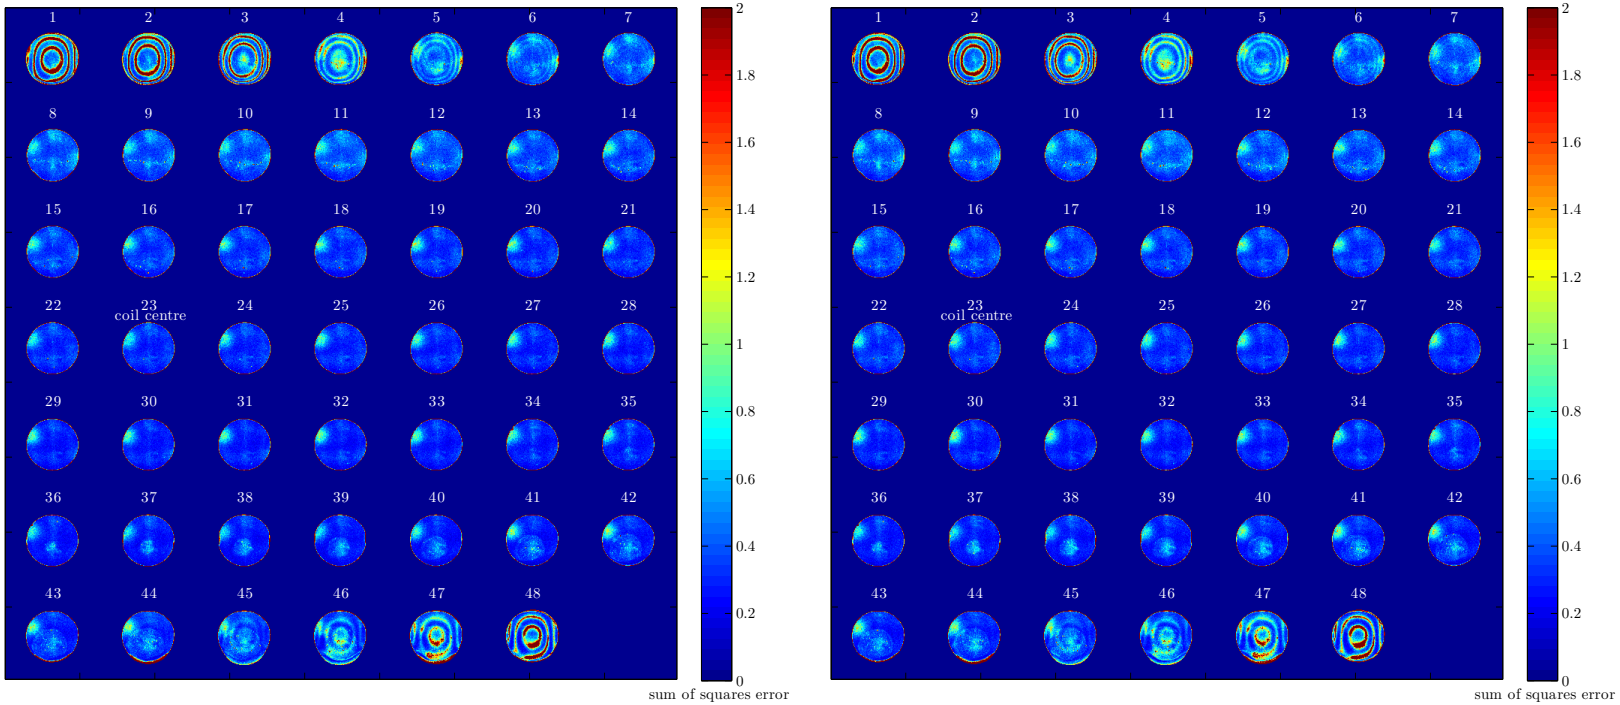

a. run 1

b. run 2

Figure A.20: Sum of squares error maps for In-Phase configuration. Colour scaled 0-2

A.4 Temperature Histograms

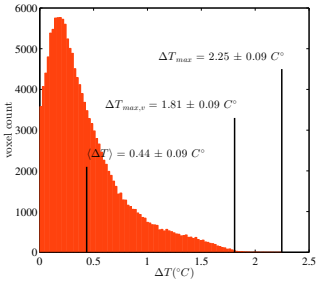

(a) Coil 1, run 1

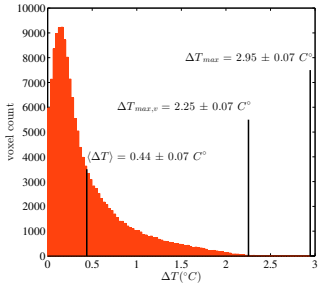

(b) Coil 2, run 1

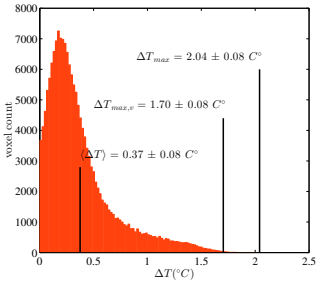

(c) Coil 1, run 2

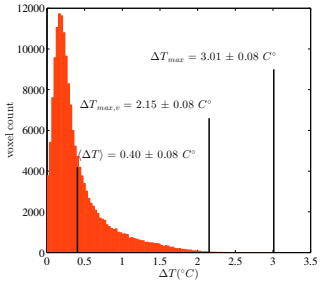

(d) Coil 2, run 2

Figure A.21: ΔT Histograms of all suitable voxels for each coil transmitting individually

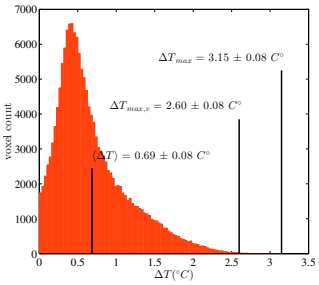

(a) Anti-Phase, run 1

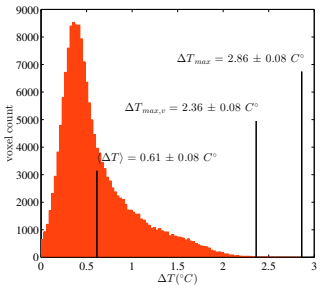

(b) Quadrature, run 1

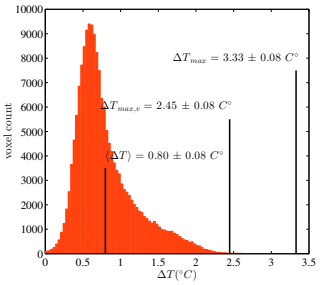

(c) In-Phase, run 1

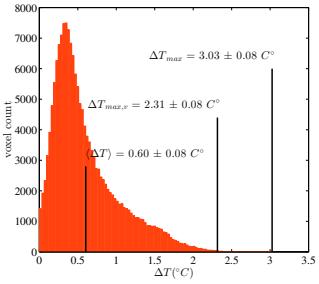

(d) Anti-Phase, run 2

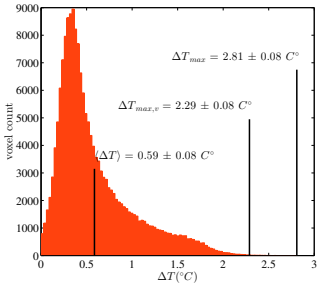

(e) Quadrature, run 2

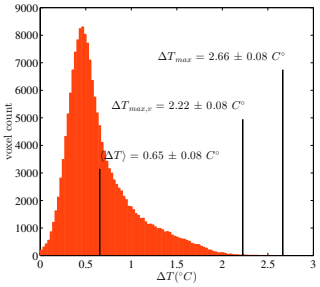

(f) In-Phase, run 2

Figure A.22: ΔT histograms of all suitable voxels for both coils transmitting simultaneously

Appendix B

Vessel Selective CASL Results

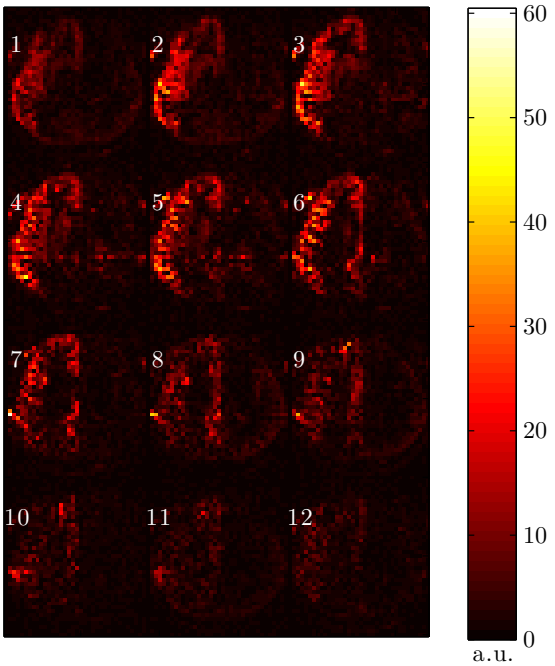

a. ΔM Left Coil

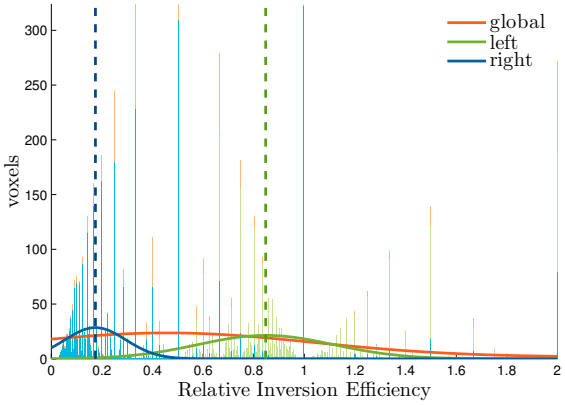

b. Histogram, Left Coil

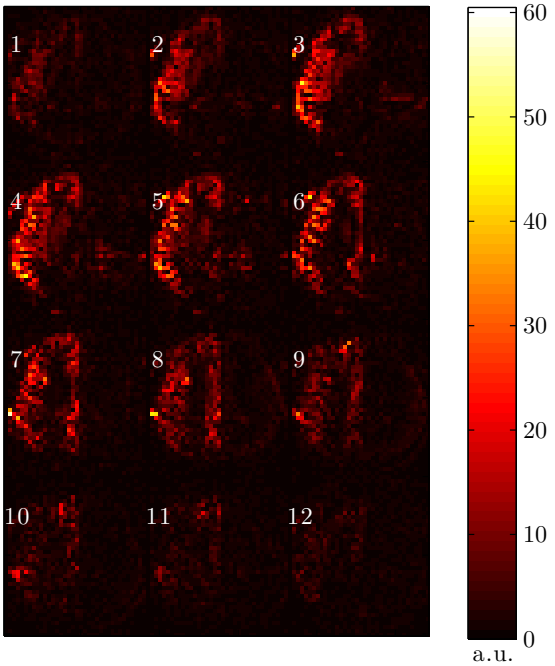

c. ΔM Parallel Transmit

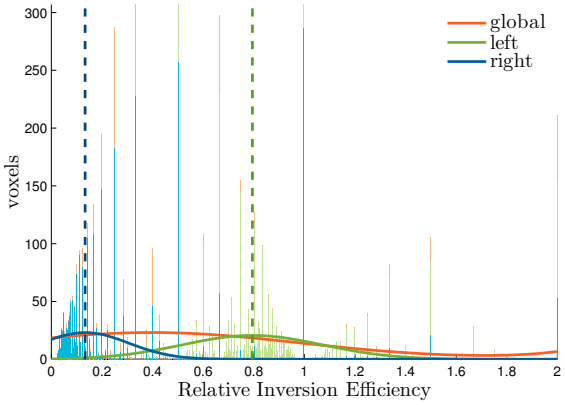

d. Histogram, Parallel Transmit

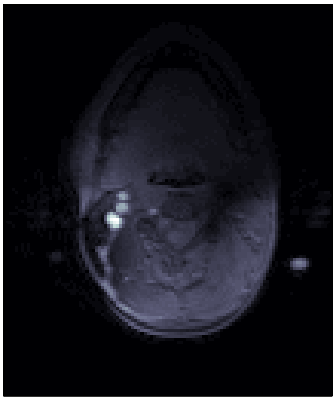

e. FLASH Image

Figure B.1: Group A, Subject 1, Vessel Selective Label Left
$A_1 = 8.5\%$, $A_2 = 100\%$, $\Delta\phi = -51.75°$

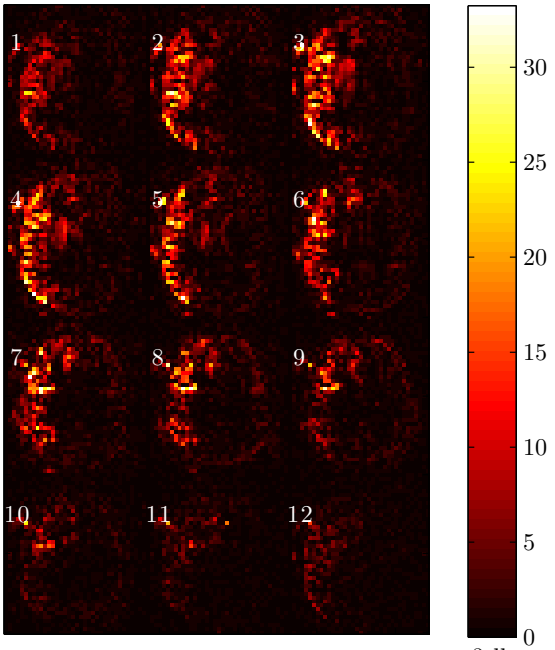

a. ΔM Left Coil

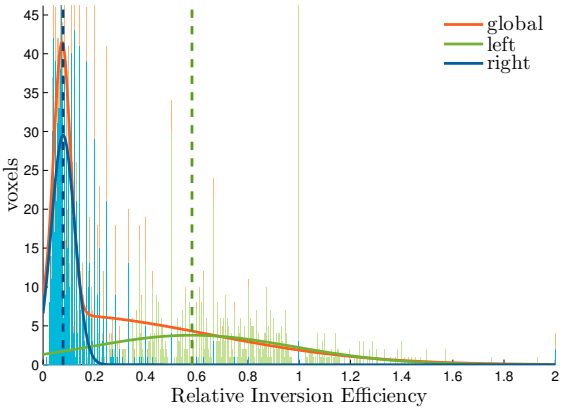

b. Histogram, Left Coil

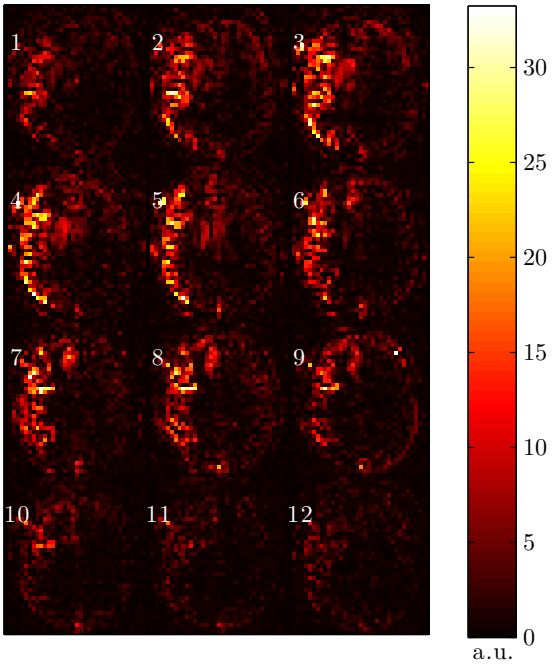

c. ΔM Parallel Transmit

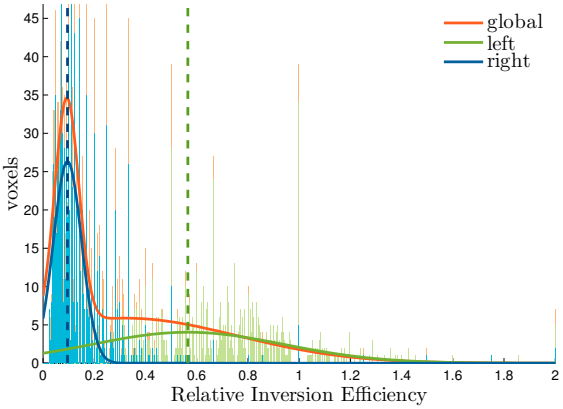

d. Histogram, Parallel Transmit

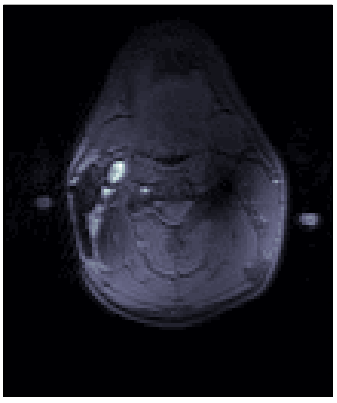

e. FLASH Image

Figure B.2: Group A, Subject 3, Vessel Selective Label Left
$A_1 = 9.04\%$, $A_2 = 100\%$, $\Delta\phi = -60.08°$

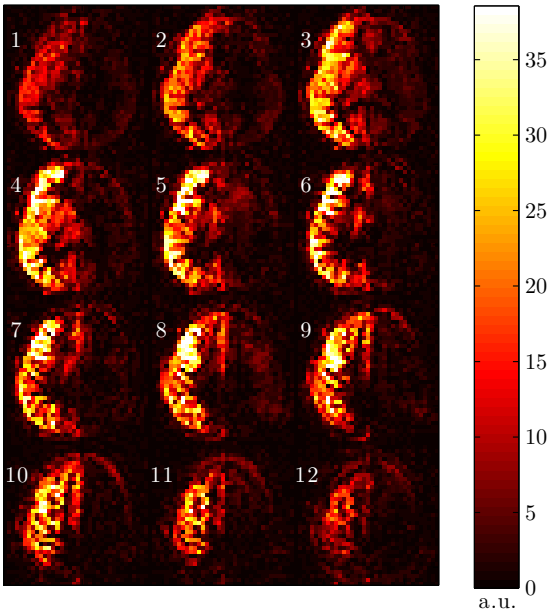

a. ΔM Left Coil

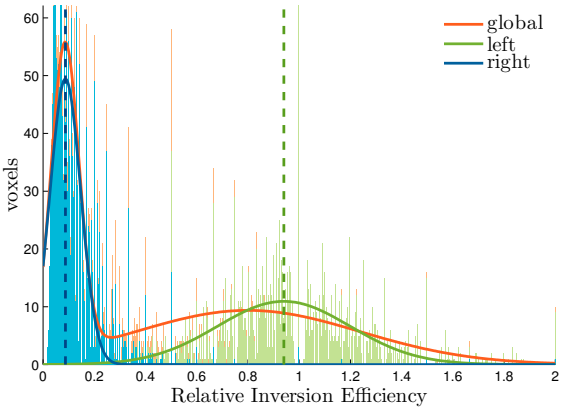

b. Histogram, Left Coil

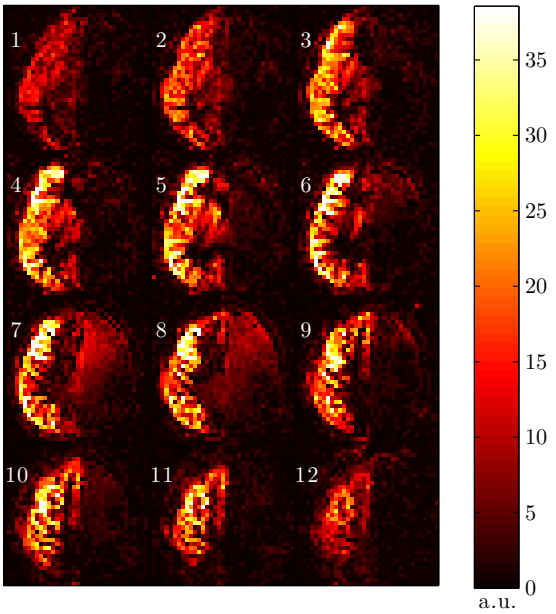

c. ΔM Parallel Transmit

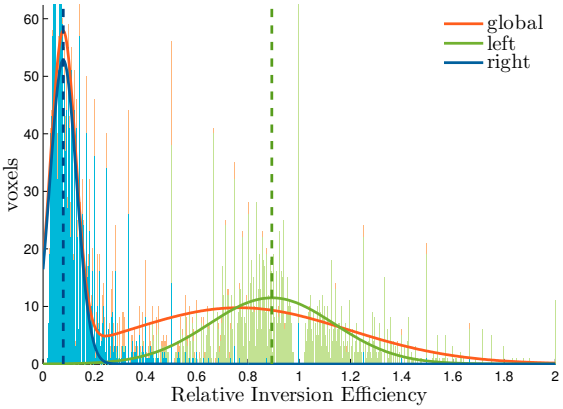

d. Histogram, Parallel Transmit

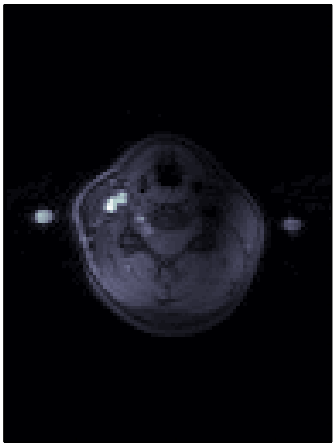

e. FLASH Image

Figure B.3: Group B, Subject 1, Vessel Selective Label Left
$A_1 = 15.93\%$, $A_2 = 100\%$, $\Delta\phi = 56.93°$

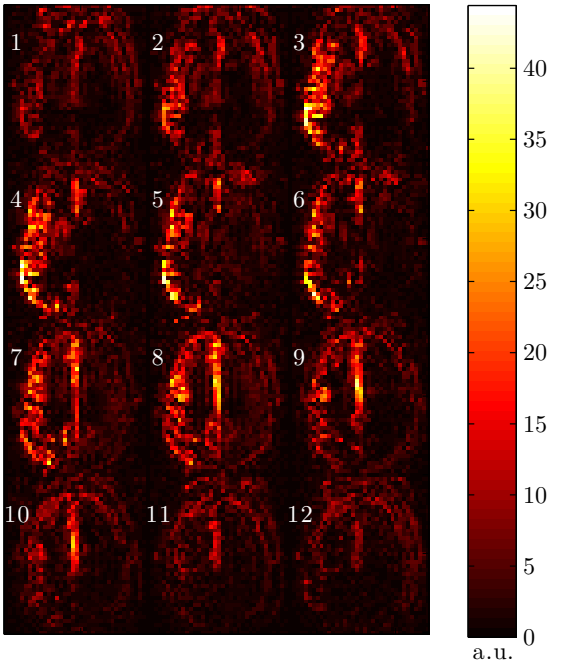

a. ΔM Left Coil

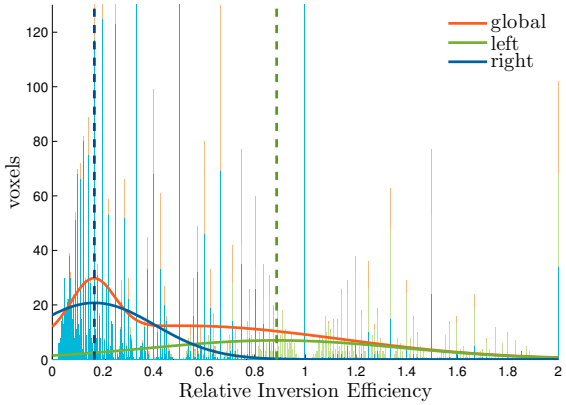

b. Histogram, Left Coil

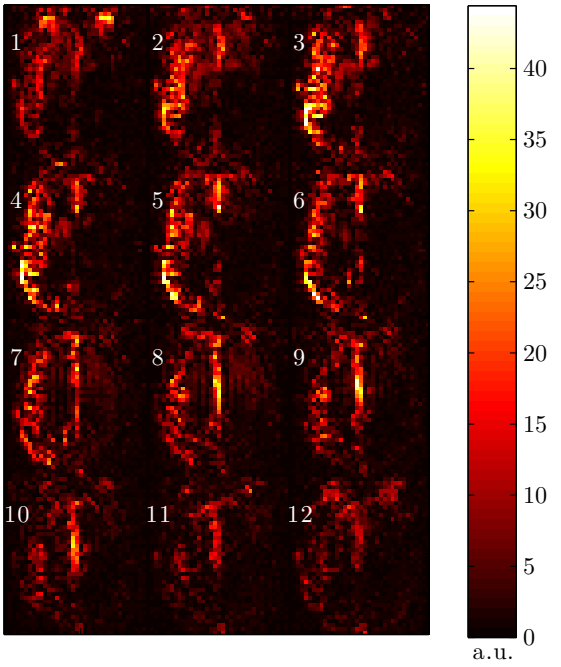

c. ΔM Parallel Transmit

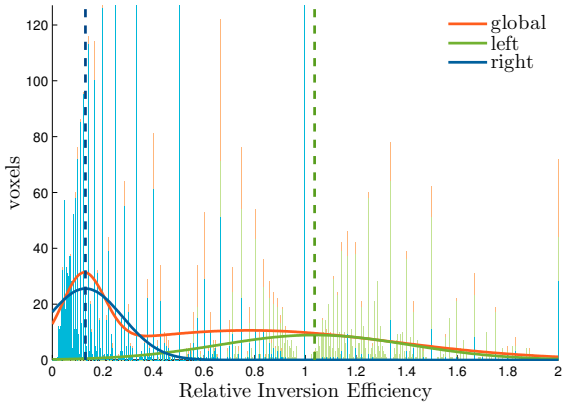

d. Histogram, Parallel Transmit

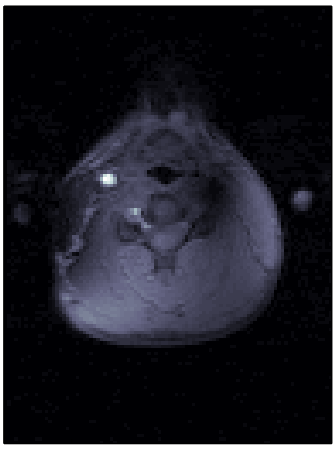

e. FLASH Image

Figure B.4: Group B, Subject 2, Vessel Selective Label Left
$A_1 = 8.11\%$, $A_2 = 100\%$, $\Delta\phi = 129.38°$

	Parallel Transmit				Single Coil			
	Left ROI		Right ROI		Left ROI		Right ROI	
Subject	$\langle\xi\rangle$	σ_ξ	$\langle\xi\rangle$	σ_ξ	$\langle\xi\rangle$	σ_ξ	$\langle\xi\rangle$	σ_ξ
A1	0.83	0.63	0.29	0.50	0.91	0.63	0.36	0.56
A2	0.96	0.51	0.14	0.26	0.97	0.44	0.22	0.43
A3	0.58	0.35	0.12	0.16	0.60	0.37	0.09	0.14
A4-1	0.85	0.34	0.25	0.36	0.97	0.28	0.23	0.25
A4-2	0.87	0.35	0.18	0.25	0.84	0.32	0.20	0.26
B1	0.91	0.49	0.11	0.19	0.90	0.44	0.12	0.14
B2	1.07	0.77	0.31	0.49	0.99	0.70	0.40	0.62
B3	0.59	0.31	0.07	0.11	0.56	0.34	0.08	0.12
Average$\langle\xi\rangle$	0.83 ± 0.17		0.18 ± 0.09		0.84 ± 0.17		0.21 ± 0.12	

Table B.1: Relative inversion efficiency ROI statistics, label left

	Parallel Transmit				Single Coil			
	Left ROI		Right ROI		Left ROI		Right ROI	
Subject	$\langle\xi\rangle$	σ_ξ	$\langle\xi\rangle$	σ_ξ	$\langle\xi\rangle$	σ_ξ	$\langle\xi\rangle$	σ_ξ
A2	0.14	0.29	0.92	0.44	0.11	0.22	0.92	0.39
A4	0.15	0.19	0.85	0.37	0.24	0.23	0.88	0.34
B3	0.06	0.08	0.61	0.24	0.08	0.08	0.62	0.26
Average $\langle\xi\rangle$	0.12 ± 0.05		0.79 ± 0.16		0.14 ± 0.08		0.80 ± 0.16	

Table B.2: Relative inversion efficiency ROI statistics, label right

	Parallel Transmit				Single Coil			
	Left ROI		Right ROI		Left ROI		Right ROI	
Subject	μ	σ	μ	σ	μ	σ	μ	σ
A1	0.79	0.28	0.13	0.17	0.85	0.27	0.11	0.12
A2	0.93	0.17	0.07	0.05	0.95	0.20	0.11	0.12
A3	0.57	0.37	0.10	0.06	0.58	0.40	0.08	0.05
A4-1	0.86	0.24	0.15	0.13	1.00	0.17	0.14	0.07
A4-2	0.89	0.23	0.11	0.07	0.85	0.25	0.13	0.08
B1	0.89	0.24	0.08	0.05	0.94	0.26	0.09	0.06
B2	1.04	0.38	0.13	0.15	0.89	0.50	0.17	0.24
B3	0.62	0.22	0.05	0.03	0.59	0.21	0.06	0.04
Average μ	0.82 ± 0.16		0.10 ± 0.03		0.83 ± 0.16		0.11 ± 0.04	

Table B.3: Relative inversion efficiency ROI gaussian fit, label left

	Parallel Transmit				Single Coil			
	Left ROI		Right ROI		Left ROI		Right ROI	
Subject	μ	σ	μ	σ	μ	σ	μ	σ
A2	0.08	0.05	0.88	0.20	0.08	0.05	0.88	0.20
A4	0.12	0.06	0.84	0.22	0.18	0.09	0.88	0.22
B3	0.05	0.02	0.66	0.14	0.06	0.04	0.66	0.16
Average μ	0.08 ± 0.04		0.79 ± 0.12		0.11 ± 0.07		0.81 ± 0.13	

Table B.4: Relative inversion efficiency ROI gaussian fit, label right

Subject	$\Delta \langle \xi \rangle_{left}$ (%)	$\Delta \mu_{left}$ (%)	$\Delta \langle \xi \rangle_{right}$ (%)	$\Delta \mu_{right}$ (%)
A1	-8.53	-6.21	-24.13	-23.24
A2	-0.33	-2.49	-65.95	-34.57
A3	-3.90	-2.78	24.97	22.03
A4-1	-12.40	-14.33	10.80	1.82
A4-2	3.79	4.33	-10.85	-11.92
B1	0.24	-4.99	-2.53	-9.78
B2	7.99	16.94	-31.51	-21.35
B3	5.61	5.61	-16.03	-11.81
Average	-0.94 ± 7.02	-0.49 ± 9.40	-14.4 ± 27.74	-11.10 ± 17.20

Table B.5: Percent change in inversion efficiency for the ROI mean, and ROI gaussian function, label left

Subject	$\Delta \langle \xi \rangle_{left}$ (%)	$\Delta \mu_{left}$ (%)	$\Delta \langle \xi \rangle_{right}$ (%)	$\Delta \mu_{right}$ (%)
A2	26.29	2.01	0.00	0.15
A4	-37.97	-35.01	-3.85	-4.50
B3	-24.75	-28.44	-0.98	-0.59
Average	-12.14 ± 33.93	-20.48 ± 19.75	-1.61 ± 2.00	-1.65 ± 2.50

Table B.6: Percent change in inversion efficiency for the ROI mean, and ROI gaussian function, label right