Lawrence Berkeley National Laboratory

Lawrence Berkeley National Laboratory

Title

SQUID-Detected MRI in the Limit of Zero Static Field

Permalink

https://escholarship.org/uc/item/47x85008

Author

Kelso, Nathan Dean

Publication Date

2009-12-14

Peer reviewed

SQUID-Detected MRI in the Limit of Zero Static Field

by

Nathan Dean Kelso

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor John Clarke, Chair
Professor Zi Q. Qiu
Professor Alexander Pines

Fall 2009

SQUID-Detected MRI in the Limit of Zero Static Field

Nathan Dean Kelso

Disclaimer

Abstract

SQUID-Detected MRI in the Limit of Zero Static Field

by

Nathan Dean Kelso

Doctor of Philosophy in Physics

University of California, Berkeley

Professor John Clarke, Chair

The magnetic gradient fields used in magnetic resonance imaging (MRI) have a component which is parallel to the uniform field $\mathbf{B_0} = B_0\hat{\mathbf{z}}$, as well as a component perpendicular to $\mathbf{B_0}$. The component parallel to $\mathbf{B_0}$ is used in spatial encoding. The component perpendicular to $\mathbf{B_0}$, called the "concomitant gradient," causes image distortions (by altering the magnitude and direction of the total field) if its magnitude approaches B_0 at any point in the field of view (FOV). In a conventional imaging sequence, the presence of the concomitant gradients limits the maximum gradient that can be used with a given $\mathbf{B_0}$ field or, conversely, limits the minimum $\mathbf{B_0}$ field that can be used with a given gradient field.

This thesis describes an implementation of the so-called "zero-field MRI" (ZFMRI) pulse sequence, which allows for imaging in an arbitrarily low $\mathbf{B_0}$ field. The ZFMRI sequence created an effective unidirectional gradient field by using a train of π pulses to average out the concomitant gradient components during encoding. The signals were acquired using a low-transition temperature dc Superconducting QUantum Interference Device (low-T_c dc SQUID) coupled to a first-order axial gradiometer. The experiments were carried out in a liquid helium dewar which was magnetically shielded with a single-layer mu-metal can around the outside and a superconducting Pb can contained within the helium space. We increased the filling factor of the custom-made, double-walled Pyrex insert by placing the liquid alcohol sample, at a temperature of approximately $-50°$C, at the center of one loop of the superconducting gradiometer, which was immersed in the helium bath.

Using the aforementioned sequence and apparatus, images were acquired in the limit of zero static field, using gradients of up to 100 μT/m over a 23 mm FOV. The change in field magnitude over the FOV due to gradients was up to 10 times larger than the magnitude of any static field present in the dewar (static fields arose from residual magnetic fields and were 1 μT or less). These images were free of concomitant gradient distortions. Images encoded using a conventional imaging sequence under similar conditions were also acquired; the conventional images were irreparably distorted.

The limitations of the present ZFMRI sequence implementation are considered, as well as how the procedure could be made more practical with regard to imaging time. The extension of the technique to unshielded operation in a uniform ambient field is discussed, as are other methods of mitigating or eliminating concomitant gradient distortions.

Contents

List of Figures

Acknowledgments

I must first and foremost thank Professor John Clarke for his many years of guidance and support. It has been a privilege to work in his group, and it would be an understatement to say that I merely received an education. I am also indebted to Professor Alex Pines for his enthusiasm not just for the zero-field MRI project, but for SQUID-detected MRI in general. I also thank Prof. Pines for serving on my dissertation and qualifying exam committees with Prof. Zi Q. Qiu, and Prof. Ivo Souza for serving on my qualifying exam committee.

I began work on SQUID-detected MRI with Dr. Robert McDermott. I learned many things from him, including how to solder, how a SQUID works, and that it is okay to break into song during liquid helium transfers. I also worked closely with Dr. Seung-Kyun Lee, who taught me a great deal about SQUIDs and MRI. Without his extensive skills, the zero-field MRI project may never have gotten off the ground. The example he set for hard work and laboratory excellence has been an inspiration.

On SQUID-detected MRI, I have also been privileged to work closely with Dr. Michael Moessle, Michael Hatridge, and Dr. Whittier Myers. I especially thank Sarah Busch, with whom I have had many thoughtful, enlightening, and entertaining discussions about physics (and just about every other imaginable topic). I also wish to thank all of the members and visitors with whom I have overlapped, especially Steve Anton, Dr. Shane Cybart, Stephen Wu, Jed Johnson, Dr. Emile Hoskinson, Monica Smith, Dr. Travis Hime, Dr. Paul Reichardt, Dr. Darin Kinion, Dr. Tim Robertson, and Dr. Bennie ten Haken. I also must express my gratitude to Barbara Salisbury for her help throughout my time in graduate school.

Over the years I have benefited from collaborations with Dr. Song-I Han, Dr. Kristie Koski, and Dr. Vasiliki Demas. I am particularly grateful to Dr. Louis Bouchard; in addition to his help with the zero-field MRI project, he has been almost supernaturally helpful with all aspects of magnetic resonance.

I also thank the staff of the Physics Department Electronics Shop and that of the Machine Shop for their expert craftsmanship and willingness to offer advice and assistance. I especially thank Joseph Kant of the Student Machine Shop; without his instruction and advice, very little of my graduate work could have been accomplished. In addition, the Physics Department staff, especially Anne Takizawa, Donna Sakima, and Anthony Vitan, have been indispensable.

This project would have come to a sudden and ignoble end were it not for Jim Breen of the Chemistry Department Glass Shop, who made the new Pyrex insert after the old one suffered, as engineers say, "an anomaly." I would also like to thank the Chemistry Department Machine Shop for making one of the phantoms for this project, as well as for a wonderful construction job for my elastography project (which is not described in this thesis).

I am, as always, grateful for the love and support of my parents. I also thank my brother, who was working toward a doctorate for much of the time I was, and able to share grad school "war stories."

Finally, my deepest and most heartfelt thanks go to my wonderful wife, Janelle, for all of her love, patience, and support, and to our beautiful daughter, Annabelle.

Chapter 1

Introduction

Magnetic resonance imaging (MRI) in ever higher fields has obvious advantages in increased signal-to-noise ratio (SNR) and resolution [1]. However, there has been recent interest in the opposite limit: that of MRI in field magnitudes on the order of microtesla [2–6]. Although we do not expect it to supplant modern high-field MRI, imaging in very low fields does have significant advantages: the potential exists for cheaper, more open systems in which field homogeneity requirements are relaxed [2], T_1-weighted contrast is enhanced (where T_1 is the longitudinal relaxation time) [3], chemical shift artifacts are eliminated [4], and that can acquire images in the presence of metal such as titanium implants [5]. One could also imagine combining low-field MRI with existing magnetoencephalography (MEG) systems in order to combine the two functions into a single device [6].

Performing MRI in microtesla magnetic fields does pose some important challenges, however. First of all, the thermal polarization of nuclei in microtesla fields, and therefore the signal, are several order of magnitudes lower than in typical high-field systems. Low magnetic fields also imply low precession frequencies; if signal acquisition is done using tuned detectors, Faraday's Law implies that the received signal will also be low compared to typical systems. Combining these two effects, we find that for Faraday-detected signals, the received signal scales as the square of the magnetic field [1]. The signal loss involved in using microtesla fields can be compensated in part by the use of a prepolarization pulse and signal acquisition with a detector such as a SQUID, which is sensitive to the magnetic flux rather than its time rate of change.

Another challenge arising in low-field MRI is that of concomitant gradient distortions. The encoding fields contain components that are parallel to $\mathbf{B_0}$ as well as components that are perpendicular to $\mathbf{B_0}$. In conventional imaging, when the uniform field is much larger than the maximum field produced by gradients, the "concomitant gradient" fields perpendicular to $\mathbf{B_0}$ can be ignored, and the gradients treated as linear and unidirectional variations in the field magnitude. As B_0 is reduced, the concomitant gradient field magnitudes become a significant fraction of B_0 and lead to distortions in the image [7]. This sets an upper limit on the gradients that can be used for a given value of B_0, or sets a lower limit on the uniform field B_0 that can be used for a given gradient.

Several techniques have been developed to correct the distortions due to concomitant gradients in low-field imaging [8–11]. In this thesis, we will investigate imaging in the limit that $B_0 \to 0$, using a pulse sequence that averages out the concomitant gradients during image encod-

ing [9, 10], and measuring the signal with a SQUID. It is not our intention to present SQUIDs or MRI in exhaustive detail, but we will present sufficient background information on these topics to describe our experiments. Additional information about SQUIDs, NMR, and MRI can be found in the references of their respective chapters.

This thesis is organized as follows. In Chapter 2 we present an overview of the dc SQUID and the associated electronics which we used in this experiment. We also discuss the concept of the flux transformer and the use of series arrays of Josephson junctions to protect the SQUID and input coil from the pulsed magnetic fields present in our experiment. Chapter 3 contains an introduction to nuclear magnetic resonance (NMR) and the physical principles upon which MRI is based. Chapter 4 covers the basics of MRI including spatial encoding, k-space, and common imaging techniques. In Chapter 5 we describe the origin and effect of concomitant gradient components, as well as the "zero-field MRI" pulse sequence that we used to acquire images without the use of a $\mathbf{B_0}$ field. Details on our experimental implementation of the sequence and the apparatus are presented in Chapter 6, along with the results of our imaging experiments. Finally, Chapter 7 presents a discussion of the experiment and results, including the imaging time, the effect of a uniform ambient field, and a comparison to other methods of mitigating or eliminating concomitant gradient distortions.

Chapter 2

Introduction to SQUIDs

Superconducting QUantum Interference Devices (SQUIDs) are extraordinarily sensitive detectors of magnetic flux; modern devices can have a flux sensitivity of approximately 1 $\mu\Phi_0/\text{Hz}^{1/2}$ where Φ_0 is the magnetic flux quantum ($\Phi_0 = 2.07 \times 10^{-15}$ T/m) [12]. As they are superconducting devices, they must be operated at cryogenic temperatures. Despite this difficulty, because of their exquisite sensitivity SQUIDs are used in many fields of research and industry, including medicine [13], geophysics [14], non-destructive evaluation [15], materials characterization [16], and even investigations of Einstein's theory of general relativity carried out in Earth orbit [17]. In this chapter we will focus on a qualitative understanding of SQUIDs and their operation. Detailed and quantitative information on SQUID design, fabrication, and operation in a variety of circumstances are readily available, for example in [18, 19] and the references therein.

2.1 SQUID basics

2.1.1 Flux quantization in a superconducting ring

The coherence length ξ of a superconductor is the distance over which the Cooper pairs can maintain phase coherence; it is commonly thought of as the effective size of a Cooper pair [20]. The coherence lengths of low-T_c superconductors typically range from tens to hundreds of nanometers (or more) [20]. Cooper pairs within ξ will overlap with one another, and a sphere of radius ξ will contain approximately 10^6 pairs. The phases of the overlapping Cooper pairs lock together, resulting in long-range quantum coherence characterized by an ensemble-average wave function [20, 21]

$$\psi(\mathbf{r}) = [n(\mathbf{r})]^{1/2} e^{i\theta(\mathbf{r})}, \tag{2.1}$$

where $n(\mathbf{r})$ is the density of Cooper pairs and $\theta(\mathbf{r})$ is the phase. Typically $n(\mathbf{r})$ is independent of position, or nearly so, inside the superconductor; we will therefore drop the $\mathbf{r}$ dependence and write it simply as n. One consequence of the macroscopic quantum coherence is flux quantization in a superconducting loop.

Consider a ring of superconducting material placed in a magnetic field $\mathbf{B}$, as illustrated in Fig. 2.1A. The thickness of the ring is taken to be much larger than the penetration depth. Although the field inside the superconductor is zero due to the Meissner effect, the magnetic vector

Fig. 2.1: (A) Superconducting loop placed in a magnetic field. A flux Φ threads the center of the loop. (B) Schematic diagram of a Josephson junction. Gray color is superconductor; the black bar represents an insulating layer. (C) Schematic of a dc SQUID. Gray color is superconductor; the black bars are Josephson junctions. A flux Φ threads the center of the loop.

potential $\mathbf{A}$ may be nonzero. The momentum of a charged particle in a vector potential is given by

$$\mathbf{p} = m\mathbf{v} + q\mathbf{A}, \tag{2.2}$$

where m is the particle's mass, $\mathbf{v}$ is its velocity, and q is its charge. The momentum can be found using the momentum operator $-i\hbar\nabla$ which yields

$$\hbar\nabla\theta = m\mathbf{v} + q\mathbf{A}. \tag{2.3}$$

The particle's velocity $\mathbf{v}$ is proportional to the current density $\mathbf{J}$. However, because the magnetic field is zero inside the superconductor, we find

$$\mu_0\mathbf{J} = \nabla \times \mathbf{B} = 0, \tag{2.4}$$

so Eq. 2.3 reduces to

$$\hbar\nabla\theta = q\mathbf{A}. \tag{2.5}$$

We now integrate each side of this equation around the path C shown in Fig. 2.1A. The phase must be 2π-periodic, so

$$\oint_C \hbar\nabla\theta \cdot d\mathbf{l} = 2\pi\hbar n. \tag{2.6}$$

Using Stokes' theorem, the right-hand side of Eq. 2.6 becomes

$$\oint_C q\mathbf{A} \cdot d\mathbf{l} = q\int_S (\nabla \times \mathbf{A}) \cdot d\mathbf{S} = q\int_S \mathbf{B} \cdot d\mathbf{S} = q\Phi, \tag{2.7}$$

where $d\mathbf{S}$ is an element of the surface S bounded by the curve C, and Φ is the flux through C. Because we are dealing with Cooper pairs, we have $q = -2e$ where e is the (positive) electron charge. We therefore find, for the flux through C,

$$\Phi = \frac{2\pi\hbar n}{q} = \frac{h}{2e}n = \Phi_0 n. \tag{2.8}$$

2.1.2   Josephson effect and Josephson junctions

Another consequence of the macroscopic wave function is the tunneling of Cooper pairs between weakly coupled superconductors, such as two superconductors separated by a thin insulating barrier (Fig. 2.1B); this phenomenon, first described by B. D. Josephson [22], is known as the Josephson effect and the barrier is known as a Josephson junction. The Josephson junction is the active component in superconducting electronics [20]. In practice, there are many kinds of "weak links" that could be used [20,23]. The most common are insulating tunnel junctions, which are the type we will consider in this thesis.

We will give here a simple derivation of the Josephson relations, following [20]. Consider two superconductors with macroscopic wave functions

$$
\begin{aligned}
\psi_1 &= n_1^{1/2} \exp[i\theta_1] \\
\psi_2 &= n_2^{1/2} \exp[i\theta_2].
\end{aligned}
\tag{2.9}
$$

If the two superconductors interact with one another, we can write the time evolution of their wave functions as

$$
\begin{aligned}
i\hbar\frac{\partial \psi_1}{\partial t} &= U_1\psi_1 + K\psi_2 \\
i\hbar\frac{\partial \psi_2}{\partial t} &= U_2\psi_2 + K\psi_1,
\end{aligned}
\tag{2.10}
$$

where U_1 and U_2 are the energies of the wavefunctions and K represents the interaction between the superconductors. In general there will be some voltage difference between the superconductors; we will call this voltage difference V, so the energy difference can be written as $U_2 - U_1 = qV$, where q is the charge of a Cooper pair. We are free to write the energies as $U_1 = -qV/2$ and $U_2 = +qV/2$. With these values for U_1 and U_2, the real parts of Eq. 2.10 yield:

$$
\frac{\partial n_1}{\partial t} = -\frac{\partial n_2}{\partial t} = \frac{2}{\hbar}K(n_1 n_2)^{1/2}\sin\delta,
\tag{2.11}
$$

where $\delta = \theta_2 - \theta_1$ is the phase difference between the two superconductors. Equation 2.11 implies a current flowing from 2 to 1 which depends on the sine of the phase difference δ,

$$
I = I_c \sin\delta.
\tag{2.12}
$$

The value I_c is the critical current of the junction. This is the DC Josephson effect. It depends only on the phase difference between the superconductors, so it indicates that a DC current can flow across the junction in the absence of a voltage difference between the superconductors.

The energy difference qV emerges if we look at the imaginary parts of Eq. 2.10. In this case, it can be shown that

$$
\frac{\partial \delta}{\partial t} = \frac{2e}{\hbar}V.
\tag{2.13}
$$

This equation describes how a voltage applied across the weak link changes the phase difference δ.

The consequences of the Josephson effects are as follows: When a current $I < I_c$ is applied across the junction, Cooper pairs can tunnel across the barrier. This constitutes a supercurrent (often called the Josephson current), so the voltage across the junction is zero. If the current

Fig. 2.2: (A) Critical current vs. flux for a dc SQUID. (B) Current through SQUID vs. voltage across SQUID, with curve shown at two values of flux. Dotted line indicates the bias current. (C) Voltage across SQUID vs. flux.

across the junction exceeds the critical current, a voltage appears across the junction and the phase difference between the superconductors follows the AC Josephson relation (Eq. 2.13). It should be noted that high-quality junctions are often hysteretic (due to their self-capacitance) [23], so that once they have switched into the voltage state the current must be reduced well below I_c before the junction switches back to the superconducting state. This behavior is often undesirable [12] and can be eliminated with the addition of a suitable shunt resistor across the junction. The criteria for this resistor, and its effects on the behavior of the junction, are discussed in [12, 23].

2.1.3 The dc SQUID

The dc SQUID comprises two Josephson junctions in parallel on a superconducting ring, as shown in Fig. 2.1C. This device has a critical current I_c due to the Josephson junctions. A magnetic flux through the ring modulates the critical current with a period of one flux quantum (see Fig. 2.2A), due to quantum interference effects between the wave functions in each arm [21]. The value of the flux through the loop can then be found by determining the critical current I_c of the SQUID, for example by applying a bias current and increasing this current until I_c is reached, at which time a nonzero voltage appears across the SQUID. However, this method is effective only for a magnetic flux which remains a fixed value.

For SQUIDs with nonhysteretic (i.e. resistively-shunted) junctions, there is a better way to measure the magnetic flux through the loop [23]. In this method, a bias current I_b which is larger than the critical current I_c is applied across the SQUID, which leads to the creation of a voltage V across the SQUID. The magnitude of the voltage depends on the difference between bias current and critical current. As the flux through the SQUID changes, the critical current I_c is modulated with a period of Φ_0, as we discussed in the previous paragraph; the voltage V across the SQUID will also therefore be a periodic function with a period of Φ_0. This effect is illustrated in Fig. 2.2B-C. When operated in this way, the SQUID acts as a flux-to-voltage transducer. For measurement of small changes in magnetic flux, the SQUID can be biased with a DC magnetic flux to a point where the transfer function (the V-Φ curve) is steepest, and the voltage change

read out directly (Fig. 2.3A). However, this direct readout scheme has a limited dynamic range; the measured voltage is only an accurate measurement of magnetic flux when the it remains on the linear part of the IV curve. For a sinusoidal V-Φ relation, this linear flux range is about Φ_0/π [24]. Figure 2.3B shows the result of a direct readout when the applied flux is larger than $\Phi_0/2$. For measuring changes in magnetic flux outside the linear region, some technique to increase the dynamic range is required. One such technique, the flux-locked loop, will be considered in the next section.

2.2 The flux-locked loop

The aim of the flux-locked loop (FLL) [24] is to keep the flux through the SQUID loop constant, by the use of negative feedback to cancel applied flux. The FLL linearizes the SQUID response and allows for the measurement of very large and very small changes in flux, greatly increasing the dynamic range. The basic operation is as follows: the voltage across the SQUID is amplified, integrated, and then fed back into the SQUID as a flux (through a feedback resistor and feedback coil coupled to the SQUID). The output of the loop is the voltage across the feedback resistor. We will examine a basic FLL in order to better understand the basic operation. For simplicity we will neglect the effects of a time delay in the feedback.

A basic FLL is shown in Fig. 2.4A. The SQUID is biased with current I_b (which is greater than the critical current) and the applied flux through the loop is Φ_a. The voltage V_{SQ} across the SQUID is input to an amplifier of gain A. The output of the amplifier next goes through an integrator with gain $-1/i\omega R_I C_I$ where ω is the angular frequency. The output voltage is then

$$V_{out} = -V_{SQ}A\frac{1}{i\omega\tau_I}, \tag{2.14}$$

where $\tau_I = R_I C_I$ and the voltage across the SQUID is given by

$$V_{SQ} = \Phi_a \left.\frac{\partial V}{\partial \Phi}\right|_I = \Phi_a V_\Phi. \tag{2.15}$$

The voltage V_{out} across the feedback resistor R_F and the feedback coil generates a feedback flux Φ_F which is given by

$$\Phi_F = -\Phi_a V_\Phi A\frac{1}{i\omega\tau_I}\frac{M_F}{R_F} = G(\omega)\Phi_a, \tag{2.16}$$

where $G(\omega)$ is the open-loop gain of the FLL.

When the feedback is taken into account, we can find the closed-loop gain from:

$$\Phi_F = G(\omega)[\Phi_a + \Phi_F] \tag{2.17}$$

$$\Phi_F[1 - G(\omega)] = G(\omega)\Phi_a \tag{2.18}$$

$$\frac{\Phi_F}{\Phi_a} = \frac{G(\omega)}{1 - G(\omega)} = \frac{1}{1 - \frac{1}{G(\omega)}}. \tag{2.19}$$

We can make the substitution

$$\Phi_F = \frac{V_{out}}{R_F}M_F \tag{2.20}$$
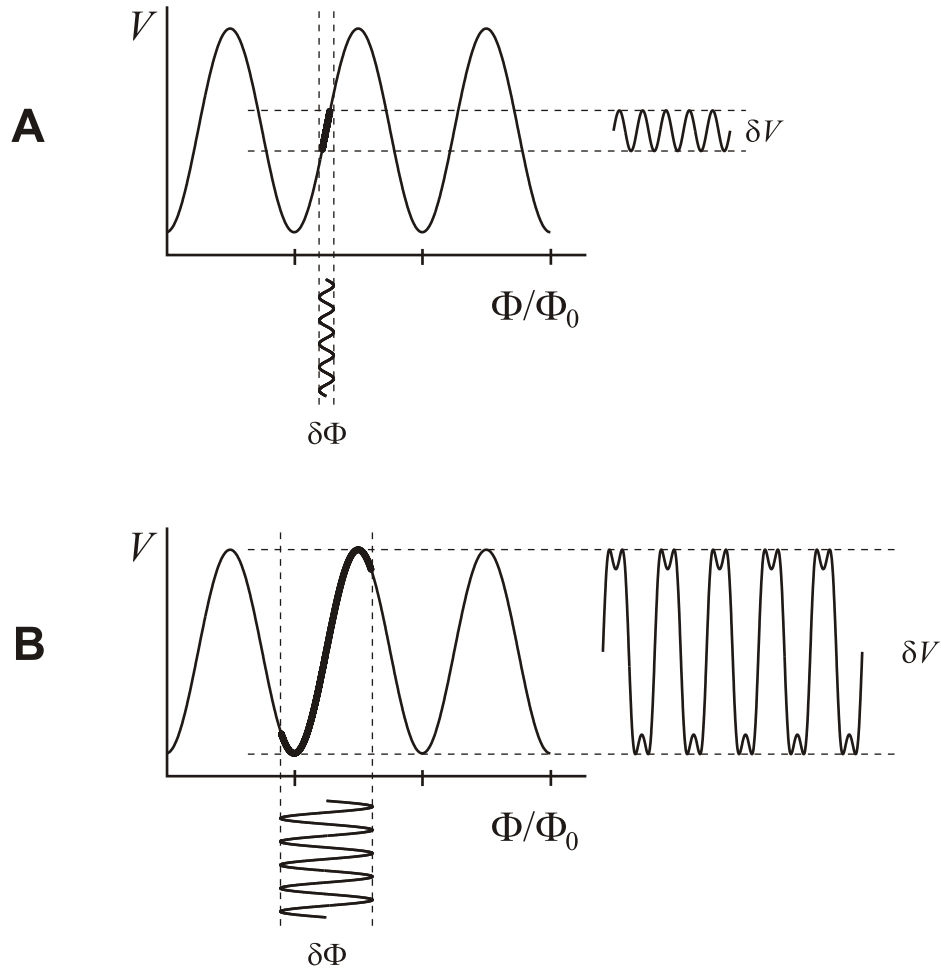
Fig. 2.3: Direct voltage readout of SQUID. (A) Small changes in flux, remaining in the linear region of the V-Φ curve. (B) Large changes in flux, which do not remain in the linear region of the V-Φ curve.
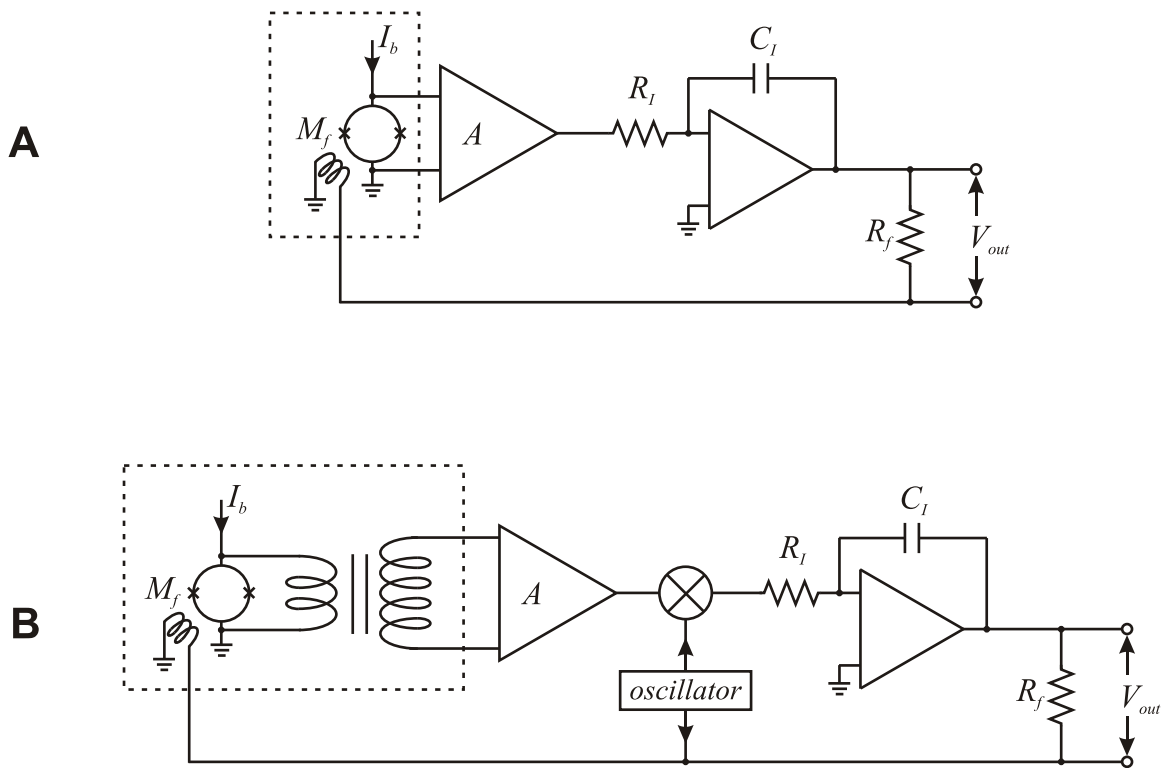
Fig. 2.4: Flux-locked loops. Components inside dotted lines are at cryogenic temperature. (A) Basic FLL. (B) FLL with modulation.

to find

$$\frac{V_{out}}{\Phi_a} = \frac{R_F}{M_F} \cdot \frac{1}{1 + (i\omega\tau_I R_F / V_\Phi A M_F)}. \tag{2.21}$$

The second term on the right-hand side of Eq. 2.21 is the transfer function for a single-pole low-pass filter,

$$H(if) = \frac{1}{1 + (if / f_1)}, \tag{2.22}$$

where the cutoff frequency f_1 is given by

$$f_1 = \frac{V_\Phi A M_F}{R_F} \frac{1}{2\pi R_I C_I}. \tag{2.23}$$

It is often the case that the noise of the preamplifier is a significant source of system noise when it is used directly across the SQUID, especially at low frequencies where $1/f$ noise of the preamplifier may be considerable. This problem can be alleviated by the addition of a step-up transformer between the SQUID and the preamplifier, and adding an AC flux modulation (which is applied via the feedback coil) [24]. The transformer serves to increase the voltage input to the preamplifier, and the turns ratio can be chosen to achieve the best impedance matching between the SQUID and the preamplifier. The modulation frequency is chosen to be well above the $1/f$ knee of the preamplifier. In this case, the modulated signal is amplified and then passed through a lock-in detector (such as a mixer) to recover the demodulated signal. The demodulated signal is then integrated and fed back to the SQUID to cancel applied flux. A FLL with flux modulation is illustrated in Fig. 2.4B.

2.3 Flux transformer

In conventional operation, the voltage across the SQUID is proportional to the magnetic flux through the SQUID loop. However, modern SQUIDs such as those used in this thesis are very small; our SQUIDs were squares approximately one millimeter on a side. A device of this size obviously cannot look at a large area. Directly measuring with the SQUID has other disadvantages as well, among which are: random signals in the environment can degrade their performance (or even damage the SQUID); the SQUID cannot discriminate between fields due to desired nearby sources and those due to distant undesired sources; a large or rapidly changing flux can create high currents which would drive the loop out of the superconducting state and may cause significant heating and damage; and the input cannot be tuned to detect at a particular frequency. Fortunately, all of these issues can be addressed by use of a separate pickup coil which is coupled inductively to the SQUID. This separate coil is called a flux transformer.

The flux transformer [25] can take many different forms. The simplest of these has a single superconducting pickup loop and a superconducting multiturn input coil which is tightly coupled to the SQUID (usually patterned directly on the SQUID chip), as shown schematically in Fig. 2.5A. When a field is applied to the pickup coil, a current flows in the flux transformer; this occurs even for static fields because the flux transformer is superconducting. The effect of the flux transformer is to increase the field sensitivity of the SQUID by increasing the effective sensing area of the loop. If the pickup loop has area A_p and inductance L_p, and the mutual inductance between the input

Fig. 2.5: Flux transformers with RC shunts in parallel with input coil. Dotted lines indicate superconducting shield. (A) Pickup coil configured as untuned magnetometer. (B) Pickup coil configured as untuned first-order axial gradiometer. (C) Pickup coil configured as tuned first-order axial gradiometer.

coil and the SQUID is $M_i = \alpha\sqrt{L_i L_{SQ}}$ (where L_i and L_{SQ} are the inductances of the input coil and SQUID, respectively, and α is the coupling constant), then the effective area of the SQUID magnetometer is

$$A_{eff} = A_p \frac{\alpha\sqrt{L_i L_{SQ}}}{L_i + L_p}. \tag{2.24}$$

Using a flux transformer, it is possible to obtain an increase in effective area of several orders of magnitude. An additional advantage of using a flux transformer, also shown in Fig. 2.5, is that the SQUID and input coil can be placed into a superconducting shield to protect them from environmental interference.

If the flux transformer instead has two coaxial loops of equal area, wound in opposite directions and separated by a baseline (Fig. 2.5B), it will be sensitive only to the first derivative of the field along the direction of its axis. This arrangement is a first-derivative axial gradiometer. The purpose of such an arrangement is to reject uniform (or nearly uniform) fields from distant sources, while remaining sensitive to fields from nearby sources, such as the sample being measured.

The use of a separate coil also allows one to insert components into it to condition the signal before it is coupled into the SQUID. For example, it is a common practice to insert an RC network in parallel with the input coil (shown in Fig. 2.5). This network acts as a shunt for undesired radiofrequency signals. In addition, the flux transformer can be tuned to a particular frequency through the addition of a capacitor in series with the pickup coil, as in Fig. 2.5C. One can also add protective components to prevent the formation of dangerously high currents; we will consider one type of protective component in the next section.

2.4 Current-limiting Josephson junction array

In some experiments, such as the NMR and MRI experiments described in this thesis, the pickup gradiometer is placed as close as possible to the sample to maximize the received signal. The nuclear spins in the sample are manipulated via magnetic field pulses; in some cases the amplitude of the pulses approach 100 mT at the pickup coil, and they are oriented along the axis of the gradiometer [26]. An untuned gradiometer subjected to such large fields would develop very large currents and lead to very large fields at the SQUID; if the flux transformer were to remain superconducting, Lee [27] estimates the current in the gradiometer would exceed 800 A and the magnetic field at the SQUID would be approximately 35 T. In reality, the flux transformer (and possibly the SQUID) would be driven normal, likely resulting in damage to the system.

A series array of Josephson junctions added to the flux transformer (shown in Fig. 2.6) is an effective means of protecting the gradiometer, input coil, and SQUID from excessive current due to large field pulses [27]. The junction array works as follows: consider a series array of Josephson junctions, each with critical current I_c which is lower than that of the flux transformer. As long as the current in the flux transformer is below I_c, such as during signal acquisition, the junctions are in the superconducting, zero-voltage state and do not interfere with the measurement. When the current increases above I_c, such as during a strong field pulse, the junctions switch into the voltage state, where each junction has a resistance R. The resistance limits and damps the current; when the current drops sufficiently, the junctions switch back to the superconducting state. Thus the junction array acts as a self-resetting fuse.

Fig. 2.6: Flux transformer with Josephson junction array in series with input coil. For simplicity, only three junctions are shown.

Chapter 3

Introduction to Nuclear Magnetic Resonance

Nuclear magnetic resonance (NMR), the interaction between an external magnetic field and a nucleus with a magnetic moment, is the phenomenon on which magnetic resonance imaging is built. Nearly all of the elements in the periodic table have at least one isotope in which the nucleus possesses a magnetic moment, and hence can be used in NMR. However, it is the hydrogen nucleus, composed of a single proton, which is most often studied; the hydrogen nucleus is the only one we will be concerned with in this thesis. This chapter is intended as a brief introduction to the NMR concepts which will be used in the zero-field MRI experiment; as such, we will omit some topics which are important in conventional NMR, but which are irrelevant for our purposes. Some of the topics we will omit are excitation pulses, chemical shift, and spin echoes; information on these topics can be found in the references [1, 28–30].

3.1 Polarization of sample in magnetic field

According to the Zeeman effect, a particle with a magnetic moment $\boldsymbol{\mu}$ placed in an external magnetic field $\mathbf{B} = B\hat{\mathbf{z}}$ has a Hamiltonian given by

$$H = -\boldsymbol{\mu} \cdot \mathbf{B} = -\gamma\hbar\mathbf{I} \cdot \mathbf{B} = -\gamma\hbar I_z B, \tag{3.1}$$

where γ is the magnetogyric ratio and $\hbar\mathbf{I}$ is the angular momentum of the nucleus [28]. The eigenvalues of this Hamiltonian are

$$E_m = -m\gamma\hbar B, \tag{3.2}$$

where m is the spin quantum number with allowed values $m = -I, -I+1, ..., I-1, I$ [28]. In particular, for a nucleus of spin $I = 1/2$ (such as that of hydrogen), the available values of the spin quantum number are $m = \pm 1/2$ and the energy levels are

$$E = \pm\frac{\gamma\hbar B}{2}, \tag{3.3}$$

as shown in Fig. 3.1. The lower-energy state corresponds to the situation in which the magnetic moment is aligned parallel with the external field; the higher-energy state corresponds to a magnetic moment aligned antiparallel with the field.

Fig. 3.1: Energy level splitting for a spin-1/2 nucleus in a magnetic field.

Because the parallel arrangement has lower energy, at thermal equilibrium the probability will be higher that the nucleus will be in this state. At a temperature T, the probability of finding the nucleus in a given state with energy E_m is given by [31]

$$P(E_m) = \frac{\exp(-E_m/k_B T)}{\sum\limits_{m} \exp(-E_m/k_B T)} \tag{3.4}$$

where k_B is Boltzmann's constant. For a large collection of such nuclei, some fraction will be in the spin-up state and some will be in the spin-down state. The ratio between the number of nuclei in each state is

$$\frac{N_{+1/2}}{N_{-1/2}} = \frac{\exp(\gamma\hbar B/2k_B T)}{\exp(-\gamma\hbar B/2k_B T)} = \exp(\gamma\hbar B/k_B T). \tag{3.5}$$

For a sample at $T = 300$ K in a field of $B = 1$ T, and with the magnetogyric ratio of hydrogen $\gamma/2\pi = 42.58$ MHz/T, this ratio is approximately 1.000007 — that is, the population difference is a few parts in 10^6. Although the difference in the number of spin-up nuclei versus spin-down nuclei is small, due to the large number of nuclei present it results in a macroscopic magnetization, which can be found using

$$M_0 \approx \frac{\rho_0 \gamma^2 \hbar^2 B}{4k_B T}, \tag{3.6}$$

where ρ_0 is the spin density in the sample [1].

3.2 Bloch equation

Although the interaction of nuclear spins with magnetic fields is inherently a quantum phenomenon, the semiclassical Bloch equation [29] provides an intuitive and easy-to-understand picture of the behavior of a magnetic moment in a magnetic field.

The Bloch equation is

$$\frac{d\mathbf{M}}{dt} = \gamma\mathbf{M} \times \mathbf{B} - \frac{(M_z - M_0)\hat{\mathbf{z}}}{T_1} - \frac{M_x\hat{\mathbf{x}} + M_y\hat{\mathbf{y}}}{T_2}, \tag{3.7}$$

where M_x, M_y, and M_z are the components of the bulk magnetization vector, M_0 is the equilibrium value of magnetization due to the external field (Eq. 3.6), and T_1 and T_2 are the so-called longitudinal and transverse relaxation times, respectively. The first term arises from the torque on the magnetic moment due to the external magnetic field; the second and third terms describe relaxation processes due to interactions between the spin and its environment, as well as interactions between spins. Each of these terms will be discussed in further detail below.

3.2.1 Torque on magnetic moment due to external field

A magnetic moment $\boldsymbol{\mu}$ placed in an external magnetic field $\mathbf{B}$ will experience a net torque which is given by [1]

$$\mathbf{N} = \boldsymbol{\mu} \times \mathbf{B}. \tag{3.8}$$

Torque is the time rate of change of angular momentum, so it follows that

$$\frac{d\mathbf{J}}{dt} = \boldsymbol{\mu} \times \mathbf{B}. \tag{3.9}$$

The magnetic moment of the nucleus is given by $\boldsymbol{\mu} = \gamma \hbar \mathbf{I} = \gamma \mathbf{J}$, which yields

$$\frac{d\boldsymbol{\mu}}{dt} = \gamma \boldsymbol{\mu} \times \mathbf{B}. \tag{3.10}$$

For a large collection of magnetic moments, bulk magnetization is found by summing all of the magnetic moments in a unit volume [28], so we find

$$\frac{d\mathbf{M}}{dt} = \gamma \mathbf{M} \times \mathbf{B}. \tag{3.11}$$

We will examine the implications of this term using a simple example — that of a sample in a uniform, static magnetic field $\mathbf{B_0} = B_0 \hat{\mathbf{z}}$. If $\mathbf{M}$ and $\mathbf{B_0}$ are not parallel, the solution to Eq. 3.11 is

$$\begin{aligned}
M_x(t) &= M_x(0) \cos \omega_0 t + M_y(0) \sin \omega_0 t \\
M_y(t) &= M_y(0) \cos \omega_0 t - M_x(0) \sin \omega_0 t \\
M_z(t) &= M_z(0),
\end{aligned} \tag{3.12}$$

where $\omega_0 = \gamma B_0$. The result is that $\mathbf{M}$ precesses about $\mathbf{B_0}$ at the Larmor frequency ω_0. This is illustrated in Fig. 3.2.

3.2.2 Longitudinal Relaxation

Our sample of spins in a magnetic field will be in thermal equilibrium — the state with the lowest energy — when the magnetization is aligned with the external field and of the magnitude given by Eq. 3.6. In any other state, the spins will give up energy to the environment in order to return to equilibrium. While the actual mechanisms for this energy transfer to the environment are due to random processes on individual nuclei and may be complicated, Bloch [29] found that a nonequilibrium magnetization M_z returns to the equilibrium value M_0 exponentially with a characteristic time constant which he called T_1:

$$\frac{dM_z}{dt} = -\frac{(M_z - M_0)}{T_1}. \tag{3.13}$$

Fig. 3.2: Precession of **M** around **B₀**.

The solution to Eq. 3.13 is given by

$$M_z(t) = M_0 + (M_z(0) - M_0) e^{-t/T_1}. \tag{3.14}$$

Figure 3.3 shows a plot of $M_z(t)$ for $M_z(0) = 0$.

3.2.3 Transverse relaxation

The individual magnetic moments interact with the local fields created by other magnetic moments. This so-called spin-spin coupling creates field inhomogeneities, with the result that each magnetic moment sees a slightly different field and therefore precesses at a slightly different frequency. Over time this leads to a loss of coherence which reduces the apparent magnetization in the transverse plane (the plane of precession), but does not affect the longitudinal component of the magnetization (and therefore does not affect the energy of the system).

The loss of transverse magnetization manifests as an exponential decay with time constant T_2:

$$\frac{dM_{x,y}}{dt} = -\frac{M_{x,y}}{T_2}. \tag{3.15}$$

Fig. 3.3: Plot of $M_z(t)$ versus t.

Fig. 3.4: Receiving coil in transverse plane.

This differential equation has solution

$$M_{x,y}(t) = M_{x,y}(0)e^{-t/T_2}. \tag{3.16}$$

It is important to note that longitudinal relaxation also affects the transverse magnetization. Magnetization that has returned to the longitudinal direction is no longer available in the transverse plane, so $T_2 \leq T_1$.

It is apparent that inhomogeneities in the external field will also cause a spread in precession frequencies and a corresponding loss of phase coherence. This leads to a decrease in the time constant T_2. In practice, the dephasing due to external factors is described by a separate time constant T_2'. The total effective time constant T_2^* is then given by

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{1}{T_2'}, \tag{3.17}$$

where T_2 is the intrinsic spin-spin relaxation time constant.

3.3   Signal acquisition

The magnitude of the transverse magnetization can be measured by placing a coil perpendicular to the transverse plane as shown in Fig. 3.4. As the magnetization vector precesses, the changing magnetic field in the coil induces a voltage across the loop which is measured versus time. In conventional NMR systems, the receiving coil is a copper coil which is part of a tank circuit tuned to the precession frequency. One can also use a superconducting magnetometer or gradiometer which is coupled to the input coil of a SQUID (see Ch. 2).

When the sample is in a uniform static field, the received time-domain signal has the form of an exponentially decaying sinusoid as shown in Fig. 3.5A. The time constant of the exponential decay is T_2^*. The time-domain signal is known as a free induction decay (FID). The FID is subsequently inverse-Fourier-transformed to find the frequency spectrum. In the case of a uniform static field, the frequency spectrum will consist of a single Lorentzian peak at the Larmor

A

B

Fig. 3.5: (A) Free induction decay showing sinusoid at ω_0 which decays exponentially with time constant T_2^*. (B) Fourier transform of the FID, showing Lorentzian peak at ω_0 with peak width proportional to $1/T_2^*$.

frequency with a width proportional to $1/T_2^*$ (see Fig. 3.5B). (It should be noted that the time-domain signal is usually acquired discretely. A discussion of discrete time sampling and its effect on the frequency-domain signal will be deferred to Ch. 4.)

3.4   How to get magnetization in the transverse plane

After a sample has been polarized in a field along $\hat{\mathbf{z}}$ the magnetization must be tipped away from the z-axis in order to observe the Larmor precession. The conventional way to tip the magnetization is to apply a resonant pulse — that is, a short-lived rotating magnetic field in the transverse plane which is tuned to the Larmor frequency [1, 28]. However, this technique was not used in the present experiments.

Another technique for tipping the magnetization into the transverse plane is simply to create it there by applying a polarizing field in the plane, rather than along the z-axis. This polarizing field is applied until a suitable magnetization develops. If the direction of the applied field changes nonadiabatically to the $\hat{\mathbf{z}}$ direction, the magnetization will precess around it. In practice this is most commonly achieved by using a polarizing field which is much larger than the precession field; when both fields are applied, the polarizing field will dominate.

The change in field must be nonadiabatic; that is, it must satisfy the requirement

$$\frac{dB}{dt} \gg \gamma B^2, \tag{3.18}$$

where B is the instantaneous magnitude of the magnetic field [30].

3.5 Quantum-mechanical operators

The semiclassical picture we have discussed provides a convenient and intuitive way of thinking about the dynamics of magnetic resonance. Magnetic resonance remains, however, an inherently quantum phenomenon as it deals with the mechanics of nuclear spins. It is instructive, therefore, to briefly examine it from a quantum-mechanical point of view. This is helpful when analyzing more complicated experiments, such as the "zero-field MRI" experiment discussed in Ch. 5.

As mentioned in Sec. 3.1, the Hamiltonian of a particle which has a magnetic moment and is in an external field is given by Eq. 3.1:

$$H = -\boldsymbol{\mu} \cdot \mathbf{B} = -\gamma \hbar \mathbf{I} \cdot \mathbf{B}.$$

For a static field, this Hamiltonian is not time-dependent, so the exponential operator is simply [28]

$$U = \exp(i\gamma t \mathbf{I} \cdot \mathbf{B}). \tag{3.19}$$

Chapter 4

Introduction to Magnetic Resonance Imaging

A spin in a magnetic field will precess, with a frequency which is proportional to the field at that point. If a spatial variation is applied to the magnetic field, the frequency and phase of the precession can be made into functions of position. The relative signal amplitudes at different frequencies and phases can then be used to form an image of the real-space spin distribution. This is the basis of MRI.

In this chapter we will give a basic introduction to MRI, including the Fourier-transform relationship between the time-domain signal and the real-space image, the concept of k-space, and common imaging techniques. We will also briefly examine the concept of the rotating reference frame.

4.1   Spatial encoding

The magnetic field is made spatially dependent by the application of magnetic field gradients. In this chapter, we will consider only static gradients which can be switched on and off, but which do not vary in time while they are on. Ideally these gradients are linear and unidirectional, so that the total field can be written as

$$\mathbf{B}(x, y, z) = \mathbf{B_0} + \mathbf{B_G}(x, y, z) = \left(B_0 + G_x x + G_y y + G_z z\right) \hat{\mathbf{z}}, \qquad (4.1)$$

where G_x, G_y, and G_z are defined as

$$G_x = \frac{\partial B_z}{\partial x} = const, \qquad (4.2)$$

$$G_y = \frac{\partial B_z}{\partial y} = const, \qquad (4.3)$$

$$G_z = \frac{\partial B_z}{\partial z} = const. \qquad (4.4)$$

Note that the ideal applied gradients change the magnitude of the field, but do not change the direction. The ideal G_z gradient is illustrated in Fig. 4.1.

Fig. 4.1: Ideal imaging gradient G_z. Length of vectors indicate magnitude of magnetic field.

It is clear then that while the gradients are applied, the precession frequencies of the spins in the sample will be dependent on their position, so that

$$\omega_0(x, y, z) = \gamma \left(B_0 + G_x x + G_y y + G_z z \right). \tag{4.5}$$

Because of the spatial dependence of precession frequency, information about the spin distribution in the sample is encoded in the frequency distribution of the acquired signal. An example of this for a single gradient G_z is shown schematically in Fig. 4.2. This is referred to as frequency encoding. If the gradient is applied only for some finite time and then removed, the spins at all points will again precess at the same frequency ω_0, but the phase of the precession will now depend on position. For example, if a static gradient G_z is applied for a time τ, the phase of precession is given by

$$\phi(z, \tau) = \int_0^\tau \gamma G_z z \, dt = \gamma G_z z \tau. \tag{4.6}$$

The different kinds of magnetic field gradients are created by different kinds of coils. Although there are many possible gradient coil geometries, the G_z gradients are commonly generated by a Maxwell pair while the G_x and G_y gradients are provided by Golay ("saddle") coils [32]. It must be noted, though, that Maxwell's equations prohibit the creation of unidirectional magnetic field gradients. A field with a gradient in one direction must be accompanied by field components of comparable magnitude in at least one other direction. The actual magnetic field created by a Maxwell pair is

$$\mathbf{B_{Maxwell}}(x, y, z) = \frac{G}{2}(-x\hat{\mathbf{x}} - y\hat{\mathbf{y}} + 2z\hat{\mathbf{z}}), \tag{4.7}$$

while the fields of the Golay pairs, which create gradients along x and y respectively, are

$$\mathbf{B_{Golay}}(x, z) = G(z\hat{\mathbf{x}} + x\hat{\mathbf{z}}) \tag{4.8}$$

Fig. 4.2: Effect of magnetic field gradients on NMR signal. (A) In a uniform magnetic field, all spins in an object will precess at the same frequency, so the frequency spectrum is a single peak. (B) When the field amplitude depends on position, the frequency spectrum yields a projection of the object along the direction of the field gradient.

and

$$\mathbf{B_{Golay}}(y, z) = G(z\hat{\mathbf{y}} + y\hat{\mathbf{z}}), \tag{4.9}$$

where G is a constant. The fields along the x- and y-axes are referred to as "concomitant gradients." The concomitant gradients affect both the magnitude and direction of the field. The concomitant gradients can, in general, be safely ignored if the field change due to the gradients is small compared to the B_0 field; the gradient fields are "truncated" by the B_0 field. We will assume this is the case throughout this chapter. The concomitant gradients and their effects will be explored in more detail in Ch. 5.

4.2 Signal equation and *k*-space

The spatial information about the sample is encoded in the frequency and phase of the precessing magnetization. However, we detect and record the signal from the precessing magnetization in the time domain. In order to convert between the two, we will use a simple physical model to look at how the received signal is constructed [33]. For simplicity, we will limit our discussion to a one-dimensional case; the extension to three dimensions is straightforward.

Consider a one-dimensional distribution of spins, which we call $m(z)$. The signal received from a single point z in this distribution can be written as

$$s_r(z, t) = m(z)e^{-i\phi(z,t)}, \tag{4.10}$$

where the phase $\phi(z,t)$ is given by the integral of the frequency:

$$\phi(z,t) = \int_0^t \omega(z,t')\,dt' = \omega_0 t + \gamma z \int_0^t G_z(t')\,dt'. \tag{4.11}$$

Assuming a uniform detector sensitivity over the sample, the total received signal is the sum of the signals received from all points, so

$$s_r(t) = \int_z m(z)e^{-i\phi(z,t)}\,dz$$

$$= \int_z m(z)\exp\left[-i\omega_0 t\right]\exp\left[-i\gamma z \int_0^t G_z(t')\,dt'\right]\,dz. \tag{4.12}$$

The $\exp[-i\omega_0 t]$ term represents a carrier frequency, which can be removed by demodulation to leave

$$s_r(t) = \int_z m(z)\exp\left[-i\gamma z \int_0^t G_z(t')\,dt'\right]\,dz. \tag{4.13}$$

The signal equation now has the form of a Fourier transform

$$s_r(t) = \int_z m(z)\exp\left[-i2\pi k_z z\right]\,dz, \tag{4.14}$$

where the Fourier coefficient k_z is

$$k_z(t) = \frac{\gamma}{2\pi}\int_0^t G_z(t')\,dt'. \tag{4.15}$$

The extension of this analysis to higher dimensions is straightforward. In three dimensions, Eq. 4.14 is

$$s_r(t) = \int_x \int_y \int_z m(x,y,z)\exp\left[-i2\pi(k_x x + k_y y + k_z z)\right]dx\,dy\,dz, \tag{4.16}$$

with

$$k_\alpha(t) = \frac{\gamma}{2\pi}\int_0^t G_\alpha(t')\,dt'. \tag{4.17}$$

for $\alpha = x$, y, or z. Spatial information is recovered from the time-domain signal by means of the Fourier transform with time-dependent k values. The demodulated time domain signal is referred to as "k-space."

In modern MRI, the k-space data is invariably digitized and fed into a computer for storage and processing. As a result, the k-space data is a collection of discrete points and the recovery of the spin distribution $m(x,y,z)$ is done via discrete Fourier transforms (DFTs). Due to the DFT,

if the k-space sampling period along the direction α (where $\alpha = x$, y, or z as in Eq. 4.17) is Δk_α, then the resulting image will have a field of view (FOV) along α given by

$$FOV_\alpha = \frac{1}{\Delta k_\alpha} = \frac{1}{\frac{\gamma}{2\pi} G_\alpha \Delta t}, \tag{4.18}$$

where G_α is the gradient along α and Δt is the sampling period. If the extent of the spin distribution is larger than the FOV, the image will be aliased [1,33]. To avoid aliasing, the Δk_α must be adjusted (by setting G_α and/or Δt) so that the FOV is larger than the extent of the spin distribution; this is simply a statement of the Nyquist sampling criterion.

The resolution of the reconstructed image along the direction α depends on the maximum acquired value of k-space, which we call $k_{\alpha,max}$. The quantity

$$\Delta\alpha = \frac{1}{k_{\alpha,max}} \tag{4.19}$$

is the spatial step size along the α direction in the image, which represents the smallest feature that can be resolved [1]. In reality, the smallest feature that can actually be resolved is likely to be larger due to one or more of the many effects which can blur the image, thus reducing the resolution. These includes intrinsic effects which broaden the NMR line, such as diffusion and T_2^* decay, as well as effects from processing such as signal filtering. For example, the actual resolution of an image is not improved by extending acquisition to k-space values well beyond the point at which the signal has decayed away, even though the spatial step size $\Delta\alpha$ will be smaller. Though it does not increase actual resolution, a similar technique known as "zero-filling" or "zero-padding" is commonly used to interpolate and smooth images. In this technique, the k-space is extended to values much larger than those actually acquired, and the extra k-space is filled with zeros. This reduces the spatial step size, but does not increase the resolution. The effect is to smooth or interpolate the image data [1].

4.3 Frequency and phase encoding: k-space perspective

Figure 4.3 shows a pulse sequence for a one-dimensional frequency-encoding experiment utilizing a prepolarizing pulse which is switched off non-adiabatically (see Sec. 3.4). After the polarizing pulse, a uniform static field B_0 is applied along with a gradient field $B_G = -G_z z$. The gradient is applied in the negative direction to form a so-called "gradient echo" [33], which is used to begin signal acquisition in negative k-space. After a time t_1, the gradient is reversed so that $B_G = G_z z$ and the data acquisition is started; at some future time t_2 the acquisition is stopped. The first data point will be at $k_z(t_1) = -(\gamma/2\pi)G_z t_1$ and the final point will be at $k_z(t_2) = (\gamma/2\pi)G_z(t_2 - t_1)$. The number of intermediate points, and therefore the value of Δk_z between points will depend on the sampling rate of the detector.

A one-dimensional phase encoding pulse sequence is displayed in Fig. 4.4. In this case, acquisition occurs in B_0 after the gradient pulses have ended. The inverse Fourier transform of the acquired signal will yield the data for the single k-space point $k_z(t_1) = (\gamma/2\pi)G_z t_1$. Multiple points in k-space can be obtained by running the sequence multiple times, changing the value of t_1 each time. The Δk_z between points is determined by the timing selected in each sequence.

Fig. 4.3: Pulse sequence for prepolarized frequency encoding. After the prepolarizing field B_p is shut off nonadiabatically, a gradient field B_G is applied until time t_1, when it is flipped. Acquisition occurs from t_1 to t_2. The uniform field B_0 is assumed to be on during the entire sequence.

Fig. 4.4: Pulse sequence for prepolarized phase encoding. After the prepolarizing field B_p is shut off nonadiabatically, a gradient field B_G is applied until time t_1. Acquisition occurs from t_1 to t_2. The uniform field B_0 is assumed to be on during the entire sequence.

Fig. 4.5: Unidirectional gradient at an angle θ with respect to the z-axis, for use in projection reconstruction imaging.

4.4 Common imaging methods: Projection reconstruction and 2D Fourier Transform

Here we will briefly introduce and discuss two common imaging schemes: projection reconstruction and 2D Fourier Transform methods.

In projection reconstruction imaging, a linear magnetic field gradient is produced along some direction. In the y-z plane, for example, a gradient G_θ in an arbitrary direction θ with respect to the z-axis is produced by using a Maxwell pair to create a gradient with magnitude $G_\theta \cos\theta$ along the z-axis, and a Golay pair to create a gradient with magnitude $G_\theta \sin\theta$ along the y-axis. This arrangement is illustrated in Fig. 4.5. This gradient is used to obtain a one-dimensional projection (usually via frequency encoding, but phase encoding can be used as we shall see in Ch. 6). After acquiring projections at multiple values of θ, the image is reconstructed by the process of filtered back-projection [1]; in short, each projection is "smeared out" along its angle θ though the imaging space. The intersections of multiple projections will have a higher value; the highest values, where all of the projections intersect, will correspond with the object that was imaged (Fig. 4.6). The resulting image is then filtered to smooth it.

Alternatively, one can perform 2D Fourier Transform imaging [1] by combining frequency encoding along one direction in k-space with phase encoding along another direction in k-space. For example, after polarization, one could apply a gradient G_y for a short time in order to phase encode to a particular value of k_y. A gradient G_z could then be applied to acquire the k-space data along k_z, as shown in Fig. 4.7. The process is repeated for as many different k_y values as desired. The result is a two-dimensional data set in the (k_y, k_z) domain; performing a two-dimensional inverse Fourier transform on this k-space data yields the real-space image (Fig. 4.8).

Fig. 4.6: Example of projection reconstruction using two projections. (A) White circles represent object to be imaged. Two projections (for $\theta = 0°$ and $90°$) are shown. (B) Reconstruction by "smearing out" the projections.

Fig. 4.7: Example of 2DFT imaging. (A) Sequence showing encoding to a particular k_y value and applying a gradient echo to read out k_z. Signal is acquired from t_1 to t_2. Uniform field B_0 is applied during entire sequence. (B) k-space trajectory for the sequence shown in A. Location in k-space at t_1 and t_2 are indicated.

Fig. 4.8: Example of k-space. (A) White circles represent object to be imaged. (B) k-space data for the real-space object shown in A.

Fig. 4.9: Laboratory reference frame (x, y, z) and rotating reference frame (x', y', z).

4.5 The rotating reference frame picture

An important concept in NMR and MRI is that of the reference frame which rotates about the z-axis with frequency Ω, as illustrated in Fig. 4.9. In this frame, there will be an effective field in the z-direction given by

$$B_{z,rot} = B_{z,lab} - \frac{\Omega}{\gamma}, \tag{4.20}$$

where $B_{z,lab}$ is the z component of the magnetic field in the lab frame (Fig. 4.10). This concept is most useful when $\Omega = \omega_0 = \gamma B_0$. When only the B_0 field is present in the lab frame, $B_{z,rot} = 0$ so the spins are stationary in the rotating frame (Fig. 4.11). When a gradient is applied which changes the magnitude of the total field (but does not change its direction), the effective field along z varies with position, so the spins precess in the rotating frame with a frequency and direction which depends on the magnitude and polarity of the effective field at that point. This is the effect of demodulating the signal equation in Sec. 4.2. In the classical picture of MRI, the change into the rotating frame is formally performed by multiplying vectors in the lab frame by a rotation matrix about the z-axis [33]

$$\mathbf{R_z}(\Omega t) = \begin{bmatrix} \cos \Omega t & \sin \Omega t & 0 \\ -\sin \Omega t & \cos \Omega t & 0 \\ 0 & 0 & 1 \end{bmatrix}. \tag{4.21}$$

Quantum mechanically, the transition into the rotating frame is accomplished in the following way. Recall that the Hamiltonian of a spin in a magnetic field $\mathbf{B}$ is

$$H = -\gamma \mathbf{I} \cdot \mathbf{B}. \tag{4.22}$$

In MRI, the field is usually a uniform static field plus additional fields due to gradients. The additional fields may in general be time-dependent, so we can write the Hamiltonian in the laboratory frame as

$$H = H_0 + H_1(t) = -\omega_0 I_z + H_1(t). \tag{4.23}$$

Fig. 4.10: Effective field in the rotating reference frame. (A) Effective field at a point where $B_{z,lab} > B_0$. (B) Effective field at a point where $B_{z,lab} < B_0$.

Fig. 4.11: (A) A magnetization vector $\mathbf{M}$ in a uniform field $\mathbf{B_0}$ in the laboratory frame. $\mathbf{M}$ precesses about $\mathbf{B_0}$ at angular frequency $\omega_0 = \gamma B_0$. (B) A magnetization vector $\mathbf{M}$ in a uniform field $\mathbf{B_0}$ in the rotating frame. $\mathbf{M}$ is stationary, indicating zero effective field.

The Hamiltonian of the system in the rotating frame is given by the transformation of the perturbing Hamiltonian $H_1(t)$ into the interaction representation [11, 34], so that

$$H_{rot} = \exp\left[\frac{iH_0t}{\hbar}\right] H_1(t) \exp\left[\frac{-iH_0t}{\hbar}\right].$$ (4.24)

The evaluation of this equation is aided by the use of the Baker-Hausdorff lemma [34]

$$\exp[iG\lambda]A\exp[-iG\lambda] = A + i\lambda[G, A] + \left(\frac{i^2\lambda^2}{2!}\right)[G, [G, A]]+$$

$$\cdots + \left(\frac{i^n\lambda^n}{n!}\right)[G, [G, [G, \ldots [G, A]]] \ldots] + \ldots$$ (4.25)

where G and A are operators and λ is a real parameter.

By way of example, we will consider two simple cases in the quantum mechanical representation of the rotating frame. First, we will examine a sample in a uniform static field. This situation is trivial; the Hamiltonian in the lab and rotating frames are

$$H_{lab} = H_0 = -\omega_0 I_z,$$ (4.26)

$$H_{rot} = 0$$ (4.27)

because $H_1(t) = 0$.

The situation becomes only slightly more complicated with the addition of a static, linear, unidirectional gradient (such as the "ideal" gradients used for imaging). If an imaging gradient along z is applied (see Eq. 4.4), then the Hamiltonians become

$$H_{lab} = H_0 - \gamma z G_z I_z,$$ (4.28)

$$H_{rot} = -\gamma z G_z I_z.$$ (4.29)

The rotating frame Hamiltonians are not as simple when **B** has components perpendicular to the z-axis; we will consider cases like this in Ch. 5.

Chapter 5

Concomitant Gradients and the Zero-Field MRI Sequence

We saw in Ch. 4 that for encoding we assume a simplified form of the gradient fields in which we ignore the field components perpendicular to the uniform field $\mathbf{B_0}$. This simplification was justified by the assumption that the magnitude of the uniform field was much larger than the maximum magnitude of the transverse components of the gradient fields. But as the uniform field magnitude is reduced, as for microtesla MRI, this assumption may not hold. In this case, the so-called concomitant gradients will affect the image. These concomitant gradients set an upper limit on the maximum gradient that can be used for a given uniform field magnitude; conversely, they can also be considered as setting a lower limit on the uniform field magnitude which can be used for a given gradient.

In this chapter we will examine the origin and effect of the concomitant gradients as well as their effect on the imaging field, and on the images. We will then discuss a pulse sequence which removes the concomitant gradients by averaging them out, allowing for undistorted images even in the limit that B_0 approaches zero.

5.1   Origin of concomitant gradients

The total magnetic field is made spatially dependent by the addition of one or more gradient fields to the uniform field $\mathbf{B_0} = B_0\hat{\mathbf{z}}$. The ideal gradient fields are

$$\mathbf{B_G^{x,ideal}} = \left(\frac{\partial B_z}{\partial x}\right) x\hat{\mathbf{z}} = G_x x\hat{\mathbf{z}}, \tag{5.1}$$

$$\mathbf{B_G^{y,ideal}} = \left(\frac{\partial B_z}{\partial y}\right) y\hat{\mathbf{z}} = G_y y\hat{\mathbf{z}}, \tag{5.2}$$

$$\mathbf{B_G^{z,ideal}} = \left(\frac{\partial B_z}{\partial z}\right) z\hat{\mathbf{z}} = G_z z\hat{\mathbf{z}}, \tag{5.3}$$

where the G's are constants. Note that these fields are all unidirectional along z, so they are parallel to the uniform field $\mathbf{B_0}$. The effect of these ideal gradient fields is to change only the magnitude

Fig. 5.1: Approximate gradient used for imaging in the experiments described in this thesis.

of the field (not its direction), and therefore the precession frequency, in a linear way along one direction.

Unfortunately, the fields described by Eqs. 5.1–5.3 cannot be realized. A magnetic field **B** in free space must have the properties

$$\nabla \cdot \mathbf{B} = 0, \tag{5.4}$$

$$\nabla \times \mathbf{B} = 0. \tag{5.5}$$

It is easy to see that the ideal imaging gradients of Eqs. 5.1–5.3 satisfy one or the other of these criteria but do not satisfy both. In order to meet the requirements on the divergence and curl, any spatially-varying magnetic field in free space must have concomitant components in at least two directions, and of comparable magnitudes. For example, common gradient fields used in MRI are those generated by coils of the Golay and Maxwell geometries (see Ch. 4):

$$\mathbf{B_{Golay}}(x, z) = G(z\hat{\mathbf{x}} + x\hat{\mathbf{z}}), \tag{5.6}$$

$$\mathbf{B_{Golay}}(y, z) = G(z\hat{\mathbf{y}} + y\hat{\mathbf{z}}), \tag{5.7}$$

$$\mathbf{B_{Maxwell}}(x, y, z) = \frac{G}{2}(-x\hat{\mathbf{x}} - y\hat{\mathbf{y}} + 2z\hat{\mathbf{z}}). \tag{5.8}$$

Other geometries are possible as well. For example, in the experiments described in Ch. 6, we use a gradient coil which creates a field (Fig. 5.1) approximately given by

$$\mathbf{B_G}(y, z) = G(-y\hat{\mathbf{y}} + z\hat{\mathbf{z}}). \tag{5.9}$$

The components of field along x and y are known as "concomitant gradients." These concomitant gradients alter the field magnitude and direction from the ideal case.

Fig. 5.2: Curves of constant frequency in the y-z plane from different gradient coils in the absence of a uniform static field. (A) Ideal gradient coil (no concomitant gradients). (B) Coil which provides field shown in Fig. 5.1. (C) Maxwell coil.

5.2 Effect of concomitant gradients on the imaging field

5.2.1 Concomitant gradients in the laboratory frame

In the absence of concomitant gradients, the application of a gradient field results in planes of constant field, and therefore constant precession frequency. For example, with an ideal gradient along z the precession frequency is

$$\omega(x, y, z) = \gamma \left(B_0 + Gz\right);$$

(5.10)

because the frequency does not vary with x or y, each plane of constant z is also a plane of constant precession frequency. This is shown in Fig. 5.2A. When concomitant gradients are considered, the planes of constant frequency bend into closed surfaces; depending on the gradient, these surfaces can be either cylinders or ellipsoids [7]. To demonstrate this, we will consider several specific examples.

We consider first the case of a gradient along the z-direction, applied by the coil which produces the field in Eq. 5.9. When added to the uniform $\mathbf{B_0}$ field, the total field becomes

$$\mathbf{B}(y, z) = -Gy\hat{\mathbf{y}} + (B_0 + Gz)\,\hat{\mathbf{z}},$$

(5.11)

where G is the gradient. The precession frequency is then given by

$$\omega(x, y, z) = \gamma \left[(Gy)^2 + (B_0 + Gz)^2\right]^{1/2}$$
$$= \gamma G \left[y^2 + \left(\frac{B_0}{G} + z\right)^2\right]^{1/2}.$$

(5.12)

From this equation, it is easy to see that the surfaces of constant frequency are no longer given by parallel planes, but rather by concentric cylinders which are centered on the z-axis at the point $(y, z) = (0, z_c)$, where z_c is given by

$$z_c = -\frac{B_0}{G}. \tag{5.13}$$

For two-dimensional imaging in the y-z plane, these cylinders appear as circles of constant precession frequency (Fig. 5.2B).

The situation is slightly different when using a gradient field with components in all three directions, such as that of the Maxwell coil (Eq. 5.8). The sum of the Maxwell coil field and the uniform field is

$$\mathbf{B}(x, y, z) = -\frac{Gx}{2}\hat{\mathbf{x}} - \frac{Gy}{2}\hat{\mathbf{y}} + (B_0 + Gz)\,\hat{\mathbf{z}}, \tag{5.14}$$

where G is the gradient. The precession frequencies are then

$$\begin{aligned}
\omega(x, y, z) &= \gamma \left[\left(\frac{Gx}{2}\right)^2 + \left(\frac{Gy}{2}\right)^2 + (B_0 + Gz)^2 \right]^{1/2} \\
&= \gamma G \left[\left(\frac{x}{2}\right)^2 + \left(\frac{y}{2}\right)^2 + \left(\frac{B_0}{G} + z\right)^2 \right]^{1/2}.
\end{aligned} \tag{5.15}$$

The surfaces of constant frequency are ellipsoids (in this particular case, prolate spheroids) which are centered at $(x, y, z) = (0, 0, z_c)$ where

$$z_c = -\frac{B_0}{G}. \tag{5.16}$$

For two-dimensional imaging in the y-z plane, these cylinders appear as ellipses of constant precession frequency (Fig. 5.2C).

We can define a convenient parameter [7]

$$\varepsilon = \frac{GL}{B_0}, \tag{5.17}$$

where L is the width of the FOV. This parameter is the ratio of the change in field magnitude due to gradients over the FOV to the uniform field magnitude. This can also be written as, in comparison with Eq. 5.13 (for example),

$$\varepsilon = \left|\frac{L}{z_c}\right|, \tag{5.18}$$

which compares the width of the FOV to the position of the center of the curves of constant frequency. In the case that $B_0 \gg GL$ (so that $\varepsilon \ll 1$), shown in Fig. 5.3A, we see that the concomitant gradients have very little effect on the lines of constant frequency within the FOV; because the circles of constant frequency are centered far outside the FOV, the circles of constant frequency within the FOV appear as straight, parallel lines (as they would be in the absence of concomitant gradients).

Fig. 5.3: Curves of constant frequency for increasing ε values. Dotted lines represent locations of planes of constant frequency in the absence of concomitant gradients. (A) $\varepsilon \ll 1$. This situation approximates an ideal gradient. (B) $\varepsilon = 0.4$. In this moderate concomitant gradient regime, bending of the lines of constant frequency is readily apparent. (C) $\varepsilon = 4$. Center of the circles of constant frequency is located within the FOV.

It is clear that as ε increases, however, the curvature of the constant frequency lines gets stronger because the center of curvature gets closer to the FOV (see Fig. 5.3B). At the edges of the FOV, the deviation of the curves from straight lines is given approximately by [7]

$$(\delta z)_{max} \approx \frac{\varepsilon L}{8} \tag{5.19}$$

for the field given by Eq. 5.9. As ε grows larger, eventually z_c will be smaller than $L/2$. In this case, there will be closed circles of constant frequency within the FOV (Fig. 5.3C).

5.2.2 Concomitant gradients in the rotating frame

In the laboratory frame, the concomitant gradients are static components in a direction perpendicular to $\mathbf{B_0}$; in the rotating frame picture, this translates into a component of field in the x'-y' plane which rotates at a frequency $-\omega_0$. This adds to the effective field along z', so that the total effective field in the rotating frame makes some angle with the z'-axis (see Fig. 5.4). The magnetization precesses about this effective field; in general the motion is complicated, because the effective field is not static in the frame.

Fortunately, a simplification can be made when the uniform field magnitude is much larger than the maximum gradient field magnitude. In this case, the magnitude of the effective field will be quite small. Therefore, in the rotating frame the transverse component of the effective field will precess at frequency $-\omega_0$, while the sample's magnetization will precess much more slowly because $\omega_0 \gg \gamma GL$. In this case, from the point of view of the slowly-precessing magnetization, the transverse component of the effective field will simply average to zero, and the concomitant gradients will have very little effect. However, if $\omega_0 \sim \gamma GL$ the rotation of the transverse field

Fig. 5.4: Effective field in rotating frame. In the presence of concomitant gradients, the effective field has a transverse component in the x'-y' plane. The transverse component of field rotates at frequency ω_0.

component will occur on a similar time scale as the precession of the magnetization; in this case, the transverse field component will no longer average away.

The same effect can be seen quantum-mechanically. For example, we will consider the use of a field of the form in Eq. 5.9 on top of the uniform static field. Including concomitant gradients, the total field in the laboratory frame is

$$\mathbf{B_{tot}}(y, z) = -Gy\hat{\mathbf{y}} + (B_0 + Gz)\,\hat{\mathbf{z}}. \tag{5.20}$$

The lab-frame Hamiltonian is therefore

$$H_{lab} = -\gamma\left[-I_y Gy + I_z\left(B_0 + Gz\right)\right]. \tag{5.21}$$

We apply the rotating frame transformation from Sec. 4.5 to find for the rotating frame Hamiltonian

$$H_{rot} = \gamma Gy\left(I_y \cos\omega_0 t + I_x \sin\omega_0 t\right) - \gamma GzI_z. \tag{5.22}$$

This Hamiltonian describes two precessions of the sample magnetization: precession about z' with frequency γGz, and precession at frequency γGy about a vector which itself rotates at frequency ω_0 in the x'-y' plane. As in the classical case, if ω_0 is much higher than γGL, the transverse field component will simply average away on the time scale of the magnetization's precession.

5.3 Effect of concomitant gradients on images

The concomitant gradients cause the magnitude and direction of the magnetic field to differ from the ideal case. If the differences are large, the result will be distortions in the image.

The ε parameter introduced in Sec. 5.2.1 can be used as a convenient measurement of the degree of distortion due to the concomitant gradients. In this section we will give a brief overview of these distortions for moderate and severe concomitant gradients; more detail can be found in [8, 35].

For small to moderate concomitant gradients, such that $\varepsilon \sim 0.1$ to 1, the lines of constant frequency in the FOV have a noticeable curve, which will be replicated in the frequency encoding direction in the image. This is visible in Fig. 5.5B and C.

In the phase-encoding direction for moderate concomitant gradients, there will also be curving due to the curvature of the lines of constant frequency. In addition to the variations in frequency, the fact that the field direction varies due to concomitant gradients introduces additional phase errors. The additional phase errors lead to blurring of the image, as can be seen in Fig. 5.5C. For modest concomitant gradients (approximately $\varepsilon < 0.5$), these artifacts can be corrected in post-processing [8].

When the concomitant gradients are large enough that the lines of constant frequency form closed curves inside the FOV, the distortions in both the frequency and phase encoding directions are severe. An example for $\varepsilon > 6$ is shown in Fig. 5.6. Distortions of this magnitude cannot be corrected in post-processing. However, it is possible to acquire undistorted images even in the limit of $\varepsilon \to \infty$ (that is, $B_0 \to 0$). In the next section we will discuss how this is achieved.

5.4 Elimination of concomitant gradient artifacts by averaging: the ZFMRI sequence

In the experiments reported in this thesis, we acquired images in severe concomitant gradients ($\varepsilon > 1$) using a pulse sequence which was designed to average out the concomitant gradients by means of a train of π pulses, with gradient field pulses applied between the π pulses [9, 10]. The sequence is based on the fact that, for small angles, the precession of a spin about an arbitrary $\mathbf{B}$ field can be represented by the sum of the precessions about each component of $\mathbf{B}$. The magnetization components which have evolved in the concomitant fields can be reversed to leave an effective unidirectional gradient. This is an example of an average Hamiltonian [36]. Because this sequence is used for imaging in the limit of zero static, uniform magnetic field, we have come to refer to it as the "zero-field MRI," or ZFMRI, sequence. We will consider here two-dimensional imaging in the y-z plane.

5.4.1 Classical picture of the ZFMRI sequence

In this analysis, we will assume a gradient field like that given in Eq. 5.9 and Fig. 5.1, as this gradient field approximates the one we used in our experiments. We further assume that all fields are switched nonadiabatically. Figure 5.7A shows the basic ZFMRI pulse sequence, while Fig. 5.7B and C show schematically the classical evolution of a proton spin at a point (y', z') subjected to this sequence.

The proton spin is first polarized along the x-axis by a large prepolarizing field B_p which is turned off at time $t = 0$. This corresponds with point a in Fig. 5.7. The gradient field is turned on for a time τ. During this time the spin precesses about the gradient field $\mathbf{B_G}(y, z)$. The time τ is chosen such that $\tau \ll 1/\gamma GL$, so that the angle of precession will be small and we can treat the precession as the sum of separate precessions about the y- and z-components of the field. The

Fig. 5.5: Simulated images with concomitant gradient distortions. Frequency encoding is along y; phase encoding is along z. (A) No concomitant gradient distortions. (B) $\varepsilon_{freq} = 0.16$, $\varepsilon_{phase} = 0.11$. (C) $\varepsilon_{freq} = 0.16$, $\varepsilon_{phase} = 0.76$. (Reprinted with permission from [8].)

Fig. 5.6: Example of severe concomitant gradient distortions. (A) High-field image showing geometry of phantom. (B) Projection-reconstruction image with $\varepsilon > 6$.
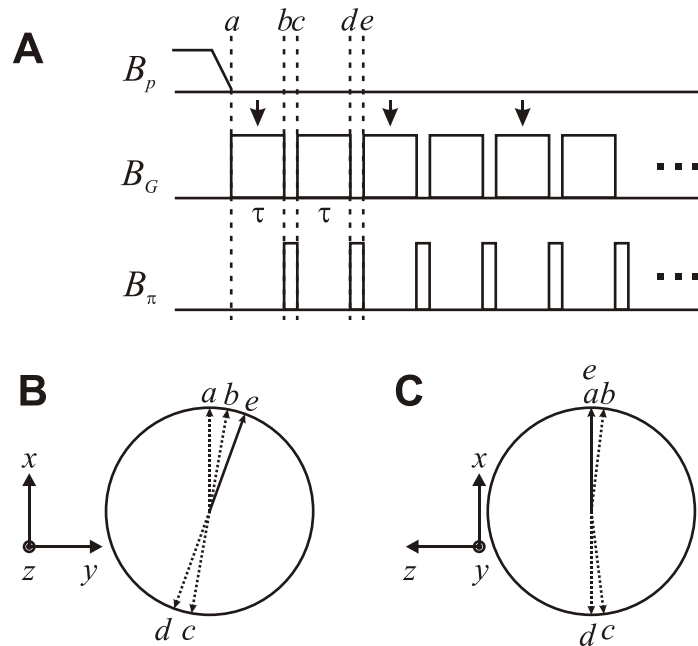
Fig. 5.7: Protocol for MRI in zero static field. (A) Basic ZFMRI pulse sequence vs. time. Arrows indicate stroboscopic data acquisition. (B) Evolution of a spin vector around the z-component of field. (C) Evolution of a spin vector around the y-component of field (the concomitant gradient component).

precession about the z-component of $\mathbf{B_G}$ is shown in Fig. 5.7B while the precession about the y-component is shown in Fig. 5.7C.

After the gradient pulse, we are at point b in Fig. 5.7. At this time a π pulse is applied along the z-axis; this is a dc pulse of uniform field, with the amplitude B_π and duration t_π chosen to satisfy $\gamma B_\pi t_\pi = \pi$. That is, the effect of the π pulse is to cause the spin to precess 180° around the z-axis. Ideally the π pulses are short, such that $t_\pi \ll \tau$. The end of the π pulse brings us to point c in Fig. 5.7. A second gradient pulse, identical to the first, is applied (ending at point d), followed by another π pulse (ending at point e). From Fig. 5.7B and C, we can see that if data are acquired (stroboscopically) after every second π pulse, we will observe a net precession around the z-component of $\mathbf{B_G}$, but no net precession about the y-component. Therefore the two π pulses average out the concomitant gradients and leave an effectively unidirectional gradient field

$$\mathbf{B_{G,eff}}(y, z) = Gz\hat{\mathbf{z}}. \tag{5.23}$$

It is convenient to describe a "pulse unit" consisting of two gradient pulses and two π pulses. Clearly the addition of more pulse units increases the angle of precession about $\mathbf{B_{G,eff}}$. In this way we traverse k-space in steps of $\Delta k = 2\gamma G\tau$.

The data can be acquired immediately after the last π pulse in the pulse unit, or it can be acquired after a small delay. It is advantageous to wait a time of $\tau/2$ before acquisition, for reasons that we will explain at the end of the next section.

5.4.2 Quantum-mechanical picture of the ZFMRI pulse sequence

If we assume that the π pulses are short enough that they can be treated as instantaneous transformations, then the time evolution operator after n pulse units is [9, 10]

$$U(2n\tau) = [U(2\tau)]^n, \tag{5.24}$$

where

$$U(2\tau) = e^{-iH\tau} e^{-i\pi I_z} e^{-iH\tau} e^{-i\pi I_z}. \tag{5.25}$$

The Hamiltonian H is given by

$$H = \gamma \mathbf{I} \cdot \mathbf{B} = \gamma \left(I_y B_y + I_z B_z \right) \tag{5.26}$$

in our situation. Inserting Eq. 5.26 into Eq. 5.25 and rearranging yields

$$\begin{aligned}
U(2\tau) &= \exp\left[-i\gamma\tau \left(I_y B_y + I_z B_z\right)\right] \exp\left[-i\gamma\tau \left(-I_y B_y + I_z B_z\right)\right] \\
&= \exp\left[-i2\gamma\tau I_z B_z + O\left(\tau^2\right)\right] \\
&= \exp\left[-i2\gamma\tau I_z Gz + O\left(\tau^2\right)\right],
\end{aligned} \tag{5.27}$$

where $O(\tau^2)$ indicates terms of second order and higher in τ. These higher-order terms are due to the fact that the operators in the two exponentials in Eq. 5.27 do not commute; the higher-order terms can be calculated by means of the Baker-Campbell-Hausdorff formula or the Magnus expansion [37]. So we see here that, at least to first order, the π pulses serve to average out the field components perpendicular to it. (The reason for the delay of $\tau/2$ before signal acquisition,

mentioned at the end of Sec. 5.4.1, is that it cancels out some of the higher-order components of the $O(\tau^2)$ term in the exponential [9, 10, 37].)

In the above analysis, we assumed that the π pulses were along the z-axis. It should be noted, however, that the π pulses could theoretically be applied along an arbitrary direction in the imaging plane. Changing the direction of the π pulse changes the direction of the effective gradient in a predictable way (which is different for each form of gradient field). This feature could be used to perform projection reconstruction, by changing the direction of the π pulse in order to acquire projections along different angles. We did not use this technique in this thesis, because we did not have the ability to apply π pulses along an arbitrary direction (see Sec. 6.1.1), and therefore we will not work this out in detail here. A more detailed analysis, for imaging in the x-z plane using a Golay coil and with π pulses along an arbitrary direction x', is given in [9, 10].

5.5   Discussion of the ZFMRI pulse sequence

There are several pertinent facts about the ZFMRI sequence which we would like to emphasize.

In a conventional sequence, the largest gradient which can be successfully used for imaging depends on the magnitude of the $\mathbf{B_0}$ field. In contrast, the maximum gradient which can be successfully used in the ZFMRI sequence depends on the hardware of the system. Larger gradients require shorter pulses to keep the precession during the gradient pulses in the small-angle regime. The duration of the gradient pulses affects the duration and amplitude of the π pulses (because one desires $t_\pi \ll \tau$). Also, since signal acquisition is performed during the gradient pulses, the detector dead time becomes an issue: if the dead time is longer than τ, stroboscopic acquisition cannot be performed after each pulse unit. In this case, a phase-encoding sequence can be used, at the expense of a longer acquisition time. See Sec. 6.1.3 for an example of such as sequence, and see Sec. 7.1 for a discussion of the imaging time when using such a sequence.

Though it was designed for imaging without a uniform static field, the ZFMRI sequence can be adapted to the presence of an ambient uniform static field. The presence of a uniform field imposes the same requirements on the hardware that are imposed by the use of strong gradients. However, a potential benefit of an ambient uniform field is that it could be used to augment the π pulses, which may reduce the requirements on the π pulse hardware. See Sec. 7.2 for a more specific discussion of the use of the ZFMRI sequence in the presence of an ambient field.

Chapter 6

The ZFMRI Experiment and Results

In this chapter we will present an experimental implementation of the ZFMRI imaging sequence. In particular, we will discuss its use in projection reconstruction imaging, and modifications which had to be made to the basic sequence to adapt it to our apparatus. We will describe our apparatus, including the difficulties we encountered when part of the original system suddenly and catastrophically ruptured. Finally we will show images which were encoded in the limit of zero static field. In this chapter, we will only be concerned with two-dimensional imaging in the y-z plane.

6.1 Pulse sequence details

6.1.1 Projection reconstruction by phantom rotation

The effective gradient created by the ZFMRI sequence is linear and unidirectional, with a direction determined by that of the π pulse; such a gradient lends itself well to projection reconstruction. Ideally one would rotate the direction of the π pulse in the y-z plane to acquire the necessary projections [9]. However, we found in our apparatus that π pulses along y could not be properly tuned due to spurious signals induced by the pulse; we were therefore limited to pulsing along only the z-direction. We acquired different projections by rotating the phantom about its axis. This approach to projection reconstruction is not ideal, as rotating the phantom may cause the center of the sample to move relative to the center of the gradient; with a small number of projections, it is often possible (and not too time-consuming) to correct the resultant frequency shifts by eye.

6.1.2 Tuning the π pulses

One troublesome detail of our system was that our pulsing magnetic fields induced currents or relaxations in other parts of our apparatus, due to imperfect shielding of the fields [27, 30]. In addition to being an issue with signal acquisition (see Secs. 6.1.3 and 6.2.5), the fields from these undesired signals may influence the spins in our sample. These fields had to be taken into account when setting up the π pulse.
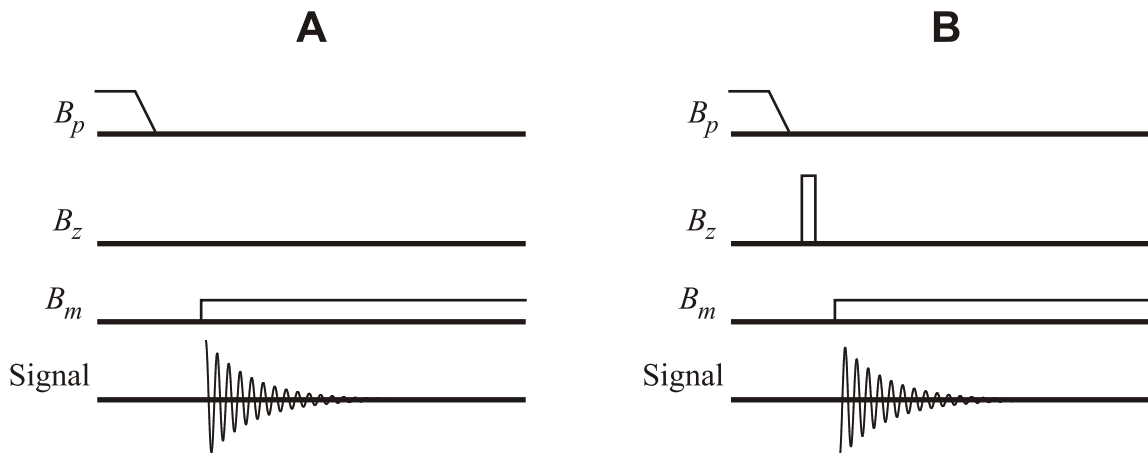
Fig. 6.1: Pulse sequences for tuning the π pulse. The phase of precession from sequence A (without a π pulse) is compared to that of sequence B (in which a π pulse is applied). The pulse is retuned if the phase difference is not sufficiently close to $180°$.

In order to tune the π pulse, we applied a polarization pulse along the x-direction, which we shut off nonadiabatically. After a small delay time (a few milliseconds), we applied a measurement field $B_m \approx 3.75 \ \mu$T ($f_m = \gamma B_m/2\pi \approx 160$ Hz) along the z-direction and acquired the FID. We then repeated the experiment, but this time we applied a π pulse in the middle of the delay between polarization and acquisition (see Fig. 6.1). In both experiments, the time between the shutoff of the polarizing field and the application of the measurement field was the same. The amplitudes and phases of the two FIDs were compared, and the amplitude of the π pulse was readjusted if the resultant phase difference was not within about $\pm2\%$ of π. We chose a limit of 2% because this accuracy was easy to achieve and resulted in suitable images.

6.1.3 Practical experimental pulse sequence

For our ZFMRI experiment, we made some modifications to the basic pulse sequence presented in Sec. 5.4.1. Those modifications are as follows.

In order to ensure that the $k = 0$ point was included in our projections, we created a gradient echo by applying negative gradient pulses in several pulse units before the first data point was acquired. This procedure guaranteed that the acquisition started in negative k-space. In our first images we used five negative pulse units; later this was reduced to one negative pulse unit so that we could reach higher points in k-space before T_2^* decay eliminated the visible signal.

The pulse sequence of Fig. 5.7 calls for stroboscopic acquisition in which one data point is acquired after every pulse unit after a delay of $\tau/2$. Using stroboscopic acquisition, all of the k-space points in the projection would be acquired after one polarizing pulse. Imaging in our apparatus, however, required an adaptation of the basic ZFMRI sequence. In our modification, we polarize the sample and apply the sequence until we have reached a desired point in k-space (after a desired number of pulse units); at this point we stop the ZFMRI encoding, switch on a uniform static measurement field $B_m \approx 3.75 \ \mu$T, and acquire the FID in the static field. The inverse Fourier

Fig. 6.2: Pulse sequence for ZFMRI with one-dimensional phase encoding and one negative pulse unit. The inverse Fourier transform of the received signal yields the real and imaginary parts of the ZFMRI signal at $k(9\tau/2) = (\tau/2)\gamma G_z$.

transform of this real-valued FID produces a complex-valued peak in frequency space, which yields the real and imaginary parts of the ZFMRI signal at the point in k-space. In short, we perform one-dimensional phase encoding with the ZFMRI sequence. An example of this sequence (including the negative gradient pulses discussed above) is illustrated in Fig. 6.2. Although our particular incarnation of the ZFMRI sequence does use a uniform static field for detection, we emphasize that this field is *not* applied during image encoding.

The phase encoding procedure described above has two advantages. First, it yields the real and imaginary part of the ZFMRI signal, so it enables quadrature detection with a single detector. Quadrature detection was necessary in this experiment because all of the spins did not precess in the same direction, in contrast to a conventional image. For a collection of precessing spins, a single detector can give information about the frequencies of precession, but not about the direction. In a conventional image in which there is a static field that is much larger than the field change due to gradients (when $\varepsilon \ll 1$), the direction of the magnetic field will be the same at all points, so all spins will precess in the same direction. The signal from a single detector can be split into two channels, one of which is appropriately phase shifted, in order to acquire the real and imaginary parts of the signal. When the field does not point in the same direction over the entire sample, however, some spins will precess clockwise while others precess anticlockwise. In this case, the signal from a single detector cannot simply be split and phase-shifted to find the real and imaginary parts of the signal.

The second advantage of the phase-encoding procedure is that it allows us to delay the acquisition of the signal, and also acquire the ZFMRI signal at an arbitrary frequency. These features of phase encoding were essential in our experiment because the magnetic field pulses generated relaxation signals in our apparatus which prevented locking of the FLL for approximately 10 ms after a field was turned on or off. Our gradient pulse duration was $\tau = 5$ ms, so stroboscopic

Fig. 6.3: Pulse sequence for a conventional gradient echo image modified for one-dimensional phase encoding. The inverse Fourier transform of the received signal yields the real and imaginary parts of the gradient echo signal at $k(9\tau/2) = (\tau/2)\gamma G_z$.

acquisition was not possible. Furthermore, the relaxation signals persisted for tens to hundreds of milliseconds and were of sufficient amplitude to saturate the acquisition system as described in Sec. 6.2.5. By increasing the frequency of the acquired signal, we were able to high-pass filter the signal to remove the relaxation and avoid saturating the electronics.

Unfortunately, the phase encoding technique has a very important disadvantage: it is slow. Using phase encoding, acquiring each point in k-space requires its own polarization, encoding, and detection steps. An image with 24 projections and 24 points per projection, for example, requires 576 repetitions of the polarization-encoding-detection cycle (versus only 24 repetitions of the cycle for stroboscopic acquisition). A more detailed examination of imaging time appears in Sec. 7.1.

6.1.4 Phase-encoded conventional gradient echo

After we established that we could acquire images using the ZFMRI sequence described in Sec. 6.1.3, we wanted to compare them to images taken with a conventional imaging sequence under similar conditions. We constructed a gradient echo sequence which used a very small B_0 and was designed for one-dimensional phase encoding.

An example of this sequence is shown in Fig. 6.3. After a polarizing pulse, a negative gradient field and a small uniform $\mathbf{B_0}$ field along the z-axis are applied. After a time 2τ (where τ has the same value as in the ZFMRI sequence), the gradient is reversed. The $\mathbf{B_0}$ and gradient fields are maintained until the desired point in k-space is reached, at which time a measurement field $\mathbf{B_m}$ is applied along z. The time-domain data are acquired and processed as described in the Sec. 6.1.3 for the zero-field experiment.

The same polarizing pulse, gradients, and encoding times are used in this sequence as in

the ZFMRI sequence. Because of this, it was easy to acquire the conventional image concurrently with the ZFMRI image by simply running both sequences for each projection.

Although both sequences use the same field gradients, the value of ε is somewhat lower in the conventional sequence because the conventional sequence requires an applied $\mathbf{B_0}$ field during encoding to establish a "preferred" gradient direction; components of the gradient perpendicular to $\mathbf{B_0}$ are the unwanted concomitant terms. (Recall that in the zero-field sequence, the preferred gradient direction is determined by that of the π pulses.)

6.2 Apparatus

6.2.1 SQUID package and gradiometer

The SQUID used for all of the results in this thesis was made by Michael Mück of the University of Gießen, Germany. It is a Nb-AlOx-Nb device with an integrated 20-turn Nb input coil. The maximum critical current of the SQUID was $I_c = 8$ μA with a shunt resistance $R_n = 7$ Ω. The flux noise of this SQUID was approximately 19 $\mu\Phi_0/\mathrm{Hz}^{1/2}$ over the frequency range of 10 Hz to at least 100 kHz. The SQUID was mounted on a circuit board as shown in Fig. 6.4. The circuit board was inserted into a Nb tube (45 mm long, 13 mm in diameter) to shield it from external magnetic fields.

We used a first-derivative axial gradiometer wound of 75-μm diameter Formvar-insulated Nb wire on a G-10 fiberglass tube. The loops had a diameter of 38 mm and were separated by a distance of 43 mm. The gradiometer leads were connected to the SQUID package by means of screw terminals on Nb foil pads which had been glued to the circuit board. The Nb pads were connected by a series RC ($R = 13$ Ω, $C = 560$ pF) shunt to prevent coupling of radiofrequency fields to the SQUID. A series array of Josephson junctions was inserted between one Nb pad and one input coil pad on the SQUID. This array consisted of 20 hysteretic junctions; the critical current of one junction was 30 μA and the normal state resistance per junction was about 50 Ω (for a total resistance of 1 kΩ). Superconducting connections between the Nb pads, junction array, and SQUID input coil were made by wire-bonding annealed Nb wire.

The SQUID package and gradiometer were the same as those used in [27].

6.2.2 Liquid helium dewar and magnetic shielding

The ZFMRI experiments were conducted in the liquid helium (LHe) dewar with dimensions shown in Fig. 6.5. The outer vessel of the dewar was aluminum; the inner vessel consisted of a G-10 fiberglass neck with a steel belly. To screen external magnetic fields, the dewar was suspended in a single-layer, open-topped mu-metal can of approximate diameter 400 mm. (This dewar was also used for measurements in [27, 30].)

The belly of the dewar contained a cylindrical Pb shield to stabilize the residual magnetic field. The original shield [27] was nearly the full diameter of the belly and was made from three sheets of Pb soldered into a cylinder, with phenolic slats as supports. This cylinder rested on a Pb plate which covered the bottom of the dewar. To be inserted into the dewar, this shield had to be collapsed to a diameter of about 150 mm to fit down the dewar neck, then re-expanded once in the dewar belly. Eventually, this shield was replaced by one with a diameter of approximately

Fig. 6.4: Photograph of SQUID package. From right: The gradiometer leads are affixed to the niobium foil pads by means of screw terminals. A series RC shunt is soldered across the pads to divert undesired radiofrequency signals. The niobium pads are connected to the input coil of the SQUID via superconducting niobium wire bonds; a Josephson junction array (center) is placed in series with the input coil. The SQUID washer is connected to copper pads with aluminum wire bonds. Copper leads soldered to the pad connect the SQUID to the room-temperature electronics.

Fig. 6.5: Liquid helium dewar used for ZFMRI experiments. Gray color represents vacuum and insulation space between inner and outer vessels. The dewar sits in a mu-metal can which is 1.22 m (48") tall. Lead shield inside belly of dewar is not shown.

150 mm, which did not need to be collapsed to insert into the dewar. The new shield was made from a single sheet of Pb (610 mm × 610 mm), which was formed into a cylinder and soldered. A tail cap was also fabricated from Pb sheet and soldered into one end. Small holes were drilled in the tail cap and wall of the cylinder to ensure that the LHe would fill the entire dewar belly.

6.2.3 Pyrex cryogenic insert

This project utilized a double-walled Pyrex glass vessel which was inserted into the dewar; the insert had a bore with room temperature access for the phantom and also served to support the magnetic field coils and SQUID package in the LHe bath around the sample. The outer surface of the insert had eight glass hooks to aid in mounting the coils.

We began these experiments with the insert described in [27,30], which is shown schematically in Fig. 6.6A. A liquid nitrogen (LN_2) space (with room temperature access) surrounded the top 0.8 m of the bore; the coils were placed around the remaining 80 mm "tail." The phantom rested at the bottom of this tail. To allow for rotating the phantom during imaging, it was taped to the end of a long G-10 rod which protruded from the top of the bore. A single continuous vacuum jacket served to insulate the bore, LN_2, and LHe spaces from one another; the glass walls inside the vacuum space were strip-silvered to reflect radiation.

Despite the insulation against conductive and radiative heat exchange, the equilibrium temperature at the bottom of the bore was approximately −120°C (about 150 K). We used methanol (melting point = −97.53°C) or ethanol (melting point = −114.14°C) [38] for our samples rather than water, because the alcohols required less heating to remain in the liquid state. The samples were kept liquid with a heater made of twisted-pair manganin wire and coated in Stycast 2850FT epoxy. This heater was not built into the glass vessel; rather, the heater was placed around the phantom before lowering it down the bore. Because of the low melting points of the alcohols, only occasional heating was required to maintain the liquid state. Therefore during imaging the heater was turned of so as to not inject noise into the measurements. We found that running the heater with 100 mA of current for 5–10 minutes or so every 30–45 minutes was usually sufficient.

Most unfortunately, the original Pyrex insert was destroyed not long into the ZFMRI project. While we did not definitively identify the cause, we speculate that fatigue at the glass-to-metal seal caused the glass to crack, which led to the implosion of the insert. In order to expedite the fabrication of a new insert, we revised the design to omit the LN_2 jacket. (We also performed the silvering of the vacuum space in our lab, following the techniques found in [39–41].) The second-generation insert is shown schematically in Fig. 6.6B. We found that, due to the design changes, the bore temperature was significantly lowered to around −220°C (approximately 50 K). This posed several challenges: we had to keep air out of the bore in order to prevent the formation of ice plugs, and our alcohol samples required nearly constant heating to keep from freezing.

We isolated the bore from the atmosphere through the use of an "air lock" fashioned from a ball valve topped by a quick-disconnect vacuum fitting (Fig. 6.7). Rather than taping the phantom to the end of a G-10 rod, we sealed one end of a phenolic tube and inserted the phantom in the other end. The tube would be placed in the vacuum fitting, then evacuated and flushed with helium gas by means of a drain port on the ball valve. A boss on the inside surface of the tube prevented the phantom from being pushed up into the evacuated tube when it was flushed with helium. Once the air had been removed from the tube, we opened the ball valve and lowered

A **B**

Brass top plate

Vacuum space

Mu-metal

Pb shield

Sample

LHe

Fig. 6.6: Schematics of Pyrex glass inserts (with sample) and superconducting Pb shields in LHe dewar of Fig. 6.5. Dark gray represents vacuum space in the dewar and inserts. Light gray represents LHe. (A) Original insert (with LN_2 jacket) and original Pb shield. (B) Second-generation insert (without LN_2 jacket) and new Pb shield.

Fig. 6.7: Photograph of dewar top plate showing ball valve with drain port, used to evacuate the upper chamber to prevent air from entering the bore of the insert.

the phantom to the bottom of the insert. We also filled the bore with a small quantity of helium gas to reduce the pressure differential between the environment and the bore (to help reduce the possibility of air leaks into the bore). In later experiments we used neon gas in the bore, because neon diffuses through Pyrex at a much slower rate than helium [42].

We modified our pulse sequence and electronics in order to run the heater during the polarizing pulses, when magnetic fields from the heater would not interfere with encoding or detection. This gave us a duty cycle of approximately 60% during imaging, which was suitable to keep the sample liquid with heater currents of around 100 mA. The heater was applied continuously between projections or during other "down time" in the experiments. Late in the experiment we added a thermocouple (type T, consisting of copper and constantan wires) to monitor the temperature of the sample; we found that ethanol samples worked best for NMR when they were at $-50°$C or higher, due to the temperature dependence of T_2.

An additional challenge was that, due to the low pressure environment, liquid alcohol samples tended to evaporate quickly unless the phantom was well-sealed. Significant quantities of the sample could be lost during imaging, or even during the evacuation of the sample tube. To combat this, we designed and fabricated a sealed phantom, which will be discussed in further detail in Sec. 6.3.3.

6.2.4   Coils

All of the field-generating coils used in our ZFMRI experiment were wound of 0.2-mm diameter Cu-clad NbTi superconducting wire with Formvar insulation (note that the heater, which was wound of manganin wire, is not considered a field-generating coil). A schematic diagram of the coils and their relation to the Pyrex insert and the sample are shown in Fig. 6.8A. The coils were bolted together into a single assembly and affixed to the Pyrex insert by tying loops of dental floss between the coil assembly and glass hooks on the outer walls of the insert (Fig. 6.8B).

Current through each of the coils was controlled by relays activated by digital logic (TTL) pulses from an Interface Technology RS-660 timing simulator. All of the coils besides the polarizing coil were powered by batteries; current amplitude adjustments for the battery-powered coils were performed with potentiometers in series with the batteries.

6.2.4.1   Polarizing coil

The polarizing coil was made of approximately 200 turns wound as a single-layer solenoid directly on the tail of the Pyrex insert. The coil diameter was 35 mm and the inductance was about 0.4 mH. The polarizing coil was the only one not powered by batteries during imaging; power for this coil was provided by a HP 6434B DC power supply. For all of the images shown in this thesis, we used a polarization current of approximately 3 A to achieve a polarization field around 10 mT. This field was switched off nonadiabatically with a L/R time of about 80 μs. A solid-state relay was used to turn off the current, and then a mechanical relay ensured that the coil was open during data acquisition to reduce noise.

Fig. 6.8: (A) Schematic diagram of insert tail, showing double-walled glass vessel, sample, gradiometer, and field coils. Colors are used to visually separate the components. (B) Photograph of coil assembly affixed to the tail of the second-generation insert. Note the glass hooks immediately above top green coil. Another row of glass hooks is located higher, above the top of the photo. The vertical brown boom on the left holds the SQUID package below the bottom of the insert to maximize the time it is in LHe.

Fig. 6.9: Calculated magnetic field in the imaging region in the y-z plane due to our planar gradient coils.

6.2.4.2 Uniform field coils (for π, $\mathbf{B_0}$, and $\mathbf{B_m}$ pulses)

Our ZFMRI technique utilized two different uniform fields: $\mathbf{B_z}$ for the π pulses and $\mathbf{B_m}$ for the measurement field pulses. The phase-encoded conventional gradient echo sequence described in Sec. 6.1.4 required a $\mathbf{B_0}$ field during encoding, in addition to the $\mathbf{B_m}$ field for acquisition. The coils to generate these fields were described in [27]. The $\mathbf{B_z}$ coils comprised a Helmholtz pair wound on G-10 frames. The mean diameter of the coils was approximately 90 mm, and each coil contained 70 turns of wire. The inductance of this Helmholtz pair was measured to be 2.1 mH, and the coil constant was 1.24 μT/mA. The $\mathbf{B_m}$ coils consisted of 10 turns wound on top of each of the $\mathbf{B_z}$ coils. The inductance of the $\mathbf{B_m}$ pair was not measured. The coil constant of the $\mathbf{B_m}$ pair was 0.28 μT/mA.

The coil assembly contained an additional Helmholtz pair $\mathbf{B_y}$ perpendicular to $\mathbf{B_z}$ and $\mathbf{B_m}$; however, it was not used in these experiments (see Sec. 6.1.1).

6.2.4.3 Gradient coils

We applied the magnetic field gradients using a pair of rectangular, counterwound coils which were designed to fit into the existing Helmholtz coil assembly (one of the gradient coils is visible on the left-hand side of the coil assembly in Fig. 6.8B). The gradient coils were machined from 5-mm thick sheets of grade LE phenolic. They were 100 mm tall and 43.5 mm wide, allowing them to just fit between the $\mathbf{B_y}$ coils mentioned in the previous section. Separation between the coils was 53 mm. Each coil contained 25 turns of wire. The inductance of the pair was measured to be approximately 0.25 mH.

The field profile of this gradient pair in the y-z plane was approximated in MATLAB by calculating the field from four infinite vertical wires. The calculated field is shown in Fig. 6.9. Because the width of the coils was not equal to the separation, there is a small anisotropy in the field

gradients. The calculated gradient versus current along the y-direction is $G_y \approx -15.7 \ \mu\mathrm{T/m/mA}$, whereas along the z-direction it is $G_z \approx 17.2 \ \mu\mathrm{T/m/mA}$. Subsequently, these estimates were confirmed experimentally (to within about 10%) by measuring the bandwidth of the signal from a vial of alcohol.

6.2.4.4 Cancellation coils

We added a pair of coils along the axis of the insert (the x-direction in our system) in an attempt to cancel the residual fields along this direction and reduce the field at the phantom as close to zero as possible. Our system already has uniform field coils along y and z, which could be used to cancel residual fields in those directions. The new cancellation pair was designed to fit on the existing coil assembly, as can be seen in Fig. 6.8B, which constrained the design somewhat. In particular, we could not use a Helmholtz pair because the separation of the cancellation pair was dictated by the sizes of the other coils while the diameter of the cancellation coils was limited by the neck of the dewar. The cancellation coils are 76 mm in diameter and separated by approximately 110 mm. Each coil has 100 turns and the coil constant is 0.35 to 0.4 $\mu\mathrm{T/mA}$.

Because these coils were coaxial with the gradiometer, we were concerned about injecting noise into the measurements; we attempted to switch the field off during acquisition but found that this resulted in a strong ringing signal picked up by the gradiometer. Fortunately, we subsequently learned that we could run the cancellation coils continuously from batteries during acquisition with little or no increase in system noise.

6.2.5 Control and acquisition electronics

The voltage output of the FLL passed through a PAR-113 preamplifier (gain = 100) and a Krohn-Hite 3320 filter configured for high-pass (cutoff frequency = 100 Hz), before being acquired and digitized using a National Instruments DAQ card in conjunction with LabVIEW. The filtering was required because magnetic field pulses in the dewar generated relaxation signals from the nearby metal in the apparatus due to imperfect shielding; the relaxation signals lasted tens of milliseconds or longer and were of sufficient amplitude to saturate the DAQ card. Putting the signal through a high-pass filter effectively removed the relaxation from the acquired signal, but also required that our NMR signal be above the filter's cutoff frequency. This is one of the reasons we used a one-dimensional phase encoding sequence in the ZFMRI experiment (as described in Sec. 6.1.3).

For the early images the DAQ card was the PCI-MIO-16XE-10; this card ceased to function during the time that the second-generation glass insert was being fabricated. We replaced the card with a new card, the PCI-6036E, which had nearly identical specifications, so that no additional changes to the acquisition system were necessary.

The triggering of magnetic field pulses, heater pulses, FLL integrator reset, and signal acquisition were controlled with TTL pulses generated by a RS-660 word generator/timing simulator. It was run in timing simulator mode, in which it featured eight TTL channels. Pulse sequences could be programmed either by hand using the RS-660 front panel, or via the General Purpose Instrument Bus (GPIB) interface. The LabVIEW acquisition software included a routine to add pulse units to the sequence via GPIB; without this ability, each new sequence would have to be programmed in by hand, making the ZFMRI experiment excessively slow.

Fig. 6.10: (A) Top view diagram of Teflon phantom used for our first ZFMRI image. Darker circles indicate holes which were filled with methanol. The phantom was 28 mm tall and the holes were 25 mm deep. (B) First ZFMRI image. Image consists of 12 projections with 20 points per projection; the projections were zero-filled to 128 points during processing.

6.3 Results

6.3.1 First image acquired in the zero-static-field limit

For our first image, we used a simple phantom which was made of four holes drilled in a 17-mm diameter Teflon cylinder. The diameter of the cylinder was chosen to fit an existing heater made of black Stycast-coated manganin twisted pair. A diagram of the phantom is shown in Fig. 6.10A. Each column was 25 mm tall and filled with approximately 0.1 mL of methanol; the open top of the phantom was sealed with masking tape and Play-Doh to hinder sample leakage and evaporation. We used methanol due to its low freezing point, as the bottom of the warm bore was significantly below 0°C. Despite the low freezing point of the sample, it still required occasional warming with the heater.

The G_z gradient for this image was approximately 40 μT/m and the gradient pulse duration was $\tau = 5$ ms. The π pulse duration was 1 ms, which corresponds to a field amplitude of approximately $B_\pi = 12$ μT and a duty cycle of 1/6. Five negative pulse units were applied to start the acquisition in negative k-space. We acquired data at 20 points in k-space per projection (with each point averaged 10 times to increase SNR), and 12 projections at 15° increments to cover the angles from 0° to 165°. Rotation of the sample was accomplished by taping it to the end of a long G-10 rod which protruded from the bore and could be grasped.

The first image taken with the ZFMRI sequence is shown in Fig. 6.10B. We estimated a residual field in the dewar of up to 1 μT, which yielded $\varepsilon = GL/B_0 \approx 2.3$ (although the residual field is not an applied $\mathbf{B_0}$ field, it is still a nonzero static field which may truncate the concomitant

Fig. 6.11: (A) New phantom, shown with a dime for scale. Phantom was filled with methanol for imaging; red liquid in photo is for visibility. (B) ZFMRI image with $G_z = 100$ μT/m. Image is 24 projections with 24 data points per projection, zero-filled to 128 points in processing.

gradients, so we treat it as $\mathbf{B_0}$ when we calculate ε). However, because the FOV was much larger than the phantom, our sample did not see that large of a field change. If we take ΔB to be the change in field across the sample, then we find $\Delta B/B_0 \approx 0.68$; clearly we needed to try again with higher gradients.

6.3.2   Improved phantom and increased gradient

The first image served as a proof-of-principle that the ZFMRI sequence could work to produce images. For subsequent experiments, we desired a phantom which would highlight distortions due to concomitant gradients in a conventionally-encoded image. To these ends, we designed a new phantom consisting of an acrylic cylinder with three parallel channels which were bisected with a perpendicular cross channel. This phantom was made by the Chemistry Machine Shop and is shown in Fig. 6.11A. We again used methanol for the sample and sealed the top of the phantom with masking tape and Play-Doh.

We increased the amplitude of the gradient pulse to 100 μT/m for the same 5-ms pulse duration. This resulted in a FOV of about 23 mm, which was much closer to the size of our phantom. In this case we took projections every 7.5° from 0° to 172.5°, for a total of 24 projections. We reduced to one the number of negative pulse units in the sequence so that we could get to higher k-space points before too much signal was lost to T_2^* decay. A total of 24 points were acquired for each projection, with each point averaged 10 times. The π pulse duration was 1 ms as in the previous image.

Over the course of this imaging experiment, we observed that the SNR dropped to approximately 1/3 of its starting value. We discovered that the loss of SNR was due to evaporation of about 2/3 of the initial sample (in later testing, we found that the methanol was able to "climb" the walls of the phantom and pool on the top surface). In order to process the data, each projection was normalized by dividing the data by the maximum amplitude in the projection; the reconstructed image is shown in Fig. 6.11B. Assuming a residual field of up to 1 μT, we estimated for this image

that $\varepsilon \approx 2.3$. Thanks to the ZFMRI sequence, this image is free of concomitant gradient distortions despite the large field change across the sample.

The loss of over half of the sample during several hours of imaging was obviously a serious problem; we took steps to remedy this by boring out the top 5 mm or so of the phantom while the Chemistry Machine Shop fabricated a tight-fitting Teflon cap to seal this hole. However, we did not get the chance to image this phantom again, as we explain below.

6.3.3 Imaging with a new phantom in a new insert

Our original insert shattered approximately two weeks after we acquired the image described in the preceding section. A description of the replacement insert and the issue of the very cold bore were presented in Sec. 6.2.3. Although we added helium gas to the bore, the low temperatures ensured that the bore was still under partial vacuum. Our methanol samples, heated to approximately -50°C, evaporated very quickly under these conditions despite the cap we had made for the phantom. We attempted to modify our acrylic phantom so that it would seal more tightly, but the modifications were ineffective; the phantom broke after increasingly aggressive methods to seal it against the vacuum in the bore.

We designed and fabricated a new phantom, a diagram and photo of which are shown in Fig. 6.12. The new phantom was made of nylon 6/6, which is chemically resistant to alcohols [43], and hermetically sealed via solvent welding. We avoided the use of adhesives out of concerns for differential thermal contraction between the nylon and the adhesive layer. We also switched the sample material from methanol to ethanol, which has a lower vapor pressure [38]. Details of the phantom construction are recounted below.

A nylon cylinder was mounted horizontally in a milling machine; material was cut away to leave a half-cylinder. A pattern was machined into the flat side of the half-cylinder. This process was repeated to make another half-cylinder. The flat sides of the half-cylinders were sanded until smooth and solvent-welded together with 96% formic acid to form a cylinder. A piece of thin nylon sheet was solvent-welded around the cylinder to strengthen the joints and add extra protection against leaks and ruptures. A small hole was drilled and tapped through the top of the phantom and into the sample space. Ethanol was injected into this hole using a syringe; a nylon screw was then solvent-welded into the hole to seal the phantom.

Although the sample evaporation was much reduced with the new phantom, it could not be completely eliminated. The ethanol had to be refilled after a few weeks due to evaporation through the nylon. Refilling was easily accomplished by drilling and tapping a new hole through the top of the phantom, injecting ethanol, and sealing the new hole with a solvent-welded nylon screw. Although the nylon is opaque, it is translucent enough to allow the sample level to be seen through the side if the phantom is held to the light; alternatively, the amount of ethanol could be monitored by weighing the phantom.

With the new phantom in hand, and other issues with the second-generation insert dealt with, we set out to resume imaging. Our imaging parameters were: gradient pulse duration of $\tau = 5$ ms, gradient amplitude of $G_z = 86\ \mu$T/m (corresponding to a current in the gradient coil of $I_G = 5$ mA), and π pulse duration $t_\pi = 1$ ms. Before we began this image, we put the heater control on a TTL line so that we could heat the sample during polarization pulses. We took 24 projections (7.5° to 180° in steps of 7.5°), with 24 k-space points per projection and 10 averages

A

B

Fig. 6.12: (A) Diagram of sealed nylon phantom. Dashed lines outlining the pattern indicate that the material is opaque. Shaded area in figure is the pattern machined into the two halves, which are joined along the dotted line. (B) Photograph of one such phantom. Sealing line is visible horizontally between the screws.

Fig. 6.13: (A) High-field image of the nylon phantom filled with water. This image was scaled and rotated to facilitate visual comparison with the ZFMRI image. (B) First ZFMRI image of nylon phantom. Image was 24 projections with 24 points per projection, and zero-filled to 128 points in processing. For this image, $\varepsilon \approx 7$.

per point.

Because the new nylon phantom was opaque, we could not see through it to determine how accurately it was represented in our images. In order to observe the actual geometry of the phantom, we drilled a hole in the phantom, drained the ethanol, refilled with water, and imaged the phantom in a high-field system. (The high-field imaging actually was performed several weeks after the ZFMRI image we report in this section.) To our mild chagrin, we found the machining work on this phantom was not as pretty as we had hoped; however, we did find that our images accurately reproduced the characteristics of the phantom. The high-field image is shown in Fig. 6.13A and the first ZFMRI image of this phantom is shown in Fig. 6.13B.

As part of the troubleshooting of the second-generation insert, we replaced the 300-mm diameter superconducting Pb shield in the dewar belly with a new shield with a diameter of 150 mm. In later measurements we found that the residual field was now less than 0.33 μT, which yields $\varepsilon \approx 7$ for this image.

6.3.4 Comparison to a conventionally-encoded image with $\varepsilon > 1$

By this point in the experiment, we had seen many examples of images encoded in the zero-static-field limit using the ZFMRI sequence, but we had never compared the results of ZFMRI to the results of a conventional sequence under similar conditions of gradient, field of view, and very low static fields. We designed the phase-encoded gradient echo sequence of Sec. 6.1.4 so that we could compare the two encoding methods. We furthermore decided that the comparison would be most valid if the conventional-sequence images were taken concurrently with the ZFMRI images; at each projection angle, we would take a conventional data set and a ZFMRI data set. In this way, we would be sure that any distortions due to errors in the projection angle would affect both images equally. However, this had the obvious disadvantage of making the imaging take about twice as long. Because of this, we often decided to break up an image to be done over two or more

Fig. 6.14: Comparison of ZFMRI sequence and a conventional sequence. (A) High-field image to show phantom geometry (this is the same image as in Fig. 6.13A, but rotated and rescaled to match the ZFMRI image). (B) Image taken with ZFMRI sequence, with $\varepsilon \approx 7$. (C) Image taken with phase-encoded gradient echo sequence (Sec. 6.1.4). In this case, $\varepsilon \approx 5$ due to the B_0 added to the estimate of the residual field.

days.

The conventionally-encoded image sequence used the same gradients as the ZFMRI image, as well as the same k-space intervals and the same $\mathbf{B_m}$ field. We applied a small $\mathbf{B_0}$ using the $\mathbf{B_z}$ coil, which was used for applying π pulses in the zero-field sequence. The only changes that had to be made to switch between conventional and ZFMRI encoding were to change the sequence loaded into the RS-660 timing simulator, and to add a resistor in series with the $\mathbf{B_z}$ coil in order to reduce the current for $\mathbf{B_0}$. We selected $B_0 \approx 0.12$ μT, which corresponds to a Larmor frequency of about 5 Hz.

The results of one such comparison are shown in Fig. 6.14. In Fig. 6.14A we show the high-field image to establish the geometry of the phantom. Figure 6.14B shows the outcome of the ZFMRI sequence, with parameters as follows: $G_z = 100$ μT/m; $\tau = 5$ ms; $t_\pi = 1$ ms; 24 projections with 24 points per projection and 10 averages per point. Figure 6.14C contains the phase-encoded conventional image with $G_z = 100$ μT/m, 10 ms of encoding time between k-space points, and $B_0 \approx 0.12$ μT. The conventional image had the same number of points, projections, and averages.

6.3.5 Cancellation of residual field along axis of insert

In an effort to reduce the residual magnetic field in the dewar as much as possible (and therefore to maximize our ε values), we constructed the cancellation coils discussed in Sec. 6.2.4.4. We needed only to construct coils along the axis of the insert, which was the x-direction, because we already had uniform field coils along the y- and z-directions.

In order to cancel the residual field in the dewar, we first had to measure it. The stray field along directions perpendicular to the gradiometer axis are rather easy to measure; for example, it can be done by simply acquiring an NMR spectrum, then switching the direction (but not the magnitude) of the current in the $\mathbf{B_0}$ coil and acquiring another spectrum. The difference in

Fig. 6.15: (A) High-field image to show phantom geometry, rotated and scaled to match ZFMRI image. (B) ZFMRI image with cancellation of stray fields ($\varepsilon > 10$). (C) Conventional sequence with cancellation of stray fields ($\varepsilon > 6.5$).

Larmor frequencies between the two spectra will be double the frequency due to the stray field. We found that the stray fields in the y-z plane of our apparatus were generally around 0.05 μT, which corresponds to a Larmor frequency of about 2 Hz.

On the other hand, measuring stray fields along the axis of the gradiometer is not so straightforward, because we cannot observe the precessing magnetization. Our typical method of determining the residual flux along the axis of the insert was to apply a small field along either y or z. The field along x will add in quadrature, so by applying a B_x we could observe a change in the precession frequency; a minimum in the precession frequency indicated the best cancellation field.

For very precise cancellation, this procedure could become tedious in our apparatus. If the residual field along x is small, then one wants to apply a small transverse field so as not to swamp the B_x field. However, if the Larmor frequency is below the cutoff frequency of our high-pass filter, the experiment has to be done using an indirect method similar to the phase-encoding techniques we used for our low-frequency images. To acquire a point in the low-frequency FID, we had to polarize the sample, let it evolve in the small field for some given time, then apply the $\mathbf{B_m}$ field to acquire an indirect FID; the Fourier transform of the indirect FID yields the real and imaginary values for one point of the low-frequency FID. This procedure had to be repeated for each point in the low-field FID.

We found that the residual fields in our apparatus could vary day to day, sometimes by as much as a factor of two. Careful cancellation of the residual field would therefore require measuring the stray field every day, possibly more than once. It is clear that if precise cancellation is desired, a new technique to measure the residual fields must be devised. Nevertheless, we did acquire images with at least partial cancellation of the residual field; these images are shown in Fig. 6.15. The sequence parameters were identical to those of the images shown in Sec. 6.3.4. Because of the cancellation of residual fields, we estimate $\varepsilon > 10$ in the ZFMRI image of Fig. 6.15B and $\varepsilon > 6.5$ in the conventional-sequence image of Fig. 6.15C.

Chapter 7

Discussion

We showed in Ch. 6 that it is possible to obtain undistorted magnetic resonance images in the absence (or near-absence) of a uniform static magnetic field. This is in itself a significant result, but the apparatus as it currently exists is unsuitable for many practical applications. In this chapter we will discuss topics relevant to improving the utility of the system, including how to reduce the imaging time, the effect of an ambient field, and power dissipated in the uniform field pulses. We will also briefly introduce other techniques for the elimination of concomitant gradient distortions.

7.1   Imaging time

Our implementation of ZFMRI, which was described in Sec. 6.1.3, involved the time-consuming technique of phase encoding to acquire the one-dimensional k-space data, that is, acquiring each point of k-space individually. The reasons to use such a technique were two-fold. First, the indirect method can detect the real and imaginary parts of the signal (known as quadrature detection) with a single sensor. Second, the magnetic field pulses used in the ZFMRI sequence generated relaxation signals from the Pyrex and metal of the experimental apparatus [27, 30], which had to be filtered out (via a high-pass filter) to avoid saturating the NI DAQ card used to record the data. The indirect detection technique permitted us to use an acquisition frequency above the cutoff frequency of the filter.

For each of the 24 projections of the image we acquired 24 k-space points, each taking 3.5 seconds (2-s polarizing pulse followed by up to 0.5-s encoding time and 1-s data acquisition time), and each point was averaged 10 times to increase the SNR, leading to a total imaging time of about 5.6 hours. The time could be reduced substantially if the system were modified to allow stroboscopic acquisition; that is, to allow the FIDs to be acquired during the gradient pulses, so that all k-space points could be acquired after a single polarization. Such a system would require two orthogonal detectors for quadrature detection and would require careful design to minimize magnetic relaxation and detector dead time. With a suitable system, we predict that the imaging time could be reduced to only a few minutes. In particular, stroboscopic acquisition would reduce imaging time by a factor of 24 (as only one repetition of the polarization-encoding-detection cycle would be needed), and eliminating signal averaging would save another factor of 10.

We now outline our calculations of the statements above. In MRI, the SNR is commonly defined as the signal amplitude divided by the standard deviation of the noise. For an acquisition time t_{acq}

$$SNR = \frac{\int\limits_0^{t_{acq}} s(t)\,dt}{\sqrt{\int\limits_0^{t_{acq}} \sigma_n^2\,dt}} = \left[\frac{\int\limits_0^{t_{acq}} s(t)\,dt}{\sqrt{t_{acq}}}\right]\frac{1}{\sigma_n}, \tag{7.1}$$

where σ_n is the standard deviation of the noise and $s(t)$ is the integrated signal from the nuclear magnetization detected by the sensor (assuming negligible noise) [1]. To compare the SNR of two acquisition methods, we estimate the value of the bracketed term from the formula

$$s(t) \propto \int\limits_y \int\limits_z m(y,z)\exp[-t/T_2^*]\exp[-i\gamma t B_G(y,z)]dy\,dz \tag{7.2}$$

where T_2^* is the transverse relaxation time, $B_G(y,z)$ is the field due to applied gradients, and $m(y,z)$ is a function representing the spin distribution in the sample. In this analysis we use a function $m(y,z)$ which is uniform over the imaging region and normalized such that

$$\int\limits_y \int\limits_z m(y,z)dy\,dz = 1 \tag{7.3}$$

The standard deviation of noise is time independent and assumed to originate from the sample, detector, and electronics, which are the same in both acquisition methods; therefore, the σ_n term has been omitted from the following comparison.

In the one-dimensional phase-encoding (pe) acquisition method described in Sec. 6.1.3, the signal at each point in k-space is acquired as a free induction decay (FID) in the uniform field $\mathbf{B_m}$. The demodulated signal equation in this case is given by

$$s_{pe}(t) \propto \int\limits_y \int\limits_z m(y,z)\exp[-t/T_2^*]dy\,dz, \tag{7.4}$$

where, in our experiments, the transverse relaxation time T_2^* was measured to be 300 ms. Using our acquisition time $t_{acq} = 1$ second,

$$SNR_{pe} \propto \frac{\int\limits_0^{t_{acq}} s_{pe}(t)\,dt}{\sqrt{t_{acq}}} \approx 0.289. \tag{7.5}$$

In a stroboscopically-detected experiment using two orthogonal detectors, the signal could be detected as precession about the gradient field $\mathbf{B_G}$ during every second gradient pulse. The

demodulated signal equation in this case is given by

$$s_{srt}(t) \propto \int\limits_y \int\limits_z m(y,z) \exp[-t/T_2^*] \exp[-i\gamma t B_G] dy\, dz, \qquad (7.6)$$

where

$$B_G = \sqrt{(Gy)^2 + (Gz)^2} \qquad (7.7)$$

is the field magnitude at the point (y, z). For our value $G = 100\ \mu$T/m and a 5-ms acquisition, we find

$$SNR_{str} \propto \frac{\sqrt{2} \int\limits_0^{t_{acq}} s_{str}(t)\, dt}{\sqrt{t_{acq}}} \approx 0.0990. \qquad (7.8)$$

The factor of $\sqrt{2}$ arises from the use of two detectors in quadrature detection.

We now compare the difference in SNR between the two acquisition methods. If the noise standard deviation σ_n is the same in both cases, we find

$$\frac{SNR_{str}}{SNR_{pe}} \propto \frac{0.0990}{0.289} \approx 0.34. \qquad (7.9)$$

The SNR in the stroboscopically-detected experiment is reduced by a factor of approximately three compared that of the phase-encoded experiment, while the imaging time is reduced by a factor of 24.

Elimination of signal averaging reduces imaging time by an additional factor of 10, at the cost of an additional factor of approximately three in SNR [1]. Thus, a 240-fold reduction in imaging time would cause the SNR to drop by approximately a factor of nine. The loss in SNR could be recovered by increasing the prepolarization field from 10 mT to 100 mT. If the full 240-fold reduction were realized, the imaging time would be less than two minutes. It must be noted that this estimate of minimum imaging time is specifically for the images shown in Ch. 6. Changes in noise level, desired FOV, or desired resolution may result in different imaging times.

7.2 Effect of an ambient uniform magnetic field

The experimental apparatus described in Sec. 6.2.3 is contained inside a dewar with a diameter of about 0.5 m, so it is fairly easily to achieve good shielding from ambient magnetic fields (for example, by putting the dewar inside a mu-metal can). Such a system is suitable for small samples such as that shown in Fig. 6.12. However, the required size of the apparatus increases with the desired sample size. For a system suitable to image parts of the human body, magnetic shielding rapidly becomes cumbersome. This may not be a problem for a stationary system, but it is inconvenient if one desires a portable imager. As discussed in Sec. 5.5, the ZFMRI sequence can be generalized to the case of unshielded imaging in the presence of a uniform ambient field $\mathbf{B_a}$.

The presence of a uniform ambient field imposes conditions on the gradient and π pulses. The maximum gradient pulse duration is limited by the need to keep the precession angle small

Fig. 7.1: Effect of an ambient uniform field on π pulse amplitude and direction.

during a gradient pulse. In practice, though, we find that the ZFMRI sequence is quite robust—in the image shown in Fig. 6.15B, the maximum precession angle during a gradient pulse was approximately $65°$. For a total field of $50~\mu$T (approximately the Earth's field), an upper bound of $65°$ limits the maximum duration of the gradient pulse to about $85~\mu$s: the angle of precession during a field pulse of amplitude B and duration τ is given by $\delta = \gamma B\tau$. For $\delta = 65°$ and $B = 50~\mu$T we find $\tau \approx 85\mu$s. Restrictions on the maximum gradient pulse duration place demands on the imaging hardware; in particular, in order to acquire k-space points stroboscopically the detector's dead time must be less than the gradient pulse duration. Limiting the gradient pulse duration also limits the time during which the signal may be acquired, which may have implications for the SNR of the image.

The presence of $\mathbf{B_a}$ also sets conditions on the amplitude and duration of the π pulse. Components of $\mathbf{B_a}$ perpendicular to the π pulse induce errors by modifying the pulse amplitude and direction. To limit the error in the total π pulse amplitude to less than 1%, the applied pulse amplitude must be approximately seven times the perpendicular component of $\mathbf{B_a}$. We explain this estimate as follows.

We define the z-axis to be along the direction of the applied π pulse. In the presence of an ambient field the total field amplitude along z is given by $B_{\pi z} + B_{az}$, where $B_{\pi z}$ is the applied pulse (generated by our field coils) and B_{az} is the component of ambient field parallel to the z-axis (see Fig. 7.1). For a component of ambient field perpendicular to the z-axis (in the x-y plane) B_{axy}, the total field during the π pulse is given by

$$B_{\pi,tot} = \sqrt{(B_{\pi z} + B_{az})^2 + (B_{axy})^2}. \tag{7.10}$$

Thus $B_{\pi,tot} \leq 1.01(B_{\pi z} + B_{az})$ for $B_{axy} \leq B_{\pi,tot}/7$.

The presence of large perpendicular components B_{axy} would therefore require large applied $B_{\pi z}$ pulses — approximately $350~\mu$T if $B_{az} = 0$. An obvious solution is to align the π pulse with $\mathbf{B_a}$. The total field during the π pulse is then

$$B_{\pi,tot} = B_{\pi z} + B_{az}, \tag{7.11}$$

where $B_{\pi z}$ is the pulse applied by the field coils. Thus, if one aligns the system so that the ambient field is aligned with the desired π pulse direction, one may be able to use the ambient field as the π pulse — that is, $B_{\pi,tot} = B_a$ and $B_{\pi,z} = 0$. As long as the Earth's field is aligned with the desired π pulse direction to within about $8°$, the $B_{axy} \leq B_{\pi,tot}/7$ condition will be satisfied and the error in the amplitude of $B_{\pi,tot}$ will be less than 1%.

7.3 Power dissipated in π pulses versus power dissipated in a static field

Although the imaging sequence we have implemented does not include a uniform static field, it does feature pulses of uniform field. These pulses are designed to cause the spins to precess by π over their duration. The amplitude of the pulse is therefore dictated by the duration. If the desired pulse time is very short, a large uniform field will be required; for example, a 100-μs pulse requires $B_\pi \approx 120~\mu$T, which corresponds to a Larmor frequency of about 5 kHz. This field is similar in magnitude to fields commonly used in SQUID-detected MRI and in which concomitant gradients have not been a problem [2, 3, 5, 26, 44, 45]. A common question is why, if such fields can be achieved in the apparatus, do we not make them static and use them as $\mathbf{B_0}$ in a conventional imaging sequence? The answer to this question depends on one's reason for pursuing ZFMRI. In the first place, one simply may not want such a large $\mathbf{B_0}$; advantages of small $\mathbf{B_0}$ fields were discussed in Ch. 1. Another reason to not use a static $\mathbf{B_0}$ is power consumption; if one is examining ZFMRI for its suitability in a portable or otherwise low-power-consumption system, one way to reduce power is to eliminate the $\mathbf{B_0}$ field. Then the question we must ask is this: If we were to use a static uniform field which consumes the same power as our uniform field pulses, what ε could we achieve? This will tell us if the apparatus would require less power to use a static uniform field suitable for imaging.

In order to consider this question, we will find an expression for ε_{static}, which is the ε value we would have for the same gradient field, but with the power from the π pulses instead used to maintain a static field. For simplicity we will assume that the π pulses are rectangular, with constant amplitude over the entire duration (that is, we neglect the current rise and fall times).

The π pulse in our apparatus is powered by a 9-volt alkaline battery and the amplitude is adjusted by means of a potentiometer in series with the coil (see Sec. 6.2.4). Because we use a superconducting coil to provide the field, nearly all of the power in the circuit is dissipated in the resistor. We will first find an expression for ε_{static} in this case, where the voltage is constant and the current for the π pulses is adjusted by changing the resistance R of the circuit.

For a constant voltage, the power dissipated in the resistor is given by the familiar

$$P = IV, \tag{7.12}$$

where V is the voltage across the resistor and I is the current thorough the resistor. Assuming a steady current I, the magnetic field can be found with the well-known Biot-Savart Law [46],

$$\mathbf{B} = \frac{\mu_0 I}{4\pi} \int \frac{d\mathbf{l} \times \mathbf{r}}{|\mathbf{r}|^3}, \tag{7.13}$$

where μ_0 is the permeability of free space. Since the field amplitude is linear in I, we will write it generally as

$$|\mathbf{B}| = B = \alpha I, \tag{7.14}$$

where α is the proportionality constant. Note that P and B are both linear in I, so we can write

$$P = \frac{V}{\alpha} B = \beta_V B, \tag{7.15}$$

where β_V is a proportionality constant (recall that V is constant). The power dissipated in a π pulse is then $P_\pi = \beta_V B_\pi$. The duty cycle of the π pulse is given by

$$D = \frac{t_\pi}{t_\pi + \tau}, \tag{7.16}$$

where t_π is the duration of the π pulse and τ is the duration of a gradient pulse (see Fig. 6.2). The average power P_{avg} during the pulse sequence is

$$P_{avg} = \beta_V B_\pi D. \tag{7.17}$$

This gives the same power dissipation as we would get for the creation of a static field,

$$B_{static} = B_\pi D = B_\pi \frac{t_\pi}{t_\pi + \tau}. \tag{7.18}$$

The pulse amplitude B_π is chosen to cause the magnetization to precess by π, so B_π is related to t_π by $\gamma B_\pi t_\pi = \pi$. If we write B_π in terms of t_π we find

$$B_{static} = \frac{\pi}{\gamma t_\pi} \frac{t_\pi}{t_\pi + \tau} = \frac{\pi}{\gamma} \frac{1}{t_\pi + \tau}. \tag{7.19}$$

In order to calculate ε_{static}, we must compare this B_{static} to GL where G is the gradient strength and L is the width of the FOV. The field of view L is defined as [1]

$$L = \frac{1}{\frac{\gamma}{2\pi} G \Delta t}, \tag{7.20}$$

where Δt is the time separation between data points; in our sequence this is $\Delta t = 2\tau$. Therefore

$$GL = \frac{2\pi}{\gamma} \frac{1}{2\tau} = \frac{\pi}{\gamma \tau}. \tag{7.21}$$

We can now calculate ε_{static}:

$$\varepsilon_{static} = \frac{GL}{B_{static}} = \frac{t_\pi + \tau}{\tau}. \tag{7.22}$$

This equation gives the value of ε_{static} based on parameters from our ZFMRI sequence. It is obvious from Eq. 7.22 that $\varepsilon_{static} \geq 1$. That is, when the current through the uniform field coil is provided by a constant voltage and varied by adjusting a resistor in series with the coil, a static uniform field which dissipates the same power as the π pulses will not be large enough to avoid severe concomitant gradient distortions. A static uniform field which is appropriate for imaging at the same gradient strengths will dissipate much more power than the ZFMRI sequence.

The preceding analysis assumed a constant supply voltage for the uniform field coils. We consider here another method of controlling the current in these coils: the use of a variable voltage supply with a fixed resistor in series with the field coils. In this case it is convenient to use the power expression

$$P = \frac{V^2}{R}. \tag{7.23}$$

By replacing I with V/R in the Biot-Savart Law we find

$$B = \frac{\alpha V}{R}, \tag{7.24}$$

where V is the voltage and R is the resistance. It is obvious from these equations that

$$P = \left(\frac{R}{\alpha^2}\right) B^2 = \beta_R B^2, \tag{7.25}$$

where β_R is a proportionality constant. The average power dissipated is then

$$P_{avg} = \beta_R B_\pi^2 D, \tag{7.26}$$

where D is the duty cycle of the π pulse. This is the same power dissipated in the creation of a static field

$$B_{static} = B_\pi \sqrt{D} = \frac{\pi}{\gamma} \sqrt{\frac{1}{t_\pi (t_\pi + \tau)}}. \tag{7.27}$$

Following Eqs. 7.21 and 7.22 we find

$$\varepsilon_{static} = \frac{GL}{B_{static}} = \frac{\sqrt{t_\pi (t_\pi + \tau)}}{\tau}. \tag{7.28}$$

If we assume that the π pulses are very short, so that $t_\pi \ll \tau$, then Eq. 7.28 can be approximated as

$$\varepsilon_{static} \approx \sqrt{\frac{t_\pi}{\tau}}. \tag{7.29}$$

From Eq. 7.29, we can see that if one uses a high B_π (and thus a short t_π) and a small duty cycle, he may be better off (from a power consumption standpoint) applying a static field rather than using the ZFMRI sequence. It should be noted, however, that a very small duty cycle ($D \sim 0.01$) is required for $\varepsilon_{static} < 0.1$.

The energy dissipated during a π pulse is given by

$$E_\pi = P_\pi t_\pi, \tag{7.30}$$

where t_π is the π pulse duration. Recall that this duration is related to B_π by

$$t_\pi = \frac{\pi}{\gamma B_\pi}. \tag{7.31}$$

We find, then, that the energy dissipated during a π pulse is simply

$$E_\pi = \frac{\pi}{\alpha \gamma} V, \tag{7.32}$$

where α is the proportionality constant between field magnitude B_π and the current I in the coils which generate the pulse. From this we see that, for a given set of uniform field coils, the energy consumed during the π pulse is proportional to the voltage used to create the pulse.

Equation 7.32 implies several tradeoffs in the design of a practical ZFMRI system. Using a low voltage to power the π pulses reduces the energy used by the system, but applying large fields requires a small R (to achieve large I), or coils of larger inductance; either of these increases the L/R time constant of the coils and makes rapid pulsing more difficult. In addition, a system with resistive coils will not have an infinitely adjustable R, which may set a lower bound on V. A versatile system, capable of a large range of B_π values (and therefore a large range of t_π times) will necessarily require a large available range of R or V. Furthermore, while short t_π times are desirable from an imaging standpoint, they may require more voltage to apply and negatively impact the energy economy of the system.

7.4 Comparison of ZFMRI to other methods of dealing with concomitant gradients

With the increasing popularity of low-field MRI [2, 3, 5, 26, 44, 45, 47–49], there has been a surge of interest in the effects of concomitant gradients [7, 35] on imaging, and on the correction (or avoidance) of concomitant gradient distortions [8–11]. In this section we will briefly examine some other methods for mitigating the problem of concomitant gradients on imaging as B_0 approaches zero. We will examine these methods mostly on their ability to correct concomitant gradient distortions, but we will take into account other pertinent factors such as computational complexity and hardware requirements. We will also consider PatLoc, a high-field imaging scheme which is not designed to correct concomitant gradient distortions but which uses an array of detectors to perform imaging in nonlinear gradients [50–52].

7.4.1 Computational corrections in post-processing

The concomitant gradient correction technique of Myers *et al.* [8] begins with a Taylor-series expansion of the magnetic field magnitude. For example, possible phase and frequency encoding fields are

$$\mathbf{B_{phase}} = \left(-\frac{1}{2}xG_z\right)\hat{\mathbf{x}} + \left(zG_y - \frac{1}{2}yG_z\right)\hat{\mathbf{y}} + (B_0 + yG_y + zG_z)\hat{\mathbf{z}} \qquad (7.33)$$

and

$$\mathbf{B_{freq}} = zG_y\hat{\mathbf{y}} + (B_0 + yG_y)\hat{\mathbf{z}}. \qquad (7.34)$$

The $\hat{\mathbf{z}}$ components of these fields are the "ideal" fields; the $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ components are the concomitant gradient terms. The Larmor frequency of a spin at some point in each of these fields is given by $\omega = \gamma B$ where B is the magnitude of the field at that point. The field magnitudes can be expanded in the y-z plane as follows:

$$B_{phase} = B_0 + G_y y + G_z z + \frac{G_y^2}{2B_0}z^2 + \frac{G_z^2}{8B_0}y^2 - \frac{G_y G_z}{2B_0}yz + O\left(\frac{G^3}{B_0^2}\right), \qquad (7.35)$$

$$B_{freq} = B_0 + G_y y + \frac{G_y^2}{2B_0}z^2 + O\left(\frac{G_y^3}{B_0^2}\right). \qquad (7.36)$$

The terms which are second order and higher in the gradient are due to concomitant gradients.

The series field expansion is used to divide the frequency and phase each into two parts, where ω^{ideal} and ϕ^{ideal} are due to the ideal fields (i.e. without concomitant gradients), while $\Delta\omega^c$ and $\Delta\phi^c$ are due to the concomitant gradient terms. The terms due to concomitant gradients are removed by defining a scaled time parameter

$$t' = \left(\frac{\omega^{ideal}}{\omega^{ideal} + \Delta\omega^c}\right)\left(t - \frac{\Delta\phi^c}{\omega^{ideal}}\right), \tag{7.37}$$

which is used to replace the time parameter in image reconstruction. Because these terms are position-dependent, the calculation is simplified by breaking the image into regions, and using the position of the region's center to evaluate t' over the whole region. Using a larger number of regions improves the concomitant gradient cancellation, but is more computationally intensive.

Equations 7.35 and 7.36 are based on the well-known binomial series

$$(1+x)^k = 1 + kx + \frac{k(k-1)}{2!}x^2 + \frac{k(k-1)(k-2)}{3!}x^3 + \ldots, \tag{7.38}$$

which converges for $|x| < 1$ [53]. We therefore cannot expect this expansion to accurately represent the field magnitude when $\varepsilon > 1$. Furthermore, simple numerical calculations indicate that the convergence of the series for $k = 1/2$ slows down as x approaches 1, so we would expect that the quality of the concomitant gradient correction will degrade as ε increases, unless higher-order terms are included in the expansion (which increases the computational complexity). These predictions have been borne out in simulations and experiments. The computational technique has been shown to be successful in both simulations and experiments with values of ε up to about 0.5. However, a simulation with $\varepsilon \approx 0.75$ was reported to retain unacceptable levels of distortion after the correction algorithm was applied [8].

An additional requirement for the correction in the phase-encode direction is that

$$\gamma B_0 T_{ramp} \gg 1, \tag{7.39}$$

where T_{ramp} is the time over which the gradient field is ramped. Clearly this criterion becomes difficult to satisfy as $B_0 \to 0$.

7.4.2 Expanding spin density as a Fourier series

In addition to developing the sequence used in our ZFMRI image, Meriles *et al.* [10] describe a second method of imaging which avoids concomitant gradient distortions. The second method, identified in the paper as "Scheme 2," uses a radically different approach to image encoding and reconstruction. Consider an object in the FOV as shown in Fig. 7.2. The spin density along a closed circular path centered at the origin must be a periodic function. More formally, in polar coordinates the spin density is a periodic function of the angle and can therefore be expressed as a Fourier series:

$$n(r, \theta) = \frac{1}{2\pi}\sum_{m=-\infty}^{\infty} c_m(r)e^{im\theta}, \tag{7.40}$$

Fig. 7.2: Spin density as a function of angle. (A) Sample (light) in field of view (dark). (B) Spin density vs. angle for the inner ring in A. (C) Spin density vs. angle for the outer ring in A.

where

$$c_m(r) = \int rn(r, \theta)e^{-im\theta} \, d\theta \qquad (7.41)$$

(note that previously in this thesis m was used to represent spin density; in this section we will follow the notation of [10]). The coefficients c_m can be determined from the signal acquired as the spins precess in the gradient field. To illustrate this, we will examine an example given in the paper [10].

The apparatus we will use in this analysis consists of a detector along the y-axis, along with coils that allow for prepolarization and π pulses to be applied along any axis. The apparatus is surrounded by a Golay pair which applies a gradient field of

$$\mathbf{B_{Golay}}(x, z) = g\left(z\hat{\mathbf{x}} + x\hat{\mathbf{z}}\right) \qquad (7.42)$$

in the x-z plane.

Assume the sample is polarized along the y-axis via a prepolarization pulse, then allowed to evolve in the gradient field. The detected signal will be

$$S_y(k_r) = \int_0^\infty \cos(k_r r)s_y(r) \, dr, \qquad (7.43)$$

where $k_r = gt$ (g is the gradient strength) and

$$s_y(r) = \int_0^{2\pi} rn(r, \theta) \, d\theta \qquad (7.44)$$

(recall that $n(r, \theta)$ is the spin density). It can be seen that the function $s_y(r)$ is related to the $m = 0$ coefficient of the Fourier series in Eq. 7.41.

If the prepolarization pulse is applied along x, then the signal received by the detector will be

$$S_x(k_r) = \int_0^\infty \sin(k_r r) s_x(r)\, dr, \tag{7.45}$$

where

$$s_x(r) = \int_0^{2\pi} rn(r,\theta) \cos\theta\, d\theta. \tag{7.46}$$

For the prepolarization pulse along the z-axis,

$$S_z(k_r) = \int_0^\infty \sin(k_r r) s_z(r)\, dr, \tag{7.47}$$

where

$$s_z(r) = \int_0^{2\pi} rn(r,\theta) \sin\theta\, d\theta. \tag{7.48}$$

These functions are related to the $m = \pm 1$ coefficients of the Fourier series (Eq. 7.41).

Higher-order coefficients can be found by preparing the initial spin state before acquisition. This can be done by applying a train of π pulses, of varying directions, after the prepolarization. The imaging experiment, then, will consist of a series of polarization-encoding-acquisition steps, as shown in Fig. 7.3. Each pulse in the train allows the determination of two higher-order coefficients (two higher values of m) as long as all of the lower-order coefficients are known. Because the spin density function $n(r,\theta)$ is real, c_m and c_{-m} are complex conjugates; therefore, a total of four additional coefficients of Eq. 7.40 can be found for each π pulse.

This Fourier-series imaging procedure, while quite different in the processing of the data, is similar in form to the procedure we used (as described in Sec. 6.1.3). Both techniques begin with a prepolarization pulse, and then spatial information is encoded by a train of π pulses, followed by detection of the precessing magnetization. As such, this Fourier-series imaging scheme shares some of the same disadvantages of our ZFMRI experiment: for example, errors due to inaccuracies in the π pulses will add, and the experiment will consist of a time-consuming repetition of the polarization-encoding-acquisition process. The Fourier-series imaging has some additional disadvantages of its own as well: as the number of coefficients increases, each new coefficient contributes less to the detected signal.

7.4.3 Rotating magnetic field gradients

As pointed out in Ch. 5, concomitant gradients lead to nonzero field components which point in some direction other than the z-axis. These field components are static in the lab frame, so they will rotate in the rotating frame. The rotating components will average away if they are sufficiently small.

A pair of orthogonal gradient coils, driven by currents 90° out of phase, can apply rotating gradient fields. If the frequency of the gradient rotation is ω_0, the field due to gradients will have

Fig. 7.3: Pulse sequence for "Scheme 2." (Reprinted with permission from [10].)

some stationary components in the rotating frame. Bouchard [11] showed that the existence of the stationary components improve the averaging, and therefore result in smaller phase errors and less image distortion than traditional static gradients for the same ε.

Bouchard considers two types of rotating frame gradients. Rotating frame gradients of the first kind are the sum of the fields created by two orthogonal sets of Golay coils, such as

$$\mathbf{B_a}(t) = a(t)\,(z\hat{\mathbf{x}} + x\hat{\mathbf{z}}) = g\cos(\omega t + \varphi)\,(z\hat{\mathbf{x}} + x\hat{\mathbf{z}}) \tag{7.49}$$

and

$$\mathbf{B_b}(t) = b(t)\,(z\hat{\mathbf{y}} + y\hat{\mathbf{z}}) = g\sin(\omega t + \varphi)\,(z\hat{\mathbf{y}} + y\hat{\mathbf{z}}), \tag{7.50}$$

where g is the gradient strength, ω is the frequency of the rotating gradient, and φ is the phase of the rotating gradient. Rotating frame gradients of the second kind are a linear superposition of the fields

$$\mathbf{B_a}(t) = a(t)\,(z\hat{\mathbf{x}} + x\hat{\mathbf{z}}) \tag{7.51}$$

and

$$\mathbf{B_b}(t) = \epsilon b(t)\,(-x\hat{\mathbf{x}} - y\hat{\mathbf{y}} + 2z\hat{\mathbf{z}}), \tag{7.52}$$

where $a(t)$ and $b(t)$ are as in Eqs. 7.49 and 7.50 and ϵ is a proportionality constant.

With either type of rotating frame gradient, the result is a region of enhanced linearity when compared with static gradients of the same strength g. The region of enhanced linearity is a square, tilted 45° relative to the original FOV. The area of the region is $1/\sqrt{2}$ of the area of the original FOV. Figure 7.4 shows simulated images taken with rotating frame gradients where $\varepsilon \sim 3.2$.

In simulations, the technique of rotating frame gradients seems effective in reducing distortions due to moderate-to-strong concomitant gradients (including $\varepsilon > 1$) at the cost of a reduction in the FOV. As this technique was not designed for $B_0 = 0$, it may be well-suited for use in the Earth's field. Because it is a phase-encoding technique, the entire image must be phase-encoded, or some other method for dealing with concomitant gradients in the frequency-encoding field must be used.

7.4.4 Signal acquisition with an array of detectors

In principle, if the field distribution over the sample is well-known it should be possible to reconstruct the image if one has a way to remove ambiguities in the encoding. For example, in the radially symmetric magnetic field shown in Fig. 5.1, measuring the precession frequency of a spin yields information on the distance from the origin to the spin, but no information on the direction. This ambiguity could be removed by the use of an array of detectors with well-known positions and sensitivity profiles. An example of such a method in high-field imaging is known as PatLoc [50–52]. It should be noted PatLoc was not developed as a technique for imaging at arbitrarily low B_0, nor for correcting concomitant gradient distortions; we examine it here as an example of imaging in nonlinear gradients using multiple detectors.

In PatLoc imaging, an array of coils is used to apply "spatial encoding fields," or SEMs, which in general are not linear. The SEMs are applied on top of a strong $\mathbf{B_0}$ field to spatially modulate the total magnetic field; in the literature, the term "SEM" is preferred to "gradient" to

Fig. 7.4: Simulated images with rotating frame gradients. (A) Proton density map. (B) Simulated image in conventional MRI gradients ($\varepsilon < 0.1$). (C) Simulated image with conventional MRI gradients ($\varepsilon \approx 3.2$). (D) Simulated image using rotating frame gradients ($\varepsilon \approx 3.2$). (Reprinted figure with permission from: L-S. Bouchard, *Phys. Rev. B* **74**, 054103 (2006).)

distinguish the SEMs from conventional linear imaging gradients. The exact form of the SEMs depends on the configuration of the coil array, but in general there is not a one-to-one mapping between spin position and the frequency/phase of its precession (that is, they are "non-bijective"). Some examples are the radial and multipolar fields shown in [51]. Imaging in these fields is possible because they are bijective (or nearly so) over local subregions; the ambiguities in the encoding can be addressed by using separate detectors in each subregion and applying parallel imaging techniques such as SENSE or GRAPPA in the reconstruction [51].

The PatLoc technique illustrates some of the disadvantages that may arise in a multi-detector method of concomitant gradient correction. First and foremost, the technique requires an array of detectors with well-known sensitivity profiles. The imaging field must also be well-known, which may not be the case in an unshielded ZFMRI system. Because of the nonlinear nature of the SEMs and the use of multiple detectors, the reconstruction procedure may be complicated and nonintuitive [51]. Finally, it has been observed in simulations that PatLoc images frequently have "holes," which are large regions where distortions and blurring render the image unusable. See Fig. 7.5 for example simulations of PatLoc imaging on a variety of phantoms; note the blurring at the centers of the images.

Fig. 7.5: Simulations of PatLoc imaging. (A,C,E) Simulations using conventional linear imaging
gradients. (B,D,F) Simulations using PatLoc gradients. (With kind permission from Springer
Science+Business Media: Magnetic Resonance Materials in Physics, Biology, and Medicine, Parallel
imaging in non-bijective, curvilinear magnetic field gradients: a concept study, vol. 21, 2008, pp.
5-14, J. Hennig et al., Fig. 7, © ESMRMB 2008.)

Bibliography

[1] M. E. Haacke, R. W. Brown, M. R. Thompson, R. Venkatesan, M. E. Haacke, R. W. Brown, M. R. Thompson, and R. Venkatesan, *Magnetic Resonance Imaging: Physical Principles and Sequence Design* (Wiley-Liss, New York, 1999).

[2] R. McDermott, N. Kelso, S.-K. Lee, M. Mößle, M. Mück, W. Myers, B. ten Haken, H. Seton, A. H. Trabesinger, A. Pines, and J. Clarke, J. Low Temp. Phys. **135**, 793 (2004).

[3] S.-K. Lee, M. Mößle, W. Myers, N. Kelso, A. H. Trabesinger, A. Pines, and J. Clarke, Magn. Reson. Med. **53**, 9 (2005).

[4] R. McDermott, A. H. Trabesinger, M. Mück, E. L. Hahn, A. Pines, and J. Clarke, Science **295**, 2247 (2002).

[5] M. Mößle, S.-I. Han, W. R. Myers, S.-K. Lee, N. Kelso, M. Hatridge, A. Pines, and J. Clarke, J. Mag. Reson. **179**, 146 (2006).

[6] W. R. Myers, Ph.D. thesis, University of California, Berkeley, 2006.

[7] D. A. Yablonskiy, A. L. Sukstanskii, and J. J. H. Ackerman, J. Mag. Reson. **174**, 279 (2005).

[8] W. R. Myers, M. Mößle, and J. Clarke, J. Mag. Reson. **177**, 274 (2005).

[9] C. A. Meriles, D. Sakellariou, A. H. Trabesinger, V. Demas, and A. Pines, Proc. Natl. Acad. Sci. U.S.A. **102**, 1840 (2005).

[10] C. A. Meriles, D. Sakellariou, and A. H. Trabesinger, J. Mag. Reson. **182**, 106 (2006).

[11] L.-S. Bouchard, Phys. Rev. B **74**, 054103 (2006).

[12] J. Clarke and A. I. Braginski, in *The SQUID Handbook, Vol. I: Fundamentals and Technology of SQUIDs and SQUID Systems*, edited by J. Clarke and A. I. Braginski (Wiley-VCH, Weinheim, 2004), pp. 1–28.

[13] J. Vrba, J. Nenonen, and L. Trahms, in *The SQUID Handbook, Vol. II: Applications of SQUIDs and SQUID Systems*, edited by J. Clarke and A. I. Braginski (Wiley-VCH, Weinheim, 2006), pp. 269–389.

[14] T. R. Clem, C. P. Foley, and M. N. Keene, in *The SQUID Handbook, Vol. II: Applications of SQUIDs and SQUID Systems*, edited by J. Clarke and A. I. Braginski (Wiley-VCH, Weinheim, 2006), pp. 481–543.

[15] H.-J. Krause and G. Donaldson, in *The SQUID Handbook, Vol. II: Applications of SQUIDs and SQUID Systems*, edited by J. Clarke and A. I. Braginski (Wiley-VCH, Weinheim, 2006), pp. 441–479.

[16] R. C. Black and F. C. Wellstood, in *The SQUID Handbook, Vol. II: Applications of SQUIDs and SQUID Systems*, edited by J. Clarke and A. I. Braginski (Wiley-VCH, Weinheim, 2006), pp. 391–440.

[17] B. Muhlfelder, J. M. Lockhart, and G. M. Gutt, Adv. Space Res. **32**, 1397 (2003).

[18] *The SQUID Handbook, Vol. I: Fundamentals and Technology of SQUIDs and SQUID Systems*, edited by J. Clarke and A. I. Braginski (Wiley-VCH, Weinheim, 2004).

[19] *SQUID Sensors: Fundamentals, Fabrication, and Applications*, edited by H. Weinstock (Kluwer, Dordrecht, 1996).

[20] T. Van Duzer and C. W. Turner, *Principles of Superconductive Devices and Circuits*, 2nd ed. (Prenctice Hall, New Jersey, 1999).

[21] C. Kittel, *Introduction to Solid State Physics*, 7th ed. (Wiley, New York, 1996).

[22] B. D. Josephson, Phys. Lett. **1**, 251 (1962).

[23] B. Chesca, R. Kleiner, and D. Koelle, in *The SQUID Handbook, Vol. I: Fundamentals and Technology of SQUIDs and SQUID Systems*, edited by J. Clarke and A. I. Braginski (Wiley-VCH, Weinheim, 2004), pp. 29–92.

[24] D. Drung and M. Mück, in *The SQUID Handbook, Vol. I: Fundamentals and Technology of SQUIDs and SQUID Systems*, edited by J. Clarke and A. I. Braginski (Wiley-VCH, Weinheim, 2004), pp. 127–170.

[25] R. Cantor and D. Koelle, in *The SQUID Handbook, Vol. I: Fundamentals and Technology of SQUIDs and SQUID Systems*, edited by J. Clarke and A. I. Braginski (Wiley-VCH, Weinheim, 2004), pp. 171–217.

[26] J. Clarke, M. Hatridge, and M. Mößle, Annu. Rev. Biomed. Eng. **9**, 389 (2007).

[27] S.-K. Lee, Ph.D. thesis, University of California, Berkeley, 2005.

[28] C. P. Slichter, *Principles of Magnetic Resonance*, 3rd ed. (Springer-Verlag, Berlin, 1990).

[29] F. Bloch, Phys. Rev. **70**, 460 (1946).

[30] R. F. McDermott, Ph.D. thesis, University of California, Berkeley, 2002.

[31] C. Kittel and H. Kroemer, *Thermal Physics*, 2nd ed. (W. H. Freeman and Company, New York, 1980).

[32] P. A. Bottomley, J. Phys. E Sci. Instrum. **14**, 1081 (1981).

[33] D. G. Nishimura, Principles of Magnetic Resonance Imaging (1996).

[34] J. J. Sakurai, *Modern Quantum Mechanics*, 2nd ed. (Addison-Wesley, Reading, MA, 1994).

[35] P. L. Volegov, J. C. Mosher, M. A. Espy, and R. H. Kraus Jr., J. Mag. Reson. **175**, 103 (2005).

[36] U. Haeberlen and J. S. Waugh, Phys. Rev. **175**, 453 (1968).

[37] U. Haeberlen, *High-Resolution NMR in Solids: Selective Averaging* (Academic Press, New York, 1976).

[38] *CRC Handbook of Chemistry and Physics*, 90th ed., edited by D. R. Lide (CRC Press, Boca Raton, 2009).

[39] R. B. Scott, J. W. Cook, and F. G. Brickwedde, J. Res. Nat. Bur. Stand. **7**, 935 (1931).

[40] W. E. Barr and V. J. Anhorn, *Scientific and Industrial Glass Blowing and Laboratory Techniques* (Instruments Publishing Co., Pittsburgh, 1959).

[41] E. L. Wheeler, *Scientific Glassblowing* (Interscience, New York, 1963).

[42] V. O. Altemose, in *7th Symposium on the Art of Glassblowing* (American Scientific Glassblowing Society, Madison, NC, 1962), pp. 61–70.

[43] See, for example, `http://www.omega.com/pdf/tubing/technical_section/chemical_chart_1.asp`.

[44] R. McDermott, S. Lee, B. ten Haken, A. H. Trabesinger, A. Pines, and J. Clarke, Proc. Natl. Acad. Sci. U.S.A. **101**, 7857 (2004).

[45] M. Mößle, W. R. Myers, S.-K. Lee, N. Kelso, M. Hatridge, A. Pines, and J. Clarke, IEEE Tran. Appl. Supercon. **15**, 757 (2005).

[46] J. D. Jackson, *Classical Electrodynamics*, 3rd ed. (Wiley, New York, 1999).

[47] S. Xu, V. V. Yashchuk, M. H. Donaldson, S. M. Rochester, D. Budker, and A. Pines, Proc. Natl. Acad. Sci. U.S.A. **103**, 12668 (2006).

[48] V. S. Zotev, A. N. Matlashov, P. L. Volegov, I. M. Savukov, M. A. Espy, J. C. Mosher, J. J. Gomez, and R. H. K. Jr., J. Mag. Reson. **194**, 115 (2008).

[49] I. M. Savukov, V. S. Zotev, P. L. Volegov, M. A. Espy, A. N. Matlashov, J. J. Gomez, and R. H. Kraus Jr., J. Mag. Reson. **199**, 188 (2009).

[50] J. Hennig, M. Zaitsev, and O. Speck, Proc. Intl. Soc. Mag. Reson. Med. **15**, 453 (2007).

[51] J. Hennig, A. M. Welz, G. Schultz, J. Korvink, Z. Liu, O. Speck, and M. Zaitsev, Magn. Reson. Mater. Phy. **21**, 5 (2008).

[52] J. Hennig, Apparatus and method for NMR tomography and acquisition with local magnetic field gradients in connection with local receiver coils, U.S. Patent No. 7,411,395 B2, 2008.

[53] R. E. Larson, R. P. Hostetler, and B. H. Edwards, *Calculus with Analytic Geometry*, fifth ed. (D. C. Heath, Lexington, 1994).