Title

Modelling, Designing, and Applying Portable, Low-Field Nuclear Magnetic Resonance

Permalink

https://escholarship.org/uc/item/03v03826

Author

Gruber, Daniel M

Publication Date

2021

Peer reviewed|Thesis/dissertation

Modelling, Designing, and Applying Portable, Low-Field Nuclear Magnetic Resonance

By

DANIEL M. GRUBER
DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Chemical Physics

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

Matthew P. Augustine, Chair

William Casey

James B. Ames

Committee in Charge

2021

MODELLING, DESIGNING, AND APPLYING PORTABLE, LOW-FIELD NUCLEAR MAGNETIC RESONANCE

BY

DANIEL GRUBER

Abstract

Without a doubt, nuclear magnetic resonance (NMR) has been an indispensable tool for structural determination in chemistry, molecular biology, and more. The power of NMR to elucidate positions and distances of atoms within molecules has allowed generations of chemists, chemical biologists, and molecular biologists to develop strategies for synthesizing new medicines, advanced materials, and to verify their syntheses. To this day, the core strategy from which this tool draws its power remains a drive to higher and higher magnetic field. For the purposes described above, it has been a necessary drive. A higher field allows for better signal to noise ratio (S/N) and narrows relative spectral linewidths, resulting in faster data acquisitions and for acquired data to be more granulated. There are, however, side effects that accompany the drive to higher field. Extremely strong magnetic fields are physically dangerous if any magnetizable materials come close enough to magnets. Equipment, tools, and some people themselves must therefore be kept a safe distance away, and instances where failures in this regard have led to deaths are not unheard of. The superconducting magnets that generate these intense fields are also quite large, are limited to tiny sample volumes, and require special maintenance, which can be expensive. Specifically, they require a constant supply of liquid helium in order to keep the coils at temperatures low enough to maintain their superconducting characteristics. This helium is not a renewable or ubiquitous resource and it gets more expensive by the year, as challenges to exploration, capture, and even political upheaval destabilize the global market for liquid helium. Meanwhile, advances in amplifier and NMR spectrometer technology have resulted in better control and sensitivity in low-field experiments. For these

reasons, and more, low-field NMR has become an ever-more attractive area of research for the advancement of NMR and magnetic resonance in general.

This work describes a body of research focused on design and data analysis for low-field NMR, which is loosely defined for the purposes of this manuscript as NMR performed in a static magnetic field, B_0, less than 0.1 T, or a Larmor frequency lower than about 4.4 MHz. Because NMR at low-field does not generally produce chemical shift-resolved spectra, data acquisition and analysis instead tends to focus on bulk magnetic relaxation rates and imaging. Furthermore, these parameters are often difficult to understand using direct mathematical models because samples investigated at low-field are often inhomogeneous and possess geometries that are not always amenable to symmetry-based mathematical simplifications. It is often better to instead use statistical correlation and categorization to understand experimental results from a higher, less granular level. This fact motivates an initial exploration into the use of partial least squares as a method to determine the electrical permittivity of aqueous electrolyte samples at extremely high pressures. While this is not a study of magnetic resonance, it serves as an example for the application of this technique to similar problems that might present themselves in NMR. The next chapter describes Matlab codes that have been written for simulating magnetic fields for arbitrary coil and magnet shapes. The final two chapters are an exploration into novel rf coil designs for use in single-sided NMR, and specifically in the presence of shielding due to conductive magnet material. Multiple coil designs are described and their performances characterized. Ultimately an unexpected "fringe coil" geometry is shown to be far superior in theoretical models and practical testing with a variety of sample types and geometries.

Acknowledgements

It seems common for many who pursue and attain their doctorates to follow a more-or-less direct path. They may not always know from a young age that they will go to graduate school, but they are generally of a type that finds success at each academic level, each in turn opening a door to the next, until one day they find themselves toiling away at papers, teaching, and eventually a dissertation. This was not my experience. For the greater portion of my youth and young-adulthood, I instead preferred to sabotage my own academic progress. It was only by the guidance, example, and sometimes punishment of others that I finally managed to get my mind right at just the last moment so I could take up this challenge and earn such an achievement. With that in mind, some acknowledgements are in order.

The first is to my parents, of course. Thank you for your patience and the amazing examples you both set. I am sorry that I put you through such trials and arguments over homework as a child. Punishing the teachers that I didn't like by not doing their homework was never a good strategy. Perhaps you will now have the satisfaction of watching me struggle with the same fights with a child (or children) who is just like me. I hope not.

My high school physics teacher, Michael Lampert, thought for some reason that I might be a good fit for a Toshiba ExploraVision Science competition team. I am not sure how he chose me, but his faith in me and our success at the national level changed how I viewed myself in science. I had never before considered myself as one that had any real potential in sciences.

Professor David Strom at the University of Oregon was my only physics professor that I felt cared much about my success. That meant an enormous amount to me and motivated me to

work harder – even earning some decent grades in his classes. To this day I still use the skills that I learned in his electronics classes with my current work. Years later, he still remembered me and wrote a letter of recommendation for my application to UC Davis, without which I could not have been accepted.

The United States Air Force, and specifically Special Tactics deserves a mention here. That three year training pipeline and deployment was one of the hardest, worst, and best times of my life. It taught me discipline that I so badly lacked and when I might otherwise have gotten down about tough times, it gave me perspective that allowed me to "embrace the suck." Memories of horrible physical punishment for failures were valuable to draw upon when my motivation would wane in the late hours of the night and I had an assignment due the next day. I should specifically throw a shout out to my former commander, Dan Schilling, as well as Roy Lofts, and John Buursma from Grey Team at the 125th Special Tactics Squadron. These men encouraged me, mentored me, and held me to high expectations.

After returning from Afghanistan, I thought I wanted to go to physical therapy school and took many post-baccalaureate classes at Sacramento City College. The faculty of the Chemistry and Biology Departments are some of the best I have ever had the pleasure of studying under. Specifically, I have to point out Dr. Bill Miller, who found out about my PT school plans and derailed them so he could steer me to UC Davis for Chemistry instead. Bill got me my first publication and continues to be a mentor to this day. I look forward to many years of collaborations with him. Without a doubt, I would not be writing this manuscript were it not for him and his colleague, Dr. Gang-Yu Liu, who advocated for me at UC Davis.

To Professor Matt Augustine and the members (both actual and de facto) of the Augustine lab: I could not have had a better place to spend my time in graduate school at UC

Davis. "Unorthodox" only partially describes this bunch of misfits. From the first day that I arrived in Matt's office to be met with the salutation, "dude!" to partying with the cover band at ENC, to teaching CHE125, to chasing hot springs, to everything in between (yes, research too), this lab was never dull. I am so lucky to have been able to spend my graduate school years with this bunch of nerds. My gratitude falls with extra weight to Dr. Gerry Ochoa who taught me the first practical NMR techniques to get me started; to Dr. Sophia Fricke who challenged me intellectually at a quantum level on a daily basis (and challenged me hiking too); and to Julia Kerr who was (is) my partner in crime when it comes to teaching, pressure studies, and inappropriate jokes. And I cannot forget Dr. Michele Martin who set the stage for pretty much all of my research and without whom I would not be moving on to NIST in Boulder. I can't wait to be working with you again. There will be many more science jogs in our future and, no, you will never figure out a way to make me mad at you.

Finally, to Ava, my wife, who has been with me at every stage leading up to and throughout this journey from college on. Watching you, I finally learned what it meant to actually study and what it took to be academically successful. You never stopped believing in me, even when it made me mad. Your unwavering confidence in me gave me my own confidence. You are the foundation of all of my achievements. Thank you.

Table of Contents

List of Figures

volume. In this case, the rf coil may be a surface loop, which generates a B_1 perpendicular to the surface as depicted, or a double-D coil, which would produce a B_1 field parallel to the surface, but perpendicular to B_0. The second common geometry is depicted on the right, in (b), in which the B_0 is generated perpendicular to the surface of the magnet, penetrating the sample volume directly, in which case the B_1 field must be produced parallel to the surface, usually with a double-D coil. 42

perpendicular to the z-direction without and with a conducting surface placed h_1 = 6 mm below the top of the rf coil complex. The top and bottom rows in (a) and (b) respectively correspond to the rf field calculated in the h-w and h-z planes. The right column shows the standoff performance of the coils obtained in the h-direction at the z = w = 0 center of the rf coil. The solid and dashed lines show the rf field value without and with the conducting surface as a function of distance h_2 from the coil surface. The vertical gray line shows the position of the homogenous static magnetic field volume where the shielding performance is evaluated. 111

Figure 5.5. Magnetic resonance two-dimensional slices, extracted from 3D magnetic resonance images obtained at a 4 MHz ^{1}H Larmor frequency for a large beaker of mineral oil placed on the surface of a similarly sized square loop coil, a fringe coil and a split fringe coil in (a) – (c) respectively. The images in the right column include an aluminum plate placed h_1 = 6 mm below the top of the rf coil while the images in the left column have the aluminum plate removed. The number inset on the lower right corner of the images in the right column is the ratio of the maximum signals obtained with and without the conducting plate. 113

Figure 5.6. Summary of T_2^{obs} values obtained for light mineral oil with the standard fringe (a) and split fringe (b) coils as a function of beaker-to-magnet displacement, h_2. The T_2^{obs} values were obtained from a Matlab™ fit of the decay of the raw CPMG echoes to a single exponential function. In both plots, the circle captures the average value of T_2^{obs} obtained at h_2, while the bars capture the error in this average. As the value of h_2 increases, the S/N drops for a constant number of signal averages, and the error in T_2^{obs} increases. Note that at h_2 = 2.1 cm, the S/N for the standard fringe coil transient was too low to obtain a T_2^{obs} estimate. 116

Figure 5.7. Spin echo intensity as a function of rf pulse length. Here, the ordinate corresponds to the length of the first rf pulse in the spin echo experiment. This ordinate value is half the length of the second rf pulse in the spin echo experiment. The abscissa captures the total spin echo signal intensity. This value was obtained by integrating the raw spin echo signal in the time domain. Five spin echoes were obtained at each rf pulse length observed. The thick solid lines correspond to the average signal intensity, while the shaded regions communicate the 95% confidence in this measured spin echo intensity. The progressively darker shading corresponds to the real, imaginary, and absolute value of the signal intensity. Results are shown for the standard fringe and split fringe coil in (a) and (b), respectively. 118

Figure 5.8. Two real examples of the split fringe coil in action. The split fringe coil inserts into the existing groove of WA-1 unilateral magnet and the split fringe coil/unilateral magnet complex is placed into a water tight Pelican™ box lined with copper as shown in the picture on the left in (a). The enclosed NMR sensor was submerged in a bucket of laundry detergent, as shown in the upper right in (a), and application of the CPMG pulse sequence produced the decay shown on the left in (b), with a T_2^{obs} = 16 ms time constant obtained from the ILT on the far right. The lower right in (a) shows the placement of a branch segment on the top of the sealed magnet. Application of the CPMG pulse sequence yielded the transient decay shown in the left of (c). In this case, the ILT reports

List of Tables

Chapter One – Relevant Principles

"Chance favors the prepared mind."

– Louis Pasteur

1.1 Introduction

The explorations laid out in the following chapters center on the use of nuclear magnetic resonance (NMR). After a brief foray into large dataset regression in Chapter 2, the arc turns toward computer-model-aided design of coils and magnets. In particular, the use of NMR in a single-sided geometry and with a weak static magnetic field is the focus. As such, it is prudent to briefly review some of the key concepts of NMR for context.

1.2 Spin

The principle of magnetic resonance is based on the quantum mechanical concept of "spin." Spin is a fundamental property of matter, possessed by all elementary particles such as electrons, quarks, and so on. Because everything else is made up of these elementary particles, spin is therefore also possessed by everything composed of those elementary particles, including composite particles and atomic nuclei. As a fundamental property, spin exists in the same way that mass or charge exist. It just does.

Spin is also the most unfortunately named of the fundamental properties. From its name, one might imagine that all particles are turning on their own axes as they go about their business, like tiny little planets. That is not the case. The term, "spin," comes from the fact that this

property behaves like angular momentum, having magnitude and direction, except that a particle with spin is not orbiting anything else. It just appears to have this angular momentum within itself alone *as if* it were spinning [1].* This angular momentum is evident from the fact that a particle with non-zero spin and non-zero charge has an observable magnetic moment [2]. Finally, spin is also a quantized property. The value of a particle's total spin may be one half, one, zero, or any multiple of these. A proton, for instance, has a total spin of one half.

Although many particles possess the property of spin, this manuscript focuses specifically on *nuclear* magnetic resonance, and so from this point, only nuclear spins will be referenced. Moreover, the hydrogen nucleus will be the most common nucleus referenced, and the term "proton" will often be used in place of hydrogen because this is common parlance in the world of NMR and because the nucleus of a typical hydrogen atom is just a proton. Now speaking more generally again, while nuclei only have one value for total spin, their observed values may be any sum of positive or negative halves, whose positive-only total would sum to the positive total spin value. For instance, a spin-½ proton, having a total spin angular momentum of ½, may be observed as -½ or ½. A spin-1 lithium nucleus may be observed to be (-½ + -½), (-½ + ½), (½ + -½), or (½ + ½), which each sum to the possible observed values of -1, 0, or 1. Thus, it can be seen that the observed values of spin always differ from each other in integer increments. These observed values are referred to as states, and specifically for spin-½ nuclei in NMR such as hydrogen nuclei, the states +½ and -½ are referred to as up and down, or $|+\rangle$ and $|-\rangle$, respectively.

---

* At the time of its discovery, scientists may have actually thought that particles were spinning, and therefore given it this unfortunate moniker. However it came about, though, the name stuck and today, like Sisyphus, we must toil against the continuous task of explaining to poor young budding scientists that a "spin" is not actually spinning.

1.3 Observing spin: Zeeman splitting and the Stern-Gerlach Experiment

Of course, theory is fine, but if a theory cannot be verified by observation, it is useless. And indeed, the existence of spin can be observed in a very straightforward manner. As stated above, a non-zero spin-possessing nucleus has charge and will therefore have a magnetic moment. Furthermore, a nucleus with a magnetic moment will gain or lose potential energy when placed in a magnetic field according to the following equation, and as depicted in Figure 1.1.

$$E = -\boldsymbol{\mu} \cdot \boldsymbol{B} \tag{1.1}$$

Here, E refers to the energy, $\boldsymbol{\mu}$ refers to the magnetic moment vector, and $\boldsymbol{B}$ refers to the magnetic field at the location of the magnetic moment. Note that minus sign means that the energy is lower when the magnetic moment and the magnetic field are parallel, and higher when the two are antiparallel.

Figure 1.1. Two non-interacting magnetic moments, μ_a and μ_b, are depicted in a magnetic field, $\boldsymbol{B}$. The moment μ_a is oriented parallel to the field and therefore has low energy, whereas μ_b is oriented anti-parallel to the field and has relatively higher energy.

The magnetic moment of a nucleus is related to the spin via the gyromagnetic ratio, given by the symbol gamma, γ.

$$\boldsymbol{\mu} = \gamma \boldsymbol{I} \tag{1.2}$$

where $\boldsymbol{I}$ refers to the spin angular momentum of the nucleus. The values of each nucleus' gyromagnetic ratio can be found in existing tables. For a hydrogen nucleus, of spin-½, the values of $\boldsymbol{I}$ may be $\frac{1}{2}\hbar$ or $-\frac{1}{2}\hbar$ for spin up and down, respectively, where $\hbar$ refers to the reduced Planck constant, approximately 1.055 x10^{-34} J s [3]. As a note, this text and NMR in general will prefer the use of $\hbar$ over h, the standard Planck constant of approximately 6.626 x10^{-34} J/Hz [4].

By combining Equations 1.1 and 1.2, it is possible to derive an equation for the energy difference between spins of different states in a magnetic field, and then to create a diagram showing the energy splitting trend depending on the strength of the magnetic field in which the nucleus is placed. This energy splitting, dependent on the strength of the external field, is called the Zeeman effect, after Dutch physicist Pieter Zeeman and is shown mathematically in Equation 1.3.

$$\Delta E = E_- - E_+ = \left(\frac{\gamma\hbar B}{2}\right) - \frac{\gamma\hbar B}{2} = \gamma\hbar B \tag{1.3}$$

Figure 1.2 shows the simplest manifestation of such an energy splitting diagram, which depicts a spin-½ system, as is the case for the hydrogen nucleus.

E

$|-\rangle$

$\triangle E = \gamma \hbar B$

$|+\rangle$

B

Figure 1.2. A diagram depicting the splitting of the energy levels between spin-½ nuclei of up and down states. The energy of each state is given by negative one times the product of the spin, the gyromagnetic ratio, and the strength of the external field. The difference in the energy for the up and down states increases linearly with increasing magnetic field strength, B.

Now, if one could pass a beam of spin-bearing nuclei through a magnetic field, which had a linear gradient oriented transverse to the direction of travel, the magnetic moments of the nuclei would feel a deflecting force according to the spatial derivative of the magnetic interaction energy. Energy can be defined as the spatial derivative of the energy of an object and combining this with Equations 1.1 and 1.2 one arrives at the following:

$$\boldsymbol{F} = -\frac{dE}{dz} = \frac{d}{dz}\gamma I_z B_0 \tag{1.4}$$

where the magnitude of the magnetic field, B_0, has been used, rather than the vector, $\boldsymbol{B}$. By convention, the applied magnetic field, B_0, sets the z-coordinate of the system under investigation and this text will follow the same tradition [1]. Since the direction of the applied magnetic field determines the orientation of the spin being observed, Equation 1.4 has also been simplified to use only the z-component of the spin, I_z, rather than the full vector spin. If the heterogeneity of the applied magnetic field is linear in the z-direction, then B_0 can described simply as a scalar

gradient value, G_z, times the position, z, and Equation 1.4 then becomes constant within the region of the magnetic field:

$$F_z = \gamma I_z G_z \qquad (1.5)$$

It is this premise upon which the famous Stern-Gerlach experiments were performed in 1922 by Walther Gerlach, after their conception in the imagination of Otto Stern [5] [6] [7]. This experiment used a beam of silver atoms, which have neutral overall charge, but a nuclear spin of ½. This ensured that any deflection of the beam would be due to magnetic dipole interaction only, and not the Lorentz force of charge moving in a magnetic field. As shown in Figure 1.3, the atomic beam was passed through a linearly heterogeneous magnetic field, after which the atoms of the beam struck a screen for observation. The result was that, rather than a *distribution* of impacts spread out in the z-direction, two discreet impact *points* were observed, having been deflected either up or down, depending on the z-component of each atom's spin. The results of this experiment showed that spin is a quantized and intrinsic property.

Figure 1.3. The basic elements of the Stern-Gerlach Experiment are the source (left), the beam of neutral silver atoms (dotted line), a linearly heterogeneous external magnetic field (dB_0/dz) generated by magnets above and below the beam (indicated with "N" and "S" labels), and a detecting screen (right). As the silver atoms pass through the magnetic field, the magnetic moments from their nuclear spins interact with the external field, B_0, which imparts a deflecting force on the individual atoms, depending on their spin states. Up spins are deflected toward the stronger positive magnetic field and vice versa for the down spins. The result is two discrete points of impact at the screen.

1.4 Beyond the Stern-Gerlach Experiment: From Rabi to modern NMR

Unfortunately, it was not immediately realized that the results of the Stern Gerlach experiment were due to spin, even by Stern and Gerlach themselves. Instead, they slightly missed the mark and attributed the splitting to something called "space quantization in a magnetic field." Nonetheless, the experiment proved to be pivotal in lending experimental evidence to a number of firsts, including the quantization of the nuclear angular momentum, the strength of a nuclear magnetic moment, and the splitting of a beam based on the total angular momentum of its nuclei, to name a few [8]. It was not until Isidor Rabi's experiments with particle beams in 1939 that true nuclear magnetic resonance was actually performed. For these efforts, Rabi was awarded the Nobel Prize in 1944.

To understand Rabi's advance over the Stern-Gerlach experiment, it is helpful to look back at the Zeeman splitting diagram (Figure 1.2), where another important insight can be

grasped. If there is an energy difference between the states of a spin system in a static magnetic field, then by the principles of quantum mechanics, the spin state may undergo a photon-mediated transition. This may happen if the interacting photon is of the same energy as the energy splitting of the spin system, also called being *resonant*. From the famous Planck-Einstein relation, given here as Equation 1.6, this can be equated to a frequency. The energy of a photon, E_p, with frequency ν, or angular frequency ω, is

$$E_p = h\nu = \hbar\omega \qquad (1.6)$$

The specific frequency of this splitting for a particular nucleus is called the Larmor frequency (ω_0) and it will be referenced often in this text.

Figure 1.4. The three ways that states are transitioned by photons in a two-level spin system are a) absorption of a resonant photon to transition to a higher energy state, b) spontaneous emission of a photon by the spin, enabling a transition to a lower energy state, or c) stimulated emission, in which a passing resonant photon elicits an emission of a coherent photon from a spin transitioning from a high energy state to a low energy state.

The three main processes by which this may happen are depicted in Figure 1.4, which shows a simple two-state system, as is typically the case in NMR. These three processes are:

1. Absorption, in which a spin in a low energy state is transitioned to a high energy state by absorbing a resonant photon.

2. Spontaneous emission, in which a spin in a high energy state transitions to its low energy state, releasing a photon of the energy described by the splitting.

3. Stimulated emission, in which a resonant photon passes a spin in its high energy state, causing it to release its own photon of the same energy, phase, and polarity (called coherent), while simultaneously transitioning to its low energy state.

From these principles, one can imagine that a spin could be placed in a static field, generating an energy level splitting, and then a second oscillating field could be introduced to provide radiofrequency (rf) energy that would be able to induce transitions in the spin energy state.

Rabi's experiment was a modification of the Stern-Gerlach experiment that used precisely this concept. As in the Stern-Gerlach experiment, it used a particle beam, but instead of a single inhomogeneous magnetic field, Rabi used two successive inhomogeneous fields, whose directions were aligned, but with opposed gradients [9]. With this configuration, a beam of particles with known spin could be passed through the contraption at an offset angle and, if there were no other modifications, the beam would be initially deflected back toward center. The beam would then overshoot in the opposite offset direction, then be deflected again back toward center in the second field gradient, to strike a detector. However, as shown in Figure 1.5, Rabi also

added a component between the two field gradients. Here, a homogeneous magnetic field, B_0, was inserted, oriented in the same direction as the gradients, as well as an oscillating magnetic field, B_1, which was oriented perpendicular to B_0. If the frequency of B_1 oscillations matched the Larmor precession frequency of the nuclei in the beam, the nuclei and B_1 would be *in resonance* and the spin state of the nuclei could undergo a transition. The function of this section of the apparatus was to disturb the magnetic moment orientation of the spin beam so that, upon entering the second field gradient, the overall returning force presented by the second gradient was no longer equal and opposite that of the first gradient. Spins which had their orientation deflected by the central resonating component would not strike the detector and the intensity of the beam at the detector would be diminished. The frequency of H_1 oscillation was able to be finely controlled and in this way, the Larmor frequency – and thus the magnetic moment – of the atoms in the beam could be probed with exquisite detail.

Figure 1.5. The Rabi experiment directed a beam of spin-possessing particles (dashed line) between two inhomogeneous magnetic fields generated by sequential magnets, labeled "Magnet A" and "Magnet B" here. With only these magnets, the beam could be aimed so the path would be equally curved back and forth in the two regions at the same rate, landing at the detector on the right. The introduction of a third homogeneous magnetic region, generated by "Mag. C," along with perpendicular oscillating field, B_1, allowed the experimenter to flip the spins in the region of Mag. C. This flip changed the direction of the particle, causing the path in region B to veer from the detector, the path represented here by the dotted line. The probability of a spin state change is highest when B_1 is resonant with the spin states in a static B_0 field. Thus, by knowing the strength of B_0 and frequency of B_1, the gyromagnetic ratio of the nuclei under observation could be determined with precision.

Today, magnetic resonance is detected using Faraday induction and without the need for a molecular beam, as depicted in Figure 1.6. This method was pioneered simultaneously by Felix Bloch's group at Stanford and Edward Purcell's group at Massachusetts Institute of Technology between 1945 and 1947 [10] [11] [12]. Bloch's papers lay out the classical math that is still used today as the simplest way of understanding nuclear magnetic resonance induction, namely the Bloch sphere. Purcell's papers described both a resonant cavity method, as well as a solenoid method. It is the second that is most commonly used today. In this method, a sample – typically liquid – is placed inside a solenoid that generates the oscillating B_1 field – the *rf coil*. This assembly is then placed inside a larger magnet that generates the permanent B_0 field. This B_0 field is most commonly generated by super-cooled coils and, at the time of this writing, are capable of generating fields in excess of 25 tesla, corresponding to a proton Larmor frequency of over a GHz. Signals may be generated and detected by a continuous wave (CW) method, in which the rf coil is used only to transmit and another receiver coil is placed, with its axis orthogonal to both B_0 and B_1, so that the experimenter may "listen" while the rf coil is transmitting. Today, however, the most common method by far is to use "pulsed" NMR. In this method, the rf coil is used for both transmitting pulses as well as for receiving.

Figure 1.6. Today, the typical NMR setup consists of a sample placed inside an RF coil, depicted here as the solid swirling line around the sample. This generates the perturbing field (B_1), which is perpendicular to the powerful, homogeneous, static magnetic field (B_0) generated by superconducting coils, which are depicted here as cutaway coils in the background.

In the simplest pulsed NMR experiment, the *free induction decay* (FID), a single rf pulse is used to deflect the bulk magnetization vector away from its equilibrium along the B_0 axis, so that it lays down into the plane perpendicular to the B_0 axis. By NMR convention, the B_0 axis assigns the z-direction, so that the perpendicular plane is then the *xy*-plane. At this point, left to itself in a static field, the bulk sample magnetization vector will precess in the *xy*-plane at the Larmor frequency. Using a bar magnet to represent the bulk sample magnetization, this process is shown in Figure 1.7.

Figure 1.7. In the FID experiment, a sample starts out magnetized by B_0 along the z-axis as shown on the left. A pulse from the rf coil of the correct power, frequency, and time will turn the bulk magnetization to lie in the *xy*-plane, after which it will precess. This precession of a magnetic field inside a coil induces a current in the coil, which is then detected by the spectrometer.

As can be seen, this precession is in the same plane that contains the axis of the rf coil and therefore the changing flux produced by the rotating magnetization induces a current in the coil. This induction obeys the Faraday law of induction, which says that an electromotive force, ε, is generated by a changing magnetic flux through the area defined by a coil, Φ_B.

$$\varepsilon = -N\frac{d\Phi_B}{dt} \tag{1.7}$$

where N refers to the number of loops that make up the coil [13]. This current is detected and recorded as the precessing bulk magnetization shrinks in the *xy*-plane and regenerates along the z-axis, returning to its equilibrium state, a process called *relaxation*. The resulting signal that is induced in the rf coil from the freely precessing and relaxing magnetization is recorded as a

decaying oscillation – a free induction decay signal. This FID can then be Fourier transformed to reveal a spectrum of the frequency components that make up the oscillations in the initial signal.

1.5 A brief summary of Bloch spin mechanics

1.5.1 Magnetization, equilibrium, and the Bloch Equation without relaxation

The math described by Bloch deserves a summary here, as it will be useful in understanding the chapters that follow. In this description, the simplification will be used that the sample is a collection of non-interacting spins that can be approximated as a single bulk magnetization, *M*. Consider the basic, modern NMR experiment. As described previously, the setup for the modern NMR experiment involves placing a sample into a large static magnetic field, B_0, which by convention we take to be oriented in the *z*-direction. Upon doing so the sample (or, more accurately, the nuclei in the sample) will be magnetized proportionally to the strength of the external magnetic field in which it is placed, according to Equation 1.8:

$$\boldsymbol{M} = \chi\boldsymbol{H} \tag{1.8}$$

where χ represents the magnetic susceptibility of the sample – the constant of proportionality that describes how much an object is magnetized when placed in the external magnetic field, *H*. This manuscript has been using *B* and not *H* when describing magnetic fields, so it is helpful to translate this into something more recognizable using the definition of the *H* field in Equation 1.9.

$$\boldsymbol{H} = \frac{\boldsymbol{B}}{\mu_0} - \boldsymbol{M} \tag{1.9}$$

where μ_0 is the vacuum permeability and all other variables are as previously defined. This equation applies generally to fields in regions of all ranges of magnetizability, but for this example, the *external* field in which a sample is placed is what is of interest. This allows $\boldsymbol{M}$ in Equation 1.9 to be ignored since the vacuum is not magnetized. Further, the external magnetic field has already been defined to have only one component, B_0 in the z-direction. Therefore, combining these assumptions with Equations 1.8 and 1.9, a new and more useful equation can be derived.

$$M_0 = \chi \frac{B_0}{\mu_0} \tag{1.10}$$

Here, the initial equilibrium magnetization of the sample, M_0, has been defined, and is typically assumed, to be in the same direction as the external magnetizing field, B_0.*

Having defined B_0 and M_0, the initial conditions of the modern, basic NMR experiment have been set. That is to say, there is a sample that has been magnetized to the initial equilibrium bulk magnetization, M_0, by an external field, B_0. At this point, the famous Bloch Equations may be introduced. Without the effects of magnetic relaxation to equilibrium (which will be discussed in more detail later), the change in magnetization of a sample in a magnetic field is described (without relaxation) by this equation.

---

* Here again is one of those times where we suffer from the poor foresight of our scientific forebearers. Strictly speaking, $\boldsymbol{B}$ is "magnetic flux density" and $\boldsymbol{H}$ is "magnetic field." However, it is common practice that both $\boldsymbol{B}$ and $\boldsymbol{H}$ are referred to as "the magnetic field," even though, by their defined relationship, they are clearly not the same thing. One would think that someone with some authority on the matter could have settled the issue by giving the fields some nice descriptive and *different* names. I'm looking at you, Maxwell.

$$\frac{d\boldsymbol{M}(t)}{dt} = \gamma\big(\boldsymbol{M}(t) \times \boldsymbol{B}(t)\big) \tag{1.11}$$

where, as before, γ is the gyromagnetic ratio.

From the initial conditions described above, this is not a very interesting case. Initially, M_0 and B_0 are parallel and the cross product of two parallel vectors is zero, so there would be no change. If the initial conditions are changed, however, some important features are revealed. First, consider if the bulk magnetization is not set initially to be parallel to B_0, but instead let $\boldsymbol{M}(t{=}0)$ be tilted away from B_0. A full derivation of this is presented in appendix A.1, but to summarize, Equation 1.11 predicts that 1) there will be no change to the z-component of $\boldsymbol{M}$, 2) the effect of $\boldsymbol{B}$ on $\boldsymbol{M}$ will be to cause $\boldsymbol{M}$ to precess about $\boldsymbol{B}$, and 3) the frequency of that precession will be the characteristic Larmor frequency, ω_0. However, the more interesting dynamics occur when a perturbing field is introduced. The perturbing field referred to here, is a pulse of rf energy as described above in the description of the free induction decay experiment and in figures 1.6 and 1.7. A more detailed description will now be presented here.

Thus far, descriptions of the magnetic moment have been from the universal reference frame, or the "lab frame." In order to more easily comprehend the effects of an rf pulse, it is beneficial to temporarily change reference frames to one that rotates at the frequency of the pulse. This is typically on resonance with, or very near resonance with the Larmor frequency, since this is the frequency required for a spin to absorb rf energy. Unlike the simple Bloch Equations above, there are many examples that can be found for how to transfer to the rotating frame and therefore an appendix has not been included in this manuscript. The important *assumptions* in this approximation are that 1) the frame rotates at very nearly the same rate as the

rf pulse and observe frequency, and that 2) the pulse of rf magnetic field is decomposed into a

positively rotating vector (the one that is being observed) and a counter-rotating vector that is so

far off resonance that it can be ignored without losing much predictive accuracy. The important

results from this reference frame change are that 1) the effective static field is equal to

$$B_{eff} = B_0 - \frac{\omega_{obs}}{\gamma} \tag{1.12}$$

so that if the rotation of observation is on resonance, $\omega_{obs} = \omega_0$, and $B_{eff} = 0$, and 2) the rf field

appears static in the new reference frame so that now the Bloch Equation may be used as before

in the rotating frame. The Bloch Equation in the rotating frame becomes

$$\frac{d\widetilde{\boldsymbol{M}}}{dt} = -\gamma\left(\widetilde{\boldsymbol{M}} \times \widetilde{\boldsymbol{B}}_{tot}\right) \tag{1.13}$$

where the tilde over the vectors indicates that this is "as seen from the rotating frame." The total

magnetic field is the vector sum of the static field and the rf pulse field. Since the effective static

field is aligned along the axis of rotation, it does not change in the transfer to the rotating frame.

The rf portion of the field in its decomposed lab frame form would appear as two counter-

rotating field vectors, which rotate at the Larmor frequency:

$$\boldsymbol{B}_1 = \frac{B_1}{2}(\cos(\omega_0 t + \phi)\hat{x} + \sin(\omega_0 t + \phi)\hat{y}) + \frac{B_1}{2}(\cos(\omega_0 t + \phi)\hat{x} - \sin(\omega_0 t + \phi)\hat{y})$$

$$\tag{1.14}$$

where ϕ is an arbitrary phase correction. For now, ϕ will be left as zero. As noted before, only the on-resonance half of the rf field will be used. Furthermore, since the rotating frame is rotating at the Larmor frequency, this on-resonance vector appears static. The total magnetic field is then

$$\widetilde{\boldsymbol{B}}_{tot} = \frac{B_1}{2}\hat{x}' + \left(B_0 - \frac{\omega_{obs}}{\gamma}\right)\hat{z} \qquad (1.15)$$

where the prime after the unit vector for the x-direction denotes that this is a rotating axis that is static in the rotating frame. The z-direction unit vector does not need a prime, since it is unchanged upon transfer into the rotating frame.

What is important to note here is that when the pulse/observe frequency matches the Larmor frequency, B_{eff} becomes zero and the Bloch Equation becomes

$$\frac{d\widetilde{\boldsymbol{M}}}{dt} = -\gamma\left(\widetilde{\boldsymbol{M}} \times \widetilde{\boldsymbol{B}}_1\right) \qquad (1.16)$$

where again, the tilde over $\boldsymbol{B}_1$ indicates that it is as viewed from the perspective of the rotating frame, that is to say, it is static. At this point, for clarity, the tildes can be dropped so long as it is clear that all descriptions and variables are in the rotating frame. Here, the denominator of 2 has also been neglected for simplicity. It is understood from here forward that B_1 is actually only half of the applied rf field magnitude.

Solving for the dynamics of this form of the Bloch Equation gives essentially the same results as before in the lab frame, but now the rotation is about B_1 and the frequency of rotation is

$$\omega_1 = \gamma B_1 \tag{1.17}$$

where B_1 is being used simply as a scalar, with the understanding that the magnetization rotates

around B_1. This is new frequency, ω_1, is sometimes called the *nutation frequency*. It can be used

to predict the result of a pulse on a magnetized sample.

1.5.2 Pulses and rotations using the Bloch Equation in a rotating frame

Earlier, the "arbitrary" phase correction, ϕ, was ignored for the simplicity of the

derivation. It turns out, however, that this phase correction sets the angle about which the

nutation rotates. From the decomposed vector picture in Figure 1.8, one can see that as ϕ is

changed, the angle around the transverse plane at which the decomposed, rotating vectors start

from and coincide at will increase at the same rate. Thus, a choice of $\phi = \pi/3$ will cause B_1 to

effectively oscillate at 60 degree angle from zero. According to the equations above, then,

subjecting a sample magnetization to a pulse with such a phase offset would cause it to rotate

about B_1, which itself is tilted by 60 degrees. By convention, $\phi = 0$ refers to B_1 oscillating along

the x-axis. Accordingly a pulse with a phase of zero (at the Larmor frequency) is taken to be an

x-pulse with any sample magnetization subjected to that pulse rotating about the x-axis as a

result. By the same logic, a pulse with a phase offset of $\phi = \pi/2$ is called a y-pulse and it causes a

sample's bulk nuclear magnetization to rotate about the y-axis.

Figure 1.8. An oscillatory B_1 field can be decomposed (and is equivalent) to a pair of counter-rotating vectors, each with half the magnitude of the original B_1 (a). In panel (b), it is shown that the pairs of counter-rotating vectors can be summed up to an equivalent B_1 field described by Equation 1.14. If the phase in Equation 14 is set to $\phi = \pi/3$, all vectors are rotated by 60 degrees (c), and the resulting summed B_1 is also rotated by 60 degrees (d).

In general these principles can be used to develop a set of rules-of-thumb for the two most important pulses, x- and y-pulses, for determining their effect on a bulk nuclear magnetization. This is easiest to do in a matrix format where a bulk magnetization vector is described by a 3-by-1 column vector. In this form, and in the rotating frame, the rotation matrices for x- and y- pulses, are the normal rotation matrices from linear algebra:

$$\mathbb{R}_x(\theta) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\theta & -\sin\theta \\ 0 & \sin\theta & \cos\theta \end{bmatrix} \qquad (1.18.a)$$

$$\mathbb{R}_y(\theta) = \begin{bmatrix} \cos\theta & 0 & \sin\theta \\ 0 & 1 & 0 \\ -\sin\theta & 0 & \cos\theta \end{bmatrix} \qquad (1.18.b)$$

where the double-stack notation has been used to indicate a matrix. These matrices are typically unnecessary, however. For the purpose of understanding pulse sequences, it is usually sufficient to simply speak of the degree and direction of rotation. For example, one might speak of a $\pi/2$ x-pulse to indicate an x-pulse that rotates the magnetization by 90 degrees about the x-axis. Common notation for this is to indicate the angle of rotation in parentheses, followed by a subscript of the axis about which the magnetization vector is being rotated. A $\pi/2$ x-pulse, for example, would be denoted as $(\pi/2)_x$.

This pulse notation is commonly incorporated into depictions of pulse sequences for the purpose of conveying the parameters of an experiment. It is thus prudent to describe this notation here. Pulse sequences are typically depicted as time plots, with pulses themselves as boxes overlayed onto a time axis. Their length corresponds to pulse length in time and their height corresponds to pulse power. Since pulse power determines the strength of B_1 generated, and Equation 1.17 shows that nutation frequency is directly proportional to B_1, this means that the area of the pulse depicted in a pulse sequence diagram – nutation frequency times time – relates directly to the angle swept out by the magnetization vector as a result of the pulse. Figure 9 shows the simplest pulse sequence, the free induction decay.

Figure 1.9. The free induction decay consists of an rf pulse, shown here as the black rectangle resting atop the horizontal time axis. The pulse power, and thus the pulse nutation frequency, is depicted by the height of the pulse, while the pulse length is depicted by the length of the pulse. Above the pulse, the angle swept out by the nutation is indicated in parentheses, with a subscript for the axis about which the nutation rotates, which itself indicates the phase. To the right of the pulse, the resulting signal is depicted by a decaying sine wave.

It is important to note here that, were it not for dephasing and relaxation, the result of the FID experiment would be an infinite sine wave. In reality, samples contain spins in slightly different local fields, both because different electronic environments within molecules shield the different nuclei in a non-uniform manner and because all NMR magnets generate B_0 fields that are a little bit inhomogeneous. These effects cause the spins in a population to precess at slightly different frequency, causing them to eventually fall out of sync. This is dephasing. Additionally, all spins exchange energy with each other and their environment, eventually causing the magnetization of the population to relax back to equilibrium. The result of these two effects is that the overall amplitude of the signal decays away with time, hence the name "free induction *decay*."

1.5.3 NMR with relaxation

The relaxation of a perturbed magnetization back to equilibrium is characterized by two flavors of relaxation, referred to by their time constants: 1) relaxation of the magnetization component parallel to the static field, B_0, as this component returns to its full equilibrium amplitude has the time constant T_1, and 2) the relaxation of components perpendicular to the static field as they revert back to zero is characterized by the time constant T_2. Respectively, T_1 and T_2 are also often called *longitudinal* and *transverse*, or *spin-lattice* and *spin-spin* relaxation.

These are incorporated into the Bloch Equations in the following way.

$$\frac{d\boldsymbol{M}(t)}{dt} = \gamma\big(\boldsymbol{M}(t) \times \boldsymbol{B}(t)\big) + \begin{bmatrix} -\dfrac{M_x}{T_2} \\[2ex] -\dfrac{M_y}{T_2} \\[2ex] \dfrac{(M_0 - M_z)}{T_1} \end{bmatrix} \tag{1.19}$$

where M_x, M_y, and M_z are simply the instantaneous magnitudes of the sample magnetization in the x-, y-, and z-directions, respectively, and M_0 is the equilibrium magnetization along the direction of the external static field, B_0. Equation 1.19 predicts that the components will relax exponentially and this comports with observation quite well, although the truth is that this is a phenomenological description and is not derived from first principles. This is not to say that the origins of relaxation are completely unknown. Some of the origins are known, especially at the macroscopic level, and they will be discussed briefly later in this introduction.

In reality, relaxation of a sample typically has multiple components, due to changing proximities of spins to each other, molecular tumbling, and other effects. The overall observed relaxation time constant is the weighted average of all of these components within a sample. For a single-component sample, the result of Equation 1.19 is that the magnetization decays according to a single exponential component. The general solution to Equation 1.19 is not one that can usually be found analytically. There are, however, known solutions for particular cases, such as the FID, inversion recovery and saturation recovery. The accepted solution for an FID as described above, and in the rotating frame, is

$$\widetilde{M_x} = 0\,\tilde{\hat{x}} \qquad\qquad (1.20.a)$$

$$\widetilde{M_y} = -M_0 e^{-\frac{t}{T_2}}\,\tilde{\hat{y}} \qquad\qquad (1.20.b)$$

$$M_z = M_0\left(1 - e^{-\frac{t}{T_2}}\right)\hat{z} \qquad\qquad (1.20.b)$$

where the tilde has been dropped from the z-component because the rotating z-component and the lab-frame z-component are identical. In this set of equations, M_x is zero because it is assumed that a $(\pi/2)_x$ pulse was used to rotate the initial magnetization along the y-axis. The same logic can be used to arrive at a similar set of equations for an FID with a $(\pi/2)_y$ pulse where M_y is zero.

Figure 1.10. The principle of relaxation is demonstrated through the example of a FID experiment, shown here in the rotating frame. The plot above shows a pulse sequence on a time axis, which is aligned to depictions, below, of the sample magnetic moment vector (black arrow), along with its longitudinal and transverse components (grey arrows). The sample starts out at equilibrium (*eqbm*), magnetized along the external B_0 field, in the z-direction. A $(\pi/2)_x$ pulse rotates the magnetization into the xy-plane, to lie in the negative y-direction ($t = 0$). As the bulk magnetic moment precesses, the transverse components shorten and the longitudinal components grow ($t = 1, 2, 3, \& 4\ T_1$). After 5 T_1 periods, the sample moment has nearly fully recovered its equilibrium state ($t = 5\ T_1$).

A visual depiction of relaxation in this experiment is shown in Figure 1.10. As can be seen, the $(\pi/2)_x$ pulse rotates the bulk magnetization around the x-axis to lie in the negative y-direction, after which, the magnetization relaxes back to its equilibrium value in both the longitudinal and transverse directions. According to Equation 1.20.b, after five T_1 periods, the longitudinal magnetization has recovered to 99.4 % of its equilibrium value and can be considered essentially recovered. The transverse magnetization is always faster than the longitudinal relaxation, and so can essentially be considered zero after five T_1 periods.

1.6 Sources of relaxation

It is critical, especially in the context of this manuscript, to consider the origins of relaxation, as this knowledge is what makes relaxometry useful. In a somewhat perverse irony, it turns out that longitudinal relaxation occurs because of fluctuations in the transverse magnetic fields of spins' local environments, whereas transverse relaxation arises from fluctuations in the longitudinal field of the spins' local environments. This can be understood by analogy to the requirement in the Rabi experiment that the oscillatory field be perpendicular to the static field in order to perturb the spin out of its state. In a similar way, longitudinal magnetization is perturbed by fluctuating fields that are perpendicular to itself (transverse), and vice versa for transverse magnetization perturbed by longitudinal fluctuations.

The next few sections will briefly detail some important contributions to relaxation rates, particularly in liquids, which is the focus of research in this manuscript. The contributions of these mechanisms add together to create the overall observed relaxation rate $1/T_{obs}$ as in Equation 1.21 below.

$$\frac{1}{T_{obs}} = \sum_i \frac{1}{T_i} \tag{1.21}$$

where the subscript, i, denotes the i^{th} contribution to relaxation.

1.6.1 Molecular motions

The first source of relaxation that will be discussed here is that of the interaction between molecular motions and the local magnetic field itself. In any sample in which the constituent spins are mobile, they will experience small fluctuations in their local magnetic fields as they pass and rotate in the vicinity of other molecules or microscopic magnetic environments. This motion is captured by a parameter called the correlation time, denoted τ_c, and is associated with a sample liquid's viscosity, η, effective volume of the molecules, V, temperature, T, and the Boltzmann constant, k. The equation that describes this is the Stokes-Einstein-Debye Equation, given here [14].

$$\tau_c = \frac{\eta V}{kT} \tag{1.22}$$

The correlation time is often defined as the time that it takes for a molecule to tumble one radian, but as stated earlier, the motions are much more complicated than simply rotating, so this definition should be taken more notionally than as a strict definition. What is more important than this strict definition is the implication that a shorter correlation time is, well, *correlated* with greater molecular motion. As Equation 1.22 shows, a liquid with greater viscosity is one in which the molecules take longer to tumble. In such a case, one can imagine that the individual

molecules would also be perturbing each other's magnetic environments less frequently. In fact, there is an equation that describes the evolution of these field fluctuations in time called the *autocorrelation function* or sometimes, the *rotational correlation function*.

$$G(\tau) = \frac{\langle B(t)B(t + \tau)\rangle}{|B^2|} = e^{-\frac{|\tau|}{\tau_c}} \qquad (1.23)$$

This equation states that the average product of the field experienced by a given molecule at time, t, and its field experienced a time, τ, after t decays at a rate defined by the correlation time – the rate at which the molecule tumbles. Essentially, the faster a molecule tumbles, the less likely it is to be oriented with where it was at some time, τ, earlier. The decaying exponential nature of this probability makes sense intrinsically, although strictly speaking it is only assumed. The correlation function may be Fourier transformed to reveal its probability of tumbling at a given angular frequency, ω_c. This is called the *spectral density function*, $J(\omega_c)$, and is typically normalized by the mean square field [15].

$$J(\omega_c) = \frac{2\tau_c}{1 + \omega_c^2\tau_c^2} \qquad (1.24)$$

From Equation 1.24, it can be seen that the probability of a molecule tumbling at a certain frequency diminishes as that frequency increases, as well as with longer correlation times. This simply means that samples in which molecules tumble more slowly on average are less likely to contain higher frequencies of rotation, which should make sense. Recall that longitudinal relaxation occurs as a result of transverse fluctuations in the magnetic field. Thus, the

longitudinal relaxation rate due to motion, as a function of correlation time, can be found by

setting the rate, $1/T_1$, proportional to both the mean squared field and the probability of

fluctuations occurring at a given frequency needed for transitions (the Larmor frequency).

Multiplying the transverse magnetization's mean squared value and the spectral density at the

Larmor frequency, this gives an expression for T_1 (or, rather, its inverse) [16].

$$\frac{1}{T_1} \propto \gamma^2 \langle B_x{}^2 + B_y{}^2 \rangle \frac{\tau_c}{1 + \omega_0{}^2 \tau_c{}^2} \tag{1.25}$$

The mean squared field is used rather than simply using the field itself because the

average of a fluctuating field is, by definition, zero. Similar logic to that above leads to a relation

for the transverse motional relaxation rate, although it will not be explored in detail here [16].

$$\frac{1}{T_2} \propto \gamma^2 \langle B_z{}^2 \rangle \tau_c + \frac{1}{2} \gamma^2 \langle B_x{}^2 + B_y{}^2 \rangle \frac{\tau_c}{1 + \omega_0{}^2 \tau_c{}^2} \tag{1.26}$$

Since motional relaxation depends primarily on the correlation time, anything that affects

the rate at which molecules move and tumble – temperature, molecular size, and viscosity – will

in turn affect the longitudinal and transverse relaxation rates. Nuclear magnetic resonance

relaxometry can thus be used to elucidate changes in these macroscopic parameters in a sample.

1.6.2 Dipolar coupling

Magnetic nuclei that are in close proximity are affected by each other's fields and this

kind of direct interaction between nuclei is called *dipolar coupling*. The source of this type of

relaxation is from the interaction between two nuclei. The field that a nucleus produces at some point a distance, r, away from itself and at some angle, ψ, to its orientation can be written in terms of the field in polar coordinates, B_r and B_ψ.

$$B_r = \frac{\mu_0}{4\pi} \frac{2\mu}{r^3} \cos(\psi) \tag{1.27}$$

$$B_\psi = \frac{\mu_0}{4\pi} \frac{2\mu}{r^3} \cos(\psi) \tag{1.28}$$

where μ_0 refers to the vacuum permeability and μ is the magnetic moment of the field-producing nucleus [14]. From these equations, it can be seen that the field decays in a cubic nature as the distance between nuclei increases, so nuclei must be very close in order to affect each other in this way. Because of this fact, this mechanism is considered to be important mostly for neighboring nuclei within the same molecule.

Since the spins within a molecule are not bound to change orientation with the molecule, the correlation time of the molecule again comes into play here. The dipolar relaxation rate is proportional to the product of the fields of the two nuclei. Since the field strengths decay as the third power of the distance between the two nuclei, this means that the relaxation rate due to this interaction falls off as the sixth power of the internuclear distance. The longitudinal dipolar relaxation rate is given by the equation

$$\frac{1}{T_1} = \frac{\left(\frac{\mu_0}{4\pi}\right)^2 \gamma_A^2 \gamma_X^2 \hbar \tau_c}{r^6} \tag{1.29}$$

This equation has a fortunate feature in that all variables are constants, with the exception of the correlation time. Thus the relaxation due to dipole coupling is directly proportional to correlation time. Moreover, the contribution to relaxation rate from this mechanism is greater with longer correlation times and vice versa. For liquids with fast tumbling molecules, this may not be much of a factor at all.

1.6.3 Paramagnetic interaction

When a solution contains unpaired electrons, those electrons' spins interact with non-zero spin containing nuclei and, similar to the dipolar effect, this interaction forms another route for relaxation. Unlike nuclear dipole coupling, however, the gyromagnetic ratios are no longer on the same order of magnitude. Instead, the gyromagnetic ratio of an electron is 28,025 MHz/T [16], compared to the gyromagnetic ratio of a proton – 42.6 MHz/T [17]. Thus the presence of paramagnetic species in a sample can have a dramatic effect on the relaxation rate. This manuscript does not explore paramagnetic interactions in particular, however, the properties described here make paramagnetic species extremely useful when increasing the relaxation rate of a sample allows for a more expedient experiment, as in the doping of water with copper sulfate.

1.7 NMR spectra

One of the most important applications of NMR is the determination of molecular structure by producing NMR spectra from the raw FID signal. While this manuscript does not focus on structural determination per se, the theoretical underpinnings of spectral features are

useful for understanding strengths and limitations of low-field NMR. Consider, then, a simple

FID experiment. Upon perturbation with a $(90^\circ)_x$ pulse (or other similar pulses that put the

sample magnetization into the transverse plane), a sample with a collection of spins at various

local fields will produce a complex signal as it rings down. This signal contains a collection of

different frequencies for each local magnetic environment. If relaxation were not present, the real

signal would simply be proportional to the sum of all of these frequencies, weighted by the

population of spins at those frequencies.

$$S(t) \propto \sum_n A_n \cos \omega_n t \tag{1.30}$$

where A_n and ω_n refer to the population and the Larmor frequency of spins in each particular

local magnetic environment, respectively.

Consider, for example, the signal from methanol. Methanol contains hydrogens in two

magnetic environments: 1) the three hydrogens attached to the carbon, and 2) the single

hydrogen attached to the oxygen. The proton signal from a sample of methanol would thus

contain two frequencies: one for each local magnetic environment. The oxygen, having much

greater electronegativity, tends to pull electron density away from the electronegatively weak

hydrogen bonded to it. This "de-shields" the hydrogen from the static field, causing it to "feel"

the static magnetic field, B_0, more. Consequently, its Larmor frequency will be higher. This

frequency shift is called *chemical shift* and is critical in the power of NMR to reveal molecular

structure. Additionally, the amplitude associated with the three hydrogens attached to the carbon

would be three times greater than that of the single hydrogen on the oxygen. At a B_0 field

corresponding to a proton Larmor frequency of 300 MHz, this signal would appear as a sine

wave with the amplitude expanding and contracting every 10 milliseconds or so as the

frequencies alternately interfere with each other constructively and destructively.

Figure 1.11. The theoretical NMR signal of methanol, without relaxation or *J*-coupling, is a sinewave with an amplitude that expands and contracts. The "beats" of the signal arise from the periodic constructive and destructive interference of the component frequencies that make up the signal. This plot was constructed using Matlab™.

To see the spectrum of frequencies that make up this signal, a Fourier transform is

applied to the recorded signal. This is done computationally, using a fast Fourier transform (FFT)

of the acquired signal, and is not attempted analytically. For the 300 MHz methanol example

above, and using Equation 1.30 for the signal, this results in two infinitely narrow peaks

separated by approximately 6 Hz, and with the lower frequency peak three times taller than the

higher.

In reality, however, relaxation is always present and this greatly affects the signal. Instead

of a sinewave that carries on forever, relaxation causes the received signal to damp with time.

This makes the Fourier transform much more difficult to perform analytically, so it will not be

derived here. Instead the solution is presented and the reader is encouraged to follow the citations to the derivation if they are so inclined. For a spin *with* relaxation, the signal is

$$S(t) \propto Ae^{-\frac{t}{T_2^*}} \cos \omega_0 t \tag{1.31}$$

where the asterisk is used for relaxation, T_2^*, to indicate that this is the *observed* relaxation of the signal due to dephasing as well as the true T_2 [1]. Otherwise, all variables are as described previously. For a population of spins in n magnetic environments,

$$S(t) \propto \sum_n A_n e^{-\frac{t}{T_2^*}} \cos \omega_n t \tag{1.32}$$

where again all variables are as described previously. As before, a Fourier transform is performed in order to reveal the spectrum. Looking again at the single spin example from Equation 1.31, the Fourier transform of this spin component can be found in tables of Fourier transforms and the real part is found to be [1]

$$S(\omega) = \frac{\left(\frac{1}{T_2^*}\right)}{\left(\frac{1}{T_2^*}\right)^2 + (\omega - \omega_0)^2} = \frac{T_2^*}{1 + (\omega - \omega_0)^2 T_2^{*2}} \tag{1.33}$$

Figure 1.12 shows a plot of this, and it can be seen that the peaks are no longer infinitely narrow, as was the case with no relaxation. The effect of relaxation on the spectrum obtained from a spin in a sample is that the infinitely narrow line describing a spin frequency in a

theoretical spectrum without relaxation becomes a peak with a width defined by $T_2{}^*$. An important takeaway from this plot is that the width of the peak is not relative to the Larmor frequency as was the case for chemical shift. To resolve most spectra, spectral resolution must be finer than 1 ppm, which per the definition of low-field at the beginning of this chapter, is less than 4.4 Hz. At the same low fields, the peak width for a sample with 100 ms T_2 (a common value of T_2) would be 20 Hz at the very best, which equates to a little more than 4.5 ppm. This illustrates why low-field NMR is not used for structural spectroscopy.

Figure 1.12. The Fourier transform of the signal obtained from a single spin after a $(\pi/2)_x$ pulse, while accounting for transverse relaxation. This theoretical plot was calculated using a $T_2{}^*$ of 100 ms and the center of the abscissa represents the on-resonance Larmor frequency. The plot reveals a peak width (sometimes called linewidth) that, at half height is equal to $2/T_2{}^*$. This plot was constructed using Matlab ™.

Longitudinal relaxation also affects the spectrum in that it determines the time required to repeat experiments thereby increasing the signal to noise ratio (S/N) of the resultant spectrum. This can be seen in the peak height above the noise floor that can be obtained for a given experiment time. If a sample has a fast T_1, its magnetization will fully recover more quickly, and the experiment may be repeated and averaged more times within the overall experiment time,

thus suppressing noise and increasing the effective peak height. This becomes very important in relaxometry, in which S/N is the more important goal, and spectral resolution is not.

1.8 Components of signal

A major focus of this manuscript is the acquisition of signal in challenging and noisy environments. Since single-sided NMR is typically in the low-field regime, any hope of attaining spectra with peak separation is lost. This was explained previously when describing the effects of relaxation on NMR spectra. Because of this, low-field NMR is concerned primarily with the ability to acquire signal *amplitude* – achieving a high S/N.

1.8.1 Signal in an ideal NMR setup

Consider the typical NMR setup – a liquid sample placed such that it fills the volume of a cylindrical solenoid rf coil, and all of this is then placed inside a large, homogeneous, static magnetic field, with the axes of the permanent and rf fields perpendicular to each other, as was described back in Figure 1.6. Of course, this setup is far from the reality of low-field, single-sided NMR, but the scenario provides a simple starting point for understanding the factors that go into producing NMR signal. In such a case, it can be assumed that the sample resides in a uniform field and is therefore magnetized uniformly. Because the sample resides entirely inside the rf coil, it can also be assumed that the pulse sequence power is applied evenly across the sample.

The signal acquired in NMR is a result of the precession of the sample's bulk magnetic moment after a pulse sequence has disturbed the moment from its equilibrium. As the magnetic moment rotates, it alternately passes between parallel and perpendicular to the long axis of the rf

solenoid. This creates a sinusoidal pattern of the sample's magnetic flux with respect to the solenoid, which in turn generates a current in the solenoid according to Faraday's Law of Induction from Equation 1.7. The current generated in the solenoid is then detected by the spectrometer for recording. The signal is therefore a combination of these factors:

- The strength of the sample's magnetic moment as determined by the strength of the magnetizing static field and the magnetic susceptibility of the sample
- The rf solenoid volume, and therefore the volume of the sample
- The sensitivity of the spectrometer

The strength of the sample's nuclear magnetic moment is a result of the number of active nuclear spins in the sample volume, as well as the Boltzmann distribution of spin states for the temperature at which the experiment is being performed. A stronger distribution leads to a greater difference between spin-up and spin-down nuclei, which in turn results in greater magnetization. Moreover, the flux due to the whole sample is, roughly speaking, the sum of the flux generated by each spin for which there is not another counter-oriented spin to cancel it out. Thus if the spin density is higher within the volume, there will be more un-canceled spins for a given Boltzmann distribution of spin states.

The Boltzmann distribution of states is determined by examining the ratio of the populations of two states whose energy difference, ΔE, is known.

$$\frac{P_\alpha}{P_\beta} = \frac{e^{-\frac{E_\alpha}{kT}}}{e^{-\frac{E_\beta}{kT}}} = e^{\frac{\Delta E}{kT}} \qquad (1.34)$$

where P_α and P_β are the populations of spins in the up and down state, respectively. E_α and E_β refer to the energies of the α and β states, respectively. The terms k and T refer to the Boltzmann constant and temperature, respectively. The energy difference, ΔE, is calculated from Equation 1.3 (in the case of a spin-½ nucleus).

From this, the spin *difference* can be calculated as a fraction of the whole.

$$\Delta P = \frac{P_\beta}{P_{tot}} - \frac{P_\alpha}{P_{tot}} = \frac{e^{-\frac{E_\beta}{kT}} - e^{-\frac{E_\alpha}{kT}}}{e^{-\frac{E_\alpha}{kT}} + e^{-\frac{E_\beta}{kT}}} \tag{1.35}$$

Conveniently, each of these terms in the exponentials are very small, so the denominator turns out to be very close to 2. Using similar assumptions, the terms in the numerator can be simplified by series approximations. Namely the well-known exponential expansion

$$e^x = 1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \cdots = \sum_{n=0}^{\infty} \frac{x^n}{n!} \tag{1.36}$$

Only the first two terms are needed since the terms in the exponents are so small. For spins-½, Equation 1.3 gives that E_α and E_β are equal to $\frac{\gamma \hbar B_0}{2}$ and $-\frac{\gamma \hbar B_0}{2}$, respectively, So with the above series expansion, Equation 1.35 can be written in a much more palatable form [1].

$$\Delta P = \frac{e^{-\frac{E_\beta}{kT}} - e^{-\frac{E_\alpha}{kT}}}{e^{-\frac{E_\alpha}{kT}} + e^{-\frac{E_\beta}{kT}}} \cong \frac{\left(1 + \frac{\gamma \hbar B_0}{2kT}\right) - \left(1 - \frac{\gamma \hbar B_0}{2kT}\right)}{2} = \frac{\gamma \hbar B_0}{2kT} \tag{1.37}$$

As this shows, the spin difference (as a fraction of the whole number of spins in the sample) depends linearly on the strength of the static field. If the spin density, ρ, is known in terms of spins per cubic centimeter, then the excess number of spins in a sample can be known:

$$N_s = \rho V \frac{\gamma \hbar B_0}{2kT} \tag{1.38}$$

where N_s refers to the number of excess spins – those for which there is not another spin of opposite state to cancel its magnetic moment. The term V refers to the volume of the sample, as usual.

Finally, the bulk magnetic moment is the sum of all excess magnetic moments in the sample. Using Equation 1.2, and understanding that the excess spins are those in the "up" state, this means that the bulk sample magnetization is given, to a very good approximation, by

$$M_0 = \mu N_s = \rho V \frac{\gamma^2 \hbar^2 B_0}{4kT} \tag{1.39}$$

where, as before, the assumption has been made that the magnetization vector is parallel to the coordinate defined by B_0, and so the term M_0 has been used to indicate that this is a magnitude, rather than using the more general bold vector notation. Note also that this shows that the magnetization of a sample depends on the *square* of the gyromagnetic ratio of the spins being investigated. The choice of a spin with a large gyromagnetic ratio is therefore very useful in pursuing more signal.

It must be remembered that signal is not just magnetization. It is what is detected by the spectrometer and can be recorded. In the typical method of pulse NMR used today, this magnetization is used to induce a current in the rf coil surrounding it by rotating the bulk magnetization into the transverse plane and allowing it to precess freely as was shown in Figure 1.7. This free precession produces an alternating flux through the turns of the rf coil, which then generates a current by Faraday's Law, as described by Equation 1.7. Assuming the best possible case in which the sample volume exactly matches the rf coil, then the entire magnetic field generated by the magnetized sample is captured by the cross-sectional area of the coils and the flux generated is equal to the bulk magnetization times this area, A.

$$|\Phi_{rf}| = M_0 A \tag{1.40}$$

As has been described previously from the Bloch equations, once the bulk magnetization has been rotated perpendicular to the static field, B_0, it will begin to precess at the Larmor frequency for that gyromagnetic ratio and field strength. The flux is therefore also time dependent:

$$\Phi_{rf}(t) = AM_0 \cos(\omega_0 t) \tag{1.41}$$

The electromotive force (voltage) induced in the rf coil can then be calculated from the Faraday Law of Induction (Equation 1.7) as the negative derivative of the flux across all turns of the rf coil.

$$\varepsilon(t) = NAM_0\omega_0 \sin(\omega_0 t) \tag{1.42}$$

The induced current in the coil circuit, calculated simply by Ohm's Law, is then simply

$$I(t) = \frac{\varepsilon}{Z_c} = \frac{NAM_0\omega_0 \sin(\omega_0 t)}{Z_c} \tag{1.43}$$

where Z_c is the impedance of the rf coil circuit, which is typically 50 Ohms and must be matched to the rest of the system as such. Substituting M_0 as in Equation 1.39, along with $\omega_0 = \gamma B_0$, reveals the following formula [10]:

$$I(t) = \frac{\varepsilon}{Z_c} = \frac{NA\rho V}{Z_c} \frac{\gamma^3 \hbar^2 B_0^2}{4kT} \sin(\omega_0 t) \tag{1.44}$$

At this point it is prudent to examine the factors, which can be controlled by the design of the experiment, that might be harnessed to increase signal. If one desires to look at a particular sample, then the nucleus is generally not an option for change. Besides, experiments are typically conducted on protons, which already have the highest nuclear gyromagnetic ratios, so that possibility is already tapped out. The other factor that immediately comes to mind is to increase the static field strength, as the signal magnitude depends on its square. Besides this, the most effective ways to increase signal in the ideal NMR experiment are to increase the sample and rf coil volume (and thus the loop cross-sectional area), or increase the number of rf coil turns.

1.8.2 Signal considerations for single-sided NMR

Single-sided NMR is of course decidedly *not* ideal and the above theoretical arguments about signal must be modified in order for the single-sided regime to be better understood. In single sided NMR, some variation of two common geometries is typically used, depicted in Figure 13. In one configuration, a permanent magnetic field, B_0, is generated parallel to the surface of the magnet and the rf coil may either generate a rf field perpendicular to the surface, as in a simple surface loop, or parallel, as in a "double-D" coil, but then perpendicular to B_0. This configuration is used by Bernard Blümich et. al. in the NMR Mouse and by Bruce Balcom, and his group at the University of New Brunswick in their three-magnet array [19] [20]. The other common geometry is to generate the B_0 field such that it is perpendicular to the magnet surface, in which case, the rf field must be generated parallel to the surface, usually with some kind of double-D coil. This geometry is used by Fukushima et. al. at ABQMR in their barrel magnet [21].

Figure 1.13. Single-sided NMR generally takes one of two geometries, depicted here. The first type, in (a) generate a B_0 field parallel to the magnet surface and transverse in the sample volume. In this case, the rf coil may be a surface loop, which generates a B_1 perpendicular to the surface as depicted, or a double-D coil, which would produce a B_1 field parallel to the surface, but perpendicular to B_0. The second common geometry is depicted on the right, in (b), in which the B_0 is generated perpendicular to the surface of the magnet, penetrating the sample volume directly, in which case the B_1 field must be produced parallel to the surface, usually with a double-D coil.

Unlike traditional NMR in a superconducting magnet, single-sided NMR B_0 fields tend to be very weak and inhomogeneous, so the trick employed by traditional NMR of reaching for higher and higher fields is generally not realistic. Therefore, the first challenge that is typically addressed when attempting single-sided NMR is to maximize the region of homogeneous, static field, B_0, produced by the permanent magnet. This increases the effective sample volume available, and as discussed above, this increases signal. Some single-sided magnets do generate gradients instead, as in some variations of the NMR-MOUSE and three-magnet arrays, but in the lateral dimensions, they still strive for maximum field homogeneity. To affect this goal, nearly all single-sided magnets incorporate some variation on the theme shown in Figure 1.14. This motif involves a split main exterior static field-generating magnet with a central, recessed magnet that serves to shim the sensitive region.
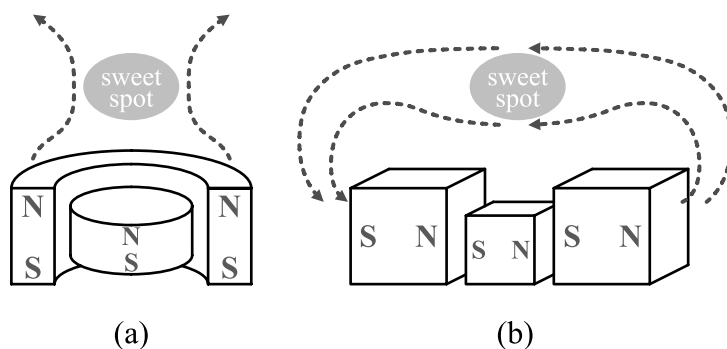
Figure 1.14. Two examples of single-sided magnets with a split main magnet and a central shimming magnet are shown. Dashed lines indicate flux lines of the B_0 field and the grey ellipse indicates roughly the location of the homogeneous "sweet spot" region. Whether the static field is perpendicular (a) or transverse (b) to the surface, the effect is the same with respect to field strength and gradient, although the direction of field is different. These diagrams are provided for conceptual purposes and elements are not drawn to scale.

In addition to concerns about the homogeneity of the static field, the shape and strength of the rf field must also be considered. As described above with Equation 1.16, the dynamics of the sample bulk magnetization depend on the cross product of the B_1 field and the sample magnetization. Thus, because the sample is also magnetized by – and parallel to – the B_0 field, the ideal, the rf field will be engineered such that it is perpendicular to the B_0 field and homogeneous in strength across the entire B_0 homogeneous region sweet spot. If the homogeneous region is small and close to the surface of the magnet, this is easier, but then signal is sacrificed from the loss in volume and usefulness is lost because of the reduced penetration depth. If the B_1 field produced by the probe coil is not homogeneous or is not perpendicular to B_0 across the B_0 region, each pulse will have the effect of dephasing the magnetization of the disparate regions of the active volume from each other and they can quickly begin to destructively interfere, again resulting in lost signal. Importantly, the overall magnitude of the B_1 strength is not nearly as important as the homogeneity of the B_1 field.

In the ideal case described above, the sample resided entirely within the rf coil, allowing the assumption to be made that all of the field generated by the sample's bulk magnetic moment would be captured in the flux of the receiving coil. In single-sided NMR, however, the sample cannot reside inside the rf coil if it is to be truly single-sided NMR. Moreover, the sample that can be observed is limited to the volume of the sensitive region of the static field, which is typically displaced from the magnet and coil assembly. The sample instead only "feels" a fraction of the flux generated by the rf coil and, by reciprocity, the coil is only sensitive to a fraction of the total magnetic moment of the sample. This fraction of the coil flux incident upon the sample, and conversely the fraction of the sample flux sensed by the cross-sectional area of the coil, is called filling factor. In a traditional NMR setup, maximizing the filling factor boils down to simply filling the rf coil with sample as much as possible. In single-sided NMR, the filling factor is dependent on B_1 field strength, shape, and distance to the sample.

The distance of the sample and sensitive region from the magnet and probe assembly also has a significant effect on signal for more reasons than just filling factor. The strength of the magnetic field due to a magnetic dipole moment falls off as the cube of distance. In polar coordinates, where r indicates the distance from the dipole to the observed point, and θ is the angle between the dipole vector and the vector from the dipole to the point of observation, the magnetic field is given by

$$\boldsymbol{B(r,\theta)} = \frac{M_0}{4\pi\mu_0 r^3}\left(2\cos(\theta)\,\hat{\boldsymbol{r}} + \sin(\theta)\hat{\boldsymbol{\theta}}\right) \qquad (1.45)$$

where μ_0 is the vacuum permeability and M_0 is the strength of the dipole moment. This has compounding effects for receiving signal in single-sided NMR. In addition to the problem of

reduced filling factor, which already diminishes flux received by the probe coil, the flux sensed

by the receive coil is further weakened because the strength of the field generated by the sample

dipole moment diminishes severely with distance.

1.9 Relaxometry

Because low-field, single-sided NMR cannot resolve chemical shift spectra, its focus

instead shifts primarily to relaxometry – the study of a sample's relaxation back to magnetic

equilibrium after being disturbed. Earlier sections have already mentioned some of the theory

that explains the causes of relaxation. This section will specifically describe two experiments

used to elucidate the two main forms of relaxation, T_1 and T_2.

1.9.1 Measuring T_1 – the inversion and saturation recovery experiments

The definition of T_1 is the time constant that defines the rate at which a sample regains its

longitudinal magnetization after being place in non-equilibrium state. Specifically, it is the time

that it takes for magnetization to recover an e^{-1} fraction of the equilibrium magnetization. As has

been described in the Bloch Equations section above, a sample magnetization in a static

magnetic field can be arbitrarily rotated using rf pulses. Thus, the simplest way to take a sample

out of longitudinal equilibrium (without concern for transverse) is to invert the magnetization

using a π-pulse. The *inversion recovery* experiment observes the recovery of a sample's *z*-axis,

or longitudinal, recovery by doing just this and is shown in Figure 1.15 [1].

Figure 1.15. The pulse sequence for an inversion recovery (a) consists of an inverting pulse, followed by a delay time that allows the longitudinal magnetization to recover, then a second pulse that rotates the magnetic moment into the transverse plane for observation. The pulse sequence is repeated n times, with a full recovery time between each repetition. The plot of the inversion recovery (b) is the series of amplitudes plotted against the recovery times, τ_n, for which they were acquired. The plot in (b) was constructed in Matlab™.

Because z-magnetization cannot be directly observed, however, the sample magnetization must be flipped into the transverse plane in order for the magnitude of the z-magnetization to be observed at that point in time. The sample is then given time to fully recover (over five times the T_1 time) and the experiment is repeated for a new observation time. The delay between the π-pulse and the $\pi/2$-pulse is called the tau time, τ. The inversion recovery experiment starts with a recovery delay time $\tau = 0$, and is repeated with τ increasing in discrete increments until the longitudinal magnetization has fully recovered within the τ time.

The resulting observed transverse magnitudes are then plotted against delay times as shown in Figure 1.16, and the points are fitted to a curve of the form

$$S(\tau) = S_0 \left(1 - 2\, e^{-\frac{\tau}{T_1}}\right) \tag{1.46}$$

where S_0 is the initial amplitude and its units are arbitrary (based on the sensitivity of the instrument, number of signal averages, etc.), and the variable of interest is T_1.

A major weakness of the inversion recovery experiment is that the sample magnetization must be given full recovery time between each iteration of the experiment, so the experiment is very time-consuming. To address this, another way to measure T_1 is to use a *saturation recovery* [22]. This experiment is similar to the inversion recovery, but where the inversion recovery uses a delay time to allow magnetization to fully recover between repetitions, the saturation recovery applies a series of fast "crushing pulses" that serve to saturate the magnetization to zero between each repetition. The pulse sequence and resulting plot of amplitudes with respect to delay time, τ_n, is shown in Figure 1.16. The magnetization can be saturated in a few milliseconds, much more quickly than the time it takes to allow the magnetization to fully recover. In this way the need to wait is all but eliminated and the experiment is made to be much, much quicker.

Figure 1.16. The pulse sequence for a saturation recovery (a) consists of a series of i saturating pulses, typically $\pi/2$-pulses, followed by a delay time that allows the longitudinal magnetization to recover. A second pulse rotates the magnetic moment into the transverse plane for observation. The pulse sequence is repeated n times, with no need for a full recovery time between each repetition. The plot of the saturation recovery (b), which was constructed in Matlab™, is the series of observed amplitudes plotted against the recovery times, τ_n, for which they were acquired.

In the case of the saturation recovery, the magnetization starts out at zero, so there is no need for an inverting pulse. Only a delay time, followed by a $\pi/2$-pulse is needed. The equation to which the resulting plot in Figure 1.16 must be fit is accordingly adjusted to

$$S(\tau) = A\left(1 - e^{-\frac{\tau}{T_1}}\right) \tag{1.47}$$

where the variables are defined in the same way as Equation 1.46. Because of the immense time savings offered by the saturation recovery over the inversion recovery, the saturation recovery is almost always preferred for T_1 measurements.

1.9.2 Measuring T_2 – the Carr-Purcell-Meiboom-Gill (CPMG) Experiment

The time constant that defines the relaxation rate of the transverse magnetization back to its equilibrium (which is zero) is T_2. Since transverse magnetization can be directly observed, it might seem like one could just apply a $\pi/2$-pulse to the sample and start acquiring, as in an FID experiment. As the signal fades, that rate could be fit for a direct observation of T_2. It turns out that this decay is not T_2, but rather T_2^*, the "*apparent T_2*". Static magnetic fields generated for NMR are never perfectly homogeneous and different micro regions can and do experience slightly different static magnetic fields. This means that, after a $\pi/2$-pulse, the magnetic moments at different regions of a sample will precess at slightly different frequencies. After a short time, the moments of the different regions will begin to slip out of phase. The signal that is acquired, however, is the sum of all the magnetic moment vectors in the active sample region, also called *isochromats*. If the isochromats in the sample are precessing out of phase, they will destructively interfere and the signal will die out. This dying out is NOT true T_2 relaxation. As described in the *Sources of relaxation* section above, T_2 relaxation is a shrinking of magnetic moments back to thermal equilibrium due to energy exchange interactions between nearby spins.

In order to observe T_2 and not T_2^*, a pulse sequence capable of refocusing these diverging isochromats must be used. Similar to the discrete scheme of the inversion and saturation recoveries, spin echoes can be used to accomplish this [18]. In a spin echo, a $\pi/2$-pulse is used to rotate the sample magnetization into the transverse plane. After some time allowed for

isochromats to de-phase, τ_{del}, a second pulse – this time a π-pulse – is applied, which has the

effect of flipping the order of the precessing isochromats. Slow isochromats are sent to the front

of the line, while fast isochromats are tossed to the back. Now, the fast isochromats at the back

start to catch up and the slow isochromats at the front lag back to the middle. After another τ_{del}

time, the isochromats all re-phase and a "spin echo" is observed as shown in Figure 1.17(a).
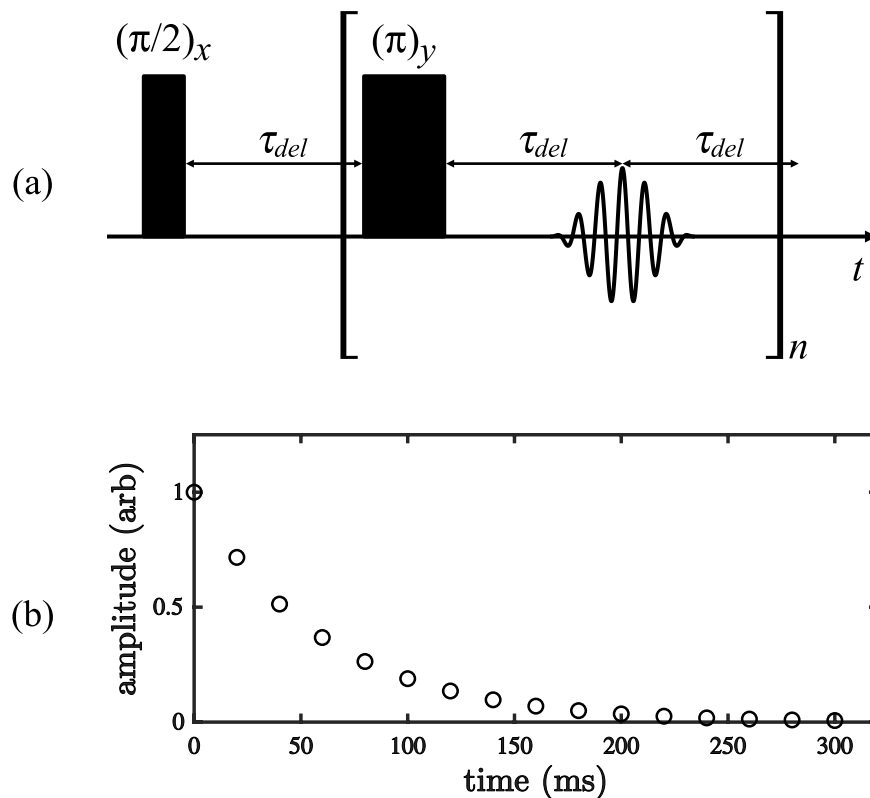
Figure 1.17. A Carr-Purcell-Meiboom-Gill (CPMG) pulse sequence (a) consists of a $\pi/2$-pulse, followed by a delay time, τ_{del}, during which the isochromats of a sample de-phase. A π-pulse flips the isochromats so that the heterogeneous precession now serves to re-focus the isochromats. After another τ_{del}, the spin echo is observed as the isochromats complete their re-focusing. The rephasing part of the pulse sequence is repeated n times, with an echo being acquired between each refocusing π-pulse. The plot of the CPMG (b) is the series of observed echo amplitudes plotted against the experiment time at which they were acquired.

If the spin echo experiment is repeated for longer and longer τ_{del} times, the amplitudes of each spin echo can be recorded. Since the all of the moments with various precession rates within the sample are refocused, the effect of de-phasing is removed from the observed decay in spin echo amplitude and any remaining decay is due to the effects of true T_2. As with the case of the inversion recovery, this strategy is quite time-consuming. Instead of repeating the experiment for longer and longer τ_{del} times, the Carr-Purcell-Meiboom-Gill (CPMG) experiment simply repeats the refocusing pulse and acquisition portion of the spin echo [19]. Shown in Figure 1.17(a), this is the part within the square brackets. With each refocusing and echo, the total time accumulated for the sample magnetization to decay is added to all the previous repetitions. The experiment is thus continually refocusing the field inhomogeneities and acquiring just the decay due to spin-spin relaxation. The amplitude of each point is then plotted against total experiment time and the series is fit to the equation for transverse relaxation:

$$S(t) = Ae^{-\frac{t}{T_2}} \tag{1.48}$$

where, as before, A is the arbitrary amplitude (at the beginning of the experiment) and T_2 is the variable of interest that is being fitted.

References

[1] M. H. Levitt, Spin Dynamics, West Sussex: Wiley, 2008.

[2] A. Carrington and A. D. McLachlan, Introduction to Magnetic Resonance, New York, NY: Harper and Row, 1967.

[3] NIST, "The NIST Reference on Constants, Units, and Uncertainty," [Online]. Available: https://physics.nist.gov/cgi-bin/cuu/Value?hbar|search_for=h-bar. [Accessed 29 April 2021].

[4] NIST, "The NIST Reference on Constants, Units, and Uncertainty," [Online]. Available: https://physics.nist.gov/cgi-bin/cuu/Value?h|search_for=h. [Accessed 29 April 2021].

[5] W. Gerlach and O. Stern, "Der experimentelle Nachweis der Richtungsquantelung im Magnetfeld," *Zeitschrift für Physik,* vol. 9, no. 1, pp. 349-352, 1922.

[6] W. Gerlach and O. Stern, "Das magnetische Moment des Silberatoms," *Zeitschrift für Physik,* vol. 9, no. 1, pp. 353-355, 1922.

[7] W. Gerlach and O. Stern, " Der experimentelle Nachweis des magnetischen Moments des Silberatoms," *Zeitschrift für Physik,* vol. 8, no. 1, pp. 110-111, 1922.

[8] H. Schmidt-Böcking, L. Schmidt, H. Lüdde, W. Trageser, A. Templeton and T. Sauer, "The Stern-Gerlach Experiment Revisited," *European Physical Journal H,* vol. 41, p. 327–364, 2016.

[9] I. I. Rabi, S. Millman and P. Kusch, "The Molecular Beam Resonance Method for Measuring Nuclear Magnetic Moments," *Physical Review,* vol. 55, pp. 526-535, 1939.

[10] F. Bloch, "Nuclear Induction," *Physical Review,* vol. 70, no. 7-8, pp. 460-474, 1946.

[11] F. Bloch, W. W. Hansen and M. Packard, "The Nuclear Induction Experiment," *Physical Review,* vol. 70, no. 7-8, pp. 474-485, 1946.

[12] E. M. Purcell, H. C. Torrey and R. V. Pound, "Resonance Absorption by Nuclear Magnetic Moments in a Solid," *Physical Review,* vol. 69, no. 1-2, pp. 37-38, 1946.

[13] D. J. Griffiths, Introduction to Electrodynamics, 3rd Edition, Upper Saddle River, New Jersey: Prentice-Hall, Inc., 1999.

[14] R. K. Harris, Nuclear Magnetic Resonance Spectroscopy, 1986.

[15] N. Bloembergen, Nuclear Magnetic Relaxation, New York: W. A. Benjamin, Inc., 1961.

[16] C. P. Poole and H. A. Farach, Relaxation in Magnetic Resonance, New York: Academic Press, 1971.

[17] NIST, "NIST Reference on Constants, Units, and Uncertainty," [Online]. Available: https://physics.nist.gov/cgi-bin/cuu/Value?gammae. [Accessed 18 June 2021].

[18] NIST, "NIST Reference on Constants, Units, and Uncertainty," [Online]. Available: https://physics.nist.gov/cgi-bin/cuu/Value?gammapbar. [Accessed 18 June 2021].

[19] B. Blümich, F. Casanova and J. Perlo, "Single-sided NMR sensor with microscopic depth resolution". United States of America Patent US7358734B2, 15 April 2008.

[20] A. E. Marble, I. Mastikhin, B. G. Colpitts and B. J. Balcom, "A Compact Permanent Magnet Array with a Remote Homogeneous Field," *Journal of Magnetic Resonance,* vol. 186, pp. 100-104, 2007.

[21] S. Utsuzawa and F. E., "Unilateral NMR with a barrel magnet," *Journal of Magnetic Resonance,* vol. 282, pp. 104-113, 2017.

[22] E. Fukushima and S. B. W. Roeder, Experimental pulse NMR: a nuts and bolts approach, London: Addison-Wesley, 1981.

[23] E. L. Hahn, "Spin Echoes," *Physics Review,* vol. 80, no. 4, pp. 580-594, 1950.

[24] S. Meiboom and D. Gill, "Modified Spin-Echo Method for Measuring Nuclear Relaxation Times," *The Review of Scientific Instruments,* vol. 29, no. 8, pp. 688-691, 1958.

[25] H. C. Helgeson and D. H. Kirkham, "Theoretical Prediction of the Thermodynamic Behavior of Aqueous Electrolytes at High Pressures and Temperatures: I. Summary of the Thermodynamic/Electrostatic Properties of the Solvent," *American Journal of Science,* vol. 274, p. 1089–1198, 1974.

Chapter Two – Determination of Dielectric Constants of Solutions at High Pressures by Partial Least Squares Regression

---

An application of partial least squares analysis to elucidate a non-ideal system

"In the midst of chaos, there is also opportunity."

– Sun Tzu

2.1 Forward

The following is an exploration in the use of multivariate regression to analyze a highly imperfect spectroscopic problem. This is not actually an application of NMR, but it serves as a good example of how multivariate, and specifically partial least squares (PLS) regression, can be used in messy environments such as those encountered in low-field and single-sided NMR. It also contains some arguments supporting the use of PLS over other multivariate regression models. A full explanation of PLS is not provided in this section, however, a brief primer is provided in Appendix 2 for a better understanding. The purpose of this section in this manuscript is to provide a discussion of one more tool in the toolbox of the spectroscopist who finds themselves in the unfortunate position of having to tease out conclusions from noisy and otherwise indecipherable environments.

2.2 Abstract

The determination of dielectric changes in liquids at high pressures is of interest to geophysicists endeavoring to understand the nature of electromagnetic (EM) waves deep within the Earth, as well as changes to chemical processes at high pressures. In this environment, the concentration of salts must also be considered. Previously, studies have experimentally described the behavior of the dielectric constant in pure water at pressures up to 2.0 GPa, as well as in ionic solutions at ambient pressures. However, the literature currently does not include an experimental investigation of the case of an ionic solution at high pressure. This study proposes a new method for the experimental determination of aqueous, ionic solutions up to 2.0 M and at pressures up to 2.0 GPa.

2.3 Introduction and background

To study the chemistry of electrolyte solutions at high pressures, an accurate depiction of the Gibbs free energy for the solvation of the constituent ions is necessary. This can be theoretically estimated using the Max Born equation for the energy of solvation of an ion:

$$\Delta G_{Solv} = \frac{z_i^2 e^2 N_A}{8\pi\varepsilon_0 r_i}\left(1 - \frac{1}{\varepsilon_r}\right) \tag{2.1}$$

where z_i is the charge number of the ion, e is the elementary charge, N_A is Avogadro's number, ε_0 is the permittivity of free space, r_i is the ion's effective radius, and ε_r is the relative permittivity of the solution that is solvating the ion.

This equation unfortunately does not account for changes in temperature or pressure. To address this limitation, Helgeson, Kirkham and Flowers introduced a modification (known as the HKF model) in which empirically derived parameters are added and the effective ionic radius, r_i, is redefined as an adjustable parameter that is a function of temperature and pressure [1]. However, the static relative permittivity of the solvent remains important for describing these states.

The permittivity of a material describes how it will be polarized in an electric field according to the equation

$$P = \varepsilon_0 \left(\frac{\varepsilon}{\varepsilon_0} - 1 \right) E \tag{2.2}$$

where P is the polarization, ε is the permittivity of the material, and E is the electric field imposed on the material. The term, "dielectric constant," by convention refers specifically to the relative permittivity, ε_r (or sometimes κ) of a material

$$\varepsilon_r = \kappa = \frac{\varepsilon}{\varepsilon_0} \tag{2.3}$$

which is unitless. Equation 2.2, shows that a higher dielectric constant will lead to a greater polarization of a material in a given electric field. Permittivity also has implications for changing electric fields, as can be inferred from Maxwell's equation for an induced magnetic field

$$\nabla \times H = J_f + \varepsilon \frac{\partial}{\partial t} E \qquad (2.4)$$

where H is the magnetic field, and J_f is the current flux.

According to Equation 2.4, the strength of an induced magnetic field, due to a changing electric field, increases in proportion to the permittivity of the medium, as well as the strength of the current flux. Therefore, in addition to its effect on chemical equilibria, an understanding of dielectric changes in ionic solutions with respect to concentration and pressure is also critical to predicting that material's response to EM radiation.

Several methods currently exist to measure the dielectric constant of materials at atmospheric pressures. For instance, one may pass an alternating current between parallel plates, and by comparing the capacitance with a dielectric versus without, calculate the dielectric constant, according to the equation

$$\varepsilon_r = \frac{C_{diel}}{C_0} \qquad (2.5)$$

where C_{diel} is the capacitance measured with the dielectric and C_0 is without. In another typical experiment, the reflected power of a band of frequencies is measured, after placing a flat end of a coaxial transmission line into a sample. The reflection coefficient, Γ, at the junction of two materials of different impedances is

$$\Gamma = \frac{Z_1 - Z_2}{Z_1 + Z_2} \tag{2.6}$$

where Z_1 is the impedance of the first medium, and Z_2 the impedance of the second medium. The impedance of a medium is itself related to permittivity, thus, in principle, allowing its calculation. Both approaches require assumptions. In the first approach, the parallel plates must be effectively infinite compared to the gap between them. In the second, the sample depth must be effectively infinite compared to the wavelength being used. For both, the sample cannot be not surrounded by a conductor.

Pressurizing samples at GPa pressures, requires materials able to withstand such pressures. Also, to prevent extrusion, the pressure cell must enclose the sample. The pressure cell being used here is composed of Berylco© 25 (BeCu), which is conductive, thus invalidating the assumption that nothing conductive surrounds the sample. Further, the sample size must be small for a hydraulic press to push the samples to such high pressures. This means that the depth of the material is not effectively infinite. Also, to analytically determine (even computationally) the reflection coefficient as a function of the dielectric constant, one must know the exact dimensions within the pressure cell. At the small scale of the sample capsule, this is extremely difficult, if not impossible. Even a half-millimeter error would drastically change the calculation results.

Because of these challenges, the reflection spectra are instead correlated to known dielectric constants using a multivariate statistical method called PLS regression, which seeks to correlate two sets of data by analyzing variances in one set of data (predictors or independent data) and calibrating those with the variances of a resulting data set (responses or dependent

data). A successful correlation will allow a determination of the dielectric constant for a material based on the reflection spectrum alone. Additionally, this same regression could also be used to predict other parameters, so long as they relate to the reflection coefficients.

Because statistical correlation by nature disregards a strict analytical link between the system and the predictions, a strong argument must be made that the system is one that indeed lends itself well to the use of a particular regression model. In this case the argument that must be made is that the changes in reflection spectra will be in response to, and consistent with, underlying changes in permittivity. Fortunately, such a link exists. Reflection spectra a samples in this experiment are recorded by connecting a 50 Ω coaxial cable to a pressure cell, which is acts as a capacitor in which the sample itself is the dielectric material. Applying Equation 2.6 to this system, the first impedance, Z_1, is assigned the value of the coaxial cable, while the second impedance, Z_2, is assigned the value of the capacitive pressure cell and dielectric sample. In an ideal capacitor, the impedance is dependent on both the frequency and the dielectric constant, and is given by

$$Z_c = \frac{1}{i\omega\varepsilon_r C_0} \tag{2.7}$$

where ω is frequency and C_0 is the value of the capacitor without dielectric. While the pressure cell is certainly not an ideal capacitor, one can still be certain that the impedance – and therefore the reflection spectrum – will change with different dielectric constants. Thus, if one can correlate different reflection spectra with different known dielectric constants, these data can be used as a map for predicting the dielectric constant of new samples outside the initial dataset.

It is also important to understand how the concepts of how PLS works and why this statistical method would be better than other multivariate modelling (MM) schemes. Two other MM strategies are very popular at the time of this writing and they are both worth briefly considering. Principle component analysis (PCA) is a way of mathematically rearranging the dimensions of a dataset (e.g. points of spectra) and finding components with the most variance. Those components can then be used as a simplified characterization of the data. If those components are then correlated with the components of greatest variance in the dependent dataset (outcomes), a link may be made for predictions. While it may be tempting to apply such a scheme to this problem, it is not clear that the largest variance in the reflection spectra will necessarily be due to changes in the dielectric constant of the sample material. Multiple linear regression (MLR) is a way of directly using an independent dataset and its associated dependent dataset to find a key, b, which can be applied to a new independent observation, such that a prediction (or estimate) vector is reached. The elements of b are called the *sensitivities* coefficients. This relationship between the sensitivities, predictors, and predictions can be mathematically represented as

$$bx = \acute{y} \tag{2.8}$$

where x is some set of predictors, such as the points of a spectrum, and $\acute{y}$ is a set of predictions. The key, b, may be a vector if only one prediction is sought, or it may be a matrix if it is to be used to make multiple predictions. This also may seem quite inviting, however the math behind MLR turns out to treat all component variances as equally important, regardless of how well they actually correlate to changes in the dependent dataset.

PLS is very similar to MLR, but instead of a rigid pseudoinverse scheme for finding **b**, it uses a system of weighting and scoring guesses based on data in both the independent and dependent dataset, and then using the outcome of that procedure to form a new guess, until a stable value is reached. In perfect data and with all components being used, this actually becomes the same as MLR. However, in noisy and imperfect data, PLS has the advantage that it rejects correlations between components in the independent and dependent data that do not vary together. Thus it is less likely to be "fooled" by variance due to noise or other conditions that are not connected to the predictions being sought. It is for this reason that PLS was chosen for this experiment.

In practice, PLS is not easily solved for all principle components. Instead, the components (and sensitivities) are calculated iteratively, starting with the components of greatest weight, in a algorithm called Non-linear Iterative Partial Least Squares, or NIPALS. A more in-depth explanation of this method is described in Appendix 2.

2.4 Methods

Before attempting high-pressure measurements, the concept was first tested with various concentration of $NaCl_{(aq)}$ in an open beaker. A section of stiff coaxial transmission line, arranged to maintain constant depth, was placed into the solution and was connected by a coaxial cable to a network analyzer that displayed the spectrum of reflection coefficients. The reflectivity spectra were then correlated to the calculated static dielectric constants of each sample. Known theoretical values for dielectric constants of solutions with known NaCl concentrations were calculated using methods described by Hasted, Ritson, and Collie [2]. A full cross-validation was performed, in which one spectrum was sequentially removed from the independent data set

before performing a regression on the remainder. This then provided a predicted value for the

dielectric constant of the removed spectrum, which was compared to the actual value in the

dependent data.

The same analysis was then repeated in a pressure cell, shown in Figure 2.1. The probe

construction and the pressurization procedure are discussed in literature [3]. An enameled wire

carrying the signal from the network analyzer was inserted through the feedthrough into the

sample chamber in the pressure capsule and an application of Stycast ™ was used to seal and

secure the wire in place. Additionally, a small ruby bead was secured inside the pressure

chamber, which was epoxied to a tail of optical fiber that was also inserted through the

feedthrough and secured with the same Stycast™. This ruby and optical fiber allowed the

experimenter to read the pressure inside the cell using fluorescence spectroscopy [4]. The sample

in the sample chamber was then placed inside the main body, which served as the electrical

ground. In this way, the enameled wire, the electrolyte sample, and the surrounding pressure cell

body formed the plate-dielectric-plate components that one would think of in a traditional

capacitor, but with radial, rather than the traditional flat geometry.

Figure 2.1: The BeCu pressure cell is shown in a blown up diagram. The
pressure capsule holds approximately 0.5 mL.

This second iteration was still without elevated pressure. As before, the samples were prepared with increasing molarity of NaCl. PLS regression and cross-validation was performed. A spectrum of a sample with a concentration different from any in the calibration set was also recorded and its dielectric constant was predicted for comparison.

In the final test, a set of measurements was obtained with pure water, at pressures from ambient up to 1.0 GPa, using the same probe and pressure cell as before. Then, using the same setup, spectra were recorded at ambient pressure, with concentrations up to 1.0 M. Once spectra were taken at both standard conditions (zero concentration and ambient pressure), PLS regression was performed on this calibration set. Known theoretical values for the dielectric concentration of pure water at elevated pressures were calculated using methods described by Floriano and Nascimento [5]. To ensure the validity of the PLS predictions, three steps were performed:

1) Cross-validation of the calibration set.
2) Predictions from outside spectra taken either at varying concentrations, or at high pressures (not both) were compared.
3) Prediction from outside measurements with both elevated pressure and concentration, were compared to expected analytical predictions.

The first two steps can be considered a "calibration phase," while the third step is the "exploratory phase." While it is imperative that the calibration tests be consistent, the third test is in fact an exploration of the desired response.

2.5 Results

The initial tests of the PLS methods were conducted in an open 50 mL beaker of varying concentrations of NaCl, and yielded promising results. Initial inspection of the spectra showed a clear trend correlated with changing dielectric constant. The results of cross-validation showed that the PLS regression predicted the dielectric constant within +/- 1.9 (p = 0.05), out of an expected theoretical range of approximately 27. For the outside sample (not included in cross-validation) a solution of 0.375 M, corresponding to a theoretical permittivity of 74.18, was prepared. This sample was predicted to have a dielectric constant of 73.81, only 0.37 below the actual value, and a percent error of 0.49%.

In the next test, spectra were taken at varying concentration again, still at ambient pressure, but now inside the pressure cell. In this trial, cross-validation yielded predictions within +/- 2.4 (p = 0.05) of the dielectric constant. Using the calibration set to predict the dielectric constant of the outside 0.375 M sample yielded a result of 72.61, 1.57 below the theoretical value, which is a 2.1% error.

Finally, spectra of pure water in the pressure cell were collected at pressures up to 0.62 GPa (estimated), as well as of NaCl solutions at ambient pressures (Figure 2.2). In this set, it was not feasible to repeat measurements at each condition as before, because the pressure cell must be sequentially ramped up for each measurement, without releasing pressure. Cross-validation from this calibration set showed that the model would predict within only +/- 10 (p = 0.05), yet modeled the overall trend of the dielectric constants well (Figure 2.3). The high variability is likely due to the inability to incorporate repeat measurements into the calibration set as before.

Figure 2.2. The reflectivity spectra of sample at either elevated pressure, or elevated ionic concentration.

Figure 2.3. A comparison of dielectric constants predicted by PLS regression (circles) versus theoretically calculated values (squares). The y-axis reflects the values calculated by the PLS algorithm, whereas the x-axis is the known value of dielectric constants from calculating from the sample pressure and NaCl concentration. The blue line is a least-squares fit to the predicted data and shows that despite poor precision of the individual predictions, the overall trend of dielectric constant change is quite accurate.

2.6 Discussion

The goal of this work has been to devise an experimental method for measuring the

dielectric constant of highly pressurized electrolyte solutions and the results show that using PLS

regression on reflectivity measurements of a sample is a feasible method. The results of the first

and second iterations (open beaker and no-pressure, respectively) show that, in a well-controlled

environment, the predictive power of this method is remarkable. In the third iteration, the samples were pressurized or had increased concentrations of NaCl, and it was showed that even in this difficult to control environment, where the spectra are seemingly more disparate, PLS is still able to model the trends quite well. This is impressive, although the ideal result would be precise as well as accurate.

In their work, Hasted, Ritson, and Collie measured dielectric constants of ionic solutions at ambient pressure within +/- 1.8 of expected (calculated) values [2]. This method would ideally replicate that level of precision.

2.7 Conclusion and potential future work

The method of PLS regression of reflectivity measurements shows promise as an experimental test of dielectric constant behavior in pressurized ionic solutions. With refinement, the accuracy, precision, and therefore the validity of this method will allow for testing the behavior of the dielectric constant at simultaneously high pressures, as well as elevated ionic concentration. In turn, these measurements and their trends will provide a way to test, not just the theories that currently govern the behaviors of the dielectric constant at high pressures, but also the ionic solvation energy as it is affected by high pressures.

The remaining work in the calibration phase will focus on improving control of the electrical and geometric environment of the sample, as well as expanding the calibration set size. The spectrum of the sample is greatly stabilized by providing electrical shielding, and the geometry of the sample test environment must be controlled with respect to Stycast™ volume protruding into the sample space, and probe wire length. Additionally, improvements to sample

preparation, and strengthening the sample set-up, will enable expanded calibration set sizes, as fewer experiments are terminated early due to failed sample containers.

The conclusions from subsequent exploratory phase tests will be applicable in extreme environments that are of increasing interest to both geophysicists and astro-chemists who encounter chemical processes at extreme pressures. In the future, similar methods may also be used for exploring the interplay of other factors such as thermal or optical properties at high pressure.

References

[1] H. C. Helgeson, D. H. Kirkham, Theoretical Prediction of the Thermodynamic Behavior of Aqueous Electrolytes at High Pressures and Temperatures: I. Summary of the Thermodynamic/Electrostatic Properties of the Solvent. *American Journal of Science*, 274 (1974) pp. 1089–1198.

[2] J. B. Hasted, D. M. Ritson, C. H. Collie, Dielectric Properties of Aqueous Ionic Solutions. Parts I and II. *The Journal of Chemical Physics*, 16 (1948) 1-21.

[3] G. Ochoa, C. D. Pilgrim, M. N. Martin, C. A. Colla, P. Klavins, M. P. Augustine, W. H. Casey, ^{2}H and ^{139}La NMR Spectroscopy in Aqueous Solutions at Geochemical Pressures. *Angew. Chem. Int. Ed.*, 54 (2015) pp. 15444–15447.

[4] O. Grasset, Calibration of the R Ruby Fluorescence Lines in the Pressure Range [0-1GPa] and the Temperature Range [250-300 K]. *High Pressure Research*, 21 (2001) pp. 139-157.

[5] W. B. Floriano, M. A. C. Nascimento, Dielectric Constant and Density of Water as a Function of Pressure at Constant Temperature. *Braz. J. Phys*. 34 (2004)

Chapter Three – Simple Magnetic Field Models

A description of codes used for designing coils and magnets

"Ten percent of nothin' is ... let me do the math here ...
nothin' into nothin' ... carry the nothin' ... "

-Jayne, Firefly

3.1 Introduction

In the process of designing coils for low-field and unilateral magnetic resonance (UMR), it became clear that a quick method for modelling coil designs would be critical. Coil design had previously involved a guess-and-check approach that was labor intensive and time consuming. This led to literal piles of test coils (Figure 3.1). While computer models (especially simple ones) are rarely able to predict reality with 100% accuracy, even a rough prediction can save significant research time if the researcher remains cognizant of the strengths and limitations of the code. There are of course commercially available software suites that are very capable in performing these exact kinds of calculations, but they are very expensive and require significant training to use. What is needed is an inexpensive, easy to use, and quick way to test ideas. This is precisely where the MagModel fits into the work flow of UMR and coil design. MagModel is a Matlab™ based code that was written by myself during my time in the Augustine Group at UC Davis between the years 2016 and 2021. The code is not included here due to length, but may be acquired by contacting myself or the Augustine Group. The following is a brief description of the MagModel code and instructions on how to use it.

Figure 3.1. Designing coils by trial and error is a time-consuming and laborious process. This shows a sample of the numerous iterations that were built and tested on the path of designing a coil for one system.

3.2 MagModel general description

MagModel is a Matlab™ code that takes a user's instructions for an array of coils or permanent magnet blocks and calculates the field due to that array, then plots the results for the user's viewing [1]. At its core, MagModel is a Biot-Savart law and magnetic dipole field solver. The Biot-Savart Law gives the magnetic field, $\boldsymbol{B}$, at a point, $\boldsymbol{r}$, due to an electrical current, I, through a line element, $d\boldsymbol{l}$ [2].

$$\boldsymbol{B}(\boldsymbol{r}) = \frac{\mu_0}{4\pi} \int_l \frac{I \cdot (d\boldsymbol{l} \times \boldsymbol{r}')}{|\boldsymbol{r}'|^3} \tag{3.1}$$

where μ_0 is the permeability of free space, $\boldsymbol{r'}$ is the vector from the line element $d\boldsymbol{l}$ to $\boldsymbol{r}$, and the integral is taken over all current-containing line elements. The code further assumes that the

current through a coil is simply a sinusoidal application of *I*, and thus the B_1 field calculated is a time-independent amplitude for an assigned given direction.

The magnetic field, **B**, at some point, **r**, from a magnetic dipole, **m**, that is considered to be at the origin is given by

$$\boldsymbol{B}(\boldsymbol{r}) = \frac{\mu_0}{4\pi}\left[\frac{3\boldsymbol{r}(\boldsymbol{m}\cdot\boldsymbol{r})}{r^5} - \frac{\boldsymbol{m}}{r^3}\right] \tag{3.2}$$

where the non-boldface r refers simply to the distance between the point of observation and the dipole [1]. This equation is the cartesian equivalent of Equation 1.45. In theory, the user could calculate the exact dipole moment per point in the magnetic block and use this to calculate the actual magnetic field of the permanent magnet in a simulation, but the usefulness of the code is really in characterizing the *shape* of the B_1 field, respective to the B_0 field.

Numerical outputs are also able to be acquired from the results of the code. The numerical outputs of MagModel are in SI units, but again, the utility of the code lies more in the qualitative field geometry than in the numerical predictions of field strength and direction. User inputs are provided by altering Excel™ files. In this way, a user may save sets of useful Excel™ files for future use, should certain designs be needed again. An example of the output for a simple loop above a magnetic block is provided in Figure 3.2.
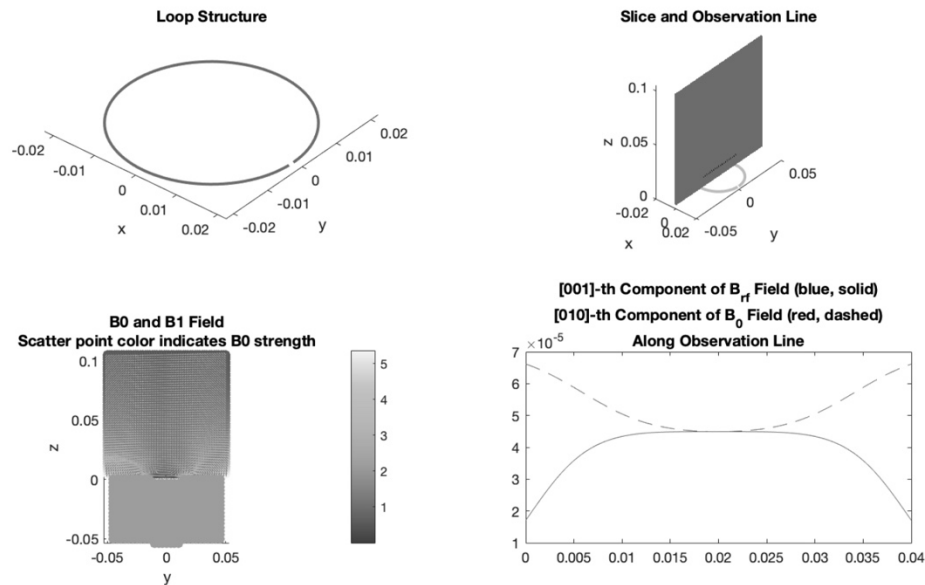
Figure 3.2. The output of MagModel displays the user-defined coil loops (upper-left); the slice in space at which the magnetic field will be calculated, as well as the line on which further magnetic field calculations will be made, overlayed onto the coil (upper-right); a depiction of both the static and rf magnetic fields plotted over the loop and magnet blocks (lower-left); and a plot of the user-defined component of the magnetic field strengths, along the user-defined line (lower-right). All plots can be turned and magnified by the user.

3.3 Using MagModel

3.3.1 Entering data

In order to perform a calculation, the user must first define the initial parameters:

1. The coil loops

2. The permanent magnet blocks

3. The slice of space upon which the resulting magnetic field will be calculated

4. The line across space upon which the resulting magnetic field will be calculated, as well as the direction component of the field that will be displayed.

72

All initial parameters are defined using Excel™ spreadsheets that are saved in the same folder as the main code and have specific names. To define parameters in a particular spreadsheet, the desired spreadsheet is opened using Excel™ and the required data is entered. Data is entered in lines for each entity (e.g. coil loops) and there should be no empty lines between entities.

Define the coil by entering data in LoopData.xlsx. An example is shown in Figure 3.3, with the entries used to make the results in Figure 3.2. Each line contains the data for one loop of the coil. Loop elements are defined by assigning a shape length and width that is assumed to initially lie in the xy-plane, and centered at the origin. Loops are drawn by defining a number of sequential points in the shape called for by the user. The position of the loop element can then be offset in the x-, y-, or z-directions. Finally, the loop element can be rotated about the y- or z-axes. Rotations are applied assuming the right hand rule, for example, a positive rotation assigned to the y-axis will rotate the loop consistent with a torque vector aligned along the positive y-axis. Operations (drawing, moving, and rotating the loop) are applied to each loop element in the order that they have been described here. These conventions hold data entered into all spreadsheets. In the order of the columns, the data is entered as follows in Table 3.1.
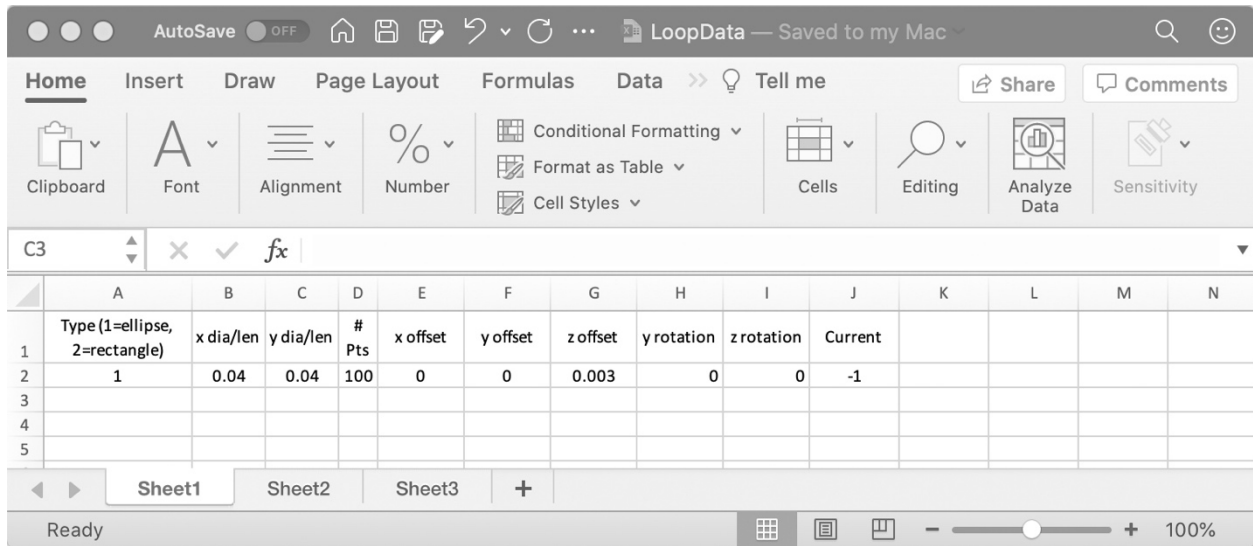
Figure 3.3. The LoopData.xlsx spreadsheet is used to define the loops of the rf coil being modeled. Shown here, LoopData.xlsx contains the entries that were used to obtain the results in Figure 3.2.

Table 3.1. Data required for LoopData.xlsx

Column(s)	Description of data	Data requirements / meanings
A	Determines whether the loop element will be an ellipse or rectangle.	Use a "1" for an ellipse, or a "2" for a rectangle.
B and C	Assigns the diameters of the ellipse or the side lengths of the rectangle in the x and y directions, respectively.	Enter a float. Entries are understood to be in meters.
D	Assigns the number of points that will be used to draw the loop.	Enter an integer.
E, F, G	Used to apply a translational offset in the x-, y-, and z-directions, respectively.	Enter a float. Entries are understood to be in meters.
H and I	Used to apply a rotation about the y- and z-axes, respectively.	Enter a float. Entries are understood to be in radians. Operations are applied about the y-axis first, then the z-axis. Rotations are applied about the loops own center (not the universal origin).
J	Assigns a current to the loop.	Enter a float. Entries are understood to be in amperes. The direction of current is clockwise, as viewed from above, when the loop is initially drawn in columns A-C.

Define the blocks of permanent magnet material using BlockData.xlsx. An example is shown in Figure 3.4, with the entries used to make the results in Figure 3.2. Each line contains the data for one block element of the overall magnet. Block elements are defined by assigning a block side length in the x-, y-, and z- direction. A density of points must then be assigned in terms of points per meter, and the position of the block element can then be rotated about the y- or z-axes assuming the right hand rule. Next, the loop element can be offset in the x-, y- or z-

directions. Finally, the dipole moment direction of the magnet block is assigned. In the order of

the columns, the data is entered as follows in Table 3.2.

x Length	y Length	z Length	Point Density (per meter)	y Rotation	z Rotation	x Offset	y Offset	z Offset	x Dipole	y Dipole	z Dipole
0.1	0.03	0.05	333	0	0	0	0.03	-0.025	0	1	0
0.1	0.03	0.05	333	0	0	0	-0.03	-0.025	0	1	0
0.1	0.02	0.05	333	0	0	0	0	-0.0285	0	1	0

Figure 3.4. The BlockData.xlsx spreadsheet is used to define the magnetic block elements of the permanent magnet, which produces B_0. Shown here, BlockData.xlsx contains the entries that were used to obtain the results in Figure 3.2. These elements are based on a three-magnet array as used in the Balcom Group at University of New Brunswick.

Table 3.2. Data required for BlockData.xlsx.

Column(s)	Description of data	Data requirements / meanings
A, B, C	Assigns the side lengths of the block in the x-, y-, and z-directions, respectively.	Enter a float. Entries are understood to be in meters.
D	Assigns the point density used to model the block.	Enter an integer. Entries are understood to be in points per meter and are applied in all cardinal directions in a uniform, cubic lattice.
E, F	Used to apply a rotation about the y- and z-axes, respectively.	Enter a float. Entries are understood to be in radians. Operations are applied about the y-axis first, then the z-axis.
G, H, I	Used to apply a translational offset in the x-, y-, and z-directions, respectively.	Enter a float. Entries are understood to be in meters.
J, K, L	Assigns the x-, y-, and z-components of the magnetic dipole for the block element.	Enter floats. The code treats the dipole components as the components for each individual point making up the larger block element (x-, y-, and z-directions for Columns J, K, and L, respectively). An accurate simulation will therefore scale these values to their density as determined by Column D.

Magnetic field vectors are calculated at an array of points defined on a square slice of space that is defined by the user. The slice is defined using SliceData.xlsx. An example is shown in Figure 3.5, with the entries used to make the results in Figure 3.2. Only one line is used for this spreadsheet, as only one slice is displayed for observation of simulation results. The slice is defined by assigning a side length and a number of points per side. This sets the density of points for the observation slice. The slice can then be rotated about the y- or z-axes assuming the right

hand rule. Next, the slice is assigned offsets in the x-, y- or z-directions. In the order of the

columns, the data is entered as follows in Table 3.3.
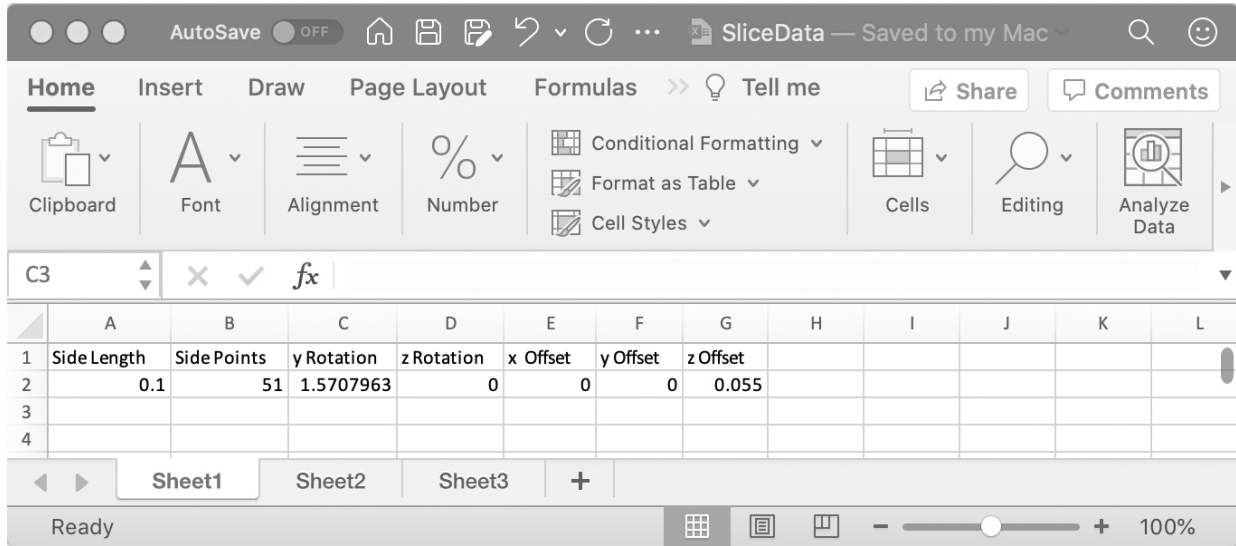
Figure 3.5. The SliceData.xlsx spreadsheet is used to define the slice upon
which magnetic field data will be displayed. Shown here, SliceData.xlsx
contains the entries that were used to obtain the results in Figure 3.2.

Table 3.3. Data required for SliceData.xlsx.

Column(s)	Description of data	Data requirements / meanings
A	Assigns the side length of the slice.	Enter a float. Entries are understood to be in meters.
B	Assigns number of points per side of the slice.	Enter an integer. The combination of this entry and the entry in Column A will define the point density for the observation slice.
C, D	Used to apply a rotation about the y- and z-axes, respectively.	Enter a float. Entries are understood to be in radians. Operations are applied about the y-axis first, then the z-axis.
E, F, G	Used to apply a translational offset in the x-, y-, and z-directions, respectively.	Enter a float. Entries are understood to be in meters.

Finally, MagModel will calculate the strength of the B_1 and B_0 magnetic fields along a line in space, defined by the user. The user also defines what direction of the magnetic field is to be plotted. The observation line is defined using LineData.xlsx. An example is shown in Figure 3.6, with the entries used to make the results in Figure 3.2. As it is for the slice, only one line is used for this spreadsheet, as only one line is displayed for observation of simulation results. The user assigns cartesian start and end points, the number of points to be calculated, and the vector component of the magnetic field to be calculated. The direction of the B_1 field is determined by assigning values in the x-, y-, and z-directions, and the magnitude of the final observation vector is normalized to a magnitude of one. In this way, if a user wishes to observe along a vector halfway between the y- and z- axis, assigning observation components of [0 1 1] will not produce an observation vector that is 1.414 in length. The results of the B_0 calculation are normalized to the maximum of the B_1 plot values. This is so that both fields can be viewed on the same plot, even with very different field strengths. It does, of course, mean that the actual B_0 field strength is not displayed for the user, but the purpose of this code is more to reveal the qualitative *shape* of the fields as they relate to each other. The abscissa in the line plot displays the distance from the line start point. In the order of the columns, the data is entered as follows in Table 3.4.
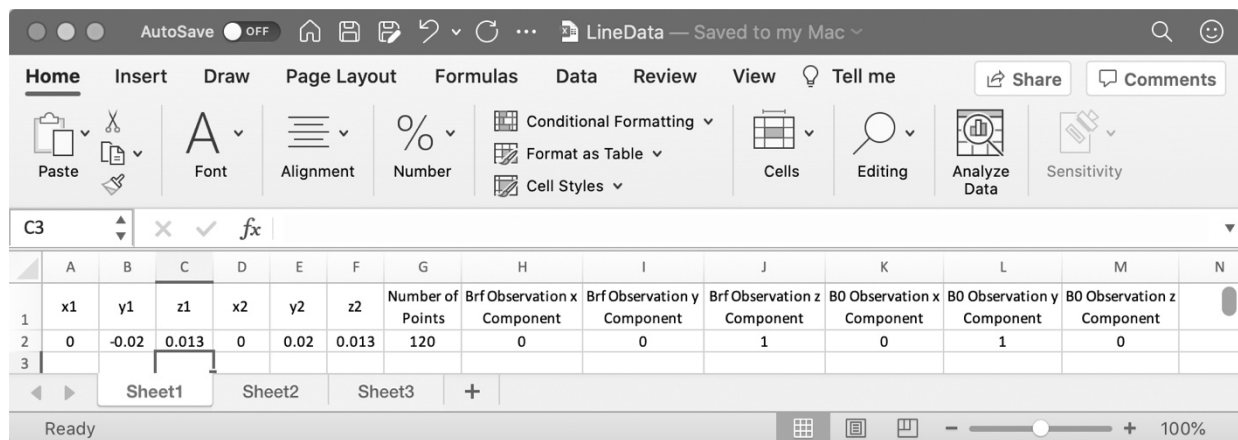
Figure 3.6. The LineData.xlsx spreadsheet is used to define the line upon which magnetic field data will be displayed. Shown here, LineData.xlsx contains the entries that were used to obtain the results in Figure 3.2.

Table 3.4. Data required for LineData.xlsx.

Column(s)	Description of data	Data requirements / meanings
A, B, C	Assigns the start point of the line along the x-, y-, and z-axes, respectively.	Enter floats. Entries are understood to be in meters, relative to the simulation space origin [0 0 0].
D, E, F	Assigns the end point of the line along the x-, y-, and z-axes, respectively.	Enter an integer. The combination of this entry and the entry in Column A will define the point density for the observation slice.
G	Assigns the number of points along the line.	Enter an integer. More points mean that more calculations will be performed and the resulting plot will be higher resolution.
H, I, J	Assigns the direction of the B_1 field that will be displayed, in terms of x-, y-, and z-directions, respectively.	Enter floats. Typically, it is sufficient to enter "1" into the cartesian direction of interest and to leave the rest as "0." If a direction other than the normal x-, y-, and z-directions is desired, the user may enter some combination of values, which will be added vectorially and normalized to 1 to determine the direction of observation.
K, L, M	Assigns the direction of the B_0 field that will be displayed, in terms of x-, y-, and z-directions, respectively.	Enter floats. Typically, it is sufficient to enter "1" into the cartesian direction of interest and to leave the rest as "0." If a direction other than the normal x-, y-, and z-directions is desired, the user may enter some combination of values, which will be added vectorially. Since the values for B_0 are scaled to the magnitude of B_1 for display, this vector is effectively normalized as well.

3.3.2 Executing the calculation

Once the coil, slice, block, and line data has been entered, the user must indicate if the permanent B_0 magnetic field is to be calculated and, if so, whether the field is to be calculated from the information in BlockData.xlsx, or if the data should be acquired from a saved file. Often, only the magnetic field of the coil is desired to be modeled. In these cases, the B_0 field calculations will only serve to clutter and confuse the resulting plots. This determination is set by altering the code at line 8 in MagModel. Simply set `B0_on = true;` if the permanent field is to be used or `B0_on = false;` if not.

The B_0 field plot requires calculating the magnetic field from every point in all of the block magnets, at each point on the slice. The calculation time, T_c, for the B_0 field on the slice thus scales roughly as

$$T_c \propto d_b{}^3 l_b{}^3 P_s{}^2 \tag{3.1}$$

where d_b, l_b, and P_s are the point density (per meter) of the block, the length of the block side (assuming a cube), and the number of points per side of the slice, respectively. As can be seen, the calculation burden can become enormous quite quickly with only moderately increasing dimensions of the block. Even an increase in the slice dimension increases computations by a square.

For this reason, every time the code executes, a .mat file that contains the calculated B_0 values on the observation slice is saved in the folder ./SavedFiguresAndData, located in the same folder as the main .m file, MagModel. Thus, once a particular magnet block configuration has been used and the calculations for it have been performed on a slice, the user can use these

previous data for future simulations using the same magnet block and observation slice, instead

of recalculating the B_0 field every time. If `B0_on` is set to `true`, indicate if a saved file is to be

used on line 10 by setting `UseSavedB0Field = true;` if a saved file is to be used and

`UseSavedB0Field = false;` if not. If a saved file is to be used, change the file path on

line 12 to the appropriate path where the previous data can be found. An example is shown in

Figure 3.7. Remember that the saved data are valid for a particular slice, so the parameters in

SliceData.xlsx must match the saved slice data. If this is not validated to be the case, the code

will give the user the option to calculate B_0 from BlockData.xlsx instead, or abort the simulation.

```
MagModel.m
1     function MagModel
2     % This is the main function for the magnetic field model. To run this
3     % program, define the loops in the LoopData.xlsx file, then run this.
4
5     % Written by Daniel Gruber (2021) Augustine Lab, UC Davis
6
7     % Calculate block magnet B0?
8     B0_on = true;
9     % Have the option to use saved B0 field data
10    UseSavedB0Field = true;
11    % If a saved .mat file is to be used, set the path here.
12    MatPath = './BackedUpConfigurations/Balcom2MHzMagnetSimData.mat';
13
14    %---------------------------------------%
15    %%% NO USER CHANGES AFTER THIS POINT %%%
16    %---------------------------------------%
17
18    % If using saved data, assign the names and locations here
19    if UseSavedB0Field == true
```

Figure 3.7. The calculations to be performed in MagModel, as well as the path to saved data, are set in MagModel, lines 8, 10 and 12. The settings shown here were the settings used for the calculation of Figure 3.2.

Once these settings are defined, no other preparations are needed. The user may execute

the simulation by running the code in MatLab.

MagModel may also be used to model the magnetic field of a magnet block alone. To do this, simply leave the CoilData data lines blank and set `B0_on = true`. All other parameters are set as described above.

3.3.2 Limitations

In order to properly interpret the results of any model, one must understand the limits within which the model is valid. In the case of MagModel, the biggest and most obvious limitation is that the simulation is static, so no time-dependent effects are modeled. Importantly, this means that MagModel cannot not model eddy-current (shielding) effects or any rf propagation. Since NMR is conducted in the near-field regime, rf propagation is of little concern. However, the effects of eddy-currents are critical to design in many environments where conductive material is near the rf coil. In the example shown above in Figure 3.2, this would significantly change the resulting field profile and a powerful software such as a finite difference (FD) code or ANSYS should be used.

Another limitation of the MagModel code is that everything is calculated from finite elements on a structured mesh. This results in long calculation times and in some cases may lead to inaccuracies when the loop and block elements are not chosen to be small compared to the distances to observation points. Inaccurate field calculations can be mitigated by increasing coil and magnet block point densities, or by increasing the observation slice distance from these elements. The user may also simply choose to ignore results at points in the observation slice (or line) that are very close to field-generating elements. Because of the long computation times that will result, users may wish to use a different software suite if large magnet blocks with high point densities are unable to be avoided.

References

[1] D. Gruber (2021) MagModel (Version 1.0) [Source code].
    http://chemgroups.ucdavis.edu/~augustine/

[2] D. J. Griffiths, Introduction to Electrodynamics, 3rd Edition, Upper Saddle River, New
    Jersey: Prentice-Hall, Inc., 1999.

Chapter Four – Bucking Coils

A probe coil element for shaping B₁ field and mitigating the effects of eddy-currents in unilateral magnetic resonance

*"The important thing in science is not so much to obtain
new facts as to discover new ways of thinking about
them."*

-William Lawrence Bragg

4.1 Abstract

A novel radio frequency (rf) coil design strategy for mitigation of eddy-current shielding in unilateral magnetic resonance (UMR) is described. The coil design exploits full the standoff distance of the unilateral magnet by allowing placement of the rf coil directly on the surface of the magnet, This strategy fails for normal UMR surface coils due to eddy-current shielding by the conductive magnet surface. The approach includes the addition of coils orthogonal to standard loop or double-D coils in order to reduce the rf field at the conducting magnet surface.

4.2 Introduction

Nuclear magnetic resonance (NMR) has proven itself to be a powerful tool in chemical, biological, and physical analysis. However, NMR today typically requires incredibly strong homogeneous magnetic fields and can only be performed on very small samples. The technique is therefore not conducive to application in many field or processing environments, where portability is a necessity and where such strong fields would be dangerous. To address this, small, single-sided (or unilateral) magnets have been designed and constructed to allow for

homogeneous regions displaced outside the magnet for polarizing sample nuclei [1][2][3]. Coils

for use in this type of NMR are typically some version of surface loops or double-D (Figure-8)

coils and, depending on the magnet that these coils are applied to, these may be vulnerable to

self-interference due to eddy-currents induced by the coils in the magnet body.

4.2.1 Eddy-currents

Single-sided magnets are nearly universally constructed from rare earth magnets, which

are conductive. When a surface coil produces an rf magnetic field incident upon this conductive

surface, it introduces eddy-currents within the surface of the magnet. These eddy-currents

produce a reactive magnetic field opposing the initial incident rf field. For a typical surface coil,

this interferes destructively with the initial surface coil fields and drastically reduces the rf power

delivered to the sample, as well as the sensitivity of the coil.

Figure 4.1. Th mechanism of eddy-currents is depicted. A coil containing an rf current, I_{coil}, positioned above a conducting plane generates a magnetic field, B_1. In the plane of the conducting surface below, according to Faraday's Law of Induction, the incident B_1 induces a counter-rotating eddy-current, I_{react}, which then produces a reactive magnetic field, B_{react}. The overall magnetic field experienced by a parcel of sample is the sum of these two fields and is usually greatly reduced.

Eddy-currents can be vexingly complex to calculate, but the task is made much simpler if the problem is constrained to simpler terms. In this spirit, the eddy-currents produced by some incident flux can be understood more easily if the incident field is normal to the conducting surface under examination and whole region of eddy-current is deconstructed into concentric loops of individual eddy-currents, as is shown in Figure 4.1 [4]. In this conceptualizing of the problem, the flux generated by a single loop of eddy-current in a magnet surface, Φ_{gen}, is equal to the product of the self-inductance of the loop, L_{loop}, and the reactive current generated in the loop, I_{react}.

$$\Phi_{gen} = L_{loop} I_{react} \tag{4.1}$$

where, by Ohm's Law, the current is just the voltage of the loop, V_{loop}, divided by the resistance in the loop, R_{loop}, and the voltage is the negative time derivative of the total sum of flux, Φ_{tot}, within the area of the loop itself.

$$I_{react} = \frac{V_{loop}}{R_{loop}} \tag{4.2}$$

$$V_{loop} = -\frac{d\Phi_{tot}}{dt} \tag{4.3}$$

Recursively, the total flux is the sum of the generated flux and the imposed flux. As usual, the imposed flux is the imposed rf magnetic field, B_{imp}, over the area circumscribed by the loop of eddy-current, A_{loop}.

$$\Phi_{tot} = \Phi_{gen} + \Phi_{imp} = L_{loop}I_{react} + B_{imp}A_{loop} \qquad (4.4)$$

Since probe coils are driven with sinusoidal currents of frequency ω_0, the time derivatives of I_{react} and B_{imp} are simply $-i\omega_0 I$ and $-i\omega_0 B_{imp}$, respectively. Therefore, combining the above equations, and differentiating the terms of the total flux,

$$\Phi_{gen} = \frac{i\omega_0 L_{loop}}{R_{loop}}\left(\Phi_{gen} + B_{imp}A_{loop}\right) \qquad (4.5)$$

Solving this for the flux generated, then, gives

$$\Phi_{gen} = \frac{i\omega_0 L_{loop}}{R_{loop} - i\omega_0 L_{loop}} B_{imp}A_{loop} \qquad (4.6)$$

For good conductors, R_{loop} is very small and, since $B_{imp}A_{loop}$ is just the imposed flux, this says that the flux generated at the surface of a good conductor, due to eddy-currents, is equal in magnitude, and opposite in direction to the flux imposed on it. The same can be said of the magnetic field itself. Substituting $B_1 A_{loop}$ for Φ_{gen},

$$B_{gen} = \frac{i\omega_0 L_{loop}}{R_{loop} - i\omega_0 L_{loop}} B_{imp} \qquad (4.7)$$

Since the strength of a coil's magnetic field decreases exponentially with distance, this also means that the closer a source coil sits to a conducting surface, the closer the generated field comes to equaling the field of the source coil itself, and thus cancelling.

If a simple, round surface coil is placed directly on a magnet's surface, B_{imp} will be equal to the full strength of B_1 and, from Equation 4.7, B_{gen} will be essentially equal to $-B_1$. Thus, this configuration produces almost no rf field due to the interference of eddy-currents. Surface coils must therefore be lifted off the surface in order to overcome the deleterious effects of eddy-currents. Integrating the Biot-Savart equation for points along the center axis of a single loop coil, the field strength, B, at a distance, r, from the center can be found as

$$B = \frac{\mu_0 R^2 I}{2 \left(r^2 + R^2\right)^{3/2}} \tag{4.8}$$

where R is the radius of the coil loop and I is the current in the coil loop. From this equation, it can be seen that B will not start to drop significantly until the distance, r, becomes greater than the radius of the coil, R. To *start* effectively mitigating the effects of eddy-currents by displacing the coil from the surface therefore requires lifting the coil *at least* a coil radius.

This introduces a new problem, however, in that signal penetration into the sample is sacrificed (in a unilateral configuration, the sample cannot get any closer to the magnet than the coil itself). In this traditional geometry, the lesson is that more signal comes at the expense of penetration depth.

4.2.2 Bucking coils

Another way to mitigate the effects of eddy-currents is to shape the field such that the field imposed at the magnet surface, B_{imp}, is minimized. The idea is to add an array of coils to cancel or "buck" the field below the main coil loop to zero. This approach can also have the benefit of adding rf field in the sample region above the coil.

Bucking coils are mounted orthogonal to the orientation of the main coils and on their outside perimeter, as shown in Figure 4.2. The field produced by the whole coil assembly is then the sum of the bucking coils and the main coil. By orienting the bucking coils in a transverse fashion, the field produced above the coil assembly interferes constructively, adding to the overall strength of the B_1 rf field. Below the coil assembly, the fields interfere destructively, ideally to the point that the field below is canceled at the magnet surface. In the coil arrays described here, the bucking coils are wound in line with the primary coil so that they also serve as observe coils.
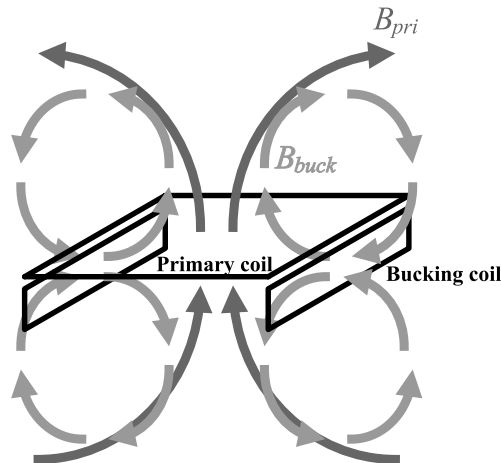
Figure 4.2. Bucking coils are positioned at the sides and orthogonal to the primary coil. The current in the bucking coils is oriented such that the field produced by the bucking coils (B_{buck}) supports the field of the primary coil (B_{pri}) above the assembly. Below the coil array, B_{buck} opposes B_{pri} so that the total field at the magnet surface is ideally zero.

Because the bucking coil field strength is stronger toward the outside of the coil array (where the bucking coils are located) and the primary field is strongest in the center, it is not possible to perfectly cancel the field at the magnet surface within the entire area of the coil array. A close approximation to this goal can be achieved by using computational software, such as ANSYS to design the array such that the sum of the flux within the area of the coil array is zero.

4.3 Materials and methods

For comparison, a traditional surface loop was compared to a bucked version of the same coil using ANSYS simulation software. Profiles for the rf magnetic fields of a simple, square surface coil and another square surface loop with bucking coils were calculated and plotted as a function of distance from the magnet surface. The simulation of the first coil was based on a 6 turn, 3.5 cm side length square coil. The inductance and resistance of this simple surface coil were $L = 2.98$ µH and $R = 0.2$ Ω, respectively. The second coil assembly simulated also included two 3.5 cm by 0.63 cm rectangular bucking coils positioned on the opposing sides of the primary square coil. The square surface coil in this case had 3 turns, while each rectangular bucking coil had 9 turns, yielding combined $L = 5.25$ µH and $R = 0.4$ Ω values.

4.4 Calculations and images

The bucking coil strategy attempts to mitigate the rf field profile changes introduced by the conducting plane, shown in Figure 4.1 for a circular loop, is to include additional coils that minimize the rf field impinging on the conducting surface. This concept is best investigated in practice by considering a square instead of a circular surface loop. If the system is oriented such that the permanent magnetic field of a three-magnet array defines the *z*-direction and the height

from the surface of the magnet defines the *y*-direction, then the surface of the magnet occupies an *xz*-plane. The field in the center of this square loop in the *xz*-plane, then, is oriented along the *y*-direction. The addition of two rectangular, counter-wound secondary loops on either side of, below, and orthogonal to the primary square loop in the *xy*-plane yields rectangular loop centered fields pointing in the positive and negative *z*-direction. It is the dipole nature of these loop fields that add to the field above the square loop and subtract from the field below the square loop. The reduced field below the two rectangular coil, bucked square coil complex concomitantly reduces eddy-current formation on the conducting surface. This geometry is shown above in Figure 4.2. The two square, counter-wound surface loops produce the primary field, B_{pri}, while the added orthogonal bucking coil yields B_{buck}. It should be clear from Figure 4.2 that the field amplitude above the bucked coil assembly is $B_1 = B_{pri} + B_{buck}$ while below it is $B_1 = B_{pri} - B_{buck}$. The size, number of turns, placement, and number of bucking coils are adjusted to minimize $B_{pri} - B_{buck}$, and thus eddy-current formation, at the conducting surface.

This simple idea was tested for a simple square loop and bucked square loop using the ANSYS modeling software and the results are provided in Figure 4.3. The two-dimensional rf field plots show the coil lying in the *xz*-plane as black lines and the magnitude of the magnetic field produced by the coils perpendicular to the *z*-direction in the *xy*-plane at $z = 0$, or the coil center. The rf field plots in the left and center columns respectively correspond to the coil residing in free space and having the top of the coil 6 mm above a flat conducting surface. The plot in the right column shows how the rf field behaves with standoff distance from the coil surface in the *y*-direction, with $x = z = 0$.

Figure 4.3. A simulation of coils is presented, both in free space and when placed atop a conducting surface. The left column qualitatively displays the coil in free space and the middle column shows the coil and field in the presence of a conducting surface. The right hand column shows the total effective field strength as a function of distance from the mounting surface (6 mm below the surface of the coil). The solid line represents the free space field strength and the dashed line shows the field strength when placed above of a conducting surface. The simple square loop (a) suffers significant degradation in effective field at all standoff distances. The bucked loop (b) suffers much less in the presence of a conducting surface – a 25 % reduction at a standoff distance of 2 cm, as compared to the 42 % reduction for the simple surface loop in (a).

As shown in Figure 4.3, introduction of the conducting surface clearly reduces the predicted field intensity provided by the coil as evidenced by the darkening of color in the plot in the center of Figure 4.3(a) in comparison to the picture in the left column. Again this purely numerical result is consistent with the conceptual prediction described above in the Introduction. The same consistent comparison is revealed in the reduced standoff performance as well. This is demonstrated by comparing the free space and shielded square coil results in the right column of Figure 4.3(a). The solid line corresponds to free space and the dashed line describes the

94

conductor-attenuated results. The vertical gray lines shown in the standoff distance plots in Figure 4.3 are a conservative estimate of the position of the homogeneous static magnetic field volume established by the unilateral three-magnet array. A comparison of the rf field gain or reduction at this position serves as one performance metric for conducting shield compensation as a function of coil geometry. In this first case of a square loop, the rf field 2 cm from the coil surface is reduced by 42% when the conducting plane is introduced. This rf field reduction is improved by almost a factor of two when two counter-wound bucking coils are introduced on opposing sides of the flat square coil as shown in Figure 4.3(b). The added bucking coils introduce two interesting features to the rf field profile. First, comparison of the first two columns in Figure 4.3(b) to (a) suggests that the rf field volume is higher in the sweet spot of the static magnetic field. Second, comparison of the standoff distance plots in the right column of Figures 4.3(a) and (b) suggest that the added bucking coils accomplish their goal, as the 2 cm standoff attenuation factor drops to 25 %, as compared to the 42 % reduction observed for the standard square surface coil.

The comparison of the performance of the standard square loop to the bucked version in Figure 4.3 strongly suggests that the core idea of reducing the rf field at the conducting surface by adding orthogonal, counter-wound coils improves the rf standoff distance by about a factor of two. Further improvements require either making the orthogonal bucking coils larger, adding more bucking coils or reducing the size of the primary square loop or DD coil to more effectively match the bucking and primary rf fields. Unfortunately, the addition of larger bucking coils underneath the primary square loop is not an option here, as that design change reduces the standoff distance because the coil complex will be thicker and further lifted from the unilateral magnet surface. Adding more bucking coils is also not an attractive option because the produced

95

fields become more complicated while not offering much improvement in metal-shielding compensation. Reducing the primary square loop coil size so that the existing bucking coils shown in Figure 4.3 more effectively match the primary loop field is also not desirable as the this results in reduced homogenous static magnetic field volume used, and thus smaller NMR signals and increased experiment time.

4.5 Conclusion

The advent of UMR magnets has provided access to non-standard sample geometries and environments that were otherwise impractical or dangerous with traditional high-field NMR. However, as large, displaced homogeneous static field regions have been pursued for the purpose of greater active sample volume and stronger signal, this has resulted in the use of large blocks of conductive magnetic material, which in turn support eddy-currents that are themselves detrimental to UMR signal. This has been avoided by using less magnetic material in the construction of the permanent magnet, which reduces field strength and thus signal, or by displacing the rf surface coil away from the surface of the permanent magnet, which sacrifices penetration depth.

This work proposes, and verifies by simulation, a solution that requires a sacrifice of neither magnetic material nor standoff distance. This is achieved through the addition of bucking coils to the traditional surface loop rf coil. Numerical simulations of the rf field provided by the square loop and bucked square coil suggest that eddy-current shielding can be mitigated by adding orthogonal bucking coils to the primary coil. Another alternative not explored here would be to consider an array of small, bucked, primary coils. Although the rf field profile off the

surface of the coil is complicated, the approach should yield useful, measurable NMR signals with an experimentally practical standoff.

Simulations of the bucked square, as compared to the traditional surface square loop yielded encouraging results. The simulated loss in B_1 strength for the simple surface loop was significant and calculated to be 42%, which comports well with experience and conceptual understanding of current loops above conducting planes. By comparison, the simulation of a bucked version of the same loop showed only a 25% reduction in field strength. These results may be reasonably expected to apply to double-D style loops, or any geometry that employs surface loops above a conducting magnet body.

References

[1] G. Eidmann, R. Savelsberg, P. Blümler, B. Blümich. The NMR MOUSE, a mobile universal surface explorer. J. Magn. Reson., Series A. 122 (1996) pp. 104-109. 10.1006/jmra.1996.0185.

[2] S. Utsuzawa, E. Fukushima. Unilateral NMR with a barrel magnet. J. Magn. Reson., 282 (2017) pp. 104-113 10.1016/j.jmr.2017.07.006.

[3] A. E. Marble, I. Mastikhin, B. G. Colpitts, B. J. Balcom. A Compact Permanent Magnet Array with a Remote Homogeneous Field. J. Magn. Reson., 186 (2007), pp. 100-104. 10.1016/j.jmr.2007.01.020.

[4] M. S. Conradi, A. P. Zens. RF shielding and eddy currents in NMR probes. J. Magn. Reson., 305 (2019), pp. 180–184. 10.1016/j.jmr.2019.06.011

Chapter Five – Fringe Coils

A surface probe coil design that harnesses the effects of eddy-currents in unilateral magnetic resonance for enhanced signal and stable tuning

"Living is worthwhile if one can contribute in some small way to this endless chain of progress."

-Paul A.M. Dirac

5.1 Forward

The following is a reproduction of a manuscript submitted to the Journal of Magnetic Resonance in 2021. The work here represents my own efforts with help from Sophia Fricke and Vanessa Lee, in collaboration with Professor Bruce Balcom at University of New Brunswick, and under the advisement of Professor Matthew Augustine. It was, in fact, during a seminar by Dr. Balcom when he visited UC Davis that I found the inspiration for bucking coils. At the time, we were trying to figure out a way to get around the deleterious effects of eddy-currents formed by surface loop probe coils in single-sided NMR. At some point, Dr. Balcom mentioned Halbach magnets in his talk and it occurred to me that we could use similar interplay of magnetic fields to reduce the field of a rf probe coil interacting with a conductive magnet body, while enhancing the field produced by the same coil in the sample volume. The result was initially bucking coils, as described in the previous chapter. Ultimately, experimentation with variations of bucking coil designs led to the flat fringe coils described in this chapter..

5.2 Abstract

Two types of radio frequency (rf) coils for linear array, unilateral magnets are described.

These coils are designed to fully exploit the standoff distance of the unilateral magnet by

placement directly on the magnet surface. This approach fails for normal surface coils used for

magnetic resonance due to eddy-current induced shielding by the conductive magnet surface.

The coil design strategy includes a rectangular cross section solenoid coil, either continuous or

split in the center, mounted with the center axis of the coil parallel to the magnet surface. These

geometries, when placed on a conducting surface, enhance the rf field produced in the sample

region, outside of the solenoid coil. The spatial homogeneity of both rf coils is characterized

using the ANSYS finite element modelling software. ANSYS modelled coil geometries led to

homogeneous, surface displaced rf fields. These coils were then constructed and characterized

with magnetic resonance imaging. Finally, two experiments that use these coils to perform large

standoff relaxation measurements are described.

5.3 Introduction

While nuclear magnetic resonance (NMR) has proven to be a powerful technique in

chemical analysis, traditional hardware typically uses high radio frequency (rf) power and

presents inflexible sample geometries that are not conducive to the analysis commonly

encountered in manufacturing, processing, or field work environments. Magnets have been

previously designed and constructed to conduct NMR outside the body of the magnet itself [1]

and some have further extended this unilateral design to produce regions of relatively

homogeneous magnetic field, well-displaced from the magnet surface that can be used for single-

sided NMR [2]. These permanent magnets are much smaller than traditional superconducting or

electro-magnets allowing the NMR instrument to be portable. However, as the magnetic field is weak and the sensors are typically deployed in dynamic, often electrically noisy environments, measurements typically have low signal-to-noise (S/N). A major impediment to unilateral magnetic resonance in connection with experiments requiring large standoff, or distance between the magnet and sample, is that the magnet interferes with the applied oscillatory rf magnetic field, B_1, by supporting eddy-currents. A surface coil may be lifted off of the magnet in order to mitigate eddy-currents, however this comes at the expense of sample standoff distance from the magnet and probe assembly surface.

In the design of a unilateral magnet-based NMR sensor, the sources of the static magnetic field, B_0, and the rf field, B_1, must be considered. Several magnets have been designed to address the challenge of producing a homogeneous B_0 field displaced from a surface. One of the most well-known and commercially available single-sided sensors is the NMR-MOUSE [1]. This configuration consists of two opposed-polarity magnets that have been separated by just enough space to allow for a probe coil. The B_0 field produced by the NMR-MOUSE is transverse to the surface, while the B_1 produced by the probe coil is perpendicular. This design suffers very little from eddy-currents in the magnet body and achieves a ^{1}H Larmor frequency of about 17 MHz, however the useful penetration depth is only a couple of millimeters.

Another attractive magnet design is the cylindrical magnet array [2]. This "barrel" magnet involves several stacked ring magnets and a cylindrical bar magnet placed in the center of the ring stack. The position of the cylindrical magnet is adjustable and used to establish the homogeneous region of the B_0 field. The static field in this configuration is perpendicular to the surface, while a figure-eight or "double D" (DD) style probe coil placed on top of the magnet provides a B_1 field transverse to the surface. This design has resulted in a homogeneous sensitive

region ca. 1.5 cm above the surface, with a [1]H Larmor frequency of ca. 5 MHz. Again, this geometry does not suffer from eddy-current formation as the majority of the DD coil is above the solid cylinder bar magnet that is depressed in the magnet rings.

A linear magnet array was created based on the idea of a solid, sideways bar magnet [3]. Here, a bar magnet is placed on its side, so that the static magnetic field is transverse to the surface. That bar magnet is then split into left and right halves and separated, leaving an inverted saddle region in the B_0 field above the surface. Then, a smaller bar magnet, aligned in the same direction as the first two pieces, is inserted between the halves, and its position may be adjusted for fine tuning the distance of the homogeneous saddle region from the surface, similar to the cylindrical magnet array described above. Unlike the magnets previously described, in which the magnetic material is separated and pushed away from the sample region, this design allows for all of the magnetic material to be collocated as one solid volume next to the sample region. This results in a stronger and larger homogenous static magnetic field volume while using less material. A further advantage of this design is offered by the fact that it is safer because the individual magnet pieces are aligned in a low-energy arrangement and are not being forced together by a frame, which would explode, should it fail. The linear magnet array has produced magnetic fields with similar characteristics as the cylindrical magnet array, however with a much simpler design that uses approximately 20% less magnetic material, making this design significantly lighter and less costly. Additionally, the solid nature of the construction of these magnets makes them less susceptible to temperature-induced field drift [4].

The linear magnet arrays are extremely attractive for industrial applications where portable sensors are needed to study materials well separated from the magnet surface by packaging, pipes and other sealed containers. However, the same feature that gives these

magnets their desirable qualities, namely that they are essentially solid blocks of magnetic

material, creates a significant problem for probe coil design when applications requiring the large

standoff that these magnets offer are encountered.

Figure 5.1. A simplified sketch showing the relative placement of the linear
unilateral magnet array, NMR detection coil and homogeneous static magnetic
field sample volume. The homogenous static magnetic field volume is indicated
by the axis system intersecting the single B_0 flux line that defines the z-direction.
The h-direction corresponds to displacement normal to the unilateral magnet
surface while the w-direction assigns the remaining axis along the width of the
magnet. The sample standoff distance from the magnet is $h_1 + h_2$ while the
standoff distance from the sample to the NMR detection coil is h_2. The direction
of the rf fields provided by a loop or DD coil are also included. This work
endeavors to develop NMR detection coils that have the distance h_1 minimized
to maximize the distance h_2 and thus improve sample standoff distance.

Consider the usual linear unilateral magnet geometry with either a single circular loop or

DD coil as shown in Figure 5.1. In this figure, and throughout the remainder of this manuscript,

the $\hat{z}$, $\hat{h}$ and $\hat{w}$ unit vectors respectively describe the direction of the static magnetic field in the

homogeneous volume, the displacement off of the magnet surface and the final orthogonal

direction across the width of the magnet. With this convention, the coil is displaced from the

magnet surface by the distance h_1 and the coil is separated from the homogeneous sample

volume by the distance h_2. A single magnetic flux line for the unilateral magnet is shown and the direction of the rf field offered by the loop or DD coil is indicated at the position of the homogeneous sample volume. It should be clear from Figure 5.1 that h_1 must be minimized in order to fully utilize the standoff and large homogeneous sample volume produced by the unilateral magnet. It is this requirement that introduces complications to the performance of the loop and DD coils typically used for physically localized imaging and single-sided NMR. Figure 5.2(a) describes what happens when a loop in the z-w plane producing an rf magnetic field is placed close to a conducting surface also in the z-w plane. Here an applied rf current, I, in the loop produces a primary rf magnetic field B_1 that induces eddy-currents in the conductor I_{eddy}. These eddy-currents produce a secondary magnetic field B_{eddy} that reduces the applied B_1 primary field. This effect diminishes both the size and homogeneity of the applied primary field. An appreciation for the severity of this effect at 5 MHz can be obtained from the magnetic field plots in Figures 5.2(b) and (c) for a simple current loop of radius R = 2 cm held above a flat copper surface [5]. The solid lines in both plots confirm expectations in the absence of a conducting surface where h_1 > 4 cm. The plots in Figure 5.2(b) as a function of distance from the center of the coil, r, calculated at h_2 = 1 cm or at half the loop radius above the loop recovers the expected field homogeneity. Both the amplitude and the radial homogeneity are spoiled when the conducting surface is moved closer to a practically useful h_1 = 2 mm displacement as shown by the dashed line in Figure 5.2(b). A similar performance reduction is also observed for the on axis field at r = 0 as a function of displacement h_2 from the surface of the coil. Again, the reduced field profile for the h_1 = 2 mm offset is shown as the dashed line. Although just the magnetic field component normal to the loop and conducting surface in the $\hat{h}$ direction is shown Figure

5.2, it should be clear that the radial component of the field suffers from the same eddy-current induced reduction and that these effects are just as severe for the DD coil.
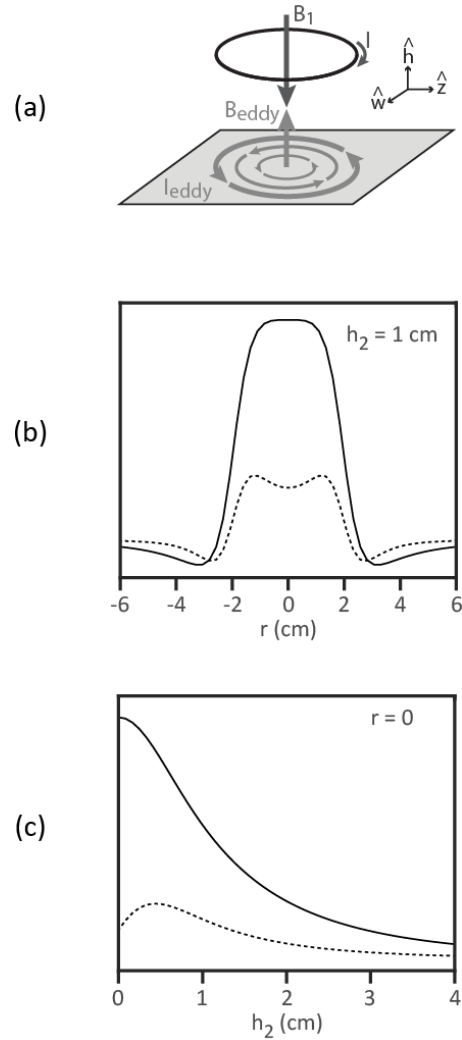
Figure 5.2. A graphical model for the generation of the eddy-current-reduced rf field is shown in (a). Here the primary B_1 field from a circular loop induces eddy-currents Ieddy on the conductor surface that concomitantly produce a field Beddy that subtracts from B_1. The plot in (b) describes the rf field strength of a loop in free space, calculated at a distance $h_2 = 1$ cm from the $R = 2$ cm circular loop coil surface, as a function of the radial distance from the center of the loop, r, shown as the solid line. The dashed line in (b) shows what happens when the conducting surface is moved closer to $h_1 = 2$ mm. The plots in (c) show the standoff distance dependence of the rf field developed by the same circular loop coil at the center $r = 0$ as a function of h_2. Again the free space rf field with $h_1 = 4$ cm is shown as the solid line while moving the conductor closer to $h_1 = 2$ mm produces the dashed line. These plots were generated from the analytical results of Moser [5].

A seemingly obvious solution to this problem that does not work simply involves placing a layer of ferrite magnet metal between the magnet array and the coil. The resistivity of ferrite is much greater than that of the neodymium alloys comprising the magnet. This should mitigate eddy-currents without significantly reducing the B_0 field strength. This idea, however, neglects the fact that resistivity is magnetic field dependent [14, 15]. Indeed, the beneficial conductive properties of ferrite experienced at zero field vanish at NMR-useful magnetic field strengths and frequencies. Under these conditions, the resistivity of ferrite more closely resembles that of copper and aluminum. These effects were briefly tested and confirmed, and the use of ferrite as a mitigation strategy was therefore abandoned.

Two other ways of circumventing the negative effects introduced by the conducting magnet surface on the primary field of a surface coil are also immediately apparent. One approach, which was mentioned above, is to sacrifice unilateral magnet penetration depth by using simple loop or DD coils displaced by $h_1 > R$ from the magnet surface. In this configuration, eddy-currents are minimized and changes to the applied primary field amplitude and homogeneity are reduced. This work instead considers rethinking the geometry of the coil itself. The approach allows h_1 to be minimized, with the goal of placing the rf coil directly on the magnet surface. In this way, the sample standoff distance of the unilateral magnet is fully exploited. The loops of these coils are arranged orthogonally to the surface to the magnet, producing primary magnetic flux lines that do not impinge the magnet surface. The term "primary" is used here to refer to field lines produced directly by the coil and which intersect the sample region. This work describes and mathematically, as well as experimentally, characterizes these "fringe" coils in detail with the goal of identifying the best coil geometry for accomplishing large standoff NMR with linear unilateral magnets.

5.4 Materials and methods

The rf magnetic field profiles shown in Figures 5.2(b) and (c) were calculated with Matlab™ using the analytical results presented in Moser [5]. The asymmetry of the fringe coils and conducting plane encountered in this study, and the complexity that they introduce into the Maxwell equations required that the ANSYS AIM (Cannonsburg, PA) finite element, three dimensional differential equation solver be used for theoretical coil performance comparisons.

Three coils were constructed to experimentally test the numerical model predictions by magnetic resonance imaging the rf excitation field above the coils with and without a conducting plane. The first coil is a 6 turn, 3.5 cm side length square coil. The inductance and resistance of this simple surface coil are L = 2.98 µH and R = 0.2 Ω. The L = 9.38 µH, R = 0.4 Ω standard "fringe" coil used here involved wrapping 33 turns onto a 2 cm long, 3.5 cm wide, 0.63 cm thick support. To be clear, the coil is wrapped around the "long" or "longitudinal" axis and the width of the coil is typically greater than the length. The coil wraps were continuous along the 2 cm long axis so that the final coil resembles a smashed solenoid coil. Finally, a "split fringe" coil, involving two counter-wound 13 turn, 1.8 cm long, 3.5 cm wide, 0.63 cm thick fringe coils separated on the long axis by 1.5 cm was constructed with L = 4.10 µH and R = 0.4 Ω.

Mineral oil and laundry detergent samples were obtained from Safeway and used as received. All NMR data were obtained with a Tecmag (Houston, TX) Redstone spectrometer and two separate magnets were used in this study. Magnetic resonance images of the rf fields produced by the four model surface coils were obtained at a 4 MHz [1]H Larmor frequency with a 0.09 T SMIS imaging electromagnet. All of the images were created using a two rf pulse spin echo, pulse sequence, with frequency encoding along the z direction and phase encoding along the x and y directions. In all cases the pulsed field gradients were adjusted to provide 128 x 48 x

48 points in the z, x and y directions ($\hat{z}$, $\hat{h}$ and $\hat{w}$) and an associated 12.8 cm x 12.8 cm x 12.8 cm field of view. The pole face gap in the SMIS magnet is large enough to allow a 4 inch diameter, 800 mL beaker of mineral oil to be placed directly on top of the fringe coils in the homogeneous region of the static magnetic field. In this way, simply sliding an aluminum plate, cut to the surface size and shape of a 3-magnet linear array, under the coil and beaker complex emulates the surface of a unilateral magnet. The portable unilateral magnet used in this study labeled as WA-1 was built by Balcom, et. al. at the University of New Brunswick [6]. It is a 18 cm x 11 cm x 7 cm, 0.05 T unilateral three permanent magnet array with a homogeneous sweet spot approximately 2.5 cm off the magnet surface that provides a 2.27 MHz ^{1}H Larmor frequency. This magnet along with the coil was sealed into a copper shielded, waterproof Pelican™ case so that the sensor can be used on a benchtop or submerged into a liquid sample. All observed transverse relaxation time constants, T_2^{obs}, were obtained the Carr-Purcell-Meiboom-Gill (CPMG) pulse sequence [7, 8]. Typical $\pi/2$ rf pulse and echo delay times were 18 µs and 1.046 ms respectively in the CPMG pulse sequence.

5.5 Results and discussion

5.5.1 Calculations and images

The coil design strategy described here considers inductors on metal surfaces that yield primary flux lines that do not intersect the conducting surface. The solenoid coil shown in Figure 5.3 satisfies this requirement. The two flux lines shown in black and labelled as B_1 describe the primary field of the solenoid coil. The flux lines above the center axis of the solenoid in the $\hat{w}$ direction should be immune to the conductor shielding as these flux lines do not intersect the unilateral magnet surface. The flux lines below the center axis of the solenoid do intersect the

conducting surface, however, because of the coil-to-conducting surface geometry, these flux lines generate eddy-currents I_{eddy} that have the surprisingly advantageous effect of generating a field B_{eddy} that adds to the field B_1 above the solenoid coil in the measurement region. The standard solenoid coil with circular cross section is not desirable for this application as the standoff distance is not maximized. To take advantage of the rf amplification offered by the conducting surface suggested by Figure 5.3 while maximizing standoff distance h_2, the solenoid coil is flattened to create an ellipsoidal or rectangular cross section solenoid coil. This smashed solenoid coil is referred to as a "fringe coil" to distinguish it from a standard circular cross section solenoid coil. The nomenclature was also chosen because the active volume is outside of the coil in the fringe case while for a standard solenoid coil in NMR the active volume is inside of the coil.

Figure 5.3. A sketch describing how the rf field from a solenoid coil placed sideways on a conducting surface increases the rf field in the homogeneous static magnetic field volume above the coil. The solid black flux line describes the primary field of the solenoid coil while the flux line for the eddy-current I_{eddy} generated rf field B_{eddy} enhances the rf field in the sample region. Again, the static magnetic field is in the z-direction.

The plots shown in Figure 5.4(a) describe the performance of a 1 cm long, 3.5 cm wide, 0.5 cm thick fringe coil with its center axis in the $\hat{w}$ direction in free space in the left column and above the conductor in the center column. The top row in Figure 5.4(a) shows the rf field profile in the h-w plane at z = 0 while the bottom row displays the rf field profile in the h-z plane at w = 0. These plots suggest that the rf field above the fringe coil is not attenuated by the conducting surface. Rather, as predicted by the standoff distance summary in the right column of Figure 5.4(a) and by Figure 5.3, the rf field in the measured sample volume at $h_2 = 2$ cm from the unilateral magnet surface increases by 43 %.

Figure 5.4. Summary of ANSYS developed numerical rf field calculations for a fringe and split fringe coil in (a) and (b) respectively. The left and center column show the rf fields perpendicular to the z-direction without and with a conducting surface placed $h_1 = 6$ mm below the top of the rf coil complex. The top and bottom rows in (a) and (b) respectively correspond to the rf field calculated in the h-w and h-z planes. The right column shows the standoff performance of the coils obtained in the h-direction at the $z = w = 0$ center of the rf coil. The solid and dashed lines show the rf field value without and with the conducting surface as a function of distance h_2 from the coil surface. The vertical gray line shows the position of the homogenous static magnetic field volume where the shielding performance is evaluated.

The standard fringe coil is designed to obtain signals directly above and outside of the coil. However, both common knowledge and the calculations shown in the top row of Figure 5.4(a) suggest that the rf field at either end of the fringe coil near the rectangular openings is

more intense than the rf field outside and above the fringe coil. It is for this reason that the opposed, counter-wound fringe coil was created. The performance of this "split fringe" coil is shown in Figure 5.4(b). The two 1.8 cm long, 3.5 cm wide, 0.63 cm thick counter-wound coils are separated by 1.5 cm. The top plots in Figure 5.4(b) show the rf field profile in the h-w plane at the $z = 0$ center of the split fringe coil while the lower plots display the rf field profile in the h-z plane at the $w = 0$ center of the split fringe coil. Finally, the plot in the far right column shows the rf standoff performance of the fringe coil. Once again, the signal at $h_2 = 2$ cm is increased, this time by about 44 %.

The theoretical predictions made in Figure 5.4 suggest that the response of the fringe and split fringe coils to the unilateral magnet conducting surface is amplification of the field and signal, rather than attenuation, as in the case for standard surface loops and double-D coils. An experimental comparison of the performance of these two fringe coils to a standard square loop is shown in Figure 5.5. Here, the images in the left column are without a conductor while an aluminum plate is placed 6 mm below the top of the coil in the right hand column. The number included in the lower right corner of the right column image is the ratio of the maximum signal obtained with the conductor in place to that same maximum signal in free space. The crescent shape in all images comes from the fact that the B_1 field is stronger near the coil center and weaker at the edge, consistent with Figure 5.2. The square loop performs as expected, providing a reduction of about a factor of two upon placement on the conducting shield. More important is the gain in rf field amplitude for the fringe and split fringe coils predicted in Figure 5.4 and experimentally observed in Figure 5. Moreover, these results indicate that the standoff distance of the fringe and split fringe coils improves slightly upon adding the conducting surface underneath the coil. Again, this is anticipated by the predictions made in Figures 5.3 and 5.4.

Figure 5.5. Magnetic resonance two-dimensional slices, extracted from 3D magnetic resonance images obtained at a 4 MHz ^{1}H Larmor frequency for a large beaker of mineral oil placed on the surface of a similarly sized square loop coil, a fringe coil and a split fringe coil in (a) – (c) respectively. The images in the right column include an aluminum plate placed h_1 = 6 mm below the top of the rf coil while the images in the left column have the aluminum plate removed. The number inset on the lower right corner of the images in the right column is the ratio of the maximum signals obtained with and without the conducting plate.

Another interesting, unexpected feature of these metal surface mounted fringe coils is that they perform as well or in some cases better than a simpler loop coil in free space. For example, the images shown in Figure 5.5 can be used to define a S/N. The image for the free space coil in the left column in Figure 5.5(a) yields S/N = 70. A S/N = 56 and 97 are obtained from the images for the metal mounted fringe and split fringe coils in the right column of Figure 5.5(b) and (c), respectively. It may be unexpected that the split fringe could surpass the "ideal" case of a loop in free space. This can be explained by recalling that in all of these experiments, the sample is placed outside of the coil, and thus the geometry of the setup is not the ideal case of

a sample inside a coil. In unilateral magnetic resonance, the sample is by nature outside of the coil, where the simple loop is not as attractive as fringe coils, that harness the conducting properties of the magnet body.

5.5.2 Unilateral magnet-based characterization

The standard fringe and the split fringe coils were further explored by conducting two additional experiments. Both experiments use the 0.05 T (WA-1, 2.27 MHz [1]H Larmor Frequency) unilateral magnet and the same fringe coils that were used to obtain the rf excitation images in Figure 5.5(c). The 3.5 cm width of these fringe coils was chosen so that they could be inserted easily into the 3.6 cm rectangular groove created by the depressed central bar magnet in the WA-1 three-magnet linear array.

In the first set of characterization experiments, a large beaker of mineral oil was placed on top of the magnet and coil, and the CPMG pulse sequence was used to measure the observed transverse relaxation time, T_2^{obs}, as a function of beaker-to-coil distance (0.6 cm > h_2 > 2.1 cm). This experiment was repeated for both coils on the same WA-1 magnet. The average and standard deviation of the T_2^{obs} values obtained for beaker displacements are shown in Figure 5.6. When S/N is low in the CPMG experiment, echo signals quickly drop below the noise floor, resulting in the appearance of a shortened decay time constant. This, along with the observed fact that S/N significantly decreases with greater displacement from the magnet and coil, means that the T_2^{obs} values appeared shorter, with larger displacement for a constant number of scans. As expected, however, T_2^{obs} increased and stabilized once a sufficiently high S/N was achieved with additional signal averaging. Figure 5.6(a) demonstrates the flat fringe coil performs well at short standoff distances, but as the sample starts to approach the edge of the magnet sweet spot, the

T_2^{obs} uncertainty increases until noise completely washes out the decay signal. Figure 5.6(b) displays the T_2^{obs} values obtained in exactly the same way, but with the split fringe coil. A comparison of Figure 5.6(a) to (b) suggests that the split fringe coil performs better than the standard fringe coil as roughly the same average T_2^{obs} value is obtained at each displacement, the error bars in the split fringe case are smaller, thus the T_2^{obs} value is more certain and, finally, the constant can be measured at the $h_2 = 2.1$ cm displacement.

Figure 5.6. Summary of T_2^{obs} values obtained for light mineral oil with the standard fringe (a) and split fringe (b) coils as a function of beaker-to-magnet displacement, h_2. The T_2^{obs} values were obtained from a Matlab™ fit of the decay of the raw CPMG echoes to a single exponential function. In both plots, the circle captures the average value of T_2^{obs} obtained at h_2, while the bars capture the error in this average. As the value of h_2 increases, the S/N drops for a constant number of signal averages, and the error in T_2^{obs} increases. Note that at $h_2 = 2.1$ cm, the S/N for the standard fringe coil transient was too low to obtain a T_2^{obs} estimate.

The second set of characterization experiments considered the strength of the B_1 rf field produced by the two fringe coils as determined from a nutation experiment. Using light mineral

oil as the sample, the spin echo intensity is tracked as a function of rf pulse length, at a fixed applied rf amplitude. The solid lines shown in Figures 5.7(a) and (b), respectively, for the fringe and split fringe coils report the average imaginary, real, and magnitude spin echo intensity, from light to dark respectively, over five experiments for separate rf pulse lengths. The shaded regions communicate the 95% confidence level uncertainty in the average spin echo intensity at each rf pulse length.

In a standard nutation experiment, the amplitude of the free induction signal is recorded as a function of pulse length. The maximum signal occurs when the magnetization, as nutated by $\pi/2$, rotates completely into the transverse plane. Here, in the single-sided case where the collection of free induction decay signals are impractical, two pulses are instead applied, producing the anticipated spin echo. The nutation of the intensity of this spin echo behaves just like the signal following just one pulse, specifically the maximum spin echo occurs for a $\pi/2$ initial pulse. This effect is demonstrated in Figure 5.7 where the maximum signal is at approximately 35 microseconds, giving a Rabi frequency of 7.1 kHz. The fact that the Rabi cycling does not continue after the first cycle is indicative of the inhomogeneous nature of both the B_1 and B_0 fields. The spins are actively dephasing during the pulse itself. There is also a phase shift introduced by the metal surface, manifested as a non-zero imaginary component when the echoes are integrated. This is due to phase shifts in the reactive B_1 rf field at the reflection plane of the magnet surface, which then sum with the primary B_1 rf field at the sample. The same effect occurs during signal detection, and is consistent with NMR observations in the presence of metal shielding [9].

Figure 5.7. Spin echo intensity as a function of rf pulse length. Here, the ordinate corresponds to the length of the first rf pulse in the spin echo experiment. This ordinate value is half the length of the second rf pulse in the spin echo experiment. The abscissa captures the total spin echo signal intensity. This value was obtained by integrating the raw spin echo signal in the time domain. Five spin echoes were obtained at each rf pulse length observed. The thick solid lines correspond to the average signal intensity, while the shaded regions communicate the 95% confidence in this measured spin echo intensity. The progressively darker shading corresponds to the real, imaginary, and absolute value of the signal intensity. Results are shown for the standard fringe and split fringe coil in (a) and (b), respectively.

5.5.3 Unilateral magnet applications

The superb performance of the split fringe coil in all characterization efforts prompted two challenging real-world applications. In the first application, the coil and magnet assembly in a copper lined, sealed Pelican™ case were submerged in a 5-gallon bucket of Gain™ detergent. This setup was chosen to emulate observation in a large format vat mixer or other industrial

setting. The photographs in Figure 5.8(a) show the unilateral magnet and coil in the open case (left), as well as submerged in Gain™ (upper-right). Application of the CPMG pulse sequence in this setup produced the transient decay signal on the left in (b), with a $T_2{}^{obs} = 16$ ms decay constant determined from the inverse Laplace transform (ILT) shown on the right. The reported signal represents the sum of 256 transients, in 2.8 minutes of data acquisition, at a S/N = 148.

Figure 5.8. Two real examples of the split fringe coil in action. The split fringe coil inserts into the existing groove of WA-1 unilateral magnet and the split fringe coil/unilateral magnet complex is placed into a water tight Pelican™ box lined with copper as shown in the picture on the left in (a). The enclosed NMR sensor was submerged in a bucket of laundry detergent, as shown in the upper right in (a), and application of the CPMG pulse sequence produced the decay shown on the left in (b), with a $T_2^{obs} = 16$ ms time constant obtained from the ILT on the far right. The lower right in (a) shows the placement of a branch segment on the top of the sealed magnet. Application of the CPMG pulse sequence yielded the transient decay shown in the left of (c). In this case, the ILT reports two T_2^{obs} values, 3.5 ms and 20 ms that correspond to water in the cells and lumen of the bulk wood, respectively.

The second application considers the in-situ exploration of natural products or crops. Unilateral magnetic resonance has previously been used to measure the moisture content of wood, using a similar magnet [10, 11]. A small piece of tree branch, approximately 1 inch in diameter, was obtained and sealed in a plastic bag between experiments to prevent drying. This wood was then placed over the sweet spot of the Pelican™ case-sealed WA-1 magnet and split fringe coil, as shown in the lower-right of Figure 5.8(a). Application of the CPMG pulse sequence in this case led to the transient decay signal shown on the left of (c). In this example, 1,735 were averages accumulated in 10 minutes of data acquisition, yielding a S/N of 1,047. The ILT reports two time constants, $T_2{}^{obs} = 3.5$ ms and 20 ms, shown in the plot on the right of (c). These two constants correspond to the water in the wood cells and lumen, respectively, and are in agreement with previous observations carried out with various wood types [12, 13]. Such a result illustrates the potential for this coil and magnet to be used in similar investigations, and is remarkable for such a small sample, that does not fully exploit the magnet's sweet spot volume, as compared to the logs that were used in the cited studies, or when completely immersing the magnet and coil in sample.

5.6 Conclusion

The availability of unilateral magnets has broadened the types of problems that can be studied with NMR spectroscopy and relaxometry. However, the primary challenge of single-sided NMR has been designing magnets that provide large homogeneous static magnetic field volumes significantly displaced $h_1 + h_2 > 2$ cm from the magnet surface while ensuring that the field strength is large enough to provide signals with minimal signal averaging. The linear unilateral magnet arrays mentioned above meet this challenge, as the magnet material is located

as close to the sample as possible. The consequence of this magnet choice is that the full standoff distance offered by the linear array is not able to be fully exploited as the use of standard surface rf coils themselves require an offset from the magnet surface due to eddy-current formation.

This work specifically addresses the eddy-current shielding problem by proposing the use of fringe coil geometries, essentially flattened solenoids, in which the signal is obtained from outside of the rf coil, in the "fringe" region. The interesting feature of this family of coils is that the conducting surface enhances the rf field at the homogenous sample volume. Two geometries of fringe coils, a standard "fringe" coil and an opposed "split fringe" coil, were constructed and two characterization experiments were performed to compare them. These characterizations involved sequential experiments with the sample successively spaced farther from the magnet surface, as well as a nutation experiment, which indirectly measured the rf power in the sample. In this way, the maximum sample standoff, T_2^{obs} values and Rabi cycling times were obtained at an $h_2 = 2$ cm standoff distance for a beaker of light mineral oil. The fringe coil geometries were predicted and observed to perform better than the traditional surface coil geometries. Moreover, in all characterization experiments, the split fringe coil outperformed the flat fringe coil. Further experiments with real-world applications were also conducted using the better performing split fringe coil. These experiments involved the immersion of the magnet and coil assemblies into laundry detergent, as well as the investigation of a small diameter wooden branch placed on top of the assembly. The split fringe coil was able to easily perform experiments in both environments, even discerning multiple relaxation components in the branch segment.

All experiments yielded positive results, typically with just over two minutes of signal averaging, which is more than adequate for field use. The split fringe coil is the simplest coil array that largely matches the spatial profiles of B_1 to B_0 within the sensitive region of the

magnet. Adding more fringe coils to make a cross or star pattern involving fringe coil building blocks of varying size will likely match the B_1 and B_0 fields better, and lead to greater standoff distance and larger signals. However, for the magnets used in this study, this would preclude the advantage of utilizing the depressed notch in the center of the magnet array, leading to slightly reduced useable standoff distance. Still, it is these larger signals that reduce experiment time, an important parameter when the rapid throughput measurements commonly encountered in industrial problems are required. Future work will involve exploring the multiple split fringe coil geometries Another future pursuit could involve engineering the surface geometry of the magnet to better exploit the eddy-current-induced field magnification effects – essentially making near-field rf reflective "lenses" that the rf field intensity and tailor its shape to produce even better NMR signal generation and reception.

References

[1] G. Eidmann, R. Savelsberg, P. Blümler, B. Blümich. The NMR MOUSE, a mobile universal surface explorer. J. Magn. Reson., Series A. 122 (1996) pp. 104-109. 10.1006/jmra.1996.0185.

[2] S. Utsuzawa, E. Fukushima. Unilateral NMR with a barrel magnet. J. Magn. Reson., 282 (2017) pp. 104-113 10.1016/j.jmr.2017.07.006.

[3] A. E. Marble, I. Mastikhin, B. G. Colpitts, B. J. Balcom. A Compact Permanent Magnet Array with a Remote Homogeneous Field. J. Magn. Reson., 186 (2007), pp. 100-104. 10.1016/j.jmr.2007.01.020.

[4] M. Martin, B. Balcom, M. McCarthy, M. Augustine. (2019). Noninvasive, Nondestructive Measurement of Tomato Concentrate Spoilage in Large-Volume Aseptic Packages. J. Food Sci., 84 (10) (2019), pp. 2898-2906

[5] J. R. Moser. Low-Frequency Shielding of a Circular Loop Electromagnetic field Source. IEEE Transactions on Electromagnetic Compatibility, 9 (1) (1967), pp. 6-18.

[6] J. C. Garcia-Naranjo, P. Guo , F. Marica, G. Liao, B. J. Balcom. Magnetic Resonance Core-Plug Analysis with the Three-Magnet Array Unilateral Magnet. Petrophysics, 55 (2014) pp. 229–239.

[7] H. Y. Carr and E.M. Purcell. Effects of diffusion on free precession in nuclear magnetic resonance experiments. Phys. Rev., 94 (1954), pp. 630-638.

[8] S. Meiboom and D. Gill. Modified spin-echo method for measuring nuclear relaxation times. Rev. Sci. Instrum., 29 (1958), pp. 688-691.

[9] M. Mohammadi, A. Jerschow. In situ and operando magnetic resonance imaging of electrochemical cells: A perspective. J. Magn. Reson., 308 (2019), pp 106600.

[10] C. Lamason, B. MacMillan, B. Balcom, B. Leblon, Z. Pirouz. Water content measurement in black spruce and aspen sapwood with benchtop and portable magnetic resonance devices. Wood Material Sci. and Eng., 10 (1) (2015), pp. 86-93.

[11] C. Lamason, B. MacMillan, B. Balcom, B. Leblon, Z. Pirouz. Field Measurements of Moisture Content in Black Spruce Logs with Unilateral Magnetic Resonance. Forest Products Journal, 67 (1-2) (2017), pp. 55-62.

[12] M. T. Riggin, A. R. Sharp, R. Kaiser, M. H. Schneider. Transverse NMR relaxation of water in wood. Journal of Applied Polymer Science, 23 (11) (1979), pp. 3147-3154.

[13] R. S. Menon, A. L. MaCkay, J. R. T. Hailey, M. Bloom, A. E. Burgess, J. S. Swanson. An NMR determination of the physiological water distribution in wood during drying. J. Applied Polymer Sci., 33 (4) (1987), pp. 1141-1155.

[14] E. C. Snelling. Soft Ferrites: Properties and Applications, 2nd Ed.; Butterworths: London, 1988; Ch. 2

[15] P. J. van der Zaag. New views on the dissipation in soft magnetic ferrites. Journal of Magnetism and Magnetic Materials, 197 (1999), p. 315.

Appendix

"Please, sir, I want some more."

-Oliver Twist

A.1 A solution to the Bloch equation without relaxation

It is strangely difficult to find a plain solution to the Bloch equation without relaxation, so I am including this one here. Others may have done it differently, but this is how I did it. The three big epiphanies from this derivation are that:

1) a **B** field imposes no change to the *z*-component of the magnetic moment.

2) the *x*- and *y*-components oscillate sinusoidally, constituting a precession about the *z*-axis set by the **B** field.

3) the rotation of **μ** about **B** occurs at the angular frequency $\omega_0 = \gamma B_0$.

Consider a spin with magnetic moment **μ** in a static external field, **B** with only one component along the *z*-direction, B_0.

Figure A.1.1. A magnetic moment, $\boldsymbol{\mu}$, is depicted in a static, external magnetic field, $\boldsymbol{B}$, which consists only of a z-component, B_0. The magnetic moment may be decomposed into μ_z, a z-component that is parallel with $\boldsymbol{B}$ and μ_{xy}, a component that lies in the xy-plane.

The Bloch equation without relaxation tells how the magnetic moment will move in response to being placed in this field, and it is given by

$$\frac{d}{dt}\boldsymbol{\mu} = \gamma(\boldsymbol{\mu} \times \boldsymbol{B}) \qquad (A.1.1)$$

where γ is the gyromagnetic ratio, an experimentally derived constant.

Applying the above condition that $\boldsymbol{B} = (0, 0, B_0)$ and expanding the cross product into its unit vector components in the usual way results in

$$\frac{d}{dt}\boldsymbol{\mu} = \begin{vmatrix} \hat{x} & \hat{y} & \hat{z} \\ \mu_x & \mu_y & \mu_z \\ 0 & 0 & B_0 \end{vmatrix} = \gamma\left(\mu_y B_0 \hat{x} - \mu_y B_0 \hat{y} + 0\hat{z}\right) = \gamma B_0\left(\mu_y \hat{x} - \mu_y \hat{y} + 0\hat{z}\right)$$

$$(A.1.2)$$

where, in the last term, B_0 has simply been undistributed. To know the state of the magnetic moment and not the change, the equation is integrated with respect to time.

$$\boldsymbol{\mu} = \gamma B_0 \left(\int \mu_y \hat{x} \, dt + \int \mu_x \hat{y} \, dt + \int 0 \hat{z} \, dt \right) \qquad (A.1.3)$$

At this point, all that is known about μ_x and μ_y is that they are unknown functions of time, and therefore, their integrals cannot be computed yet. The z-component of the magnetic moment, however, can easily be solved. Separating the last term from the equation, it can be seen that μ_z is a constant, so whatever μ_z was at $t = 0$, so it will remain until it is otherwise disturbed. Now, the x- and y-components may be solved for. The equation without the z-component is

$$\boldsymbol{\mu}_{xy} = \gamma B_0 \left(\int \mu_y dt \, \hat{x} + \int \mu_x dt \, \hat{y} \right) \qquad (A.1.4)$$

where $\boldsymbol{\mu}_{xy}$ is a vector that refers only to the x- and y-components of the magnetic moment. Recall now that the total magnetic moment of a spin is constant, and also that the z-component has already been shown to be constant. Therefore the magnitude in the x-y plane must also be constant. The magnitude squared, in the x-y plane, can be calculated by taking the dot product of $\boldsymbol{\mu}_{xy}$ with itself. It is also convenient to divide the through by other two constants, γ and B_0, so that all constants are on the same side.

$$\frac{|\mu_{xy}|^2}{\gamma^2 B_0^2} = \left(\int \mu_x dt \right)^2 + \left(\int \mu_y dt \right)^2 \qquad (A.1.5)$$

This equation, a constant being equal to the sum of two squares, is of a familiar form. Recall the trigonometric identity

$$A = A \cos^2 \theta + A \sin^2 \theta \qquad (A.1.6)$$

And recall also that an angle is equivalent to the product of an angular velocity and a time. In this case, the notation ω_0 will be used to indicate that the natural angular velocity of this system is being referenced.

$$\theta = \omega_0 t + \phi \qquad (A.1.7)$$

where t is time, as usual, and ϕ is an arbitrary phase.

Letting equations A.1.5 and A.1.6 be equivalent, while using the substitution in A.1.7, this implies that

$$A = \frac{|\mu_{xy}|^2}{\gamma^2 B_0^2} \qquad (A.1.8)$$

and further, that

$$\begin{cases} \left(\int \mu_x dt \right)^2 = \dfrac{|\mu_{xy}|^2}{\gamma^2 B_0^2} \cos^2(\omega_0 t + \phi) \\[4mm] \left(\int \mu_y dt \right)^2 = \dfrac{|\mu_{xy}|^2}{\gamma^2 B_0^2} \sin^2(\omega_0 t + \phi) \end{cases} \qquad (A.\,1.9)$$

Taking the square roots of both sides, this gives

$$\begin{cases} \int \mu_x dt = \dfrac{|\mu_{xy}|}{\gamma B_0} \cos(\omega_0 t + \phi) \\[4mm] \int \mu_y dt = \dfrac{|\mu_{xy}|}{\gamma B_0} \sin(\omega_0 t + \phi) \end{cases} \qquad (A.\,1.10)$$

These can now be differentiated with respect to time, giving us our answers for μ_x and μ_y as functions of time.

$$\begin{cases} \mu_x = -\dfrac{|\mu_{xy}|}{\gamma B_0} \omega_0 \sin(\omega_0 t + \phi) \\[4mm] \mu_y = \dfrac{|\mu_{xy}|}{\gamma B_0} \omega_0 \cos(\omega_0 t + \phi) \end{cases} \qquad (A.\,1.11)$$

These equations represent a circular rotation about the z-axis, which by convention is set by the direction of B_0. Thus it is shown that a magnetic moment, μ, in a static external field, $\boldsymbol{B}$, will precess about $\boldsymbol{B}$, without any change to the z-component of the magnetic moment. Now the frequency of the precession will be derived by using the Pythagorean Theorem in conjunction

with the identity in equation A.1.6. From Pythagoras, it is known that the sum of the squares of

two orthogonal elements is equal to the squared magnitude of the elements' vector sum.

$$\mu_x^2 + \mu_y^2 = \left|\mu_{xy}\right|^2 \tag{A.1.12}$$

Equation A.1.6 also says that the sum of the squares in this case are equal to a known

constant.

$$\mu_x^2 + \mu_y^2 = \frac{\left|\mu_{xy}\right|^2}{\gamma^2 B_0^2} \omega_0^2 \sin^2(\omega_0 t + \phi) + \frac{\left|\mu_{xy}\right|^2}{\gamma^2 B_0^2} \omega_0^2 \cos^2(\omega_0 t + \phi) = \frac{\left|\mu_{xy}\right|^2}{\gamma^2 B_0^2} \omega_0^2$$

$$\tag{A.1.13}$$

Then combining the results of these two equations, this must mean that

$$\left|\mu_{xy}\right|^2 = \frac{\left|\mu_{xy}\right|^2}{\gamma^2 B_0^2} \omega_0^2 \tag{A.1.14}$$

With a little cancelling, square rooting, and rearranging, this becomes

$$\omega_0 = \gamma B_0 \tag{A.1.15}$$

proving that the frequency of precession, ω_0 (the Larmor frequency), is proportional to the

strength of the external static magnetic field, B_0, and that the constant of proportionality is the

gyromagnetic ratio, γ. As a last comment, it is important to note that, this solution does not

depend on the strength of the magnetic moment. Therefore, whether speaking about the moment of a spin $\boldsymbol{\mu}_{xy}$, or the bulk magnetic moment $\boldsymbol{M}_{xy}$, the solution remains valid in predicting the precession and Larmor frequency.

A.2 A primer on Non-linear Iterative Partial Least Squares

To concisely and clearly explain the theory of Partial Least Squares can be confusing in itself: The idea is to decompose a set of dependent (prediction) data into a descending weighted sum of components, each of which are functions of similarly weighted components of the independent (input) data, and to do this based on the common variances of both sets of data. If done right, the outcome is simply a matrix that can be multiplied against any new set of input data, whose product is the set of desired predictions. Since the components used to build this matrix are ordered from most to least weight, this may be done with just a few components, while still maintaining a high level of predictiveness. Simple enough, right? But when it comes to the practicalities of determining these components and their sensitivities, however, the task can get downright daunting. Fortunately, there is an algorithm that determines, in a stepwise fashion, the most important (principal?) components and their sensitivities. This algorithm is called Non-linear Iterative Partial Least Squares, or NIPALS. No kidding. I will attempt to present it in a clear manner below.

First recall that a solution is desired in which a prediction of $\acute{y}$ values may be made from some independent observation, x, and that this prediction is found by mining calibration sets of X and their resulting sets of parameters, Y. The solution would thus look like

$$\acute{y} = Bx \qquad\qquad (A.2.1)$$

where B is a matrix of sensitivities of dimensions such that Bx yields the predicted parameters in the order that they were presented in the calibration sets.

NIPALS works by comparing a component in the independent (predictor) dataset to a component in the dependent (prediction) dataset, and looking for high overlap in both datasets. When a component is found with high overlap in both datasets, it is stored as the first principle component. The iterative part comes in the way that the component is initialized, compared and altered until it is found. The process is as follows, and is summarized in Figure A.2.1.

1) *Initialize*. Pick an initial component. This is typically done by selecting a random component of the independent dataset. For a comparison of frequency spectra, this would be selecting a particular frequency and making a component vector of the amplitudes of all calibration spectra at that frequency. This vector is referred to as u.

2) *Weight u in X.* This means to take the overlap of u and each column of X, but what it *does* is to create a new vector, w, the elements of which are the dot products of each column of X with u. Thus columns in X that overlap more with u will generate larger values for elements of w. Furthermore, the resulting overlap vector is normalized by the magnitude of u and then again it is normalized to a magnitude of one. Using the spectrum example, this would be taking looking at the overlap of amplitudes between the particular frequency component vector and each frequency component set in the calibration set. Mathematically, this is

$$w = \frac{X'u}{u'u} \qquad\qquad (A.\,2.2)$$

$$w = \frac{w}{\sqrt{w'w}} \qquad\qquad (A.\,2.3)$$

3) *Score **w** in **X***. Just like above in step 2, this just means overlapping **w** and **X**, but this time, the vector **w** is overlapped onto each *row* of **X**, instead of columns. The result is a vector, **t**, that has the same number of elements as the rows of both **X** and **Y**, and the elements of which represent the rows that overlap the most with the overlap of the previous component vector **u** on each column of **X**. In the example of frequency spectra, this is taking the overlap of the weight vector onto each *spectrum* in the calibration set. The result here is also normalized to the self-overlap of **w**. This is mathematically represented as

$$t = \frac{Xw}{w'w} \qquad\qquad (A.2.4)$$

4) *Weight **t** in **Y***. Now the result of the previous steps in **X** are overlapped with the components of **Y**. If at this point, the results of the previous two steps were instead overlapped (regressed) back into **X** and the steps repeated, this would converge to the component of greatest variance in **X**. That is the procedure for PCA in which the goal is to find the components of greatest variance in **X** and **Y**, regardless of their connection to each other. In PLS, the algorithm instead switches between regressing through *both* datasets and thus component of greatest *mutual* variance are constructed. The mathematical expression for this step, similar to step 2 above, is

$$c = \frac{Y't}{t't} \qquad\qquad (A.2.5)$$

135

5) *Score c in Y*. Again, the result of the weighting step is regressed back through **Y** to

   arrive at a new initial component, **u**. Mathematically, this will look familiar:

$$u = \frac{Yc}{c'c} \qquad (A.2.6)$$

   From this point, steps 2 through 5 are repeated and a comparison is made in each

   repetition between steps 3 and 4.

If the score of the previous repetition, t_{old}, and the new repetition, t_{new}, are close, within a set

threshold (e.g. 99%), all of the vectors (**u**, **w**, **t**, **c**) are stored collectively as the i^{th} PLS

component – *i* representing whatever number iteration is being performed. The algorithm at this

point breaks out to calculate the sensitivities (which will be used to make predictions later) and

remove the variance described by this component, or "deflate," before returning to the cycle

using a new, deflated, dataset. The steps for this procedure are described below.

6) *Calculate the sensitivities, **B***. If **t** has converged, the first element of **b**, b_i, is found by

   taking the product of t_i with u_i:

$$b_i = \frac{X_i' t_i}{t_i' t_i} \qquad (A.2.7)$$

   This is the *sensitivities* matrix.

7) *Calculate the residuals of X and Y, E and F.* The residuals of Y are calculated

   directly:

$$F = Y - b_i t_i c_i'$$

<div align="right">(A. 2.8)</div>

The residuals of X, are calculated using the weights recorded from the last iteration.

$$E = X - t_i w_i'$$

<div align="right">(A. 2.9)</div>

8) *Reenter the iterative cycle using the residuals, E and F, as the new datasets, X and Y.*

   The cycle of steps 2 through 5 is reentered using E and F as the new X and Y,

   respectively. A new t is required here, and it is found (typically) using a scaled t_{old}:

$$t_{new} = t_{old} |w_{old}|$$

<div align="right">(A. 2.10)</div>

Once a sufficient number of components have been found to explain the variance in X and make

predictions, $\acute{y}$, on future data, x, can be made.

$$\acute{y} = B' x$$

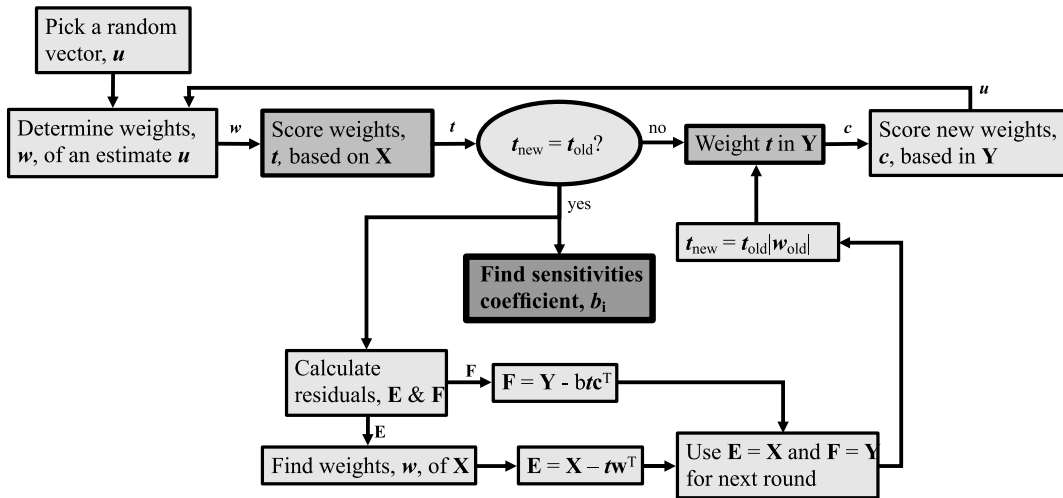<div align="right">(A. 2.11)</div>

NIPALS Algorithm in PLS

Figure A.2.1. The NIPALS algorithm described in steps 1 through 8 is summarized here in a flow chart.

A.3 A solved equation for predicting the field of a split fringe coil, perpendicular to a conducting surface

In the course of designing and characterizing the properties of fringe coils for UMR, it became clear that it would be useful to have an equation that would analytically describe the magnetic field produced by an opposed fringe coil. Presented here is such an equation. Furthermore, since the goal is to produce a homogeneous B_1 field at a prescribed distance above the probe coil, one would benefit from an equation that describes the second derivative of the field strength at the standoff distance, as a function of lateral offset from center. This was calculated using Mathematica™ by using the Biot-Savart Law for a 2D surface and integrating across the surface. Only the surface of the opposing coils that is displaced from the magnet body is needed, since the lower surface lied directly on the conducting magnet and is therefore completely negated by the effects of eddy-currents (see Chapter 4 for the derivation that shows this).

This is entirely based on solving the Biot-Savart Law equation for a surface current. In this formulation, the magnetic field at a point r is:

$$\boldsymbol{B}(\boldsymbol{r}) = \frac{\mu_0}{4\pi} \int \frac{\boldsymbol{K}(\boldsymbol{r}') \times \widehat{\Delta \boldsymbol{r}}}{\Delta r^2} da' \qquad (A.3.1)$$

in which μ_0 is the permeability of free space, $\boldsymbol{K}(\boldsymbol{r'})$ is the surface current at the location $\boldsymbol{r'}$, $\Delta\boldsymbol{r}$ is the vector from $\boldsymbol{r'}$ to $\boldsymbol{r}$.

The steps for solving this were, generally:

1. Solve by hand the cross-product of the surface current and the displacement unit vector, in vector form.

2. Using Mathematica, integrate generally the B_x and B_y components separately, each with respect to *da'* by integrating with respect to *dx'* first, then *dy'*.

3. Use the resulting integral to define a definite integral based on additional new variables, *d* (one-half the separation distance between the opposing coils), *l*, (the length of each opposing coil in the x-direction), and *w* (the width of the coils in the y-direction).

This solves for the z-component of the magnetic field at an arbitrary point, **r**, but to find the optimal height, the curvature (at least in the *x*-direction) must be determined as a function of height, *z*. This adds an extra step of taking the second derivative of the magnetic field, with respect to *x*-direction.

4. Use Mathematica to calculate the second derivative with respect to **x**, $d^2/dx^2\ B_z$, and plot as a function of height, *z*.

The resulting calculations nearly broke my computer, but in the actual calculation of them, as well as in the attempt to include the resulting equations in this manuscript. This is the reason for including the curvature equation as an image, rather than as true math-type equations.

The field strength was found to be:

$$B_z(r) = -\frac{1}{4\pi} K y \mu_0 \left(\frac{\text{ArcTan}\left[\frac{-\frac{w}{2}+y}{\sqrt{d^2 - 2dx + x^2 + z^2}}\right]}{\sqrt{d^2 - 2dx + x^2 + z^2}} - \frac{\text{ArcTan}\left[\frac{\frac{w}{2}+y}{\sqrt{d^2 - 2dx + x^2 + z^2}}\right]}{\sqrt{d^2 - 2dx + x^2 + z^2}} \right.$$

$$+ \frac{\text{ArcTan}\left[\frac{-\frac{w}{2}+y}{\sqrt{d^2 + 2dx + x^2 + z^2}}\right]}{\sqrt{d^2 + 2dx + x^2 + z^2}} - \frac{\text{ArcTan}\left[\frac{\frac{w}{2}+y}{\sqrt{d^2 + 2dx + x^2 + z^2}}\right]}{\sqrt{d^2 + 2dx + x^2 + z^2}}$$

$$- \frac{\text{ArcTan}\left[\frac{-\frac{w}{2}+y}{\sqrt{(d+l)^2 - 2(d+l)x + x^2 + z^2}}\right]}{\sqrt{(d+l)^2 - 2(d+l)x + x^2 + z^2}}$$

$$+ \frac{\text{ArcTan}\left[\frac{\frac{w}{2}+y}{\sqrt{(d+l)^2 - 2(d+l)x + x^2 + z^2}}\right]}{\sqrt{(d+l)^2 - 2(d+l)x + x^2 + z^2}}$$

$$- \frac{\text{ArcTan}\left[\frac{-w+2y}{2\sqrt{d^2 + l^2 + 2lx + x^2 + 2d(l+x) + z^2}}\right]}{\sqrt{d^2 + l^2 + 2lx + x^2 + 2d(l+x) + z^2}}$$

$$\left. + \frac{\text{ArcTan}\left[\frac{w+2y}{2\sqrt{d^2 + l^2 + 2lx + x^2 + 2d(l+x) + z^2}}\right]}{\sqrt{d^2 + l^2 + 2lx + x^2 + 2d(l+x) + z^2}} \right)$$

$$(A.3.2)$$

while the x-direction curvature was found to be:

$$-\frac{1}{2\pi}$$

$$\mu\left(\frac{1}{2}\,Ky\left|\frac{(-2\,d+2\,x)^2\,\left(-\frac{w}{2}+y\right)}{2\,\left(d^2-2\,d\,x+x^2+z^2\right)^3\,\left(1+\frac{\left(-\frac{w}{2}+y\right)^2}{d^2-2\,d\,x+x^2+z^2}\right)}+\frac{1}{\sqrt{d^2-2\,d\,x+x^2+z^2}}\left|-\frac{(-2\,d+2\,x)^2\,\left(-\frac{w}{2}+y\right)^3}{2\,\left(d^2-2\,d\,x+x^2+z^2\right)^{7/2}\,\left(1+\frac{\left(-\frac{w}{2}+y\right)^2}{d^2-2\,d\,x+x^2+z^2}\right)^2}+\right.\right.\right.$$

$$\left.\frac{3\,(-2\,d+2\,x)^2\,\left(-\frac{w}{2}+y\right)}{4\,\left(d^2-2\,d\,x+x^2+z^2\right)^{5/2}\,\left(1+\frac{\left(-\frac{w}{2}+y\right)^2}{d^2-2\,d\,x+x^2+z^2}\right)}-\frac{-\frac{w}{2}+y}{\left(d^2-2\,d\,x+x^2+z^2\right)^{3/2}\,\left(1+\frac{\left(-\frac{w}{2}+y\right)^2}{d^2-2\,d\,x+x^2+z^2}\right)}\right)+$$

$$\left.\left(\frac{3\,(-2\,d+2\,x)^2}{4\,\left(d^2-2\,d\,x+x^2+z^2\right)^{5/2}}-\frac{1}{\left(d^2-2\,d\,x+x^2+z^2\right)^{3/2}}\right)\text{ArcTan}\left[\frac{-\frac{w}{2}+y}{\sqrt{d^2-2\,d\,x+x^2+z^2}}\right]\right)-\frac{1}{2}\,Ky$$

$$\left(\frac{(-2\,d+2\,x)^2\,\left(\frac{w}{2}+y\right)}{2\,\left(d^2-2\,d\,x+x^2+z^2\right)^3\,\left(1+\frac{\left(\frac{w}{2}+y\right)^2}{d^2-2\,d\,x+x^2+z^2}\right)}+\frac{1}{\sqrt{d^2-2\,d\,x+x^2+z^2}}\left|-\frac{(-2\,d+2\,x)^2\,\left(\frac{w}{2}+y\right)^3}{2\,\left(d^2-2\,d\,x+x^2+z^2\right)^{7/2}\,\left(1+\frac{\left(\frac{w}{2}+y\right)^2}{d^2-2\,d\,x+x^2+z^2}\right)^2}+\right.\right.$$

$$\left.\frac{3\,(-2\,d+2\,x)^2\,\left(\frac{w}{2}+y\right)}{4\,\left(d^2-2\,d\,x+x^2+z^2\right)^{5/2}\,\left(1+\frac{\left(\frac{w}{2}+y\right)^2}{d^2-2\,d\,x+x^2+z^2}\right)}-\frac{\frac{w}{2}+y}{\left(d^2-2\,d\,x+x^2+z^2\right)^{3/2}\,\left(1+\frac{\left(\frac{w}{2}+y\right)^2}{d^2-2\,d\,x+x^2+z^2}\right)}\right)+$$

$$\left(\frac{3\,(-2\,d+2\,x)^2}{4\,\left(d^2-2\,d\,x+x^2+z^2\right)^{5/2}}-\frac{1}{\left(d^2-2\,d\,x+x^2+z^2\right)^{3/2}}\right)\text{ArcTan}\left[\frac{\frac{w}{2}+y}{\sqrt{d^2-2\,d\,x+x^2+z^2}}\right]\right)+\frac{1}{2}\,Ky$$

$$\left(\frac{(2\,d+2\,x)^2\,\left(-\frac{w}{2}+y\right)}{2\,\left(d^2+2\,d\,x+x^2+z^2\right)^3\,\left(1+\frac{\left(-\frac{w}{2}+y\right)^2}{d^2+2\,d\,x+x^2+z^2}\right)}+\frac{1}{\sqrt{d^2+2\,d\,x+x^2+z^2}}\left|-\frac{(2\,d+2\,x)^2\,\left(-\frac{w}{2}+y\right)^3}{2\,\left(d^2+2\,d\,x+x^2+z^2\right)^{7/2}\,\left(1+\frac{\left(-\frac{w}{2}+y\right)^2}{d^2+2\,d\,x+x^2+z^2}\right)^2}+\right.\right.$$

$$\left.\frac{3\,(2\,d+2\,x)^2\,\left(-\frac{w}{2}+y\right)}{4\,\left(d^2+2\,d\,x+x^2+z^2\right)^{5/2}\,\left(1+\frac{\left(-\frac{w}{2}+y\right)^2}{d^2+2\,d\,x+x^2+z^2}\right)}-\frac{-\frac{w}{2}+y}{\left(d^2+2\,d\,x+x^2+z^2\right)^{3/2}\,\left(1+\frac{\left(-\frac{w}{2}+y\right)^2}{d^2+2\,d\,x+x^2+z^2}\right)}\right)+$$

$$\left(\frac{3\,(2\,d+2\,x)^2}{4\,\left(d^2+2\,d\,x+x^2+z^2\right)^{5/2}}-\frac{1}{\left(d^2+2\,d\,x+x^2+z^2\right)^{3/2}}\right)\text{ArcTan}\left[\frac{-\frac{w}{2}+y}{\sqrt{d^2+2\,d\,x+x^2+z^2}}\right]\right)-\frac{1}{2}\,Ky$$

$$\left(\frac{(2\,d+2\,x)^2\,\left(\frac{w}{2}+y\right)}{2\,\left(d^2+2\,d\,x+x^2+z^2\right)^3\,\left(1+\frac{\left(\frac{w}{2}+y\right)^2}{d^2+2\,d\,x+x^2+z^2}\right)}+\frac{1}{\sqrt{d^2+2\,d\,x+x^2+z^2}}\left|-\frac{(2\,d+2\,x)^2\,\left(\frac{w}{2}+y\right)^3}{2\,\left(d^2+2\,d\,x+x^2+z^2\right)^{7/2}\,\left(1+\frac{\left(\frac{w}{2}+y\right)^2}{d^2+2\,d\,x+x^2+z^2}\right)^2}+\right.\right.$$

$$\left.\frac{3\,(2\,d+2\,x)^2\,\left(\frac{w}{2}+y\right)}{4\,\left(d^2+2\,d\,x+x^2+z^2\right)^{5/2}\,\left(1+\frac{\left(\frac{w}{2}+y\right)^2}{d^2+2\,d\,x+x^2+z^2}\right)}-\frac{\frac{w}{2}+y}{\left(d^2+2\,d\,x+x^2+z^2\right)^{3/2}\,\left(1+\frac{\left(\frac{w}{2}+y\right)^2}{d^2+2\,d\,x+x^2+z^2}\right)}\right)+$$

$$\left(\frac{3\,(2\,d+2\,x)^2}{4\,\left(d^2+2\,d\,x+x^2+z^2\right)^{5/2}}-\frac{1}{\left(d^2+2\,d\,x+x^2+z^2\right)^{3/2}}\right)\text{ArcTan}\left[\frac{\frac{w}{2}+y}{\sqrt{d^2+2\,d\,x+x^2+z^2}}\right]\right)-$$

$$\frac{1}{2}\,Ky\left(\frac{(-2\,(-d-l)+2\,x)^2\left(-\frac{w}{2}+y\right)}{2\left((-d-l)^2-2\,(-d-l)\,x+x^2+z^2\right)^3\left(1+\frac{\left(-\frac{w}{2}+y\right)^2}{(-d-l)^2-2\,(-d-l)\,x+x^2+z^2}\right)}+\right.$$

$$\frac{1}{\sqrt{(-d-l)^2-2\,(-d-l)\,x+x^2+z^2}}\left(-\frac{(-2\,(-d-l)+2\,x)^2\left(-\frac{w}{2}+y\right)^3}{2\left((-d-l)^2-2\,(-d-l)\,x+x^2+z^2\right)^{7/2}\left(1+\frac{\left(-\frac{w}{2}+y\right)^2}{(-d-l)^2-2\,(-d-l)\,x+x^2+z^2}\right)^2}+\right.$$

$$\frac{3\,(-2\,(-d-l)+2\,x)^2\left(-\frac{w}{2}+y\right)}{4\left((-d-l)^2-2\,(-d-l)\,x+x^2+z^2\right)^{5/2}\left(1+\frac{\left(-\frac{w}{2}+y\right)^2}{(-d-l)^2-2\,(-d-l)\,x+x^2+z^2}\right)}-$$

$$\left.\frac{-\frac{w}{2}+y}{\left((-d-l)^2-2\,(-d-l)\,x+x^2+z^2\right)^{3/2}\left(1+\frac{\left(-\frac{w}{2}+y\right)^2}{(-d-l)^2-2\,(-d-l)\,x+x^2+z^2}\right)}\right)+$$

$$\left(\frac{3\,(-2\,(-d-l)+2\,x)^2}{4\left((-d-l)^2-2\,(-d-l)\,x+x^2+z^2\right)^{5/2}}-\frac{1}{\left((-d-l)^2-2\,(-d-l)\,x+x^2+z^2\right)^{3/2}}\right)$$

$$\left.\mathrm{ArcTan}\left[\frac{-\frac{w}{2}+y}{\sqrt{(-d-l)^2-2\,(-d-l)\,x+x^2+z^2}}\right]\right)+$$

$$\frac{1}{2}\,Ky\left(\frac{(-2\,(-d-l)+2\,x)^2\left(\frac{w}{2}+y\right)}{2\left((-d-l)^2-2\,(-d-l)\,x+x^2+z^2\right)^3\left(1+\frac{\left(\frac{w}{2}+y\right)^2}{(-d-l)^2-2\,(-d-l)\,x+x^2+z^2}\right)}+\right.$$

$$\frac{1}{\sqrt{(-d-l)^2-2\,(-d-l)\,x+x^2+z^2}}\left(-\frac{(-2\,(-d-l)+2\,x)^2\left(\frac{w}{2}+y\right)^3}{2\left((-d-l)^2-2\,(-d-l)\,x+x^2+z^2\right)^{7/2}\left(1+\frac{\left(\frac{w}{2}+y\right)^2}{(-d-l)^2-2\,(-d-l)\,x+x^2+z^2}\right)^2}+\right.$$

$$\frac{3\,(-2\,(-d-l)+2\,x)^2\left(\frac{w}{2}+y\right)}{4\left((-d-l)^2-2\,(-d-l)\,x+x^2+z^2\right)^{5/2}\left(1+\frac{\left(\frac{w}{2}+y\right)^2}{(-d-l)^2-2\,(-d-l)\,x+x^2+z^2}\right)}-$$

$$\left.\frac{\frac{w}{2}+y}{\left((-d-l)^2-2\,(-d-l)\,x+x^2+z^2\right)^{3/2}\left(1+\frac{\left(\frac{w}{2}+y\right)^2}{(-d-l)^2-2\,(-d-l)\,x+x^2+z^2}\right)}\right)+$$

$$\left(\frac{3\,(-2\,(-d-l)+2\,x)^2}{4\left((-d-l)^2-2\,(-d-l)\,x+x^2+z^2\right)^{5/2}}-\frac{1}{\left((-d-l)^2-2\,(-d-l)\,x+x^2+z^2\right)^{3/2}}\right)$$

$$\left.\mathrm{ArcTan}\left[\frac{\frac{w}{2}+y}{\sqrt{(-d-l)^2-2\,(-d-l)\,x+x^2+z^2}}\right]\right)-$$

$$\frac{1}{2}\,Ky\left|\frac{(-2\,(d+l)+2\,x)^2\left(-\frac{w}{2}+y\right)}{2\left((d+l)^2-2\,(d+l)\,x+x^2+z^2\right)^3\left(1+\frac{\left(-\frac{w}{2}+y\right)^2}{(d+l)^2-2\,(d+l)\,x+x^2+z^2}\right)}+\frac{1}{\sqrt{(d+l)^2-2\,(d+l)\,x+x^2+z^2}}\right.$$

$$\left|-\frac{(-2\,(d+l)+2\,x)^2\left(-\frac{w}{2}+y\right)^3}{2\left((d+l)^2-2\,(d+l)\,x+x^2+z^2\right)^{7/2}\left(1+\frac{\left(-\frac{w}{2}+y\right)^2}{(d+l)^2-2\,(d+l)\,x+x^2+z^2}\right)^2}+\right.$$

$$\frac{3\,(-2\,(d+l)+2\,x)^2\left(-\frac{w}{2}+y\right)}{4\left((d+l)^2-2\,(d+l)\,x+x^2+z^2\right)^{5/2}\left(1+\frac{\left(-\frac{w}{2}+y\right)^2}{(d+l)^2-2\,(d+l)\,x+x^2+z^2}\right)}-$$

$$\left.\frac{-\frac{w}{2}+y}{\left((d+l)^2-2\,(d+l)\,x+x^2+z^2\right)^{3/2}\left(1+\frac{\left(-\frac{w}{2}+y\right)^2}{(d+l)^2-2\,(d+l)\,x+x^2+z^2}\right)}\right)+$$

$$\left(\frac{3\,(-2\,(d+l)+2\,x)^2}{4\left((d+l)^2-2\,(d+l)\,x+x^2+z^2\right)^{5/2}}-\frac{1}{\left((d+l)^2-2\,(d+l)\,x+x^2+z^2\right)^{3/2}}\right)$$

$$\left.\text{ArcTan}\left[\frac{-\frac{w}{2}+y}{\sqrt{(d+l)^2-2\,(d+l)\,x+x^2+z^2}}\right]\right]+$$

$$\frac{1}{2}\,Ky\left|\frac{(-2\,(d+l)+2\,x)^2\left(\frac{w}{2}+y\right)}{2\left((d+l)^2-2\,(d+l)\,x+x^2+z^2\right)^3\left(1+\frac{\left(\frac{w}{2}+y\right)^2}{(d+l)^2-2\,(d+l)\,x+x^2+z^2}\right)}+\frac{1}{\sqrt{(d+l)^2-2\,(d+l)\,x+x^2+z^2}}\right.$$

$$\left|-\frac{(-2\,(d+l)+2\,x)^2\left(\frac{w}{2}+y\right)^3}{2\left((d+l)^2-2\,(d+l)\,x+x^2+z^2\right)^{7/2}\left(1+\frac{\left(\frac{w}{2}+y\right)^2}{(d+l)^2-2\,(d+l)\,x+x^2+z^2}\right)^2}+\right.$$

$$\frac{3\,(-2\,(d+l)+2\,x)^2\left(\frac{w}{2}+y\right)}{4\left((d+l)^2-2\,(d+l)\,x+x^2+z^2\right)^{5/2}\left(1+\frac{\left(\frac{w}{2}+y\right)^2}{(d+l)^2-2\,(d+l)\,x+x^2+z^2}\right)}-$$

$$\left.\frac{\frac{w}{2}+y}{\left((d+l)^2-2\,(d+l)\,x+x^2+z^2\right)^{3/2}\left(1+\frac{\left(\frac{w}{2}+y\right)^2}{(d+l)^2-2\,(d+l)\,x+x^2+z^2}\right)}\right)+$$

$$\left(\frac{3\,(-2\,(d+l)+2\,x)^2}{4\left((d+l)^2-2\,(d+l)\,x+x^2+z^2\right)^{5/2}}-\frac{1}{\left((d+l)^2-2\,(d+l)\,x+x^2+z^2\right)^{3/2}}\right)$$

$$\left.\left.\text{ArcTan}\left[\frac{\frac{w}{2}+y}{\sqrt{(d+l)^2-2\,(d+l)\,x+x^2+z^2}}\right]\right]\right)$$

$$(A.3.3)$$

The curvature equation can be vastly simplified by setting the surface current, K_y, as well as the permeability of free space, μ_0, to 1. If components x and y are set static at zero, so that the

equation only describes the curvature as a function of height, z, then the equation becomes almost manageable:

$$\frac{d^2}{dx^2}Bz = \frac{1}{2\pi}\Big(\frac{1}{w^3}\,4\big(\frac{3d^2w^2}{(d^2+z^2)^2} - \frac{4d^2+w^2}{d^2+z^2} - \frac{3(d+l)^2w^2}{((d+l)^2+z^2)^2} + \frac{4(d+l)^2+w^2}{(d+l)^2+z^2}$$

$$-\frac{32d^2w^2}{(4d^2+w^2+4z^2)^2} + \frac{4(4d^2+w^2)}{4d^2+w^2+4z^2} + \frac{32(d+l)^2w^2}{(4(d+l)^2+w^2+4z^2)^2}$$

$$+\frac{-16(d+l)^2-4w^2}{4(d+l)^2+w^2+4z^2}\big) - \frac{2(-2d^2+z^2)\mathrm{ArcTan}[\frac{w}{2\sqrt{d^2+z^2}}]}{(d^2+z^2)^{5/2}}$$

$$+\frac{2(-2(d+l)^2+z^2)\mathrm{ArcTan}[\frac{w}{2\sqrt{(d+l)^2+z^2}}]}{((d+l)^2+z^2)^{5/2}}\Big)$$

$$(A.3.4)$$

From all of this, a useful equation would be one that describes the optimal height at which the curvature of B_z with respect to x would be zero. This can be done by varying values of w, l, and d, and fitting a function to the behavior of the changes in optimal height. For this equation, this was done by asserting that the resulting function would be separable into functions of w, l, and d, then multiplying them all together. The function found from this procedure is as follows.

$$z_{opt} = 2.1 \times (0.36(1 - e^{-1.2l}) + 0.47) \times (1.11d + 0.14) \times \Big(0.15e^{-\frac{w^2}{3.1}} + 0.62\Big)$$

$$(A.3.5)$$

where again, *w*, *l*, and *d* are the width, length, and ½-separation distance of the opposing coils, and *z* is height from the top surface of the coils.

A strong caution must be advised when using this equation, in that it only describes the *z*-component of B_1. Strong deviances from zero in the *x*- and *y*-directions will severely compromise the validity and usefulness of this equation when designing opposed fringe coils.