

Abstract

Using Novel Pulse Sequences for Magnetic Resonance Imaging of ^{31}P Phosphorus in Hard and Soft Solids

Merideth A. Frey

2013

Since its invention in 1973, magnetic resonance imaging (MRI) has become an invaluable tool for clinical medicine, fundamental biomedical research, the physical sciences, and engineering. The vast majority of all MRI studies, in medicine and beyond, detect only the signal from a single nuclear isotope, ^1H , in liquid water. Extending the reach of MRI to the study of other elements, and to hard or soft solids, opens new frontiers of discovery. In practice, however, the slower motion of the nuclei in solid environments compared to ^1H in water results in much broader magnetic resonance (MR) spectra, limiting both the attainable spatial resolution and the signal-to-noise. Our lab recently discovered a novel nuclear magnetic resonance (NMR) pulse sequence while doing fundamental research related to the ‘spins in semiconductors’ approach to quantum computing. This sequence can greatly narrow the MR linewidth of solids, and it opens a new path to do high-resolution MRI of various nuclei in solids. In this thesis work, I use our quadratic echo line-narrowing pulse sequence to take the highest resolution MR images of ^{31}P in hard and soft solids using a conventional animal MRI system. I also discuss strategies to accelerate the imaging speed by making use of sparse MRI techniques as well as a new algorithm developed in our lab to do fast and accurate image reconstruction from sparse data. For future work, I propose ways to enhance spatial resolution and speed up imaging as well as discuss the potential applications of this work to a wider range of scientific problems.

**Using Novel Pulse Sequences for Magnetic
Resonance Imaging of ^{31}P Phosphorus in
Hard and Soft Solids**

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by
Merideth A. Frey

Dissertation Director: Professor Sean E. Barrett

December 2013

© 2013 by Merideth A. Frey

All rights reserved.

Contents

Table of Contents	iii
List of Figures	vi
List of Tables	ix
Words of Acknowledgment	x
1 Introduction	1
1.1 Review of the Field Prior to this Research	3
1.1.1 Living with the Short T_2 of Solids	3
1.1.2 Making Solids Look More Like Liquids	5
1.1.3 Previous Work Imaging Bone	6
1.1.4 Our Method	7
1.2 Products of this Research	7
1.3 Organization of the Remaining Chapters	8
Introduction	1
2 NMR and MRI Basics	9
2.1 Basics of NMR	9
2.1.1 Larmor Frequency and Precession	9

2.1.2	Classical Approach to NMR	11
2.1.3	Relevant Spin Interactions	16
2.2	Basics of MRI	18
2.2.1	k Space	20
2.2.2	Traversing k Space	21
2.2.3	Sampling, Field-Of-View, and Spatial Resolution	21
2.2.4	Exploiting Hermitian Symmetry	24
2.3	Equipment	25
2.3.1	NMR Equipment	25
2.3.2	Imaging Equipment	25
2.4	Challenges of MRI of Solids	27
3	Quadratic Echo Line-Narrowing and Imaging Pulse Sequences	30
3.1	Brief History of the Development of the Line-Narrowing Sequence . . .	31
3.1.1	Average Hamiltonian Theory	34
3.1.2	Testing Our Model	37
3.1.3	Designing the Quadratic Echo Line-Narrowing Pulse Sequence .	41
3.2	Imaging Sequences	43
3.2.1	Second Generation Pulse Block	45
3.2.2	Low Resolution Imaging Sequence	48
3.2.3	High Resolution Fast Switch Imaging Sequence	48
3.2.4	High Resolution Slow Switch Imaging Sequence	52
3.3	Important Considerations for MRI of Solids	54
3.3.1	Lowering Pulse Power	55
3.3.2	Imaging Sequence Used for MRI Results	55

4	High Resolution MRI of Solids	58
4.1	MRI of Solids Using Cartesian Sampling	59
4.1.1	Cartesian Sampling of k Space	59
4.1.2	Cartesian Sampling Analysis	62
4.1.3	Cartesian Sampling Results	69
4.2	MRI of Solids Using Radial Sampling	79
4.2.1	Radial Sampling of k Space	80
4.2.2	Radial Sampling Analysis	90
4.2.3	Radial Sampling Results	99
5	Imaging Solids Faster	101
5.1	Sparse MRI	102
5.1.1	Sparse Sampling	102
5.1.2	Sparse Sampling Results	104
5.2	Reconstruction Algorithms	108
5.2.1	Notation	110
5.2.2	Projections	111
5.2.3	Alternating Projections	112
5.2.4	Difference Map	114
5.2.5	What are the differences between the two algorithms?	116
5.2.6	Metrics to Monitor Convergence	117
5.2.7	Choosing Error Thresholds	121
5.2.8	k Space Error Handling	123
5.2.9	Image Space Error Handling	128
5.3	Image Reconstruction Results	131
5.3.1	Reconstruction Results on Sparse MRI Data	131

5.3.2	Reconstruction Results on 2D NMR Data	139
5.3.3	Further Comments	149
6	Future Work	156
6.1	Enhancing Resolution via Decoupling	156
6.1.1	Double Resonance Probe	157
6.1.2	Preliminary Decoupling Results	160
6.2	Future Applications of MRI of Solids	163
6.2.1	Probing Granular Material	164
6.2.2	Other Potential Samples of Interest	165
6.3	Future Applications of Image Reconstruction Algorithms	166
6.4	Final Thoughts	168
A	Measurements of Gradient Ramp Times	169
B	Practical Applications for Radial Sampling	171
C	Deriving 3D Density Weighting Function	173
D	Double-Resonance Probe Design Details	176
D.1	Proton Trap	176
D.2	High Frequency Circuit	177
D.3	Low Frequency Circuit	178
	Bibliography	181

List of Figures

1.1	Liquid vs. solid signal	2
2.1	Larmor Precession	12
2.2	Hahn Echo	17
2.3	Lauterbur's MRI Figure	19
2.4	Resonant Tank Circuits	26
2.5	Comparison of liquid and solid linewidths	28
3.1	Comparison of Hahn Echo and CPMG experiments in silicon	33
3.2	Testing our theoretical model for spin coherence control in C ₆₀	39
3.3	Three experiments inspired by the magic echo	40
3.4	Explanation of the quadratic echo line-narrowing pulse block	43
3.5	Line-narrowing sequence on a sample of Si:Sb	44
3.6	Second Generation Pulse Block	46
3.7	High resolution fast switch pulse sequence	50
3.8	Fast switch box top on ³¹ P in bone	50
3.9	High resolution slow switch pulse sequence	53
3.10	Slow switch box top on ³¹ P in bone	53
4.1	Schematic of our imaging sequence	60
4.2	Mapping acquired data onto \vec{k} space and the resulting image	63

4.3	Schematic of the gradient transient	67
4.4	Measurement of the gradient transient	68
4.5	Bone structure	70
4.6	Images of <i>ex vivo</i> pork rib sample	73
4.7	MRI and Micro-CT Comparison	74
4.8	Images of <i>ex vivo</i> rabbit femoral head sample	76
4.9	Phosphorus in a human cell	77
4.10	Images of <i>ex vivo</i> mouse brain	78
4.11	Pulse sequence diagram using radial sampling	81
4.12	Radial 1D frequency spectra with x and y effective frequency offsets	82
4.13	Radial 1D frequency spectra with common frequency offset	84
4.14	Uniformly distributing points on the surface of a sphere	85
4.15	2D depiction of radial sampling \vec{k} space	86
4.16	Comparing density compensation factors	96
4.17	Comparing radial images with and without interpolation	98
4.18	Comparing radial and Cartesian sampled images	100
5.1	Sparse Sampling in MRI	105
5.2	Sparse sampling of \vec{k} space and resulting image	106
5.3	Alternating Projections Reconstruction	113
5.4	Difference map runaway behavior near gap	115
5.5	Plots of various convergence measures	120
5.6	Choosing 3D Sparse Trajectories	133
5.7	Plot of l_2 for different sparse sampling strategies	134
5.8	Reconstruction of sparsely sampled solid-state MRI data	136
5.9	Comparing image reconstruction from different 3D sampling strategies	138

5.10	2D slices of sparse images from different 3D sampling strategies	139
5.11	Mapping the four time-domain quadrants	142
5.12	Comparing reconstruction performance using different \hat{P}_1 masks for IGPS data	145
5.13	Reconstucting liquid-state 2D NMR data of LuxU sample	151
5.14	Reconstucting liquid-state 2D NMR data of LuxU sample with 50 t_1 rows	152
5.15	Reconstucting liquid-state 2D NMR data of LuxU sample with 35 t_1 rows	152
5.16	Reconstucting liquid-state 2D NMR data of LuxU sample with 20 t_1 rows	153
5.17	Reconstructing liquid-state 2D NMR data of IGPS sample	154
5.18	Reconstructing solid-state 2D NMR data of GB1 sample	155
6.1	Proposed layout of double-resonance circuit	158
6.2	Final layout of double-resonance circuit	160
6.3	Proton Decoupling in Bone Mineral	161
6.4	T_1 Pressure Dependence in Rubber	163
6.5	Probing Local Contact Forces in Granular Material	164
D.1	Layout for effective high frequency circuit	177
D.2	Layout for effective low frequency circuit	179

List of Tables

3.1	Table of Effectiveness of Line-Narrowing with Lower Pulse Power . . .	54
6.1	Table of values for double-resonance circuit components	160
A.1	Table of Gradient Ramp Times	169
D.1	Table of theoretical values for double-resonance circuit components . .	180

Words of Acknowledgment

This work builds from the previous PhD thesis work by Yanqun Dong and Rona Ramos with contributions from many talented Yale undergraduates: Zach Sethna (the difference map discoverer), Peter Zhan, Andrew Wang, Ben Deen, and Suyog Bhandari.

I greatly appreciate the support of colleagues at the Medical Resonance Research Center, including Gigi Galiana (who taught me Paravision), Peter Brown (who taught me how to build MRI probes), Terry Nixon (who patiently dealt with all our equipment requests), Robin de Graaf (who freely gave pulse sequence programming advice), Todd Constable (who gave us constant support), and Doug Rothman (who gave us many ‘free’ hours on the mini-4T animal magnet system). I am grateful for the help of our biomedical collaborators, especially Michael Michaud (who provided all the *ex vivo* soft tissue samples), Joshua VanHouten (who took micro-CT images of our bone samples), Karl Insogna (the bone expert), and Joseph Madri (the cell expert).

I want to thank my thesis committee: Michel Devoret, Meg Urry, and Steve Girvin, for their advice and time. Extra thanks to my advisor, Sean Barrett, whose boundless optimism, natural curiosity, and utmost support of my career development has made me a confident experimental physicist.

Finally, I want to thank my family, who have supported all my scientific aspirations since childhood, and my husband, M. E. Irizarry-Gelpí, who has dealt with late nights in the lab and erratic schedules with patience and grace.

Chapter 1

Introduction

Since Paul Lauterbur's first attempt at magnetic resonance imaging (MRI) in 1973, MRI has sky-rocketed in popularity to become a cornerstone for many diagnostic biomedical procedures and remains the best technique to do non-invasive *in vivo* imaging of soft tissue. MRI still provides an active research arena which attracts a wide-range of researchers: from chemists developing contrast agents that highlight specific organs to mathematicians using information theory to develop faster imaging methods by making use of randomly sparse data sets. Despite the plethora of MRI research being done, the vast majority of imaging is still looking at just a single nuclear isotope, ^1H , in water.

The reasons for such a seemingly single-minded pursuit are as follows: (a) water is abundant in biological tissues, (b) ^1H has the largest magnetic moment of all naturally occurring nuclear isotopes, and (c) the motion of liquid water naturally narrows the ^1H spectra as well as speeds up the imaging time, thus enabling faster and higher spatial resolution imaging. There are many other nuclear isotopes which exhibit magnetic resonance that could give different and complementary information about the particular sample, but imaging using non-proton nuclei still requires several challenges to be

overcome.

Along with having lower concentrations (e.g., in comparison to hydrogen in biological samples), these other nuclei are often found in solid or soft-solid structures whose slower molecular motion results in the signal decaying very rapidly (i.e., short T_2) and much broader magnetic resonant spectra, which hurt the available spatial resolution and the signal-to-noise. The lack of fluctuating magnetic fields also leads to a long wait time between experiments for the spins to equilibrate along the applied magnetic field (i.e., long T_1). Fortunately, the novel pulse sequence developed in our lab is well-suited to make the magnetic resonance spectrum of a solid more like that of a liquid, and appears to help remove the hurdle to do high spatial resolution imaging. Our lab has also developed algorithms that can reconstruct very accurate images from sparse data, offering a route to speed-up imaging of solids as well. Both these techniques combined make MRI of nuclei in solid structures much more practical.

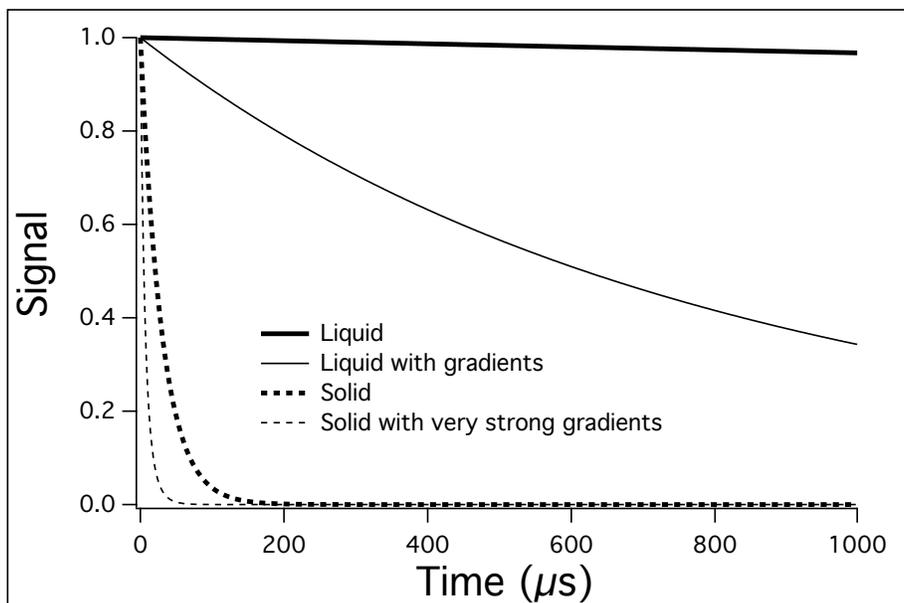


Figure 1.1: A plot of the signal from a liquid-state (solid lines) or solid-state sample (dashed lines), with and without applied magnetic field gradients.

1.1 Review of the Field Prior to this Research

Due to all these challenges, there is not yet a standard method to use for MRI of solids. Good reviews for the different techniques can be found in [1] and [2]. These different imaging methods can be understood if we keep in mind a picture of what the signal versus time looks like for a solid sample compared with a liquid sample, as shown in figure 1.1. Here I will briefly discuss the various methods that have been pursued for imaging solids prior to this work and how our proposed method differs. Details of the relevant MR theory and experiments will be explored in later chapters.

1.1.1 Living with the Short T_2 of Solids

The most common approach to MRI of solids try to deal with the short decay time (T_2) of the solid signal (as shown by the dashed lines in figure 1.1) by changing experimental parameters and/or equipment to work around its inconvenience.

Single Point Imaging

Single point imaging (SPI) is the most common technique of imaging solids using MRI - because it is a fairly straightforward extension of conventional liquid-state MRI. SPI was first developed by Emid and Creyghton [3] and acquires a single point of the data for each experiment as soon as possible before the solid signal disappears (which is often $< 500\mu s$). Understandably, this method of imaging is very slow since only a single point of data is being taken for each experiment, one has to wait the long T_1 of solids before repeating experiments, and images require thousands of data points. Also, in order to get a nicely resolved image, very high gradient amplitudes and rapid on/off switching of gradients need to be used. To speed up imaging as well as to avoid the need for very fast gradient switching, the single point ramped imaging with T_1 enhancement

(SPRITE) method was developed by Balcom, et al. [4]. In this method, the gradient is ramped up in discrete steps so more points can be acquired in single experiment, and the need for on/off gradient switching is limited. Small-angle pulses are also used so one does not have to wait a full T_1 before repeating experiments. However, the short T_2 of the sample ultimately determines how well this imaging method functions.

Less Commonly Used Methods

Other imaging of solids methods have been developed which try to get better resolved images of solids by pushing the limits of the equipment and living with the short T_2 of solids. One such method includes using much larger gradients so that the effect of gradients can be seen on the rapidly decaying signal (as shown in the dashed, thin line in figure 1.1). Stray-field imaging (STRAFI) makes use of the fringe field of large superconducting magnets to provide the large static field gradients required. Initially, the sample would have to be mechanically moved through the fringe field to build up an image along one spatial dimension. More recently, this technique has been applied with pulsed gradients [5] to eliminate the need for moving the sample. However, this technique is still limited to planar samples where high-resolution is only desired along one spatial dimension.

Some techniques use completely different equipment when imaging solids than when imaging liquids. For example, continuous wave (CW) imaging makes use of continuous radio frequency (rf) irradiation and detection in the presence of continuously applied gradients while the external magnetic field is swept slowly through resonance [2]. Here one does not need fast gradient switching times or high-power rf, but one would need a customized system to sweep the magnetic field (instead of just using a large, static field with smaller magnetic field gradients).

1.1.2 Making Solids Look More Like Liquids

Unlike the previous methods, these methods aim to make the solid signal look more like liquid signal by extending the effective T_2 . The dashed lines then become more like the solid lines in figure 1.1, and conventional liquid-state imaging techniques can then be used. In order to make the solid-state signal more liquid-like, one must get rid of the dipolar interactions (H_{ZZ}) in solids which naturally average to zero in liquids due to molecular motion.

Magic Angle Methods

The ‘magic angle’ imaging methods were derived from NMR of solids techniques which eliminate or decrease the dipolar interactions in the solid - the primary cause for the short T_2 time. These techniques make use of the ‘magic angle’ where the dipolar interactions between two neighboring spins becomes zero.

The most common method, magic angle spinning (MAS), rapidly rotates the solid sample at the magic angle so the dipolar interactions, as well as the anisotropic Zeeman interactions, average out to zero [6]. The isotropic Zeeman interactions, however, remain unaffected. In order to use MAS for imaging, one must have rotor synchronized rotating gradients, which requires technically demanding hardware and limits on the sample one can use.

Another method that makes use of the magic angle is magic-angle rotating-frame imaging. Instead of mechanically rotating the sample, the dipolar and anisotropic Zeeman interactions are eliminated by evolving spins around an effective radiofrequency field in the magic-angle rotating frame [7]. However, this method also is technically demanding of the equipment, because very precise gradients are needed to give the magic angle condition. Due to the difficult set-up, thus far imaging has only been done in one or two dimensions using this method.

Pulsed Line-Narrowing Approaches

Other imaging of solids approaches derived from solid-state NMR use pulsed line-narrowing sequences to get rid of the effects of the dipolar interactions in solids. These sequences are usually made of multiple 90° pulses to generate effective fields that make the net dipolar interactions zero by using the fact that $H_{XX} + H_{YY} + H_{ZZ} = 0$. The most well-known pulse line-narrowing approaches are WAHUHA [8], MREV-8 [9], and BR-24 [10]. Magic-Echo Imaging makes use of the magic sandwich echo pulse sequence [11], which gets rid of the effect of dipolar interactions by reversing the sign of the dipolar Hamiltonian.

These line-narrowing pulse sequences work best when there are no other interactions other than dipolar interactions. These pulse sequences are less effective for samples where other interactions are also prominent, for example, Zeeman interactions (H_Z). In order to do imaging, gradients need to be applied in between the line-narrowing pulse blocks, so large amplitude gradients with very fast gradient switching are required.

1.1.3 Previous Work Imaging Bone

In this work, many of our results are imaging ^{31}P in bone mineral. In bone there are several different approaches imaging the short T_2 proton signal from collagen and tightly bound water in the organic bone matrix with a spatial resolution typically $>400\mu\text{m}$ but one recent article got resolution down to $56\mu\text{m}$ [12]. These approaches are various modifications of the SPI approach, where data is acquired as soon as possible before the solid signal disappears. These include WASPI developed by Wu, Ackerman, et al. [13, 14, 15], UTE developed by Bydder et al. [16, 17, 18, 19], SWIFT developed by Garwood et al. [20, 21], and ZTE developed by Weiger et al. which has given the highest resolution thus far for hydrogen in *ex vivo* bone [12]. Only a few groups

have reported direct ^{31}P imaging in bone, with one group reporting spatial resolution of 0.5mm (but 2-5mm is more typical) [22, 23, 24, 25, 26]. In all these methods, the spatial resolution is limited by the broad linewidth caused by the short T_2 of ^{31}P in bone ($< 200\mu\text{s}$). One group has used a line-narrowing approach for ^{31}P MRI in bone by using a ‘solid echo’ to partially refocus the dipolar interactions and lengthen the effective T_2 to a few times longer than the actual T_2 [27, 28]. This approach is limited by the ineffectiveness of the solid echo to completely refocus the dipolar Hamiltonian in bone mineral.

1.1.4 Our Method

Our method adopts the strategies of earlier line-narrowing approaches to the MRI of solids, where one makes the spin dynamics of the solid look more like a liquid. The key difference is that our pulse block uses a quadratic echo to refocus both the Zeeman (H_Z) and dipolar interaction (H_{ZZ}) terms in the spin Hamiltonian, and works best in the limit where $H_Z \geq H_{ZZ}$. Our pulse sequence is thus a better choice for certain samples (e.g. ^{31}P in bone mineral, where $H_Z \geq H_{ZZ}$), and there is less need for strong gradients and very fast gradient switching, since the gradients can be left on during the pulse burst.

1.2 Products of this Research

In this thesis work, our new approach to the MRI of solids was implemented on the Bruker 4 Tesla animal MRI system. We can do imaging using both Cartesian and radial sampling, and have successfully imaged ^{31}P in bone mineral in various *ex vivo* bone samples as well as ^{31}P in many types of soft tissue. We currently have the highest spatial resolution image of ^{31}P in bone mineral, and the first image of ^{31}P in soft tissue

where the signal is predominately from the cell membranes. Our lab also developed a new, faster, and computationally efficient algorithm to reconstruct images from sparse data to speed up imaging, and we currently are exploring its use beyond MRI.

1.3 Organization of the Remaining Chapters

The second chapter covers nuclear magnetic resonance (NMR) and MRI basics that will be very useful for understanding the challenges of MRI of solids and our approaches to overcoming these challenges.

The third chapter will briefly discuss the development of the quadratic echo line-narrowing sequence and various designs of the pulse sequence to be used particularly for MRI of solids.

The fourth chapter discusses in depth our application of the quadratic echo line-narrowing sequence to provide high spatial resolution MR images of solids. Here I also present our results imaging ^{31}P in bone mineral and soft tissue.

The fifth chapter discusses our approach to overcome the slow imaging time by making use of sparse sampling techniques and a novel algorithm to reconstruct high quality data from sparse data.

The sixth and final chapter discusses potential uses of our technique for MR imaging of solids, along with some preliminary data and final thoughts.

Chapter 2

NMR and MRI Basics

In this chapter, I introduce the theoretical and experimental knowledge of NMR and MRI that will be most helpful in understanding this dissertation work.

2.1 Basics of NMR

2.1.1 Larmor Frequency and Precession

Let's begin by looking at the simple quantum system of a single nuclear spin (with nuclear spin operator, \vec{I}) in the presence of a static magnetic field, \vec{B}_0 . This nuclear spin has an associated magnetic moment, $\vec{\mu}$ given by

$$\vec{\mu} = \hbar\gamma\vec{I} \tag{2.1}$$

where γ is the gyromagnetic ratio of the nucleus. The Hamiltonian for this system is given by,

$$H = -\vec{\mu} \cdot \vec{B}_0. \tag{2.2}$$

We will choose a coordinate system such that the static magnetic field will be in the z direction and only concern ourselves with spin- $\frac{1}{2}$ nuclei, so the allowed eigenvalues of I_z are $m_I = \pm\frac{1}{2}$. This results in the iconic two-level system, with energy splitting given by

$$\Delta E = \hbar\gamma B_0. \quad (2.3)$$

The distribution of nuclear spins between the two energy levels can be approximated by Boltzmann statistics. In thermal equilibrium, the number of spins in the lower energy level, $N_{m_I=+\frac{1}{2}}$ compared to the higher energy level, $N_{m_I=-\frac{1}{2}}$ is given by

$$\frac{N_{m_I=-\frac{1}{2}}}{N_{m_I=+\frac{1}{2}}} = \exp(-\Delta E/k_B T), \quad (2.4)$$

where k_B is Boltzmann's constant and T is the temperature in Kelvin. For nuclear spins, ΔE is so small compared to $k_B T$ for typical temperatures that there is only a very slight bias in the spin distribution at thermal equilibrium. This leads to a very small spin polarization, given by

$$P = \frac{N_{m_I=+\frac{1}{2}} - N_{m_I=-\frac{1}{2}}}{N_{m_I=+\frac{1}{2}} + N_{m_I=-\frac{1}{2}}}. \quad (2.5)$$

Typical polarization values range from 10^{-4} to 10^{-5} .

Photons with angular frequency equal to

$$\omega_0 = \gamma B_0 \quad (2.6)$$

will resonate with this two-level system and cause transitions between the levels. This special frequency is called the Larmor frequency, and note that it is directly proportional to the magnetic field felt by the nuclear spin. This is the essential point behind

MRI. We can encode spatial information into the measured frequencies by applying spatially-dependent magnetic fields (usually linear field gradients).

2.1.2 Classical Approach to NMR

The dynamics of the magnetic moments $\vec{\mu}$ of nuclei in the presence of a static magnetic field \vec{B}_0 can be easily understood through a simple classical analogue where the magnetic field exerts a torque on the magnetic moment, which we will treat as a simple bar magnet. This torque is given by the expression

$$\vec{\tau} = \vec{\mu} \times \vec{B}_0 \quad (2.7)$$

If the magnetic moment is allowed to rotate freely, it would tend to align itself with the applied magnetic field. However, our nuclear magnetic moments have intrinsic angular momentum, given by $\vec{J} = \hbar\vec{I}$. In classical mechanics, the change in the angular momentum must be equal to the applied torque. The equation of motion for the nuclear magnetic moment then becomes,

$$\frac{d\vec{J}}{dt} = \vec{\mu} \times \vec{B}_0. \quad (2.8)$$

Using $\vec{\mu} = \hbar\gamma\vec{I} = \gamma\vec{J}$, we can rewrite this equation of motion to be

$$\frac{d\vec{\mu}}{dt} = \gamma\vec{\mu} \times \vec{B}_0. \quad (2.9)$$

When $\vec{\mu}$ is not perfectly aligned with \vec{B}_0 , the magnetic moment moves in a direction perpendicular to both the direction of $\vec{\mu}$ and \vec{B}_0 . We can think of the magnetic moment of the nuclear spin moving on a cone with a constant angle (Φ) between $\vec{\mu}$ and \vec{B}_0 (see figure 2.1). The angular frequency of this precession is given by $\omega_0 = \gamma B_0$. Thus the

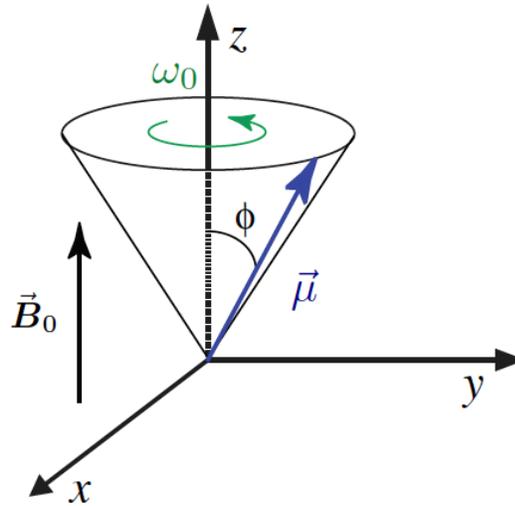


Figure 2.1: Depicts the precession of the magnetic moment $\vec{\mu}$ about the static magnetic field \vec{B}_0 on a cone at a fixed angle Φ measured from the z axis (direction of the applied magnetic field). Figure courtesy of Rona Ramos.

magnetic moment precesses around the applied magnetic field at the Larmor frequency we found in the quantum mechanical description.

In NMR we measure this precession frequency in order to learn about the local magnetic field environments of the nuclei we are observing. We measure this signal by using a solenoidal coil positioned perpendicular to \vec{B}_0 . The sum of the magnetic moments in a unit volume gives some total magnetization, \vec{M} . For example, if there is a net polarization of the spins in equilibrium, then the total \vec{M} is non-zero, and aligned along the applied field. Once \vec{M} is rotated away from the \vec{B}_0 direction by a pulse, equation 2.9 tells us that it will precess about \vec{B}_0 . We can measure the precession of \vec{M} by looking for the induced alternating voltage in our coil, due to the time-varying magnetic flux through the loops. We get the strongest signal when we have the largest \vec{M} precessing in the xy plane. The typical experiment then goes as follows: (1) Wait some time for spins to thermally equilibrate and polarize along \vec{B}_0 . (2) Use a 90° radio frequency (rf) pulse (often called the 'excitation pulse') that knocks

down the polarized spins into the xy plane. (3) Measure the induced voltage in the coil caused by the precession of the spin magnetic moments. This is a free induction decay (FID) experiment. Once the signal is acquired in the time domain, we take the fast Fourier transform (FFT) in order to get a frequency spectrum of the signal. If all spin magnetic moments are precessing at the same frequency, we see a very narrow spike at the Larmor frequency. Variations on this basic frequency spectrum tells us about the different local magnetic field environments of the spins we are observing.

Rotating Frame

To understand the effects of rf pulses on the spin magnetic moments, it is convenient to analyze the spin dynamics in the frame rotating at the Larmor frequency [29, 30]. Since the externally applied field, \vec{B}_0 , is much larger in scale than the applied pulses and residual internal fields of interest, it is very convenient to work in the frame that removes the rather trivial dynamics caused by \vec{B}_0 . In this rotating frame, spins precessing at the Larmor frequency now appear to stand still.

In order to reorient spins polarized along the z axis into the xy plane (to maximize signal measured in the xy plane), one can imagine applying a magnetic field (\vec{B}_1) in a direction perpendicular to z in this rotating frame. This static magnetic field in the rotating frame is generated in the lab frame by a linearly polarized alternating magnetic field caused by applying an alternating rf voltage through the solenoidal coil at the Larmor frequency. In the lab frame, the linear magnetic field produced by the coil can be regarded as two counter rotating fields with the same frequency. When viewed in the rotating frame, the component that rotates in the same sense as the rotating frame will look like a constant field, while the component that rotates in the opposite sense will appear to precess at $2\omega_0$. This counter-rotating field oscillates far enough away from resonance that it typically has little effect on the magnetic moments

and is therefore ignored.

The reorientation of the spins due to the magnetic field produced by the pulse, \vec{B}_1 , is called nutation. The amount of nutation depends on the pulse length (i.e. the time the rf alternating voltage is applied) as well as the pulse power (determined by the amplitude of the rf voltage and the quality factor of the coil circuit). We will denote the most relevant pulse times through the variables T_{90} and T_{180} , which gives the pulse length for a 90° and 180° nutations, respectively. One can also control the direction of the pulse in the xy plane by controlling the phase of the applied rf voltage. We often will denote a pulse using an expression such as 90_Y , which means the spins will be nutated 90° about the y axis in the rotating frame.

Spin Relaxation

In the simple system discussed above, the precession of the magnetic moment would go on indefinitely because there are no other sources for relaxation. In reality, there are many possible sources to cause these quantum magnetic moments to decohere from each other and eventually relax to the lowest energy state. These different relaxation processes are often characterized by two relaxation times, T_1 and T_2 .

The T_1 relaxation time characterizes the relaxation of the total magnetization \vec{M} back to the thermal equilibrium value for the system. This process involves spins transferring energy to the environment to relax to the lower energy level. Thus, the speed at which the magnetization along the direction of the applied magnetic field (M_z) returns to its equilibrium value depends on the mechanisms available to transfer energy from the spin system out to a possible energy reservoir (e.g. other nearby spin species, translations, rotations, and vibrations of atoms, etc.) For these reasons, T_1 is often referred to as the spin-lattice relaxation time. In an FID experiment, this relaxation

process is described by the simple equation,

$$M_z(t) = M_0(1 - \exp(-t/T_1)), \quad (2.10)$$

where M_0 is the longitudinal magnetization at thermal equilibrium and is given by

$$M_0 = \frac{1}{2} \hbar \gamma (N_{m_I=+\frac{1}{2}} - N_{m_I=-\frac{1}{2}}). \quad (2.11)$$

Thus, after an FID experiment, if you do not wait several T_1 's for full re-equilibration, the amplitude of the next FID signal (and the corresponding spectrum) will be reduced below the equilibrium value. As a result, experiments in long T_1 samples are either slow, suffer from poor signal-to-noise, or both.

The T_1 relaxation time is often measured using either an inversion recovery experiment (where spins are first initialized in the higher energy state through use of a 180° pulse and then allowed to relax back to thermal equilibrium) or saturation recovery experiment (where the net magnetization is scrambled to be approximately zero and then allowed to recover) [31].

The second relaxation time constant, T_2 , characterizes the decoherence of spins in the transverse (xy) plane due to interactions in which there is no net energy transfer from the spin system to the environment. One mechanism for T_2 relaxation comes from local fluctuations in the magnetic field at the site of the nucleus, often due to spin-spin dipolar interactions or interactions with the local electronic environment. Due to these fluctuations in the local microscopic magnetic fields, the nuclei precess at slightly different rates. As the spins get more and more out of step with each other, they fan out and point in all directions in the xy plane, causing the vector sum of the transverse magnetization measured in our coil to decay to zero. This decay in the time domain signal leads to equivalent broadening of the linewidth (proportional to $1/T_2$) in the

corresponding frequency spectrum.

Given our description of T_2 above, one plausible way to measure it is to look at the decay of the signal after an FID experiment. However, the measured decay is often faster than the pure T_2 decay that tells us only about interactions with local microscopic magnetic fields. This faster decay can be due to macroscopic field inhomogeneities in the static field \vec{B}_0 itself or magnetic susceptibility effects in the sample. We label this faster decay time measured in a FID experiment T_2^* . One way to get rid of the effects of these macroscopic field inhomogeneities is to start the FID experiment, but then add a 180° pulse some time τ after the excitation pulse which has the effect of refocusing the effects of inhomogeneous magnetic fields to produce an echo at 2τ after the excitation pulse. A depiction of this refocusing in the rotating frame is shown in figure 2.2. This sequence is called the Hahn echo sequence, and is the conventional way to measure T_2 by increasing τ and watching the subsequent decay in the echo peak amplitude [31, 30].

2.1.3 Relevant Spin Interactions

Here I will give the two relevant spin interaction Hamiltonians leading to T_2 decay that we will focus on in this work. The first is the Zeeman Hamiltonian, which is the interaction of the nuclear spin with the local magnetic field due to both the external field, \vec{B}_0 , and any local magnetic field fluctuations at nuclear site i , δB_i . In the lab frame, the Zeeman Hamiltonian has the form

$$H_Z^{\text{lab}} = \sum_{i=1}^N -\hbar\gamma(B_0 + \delta B_i)I_{z_i} \quad (2.12)$$

where I_{z_i} is the z -component nuclear spin operator for the i th nucleus. In the rotating frame, this Hamiltonian reduces to only a single term containing the small Zeeman

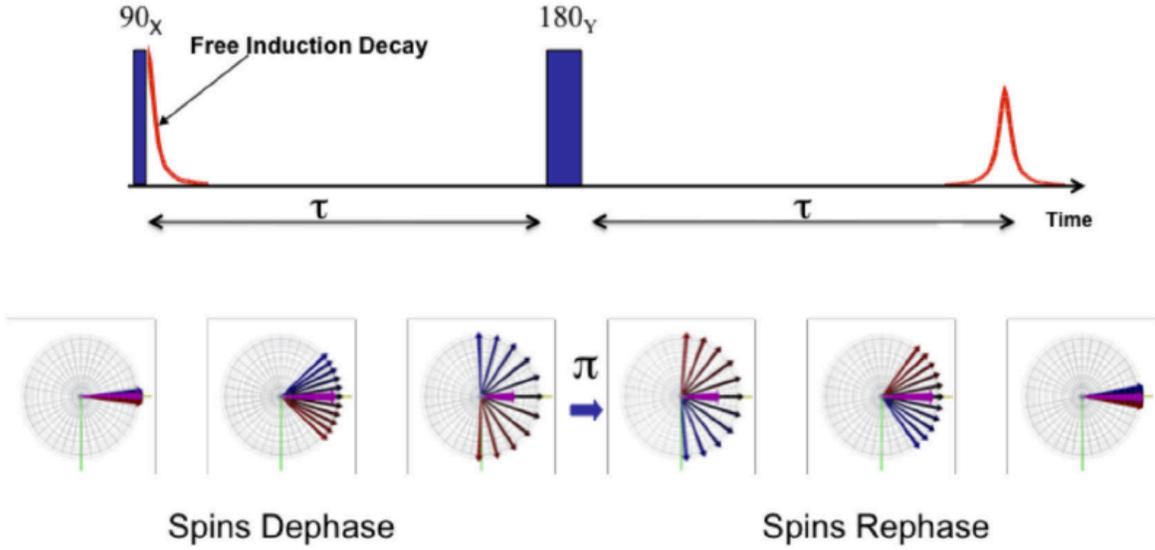


Figure 2.2: The Hahn echo pulse sequence and its resulting spin dynamics in the rotating frame. After an initial 90_X pulse, spins are precessing at frequencies evenly distributed around the Larmor frequency. Here, the black spin is on-resonance, while higher frequency spins are more blue, and lower frequency spins are more red. The resulting decoherence leads to signal decay in the measured FID (which can also be seen in the decreased size of the pink arrow, which is the total transverse magnetization, M_y). After the 180_Y pulse (or π pulse), the spins are flipped like a pancake about the y axis and continue moving with the same sense as before, leading to refocusing of the transverse magnetization. Increasing τ and watching the subsequent decay in the echo peak amplitude gives a measurement for T_2 . Figure courtesy of Yanqun Dong.

shift, $\Omega_{Z_i} = -\hbar\gamma\delta B_i$. For the systems we are dealing with, Ω_{Z_i} is the same for a large number of neighboring spins, so we can drop the index i . The Zeeman Hamiltonian in the rotating frame is then given by

$$H_Z = \sum_{i=1}^N \Omega_Z I_{z_i} = \Omega_Z I_{z_T}, \quad (2.13)$$

where $I_{z_T} = \sum_{i=1}^N I_{z_i}$ is the total I_z spin operator.

The second relevant interaction Hamiltonian is the homonuclear dipolar interaction term, which is the spin-spin interaction between nuclei of the same type. This interaction can be written in terms of the magnetic moments $\vec{\mu}_i$ and $\vec{\mu}_j$ correspond-

ing to nuclei at the i th and j th site, respectively. In the lab frame, the full dipolar Hamiltonian is

$$H_d^{\text{lab}} = \sum_{i=1}^N \sum_{j>i}^N \left[\frac{\vec{\mu}_i \cdot \vec{\mu}_j}{|\vec{r}_{ij}|^3} - \frac{3(\vec{\mu}_i \cdot \vec{r}_{ij})(\vec{\mu}_j \cdot \vec{r}_{ij})}{|\vec{r}_{ij}|^5} \right]. \quad (2.14)$$

In the rotating frame, this Hamiltonian can be greatly simplified [30] to give the secular dipolar Hamiltonian

$$H_{ZZ} = \sum_{i=1}^N \sum_{j>i}^N B_{ij} (3I_{z_i} I_{z_j} - \vec{I}_i \cdot \vec{I}_j), \quad (2.15)$$

where the dipolar coupling constant, B_{ij} , is given by

$$B_{ij} = \frac{1}{2} \frac{\gamma^2 \hbar^2}{r_{ij}^3} (1 - 3 \cos^2 \theta_{ij}). \quad (2.16)$$

Here r_{ij} is the distance between the i th and j th spins (with distance vector \vec{r}_{ij}) and θ_{ij} is the angle between \vec{r}_{ij} and \vec{B}_0 .

The rotating frame expressions for the Zeeman (H_Z) and dipolar (H_{ZZ}) Hamiltonians will be the form assumed throughout this dissertation.

2.2 Basics of MRI

MRI uses the well-known techniques of NMR along with applied magnetic field gradients - spatially varying magnetic fields - to encode spatial information about the sample into the measured signal. In 1973, Paul Lauterbur published the first paper [32] to show an MRI image (see figure 2.3). His sample was two test tubes of water inside a beaker of D_2O . By applying a magnetic field gradient ($\vec{G} \parallel \hat{r}$), the 1D frequency spectrum became a map of the total proton spin density along \hat{r} since the frequency at which the spins were precessing depended directly on where the spins were located in the magnetic field gradient. By taking many 1D spectra with the applied gradient in

different directions (in a 2D plane), Lauterbur could then infer the 2D spatial locations of his proton-dense test-tubes, giving him a 2D image.

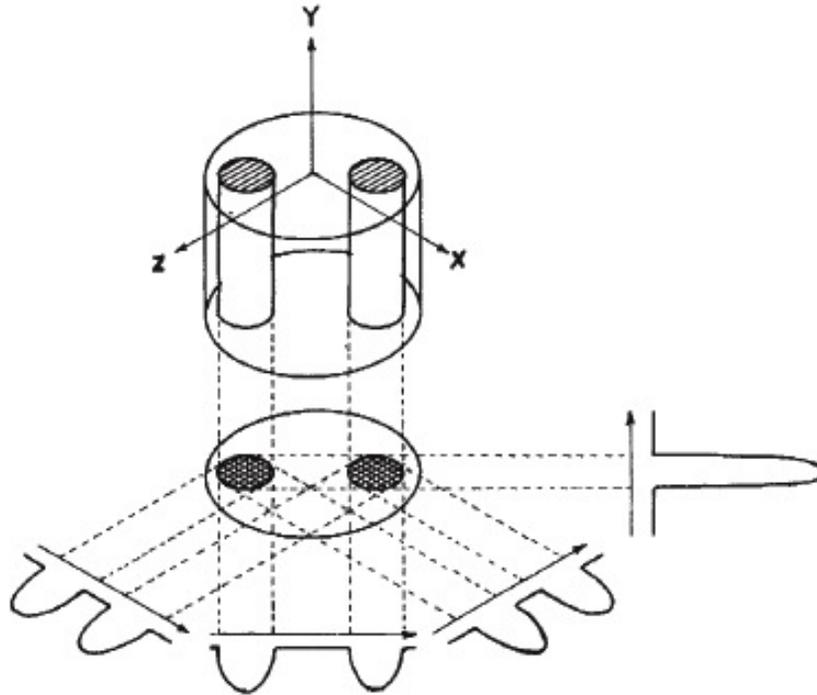


Fig.1 Relationship between a three-dimensional object, its two-dimensional projection along the Y-axis, and four one-dimensional projections at 45° intervals in the XZ-plane. The arrows indicate the gradient directions.

Figure 2.3: Figure from Lauterbur's original paper which had the first MRI image. This shows his approach to resolve a 2D image via back-projection.

Today, the main idea is essentially the same, but now we have more sophisticated ways to explain image acquisition and reconstruction, borrowed from solid-state physics and introduced by Peter Mansfield. The most important techniques pertinent to this work are explained in detail below.

2.2.1 k Space

The biggest element borrowed from solid-state physics is the idea of sampling \vec{k} space, which is the reciprocal space to the spatial image domain (i.e., \vec{k} space corresponds to the image space via Fourier transformation). While the data is still sampled in the time domain (as in NMR), a change of variables leads to the picture that in MRI we use magnetic field gradients to sample the \vec{k} space domain. How this is done can be easily understood by recalling the complex signal we are measuring in the rotating frame can be described by

$$S(t) \propto \int \vec{M}_\perp(\vec{r}) e^{-i\phi(\vec{r},t)} d^3r, \quad (2.17)$$

where $\vec{M}_\perp(\vec{r})$ is the total magnetization in the transverse plane in the neighborhood of point \vec{r} and $\phi(\vec{r}, t)$ is the phase factor picked up by that magnetization (in the rotating frame). This phase factor is given by,

$$\phi(\vec{r}, t) = \gamma \int_0^t \vec{r} \cdot \vec{G}(t') dt', \quad (2.18)$$

where $\vec{G} = (dB_z/dx, dB_z/dy, dB_z/dz)$ is the applied magnetic field gradient (with units of magnetic field over distance). Equivalently we can take this phase factor to be $2\pi\vec{k}(t) \cdot \vec{r}$ to write it strictly in terms of the conjugate variables we want to use for imaging (using the crystallographer's convention for defining the reciprocal lattice). That makes $\vec{k}(t)$ defined to be

$$\vec{k}(t) = \frac{\gamma}{2\pi} \int_0^t \vec{G}(t') dt', \quad (2.19)$$

and we can rewrite the measured signal in terms of \vec{k} as

$$S(t) \propto \int \vec{M}_\perp(\vec{r}) e^{-i2\pi\vec{k}(t)\cdot\vec{r}} d^3r. \quad (2.20)$$

As can be seen from equation 2.19, the integral of the gradient over time determines our current location in \vec{k} space.

2.2.2 Traversing \mathbf{k} Space

The \vec{k} -space trajectory is the path traced out by $\vec{k}(t)$. In NMR and MRI experiments, data are sampled at discrete time values. We are then only filling discrete \vec{k} -space points along this trajectory when we are actively acquiring (sampling) data, but we can still traverse \vec{k} space when not acquiring data. This leads us to the two ways of traversing \vec{k} space which can be used together in a variety of different ways to acquire the MRI data. The first is called phase encoding, where gradients are applied to move to a particular position in \vec{k} space, but no sampling of data has taken place. The second is called frequency encoding, where gradients are applied while acquisition is taking place, so we are sampling the magnetization at discrete \vec{k} space points along the \vec{k} -space trajectory. One does not need to use phase encoding to acquire the \vec{k} -space data, but many imaging sequences use it to choose an initial point to start the \vec{k} -space trajectory.

2.2.3 Sampling, Field-Of-View, and Spatial Resolution

Since we are making use of the Fourier transform to get the image from the acquired \vec{k} -space data, let's review the various relations between reciprocal spaces.

We are sampling N points in the time domain every δt (often called the dwell time) for a total acquisition time $\Delta t = N\delta t$. The corresponding frequency domain will have

N points with spacing $\delta f = 1/N\delta t = 1/\Delta t$ and a total bandwidth $\Delta f = N\delta f = 1/\delta t$. If we expect the spectrum to be symmetric, centered at $f = 0$ with a width L_f in Hertz, then the Nyquist-Shannon sampling theorem states we need to sample such that

$$\delta t \leq \frac{1}{L_f}, \quad (2.21)$$

otherwise the spectrum will be larger than the bandwidth and any signal outside the bandwidth will be aliased back inside the bandwidth. This is because the Fourier transform replicates the measured spectrum after every $\Delta f = 1/\delta t$, so if L_f is wider than Δf , then the replicates overlap inside the bandwidth (aliasing the results).

We can write the equivalent formulas to the ones above for \vec{k} space and image space. For simplicity, I will only give formulas below for one-dimension, but equivalent formulas exist for y and z directions as well. For uniformly-spaced frequency-encoded \vec{k} -space sampling along the x direction, the spacing between sampled \vec{k} -space points is given by

$$\delta k_x = \frac{\gamma}{2\pi} G_x \delta t, \quad (2.22)$$

where G_x is the gradient amplitude in the x direction. If we have a sample of width L_x in the x direction, then the Nyquist-Shannon sampling theorem becomes

$$\delta k_x \leq \frac{1}{L_x}. \quad (2.23)$$

The total field-of-view (FOV - the spatial extent of the image), Δx , is given by

$$\Delta x = \frac{1}{\delta k_x} = \frac{2\pi}{\gamma G_x \delta t}. \quad (2.24)$$

The spatial resolution of the image, δx , is given by

$$\delta x = \frac{\Delta x}{N} = \frac{1}{N\delta k_x} = \frac{2\pi}{\gamma G_x N \delta t}. \quad (2.25)$$

One strategy to boost the spatial resolution (i.e. decrease δx) is to increase the gradient amplitude, G_x . However, there is a maximum gradient amplitude allowed dictated by the Nyquist-Shannon sampling theorem. A similar limit exists for attempting to boost spatial resolution by increasing δt . The surest way to boost spatial resolution (without any limit due to the sampling theorem) is by taking more points N , while choosing a δk that is at the maximum value allowed by the Nyquist-Shannon sampling theorem.

Ultimately, even if one takes more and more points in each direction of \vec{k} space, the spatial resolution will be ultimately determined by the linewidth of the frequency spectrum of the sample. We can not resolve any features of the sample within this linewidth because we no longer are sure from where the measured signal originates. The natural linewidth of the spectrum is proportional to $1/T_2$ of the given sample, so one needs to increase the effective T_2 time in order to increase the possible spatial resolution.

A method commonly used to give (artificial) higher resolution, is to just add points to the raw \vec{k} -space data in all directions. These points are given the value zero, leading to the name of zero filling or zero padding. Occasionally, zero filling is used along with multiplication of a broadening function (usually a Gaussian or Lorentzian, whose decay is slow as to not greatly broaden the spectrum linewidth significantly). This is done to ensure the acquired data decays to zero within the \vec{k} -space points measured so zero filling does not add artifacts to the image (due to the signal suddenly becoming zero in the padded areas). Zero filling interpolates the image to add points in between

original image points and generally smooths the image as a result. Of course, zero filling is not adding any more information to the image, so this gives artificial higher resolution. The true spatial resolution is still determined by the measured data. If a broadening function is used, this true spatial resolution will actually be decreased somewhat because δx will be increased by the extra linewidth caused by the broadening function.

2.2.4 Exploiting Hermitian Symmetry

Another common strategy that we will employ in our image reconstruction is to exploit Hermitian symmetry to speed up data acquisition by a factor of two by only sampling half of \vec{k} space. We are able to do this because the complex image is expected to have no imaginary part (since it is a measurement of a real physical property, in our case the density of spins at a particular spatial location). If our $S(\vec{r})$ is real-only, that means its inverse Fourier transform, $S(\vec{k})$, has Hermitian symmetry about the origin of \vec{k} space. For us, that means our acquired \vec{k} -space data obeys the following relation,

$$S(-k_x, -k_y, -k_z) = S^*(k_x, k_y, k_z), \quad (2.26)$$

where $S^*(\vec{k})$ is the complex conjugate of $S(\vec{k})$. This means we only need half of \vec{k} space to fully reconstruct our real image. This can only be done on data that has no unwanted phase shifts due to hardware group delays, eddy currents, etc. Our imaging sequence gives us very good and accurate phasing of the data to produce a purely real image.

2.3 Equipment

Here I will briefly discuss the equipment we have used to perform the various experiments described in this dissertation work.

2.3.1 NMR Equipment

The NMR experiments in this work were performed on an Oxford Instruments Teslatron superconducting magnet with an 88-mm bore at 12 Tesla. We used a Tecmag Apollo spectrometer controlled via the NTNMR software. The signal was detected using home-built solenoidal coils matched to sample sizes and tuned to the desired Larmor frequency via a resonant tank circuit in order to boost transmitter pulse power and receiver signal sensitivity.

The tank circuit includes the coil with inductance L and two variable capacitors, C_T and C_M . Any resistance is very small and due to the wiring of the circuit. The tuning capacitor, C_T , plays the role of the capacitor in conventional RLC circuits and determines the resonant frequency of the circuit, $\omega_0 = 1/\sqrt{LC_T}$. In order to match the impedance of the circuit to the 50Ω cable attached to the spectrometer, a variable matching capacitor (C_M) is added to the circuit to control circuit impedance. The two tank circuits we use (depending on the desired tuning frequency) are given in figure 2.4.

2.3.2 Imaging Equipment

The MRI data was acquired on a Bruker Avance 4.0 Tesla/31 cm animal system running ParaVision 3.0.1. The 15-cm bore, actively shielded Magnex gradient coil set has a maximum gradient strength of 150mT/m for all three axes (although the largest

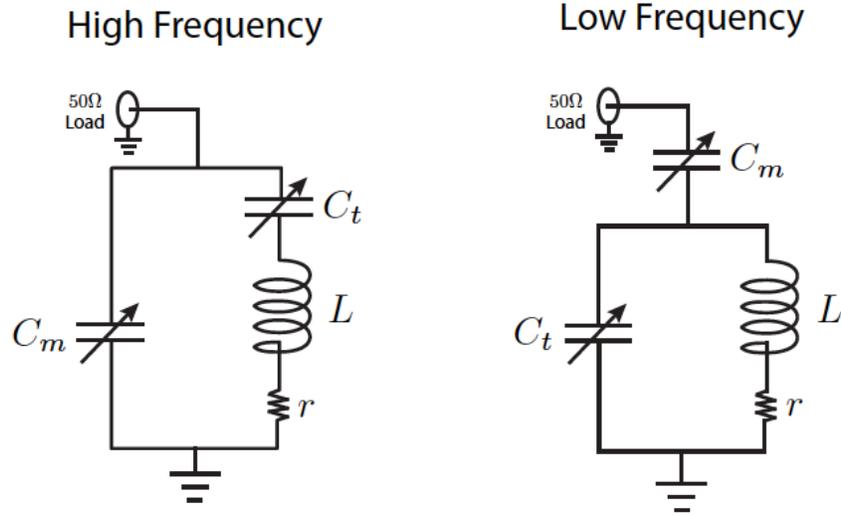


Figure 2.4: High frequency ($\omega_0 > 100\text{MHz}$) and low frequency ($\omega_0 < 100\text{MHz}$) circuits used for tuning the probe.

gradient magnitude we have used for imaging has been 30mT/m due to the need to confine the signal within a single octant of our 3D, eight-octant spectral FOV). Each gradient coil is connected to its own AE Techron 8607 gradient amplifier ($150\text{V}/130\text{A}$). The gradient ramp mode was set to “ramp off” to achieve the fastest possible transients.

The signal was detected using home-built solenoidal coils (matched to the sample sizes). These same coils were used to generate the rectangular pulses used in the imaging pulse sequence (with T_{90} in the range of $4\text{-}12\mu\text{s}$). The probe circuit was tuned using the low frequency tank circuit shown above. The imaging data were acquired stroboscopically using the Bruker’s analog acquisition mode. Short bursts of 5 complex points were sampled using an external dwell time of $2\mu\text{s}$ around the end of each pulse block. The built-in 125-kHz low-pass filter of the Bruker Avance system was supplemented by a second inline filter (Krohn-Hite Model 3940, dual-channel, 15-kHz low-pass Bessel) inserted just before the analog-to-digital converter.

2.4 Challenges of MRI of Solids

The main challenges between imaging ^1H in water and other nuclear isotopes - like ^{31}P , which our research will be focused on in this work - derives mainly from the basic differences between liquids and hard or soft solids. In MRI, the possible spatial resolution is going to be ultimately limited by the MR linewidth of the signal, which itself is inversely proportional to the effective T_2 of the sample. As was discussed above, T_2 is the characteristic time for spins to dephase in the transverse plane due to differing local magnetic fields. These differing local magnetic fields can be due, for example, to spin-spin (dipolar) interactions between the nuclei (which shall be denoted by H_{ZZ}), to chemical or magnetic shifts arising from variations in electronic configurations around each nucleus (denoted by H_Z), or various small time-dependent fluctuating fields which can be in any direction (denoted by $\delta H_i(t)$). In a liquid, the molecules move freely on fast time-scales, so the H_Z and H_{ZZ} terms average to $\langle H_Z \rangle$ and zero, respectively, leaving T_2 to be determined predominately by the small fluctuating fields, $\delta H_i(t)$. The resulting T_2 relaxation time in liquids is on the order of seconds. Immobile nuclei in a the crystalline lattice of a solid, on the other hand, feel the full effects of Zeeman (H_Z) and dipolar (H_{ZZ}) terms and, as a result, have much faster decay times, where T_2 is often less than a millisecond. This leads to much wider linewidths for solids and poorer spatial resolution when attempting imaging. See figure 2.5 for a good comparison of the linewidth difference between solids and liquids.

Of course one possible solution for short T_2 samples is to just repeat the experiment multiple times, taking a single point in each experiment, referred to as ‘single point imaging’ [3]. Here the longitudinal relaxation time T_1 plays a large role, since one typically waits longer than T_1 before repeating an experiment in order for the spins to polarize along the direction of the applied external magnetic field. Longitudinal

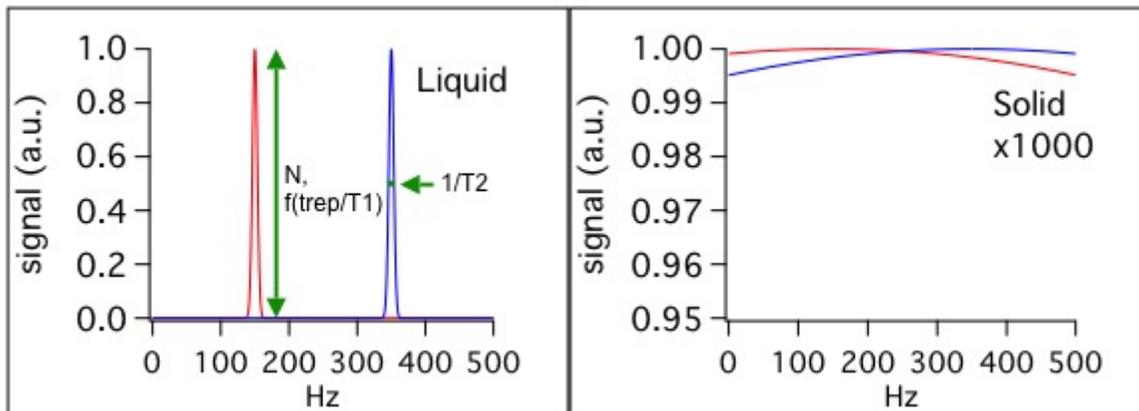


Figure 2.5: (*Left*) Example of two peaks located at 150Hz (red) and 350Hz (blue) with a typical liquid linewidth (≈ 8 Hz). The amplitude is normalized to one in this figure, but is experimentally determined by the number of nuclei in the sample (N) and a function of the wait time between experiments (t_{rep}) over T_1 . The FWHM is inversely proportional to the effective T_2 of the sample. (*Right*) The same two peaks are shown with a typical solid linewidth (≈ 5 kHz) and the signal has been multiplied by 1000 to have the same scaling as that of the liquid. The two peaks overlap badly and are virtually indistinguishable (note the y-axis starts at 0.95, not 0), leading to worse spatial resolution when imaging solids.

relaxation times can vary widely from sample to sample since they are determined by the fluctuating field term ($\delta H_i(t)$) in the spin Hamiltonian (with drier samples often having more lengthy T_1 's since fewer protons from water are present to help induce relaxation). Since one must wait longer than T_1 before repeating the experiment, using single point imaging on a living sample is often impractical due to the very long imaging times required.

Our sequence can get rid of the effects of the Zeeman (H_Z) and dipolar terms (H_{ZZ}) in the spin Hamiltonian, making the effective T_2 of the solid similar to that of a liquid, providing the narrow linewidths required for higher spatial resolution. In other words, we convert the ^{31}P spectra in bone from figure 2.5(*right*) to figure 2.5(*left*). Details about how this line-narrowing pulse sequence works are left for chapters 2 and 3. Our pulse sequence has no effect on T_1 , so imaging times can still be very long. However, we are still much faster than single-point imaging because we are able take many data

points for a given experiment before the signal decays. To speed up imaging further, we make use of sparse sampling techniques and image reconstruction algorithms, which will be discussed in chapter 4.

Chapter 3

Quadratic Echo Line-Narrowing and Imaging Pulse Sequences

The use of complex pulse sequences is very common in state-of-the-art NMR. Advanced pulse sequences can be used to control the most important terms in the spin Hamiltonian of the system, including the Hamiltonian term H_{ZZ} which describes the dipole-dipole interactions between spins and H_Z which describes the Zeeman shifts due to the local magnetic fields felt by the nuclei. Typically, pulse sequences cannot control both H_Z and H_{ZZ} at the same time. The pulse sequences developed in our group differ from existing NMR pulse sequences in several ways. For example, most prior pulse methods to control the dipolar term in the spin Hamiltonian assume the conventional Dirac-Delta approximation for pulses and work best for spin systems with small Zeeman terms. Our technique uses the quantum dynamics inside of the pulses themselves, and it works well in the regime where both Zeeman and dipolar interaction terms are relevant [33, 34, 35].

In this chapter I discuss the line-narrowing pulse sequences designed by our lab to reverse signs of certain terms of the spin Hamiltonian - either the Zeeman off-resonance

term, the dipolar coupling interaction term, or both - so that measurements can be made when certain terms of the Hamiltonian have acted, and others, notably the interaction terms, have not. This control over the effect of interactions on the time evolution of the spin system can lead to very long nuclear coherence times - up to 70,000 times longer [34] in one pure silicon sample - and can theoretically be used for any spin-1/2 system with a similar Hamiltonian, whether it be a system of qubits or phosphorus nuclei in a chunk of bone.

Section 3.1 is primarily a summary of the doctoral work by Dale Li [33], Yanqun Dong [34, 36], and Rona Ramos [35], which led to the development of the quadratic echo line-narrowing sequence. I will continue using ‘we’ and ‘our’ in this section, but this refers to the work and ideas of Dale, Yanqun, and Rona based on my discussions with them and the content of their theses and papers. Section 3.2 discusses imaging techniques developed by Yanqun and Rona, which were then further tested and perfected by undergraduates Ben Deen and Suyog Bhandari. I used these techniques to study ^{31}P in bone, and all the data shown are my results from these experiments. The last section (3.3) and the remaining chapters focus on my doctoral research, so ‘we’ will refer to my own personal experiences.

3.1 Brief History of the Development of the Line-Narrowing Sequence

In April 2001, the Barrett lab set out to do basic NMR measurements of phosphorus spins in silicon to test the applicability of Bruce Kane’s seminal proposal [37] to build a quantum computer using spins in semiconductors. The first goal was to measure the transverse relaxation time, T_2 . This is conventionally measured using a spin echo experiment which refocuses the Zeeman terms in the spin Hamiltonian and leaves only

the dipolar terms. The most common spin echo experiment is the Hahn spin echo experiment, where one excites the spins with a 90° pulse, waits some time delay τ , then applies a 180° pulse, and after another delay τ the peak of the echo can be acquired (see figure 2.2). One then repeats the experiment multiple times using a series of τ values and looks at how the peak of the echo decays with increasing τ to learn about relaxation due to dipolar terms in the spin Hamiltonian. A faster way to get this decay (without having to repeat so many experiments) is to keep on applying 180° pulses while acquiring the signal of the echo after each pulse. This is called the Carr-Purcell-Meiboom-Gill (CPMG) spin echo pulse sequence, and can be written as: $90_X - \{\tau - 180_Y - \tau\}^n$, where acquisition is made at the end of the second τ period before repeating the block in the curly brackets. The decay of the CPMG echo train is another conventional measure of T_2 [30], and the case $n = 1$ is the Hahn echo sequence.

In order to test the accuracy of the T_2 measuring methods before looking straight at the phosphorus nuclei in phosphorus-doped silicon, the lab switched from observing the ^{31}P nucleus to the more abundant spin-1/2 isotope in the sample - the ^{29}Si nuclear spins of the host lattice. Almost immediately, the lab stumbled upon a very surprising result. The CPMG sequence was giving a very slowly decaying signal that lasted much longer than the Hahn echo signal decay (see figure 3.1). When one treats the pulses as perfect, delta-function 180° -pulses, the two experiments ought to give exactly the same results. In fact, the Hahn echo signal agreed very well with the expected decay (black curve), but we saw *more* signal than theory predicts in the CPMG measurements. Similar results have been found in C_{60} and Y_2O_3 samples [38, 39, 40, 41, 33].

A reasonable first guess to explain this puzzle is that the CPMG experiment suffered from imperfections (e.g. 180° -pulse errors, phase transients, spatial inhomogeneity of the pulses, etc.), but after changing various experimental parameters and greatly

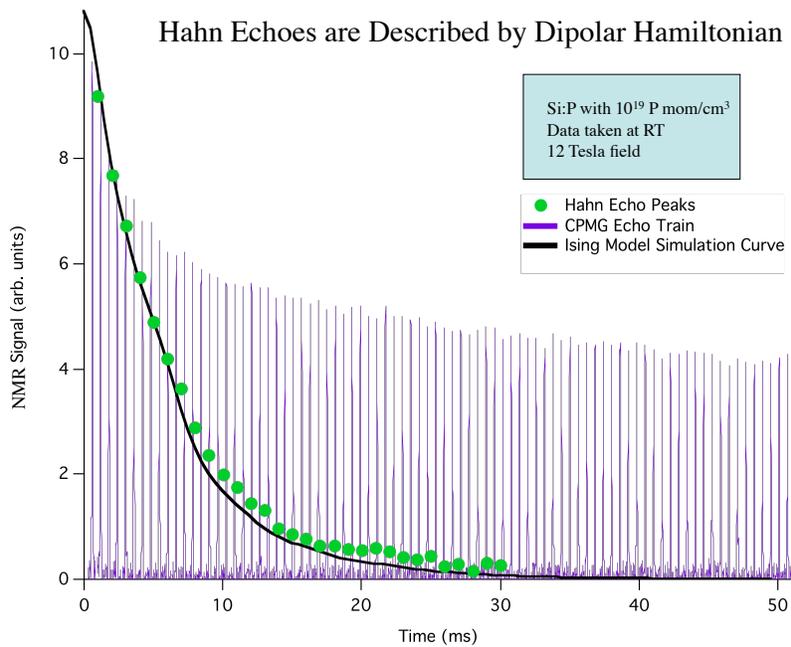


Figure 3.1: Two ^{29}Si NMR experiments to measure T_2 in doped silicon disagree sharply. Each of the Hahn echo peaks (green dots) are generated with a single 180° pulse, and they track the homonuclear dipolar decay (black curve) calculated in the Ising model limit. On the other hand, the train of CPMG echoes (purple) are generated with multiple 180° pulses spaced with delay $2\tau = 592\mu\text{s}$, and they decay much more slowly, after many 180° pulses are applied. Figure adapted from Dale Li's work [33].

improving the pulses, this discrepancy remained. Clearly, something strange happens when many 180° pulses are applied versus just a single 180° pulse. This forced the lab to consider a possible intrinsic effect - real pulses are never Dirac-delta functions (since they have some non-zero duration). 50 years worth of NMR experiments have used the Delta-function approximation for pulses without manifesting the large discrepancies seen here. We also were using very ‘hard’ (high-power, short-duration) pulses that deviates only slightly from the Dirac-delta approximation. However, the lab eventually realized that this intrinsic effect is not a random error, but is instead a coherent effect. Thus, for pulse sequences with many pulses of particular phases and spacing (like CPMG), these tiny, coherent ‘errors’ can add constructively to produce surprisingly large effects on our observables. But the question still remained if the lab could actually understand the physics of the system well enough to make use of this knowledge. For these answers, the Barrett lab turned to the approximate analytic expressions offered by average Hamiltonian theory (also known as coherent averaging theory).

3.1.1 Average Hamiltonian Theory

Here I will briefly introduce coherent averaging (now more commonly called average Hamiltonian theory), which the lab used to give an analytic approximation of pulse sequences with finite pulses. See the theses of Yanqun Dong [36] and Rona Ramos [35] for a far more in-depth description of these calculations and the development of the quadratic echo line-narrowing sequence.

Average Hamiltonian theory is a very useful analytical tool to describe systems evolving under the influence of a time-dependent periodic perturbation [42, 43, 43]. For our case, the perturbation is due to the multiple-pulse block being applied periodically. One chooses an interaction frame (also known as the toggling frame) such that this perturbation does not explicitly appear in the resulting average Hamiltonian. We first

write the total Hamiltonian as a sum of a time-independent and a time-dependent (perturbing) term:

$$H = H_0 + H_1(t), \quad (3.1)$$

where H_0 is the unperturbed Hamiltonian and $H_1(t)$ is the perturbation. The general time evolution operator is give by:

$$U(t) = \mathbf{T} \exp \left[-\frac{i}{\hbar} \int_0^t dt' (H_0 + H_1(t')) \right], \quad (3.2)$$

where \mathbf{T} is the Dyson time ordering operator. In order to separate the effects of H_0 and $H_1(t)$, we can divide the general time evolution operator into two factors such that,

$$U(t) = U_1(t)U_0(t), \quad (3.3)$$

where

$$U_1(t) = \mathbf{T} \exp \left[-\frac{i}{\hbar} \int_0^t dt' H_1(t') \right] \quad (3.4)$$

$$U_0(t) = \mathbf{T} \exp \left[-\frac{i}{\hbar} \int_0^t dt' \tilde{H}_0(t') \right] \quad (3.5)$$

and $\tilde{H}_0(t)$ is the Hamiltonian in the toggling frame with respect to $H_1(t)$ given by:

$$\tilde{H}_0(t) = U_1^{-1}(t)H_0U_1(t). \quad (3.6)$$

If we require the time-dependent perturbation Hamiltonian, $H_1(t)$, to be periodic with period t_c (the cycle time),

$$H_1(t + nt_c) = H_1(t) \text{ for } n = 0, 1, 2, \dots \quad (3.7)$$

as well as cyclic so that

$$U_1(t_c) = 1, \quad (3.8)$$

H_1 has no direct effect after one full cycle. This simply means that the toggling frame returns to its original orientation (and \tilde{H}_0 returns to its original value) after each t_c . Provided we only look stroboscopically at interval times of t_c , the general time evolution operator is then given simply by

$$U(nt_c) = U_0(t_c)^n. \quad (3.9)$$

It is convenient to express this time evolution in terms of a single average Hamiltonian \bar{H} such that

$$U(t_c)^n = \exp \left[-\frac{i}{\hbar} \bar{H} n t_c \right]. \quad (3.10)$$

This average Hamiltonian can be expanded using the Magnus expansion to be:

$$\bar{H} = \bar{H}^{(0)} + \bar{H}^{(1)} + \bar{H}^{(2)} \dots \quad (3.11)$$

where the first two terms (which are the only ones typically used) are given by:

$$\bar{H}^{(0)} = \frac{1}{t_c} \int_0^{t_c} dt \tilde{H}_0(t) \quad (3.12)$$

$$\bar{H}^{(1)} = -\frac{i}{2\hbar t_c} \int_0^{t_c} dt_2 \int_0^{t_2} dt_1 [\tilde{H}_0(t_2), \tilde{H}_0(t_1)]. \quad (3.13)$$

The commutators between Hamiltonian operators at different times involved in the higher-order terms become smaller for shorter cycle times t_c , so a faster multiple-pulse sequence in general will lead to better averaging where one can disregard the effects of higher order terms. In practice for our pulse sequences, $H_1(t)$ consists of a small number of piecewise constant time periods of pulses or free evolution. See references

[36] and [35] for detailed calculations of the average Hamiltonians for various pertinent pulse blocks.

3.1.2 Testing Our Model

The Barrett lab used average Hamiltonian theory to take into account the effects of finite pulses and were able to understand the results of many simple multi-pulse sequences. For example, the experimental results for CPMG depend greatly on the phases of the 180° pulses used, as shown in figure 3.2*a*. Here the gray signal comes from using only 180_Y pulses (CPMG), and the blue signal comes from alternating between 180_Y and 180_{-Y} pulses (APCPMG). In the Dirac-delta limit of pulses, these two pulse sequences should be equivalent. However, when taking into account the finite width of the pulses, average Hamiltonian theory shows that the the APCPMG pattern has an extra transverse field term, $-\lambda\Omega_Z I_{X_T}$ where λ is a dimensionless constant and Ω_Z is the Zeeman energy shift such that $H_Z = \Omega_Z I_{Z_T}$. This extra term produces effects similar to T_2^* dephasing in the xy plane for the normal FID, except that in this case it dephases magnetization in the yz plane. In order to calculate the average Hamiltonian terms for CPMG and APCPMG, we took advantage of the periodic nature of the pulses where CPMG is just repeated units of $\{Y, Y\} = (\tau - 180_Y - 2\tau - 180_Y - \tau)$ and APCPMG is repeated units of $\{Y, -Y\} = (\tau - 180_Y - 2\tau - 180_{-Y} - \tau)$.

Since we have identified the effective transverse field term leading to a ‘FID’ of the CPMG echo train, it is natural to try to manipulate it in order to make an ‘echo of the echo train’ by inserting a single 180_Y pulse (see figure 3.2*a*) or by reversing the APCPMG pattern for the second half of the acquisition (see figure 3.2*c*). As our model predicted, these two changes to the sequence give the same exact results. These results also suggest that these observations cannot be due to pulse error, because we add one pulse in one case and change many pulses in the other. As further test of the model,

adding a 180_X pulse does not do anything because the effective field term is in the x-direction (figure 3.2*b*). After playing with a single ‘echo of echoes’, one can then try to extend this idea to the ‘CPMG of the echoes’, which was done in figure 3.2*d*.

This shows that the approximate average Hamiltonian model works very well at telling us what to analytically expect from simple pulse sequences and the effective fields due to the finite nature of the pulses. The experiments above used just the zeroth order of the average Hamiltonian expansion to control the Zeeman interactions. Next the lab tried to manipulate the dipolar interactions as well, as shown in figure 3.3. Here we were inspired by the magic echo experiments [44], where the dipolar interaction is rephased through use of a long rf pulse. In our variant on the original magic echo experiment, the role of the long rf pulse is played by the effective transverse field term that appears in the zeroth-order average Hamiltonian. Here both Zeeman and dipolar phases are refocused using a well-placed 180_Y pulse and a negative dipolar-Hamiltonian term in the pulse burst. This complexity explains why the green echo has a larger amplitude than the blue echo in figure 3.3*d*, despite the fact that it forms later in the free evolution period. This control over both Zeeman and dipolar phases is very different from the original magic echo, which works best if the total Zeeman energy shift due to both internal and external Zeeman spin Hamiltonians ($\Omega_Z^{\text{net}} = \Omega_Z^{\text{int}} + \Omega_Z^{\text{ext}}$) is zero.

3.1.3 Designing the Quadratic Echo Line-Narrowing Pulse Sequence

As can be seen from the echo experiments in figures 3.2 and 3.3, we can take advantage of the transverse field term appearing in the zeroth order of the average Hamiltonian

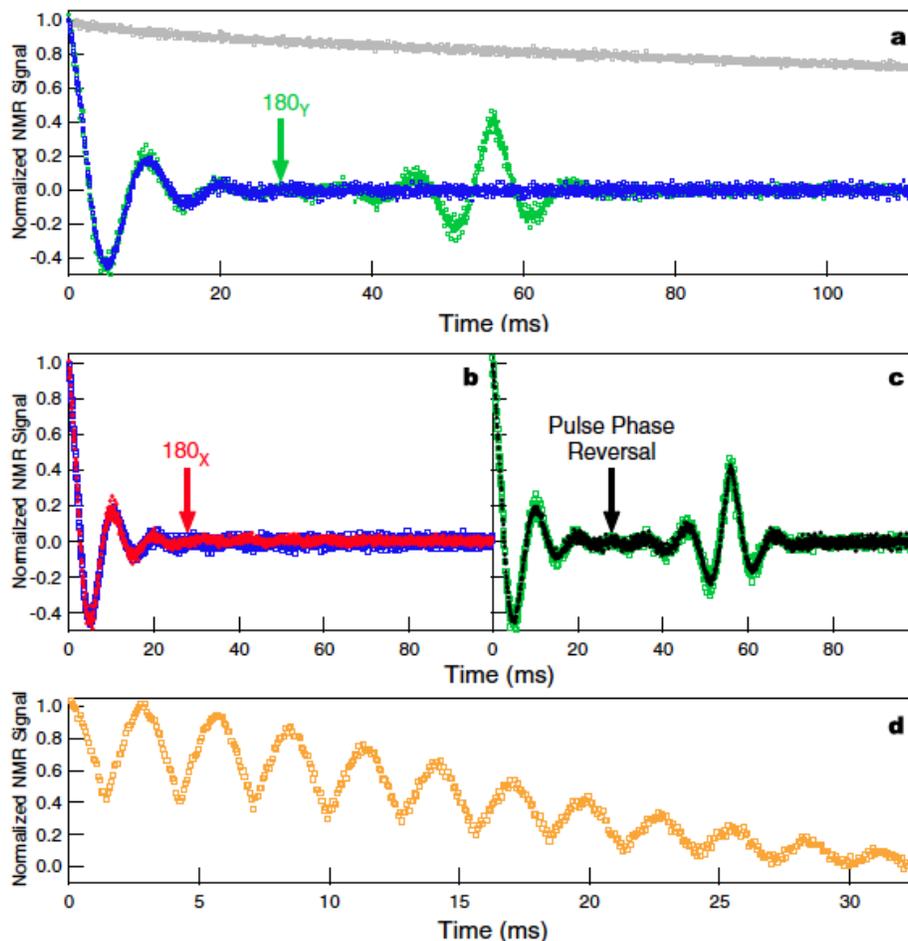


Figure 3.2: Data for ^{13}C nuclei in C_{60} at room temperature and 12 Tesla where each point here is the peak of an echo, $\tau = 25\mu\text{s}$ and the signal is normalized to the amplitude of the C_{60} FID. (a) Comparison of the CPMG sequence (gray) to the APCPMG sequence (blue). Inserting a single 180_Y pulse into APCPMG (green) induces an echo of the echo train. (b) Inserting a single 180_X pulse into APCPMG (red) has no effect. (c) Reversing the APCPMG phase pattern (black) has the same effect as inserting a single 180_Y pulse (green). (d) A CPMG of the echo train is induced by using $90_X\{-Y, Y\}^{10}(\{Y, -Y\}^{20}\{-Y, Y\}^{20})_{\text{repeat}}$ where $\{-Y, Y\}^N$ represents the sequence $(\tau - 180_Y - 2\tau - 180_Y - \tau)$. From reference [34].

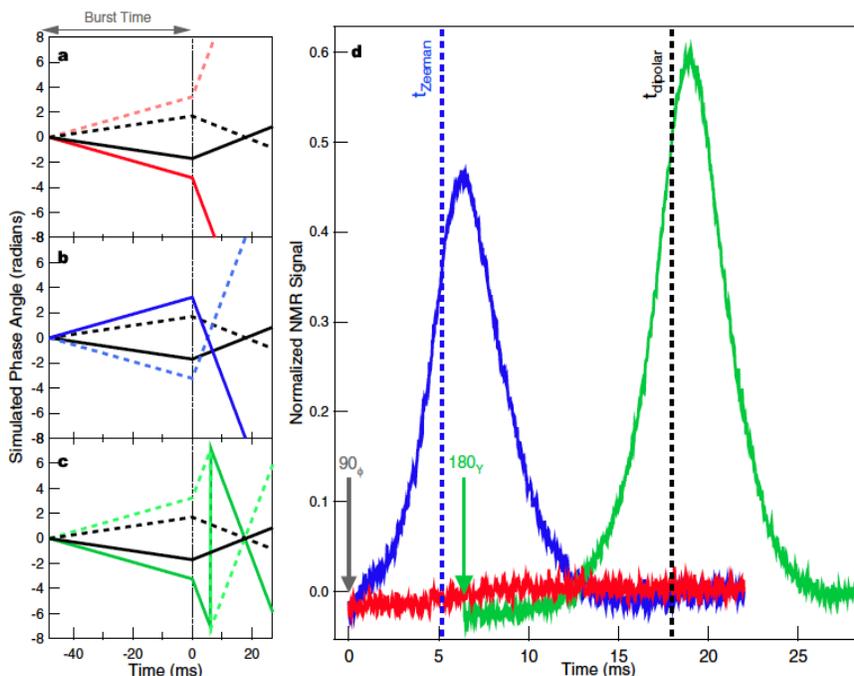


Figure 3.3: Three experiments inspired by the magic echo using the C_{60} sample at room temperature and 12 Tesla. All experiments start with $\{-X, X\}^{200}$ using $\tau = 50\mu s$, but have distinctly different results. (a)-(c) Simulations of the phase angle evolution with the dipolar phase ($\Phi_{ZZ}(t)$) in black and the Zeeman phase ($\Phi_Z(t)$) in the color of the experimental trace. The crossing of the lines gives the expected location of an echo. In (a), when a 90_{-X} follows the repeated block, the $\Phi_Z(t)$ lines fail to cross and the large Zeeman dephasing spoils the dipolar echo, so no echo appears (red in (d)). In (b), when a 90_X follows the repeated block, the $\Phi_Z(t)$ lines do cross and an echo is observed (blue in (d)) close to the expected t_{Zeeman} location (blue dashed line in (d)). Large Zeeman dephasing spoils any later dipolar echo from appearing. In (c), the failed sequence in (a) is repaired by applying a 180_Y pulse at a time t_f after the burst which forces the $\Phi_Z(t)$ and $\Phi_{ZZ}(t)$ lines to cross simultaneously, resulting in an optimized echo (green in (d)) near the expected location, $t_{dipolar}$. From reference [34].

of the ‘alternating phase’ $\{\phi, -\phi\}$ family of pulse sequences. When these experiments worked, the lab wondered if we could make use of the transverse field that appears in the ‘same phase’ $\{\phi, \phi\}$ family of pulse sequences (like CPMG $\{Y, Y\}$ or CP $\{X, X\}$). However, our first attempts to make use of CP failed miserably when we found that the signal was quickly destroyed by the ugly terms in the $\bar{H}^{(1)}$ term that we had hoped to ignore. However, inspired by Solomon’s rotary echo experiment [45], we tried using an interval of CP followed by an equal interval of flip-CP $\{-X, -X\}$, to cancel out the effects of the unwanted $\bar{H}^{(1)}$ term. This composite block can then be approximated by a fairly simple analytic expression for the average Hamiltonian [34, 36, 35],

$$-\frac{H_{XX}}{2} - 2\kappa^2 \Omega_Z^{\text{int}} \Omega_Z^{\text{ext}} I_{X_T}, \quad (3.14)$$

where Ω_Z^{int} is any internal Zeeman shift (e.g., due to sample diamagnetism, offset from the magnet’s isocenter, and the chemical shift) and Ω_Z^{ext} is any external offset imposed by the experimentalist (e.g., due to the applied magnetic field gradient and any pulse frequency offset). We have shown [34] that this can be well approximated further (through second averaging in the toggling frame of the field given by the second term) by $-\frac{1}{2}H_{XX}$. This composite block can then be used to refocus the dipolar Hamiltonian. Since the first-order effective Hamiltonian term providing this refocusing arises from a commutator, it is proportional to $(\Omega_Z^{\text{net}})^2$ and these two-spin interactions can be used to cancel the two-spin dipolar interaction. To illustrate this difference compared to other pulse sequences (which make use of effective Hamiltonian terms linearly proportional to Ω_Z^{net}), we named this pulse sequence the quadratic echo.

Figure 3.4 shows a simple version of the line-narrowing pulse sequence which uses the quadratic echo to refocus the (time-independent) dipolar Hamiltonian and a hidden 180_Y to refocus the (time-independent) Zeeman Hamiltonian at the end of the 6Δ time

period. Here the quadratic echo pulse block is surrounded by 90_{-Y} ‘external wrapper’ pulses which have, in the delta-function approximation, the effect of rotating the H_{XX} term given above to H_{ZZ} and also provides the hidden 180_Y pulse (because the second 90_{-Y} can be thought of as the 90_{+Y} required to ‘rotate’ H_{XX} into H_{ZZ} , followed immediately by a 180_{+Y} , which provides the ‘flip’ of H_Z , leaving H_{ZZ} alone). As a result, if we are only acquiring at the beginning and end of the pulse block, the spins should see no net time evolution due to Zeeman and dipolar spin Hamiltonians. This is then essentially a ‘time-suspension’ sequence that leads to extreme line-narrowing.

It is important to note that we are controlling both Zeeman (H_Z) and dipolar interactions (H_{ZZ}) in order to achieve extreme line-narrowing. This is unlike other line-narrowing methods, which only work in the regime dominated by either dipolar interactions or Zeeman interactions. Our technique fills the gap to work with samples where both Zeeman and dipolar interaction terms are relevant, ($H_Z \geq H_{ZZ}$). This is particularly helpful for imaging, where magnetic field gradients are applied to encode spatial information in the frequency of the signal via non-zero Zeeman interactions with the nuclear spins. Other line-narrowing pulse sequences (e.g. MREV or magic echo) must turn off the gradients during the pulse block for their pulse sequence to be most effective.

The effectiveness of this line-narrowing sequence can be seen in figure 3.5, where the effective T_2 of ^{29}Si in a sample of silicon powder doped with antimony was pushed all the way out to nearly $1/3 T_1$. This narrowed the spectrum by a factor of nearly 70,000. Once we have the ability to turn off both the Zeeman and dipolar terms in the internal spin Hamiltonian, we can then turn on other terms to design a spin Hamiltonian to achieve a specific experimental goal. For MR imaging, we would like spins to ‘see’ the external Zeeman Hamiltonian from applied magnetic field gradients,

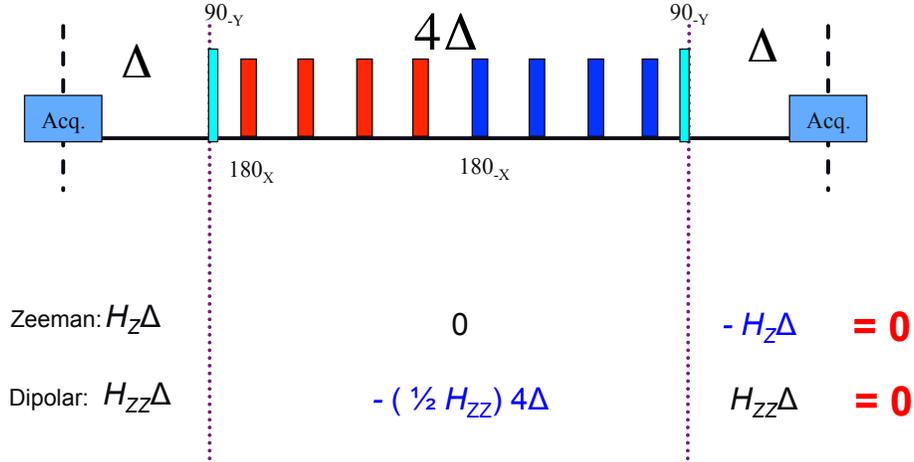


Figure 3.4: Depiction of the quadratic echo line-narrowing pulse block. The pulse block is made up of three time periods between consecutive acquisition periods (labelled ‘Acq.’). The Zeeman (H_Z) and dipolar (H_{ZZ}) terms of the spin Hamiltonian acting over each of the three time periods is given beneath the pulse sequence. The net effect of both Zeeman and dipolar Hamiltonian terms over the entire $6\Delta + T_{180} \approx 6\Delta$ of the pulse block is zero, leading to time-suspension or extreme line-narrowing.

without the internal Zeeman and dipolar Hamiltonians broadening the spectrum. In the next section, we talk about the various ways of implementing this line-narrowing sequence particularly for use in MRI.

3.2 Imaging Sequences

In order to do imaging, along with using line-narrowing to get rid of unwanted internal spin Hamiltonians, we also want our sequence to leave alone the effects of any externally applied Hamiltonians (e.g., due to magnetic field gradients or a pulse frequency offset). For these externally applied Hamiltonians to not be cancelled by our line-narrowing pulse sequence, we must change the sign of the externally applied fields in sync with our pulse sequence. Notably, if the externally applied gradient or offset is positive in the first Δ of the pulse sequence, than it should be negative in the final Δ in order to not be refocused by the hidden 180_Y pulse at the end of our pulse burst (see figure 3.4).

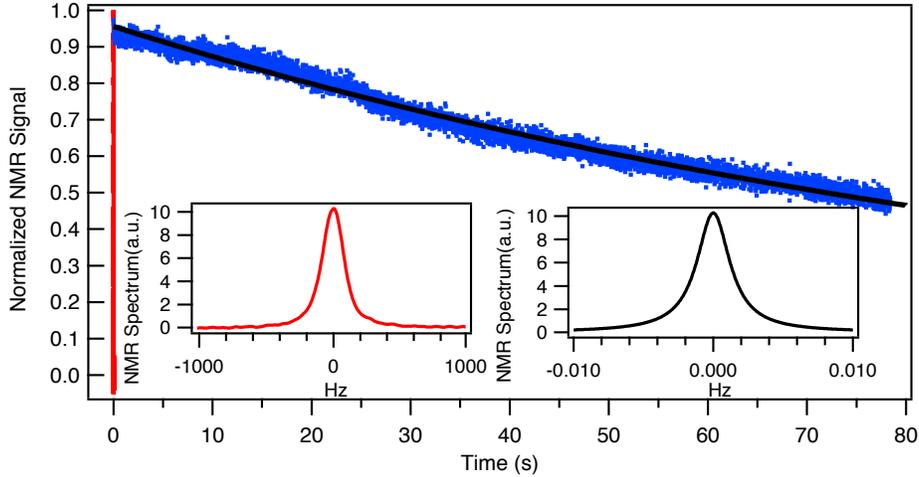


Figure 3.5: The ^{29}Si time-suspension (line-narrowing) data (blue) on a sample of Si:Sb using the sequence $90_X\{2, 0, -Y, -Y\}^{84000}$, with $\tau = 60\mu\text{s}$, $\nu_{\text{offset}} = 2.5\text{kHz}$, and the corresponding fitting curve (black) extend far beyond the normal ^{29}Si FID with $\nu_{\text{offset}} = 0\text{Hz}$ (red). (inset) The 200Hz normal spectrum (red) is narrowed to 0.003Hz (black, Fourier transformation of the fitting curve), centered at ν_{offset} , this is a line-narrowing by a factor of nearly 70,000.

Of course, there are many different possible patterns to do this, and further below we explore a few of the possibilities. Before exploring these various versions of our imaging pulse sequence, I first discuss the modification of our first-generation line-narrowing pulse block to be better suited for imaging experiments.

3.2.1 Second Generation Pulse Block

For conventional high-resolution imaging, one would like to increase the sampling dwell time to reduce the spectral bandwidth to ‘zoom-in’ on a smaller spatial area for the image FOV. In our case, the ‘dwell time’ is now replaced with what we call the ‘sparse dwell time’ which is the time between acquisition windows. For the first-generation quadratic echo line-narrowing pulse block given above, this sparse dwell time is equal to $6\Delta + T_{180}$ where $\Delta = 2\tau + T_{180}$ using the quadratic echo line-narrowing pulse sequence given above. One can increase the sparse dwell time by simply increasing τ , but this

also reduces the effectiveness of the pulse sequence itself. (It is also worthwhile to note that one can use a τ that is too short so that the transverse field terms we are utilizing can become comparable to the large external magnetic field for large offset frequencies. This causes spins to precess out of the xy-plane and reduces the measured signal. We typically keep τ comparable to T_{180} or larger.) However, we found that we can get essentially the same effectiveness with double the sparse dwell time by keeping a short τ and using two composite blocks of CP and flip-CP with added ‘internal wrappers’ $90_{\pm X}$ pulses (see figure 3.6). The sparse dwell time then becomes $6\Delta + 2T_{180}$ where $\Delta = 2(2\tau + T_{180})$.

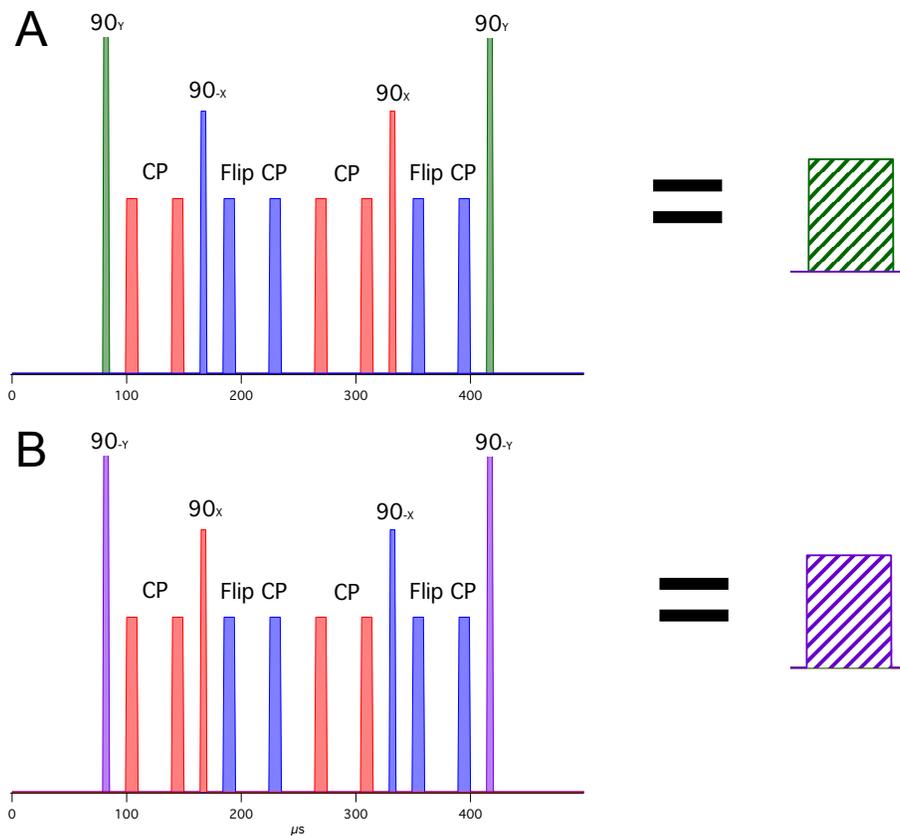


Figure 3.6: Versions of the second generation line-narrowing pulse block with external (tallest height) and internal (medium height) 90° wrapper pulses with different phases. (The heights of the pulses were changed to emphasize the locations of the wrapper pulses; in application, all pulses had the same amplitude.)

The second generation pulse block is now beyond our theoretical understanding (because average Hamiltonian theory on this large composite block introduces many complicated terms without a way to gauge which are the most important). This particular design was instead determined through many experiments. Undergraduates Ben Deen and Suyog Bhandari both had independent research projects where they changed different aspects of this second-generation pulse sequence to determine what parameters produced the best line-narrowing in PbTiO_3 , particularly for use in an electric dipole moment experiment. They tested the best number of composite blocks to include, the pattern of arranging these blocks, as well as the phases of the external and internal wrappers. They found that increasing the number of composite blocks between acquisition windows increased the linewidth and the pattern of the composite block (either ‘alternating’ or ‘blocked’ patterns) did not matter for the most part. A summary of some of their findings regarding external and internal wrapper phasing are given below.

External Wrapper Pulses

For both line-narrowing and imaging sequences, repeated pulse blocks are used. One can choose to always have the same phase external wrappers on all pulse blocks or to change the phase in different patterns. The simplest pattern would be to simply alternate between using 90_Y and 90_{-Y} as external wrapper phases in consecutive pulse blocks. A similar alternating pattern (where phase was changed after every *two* blocks instead) was found to give a more symmetric response when either a 90_X or 90_Y excitation pulse is being used. We expect this is the case because the train of 90_Y (or 90_{-Y}) pulses when all pulse blocks have the same external wrapper pulse phase introduces an effective H_{YY} term. This term dephases the spins when they are directed along the x axis (resulting from a 90_Y initial excitation pulse), but does not dephase

the spins when they are directed along the y axis (resulting from a 90_X excitation pulse). One can get rid of this effective H_{YY} term by simply changing the phase of the external wrappers.

Internal Wrapper Pulses

Internal wrappers were added to the second-generation line-narrowing pulse block in order to help the second averaging which suppresses the second term in the composite block Hamiltonian (see equation 3.14). In experiment, we did find that internal wrappers helped with line-narrowing in the pulse blocks containing multiple composite blocks. Similar to what was found for external wrappers, reversing the phase of all internal wrappers along with external wrappers after every two pulse blocks provided good line-narrowing as well as a symmetric response when either a 90_X or 90_Y excitation pulse was being used. One can imagine using more internal wrappers than the two shown in figure 3.6, but we found similar linewidths were attainable when using six internal wrappers compared with two, so we stayed with two for simplicity.

3.2.2 Low Resolution Imaging Sequence

One possible imaging sequence which could be useful in the future (despite its lower resolution) is to leave all Zeeman Hamiltonian terms acting by switching the phase of second external wrapper pulse to be opposite the first external wrapper pulse in figure 3.6. This change means there is no longer a hidden 180_Y pulse after the pulse burst, so the internal (and any applied external) Zeeman Hamiltonian terms are not refocused. This would result in worse resolution, but would keep any chemical shift information that might be important for certain types of experiments. This would also be helpful for imaging systems without fast-switching gradients, as the gradients would no longer need to be changing signs after every pulse block.

3.2.3 High Resolution Fast Switch Imaging Sequence

In order to do high resolution imaging, we want to make full use of the quadratic echo line-narrowing sequence to get rid of all internal spin Hamiltonians (both dipolar and Zeeman), but keep the effect due to external Zeeman Hamiltonians (due to applied magnetic field gradients needed for imaging or pulse frequency offsets). To keep the pulse sequence from cancelling the effects of the external Zeeman Hamiltonians, we need to change the sign of the applied gradients or frequency offsets in sync with the pulse sequence. Basically, the sign of the applied gradients or frequency offsets for the final Δ needs to be opposite the sign of the applied gradients or frequency offsets in the initial Δ in each pulse block. One simple way to do this would be to change the sign of the gradients or pulse frequency offsets at every free-evolution period, Δ . In fact this was done in our actual MRI experiments, because it was simpler to program with the pulse program length constraint of the Bruker MRI system (see the next chapter for the final version used for imaging). However, similar to the findings above, one gets better line-narrowing performance by changing the periodicity some what. Figure 3.7 shows the fast switch pulse sequence we perfected to get very good linewidth across a wide range of offset frequencies (used to mimic the effect of magnetic field gradients). This sequence makes use of changing external and internal wrapper phases after every two pulse blocks, as we found worked in previous experiments.

In order to mimic the effect of applying magnetic field gradients (which we do not have installed in the 12T system in the Barrett physics laboratory), we added a pulse frequency offset to the transmitter and receiver of our spectrometer. A magnetic field gradient causes spins in the sample to have slightly different precession frequencies depending on their spatial location in the gradient. If a sample is uniform along the gradient, we would expect a ‘box top’ shape. To mimic this effect, we ran our fast

High Resolution Fast Switch Pulse Sequence

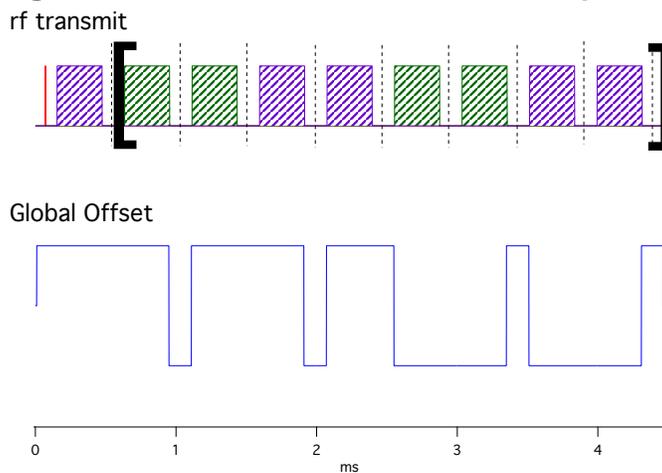


Figure 3.7: Pulse sequence for high resolution imaging with fast switching of the externally applied global Zeeman offset. This global offset is due to applied magnetic field gradients, for imaging, or adding an offset frequency to the transmitter and receiver of the spectrometer to mimic the effects of gradients, as was done in our box top data. Acquisition is taken at the black dotted lines and the entire sequence between the square brackets is looped until the desired amount of data points have been acquired.

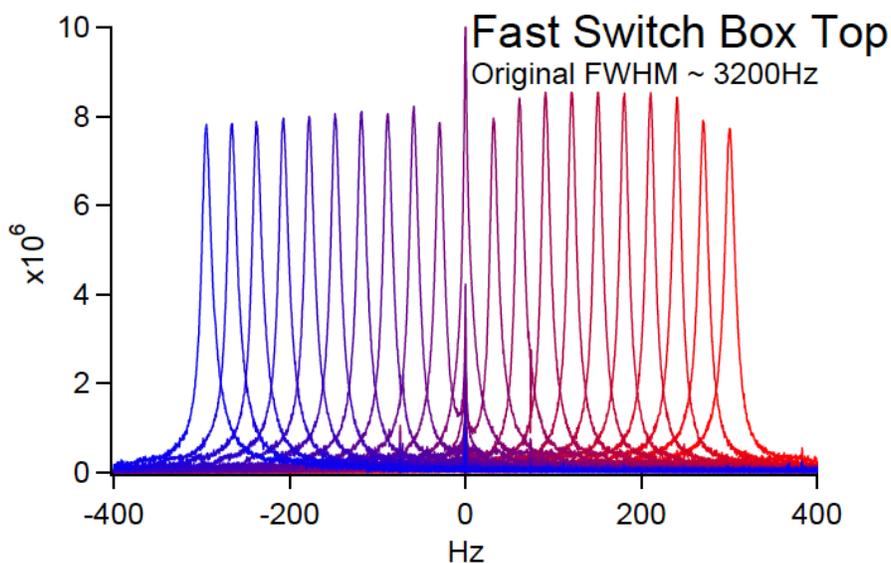


Figure 3.8: Box top data using the larger loop fast switch imaging sequence shown in figure 3.7 with $\tau = 5\mu s$ and $T_{180} = 10.1\mu s$. This data was taken looking at ^{31}P in a powdered bovine bone sample.

switch sequence using a set of uniformly-spaced offset frequencies to test the sequence's effectiveness across a wide range of frequencies to give a discretized 'box top' shape. If our sequence gives spectra with the same amplitude and linewidth across a wide range of frequencies, then we can be confident that no added distortions would be added to an MRI image from the pulse sequence.

Figure 3.8 shows the 'box top' data we took using the fast switch pulse sequence shown in figure 3.7 (looking at ^{31}P in a powdered bone sample). It is important to note that even though the applied offset frequencies were -1kHz to 1kHz in 100Hz steps, the observed frequencies are much smaller (by a factor of $\approx 1/3$) because the external Zeeman Hamiltonian only acts over 2Δ out of the total $\approx 6\Delta$ pulse block. The box top is fairly uniform across this range of frequencies, with the notable exception of 0Hz (no frequency offset). In this case the spectrum is slightly narrower and has a higher amplitude with the help of spin-locking (where spins align with effective fields, in this case due to the pulse sequence itself). Running with a frequency offset however keeps spins rotating in the toggling frame, limiting spin-locking from occurring. Despite its line-narrowing effect on the spectrum, spin-locking is undesirable because it keeps some spins from evolving as we desire (e.g., with a given offset frequency), so some signal from the each offset frequency peak leaks into a small peak at 0Hz. This effect gets worse as the offset frequency approaches 0Hz, so, in application for imaging, we often purposely use offset frequencies to ensure the entire spectrum (spread over many frequencies due to applied gradients) is far from the 0Hz region.

3.2.4 High Resolution Slow Switch Imaging Sequence

As its name suggest, the fast switch imaging sequence requires switching the gradients fairly frequently and rapidly, which is a challenge for many conventional MRI systems where the gradient ramp time is usually fairly long (hundreds of microseconds). In

order to fit in the gradient change, the timing of the pulse sequence might have to be lengthened and its effectiveness reduced. Anticipating this difficulty, we also developed the slow switch imaging sequence shown in figure 3.9, which has the gradient change after several pulse blocks. The total Zeeman Hamiltonian (both internal and external) is allowed to act over these pulse blocks by changing the phase of the final external wrapper, just like the low resolution pulse sequence. These pulse blocks are then followed by a $180_{\pm Y}$ pulse to change the sign of all the Zeeman Hamiltonian terms. The externally applied gradient or offset frequency changes sign at this time as well, so its effect is *not* refocused. After the same number of pulse blocks, one can then acquire the echo where both Zeeman and dipolar Hamiltonians should be refocused. Note, we are no longer acquiring after each pulse block, so now the sparse dwell time is increased by a factor equal to the number of pulse blocks between the 180° pulses.

Looking at P^{31} in a powdered bone sample, figure 3.10 shows the ‘box top’ data using the slow switch pulse sequence shown in figure 3.9. In comparison with the fast switch results, the frequency bandwidth is greatly reduced due to the increase in dwell time, so the box top fills a much larger portion of the total bandwidth. There is also more variation in the spectra for different offset frequencies (e.g., phasing is slightly off, 0Hz peaks from non-zero offset frequencies are larger, and the overall linewidth of each peak is wider). Clearly, fast switch is preferable and one would use this only if needed to fit in slow-changing gradients.

3.3 Important Considerations for MRI of Solids

In order to to implement the pulse sequences above for solid-state imaging, there are some other factors one must consider. The methods above involve very high-power

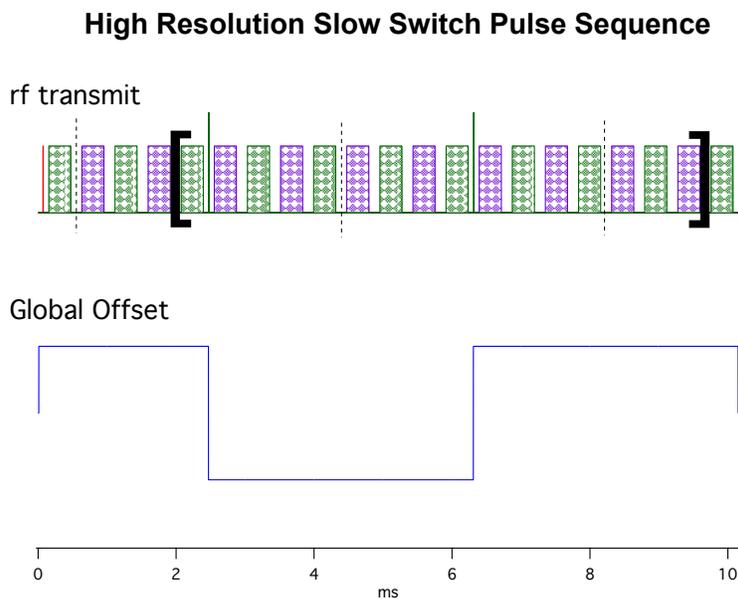


Figure 3.9: Pulse sequence for high resolution imaging with slow switching of the externally applied global Zeeman offset (Ω_{glob}). Here the pulse blocks are similar to the pulse blocks of the same color shown in 3.6 *except* the final external wrapper has the opposite phase as the initial external wrapper so the Zeeman Hamiltonian is left acting over the pulse block. To refocus just the unwanted internal Zeeman Hamiltonian, we add in 180° pulses (made larger so you can see their locations) in sync with changing Ω_{glob} .

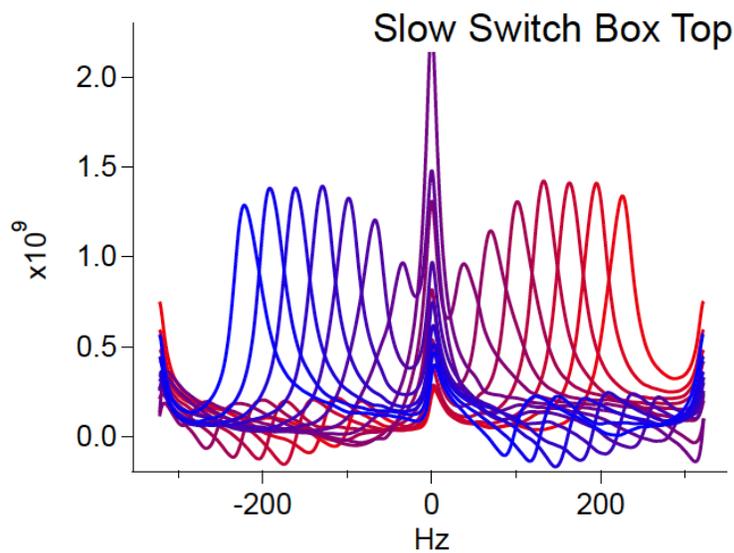


Figure 3.10: Box top data using the slow switch imaging sequence shown in figure 3.9 with $\tau = 3\mu\text{s}$ and $T_{180} = 8.7\mu\text{s}$. This data was taken looking at ^{31}P in a powdered bovine bone sample.

pulses typical for solid-state NMR but not used in MRI due to heating concerns for patients. We plan to focus first on *ex vivo* samples, but if the quadratic echo imaging pulse sequence were to be used on human patients, the pulse power would need to be reduced significantly. Below we explore the effectiveness of the quadratic echo imaging pulse sequence at lower pulse powers. There are also other considerations to think about when converting a conventional MRI system (used mainly to image ^1H in water) for use in solid-state MRI. Further below we discuss some of the modifications we made to a liquid-state MRI system in order to do our solid-state imaging.

Table 3.1: Table of the change in FWHM and magnitude of spectral peak with lowering pulse power (increasing T_{180}). The T_{180} values below the horizontal line were too long to excite the entire spectrum. For consistency, the values are taken from spectra where no global frequency offset was used to get rid of any effects due to spectra being in different locations inside the FOV. The magnitude of the spectral peak was also corrected for the different sparse dwell times of the experiments. *Hard excitation pulse and no internal wrappers were used.

T_{180} (μs)	FWHM (Hz)	Magnitude of Peak
10.1	7	2.98×10^3
18.6	12	2.05×10^3
46.0	20	9.49×10^2
109.9*	25	7.75×10^2
225.4*	75	2.45×10^2
267.1*	87	1.50×10^2
307.9*	104	1.12×10^2

3.3.1 Lowering Pulse Power

The data reported in this thesis work were measured using strong pulses ($T_{90} \approx 10\mu\text{s}$). However, if one wanted to use quadratic echo imaging pulse sequences on a human patient, the pulse power would need to be significantly lowered (and/or duty cycle greatly reduced) to reduce possible heating in the patient. Thus, we were interested in the effectiveness of the quadratic echo imaging pulse sequences with lower pulse powers.

Results from these experiments are given in table 3.1. Compared to the 3200Hz original FWHM linewidth of the data, the quadratic echo imaging sequence works surprisingly well for even very low pulse powers. Its effectiveness understandably gets dramatically worse in the regime where the pulses were too long to actually excite the entire width of the FID spectrum.

During these experiments, we discovered that to get the best results for weaker pulses, it was better to use an imaging pulse sequence without internal wrapper pulses as well as a ‘hard’ initial excitation pulse. Finding the correct T_{180} times was also challenging for weaker pulses because the nutation curves were not ideal. We instead estimated the ‘correct’ low power T_{180} times from the high-power pulse T_{180} time by measuring the ratio of the peak-to-peak voltages of the pulses (V_{pp}) and using the equation,

$$\frac{T_{180}^{\text{low power}}}{T_{180}^{\text{high power}}} = \frac{V_{pp}^{\text{high power}}}{V_{pp}^{\text{low power}}}. \quad (3.15)$$

3.3.2 Imaging Sequence Used for MRI Results

All the data shown thus far was taken on a 12T solid-state NMR magnet in the Barrett lab. Actual MRI needs to be done on a system with magnetic field gradients, and so we used the 4T Bruker Avance system at Yale Medical School’s Magnetic Resonance Research Center. This MRI system has an out-of-date spectrometer control system that is designed primarily for liquid-state imaging of ^1H . As a result, the nicely perfected imaging pulse sequences discussed above were not easy (or practically impossible) to implement due to system constraints. For example, due to memory constraints on the length of the pulse sequence program, I had to use pulse blocks with the same phase external and internal wrappers in order to keep the loop of the sequence as short as possible. This leads to non-ideal results (e.g. transverse field terms in the H_{YY} direction leading to spin-locking that results in an even worse artifact at $f = 0\text{Hz}$

compared to the possible results if one could implement the more optimized pulse sequences discussed above). Even using a simplified pulse sequence to shorten the pulse program as much as possible, I still ran into memory issues on the Bruker system when I tried running imaging experiments with too many averages and too many scans. At times, higher-resolution was possible (by taking more \vec{k} -space points, and thus more scans per experiment), but was inhibited because the system could not run the experiment due to memory issues. In some cases, I broke the imaging experiment into multiple shorter runs (e.g., a separate experiment for each octant) to get around this system constraint.

To further reduce noise on the system, we added an in-line Bessel low-pass filter at 15kHz (because our ^{31}P signal in solids is a lot smaller than the signal one gets from ^1H in water). The use of the filter was particularly useful for our soft tissue samples, where the signal was very hard to distinguish from the noise without the extra filter. We also used analog detection mode (rarely used in liquid-state imaging) in order to incorporate our stroboscopic detection of the data and had to incorporate a correction for the time delay introduced by our added in-line filter. In order to accurately determine the gradient ramp shape (which will be discussed in more detail in the next chapter), we also needed to know about the presence of any other filters built into the MRI system. These filter effects are very often insignificant for liquid-state imaging, but proved important for correctly scaling our solid-state images.

Chapter 4

High Resolution MRI of Solids

Our MRI of solids technique adopts the spatial encoding strategy pioneered by earlier line-narrowing approaches to the MRI of solids: repeated pulse blocks null out the internal spin Hamiltonian, while applied magnetic field gradients add in the external spin Hamiltonian [46, 47, 48, 1]. As was discussed in the previous chapter, the quadratic echo line-narrowing sequence has the unique advantage of working best in the regime where $H_Z \geq H_{ZZ}$. Thus, our pulse block is a better choice to null out the internal spin Hamiltonian of ^{31}P (and other nuclei where $H_Z \geq H_{ZZ}$ [49]), and stronger gradients are more readily applied since they can be left on during the pulse burst. Furthermore, our line-narrowing pulse block removes constant resonance offsets, so we do not need to worry about susceptibility broadening, chemical shifts, and shimming of the magnet, which are serious problems for many MRI approaches.

In this chapter, we discuss how we implemented the quadratic echo line-narrowing sequence to do high resolution ^{31}P MRI of solids using both Cartesian and radial sampling of \vec{k} space, as well as show our results imaging hard solids (e.g. bone mineral) and soft solids (e.g. mouse brain). Most of this work is featured in our 2012 article in *Proceedings of the National Academy of Sciences, USA* [50].

4.1 MRI of Solids Using Cartesian Sampling

A large majority of MRI is done using a uniformly-spaced grid of points in \vec{k} space, so-called Cartesian sampling, because of its straightforward processing with FFT algorithms. Our pulse sequence (which is naturally periodic and acquires points after every $\approx 6\Delta$) can also be easily adjusted to take points along the Cartesian grid by just applying gradients in one direction at a time. Unlike other MRI pulse sequences, which often use phase-encoding gradients to start sampling at the far corner in \vec{k} space, we always start at the origin of \vec{k} space and our sequence gives a very accurately-phased $\vec{k} = (0, 0, 0)$ point. More detailed information about how we applied our line-narrowing pulse sequence to sample \vec{k} space along the Cartesian grid is given below.

4.1.1 Cartesian Sampling of k Space

We use the 2nd generation quadratic echo line-narrowing pulse blocks discussed in the previous chapter to build a 3D MRI pulse sequence that maps out \vec{k} space along a Cartesian grid, shown in figure 4.1B. Here we use the simplest version of the 2nd generation pulse sequence (with very regular gradient and frequency offset switching, as well as simple phasing of pulses) because the Bruker MRI system we were using had limited memory for programming the pulse sequence, so we needed the smallest repeatable unit possible. For other systems, we anticipate this not being as big of an issue and the more complicated versions of the pulse sequence discussed in the previous chapter can then be tried.

In the imaging pulse sequence shown in 4.1B, a non-selective 90° pulse excites the full sample volume, and we acquire a single point in \vec{k} space after each pulse block. Since this pulse sequence effectively gets rid of time-independent Zeeman and dipolar interactions, these stroboscopically detected points oscillate solely due to the effec-

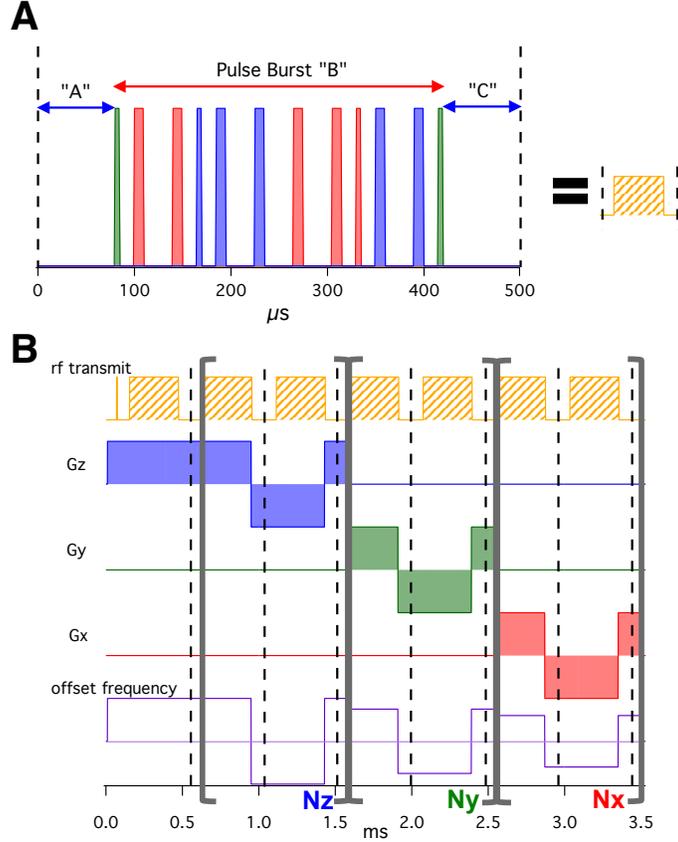


Figure 4.1: Schematic of our imaging sequence. (A) The main pulse block used in our imaging sequences with 90° (thin) and 180° (thick) pulses along the Y (green), X (red), and $-X$ (blue) directions. (B) Diagram depicting the 3D MRI sequence with our pulse block where the loop counters: N_z , N_y , and N_x determine the \vec{k} -space trajectory. The offset frequency can be set for each loop to place the image center off-center in the FOV. A single \vec{k} -space point is sampled at each dashed line.

tive applied field gradient and pulse frequency offset, f_{offset} , like a pseudo-FID. This sequence is thus a hybrid of echo (for internal fields) and FID (for external fields) imaging approaches [51], which enables high spatial resolution. In order to encode spatial information with the gradients, as well as position our image inside our FOV through use of frequency offsets, we synchronize the applied field gradient and frequency offsets with the pulse block. The gradient and frequency offset change signs for period “C” (relative to period “A”) right after the hidden 180° pulse at the end of the pulse block (as shown in figure 4.1A). The external Zeeman contributions to the spin Hamiltonian

due to the applied gradient and frequency offset effectively act over both time periods “A” and “C”.

Starting at $\vec{k} = (0, 0, 0)$, we sample \vec{k} space uniformly along Cartesian grid trajectories by systematically varying the loop parameters, N_x , N_y , and N_z (with $N_x + N_y + N_z = \text{constant}$). Each black dashed-line in figure 4.1B shows where we acquire a single \vec{k} -space point. Both G_z and f_{offset} are held constant through the first dashed-line, making this an excellent $\vec{k} = (0, 0, 0)$ point, which is crucial for proper phasing (and co-adding) of the signals at this point. Next, to map out a particular Cartesian trajectory in \vec{k} space, just G_z is modulated during the N_z loop, then G_y is modulated during the N_y loop and finally G_x is modulated during the N_x loop. We refer to this loop ordering as $\{Z, Y, X\}$. Note that the f_{offset} value can be different in each loop, to center the sample at a particular (x, y, z) spatial location in the FOV. This is particularly useful because we position our sample to take up only one single octant to avoid a known artifact at the center of the FOV. This artifact is a result of a complicated kind of spin-locking (see [52] for good explanation of spin-locking) due to spins staying aligned with the effective fields resulting from our pulse sequence, and similar effects are encountered in most line-narrowing pulse sequences.

The real parts of two measured pseudo-FIDs are plotted in figure 4.2A, (although only the first 17 points of each 32-point pseudo-FID are plotted here). They are the same through the N_z loop (equal steps along $+k_z$), but they diverge in the N_y loop, with steps along $+k_y$ ($-k_y$) for the yellow (blue) points due to flipped G_y modulation patterns. They have the same G_x modulation in the final N_x loop (steps along $+k_x$). The two trajectories through 3D \vec{k} -space are depicted in figure 4.2B and for simplicity we will refer to their loop patterns as $\{Z, Y, X\}$ (yellow) and $\{Z, -Y, X\}$ (blue).

Figure 4.2C shows an octahedron of 3D \vec{k} -space points filled in this manner. The upper half of the octahedron lying in octants 1-4 is directly filled using the loop pat-

terns: $\{Z, Y, X\}, \{Z, Y, -X\}, \{Z, -Y, X\}, \{Z, -Y, -X\}$, respectively. The same data is used to complete the lower half of the octahedron (in octants 5-8), using the Hermetian symmetry of \vec{k} -space [51]. A total of $N_{\text{pnts}} = 32$ points were acquired in each of the 1D pseudo-FIDs, which defines the octahedron's surface ($0 \leq N_z, N_y, N_x \leq 31$ for each trajectory, with the constraint $N_z + N_y + N_x = 31$). This loop pattern uses a total of $(N_{\text{pnts}}/2)(N_{\text{pnts}} + 1) = 528$ pseudo-FIDs to fill each octant. Fourier transformation of the \vec{k} -space data yields a 3D ^{31}P MR image (figure 4.2D-E).

4.1.2 Cartesian Sampling Analysis

The previous section provides a general overview of how we sample \vec{k} space along the Cartesian grid. This section gets into the specifics of how we collect and analyze the data to process the final image.

Acquisition of \vec{k} -space Points

Above we referred to a single point of \vec{k} space being acquired after each pulse block in figure 4.1B. This single complex point needs to be acquired right at the peak of the quadratic echo. In reality, we acquire five complex points (with dwell times of $2\mu\text{s}$) for every acquisition window centered about the time where we expect the peak of the echo. In order to accurately determine when the peak of echo occurs in time, we needed to take into account the group delay of the echo due to the 15kHz low-pass Bessel in-line filter we used to acquire our ^{31}P data. In offline processing, we use three applications of three-point binomial smoothing and then take the center point of each 5-point acquisition window to make the pseudo-FID. Each pseudo-FID dataset is phased using the first $\vec{k} = 0(t = 0)$ point, which has very consistent phase across all the pseudo-FIDs. The sign of the imaginary points in every odd-numbered window is then

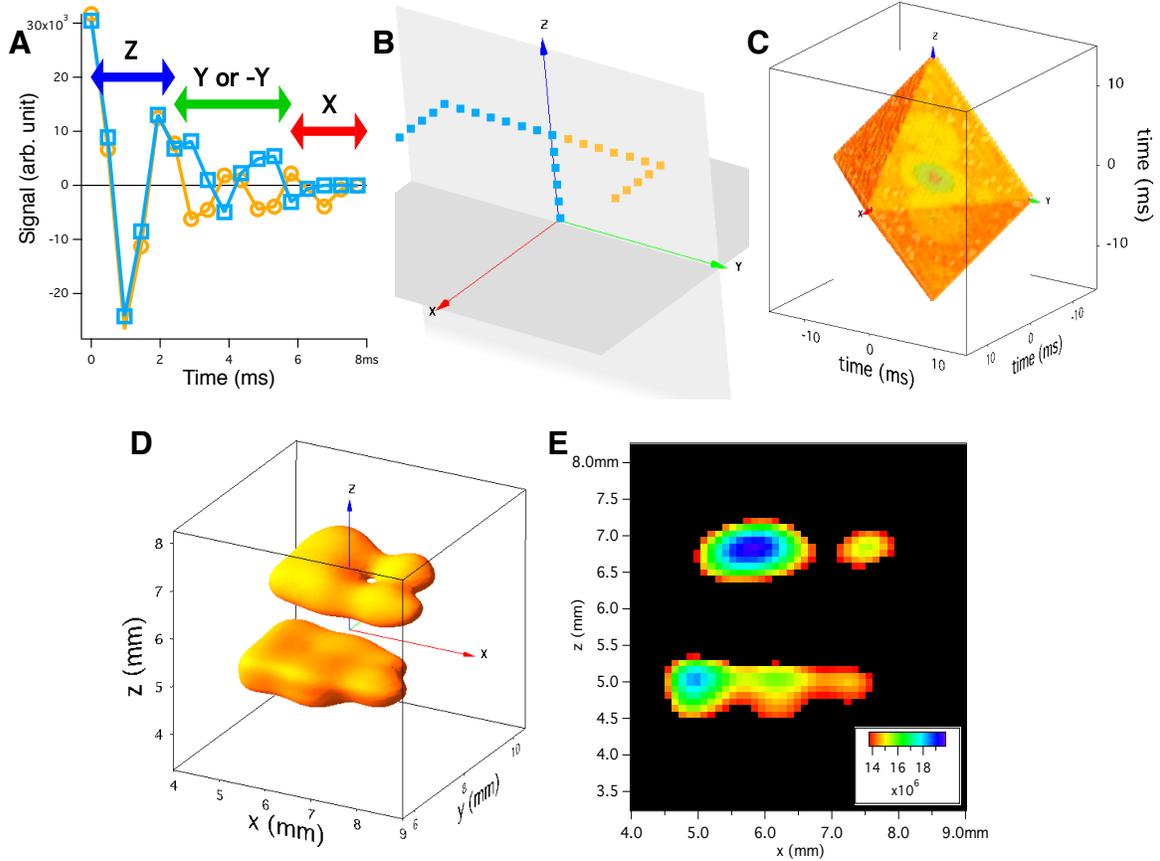


Figure 4.2: Figure showing the mapping of acquired data onto \vec{k} space and the resulting image. (A) Plot showing real data acquired (from the dry bovine bone sample) in the first 17 windows for two different trajectories in \vec{k} -space. Starting at $\vec{k} = (0, 0, 0)$, both the yellow-circle and blue-square points take 5 steps forward along the positive k_z -direction ($N_z = 5$). Next the yellow-circles (blue-squares) move 7 steps forward (backward) along the positive k_y -direction ($N_y = 7$). Finally, both the yellow-circles and blue-squares step forward along the positive k_x -direction ($N_x = 19$), for a total of 32 points. (B) A 3D-plot of these two trajectories in \vec{k} -space. (C) 3D octahedron of \vec{k} -space points measured with our imaging scheme (plotting the magnitude of the complex points in the time-domain). The imaging time was just under 47 hours. (D) Isosurface rendering of the 3D image of ^{31}P in two dry bovine bone blocks separated by a 1.1 mm gap obtained by Fourier transformation of C. The spatial resolution is $0.428 \times 0.428 \times 0.353 \text{ mm}^3$. Note the thru-hole in the top block which was drilled using a 0.343 mm diameter drill bit. The sample dimensions are $4.9 \times 2.6 \times 1.4 \text{ mm}^3$ (bottom bone) and $4.2 \times 2.4 \times 1.4 \text{ mm}^3$ (top bone). The isosurface value was chosen to show the presence of the thru-hole and is 65% of the maximum signal value. (E) A 2D slice of the 3D data (zero-filled by a factor of four) with thickness of 0.107 mm. The cut-off for the minimum of the color-scale is the isosurface value used in D to clearly show the thru-hole in the top bone and two partial holes in the bottom bone, made by the same 0.343 mm-diameter drill bit.

flipped, to account for the hidden 180° pulse in each pulse block. Each pseudo-FID then has a “sparse dwell time” of the total block duration ($6\Delta + 4T_{90} \approx 0.5\text{ms}$).

From these pseudo-FIDs (which vary due to different values of N_z , N_y , and N_x), we then fill an octahedron inside a $(64 \text{ pt})^3$ Cartesian grid in 3D \vec{k} -space, with the $\vec{k} = (0, 0, 0)$ point in the center of the grid. As a result of our sampling trajectories, many points (particularly along the z-axis and zy-plane) are sampled multiple times. For a particularly dramatic example, the trajectory with $N_z = 31, N_y = 0, N_x = 0$ will have all but the final point in common to the trajectory with $N_z = 30, N_y = 1, N_x = 0$. To correct for this oversampling of points, we first co-add the ‘ p ’ complex data points that should correspond to the same \vec{k} -value. Second, we divide the sum at each \vec{k} -value by the corresponding p (e.g., in figure 4.2B, $p = 2$ for the co-added data at $\vec{k} = (0, 0, 0)$, while $p = 1$ at the other end of the trajectories).

Once we have our 3D grid of \vec{k} -space data, we simply need to take the Fourier transform to produce the desired image. To provide smoothing in the image, we often first zero fill our \vec{k} -space data by a factor of two or four, which has the effect of adding points through interpolation in between points in the image. This zero padding is done by simply putting the original data cube at the center of a larger (by a factor of two or four) data cube. Every non-measured point is then assumed to be zero. This is a fair approximation since we typically choose parameters so that the pseudo-FID signal decays down to zero within the 32 acquired points.

Calculating Effective Gradient Factors

In order to scale the resulting image to give the correct spatial positions, we need to know the net effect of the applied gradients. The image FOV in $(\text{Hz})^3$ can be converted to the FOV in $(\text{meters})^3$ by dividing each frequency axis (e.g., f_z) by the corresponding gradient factor (e.g., $\alpha_z \gamma G_z / 2\pi$, where α_z is a scale factor less than 1 that compares

the time that the gradient “is effective” to the total sparse dwell time between points of the pseudo-FID). The spatial resolution in each direction (e.g. x is shown here, but results for y and z have the same form) is then given by:

$$\delta x = \frac{\pi \delta f}{N_{\text{points}} \alpha_x \gamma G_x} \quad (4.1)$$

where $N_{\text{points}} = 32$ for the examples above (and the total image size is $2N_{\text{points}} \times 2N_{\text{points}} \times 2N_{\text{points}}$) and $\delta f = 1/2t_{\text{acq}}$ is the frequency resolution determined by the total pseudo-FID acquisition time, $t_{\text{acq}} = N_{\text{points}} \times \text{‘sparse dwell time’}$.

As a first approximation, $\alpha_z = \frac{1}{3}$ if we can apply $+G_z$ during intervals “A” and “B”, followed by $-G_z$ during interval “C” (see blue dotted trace in figure 4.3), since the effective gradient over the sparse dwell time is $\approx \frac{2\Delta}{6\Delta} G_z = \frac{1}{3} G_z$. Of course, the current through the gradient coils cannot change instantaneously, so a better value for α_z is obtained by measuring the output of the gradient amplifiers on a digital oscilloscope (see green dashed trace in figure 4.3). This includes a slight time delay (t_{del}) for the gradients to respond after the gradient command is given, as well as a ramp time (t_{ramp}) for the gradients to change signs. An even better value for α_z takes into account the in-line low-pass isolation filter located en route to the gradient coil (buried deep within the wall of the magnet room), which adds an additional $\approx 11\mu\text{s}$ time delay to the schematic gradient waveform (see the black solid trace in figure 4.3). This delay proved elusive at first and took about six months to track down. We first noticed our images were consistently smaller in scale compared to $\mu\text{-CT}$ images of the same sample. After double-checking all of our scaling calculations, I found that these scale discrepancies could be explained by a single gradient delay of around $12 \pm 1\mu\text{s}$. When we asked the technician at the MRRC if there was anything that could cause this delay, he mentioned the in-line filter for the gradients and fortunately had an extra

filter we could test. We measured the filter to have an $11\mu\text{s}$ delay, which perfectly matched our hypothesis, and ultimately fixed the image-scaling problem.

Including all of these important time delays, the scale factor α_i (where $i = x, y, z$) is then given by the equation:

$$\alpha_i = \frac{2\Delta - 2t_{\text{del}} - t_{\text{ramp}}}{\text{Sparse Dwell}} \quad (4.2)$$

where t_{del} includes an $\approx 6\mu\text{s}$ time delay for the gradient system to respond to the programmed command as well as $\approx 11\mu\text{s}$ delay due to the presence of the in-line low pass isolation filter. The time it takes for the gradient to go from $-G_i$ to $+G_i$ (t_{ramp}) is dependent on the gradient amplitude, as well as the direction of the gradient (because each gradient direction is generated by a different gradient coil that can respond slightly differently to the applied current). In order to calculate α_i , we thus need to measure the gradient ramp times for the gradient amplitudes we used for imaging in all three gradient directions (see appendix A for a table of the measured gradient ramp time values).

Of course, the actual waveform is more complicated, with curvature during the ramp, overshoot at the top, and then decaying oscillations (see figure 4.4). The quantitative comparison of micro-CT data to MR images that assume our schematic model (see the next section) suggests that these high frequency oscillations may be safely ignored for our conditions. In our model, we also ignore the small displacements from the ideal \vec{k} -space grid that occur when one gradient component turns off, and another turns on (e.g., while switching from G_z to G_y). Since our current method of Cartesian sampling only requires at most 2 switches of gradients (i.e. from G_z to G_y and from G_y to G_x), these are fairly small perturbations that we can safely ignore. We have seen evidence that these displacements cannot be ignored when we switch the gradients

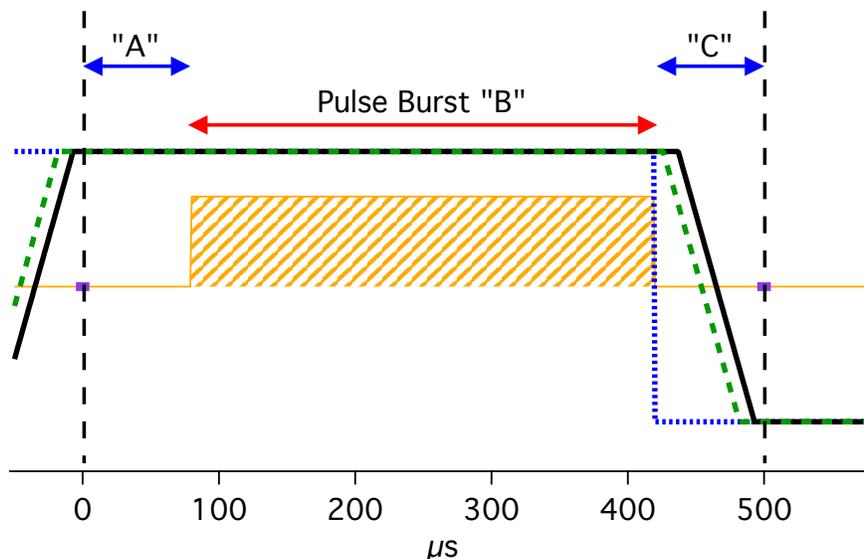


Figure 4.3: A schematic of the gradient transient. For best results, the gradient should be constant during “B”, and during each $10\mu\text{s}$ data acquisition window (the small purple rectangles at the beginning of “A” and the end of “C”). Ideally the gradient transient would be like the dotted blue line and switch instantaneously immediately after the pulse burst “B”. A better approximation of the gradient transient is given by the dashed green line where there is some delay before the gradient transient starts (approximately $6\mu\text{s}$) after it is called and there is also some time for the gradient ramp to take place ($< 60\mu\text{s}$ for the “ramp-off” mode and small gradient amplitudes we typically used). An even better model for the gradient transient takes into account the gradient low-pass isolation filter through which the current flows before arriving at the gradient coils. This adds an extra $11\mu\text{s}$ delay and is given by the black solid line. This is the final model used to calculate the effective gradient factors for scaling the images presented here.

multiple times (which will be discussed in the next chapter).

Imaging Time

Another important factor of MRI of solids is the imaging time. This is much longer than conventional MRI of ^1H in water due to the much longer T_1 in solids. Each pseudo-FID we acquire is a distinct experiment requiring time, T_{rep} , per repetition. For ^{31}P MRI this repetition time can be from 4s (in soft tissue) to 60s (in dry bone mineral), compared to T_{rep} in milliseconds for liquid samples. The number of pseudo-FID’s we

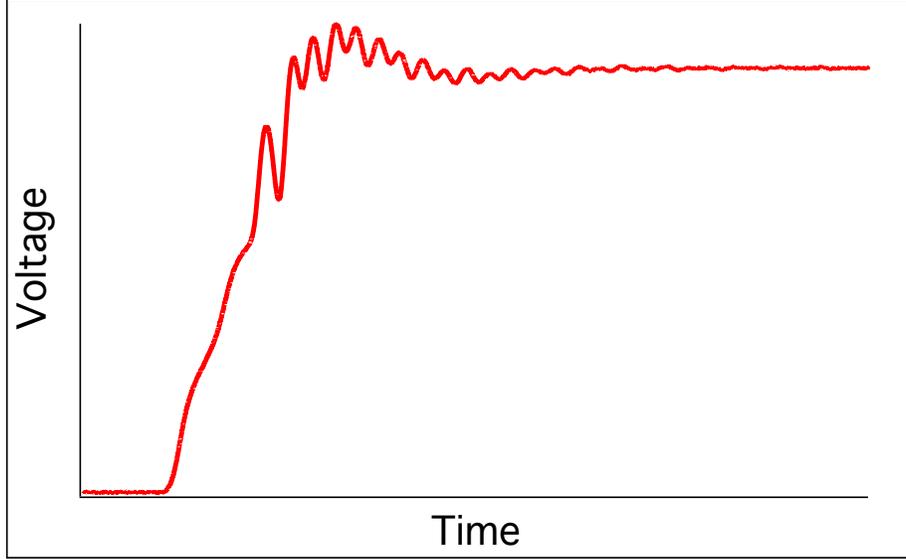


Figure 4.4: A measurement of the gradient waveform looking at the current coming from the gradient amplifier. This data was taken from a screenshot of a digital oscilloscope and shows the complicated waveform we are approximating in figure 4.3.

need to acquire in order to satisfy the Nyquist-Shannon sampling theorem (which we will call ‘dense sampling’) is given by the equation:

$$N_{\text{traj}} = N_{\text{oct}} \left(\frac{N_{\text{pnts}}}{2} \right) (N_{\text{pnts}} + 1) \quad (4.3)$$

where N_{oct} is the number of octants being sampled and N_{pnts} is the number of points in each pseudo-FID. If we are sampling all eight octants of \vec{k} space with $N_{\text{pnts}} = 32$, we then need $N_{\text{traj}} = 4224$ distinct experiments and have to wait approximately 60s between experiments. This imaging time would then be over 70 hours. In order to speed up the imaging time by a factor of two, we often only sample the first four octants (so $N_{\text{oct}} = 4$) and use Hermitian symmetry to fill-in the remaining four octants. In the next chapter we will explore speeding up imaging further by disobeying the Nyquist-Shannon sampling criterion by intentionally undersampling \vec{k} space. In order to get a high quality image from undersampled data, an alternative to simple Fourier

transformation must be used for the image reconstruction process.

4.1.3 Cartesian Sampling Results

The techniques described above were implemented on a animal MRI system located at Yale University Medical School’s Magnetic Resonance Research Center (MRRC). Here we share important ^{31}P imaging results in both hard bone mineral and soft tissue. We decided to image ^{31}P since it is a 100% naturally abundant spin-1/2 nucleus with a spin Hamiltonian in solids very similar to that of ^{29}Si in crystalline silicon and ^{13}C in C_{60} (with $H_Z \geq H_{ZZ}$). ^{31}P also is an important constituent to many biomedical samples (e.g. bone mineral and cell membranes), which provided very interesting samples to image. Unlike conventional ^1H MRI - which often uses T_1 or T_2 weighting to provide contrast - all the ^{31}P MRI shown here reveal the ^{31}P density in each sample. The addition of various contrast mechanisms will be explored in future work. It is also important to note that our samples are all *ex vivo* to date because, for reasons of animal and human patient safety, *in vivo* MRI cannot use the strong rf pulses and rapid gradients changes shown in figure 4.1B. Consequently, we are initially focusing on applications that can use *ex vivo* methods, such as our ^{31}P MRI of bone. A recent review [53] summarized 15 state-of-the-art methods to assess bone quality and listed 12 as primarily or exclusively *ex vivo* techniques. Our data indicate that determining the ^{31}P microstructure will provide important new information that complements what can be learned from these existing techniques.

Imaging Bone Mineral

Bone is a composite material [54], containing approximately 45% bone mineral by volume [55]. Bone mineral is similar to calcium hydroxyapatite (i.e., $\text{Ca}_{10}(\text{OH})_2(\text{PO}_4)_6$), but it is less crystalline, and it has a unique stoichiometry [23]. The spatial distri-

bution, composition, and quantity of bone mineral are primarily responsible for the compressive strength and stiffness of bone [54, 55, 23]. While a few ^{31}P MRI studies have successfully targeted *in vivo* [23, 24, 25] and *ex vivo* [23, 22, 56, 57] bone, the broad MR spectra have limited the achievable spatial resolution to no better than 0.5 mm [57], and more typically in the range of 2 mm. There is great interest in probing the internal composition of bone on the sub-0.1 mm length scale [53, 58], both to study normal features (see figure 4.5) and to look for signs of disease.

Compact Bone & Spongy (Cancellous Bone)

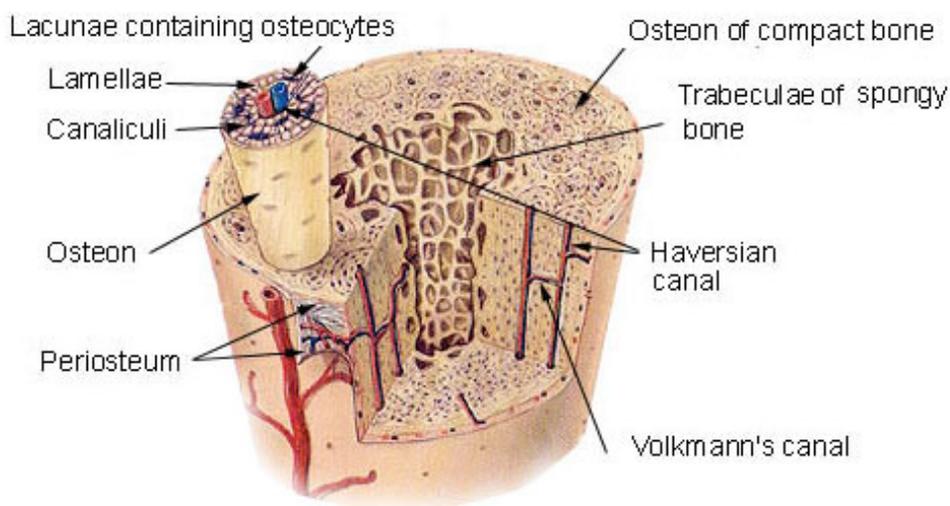


Figure 4.5: Structure of the bone, courtesy of <http://training.seer.cancer.gov/anatomy/skeletal/tissue.html>

Despite the obvious importance of the mineral component to the biomechanical properties of skeletal tissue, few useful non-destructive technologies are available to evaluate changes in its chemical structure. Micro-computed tomography (micro-CT) measures the X-ray linear attenuation coefficient, which provides high spatial resolution 3D imaging of the electron density (dominated by calcium in the bone mineral) but does not reveal the chemical information provided by MRI. In bone mineral, up to 14% of the phosphate groups are replaced by carbonate groups, and several recent Fourier

transform infrared (FTIR) and Raman spectroscopy studies [59, 60] have reported a spatial dependence of the $\text{CO}_3:\text{PO}_3$ ratio that correlates with bone tissue age, as well as with nano-indentation measurements of indentation modulus and hardness. In the imaging results below, we have pushed the 3D ^{31}P MRI spatial resolution down to the sub-0.4 mm length-scale. With the improvements discussed in the final chapter of this thesis, we fully believe our approach has the potential for sub-0.1 mm spatial resolution for *ex vivo* bone samples.

Dry Bovine Bone

The images in figure 4.2D-E, obtained using our pulse sequence, show the features of a sample composed of two dry blocks of bovine cortical bone separated by a 1.1mm masking tape spacer. These blocks were cut from a segment of white, cleaned bovine femur (a natural dog chew purchased at the local Petco). The upper block has one thru-hole, and the lower block has two partial holes made by a #80 drill bit (0.343mm-diameter). Teflon tape was wrapped around the outside of the two-block assembly to hold it together and to center it in the solenoidal coil. The threshold value chosen for the isosurface plot and 2D slice in figure 4.2 gives a clear view of the hole but does not depict the outer-surface of the blocks. However, after checking various 2D slices and different isosurface values, the image does match the spatial dimensions of the sample. These findings are consistent with our estimate of the spatial resolution of $0.428 \times 0.428 \times 0.353\text{mm}^3$, the best yet reported for ^{31}P MRI in bone.

This demonstrates that the achievable spatial resolution of the image is no longer limited by the natural linewidth. For comparison, if we were to implement ‘solid-state ^{31}P MRI’ [57] on our system using the same 30mT/m gradients, the natural ^{31}P linewidth would limit the spatial resolution to $(5.9\text{mm})^3$, which is larger than the entire volume shown in Fig.4.2D. Since quadratic echo line-narrowing overcomes this

limitation, the achievable spatial resolution of our technique is limited only by the specifications of our MRI system.

Wet Pork Rib

Our approach works just as well on wet, marrow-filled bones. In fact, the ^1H in the water produces fluctuating magnetic fields which make the measured T_1 of ^{31}P in the wet samples shorter, decreasing the imaging time by a factor of two compared to dry bone. Our ^{31}P MRI has been applied to store-bought sections of a pork rib (mostly cortical bone), shown in figure 4.6A. The pork rib was imaged in a sealed plastic cryotube filled with phosphate buffered saline (PBS) solution to keep the sample hydrated. The thick outer shell of cortical bone that surrounds the marrow space dominates both the 3D surface plot (figure 4.6B) and the 2D slice (figure 4.6C). The spongy trabecular bone in the marrow space is not visible because the density of ^{31}P in the trabecular bone is smaller than the isosurface value of the phosphorus density plotted. Given the relatively low spatial resolution, $(1.19\text{mm})^3$, using a smaller value of phosphorus density for the isosurface just fills in the marrow space entirely and the microarchitecture of the trabecular bone cannot be resolved.

As a validation of our approach and analysis, the same pork rib sample was imaged using micro-CT, which is considered the gold standard for obtaining a 3D volumetric map of the bone mineral. The micro-CT data was measured by a μCT 35 (Scanco Medical), using 55kVp for the peak energy, 500ms integration time, and $37\mu\text{m}$ (isometric) resolution/voxel size. We then Gaussian broadened this image to match the ^{31}P MRI resolution, which blurs out the trabecular bone network in the marrow space. Figure 4.7 shows a coregistration of our ^{31}P MRI isosurface plot (figure 4.6B) with an isosurface plot of the broadened micro-CT image. These look fairly similar (suggesting

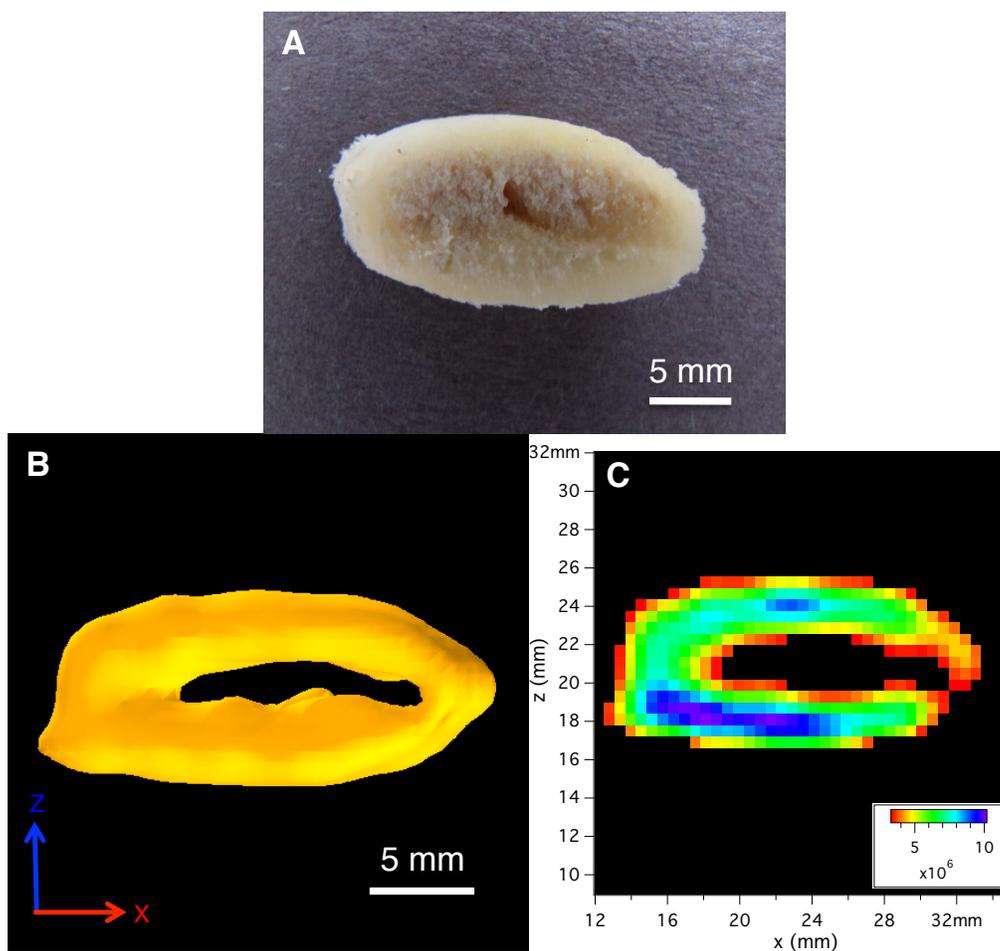


Figure 4.6: Images of the *ex vivo* pork rib sample. (A) Photo of the marrow-filled pork bone sample. (B) Isosurface rendering of the 3D image of ^{31}P in pork rib in PBS solution. The isosurface value is 33% of the maximum signal value and shows the thick cortical bone ring. The spatial resolution is $(1.19\text{mm})^3$ and the imaging time was 35.2 hours. (C) A 2D slice of the 3D data shown in B (zero-filled by a factor of two) with thickness of 0.595mm.

our calculation of the effective gradient factors, with the $11\mu\text{s}$ delay from the in-line filter, scales the image correctly), but they are not identical, which is to be expected at this early stage. Further improvements in the spatial resolution of our technique should start to reveal local differences between these two maps.

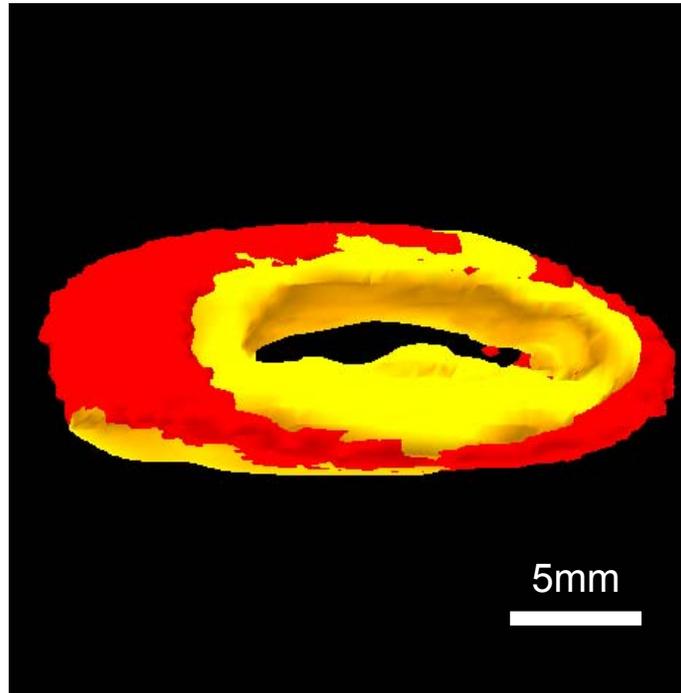


Figure 4.7: Isosurface rendering of the Gaussian-broadened 3D micro-CT data (red) registered with our ^{31}P MRI data (yellow) as shown in figure 4.6B. The isosurface value chosen for the micro-CT data is 10% of the maximum signal value, while for the ^{31}P MRI data the isosurface value is 33% of the maximum signal value, and shows the thick cortical bone ring of the wet pork rib.

Rabbit Femoral Head

We wanted to see if we could resolve the trabecular bone if we had higher resolution, and thus chose a smaller sample (to make better use of our FOV) as well as a sample that is mostly trabecular bone. The rabbit femoral head (the “ball” at the top of the femur), shown in Figure 4.8A, met both these requirements. It has a very thin,

nearly transparent layer of cortical bone covering the pink bone marrow region which is primarily trabecular bone. This sample was also imaged in a sealed plastic cryotube filled with PBS solution. The interconnected plates of trabecular bone on the interior of the rabbit femoral head (seen in the micro-CT 2D slice shown in figure 4.8B) are clearly visible in both the 3D surface plot (figure 4.8C) and a 2D slice (figure 4.8D), given the higher spatial resolution of this data set, approximately $(0.46\text{mm})^3$. A series of 2D slices (figure 4.8E-J) show a virtual sectioning of the 3D rabbit femoral head data set, which demonstrates the potential of this technique to provide non-destructive quantitative maps of phosphorus on the interior of complex 3D samples. There are good reasons the trabecular bone as seen in the MRI images do not replicate exactly the microarchitecture seen in the micro-CT images. The micro-CT has greater resolution (by a factor of 10), is imaging the electron density of the sample (dominated by the calcium), and has a different orientation than the ^{31}P MRI images.

Imaging Soft Tissues

As another potential biomedical target, soft tissues have phosphorus concentrated in the membranes, metabolites, RNA and DNA of cells (figure 4.9). This leads to a complicated, multi-peak ^{31}P MR spectrum (including a broad membrane peak and narrow metabolite peaks), which would ordinarily be a poor choice for high-resolution MRI [64]. However, our quadratic echo pulse block narrows the entire ^{31}P spectrum into a single peak, enabling high-resolution imaging. Using our sequence, we have carried out the first, to the best of our knowledge, 3D ^{31}P MRI on a variety of *ex vivo* soft tissue samples, including fixed neural stem cell-endothelial cell hydrogel co-cultures, mouse liver, mouse heart, and a variety of mouse brains. Figure 4.10 shows the *ex vivo*

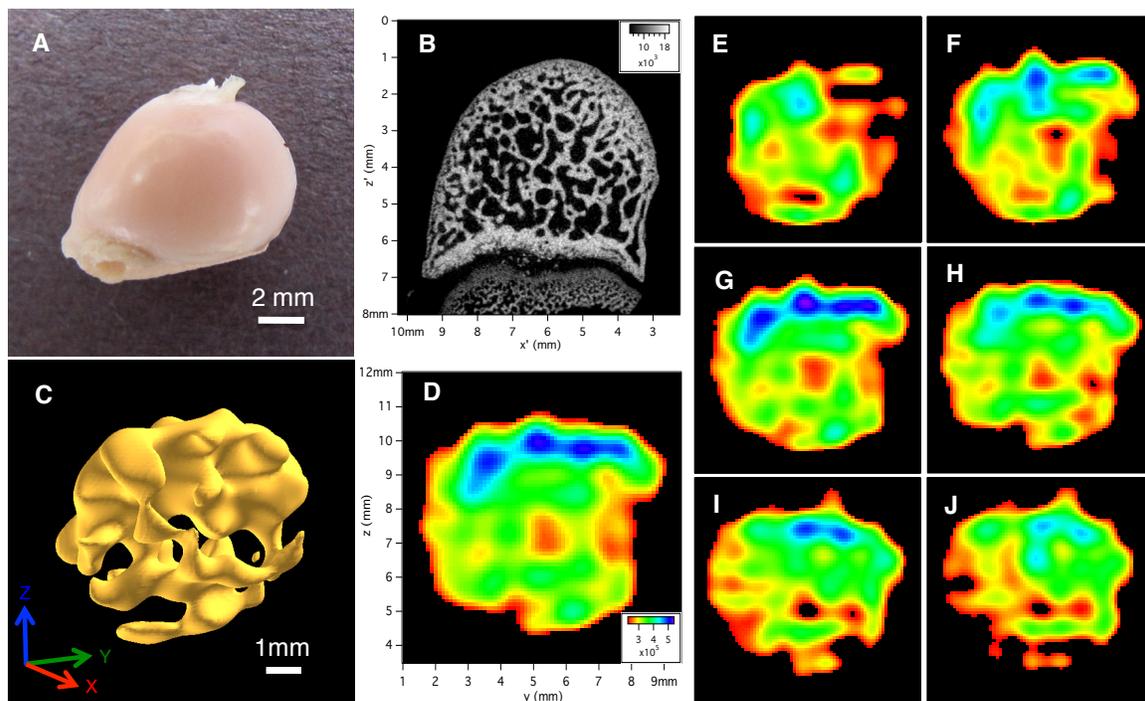


Figure 4.8: Images of the *ex vivo* rabbit femoral head sample. (A) Photo of the rabbit femoral head sample. (B) 2D slice of micro-CT data. The 2D resolution is $(0.0185\text{mm})^2$ and the slice thickness is 0.0185mm . (C) Isosurface rendering of the 3D image of ^{31}P in rabbit femoral head in PBS solution. The isosurface value is 60% of the maximum signal value and shows trabecular bone. The spatial resolution is $0.458 \times 0.458 \times 0.422\text{mm}^3$ and the imaging time was 70.4 hours. (D) A 2D slice of the 3D data shown in figure 4.6C (zero-filled by a factor of four) with thickness of 0.115mm . Note that these axes are different from those used in B and the orientation of the bone is different than in both A and B. Here the ‘flat end’ of the bone is on the right edge. (E) - (J) Multiple 2D slices (each 0.115mm thick) cutting along the x-axis going in the positive x-direction (with a 0.458mm step size) through the 3D data set shown in C, using the same color scale and FOV as in D. The 2D slice shown in D is between slices shown in G and H.

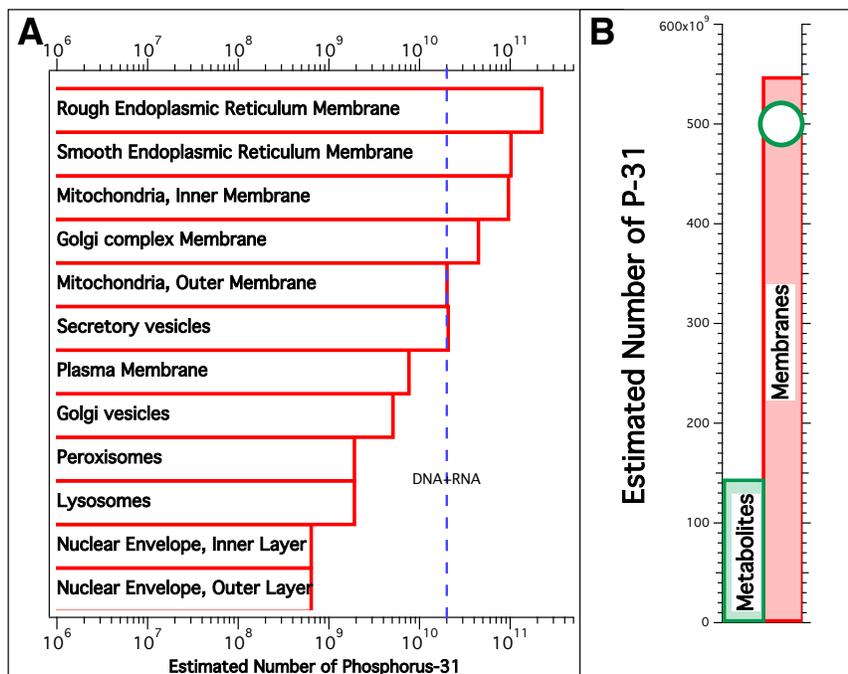


Figure 4.9: Estimated number of ^{31}P nuclei in a ‘typical’ $20\mu\text{m}$ human cell. (A) The estimated number of ^{31}P in membranes, by location [61]. The blue dashed line is the estimated number of ^{31}P in DNA and RNA ($\approx 2 \times 10^{10}$) [62]. (B) The sum of all parts in A is the total number of ^{31}P in membranes (red bar) for typical cell ($\approx 5.5 \times 10^{11}$), which is more than a factor of three times the ^{31}P content measured in brain metabolites (green bar [63]). An independent estimate of the brain membrane ^{31}P content (green circle [63]) is consistent with ours. As comparisons, the ^1H content of mobile water in the cell [61] is $\approx 9 \times 10^{13}$, which matches the ^{31}P content of the same amount of bone mineral.

3D ^{31}P MR image of a mouse brain in PBS. This is a functionally different kind of MR image, since conventional ^1H MRI probes the intracellular and extracellular free water [51]. A rough calculation of the phosphorus reservoirs in a typical cell suggests that the membrane signal represents approximately 75% of the total ^{31}P signal for our MRI of solids approach, depending on the density of mitochondria (figure 4.9). This membrane signal is rarely studied due to its broad MR linewidth but may provide new insights into cellular and tissue function that compliments the information revealed by ^1H MRI. It is important to note that the total ^{31}P content of a cell is approximately

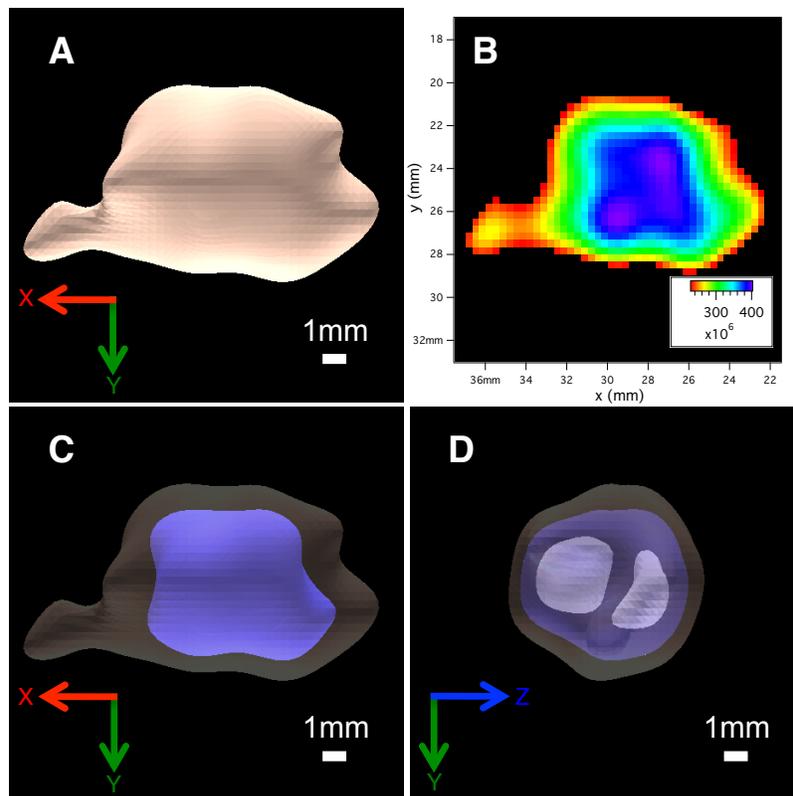


Figure 4.10: 3D ^{31}P MR image of an *ex vivo* mouse brain. (A) Isosurface rendering of the 3D image of ^{31}P in mouse brain in PBS solution. The isosurface value is 50% of the maximum signal value and shows the brain stem on the left. The spatial resolution is $(1.33 \text{ mm})^3$ and the imaging time was 88.7 hours. (B) A 2D slice of the 3D data shown in A (zero-filled by a factor of four) with thickness of 0.332 mm. The cut-off for the minimum of the color-scale is the isosurface value used in A. (C) Two isosurface renderings of the 3D image of ^{31}P in mouse brain in PBS solution. The highlighted inner isosurface value is 70% of the maximum signal value and the outer isosurface is the same as in A. (D) Three isosurface renderings of the 3D image of ^{31}P in mouse brain in PBS solution, now rotated by 90° relative to those shown in A and C. The highlighted most inner isosurface value is 90% of the maximum signal value, the middle isosurface is the same as in C, and the outer isosurface is the same as in A.

1% of the ^1H in free water, so practically speaking, this means that less dense solid samples (i.e. soft tissues) will require larger voxel dimensions or more signal averaging. Fortunately, the T_1 of ^{31}P soft tissues is dramatically shorter (by over a factor of 10) compared to the dry bone samples, so extra averaging can be done without greatly enhancing imaging time.

4.2 MRI of Solids Using Radial Sampling

In this section we explore using radial sampling of \vec{k} space instead of the conventional Cartesian sampling. Radial sampling has many potential benefits compared to our Cartesian sampling method discussed above.

One potential benefit of the radial sampling approach is that we do not change the gradient directions for a given experiment (which provides points along a radial ‘spoke’). Our Cartesian mapping changes the gradients at most two times (eg. switching from z gradient to y gradient and then y gradient to x gradient, as in figure 4.1). Every time the gradient is switched, there is an extra displacement in \vec{k} space due to ramping down the previous gradient and ramping up the next gradient (see figure 4.3). This small displacement is not taken into account in our assumption of sampling evenly spaced grid points in \vec{k} space. Fortunately, it appears this assumption appears to work fairly well in our current Cartesian sampling images, but it could introduce a small artifact that radial sampling will not have. For example, we have tried a random walk around the Cartesian grid of \vec{k} space which had many more of these small displacements due to changing gradient directions. This resulted in noticeable destructive interference as the distinct 1D pseudo-FIDs were co-added to fill the \vec{k} -space grid (because each set of data were sampling slightly different \vec{k} -space points) and no image could be seen even from very high SNR data.

Another potential benefit of radial sampling compared to our Cartesian sampling method, is that we are not oversampling particular points in \vec{k} space (other than the origin, where we should have the most signal anyways). In our Cartesian sampling method, we oversample both the z axis and the yz plane - leaving well-defined regions with much higher SNR. This could cause artifacts in the image (which is very sensitive to any regular structures in the \vec{k} space data).

One final potential benefit of radial sampling is the ability to produce high quality images even when not taking all the required \vec{k} -space points according to Nyquist-Shannon sampling theorem (i.e. sparse sampling). We discuss this further below in the section about dense versus sparse sampling. This could prove useful as part of our strategy to speed up imaging, which we will explore further in the next chapter.

4.2.1 Radial Sampling of k Space

Much like Lauterbur's original MRI approach [32], radial sampling of \vec{k} space is done by applying gradients that are linear combinations of the x , y , and z directions (see figure 4.11) such that

$$G_x^2 + G_y^2 + G_z^2 = G^2, \quad (4.4)$$

where G is the desired gradient amplitude. We still sample starting from $\vec{k} = (0, 0, 0)$ but now we are acquiring equally spaced points along some radial spoke in \vec{k} space whose direction is determined by the particular linear combination of gradients applied. Once the desired number of points are acquired along a given radial spoke, we wait some time longer than T_1 for the spin magnetic moments to equilibrate and then take the next data set with a different radial gradient direction (determined by a new set of G_x , G_y , and G_z).

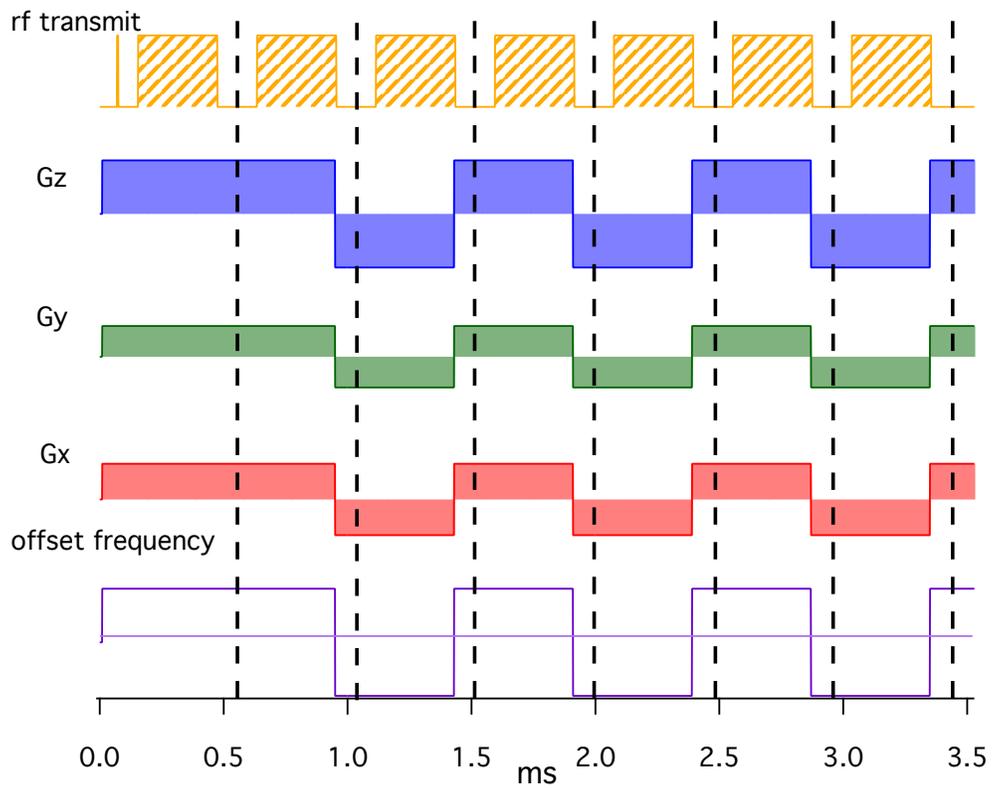


Figure 4.11: Example pulse sequence diagram for a single scan using radial sampling, in this case using the linear combination $G_z = 0.75G$, $G_y = 0.43G$, and $G_x = 0.5G$, where G is the desired gradient amplitude.

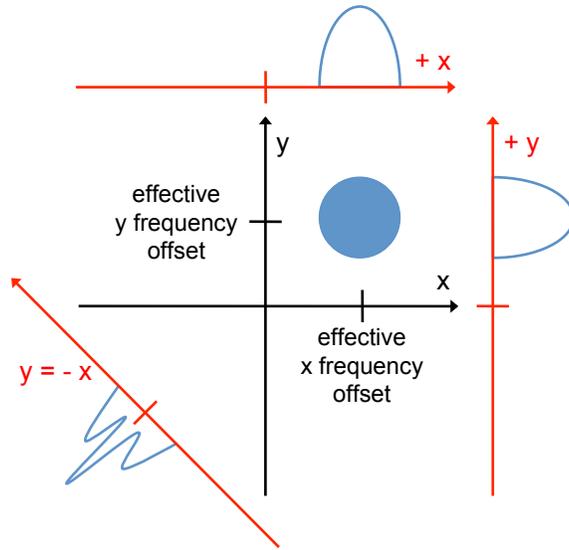


Figure 4.12: A 2D example of what three radial 1D frequency spectra would look like for an image offset by x and y effective frequency offsets using our line-narrowing method on a circular phantom. (These ‘effective’ frequency offsets are different from the applied frequency offsets due to our pulse sequence leaving only $\approx 1/3$ of the frequency effective over each pulse block). The ‘artifact’ region is at the origin of the frequency spectrum and marked by red lines in each of the 1D spectra. The avoidance of the artifact region by offsetting the image works well when applying gradients in only the x or y directions (as in Cartesian sampling), but fails when applying gradients in directions like $y = -x$. An approximation of what the 1D spectrum would look like when overlapping the artifact region is given in this case. No matter how we offset the image, there will always be some gradient directions where the image will be overlapping the artifact region.

In addition to following radial spokes in \vec{k} space instead of ‘bent’ Cartesian trajectories, another main difference of radial sampling compared to our Cartesian sampling method is our use of the offset frequency. In the Cartesian case, having the offset frequency (f_{offset}) is useful for positioning the center of our image in artifact-free regions of our FOV because this position is mainly determined by the magnitude of the frequency offsets in the x , y , and z directions if the sample is originally placed in the center of the magnet. For radial sampling, we can imagine doing a similar positioning of the image in the FOV to avoid artifact regions by now applying the offset frequency whose magnitude is given by

$$\text{offset frequency} = \left(\frac{G_x}{G}\right) f_{\text{offset},x} + \left(\frac{G_y}{G}\right) f_{\text{offset},y} + \left(\frac{G_z}{G}\right) f_{\text{offset},z}. \quad (4.5)$$

However, if we do this offset, there will still always be some 1D scans of the sample where the image will be crossing high-artifact regions, since now we are not always only going in just x , y , or z directions, but taking diagonal slices as well (see figure 4.12). A solution to this problem is to just have *all* the radial 1D spectra shifted by the same common offset frequency (as shown in the pulse sequence diagram in figure 4.11). This way no matter what gradient direction is used, *every* spectrum is offset from the artifact region by the same amount. However, we have to correct for this offset frequency in our reconstruction before combining all the 1D spectra because we need the image location to be consistent throughout all the 1D spokes. Without correction, each spectrum is offset by the effective offset frequency in the direction determined by the applied gradients (see figure 4.13), which means the spatial location of the image appears to be different for each 1D spectrum. To do this correction, before reconstruction we multiply each radial spoke’s time-domain dataset by a phase factor which will reverse the effect of the known effective offset frequency, causing all

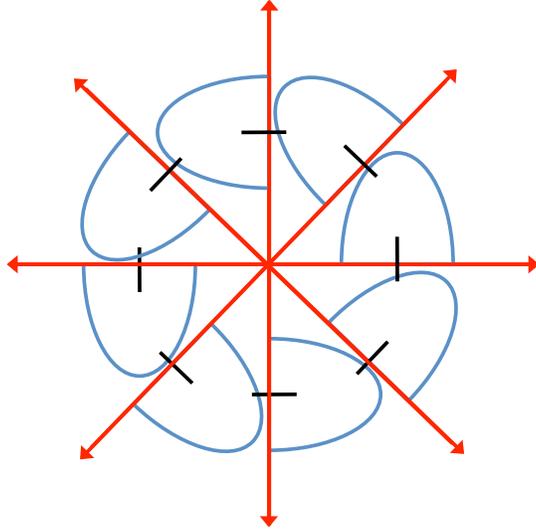


Figure 4.13: A 2D example of what eight radial 1D frequency spectra would look for a circular phantom if using a common frequency offset, as done during our radial sampling method. Each 1D spectrum is offset by the same effective frequency (given by the black line). In this way the artifact at 0 Hertz will be avoided no matter what gradient direction is used. However, we must reverse the effect of this offset before image reconstruction so that all 1D spectra share the same center.

the artifact-free radial 1D spectra to now be centered about their common origin.

Isotropic Sampling of 3D \vec{k} Space

To ensure that we are sampling \vec{k} space isotropically, we choose the gradient directions such that the end-points are uniformly distributed on the surface of a unit-sphere centered at the origin of \vec{k} space (see figure 4.14). However, since we are always sampling points with equal spacing along the radial direction, the sampling of 3D \vec{k} space goes as $1/\vec{k}^2$ in the radial direction (see figure 4.15 for 2D depiction). This actually works to our advantage because most of our signal is at low \vec{k} values and the noise is at higher \vec{k} values. However, we must take this oversampling into account in our regridding process when we reconstruct the data (which will be discussed later). The following discussion shows how we determine the gradient directions by approximately

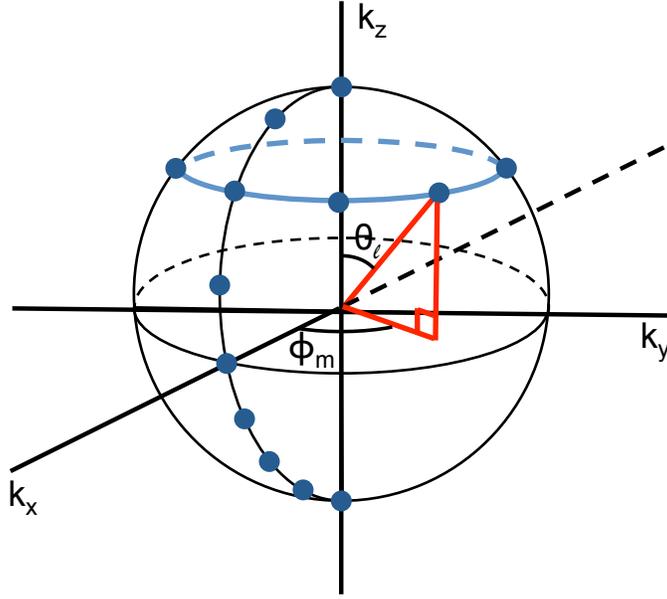


Figure 4.14: A depiction of how we are approximately uniformly distribute points on the surface of a sphere in \vec{k} space. These points will be used to determine the gradient directions applied for radial sampling of \vec{k} space to ensure we are sampling isotropically. Each gradient direction is then given by the vector pointing from $\vec{k} = (0, 0, 0)$ to the given point on the sphere.

uniformly distributing points on a sphere.

First we are going to mark off L equal segments along the prime meridian of the sphere using $L - 1$ points plus the two endpoints at each of the poles. This means that $\delta\theta = \pi/L$ where θ is the angle measured from the z axis (parallel with the ‘north pole’ in this description and in figure 4.14).

At each of these points along the meridian, (eg. at a particular $\theta_\ell = \ell\pi/L$), we can imagine swinging the radial vector around in the ϕ direction to create a circle with radius $r \sin \theta_\ell$ (see figure 4.14). In order to have an approximately uniform distribution of points along the surface of the sphere, we need to determine the number of points along this circle, M_ℓ , needed for each given θ_ℓ , since we will want fewer as the circle gets smaller (eg. we approach the poles) and more as the circle gets larger (eg. we

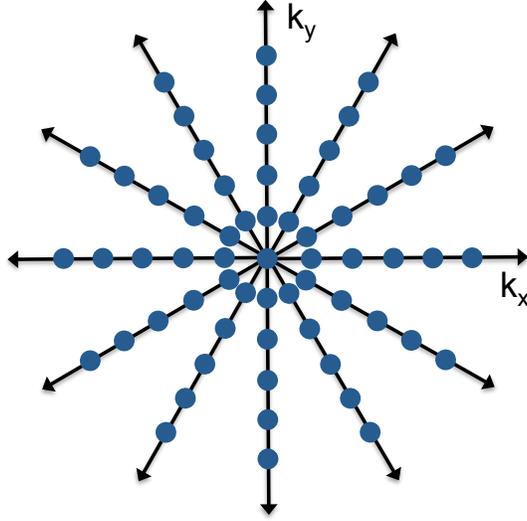


Figure 4.15: A depiction in 2D of how we would be sampling \vec{k} space. Even though our gradient directions will be isotropic, since we acquire points with equal spacing in the radial direction, the sampling of \vec{k} space goes as $1/k$ in the radial direction for the 2D radial sampling depicted here, and $1/k^2$ for 3D radial sampling. This oversampling will be taken into account during our reconstruction process.

approach the equator). These points will each have a different ϕ spherical coordinate and if we have them equally spaced then $\delta\phi_\ell = 2\pi/M_\ell$. Since the arc-length of the circle at each θ_ℓ scales with $\sin\theta_\ell$, M_ℓ must also scale with $\sin\theta_\ell$. We also know that in order for the density of points to be the same in both the ϕ and θ directions, $M_\ell = 2L$ at the equator (when $\theta_\ell = \pi/2$), so the general formula for M_ℓ is given by:

$$\begin{aligned}
 M_\ell &= 2L \sin\theta_\ell \\
 &= 2L \sin\left(\frac{\ell\pi}{L}\right).
 \end{aligned}
 \tag{4.6}$$

The $M_\ell = 2L$ value at the equator ($\theta_\ell = \pi/2$) comes from the following calculation which ensures that the number of points around the circumference of the sphere in the ϕ direction matches that in the θ direction. Above we noted there are $L - 1$ points

plus the two endpoints along the full π extent of the prime meridian of the sphere, for a total of $L + 1$ points. Doubling this to find the total number of points around the full 2π of the circumference of the sphere in the θ direction (and subtracting 2 because we have double-counted the points at the poles) gives the $2L$ value for the number of points around the entire circumference of the sphere.

If we desire to know the total number of points on the sphere we need to sum M_ℓ over all $L - 1$ values of θ_ℓ and add in the 2 points at the poles to get:

$$\begin{aligned} \text{Total number of Points} &= 2 + \sum_{\ell=1}^{L-1} M_\ell \\ &= 2 + \sum_{\ell=1}^{L-1} 2L \sin\left(\frac{\ell\pi}{L}\right). \end{aligned} \tag{4.7}$$

There is also a degree of freedom not discussed above where you can arbitrarily rotate the M_ℓ points at a given θ_ℓ by some δ_ℓ angle to try to spread the points more uniformly over the sphere (and get rid of the pattern of always having points along the meridian at $\phi = 0$). However, when you get rid of patterns in one place, you are probably causing patterns to emerge elsewhere, so for simplicity's sake, we just had the first M_ℓ point always start at $\phi = 0$. We could contemplate having δ_ℓ be some randomly generated angle we keep track of for every data set, but this might be a level of complexity that is unnecessary unless we see artifacts stemming from oversampling the $\phi = 0$ meridian.

Dense versus Sparse Sampling

As of yet, there have been no limits set on what values L can have (though presumably it should be a positive integer greater than 1). Eventually, this radially sampled data will be regridded onto a Cartesian grid with isotropic spacing given by the average of

the (in general, anisotropic) spacing along each axis of \vec{k} space,

$$\delta k = \frac{\delta k_x + \delta k_y + \delta k_z}{3} \quad (4.8)$$

where

$$\delta k_i = \frac{\alpha_i \gamma G}{2\pi} \delta t \quad (4.9)$$

where α_i is the effective gradient factor in the i -direction, γ is the gyromagnetic ratio, G is the maximum gradient amplitude applied, and δt is the sparse dwell time. Note, even if G and δt are held constant, the α_i parameter usually introduces anisotropy.

For ‘dense’ sampling, we want to have at least one measured \vec{k} -space point from our radial sampling map onto every \vec{k} -space point along the chosen Cartesian grid. In order for this to happen, the maximum arclength between two adjacent rays along the sphere is the grid-spacing of the Cartesian grid, δk . The angle between two adjacent rays along the θ -direction is given by π/L and we can approximate the length of a ray to be $\approx (N_{\text{points}} - 1)\delta k$. This gives an approximate limit on how small L can be:

$$\delta k \gtrsim \frac{\pi}{L}(N_{\text{points}} - 1)\delta k, \quad (4.10)$$

$$\rightarrow L \geq \pi(N_{\text{points}} - 1), \quad (4.11)$$

where N_{points} is the number of sparse points taken along each ray (including the $t = 0$ point) and the image size is then $2N_{\text{points}} \times 2N_{\text{points}} \times 2N_{\text{points}}$. Ultimately, spatial resolution, δx , is given by:

$$\delta x = \frac{1}{2N_{\text{points}}\delta k}. \quad (4.12)$$

For higher spatial resolution (smaller δx), increasing N_{points} requires a larger L for dense sampling, which will increase imaging time. Increasing δk by increasing the

gradient amplitude, G , is limited by fitting the entire sample in the given FOV (or, in our case, a single octant of the FOV because we try to avoid the artifact at the center of the FOV) and also how fast we can change the gradients between our pulse blocks. Increasing the sparse dwell time has been done (often to help fit in larger gradients), but this involves changing timing in our pulse sequence (usually by increasing τ) which becomes less effective the larger the space between pulses, so finding a good balance is necessary.

With the dense sampling criterion given by equation 4.11, taking a dense image using all eight octants of \vec{k} space would take about 1.86 times longer than using our current Cartesian sampling procedure. However, there are good reasons to suggest that radial sampling would be a prime candidate for sparse sampling of \vec{k} space. Sparse sampling will be the case when $L < \pi(N_{\text{points}} - 1)$, so not every \vec{k} -space point on the Cartesian grid will have a corresponding measured \vec{k} -space point in the radially sampled space. The gaps will be in the large \vec{k} regions, where we should have less signal and the lack of dense sampling should not be so apparent. These gaps will also be isotropically spread about, so we will not have to worry about large coherent artifacts due to particular regions being sampled more than others. For more information about practically implementing these radial sampling procedures, see appendix B.

4.2.2 Radial Sampling Analysis

When doing radial sampling of \vec{k} space, a slight modification to the analysis procedure is needed since now we are no longer uniformly sampling onto a rectilinear grid (required if we desire to use FFT reconstruction of the image). The most common way to modify the analysis is to ‘grid’ the radially acquired data to a uniformly sampled rectilinear grid by convolving the measured data by some smooth, finite function [51]. Since the expected image is compact (i.e., it is zero outside some finite region), the sampling

theorem states that the \vec{k} -space value at any location can be calculated exactly by convolving the measured values with a sinc function, given that the \vec{k} -space values are sampled at or greater than the Nyquist frequency. However, due to the infinite extent of the sinc function, computing the \vec{k} -space data at every new \vec{k} -space point requires multiplication of the sinc function by all measured \vec{k} -space values. Computationally, this is not very practical. In practice, the sinc function is then usually replaced by a finite-width convolution function [51]. Since we are now convolving with a finite convolution function, there will be some aliasing inside the FOV in the resulting image. This aliasing is greatly reduced by interpolation of the \vec{k} space by some interpolation factor, β , (usually a factor of two to keep the power of two size requirements for FFT purposes) so that the image FOV is increased by the interpolation factor and the excess amount (with most of the aliasing features) can then be discarded. This section will discuss gridding in general and then the practical implementations we use for the convolution function, density compensation, and interpolation in particular.

Gridding Basics

A convolution-based interpolation scheme is popular for regridding since it is relatively fast to execute while reducing artifacts. To discuss how this gridding is done using convolution, let's begin with the 1D definition of convolution,

$$(f * c)(x) = \int_{-\infty}^{\infty} f(x')c(x - x')dx', \quad (4.13)$$

where some function $f(x)$ is being convolved with a convolution function $c(x)$. Note that this integral need only be evaluated over the region where $c(x)$ is non-zero. This convolution function is usually chosen to be a smooth window-function of finite full-width, w . For some discretely sampled nonuniformly spaced \vec{k} -space data $s(k_j)$, con-

volution can be approximated by replacing the integral with a sum such that the convolved, uniformly spaced data set $s_c(k)$ is given by:

$$s_c(n\delta k) = \sum_j s(k_j)c(n\delta k - k_j)\delta k_j^{(m)}, \quad (4.14)$$

where n is an integer and δk is the grid spacing for the new uniformly sampled data. The factor $\delta k_j^{(m)}$ is the density compensation factor for the measured data that corresponds to the discrete representation of the differential length element dk for the nonuniformly sampled data. Since we are using a finite-width convolution function, we only need to sum over measured points k_j that are located within $w/2$ of each regridded point $n\delta k$.

Once we have our uniformly sampled regridded \vec{k} -space points, we can do an FFT to get the spatial dataset, $S_c(n\delta x)$. However, this is still not our final image since convolving with $c(k)$ in the \vec{k} -space domain has the effect of multiplying our desired final image with the Fourier transform of the convolution function, $C(x)$. We then get our final image by dividing $S_c(n\delta x)$ by $C(n\delta x)$.

Kaiser-Bessel Convoluting Function

Much work has been done on finding an optimal convolution function [65], with the consensus being the Kaiser-Bessel function which, in 1D, can be defined as:

$$\text{KB}(k) = \frac{I_0(B\sqrt{1 - (2k/w)^2})}{I_0(B)} \text{RECT} \left(\frac{2k}{w} \right), \quad (4.15)$$

where $I_0(x)$ is the zeroth order modified Bessel function of the first kind, w is the width of the function (usually as a integer multiple of δk , B is a dimensionless free scaling

parameter, and $\text{RECT}(u)$ is a rectangle function given by:

$$\text{RECT}(u) = \begin{cases} 1 & |u| \leq 1 \\ 0 & |u| > 1 \end{cases} \quad (4.16)$$

For our 3D radial sampling case, we used a spherical Kaiser-Bessel function defined as

$$\text{KB}(k_x, k_y, k_z) = \frac{I_0\left(B\sqrt{1 - 4(k_x^2 + k_y^2 + k_z^2)/w^2}\right)}{I_0(B)} \text{RECT}\left(\frac{2\sqrt{k_x^2 + k_y^2 + k_z^2}}{w}\right). \quad (4.17)$$

B mostly controls the shape of the Kaiser-Bessel function and its decay towards zero. One can find optimal values for B for given widths w and interpolation factors β [65, 66, 67] which reduce image artifacts. Since Beatty [67] gives a nice analytic expression for B for any interpolation factor, β , and also gives very similar results to those of Jackson [65] and O'Sullivan [66], we used their expression,

$$B = \pi \sqrt{\frac{w^2}{(\Delta k)^2 \beta^2} \left(\beta - \frac{1}{2}\right)^2 - 0.8}. \quad (4.18)$$

The continuous Fourier transform of the Kaiser-Bessel convolution function (needed to get the final image) has a simple analytical form, but in practice we used the discrete FFT of the analytic expression of the Kaiser-Bessel function, since the discrete FFT is what is really being used during our reconstruction.

Density Compensation

Here we will discuss how to find the appropriate density compensation factor for the measured data, $\delta k_j^{(m)}$, to make up for non-uniform sampling in the radial direction. The density compensation is, in effect, the volume surrounding each sampled \vec{k} -space point in 3D. In spherical coordinates, this differential volume element is $k^2 \sin \theta d\theta d\phi dk$.

For the discrete case, this can be approximated by

$$\delta k_j^{(m)} = k_j^2 \sin \theta_j \delta \theta_j \delta \phi_j \delta k_j. \quad (4.19)$$

From previous discussions, we have both $\delta \theta_j = \delta \theta = \pi/L$ and δk_j being constant for all sampled points. Since we do not care about overall scale-factors, we can thus ignore their input into the density compensation factor. However, $\delta \phi_j = 2\pi/M_j = 2\pi/(2L \sin \theta_j)$ is dependent on the sampled point. We then get for the density compensation factor:

$$\Delta k_j^{(m)} = k_j^2 \sin \theta_j \frac{2\pi}{2L \sin \theta_j} = k_j^2 \frac{\pi}{L}. \quad (4.20)$$

Since constants do not matter, the most commonly cited density compensation factor for 3D radial sampling is k_j^2 .

However, this density compensation factor has a big problem in practice. One can easily see that by multiplying the radially sampled data by this factor of k_j^2 , we will get zero for our $\vec{k} = (0, 0, 0)$ point, which is the point with our highest signal. This comes from taking a continuous distribution and using it for discretely sampled data. For continuous data, the density of points does tend to ∞ as you approach $\vec{k} = 0$, but that is clearly not the case for discretely sampled data (see, for example, figure 4.15). The big question is then: how should we then correctly determine this density compensation factor?

A majority of articles never mention or gloss over this ' $\vec{k} = 0$ problem' and just state the common k_j^2 density compensation factor, but there are a few people who have directly confronted it [68, 69, 70, 71, 72, 73, 74]. Many just treat the $\vec{k} = 0$ point separately and make the density compensation factor at that point some constant adjusted to better approximate the convolution integral [69]. Others find a more accurate density correction array by calculating the analytical expressions which should approximately

give the continuous theoretical value of k_j^2 when using the discrete Fourier transform (done in 2D for CT reconstruction by Ramachandran and Lakshminarayanan [68] and applied to 2D MRI by Joseph [71]). We have used this method to do the same in 3D (see appendix C for calculations), and this has the effect of approaching k_j^2 for higher \vec{k} -values, but then converges to a non-zero value for lower \vec{k} -values. In practice, we found using this density compensation factor gave a spherical shell of data with higher-values than expected at the edge of the sampled \vec{k} -space. This was most likely due to our data not being densely sampled (which this density compensation factor assumes) and the large \vec{k} -space points were then being overly compensated. For our sparsely sampled data, it would make sense that this density compensation factor would flatten-out at the point where the spokes are no longer within an Δk distance from each other. We thus needed a more sophisticated way to take into account our sparse radial sampling of \vec{k} space.

There have been other sophisticated methods to calculate the density compensation factor for any arbitrary sampling of \vec{k} space, including using Voronoi diagrams for 2D imaging to estimate the 2D area about each sampled point [73] and numerical iterative methods [72, 74]. We looked more closely at the more recent numerical conjugate gradient method used by Bydder, *et al.* [74]. This approach looked promising because it had been used for sparse radial datasets and had the expected ‘flattening’ of the density compensation factor that would get rid of the spherical shell we were seeing using the previous method. This method had previously been applied only to 2D images, so starting there, we applied their algorithm and soon realized its impracticality for 3D datasets, since it involves iteratively manipulating matrices whose sizes are dependent on the number of measured points and the number of regridded points. Memory problems were reached just doing a fairly low-resolution densely sampled simulated dataset.

Stepping back, we decided on a far simpler way of determining the density compensation factor, in lieu of how we fixed a similar ‘oversampling’ issue we had with our Cartesian approach. The density compensation factor is supposed to make up for the fact that more measured points (weighted by the convolution function) contribute to the regridded point for small \vec{k} values than for larger \vec{k} values. So while doing regridding, for each regridded point, we calculated the ‘density’ by summing the number of points that contributed to the regridded point, weighted by the convolution function being used (which gives higher weight to nearby points). After regridding is completed, we then divide each of these regridded points by the calculated density. This gives the desired ‘flattened’ shape of the density compensation factor of large \vec{k} values, as well as a non-zero value at $\vec{k} = 0$. One can imagine using this same method for any sampling of \vec{k} space for virtually no added computed time, which appears highly preferable to the arduous numerical iterative schemes previously discussed. For a view of the difference of our calculated density compensation factor ($= 1/\text{calculated density}$) compared with the one calculated using Ramachandran and Lakshminarayanan’s technique [68], see figure 4.16.

Interpolation and Other Corrections

A common feature of regridding reconstruction is to use interpolation by a certain factor (usually a factor of two to be able to use the FFT algorithm) [51]. This interpolation gets rid of the aliasing artifacts at the edges of the FOV in the image without interpolation (see figure 4.17) and is done by cleverly making use of the properties of the Fourier transform. Before regridding, each 1D radial spoke of data can be Fourier transformed to give a 1D projection of the image along the direction of the applied gradient for that given spoke. One can then increase the total number of points of

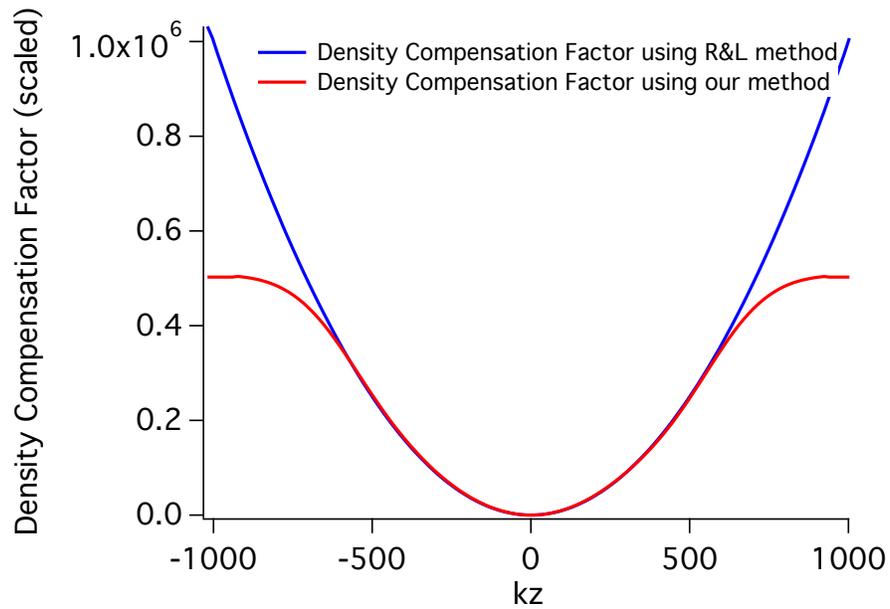


Figure 4.16: A view of the density compensation factor along the k_z axis. Our method (in red) is compared to the density compensation factor in 3D calculated using the method devised by Ramachandran and Lakshminarayanan [68]. Note that our technique flattens out for large \vec{k} values, which makes sense for our sparsely sampled data.

the spectrum by adding zeros to both sides. This is the same process as ‘zero filling’ as discussed in NMR basics, which is commonly used in the time domain in order to interpolate and smooth a frequency spectrum. Here we are zero-filling in the spatial domain, thus interpolating the original \vec{k} -space data. Adding these zeros do not add any information, but when you take the inverse Fourier transform to the original \vec{k} space domain, thanks to the wonders of Fourier transforms, you will now have double the points, where the added points are now interpolations between the measured values.

For our data, we have known artifacts that occur at the very center of our image space and the edges of the FOV, even in the 1D projections. To reduce these artifacts in the final regridded data (usually leading to spherical-shell-like features in the final image), we correct these values while doing the interpolation on the 1D spokes. While we are padding the 1D projection of the image for each spoke, we get rid of any artifacts at the origin and the edges of our original FOV by replacing these points by the average of its neighbors.

Also before regridding, we can do our correction for the applied offset frequency. We do this by multiplying the \vec{k} -space data by a factor of $\exp(-i2\pi f_{\text{eff,offset}}t)$, where $f_{\text{eff,offset}}$ is the effective offset frequency. This frequency is found by looking at a few of the 1D projections in different directions, and finding the common frequency to shift these projections so that their spectra share a common center. This factor is $\approx f_{\text{offset}}/3$, since the frequency offset applied, f_{offset} , is effective for only $\approx 1/3$ of each sparse dwell time due to our pulse sequence.

Some care needs to be taken when doing interpolation to make sure you do not add erroneous features to your data by padding in the image domain. When you find the interpolated \vec{k} -space data by taking the inverse Fourier transform, it might add some

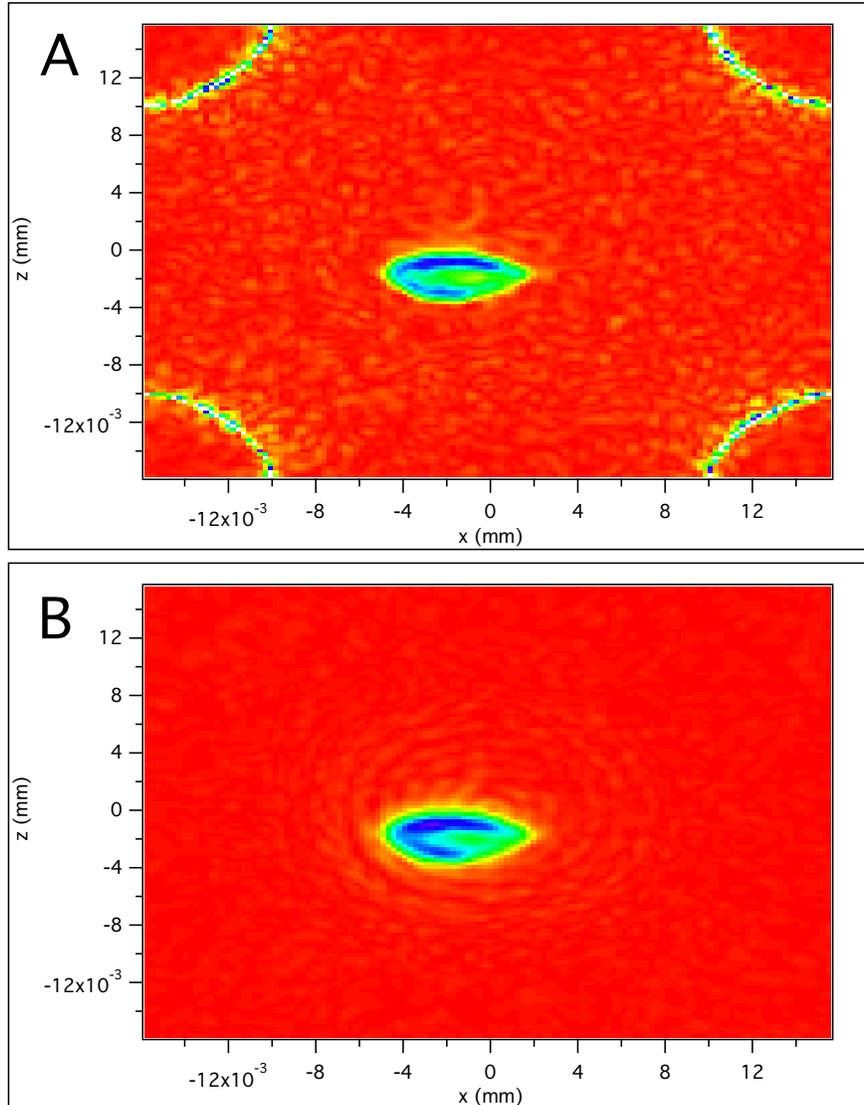


Figure 4.17: A 2D slice of 3D MR image of small pork rib sample showing the image (A) without interpolation and (B) with interpolation. With interpolation we get rid of aliasing artifacts at the edges, but this interpolation also includes some line-broadening of the image.

signal in the high- \vec{k} regions, where it is known that the signal should be near zero. In order to get rid of these effects, padding is often done along with a very slight line-broadening of the interpolated data (i.e. multiplying the \vec{k} -space data by a decaying exponential). The decay constant is chosen so that it gives a broadening well below the known spatial resolution of the image, so it will predominately affect the signal at large \vec{k} values which should naturally be close to zero.

4.2.3 Radial Sampling Results

The ^{31}P MR images using radial sampling were taken on the 4T animal MRI system at the Yale MRRC. As a proof of principle, we imaged a small pork rib sample (prepared in the same way as the pork rib sample above, and seen in figure 4.17) using both dense Cartesian and sparse radial sampling. The value chosen for L for radial sampling was chosen to more closely match the imaging time of the radial sampling (with a single average, but same repetition time) to that of the Cartesian sampling (which was only sampling four octants and then averaged twice). The radially sampled image still took a bit longer (with 5218 separate experiments compared to 4224 for Cartesian sampling). All other parameters of the pulse sequence were kept the same. The analysis of the radial sampled image included interpolation by a factor of two, regridding using a window width of $6\delta k$ and zero-padding by another factor of two. Considering the different sampling and different analysis procedures, it is difficult to do a direct comparison of the two plots except to see that the resulting image looks similar (see figure 4.18). The radial image is broadened (as part of interpolation), so the features are more smoothed out and the image has slightly lower spatial resolution than the Cartesian image. However, considering that the radial image is only 43% of the number of points required for dense radial sampling, this image looks very artifact

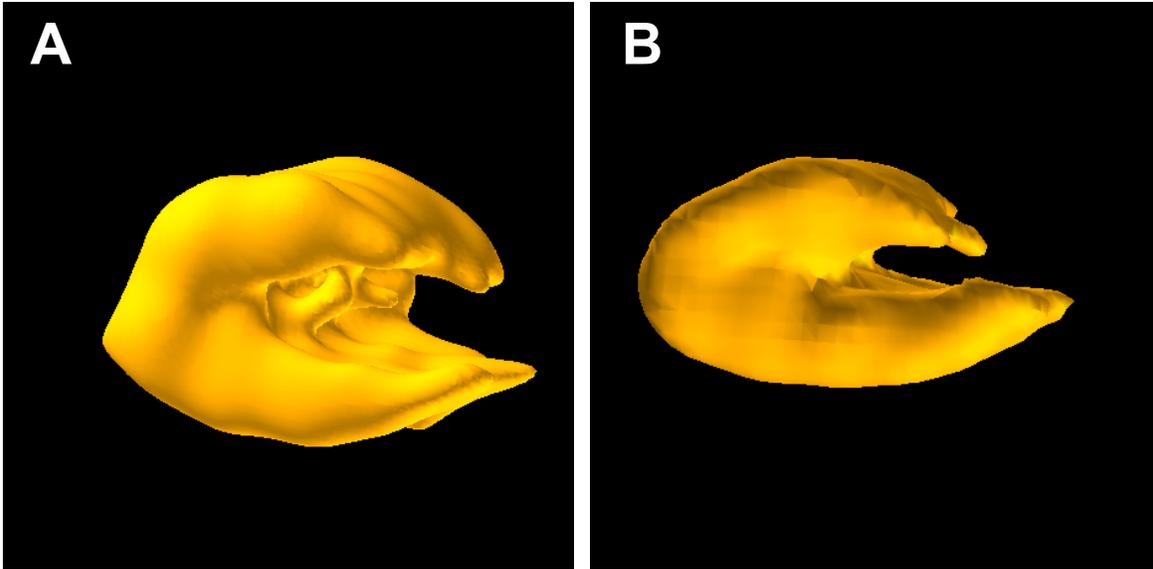


Figure 4.18: Comparisons of Cartesian and radially sampled images of a small pork rib sample. These images were taken at different times. The FOV, orientation, and zoom-factor were chosen by eye to show a rough comparison. (A) Isosurface plot of Cartesian sampled image. (B) Isosurface plot of radially sampled image.

free.

Chapter 5

Imaging Solids Faster

In the previous chapter, we discussed our approach to doing high resolution MRI of solids through the use of our quadratic echo line-narrowing sequence. This essentially increases the effective T_2 of the solid out closer to the T_1 time (which is often in the range of several seconds to minutes), enabling high spatial resolution. Unfortunately, the long T_1 time in solids, combined with the practical requirement to wait $T_{\text{rep}} \approx T_1$ between experiments, means that the imaging time for solids is very long (e.g. the image in figure 4.2 took 47.6 hours to acquire). One way to reduce the image acquisition time would be to use $T_{\text{rep}} \ll T_1$ at the cost of signal-to-noise. This is definitely not desirable if we are starting with smaller signal due to the smaller concentrations and/or gyromagnetic ratios of the target nuclei (as compared to ^1H in water). A second method extends the pseudo-FID sequence of figure 4.1 to form a pseudo-echo, which is then hit by a pulse to return spins back along the applied magnetic field. This driven equilibrium technique [75] works best if the effective T_2 is long enough so that the spins remain fairly coherent throughout the image acquisition. We have some preliminary results that this technique may work using the quadratic echo line-narrowing technique, but have yet to implement this technique for imaging (where all the gradients would have to be

reversed as well to generate the pseudo-echo). There is certainly potential to explore this technique further in future work.

The technique we believe may be the most useful method to speed up imaging is the compressed-sensing approach to MRI or sparse MRI [76, 77]. Here one essentially acquires only some of the points in \vec{k} space that the Nyquist-Shannon sampling formula demands. However, the \vec{k} -space points that are measured are sampled in a pseudo-random way such that the image artifacts due to undersampling add incoherently to look more like white noise. One can then use clever image reconstruction techniques to get the desired signal out of the resulting ‘noisy’ image. In this chapter, we discuss our work implementing sparse MRI (section 5.1) as well as the development of a novel reconstruction algorithm to reconstruct the desired image in a fast and computationally efficient way (section 5.2 - which will be more technical). The final section (5.3) will show the results using our reconstruction algorithm on sparsely-sampled MRI of solids and multi-dimensional NMR data.

5.1 Sparse MRI

In this section we discuss the ideas behind sparse sampling (compressed sensing) and how we implemented this sparse MRI approach to speed up image acquisition.

5.1.1 Sparse Sampling

To understand how sparse sampling (compressed sensing) works, consider the simulated 2D MR imaging experiment shown in figure 5.1. Here we have a physical object we want to image (in this case, the Yale College crest shown in figure 5.1A). We then apply gradients across the physical object to encode spatial information in the NMR signal which we then discretely sample to fill up a Cartesian grid of \vec{k} -space points

(figure 5.1*B*). The data is only acquired in the diamond shown (the rest is zero filled) to mimic how we would actually sample the points in 2D by using the same sampling strategy we use to acquire the octahedron of \vec{k} -space points in 3D (as discussed in the previous chapter). A 2D FFT converts this \vec{k} -space data into the image shown in 5.1*E*. This image is a good approximation of the original object, but is not perfect since we only sampled a finite number of \vec{k} -space points.

Often in conventional 2D MRI, only one dimension of the $N \times N$ \vec{k} -space data is sampled in each T_{rep} . Thus, filling this 2D grid of data requires N experiments that each require another T_{rep} waiting time. Thus, one way to go faster is to measure only one out of every four columns of \vec{k} -space data and leave the unmeasured points at zero, as in figure 5.1*C*. This yields a factor of four speed up, but, unfortunately, the FFT of this under-sampled data has very pronounced, coherent artifacts (figure 5.1*F*) which compete with the desired image. To do better, we can break up the pattern of zeroes by placing them randomly across the k-space data, as in figure 5.1*D*. As figure 5.1*G* shows, the resulting image artifacts are now incoherent and look like noise in the background of the desired image. It is important to note that in this is a noise-free simulation, so these ‘noise-like’ features in the image come purely from the incoherent artifacts due to sparse sampling. If the incoherent artifacts are not too prominent, a simple FFT of a sparsely sampled data set gives you a reasonable approximate image in much less time.

These incoherent artifacts are less prominent if the image you are reconstructing is fairly ‘sparse’ itself (meaning there is a small fraction of non-zero voxels in the image that carry all the signal information). In our quadratic echo MR images, the signal is purposely placed inside a single octant of the 3D eight-octant FOV, so we can make use of simple FFT image reconstruction without introducing huge image artifacts (as we will see in the results given in the next section). One can also potentially reconstruct

the image without these artifacts using other, more elaborate reconstruction strategies, which we will discuss further in the second half of this chapter.

5.1.2 Sparse Sampling Results

These results were taken using the same 4T animal MRI system as was used for previous imaging results at Yale University Medical School’s MRRC. Figure 5.2 shows the sparse-sampling approach applied to the bovine bone sample from figure 4.2. 66 trajectories, each with 32 acquisition windows, sample octant I of the original octahedron in \vec{k} space. To avoid artifacts due to anisotropic sampling, all three planes (xy, yz, zx) were sampled equally by using three different loop patterns, $\{Z, Y, X\}$, $\{Y, X, Z\}$, $\{X, Z, Y\}$, each making up 1/3 of the 66 trajectories. For each loop pattern, the (N_z, N_y, N_x) values are chosen randomly, subject to the total window number constraint. The trajectories followed are depicted in figure 5.2A, and the corresponding surface plot of the image following Fourier transformation (and zero-filling by a factor of four) is shown in B. For this case, the $\vec{k} = (0, 0, 0)$ point is at one corner of a 32 point \vec{k} -space data cube, and the real part of the complex image is shown. This image was acquired in only 88 minutes, a dramatic improvement over the 2,856 minutes of figure 4.2D. Despite our efforts to sample all three planes (xy, yz, zx) equally, it is still true that the points along these planes in \vec{k} -space are sampled more than points off these planes. We believe these are the primary causes for the boundaries of the image to be slightly smeared or elongated in all three directions, which you can mostly notice along the x -direction in the image because this is the longest dimension of the sample.

Since we only sampled octant I in this case, we could not make use of Hermitian symmetry for a purely absorptive, real signal. Instead there is a large dispersive imaginary part as well. As a result, only the real part of the complex image data is shown in

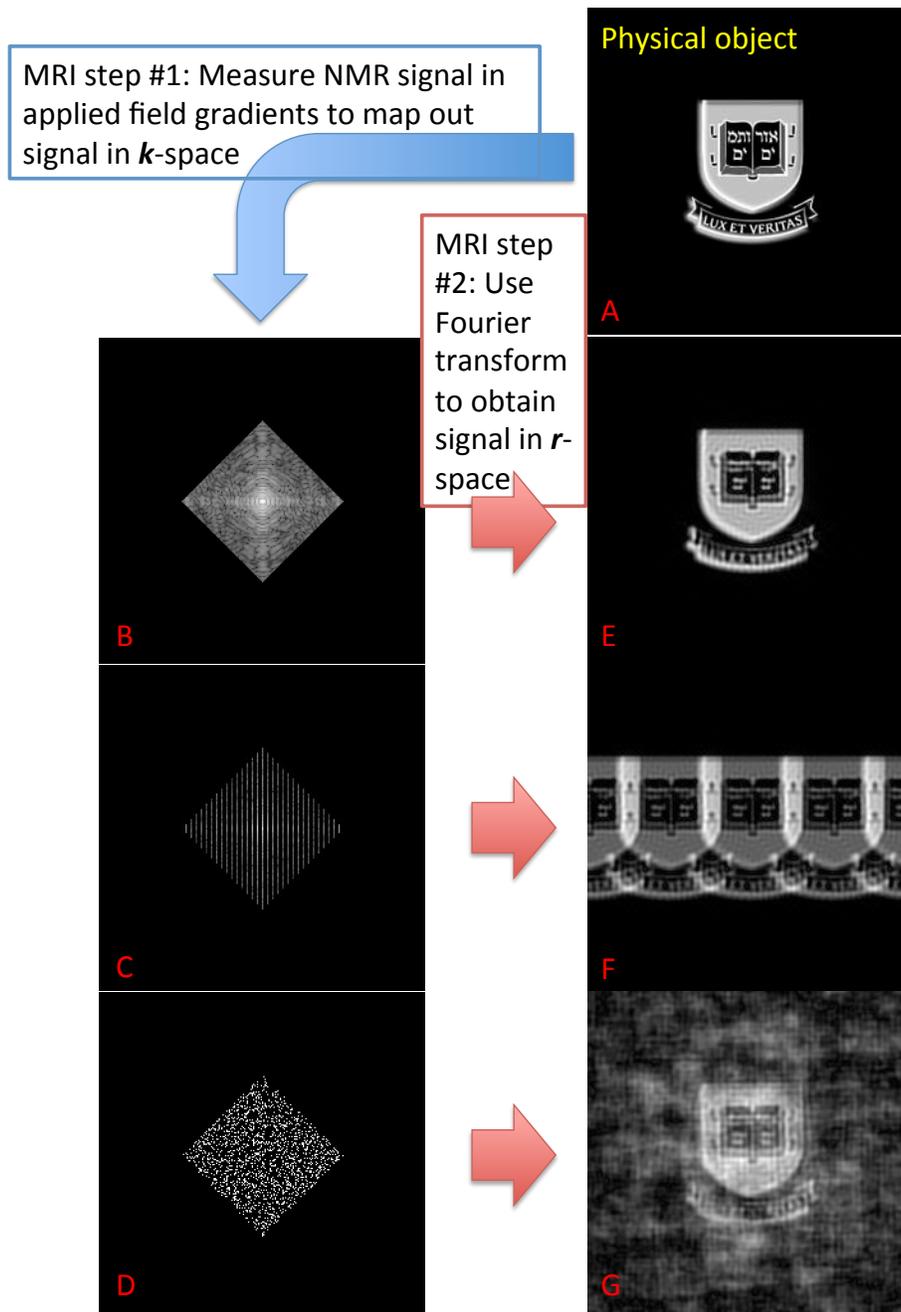


Figure 5.1: A simulation of 2D MRI, explained in the text. For a physical object (A), the conventional \vec{k} -space data set (B) yields a nice image in \mathbf{r} space (E). We can go faster by measuring only 25% of the columns in \vec{k} space (C), but this leads to pronounced, coherent artifacts in the image (F). Using a 25% random sampling of \vec{k} -space points (D) makes the artifacts due to zeroes appear as incoherent noise over the desired image (G).

figure 5.2, since the magnitude is distorted due to the dispersive imaginary component. In later experiments, we sparsely sampled in all eight octants and made use of Hermitian symmetry so the image was purely real and absorptive. In those experiments, smearing artifacts still appeared in the image, further suggesting that these artifacts result primarily from our undersampling of \vec{k} space.

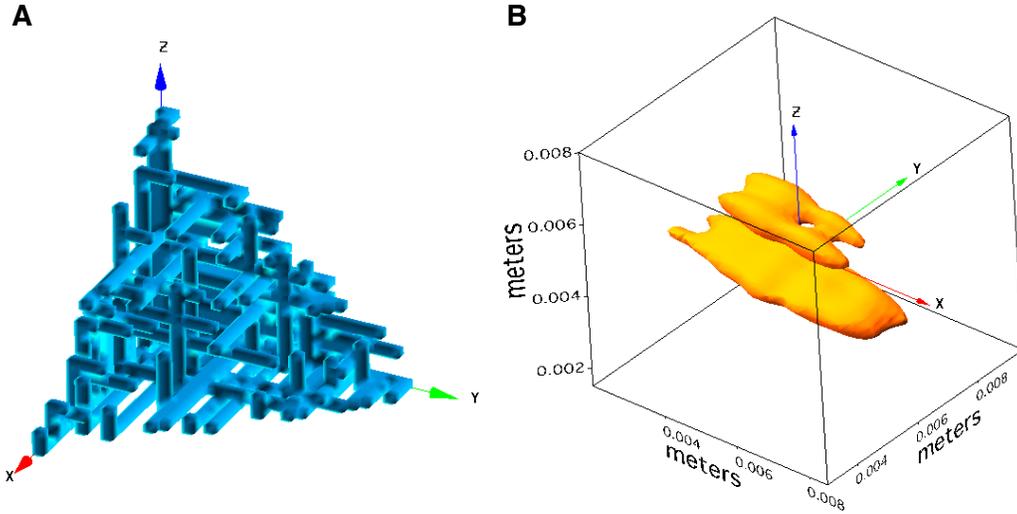


Figure 5.2: Sparse \vec{k} -space mapping and the resulting image of the dry bovine bone sample. (A) Depiction of the 66 randomly chosen sparse trajectories taken through \vec{k} -space to produce a sparse image. The three different loop patterns: $\{Z, Y, X\}$, $\{Y, X, Z\}$, and $\{X, Z, Y\}$ were used equally amongst the sparse trajectories. The loop pattern $\{Z, Y, X\}$ means after starting at $\vec{k} = (0, 0, 0)$ we map N_z points in the k_z direction, then N_y points in the k_y direction and N_x points in the k_x direction. For each loop pattern, (N_x, N_y, N_z) were chosen randomly subject to the constraint that $N_x + N_y + N_z = 31$ so that there are a total number of 32 acquisition points per trajectory. (B) The resulting image of the same sample imaged in figure 4.2 by taking the Fourier transform of the sparsely sampled \vec{k} -space. The isosurface value is 57% the maximum signal value. The imaging time was 88 minutes, compared to 47 hours for the dense sampling image in figure 4.2. The ‘smearing’ of the image is primarily due to our pseudo-random sampling of \vec{k} -space oversampling the $k_x = 0$, $k_y = 0$, and $k_z = 0$ planes compared to the rest of \vec{k} -space.

Since we believe these ‘elongation’ artifacts are primarily due to our sampling not being completely random, (e.g. certain points are sampled many times compared to

other nearby points), one potential solution is to diverge from our Cartesian trajectories where we always start along one axis, then move along another, and then finally step out in the third direction. Instead, one might consider trajectories that are 32-point long random walks along the Cartesian grid, random walks that never retrace steps, or trajectories that mimic radial sampling but stay along Cartesian grid points. We tried all these strategies, but quickly realized that the resulting images were useless because the signal from the different trajectories were not adding coherently due to the ignored shift in \vec{k} space caused by switching gradient directions. We can ignore this shift in the Cartesian sampling done in the previous chapter because we only switch gradient directions for (at maximum) two times in a given trajectory. With these other trajectories, one is potentially changing gradient directions at every step. We calculated that the shift in the \vec{k} -space position can be as much as $\delta k/4$ in the direction of the turning-off gradient, as well as a factor of $-\delta k/4$ in the direction of the turning-on gradient. If these shifts occur at nearly every step and is unique for every different trajectory, these shifts become impossible to ignore, and we are no longer sampling along a uniform Cartesian grid, but instead points nearby Cartesian grid points.

Trying to avoid these issues with sparse Cartesian sampling is what started us on the path to try radial sampling of \vec{k} space. Sparse radial sampling also causes artifacts, but the sparse sampling is done in a uniform way such that these distortions are spread uniformly along all directions and have less of a visible impact on the image (as seen in the radial results in the previous chapter). For future work, one can consider re-analyzing the more random Cartesian trajectory data taking into account all the \vec{k} -space shifts using our model of the gradient transients (discussed in the previous chapter). We could then potentially use a regridding method similar to the radial sampling analysis procedure to reconstruct the image. This may allow us to take advantage of a more ‘randomly’ sampled data set, without the corresponding errors

due to not sampling exactly along the Cartesian grid.

5.2 Reconstruction Algorithms

It is the use of the computationally efficient fast Fourier transform (FFT) for image reconstruction (as described above) that introduces the extra sampling requirements of the Nyquist-Shannon sampling theorem to produce an artifact-free image. Theoretically, using information theory, one only needs to measure as many points as are non-zero in the image we want to reconstruct. For example, in the noise-free case, if the image you are acquiring only has non-zero signal in a small number of the FOV voxels, $N_{\text{signal}} \ll N_{\text{FOV}}$ (where N_{FOV} is the total number of voxels in the FOV), one need only sample N_{signal} points in \vec{k} space in order to accurately reconstruct the image. We just need an image reconstruction algorithm to uniquely determine the desired image from the sparsely sampled data.

In MRI, there has been renewed interest in using compressive sensing (or sparse sampling) approaches [76, 78] to speed up the acquisition of imaging experiments because of the success of non-Fourier reconstruction algorithms like l_1 minimization [79, 80, 81]. Although l_1 minimization can be recast as a linear program that can be solved efficiently, these methods still require multiple matrix manipulations which greatly reduces reconstruction speed compared to the FFT algorithm one can use with a densely sampled data set. Similar reconstruction algorithms have been developed for sparse sampling methods in multidimensional NMR [82]. These include l_1 minimization methods [83], as well as other regularization methods (e.g. maximizing entropy through MaxEnt [84]). Despite the wide varieties of approaches, all these reconstruction methods yield very similar results with the same non-uniform sampled data set, so the particular sampling scheme used remains a very important determinant of the re-

construction quality [82]. These methods are also still very computationally expensive compared to Fourier reconstruction, with reconstruction times on the order of N_{FOV}^2 . In this section, we discuss a reconstruction method we developed to accurately reconstruct MR images from sparse data in a computational efficient manner by making use of the FFT and inverse FFT (IFFT) - so reconstruction time scales as $N_{\text{FOV}} \log N_{\text{FOV}}$ - as well as known constraints on the image to fill-in the data we did not sample in \vec{k} space.

To develop such an algorithm for MR data, we were inspired by the iterative-projections approach of Veit Elser [85] to solve the problem of determining the phase of a complex signal when one only has the modulus of the Fourier transform and a set of *a priori* constraints that the Fourier reconstruction of the object must satisfy. This has since been applied to a variety of other reconstruction problems [86]. Elser’s iterative ‘difference map’ algorithm unifies phase-retrieval strategies originally developed by Gerchberg [87] and Fienup [88] into a single generalized algorithm. The mathematical framework for these reconstruction algorithms has been developed only recently, but these algorithms have been shown to correspond to classical convex optimization methods [89]. The particular form of Elser’s difference map used in this work reproduces Fienup’s hybrid input-output map [88], but the more general form can be potentially useful to optimize the difference map for different uses [85, 86, 90, 91].

An alternating projections algorithm called POCS (projections onto convex sets) [92] has been used early in MRI to reconstruct phase information from limited, asymmetric \vec{k} -space data. Reconstructing sparsely-sampled MRI data has been done more recently [93, 94] using a non-iterative generalized form of Gerchberg’s alternating projections algorithm [95]. This generalized approach finds the solution to a regularized minimization problem, very similar to the minimization methods discussed above. The benefits of our iterative projections approach to sparse image reconstruction are many

fold: (1) it uses FFT and IFFT so large datasets can be handled very quickly, (2) once you set up the projections, the image reconstruction is completely deterministic, and (3) our algorithm is robust even for large noise and/or image artifacts.

5.2.1 Notation

Here we define some useful notation that will be used throughout this section. We will have the target, dense $\vec{\mathbf{k}}$ -space data represented by the complex-valued vector, $|T(\vec{\mathbf{k}})\rangle$. This vector has a total of N_{FOV} components for each of the voxels/points in the data. For example, if we have a 3D $N_x \times N_y \times N_z$ data set, $|T(\vec{\mathbf{k}})\rangle$ will have a total of $N_{\text{FOV}} = N_x * N_y * N_z$ complex components, $T(\vec{\mathbf{k}}_j)$, for each j th point in the dataset. Here $-N_i \delta k_i \leq k_i \leq (N_i - 1) \delta k_i$ for each i th dimension. The corresponding target, dense image is represented by the complex-valued vector, $|T(\vec{\mathbf{r}})\rangle$. One converts from $|T(\vec{\mathbf{k}})\rangle$ to $|T(\vec{\mathbf{r}})\rangle$ via the FFT and back again via the IFFT.

Our image domain projection expects a purely real image given by $|\tilde{T}(\vec{\mathbf{r}})\rangle$, which can be obtained from $|T(\vec{\mathbf{r}})\rangle$ by multiplication of a phase factor, as determined from the Fourier shift theorem (to account for the position of the $\vec{\mathbf{k}} = 0$ point in $|T(\vec{\mathbf{k}})\rangle$). For the general case, the (p_1, p_2, p_3) component is given by:

$$\tilde{T}(\vec{\mathbf{r}})_{p_1, p_2, p_3} = \exp(-i(\phi_{p_1} + \phi_{p_2} + \phi_{p_3})) T(\vec{\mathbf{r}})_{p_1, p_2, p_3} \quad (5.1)$$

where

$$\phi_{p_i} = \frac{2\pi(p_i - \frac{N_i}{2})(\frac{N_i}{2} - c_i)}{N_i}, \quad (5.2)$$

and $p_i = (0, 1, 2, \dots, N_i - 1)$, and $c_i = 0$ or 1 when the first point of the i^{th} dimension is acquired at $k_i = 0$ or $\frac{\delta k_i}{2}$, respectively. For our MRI data, which always samples the $\vec{\mathbf{k}} = 0$ point in all directions, $c_i = 0$. For multidimensional NMR data (which we will discuss later), c_i can be non-zero depending on how the data was acquired.

The sparsely-sampled $\vec{\mathbf{k}}$ -space data will be represented by the complex-valued vector $|S^0(\vec{\mathbf{k}})\rangle = \hat{P}_0|T(\vec{\mathbf{k}})\rangle$, where \hat{P}_0 is a projection that leaves alone the chosen sampled points and sets all other points to zero. The corresponding sparse image-domain data is given by $|S^0(\vec{\mathbf{r}})\rangle$, with the properly-phased sparse data in the image-domain given by $|\tilde{S}^0(\vec{\mathbf{r}})\rangle$.

The sparsely-sampled data is our input into our reconstruction algorithm, which runs iteratively. The reconstructed data at the n th iteration will be represented by $|S^n(\vec{\mathbf{k}})\rangle$ in the $\vec{\mathbf{k}}$ -space domain and $|S^n(\vec{\mathbf{r}})\rangle$ in the image domain. The properly-phased sparse data in the image-domain is given by $|\tilde{S}^n(\vec{\mathbf{r}})\rangle$. The final output of the algorithm after N_{iter} is given by $|F^{N_{\text{iter}}}\rangle = \hat{P}_2|S^{N_{\text{iter}}}\rangle$ (where the \hat{P}_2 projection will be explained in further detail below) and the properly-phased output image is represented by $|\tilde{F}^{N_{\text{iter}}}(\vec{\mathbf{r}})\rangle$.

5.2.2 Projections

In order to understand the algorithms given below, some simple definitions should be stated. First, a projection is any map \hat{P} of a set such that $\hat{P} \circ \hat{P} = \hat{P}$. If we have a projection, we can talk about the space of points that are unchanged under the action of the projection, or fixed points. For example, we can look at the action that maps a point onto the x axis. We can see that this action is indeed a projection because once a point is mapped onto the x axis, application of the action once more does not change the result. Points along the x axis are examples of fixed points for this projection.

We use *a priori* knowledge to set up the two projections we use for our algorithms. The first projection (\hat{P}_1) makes use of the fact that the MR image is a measurement of a physical quantity, so properly-phased, noise-free data should give an image which is real and positive. We also know the sample we are imaging has some known extent within our FOV (particularly if we are confining the signal to within a single octant of

our FOV, which we do in our 3D MRI in order to avoid spin-locking artifacts). We can then set-up a ‘mask’ support region where inside we expect the signal to be non-zero and outside the signal should be zero. We can even go farther to define known artifact regions, and whether signal should be positive or negative within the signal regions. All of these constraints on the image are enforced through the \hat{P}_1 projection in the image domain. The second projection (\hat{P}_2) enforces the constraints in the \vec{k} -space domain. This projection simply resets any measured points in \vec{k} space to their measured values, leaving any unmeasured points alone.

5.2.3 Alternating Projections

The simplest algorithm of iterative projections, called alternating projections, is simply to enforce one projection after another (e.g. $\hat{P}_1, \hat{P}_2, \hat{P}_1$, etc.) until applying projections no longer changes the output image (so the algorithm has converged on a fixed point for both projections). The alternating projection algorithm has a long history and is also known as the Gerchberg-Saxton algorithm [87] or error-reduction algorithm [88]. A depiction of the alternating projections algorithm is shown in figure 5.3 where the image we use as input to the reconstruction algorithm is the image acquired from sparse sampling the simulated 2D data set shown in figure 5.1. The output image matches nicely with the image acquired using dense sampling (5.1E) and actually may have even sharper features, because the use of constraints fills in \vec{k} -space points beyond the diamond of the dense dataset into the zero-filled region.

The alternating projections algorithm deterministically converges towards the nearest point of closest approach or intersection of the two projection spaces. The projections we use for MR image reconstruction result in convex fixed point spaces, so there will only be a single point of closest approach or intersection. For noise-free data, this

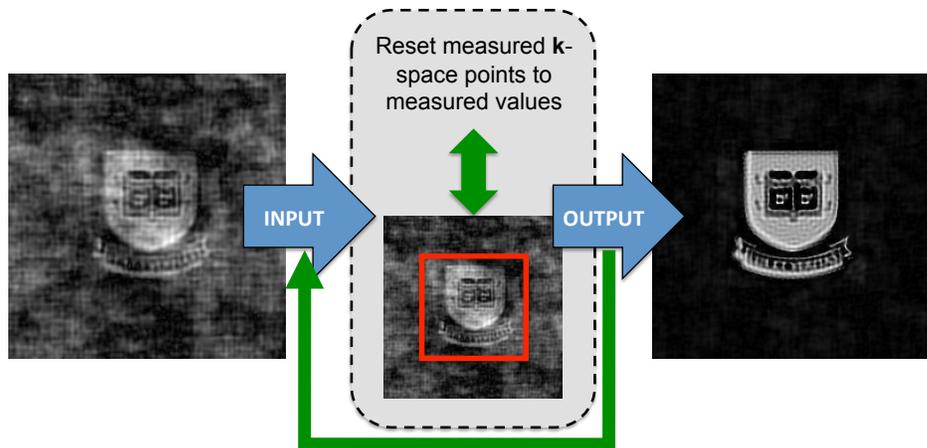


Figure 5.3: How one would use the alternating projections reconstruction algorithm to produce a high-quality image from a sparse image (such as, from figure 5.1). In each iteration, we apply the known constraints in both \mathbf{r} and \vec{k} space, and use the resulting output image as input for the next iteration. This deterministically evolves towards a final image (right).

is the unique fixed point that is in common for both projections and gives the correct image solution. Thus, the alternating projections algorithm will converge to the correct solution with noiseless data. If there is noise in the data, then a ‘gap’ opens up at the correct solution (because noise will move the fixed point from one projection away from the other), but the alternating projections algorithm still converges on the nearest point of closest approach. Since the alternating projections algorithm converges to a valid solution even in the presence of noise, there is no need to use error handling to ensure convergence. However, use of error handling (discussed later) may give a solution that provides a better fit to the ideal noise-free image. If one chooses more complicated non-convex projections (e.g. with many points of closest approach, or ‘kinks’ in the lines/planes of fixed points), the alternating projections algorithm will converge on the point of closest approach nearest to the initial starting point. In these more complicated projection spaces, the alternating projections algorithm may converge on the ‘wrong’ solution.

5.2.4 Difference Map

Elser’s difference map [85] gives a general form to combine two projections which converges on the intersection of the convex sub-spaces defined by the projections [89]. The particular form of the difference map that works best depends on the projections chosen. We found the following fairly simple combination of the two projections (which happens to be equivalent to Fienup’s hybrid input-output algorithm [88]) to be the most useful for our MR image reconstruction:

$$\hat{D} = \mathbb{I} + 2\hat{P}_1\hat{P}_2 - \hat{P}_1 - \hat{P}_2. \quad (5.3)$$

The difference map operator can be rewritten in another form closer to what we actually program (for better efficiency):

$$\hat{D} = \mathbb{I} + \hat{P}_1(2\hat{P}_2 - \mathbb{I}) - \hat{P}_2. \quad (5.4)$$

The iterative mapping is given by,

$$|S^{n+1}\rangle = \hat{D}|S^n\rangle, \quad (5.5)$$

and, after N_{iter} iterations, the final output is given by $|F^{N_{\text{iter}}}\rangle = \hat{P}_2|S^{N_{\text{iter}}}\rangle$. This final application of the \hat{P}_2 projection is used to ensure the final output still satisfies the measured points. Four iterations of the difference map is shown in figure 5.4 near a ‘gap’ between \hat{P}_1 and \hat{P}_2 fixed points.

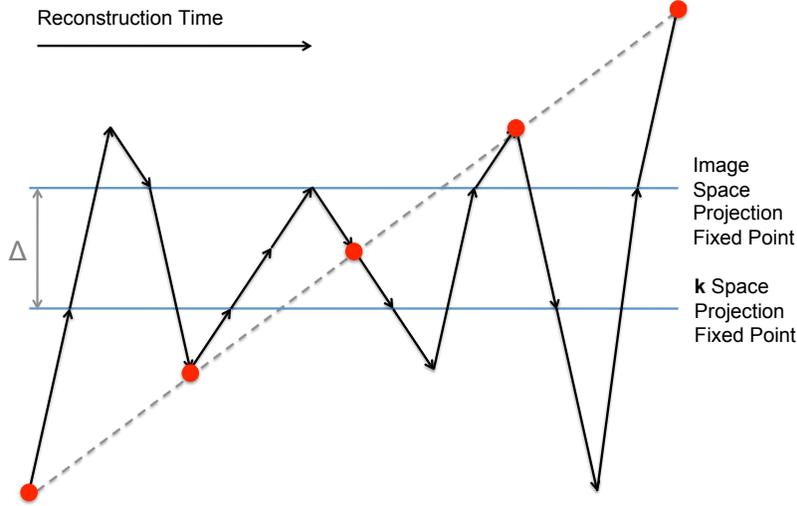


Figure 5.4: A depiction of the action of difference map algorithm without error handling near a gap. The image space projection is \hat{P}_1 and the \vec{k} space projection is \hat{P}_2 . The slope of the resulting divergence tells us something about the gap between \hat{P}_1 and \hat{P}_2 introduced due to the projections being mutually inconsistent (often due to noise or incorrect projections).

Divergence of the Difference Map

The difference map algorithm for iterative projections is designed to avoid points of closest approach and only converge on shared fixed points (the correct solution). We can thus be confident that the converged solution is the desired solution, even when more complicated, non-convex projection spaces are used to reconstruct the MR images. However, since the difference map algorithm avoids gaps, the requirements for convergence are more stringent than the alternating projections algorithm. Here we discuss the cause of divergence for the difference map and how to ensure convergence.

If the projections are mutually inconsistent, this opens a ‘gap’ between the subspaces. The difference map was designed to ‘run away’ from gaps between the subspaces in order to find a point of intersection [91]. Thus, even in the absence of noise, a mistake in the definition of either \hat{P}_1 or \hat{P}_2 will show up as divergence of $|S^n(\vec{k})\rangle$ while using the difference map. Similarly, the presence of noise in the data opens up a gap

between the two projection sub-spaces, since \hat{P}_1 and \hat{P}_2 projections that ignore noise are no longer consistent with each other.

In the limit of low, but non-zero, noise, the difference map seems to find the smallest gap Δ between the two projection sub-spaces, and then runs away in a perpendicular direction (as shown in figure 5.4). In the low-noise limit, the run away in $|S^n(\vec{k})\rangle$ doesn't seem to present a problem, apparently because the final \hat{P}_2 projection used to obtain $|F^n(\vec{k})\rangle$ brings us right back to the edge of the gap. However, in the large noise limit (e.g. our MRI of solids data), ignoring the noise produces noticeably worse results. Instead, we try to avoid the run away altogether by modifying \hat{P}_1 , \hat{P}_2 , or both to handle noise, as is explained below. Conceptually, this fills in the gap to provide a region of intersection between the two projection spaces.

5.2.5 What are the differences between the two algorithms?

There are crucial differences between the alternating projections and difference map algorithms which may determine which algorithm is best for the desired reconstruction. Alternating projections will converge to the nearest local minimum of the distance function between the two projection spaces (i.e. the nearest point of closest approach). As a result, if the distance function between the two projection spaces has multiple points of closest approach along with the desired fixed point shared by both projections, alternating projections may converge to the wrong answer (the nearest point of closest approach). The difference map, on the other hand, was originally designed by Elser to avoid these local 'traps'. The difference map thus runs away from places of close approach but converges for intersecting points. Of course, in the presence of noise, a gap opens at the point of intersection, so error thresholds are needed to 'inflate' the size of the projection spaces to give a region of overlap, otherwise the difference map will diverge. Since alternating projections already converges to the nearest point of closest

approach, the opening of gaps due to noise does not affect its convergence properties. It is worth pointing out that Elser gets rid of the effects of divergence of the difference map by applying a final projection before outputting the final reconstructed image, this brings back the diverging image to the nearest fixed point along that particular projection. As mentioned above, in our algorithm we do one final \hat{P}_2 projection before outputting the final reconstructed image. This tends to give reasonable answers in the low-noise limit without the use of error handling, but in the large-noise limit error-handling is needed for better results.

We know the projections we use for reconstructing our MR data result in convex fixed point spaces, resulting in only a single region of closest approach. Thus, for our case, both alternating projections and difference map should converge to the same point (without noise) and the same overlap region (with noise and adequate error thresholds). However, one good reason to use the difference map is the fact that it *does* diverge when the two projections are inconsistent. This actually allows us to discover if our projections are accurate and to determine good values to use for error thresholds (namely, if it converges, it is close to being a good value). This also gives us a better sense of how accurate our reconstruction is.

5.2.6 Metrics to Monitor Convergence

For large enough n (number of iterations), additional iterations of the map no longer modify the output, so any measure related to the output $|F^n\rangle$ will approach a constant as the algorithm converges. However, some measures give more useful insight into how well the output image matches the target image. Here we discuss some convergence criterion that give a helpful measure of how good the output is compared to the target output.

In the case that one has a densely-sampled data set with which to compare - as

we had when we were exploring the convergence properties of our algorithm - one can consider looking at the the difference between the output image of the algorithm at each n th iteration, $|\tilde{F}^n(\vec{r})\rangle$ and the target image resulting from the properly-phased densely-sampled data set, $|\tilde{T}(\vec{r})\rangle$. This difference can be calculated using the Euclidean (or l_2 norm) in the image domain,

$$l_2(|\tilde{T}^n(\vec{r})\rangle - |\tilde{F}^n(\vec{r})\rangle) = \sqrt{\sum_{j=1}^{N_{\text{FOV}}} |\tilde{T}(\vec{r}_j) - \tilde{F}^n(\vec{r}_j)|^2}. \quad (5.6)$$

Often times, we limit this calculation to be only over the N_{\pm} points inside the positive and negative support regions of our image domain mask. Our algorithm converges once the l_2 norm remains constant over several iterations. Ideally, the l_2 norm becomes zero once the algorithm has converged on the exact solution. Of course, in the case of noisy data, convergence to zero is nearly impossible, but the reconstructed image with the smallest l_2 norm is determined to be the best match to the target dense image. See figure 5.5D for a plot of this l_2 norm versus iteration number n for reconstruction of 2D NMR data looking at a LuxU sample with various error-handling thresholds. We will see this 2D NMR data set in more detail in later sections. This data is similar to MRI data, but instead of acquiring \vec{k} -space data, one acquires time-domain data and the image space is now frequency space.

As another measure of output quality, we take $|\tilde{T}(\vec{r})\rangle$ as a target vector, and then calculate how ‘parallel’ and ‘perpendicular’ $|\tilde{F}^n(\vec{r})\rangle$ is to that target. Specifically, the ‘parallel’ metric is given by,

$$\tilde{F}_{\parallel\tilde{T}}(n) = \frac{\langle\tilde{T}(\vec{r})|\tilde{F}^n(\vec{r})\rangle}{\langle\tilde{T}(\vec{r})|\tilde{T}(\vec{r})\rangle} \quad (5.7)$$

which should approach one, and the ‘perpendicular’ metric is given by,

$$\tilde{F}_{\perp\tilde{T}}(n) = \sqrt{\frac{\langle \tilde{F}^n(\vec{r}) | \tilde{F}^n(\vec{r}) \rangle}{\langle \tilde{T}(\vec{r}) | \tilde{T}(\vec{r}) \rangle} - \left(\tilde{F}_{\parallel\tilde{T}}^n \right)^2} \quad (5.8)$$

which should approach zero, as $|F^n(\vec{r})\rangle \rightarrow |T(\vec{r})\rangle$. For both metrics we often do the inner product only over the N_{\pm} points inside the positive and negative support regions of our image domain mask (so as not to be dominated by large artifacts). See figure 5.5C for a plot of $\tilde{F}_{\parallel\tilde{T}}(n)$ and $\tilde{F}_{\perp\tilde{T}}(n)$ versus iteration number n for the LuxU sample with various error-handling thresholds.

In actual use, one would presumably *not* have access to the dense image. In that case, other metrics can be used to determine convergence of the algorithm. One can use the l_1 norm of the reconstructed image, $|\tilde{F}^n(\vec{r})\rangle$, after each iteration, which is defined as,

$$l_1|\tilde{F}^n(\vec{r})\rangle = \sum_{j=1}^{N_{\text{FOV}}} \left| \tilde{F}^n(\vec{r}_j) \right| \quad (5.9)$$

or one can look at the entropy H of the image as defined in the MaxEnt algorithm [84]. The SIFT algorithm [96] determines the convergence by looking at the signal-to-noise ratio along particular columns of the reconstructed multidimensional NMR data. Here ‘noise’ is defined loosely to include the aliasing artifacts due to undersampling.

We also make use of examining the Euclidean (l_2) distance in \vec{k} space between the measured values and the values at those points in the current iterations:

$$R(n) = \sqrt{\sum_{j \in N_k} \left| S^n(\vec{k}_j) - S^0(\vec{k}_j) \right|^2} \quad (5.10)$$

where N_k is the set of sampled \vec{k} -space points ($N_k \leq N_{\text{FOV}}$).

It is important to note that the difference map algorithm does not try to optimize

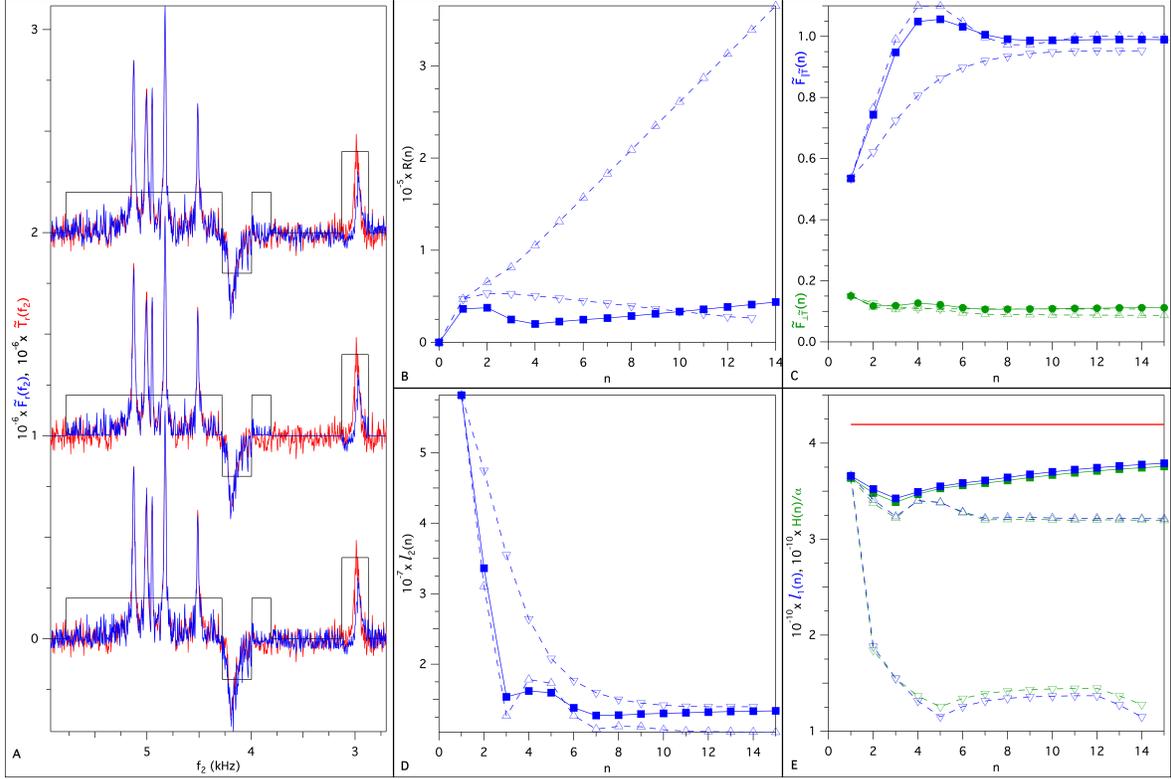


Figure 5.5: Using 2D NMR data of a LuxU sample, we compare reconstructed spectra and convergence metrics for three cases of noise-handling: (i) $f_{\text{noise}} = 45000$, $R_{\text{hyp}} = 0$, (ii) $f_{\text{noise}} = 0$, $R_{\text{hyp}} = 29185$, and (iii) $f_{\text{noise}} = 0$, $R_{\text{hyp}} = 0$ (i.e., no noise-handling). Here f_{noise} is equivalent to the 'image error threshold' discussed in the error-handling section. For each case, $\tilde{F}^{15}(f_2)$ (blue) is plotted on top of the target $\tilde{T}(f_2)$ (red), and the positive ($+2 \times 10^5$), negative (-2×10^5), and artifact ($+4 \times 10^5$) support regions are indicated (black). The vertical offset for each case is (i) 0, (ii) 1×10^6 , and (iii) 2×10^6 . Note that inside the positive and negative supports, there is very good agreement between $\tilde{T}(f_2)$ and $\tilde{F}^{15}(f_2)$ for cases (i-iii). With additional samples, the agreement within the artifact support region improves (not shown). The biggest difference is outside the supports, due to the details of the noise-handling cases. (B-E) Filled squares are case (i), down-pointing triangles are case (ii), up-pointing triangles are case (iii). (B) Examples of $R(n)$ (iii) running away linearly with iteration number due to large gap Δ , (i) rising very slowly with a greatly reduced Δ , and (ii) falling, with no gap, so the iterations terminate (at $n = 13$). (C) Plot of $\tilde{F}_{\parallel \tilde{T}}(n)$ (blue) and $\tilde{F}_{\perp \tilde{T}}(n)$ (green) versus n , for each case. (D) Plot of $l_2(n)$ distance between the output and target, as described in *Metrics to monitor convergence of $\tilde{F}^n(\vec{f})$* , for each case. (E) Plot of $l_1(n)$ norm of the output (blue) versus n for each case. The $l_1(n)$ norm of the target $|\tilde{T}(\vec{f})\rangle$ is also shown (red). Plotting $H(n)/\alpha$ for each case (green), shows the close connection between the $l_1(n)$ norm and the entropy $H(n)$ (used in the MaxEnt algorithm [84]). The value $\alpha = -590$ was picked to make $l_1(n) = H(n)/\alpha$ for the target (red).

any of these metrics, as many sparse reconstructive methods do, yet the behavior of these metrics can be illuminating. For example, figure 5.5 shows satisfactory outputs of the difference map do not necessarily minimize the l_1 norm (unlike l_1 minimization reconstruction techniques [76]) or maximize the entropy (unlike the MaxEnt algorithm [84]), but the l_1 norm does decrease and the entropy increases. The solution which minimizes the l_1 norm and maximizes the entropy is the solution which introduces no more signal to the image than is absolutely necessary to satisfy the measured data. This solution should help decrease aliasing artifacts due to undersampling, but could also possibly miss important features in the image not captured by the particular points sampled in \vec{k} space. The difference map algorithm converges on a solution that satisfies the *known* constraints on the image without making any further assumptions about what the solution should look like.

5.2.7 Choosing Error Thresholds

One can imagine including error handling in the \vec{k} -space domain, the image domain, or both. Correct error thresholds are often crucial to ensure convergence of the difference map algorithm for very noisy data sets (like our MRI data). Although alternating projections will always converge, the resulting fit can be greatly enhanced by including correct error thresholds as well into the projections. There is no good way (we have found yet) that allows you to calculate the best values for error thresholds to use for a particular data set that will guarantee you the best fit (i.e. reconstructed values for each point are as close as possible to *actual values* in the presence of no noise). Here we will address best practices of estimating good values to use for error thresholds to achieve convergence for the difference map, and then these values can be tweaked a bit to find the best fit for either difference map or alternating projection algorithms. One such method is outlined below:

1. For a given sparsely sampled \vec{k} -space data set, determine a conservative estimate for the noise of a point sampled once. Use this estimate as the individual sphere radius (defined below) for error handling in \vec{k} space.
2. Take the FFT of the sparsely sampled \vec{k} -space data to give the image before difference map reconstruction and find the average value of the noise (e.g. amplitude of signal outside the positive, negative, and artifact support regions). Use this value for the error threshold for error handling in image space.
3. The error thresholds found by the methods above should ensure convergence of the difference map algorithm within a reasonable number of iterations (often less than 20 iterations are needed). If you have issues getting the difference map to converge, try increasing either error threshold slightly. If the difference map still does not converge, check your image mask to make sure you are not missing any possible signal or artifact regions.
4. Once you have convergence, you can lower both error thresholds until the desired convergence criterion is minimized or maximized and the algorithm does not diverge. In practice, we have found that greatly lowering the \vec{k} -space error threshold often has more benefits than doing the same for the image space error threshold.

Another way to determine if the error thresholds you have chosen seem reasonable is to find the ‘gap’ distance between the fixed points of each projection at the desired fixed point. Note that this fixed point would be the same for both projections without the presence of noise. This gap size (Δ) can be found by running the difference map with no error handling (both thresholds equal to zero) and calculating $R(n)$, as defined above. The plot of $R(n)$ versus iteration number should become linear as the algorithm approaches the gap (see schematic in figure 5.4 and linear divergence in figure 5.5A).

The slope of the resulting line that sets in as the iteration number increases gives the size of the gap due to the presence of noise in the data. One can also use this method to find the size of the gap with alternating projections. When alternating projections converges (without error handling), $R(n)$ will hop between two values, and the difference between these two values gives the gap size (i.e. the difference in applying one projection versus the other). Since this gap size needs to be accounted for using error handling, the effective hypersphere radius (calculated below) and/or the image-space error threshold must be at least the size of the gap (Δ) for convergence of the difference map.

Below are more specific calculations pertaining to the use of \vec{k} -space and image-space error handling that suggest best practices in order to arrive at good fits with the difference map algorithm.

5.2.8 k Space Error Handling

Due to the presence of noise, we know the actual complex value of a given \vec{k} -space point (in the presence of no noise) should be within a sphere centered around the noisy measured value at that point. We shall call the radius of this sphere the ‘individual sphere radius’. If there is no noise in the measured data, the correct individual sphere radius to use would be zero. One might then assume that individual sphere radius should just be equal to the absolute noise level of that particular point. Interestingly, we have found that the best value of individual sphere radius to use depends on the number of points measured in the sparse data set as well. This suggests that individual sphere radius is not *just* the noise level of a sampled point, but may also depend on the reconstructed image noise (which *is* dependent on the number of \vec{k} -space points sampled, as shown later). Our reconstruction technique goes back and forth between the \vec{k} -space and image domains, so any changes in one domain affects the other in a

rather complicated fashion dependent on the \vec{k} -space points sampled as well as the image domain mask. That being said, the individual sphere radius is closely connected to the absolute noise for each measured point (and thinking about it as such can be very convenient), so knowing something about the noise in the data set can help us determine good values for the individual sphere radius.

Error handling in \vec{k} space is done by inputting a value for the individual sphere radius for the given sparsely-sampled data set. This value will quantify the distance we will allow our reconstructed data to diverge away from the measured values. Below we discuss two ways of error handling in \vec{k} space with pros and cons for each approach.

Using Individual Spheres for Error Handling

In this error handling approach, the basic idea is to relax the \hat{P}_2 projection that resets each measured point to its measured value. Instead we are happy to leave the point alone whenever the point lies within a sphere, centered on each measured value. For each \vec{k} -space projection (\hat{P}_2) we find the difference between current reconstructed complex values and the measured complex values for each sampled point. We then compare this difference to the individual sphere radius taking into account the number of times each point has been measured. The individual sphere radius is defined in a unique way taking into account that some points may be sampled more times than others, thereby reducing the absolute noise for those points by a factor of $\sqrt{m_j}$, where m_j is the number of times the j th point has been measured. The individual sphere radius then is given by $r_{\text{ind}}/\sqrt{m_j}$, where r_{ind} is the sphere radius of a single complex point sampled once. Many datasets have uniform sampling, so m_j is the same constant for all sampled points. We introduced this factor for our MRI data because some points are measured many more times compared to others, depending on the \vec{k} -space sampling method use.

A simple outline of error handling in the $\vec{\mathbf{k}}$ -space \hat{P}_2 projection using individual spheres is given below. Here $S^0(\vec{\mathbf{k}}_j)$ is the complex measured value, $S^n(\vec{\mathbf{k}}_j)$ is the complex value for each j th sampled point in the n th iteration of the reconstruction, and $|S^{n(2)}(\vec{\mathbf{k}})\rangle = \hat{P}_2|S^n(\vec{\mathbf{k}})\rangle$.

For each of the N_k sampled points in $\vec{\mathbf{k}}$ space,

1. Find $\delta(\vec{\mathbf{k}}_j) = \left| S^0(\vec{\mathbf{k}}_j) - S^n(\vec{\mathbf{k}}_j) \right|$.
2. If $\delta(\vec{\mathbf{k}}_j) > (r_{\text{ind}}/\sqrt{m_j})$ then

$$S^{n(2)}(\vec{\mathbf{k}}_j) = S^0(\vec{\mathbf{k}}_j) + \left(\frac{S^n(\vec{\mathbf{k}}_j) - S^0(\vec{\mathbf{k}}_j)}{\delta(\vec{\mathbf{k}}_j)} \right) \left(\frac{r_{\text{ind}}}{\sqrt{m_j}} \right) \quad (5.11)$$

3. If $\delta(\vec{\mathbf{k}}_j) < (r_{\text{ind}}/\sqrt{m_j})$ then $S^{n(2)}(\vec{\mathbf{k}}_j) = S^n(\vec{\mathbf{k}}_j)$.

Since all points must satisfy the criterion that $\delta(\vec{\mathbf{k}}_j) < (r_{\text{ind}}/\sqrt{m_j})$ for this algorithm to converge (i.e., none of the $\vec{\mathbf{k}}$ -space points have changed from one iteration to the next), divergence can happen simply by having at least one point never satisfy this criterion (e.g. points that are far away from the average due to statistical fluctuations in the data). For this reason, to ensure convergence one must choose $r_{\text{ind}} \geq \sqrt{m_{\text{max}}}r_{\text{max}}$ where r_{max} is the maximum individual sphere radius for the data set and m_{max} is the number of times the point with the maximum individual sphere radius has been sampled. Finding a good value to satisfy this convergence criterion must be done through estimation, trial, and error, since we do not know exactly what the actual measured values should be without the presence of noise, and we also do not know exactly how the reconstruction itself effects the sphere radius of each point. However, if there is convergence using individual spheres error handling, one then has a strong knowledge about the final reconstruction's accuracy on a point-by-point basis. Similarly, if one could pinpoint the points that are diverging, one can then change r_{ind} to r_j which can

be different for each point (e.g. r_j larger for diverging points and smaller for others). This change should help the individual spheres error handling method to achieve convergence, as well as provide a more accurate fit. However, if one wants convergence without so much sensitivity to the point-by-point noise, the next method of error handling discussed below may be more enticing.

Using Hypersphere for Error Handling

One can also do \vec{k} -space error handling using the ‘hypersphere’ which involves using the radius of the complex hypersphere made by adding together all the individual spheres used above. Essentially, the hypersphere radius is given by the square-root of the sum of the squares of the individual sphere radii. In the case of unevenly sampled data (and allowing for the absolute noise for each point to be different) the hypersphere radius is then given by:

$$R_{\text{hyp}} = \sqrt{\sum_{j=0}^{N_k} \left(\frac{r_j}{\sqrt{m_j}} \right)^2}, \quad (5.12)$$

where N_k is the number of sampled points in \vec{k} space and r_j is the noise radius of the j th complex point sampled only once. This can be rewritten as:

$$R_{\text{hyp}} = \sqrt{N_k (r^2)_{\text{avg}}} \text{ where } (r^2)_{\text{avg}} = \frac{1}{N_k} \sum_{j=0}^{N_k} \left(\frac{r_j}{\sqrt{m_j}} \right)^2. \quad (5.13)$$

A simple outline of error handling in \vec{k} space using the hypersphere is given below.

1. Calculate

$$R(n) = \sqrt{\sum_{j \in N_k} \left| S^n(\vec{k}_j) - S^0(\vec{k}_j) \right|^2} \quad (5.14)$$

2. If $R(n) > R_{\text{hyp}}$ then

$$S^{n(2)}(\vec{\mathbf{k}}_j) = S^0(\vec{\mathbf{k}}_j) + \left(\frac{S^n(\vec{\mathbf{k}}_j) - S^0(\vec{\mathbf{k}}_j)}{R(n)} \right) R_{\text{hyp}}. \quad (5.15)$$

3. If $R(n) < R_{\text{hyp}}$ then $|S^{n(2)}(\vec{\mathbf{k}})\rangle = |S^n(\vec{\mathbf{k}})\rangle$ (i.e. projection does not do anything and has converged on a solution).

Using the hypersphere radius for $\vec{\mathbf{k}}$ -space error handling allows one to have a smaller individual sphere radius for convergence compared to using individual spheres for error handling in $\vec{\mathbf{k}}$ space. Mathematically, the equivalent r_{ind} required for convergence of the hypersphere case (if all points are evenly sampled m number of times) would be $r_{\text{ind}} = \sqrt{m(r^2)_{\text{avg}}}$, whereas for convergence of the individual sphere case, $r_{\text{ind}} \geq \sqrt{m}r_{\text{max}}$. This makes sense because the convergence criterion here is much more flexible than in the case of individual spheres: you just need $R(n)$ to be smaller than R_{hyp} . This condition can be satisfied when the individual sphere radius is too small for some points (making the individual sphere error handling fail to converge), as long as it is also too large for other points to compensate.

Although using hypersphere handling allows you to converge on a solution for a smaller value of the individual sphere radius, this does not ensure that the hypersphere reconstruction will be better than the individual sphere reconstruction. If all the constraints set for the individual sphere radius are true (i.e. you are fairly certain of the absolute noise level for each point and there are not large deviations from this value), you can get a more accurate fit using individual sphere error handling. However, if you cannot make such assertions about the noise in your data, the hypersphere error handling method gives you more flexibility for convergence.

A Slight Tangent on Ideal Sampling

Presumably, if one can use a smaller hypersphere radius for the same number of sampled points, N_k , and still converge, the resulting reconstruction should be more accurate than that found using a larger hypersphere radius. If so, then it might be interesting to choose m_j to minimize the hypersphere radius needed in order to achieve convergence. For simplicity, let's assume $r_j = r$ for all sampled points, which is a good approximation if you expect all your sampled \vec{k} -space points to have the same absolute noise error if only sampled once. What values should we pick for m_j to get the minimal size of R_{hyp} if we want to take a total of $M = \sum_{j=0}^{N_k} m_j$ measurements for the desired imaging time? The constraint on m_j is a strong one that hints at the correct answer: that m_j should be the same for every point, namely a uniform sampling where $m_j = M/N_k = m$. One can see this makes sense when considering the common desire of minimizing the total effective noise determined by taking the root mean square of the individual noise values. The one point that has the highest noise will dominate the 'mean square', so one desires to have all the noise values be the same in size and as small as possible. The most sensible solution then is to evenly distribute the sampling ($m_j = m = M/N$) so that each term in the sum contributes the same amount. This reasoning also applies to the more general case when r_j can be different for each point, now we just want each term in the sum of r_j^2/m_j to contribute the same amount. This would be given by choosing m_j to be directly proportional to r_j^2 .

5.2.9 Image Space Error Handling

Noise in the measured \vec{k} -space data also introduces noise in the image domain. The resulting image noise depends only on the N_k points sampled in \vec{k} space because the unsampled points do not contribute noise (but do not contribute valuable signal either!)

As a result, the image noise should be smaller for a sparsely sampled data set (with total of N_k measured points) than the densely sampled data set (with total number of measured points N_{FOV}) by a factor of $\sqrt{N_k/N_{\text{FOV}}}$.

To handle the noise in the image-domain \hat{P}_1 projection, we use an image error threshold that determines the maximum allowed magnitude for complex points in the image *outside* the image mask as well as points *inside* the mask region that do not obey the criteria of the mask (i.e. real part of image data is positive for positive regions and negative for negative regions). Points inside the artifact region are left alone, since artifacts result from signal in the measured data and not simply noise. For many MR data sets, image noise can be constrained even further because the data is processed to give a *purely real* image (i.e. Hermitian symmetry has been enforced in the \vec{k} -space domain, and we have multiplied the image by the appropriate phase factor). One can then zero the all imaginary parts in the image as part of the image domain \hat{P}_1 projection. We will discuss here the more general case where we can expect noise in both imaginary and real parts of the image. The image after enforcing the \hat{P}_1 projection will be represented by $|S^{n(1)}(\vec{r})\rangle = \hat{P}_1|S^n(\vec{r})\rangle$.

For each j th point in the properly-phased, image-domain data,

1. Determine if the point is in an artifact region, positive support region, negative support region, or outside the image mask.
2. If point is located in an artifact region, then $\tilde{S}^{n(1)}(\vec{r}_j) = \tilde{S}^n(\vec{r}_j)$.
3. If $(\text{Re}(\tilde{S}^n(\vec{r}_j))) > 0$ and this point is located in positive mask region)
or $(\text{Re}(\tilde{S}^n(\vec{r}_j))) < 0$ and this point is located in negative mask region)
then

$$\tilde{S}^{n(1)}(\vec{r}_j) = \text{Re}(\tilde{S}^n(\vec{r}_j)). \quad (5.16)$$

4. If $\tilde{S}^n(\vec{r}_j)$ does not satisfy any of the above conditions, then

(a) If $|\tilde{S}^n(\vec{r}_j)| > \text{image error threshold}$ then

$$\tilde{S}^{n(1)}(\vec{r}_j) = \tilde{S}^n(\vec{r}_j) \left(\frac{\text{image error threshold}}{|\tilde{S}^n(\vec{r}_j)|} \right). \quad (5.17)$$

(b) If $|\tilde{S}^n(\vec{r}_j)| < \text{image error threshold}$ then $\tilde{S}^{n(1)}(\vec{r}_j) = \tilde{S}^n(\vec{r}_j)$.

The gap between the fixed points of the two projections (which, without noise, would be the same fixed point) can be overcome by either introducing error thresholds in the \vec{k} -space domain, the image domain, or both. We have found that if one increases one error threshold in one domain, they can then decrease the error threshold in the other domain to give similar results. If one uses an image error threshold smaller than the actual noise level, noise in the image can be reduced, but the boundary of the image mask will also become more apparent. Since the actual image does not have such an arbitrary boundary, lowering the image error threshold too low may cause the resulting fit to be not as accurate.

We have found that including both error thresholds can provide slightly better fits than just dealing with error in a single domain. We have also found that having a larger error threshold in image space compared to the effective hypersphere radius in \vec{k} space is more conducive to convergence than vice-versa. This may make sense because error handling in image space impacts \vec{k} space more so than error handling in \vec{k} space impacts the image. This is because the projection in image space alters more points in the image space (at minimum, all the points outside the signal and artifact regions) than the N_k points in \vec{k} space that are altered for the \vec{k} space projection. Thus, setting a slightly wrong constraint (i.e. too small an error threshold) in image space will have more of an impact that having a slightly wrong error constraint in \vec{k} space.

5.3 Image Reconstruction Results

While refining the reconstruction algorithms above, many tests were done on simulated 1D, 2D, and 3D data sets, with or without noise. Below are results using the difference map algorithm on real, noisy MRI of solids data and multi-dimensional NMR data.

5.3.1 Reconstruction Results on Sparse MRI Data

To test our reconstruction algorithms on our MRI of solids data, we used the dense image data taken of the wet pork rib sample shown in figure 4.6. We then only used a certain number of Cartesian trajectories from this data to test different \vec{k} -space sparse sampling strategies. In order to choose the particular trajectories to sample for sparse imaging, we wanted to ensure that each of the three different loop patterns: $\{Z, Y, X\}$, $\{Y, X, Z\}$ and $\{X, Z, Y\}$ were used equally amongst the sparse trajectories. Figure 5.6(*left*) shows 1584 different ways of sampling Octant I using these three different loop patterns. We then chose particular trajectories to use for sparse sampling of the data and tested to see which sampling strategies best matched the dense image after reconstruction.

The first strategy we tried was to randomly choose trajectories from the 1584 different ways to sample a single octant. This strategy was used in acquiring the sparse data set of the dry bovine sample shown in figure 5.2. Randomly choosing trajectories (using different random seeds), however, can lead to big gaps in the acquired data, which our reconstruction algorithm may find hard to fill because there are no nearby measured points. To limit any gaps, we tried a strategy to mimic uniform under-sampling of \vec{k} space, but added in some random ‘jitter’ to still enjoy the benefits of non-uniform sampling (i.e. incoherent artifacts). For each octant, the desired number of rows, N_{traj} , were uniformly distributed among the $N_{\text{dense}} = 1584$ trajectories shown

in figure 5.6(left) with uniform spacing given by $N_{\text{traj}}/N_{\text{dense}} + 1$. To add jitter, each row was randomly placed within \pm uniform spacing/2 (and we ensured no rows were repeated). This ‘quasi-uniform + jitter’ sampling strategy for choosing trajectories is equivalent to the QUEST strategy used in choosing t_1 rows in multidimensional NMR data sets, which will be explained in more detail in the multidimensional NMR section. In the MRI case, we are choosing particular 3D trajectories, so the overall coverage of \vec{k} space will be different than the sampling resulting from using QUEST to choose particular rows in multidimensional NMR data sets. Figure 5.6(*middle*) shows the 45 of the 1584 trajectories chosen to sample octant I, using this sparse sampling strategy. The sparse trajectories chosen for all eight octants is shown in figure 5.6(*left*). We also learned from our experiments that more sampling was needed along a dimension when the image support region was larger in that dimension (which matches what one would predict using information theory). Since our pork rib sample was about twice as long in the x direction compared to the y and z directions, a third sampling strategy, ‘weighted quasi-uniform + jitter’, involved weighting the number of trajectories taken from the $\{X, Z, Y\}$ loop pattern by a factor of two compared to the $\{Z, Y, X\}$ and $\{Y, X, Z\}$ loop patterns.

We can summarize the three different sampling strategies by: (1) randomly choosing trajectories with different random seeds, (2) taking trajectories in a quasi-uniform way to avoid large gaps, but then adding random jitter to increase incoherence (‘quasi-uniform + jitter’), and (3) using the same quasi-uniform strategy while purposely taking more points along the x -direction where the sample is widest (‘weighted quasi-uniform + jitter’. For a given image mask, the same total number of \vec{k} -space trajectories, and the same error thresholds: (3) gave a slightly better fit than (2), which gave a better fit than (1). However, the variation among the three sampling strategies was

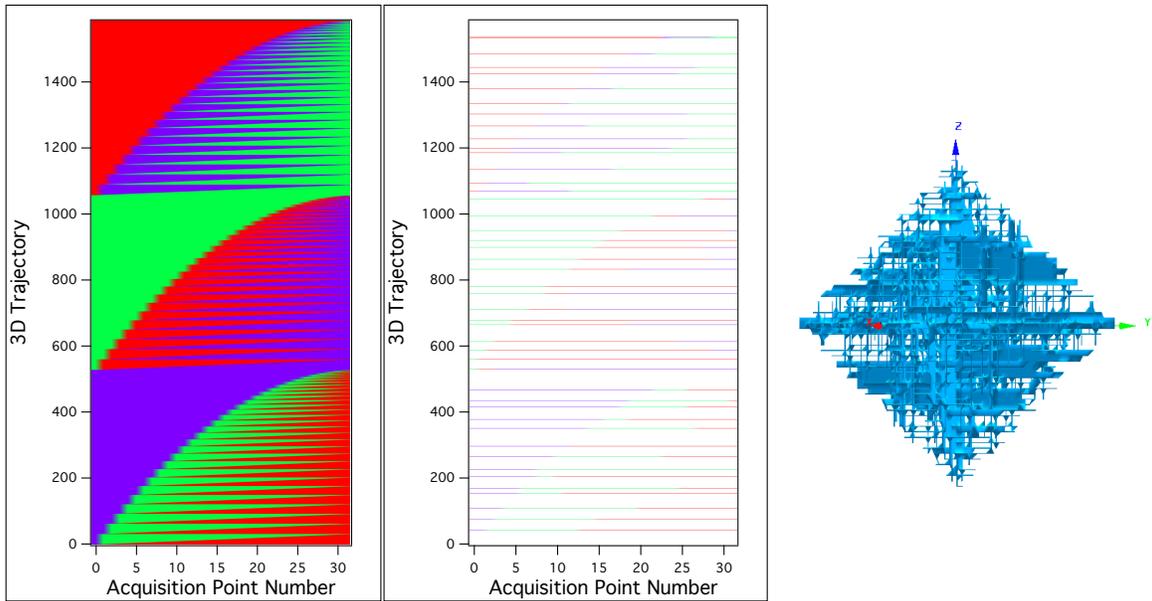


Figure 5.6: (Left) 2D image showing the different trajectories used to sample octant 1 using triple-coverage (so all directions sampled evenly). Steps in the x direction are colored red, y direction are colored green, and z direction are colored purple. Only $1/3$ of the $N_{dense} = 1584$ trajectories shown would be sufficient to densely sample Octant 1. (Middle) The $N_{traj} = 45$ trajectories chosen to sample Octant 1 using 'quasi-uniform + jitter' sampling. (Right) The full set of 360 sampled trajectories in 3D \mathbf{k} space using 'quasi-uniform + jitter' sampling.

minimal and about the same as the variation one would get simply from using different random seeds for randomly choosing trajectories or changing the error thresholds in the reconstruction procedure (see figure 5.7).

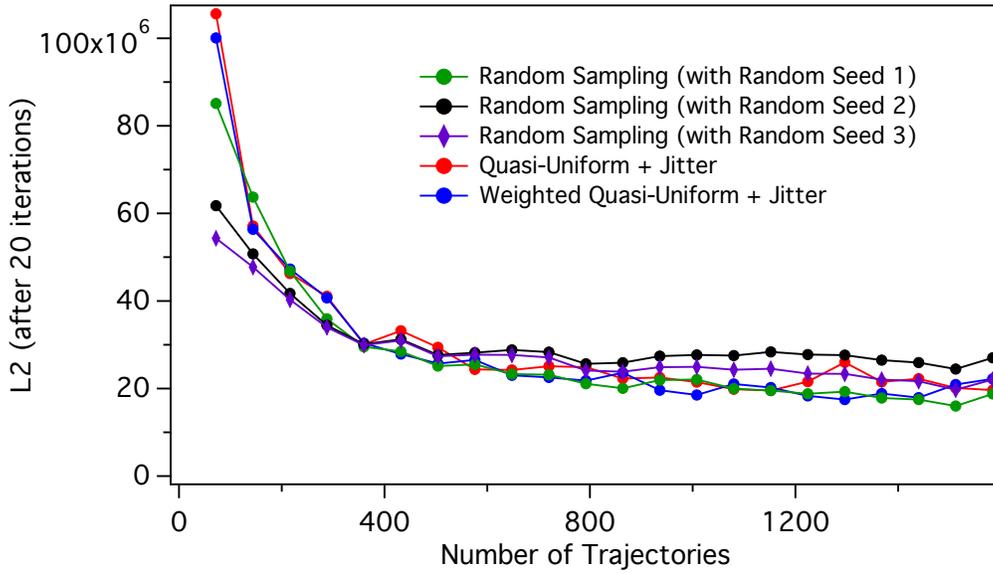


Figure 5.7: Plot of $l_2(|\tilde{T}(\mathbf{r}) - \tilde{F}^{20}(\mathbf{r})|)$ versus number of 3D trajectories used to sample the 3D MRI pork rib data set. Different strategies for choosing the trajectories were tested, including random selection, quasi-uniform sampling with jitter, and weighted quasi-uniform sampling with jitter.

These results suggest that optimizing other factors (such as error thresholds) may be more helpful than having a more specialized sampling strategy for our 3D MRI data. However, we have found some general \vec{k} -space sampling strategies help maximize performance of the reconstruction algorithm. These include (in order of importance): (1) minimizing the number of large unsampled areas in the data, (2) minimizing coherence in the sampling pattern, and (3) using knowledge of the sample dimensions to determine which directions (if any) should be sampled preferentially. These findings match well with what others have found when examining different sampling options and their corresponding artifacts [97].

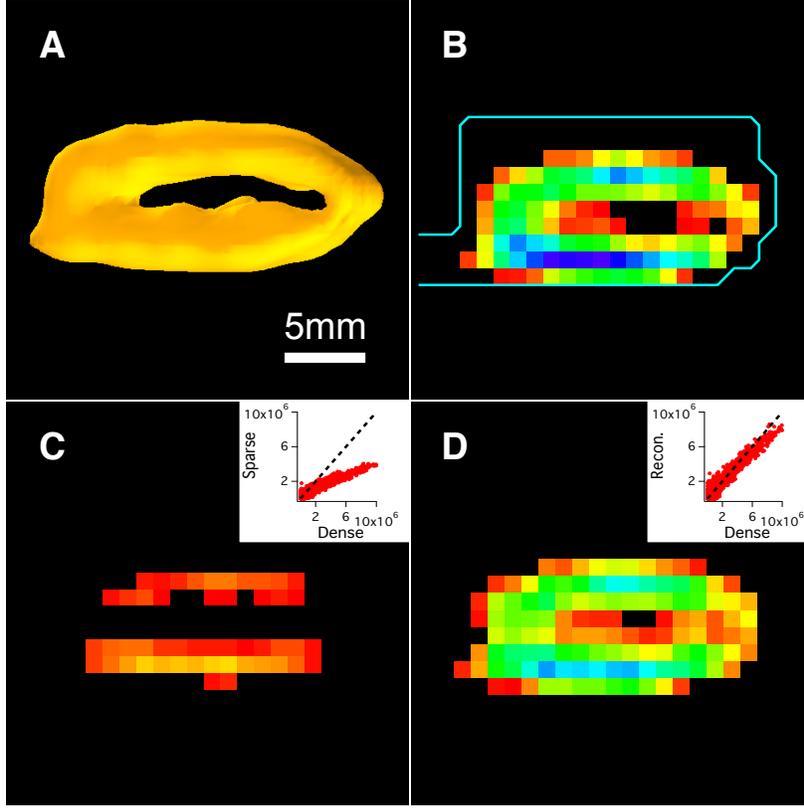


Figure 5.8: Reconstruction of sparsely-sampled solid-state MRI data. (A) Isosurface rendering of a portion of the 3D ($64 \times 64 \times 64$) $|\tilde{T}(\vec{r})\rangle$ showing ^{31}P density in pork rib in PBS solution. The isosurface value is 33% of the MDS and shows the thick cortical bone ring. The spatial resolution is $(1.19\text{mm})^3$ and T_{acq} was 35.2h. (B) A 2D slice of $|\tilde{T}(\vec{r})\rangle$ with thickness of 0.595mm. The support region for this 2D slice (where signal is expected to be positive) is outlined in blue. (C) The same 2D slice from $|\tilde{S}^0(\vec{r})\rangle$ which provided a factor of 6 speed-up in imaging time. (Inset) Plot of $\tilde{S}^0(\vec{r})$ vs. $\tilde{T}(\vec{r})$ for voxels within the support region. The thick, dashed line $y = x$ shows the poor-quality fit prior to DM reconstruction. (D) The same 2D slice of $|\tilde{S}^{20}(\vec{r})\rangle$ after DM reconstruction with \hat{P}_1 and \hat{P}_2 error handling. (Inset) Plot of $\tilde{S}^{20}(\vec{r})$ vs. $\tilde{T}(\vec{r})$ for voxels within the support region. The thick, dashed line $y = x$ shows the high-quality fit after reconstruction, where most of the points are within the 10% noise level for the measured data.

Figure 5.8 shows how well the difference map reconstruction process works on the ^{31}P MRI pork rib data (figure 5.8A). The total image reconstruction time was between two to three minutes for the various sampling strategies tried, which shows how fast our algorithm works. Figure 5.8B shows this 2D slice for the densely sampled data, $|\tilde{T}(\vec{r})\rangle$, with color thresholds chosen to highlight the thick cortical bone ring. The blue line outlines the positive support region of the image mask. Figure 5.8C shows the same 2D slice (with the same color thresholds) of the sparse image ($|\tilde{S}^0(\vec{r})\rangle$) when we use quasi-uniform sampling plus jitter to choose 1/6 of all the possible \vec{k} -space trajectories. The image signal amplitude is decreased as a result of sampling fewer \vec{k} -space points, and many features of the dense image are lost as a result. This poor correlation between the dense and sparsely-sampled images can be seen in the plot of the image point amplitudes shown in figure 5.8C(inset). We then use this image as the input to our image reconstruction procedure (in this case, using the difference map algorithm and error handling in both image and \vec{k} -space). With appropriate error thresholds, the reconstruction procedure converges and outputs the final reconstructed 3D image after 20 iterations, $|\tilde{F}^{20}(\vec{r})\rangle$, and 5.8D shows the corresponding 2D slice of this reconstructed image. Plotting the signal amplitudes of the target dense image versus our reconstructed image (figure 5.8D(inset)), we see a very good correspondence within the large noise level of this data set ($\approx 10\%$).

For figure 5.8, we used quasi-uniform sampling plus jitter to choose a total of 360 trajectories to sample all eight octants of \vec{k} space (as shown in figure 5.6(right)). The densely sampled data set (which only sampled the first four octants and used Hermitian symmetry to fill the others) took 2112 trajectories and 35.2h to acquire. Using just 360 trajectories reduces the imaging time by another factor of six. With use of our reconstruction algorithm, we have essentially gotten the same quality image as if we had acquired for six times longer.

Even more surprising, when we compare the reconstruction results of the pork rib data using our bent Cartesian trajectories versus completely random sampling of approximately the same number of \vec{k} -space points, we see nearly identical reconstruction results (see figure 5.9B). Even though our 3D sampling strategy leads to more coherent artifacts compared to completely random sampling (see figure 5.10), our reconstruction algorithm will output equally accurate images. In fact, the sparse and reconstructed image from our 3D Cartesian trajectories actually better match the dense image inside our positive support region (as can be seen in figure 5.9). This is likely because we sample more high-signal points near $\vec{k} = 0$ than random sampling, so our signal starts off higher in the corresponding sparse image (as can be seen in figure 5.10). Both completely random sampling and our Cartesian trajectories give much better results than quasi-uniform sampling *without* jitter, (see figure 5.9). Quasi-uniformly sampling the data leads to very coherent artifacts showing up in the artifact regions inside our image mask (see figure 5.10D), limiting the ability of the image space projection to suppress these artifacts.

Along with expected differences due to noise, remaining discrepancies between the reconstructed and dense images appear to be due to undersampling artifacts appearing within our (rather large) artifact region in the image mask. These points are left unchanged in the image projection, so the undersampling artifacts are never fully suppressed. The signal from these undersampling artifacts are then not fully ‘pumped’ back into the points inside the positive support region of the image mask, leading to slightly smaller signal values compared to the dense image. We see more discrepancy in high-signal points because these points produce the highest-signal undersampling artifacts (and so more signal is lost in these points due to their corresponding artifacts not being completely suppressed). The ‘tighter’ we can make the image mask (and the

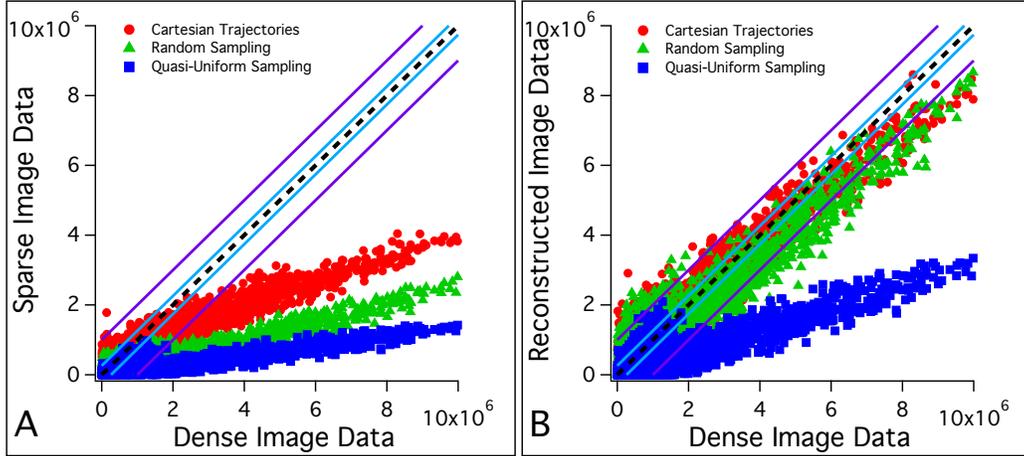


Figure 5.9: (A) Plot of the sparse image point values versus the corresponding dense image point values for different sampling strategies of 6000 points within the 3D \vec{k} -space octahedron. Only points within the positive support region of the image mask were plotted. (B) Plot of the reconstructed image point values versus the corresponding dense image point values for same sampling strategies and points shown in A. Points along the dashed black $x = y$ line perfectly match the dense image points. The blue lines show the image error threshold used in these reconstructions, and the purple lines show the 10% noise-level in the dense image.

smaller we can make the artifact regions), the smaller these discrepancies will become.

With these promising preliminary results, we are confident that in future work this reconstruction process will help speed up MRI imaging dramatically, and potentially be useful for all sorts of imaging modalities. We explore the usefulness of this reconstruction algorithm for multi-dimensional NMR data below.

5.3.2 Reconstruction Results on 2D NMR Data

Iterative-projection approaches have been developed for fast acquisition of multidimensional NMR data, most notably by Herzfeld and Matsuki who developed SIFT (Spectroscopy by Integration of Frequency and Time Domain Information) [96, 98, 99]. The SIFT approach uses the alternating projections algorithm without any form of artifact or error handling and weaker image-space projections than we have discussed above

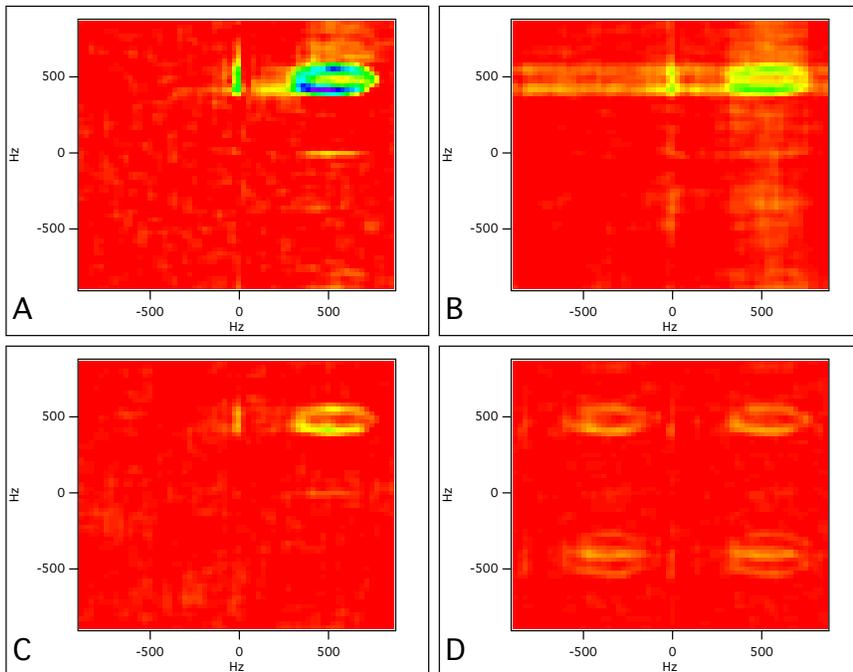


Figure 5.10: The same 2D slice from 3D images of the pork rib data resulting from different sampling strategies. The same color scale was used for all four plots. (A) 2D slice of dense image from sampling all points in \vec{k} -space octahedron. (B) 2D slice of sparse image from sampling ≈ 6000 points in \vec{k} -space octahedron via 360 Cartesian trajectories. (C) 2D slice of sparse image from randomly sampling ≈ 6000 points in \vec{k} space. (D) 2D slice of sparse image from quasi-uniform sampling ≈ 6000 points in \vec{k} space.

(because it does not assume properly-phased data with purely absorptive, real spectra). As a result, our algorithm is more likely to work better with noisier and sparser data sets, although proper phasing of the NMR data is required in order to use our stronger image/frequency-space projections. The use of the difference map algorithm may also prove more useful when non-convex projections are used that may have multiple fixed points. The difference map algorithm was designed to avoid gaps between the projection spaces and only converge at a point of intersection, whereas the alternating projections algorithm will always converge on the nearest point of closest approach.

Below we explain how we process the multidimensional NMR time-domain data to

give a purely real frequency spectrum in order to make use of our strong \hat{P}_1 projection. Note, to match the NMR data conventions, in this section we will refer to time and frequency domains instead of \vec{k} space and image domains, respectively. In the following section, we discuss advantages of the 'QUasi-Even Sampling, plus jiTter' (QUEST) sampling schedule to sparsely sample 2D NMR, and in the final section, we show results from reconstructing sparsely sampled 2D NMR of liquids and solids data.

Processing Properly-Phased Multidimensional NMR Data

We are able to use a stronger frequency domain projection (\hat{P}_1) compared to SIFT if the time-domain data is properly phased to give a purely real (i.e. absorptive) signal in the frequency domain. The 'weaker' \hat{P}_1 projection used by SIFT and others do not require the frequency-domain signal to be purely real or absorptive, so proper phasing of the time-domain data is not required. However, any hassle in properly phasing the NMR data is vastly outweighed by the many advantages of using a stronger \hat{P}_1 projection. Here we shall only discuss the 2D case, but these techniques can be easily extended to higher dimensions. Instead of the conventional purely-absorptive data reconstruction using the 'States' method [100] (which isn't easily reversible), we were inspired by \mathbf{k} -space mapping in MRI to construct a $|T(\vec{\mathbf{t}})\rangle$ that filled all four quadrants of the time domain with Hermetian symmetry about the origin, such that each component satisfies $T(\vec{\mathbf{t}}) = T^*(-\vec{\mathbf{t}})$. A 2D complex Fourier transform (followed by a phase correction) yields a purely-real spectrum. The same idea is used by the pseudo-echo transformation in 2D NMR [29].

After adjusting the spectrometer phase, data acquired using a 'States'-like 2D NMR

experiment looks something like this:

$$S_{\cos}(t_1 \geq 0, t_2 \geq 0) = \sum_{\ell} a_{\ell} \cos(\Omega_{\ell}^{(1)} t_1) \exp(i\Omega_{\ell}^{(2)} t_2)$$

$$S_{\sin}(t_1 \geq 0, t_2 \geq 0) = \sum_{\ell} a_{\ell} \sin(\Omega_{\ell}^{(1)} t_1) \exp(i\Omega_{\ell}^{(2)} t_2)$$

where ℓ is the index for an individual nuclear spin, and the signal we acquire is a sum over all these spins. To map this data into the four time-domain quadrants, we will first define:

$$\mathcal{A}(t_1, t_2) = \text{Re}(S_{\cos}(|t_1|, |t_2|))$$

$$\mathcal{B}(t_1, t_2) = \text{Im}(S_{\cos}(|t_1|, |t_2|))$$

$$\mathcal{C}(t_1, t_2) = \text{Re}(S_{\sin}(|t_1|, |t_2|))$$

$$\mathcal{D}(t_1, t_2) = \text{Im}(S_{\sin}(|t_1|, |t_2|))$$

and then the four quadrants of complex time-domain data are filled using

$$S_{\text{I}}(t_1 \geq 0, t_2 \geq 0) = \mathcal{A} - \mathcal{D} + i(\mathcal{B} + \mathcal{C}),$$

$$S_{\text{II}}(t_1 \geq 0, t_2 \leq 0) = \mathcal{A} + \mathcal{D} + i(-\mathcal{B} + \mathcal{C}),$$

$$S_{\text{III}}(t_1 \leq 0, t_2 \leq 0) = \mathcal{A} - \mathcal{D} + i(-\mathcal{B} - \mathcal{C}),$$

$$S_{\text{IV}}(t_1 \leq 0, t_2 \geq 0) = \mathcal{A} + \mathcal{D} + i(\mathcal{B} - \mathcal{C}).$$

as shown in figure 5.11. In order to fill a uniformly-spaced grid across all four quadrants, we require data acquisition to start at either $t_i = 0$ or $t_i = \frac{\delta t_i}{2}$, in both dimensions ($i = 1, 2$). Whenever States acquisition starts at $t_i = 0$, adjacent quadrants ‘double-fill’ the $t_i = 0$ axis, so those points need to be halved before loading the four quadrants

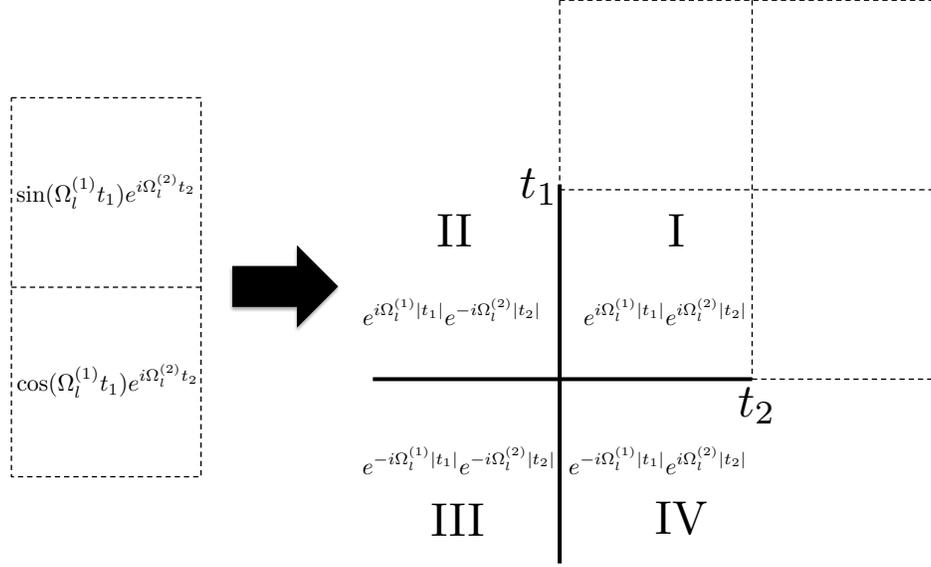


Figure 5.11: This shows how the complex 2D time-domain data acquired in a ‘States’-like experiment can be arranged into the four quadrants of the time domain, $|T(\vec{t})\rangle$. 2D NMR experiments typically acquire two sets of the $(t_1 \geq 0, t_2 \geq 0)$ first quadrant data (one with sine amplitude modulation and one with cosine amplitude modulation) and zero-fills up to the size of the four quadrants (shown in dashed lines).

into $|T(\vec{t})\rangle$.

A 2D, complex, discrete Fourier transformation of $|T(\vec{t})\rangle$ yields $|T(\vec{f})\rangle$, which can then be converted to a purely real spectrum $|\tilde{T}(\vec{f})\rangle = Ph(|T(\vec{f})\rangle)$, where the phase factor is obtained from the Fourier shift theorem (to account for the position of the $\vec{t} = 0$ point in $|T(\vec{t})\rangle$). For the 2D NMR case, the p_1, p_2^{th} component is given by:

$$\tilde{T}(\vec{f})_{p_1, p_2} = \exp(-i(\phi_{p_1} + \phi_{p_2}))T(\vec{f})_{p_1, p_2}$$

where

$$\phi_{p_i} = \frac{2\pi(p_i - \frac{N_i}{2})(\frac{N_i}{2} - c_i)}{N_i}$$

$p_i = (0, 1, 2, \dots, N_i - 1)$, and $c_i = 0$ or 1 when the first point of the i^{th} dimension is acquired at $t_i = 0$ or $\frac{\delta t_i}{2}$, respectively. All of these steps are reversible operations: $IFFT(Ph^{-1}|\tilde{T}(\vec{f})\rangle) = |T(\vec{t})\rangle$, and we make use of this fact during our iterated maps

approach.

Use of the QUEST Sampling Schedule

In many 2D NMR measurements, the experimenter waits a repetition time between acquiring rows of data (along the t_2 dimension), so we mimicked different speed up factors in experimental acquisition time by choosing a different number of t_1 rows for sparsely sampling the data. We developed the ‘QUasi-Evenly Spaced, plus jITter’ (QUEST) sampling schedule in order to pick particular t_1 rows to include in $|S^0(\vec{t})\rangle = \hat{P}_0|T(\vec{t})\rangle$ to simulate skipping particular experiments in 2D NMR acquisition. Instead of biasing samples towards low or high t_1 values, the QUEST undersampling pattern is closer to uniform spacing, which appears to have some advantages (i.e. gaps between unsampled points are minimized, which allows the iterated map algorithm to more easily fill in these gaps). To make the undersampling artifacts less coherent, we add a small amount of random jitter to determine the final row pattern. To be specific, if a dense 2D data set has N_{dense} rows with $t_1 \geq 0$, we could label them with a counter $j_{\text{dense}} = (0, 1, 2, \dots, N_{\text{dense}} - 1)$. The subset of those rows that are sampled by QUEST are described by the formula:

$$j_{\text{QUEST}} = \text{Round} \left(\frac{pN_{\text{dense}}}{N_{t_1}} \right) + \text{Round} \left(\text{Enoise} \left(\text{Trunc} \left(\frac{N_{\text{dense}}}{2N_{t_1}} \right) \right) \right)$$

Here N_{t_1} is the number of positive t_1 rows used by QUEST, $p = (0, 1, \dots, N_{t_1} - 1)$, $\text{Trunc}(x)$ keeps only the integer part of x , $\text{Enoise}(y)$ gives a random number evenly distributed within $\pm y$, and $\text{Round}(z)$ rounds z to the nearest integer. The set of j_{QUEST} values is generated as we increase p from 0. If the j_{QUEST} formula yields a row number that has already been used, we add one to that value, ensuring that no row numbers

are repeated. For the case of Figure 1, we have $N_{\text{dense}} = 128$ and $N_{t_1} = 65$. Any values of j_{dense} that are not found in j_{QUEST} are set to zero by \hat{P}_0 , and the matching pattern is used for $t_1 \leq 0$.

In order to approximate the aliasing due to QUEST sampling, we treat the simpler case of uniform undersampling of t_1 rows. Specifically, if we consider the aliases above and below the center bandwidth, every support mask feature has multiple copies, at frequencies:

$$\begin{aligned} f_{+M} &= f_1 + \text{Round}\left(\frac{N_{t_1}}{N_{\text{dense}}} \times M \times BW_1\right) \\ f_c &= f_1 \\ f_{-M} &= f_1 - \text{Round}\left(\frac{N_{t_1}}{N_{\text{dense}}} \times M \times BW_1\right) \end{aligned}$$

where BW_1 is the bandwidth for dense sampling (with N_{dense} rows), $-BW_1/2 \leq f_1 < +BW_1/2$, and $N_{t_1} \leq N_{\text{dense}}$ is the number of sparsely-sampled t_1 rows. In the figures given below, only the $M = 1$ aliases were shown for clarity. As N_{t_1} drops below N_{dense} , the copies fold over into the central bandwidth, causing aliasing.

For a given support mask, we can use this approximation to estimate the number of t_1 rows we need to take to avoid overlapping support mask regions that would limit reconstruction efficacy. For example, in figure 5.12A we plot the Euclidean distance between the reconstructed and dense spectra ($l_2[|\tilde{F}^n(\vec{f})\rangle - |\tilde{T}(\vec{f})\rangle|]$) for two different \hat{P}_1 masks, plotted versus the number of positive t_1 rows sampled (N_{t_1}). As expected, a looser mask (red curve) results in a larger difference between the reconstructed and dense spectra (larger l_2) than a tighter mask (blue curve), at each N_{t_1} . To understand the striking peaks in figure 5.12A at $N_{t_1} = 65$ and 85, we pick a \hat{P}_1 mask and study the aliasing that results from uniform undersampling (since this nicely approximates the

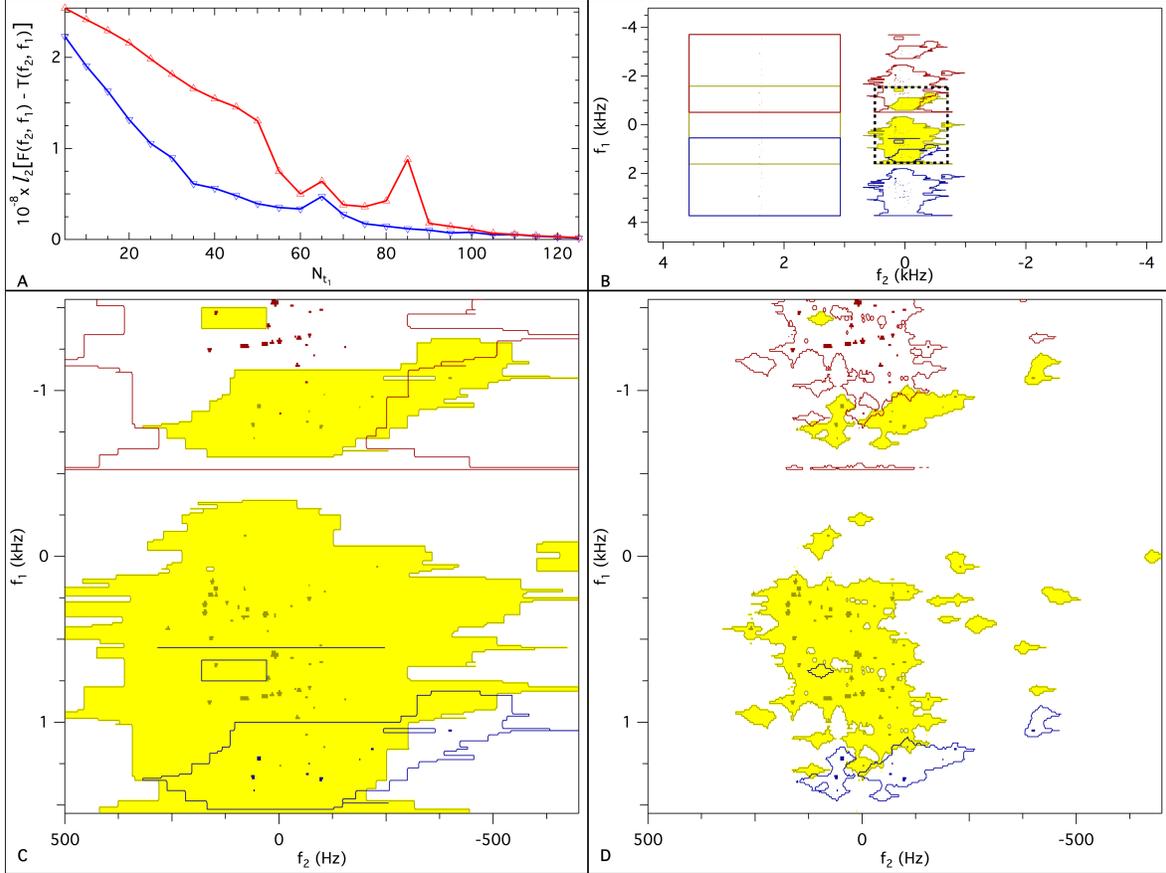


Figure 5.12: Comparing performance of the difference map algorithm using two different \hat{P}_1 masks for the IGPS data (from Figure 4.2). (A) The l_2 difference between the reconstructed and dense spectra (inside the positive support) plotted versus the number of positive t_1 rows sampled (N_{t_1}), where $N_{t_1} \leq 128$, for both looser ($> 0.14\%$ of MDS, red triangles) and tighter ($> 0.70\%$ of MDS, blue triangles) \hat{P}_1 masks. The tighter mask (blue) was used in Figure 4.2. Both traces have a peak in l_2 at $N_{t_1} = 65$, but only the looser mask (red) has a peak at $N_{t_1} = 85$. (B) The aliasing of the looser mask due to uniform sampling at $N_{t_1} = 85$. (C) Zooming into the region outlined by the instruction quality and explaining the disappearance of the peak in l_2 at $N_{t_1} = 85$ for this mask in A (blue).

artifacts from QUEST). Figure 5.12B shows the aliasing of the loose mask expected for $N_{t_1} = 85$. Zooming into the black dashed line region (in figure 5.12C) shows that many strong peaks push into the support region (highlighted in yellow) as they are aliased from above (red) and below (blue). These aliasing artifacts will not be suppressed by our \hat{P}_1 projection, resulting in a worse reconstruction (larger l_2). However, when we do

the same analysis with the tight mask (see figure 5.12D), we see that most aliased peaks miss the support region so our \hat{P}_1 projection will accurately suppress these artifacts, leading to a better reconstruction and the disappearance of the $N_{t_1} = 85$ peak for the tighter mask (red) data in figure 5.12A. In this picture, if strong aliased peaks push into the support of the central mask (as N_{t_1} is lowered), then reconstruction quality suffers. The aliased masks are a convenient proxy for the aliased peaks, and they can be used to quickly estimate the minimum N_{t_1} for excellent reconstruction using QUEST. To do this, one can plot the overlap of the aliased masks with the central mask while lowering N_{t_1} ; all overlaps close to zero will yield excellent reconstructions.

In some experiments, lower quality of the output may be acceptable in exchange for having a shorter experiment acquisition time. When pushing into that regime, QUEST can also help to identify which regions of the output spectrum will degrade first, due to aliasing artifacts. To illustrate this feature of QUEST, we analyzed a liquid-state 2D NMR data set with a LuxU protein sample (taken by our collaborators in the Pat Loria lab). We mimicked different speed up factors in experimental acquisition time by using a different number of t_1 rows: $N_{t_1} = 65$ (figure 5.13), $N_{t_1} = 50$ (figure 5.14), $N_{t_1} = 35$ (figure 5.15), $N_{t_1} = 20$ (figure 5.16). The resulting undersampling artifacts appear along the f_1 direction in the sparse spectrum as expected (see panel B). The sparse spectrum before reconstruction has a fairly poor fit with the dense spectrum, as can be seen in the red points of panel C. The most obvious discrepancies are for points that have nearly zero amplitude in the dense spectrum, but have non-zero amplitude in the sparse spectrum. These are the undersampling artifacts in regions where we expect there to be very little signal, so use of a support constraint is particularly helpful to confine these points. The mask in the frequency domain we used for our difference map reconstruction is shown in panel B. The resulting filled-in time-domain data after 15 iterations of our algorithm with frequency-domain (\hat{P}_1) noise-handling is shown in

panel *D*. The resulting reconstructed spectrum is shown in panel *E*. The reconstructed spectrum matches very nicely with the dense spectrum, as shown in both panels *C* (blue points) and *F*. The contour levels range from 6% to 72% of the maximum signal, so reconstruction is accurate over a very wide-range of signal amplitudes. Panel *G* shows how the output vector becomes more ‘parallel’ and less ‘perpendicular’ to the target vector when the algorithm works as desired. Lastly, panel *H* shows a plot in the style of figure 5.12*C*, showing the central LuxU mask (dark yellow), the $-1 \times BW_1$ mask (dark red) poking in from above, and the $+1 \times BW_1$ mask (dark blue) poking in from below.

Comparing figures 5.13-5.16, we can summarize how the algorithm seems to work, and what happens as the input data is reduced. The red points in panel *C* show that most large features start at only $\approx x\%$ of their dense value, where $x\%$ is the sparse sampling percentage; at the same time, pixels that should be zero in the dense signal are non-zero, due to the artifacts of sparse sampling. As the algorithm iterates, the artifacts are driven towards zero, while the true signals push up toward their dense values. The total area under the signal is a conserved quantity, since we always include the data point at $\vec{t} = 0$. Panel *H* shows that artifacts (blue features) that survive to the end of the algorithm are located in the portions of the frequency domain where the aliased and central masks overlap. The consequence of artifact survival is a poorer output quality, since it is also correlated with additional undershoot (red features) of the true signal amplitude in panel *H*. For clarity, panel *H* only shows the location of the $M = 1$ aliases at specific N_{t_1} values, but the formulas given earlier indicate when $M = 2$ and $M = 3$ aliases will start to matter as well. For example, at $N_{t_1} \approx 33$ the $M = 2$ aliases will just start to poke into the central mask (similar to what is seen in

figure 5.13H for $M = 1$). On the other hand, at $N_{t_1} = (50/2) = 25$ the $M = 2$ aliases will just touch in the central mask (as seen in figure 5.14H for $M = 1$). Finally, at $N_{t_1} \approx 22$ the $M = 3$ aliases will just start to poke into the central mask (as seen in figure 5.13H for $M = 1$). These ‘special’ numbers are consistent with the trends in figures 5.13-5.16, with reasonably high quality fits from $N_{t_1} = 65$ down to $N_{t_1} = 35$, and a noticeably lower quality fit for $N_{t_1} = 20$.

More Multidimensional NMR Reconstruction Results

Figure 5.17 shows our reconstruction method applied to another liquid-state 2D NMR data set from the Pat Loria lab. Using an Isoleucine, Leucine, Valine (ILV) ^{13}C -methyl labeled sample of imidazole glycerol phosphate synthase (IGPS), we were able to very accurately reconstruct the entire spectrum (figure 5.17B-C) over a wide dynamic range, starting with just 58.6% of the time-domain data. In this case, no noise-handling is used for either \hat{P}_1 or \hat{P}_2 . Note that the artifact domain (surrounded by blue in figure 5.17A) is much larger than that in figure 5.13B. In addition, this \hat{P}_1 mask has no negative support regions, and the positive support regions (surrounded by black in figure 5.17A) are the portions of the dense spectrum with $\geq 0.7\%$ of the maximum dense signal (MDS). Constructing such a tight mask would typically require the full dense spectrum, which is not available in all situations. However, in a serial experiment such as an NMR relaxation rate measurement, one dense data set (at maximum signal amplitude, to optimize the mask) could be followed by sparsely-sampled ($x\%$) data sets acquired at various relaxation delays, requiring only $\sim x\%$ of the normal experimental time, as was recently demonstrated using SIFT [99].

Figure 5.18 shows a very accurate reconstruction of a solid-state 2D NMR spectrum from a NCGB1 sample, despite starting with just 34.3% of the time-domain data and

no noise-handling. This is a very large (4096×4096 complex points) 2D data set, with sparse spectral features (including artifacts of the magic angle spinning) that span a wide range of amplitudes. For the most-part, the \hat{P}_1 positive support region is a coarse set of blocks, as we used in figure 5.13. In a crucial refinement, a ‘keyhole’ was cut into a signal-free region of the mask, to make room for an aliased rotational side-band (see figure 5.18A). Adding the tiny keyhole drove the minimum N_{t_1} from $1100 \rightarrow 702$ (53.7% \rightarrow 34.3%), while maintaining excellent quality of reconstruction over a wide dynamic range, as seen in figure 5.18B-F.

5.3.3 Further Comments

In the reconstructed images shown, we aimed for nearly ideal outputs $|\tilde{F}^n\rangle \approx |\tilde{T}\rangle$, since results of that quality can be used for any application. In the case of 2D NMR, we found that QUEST helped us to determine the minimum number of samples consistent with that goal. At first glance, our sparse sampling percentages may not seem that low, but in fact they appear to be quite close to the minimum necessary for a constrained linear system. For example, the corresponding sparse sampling percentages used in figures 5.13,5.17,5.18 are: (50.8%, 58.6%, 34.3%), which are similar to the largest percentage of positive or negative support pixels along the f_2 columns for each mask: (55.1%, 44.5%, 32.6%). Since our current 2D NMR experiments use sparse sampling in the single indirect dimension (by ‘skipping’ some t_1 values), the required sparse sampling percentages should drop quickly as this method is applied in 3D NMR, 4D NMR, etc.

The reconstructed images $|\tilde{F}^n\rangle$ also have better signal-to-noise ratios than $|\tilde{T}\rangle$, since fewer noisy samples are used at the input. The ability to sample at very long times, without requiring the acquisition of all intermediate times, can be used to achieve

higher spectral resolution in less acquisition time. The time savings offered by iterated maps could be leveraged to allow practical acquisition of higher dimensional (5D and 6D) NMR experiments, which are currently not feasible in all but the most ideal circumstances. Moreover, iterated maps offer another approach to further accelerate ultrafast 2D NMR [101, 102].

The iterative maps approach has many characteristics that should be familiar to magnetic resonance practitioners, such as its use of on-grid sampling, the FFT/IFFT, and final outputs which look the same as dense data sets in both the time and frequency domains. It complements existing methods to reconstruct spectra from sparsely-sampled data, and should find further applications in NMR and MRI of solids. More generally, any data acquisition and image modalities which make use of two reciprocal spaces related by a transformation can use this technique to harness *a priori* knowledge to fill in undersampled data. The speed, simplicity, and robustness to error makes this an ideal technique for fast analysis of noisy experimental data.

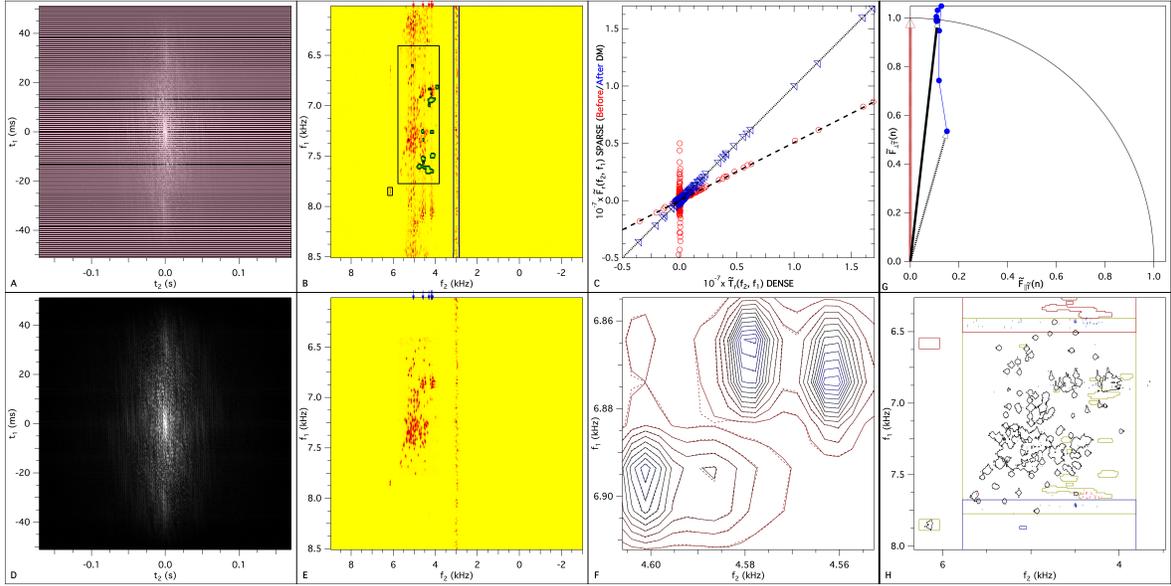


Figure 5.13: Reconstructing liquid-state 2D NMR data of the ^{15}N -LuxU sample. (A) Magnitude plot of 50.8% sparsely-sampled (t_2, t_1) -domain NMR data, $|S^0(\vec{t})|$, (256 rows \times 4096 columns). Pink shows the t_1 rows set to zero by \hat{P}_0 . (B) Real part of phase-corrected FFT of the sparsely-sampled data shown in A, $|\tilde{S}^0(\vec{f})|$, which shows the resulting aliasing artifacts along f_2 columns. Red pixels are $\geq 6\%$ of the maximum dense signal (MDS) in $|\tilde{T}(\vec{f})|$ (not shown). The \hat{P}_1 mask is shown, where blue surrounds the artifact region, black surrounds the positive support regions, and green surrounds the negative support regions. (C) Plot of pixel values along the six columns marked by red arrows in B and the blue arrows in E versus corresponding pixels from $|\tilde{T}(\vec{f})|$. The long dashed line $y = (0.508)x$ shows the poor quality of the fit before reconstruction (red open circles are $\tilde{F}^0(\vec{f})$ vs. $\tilde{T}(\vec{f})$). The short dashed line $y = x$ shows the excellent agreement after reconstruction (blue open triangles are $\tilde{F}^{15}(\vec{f})$ vs. $\tilde{T}(\vec{f})$), over the full range of positive and negative pixel values. (D) The resulting time-domain data after 15 iterations of our difference map (DM) algorithm, $|\tilde{F}^{15}(\vec{t})|$, using a small value of \hat{P}_1 noise-handling ($\pm 0.3\%$ of the MDS). (E) Real part of phase corrected FFT, $|\tilde{F}^{15}(\vec{f})|$, of the DM reconstructed data shown in D. Red pixels are $\geq 6\%$ of the MDS. (F) A portion of a contour plot showing the close match between the real parts of the dense spectrum (dashed) and the reconstructed spectrum (solid), using the same color scale and contour values (6% \rightarrow 72% of the MDS, in 6% steps) for both $|\tilde{T}(\vec{f})|$ and $|\tilde{F}^{15}(\vec{f})|$. (G) The resulting parametric plot of $\tilde{F}_{\parallel\tilde{T}}(n)$ vs. $\tilde{F}_{\perp\tilde{T}}(n)$, from the $n = 0$ point (thin black arrow), to $n = 15$ (thick black arrow), shows the approach to the target (red arrow). (H) A plot similar to figure 5.12C. The black contours are at $n_{1\%} = 1\% \times \text{MDS}$ of $|\tilde{T}(\vec{f})|$. The blue pixels are where $|\tilde{F}^{15}(\vec{f})| - |\tilde{T}(\vec{f})| \geq n_{1\%}$, and the red pixels are where $|\tilde{T}(\vec{f})| - |\tilde{F}^{15}(\vec{f})| \geq n_{1\%}$.

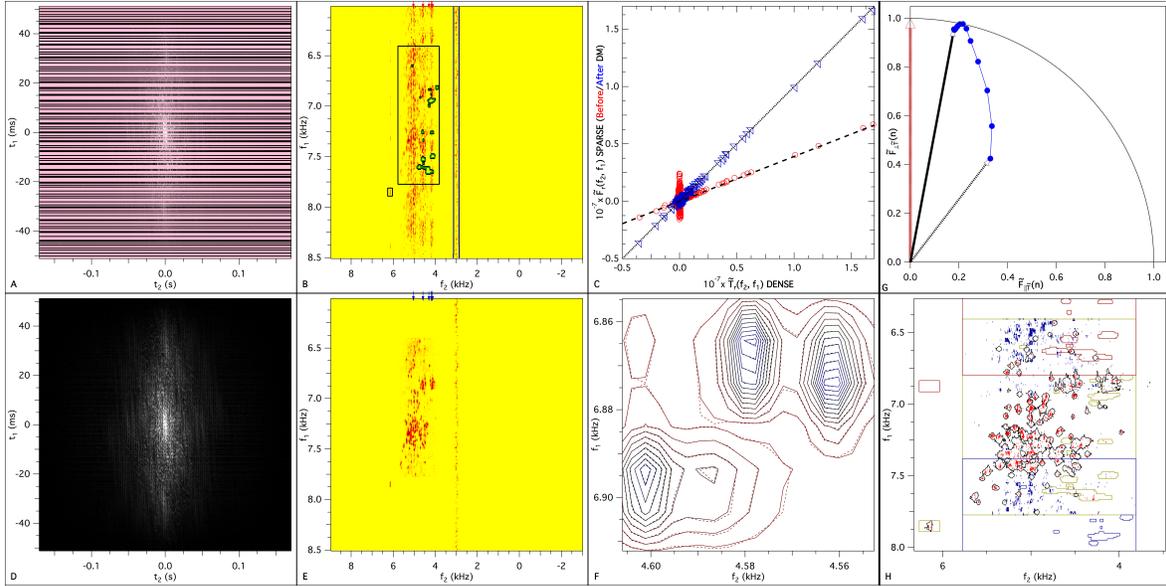


Figure 5.14: Same as figure 5.13, if we drop to $N_{t_1} = 50$ (note the increase in pink row density in *A*). The long dashed line in *C* has slope $50/128$. Still a high quality output. *H* shows more blue pixels near the overlap of masks (top and bottom of yellow box), the residual traces of sparse sampling artifacts. The red pixels are concentrated in the black contours.

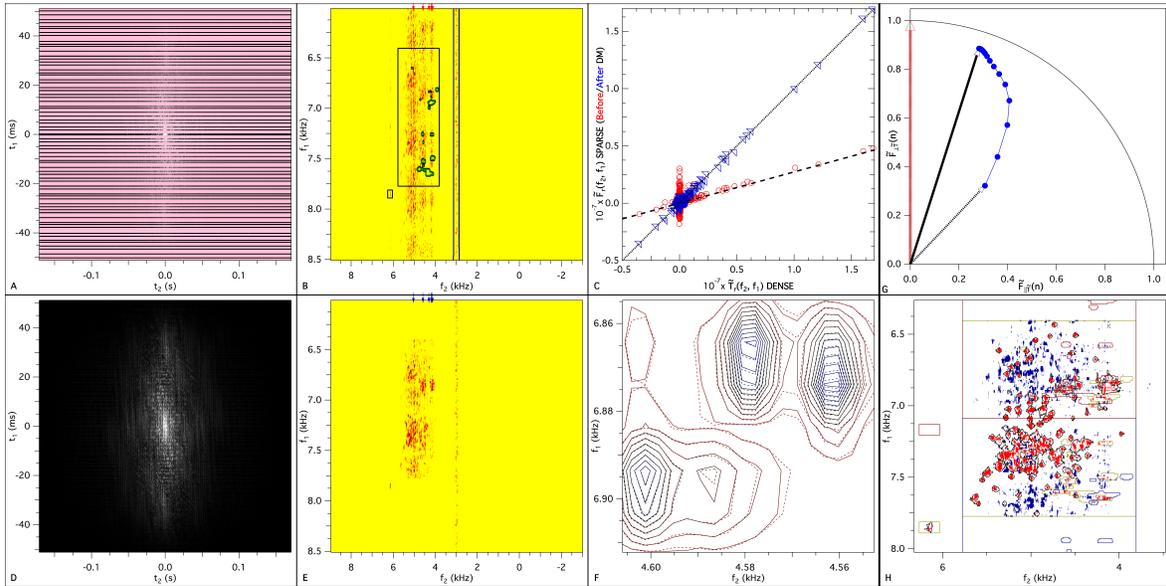


Figure 5.15: Same as figure 5.13, if we drop to $N_{t_1} = 35$. The long dashed line in *C* has slope $35/128$. This is approximately the lower end of the N_{t_1} range for high quality outputs. In (*H*), we see that the dark red and dark blue masks just touch, and so blue pixels are scattered across the central mask. Red pixels fill more of the black contours. This sparse sampling of the data requires only 27% of the normal acquisition time of the densely-sampled data set.

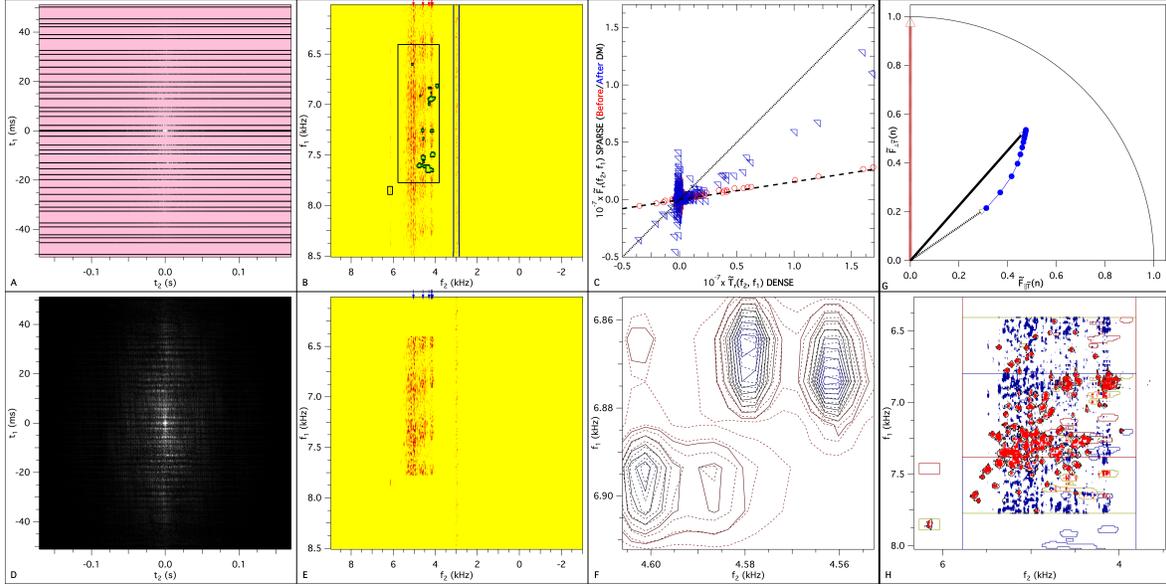


Figure 5.16: Same as figure 5.13, if we drop to $N_{t_1} = 20$, which is a lower quality output. The long dashed line in (C) has slope $20/128$. Compared to figure 5.13, most blue points in C fail to reach the short dashed line (with a slope of 1). Note also that this is the first case where many blue points in C that should be zero to match the dense image have instead grown larger than the red point values they had at the start of the algorithm (look at points at zero on the horizontal axis, in figures 5.13C-5.16C). The ‘perpendicular’ component in G grows monotonically for $n = 0$ to $n = 15$. In D, dark stripes are noticeable, along with residual aliasing within the mask in E. For this $N_{t_1} = 20$, H shows that the dark red and dark blue masks overlap, and the $\pm 2 \times BW_1$ masks (not shown) are poking into the central mask. As a result, more blue pixels are scattered across the central mask. Red pixels fill even more of the black contours. Still, C shows that the final output is better than $|\tilde{S}^0(\vec{f})|$, and E-F shows that some of the target features are recognizable, in just 16% of the normal acquisition time.

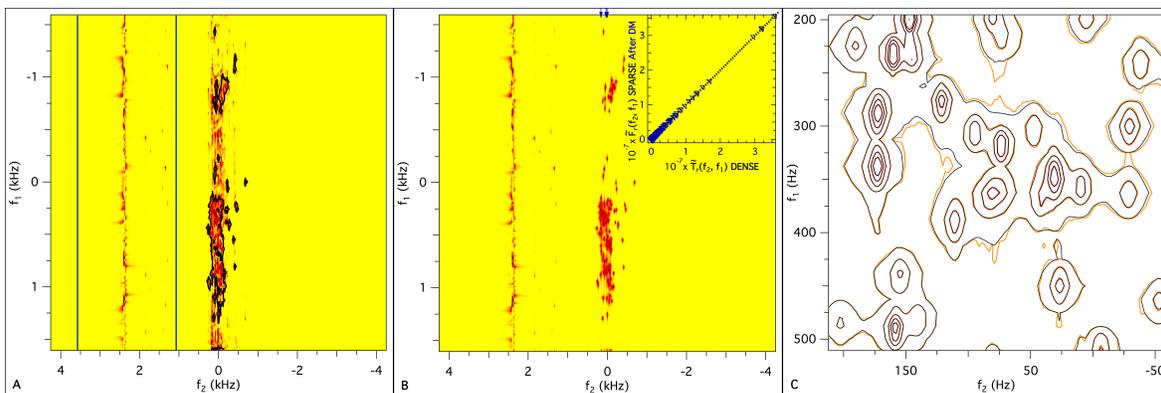


Figure 5.17: Reconstructing liquid-state 2D NMR data of the IGPS sample. (A) Real part of phase-corrected FFT $|\tilde{S}^0(\vec{f})\rangle$ of 58.6% sparsely-sampled (t_2, t_1) -domain data, which shows the resulting aliasing artifacts along the f_2 columns (256 rows \times 4096 columns). Red pixels are $\geq 6\%$ of the MDS. The \hat{P}_1 mask is shown, where blue surrounds the artifact region, and black surrounds the positive support regions ($\geq 0.7\%$ of the MDS). (B) The reconstructed (f_2, f_1) -domain data after 15 iterations of our difference map algorithm, $|\tilde{F}^{15}(\vec{f})\rangle$, without using any noise-handling. (B, Inset) Plot of pixel values along the three columns marked by blue arrows in B versus corresponding pixels from full dense spectrum. The short dashed line $y = x$ shows the excellent agreement after reconstruction (blue open triangles are $\tilde{F}^{15}(\vec{f})$ vs. $\tilde{T}(\vec{f})$). (C) A contour plot of a region comparing the real parts of the full dense spectrum $|\tilde{T}(\vec{f})\rangle$ (black) to the reconstructed spectrum $|\tilde{F}^{15}(\vec{f})\rangle$ (colored), using identical contour values (3%, 8%, 17%, 25%, 34%, 67% of the MDS) for both spectra.

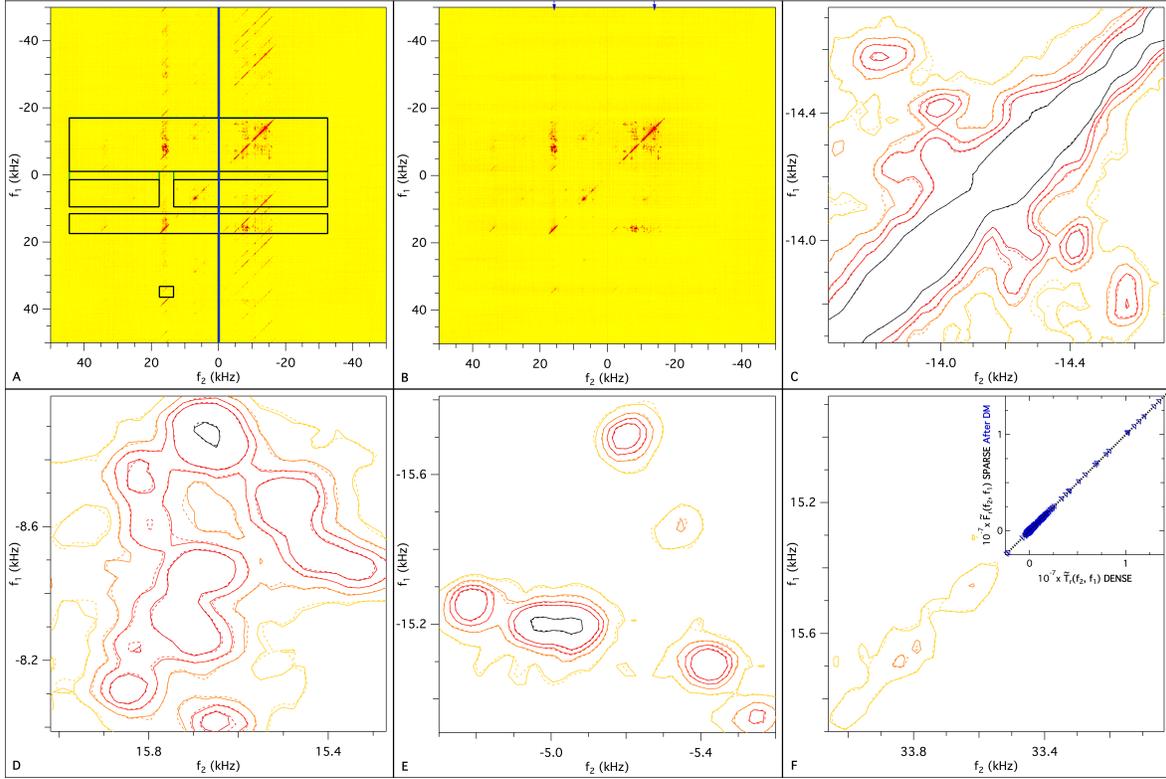


Figure 5.18: Reconstructing the ^{13}C - ^{13}C 2D MAS correlation spectrum for the $^{13}\text{C}/^{15}\text{N}$ enriched GB1 solid state sample. (A) Real part of phase-corrected FFT $|\tilde{S}^0(\vec{f})\rangle$ of 34.3% sparsely-sampled (t_2, t_1) -domain data, which shows the resulting aliasing artifacts along the f_2 columns (4096 rows \times 4096 columns). Red pixels are $\geq 2.7\%$ of the MDS. The \hat{P}_1 mask is shown, with artifact regions surrounded by blue, negative support regions by green, and positive support regions surrounded by black. (B) The reconstructed (f_2, f_1) -domain data after 15 iterations of our difference map algorithm, $|\tilde{F}^{15}(\vec{f})\rangle$, without using any noise-handling. (C - F) Zoom-in contour plots of various regions comparing the full dense spectrum (dashed) to the reconstructed spectrum (solid), using identical color scales and contour values (1%, 2%, 3%, 4%, 8% of the MDS) for both spectra. C shows a high-signal region near the main diagonal; D shows a low-signal region far from the main diagonal; E shows another low-signal region far from the main diagonal; and F shows a low-signal region where there is a rotary resonance (MAS rate ≈ 18.2 kHz). (F, Inset) Plot of pixel values along the two columns marked by blue arrows in B versus corresponding pixels from the full dense spectrum. The short dashed line $y = x$ shows the high-quality fit over a wide dynamic range after reconstruction (blue open triangles are $\tilde{F}^{15}(\vec{f})$ vs. $\tilde{T}(\vec{f})$).

Chapter 6

Future Work

As was shown in the previous chapters, we have certainly made great advances in doing higher spatial resolution ^{31}P MRI of solids, as well as explored ways to speed up the imaging process considerably. In this chapter I will discuss where we can go from here, particular with regards to enhancing resolution further using proton decoupling methods and exploring future uses of our MRI of solids technique and sparse image reconstruction algorithms.

6.1 Enhancing Resolution via Decoupling

Decoupling is a technique used to lessen the interactions between different nuclei (eg. the interactions between hydrogen nuclei and phosphorus nuclei). Hydrogen, with the largest naturally abundant magnetic moment, couples the most with other nuclei, so most decoupling schemes involve hydrogen. We found that our line-narrowing sequence in complicated materials - such as bone mineral and cell membranes - works about 1000-fold worse than in the pure silicon or buckyball samples originally examined. One important contributing factor to this lower resolution is the fact that bone mineral and cell membranes contain a lot of hydrogen which can interact with the phosphorus nuclei,

and these heteronuclear spin interactions are not effected by our pulse sequence. One way to get rid of the effects of these hydrogen-phosphorus interactions is to decouple the spins by exciting both sets of nuclei. For example, we can do our regular line-narrowing sequence on the phosphorus nuclei while constantly exciting the hydrogen nuclei, so the interactions averaged over time go to zero. This form of decoupling is often termed ‘continuous wave’ (cw) decoupling, and is the simplest decoupling method at the expense of sample heating. Other decoupling pulse sequences have been developed to make decoupling more effective and also reduce sample heating. We will first try the easy to implement cw decoupling, but further work can certainly be done exploring other more complicated, and potentially more effective, decoupling schemes.

6.1.1 Double Resonance Probe

In order to do decoupling, we first need a double resonance probe tuned to the two desired Larmor frequencies (for ^1H and ^{31}P , in our case). One can imagine just having two completely separate tuning circuits with two perpendicular sample coils around the sample (to avoid current in one coil inducing a current in the other). However, in this scheme only one coil has an optimal filling factor and the spatial distribution of the two rf fields will be different [103]. Thus, it is often desirable to doubly-tune a single coil. For our imaging case, a single coil is also preferable because then we could theoretically take both hydrogen and phosphorus images of the sample in the same exact location in the coil for direct comparison.

Overall Design

The different designs for a double resonance circuit share the common features of using frequency traps and isolation components to keep the high and low frequency parts of the circuit isolated from each other. When designing our own solid state probe for

hydrogen and phosphorus decoupling, we were ultimately deciding between three very similar probe designs [104, 105, 106, 107]. We decided to base our design off of one given by Tang, *et al.* [107], see figure 6.1, because it is supposed to be an improvement upon the conventional design and the article explained well how tuning the final circuit was done. What just remained was to adapt this design for our conditions (namely, switch from the original C^{13} nucleus to P^{31} , as well as lower the external magnetic field). See appendix D for an analysis of the main components of Tang *et al.*'s circuit design as well as theoretical estimates for the elements needed in our own probe.

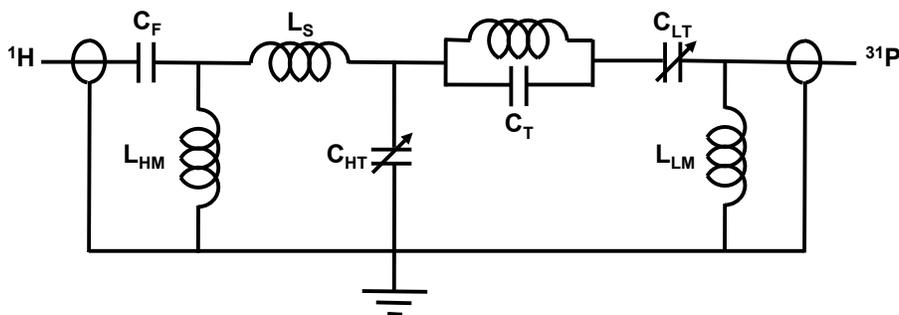


Figure 6.1: The layout of our proposed MRI of solids double-resonance circuit, based on the design of Tang, *et al.*

Final Version of Double Resonance Probe

Putting the double resonance probe together, we quickly encountered some problems not anticipated by the theory or discussed in Tang's paper. First, with such a large sample coil, the stray capacitance of the circuit (and the sample coil itself) was too high to get the high frequency circuit to tune. For this reason, we had to reduce the sample coil inductance (going from 8 turns down to 5). Also, the overall length of the wire making up the sample coil was still not a small fraction of the short wavelength at the hydrogen frequency. Our probe-designer collaborator, Pete Brown, suggested keeping this length to less than $1/16$ of the wavelength so that the voltage will be

fairly homogeneous over the entire sample coil and reducing any ‘dark spots’ where the voltage along the coil is low. To do this, we needed to break up the sample coil with a capacitor (C_{ST}). With these additions, we were able to have a good tuning range around our desired frequency. Matching was done mostly by leaving L_{HM} alone and having C_F be a variable capacitor (which also adjusted the tuning). For the low frequency circuit, tuning was now limited due to the presence of C_{ST} , so we added an inductor in parallel for the low frequency. To match the low frequency, L_{HM} was lowered. We also added, as suggested by Tang, *et al.*, another inductor to ground at the proton input for further isolation of the low frequency. To increase sensitivity in our sample coil, the high frequency trap inductance, L_T , was lowered from the original theoretical prediction, requiring C_T to be raised. Once the whole circuit looked well tuned and matched for both frequencies, C_F was fixed at the best value. Tuning with the tuning capacitors for both frequencies should work fairly well, but different loads may require matching adjustments which would have to be done manually by changing L_{HM} or C_F for the high frequency and L_{LM} for the low frequency. Figure 6.2 shows the final schematic for our double-resonance circuit. See table 6.1 for the circuit component values we used.

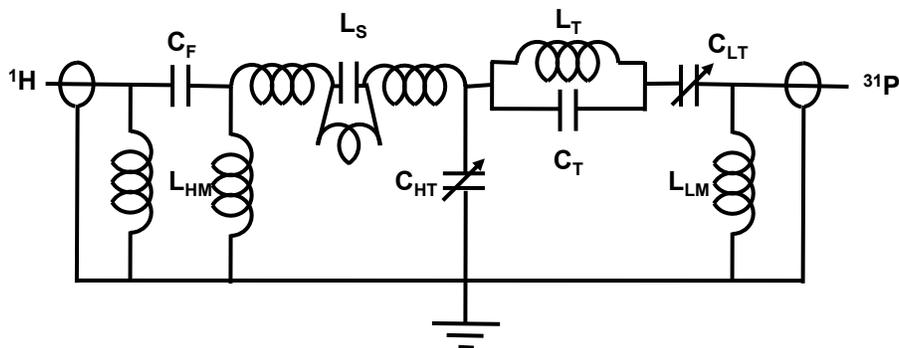


Figure 6.2: The layout of our final MRI of solids double-resonance circuit.

Table 6.1: Table of the actual (approximate) values used for our double-resonance circuit components

L_T	60nH
C_T	14.3pF
C_F	15.6pF
L_{HM}	20nH
L_S	120nH
C_{ST}	1pF
L_{ST}	20nH
C_{HT}	1-10pF variable
L_{LM} & extra isolation inductor	30nH
C_{LT}	5-25pF variable

6.1.2 Preliminary Decoupling Results

Using this probe, we were able to get high-power pulses ($T_{90} < 10\mu s$) for both ^{31}P and ^1H channels. With appropriate filters in place - a phosphorus notch filter along the phosphorus line between the probe and transmit/receive switch and a ^1H pass filter on the proton line after the proton amplifier - we were able to test our line-narrowing sequence on the wet small pork rib sample with and without cw proton decoupling. As can be seen in figure 6.3, cw decoupling extended the effective T_2 of the line-narrowing sequence, and this resulted in a spectrum narrowed by nearly a factor of three. The overall peak amplitude was also boosted by approximately a factor of four because of the line-narrowing factor and the initial $t = 0$ point starting off with a higher signal.

The $t = 0$ point determines the total integral underneath the frequency spectrum, so if both $t = 0$ points had the same amplitude, than we would expect a factor of three peak amplitude increase to go along with the factor of three narrowing. The extra boost in amplitude in the frequency spectrum comes from the time domain signal having a higher signal at $t = 0$ when decoupling is applied. This makes sense because the ‘ $t = 0$ ’ point we measure is not the *true* $t = 0$ point; it is actually taken after the first pulse block of our line-narrowing sequence ($\approx 6\Delta$ after the excitation pulse).

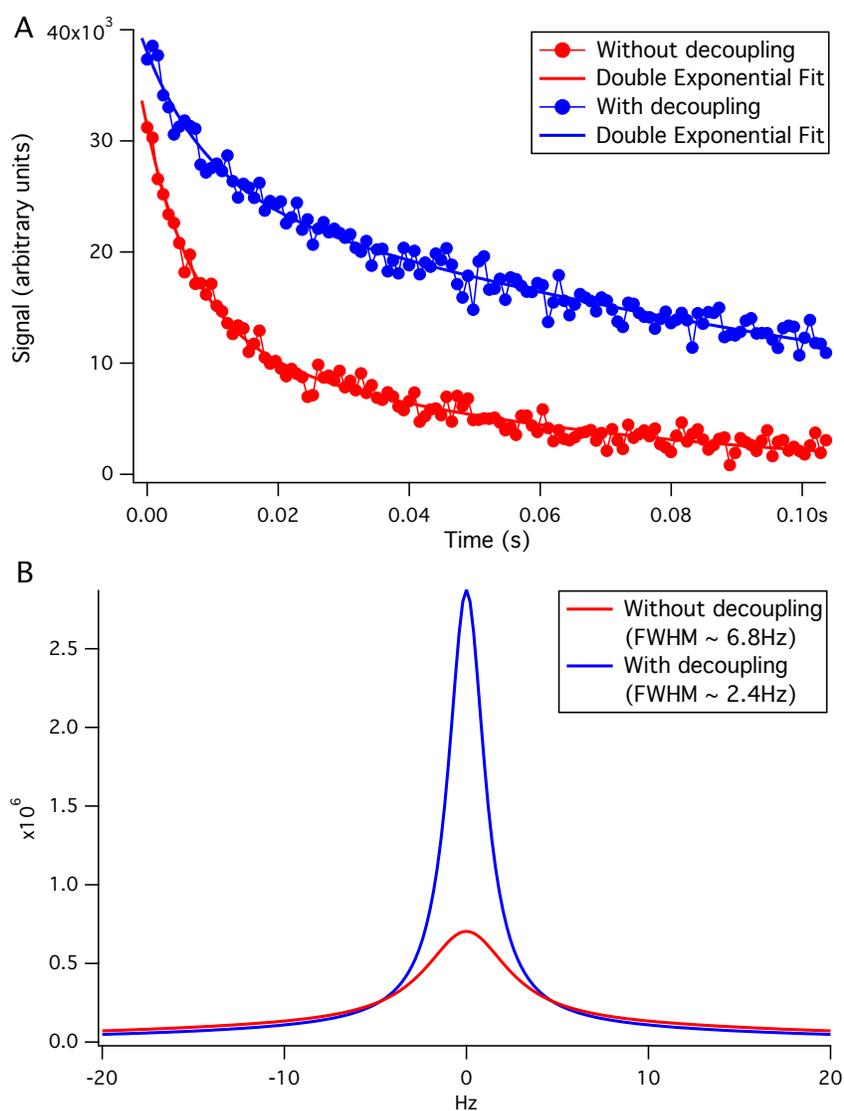


Figure 6.3: Effects of proton decoupling to enhance our line-narrowing of ^{31}P in bone mineral. (A) Time domain data using our quadratic echo line-narrowing sequence, either with (blue) or without (red) simultaneous cw proton decoupling. (B) The effective line-narrowed spectrum, either with (blue) or without (red) the use of cw proton decoupling, obtained by Fourier transformation of the double exponential fits (red and blue thick lines) to the time-domain data in A.

If we just had homonuclear spin interactions (Zeeman and dipolar interactions), then our line-narrowing pulse sequence would make the ‘ $t = 0$ ’ point taken after the first

pulse block have the same amplitude as right after the excitation pulse (*true* $t = 0$). However, in the complicated spin system of bone mineral, the interactions between ^1H and ^{31}P are also acting and left unaffected by the line-narrowing pulse block. The difference in amplitude in the two time-domain plots in the measured ' $t = 0$ ' point thus shows that decoupling is improving the natural decay of the signal due to these heteronuclear interactions over the short 6Δ time period before we acquire our ' $t = 0$ ' point.

This preliminary data supports our hypothesis that we can enhance spatial resolution through use of proton decoupling in our bone mineral samples. The next steps would be to use cw decoupling for imaging. To make use of the longer effective T_2 to provide higher spatial resolution, we would need to take more points for each pseudo-FID in the time domain. This would greatly increase the imaging time (e.g. if one doubles the number of points in the pseudo-FID to double the image resolution, one must take eight times as many points for a 3D image). Here we would need to make use of sparse imaging and reconstruction techniques to image in a reasonable length of time. Future steps to enhance spatial resolution would be to use more complicated decoupling sequences to push the line-narrowing further. A factor of three enhancement is far from the theoretical factor of 1000. However, this preliminary data was taken using the simplest form of decoupling, and many future experiments can be done exploring other decoupling schemes.

6.2 Future Applications of MRI of Solids

This thesis work focused predominately on biomedical samples since we were interested in enhancing current MRI techniques predominately used in medicine. However, our MRI of solids technique can be used on a wide variety of samples in a wide variety of

scientific disciplines. In this section, we discuss some possible future applications of our MRI of solids technique.

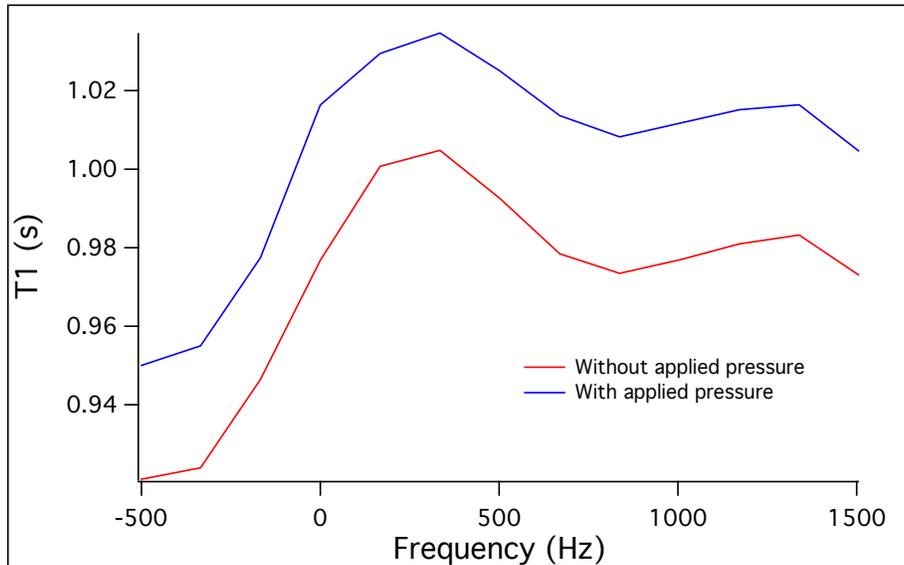


Figure 6.4: Change in T_1 with pressure in a sample of orange PDMS-based rubber inside glass cylinder at 4 Tesla. The measured T_1 is plotted over the portion of the frequency spectrum where the peak of the signal is located, before applying pressure (red) and after applying pressure (blue). T_1 for each frequency was determined from 15 saturation recovery experiments using a wide range of delay times.

6.2.1 Probing Granular Material

Granular materials display complex spatial response to applied pressures including anisotropic stress propagation along force chain networks. Imaging the response of granular media at the particle scale is crucial for understanding their mechanical behaviour during loading. Two-dimensional (2D) imaging techniques have been used for many years, but 3D imaging remains a significant challenge. Most imaging techniques rely on resolving deformation from stress-induced birefringence in an elastic medium. These techniques can not be easily extended to 3D packings. Our MRI of solids technique may be helpful in imaging the location of particles as well as contact forces

between particles in 3D granular material.

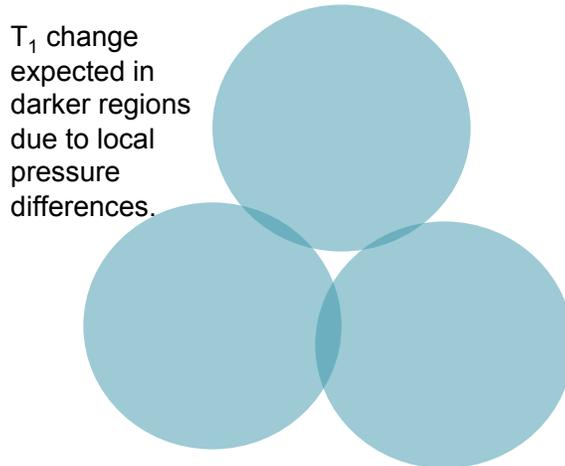


Figure 6.5: Imaging the contact forces between granular particles in 3D by using T_1 contrast and our MRI of solids technique. In certain samples (e.g. C_{60} and rubber), T_1 has been shown to change with pressure and we can use this as a local reporter of forces between granules.

Contrast in conventional MRI images of ^1H result from different T_1 or T_2 relaxations in the sample, and these relaxation times depend on the local environments of the nuclei being imaged. Studies have already shown pressure dependence of the T_1 relaxation of ^{13}C in C_{60} samples [108] as well as for ^1H in rubber [109]. For even rather modest applied pressures, we have seen similar dependence for ^1H in natural and PDMS-based rubber (see figure 6.4). To boost the change in T_1 with pressure, we plan to add super paramagnetic particles to the rubber during mixing. The pressures between granules can also be enhanced by making granules with a rubber coating about a hard core. One can then theoretically use the local T_1 changes of the granular particles to report on locally applied forces (see figure 6.5) and map out the spatial force distributions throughout the granular material. The smaller dipolar couplings of ^1H in rubber (due to the polymer dynamics) means our line-narrowing technique may be effective at providing higher-resolution images. Thus, there is potential to resolve both positions

and forces in granular media in 3D using our MRI of solids technique.

6.2.2 Other Potential Samples of Interest

Some other potential samples and applications for our MRI of solids technique include mapping out the porous channels in rock or using T_1 or T_2 contrast imaging in silicon to find broken circuit connection in 3D stacks of silicon chips.

Mapping the pore-network inside carbonate or sandstone rock is very important for determining promising oil and natural gas sites, and for geosciences in general [110, 111]. Many groups are exploring the use of micro-CT for determining these pore networks [111, 112], but the use of NMR has been limited to measuring the diffusion of water through the sample [110, 113]. Our MRI of solids technique may offer another way to image the material itself through the signal from ^{13}C in carbonate rocks or ^{29}Si in sandstone rocks.

In order to continue the Moore's Law trend to fit more and more transistors on smaller and smaller microchips, a new strategy for filling space with transistors will need to be implemented. One implementation is to move beyond 2D chips into 3D stacks of chips. However, ensuring good connections between chips then becomes vital, and finding this missed connection inside a 3D opaque object proves challenging. Currently, engineers recursively cut the faulty piece into four pieces until they determine the location of the faulty connection. If one could do this non-destructively by just looking at a 3D MRI image, a lot of physical labor could be saved. The currents through connected wires induce magnetic fields in the surrounding silicon. There are currently groups attempting to mapping these very small induced magnetic fields using superconducting quantum interference devices (SQUID) [114]. However, measuring these magnetic fields inside a 3D chip becomes more challenging the larger the chip becomes since the SQUID is measuring fields at the outside surface. It would be very

helpful to see ‘inside’ the 3D chip, for which MRI of solids techniques will be helpful. The induced magnetic fields from currents through the wires would change the local relaxation properties in the nearby silicon and could possibly be imaged using T_1 or T_2 contrast imaging. A broken connection will have no current flowing and no induced magnetic fields, so one should see a difference from its working counterparts.

Of course, these potential applications have not yet been demonstrated, but they still serve as examples of how this work can be applied in a variety of different fields to answer a variety of interesting scientific questions.

6.3 Future Applications of Image Reconstruction Algorithms

Along with the MRI of solids techniques discussed in this work, there are many broader applications of the sparse image reconstruction algorithms we have developed. In this section, we discuss just a few of the potential uses of these reconstruction algorithms for a wider variety of imaging modalities.

Along with sparse Cartesian sampling in MRI, one can imagine using our reconstruction algorithms to help regrid radially-sampled data (sparse or dense). Most fast image reconstruction algorithms involve regridding to a uniformly-spaced Cartesian grid, but this introduces streaking artifacts to the image. Using image constraints to help fill in the gaps in the regridded data may lead to higher quality images. One would need to first figure out how to satisfy the projection to return measured points to their measured values, since the regridded points no longer correlate to directly measured points. However, one could use a model of how the measured points contribute to each point on a grid (weighted by the distance of the grid point to each measured point) and allow for noise variation using methods similar to the error handling meth-

ods discussed in this thesis work. Radial sampling also does not only pertain to MR imaging. Computed tomography (CT) acquires data through 1D projections just like radial sampling, and so these better image reconstruction methods could potentially help CT imaging as well.

We have already shown how useful our reconstruction technique is for speeding up acquisition of multi-dimensional NMR data. An obvious next step is to apply this reconstruction technique for multi-dimensional NMR experiments measuring various relaxation times (usually taking a 2D spectrum with different delay times in the 3rd dimension), as was done with SIFT [99]. These 2D spectra all look the same, but the amplitudes of the peaks change. One can then imagine densely sampling the first 2D spectra to get a very accurate image mask, and using that mask to reconstruct the rest of the slices sampled sparsely to speed-up total acquisition time. With faster acquisition time, higher-dimensional (4D, 5D) NMR experiments can be taken without a significant time constraint.

Our reconstruction techniques are designed for imaging modalities that make use of Fourier transforms from the measured reciprocal space to the image domain. Thus another imaging modality that can make use of these algorithms is x-ray crystallography. Crystallography naturally has many symmetries that could be used to form tight constraints on the data and enable very sparse sampling to be done.

6.4 Final Thoughts

The goal of this thesis work was to implement the quadratic echo line-narrowing pulse sequence for use in MRI of solids. In this process we have uncovered more potential uses for this technique, as well as developed a sparse image reconstruction algorithm that is widely applicable. At the start of this project, we knew in theory that the

quadratic echo line-narrowing technique should work for imaging, but we had no idea how hard it would be to implement this idea in practice. We relied on the ^{31}P in bone mineral and soft tissues to behave similarly to ^{13}C in buckyballs or ^{29}Si in crystalline silicon, despite its more complex environment. We needed the animal imaging system to be capable of handling the pulse sequences and fast gradient transients we required. We encountered many unexpected challenges along the way that could have greatly hindered the applicability of our technique. Fortunately, for every problem we encountered, we eventually found a solution and learned much more in the process. Ultimately, this thesis work has successfully built a strong foundation for many exciting future experiments involving MRI of solids.

Appendix A

Measurements of Gradient Ramp Times

In the table below, I provide the measured gradient ramp times for different gradient amplitudes in the three gradient directions using the “ramp off” gradient ramp mode on the Bruker system. We only measured gradient amplitude values between 2.5% and 25% because we never used higher gradient amplitudes due to the constraints of fitting our sample within a single octant of our FOV.

Table A.1: Table of gradient ramp times for all 3 directions with different gradient amplitudes

Gradient Amp. (%)	G_x Ramp (μs)	G_y Ramp (μs)	G_z Ramp (μs)
2.5	49	54	54
5	40	46	39
10	39	36	32
15	40	45	30
20	51	53	37
25	61	62	47

Since the x and y gradients behaved very similarly, we used the quadratic fit of the x gradient ramp times to the gradient amplitude, G_x , for calculating both α_x and α_y .

The quadratic fits used to calculate the gradient scale factors, $\alpha_{x/y/z}$, are then given by:

$$\text{X/Y Gradient Ramp Time} = 52.5 - 2.5G_x + 0.12G_x^2 \mu\text{s}, \quad (\text{A.1})$$

$$\text{Z Gradient Ramp Time} = 61.3 - 4.5G_x + 0.16G_x^2 \mu\text{s}. \quad (\text{A.2})$$

Appendix B

Practical Applications for Radial Sampling

In this appendix, I discuss details of implementing radial sampling using methods programming on the Bruker system. In my radial sampling method program (`meri_rad.m`) L would be the parameter entered in the “Number of radial slices in theta” parameter (or ‘`slice_num`’ variable), and this number would obey equation 4.11 if we desire dense sampling of \mathbf{k} space. Since we are actually using the ‘points on the surface of the sphere’ as our gradient directions, the total number of points on the surface of the sphere (given by equation 4.7) would also give the length of my gradient array and thus the number of separate scans we would need to take to complete the image. Equations 4.6 and 4.7 are used to determine the gradient directions for dense radial sampling imaging procedures. Of course in a computer program we cannot have a non-integer number of points (or gradient directions), so M_ℓ is rounded to the nearest integer in both equations. Since the imaging method programming is in C, we make use of the ‘`cmath`’ library which has only ‘`floor()`’ and ‘`ceil()`’ options for rounding (<http://www.cplusplus.com/reference/clibrary/cmath/>). Using these functions we can

create our own rounding function given by:

$$\text{round}(num) = \text{ceil}\left(\frac{\text{floor}(2 \cdot num)}{2}\right), \quad (\text{B.1})$$

where the ordering of ‘ceil()’ and ‘floor()’ as given makes any 0.5 decimal round up (as we might prefer to be sure we are densely sampling) and switching the ordering would make it round down. The equations we actually use in our radial imaging program then become:

$$M_\ell = 2L \cdot \text{ceil}\left(\frac{\text{floor}(2 \sin(\frac{\ell\pi}{L}))}{2}\right) \quad (\text{B.2})$$

and

$$\text{Total number of Points} = 2 + \sum_{\ell=1}^{L-1} 2L \cdot \text{ceil}\left(\frac{\text{floor}(2 \sin(\frac{\ell\pi}{L}))}{2}\right). \quad (\text{B.3})$$

Appendix C

Deriving 3D Density Weighting Function

The widely accepted density compensation function of k^2 for 3D radial sampling has the problem of becoming zero at the very important $\vec{k} = 0$ point where we have our highest signal. This ‘ $\vec{k} = 0$ ’ problem is addressed in the work of Ramachandran and Lakshminarayanan [68] for radiographs’ and electron micrographs’ reconstruction and applied to MRI by Joseph [71]. Having k^2 as the density compensation is correct if we would be sampling the entirety of \vec{k} space. Since we are discretely sampling \vec{k} space over a finite range, the k^2 density compensation introduces errors, particularly for low \vec{k} values. To determine a more accurate density compensation function we follow the method laid out by Ramachandran and Lakshminarayanan [68]. For radial sampling, we are interested in the density compensation along the radial dimension, so this calculation becomes a 1D problem along the radial direction of each acquired ‘spoke’ in \vec{k} space.

We would like to find the coefficients, $q(r)$, along the radial direction in image space

such that

$$q(r) = \int_{-A/2}^{A/2} k^2 \exp(2\pi ikr) dk \quad (\text{C.1})$$

where A is the total bandwidth in \mathbf{k} space, given by $2N_{\text{points}}\delta k = 1/\delta r$, and N_{points} are the number of sparse dwell points in each pseudo-FID (equal to the image size divided by two). Evaluating the integral you get

$$q(r) = \frac{2A\pi r \cos(A\pi r) + (A^2\pi^2 r^2 - 2) \sin(A\pi r)}{4\pi^3 r^3}. \quad (\text{C.2})$$

Now we want to discretize so that $r = n\delta r$ and we use $A = 1/\delta r$, which gives

$$q(n\delta r) = \frac{2\pi n \cos(\pi n) + (\pi^2 n^2 - 2) \sin(\pi n)}{4\pi^3 (n\delta r)^3}. \quad (\text{C.3})$$

Since n is an integer, $\cos(\pi n) = (-1)^n$ and $\sin(\pi n) = 0$, this becomes,

$$q(n\delta r) = \frac{(-1)^n}{2\pi^2 n^2 (\delta r)^3}. \quad (\text{C.4})$$

For $n = 0$, this compensation factor explodes, so to find the limiting value as n approaches zero, we can go back to C.2 and take the limit as r goes to zero. This gives

$$\lim_{r \rightarrow 0} q(r) = \lim_{r \rightarrow 0} \frac{2A\pi r \cos(A\pi r) + (A^2\pi^2 r^2 - 2) \sin(A\pi r)}{4\pi^3 r^3} \quad (\text{C.5})$$

$$= \lim_{r \rightarrow 0} \frac{2A\pi r \cos(A\pi r) - 2A^2\pi^2 r \sin(A\pi r)}{12\pi^3 r^2} \quad (\text{C.6})$$

$$+ \frac{2A^2\pi^2 r \sin(A\pi r) + A^3\pi^3 r^2 \cos(A\pi r) - 2A\pi \cos(A\pi r)}{12\pi^3 r^2} \quad (\text{C.7})$$

$$= \lim_{r \rightarrow 0} \frac{A^3}{12} \quad (\text{C.8})$$

$$= \frac{1}{12(\Delta r)^3} \quad (\text{C.9})$$

where we use L'Hospital's rule in the second line.

Overall, the desired coefficients are

$$\begin{aligned}
 q(n\Delta r) &= \frac{1}{12(\Delta r)^3} & n = 0 \\
 &= \frac{-1}{2\pi^2 n^2 (\Delta r)^3} & n \text{ odd} \\
 &= \frac{1}{2\pi^2 n^2 (\Delta r)^3} & n \text{ even}
 \end{aligned}
 \tag{C.10}$$

We take the discrete inverse Fourier transform of $q(n\Delta r)$ in order to find the correct density compensation array to use for our \vec{k} -space data. Taking the discrete Fourier transform,

$$\delta r \sum_{-\infty}^{\infty} q(n\delta r) \exp(-i2\pi n\delta r)
 \tag{C.11}$$

gives the desired k^2 relation.

Appendix D

Double-Resonance Probe Design

Details

In this appendix we discuss the details of Tang et al.'s [107] double-resonance circuit design and how we theoretically estimated the component values needed for our own double-resonance probe using ^{31}P and ^1H in a 4 Tesla magnet.

D.1 Proton Trap

The simplest component in the double resonance circuit is the proton trap, made by having a capacitor (C_T) and inductor (L_T) in parallel tuned to the proton frequency (for our 4T imaging magnet, this is 170 MHz). This would then act as a very high impedance at the proton frequency (and low impedance at the lower phosphorus frequency), essentially isolating the low frequency part of the circuit from the proton frequency.

The complex impedance of the trap is given by

$$Z_{\text{Trap}} = \frac{X_{C_T}^2 R_T + i[X_{C_T} R_T^2 + (X_{L_T} + X_{C_T}) X_{L_T} X_{C_T}]}{R_T^2 + (X_{L_T} + X_{C_T})^2} \quad (\text{D.1})$$

where $X_{C_T} = -1/\omega C_T$ is the reactance of the capacitor, $X_{L_T} = \omega L_T$ is the reactance of the inductor, and R_T is the resistance of the inductor. At the resonant frequency, $\omega_0 = \sqrt{\frac{1}{L_T C_T} - \frac{R_T^2}{L_T^2}}$, this impedance becomes $\frac{L_T}{C_T R_T}$. We can then choose our values for L_T and C_T by making sure the resonance frequency is at the proton Larmor frequency, and also have $L_T \gg C_T$ since we want the impedance to be as large as possible at resonance. However, if the inductance is too high, this will decrease the low frequency circuit efficiency.

D.2 High Frequency Circuit

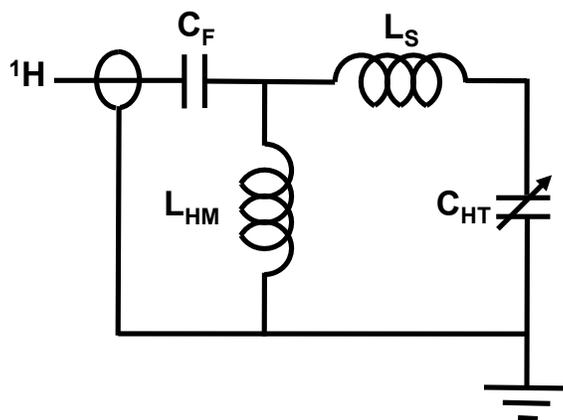


Figure D.1: Approximation of the high frequency circuit taking the proton trap as a very high impedance to ground.

The high frequency part of the double resonance circuit can be approximated as shown in Figure D.1 if we treat the proton trap as a very high impedance (or open

circuit). The complex impedance of this high frequency circuit is then given by

$$Z_{\text{High Freq}} = i \left[X_{C_F} + \frac{X_{L_{HM}}(X_{L_S} + X_{C_{HT}})}{X_{L_S} + X_{C_{HT}} + X_{L_{HM}}} \right] \quad (\text{D.2})$$

using the same reactance convention as above, and the T 's and M 's in the subscripts refer to ‘tuning’ and ‘matching’, respectively. Those that are tuning elements mostly determine the resonance frequency of the circuit (and involve variable capacitors to help with tuning), whereas matching elements mostly contribute to the overall impedance of the circuit, which we want to eventually match to 50Ω . This impedance is given not taking into account the resistances of the inductors for simplicity’s sake. However we use the full complex impedance, including resistances, in order to find approximate values for C_F , L_{HM} , and C_{HT} to satisfy both the resonant and matching condition. To do this, we estimate appropriate values for the inductor resistances and the inductance of the sample coil, L_S , which is determined by our sample dimensions.

For further isolation from the low frequency circuit, one can add an inductor to ground at the proton input (as Tang, *et al.* suggest). We also want $X_{C_F} \gg X_{L_{HM}}$ so that C_F has nearly infinite impedance and L_{HM} has nearly zero impedance for the lower phosphorus frequency.

D.3 Low Frequency Circuit

The low frequency part of the double resonance circuit can be approximated as shown in Figure D.2 if we treat C_F to have infinite impedance and L_{HM} to have nearly zero impedance at the low phosphorus frequency. The complex impedance of this low

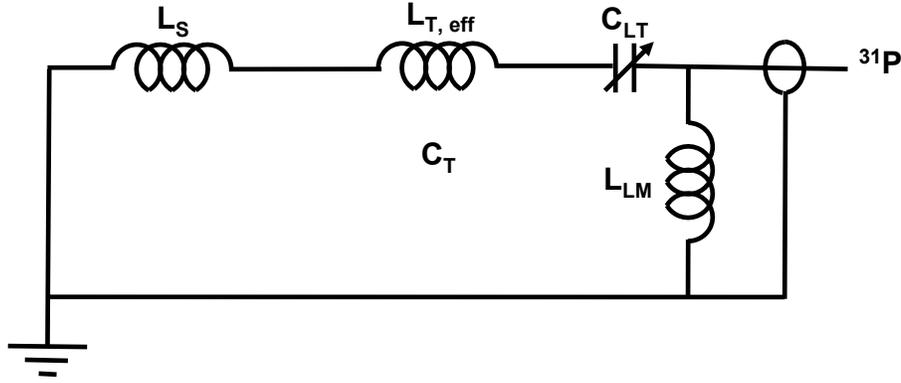


Figure D.2: For this low frequency circuit approximation we are taking C_F to be of infinite impedance and L_{HM} to be zero.

frequency circuit is then given by

$$Z_{\text{Low Freq}} = \frac{iX_{L_{LM}}(X_{L_S} + X_{L_{T, \text{eff}}} + X_{C_{LT}})}{X_{L_{LM}} + X_{L_S} + X_{L_{T, \text{eff}}} + X_{C_{LT}}} \quad (\text{D.3})$$

using the same conventions as above. Again, this impedance is given not taking into account the resistances of the inductors for simplicity's sake. However we use the full complex impedance, including resistances, in order to find approximate values for C_{LT} and L_{LM} to satisfy both the resonant and matching condition. To do this: we estimated $L_{T, \text{eff}}$ by using previously found values for the trap components, we used the known sample coil inductance, and estimated appropriate values for the inductor resistances. We also found the values needed if we wanted to tune this low frequency circuit to other nuclei (like ^{13}C and ^{29}Si), and conveniently found this could be achieved simply by increasing the capacitance of C_{LT} without changing any of the other components.

The final calculated theoretical values for the components is given in Table D.1.

Table D.1: Table of the theoretical calculations for the values of our double resonance circuit components

L_T	100nH
C_T	8.7pF
C_F	18.7pF
L_{HM}	100nH
L_S	300nH
C_{HT}	4.1pF (1-10pF variable)
L_{LM}	200nH
C_{LT}	9.6pf (P ³¹), 23pf (C ¹³), 38pf (Si ²⁹) (5-25pF variable)

Bibliography

- [1] D. E. Demco and B. Blümich. Solidstate NMR imaging methods. Part II: Line narrowing. *Concepts in Magnetic Resonance*, **12**, 269–288 (2000).
- [2] A. J. Fagan and D. J. Lurie. Continuous-Wave NMR Imaging in the Solid State. *Annual Reports on NMR Spectroscopy*, **56**, 98–140 (2005).
- [3] S. Emid and J. Creyghton. High resolution NMR imaging in solids. *Physica B*, **128**, 81–83 (1985).
- [4] B. J. Balcom, R. P. Macgregor, S. D. Beyea, D. P. Green, R. L. Armstrong, and T. W. Bremner. Single-point ramped imaging with t1 enhancement. *Journal of Magnetic Resonance - Series A*, **123**, 131–134 (1996).
- [5] J. Godward, E. Ciampi, M. Cifelli, and P. J. McDonald. Multidimensional Imaging Using Combined Stray Field and Pulsed Gradients. *Journal of Magnetic Resonance*, **155**, 92–99 (2002).
- [6] S. Hafner and H. W. Spiess. Advanced solid-state NMR spectroscopy of strongly dipolar coupled spins under fast magic angle spinning. *Concepts in Magnetic Resonance*, **10**, 99–128 (1998).
- [7] N. Lugeri, F. D. Luca, and B. Maraviglia. *Magnetic resonance microscopy*. VCH, Weinheim (1992).

- [8] J. S. Waugh, L. M. Huber, and U. Haeberlen. Approach to High-Resolution NMR in Solids. *Physical Review Letters*, **20**, 180–182 (1968).
- [9] P. Mansfield. Symmetrized pulse sequences in high resolution NMR in solids. *Journal of Physics C: Solid State Physics*, **4**, 1444 (1971).
- [10] D. P. Burum and W. K. Rhim. Analysis of multiple pulse NMR in solids. *Journal of Chemical Physics*, **71**, 944–956 (1979).
- [11] W. K. Rhim, A. Pines, and J. S. Waugh. Time-Reversal Experiments in Dipolar-Coupled Spin Systems. *Physical Review B*, **3**, 684–696 (1971).
- [12] M. Weiger, M. Stampanoni, and K. P. Preussmann. Direct Depiction of Bone Microstructure Using MRI with Zero Echo Time. *Bone* (2013).
- [13] H. Cao, A. Nazarian, J. L. Ackerman, B. Snyder, A. Rosenberg, R. Nazarian, M. Hrovat, G. Dai, D. Mintzopoulos, and Y. Wu. Quantitative ^{31}P NMR spectroscopy and ^1H MRI measurements of bone mineral and matrix density differentiate metabolic bone diseases in rat models. *Bone*, **46**(1582) (2010).
- [14] Y. Wu, G. Dai, J. Ackerman, M. Hrovat, M. Glimcher, B. Snyder, A. Nazarian, and D. Chesler. Water and Fat-Suppressed Proton Projection MRI (WASPI) of Rat Femur Bone. *Magnetic Resonance in Medicine*, **57**(554) (2007).
- [15] Y. Wu, G. Dai, M. Hrovat, J. Ackerman, T. Reese, H. Cao, K. Ecklund, and M. Glimcher. Bone Matrix Imaged In Vivo by Water and Fat-Suppressed Proton Projection MRI (WASPI) of Animal and Human Subjects. *Journal of Magnetic Resonance Imaging*, **31**(954) (2010).

- [16] I. Reichert, M. Robson, P. Gatehouse, T. He, K. Chappell, J. Holmes, S. Girgis, and G. Bydder. Magnetic resonance imaging of cortical bone with ultrashort TE pulse sequences. *Magnetic Resonance Imaging*, **23**(611) (2005).
- [17] M. Robson and G. Bydder. Clinical ultrashort echo time imaging of bone and other connective tissues. *NMR in Biomedicine*, **19**(765) (2006).
- [18] D. Tyler, M. Robson, R. M. Henkelman, I. Young, and G. Bydder. Magnetic resonance imaging with ultrashort TE (UTE) pulse sequences: technical considerations. *Journal of Magnetic Resonance Imaging*, **25**(279) (2007).
- [19] J. Du, G. Hamilton, A. Takahashi, M. Bydder, and C. Chung. Ultrashort Echo Time Spectroscopic Imaging (UTESI) of Cortical Bone. *Magnetic Resonance in Medicine*, **58**(1001) (2007).
- [20] D. Idiyatullin, C. Corum, J. Park, and M. Garwood. Fast and quiet MRI using a swept radiofrequency. *Journal of Magnetic Resonance*, **181**(342) (2006).
- [21] D. Idiyatullin, C. Corum, S. Moeller, and M. Garwood. Gapped pulses for frequency-swept MRI. *Journal of Magnetic Resonance*, **193**(267) (2008).
- [22] Y. Wu, J. L. Ackerman, D. A. Chesler, J. Li, R. M. Neer, J. Wang, and M. J. Glimcher. Evaluation of Bone Mineral Density Using Three-Dimensional Solid State Phosphorus-31 NMR Projection Imaging. *Calcified Tissue International*, **62**(512) (1998).
- [23] Y. Wu, D. A. Chesler, M. J. Glimcher, L. Garrido, J. Wang, H. J. Jiang, and J. L. Ackerman. Multinuclear solid-state three-dimensional MRI of bone and synthetic calcium phosphates. *Proceedings of the National Academy of Sciences USA*, **96**(1574) (1999).

- [24] Y. Wu, J. L. Ackerman, D. A. Chesler, J. Wang, and M. J. Glimcher. three-dimensional MRI of bone and synthetic calcium phosphates. *Proceedings 7th Annual Meeting ISMRM, Philadelphia*, **313** (1999).
- [25] M. Robson, P. Gatehouse, G. M. Bydder, and S. Neubauer. Human imaging of phosphorus in cortical and trabecular bone in vivo. *Magnetic Resonance*, **51**, 888–892 (2004).
- [26] Y. Wu, T. G. Reese, H. Cao, M. I. Hrovat, S. P. Toddes, R. A. Lemdiasov, and J. L. Ackerman. Bone mineral imaged in vivo by ^{31}P solid state MRI of human wrists. *Journal of Magnetic Resonance Imaging*, **34**, 623–633 (2011).
- [27] L. Li. ^{31}P NMR imaging of solid bone with solid echoes combined with refocused gradients. *Physics in Medicine and Biology*, **35**, 1153–1158 (1990).
- [28] L. Li. A new technique for solid NMR imaging and application to phosphorus imaging in solid bone. *Physics in Medicine and Biology*, **36**, 199–206 (1991).
- [29] R. R. Ernst, G. Bodenhausen, and A. Wokaun. *Principles of Nuclear Magnetic Resonance in one and two Dimensions*. Clarendon Press, Oxford, England (1983).
- [30] C. P. Slichter. *Principles of Magnetic Resonance*. Springer, New York, 3rd edition (1990).
- [31] E. Fukkushima and S. B. Roeder. *Experimental Pulse NMR: a Nuts and Bolts Approach*. Addison-Wesley Pub. Co., Reading, MA (1981).
- [32] P. Lauterbur. Image Formation by Induced Local Interactions: Examples Employing Nuclear Magnetic Resonance. *Nature*, **242**, 190–191 (1973).

- [33] D. Li, Y. Dong, R. Ramos, J. Murray, K. MacLean, A. Dementyev, and S. Barrett. Intrinsic Origin of Spin Echoes in Dipolar Solids Generated by Strong π pulses. *Physical Review B*, **77**(21), 214306–26 (2008).
- [34] Y. Dong, R. Ramos, D. Li, and S. Barrett. Controlling Coherence Using the Internal Structure of Hard π Pulses. *Physical Review Letters*, **100**(24), 247601–4 (2008).
- [35] R. Ramos. *Novel Pulse Sequences Exploiting the Effects of Hard π -pulses*. Ph.D. thesis, Yale University (2010).
- [36] Y. Dong. *Effects of Hard but Finite π Pulses: From Uncontrolled Coherence Flow to Extreme Line-narrowing and MRI of Solids*. Ph.D. thesis, Yale University (2009).
- [37] B. Kane. A silicon-based nuclear spin quantum computer. *Nature*, **393**, 133–137 (1998).
- [38] S. Watanabe and S. Sasaki. ^{29}Si Nuclear-Spin Decoherence Process Directly Observed by Multiple Spin-Echoes for Pure and Carrier-Less Silicon. *Japanese Journal of Applied Physics*, **42**, L1350 (2003).
- [39] M. B. Franzoni and P. R. Levstein. Manifestations of the absence of spin diffusion in multipulse NMR experiments on diluted dipolar solids. *Physics Review B*, **72**, 235410 (2005).
- [40] T. D. Ladd, D. Maryenko, Y. Yamamoto, E. Abe, and K. M. Itoh. Coherence time of decoupled nuclear spins in silicon. *Physics Review B*, **71**, 014401 (2005).
- [41] S. Plevritis and A. Macovski. Generating Unexpected Spin Echoes in Dipolar Solids with Pi Pulses. *Physics Review Letters*, **98**, 190401 (2007).

- [42] U. Haeberlen. *High Resolution NMR in Solids: Selective Averaging*. Academic Press, New York, volume 1 edition (1976).
- [43] M. Mehring. *Principles of High-Resolution NMR in Solids*. Springer-Verlag, Berlin, 2nd edition (1983).
- [44] K. Takegoshi and C. A. McDowell. A “magic echo” pulse sequence for high-resolution NMR spectra of abundant spins in solids. *Chemical Physics Letters*, **116**, 100 (1985).
- [45] I. Solomon. Rotary spin echoes. *Physics Review Letters*, **2**, 301 (1959).
- [46] J. B. Miller, D. G. Cory, A. N. Garroway, P. Mansfield, D. J. Greenslade, K. J. Packer, R. D. Farrant, J. H. Strange, J. Frahm, and E. L. Hahn. Line-Narrowing Approaches to Solid State NMR Imaging: Pulsed Gradients and Second Averaging. *Philosophical Transactions of the Royal Society of London A*, **333**, 413–426 (1990).
- [47] D. G. Cory, J. B. Miller, R. Turner, and A. N. Garroway. Multiple-pulse methods of ^1H NMR imaging of solids: second-averaging. *Molecular Physics*, **70**, 331–345 (1990).
- [48] S. Matsui. Solid-state NMR imaging by magic sandwich echoes. *Chemical physics letters*, **179**, 187–190 (1991).
- [49] Y. Wu, J. L. Ackerman, H.-M. Kim, C. Rey, A. Barroug, and M. J. Glimcher. Nuclear Magnetic Resonance Spin-Spin Relaxation of the Crystals of Bone, Dental Enamel, and Synthetic Hydroxyapatites. *Journal of Bone and Mineral Research*, **17**, 472–480 (2002).

- [50] M. Frey, M. Michaud, J. VanHouten, K. Insogna, J. Madri, and S. Barrett. Phosphorus-31 MRI of hard and soft solids using quadratic echo line-narrowing. *Proceedings of the National Academy of Sciences, USA*, **109**, 5190–5195 (2012).
- [51] M. A. Bernstein, K. F. King, and X. J. Zhou. *Handbook of MRI Pulse Sequences*. Elsevier Academic Press (2004).
- [52] R. Freeman. *Spin Choreography: Basic Steps in High Resolution NMR*. Oxford University Press, New York (1998).
- [53] E. Donnelly. Methods for Assessing Bone Quality: A Review. *Clinical Orthopaedics and Related Research*, pages 10.1007/s11999-010-1702-0 (2010).
- [54] R. O. Ritchie, M. J. Buehler, and P. Hansma. Plasticity and toughness in bone. *Physics Today*, **62**, 41–47 (2009).
- [55] F. Wehrli. Structural and functional assessment of trabecular and cortical bone by micro magnetic resonance imaging. *Journal of Magnetic Resonance Imaging*, **25**, 390–405 (2007).
- [56] S. Anumula, J. Magland, S. L. Wehrli, H. Zhang, H. Ong, H. K. Song, and F. W. Wehrli. Measurement of phosphorus content in normal and osteomalacic rabbit bone by solidstate 3D radial imaging. *Magnetic Resonance in Medicine*, **56**, 246–252 (2006).
- [57] S. Anumula, J. Magland, S. L. Wehrli, H. Ong, H. K. Song, and F. W. Wehrli. Multi-modality study of the compositional and mechanical implications of hypomineralization in a rabbit model of osteomalacia. *Bone*, **42**, 405–413 (2008).
- [58] M. Brandi. Microarchitecture, the key to bone quality. *Rheumatology*, **48**, iv3–iv8 (2009).

- [59] S. GourionArsiquaud, J. C. Burket, L. M. Havill, E. Dicarolo, S. B. Doty, R. Mendelsohn, M. C. H. van der Meulen, and A. L. Boskey. Spatial variation in osteonal bone properties relative to tissue and animal age. *Journal of Bone and Mineral Research*, **24**, 1271–1281 (2009).
- [60] E. Donnelly, A. L. Boskey, S. P. Baker, and M. C. H. van der Meulen. Effects of tissue age on bone tissue material composition and nanomechanical properties in the rat cortex. *Journal of Biomedical Material Research Part A*, pages 1048–1056 (2009).
- [61] R. Frietas. *Nanomedicine, Vol. 1: Basic Capabilities*. pages 256, Table 8.17 (1999).
- [62] R. W. Sterner and J. J. Elser. *Ecological Stoichiometry: The Biology of Elements from Molecules to the Biosphere*. Princeton University Press, Princeton, NJ (2002).
- [63] R. Buchli, C. O. Duc, E. Martin, and P. Boisiger. Assessment of absolute metabolite concentrations in human tissue by ^{31}P MRS in vivo. Part 1: Cerebrum, cerebellum, cerebral gray and white matter. *Magnetic Resonance in Medicine*, **32**, 447–452 (1994).
- [64] J. A. Stanley and J. W. Pettegrew. Postprocessing method to segregate and quantify the broad components underlying the phosphodiester spectral region of in vivo ^{31}P brain spectra. *Magnetic Resonance in Medicine*, **35**, 390–396 (2001).
- [65] J. I. Jackson, C. H. Meyer, and D. G. Nishimura. Selection of a convolution function for Fourier inversion using gridding. *IEEE Transactions on Medical Imaging*, **10**(3), 473–478 (1991).

- [66] J. D. O’Sullivan. A fast sinc function gridding algorithm for Fourier inversion in computer tomography. *IEEE Transactions on Medical Imaging*, **MI-4**(4), 200–207 (1985).
- [67] P. J. Beatty and D. G. Nishimura. Rapid gridding reconstruction with a minimal oversampling ratio. *IEEE Transactions on Medical Imaging*, **24**(6), 799–808 (2005).
- [68] G. N. Ramachandran and A. V. Lakshminarayanan. Three-dimensional reconstruction from radiographs and electron micrographs: application of convolutions instead of Fourier transforms. *Proc. Nat. Acad. Sci. USA*, **68**(9), 2236–2240 (1971).
- [69] H. Schomberg and J. Timmer. The gridding method for image reconstruction by Fourier transformation. *IEEE Transactions on Medical Imaging*, **14**(3), 596–607 (1995).
- [70] M. L. Lauzon and B. K. Rutt. Effects of polar sampling in kspace. *Magnetic Resonance in Medicine*, **36**, 940–949 (1996).
- [71] P. M. Joseph. Sampling errors in projection reconstruction MRI. *Magnetic Resonance in Medicine*, **40**, 460–466 (1998).
- [72] J. Pipe and P. Menon. Sampling density compensation in MRI: rationale and an iterative numerical solution. *Magnetic Resonance in Medicine*, **41**, 179–186 (1999).
- [73] V. Rasche, R. Proksa, R. Sinkus, P. Börnert, and H. Eggers. Resampling of data between arbitrary grids using convolution interpolation. *IEEE Transactions on Medical Imaging*, **18**(5), 385–92 (1999).

- [74] M. Bydder, A. Samsonov, and J. Du. *Magnetic Resonance Imaging*, **25**, 695–702 (2007).
- [75] C. V. Uijen and J. H. D. Boef. Drivenequilibrium radiofrequency pulses in NMR imaging. *Magnetic Resonance in Medicine*, **1**, 502–507 (1984).
- [76] M. Lustig, D. Donoho, and J. M. Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine*, **58**(6), 1182–95 (2007).
- [77] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly. Compressed sensing MRI. *Signal Proceedings IEEE*, **25**, 72–82 (2008).
- [78] E. Candes and M. Wakin. An introduction to compressive sampling. *IEEE Signal Proc Mag*, **25**, 21–30 (2008).
- [79] E. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans Inf Theory*, **52**, 489–509 (2006).
- [80] E. J. Candès, M. B. Wakin, and S. P. Boyd. Enhancing sparsity by reweighted l_1 minimization. *J. Fourier Anal. Appl.*, **14**(5-6), 877–905 (2008).
- [81] S. Hu, M. Lustig, A. Balakrishnan, P. E. Z. Larson, R. Bok, J. Kurhanewicz, S. J. Nelson, A. Goga, J. M. Pauly, and D. B. Vigneron. 3D compressed sensing for highly accelerated hyperpolarized (^{13}C) MRSI with in vivo applications to transgenic mouse models of cancer. *Magnetic Resonance in Medicine*, **63**(2), 312–21 (2010).

- [82] M. Mobli, M. Maciejewski, A. S. AD, A. Stern, and J. Hoch. Sparse sampling methods in multidimensional NMR. *Phys Chem Chem Phys*, **14**, 10835–10843 (2012).
- [83] A. Stern, D. Donoho, and J. Hoch. NMR data processing using iterative thresholding and minimum $l(1)$ -norm reconstruction. *J Magn Reson*, **188**, 295–300 (2007).
- [84] J. Hoch and A. Stern. *Encyclopedia of NMR*. John Wiley and Sons, New York (1996).
- [85] V. Elser. Phase retrieval by iterated projections. *Journal of the Optical Society of America A, Optics, image science, and vision*, **20**, 40–55 (2003).
- [86] V. Elser, P. Thibault, and I. Rankenburg. Searching with iterated maps. *Proceedings of the National Academy of Sciences, USA*, **104**, 418–423 (2007).
- [87] R. Gerchberg and W. Saxton. A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik (Stuttg)*, **35**, 237–246 (1972).
- [88] J. Fienup. Phase retrieval algorithms: a comparison. *Appl Opt*, **21**, 2758–2769 (1982).
- [89] H. Bauschke, P. Combettes, and D. Luke. Phase retrieval, error reduction algorithm, and Fienup variants: a view from convex optimization. *J Opt Soc Am A Opt Image Sci Vis*, **19**, 1334–1345 (2002).
- [90] S. Gravel and V. Elser. Divid and concur: a general approach to constraint satisfaction. *Physical Review E*, **78**, 036706 (2008).

- [91] J. Yedidia. Message-Passing Algorithms for Inference and Optimization: Belief Propagation and Divide and Concur. *Journal of Statistical Physics*, **145**, 860890 (2011).
- [92] E. Haacke, E. Lindskog, and W. Linz. A Fast, Iterative, Partial-Fourier Technique Capable of Local Phase Recovery. *Journal of Magnetic Resonance*, **92**, 126–145 (1991).
- [93] S. Plevritis and A. Macovski. MRS imaging using anatomically based K-space sampling and extrapolation. *Magnetic Resonance in Medicine*, **34**, 686–693 (1995).
- [94] S. Plevritis and A. Macovski. Spectral Extrapolation of Spatially Bounded Images. *IEEE Transactions on Medical Imaging*, **14**, 487–497 (1995).
- [95] A. K. Jain and S. Ranganath. Extrapolation Algorithms for Discrete Signals with Application in Spectral Estimation. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, **ASSP-29**, 830–845 (1981).
- [96] Y. Matsuki, M. T. Eddy, and J. Herzfeld. Spectroscopy by Integration of Frequency and Time Domain Information (SIFT) for Fast Acquisition of High Resolution Spectra. *Journal of the American Chemical Society*, **131**, 4648–4656 (2009).
- [97] K. Kazimierczuk, A. Zawadzka, and W. K. zmiński. Lineshapes and artifacts in Multidimensional Fourier Transform of arbitrary sampled NMR data sets. *Journal of Magnetic Resonance*, **188**, 344–356 (2007).
- [98] Y. Matsuki, M. T. Eddy, R. G. Griffin, and J. Herzfeld. Rapid Three-Dimensional MAS NMR Spectroscopy at Critical Sensitivity. *Angewandte Chemie International Edition*, **49**, 9215–9218 (2010).

- [99] Y. Matsuki, T. Konuma, T. Fujiwara, and K. Sugase. Boosting Protein Dynamics Studies Using Quantitative Nonuniform Sampling NMR Spectroscopy. *The Journal of Physical Chemistry B*, **115**, 13740–13745 (2011).
- [100] D. States, R. Haberkorn, and D. Ruben. *Journal of Magnetic Resonance*, **48**, 286 (1982).
- [101] Y. Shrot and L. Frydman. Compressed sensing and the reconstruction of ultrafast 2D NMR data: Principles and biomolecular applications. *Journal of Magnetic Resonance*, **209**, 352–358 (2011).
- [102] L. Frydman, T. Scherf, and A. Lupulescu. The acquisition of multi-dimensional NMR spectra within a single scan. *Proceedings of the National Academy of Sciences, USA*, **99** (2002).
- [103] F. D. Doty. Probe Design and Construction. *Encyclopedia of Magnetic Resonance* (2007).
- [104] F. D. Doty, R. R. Inners, and P. D. Ellis. A Multinuclear Double-Tuned Probe for Applications with Solids or Liquids Utilizing Lumped Tuning Elements. *Journal of Magnetic Resonance*, **43**, 399–416 (1981).
- [105] F. D. Doty, T. J. Connick, X. Z. Ni, and M. N. Clingan. Noise in High-Power, High-Frequency Double-Tuned Probes. *Journal of Magnetic Resonance*, **77**, 536–549 (1988).
- [106] E. Najim and J.-P. Grivet. Efficiency Estimation for Single-Coil, Separate-Input, Double-Tuned NMR Probes. *Journal of Magnetic Resonance*, **93**, 27–33 (1990).

- [107] P. Tang, W.-J. Chien, and G. S. Harbison. Double-resonance circuit for nuclear magnetic resonance spectroscopy. *Solid State Nuclear Magnetic Resonance*, **2**, 343–348 (1993).
- [108] R. Kerkoud, P. Auban-Senzier, J. Godard, D. Jerome, J.-M. Lambert, and P. Bernier. Effects of pressure on the thermodynamic properties of C₆₀ Studied by ¹³C NMR. *Advanced Materials*, **6**, 782–786 (1994).
- [109] A. Asano, S. Hori, M. Kitamura, C. Nakazawa, and T. Kurotsu. Influence of magic angle spinning on T_1^H of SBR studied by solids state ¹H NMR. *Polymer Journal*, **44**, 706–712 (2012).
- [110] Y.-Q. Song, S. Ryu, and P. N. Sen. Determining multiple length scales in rocks. *Nature*, **406**, 178–181 (2000).
- [111] V. Cnudde, B. Masschaele, M. Dierick, J. Vlassenbroeck, L. VanHoorebeke, and P. Jacobs. Recent progress in X-ray CT as a geosciences tool. *Applied Geochemistry*, **21**, 826–832 (2006).
- [112] H. Dong and M. J. Blunt. Pore-network extraction from micro-computerized-tomography images. *Physical Review E*, **80**, 1–5 (2009).
- [113] C. H. Arns. A comparison of pore size distributions derived by NMR and X-ray-CT techniques. *Physica A*, **339**, 159–165 (2004).
- [114] L. Knauss, A. Orozco, and S. Woods. Advances in magnetic-based current imaging for high resistance defects and sub-micron resolution. In *Proceedings of the 11th International Symposium on the Physical and Failure Analysis of Integrated Circuits*, pages 267–270 (July).