LBL-36616
UC-408

Lawrence Berkeley Laboratory
UNIVERSITY OF CALIFORNIA

Three Dimensional Nuclear Magnetic Resonance Spectroscopic Imaging of Sodium Ions Using Stochastic Excitation and Oscillating Gradients

B.deB. Frederick
(Ph.D. Thesis)

December 1994

Donner Laboratory
Biology & Medicine Division

REFERENCE COPY
Does Not
Circulate
Bldg. 50 Library.
LBL-36616
Copy 1

Prepared for the U.S. Department of Energy under Contract Number DE-AC03-76SF00098

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor the Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or the Regents of the University of California.

LBL-36616

Three Dimensional Nuclear Magnetic Resonance Spectroscopic Imaging of Sodium Ions Using Stochastic Excitation and Oscillating Gradients

Blaise deBonneval Frederick
Ph.D. Thesis

Graduate Group in Biophysics
University of California, Berkeley

and

Life Sciences Division
Lawrence Berkeley Laboratory
University of California
Berkeley, CA 94720

December 1994

This work was supported in part by the National Institutes of Health under Grant No. HL07367 and by the Director, Office of Energy Research, Office of Health and Environmental Research, Medical Applications and Biophysical Research Division, of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

Three Dimensional Nuclear Magnetic Resonance Spectroscopic Imaging of Sodium Ions Using Stochastic Excitation and Oscillating Gradients

Copyright © 1994

by

Baise deBonneval Frederick

The U.S. Department of Energy has the right to use this document
for any purpose whatsoever including the right to reproduce
all or any part thereof

recycled paper

Abstract

Three Dimensional Nuclear Magnetic Resonance
Spectroscopic Imaging of Sodium Ions
Using Stochastic Excitation and Oscillating Gradients

by

Blaise deBonneval Frederick
Doctor of Philosophy in Biophysics

University of California at Berkeley

Professor Thomas F. Budinger, Chair

Nuclear magnetic resonance (NMR) spectroscopic imaging of ^{23}Na holds promise as a non-invasive method of mapping Na^+ distributions, and for differentiating pools of Na^+ ions in biological tissues. However, due to NMR relaxation properties of ^{23}Na *in vivo*, a large fraction of Na^+ is not visible with conventional NMR imaging methods. An alternate imaging method, based on stochastic excitation and oscillating gradients, has been developed which is well adapted to measuring nuclei with short T_2.

Contemporary NMR imaging techniques have dead times of up to several hundred microseconds between excitation and sampling, comparable to the shortest *in vivo* ^{23}Na T_2 values, causing significant signal loss. An imaging strategy based on stochastic excitation has been developed which greatly reduces experiment dead time by reducing peak radiofrequency (RF) excitation power and using a novel RF circuit to speed probe recovery. Continuously oscillating gradients are used to eliminate transient eddy currents. Stochastic ^{1}H and ^{23}Na spectroscopic imaging experiments have been performed on a small animal system with dead times as low as 25μs, permitting spectroscopic imaging with 100% visibility *in vivo*. As an additional benefit, the encoding time for a 32x32x32 spectroscopic image is under 30 seconds.

The development and analysis of stochastic NMR imaging has been hampered by limitations of the existing phase demodulation reconstruction technique. Three dimensional imaging was impractical due to reconstruction time, and design and analysis of proposed experiments was limited by the mathematical intractability of the reconstruction method. A new reconstruction method for stochastic NMR based on Fourier interpolation has been formulated combining the advantage of a several hundredfold reduction in reconstruction time with a straightforward mathematical form. This permits the determination of important image parameters, such as point spread function and noise propagation.

The reduction in image reconstruction time from over 1700 hours to under 3 hours for a 32^3 image has made stochastic spectroscopic imaging practical. In addition, the active probe Q-spoiling circuits developed for this experiment can be added to any RF probe and used with any imaging sequence to significantly reduce dead time between RF excitation and sampling for short T_2 experiments.

Thomas F. Budinger

Professor Thomas F. Budinger
Dissertation Committee Chair

For Clancy, who taught me so much.

By teaching me how to play the guitar, he showed me through music how important it is to create, and to strive to always improve.

By teaching me how to fix telephones, he helped me along the path to being a scientist by showing me how to improve things by breaking them first.

By being my friend, he showed me how to value myself, and how important it is to remember.

Contents

List of Figures vii

List of Tables x

Symbol Table xi

1 Introduction 1
- 1.1 Interest in Ions . . . 1
- 1.2 Medical Applications of Sodium Imaging . . . 2
- 1.3 NMR Imaging of Ions . . . 3
- 1.4 Goals . . . 5

2 Stochastic NMR 7
- 2.1 Introduction . . . 7
- 2.2 Background . . . 7
- 2.3 Spectroscopy . . . 10
- 2.4 Advantages of Stochastic Excitation . . . 13
- 2.5 Extension to Imaging . . . 14

3 Materials and Methods 16
- 3.1 Introduction . . . 16
- 3.2 Experiment Parameters . . . 17
  - 3.2.1 Pulse Sequence . . . 17
  - 3.2.2 RF Noise Sequence . . . 18
  - 3.2.3 Gradient Trajectories . . . 21
    - 3.2.3.1 Truncated Square Wave (TSW) Gradients . . . 22
    - 3.2.3.2 Amplitude Modulated Sinusoidal (AMS) Gradients . 22
    - 3.2.3.3 AMS Modulation Function . . . 24
    - 3.2.3.4 Gradient Oscillation Frequency . . . 26
    - 3.2.3.5 Extension to Three Dimensions . . . 26
  - 3.2.4 Sampling Frequency . . . 28
- 3.3 Apparatus . . . 30

3.3.1 Spectrometer 30
3.3.2 Probes 30
3.3.3 Active Probe Q-spoiling 31
3.4 Calibrations 37
3.4.1 K Space Trajectory 37
3.4.2 Receiver Filter Delay 38

4 Reconstruction Algorithm **41**
4.1 Introduction 41
4.2 Mathematical Description of the Experiment 42
4.3 Phase Demodulation Reconstruction 43
4.4 Fourier Gridding Reconstruction 44
4.5 Point Spread Function Analysis 50
4.5.1 One Dimensional Periodic Oscillating Gradients 51
4.5.2 Rotating Oscillating Gradients 52
4.5.3 Incommensurate Frequency Oscillating Gradients 55
4.5.4 Point Spread Functions For Spectroscopic Imaging 55
4.6 Noise Analysis 60
4.6.1 Single Lag Noise Analysis 60
4.6.2 One Dimensional Oscillating Gradient 62
4.6.3 Rotating Oscillating Gradients 63
4.6.4 Incommensurate Frequency Oscillating Gradients 64
4.6.5 Multi-lag Noise Characteristics 64
4.7 Implementation 67
4.8 Discussion 68

5 Simulations **70**
5.1 Introduction 70
5.2 Simulation Program 71
5.3 Experiment Bandwidth 71
5.4 Point Spread Function 74
5.5 Systematic Noise Reduction 74
5.5.1 Repeating versus Nonrepeating k Trajectories 79
5.5.2 Analytic versus Empirical Sampling Density Correction 82
5.6 Monte Carlo Simulations of Noise Variance 83

6 Experimental Studies **86**
6.1 Introduction 86
6.2 Spectroscopic Imaging 86
6.2.1 ^{1}H Phantom Construction 86
6.2.2 Images 87
6.3 T_2 Sensitivity 93

6.3.1 ^{23}Na T_2 Phantom Construction 93
6.3.2 Images 96

7 Summary **102**

Bibliography **107**

A Reconstruction Pseudocode **112**

List of Figures

2.1 Schematic representation of the continuous wave (CW) NMR experiment . . . 10
2.2 Schematic representation of the Fourier Transform (FT-NMR) experiment . . . 12
2.3 Schematic representation of the stochastic NMR experiment . . . 12

3.1 Stochastic imaging experiment pulse sequence . . . 17
3.2 Stochastic imaging experiment pulse sequence (detail) . . . 19
3.3 One dimensional gradient waveforms . . . 23
3.4 Comparison of normalized sampling density of modulated and unmodulated sinusoidal gradients . . . 25
3.5 Sine wave modulated by periodic AMS modulation function . . . 25
3.6 Modulation functions for x, y and z. . . . 28
3.7 Dual tuned ($^{1}H/^{23}Na$) probe . . . 31
3.8 Circuit diagram, PIN diode Q-spoiler. . . . 32
3.9 Physical diagram, PIN diode Q-spoiler. . . . 33
3.10 RF System Diagram . . . 35
3.11 Demonstration of Q-spoiler . . . 36
3.12 k trajectory calibration pulse sequence . . . 38
3.13 k trajectory maps . . . 39

4.1 Composition and reconstruction of the stochastic signal in the absence of gradients . . . 45
4.2 Diagram of the gridding reconstruction procedure . . . 47
4.3 k-space trajectories arising from different RF pulses . . . 48
4.4 k_{max} as a function of correlation lag q . . . 52
4.5 Comparison of projection angles for rotating oscillating gradient trajectories . . . 53
4.6 Single lag point spread function for the rotating oscillating gradient experiment . . . 54

4.7 Single lag point spread function for the incommensurate frequency sinusoidal oscillating gradient experiment 56
4.8 Single lag point spread function for the incommensurate frequency three component truncated square wave oscillating gradient experiment 56
4.9 Overall point spread function for the rotating oscillating gradient experiment 58
4.10 Overall point spread function for the incommensurate frequency sinusoidal oscillating gradient experiment 59
4.11 Overall point spread function for the three component incommensurate frequency oscillating gradient experiment 59
4.12 Single lag root noise power spectral density for the rotating oscillating gradient experiment 64
4.13 Single lag root noise power spectral density for the incommensurate frequency oscillating gradient experiment 65
4.14 Overall root noise power spectral density for the rotating oscillating gradient experiment 66
4.15 Overall root noise power spectral density for the incommensurate frequency sinusoidal oscillating gradient experiment 66

5.1 Power spectrum, one dimensional sinusoidal gradient experiment 72
5.2 Power spectrum, amplitude modulated one dimensional sinusoidal gradient experiment 73
5.3 Power spectrum, one dimensional three component truncated square wave gradient experiment 73
5.4 Single lag point spread function for the rotating oscillating gradient experiment 75
5.5 Overall point spread function for the rotating oscillating gradient experiment 75
5.6 Single lag point spread function for the incommensurate frequency sinusoidally oscillating gradient experiment 76
5.7 Overall point spread function for the incommensurate frequency sinusoidally oscillating gradient experiment 76
5.8 Single lag point spread function for the incommensurate frequency three component oscillating gradient experiment 77
5.9 Overall point spread function for the incommensurate frequency three component TSW oscillating gradient experiment 77
5.10 Single lag point spread function for the amplitude modulated incommensurate frequency sinusoidal oscillating gradient experiment 78
5.11 Overall point spread function for the amplitude modulated incommensurate frequency oscillating gradient experiment 78
5.12 Single lag noise power spectral density for the sinusoidal oscillating gradient experiment 84

5.13 Single lag noise power spectral density for the three component truncated square wave oscillating gradient experiment 85
5.14 Single lag noise power spectral density for the amplitude modulated sinusoidal oscillating gradient experiment 85

6.1 ^{1}H chemical shift phantom 87
6.2 Spectroscopic image of the ^{1}H tube phantom using sinusoidal gradients 89
6.3 Spectroscopic image of the ^{1}H tube phantom using sinusoidal incommensurate frequency oscillating gradients 90
6.4 Spectroscopic image of the ^{1}H tube phantom using three component TSW gradients 91
6.5 Spectroscopic image of the ^{1}H tube phantom using three component TSW gradients 92
6.6 ^{23}Na T_2 phantom 93
6.7 ^{23}Na T_2^* relaxation rates of 100mM NaCl in glycerol/water solutions in a 4cm spherical sample at 2.35T 95
6.8 Images of the ^{23}Na T_2 phantom using incommensurate frequency sinusoidally oscillating gradients at different FID time points 97
6.9 Spatially resolved FID image of the ^{23}Na T_2 phantom from an incommensurate frequency sinusoidally oscillating gradient experiment . . . 98
6.10 Images of the ^{23}Na T_2 phantom using incommensurate frequency three component truncated square wave gradients at different FID time points 99
6.11 Spatially resolved FID image of the ^{23}Na T_2 phantom from an incommensurate frequency three component TSW oscillating gradient experiment 100

List of Tables

1.1 Comparison of NMR properties of ^{1}H and ^{23}Na 2

6.1 ^{23}Na T_1 and T_2 relaxation times 94
6.2 ^{23}Na T_2^* relaxation times (measured by stochastic NMR spectroscopy) 95
6.3 ^{23}Na T_2^* relaxation times (measured by stochastic NMR imaging) . . 101

Symbol Table

Symbol	Definition
T_1	Spin-lattice relaxation time constant
T_2	Spin-spin relaxation time constant
T_2^*	Effective spin-spin relaxation time constant
T_R	Cycle time of the stochastic experiment
$\hat{\rho}$	Estimate of the spin density
ρ	Spin density as a function of $\mathbf{x}$ and σ
ρ_k	Fourier transform of ρ with respect to $\mathbf{x}$
ρ_q	Fourier transform of ρ with respect to σ
ρ_K	Fourier transform of ρ with respect to $\mathbf{x}$ and σ
$\hat{\rho}_K$	Expectation of ρ_K
$\hat{\rho}_q$	Expectation of ρ_q
m_n	Transverse magnetization at time n
$\mathbf{x}$	Spatial coordinate
$\mathbf{k}_{n,q}$	Position in k-space at time n of magnetization created q time intervals in the past
$\mathbf{G}_p$	Magnetic field gradient at time p
y_n	Received RF signal at time n
s_n	RF excitation at time n
σ	Chemical shift coordinate
η	Correlation lag coordinate
α	RMS flip angle of the RF excitation
α_{Ernst}	RMS flip angle for maximum signal
w	Phase demodulation kernel weighting function
W	k-space sampling density
γ	Gyromagnetic ratio
ς^2	Measurement noise variance
B_0	Applied magnetic field
Ω	Experiment bandwidth

Acknowledgements

I would like to give my warm thanks to my thesis advisor, Dr. Thomas Budinger not only for all of his support and guidance with my thesis, but for helping me keep my eye on the big picture.

I'd like to thank Sam Wong and Mark Roos for setting me on the road to stochastic NMR, for sharing their extensive knowledge of stochastic (and other) NMR, and for their advice and direction in my research.

Josh Lack has provided important consultation on the vagaries of oscillating gradient NMR experiments and the reconstruction thereof (especially the importance of time gridding), and has been a steadfast comrade in the struggle against various computers and other infernal machines.

John Keltner has helped me with a great deal of spirited and enthusiastic discussion about the value of stochastic excitation, the particulars of sodium NMR, and has always been generous with his advice and affection.

The entire Research Medicine group at LBL has been very supportive both technically and personally, especially the animals.

My family has been very helpful throughout my graduate school career and before, giving me both the tools and the encouragement I have needed to pursue my education, and by always being there for me.

Most of all I want to thank Leslie Osborne. She has been a constant source of support and encouragement in all aspects of my life; convincing me to take up guitar again, staying up late helping me prepare talks, sharing my interests and giving me new ones, and making it possible for me to finish my thesis through her boundless patience and heroic assistance. She has had confidence in me when I have not, and has kept me going through the hardest times. But more important, her love and companionship have made the good times wonderful.

Chapter 1

Introduction

1.1 Interest in Ions

The principal medical use of nuclear magnetic resonance (NMR) imaging is to glean information about tissue structure by using contrast between various types of tissue in ^{1}H images. This contrast arises from differences in ^{1}H density, T_1 and T_2 relaxation times, and differences in flow and diffusion rates of bodily fluids. ^{1}H is the nucleus of choice for most NMR imaging because its high concentration, 100% natural isotopic abundance, and high NMR sensitivity give a significantly higher signal to noise ratio than any nucleus found in the body. There are, however, many other NMR sensitive nuclei in the body, some of which carry functional information about the activity and health of cells. NMR imaging and spectroscopy experiments that observe these nuclei offer a non-invasive technique for monitoring metabolism and pathology locally and dynamically in a clinically useful way. ^{23}Na is an NMR active nucleus with 100% natural isotopic abundance and a high gyromagnetic ratio; this and its relatively high concentration in the body give it the second highest NMR signal to noise ratio in the body, making it a good candidate for NMR imaging *in vivo*. The NMR properties of ^{1}H and ^{23}Na are compared in Table 1.1.

The relative concentrations of certain ions inside and outside of the cell are critical to many of the processes of cell function. Ion balance affects the osmotic pressure across the cell membrane, which determines cell volume and constituent

	^{1}H	^{23}Na
Concentration intracellular extracellular	111M	39mM 11mM (14mM x 80% v.f.) 28mM (142mM x 20% v.f.)
Natural Abundance	99.98%	100%
Gyromagnetic Ratio	42.57 MHz/T	11.26 MHz/T
Spin	1/2	3/2
Sensitivity (relative to ^{1}H) Raw *In vivo*	 1 1	 0.35 0.00013
T_2 Relaxation times	30ms-2000ms	300μs-50ms

Table 1.1: Comparison of NMR properties of ^{1}H and ^{23}Na

concentrations, and the balance between the osmotic pressure and electrostatic forces establishes the membrane potential. Two of the most important ions in animal cells are Na^+ and K^+ . Fully one third of the energy used in animal cells is devoted to the ATPase which pumps Na^+ out of the cell while pumping K^+ in (70% in electrically active cells) to maintain Na^+ and K^+ gradients across the cell membrane[1]. These Na^+ and K^+ gradients are responsible for driving many of the transport systems which bring substances in and out of the cell. Any stress on the cell which inhibits metabolism will cause this gradient to change, and a method of determining the relative and absolute concentrations of these ions inside and outside of the cell has great potential for indicating pathology in tissue. Of the two nuclei, ^{23}Na is more practical to study clinically with NMR, as it has a much higher sensitivity *in vivo*.

1.2 Medical Applications of Sodium Imaging

The potential of ^{23}Na imaging in medicine has already been demonstrated in the study of the brain, heart, and kidney. Clinical ^{23}Na images have shown significant promise in detecting brain lesions which are essentially invisible in ^{1}H images by seeing dramatic changes in the amount of extracellular Na^+ in damaged tissue[2]. Spectroscopic studies which monitored intracellular/extracellular Na^+ concentration ratios in perfused hearts dynamically during and after ischemia[3] and statically

postmortem[4] have shown that these concentration ratios can be used to predict whether or not reperfused tissues will recover from ischemia and regain full contractile function. Sodium imaging has also been used to track the 3D distribution of Na^+ in the kidney *in vivo*[5], which allows the monitoring of kidney Na^+ under different physiological stresses.

The clinical utility of ^{23}Na imaging has been limited by relatively long imaging times which are due to low signal to noise ratios and the low (and sometimes indeterminate) visibility parameter - the ^{23}Na signal in many physiological environments relaxes away so quickly it is "invisible"or greatly attenuated in conventional images. This makes these images somewhat qualitative; anomalous features can be detected but not characterized. While the technique is still useful without quantitative measurement, there is another area where quantitative Na^+ imaging could provide information impossible to obtain by any other means. There are indications that brain Na^+ concentration ratios are disturbed in some types of mental illness, such as bipolar illness[6]. There has been little experimental exploration of this area, however, because there are no non-invasive techniques for measuring these concentrations *in vivo*. Therefore this hypothesis has not been unequivocally tested, and is currently of no use clinically. If this hypothesis could be confirmed or disproved, this would increase the understanding of the physical bases of mental illness, and would provide a non-invasive physical diagnostic tool for mental illness. These conditions can currently only be diagnosed behaviorally; there are no direct diagnostic methods. Therefore there is a strong incentive to develop a quantitative method of measuring ^{23}Na concentration.

1.3 NMR Imaging of Ions

There are two major difficulties in measuring the concentrations of these physiological electrolytes *in vivo* with NMR. The first is that the nuclear quadrupole moments of both ^{39}K and ^{23}Na give them short relaxation times when they are motionally restricted, as they often are in the body. The short relaxation times make it difficult to "see" much of the ^{23}Na and ^{39}K *in vivo*. In most conventional imag-

ing pulse sequences, many T_2 time constants have passed before any NMR signal is measured; in fact, the relaxation times of these nuclei are so short that the actual ranges of relaxation rates *in vivo* are unknown; for many tissues, only an upper bound for T_2 is known. There is, however, some evidence from data on packed, Na^+ loaded yeast cells[7], that 300μs may be the lower bound for intracellular ^{23}Na T_2. For an NMR technique to be able to estimate the distribution of Na^+ in the body, the experimenter must either have 100% ^{23}Na visibility or be able to confidently calculate what the visibility factor is for a given experiment for each compartment of physiological interest. Rooney and Springer have calculated and verified that 100% visibility of ^{23}Na with a T_2 of 300μs requires a dead time before sampling of <25μs[8].

The second major difficulty is in differentiating the NMR signatures of these electrolytes. The resolution of NMR imaging techniques depends on many factors, but the major determinant of voxel size is the signal to noise ratio attainable in a reasonable amount of time. Imaging ^{23}Na in a fair sized animal *in vivo* cannot approach cellular resolution by many orders of magnitude (a typical reported voxel size is 0.36cm^3)[2]. Therefore, intracellular and extracellular Na^+ must be distinguished by NMR parameters rather than by directly resolving them spatially. Because the electrolytes exist as free ions, there is no chemical shift information in the spectra of the ions in normal physiological conditions. One way to differentiate ions in different environments (i.e. between the intracellular and extracellular compartments) is through the use of paramagnetic shift reagents such as dysprosium. Ions that can bind temporarily to these reagents experience a large magnetic field due to the unpaired electron in the dysprosium, changing their average resonance frequency. These reagents do not penetrate the cell membrane, however, so only extracellular Na^+ is affected, offering a way to distinguish the two compartments. Unfortunately, current shift reagents are cardiotoxic, limiting their use in humans. Also, they are not likely to be isotropically distributed in the extracellular spaces in intact tissue. This is especially true in the brain, due to the blood-brain barrier. Therefore, an experimental protocol which did not rely on shift reagents would be preferable.

There are two other proposed methods for separating various Na^+ compartments.

The first is the use of multiple quantum filters, which detect the coherences between non-adjacent energy levels in nuclei with more than two levels (^{23}Na and ^{39}K have spin 3/2; they have 4 energy levels when the nuclei are placed in a magnetizing field). These coherences develop when the quadrupolar field of the nucleus interacts with electric fields that do not average to zero in a time short compared to the nucleus' NMR precession period. This happens when nuclei are motionally restricted, as they are in the intracellular medium. However, these signals are as much as ten times smaller than the single quantum signals, and require significant signal averaging to discern, which may limit the clinical utility of this technique.

The last method for distinguishing ^{23}Na in different environments is relaxometry - measuring the distribution of relaxation times and estimating the amount of ^{23}Na with each T_2. ^{23}Na in the intracellular medium in general has much shorter relaxation times than the extracellular environment because of longer correlation times for the Na^+ ions and larger electric field gradients, although exactly what the distributions of T_2 are is not known[9]. An imaging method which preserves the free induction decay signal (FID) at every voxel allows the experimenter to estimate the populations with various relaxation times within the voxel, and therefore to form spatially resolved maps of these populations. These maps, coupled with information about the extracellular volume fraction can be used to estimate the intracellular and extracellular Na^+ concentration. The feasibility of "relaxographic" imaging has been demonstrated (for ^{1}H) by Labadie, *et al*[10]. This technique has the advantage of making use of the single quantum signal, and is perhaps the most straightforward way to do the experiment, as it measures all the ^{23}Na simultaneously.

1.4 Goals

Current imaging methods are not well adapted to quantitative sodium imaging. The goal of this thesis is to develop an NMR imaging method which can be used for clinical relaxographic ^{23}Na imaging. An ideal imaging method for this purpose must satisfy the following criteria:

1. The imaging method should allow the collection of spatially resolved T_2 relaxation data.

2. The imaging method must have a very short dead time before sampling (less than $25\mu s$).

3. The imaging method must be implementable on a clinical system and have a sufficient signal to noise ratio to complete an experiment in a tolerable amount of time (less than one hour).

Spectroscopic NMR imaging with stochastic excitation and oscillating gradients can fulfill all of these requirements. Stochastic NMR imaging can measure the time evolution of spins resolved in three dimensions. The data contain all of the information necessary to reconstruct spatially resolved spectra or free induction decay signals. The latter can be used to determine the T_2 relaxation components in each voxel of an object, much like a chemical shift imaging (CSI) experiment. Unlike CSI, however, the very low dead time to sampling required for ^{23}Na imaging is easily obtained. With a few modifications, a clinical spectrometer can perform the stochastic experiment. Experiments on physiological phantoms indicate that acceptable signal to noise levels are obtainable in reasonable examination times.

This thesis will demonstrate the feasibility of implementing a clinical ^{23}Na imaging experiment based on stochastic NMR. Chapter 2 will review the theoretical basis for the stochastic imaging experiment, Chapter 3 will describe the particular implementation of the experiment. Chapter 4 will develop and analyze a reconstruction algorithm for producing three dimensional spectroscopic images from stochastic data sets. Chapter 5 will present the results of simulations to validate the experiment, and Chapter 6 will present the results of experiments performed on a 2.35T horizontal bore imaging system, demonstrating the real-world performance of the technique.

Chapter 2

Stochastic NMR

2.1 Introduction

This chapter describes the concepts underlying the stochastic NMR experiment. Section 2.2 provides some background on NMR fundamentals, describing the physical phenomenon and how it can be used to determine chemical and physical properties of samples by observing differences in NMR spectra. Section 2.3 describes the various experimental methods for deriving NMR spectra. Section 2.4 examines the motivation for performing a stochastic NMR experiment. Finally, Section 2.5 describes how the stochastic spectroscopic experiment can be extended into an imaging technique.

2.2 Background

The nuclei of many isotopes have a non-zero spin angular momentum quantum number, I. Because nuclei are charged, this spin leads to a nuclear magnetic moment, μ. When these nuclei are placed in a magnetic field, the projection of the magnetic moment along the applied field direction (which by convention is the z axis of the system) is quantized into $2I + 1$ states which differ in energy by $2\mu B_0$, where B_0 is the applied magnetic field.

It is usually more convenient in NMR to characterize a nucleus' magnetic moment

in terms of the gyromagnetic ratio γ:

$$\gamma = \frac{2\pi\mu}{Ih} \tag{2.1}$$

where h is Planck's constant. Because of their angular momentum, nuclei placed in a magnetic field do not simply align with the applied field, they instead precess around the field at a characteristic frequency γB_0, known as the Larmor precession frequency. This nuclear precession is the basis of NMR spectroscopy.

Nuclear magnetic resonance experiments measure the properties of a spin system by determining the spectrum of NMR resonances near the Larmor frequency of the nucleus being studied. When the spin system is in equilibrium in a static magnetic field, all of the spins have random projections in the xy plane; the net magnetization is aligned with the z axis. Energy is put into the system by exciting it with a radiofrequency (RF) magnetic field; this induces transitions between the magnetic energy levels which rotates the net magnetization vector around the applied field direction. The resultant magnetization vector precesses in the xy plane at the Larmor frequency. The xy component of the field can be detected by an RF resonator. The spin system radiates radio frequency energy as the spins relax back to the low energy configuration. This radiated signal is digitized and processed to determine the properties of the spin system.

The chemical and physical environment of each nucleus will affect its Larmor precession. Modifications of the density of the electron cloud around a nucleus (by chemical bonds between atoms) will change the local B_0 that impinges on the nucleus. This is manifested as slight differences in the resonance frequencies (on the order of parts per million) of populations of nuclei in different chemical residues. This is known as chemical shift (denoted σ). For instance, hydrogen nuclei in lipids have a resonance frequency different from the nuclei in water molecules.

Also, the coupling of pools of nuclei to the lattice (a term referring generically to other energy storage modes in surrounding medium) and to each other will determine how long it will take the nuclei to relax back to their ground state so that they can be reexcited, and how long a group of nuclei that are precessing will remain coherent in their precession so that they produce a net magnetic field that can be observed.

These properties are known respectively as longitudinal (T_1) and transverse (T_2) relaxation times.

Chemical shift, relaxation times, and other properties are usually examined in the spectral dimension and are reflected in the amplitude, phase, relative position and width of resonance lines. These resonance characteristics in turn can indicate chemical and physical properties of the system being studied, such as chemical composition and concentration, physical environment of chemical species, rates of dynamic processes like proton exchange, flow and diffusion rates, *et cetera.*

Spatial information can be encoded into the NMR signal with the addition of spatial magnetic field gradients. These gradients alter the local polarizing field B_0, which in turn changes the Larmor frequency of the nuclei in different parts of the sample. This spatial information encoding can occur during either the excitation of the sample, the measurement, or both. For example, a two dimensional region of an object can be selectively excited by applying a magnetic field gradient while exciting the sample with a narrowband RF pulse. When the gradient field is applied, the resonance frequency of the spins in the sample will vary over a range determined by the strength of the gradient. By applying a narrowband RF pulse, only a fraction of these spins which experiences a particular range of local B_0 will be resonant with the applied field; therefore only a slab of magnetization will be excited. Similarly, if a gradient is applied to a group of spins which is already excited, and while the signal is being measured, the resonance frequency change will affect the signal measured from the sample; the resonance frequency of a spin indicates its position along the gradient (this is known as frequency encoding). If the gradient is applied transiently between excitation and measurement, the relative phase of the spins will be altered in the recorded signal (this is called phase encoding). By applying the proper set of gradient values during an NMR experiment, the NMR parameters described above (and the underlying physical and chemical properties that cause them) can be resolved into images.

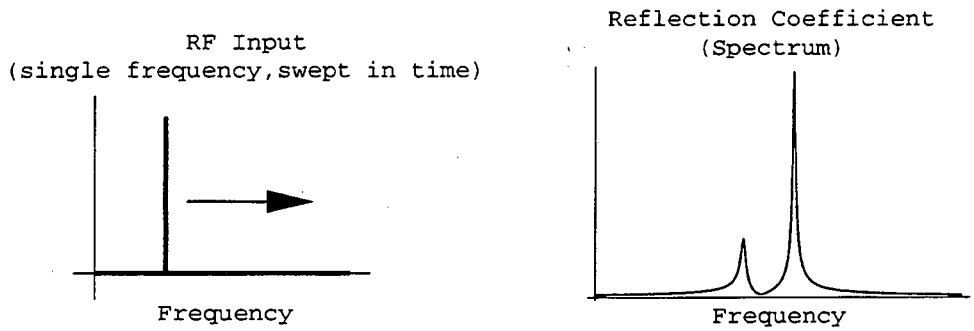

Figure 2.1: Schematic representation of the continuous wave (CW) NMR experiment. The RF input to the apparatus is a continuous signal which is slowly swept in frequency. The reflection properties of the probe are recorded over the frequency range and measure the NMR spectrum directly.

2.3 Spectroscopy

NMR spectra are measured by placing a sample in an RF resonator (probe) which can transmit and receive RF energy polarized in the plane perpendicular to the polarizing field B_0. A strong static magnetic field B_0 (typically 0.5 - 14T) is imposed on the apparatus to polarize the nuclei in the sample and causes them to precess at the Larmor frequency. The RF probe is tuned to this frequency, and the RF properties of the sample are measured.

The first method employed to determine NMR spectra was a direct frequency domain technique. A continuous wave (CW) single frequency RF signal is transmitted into a probe, and the reflection characteristics of the probe-sample system are measured. The excitation frequency is slowly scanned over the range of interest, and the resonator reflection coefficient changes as the spin system comes into resonance with the RF. This method, while straightforward, has a number of disadvantages. One is technical; it is difficult to average signals to increase sensitivity. The other limitation is the slowness of the technique - the scanning must be slow enough to give an accurate representation of the spectrum (the "slow passage condition"). The result is that a large fraction of the experiment time is spent traversing regions of the spectrum that have no resonances, which makes the technique inefficient. The CW experiment is shown schematically in Figure 2.1.

In 1966 Ernst[11] introduced Fourier transform NMR (FT-NMR) as a more efficient method of determining a spin system's characteristics. A broadband pulse of RF energy (usually 10 to 100 microseconds long) is transmitted to the probe to excite the nuclei; the probe is then switched to receive mode, and the response of the system (the Free Induction Decay, or FID) is recorded. This time domain signal is the impulse response of the spin system - which is the Fourier transform of the spectrum. Therefore the spectrum can be derived from the measured signal by a Fourier transform operation. FT-NMR excites all of the resonances in parallel, and records all of the responses simultaneously. The Fourier transform separates the responses into distinct spectral lines. This is a much more efficient process than single frequency scanning, as no additional time is spent exciting the spin system in frequency bands where there is no response to measure. The FT-NMR experiment is shown schematically in Figure 2.2.

When discussing pulsed RF NMR techniques, it is convenient to introduce the concept of the "flip angle" (α). The flip angle is a measure of how much magnetization is rotated from the z axis into the xy plane by an RF pulse. The ratio of the xy magnetization resulting from an RF pulse to the starting magnetization along the z-axis is the sine of the flip angle. A 90° pulse rotates all of the available magnetization along the z direction into the xy plane, therefore producing the maximum measurable signal. Once magnetization is rotated into the xy plane, the z magnetization grows back as e^{-t/T_1}. Therefore, subsequent pulses will have less initial magnetization to rotate into the xy plane for detection. The optimum signal to noise ratio per unit experiment time is achieved when the time between RF pulses (T_R) is comparable to the T_1 of the spins being studied. The flip angle which produces the maximum signal to noise ratio per unit time is called the Ernst angle:

$$\alpha_{\text{Ernst}} = \cos^{-1}\left(e^{-T_R/T_1}\right). \tag{2.2}$$

Ernst[12] and Kaiser[13] independently proposed yet another way to obtain NMR spectra in 1970; the use of stochastic excitation. The impulse response of a linear system can be derived by exciting the system with Gaussian white noise and cross-correlating the input and output signals. Although the NMR system response is not

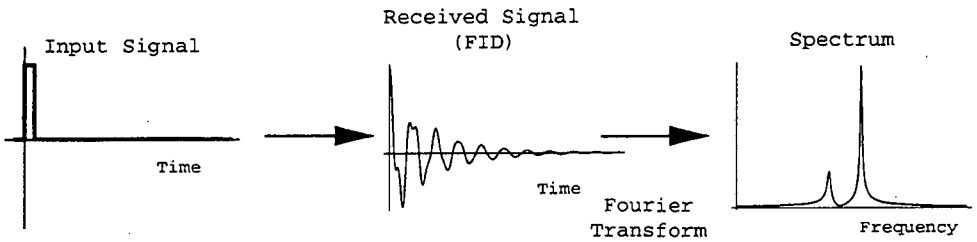

Figure 2.2: Schematic representation of the Fourier Transform (FT-NMR) experiment. The RF input to the apparatus consists of broadband pulses. After each pulse, the RF signal (the FID) from the sample is measured in the time domain. The FID is Fourier Transformed to give the NMR spectrum.

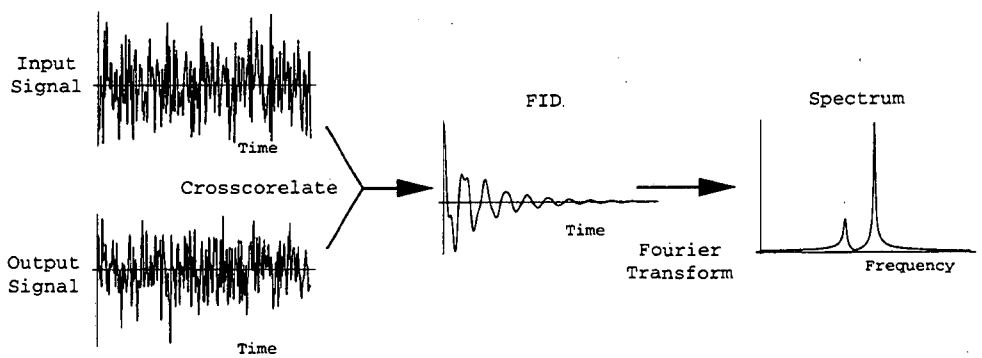

Figure 2.3: Schematic representation of the stochastic NMR experiment. The RF input to the apparatus consists of a train of random broadband RF pulses. The voltage on the probe is sampled after each pulse. The input signal is crosscorrelated with the output signal to give the FID, which is then Fourier transformed to give the NMR spectrum.

truly linear, it can be treated as linear for sufficiently small excitations[12, 13, 14]. As with FT-NMR, all resonances are excited simultaneously and encoded into a time domain signal. The stochastic NMR experiment is shown schematically in Figure 2.3. The total excitation power integrated over the course of the experiment is comparable to that of an FT-NMR experiment. The expression for the average pulse excitation angle for optimum signal to noise ratio in the stochastic experiment is given in Equation 4.4. This equation is similar to 2.2, with the difference that α_{Ernst} refers to the root mean square flip angle in the stochastic case.

Mathematically, measuring the impulse response of a linear system directly or through crosscorrelation is exactly equivalent. Ernst showed that the the signal to noise ratio of the optimum stochastic NMR experiment is equal to that of the

optimum FT-NMR experiment[12]. Thus at first sight it would seem that the significantly higher level of complexity of the stochastic experiment is unwarranted. In many practical situations, however, the optimum FT-NMR experiment is impossible to perform for a number of reasons, such as hardware or safety limitations on allowable peak RF power, or dead time in the receiver hardware following transmission of an RF pulse. Stochastic NMR experiments have a very different set of tradeoffs that make the technique very well suited to some spin systems that are difficult to study with conventional FT-NMR.

2.4 Advantages of Stochastic Excitation

The most obvious difference between stochastic excitation and FT-NMR is that while the total RF excitation energy applied during the experiment is about the same, the excitation duty cycle is very much higher in the stochastic experiment since the RF excitation is essentially continuous. As a result, the instantaneous power applied to the system can be very much smaller using stochastic excitation. This has the benefits of reducing the demands on RF components, such as probe breakdown voltage and peak amplifier power. The largest benefits, however, result from the reduction of system dead time.

The RF resonators used in NMR are high Q structures which store RF energy quite efficiently. As a result it takes some time for residual RF energy to dissipate after a pulse is applied to them. The peak power transmitted to the probes is often on the order of kilowatts for *in vivo* FT-NMR. This means that several thousand volts can be resonating in the probe after the excitation pulse is delivered. Typical induced voltages from the nuclei in the probe are on the order of microvolts, many orders of magnitude less than the residual RF voltages left over from the excitation. Also, high gain preamplifiers which are used in NMR imaging are easily saturated by signals much larger than the expected NMR signals. It can take hundreds of microseconds for the RF probe and receiver electronics to recover from a high power transmitted pulse to the point where the induced NMR signal can be observed. For many nuclei this is not a problem; however, for nuclei such as ^{23}Na, where the T_2

relaxation time is short (a few milliseconds and below), the dead time in the RF electronics can have a serious impact on the signal to noise ratio of the experiment. Much of the strongest NMR signal will have decayed away before the data are recorded.

NMR with stochastic excitation reduces the peak RF power by approximately three orders of magnitude relative to an FT-NMR experiment, which dramatically reduces the recovery time of the RF system. The dead time from excitation to sampling in the sodium experiments presented here are on the order of 20μs to 30μs, which means that even for the shortest T_2 expected *in vivo* there will be negligible signal loss before the onset of sampling. As a result, the imaging method is expected to have 100% visibility factor for sodium *in vivo*.

2.5 Extension to Imaging

The stochastic NMR experiment is not limited to performing spectroscopy. In 1984 Blümich proposed extending the stochastic experiment to include time varying gradients to perform imaging experiments[15]. He initially proposed using random gradients as a method of covering k-space while acquiring data, and developed a method of reconstructing images from the data sets using a crosscorrelation and averaging technique. This imaging method, however, was somewhat impractical due to transient eddy current response; he later modified the proposed experiment to use nonrandom sinusoidally oscillating gradients in three orthogonal dimensions[16].

In 1987, Roos and Wong extended the analysis of stochastic NMR to demonstrate how a spectroscopic imaging experiment could be performed using stochastic excitation and oscillating gradients[17]. This involved modification of the crosscorrelation method to include a weighting function which would correct for the sampling density of the gradient trajectory, and allow the signals with different chemical shift evolution times to be separated. They also performed a detailed analysis of the relationship between T_2, gradient parameters, and the resulting image quality.

The most serious obstacle to practical application of stochastic NMR imaging at this point was the difficulty of reconstruction. The weighted crosscorrelation method

is computationally expensive, and the reconstruction time grows linearly with the number of image voxels. In addition, the mathematical form of this reconstruction algorithm makes analysis of reconstructed image parameters quite tedious for all but the simplest gradient trajectories.

A much more efficient and easily analyzed reconstruction method for stochastic NMR experiments has been developed which is based on Fourier gridding, a technique for resampling arbitrarily spaced Fourier domain data onto Cartesian grid. Originally developed for radioastronomy, Fourier gridding has been applied to conventional NMR data in the last few years to allow the use of new sampling trajectories[18, 19]. The technique has been extended here to reconstruct stochastic NMR data. The new reconstruction method reduces reconstruction time by almost three orders of magnitude, and is quite amenable to detailed analysis of the relation between the choice of experimental parameters and the resulting images. The new reconstruction algorithm is described in detail in Chapter 4.

Chapter 3

Materials and Methods

3.1 Introduction

This chapter will describe the procedures and apparatus used for performing the stochastic NMR imaging experiment. In order to perform a spectroscopic stochastic NMR imaging experiment, there are a number of criteria that need to be satisfied by the RF and gradient waveforms. The major requirements on the RF sequence are that the individual RF pulses be broadband enough to excite all of the resonances in the spin system, and that the correlation properties of the train of pulses in the noise sequence allow for unambiguous reconstruction of an estimate of the spin density of the sample from the received signal. The gradients used in the experiment must fully sample the four dimensional k,t-space representation of the sample over the course of the experiment. The specific choice of RF sequence and gradient trajectory are independent and quite flexible. This flexibility allows the experimenter to base the choice of RF and gradients on practical considerations such as ease of generation, eddy current minimization, hardware limitations, and ease of reconstruction. Section 3.2 will describe the RF and gradient waveforms tailored to observing short T_2 nuclei, which will be used for the experiments presented. The effect of the physical properties of the spin system being studied on the choice of experimental parameters will be discussed. Section 3.3 will describe the hardware used to carry out the experiments, including descriptions of special hardware required, and Section 3.4

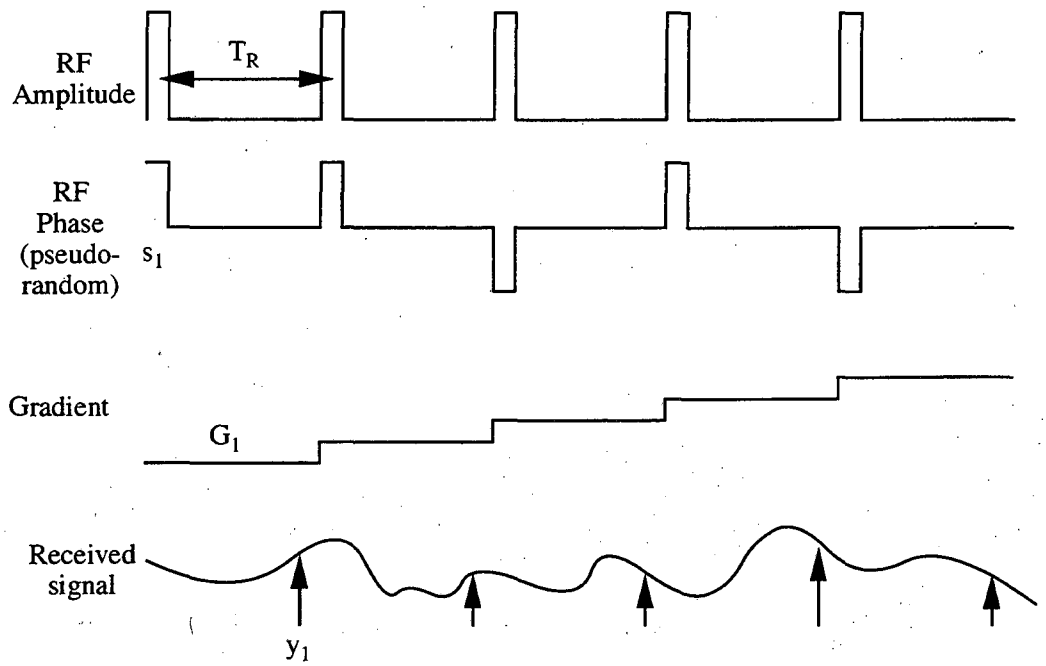

Figure 3.1: Pulse sequence of the stochastic experiment. s_n is the nth pseudorandom RF pulse, $\mathbf{G}_n$ is the gradient under which the magnetization evolves, and y_n is the NMR signal at the sampling time $n\mathrm{T_R}$.

will discuss special spectrometer calibrations which must be performed in order to do the stochastic imaging experiment.

3.2 Experiment Parameters

3.2.1 Pulse Sequence

The stochastic NMR experiment is diagrammed in Figure 3.1. The basic pulse sequence repeats throughout the course of the experiment (Figure 3.2): a new gradient value is asserted along all three axes, a short RF pulse is transmitted, and a data sample is taken after a period (e.g. 10-30μs) to allow for active and passive RF hardware recovery. This sequence repeats with a period $\mathrm{T_R}$ (usually 50-100μs) determined by the bandwidth of the signal from the spin system evolving under the gradients (Section 3.2.4). The RF pulse amplitudes and phases and the sequence of

gradient values determine the characteristics of the experiment.

3.2.2 RF Noise Sequence

The RF noise sequence used in the stochastic experiment is assumed to meet three conditions:

1. Each pulse excites all spins in the system
2. The repetition time of the noise sequence is long compared to T_1.
3. The autocorrelation of the excitation sequence is a delta function in the time domain.

The first requirement for the RF excitation is that each pulse be broadband enough to excite the entire spin system. This is accomplished by using short hard pulses. Typical RF pulses are 3μs to 15μs long, giving single pulse excitation bandwidths of 67kHz to 333kHz. These bandwidths are typically 3 to 60 times the bandwidth of the spin system evolving under the gradients which are used for spectroscopic imaging.

The second and third conditions are on the properties of the overall sequence of pulse amplitudes and phases. In order to uniquely reconstruct the k,t-space representation of the spin system using the Fourier gridding crosscorrelation method described in Chapter 4, the noise sequence must satisfy two conditions. First, for a noise sequence which repeats, the repeat time must be longer than 5 times the longest T_1 in the spin system to avoid setting up steady state magnetization which would lead to artifacts in the reconstruction. This consideration is easily met for even long T_1 values and short T_R values. For example, a sample with a T_1 of 5 seconds and a T_R of 50μs (an extreme case) would require a noise sequence at least 500,000 points long, which is quite practical. In an *in vivo* sodium experiment, T_1 values will be less than 100 ms, so a repetition time of >10,000 points is sufficient.

The last condition on the RF noise sequence is that the autocorrelation function of the sequence should be a delta function. This condition is necessary to be able

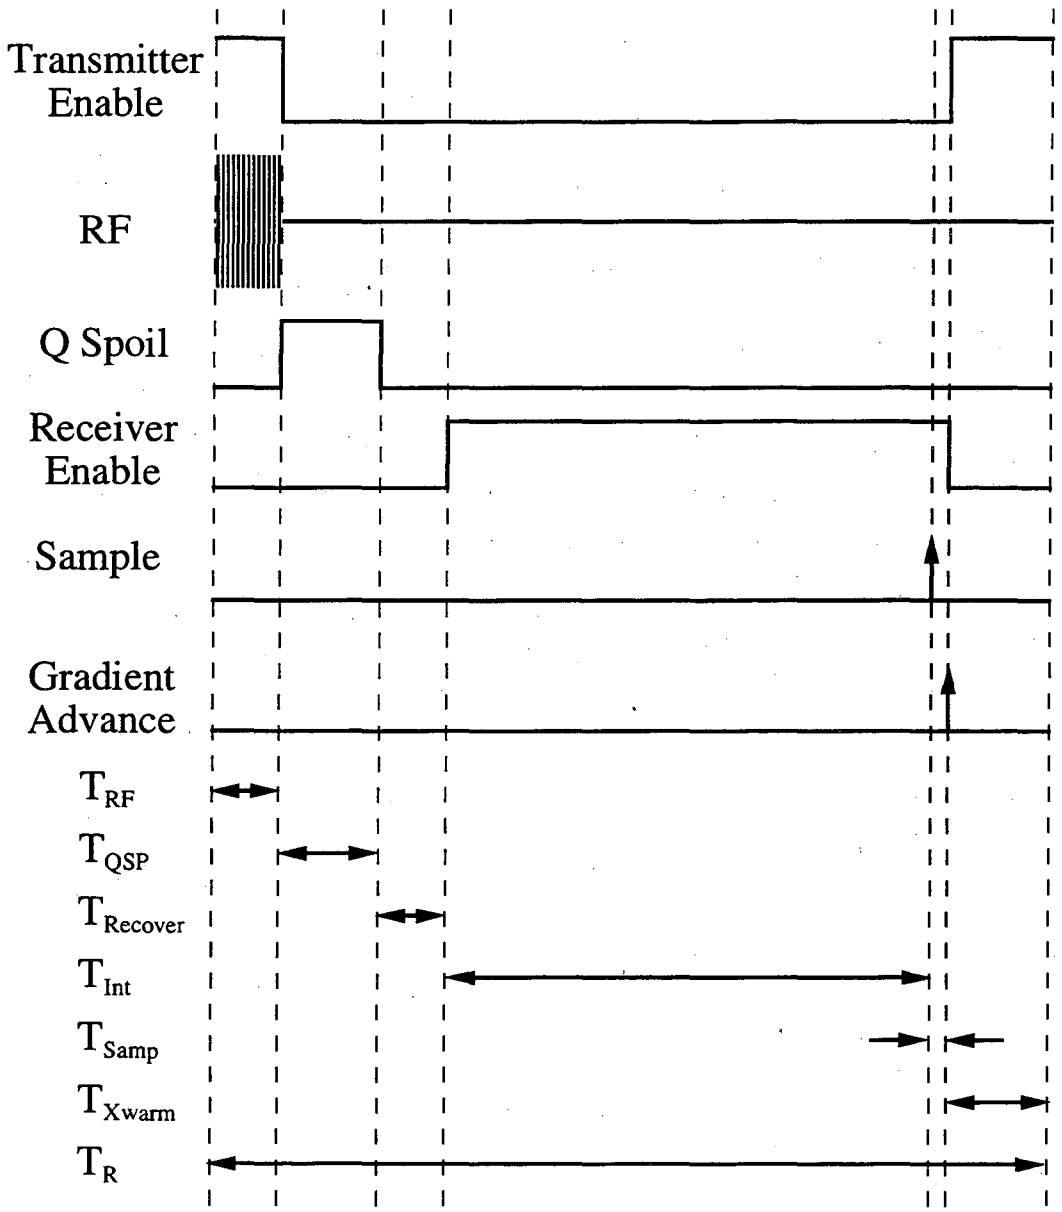

Figure 3.2: Single cycle of the pulse sequence for the stochastic imaging experiment. For a typical sodium experiment, $T_{RF} = 15\mu s$, $T_{QSP} = 14\mu s$, $T_{Recover} = 10\mu s$, $T_{Int} = 32\mu s$, $T_{Samp} << 1\mu s$, $T_{Xwarm} = 4\mu s$, $T_R = T_{RF} + T_{QSP} + T_{Recover} + T_{Int} + T_{Xwarm} = 75\mu s$.

to uniquely reconstruct the object density from the received signal. Each data sample contains information from different chemical shift evolution times and k-space positions; the validity of the reconstruction method for forming a k,t-space image estimate relies on the fact that the noise sequence has a delta function autocorrelation (see Section 4.4). In practice, this last condition is very difficult to attain exactly. The T_2 of the spins in the sample will weight the autocorrelation function so that spins with different T_2 values will see different noise sequences from the same excitation. Also the gradients will affect the excitation values included in the autocorrelation differently for every position in k-space. At best one can choose a noise sequence which is relatively free from high order autocorrelations, with the result that for any given combination of points from the noise sequence the autocorrelation is "delta function like".

One class of RF modulation functions which satisfies these conditions is quadrature phase modulation using pairs of binary maximum length sequences (MLS). MLS generators provide a sequence of pseudo random bits with desirable autocorrelation properties. An RF pulse of constant amplitude is phase modulated between 45°(+X+Y), 135°(-X+Y), 225°(+X-Y), and 315°(-X-Y). The sign of the phase along each axis is selected by a bit output from a separate MLS generator. This modulation scheme was examined in detail by Wong [14], and found to be quite effective for stochastic NMR. MLS's are attractive because they are very easily generated in simple hardware using a shift register, or in software using bit manipulation operations. Digital signal processing chips can generate the two random bit values needed for the RF envelope in real time at the highest excitation frequencies used for the experiments presented (20,000 phase values/second).

Maximum length sequences are specified by three integers: the number of terms in the generating polynomial, the "mask value" (which specifies compactly which terms of the generating polynomial have non-zero coefficients), and the seed, which specifies the starting state of the generator. The number of terms in the generating polynomial determines the repeat time of the sequence; an MLS with N terms repeats after $2^N - 1$ points. The experiments described here use 19 term MLS sequences, which repeat in 524287 points, which will always be longer than 5 T_1 values for any

practical sodium experiment. Different mask values specify different bit sequences of this length, and the seed determines the starting point within a given sequence. For simplicity these experiments use the same mask value for both bits of the phase selection, but start the generator at different seed values which are 2^{N-1} points apart in the bit sequence.

3.2.3 Gradient Trajectories

The gradient trajectory used in the stochastic experiment must satisfy the following conditions:

1. k-space must be critically sampled over a simply connected region.
2. Each point in k-space must be critically sampled in the time domain to allow spectral reconstruction.
3. Gradient slew rates must be within hardware and safety limitations.
4. Eddy currents should be minimized or fully characterized to eliminate dead time due to gradients.

All of these conditions can be met through the use of gradient waveforms based on one dimensional periodic oscillating gradients. Periodic waveforms retrace themselves over a region of k-space through the origin, providing k-space samples of arbitrary density along a line. They also revisit the same region in k-space at least once per cycle, so if the oscillation frequency is at least equal to the chemical shift bandwidth of the system, k,t-space will be adequately sampled to reconstruct a spectroscopic image. Periodic waveforms are easily generated by resonant hardware, so very large gradient slew rates can be achieved with relatively low power input. Finally, periodic gradients by necessity have steady state eddy currents.

Roos and Wong[17] examined the special case of sinusoidally oscillating gradients, which have very desirable characteristics from the point of view of eddy currents and ease of generation. Unfortunately, the k-space sampling density which results

from using a sinusoidal gradient and equal time sampling is not uniform; the sampling density is strongly peaked at high spatial frequencies. This leads to uneven distribution of measurement noise power in the reconstructed image, and uneven cancellation of systematic noise arising from stochastic excitation. Two methods were considered for modifying sinusoidal gradients to correct the sampling density while preserving the ease of generation and the steady state eddy current behavior.

3.2.3.1 Truncated Square Wave (TSW) Gradients

Because periodic gradients will always have steady state eddy currents, a periodic gradient can have an arbitrary shape depending on the sampling density required. In particular, a much more uniform sampling density can be obtained by adding odd harmonics to the fundamental sinusoid to generate an approximation to a square wave gradient, as shown in Equation 3.1.

$$\mathbf{G}_p = \sum_{i=1}^{i_{max}} \frac{G\cos\left(2(2i-1)\pi f_o p \Delta T\right)}{(-1)^i(2i-1)}. \tag{3.1}$$

The peaking of the sampling density at high spatial frequencies is greatly reduced by adding more harmonics to the gradient, but it also increases the maximum gradient slew rate. The experiments analyzed and presented in Chapters 5 and 6 using truncated square wave gradients use the first three components ($i_{max} = 3$). This number of harmonics represents a good compromise between the even sampling density of square wave gradients and the limited slew rate of sinusoidal gradients. The sampling density of this class of one dimensional periodic gradients will be analyzed in more detail in Section 4.6.2.

3.2.3.2 Amplitude Modulated Sinusoidal (AMS) Gradients

A different approach to correcting the sampling density of a sinusoidal gradient is to slowly modulate the amplitude of the oscillating gradients [21]. Changing the gradient amplitude gives sampling patterns of the same shape but covering different extents in k-space. The lower sampling density near the origin of k-space can be corrected by emphasizing smaller gradient amplitudes in the modulation function.

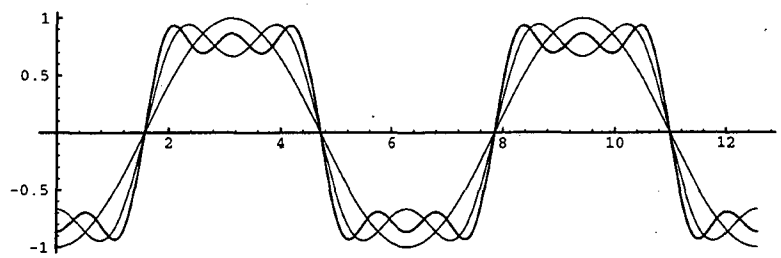

(a) Gradient waveforms

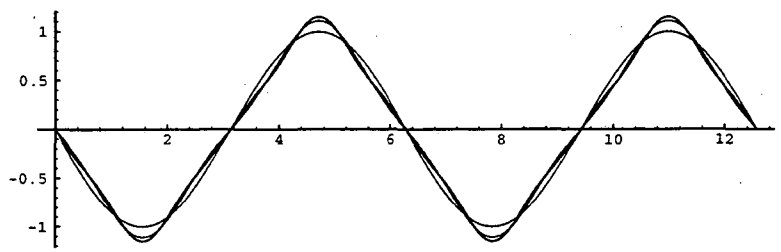

(b) k-space trajectories

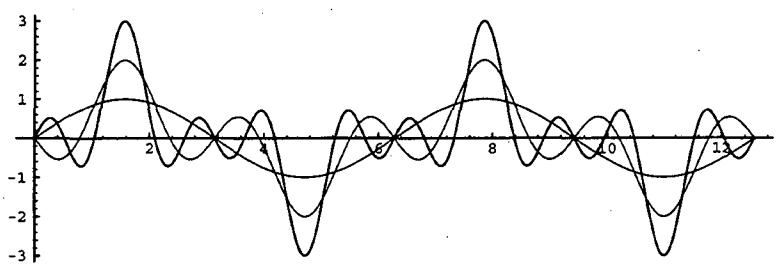

(c) Gradient slew rates

Figure 3.3: One dimensional gradient waveforms generated by Equation 3.1. (a) shows the gradient waveforms for i_{max} of one, two and three. (b) shows the resulting k-space trajectories (in arbitrary units) and (c) shows the gradient slew rates (again in arbitrary units).

The modulation must be slow relative to the period of the sinusoid to keep the shape of the sinusoidal sampling density constant. In a typical experiment, the oscillation frequency will be 200-1000Hz, and the modulation period will be on the order of a few seconds. This yields a flat spatial noise power spectrum in the reconstructed images. Because the modulation is very slow compared to the time for eddy currents to come into equilibrium, steady state eddy currents are preserved, and resonant gradients can still be employed.

3.2.3.3 AMS Modulation Function

A stochastic NMR experiment has (in general) a different sampling pattern for every time point q in the reconstructed FID image (see page 51). For a stochastic experiment with a sinusoidal gradient and a constant sample period, the number of samples at a location in k space for a time point q is proportional to:

$$W^q(k) = \frac{1}{\sqrt{1 - \left[\frac{k}{k^q_{max}}\right]^2}}, \tag{3.2}$$

where k^q_{max} is the maximum extent in k-space for a given time point t_q. This sampling pattern is shown as the dashed line in Figure 3.4.

To construct an envelope function which fills in the center of k-space, we make k^q_{max}, which is proportional to the amplitude of the sinusoidal gradient (G), a function of time. A $G(t)$ of the form:

$$G(t) = \begin{cases} \sqrt{1.0 - \left[\frac{2t}{t_{max}}\right]^2} & |t| \leq \frac{t_{max}}{2} \\ 0 & \text{otherwise} \end{cases} \tag{3.3}$$

yields a constant sampling density in k when integrated from 0 to t_{max}, where t_{max} is long compared to one period of the sinusoidal gradient. This function can be made periodic by convolving it with a series of delta functions spaced t_{max} apart. If the total experiment time is a half integral number of t_{max} long, the sampling density will be constant over the range $-k_{max}$ to k_{max}. The modulated waveform is shown in Figure 3.5, and the resulting sampling density is shown as the solid line in Figure 3.4.

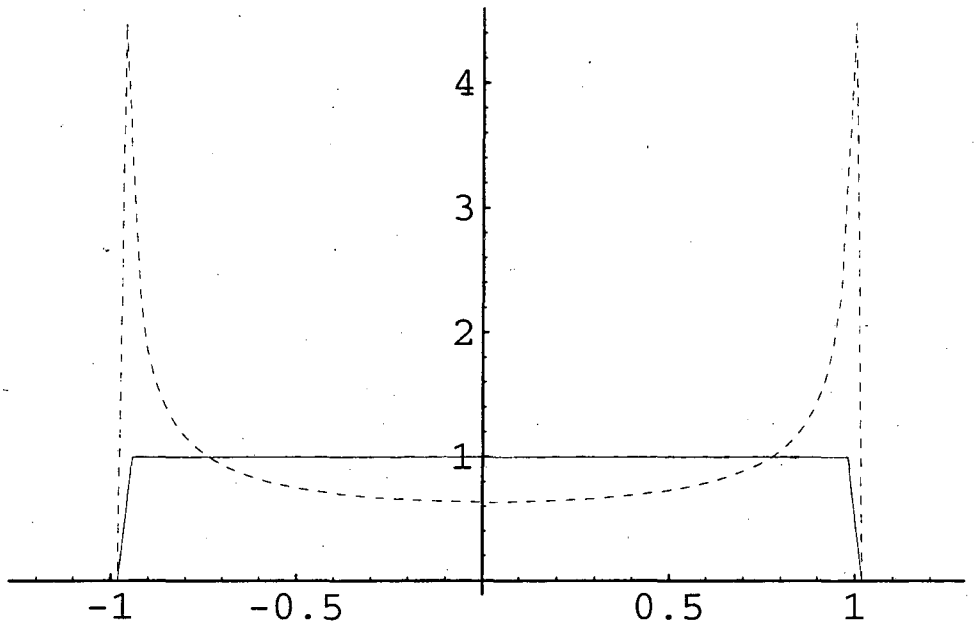

Figure 3.4: Normalized sampling density over range $-k_{max}, k_{max}$ with sinusoidal gradient (dashed) and modulated sinusoid (solid). Histogram shows the distribution of 1048576 sample points distributed over 104 bins.

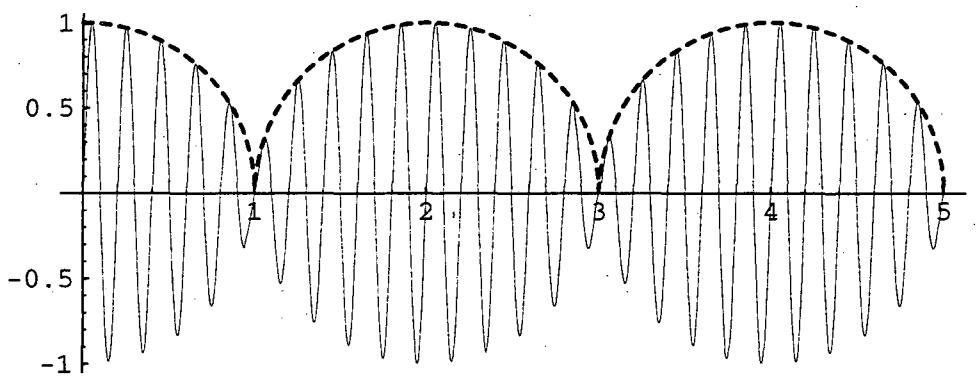

Figure 3.5: Sine wave modulated by periodic G(t) from time 0 to $5t_{max}$ (solid line). The frequency of the sinusoid is reduced for clarity. In an actual experiment, t_{max} >> one oscillation period. The modulation function is the dashed line.

3.2.3.4 Gradient Oscillation Frequency

The frequency of the gradient is determined by the chemical shift bandwidth of the sample. ^{23}Na has no natural chemical shift *in vivo*. However, there are two sources of chemical shift bandwidth to the sodium system: T_2 relaxation and shift reagents. As mentioned before, the shortest relaxation component of sodium may be as low as 300μs. This leads to a linewidth of approximately 1000Hz. Therefore the oscillating gradient frequency may need to be as high as 1000Hz for some *in vivo* experiments to fully capture the short relaxation signal. The effect of shift reagents is usually somewhat less; on the order of a few hundred Hz, so the T_2 value will usually determine the minimum useful gradient frequency. The experiments presented here use gradient oscillation frequencies of approximately 500 to 600Hz.

3.2.3.5 Extension to Three Dimensions

There are many ways to extend a one dimensional oscillating spectroscopic imaging experiment to three dimensions; two methods will be examined here. These trajectories are chosen for ease of generation, and for the fact that they allow the oscillating part of the gradient and the method of extension to three dimensions to be considered separately; the requirements on gradient frequency and strength are the same for a one dimensional and a three dimensional experiment.

The first method is to rotate the direction vector of the oscillating gradient so that it slowly traces out a spiral on the surface of a sphere. This is essentially a three dimensional projection experiment. Over short time scales, this trajectory repeatedly samples a line in k-space through the origin. Over the course of the experiment, the direction of this line rotates to sample all solid angles. The resulting sampling pattern is a sphere in k-space. If the rotation is very slow compared to both the oscillation frequency of the gradient and to the longest time constants of eddy currents in the magnet bore, the rotation does not invalidate the assumption of steady state eddy currents.

The second method of extending oscillating gradients to three dimensions consists of simultaneously applying periodic oscillating gradients which are not harmonically

related along the three Cartesian axes (called "incommensurate frequency oscillating gradients", after Blümich[16]). The incommensurate frequency gradients trace out a three dimensional Lissajous pattern, giving a cubical sampling patterns in k-space. By proper choice of the relative frequencies of the oscillating gradients and the sampling frequency, the sampling density in k,t-space can be made very uniform. The relation of the gradient trajectory to the point spread function and the noise characteristics of the reconstructed image are examined in Sections 4.5 and 4.6.

As a practical matter, the second method, incommensurate gradients, is more desirable for these experiments, for a number of reasons. The first is that the overall k-space sampling density can be made uniform in three dimensions for all time points, which is very difficult to do using the rotating oscillating gradient trajectory; also, the point spread function is the familiar three dimensional sinc, which has well understood sidelobe properties. Also, because the gradient functions along the three axes are independent, their eddy currents and phase shifts can be mapped and compensated independently (using the procedure described in Section 3.4.1).

The sampling density of a one-dimensional stochastic imaging experiment with an unmodified sinusoidal gradient varies by a factor of 8 from the center to the edges of k-space. In a three dimensional incommensurate frequency oscillating gradient experiment, most of the sampling density is concentrated in the eight corners of a cube in k-space, where there is little object information; the sampling density at the origin of k-space is only 25% of the average value, and 1/64th of the peak value. Flattening the sampling pattern using the three dimensional AMS gradient leads to significantly reduced systematic and measurement noise in the reconstructed image for a given imaging time.

If the AMS technique is used for a three dimensional experiment, care must be taken to keep the gradient modulation functions incommensurate in addition to having the oscillation functions incommensurate. If the gradients are modulated synchronously, the gradients cannot be considered fully independent in terms of sampling pattern, and the flat sampling density in one dimension will not be preserved in three dimensions. The resulting sampling pattern is concentrated along diagonal lines radiating from the origin. Since the modulation function can be made periodic,

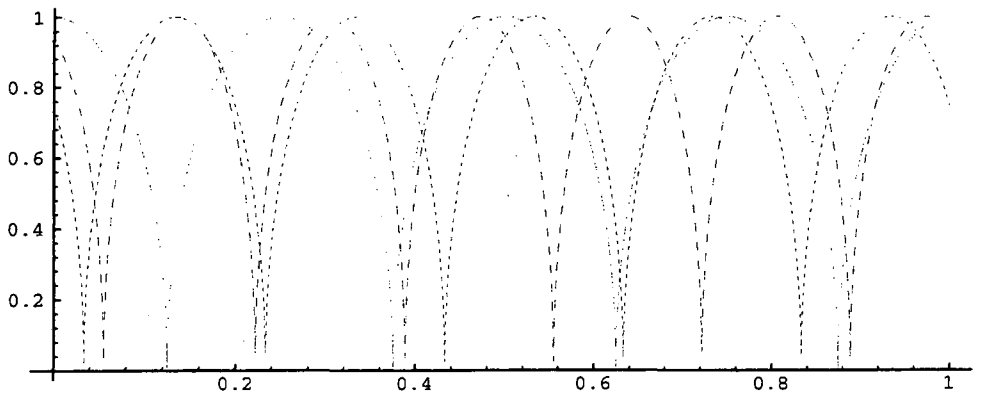

Figure 3.6: Modulation functions for x, y and z. The number of modulation periods has been chosen so that the gradients along x, y and z are modulated asynchronously. In this example the three modulation function goes through 8, 10 and 12 t_{max} along x, y, and z during the course of the experiment.

the number of periods along each axis during the experiment can be different, which eliminates the concentration of sampling along the diagonals of k-space.

3.2.4 Sampling Frequency

The sampling frequency required for the stochastic experiment depends on the amplitude of the gradient, the field of view covered, the oscillation frequency of the gradient, and to a lesser extent, the shape of the gradient.

Oscillating gradients have the effect of frequency modulating the NMR signal produced by the spin system. This modulation produces sidebands spaced at multiples of the oscillation frequency, with an envelope function on the sideband amplitudes determined by the shape of the object. The sideband amplitudes never fall exactly to zero; however, there is an effective bandwidth where most of the energy is concentrated. A common rule adopted in communications is that sideband is significant if its amplitude is greater than or equal to one percent of the unmodulated signal[22]. In analogy to the terminology of communications, the oscillating gradient amplitude and the field of view are expressed as the quantity β, the "modulation index":

$$\beta = \frac{\gamma \mathbf{G} r_{\max}}{f_0} \tag{3.4}$$

where γ is the gyromagnetic ratio, $\mathbf{G}$ is the gradient amplitude, $r_{\max}$ is the maximum extent of the object relative to the origin along the direction of the gradient, and f_0 is the gradient oscillation frequency.

For a sinusoidal oscillating gradient, the effective signal bandwidth is determined using Carson's rule[17]. The approximate bandwidth of the stochastic experiment with a sinusoidally oscillating gradient is therefore:

$$\Omega = 2f_0(1 + \beta) \tag{3.5}$$

The bandwidth of a three dimensional rotating oscillating gradient is the same as the sinusoidal gradient case, because the rotation frequency is very slow compared to the oscillation frequency the rotation makes a negligible contribution to the bandwidth. In the case of the incommensurate frequency gradient, the gradients can be considered independently, and the experiment bandwidth is the maximum of the bandwidth determined by Equation 3.5 along each of the three axes.

Because the modulation frequency of the AMS gradients is very slow compared to the sinusoidal oscillation frequency, the modulation has no observable effect on the overall bandwidth, and Carson's rule holds using the maximum amplitude of the gradient to calculate the bandwidth.

Frequency modulation is a nonlinear process, so the bandwidth resulting from multiple component gradients cannot be simply determined; the expression for sidelobe amplitude is rather unwieldy. Simulations and experiments have shown that in the regime of $\beta >> 1$, which is the usual case for stochastic NMR, the bandwidth of the sinusoidal experiment and the multiple component gradient is very similar, and Equation 3.5 can be used to determine the necessary sampling frequency. Power spectra from simulated experiments are compared in Section 5.3.

3.3 Apparatus

3.3.1 Spectrometer

All experiments were performed on a 2.35T imaging system. The imager is homebuilt around a 2.35T Bruker horizontal bore magnet with a 40cm bore. The gradients employed are custom designed actively shielded gradients with a clear bore of 25cm, powered by six Techron 7700 series amplifiers. The gradients have a maximum field strength of 60mT/m and a maximum slew rate of 300T/m/s.

The data acquisition system is divided into two parts; all realtime experimental tasks are handled by a dedicated multiprocessor VME-based computer system. Gradient and RF waveforms can be precomputed or generated in real time using two TMS320C30 digital signal processors on a Sky Challenger processor board; averaging and control of the pulse programmer and RF hardware is handled by a Motorola MVME147 68030 processor. All display and interface tasks are handled by Sun Microsystems SPARCstation-10 which is connected to the realtime system using socket connections over an Ethernet connection.

3.3.2 Probes

The probe used for the ^{23}Na experiments was a homebuilt dual tuned (^{1}H/^{23}Na) four-ring lowpass-lowpass birdcage resonator (Figure 3.7). The probe layout was adapted from the ^{1}H/^{31}P probe devised by Murphy-Boesch and coworkers.[23]. The ^{1}H mode of the probe is used for shimming and sample positioning; it is driven linearly. The ^{23}Na mode of the birdcage transmits and receives in quadrature.

An NMR probe is designed to be a high Q resonant structure for maximum signal to noise ratio. In order to minimize the preamp recovery time, an active Q-spoiling circuit was developed and added to the RF probes used for stochastic NMR experiments. This circuit provides a means of actively dissipating the power remaining in the probe after transmission by lowering the ^{23}Na probe Q by a factor of 6.7 for several microseconds. The Q-spoiler circuit is mounted on two panels 90° apart on the probe body (one for each linear mode in the birdcage). The Q-spoiling

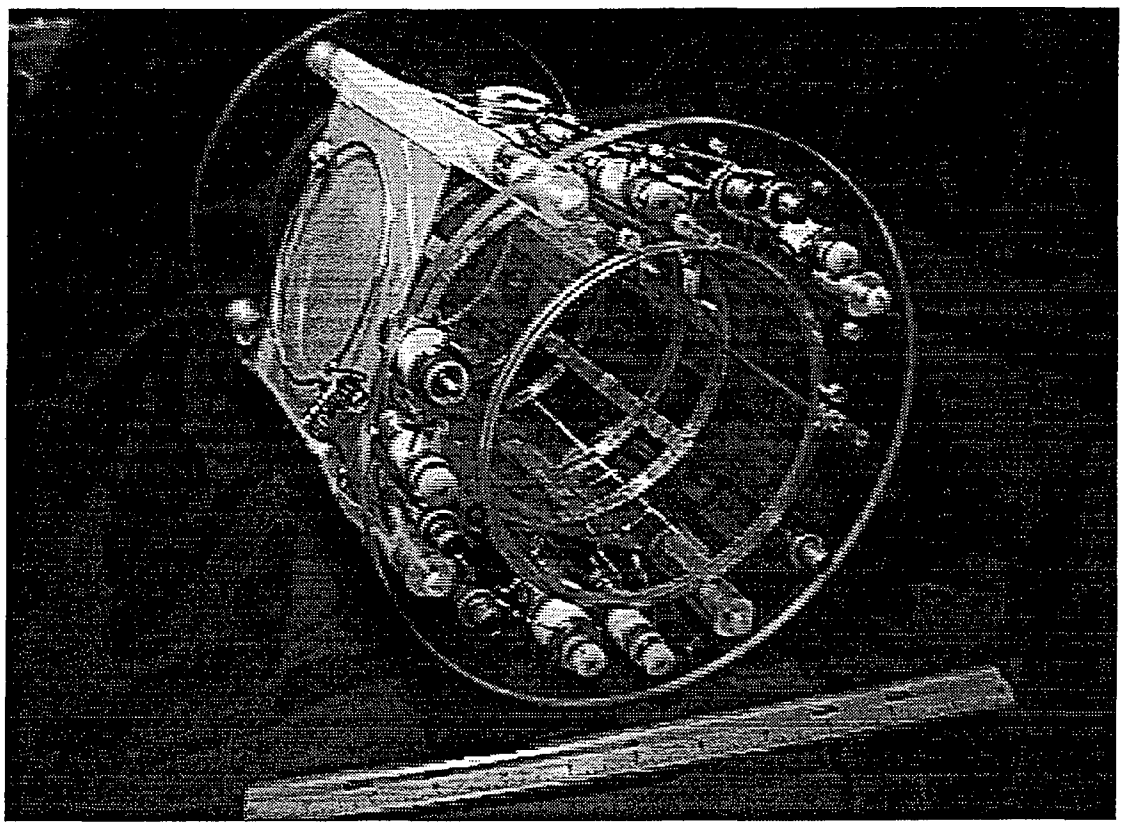

Figure 3.7: Dual tuned four-ring lowpass-lowpass ($^1H/^{23}Na$) birdcage resonator with quadrature Q-spoiler. One of the two Q-spoiler loops can be seen mounted on the left side of the probe; the second loop is on the bottom of the probe.

circuitry is described in the next section.

The ^{1}H experiments were performed on a single frequency Alderman-Grant probe with an added Q-spoiler circuit. The probe is not shown.

3.3.3 Active Probe Q-spoiling

To maximize the signal to noise ratio, especially with short T_2 samples, the receiver gate should be open as long as possible during the sampling period to allow signal to integrate before sampling. To prevent the preamplifier from saturating, the probe must be given time to ring down after transmission of an RF pulse. This ringdown time can be quite long for a high Q probe, which is necessary for sensitive detection of the NMR signal from the sample. To resolve this apparent paradox, a

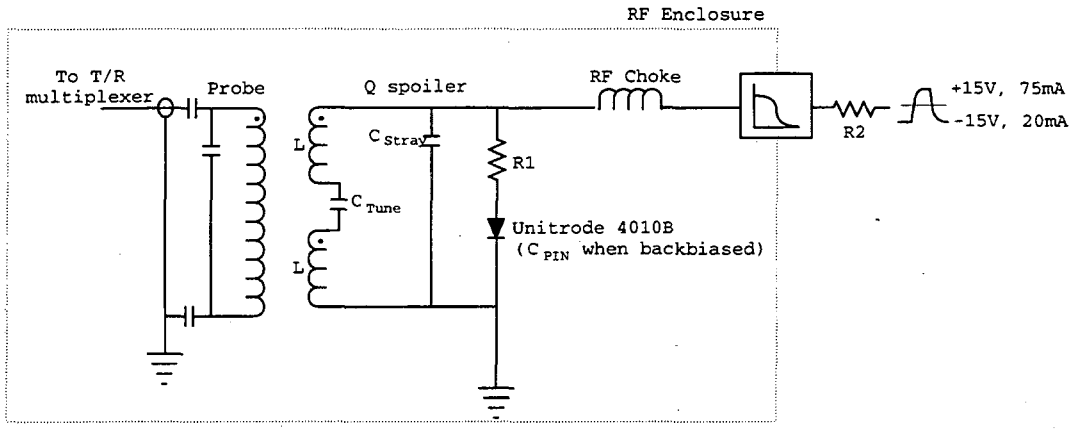

Figure 3.8: Circuit diagram, PIN diode Q-spoiler. The Q-spoiler is inductively coupled to the main tank circuit of the RF probe. When the PIN diode is backbiased, the resonant frequency of the Q-spoiler is much higher than that of the tank circuit, and does not affect it. When the diode is forward biased, the Q-spoiler is resonant at the probe's frequency and couples strongly to the main tank circuit, lowering its effective Q.

circuit was designed to give the RF probe a switchable Q, so that the residual power from the excitation pulse could be dissipated prior to opening the receiver gate.

PIN diode circuits have been employed previously to reduce probe recovery time [24, 25, 26] and to eliminate interaction of multiple coils [27]. These existing circuits, however, either generate noise in the high Q state or require extremely high voltage diodes, which in turn have high capacitance and slow switching times. Therefore an inductively coupled PIN diode circuit for rapid switching of an imaging probe between a high Q and low Q state was developed [28]. The circuit described allows the dead time between the end of RF transmission and receiver gate opening to be reduced significantly, which is essential for detection of spins with short T_2 and for stochastic NMR.

The Q-spoiling is achieved by inductively coupling a tuned split resonator to the probe, as shown in Figure 3.8. During transmit and receive, the PIN diode is reverse biased, forming a low Q resonator with resonance frequency:

$$\omega = \sqrt{\frac{C_{Tune} + C_{Stray} + C_{PIN}}{L(C_{PIN} + C_{Stray})C_{Tune}}} \approx \frac{1}{L(C_{PIN} + C_{Stray})}$$
$$C_{Tune} >> C_{PIN} > C_{Stray} \qquad (3.6)$$

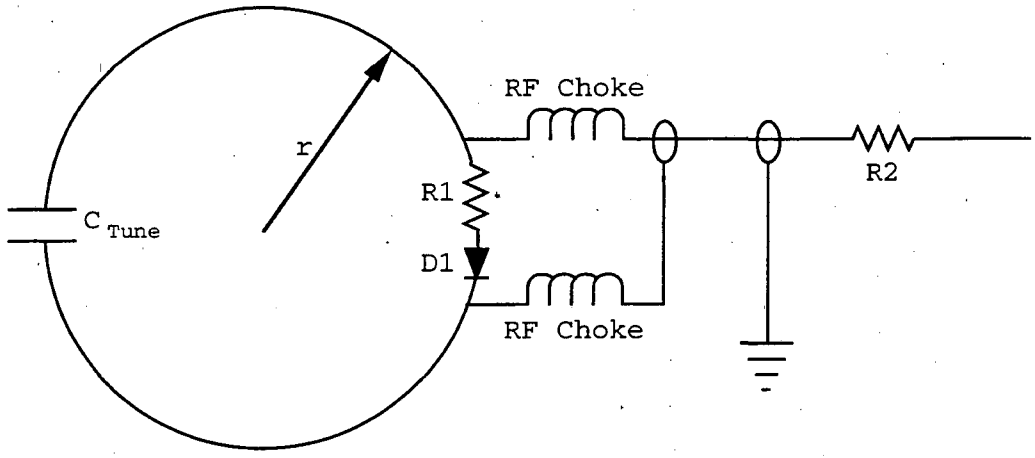

Figure 3.9: The Q-spoiler used for the birdcage coil used in this experiment consists of two loops 90° apart on the probe body, each coupled over a window of the birdcage. The loop is constructed of 14 AWG wire, and has an inductance L. For the 26.46MHz probe, r = 5.5cm, C_{Tune} = 130pF, R_1 = 9Ω, and R_2 = 120Ω. D_1 is a Unitrode 4010B PIN diode.

This frequency is much higher than the probe resonance frequency, so the effect on the main probe resonator is quite small. Also, because the coupling is not resonant, the voltage across the diode is much lower than the voltages on components in the tank circuit. When the diode is forward biased during the Q-spoil interval, the loop's resonance frequency becomes:

$$\omega = \frac{1}{\sqrt{LC_{Tune}}} \tag{3.7}$$

The loop components are chosen so that this frequency corresponds to the resonance frequency of the main probe. Adding this coupled resonator to the probe circuit splits the resonance and can lower the Q significantly, depending on the value of resistor R1 and the coupling constant k. The use of a split inductor makes the effect of switching the PIN diode on the loop's frequency quite large; switching on the diode doubles the loop inductance. This causes a much greater effect on the loop's resonance than relying on the capacitance change of the diode.

The actual design of the Q-spoiling loop is accomplished using an heuristic approach. The loop geometry is chosen much the same way as an inductive feed would

be designed; the coupling constant of the tuned loop is relatively high. A one turn loop split into two halves was used for this design. When the diode is reverse biased, the inductance of this loop is half of the value of when it is conducting. The capacitor C_{tune} is then selected to make the loop resonate at the probe's primary resonance frequency when the diode is on; (which corresponds to the loop inductor being joined into one continuous inductor). In this state, the resistor R1 is part of the circuit, lowering the Q of the loop, and hence the Q of the probe. If the loop Q is too low, its coupling to the main tank circuit is small, and it does not affect the probe Q strongly. If the Q is too high, the probe's resonance is split in two, but the Q is not decreased significantly. The resistor R1 can be selected to fall between these two extremes with the use of a variable resistor.

When the PIN diode is back biased, the loop's inductance is halved, and the Q-spoiler's resonance frequency increases to a high value determined by the lower inductance and the stray capacitance of the loop. Unless the stray capacitance is on the order of the tuning capacitance (in which case the loop is too big), the loop resonance frequency will be far above the probe tuning frequency, and its effect on probe operation will be very small.

The driver design is straightforward. The PIN diode requires 75mA to switch on, and it only draws this current for a few microseconds. However, it must be able to switch quickly while driving the capacitance of the probe loop and the lowpass filter, so when designing the driver the peak current requirements must also be considered. The resistor R2 is a current limiting resistor chosen small enough to fully forward bias the PIN diode into a low resistance state. To limit the transients caused by switching the diode on and off, the driver circuit for the diode was given a $\sim 1\mu s$ time constant to eliminate sharp transitions which could excite the main probe resonator.

The addition of the Q-spoiling circuit caused no observable decrease in the Q of the probe (relative to the probe without the circuit installed) when the diode was reverse biased. When the Q-spoiler was forward biased, however, the unloaded Q dropped by a factor of 6.7. When the probe was installed in the system shown in Figure 3.10 (which includes a multiplexer, preamplifier, and receiver), the probe dead time could be decreased as shown without transmitter breakthrough. The

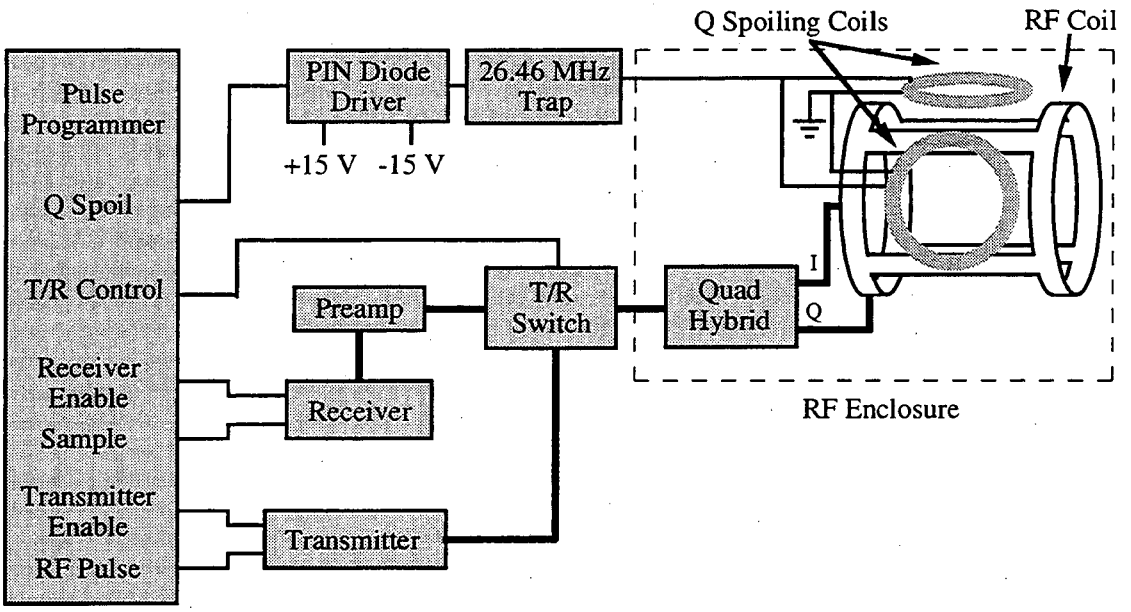

Figure 3.10: RF System Diagram.

prediction of the system dead time with and without using the spoiler is not as simple as directly comparing the Q's of the two systems, however. The matching of the probe to its cable varies during the course of the probe ringdown as the PIN diode is switched and the blocking diodes in the receiver protection circuits turn off, making predictions about ringdown based solely on Q values impossible. Also, when the probe is not matched to its cable, and the preamp can begin to recover from saturation before the probe ringdown is complete, since it is isolated from the full probe voltage, allowing parts of the system to recover in parallel rather than sequentially.

Figure 3.11 shows the Q-spoiler in use in the stochastic experiment diagrammed in Figure 3.2. In both oscilloscope photographs, the top trace shows the RF pulse, and the bottom is the analog input to the spectrometer's digitizer. In both cases, there is a 24μs delay between the end of the RF pulse and the opening of the receiver gate. On the left, the Q-spoiler is not used and the preamplifier is still saturated during the sampling period; on the right, the Q-spoiler is on for 12μs. In these pictures, the analog receiver filter before the digitizer has been removed to show the

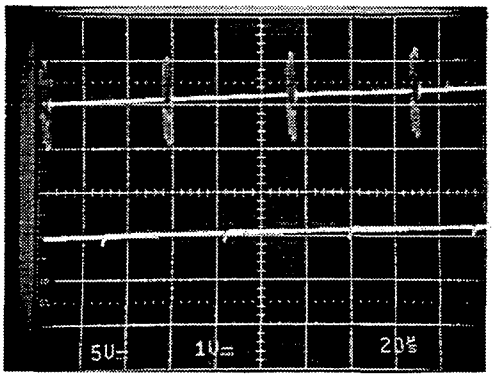

(a) Q-spoiler off. The preamplifier is saturated from the transmitted pulse.

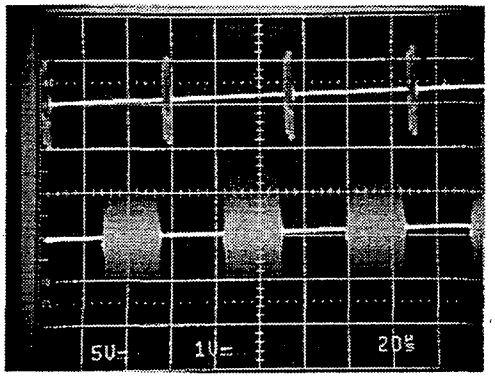

(b) Q-spoiler on for 12μs after end of RF. The preamplifier has fully recovered and is receiving normally.

Figure 3.11: Demonstration of Q-spoiler.

output of the preamp more clearly. In both pictures, the experiment was run for 65536 pulses, so the signals were summed on the film.

The receiver dead time depends not only on the probe ringdown, but on the preamp saturation characteristics and the recovery time of the T/R multiplexer. This method of decreasing probe recovery time takes all of these factors into account. The inductively coupled design allows a high degree of decoupling between the driver circuit and the transmitter and receiver; the need to protect the driver from the transmitted pulse is greatly reduced over designs which place the diode in the main probe circuit. Also, the transients arising from the switching of the diode circuit are only weakly coupled into the receiver circuitry, since very little of their energy is within the probe's bandwidth. Various other designs were tried, such as putting the PIN diode directly across the tuning or matching capacitors of the tank circuit of the probe. While these designs did give switchable Q values in low power testing, in actual circuits they exhibited decreased signal to noise and oscillating matching impedance. This is most likely due to nonideal behavior of the PIN diodes when exposed to the full RF voltage in the tank circuit. Under these conditions, the model of the PIN diode as an "RF resistor" is probably a poor approximation to its true behavior. The inductively coupled circuit does not exhibit these characteristics.

3.4 Calibrations

3.4.1 K Space Trajectory

Like echo planar imaging (EPI), stochastic NMR requires sampling in the presence of a time varying gradient. As a result, phase shifts and delays in the gradient waveform affect the actual k-space location where sampling occurs. Ignoring this effect will have deleterious effects on image quality, as the the Fourier transform of the object density will be distorted, perhaps irrecoverably, if image reconstruction is performed using incorrect **k** locations. If, however, the **k** location is known throughout the imaging sequence, the **k** positions can be corrected during reconstruction, or precompensated during gradient generation.

The k-trajectory calibration is performed using the method described by Takahashi [29]. A large, well shimmed, spherical, undoped water sample on resonance is placed in the center of the magnet. The sample is excited by a 90° pulse, and a "self encode" gradient, G_{se} is applied along one axis, followed by the component along the phase encode direction of the gradient being mapped, G_m (see Figure 3.12). Whenever the time integral of the gradient waveform starting from the RF pulse is zero, there will be an echo; therefore if the integral under the self encode pulse is known, the value of the integral of G_m at the echo time, which is the k-space position, is known. This is performed over a range of phase encode values, and for each time point in G_m the value of G_{se} producing the maximum echo amplitude is determined. The result is a measurement of the k-space position as a function of time.

The gradient trajectory is mapped for each of the three Cartesian axes. Because a periodic waveform is employed for the experiments described, only a representative sample of the gradient waveform along each axis is required. This technique assumes that the gradient waveforms along the three axes are independent, and that the eddy currents reach steady state well before sampling ends, so that the later parts of the waveform represent the true steady state k trajectory.

Measured phase shifts for sine wave gradients at 500Hz were approximately 4.7, 3.5, 9.2 degrees along x, y and z respectively. These phase shifts are enough to

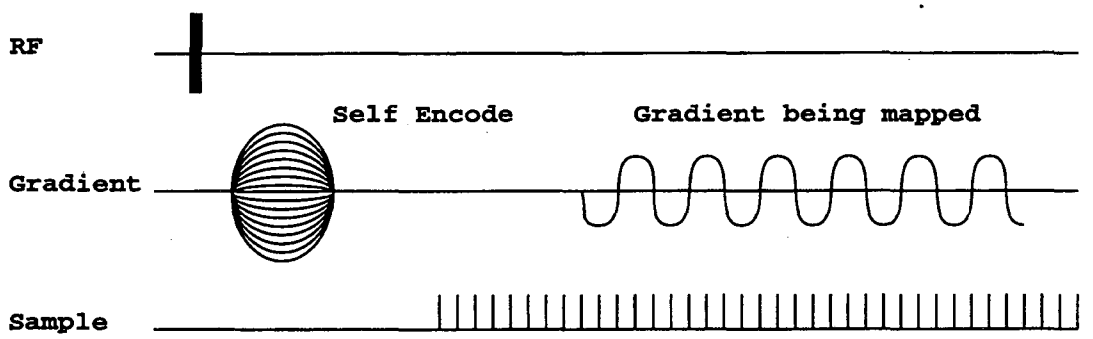

Figure 3.12: k trajectory calibration pulse sequence

cause significant image distortion if left uncorrected. Therefore, all gradient waveforms used in the experiments described were precompensated during generation to eliminate phase shifts and distortions in the k-space trajectory.

3.4.2 Receiver Filter Delay

Another important calibration that must be performed is measurement and compensation for the group delay of the receiver filter. Most NMR spectrometers have an analog filter in the RF receiver used to limit the analog signal bandwidth going to the digitizer. This filter will introduce a delay in the analog signal being generated. This delay time depends on the type of filter used, but can be quite significant, often on the order of a sampling period. In conventional imaging this filter delay will cause a phase roll in the received FID, which can be compensated for after the fact; however, in stochastic imaging, because the receiver gate is being switched throughout the experiment, the result of ignoring the filter delay may be that the digitizer samples during a time when there is little or no signal at its input.

The solution to this is to modify the pulse program to delay the actual sampling time of the ADC from the end of the experiment cycle shown in Figure 3.2 by the amount of the receiver delay. This receiver delay can be measured directly by watching the analog input to the digitizer while triggering off of the RF gate pulse and determining when the analog signal response is at a maximum. However, in many cases the receiver filter group delay is characterized by the manufacturer and included

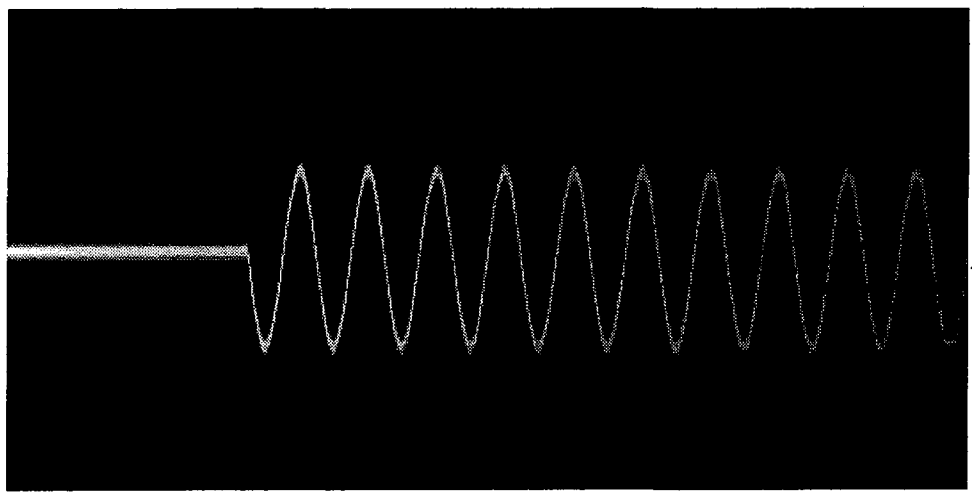

(a) Sinusoidal gradient

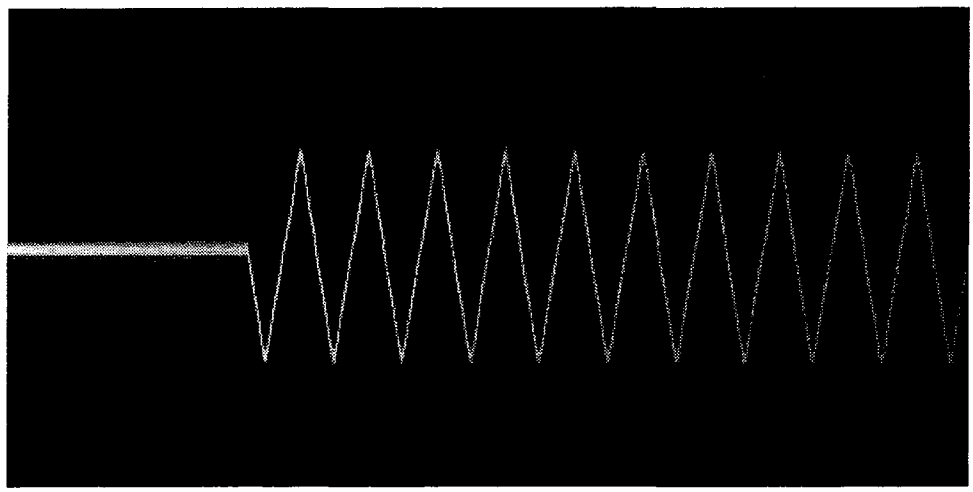

(b) Three component gradient

Figure 3.13: One dimensional k trajectory maps of gradients generated by Equation 3.1. (a) $i_{max} = 1$ *(Sinusoidal gradient) (b)* $i_{max} = 3$

in the data sheet of the filter. Since this was the case with our spectrometer, the stochastic imaging experiment setup macro used for these experiments determines the filter delay after the experiment parameters have been set and automatically generates a pulse sequence which delays the sampling time to compensate.

Chapter 4

Reconstruction Algorithm

4.1 Introduction

This chapter describes the algorithm used to reconstruct spectroscopic images from stochastic data sets. Previous image reconstruction methods have used a weighted crosscorrelation of the excitation sequence and the received signal, modified by a phase demodulation kernel. Images are reconstructed voxel by voxel, with a different phase demodulation kernel for each voxel[17]. This is computationally expensive, and the derivation of the phase demodulation kernel imposes rigid constraints on the gradient encoding waveforms that can be used if practical reconstruction times are to be obtained, even for encoding in two spatial dimensions.

Section 4.2 provides a mathematical description of the stochastic experiment, and defines some concepts which will be used throughout the Chapter. The phase demodulation reconstruction method is described briefly in Section 4.3. Section 4.4 describes and derives a new crosscorrelation reconstruction method based on interpolation onto a grid in k,t-space, or "Fourier gridding", which provides a dramatic reduction in reconstruction time and removes most of the mathematical constraints that limit the choice of image encoding process, providing much greater flexibility in experiment design. The point spread function and noise characteristics of the reconstruction are analyzed in Sections 4.5 and 4.6 respectively. Section 4.7 discusses implementation issues.

4.2 Mathematical Description of the Experiment

A typical stochastic NMR experiment is diagrammed in Figure 3.1. The RF excitation employed is a train of low flip angle (on the order of one degree) broadband RF pulses of random phase and/or amplitude, as described in Section 3.2.2. The bandwidth of each individual RF pulse is great enough to excite all of the spins in the system even in the presence of gradients. After each RF pulse, one data sample is taken, and the next gradient value is asserted. There is no spatial preselection in the stochastic experiment; the entire region within the RF probe is excited. We impose the condition that the RF excitation be small enough that the response of the transverse magnetization to the excitation sequence is linear. Then, following the analysis of Roos and Wong[17], the transverse magnetization $m_n(\mathbf{x}, \sigma)$ at time n arising from a stochastic RF excitation sequence in the presence of applied magnetic field gradients can be shown to be:

$$m_n(\mathbf{x}, \sigma) = \sum_{q=0}^{\infty} \rho(\mathbf{x}, \sigma) e^{-i(q+1)\sigma T_R} e^{-(q+1)T_R/T_2} e^{-i\mathbf{x}\cdot\mathbf{k}_{n,q}} s_{n-q}, \tag{4.1}$$

where $\rho(\mathbf{x}, \sigma)$ is the spin density as a function of space and chemical shift, s_n is the excitation sequence, T_R is the sampling interval, and

$$\mathbf{k}_{n,q} \equiv -\gamma \sum_{p=n-q}^{n} \mathbf{G}_p T_R \tag{4.2}$$

indicates the position in k-space (the spatial frequency domain[30]) at time nT_R for the magnetization created at time $(n-q)T_R$, after evolving in the presence of a magnetic field gradient $\mathbf{G}_p$. γ is the gyromagnetic ratio of the nucleus being imaged. Throughout this derivation, constants of proportionality have been suppressed for clarity.

Integrating the magnetization over space and chemical shift yields the received signal y_n:

$$y_n = \int_\sigma \int_{\mathbf{x}} \rho(\mathbf{x}, \sigma) \sum_{q=0}^{\infty} e^{-i(q+1)\sigma T_R} e^{-(q+1)T_R/T_2} e^{-i\mathbf{x}\cdot\mathbf{k}_{n,q}} s_{n-q} \, d\mathbf{x} \, d\sigma. \tag{4.3}$$

In addition to linearity, we assume that the longitudinal magnetization is in steady state. These assumptions are made for convenience. A Bloch equation analysis of the stochastic experiment after a brief interval for equilibration leads to

essentially the same result, with an added factor that accounts for T_1 saturation. The added factor indicates that the signal will be maximized by choosing the RMS flip angle to be the Ernst angle, as in a conventional experiment (note that in the conventional expression, α_{Ernst} is assumed to be constant rather than an RMS value).

$$\alpha_{\text{Ernst}} = \cos^{-1}\left(e^{-T_R/T_1}\right). \tag{4.4}$$

4.3 Phase Demodulation Reconstruction

One way to estimate the spin density from the received signal is to perform a weighted crosscorrelation of the received signal with the excitation sequence, using a phase demodulation kernel $e^{i\mathbf{x}\cdot\mathbf{k}_{n,q}}w(n,q)$:

$$\hat{\rho}(\mathbf{x},q) = \sum_{n=1}^{N} y_n e^{i\mathbf{x}\cdot\mathbf{k}_{n,q}} w(n,q) s^*_{n-q}. \tag{4.5}$$

where N is the total number of samples in the received signal y_n, and $w(n,q)$ is a weighting function.

The demodulation kernel selects for magnetization having a particular phase evolution due to the gradients - i.e. the magnetization arising from a particular region in space. The weighting function $w(n,q)$ compensates for the uneven sampling of k-space by the scanning trajectory. Each correlation lag q corresponds to a time point in the FID of the NMR signal (the Fourier transform of the spectrum). This method requires that we calculate a weighted summation of the entire received signal with the excitation sequence for each lag q for every voxel in the image, so the reconstruction scales in time as the product of the received signal length and the number of voxels reconstructed (and as $n\log(n)$ with the number of correlation lags reconstructed). This scaling makes three dimensional image reconstruction computationally impractical. Reconstruction speed is improved if crosscorrelations calculated with fast Fourier transforms can be used, but this requires that the weighting function $w(n,q)$ must be expressible as a product of two functions $w_1(n)w_2(n-q)$.

4.4 Fourier Gridding Reconstruction

We denote the spin density as a function of $\mathbf{x}$ and σ as ρ. We will use ρ_k to denote the Fourier transform of ρ with respect to $\mathbf{x}$; ρ_q to denote the Fourier transform of ρ with respect to σ; and ρ_K to denote the Fourier transform of ρ with respect to $\mathbf{x}$ and σ. Therefore:

$$\rho(\mathbf{x},\sigma) = \int_{\mathbf{k}} \rho_k(\mathbf{k},\sigma) e^{i\,\mathbf{k}\cdot\mathbf{x}}\, d\mathbf{k} = \int_{\eta}\int_{\mathbf{k}} \rho_K(\mathbf{k},\eta) e^{i\eta\sigma} e^{i\,\mathbf{k}\cdot\mathbf{x}}\, d\mathbf{k}\, d\eta = \int_{\eta} \rho_q(\mathbf{x},\eta) e^{i\eta\sigma}\, d\eta \tag{4.6}$$

which allows us to rewrite Equation 4.3 as:

$$y_n = \int_{\sigma}\int_{\mathbf{x}}\int_{\mathbf{k}} \rho_k(\mathbf{k},\sigma) \sum_{q=0}^{\infty} e^{-i\,(q+1)\sigma T_R} e^{-(q+1)T_R/T_2} e^{-i\,(\mathbf{k}_{n,q}-\mathbf{k})\cdot\mathbf{x}} s_{n-q}\, d\mathbf{k}\, d\mathbf{x}\, d\sigma. \tag{4.7}$$

Rearranging the integration and summation, and using the definition of ρ_K, we see that

$$y_n = \sum_{q=0}^{\infty} e^{-(q+1)T_R/T_2} s_{n-q} \int_{\mathbf{k}} \rho_K(\mathbf{k},(q+1)T_R)\delta^3(\mathbf{k}-\mathbf{k}_{n,q})\, d\mathbf{k} \tag{4.8}$$

$$= \sum_{q=0}^{\infty} e^{-(q+1)T_R/T_2} s_{n-q} \rho_K(\mathbf{k}_{n,q},(q+1)T_R). \tag{4.9}$$

For convenience in notation, from now on we will use the convention that $\rho_K(\mathbf{k},q)$ represents the function $\rho_K(\mathbf{k},\eta)$ evaluated at $\eta = (q+1)T_R$.

The received signal y_n is now cast in a new form, as a function of the excitation sequence and the Fourier transform of the spin density (or k-space representation) of the object. This leads to a new interpretation of the stochastic NMR signal. Each time sample of the received signal has contributions from an ensemble of signals arising from every RF pulse in the past. Each signal in the ensemble samples the four dimensional k-space representation of the object density, and each is weighted by the RF pulse that created it. The signal composition and reconstruction in the absence of gradients is illustrated in Figure 4.1. Using this fact and knowing that the excitation sequence is a white noise sequence, we can estimate the spin density $\hat{\rho}_K(\mathbf{k},q)$ in a straightforward manner.

To reconstruct an image from the input signal y_n we first form a rectilinear four dimensional array (three k-space axes and one lag axis) of $\rho_K(\mathbf{k},\eta)$ which is an

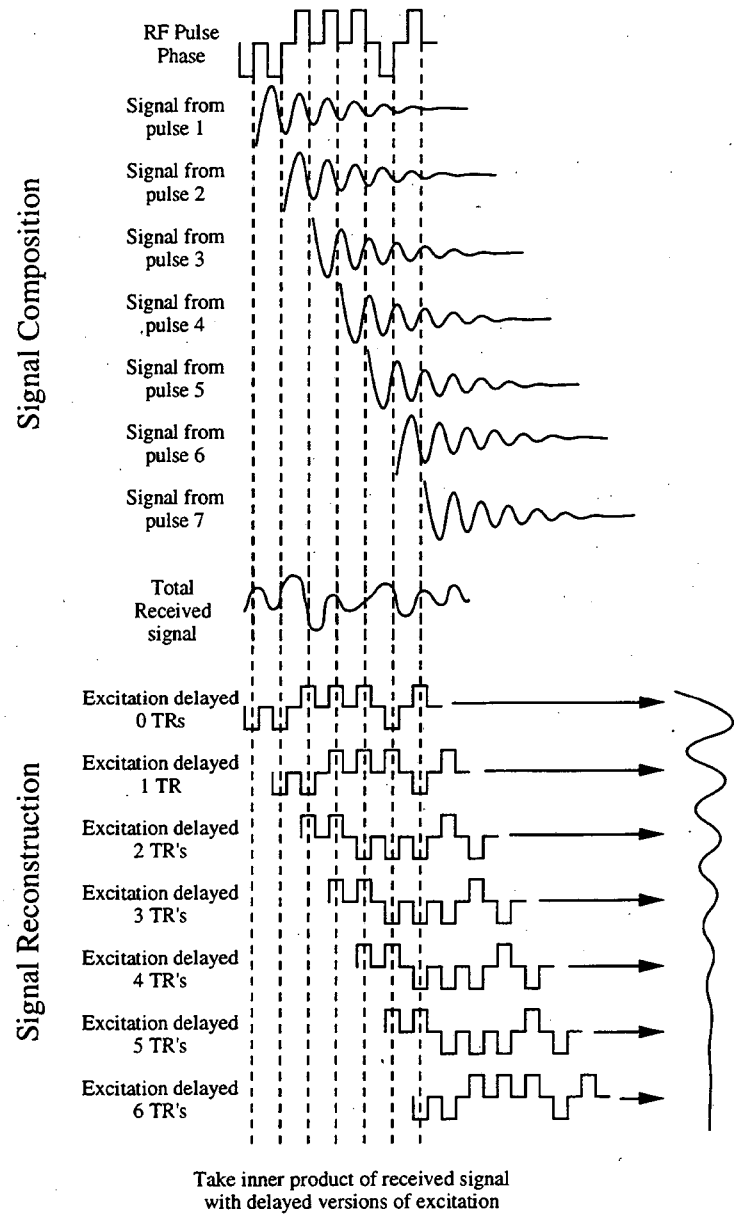

Figure 4.1: Composition and reconstruction of the stochastic signal in the absence of gradients. The top of the diagram shows how the stochastic excitation produces a signal which is a composite of many individual FIDs. Each time point of the received signal, y_n, is a summation of signals arising from magnetization created in the past. The magnetization from the RF pulse s_n has evolved through one TR, the magnetization from pulse s_{n-1} has evolved two periods, etc. Each component of the sum has initial phase. The lower section of the diagram on the shows how crosscorrelating the signal produced in this manner regenerates the FID of the spin system.

estimate of the 4D Fourier transform of the object $\rho(\mathbf{x}, \sigma)$. To do this we will make use of the autocorrelation property of the white, random excitation, $\left\langle s_q s_{q'}^* \right\rangle = \delta_{q,q'}$ to separate the signal components that are summed together in each data sample.

By way of example, consider the stochastic NMR signal in the absence of gradients,

$$y_n = \sum_{q=0}^{\infty} e^{-(q+1)T_R(\sigma+1/T_2)} s_{n-q}. \tag{4.10}$$

The FID of a spin system can be estimated from a stochastic experiment by crosscorrelating the received signal with the excitation sequence.

$$\hat{\rho}(q) = \sum_{n=1}^{N} y_n s_{n-q}^* \tag{4.11}$$

In this case, $\hat{\rho}(q)$ is the estimate of the spin density integrated over all space (the inverse Fourier transform of the spectrum). Each received signal point has contributions from many RF excitations in the past; however, the crosscorrelation operation with the conjugate of the excitation sequence with a given delay acts as a matched filter which selects only for magnetization from the appropriate time lag in the past.

In the presence of gradients $y_n s_{n-q}^*$ is now also a function of $\mathbf{k}$. While samples are acquired with uniform spacing in time, the $\mathbf{k}$ location of the samples in $\mathbf{k}, q$ space are not the same for each q (see Figure 4.3), and are not uniformly spaced in $\mathbf{k}$. To develop an extension of Equation 4.11, one approach is to interpolate these samples in the $\mathbf{k}$ dimension onto a rectilinear grid before computing the time crosscorrelation with the excitation. This crosscorrelation then yields one FID per k-space point. A four dimensional FFT then gives $\rho(\mathbf{x}, \sigma)$.

The interpolation problem in the $\mathbf{k}$ domain is the same as that encountered in imaging with arbitrary $\mathbf{k}$ trajectories and deterministic excitation[18]. The reconstruction operation is derived in two steps: first, a continuous function of $\mathbf{k}$ is constructed for each lag value q from the discrete data samples y_n; then an estimate analogous to Equation 4.11 is defined.

To begin, multiply the scaled datum $y_n s_{n-q}^*$ by a sampling function $\delta^3(\mathbf{k} - \mathbf{k}_{n,q})$ to form the function:

$$M_S^{n,q}(\mathbf{k}) \equiv y_n s_{n-q}^* \delta^3(\mathbf{k} - \mathbf{k}_{n,q}). \tag{4.12}$$

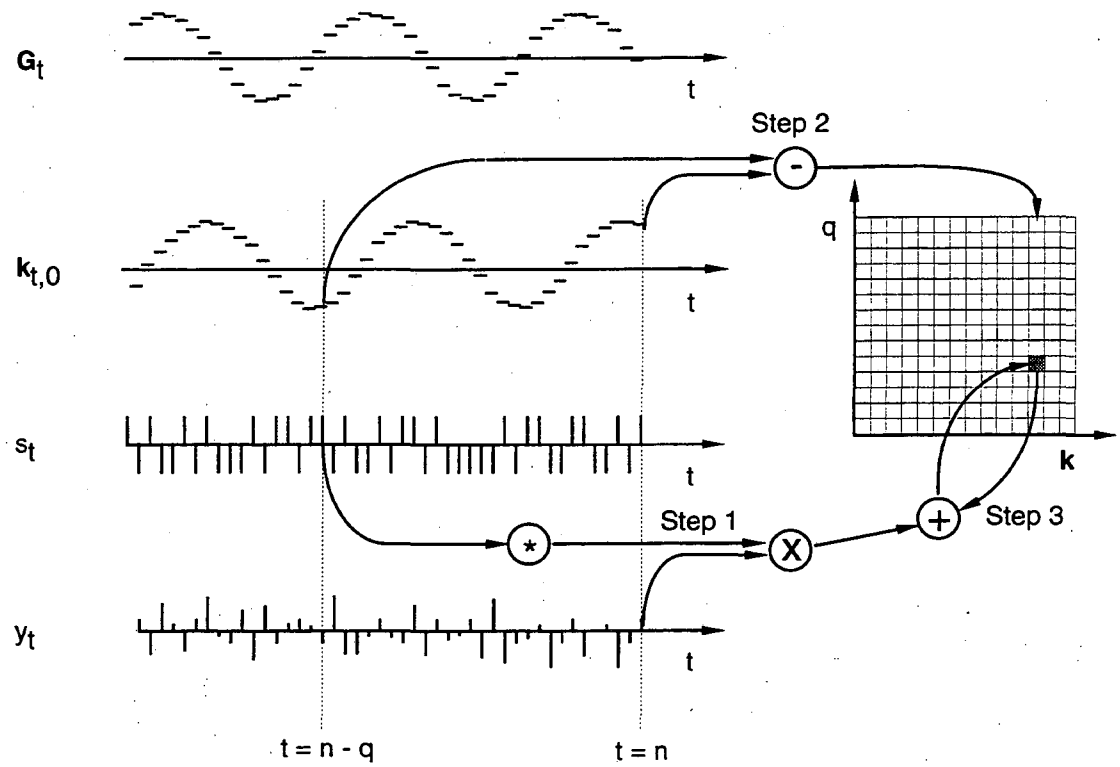

Figure 4.2: Diagram of the gridding reconstruction procedure. For all n: Step 1: multiply the received data point y_n with the conjugate of the excitation s_{n-q}. Step 2: subtract the initial k-space position from the current k-space position to determine the k location of the detected magnetization in the output grid. Step 3: add the product from step one into the grid position determined by the k location found in step two and the lag value q.

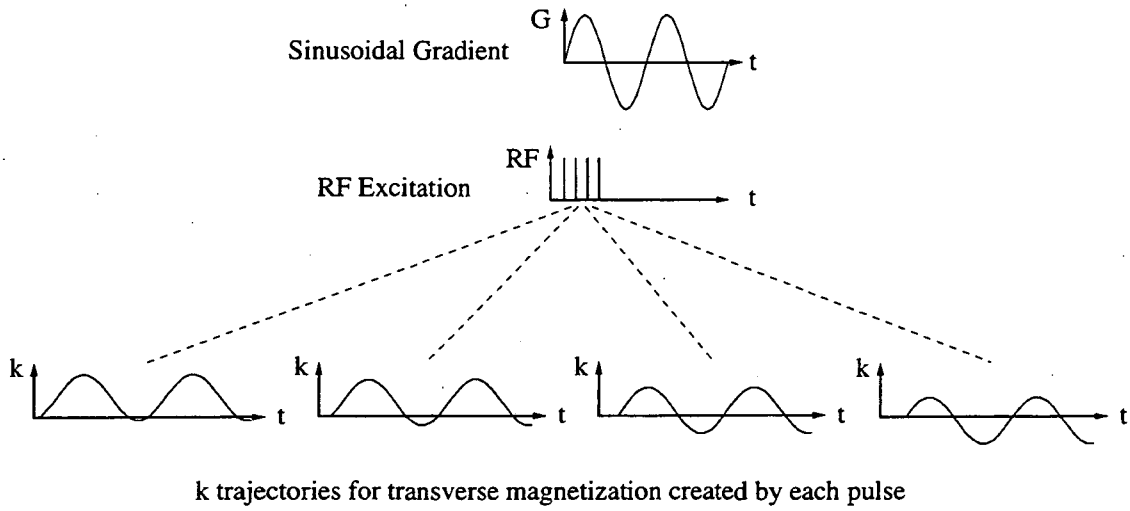

Figure 4.3: The magnetization from each RF pulse follows its own k-space trajectory in the presence of gradients. Therefore each point in the received signal is made up of signals that in general have different initial phase and k-space position.

Convolve $M_S^{n,q}(\mathbf{k})$ with a kernel C to produce a continuous function which will be resampled onto a Cartesian grid. Following Jackson, *et al*[19], we find that the convolved, weighted, discrete function (prior to resampling) is represented by:

$$M_{SWC}^{n,q}(\mathbf{k}) = \frac{M_S^{n,q}(\mathbf{k})}{W^q(\mathbf{k})} * C(\mathbf{k}). \tag{4.13}$$

where $C(\mathbf{k})$ is the convolution kernel used in the sampling operation, and

$$W^q(\mathbf{k}) = \sum_{n=1}^{N} \int_{\mathbf{k}'} \delta^3(\mathbf{k} - \mathbf{k}_{n,q}) C(\mathbf{k} - \mathbf{k}')\, d\mathbf{k}' \sum_{n=1}^{N} C(\mathbf{k} - \mathbf{k}_{n,q}) \tag{4.14}$$

is the "area density function" that accounts for the uneven sampling density.

This interpolated version of $y_n s_{n-q}^*$ can be summed over n to give an estimate of $\hat{\rho}_K(\mathbf{k}, q)$ similar to Equation 4.11.

$$\begin{aligned}
\hat{\rho}_K(\mathbf{k}, q) &= \sum_{n=1}^{N} \frac{M_S^{n,q}(\mathbf{k})}{W^q(\mathbf{k})} * C(\mathbf{k}), && (4.15)\\
&= \sum_{n=1}^{N} \int_{\mathbf{k}'} \frac{y_n s_{n-q}^* \delta^3(\mathbf{k}' - \mathbf{k}_{n,q}) C(\mathbf{k} - \mathbf{k}')}{W^q(\mathbf{k}')}\, d\mathbf{k}' && (4.16)\\
&= \sum_{n=1}^{N} \frac{y_n s_{n-q}^* C(\mathbf{k} - \mathbf{k}_{n,q})}{W^q(\mathbf{k}_{n,q})}. && (4.17)
\end{aligned}$$

This function can subsequently be evaluated at any $\mathbf{k}$ position; in particular, we can evaluate it on the points of a Cartesian grid. We will use $\mathbf{k_{III}}$ to represent the continuous variable $\mathbf{k}$ evaluated on Cartesian grid points.

Using y_n from Equation 4.8, we find:

$$\hat{\rho}_K(\mathbf{k}, q) = \sum_{n=1}^{N} \sum_{q'=0}^{\infty} e^{-(q'+1)T_R/T_2} s^*_{n-q} s_{n-q'} \frac{\rho_K(\mathbf{k}_{n,q'}, q')C(\mathbf{k} - \mathbf{k}_{n,q})}{W^q(\mathbf{k}_{n,q})}. \quad (4.18)$$

The excitation sequence s_n has a delta function autocorrelation, so that

$$\left\langle s^*_{n-q} s_{n-q'} \right\rangle = \delta_{q,q'}; \quad (4.19)$$

Employing this relation, we see that

$$\langle \hat{\rho}_K(\mathbf{k}, q) \rangle = e^{-(q+1)T_R/T_2} \sum_{n=1}^{N} \frac{\rho_K(\mathbf{k}_{n,q}, q)C(\mathbf{k} - \mathbf{k}_{n,q})}{W^q(\mathbf{k}_{n,q})}. \quad (4.20)$$

Evaluating Equation 4.20 at a grid point $\mathbf{k_{III}}$, we find that

$$\langle \hat{\rho}_K(\mathbf{k_{III}}, q) \rangle = e^{-(q+1)T_R/T_2} \sum_{n=1}^{N} \frac{\rho_K(\mathbf{k}_{n,q}, q)C(\mathbf{k_{III}} - \mathbf{k}_{n,q})}{W^q(\mathbf{k}_{n,q})}. \quad (4.21)$$

Thus the expectation of the spin density estimate is a smoothed version of the true density function with weighting dependent on the sampling density.

The use of the convolution function C in the *mtxk* domain introduces a small, correctable "doming"in the spatial domain[19, 18]. This can be approximately eliminated by dividing the spin density estimate ρ by a periodic function c which is equal to the inverse Fourier transform of C over the reconstructed field of view in $\mathbf{x}$.

The function $W^q(\mathbf{k})$ can be computed prior to or during the reconstruction if the sampling density is known analytically; however, for sufficiently smooth sampling densities it is well approximated by $W^q(\mathbf{k_{III}})$. This function may be estimated at grid points in parallel with the reconstruction operation, using:

$$W^q(\mathbf{k_{III}}) = \sum_{n=1}^{N} C(\mathbf{k_{III}} - \mathbf{k}_{n,q}). \quad (4.22)$$

In this case the weighting function can be pulled out of the summation in Equation 4.21, and the weighting applied after the gridding operation. Using a gridded version of the weighting function can speed reconstruction, but it introduces some artifacts into the image and should be avoided if possible.

4.5 Point Spread Function Analysis

The point spread function is determined by the initial sampling function and the characteristics of the correction weighting function W^q. Any useful gradient trajectory used for a stochastic experiment will be designed to cover some region of k-space densely enough so as not to have major discontinuities or holes. The point-spread function will be primarily determined by the shape of this region. Over the sampled region, after correction by W^q, the sampling density will be uniform. In addition, the point spread function can be tailored for different characteristics by multiplying the k-space spin density estimate by some filter function to enhance or deemphasize particular spatial frequencies. We will consider the case of uniform density over the sampled region for two useful three dimensional scanning trajectories.

We represent the shape of the uniformly sampled region with a window function $B^q(\mathbf{k})$:

$$B_q(\mathbf{k}) = \begin{cases} 1 & \forall\ \mathbf{k} \text{ inside sampled region} \\ 0 & \text{everywhere else} \end{cases} \tag{4.23}$$

If we perform the summation in Equation 4.15, and perform the inverse Fourier transform with respect to $\mathbf{x}$, we see that the estimate of the spin density is given by the relation

$$\hat{\rho}_q(\mathbf{x}, q) = \left\{ e^{-(q+1)T_R/T_2} \rho(\mathbf{x}, q) * b^q \right\} c \tag{4.24}$$

where b^q, and c, are the inverse Fourier transforms of B^q, and C, respectively.

After deconvolving or "undoming" by dividing out c:

$$\hat{\rho}_q(\mathbf{x}, q) = e^{-(q+1)T_R/T_2} \rho(\mathbf{x}, q) * b^q, \tag{4.25}$$

so the point spread function becomes

$$h^q(\mathbf{x}) = e^{-(q+1)T_R/T_2} b^q(\mathbf{x}). \tag{4.26}$$

Both of the scanning trajectories considered here are based on periodic oscillating gradients which are sums of sinusoids. Although the gradient trajectory is arbitrary, periodic gradients are easily generated and have the advantage of producing steady state eddy currents, which result only in a correctable phase shift of the gradient field (see Section 3.2.3).

4.5.1 One Dimensional Periodic Oscillating Gradients

We start by considering a one dimensional oscillating gradient which is a truncated Fourier series of a square wave:

$$\mathbf{G}_p = \sum_{i=1}^{i_{max}} \frac{G\cos\left(2(2i-1)\pi f_o p T_R\right)}{(-1)^i(2i-1)}. \tag{4.27}$$

We can rewrite the definition of $\mathbf{k}_{n,q}$ in Equation 4.2 as

$$\mathbf{k}_{n,q} = -\gamma T_R \left[\sum_{p=0}^{n} \mathbf{G}_p - \sum_{p=0}^{n-q-1} \mathbf{G}_p\right]. \tag{4.28}$$

Substituting the definition of $\mathbf{G}_p$ into Equation 4.28,

$$\mathbf{k}_{n,q} = \sum_{i=1}^{i_{max}} \frac{-\gamma G T_R}{(-1)^i(2i-1)} \left[\sum_{p=0}^{n} \cos\left(2(2i-1)\pi f_o p T_R\right) - \sum_{p=0}^{n-q-1} \cos\left(2(2i-1)\pi f_o p T_R\right)\right] \tag{4.29}$$

and using the fact that

$$\sum_{k=0}^{n-1} \cos(x+ky) = \cos(x + \frac{(n-1)y}{2}) \sin(\frac{ny}{2}) \csc(\frac{y}{2}), \tag{4.30}$$

we find that

$$\mathbf{k}_{n,q} = \gamma G T_R \sum_{i=1}^{i_{max}} \frac{\sin\left((2i-1)\pi f_o T_R(2n+1)\right) - \sin\left((2i-1)\pi f_o T_R(2n-2q-1)\right)}{(-1)^i 2(2i-1)\sin\left((2i-1)\pi f_o T_R\right)}. \tag{4.31}$$

For any given value of q, the phase between the lowest order ($i=1$) sine functions in the difference term is $\phi = 2(q+1)\pi f_o T_R$, which gives a maximum value for the difference term of:

$$2\sin(\phi/2) = 2\sin\left((q+1)\pi f_o T_R\right), \tag{4.32}$$

yielding a maximum k value

$$k_{max}^q = \gamma G T_R \sum_{i=1}^{i_{max}} \left| \frac{\sin\left((q+1)(2i-1)\pi f_o T_R\right)}{(2i-1)\sin\left((2i-1)\pi f_o T_R\right)} \right|. \tag{4.33}$$

This function is maximized when the two sine functions are out of phase by an odd integral multiple of π. This gives an overall maximum k value

$$k_{max} = \gamma G T_R \sum_{i=1}^{i_{max}} \frac{1}{(2i-1)\sin\left((2i-1)\pi f_o T_R\right)}, \tag{4.34}$$

$$\left(q = \left[\frac{p}{2 f_o T_R} - 1\right], p = 1, 3, 5, \ldots\right).$$

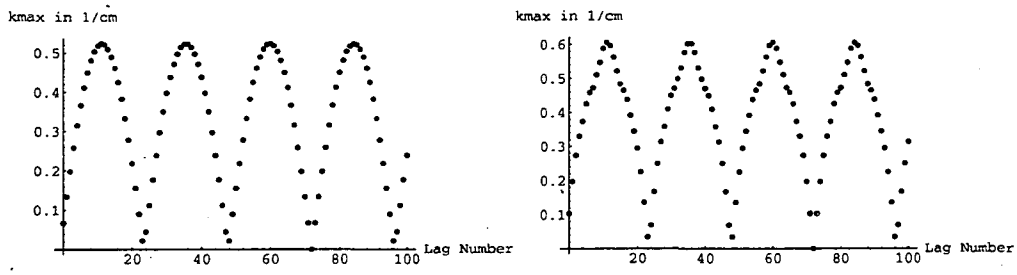

(a) One dimensional sinusoidal gradient (i = 1)

(b) One dimensional truncated square wave gradient (i = 3)

Figure 4.4: Plot of k_{max} as a function of lag q for a typical stochastic ^{23}Na experiment. Gradients are generated using Equation 3.1, and k_{max} is calculated from Equation 4.33. γ is 11.24 MHz/T, G is 8 mT/m, T_R is 75μs and f_o is 548.00846Hz

We see from this analysis that the one dimensional truncated square wave gradient samples a line in k-space which extends from $-k_{max}$ to k_{max}, and that the magnitude of k_{max} depends on the correlation lag q. The window function is therefore:

$$B_q(k) = \Pi(k^q_{max}) = \begin{cases} 1 & |k| \leq k^q_{max} \\ 0 & |k| > k^q_{max} \end{cases} \tag{4.35}$$

yielding the one dimensional sinc function point spread function

$$h^q(\mathbf{x}) = e^{-(q+1)T_R/T_2} \frac{\sin(2\pi k^q_{max} x)}{2\pi k^q_{max} x}. \tag{4.36}$$

4.5.2 Rotating Oscillating Gradients

For fast three dimensional spectroscopic imaging, one possible gradient waveform is based on the trajectory proposed by Norton[31], an oscillating gradient which rotates with a direction vector $\mathbf{\Theta}_n$ that sweeps out a spiral on the surface of a sphere. This gradient waveform samples a sphere in k-space for each time lag. The oscillation frequency of the gradients f_o must be higher than the chemical shift bandwidth of the object under study to prevent aliasing in the spectral dimension.

Norton's trajectory can be modified with the proper choice of the time dependence of the θ and ϕ values of $\mathbf{\Theta}_n$ so that k-space will be sampled uniformly over all

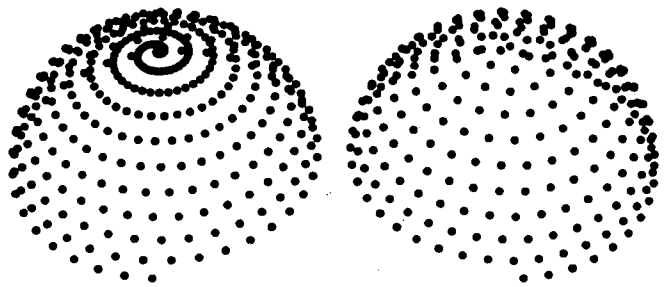

Figure 4.5: Sample point distributions obtained with Norton's original sampling (left) and with the new technique (right). The distribution on the left requires 54% more points to achieve similar density near the equator.

solid angles, leaving only a radial dependence in the k-space sampling density[32], which is corrected by the weighting function W^q. This isotropic solid angle rotating, oscillating gradient is given parametrically by:

$$\begin{aligned}
\mathbf{G_n} &= G\Theta_n \cos(2\pi f_o n T_R) = G(x_n\hat{\mathbf{x}} + y_n\hat{\mathbf{y}} + z_n\hat{\mathbf{z}})\cos(2\pi f_o n T_R) \\
z_n &= (2n - N - 1)/N \\
x_n &= \cos(\sqrt{f_0 \mathrm{T_R} N\pi}\sin^{-1} z_n)\sqrt{1 - z_n^2} \\
y_n &= \sin(\sqrt{f_0 \mathrm{T_R} N\pi}\sin^{-1} z_n)\sqrt{1 - z_n^2}
\end{aligned} \tag{4.37}$$

where N is the total number of points in the experiment, and $\mathrm{T_R}$ is the sampling period.

A useful approximation is that the gradient oscillation frequency f_o is much higher than the θ and ϕ rotation frequencies of the direction vector, so that the direction vector is essentially constant over one period of the oscillation ($\Theta_n = \Theta$); experimental parameters can be selected so that this approximation holds extremely well. So assuming that $\mathbf{G}_p = G\Theta\cos(2\pi f_o p T_R)$, and $\pi f_o T_R < \pi/2$, we can assume that the radial and angular parts of the k-space density are separable. Since all solid angles are sampled isotropically by this sampling trajectory, the density correction will be spherically symmetric, and after it is applied, the sampled region of k-space will be a uniform sphere.

If the sampling density correction is applied independently for each lag q, the resulting estimate of the spin density is convolved with a three dimensional window

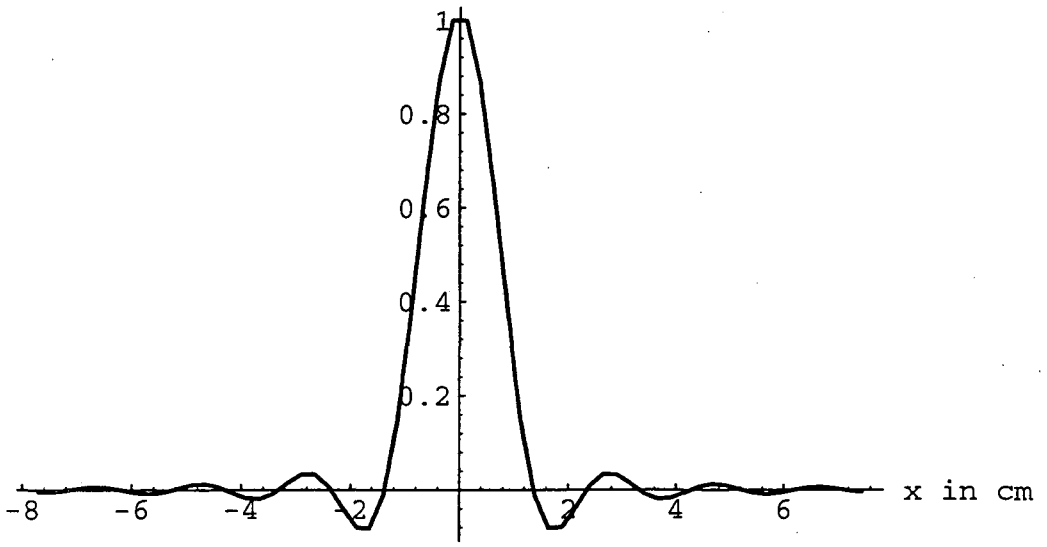

Figure 4.6: Calculated point spread function for a single lag using the rotating oscillating gradient. The line is a graph of equation 4.40 using the the gradient parameters for a typical ^{23}Na experiment described in Figure 4.4. The lag chosen is at the first maximum of $k_{max}(q)$, lag 11. $k_{max}(11) = 0.5237$.

function in the **k** domain:

$$B_q(\mathbf{k}) = \sqcap(k^q_{max}) = \begin{cases} 1 & |\mathbf{k}| \leq k^q_{max} \\ 0 & |\mathbf{k}| > k^q_{max} \end{cases} \tag{4.38}$$

which has the inverse Fourier transform

$$b_q(\mathbf{x}) = \frac{\sin\left(2\pi k^q_{max}\,|\mathbf{x}|\right) - 2\pi k^q_{max}\,|\mathbf{x}|\cos\left(2\pi k^q_{max}\,|\mathbf{x}|\right)}{2\pi^2 {k^q_{max}}^3\,|\mathbf{x}|^3}, \tag{4.39}$$

where k^q_{max} is a function of the lag q.

Combining Equations 4.26, 4.39, and 4.33, we find that for a rotating oscillating gradient

$$h^q(\mathbf{x}) = e^{-(q+1)T_R/T_2}\,\frac{\left[\sin\left(2\pi k^q_{max}\,|\mathbf{x}|\right) - 2\pi k^q_{max}\,|\mathbf{x}|\cos\left(2\pi k^q_{max}\,|\mathbf{x}|\right)\right]}{2\pi^2 {k^q_{max}}^3\,|\mathbf{x}|^3}. \tag{4.40}$$

This function is plotted in Figure 4.6. The halfwidth of this function is $0.795/k^q_{max}$, and the maximum sidelobe amplitude is 8.6%.

4.5.3 Incommensurate Frequency Oscillating Gradients

Another way to extend the oscillating gradient to sample a three dimensional region in k-space is to employ 3 independent oscillating gradients along the three Cartesian axes which have frequencies chosen such that:

$$if_1 \neq jf_2 \neq kf_3 \{i, j, k \in \text{ integers}\} \tag{4.41}$$

The resulting k-space trajectory is a three dimensional Lissajous pattern which never retraces its path. This trajectory (using sinusoidal oscillations) was initially proposed by Blümich[16]. For any given time lag, the resulting sampling density over all data points samples a right rectangular prism in k-space, each side having a k^q_{max} determined by Equation 4.33. If the gradients are close in frequency and amplitude, this region will be approximately cubic, giving the same spatial resolution along all axes. The result of sampling over a rectilinear region and performing a density correction gives the familiar three dimensional sinc shape to the point spread function:

$$b_q(\mathbf{x}) = \frac{\sin\left(2\pi k^q_{x(max)}x\right)\sin\left(2\pi k^q_{y(max)}y\right)\sin\left(2\pi k^q_{z(max)}z\right)}{8\pi^3 k^q_{x(max)}k^q_{y(max)}k^q_{z(max)}xyz}, \tag{4.42}$$

Combining Equations 4.26, 4.42, and 4.33, we find that for a incommensurate frequency truncated square wave gradients

$$h^q(\mathbf{x}) = e^{-(q+1)T_R/T_2}\frac{\sin\left(2\pi k^q_{x(max)}x\right)\sin\left(2\pi k^q_{y(max)}y\right)\sin\left(2\pi k^q_{z(max)}z\right)}{8\pi^3 k^q_{x(max)}k^q_{y(max)}k^q_{z(max)}xyz}. \tag{4.43}$$

This function is plotted in Figure 4.7. The halfwidth of this function is $0.603/k^q_{max}$, and the maximum sidelobe amplitude is 21.7%.

4.5.4 Point Spread Functions For Spectroscopic Imaging

The relations above describe the point spread function for a single reconstructed lag. If chemical shift information is not desired, then only one lag need be reconstructed; a value of the lag q can be chosen for maximum resolution by the criterion of Equation 4.35, with the value of p selected to give the desired T_2 contrast. If however, one wants to combine the information from various time lags, for example

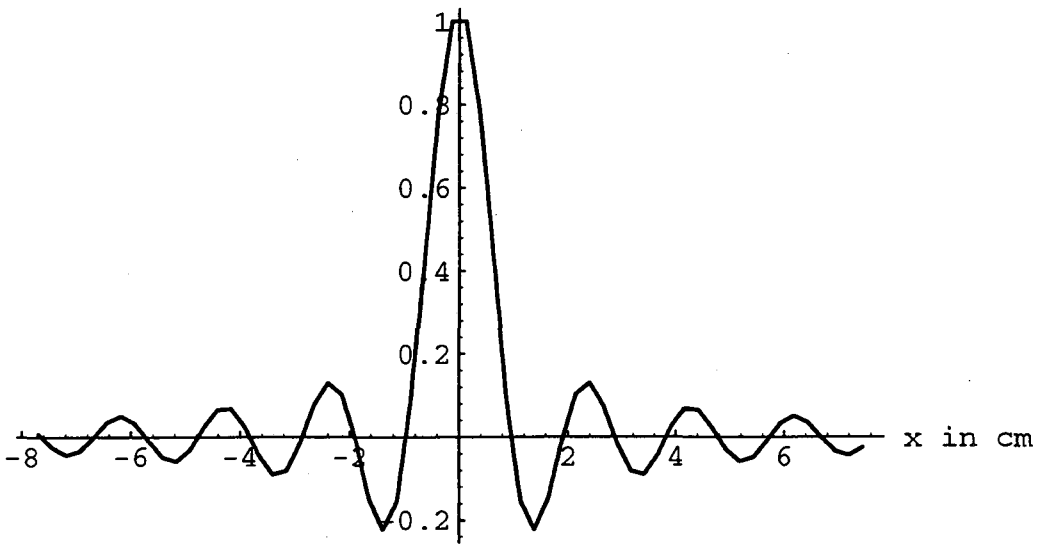

Figure 4.7: Calculated point spread function for a single lag, incommensurate frequency sinusoidal oscillating gradients. The line is a graph of equation 4.40 using the the gradient parameters for a typical ^{23}Na experiment described in Figure 4.4. The lag chosen is at the first maximum of $k_{max}(q)$, lag 11. $k_{max}(11) = 0.5237$.

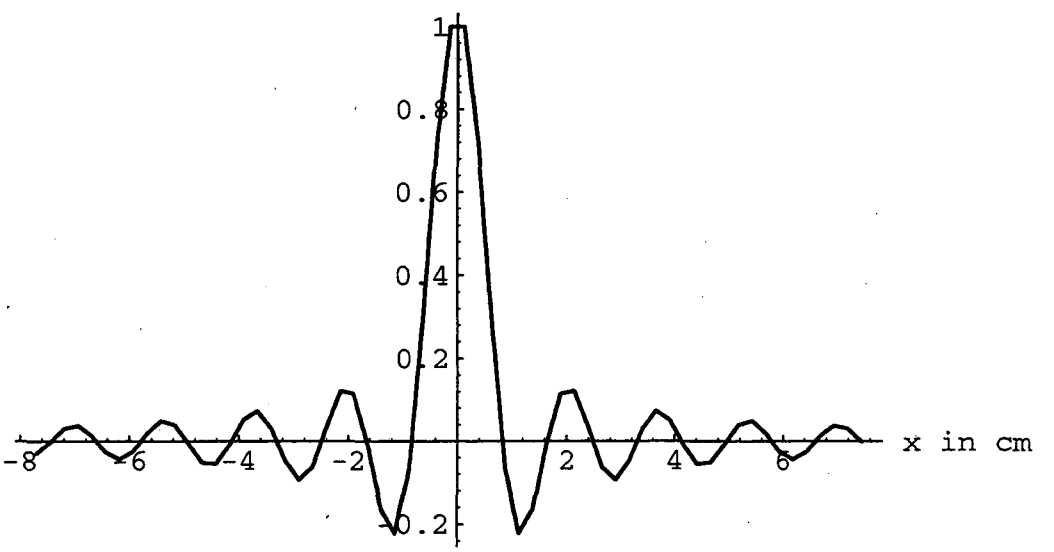

Figure 4.8: Calculated point spread function for a single lag, incommensurate frequency three component truncated square wave oscillating gradients. The line is a graph of equation 4.40 using the the gradient parameters for a typical ^{23}Na experiment described in Figure 4.4. The lag chosen is at the first maximum of $k_{max}(q)$, lag 11. $k_{max}(11) = 0.6053$.

to reconstruct a spectroscopic data set, we have to consider the k-space sampling density of the entire experiment, rather than for a single lag, to determine the point spread function. Note that any combination of lags can be used depending on the information desired in the reconstructed image; different combination weights can be used to tailor spatial sideband amplitudes and main lobe width in the point spread function, or to change the spectral lineshape. For simplicity, we will only consider the cases of a single lag and the case of all lags from 0 to some maximum lag $Q-1$.

The overall point spread function of the experiment is derived by considering how the k-space densities from each time lag combine to form an overall k-space sampling density for the chemical shift reconstruction. As can be seen from Equation 4.33, graphed in Figure 4.4, the maximum extent of the region sampled in k-space oscillates with twice the frequency of the oscillating gradient f_o. Because the oscillating gradient frequency is chosen to exceed the chemical shift bandwidth of the spin system, the fourth k-space dimension, the lag dimension, is sampled adequately to reconstruct a spectroscopic image with the full three dimensional resolution indicated by the k_{max} of the experiment given in Equation 4.35 if the proper weighting is applied.

As described above, the reconstructed spin density for each lag has been corrected to have unit sampling density over the volume in k-space that is sampled for that time lag. To generate a chemical shift image from the lag data, the data are Fourier transformed along the lag dimension (to go from the time to the frequency domain). Because the oscillating gradient makes the reconstructed spectral dimension periodic, only the baseband is considered. In performing the Fourier transform, data from all of the reconstructed lags contribute to each of the k-space positions in the object.

It is easily seen that the overall sampling density in k-space for the rotating, oscillating gradient experiment, summed over the lag dimension, is:

$$B(\mathbf{k}) = \sum_{q=0}^{Q-1} \sqcap(k_{max}^{q}) \tag{4.44}$$

The point spread function of the spectral data set reconstructed from lags 0 to $Q-1$

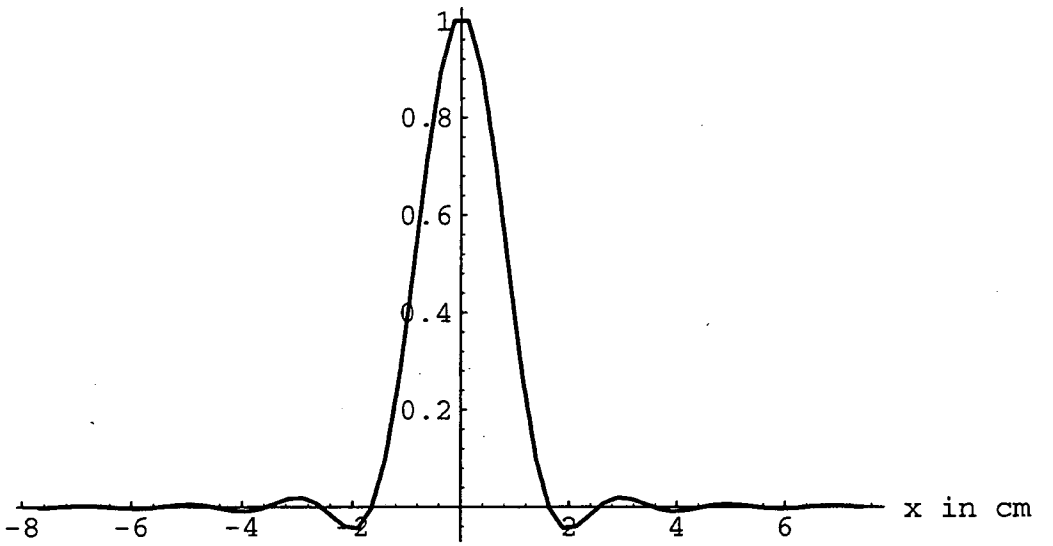

Figure 4.9: Overall point spread function for the rotating oscillating gradient experiment, 1024 lags; this is the spatial point spread function for the chemical shift reconstruction of the same experiment as in Figure 4.6; the overall k_{max} is 0.523813

is therefore:

$$h(\mathbf{x}) = \sum_{q=0}^{Q-1} e^{-(q+1)T_R/T_2} \frac{\left[\sin\left(2\pi k_{max}^q |\mathbf{x}|\right) - 2\pi k_{max}^q |\mathbf{x}| \cos\left(2\pi k_{max}^q |\mathbf{x}|\right)\right]}{2\pi^2 {k_{max}^q}^3 |\mathbf{x}|^3}. \quad (4.45)$$

The halfwidth of this function is $0.898/k_{max}$, and the maximum sidelobe amplitude is 4.5%. This function is plotted in Figure 4.9.

Similarly, the point spread function of the spectral data set reconstructed from lags 0 to $Q-1$ for the incommensurate gradient experiment is:

$$h(\mathbf{x}) = \sum_{q=0}^{Q} e^{-(q+1)T_R/T_2} \frac{\sin\left(2\pi k_{x(max)}^q x\right) \sin\left(2\pi k_{y(max)}^q y\right) \sin\left(2\pi k_{z(max)}^q z\right)}{8\pi^3 k_{x(max)}^q k_{y(max)}^q k_{z(max)}^q xyz}. \quad (4.46)$$

The shape of the function k_{max}^q depends on the number of components used in the gradient, so the shape of the overall point spread function will also vary with the number of gradient components. For the sinusoidal case, the halfwidth of this function is $0.674/k_{max}$, and the maximum sidelobe amplitude is 16.1%. This function is plotted in Figure 4.10. When three component TSW gradients are employed, the halfwidth of the point spread function is $0.613/k_{max}$, and the maximum sidelobe amplitude is 12.9%. This function is plotted in Figure 4.11.

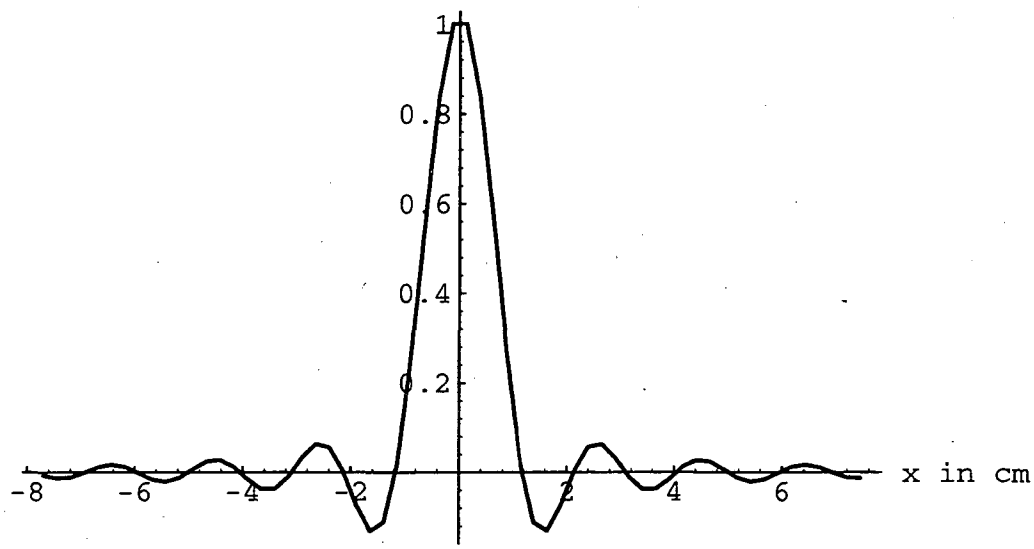

Figure 4.10: Overall point spread function for the incommensurate frequency sinusoidal oscillating gradient experiment, 1024 lags; this is the spatial point spread function for the chemical shift reconstruction of the same experiment as in figure 4.7; the overall k_{max} is 0.523813

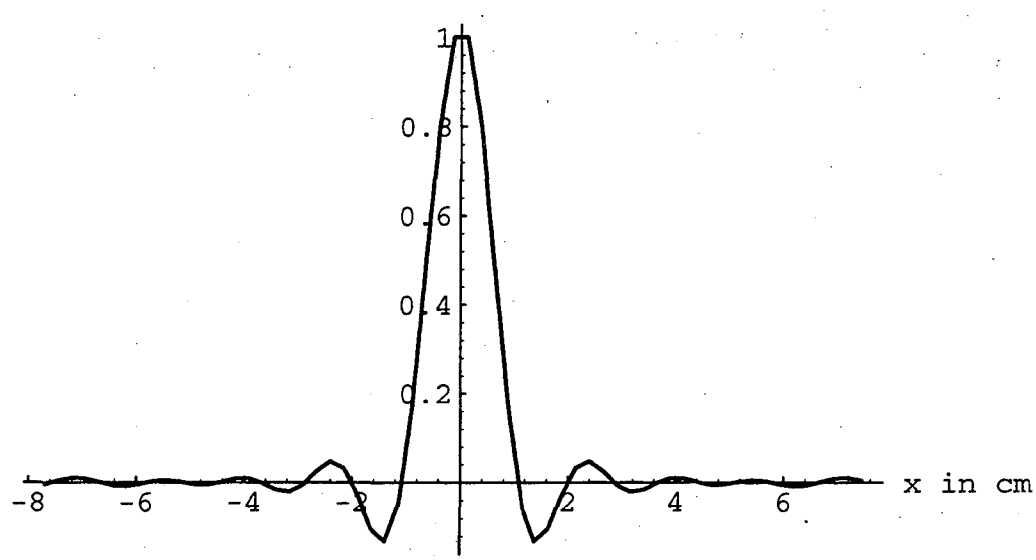

Figure 4.11: Overall point spread function for the incommensurate frequency oscillating gradient experiment ($i_{max} = 3$), 1024 lags; this is the spatial point spread function for the chemical shift reconstruction of the same experiment as in figure 4.8; the overall k_{max} is 0.602898

4.6 Noise Analysis

The initial sampling of k-space in the experiment is non-uniform, which means that a different degree of averaging occurs for different spatial frequency measurements. Therefore, there is a variation in the signal to noise ratio of the different spatial frequencies in the reconstructed image. In a real stochastic NMR experiment, there will be additive white measurement noise introduced into the received signal arising from Johnson noise in the probe and preamplifier which is independent of the gradient and RF waveforms. To determine the signal to noise ratio in the reconstructed image as a function of spatial frequency, we will consider the reconstruction of a band-limited white noise signal $\tilde{y}_n$ of variance ς^2, to determine the power spectral density of noise in the reconstructed image[33].

This noise analysis will only treat the noise power in the reconstructed image due to measurement noise. There is also systematic noise in the reconstruction, which arises from the random nature of the excitation[34].

4.6.1 Single Lag Noise Analysis

We perform the reconstruction on the noise signal $\tilde{y}_n$ for a given lag q to generate the noise image $\tilde{\rho}_q(\mathbf{x}, q)$. The noise power spectral density of such a reconstructed image is defined as:

$$S^q(\mathbf{k}) = \left\langle \int_{\mathbf{x}} \tilde{\rho}_q(\mathbf{x}, q) e^{-i\mathbf{k}\cdot\mathbf{x}} d\mathbf{x} \left[\int_{\mathbf{x}} \tilde{\rho}_q(\mathbf{x}, q) e^{-i\mathbf{k}\cdot\mathbf{x}} d\mathbf{x} \right]^* \right\rangle . \tag{4.47}$$

This is the ensemble average from the reconstructions of many images with different $\tilde{y}_n$ with identical signal statistics. The Fourier transforms inside the angle brackets can be identified as the definition of $\rho_K(\mathbf{k}, q)$.

$$S^q(\mathbf{k}) = \langle \tilde{\rho}_K(\mathbf{k}, q) \tilde{\rho}_K^*(\mathbf{k}, q) \rangle \tag{4.48}$$

Using the definition of $\hat{\rho}_K(\mathbf{k}, q)$ from Equation 4.17, and remembering that s_n and y_n have delta function autocorrelations, we see that:

$$\begin{aligned} S^q(\mathbf{k}) &= \left\langle \sum_{n=1}^{N} \int_{\mathbf{k}'} \frac{\tilde{y}_n s^*_{n-q} \delta^3(\mathbf{k}' - \mathbf{k}_{n,q}) C(\mathbf{k} - \mathbf{k}')}{W^q(\mathbf{k}')} d\mathbf{k}' \right. \\ &\qquad \left. \times \sum_{m=1}^{N} \int_{\mathbf{k}''} \frac{\tilde{y}^*_m s_{m-q} \delta^3(\mathbf{k}'' - \mathbf{k}_{m,q}) C^*(\mathbf{k} - \mathbf{k}'')}{(W^q(\mathbf{k}''))^*} d\mathbf{k}'' \right\rangle \quad (4.49) \\ &= \sum_{n=1}^{N} \sum_{m=1}^{N} \frac{\langle \tilde{y}_n \tilde{y}^*_m \rangle \left\langle s_{m-q} s^*_{n-q} \right\rangle C(\mathbf{k} - \mathbf{k}_{n,q}) C^*(\mathbf{k} - \mathbf{k}_{m,q})}{W^q(\mathbf{k}_{n,q})(W^q(\mathbf{k}_{m,q}))^*} \quad (4.50) \\ &= \varsigma^2 \sum_{n=1}^{N} \frac{\left\langle s_{n-q} s^*_{n-q} \right\rangle C(\mathbf{k} - \mathbf{k}_{n,q}) C^*(\mathbf{k} - \mathbf{k}_{n,q})}{W^q(\mathbf{k}_{n,q})(W^q(\mathbf{k}_{n,q}))^*} \quad (4.51) \end{aligned}$$

In the usual case, the reconstruction employs a real, symmetric convolution function c, so its Fourier transform C, and consequently W^q, are real.

$$S^q(\mathbf{k}) = \varsigma^2 \sum_{n=1}^{N} \frac{(C(\mathbf{k} - \mathbf{k}_{n,q}))^2}{(W^q(\mathbf{k}_{n,q}))^2} \quad (4.52)$$

The noise characteristics for any choice of sampling trajectory and convolution function can be calculated numerically using Equation 4.52. To evaluate the general characteristics of the two sampling trajectories presented analytically, we need an explicit form for the convolution kernel $C(\mathbf{k})$.

We have made two assumptions throughout this derivation; the first is that the sampling pattern covers a contiguous region in k-space adequately and the true density function is smooth enough that we can perform the convolution gridding operation; i.e., that we have approximated the true k-space estimate of the object density well enough that we can resample the density onto a Cartesian grid of a certain granularity. The second is that the extent of the convolution function in k-space is small compared to the extent of the object density ρ_K (so that the convolution function does not alter the shape of the reconstructed density too much).

To make the derivation clear, we will use the case of nearest neighbor interpolation, where all of the energy of a particular $y_n s_{n-q}$ is gridded into the nearest grid point in an oversampled version of the k-space density estimate $\hat{\rho}_K(\mathbf{k}, q)$. Nearest

neighbor interpolation produces greater aliasing artifacts than using other convolution kernels, such as the Kaiser-Bessel function, however the effect on the shape of the noise distribution in small, so it is useful for demonstration purposes.

The nearest neighbor convolution function is:

$$C(\mathbf{k}) = \sqcap(\epsilon) = \begin{cases} 1 & |k| \leq \epsilon/2 \\ 0 & |k| > \epsilon/2 \end{cases} \tag{4.53}$$

where ϵ is the spacing between points on the oversampled Cartesian grid. If ϵ is sufficiently small and W is sufficiently smooth that W is approximately constant over ϵ, C can be treated like a delta function, and the weighting factor can be pulled out of the summation. Therefore:

$$\varsigma^2 \sum_{n=1}^{N} \left| \frac{\sqcap(\mathbf{k} - \mathbf{k}_{n,q})}{W^q(\mathbf{k}_{n,q})} \right|^2 \approx \left(\frac{\varsigma}{W^q(\mathbf{k}_{n,q})} \right)^2 \sum_{n=1}^{N} \sqcap(\mathbf{k} - \mathbf{k}_{n,q})^2 = \left(\frac{\varsigma}{W^q(\mathbf{k})} \right)^2 \sum_{n=1}^{N} \sqcap(\mathbf{k} - \mathbf{k}_{n,q}) \tag{4.54}$$

The term inside the last summation is just the definition of the weighting function W. Therefore, the noise estimate is found to be (to good approximation):

$$\varsigma^2 \sum_{n=1}^{N} \left| \frac{C(\mathbf{k} - \mathbf{k}_{n,q})}{W^q(\mathbf{k}_{n,q})} \right|^2 \approx \frac{\varsigma^2}{W^q(\mathbf{k})} \tag{4.55}$$

One dimensional Monte Carlo simulations of noise reconstruction have verified that this approximation holds quite well for imaging experiments with typical parameters. These simulations are shown in Section 5.6.

To determine the noise power spectral density we must express $W^q(\mathbf{k})$ explicitly.

4.6.2 One Dimensional Oscillating Gradient

For a one dimensional sinusoidal gradient (i = 1), it can be seen from inspection of Equation 4.31 that the k trajectory for any lag q is sinusoidal with an amplitude k^q_{max}. The k-space density along the gradient is proportional to the reciprocal of the velocity of the traversal of k-space $\left|\frac{d\mathbf{k}}{dt}\right|^{-1}$. The sampling density as a function of k (normalized over the range $-k_{max} \leq k \leq k_{max}$) is therefore (if the discrete nature of

the sampling is ignored):

$$W^q(k) = \frac{N}{\pi k^q_{max}\sqrt{1-\left[\frac{|k|}{k^q_{max}}\right]^2}}. \tag{4.56}$$

This gives a noise power spectral density of:

$$S^q(k) = \frac{\varsigma^2 \pi k^q_{max}}{N}\sqrt{1-\left[\frac{k}{k^q_{max}}\right]^2}. \tag{4.57}$$

A general analytic expression for the sampling density arising from multiple component gradients is not easy to formulate. Unlike the sinusoidal gradient, the shape of the k-space trajectory, and therefore the sampling pattern, changes with different correlation lags. The sampling density can be obtained numerically using Equation 4.31.

As discussed in Section 3.2.3.3, the Amplitude Modulated Sinusoidal Gradients have a flat sampling density between $-k^q_{max}$ and k^q_{max}.

4.6.3 Rotating Oscillating Gradients

As mentioned previously, the trajectory of the direction vector of the oscillating gradient has been designed to sample all solid angles uniformly, so we need only consider the variation in sampling density along the radial direction. The effect of extending the imaging technique to multiple dimensions by rotating the gradient direction vector is to add a factor that accounts for the constant distribution of samples over 2π radians for two dimensions, or 4π steradians for three dimensions. This factor is $1/|\mathbf{k}|$ for two dimensions, and $1/|\mathbf{k}|^2$ for three dimensions, giving a normalized three dimensional sampling density of:

$$W^q(\mathbf{k}) = \frac{N}{2\pi^2 k^q_{max}|\mathbf{k}|^2\sqrt{1-\left[\frac{|\mathbf{k}|}{k^q_{max}}\right]^2}}. \tag{4.58}$$

The noise power spectral density then becomes:

$$S^q(\mathbf{k}) = \frac{2\varsigma^2\pi^2 k^q_{max}|\mathbf{k}|^2}{N}\sqrt{1-\left[\frac{|\mathbf{k}|}{k^q_{max}}\right]^2}. \tag{4.59}$$

This noise density is shown in Figure 4.12.

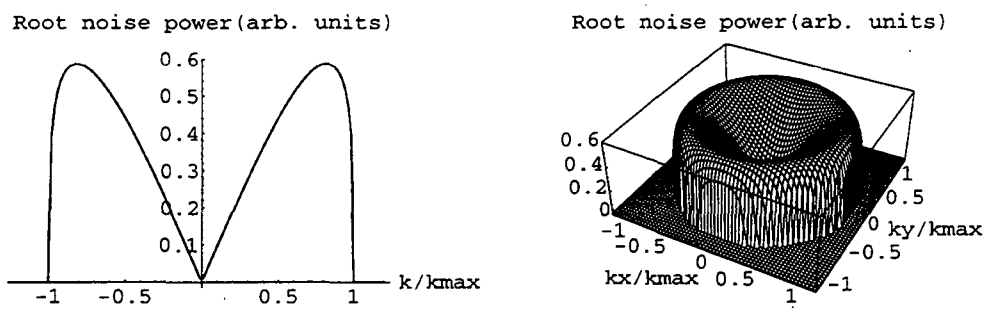

Figure 4.12: Square root of the single lag noise power spectral density for the three-dimensional stochastic experiment using rotating oscillating gradients as a function of k/k_{max}. The figure on the left shows a line through the origin of k-space; the figure on the right shows two k-space dimensions.

4.6.4 Incommensurate Frequency Oscillating Gradients

Because the gradients used in the incommensurate frequency experiment are not harmonically related, the weighting function is simply the product of the individual sampling densities along each axis.

$$W^q(\mathbf{k}) = \frac{N}{\pi^3 k^q_{x(max)} k^q_{y(max)} k^q_{z(max)} \sqrt{\left(1-\left[\frac{|k_x|}{k^q_{x(max)}}\right]^2\right)\left(1-\left[\frac{|k_y|}{k^q_{y(max)}}\right]^2\right)\left(1-\left[\frac{|k_z|}{k^q_{z(max)}}\right]^2\right)}}. \tag{4.60}$$

The noise power spectral density for one lag then becomes:

$$S^q(\mathbf{k}) = \frac{\varsigma^2 \pi^3 k^q_{x(max)} k^q_{y(max)} k^q_{z(max)}}{N} \sqrt{\left(1-\left[\frac{|k_x|}{k^q_{x(max)}}\right]^2\right)\left(1-\left[\frac{|k_y|}{k^q_{y(max)}}\right]^2\right)\left(1-\left[\frac{|k_z|}{k^q_{z(max)}}\right]^2\right)}. \tag{4.61}$$

This noise density is shown in Figure 4.13.

4.6.5 Multi-lag Noise Characteristics

The analysis of the overall noise characteristics of the reconstructed signal is somewhat more complicated, since there may be cross terms between the noise

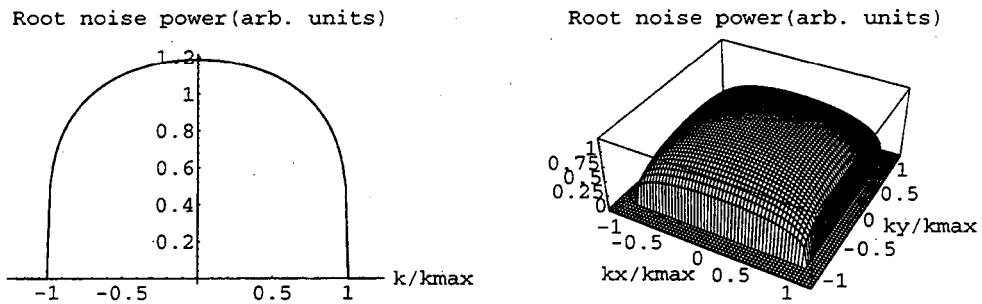

Figure 4.13: Square root of the single lag noise power spectral density for the three-dimensional stochastic experiment using incommensurate frequency sinusoidal oscillating gradients as a function of k/k_{max}. The figure on the left shows a line through the origin of k-space; the figure on the right shows two k-space dimensions.

contribution of various reconstructed lags. However, these are seen to be eliminated by the autocorrelation properties of the excitation sequence. Closely following the single lag derivation, the reconstructed noise for the full reconstruction is:

$$S(\mathbf{k},\sigma) = \left\langle \sum_{q=0}^{Q-1} e^{i(q+1)\sigma T_R} \tilde{\rho}_K(\mathbf{k},q) \sum_{q'=0}^{Q-1} e^{-i(q'+1)\sigma T_R} \tilde{\rho}_K^*(\mathbf{k},q') \right\rangle \tag{4.62}$$

$$= \sum_{q=0}^{Q-1}\sum_{q'=0}^{Q-1} e^{i(q-q')\sigma T_R} \left\langle \tilde{\rho}_K(\mathbf{k},q)\tilde{\rho}_K^*(\mathbf{k},q') \right\rangle \tag{4.63}$$

$$= \sum_{q=0}^{Q-1}\sum_{q'=0}^{Q-1} e^{i(q-q')\sigma T_R} \sum_{n=1}^{N}\sum_{m=1}^{N} \frac{\langle \tilde{y}_n \tilde{y}_m^* \rangle \left\langle s_{m-q'} s_{n-q}^* \right\rangle C(\mathbf{k}-\mathbf{k}_{n,q}) C^*(\mathbf{k}-\mathbf{k}_{m,q'})}{W^q(\mathbf{k}_{n,q})(W^{q'}(\mathbf{k}_{m,q'}))^*} \tag{4.64}$$

$$= \varsigma^2 \sum_{q=0}^{Q-1}\sum_{q'=0}^{Q-1} e^{i(q-q')\sigma T_R} \sum_{n=1}^{N} \frac{\left\langle s_{n-q'} s_{n-q}^* \right\rangle C(\mathbf{k}-\mathbf{k}_{n,q}) C^*(\mathbf{k}-\mathbf{k}_{n-q+q',q'})}{W^q(\mathbf{k}_{n,q})(W^{q'}(\mathbf{k}_{n-q+q',q'}))^*} \tag{4.65}$$

$$= \varsigma^2 \sum_{q=0}^{Q-1}\sum_{n=1}^{N} \frac{C(\mathbf{k}-\mathbf{k}_{n,q}) C^*(\mathbf{k}-\mathbf{k}_{n,q})}{W^q(\mathbf{k}_{n,q})(W^q(\mathbf{k}_{n,q}))^*} \tag{4.66}$$

$$= \sum_{q=0}^{Q-1} S^q(\mathbf{k}) \tag{4.67}$$

Therefore, the multilag noise power is simply the sum of the individual lag noise powers. The multi-lag noise powers are shown for the rotating oscillating gradient experiment and the incommensurate gradient experiment in Figures 4.14 and 4.15 respectively.

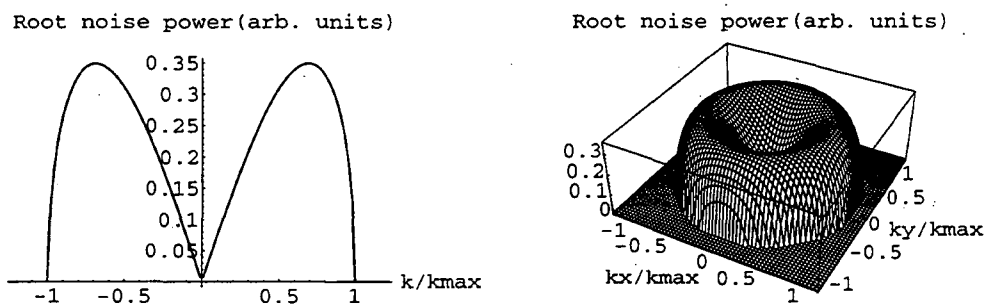

Figure 4.14: Square root of the overall noise power spectral density for the three-dimensional stochastic experiment using rotating oscillating gradients as a function of k/k_{max}. The figure on the left shows a line through the origin of k-space; the figure on the right shows two k-space dimensions.

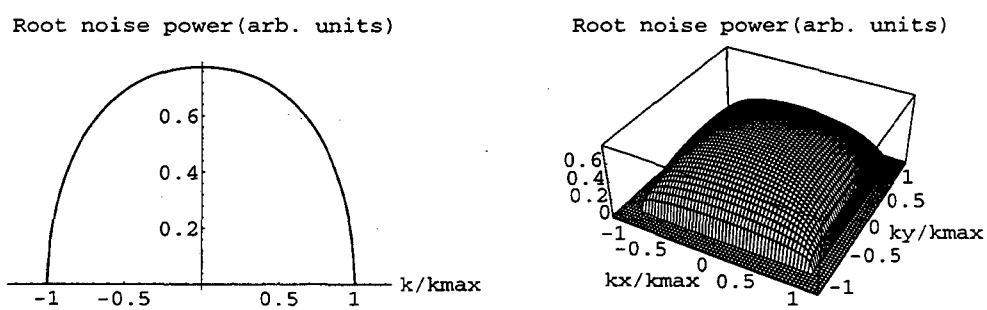

Figure 4.15: Square root of the overall noise power spectral density for the three-dimensional stochastic experiment using incommensurate frequency sinusoidal oscillating gradients as a function of k/k_{max}. The figure on the left shows a line through the origin of k-space; the figure on the right shows two k-space dimensions.

4.7 Implementation

This reconstruction technique has been implemented in ANSI C to run on a variety of platforms, and the specific implementation details have been tuned to efficiently use the resources of a scientific workstation. The order of various steps in the procedure can be interchanged for more efficient computation at the cost of requiring more memory or mass storage.

The first step in the reconstruction is to calculate and store the integral of the gradients and the conjugate of the excitation signal for every time point recorded in the experiment. The image reconstruction is then treated one correlation lag at a time; for a given lag q the k-space position of each data point is determined from the difference of the gradient integral at the time the signal was recorded and q lags in the past when it was created. The signal is multiplied by the conjugate of the excitation q lags in the past, and convolved with the spreading function C. The resulting continuous function is sampled into an oversampled three dimensional k-space array. Two arrays are maintained; the sampled function values are summed into a weight array that keeps an estimate of the sampling density, and the function values, multiplied by the complex conjugate of the excitation, are summed into an array that stores the FT of the spin density estimate. After all the points have been gridded, the k-space array is divided by the weight array for all non-zero locations in the weight array to density correct the k-space estimate. An inverse Fourier transform and downsampling into the spatial domain is performed, and the doming effect of the spreading function C is removed. The result is a spatial reconstruction for one time lag.

When reconstructing spectra, the array for each lag is multiplied by an exponential weighting factor to perform apodization, if desired, and the resulting weighted array is multiplied by the appropriate Fourier phase rotations to transform it to the spectral domain and summed into the estimate of the chemical shift reconstruction. The set of operations is repeated until all the lags have been calculated.

For time domain (FID) reconstructions, the data array for each lag is also multiplied by an apodization function, and then the data are interpolated and down-

sampled into the destination time array. The time point spacing in the destination array is the reciprocal of the lowest gradient frequency. This ensures that every **k** position in k,t-space is critically sampled in time. As with spectral reconstructions, this operation is repeated for all time lags.

Although it is more computationally efficient (and straightforward) to perform the Fourier transform into the spectral domain after all of the lags have been calculated, the memory requirements of this technique make that impractical on most workstations. As an example, to perform a 32x32x32 chemical shift reconstruction, oversampled by a factor of 2 (which is sufficient to remove any obvious artifacts of the gridding process) requires 32K complex data points per lag (256KB assuming 4 byte floating point representation). Generally at least 1024 time lags are used in calculating a spectral data set, which would require at least 256MB of storage to keep the entire k-space array available, which requires disk storage of the k-space array. This estimate neglects the storage needed for the input data, gradient integral, and excitation signals, which can be quite large (10's of megabytes), and must be kept in memory for efficient computation. The time savings from keeping the k-space array in memory and performing the spectral Fourier transform at the end are more than offset by having to go to disk to read and write the lag data.

Pseudocode for the reconstruction algorithm is shown in the appendix.

4.8 Discussion

There are three major advantages of the Fourier gridding reconstruction technique over weighted crosscorrelation in the spatial domain. First, it is much faster than latter method for large problems. Second, because the k-space sampling density is estimated explicitly, the compensation for uneven k-space sampling can be performed almost exactly. Third, the derivation of this reconstruction is not dependent on the k-space scanning trajectory, and the scanning trajectory can be chosen without concern for the mathematical tractability of deriving an appropriate weighting function.

The gradient trajectories analyzed here have very different noise characteristics,

arising from their very different sampling densities. The noise characteristics can be further shaped by modifying the density weighting correction, which also allows the point spread function to be tailored. There have recently been suggestions that preferential weighing of the low-frequency regions of k-space (as in the rotating oscillating gradient experiment) is desirable for chemical shift imaging [35, 36, 37].

As an example of the time savings that this new algorithm affords, consider the reconstruction of a 32x32x32 voxel chemical shift image with 1024 correlation lags for each voxel from a stochastic data set of 4,194,304 data points (a typical experiment length using the old reconstruction algorithm). On a Silicon Graphics Crimson Elan workstation, using the old algorithm, the reconstruction would take approximately 192 seconds/voxel, or 1748 hours; the same reconstruction using the Fourier gridding technique takes approximately 26 seconds per correlation lag, or 7.4 hours, an improvement of greater than two orders of magnitude. Using a repeating k-space trajectory (described in Section 5.5.1) allows for further speed improvements; a 32^3 voxel experiment with 6,745,200 data points can be reconstructed in about 4 hours.

It should also be noted that since each correlation lag can be calculated independently, the new reconstruction method has a high degree of parallelism which is not currently being exploited. Performing the reconstruction on a machine with a parallel architecture (or many machines) will lead to further speed improvements.

Chapter 5

Simulations

5.1 Introduction

This chapter will present the results of Bloch equation simulations of the stochastic experiment. The experiment simulations were used for three major purposes. The first was to validate the theoretical predictions of the experimental and reconstructed image parameters, such as the bandwidth of the experiment with different gradient trajectories (Section 5.3), the point spread function for various gradient trajectories (discussed in Section 5.4), and the predictions of how measurement noise propagates from the received signal to the reconstructed image (Section 5.6). The second purpose of simulations, treated in Section 5.5, was to examine how experiment and reconstruction parameters affect the systematic noise in the reconstructed image due to the stochastic excitation. This is a very difficult area to treat analytically; simulations gave a simple way to test various ideas for reducing the systematic noise without the complication of additional real measurement noise which would have accompanied actual experiments. Finally, simulations provided a method for performing validity checks during development of the reconstruction software.

5.2 Simulation Program

Simulations have been very important for developing and testing the parameters of the stochastic experiment. This has been especially true since the development of the stochastic imaging experiment has required major modifications to the spectrometer used for the experiments; simulation has helped diagnose hardware problems in addition to finding optimum parameters for the experiments and evaluating theoretical predictions of performance.

All simulations shown were performed using a simulation program which integrates the Bloch equations for an arbitrary number of spins with adjustable T_1, T_2, chemical shift, amplitude, and position. Flip angle, RF pulse duration and measurement noise level are adjustable, as are excitation noise type and gradient trajectory. Simulated data sets have been used to check theoretical predictions of point spread function, noise power spectral density, and systematic noise. The simulation program generates data files in the same format as the spectrometer, which can be processed as if they were actual data. As a result, the simulated data sets can be used to test all of the reconstruction software.

5.3 Experiment Bandwidth

The signal bandwidths of the sinusoidal oscillating gradient stochastic experiment and the amplitude modulated sinusoidal gradient experiment are easily calculated using Carson's rule (Equation 3.5). Unfortunately, due to the nonlinear modulation process, a general derivation of the bandwidth of the multiple component truncated square wave gradient is not practical. However, it is easily demonstrated with simulations that the experiment bandwidth is quite similar to the sinusoidal case; close enough that Carson's rule is an appropriate approximation. This observation was made repeatedly during experiments and simulations.

As a demonstration of the relative signal bandwidths arising from the different gradient trajectories, a one dimensional experiment was simulated on a uniform density one dimensional object of known extent. The sampling frequency was delib-

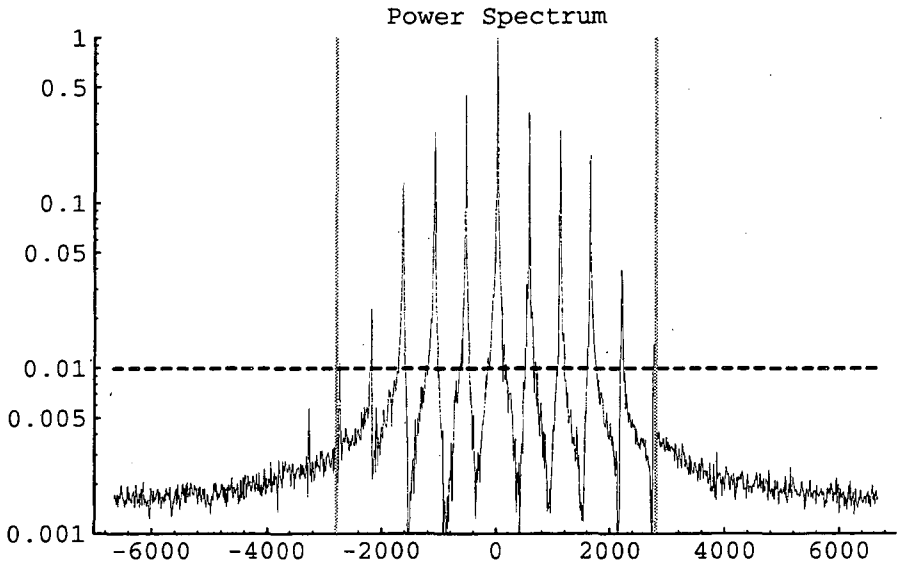

Figure 5.1: Power spectrum, one dimensional sinusoidal gradient experiment. Overall experimental bandwidth is 13334Hz. The bandwidth predicted by Equation 3.5 for the simulated object is 5590Hz (shown by the vertical bars). The 1% contour is the dashed horizontal line. $T_R = 75\mu s$, $r_{max} = 2.5cm$, $G = 8.0mT/m$, $\gamma = 11.24MHz/T$, $f_0 = 546.008Hz$.

erately chosen to far exceed the signal bandwidth arising from the object, so that a good representation of the signal power spectrum could be obtained. The power spectra of the simulated NMR signals were calculated and are presented in Figures 5.1, 5.2, and 5.3.

The simulation was of a one dimensional sodium experiment, $T_R = 75\mu s$, $G = 8.0mT/m$, $\gamma = 11.24MHz/T$, $f_0 = 546.008Hz$. The overall experimental bandwidth is 13334Hz. A uniform object which extended from -2.5cm to 2.5cm was simulated with various gradients. Carson's rule predicted that all significant sidebands would be contained within 5590Hz. The resulting power spectra are compared to the predicted bandwidth. In all cases, Carson's rule accurately predicted the bandwidth which includes all sidebands with amplitudes greater than or equal to 1% of the amplitude of the fundamental.

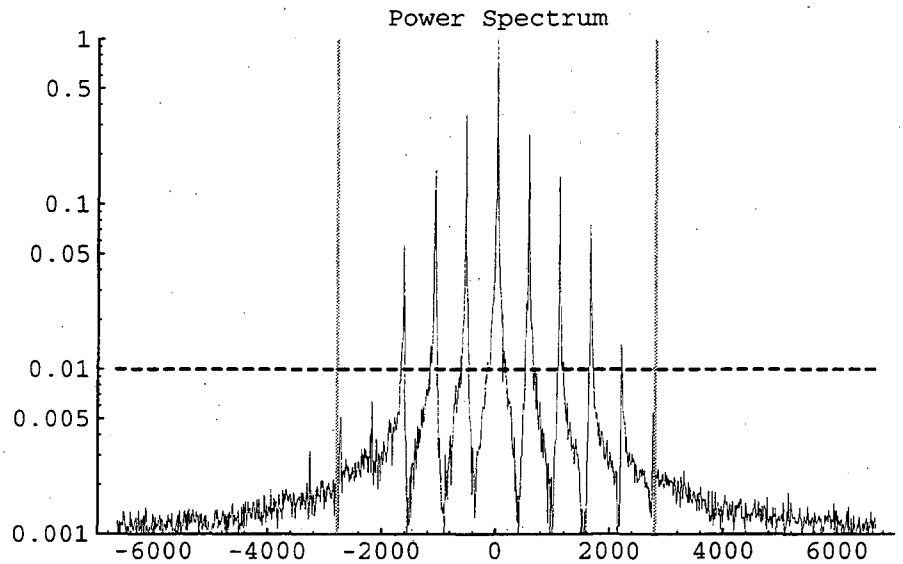

Figure 5.2: Power spectrum, amplitude modulated one dimensional sinusoidal gradient experiment. Overall experimental bandwidth is 13334Hz. The bandwidth predicted by Equation 3.5 for the simulated object is 5590Hz (shown by the vertical bars). The 1% contour is the dashed horizontal line. $T_R = 75\mu s$, $r_{max} = 2.5cm$, $G = 8.0mT/m$, $\gamma = 11.24MHz/T$, $f_0 = 546.008Hz$.

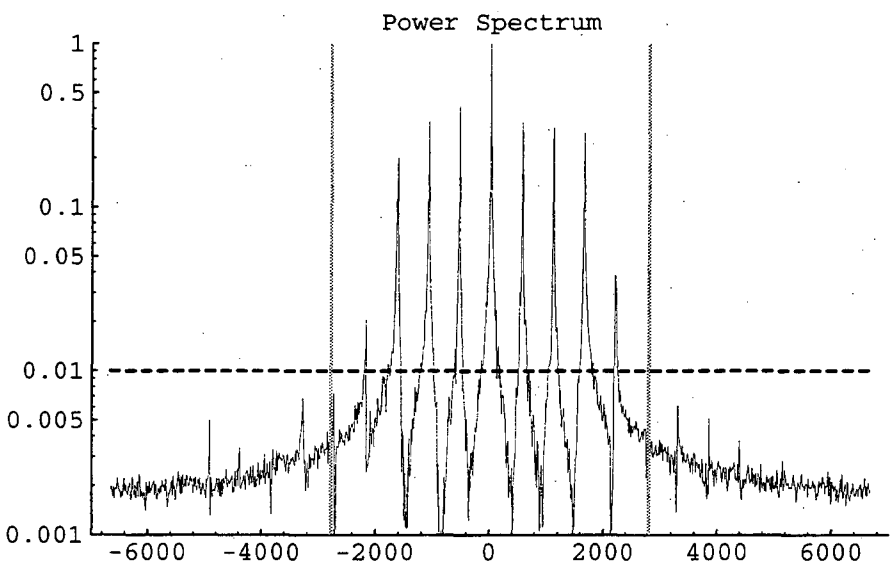

Figure 5.3: Power spectrum, one dimensional three component truncated square wave gradient experiment. Overall experimental bandwidth is 13334Hz. The bandwidth predicted by Equation 3.5 for the simulated object is 5590Hz (shown by the vertical bars). The 1% contour is the dashed horizontal line. $T_R = 75\mu s$, $r_{max} = 2.5cm$, $G = 8.0mT/m$, $\gamma = 11.24MHz/T$, $f_0 = 546.008Hz$.

5.4 Point Spread Function

A set of simulations was performed to validate the theoretical expressions for the point spread functions derived in Subsection 4.5. Figures 4.6, 4.9, 4.7 4.10, 4.8 and 4.11 show the theoretical PSF's for the rotating oscillating gradient and incommensurate frequency oscillating gradient experiments with typical imaging parameters. For validation, simulations were performed of one spin at the origin with the experimental parameters used in the theoretical expressions.

In most cases the agreement between theoretical expressions and the simulations was excellent, indicating both that the analysis was correct and that the reconstruction program was functioning properly. In some of the chemical shift reconstructions, however, the first sidelobe amplitude was attenuated. This is because the maximum extent of the sampled region in k-space for any lag is discretized in the gridding and density correction process, and is always an integral number of k-space pixels. Therefore, when many lags are combined, the resulting PSF is a sum of rect functions with a discrete rather than continuous set of sizes. The degree of divergence of the point spread function can be reduced by increasing the oversampling factor.

Figures 5.4 and 5.5 show the single lag and multilag point spread functions for a three dimensional experiment with rotating oscillating gradients. The theoretical expressions for the PSF's are shown on the same plot. Figures 5.6 and 5.7 show the single lag and multilag PSF's for the incommensurate frequency sinusoidally oscillating gradient experiment. Figures 5.8 and 5.9 show the single lag and multilag PSF's for the incommensurate frequency three component oscillating gradient experiment. Figures 5.10 and 5.11 show the single lag and multilag PSF's for the three dimensional amplitude modulated incommensurate frequency sinusoidally oscillating gradient experiment.

5.5 Systematic Noise Reduction

One consequence of using stochastic excitation is that it introduces systematic noise into the reconstructed image arising from the excitation and reconstruction

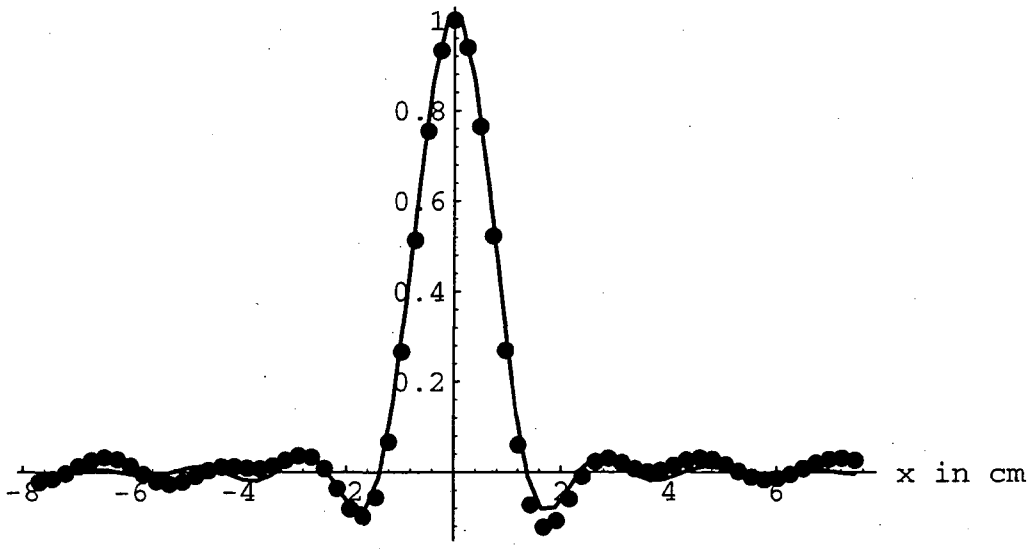

Figure 5.4: Normalized point spread function for a single lag using the rotating oscillating gradient; points are from the reconstruction of a simulated stochastic experiment with the experimental parameters given in Figure 4.4a. This is the reconstruction at the first maximum of $k_{max}(q)$, lag 11. $k_{max}(11) = 0.5237$; the line is a graph of equation 4.40 with the same parameters.

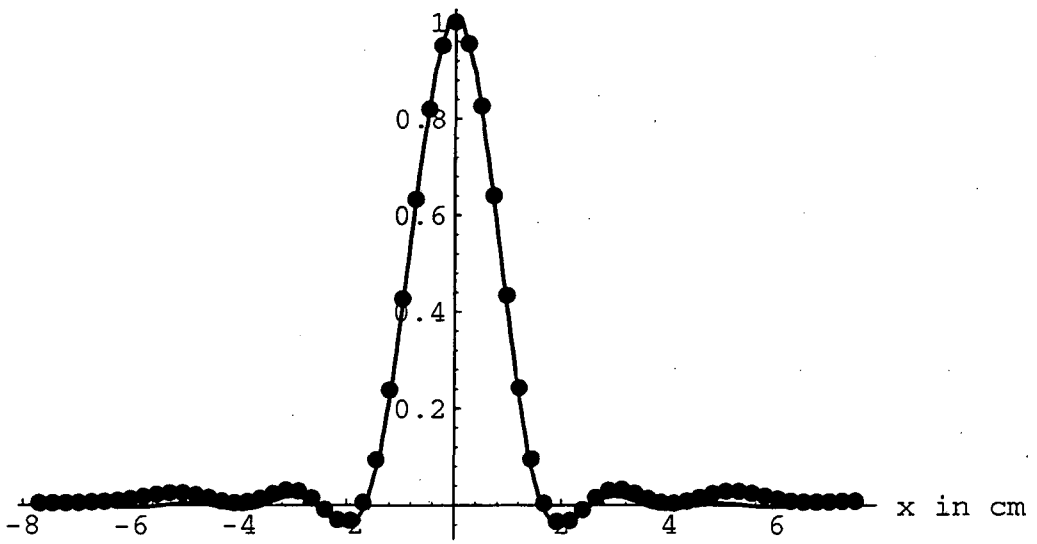

Figure 5.5: Overall point spread function for the rotating oscillating gradient experiment, 1024 lags; this is from the chemical shift reconstruction of the same data set as in figure 4.6; the overall k_{max} is 0.523813

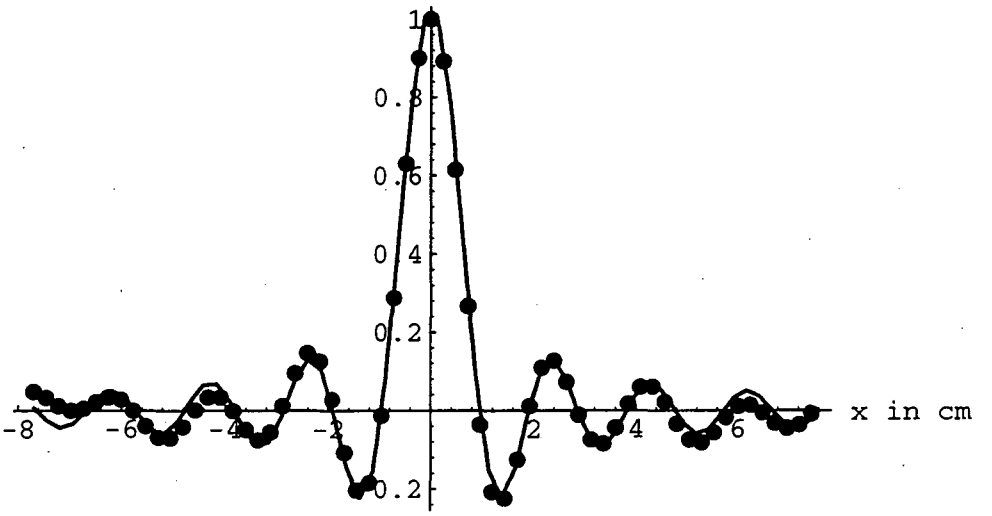

Figure 5.6: Normalized point spread function for a single lag, incommensurate frequency sinusoidal gradients ($i_{max} = 1$); points are from the reconstruction of a simulated stochastic experiment with the experimental parameters given in Figure 4.4a. This is the reconstruction at the first maximum of $k_{max}(q)$, lag 11. $k_{max}(11) = 0.5237$; the line is a graph of equation 4.43 with the same parameters.

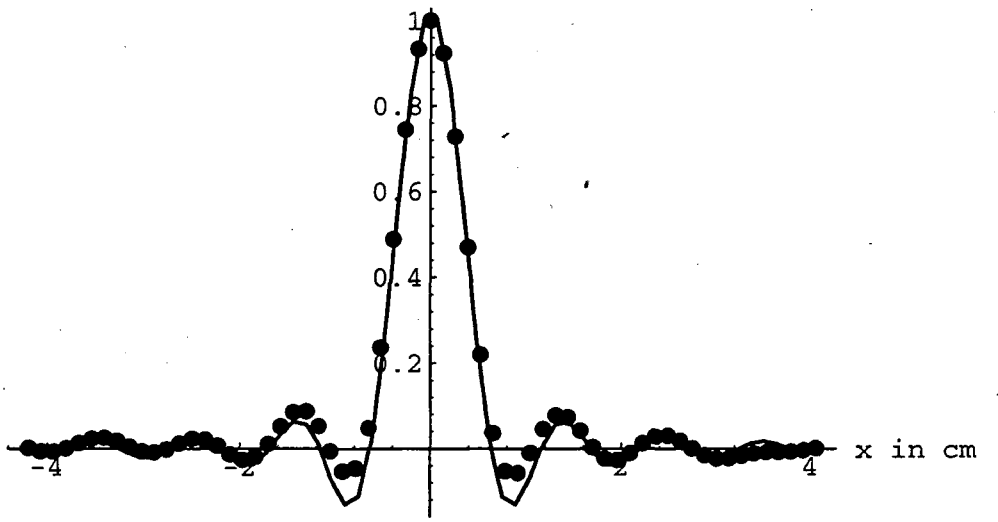

Figure 5.7: Overall point spread function for the incommensurate frequency sinusoidally oscillating gradient experiment ($i_{max} = 1$), 1024 lags; this is from the chemical shift reconstruction of the same data set as in figure 4.7; the overall k_{max} is 0.523813

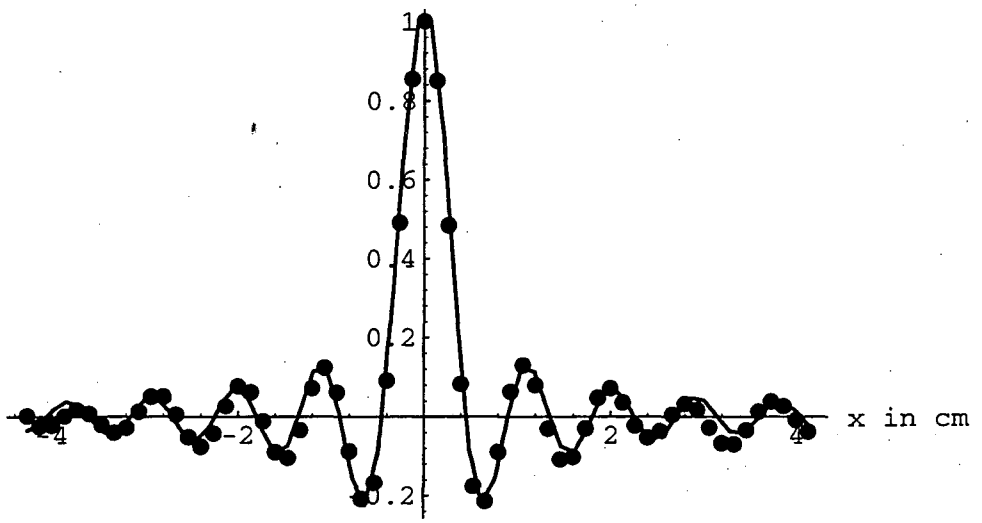

Figure 5.8: Normalized point spread function for a single lag, incommensurate frequency three component TSW gradients ($i_{max} = 3$); points are from the reconstruction of a simulated stochastic experiment with the experimental parameters given in Figure 4.4b. This is the reconstruction at the first maximum of $k_{max}(q)$, lag 11. $k_{max}(11) = 0.5237$; the line is a graph of equation 4.43 with the same parameters.

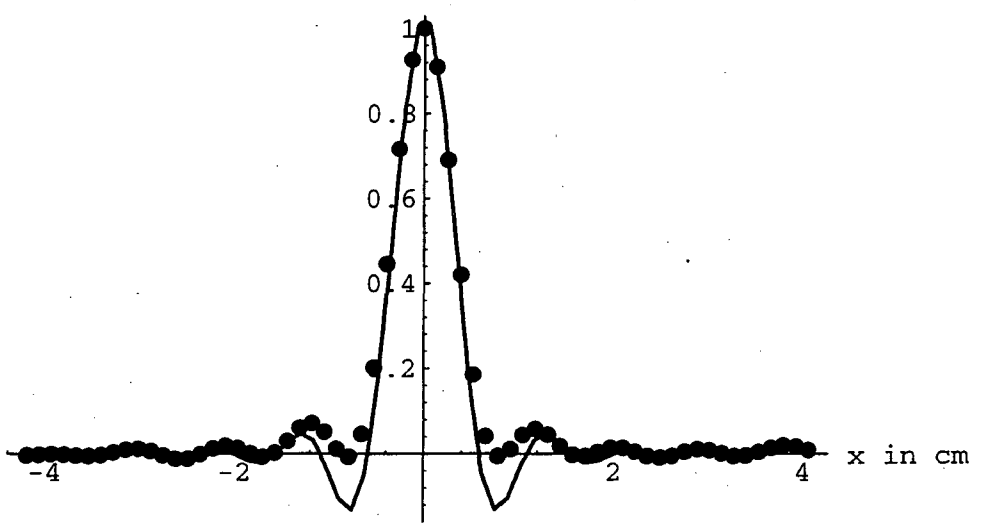

Figure 5.9: Overall point spread function for the incommensurate frequency three component TSW oscillating gradient experiment ($i_{max} = 3$), 1024 lags; this is from the chemical shift reconstruction of the same data set as in figure 4.8; the overall k_{max} is 0.523813

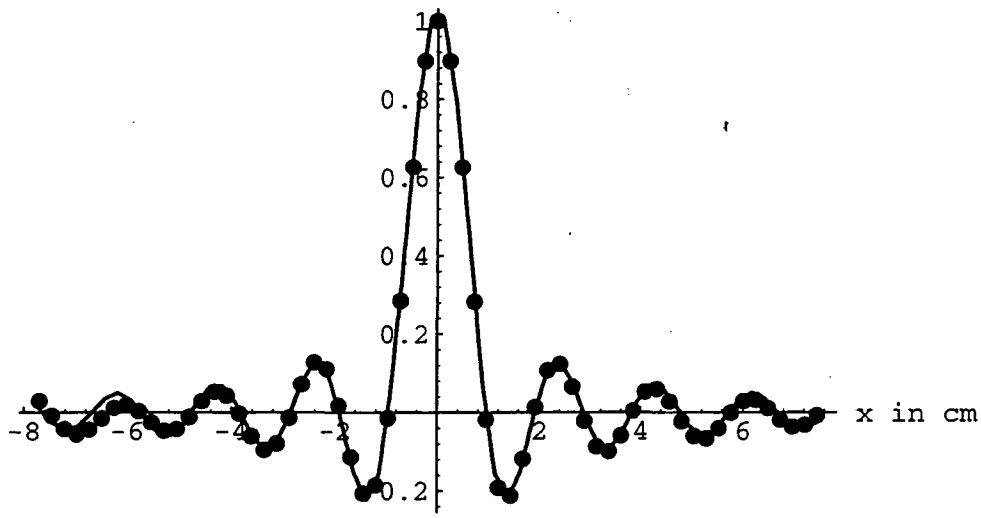

Figure 5.10: Normalized point spread function for a single lag, amplitude modulated incommensurate frequency sinusoidal gradient experiment; points are from the reconstruction of a simulated stochastic experiment with the experimental parameters given in Figure 4.4a. This is the reconstruction at the first maximum of $k_{max}(q)$, lag 11. $k_{max}(11) = 0.5237$; the line is a graph of equation 4.43 with the same parameters.

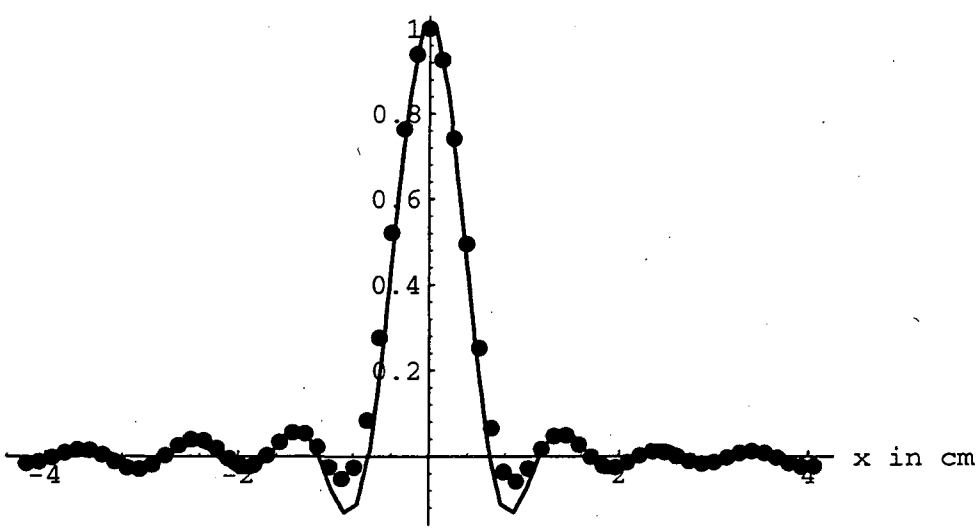

Figure 5.11: Overall point spread function for the amplitude modulated incommensurate frequency oscillating gradient experiment, 1024 lags; this is from the chemical shift reconstruction of the same data set as in figure 4.7; the overall k_{max} is 0.523813

process, in addition to the random thermal noise present in any physical measurement. This systematic noise comes from the variance in the estimate of the object density in each position in k-space. There are many factors which affect the magnitude of the variance of the estimate; these include the total length of the excitation sequence, the sampling density in k-space, higher order correlations in the noise sequence, and object dependent cross-terms between the object density and the excitation sequence. This is an extremely complicated subject to treat analytically, and such a treatment will be avoided here. For a detailed examination of the subject see Wong[14]. There are some general statements which can be made regarding the factors that affect systematic noise, however.

The most important consideration in reducing the systematic noise in the reconstructed image is the number of samples at any point in k-space. The variance in the estimate of the object density is inversely proportional to the number of samples at that point. There is therefore an intimate relationship between gradient trajectories and systematic noise, since the sampling density is determined by the gradient trajectory used.

5.5.1 Repeating versus Nonrepeating k Trajectories

One of the most useful (and surprising) results derived from experimental simulations was the discovery that repetitious sampling in k-space leads to a larger than expected reduction in the amount of systematic noise in the reconstructed image arising from a stochastic experiment for experiments based on incommensurate frequency gradients. Initially, care was taken to keep the sample rate of the experiment unrelated to the oscillation frequencies of the gradients, based on the assumption that the artifact introduced by gridding would be smallest if the k-space sampling along each axis before gridding was very fine. However, this strategy introduces two difficulties. The first is that when the gradient frequencies along each axis are completely unrelated, the three dimensional sampling density is very hard to characterize. While the sampling density will converge to the expression in Equation 4.60 as the number of data points goes to infinity, it is very difficult to know *a priori*

what the overall sampling density in any three dimensional k-space position is for an experiment of finite length. The density can be tabulated for any experiment, but there is no simple relationship that will predict the actual number of sample points that will land in any particular voxel in k-space for arbitrary gradient frequencies. When the gradient frequencies are chosen solely to be incommensurate with each other and the sample frequency for experiments of finite length, there are holes in the three dimensional sampling density, which can mean that the object density is undersampled. This leads to artifacts in the reconstructed image.

The second consideration is that this variation in sampling density means that the degree of cancellation of the excitation sequence at every voxel in k-space can vary quite a bit. As a result, the assumption of Equation 4.19, that the autocorrelation of the excitation sequence is a delta function, is not a very good one for many voxels in the image. This in turn introduces systematic object dependent noise into the image.

Simulations revealed that both of these considerations can have an effect on final image quality. The first task was to investigate the effect of the choice of gradient frequency directly on sampling density. Simulations were performed of three dimensional experiments with numerous choices of incommensurate gradient frequencies, and the sampling density was tabulated for each experiment. In each case, there were a number of holes in the sampling distribution, which sampling distribution, indicating that k-space was not adequately sampled. No obvious relationship was found that predicted how many holes there were or the deviation of the sampling density from the average value from the gradient frequencies.

To examine the effect of noise sequence cancellation on systematic noise, two sets of simulations were performed. In the first set three dimensional experiment ≈ 8 million points long was simulated and reconstructed. The gradient frequencies had been determined from the previous simulations to have very few holes in the k-space sampling density even for relatively short (≈ 1 million points) experiments. Then 8 experiments were simulated and reconstructed which covered the first 1/8th of the k-space trajectory of the long experiment, and the resulting images were added. Each of these simulations were started with different seeds in the noise sequence. The two

images had the same total number of samples, and both fully sampled k-space. The second image using the multiple noise seeds had significantly lower systematic noise than the single long experiment, indicating that sampling the exact same locations in k-space times resulted in observably better images than sampling for the same total amount of time with the same average distribution of k-space sample positions. Because the total experiment time and average sampling density are the same, the measurement noise variance in the image (as distinct from the systematic noise) is the same.

A technique which addresses both of these issues is to choose gradient frequencies such that the gradient along each axis repeats in an integral number of samples. By properly choosing the number of samples in which each gradient repeats (denoted a, b, and c), a k-space trajectory can be chosen which:

1. Adequately samples three dimensional Cartesian k-space at a chosen resolution for all lags in a known number of data points.

2. Revisits the same locations in k-space repeatedly as experiment time increases with different parts of the excitation sequence to better estimate the object density.

To guarantee that k-space is adequately sampled, the repeat times along the three axes must be chosen to satisfy certain conditions. The first is that the three integers a, b and c should share no common prime factors. This is a slight relaxation of the requirement in Equation 4.41; because a, b and c are integers, the three gradients will return to the same relative phase every $N = a \times b \times c$ points, but will not do so before that. If an experiment runs for N points, it will have exactly one sample in each of N positions in a right rectangular prism in k-space. If the gradient frequencies and amplitudes are kept close, this region will be approximately cubical. This first condition can be met by choosing a, b and c to be consecutive odd integers.

The second condition is the sampling density must be sufficient so that when the k-space samples are interpolated onto a Cartesian grid, there will be no holes in the sampled distribution. This can be assured if the greatest distance between samples

on each axis is less than or equal to the spacing of points on the Cartesian grid. Of the gradient waveforms proposed, the sinusoidal gradient has the most uneven sample spacing. A sinusoidal gradient which repeats after an odd number of points a will sample a locations along an axis in k-space. The greatest distance between sample points will be:

$$\Delta k \approx \frac{2\pi k_{max}}{a}. \tag{5.1}$$

Therefore, for a sinusoidal gradient to adequately sample p points on a Cartesian grid extending from $-k_{max}$ to k_{max},

$$a \geq \frac{\pi p}{2}. \tag{5.2}$$

The spectroscopic reconstructions considered here are usually on the order of 32x32x32, which requires repeat times of greater than 51 points, or a minimum repeat time of $51 \times 53 \times 55 = 148665$ points. If the repeat time of the noise sequence is not equal to the gradient repeat time, running the experiment for an integral number of full gradient cycles will be equivalent to running the experiment with multiple noise seeds, as described above.

As an additional benefit, the use of repeating gradients speeds up reconstruction time severalfold by allowing the use of pretabulated k-space positions so that the gridding operation becomes a table lookup. It also decreases memory access since several points end up in the same k-space position and can be processed simultaneously. The combined speedup from these two factors is approximately a factor of three relative to nonrepeating gradient trajectories for reconstruction of normal spectroscopic data sets of 5 to 10 million points in the current implementation of the reconstruction program.

5.5.2 Analytic versus Empirical Sampling Density Correction

Another issue which affects image quality in a systematic way is the sampling density correction scheme used. As discussed in Chapter 4, sampling density correction can be performed either with empirically calculated sampling patterns or

the analytically derived expressions (as shown in 4.6) or with the tabulated sampling density at each time lag. Simulations with sinusoidal oscillating gradients showed that analytic density correction was preferable in terms of image quality. This is because the analytic density correction is applied before the k-space position is discretized and is correct for every data point dropped into the k,t-space object estimate, whereas the empirical sampling density correction is applied afterwards, and makes no distinction between points within a discrete bin in k-space. For sharply peaked sampling densities such as that of the sinusoidal gradient, there can be significant density variation across a k-space voxel. For certain sampling trajectories, such as the TSW gradients, the analytic sampling density is difficult to calculate analytically; in all TSW gradient experiments and simulations, empirical sampling density was employed out of necessity.

5.6 Monte Carlo Simulations of Noise Variance

Because the sampling density in stochastic experiment is non-uniform for most of the proposed trajectories, the measurement noise in the experiment will lead to colored noise in the reconstructed image. The propagation of noise through the reconstruction is analyzed theoretically in section 4.6. This analysis was validated for single lag reconstruction using Monte Carlo analysis.

A one dimensional oscillating gradient experiment was simulated for the noise analysis. The one dimensional results serve to validate the derivation of the basic expression for noise propagation, and the relationship between sampling density and noise power spectral density. By comparing the one dimensional simulation results to equation 4.55, the accuracy of the expression can be shown, and the validity of 4.55 is inferred. Due to the reconstruction time, it was impractical to simulate the three dimensional experiments repeatedly.

The noise power was estimated by generating and reconstructing data sets consisting solely of measurement noise. The experiment parameters were chosen to reflect typical values which would be employed in a ^{23}Na imaging experiment. The noise-only experiment was performed 2048 times and the mean power spectral den-

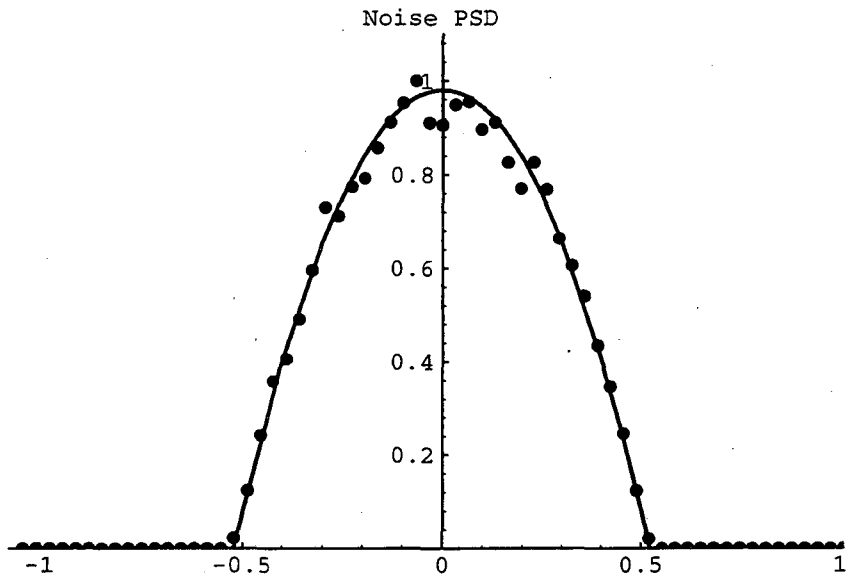

Figure 5.12: Single lag noise power spectral density in one dimension, sinusoidal oscillating gradients. The mean power spectral density of 2048 reconstructions of lag 12 of 128Kpoint data sets containing only measurement noise is shown as points. The line is a plot of Equation 4.57 with the same parameters. γ is 11.24 MHz/T, G is 8 mT/m, T_R is 75μs f_o is 546.977Hz.

sity of the reconstructed k-space object density was calculated. The resulting power spectral densities are shown with the noise power distributions predicted by equation 4.55. In all cases the agreement is excellent. Note that for the three component square wave gradient, the sampling density used in the calculated noise power was derived numerically by tabulating the k-space positions sampled in the experiment.

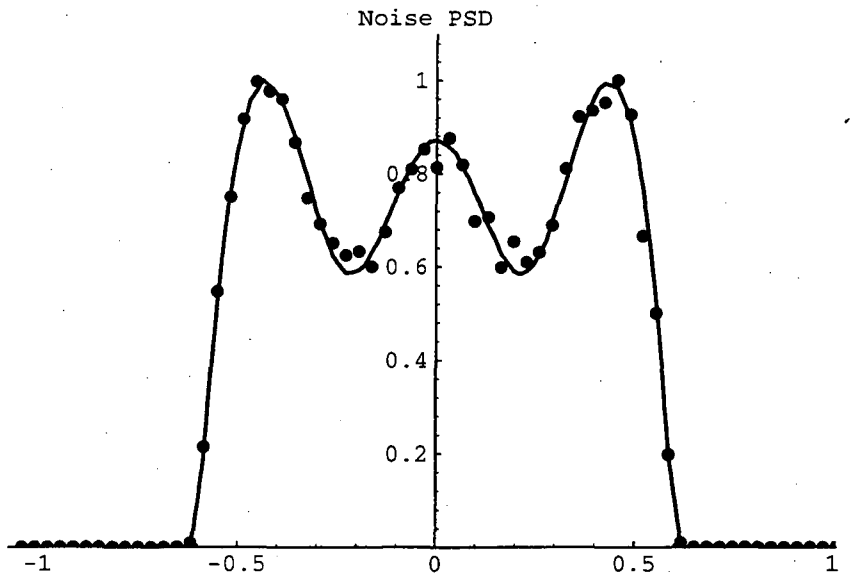

Figure 5.13: Single lag noise power spectral density in one dimension, three component truncated square wave oscillating gradients. The mean power spectral density of 2048 reconstructions of lag 12 of 128Kpoint data sets containing only measurement noise is shown as points. The line is a plot of Equation 4.55 using the tabulated sampling density for the experiment. γ is 11.24 MHz/T, G is 8 mT/m, T_R is 75μs f_o is 546.977Hz.

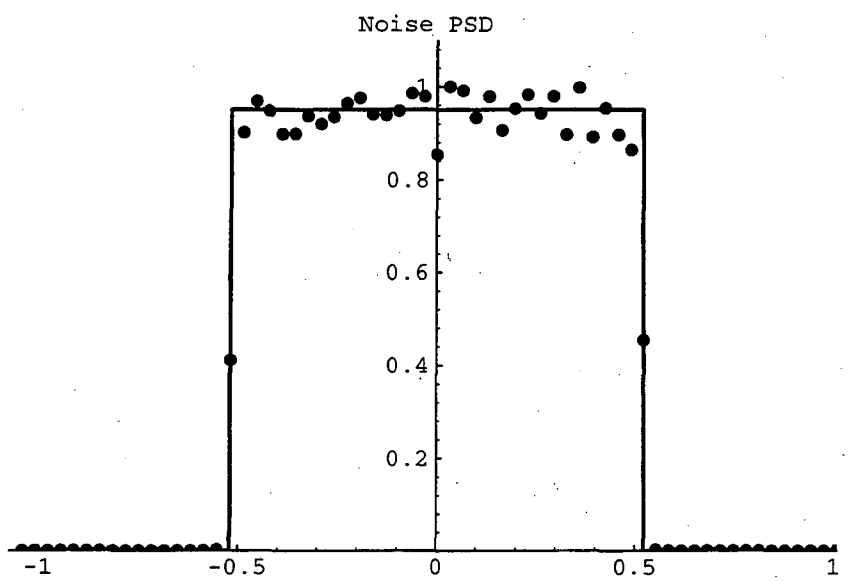

Figure 5.14: Single lag noise power spectral density in one dimension, modulated sinusoidal oscillating gradients. The mean power spectral density of 2048 reconstructions of lag 12 of 128Kpoint data sets containing only measurement noise is shown as points. The sampling density of the modulated gradient is designed to be flat over the sampled region, so the noise PSD should also be flat. γ is 11.24 MHz/T, G is 8 mT/m, T_R is 75μs f_o is 546.977Hz.

Chapter 6

Experimental Studies

6.1 Introduction

This chapter will present the results of stochastic NMR experiments on phantoms to demonstrate the performance of various aspects of the technique. The first set of phantom experiments, described in Section 6.2, demonstrates spectroscopic imaging on a ^{1}H phantom. The second set of experiments, shown in Section 6.3 shows the performance of the stochastic relaxometric imaging experiment over a range of relaxation rates in a ^{23}Na phantom.

6.2 Spectroscopic Imaging

6.2.1 ^{1}H Phantom Construction

To demonstrate spectroscopic stochastic NMR imaging, a ^{1}H phantom with two chemical shift species was constructed. The outer flood field contains a 4% weight/weight agarose gel, which gives a water signal, and the tubes contain peanut oil, which gives a lipid signal. The resonances are separated by 380Hz at 2.35T (3.8ppm). The flood field is used to show variation in image intensity across the field of view, and the small tubes demonstrate the spatial resolution of the technique. The phantom is diagrammed in Figure 6.1.

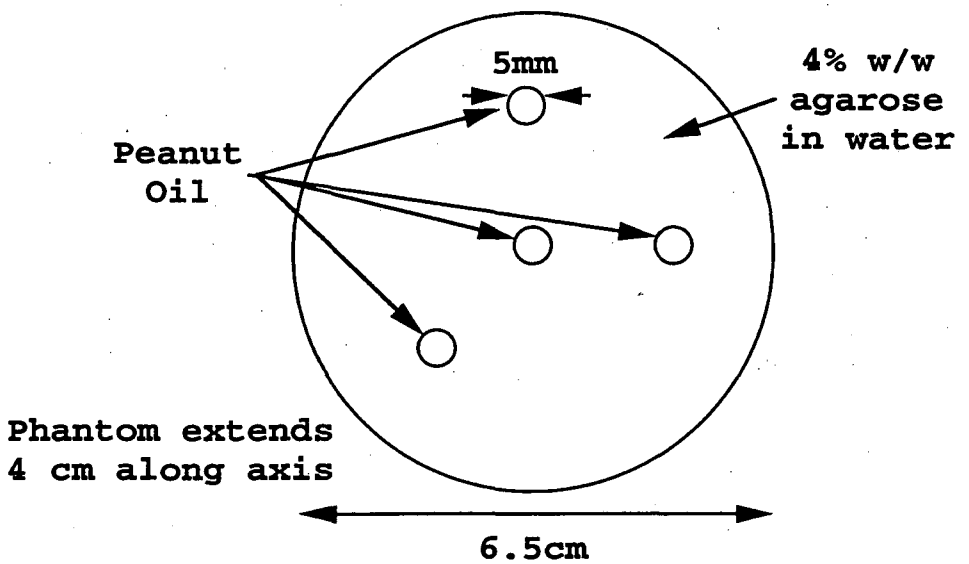

Figure 6.1: ^1H chemical shift phantom

6.2.2 Images

The ^{1}H phantom was imaged using a stochastic imaging sequence with incommensurate frequency oscillating gradients. Both sinusoidal and three component truncated square wave (TSW) gradient experiments are shown. The water and lipid images shown are integrated in the spectral dimension over the water and lipid lines, starting from a spectroscopic image reconstructed on a 64x64x32 grid. The spectra from the central z plane of a 32x32x32 reconstruction are also shown.

The first set of images shows an experiment employing sinusoidal incommensurate frequency gradients. The gradient oscillation frequencies were 547.945, 533.333, and 519.480Hz along x, y and z respectively, with repeat times of 73, 75, and 77 points. The gradient amplitudes along x, y, and z were G = 4.00, 3.89, and 3.79mT/m respectively, giving a point spread function halfwidth of 0.68cm along each axis. The voxel volume resulting from the PSF halfwidth is 0.314cm^3. The RF sequence employed quadrature phase modulation, with the phase selected by two 19 bit MLS's, and the T_R for the experiment was 50μs. The 421,575 point k trajectory was repeated 16 times, for a total of 6,745,200 data points. The total imaging time was 337.26 seconds.

The spectral field of view (FOV) equals the lowest gradient frequency, 519.48Hz. 1024 lags were reconstructed, giving a spectral resolution of 19.6Hz. The water and lipid images are shown in Figure 6.2, reconstructed on a 64x64x32 grid (FOV = 8.68cm in x and y, and 13.13cm in z). The spectra of the central z plane of a 32x32x32 reconstruction (FOV = 4.92cm in x, y, and z) are shown in Figure 6.3.

The second set of images is of the same phantom, using incommensurate frequency three component TSW gradients. The same gradient oscillation frequencies were used, 547.945, 533.333, and 519.480Hz along x, y and z respectively, with repeat times of 73, 75, and 77 points. $G = 4.00$, 3.89, and 3.79mT/m along x, y, and z. The point spread function halfwidth using the TSW gradients is 0.54cm along each axis, giving a voxel volume of 0.155cm^3. The same RF sequence and T_R were used. The 421,575 point k trajectory was repeated 8 times, for a total of 3,372,600 data points. The total imaging time was 168.63 seconds.

The reconstruction parameters were the same as in the first experiment. The spectral field of view is 519.48Hz, 1024 lags were reconstructed, and the spectral resolution is 19.6Hz. The water and lipid images (64x64x32 voxels, (FOV = 8.68cm in x and y, and 13.13cm in z) are shown in Figure 6.4, and the spectra of the central z plane are shown in Figure 6.5 (32x32x32, FOV = 4.92cm along each axis).

The difference in subjective image appearance between the two sets of ^{1}H images arises from differences in both the gradient trajectories employed and in experimental conditions. As seen in Section 5.6, the noise characteristics of the two experiments are quite different; the sinusoidal gradient experiment has measurement and systematic noise concentrated in the low spatial frequencies compared to the TSW experiment, which has a much flatter noise spectrum. No spatial filtering or overall sampling density correction was performed on these images, so the differences in noise power distribution are apparent.

In addition, the second experiment (using TSW gradients) was performed after an enhanced shimming method was developed for our spectrometer based on spectroscopic imaging. The differences in the quality of the shimming in the outer voxels can be seen by comparing the spectral plots of the two images. Intravoxel dephasing in the outer voxels is likely to be responsible for some of the inhomogeneity in the si-

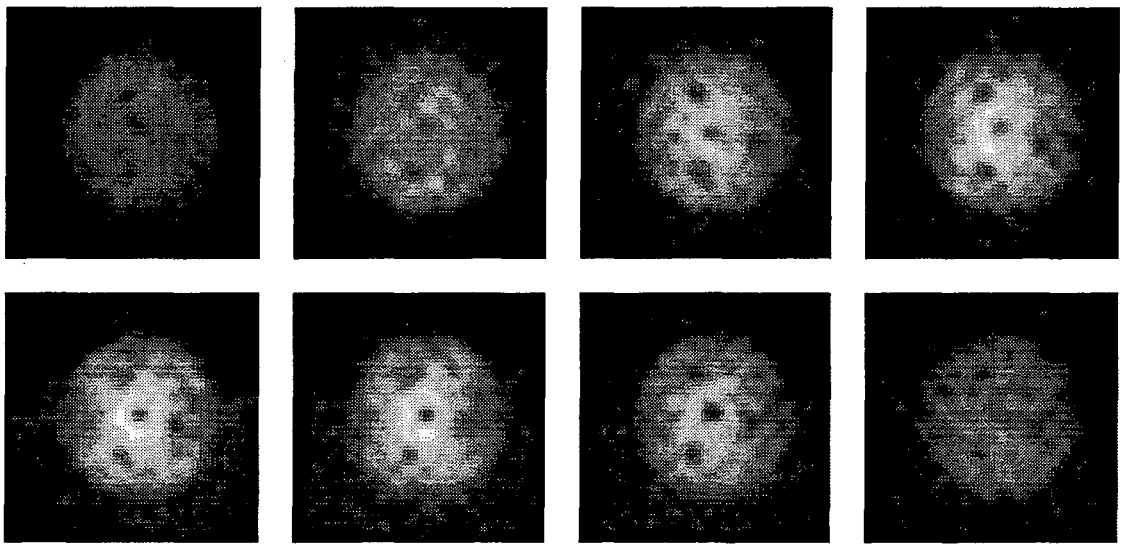

(a) Images integrated over the water line

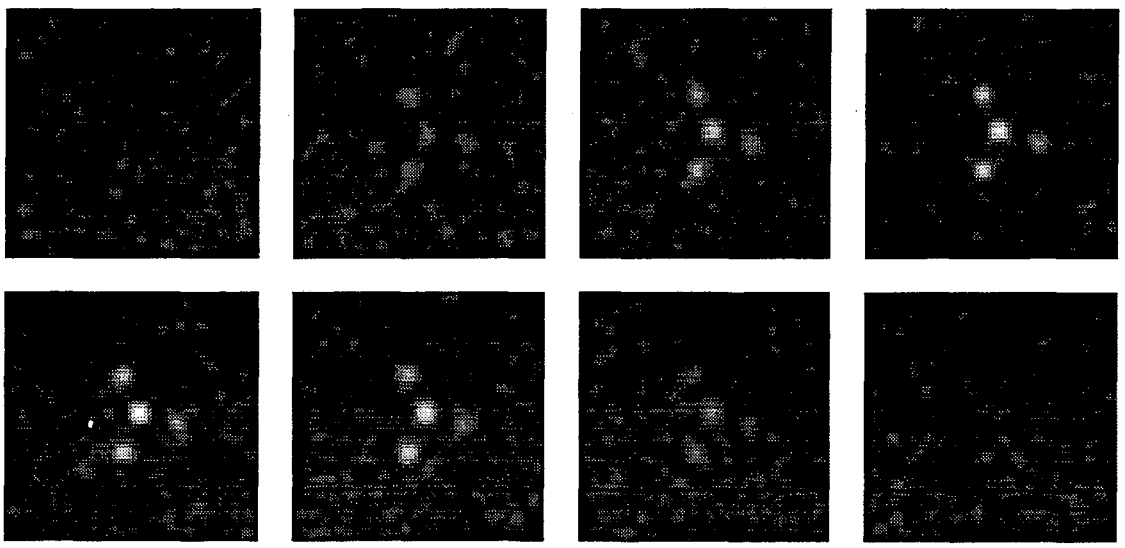

(b) Images integrated over the lipid line

Figure 6.2: Spectroscopic images of the ^1H tube phantom using sinusoidal gradients. The 8 central planes in z are shown left to right, top to bottom of the water (top) and lipid (bottom) lines from the phantom. The field of view is 8.68cm in x and y. The voxel halfwidth is 0.68cm.

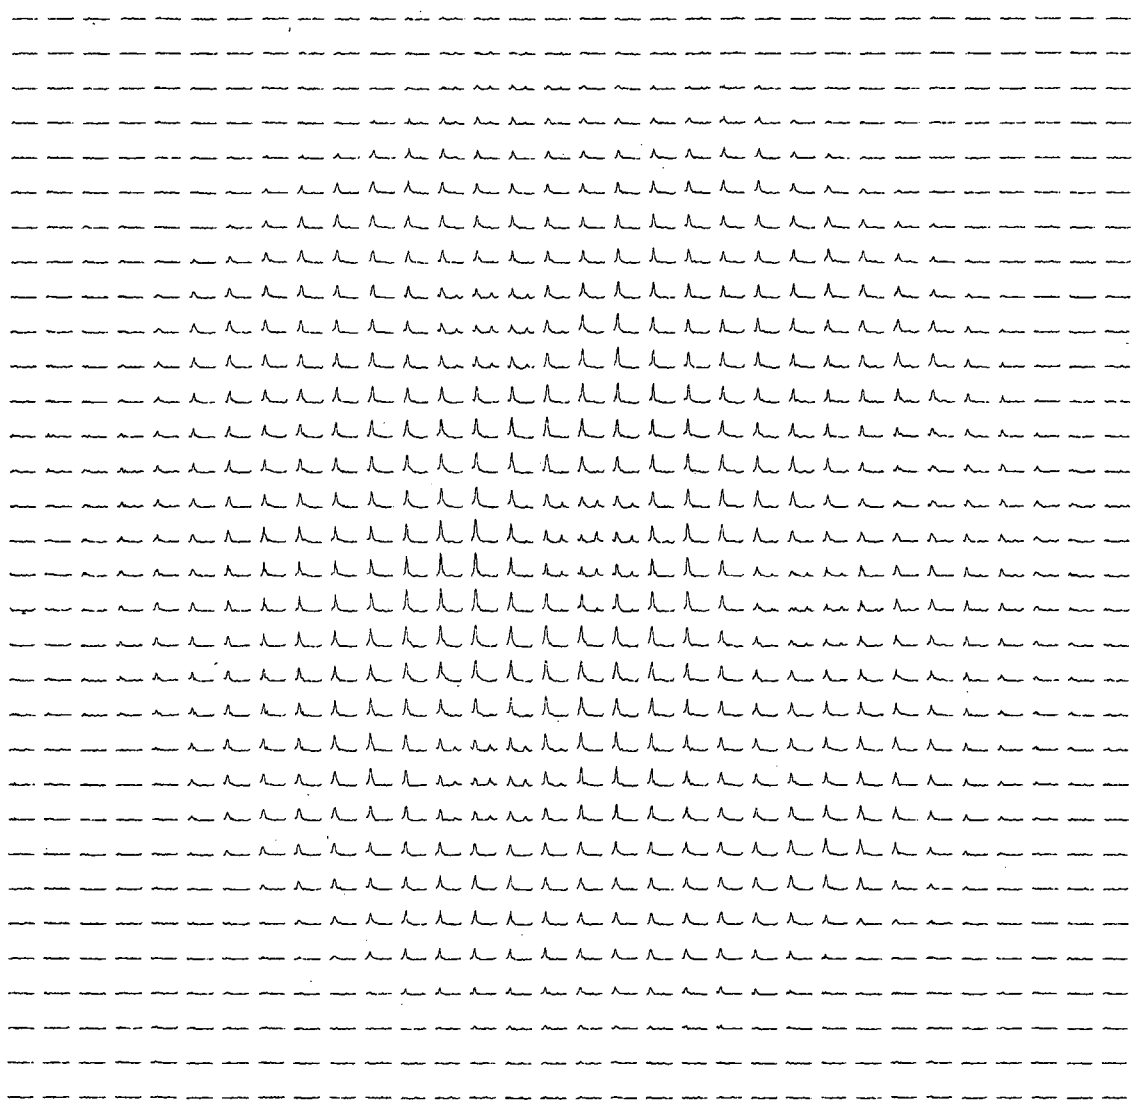

Figure 6.3: Spectroscopic image of the central z plane of the ^1H tube phantom using sinusoidal incommensurate frequency oscillating gradients. The data are presented as a 32 by 32 grid of spectra. FOV = 4.92cm in x and y.

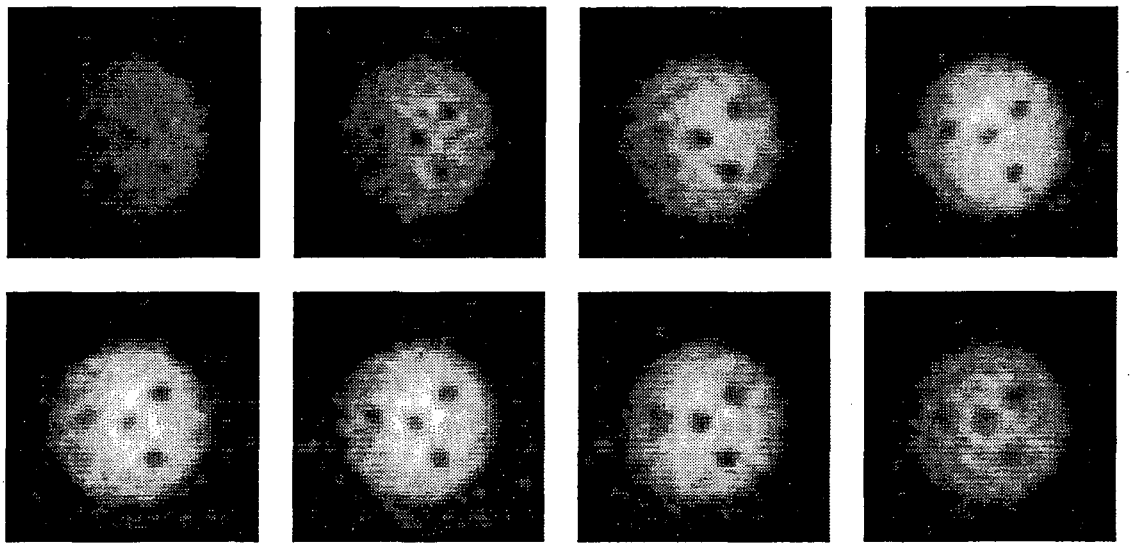

(a) Images integrated over the water line

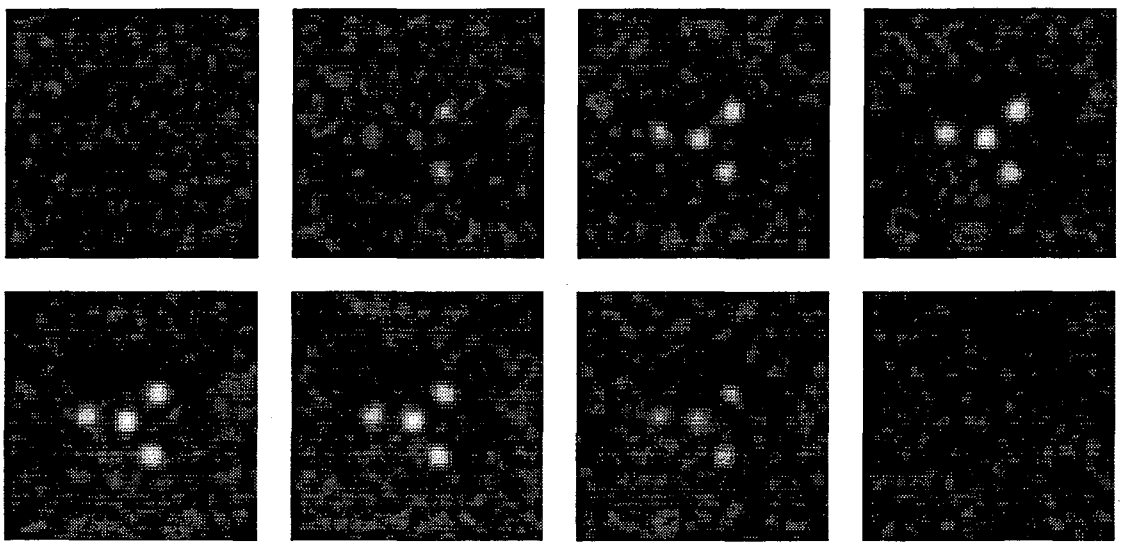

(b) Images integrated over the lipid line

Figure 6.4: Spectroscopic images of the ^1H tube phantom using three component TSW gradients. The central 8 slices in z are shown left to right, top to bottom of the water (top) and lipid (bottom) lines from the phantom. The field of view is 8.68cm in x and y. The voxel halfwidth is 0.54cm.

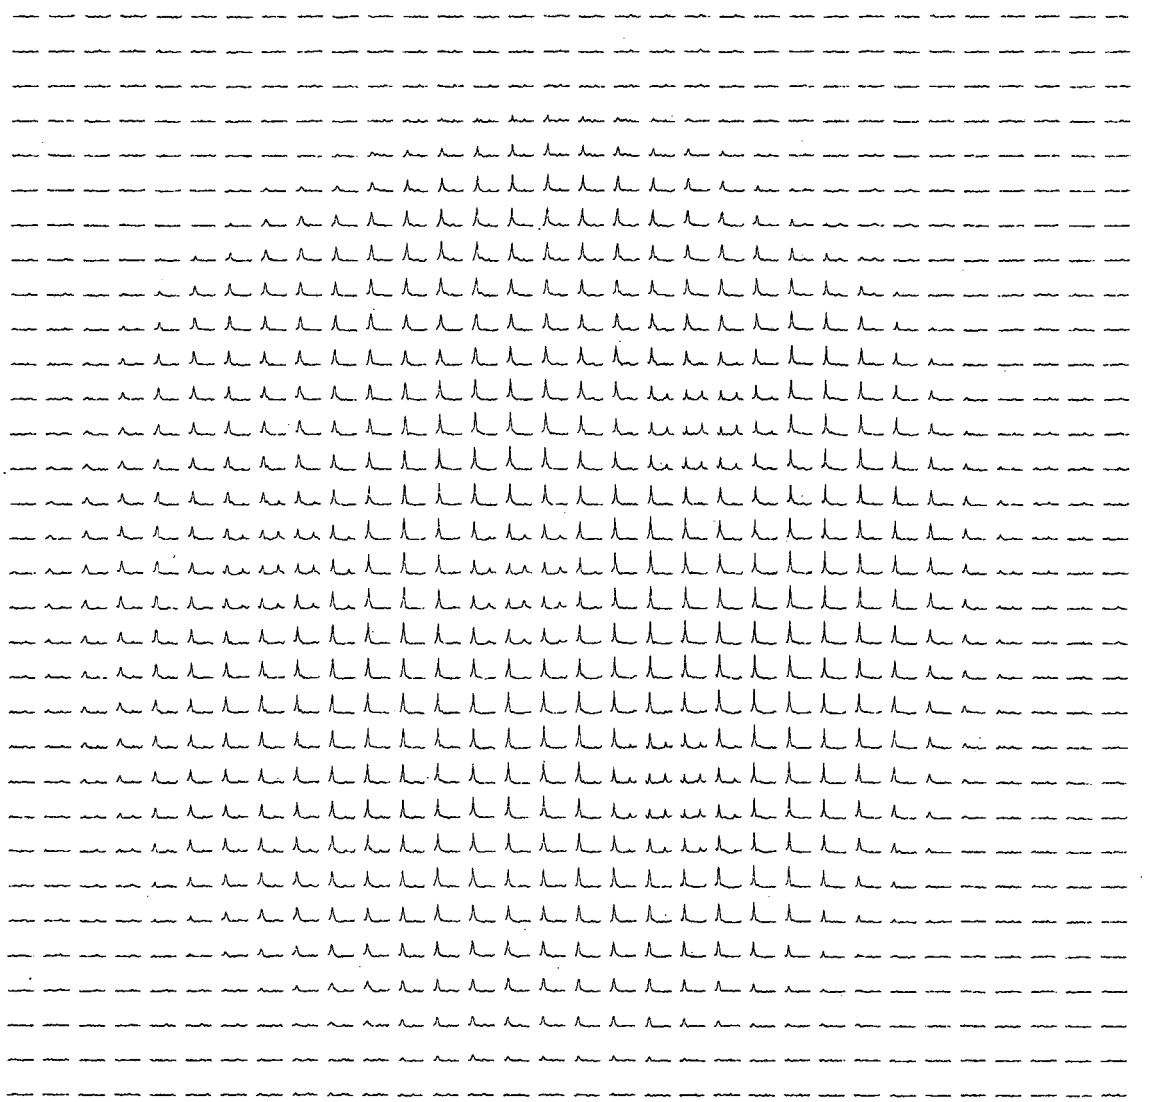

Figure 6.5: Spectroscopic image of the central z plane of the ^1H tube phantom using three component TSW gradients. The data are presented as a 32 by 32 grid of spectra. FOV = 4.92cm in x, y and z.

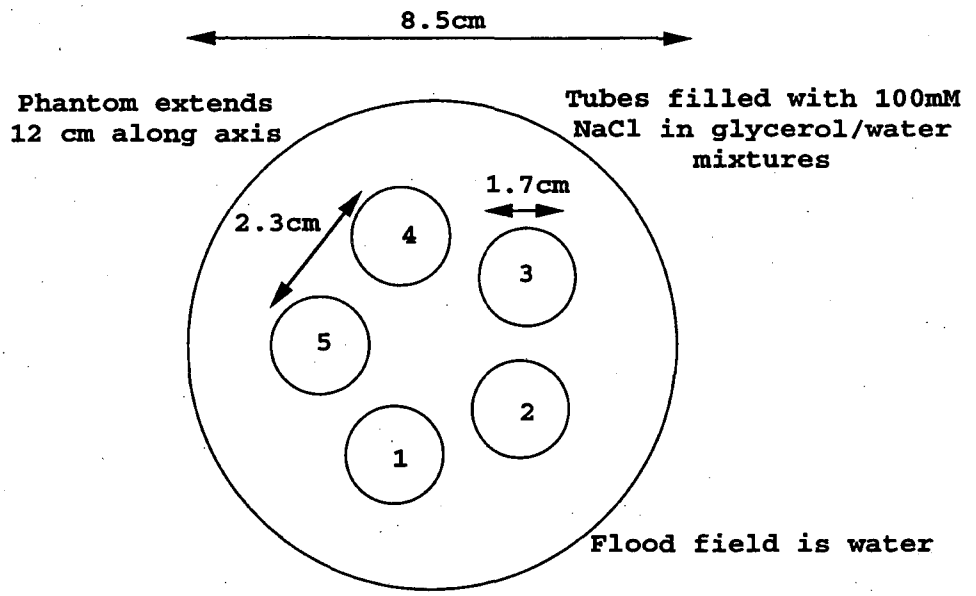

Figure 6.6: ^{23}Na T_2 phantom. The phantom is cylindrically symmetric along the z axis. Tubes 1-5 contain 100mM NaCl in 86%,77%,61%,45%, and 0% glycerol v/v in water respectively.

nusoidal gradient image. However it is important to note that for the same gradient amplitude, the voxel size of the TSW experiment is approximately half that of the similar sinusoidal gradient experiment, so the TSW experiment is less susceptible to intravoxel dephasing *a-priori*.

6.3 T_2 Sensitivity

6.3.1 ^{23}Na T_2 Phantom Construction

A ^{23}Na phantom was constructed to determine the performance of the stochastic experiment sensitivity over a range of relaxation rates. The cylindrical phantom consists of 5 1.7cm diameter tubes arranged in a pentagonal formation, separated by 2.3cm, each containing 100mM NaCl with different T_2 values. The flood stage was filled with water to reduce susceptibility artifacts. The phantom is diagrammed in Figure 6.6.

The solvent in each tube is a volume/volume mixture of glycerol and water. In-

Percentage glycerol (v/v)	Tube	T_1(ms)	T_2(ms)
0%	5	60.6	59.3
45%	4	11.3	11.2
61%	3	4.4	4.4
77%	2	1.2	1.3
86%	1	0.60	0.64

Table 6.1: ^{23}Na T_1 and T_2 relaxation times of 100mM NaCl in glycerol/water solutions at 2.35T.

creasing the concentration of glycerol increases the viscosity of the medium, and decreases the T_2 of the ^{23}Na in the solution. The sodium spectrum of all the solutions used is in the extreme narrowing limit, which means that the T_2 relaxation is monoexponential. The five solutions were formulated to span a wide range of T_2 values. The T_2 values of the solutions was measured using a spin-echo experiment (90_x-τ-180_y-τ-acq). Subsequent acquisitions employ different τ values to map out the FID. The T_1 was measured using an inversion recovery experiment (180_x-τ-90_x-acq), again ranging over various τ values, and both curves were fit to obtain the T_1 and T_2.

The results of these measurements are presented in Table 6.1. As expected for ^{23}Na in the extreme narrowing limit, the T_1 and T_2 values are identical for all of the solutions. Due to finite pulse lengths and delays in the receiver system, the measurement technique breaks down at very short T_1 and T_2 values, leading in some cases to measured T_1 values that are shorter than T_2. The five tubes in the phantom cover two orders of magnitude in T_1 and T_2. The relaxation times are likely to be somewhat shorter *in vivo*.

The stochastic imaging experiment measures T_2^* rather than T_2. Because the T_2 values of ^{23}Na are rather short to begin with, and the gyromagnetic ratio is not very high, the difference between T_2 and T_2^* is not very significant at many of the T_2 rates observed *in vivo*. T_2^* values of a 4cm diameter spherical sample of the glycerol water solutions were obtained by fitting the FIDs obtained using a stochastic spectroscopy experiment. The reconstructed FIDs are shown in Figure 6.7, and the T_2^* values are presented in Table 6.2.

Percentage glycerol (v/v)	T_2^*(ms)
0%	49.6
45%	10.5
61%	4.3
77%	1.2
86%	0.67

Table 6.2: ^{23}Na T_2^* relaxation rates of 100mM NaCl in glycerol/water solutions in a 4cm spherical sample at 2.35T measured using stochastic spectroscopy.

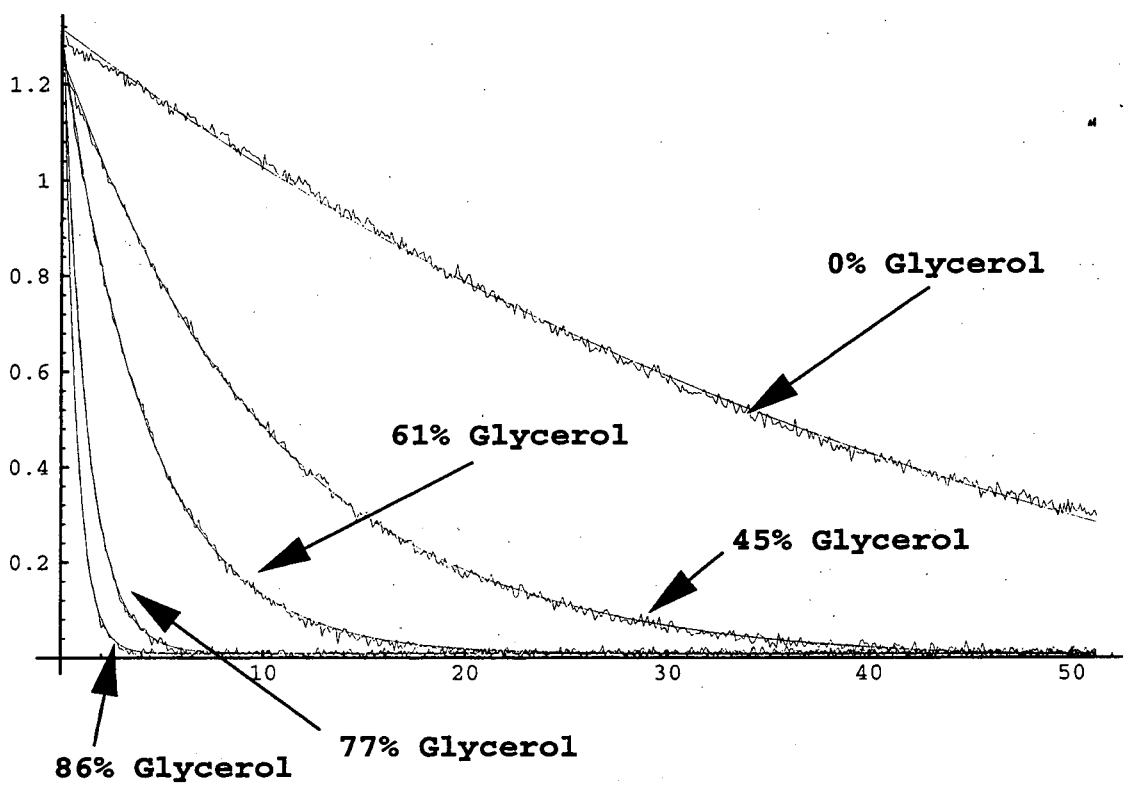

Figure 6.7: ^{23}Na T_2^* relaxation rates of 100mM NaCl in glycerol/water solutions in a 4cm spherical sample at 2.35T. These FIDs were obtained using a stochastic spectroscopy experiment. $T_R = 75\mu s$, 524,287 points were recorded. Time axis is in milliseconds. The FID amplitudes are scaled separately for clarity.

6.3.2 Images

The T_2 phantom was imaged using a stochastic image sequence with incommensurate frequency oscillating gradients. Again, both sinusoidal and three component truncated square wave (TSW) gradient experiments are shown. All reconstructions were done on a 32x32x32 grid, and the data are reconstructed in the time domain as FIDs rather than spectra.

The first set of ^{23}Na images shows an experiment employing sinusoidal incommensurate frequency gradients. The gradient oscillation frequencies were 547.945, 533.333, and 519.481Hz along x, y and z respectively, with repeat times of 73, 75, and 77 points. The gradient amplitudes along x, y, and z were $G = 8.00$, 7.79, and 7.59mT/m respectively, giving a point spread function halfwidth of 1.28cm along each axis. The voxel volume resulting from the PSF halfwidth is 2.11cm^3. The RF sequence employed quadrature phase modulation, with the phase selected by two 19 bit MLS's, $\alpha = 5^o$, and the T_R for the experiment was 75μs. The 421,575 point k trajectory was repeated 16 times, for a total of 6,745,200 data points. The total imaging time was 505.89 seconds.

Two reconstructions were performed. Figure 6.8 shows spatial images of the central z plane of the phantom at sequential time points of the FID. 1046 time lags were reconstructed, with a time resolution of 1.82ms. The first time point is centered at 912μs. The total field of view along the time axis is 77.35ms. 32x32x32 spatial voxels were reconstructed, with a field of view of 6.19cm along each axis. Figure 6.9 shows the FIDs of the central z plane of 32x32x32 voxels. 523 time lags were reconstructed, with a time resolution of 912μs. The total field of view along the time axis is 38.68ms. The first time point is centered at 456μs. The spatial field of view is again 6.19cm along each axis.

The same experiment was performed with incommensurate TSW gradients. Identical gradient frequencies and G values were used. The voxel halfwidth was 1.11cm, giving a voxel volume of 1.39cm^3. The 421,575 point k trajectory was repeated 12 times, for a total of 5,058,900 data points. The total imaging time was 379.42 seconds. Identical reconstruction parameters were used. The resulting FID time point

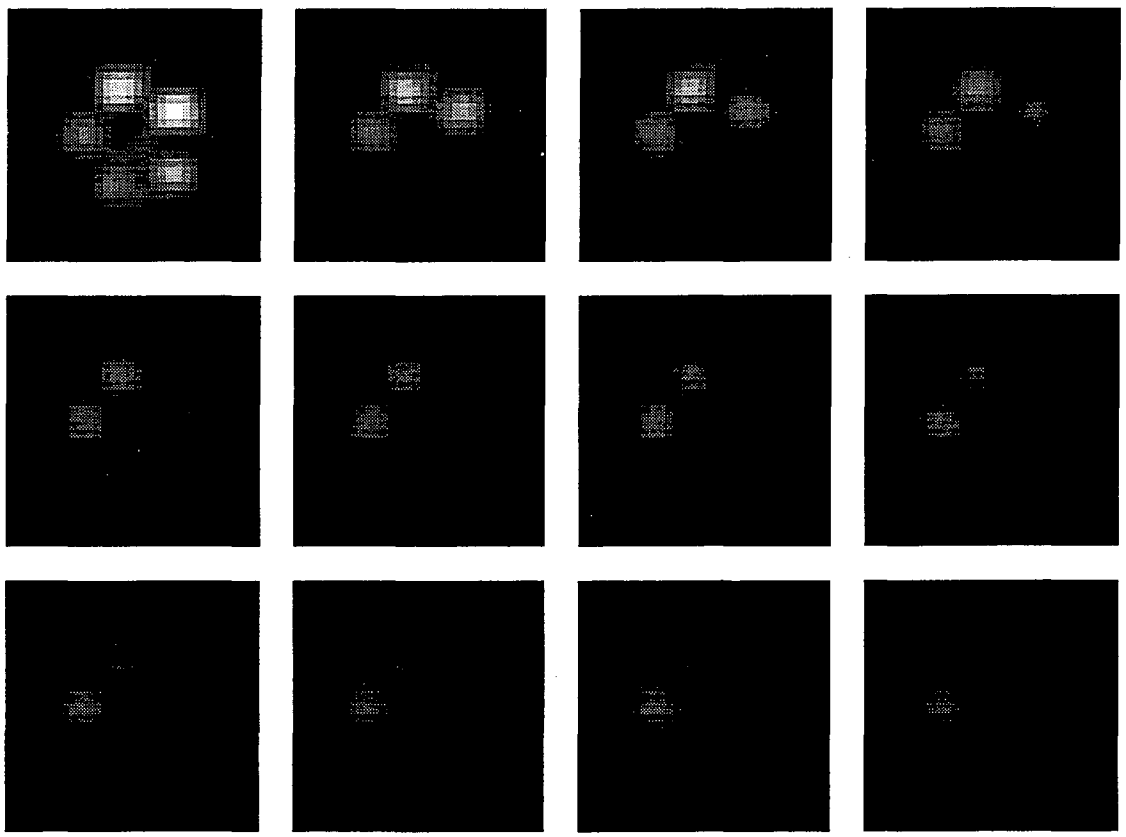

Figure 6.8: Images of the central z plane of the ^{23}Na T_2 *phantom using incommensurate frequency sinusoidally oscillating gradients. The images show different time points of the FIDs, presented as images, left to right, top to bottom. The first image is of the time point centered at 912μs. The time spacing between succeeding images is of 1.825ms. FOV = 6.19cm along each axis.*

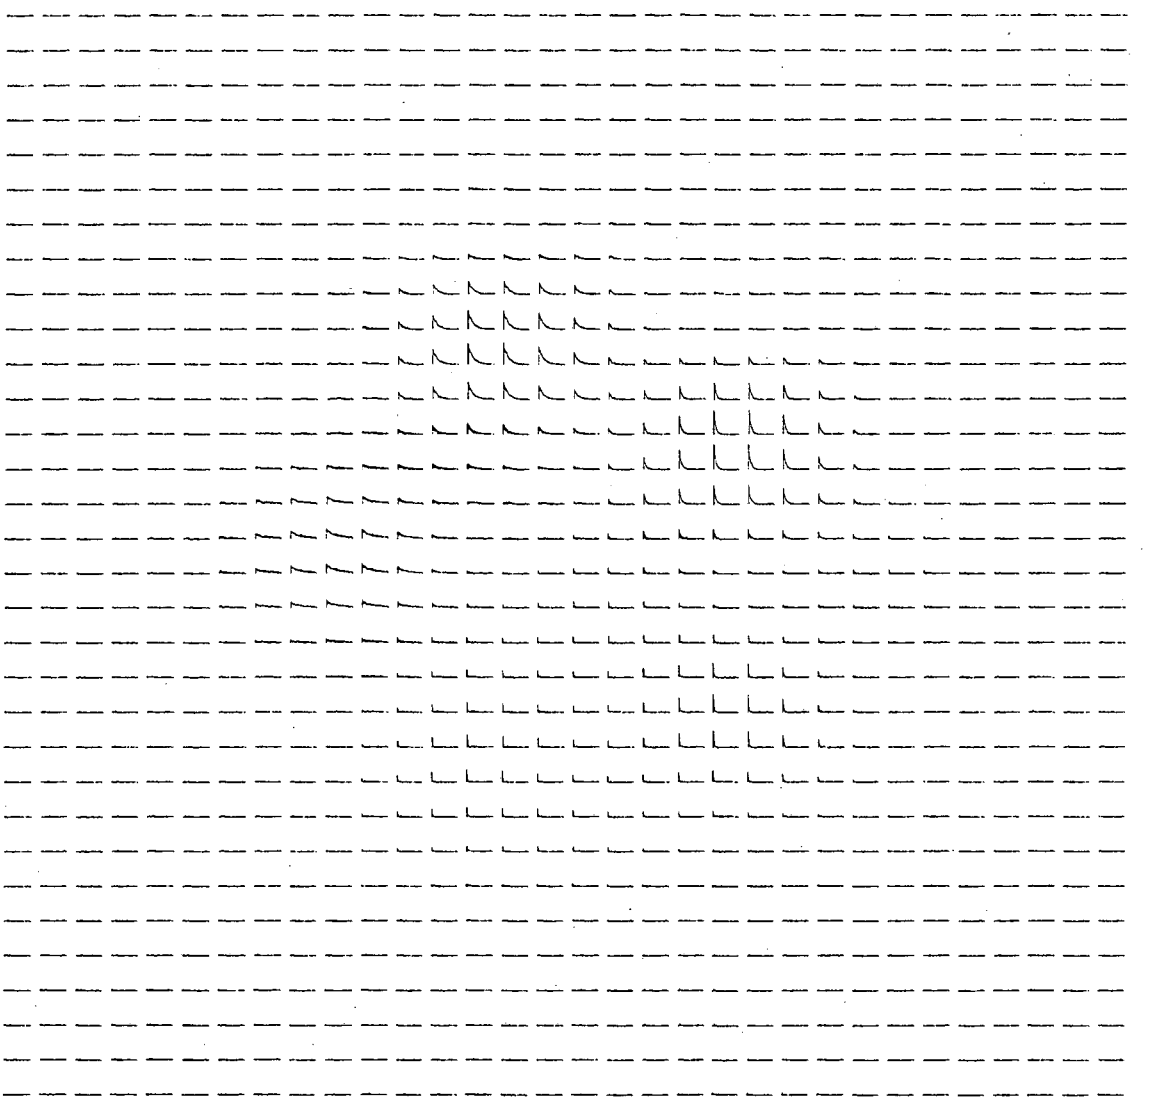

Figure 6.9: FID image of the central z plane of the ^{23}Na T_2 phantom using incommensurate frequency sinusoidally oscillating gradients. The data are presented as a 32 by 32 grid of FIDs. The first time point is at 456μs, and the spacing between time points is 912μs. The spatial FOV = 6.19cm along each axis.

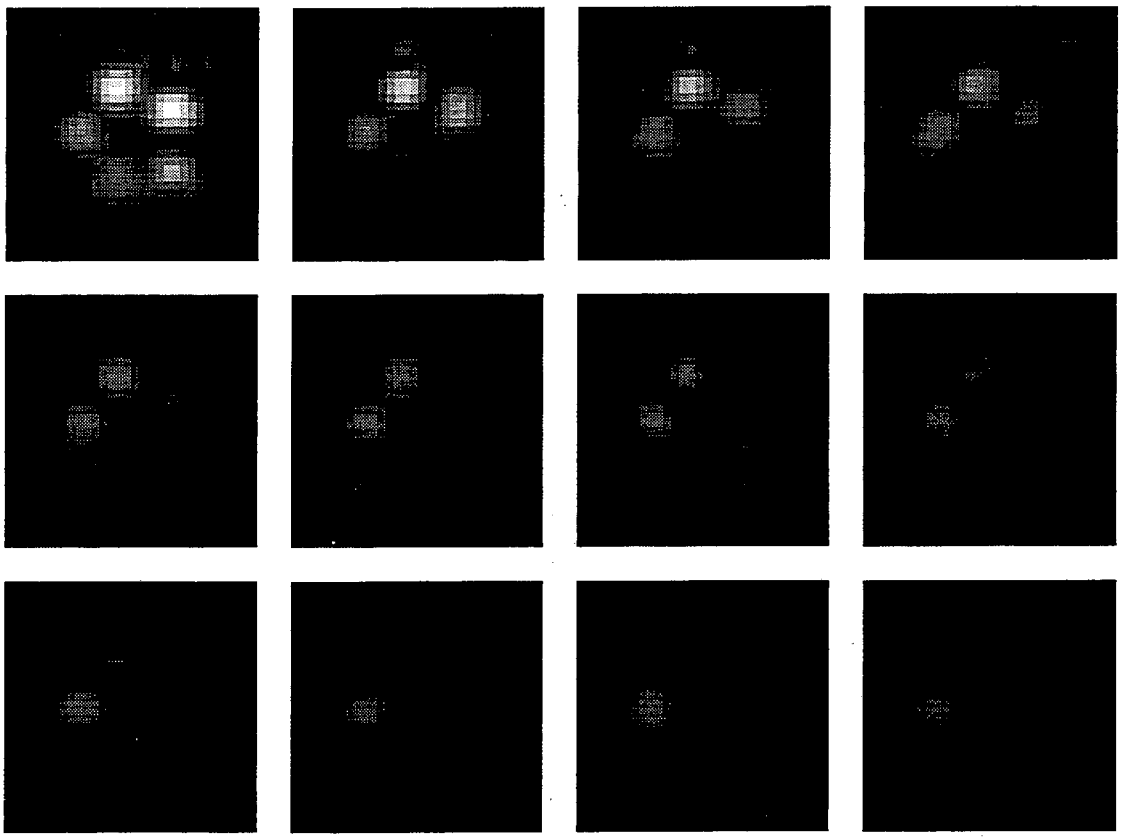

Figure 6.10: Images of the central z plane of the ^{23}Na T_2 phantom using incommensurate frequency three component truncated square wave gradients. The images show different time points of the FIDs, presented as images, left to right, top to bottom. The first image is of the time point centered at 912μs. The time spacing between succeeding images is of 1.825ms. FOV = 6.19cm along each axis.

images are shown in Figure 6.10, and the array of FIDs is shown in Figure 6.11.

A large fraction of the overall intensity variation of the tubes in the image is due to T_1 contrast (since this is a spectroscopic image method, the T_2 contrast can be determined *a postiori*). As shown in Equation 4.4, the overall system response depends on T_1, T_R, and the RMS flip angle α. The T_1 values in the phantom span a wide range. α was chosen to optimize the experiment for a T_1 in the middle of this range. Tubes 3 and 4, which have intermediate T_1 values are therefore the brightest. Another source of apparent intensity variation is the decrease in intensity for very short T_2 species. This is an image artifact; all of the signal from the very short T_2 nuclei is captured by the stochastic experiment; however the density correction

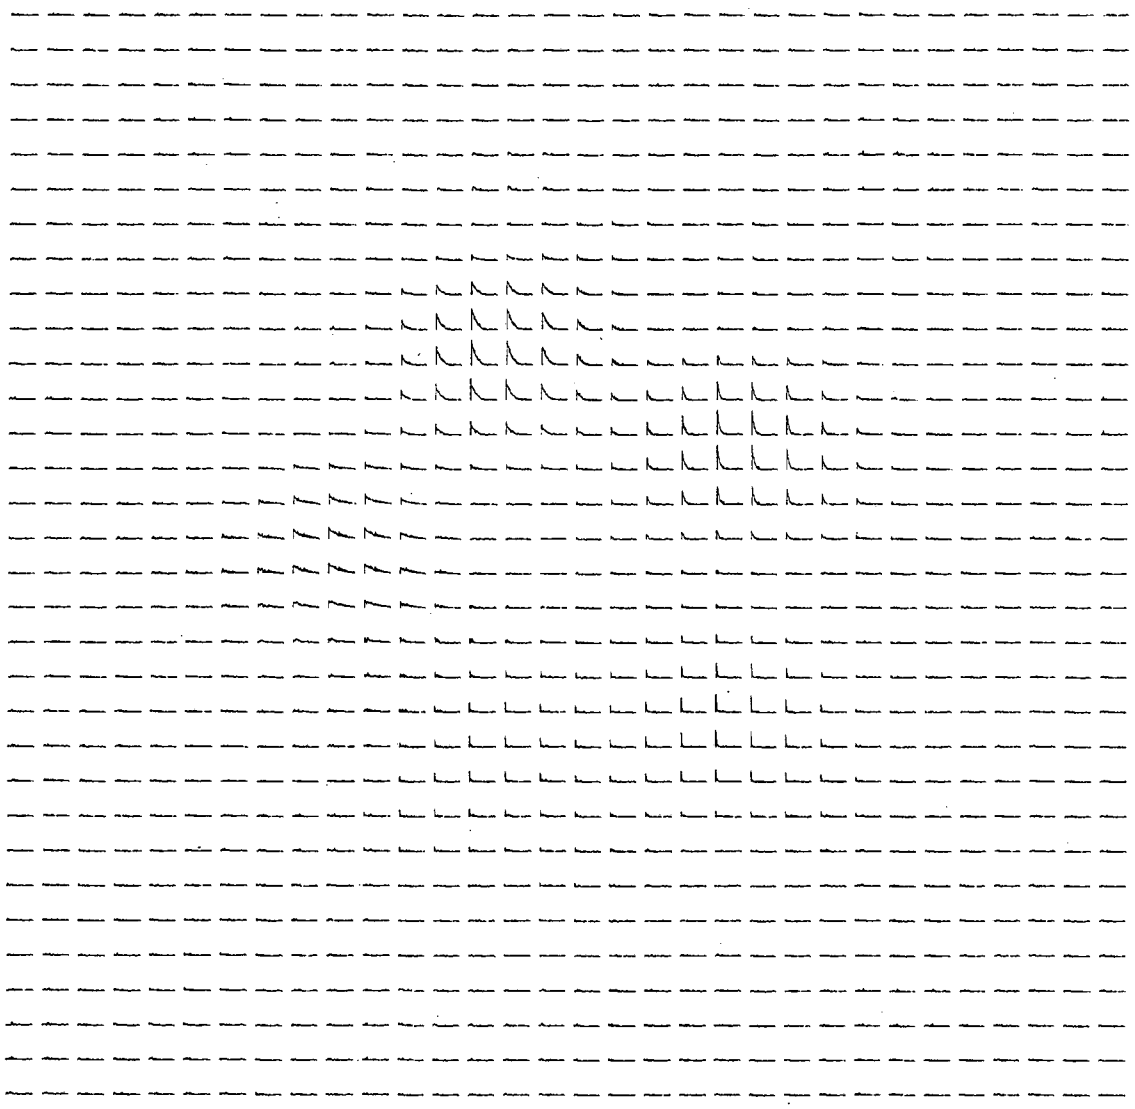

Figure 6.11: FID image of the central z plane of the ^{23}Na T_2 phantom using incommensurate frequency three component TSW oscillating gradients. The data are presented as a 32 by 32 grid of FIDs. The first time point is at 456μs, and the spacing between time points is 912μs. The spatial FOV = 6.19cm along each axis.

Percentage glycerol (v/v)	T_2^*(ms) from stochastic spectroscopy	T_2^*(ms) from stochastic imaging
0%	49.6	22.0
45%	10.5	7.6
61%	4.3	3.1
77%	1.2	1.1
86%	0.67	0.68

Table 6.3: Comparison of the ^{23}Na T_2^* relaxation rates of 100mM NaCl in glycerol/water solutions at 2.35T measured using stochastic spectroscopy and measured from the stochastic image in Figure 6.9.

algorithm employed for these images has a tendency to underemphasize the data from the center of k-space, which is where the signal energy from nuclei with T_2 shorter than one half of a gradient period is concentrated.

These uncorrected effects make it difficult to perform proper quantitative relaxometry. However, a "quick look" T_2 fit to the data indicates the promise of the technique. The FID from the central voxel of each of the tubes was selected from the data set shown in Figure 6.9, and fit to determine the T_2 values. The results are shown in Table 6.3. The anomalously low value of the 0% glycerol T_2^* may be due to shimming.

Chapter 7

Summary

This dissertation presents a fast three dimensional spectroscopic NMR imaging experiment which provides spatially resolved information on ^{23}Na relaxation parameters. This technique may prove useful for medical investigations of brain function. An experiment based on stochastic RF excitation and employing oscillating gradients provides both the fast encoding time necessary to make high resolution *in vivo* spectroscopic images and the sensitivity to the wide range of T_2 values exhibited by ^{23}Na *in vivo*.

Stochastic NMR is not a new technique; it was originally proposed in 1970, and has been developed over the years by a number of researchers. However, the development of spectroscopic imaging with oscillating gradients by Roos and Wong opened the door to its use in biological imaging. My contribution in this thesis has been to address the issues necessary to adapt the technique to imaging of short T_2 nuclei.

The important theoretical contribution of this thesis is the development of the fast reconstruction algorithm which makes three dimensional imaging practical. The speed increase of the Fourier interpolation stochastic reconstruction algorithm reduces the reconstruction time for typical spectroscopic data sets to reasonable values, a few hours on a general purpose scientific workstation. Further optimization of the reconstruction code, or the use of faster or more specialized processors can be expected to provide further speed increases.

Another benefit of the new reconstruction algorithm is that it has greatly simplified the analysis of sampling density and noise propagation in the stochastic experiment. This led to the development of the new gradient trajectories described in Chapter 3.2.3, the truncated square wave gradients and the amplitude modulated sinusoidal gradients. These gradient trajectories provide more uniform k-space sampling densities than those afforded by ordinary sinusoidal gradients, while still preserving the advantages of steady state eddy currents. The AMS gradients have uses beyond stochastic NMR; echo planar trajectories based on sinusoidal readout gradients can employ the AMS modulation function to gain a uniform k-space sampling density without resorting to nonlinear time sampling.

The reduced reconstruction time afforded by the new reconstruction algorithm permitted extensive simulation of the experiment, which in turn led to the exploration factors that affect image quality, and discoveries such as the improvement in reconstructed image quality through the use of repeating versus non-repeating k-space trajectories, and the use of analytic instead of empirical sampling density correction.

Implementing the stochastic experiment for short T_2 nuclei led to the development of a novel Q-spoiler circuit. The probe Q-spoiler can be used in any type of NMR experiment where dead time caused by probe ringdown time needs to be minimized. Unlike previously designed PIN diode circuits for Q-spoiling, the inductive coupled design can be added to any existing probe with a minimum of modification. Also, in the non-energized state, the circuit causes no significant changes to the probe performance. Previous designs required active high voltage back-biasing of the PIN diode to keep the probe in the high Q state, which injects unwanted noise into the probe. Q-spoilers of the design described in this thesis are already in use in other non-stochastic short T_2 experiments being done in our lab.

In the course of implementing the experiment, other important practical lessons were learned, such as the necessity of measuring and compensating for gradient phase shifts and receiver delay times, factors which can usually be ignored for conventional imaging experiments, but which can significantly degrade stochastic images if left uncorrected.

The imaging method developed here combines short encoding time, on the order of 20 to 30 seconds, with sensitivity to a wide range of T_2 values. These qualities make it a promising method for imaging ^{23}Na *in vivo*. The gradient, RF, and sampling frequency requirements are within the capabilities of modern clinical echo-planar imaging systems.

As an example of the requirements of a clinical experiment, consider the problem of making a three dimensional spectroscopic sodium image of the human head at with 1cm^3 voxels. If we design an incommensurate frequency truncated square wave gradient experiment (which is very close to a trapezoidal gradient, available on many echo-planar machines) using 1kHz gradient frequencies, we require a gradient amplitude of $G = 14.7\text{mT/m}$. To provide a field of view of 24cm requires a sampling frequency of 41,700Hz. Because of the high sampling rate, extremely small RF pulses (less than 1°) can be employed. These hardware requirements are not unreasonable for a clinical EPI machine. Whether or not the software requirements for RF and gradient waveform generation can be met depend on the flexibility of the control console. If necessary, the waveforms could be generated and fed in from an external processor.

There are a number of ways in which the stochastic experiment can be further improved. Image reconstruction could be improved by modifying the density correction algorithm that is applied to the gridded data, as mentioned in Section 5.5.2. First, I have not yet found a computationally efficient way to calculate a general analytic expression for the sampling density of the truncated square wave gradients for anything other than the $i_{max} = 1$ (sinusoidal gradients) case. The empirical sampling density correction introduces artifacts into the single lag point spread function. Finding an analytic expression for the sampling density would improve the resulting image quality.

Perhaps more important is the fact that the current reconstruction technique applies density correction in the spatial domain alone. Because different time lags in the reconstruction cover different extents in k-space, there are different numbers of time samples taken at different k-space locations. Because the gradients are periodic, and k,t-space is critically sampled in time, this results primarily in a change in the

spatial point spread function (discussed in Section 4.5.4). However, it also tends to deemphasize the first lags in the reconstruction, which may affect the overall signal intensity of very short T_2 nuclei, especially for the first point of the FID. This effect must be taken into account for accurate relaxometry. Performing the time-space density correction requires development of an analytic function for the overall time-space density, a task which has so far eluded me.

Accurate, quantitative analysis of concentrations and relaxation parameters in a voxel will require a method for calibrating RF flip angle throughout the image, so that the effect of T_1 contrast can be considered. In a sample where there is a wide range of T_1, as was the case in the ^{23}Na tube phantom, the intensity variation can be large, and must be taken into account.

There are also of ways the implementation of the experiment can be improved. Paff has shown for spectroscopy that oversampling in the time domain (i.e. having a T_R much smaller than that required by the spectral bandwidth of the system) can be used to reduce the systematic noise in stochastic spectra while leaving the measurement noise unchanged by improving the cancellation of the excitation noise sequence[43]. This technique can also be applied to stochastic imaging, and would be expected to greatly improve image quality, at the cost of additional reconstruction time.

Finally, a promising approach is to develop a hybrid technique which takes many data samples for each RF pulse. If capturing extremely short T_2 signals is not crucial, spacing the excitation pulses farther apart would allow for extremely fast sampling, since there is no need to wait for probe relaxation between each sample. This would allow finer sampling of k-space. It would also allow the use of excitation sequences with much shorter repetition cycles, since there would be fewer RF pulses within $5T_1$ (see Section 3.2.2). As a result, it may be possible to use a full length excitation sequence at each point in k-space. Wong showed[14] that this causes a significant decrease in the amount of systematic noise (on the order of 20dB) in the resulting spin density estimate, a result which has been demonstrated in stochastic imaging with constant gradients (Wong, personal communication). Also, the reconstruction for such images should be much faster than for "conventional" stochastic imaging,

because many of the calculations could be done in parallel. This could result in a significantly enhanced fast spectroscopic imaging method.

Bibliography

[1] B. Alberts and Others. *Molecular Biology of the Cell*, page 291. Garland Publishing, 1983.

[2] S.K. Hilal, A.A. Maudsley, J.B. Ra, H.E. Simon, P. Roschmann, S. Wittekoek, Z.H. Cho, and S.K. Mun. In vivo NMR imaging of sodium-23 in the human head. *Journal of Computer Assisted Tomography*, 9:1–7, 1985.

[3] G. A. Elgavish. Shift reagent aided ^{23}Na NMR spectroscopy. *Investigative Radiology*, 24:1028–1033, 1989.

[4] P. J. Cannon, A. A. Maudsley, S. K. Hilal, H. E. Simon, and F. Cassidy. Sodium nuclear magnetic resonance imaging of myocardial tissue of dogs after coronary artery occlusion and reperfusion. *JACC*, 7(3):573–579, 1986.

[5] S. D. Wolff, J. Eng, B. A. Berkowitz, S. James, and R. S. Balaban. Sodium-23 nuclear magnetic resonance imaging of the rabbit kidney in vivo. *American Journal of Physiology*, 258:F1125–F1131, 1990.

[6] F.K. Goodwin and K.R. Jamison. *Manic-Depressive Illness.* Oxford Univ. Press, 1990.

[7] W.D. Rooney and C.S. Springer. The molecular environment of intracellular sodium: ^{23}Na NMR relaxation. *NMR Biomed.*, 4:227–245, 1991.

[8] W.D. Rooney and C.S. Springer. A comprehensive approach to the analysis and interpretation of the resonances of spins 3/2 from living systems. *NMR Biomed.*, 4:209–226, 1991.

[9] D. Burstein and E. T. Fossel. Intracellular sodium and lithium NMR relaxation times in the perfused frog heart. *Magn. Reson. Med.*, 4:261–273, 1987.

[10] C. Labadie, T.M. Button, W.D. Rooney, J.H.L. Lee, and C.S. Springer. Relaxographic imaging. *Book of Abstracts, Soc. Magn. Reson. Med.* 10th *Annual Meeting*, works in progress:1218, 1991.

[11] R. R. Ernst and W. A. Anderson. Application of Fourier transform spectroscopy to magnetic resonance. *Rev. Sci. Instrum.*, 37:93–102, 1966.

[12] R. R. Ernst. Magnetic resonance with stochastic excitation. *J. Magn. Reson.*, 3:10–27, 1970.

[13] R. Kaiser. Coherent spectrometry with noise signals. *J. Magn. Reson.*, 3:28–43, 1970.

[14] S.T.S. Wong. *Discrete Analysis of Stochastic NMR.* PhD thesis, University of California, Berkeley, Dec. 1988.

[15] B. Blümich. Stochastic NMR imaging. *J. Magn. Reson.*, 60:37–45, 1984.

[16] B. Blümich and H. W. Spiess. NMR imaging with incommensurate sampling and gradient modulation rates. *J. Magn. Reson.*, 66:66–73, 1986.

[17] M.S. Roos and S.T.S. Wong. Spatial localization in stochastic NMR imaging with oscillating gradients. *J. Magn. Reson.*, 87:554–566, 1990.

[18] J.D. O'Sullivan. A fast sinc function gridding algorithm for Fourier inversion in computer tomography. *IEEE Trans. Med. Imag.*, 4(4):200–207, 1985.

[19] J.I. Jackson, C.H. Meyer, D.G. Nishimura, and A. Macovski. Selection of a convolution function for Fourier inversion using gridding. *IEEE Trans. Med. Imag.*, 10(3):473–478, 1991.

[20] F.A. Bovey, L. Jelinski, and P.A. Mirau. *Nuclear Magnetic Resonance Spectroscopy.* Academic Press, second edition, 1988.

[21] B.deB. Frederick, S.T.S. Wong, and M.S. Roos. A gradient trajectory for 3D stochastic NMR imaging with uniform sampling density. *Book of Abstr., Soc. Magn. Reson.* 2nd *Ann. Meeting*, page 412, 1994.

[22] F.G. Stremler. *Introduction to communication systems.* Addison-Wesley, second edition, 1982.

[23] J. Murphy-Boesch, R. Srinivasan, L. Carvajal, and T.R. Brown. Two configurations of the four-ring birdcage coil for 1H imaging and 1H-decoupled 31P spectroscopy of the human head. *J. Magn. Reson.*, 103(2):103–14, 1994.

[24] K. E. Kisman and R. L. Armstrong. Coupling scheme and probe damper for pulsed nuclear magnetic resonance single coil probe. *Rev. Sci. Instrum.*, 45(9):1159–1163, 1974.

[25] D. I. Hoult. Fast recovery with a conventional probe. *J. Magn. Reson.*, 57:394–403, 1984.

[26] C-N. Chen, , and D. I. Hoult. *Biomedical Magnetic Resonance Technology.* Adam Hilger, 1989.

[27] W. A. Edelstein, C. J. Hardy, and O. M. Mueller. Electronic decoupling of surface-coil receivers for NMR imaging and spectroscopy. *J. Magn. Reson.*, 67:156–161, 1986.

[28] B. deB. Frederick, M.S. Roos, and S.T.S. Wong. An inductively coupled PIN diode Q-spoiler for fast probe recovery. *Book of Abstracts, Soc. Magn. Reson. Med.* 10th *Annual Meeting*, page 726, 1991.

[29] A. Takahashi and T. M. Peters. Measurement of k-space trajectories for RF pulse design. *Proc. Soc. Magn. Reson. Med.* 12th *Annual Meeting*, page 424, 1993.

[30] D.B. Twieg. The k-trajectory formulation of the NMR imaging process with applications in analysis and synthesis of imaging methods. *Med. Phys.*, 10:610–621, 1983.

[31] S.J. Norton. Fast magnetic resonance with simultaneously oscillating and rotating field gradients. *IEEE Trans. Med. Imag.*, 6:21–31, 1987.

[32] S. T. S. Wong and M. S. Roos. Strategy for sampling on a sphere withe applications to 3D selective pulse design. *Proc. Soc. Magn. Reson. Med.* 12th *Annual Meeting*, page 1178, 1993.

[33] K. M. Hanson. Detectability in computed tomagraphic images. *Med. Phys.*, 6(5):441–451, 1979.

[34] S.T.S. Wong, M.S. Roos, R.D. Newmark, and T.F. Budinger. Discrete analysis of stochastic NMR. II. *J. Magn. Reson.*, 87:265–286, 1990.

[35] A.G. Webb, T.H. Mareci, and R.W. Briggs. Relative efficiencies of weighting methods for phase encoded localized NMR. *Proc. Soc. Magn. Reson. Med.* 12th *Annual Meeting*, page 900, 1993.

[36] S.L. Ponder and D.B. Tweig. A novel sampling method for spectroscopic imaging with improved sensitivity and spatial resolution. *Proc. Soc. Magn. Reson. Med.* 12th *Annual Meeting*, page 899, 1993.

[37] T.L. Chenevert and J.G. Pipe. Dynamic 3D imaging at high temporal resolution via reduced k-space sampling. *Proc. Soc. Magn. Reson. Med.* 12th *Annual Meeting*, page 1262, 1993.

[38] A. Maeda, K. Sano, and T. Yokoyama. Reconstruction by weighted correlation for MRI with time varying gradients. *IEEE Trans. Med. Imag.*, 7:26–31, 1988.

[39] A. Macovski. Volumetric NMR imaging with time-varying gradients. *Magn. Reson. Med.*, 2:29–40, 1985.

[40] M.S. Roos, S.T.S. Wong, and R.D. Newmark. Stochastic NMR imaging: Spatial localization with oscillating gradients. *Book of Abstracts, Soc. Magn. Reson. Med.* 7th *Annual Meeting*, 1:956, 1988.

[41] M.S. Roos, S.T.S. Wong, and B.deB. Frederick. Fast spectroscopic imaging with oscillating gradients and stochastic excitation. *Book of Abstracts, Soc. Magn. Reson. Med.* 9th *Annual Meeting*, page 135, 1990.

[42] S.T.S. Wong, M.S. Roos, R.D. Newmark, and T.F. Budinger. Discrete analysis of stochastic NMR. I. *J. Magn. Reson.*, 87:242–264, 1990.

[43] J. Paff, R. Freeman, and B. Blümich. Reduction of systematic noise in stochastic-excitation NMR by oversampling. *J. Magn. Reson. Ser. A*, 102(3):332–343, 1993.

Appendix A

Reconstruction Pseudocode

```
read signal s[n] into memory
generate and store y[n]      /* excitation sequence */
generate and store kx[n]     /* integral of the x gradient */
generate and store ky[n]     /* integral of the y gradient */
generate and store kz[n]     /* integral of the z gradient */
initialize rhok[x][y][z][q] /* object kspace estimate */
initialize density[x][y][z] /* number of samples at each
position in rhok */
for each lag q
    {
    for each time point n
        {
        /* rotate data value by the phase of the excitation pulse
            q lags in the past */
        gridval = conjugate(y[n - q]) * s[n];

        /* find the current position in k-space for magnetization
            from time n - q */
        kxpos = kx[n] - kx[n - q];
        kypos = ky[n] - ky[n - q];
        kzpos = kz[n] - kz[n - q];

        /* convolve each sample with the spreading function to find
            it's contribution to neighboring points in k-space */
        for all x,y,z in the neighborhood of kxpos,kypos,kzpos
            {
            contrib = C((x,y,z) - (kxpos,kypos,kzpos));
```

```
    if(densitycorrection == empirical)
{
            density[x][y][z] += contrib;
            rhok[x][y][z][q] += gridval * contrib;
}
    if(densitycorrection == analytic)
{
            rhok[x][y][z][q] +=
    gridval * contrib/analyticdensity[x][y][z];
}
           }
        }
    /* correct for the sampling density if doing empirical
density correction */
    if(densitycorrection == empirical)
{
     for all x,y,z
         {
         if density[x][y][z] != 0
                {
                rhok[x][y][z][q] /= density[x][y][z];
                }
         }
     reinitialize density[x][y][z];
}

    /* transform rho estimate to the spatial domain */
    rhok[x][y][z][q] = 3DIFT(rhok[x][y][z][q], w.r.t x,y,z);
    }
/* transform rho estimate to spectral domain */
for all x,y,z
    {
    rhok[x][y][z][q] = FT(rhok[x][y][z][q], w.r.t q);
    }
}
store rho estimate
```

LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA
TECHNICAL INFORMATION DEPARTMENT
BERKELEY, CALIFORNIA 94720