

# *Near Real-Time 2D Non-Fourier Basis*

## *MRI*

by

Dimitrios Mitsouras

B.S., Brown University (1998)

Submitted to the  
Department of Electrical Engineering and Computer Science  
in partial fulfillment of the requirements for the degree of  
Master of Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY.


May 2000

*[Faint stamp: JUN 2 2000]*

© Dimitrios Mitsouras, MM. All rights reserved.

*The author hereby grants to MIT permission to reproduce and distribute  
publicly paper and electronic copies of this thesis document in whole or in  
part, and to grant others the right to do so.*

**ENG**



Author \_\_\_\_\_

Department of Electrical Engineering and Computer Science  
May 5, 2000

Certified by \_\_\_\_\_

Alan Edelman  
Associate Professor of Applied Mathematics  
Thesis Supervisor

Certified by \_\_\_\_\_

Gary P. Zientara  
Assistant Professor of the Harvard Medical School  
Thesis Supervisor

Accepted by \_\_\_\_\_

Arthur C. Smith  
Chairman, Departmental Committee on Graduate Students

# *Near Real-Time 2D Non-Fourier Basis MRI*

by  
Dimitrios Mitsouras

*Submitted to the  
Department of Electrical Engineering and Computer Science  
on May 5, 2000 in partial fulfillment of the  
requirements for the degree of Master of Science.*

## ***Abstract***

This thesis presents a parallel implementation of recently proposed mathematical methods [55, 58] which increase the efficiency of standard Magnetic Resonance Imaging (MRI) signal acquisition and reconstruction by an order of magnitude over an equivalent standard pulse sequence.

Standard 2D Magnetic Resonance Imaging relies on data acquisition from the Field Of View (FOV) via sequential repetitive excitations of consecutive single spatial Fourier components from the entire FOV. In a ground-breaking approach, based on Maudsley's usage of shaped RF pulses,[25] Weaver and Healy [51] noted that much more efficient encoding can be achieved using wavelets if specially shaped Radio Frequency (RF) pulses were used to generate spatial excitation profiles. It has since been realized that any orthogonal basis set may be used to encode an MR image, yielding much more efficient encoding [55, 58]. We have designed and implemented a system which, for the first time, will produce 2D clinical MR images in near real-time, at least an order of magnitude higher temporal resolution improvement over currently possible for an equivalent Spin Echo pulse sequence. Our spatial encoding methods derive basis sets from linear-algebraic approaches such as the SVD and Krylov subspace methods, obtaining an optimal encoding basis set from the current FOV. The large amount of computing required is handled by a heterogeneous RACEway multi-processing system directly connected to the MR scanner via a fiber channel interface.

This document is composed of two parts. The first half is of *pedagogical* nature, aiming to introduce the reader to MR Imaging from first principles and furthermore explain the recent advances in spatial encoding methods. The second half aims to explain our contributions, giving a detailed explanation of the implementation, system integration and near real-time results obtained on a standard clinical General Electric 1.5T MR scanner. Finally, we present a new, promising approach of MR image acquisition and reconstruction derived from non-orthogonal vector sets.

Thesis Supervisor: Alan Edelman  
Title: Associate Professor of Applied Mathematics

Thesis Supervisor: Gary P. Zientara  
Title: Assistant Professor of the Harvard Medical School

# *Acknowledgments*

First and foremost, I would like to thank both my advisors, Professors Gary Zientara and Alan Edelman. It is because of their insight, excitement for their respective fields, and will to help me gain some of their understanding, that this work has flourished. At times, when things got astray, it was their will to help, motivate and push me to my limits, sometimes even going as far as bringing me cookies at nine o'clock in the morning (many thanks to Ellen Zientara!) that allowed me to successfully complete this project. Many times, I found the answer to problems by recalling their words, which at the time, due to ignorance, had seemed irrelevant. The form of this thesis, as an introduction to the field of fast MRI, was shaped by those realizations, and the wish to pass them on. This work is proof of my appreciation and gratitude for and towards them. I would also like to thank Professor Larry Panych, for offering his extremely valuable and useful opinion and knowledge whenever I was in need.

I would not have reached my goals without the unlimited energy, support and love of my family, to whom I owe everything, beyond what words can convey. I would like to dedicate this thesis to my family, especially my dearest aunt who spent countless hours step by step alongside me in spirit (and along phone lines), trying to help me any way she could, mentally, spiritually and even technically. I would also like to thank my dearest sister, always caring, supporting and standing up for me, but also for keeping my bank accounts low enough so that I had not much choice other than work on my research! My mother has dedicated her life and work to myself and my sister, and although I will never be able to repay her, I can only dedicate this work to her. My father provided for me during my early years as a foreign college student, and I would not have been here writing this had he not done so.

This work was in part supported by NIH contract R01-CA78299 and a BWH Research Council Interdisciplinary Seed Grant as well as an Edgerly Science Partnership Fund Award.

Finally, I would like to thank all the dwellers of MIT's Laboratory for Computer Science, and those of the Brigham And Women's Hospital Surgical Planning



Lab for making my office spaces an intellectually stimulating environment, and more important, a very friendly one. Thanks to Lou Arruda, Harald Prokop, Sarfraz Khurshid, Walid Kyriakos, George Hadjiyiannis, Seung-Schik Yoo and Nitin Thaper for standing by me and tolerating me!

DIMITRIS MITSOURAS  
*Cambridge, Massachusetts*  
*May 10, 2000*

# Contents

<b>1</b>	<b>Introduction</b>	<b>8</b>
<b>2</b>	<b>MRI Basics</b>	<b>12</b>
2.1	Overview of the MRI process . . . . .	12
2.2	NMR & Physics Essentials . . . . .	13
2.2.1	Classical view of orbiting charged particles . . . . .	14
2.3	Quantized view . . . . .	15
2.4	Quantum Mechanics & NMR . . . . .	17
2.5	Observing $\vec{\mu}$ . . . . .	18
2.5.1	Classical description of precession of a single spin . . . . .	20
2.5.2	Motion of the net magnetization vector $\vec{\mu}$ . . . . .	23
2.5.3	Generation of the MR signal: Induction . . . . .	23
2.6	$\vec{\mu}$ Under a Pulsed Magnetic Field . . . . .	24
2.6.1	Rotating frame of reference . . . . .	26
2.6.2	Flip angle . . . . .	28
2.7	Relaxation Processes . . . . .	30
2.7.1	Spin-lattice ( $T_1$ ) relaxation . . . . .	31
2.7.2	Spin-spin ( $T_2$ ) relaxation . . . . .	32
2.8	The Bloch Equations . . . . .	34
2.8.1	The Bloch equations in a rotating frame . . . . .	34
2.9	Spatial Localization of the MR Signal . . . . .	36
2.9.1	Slice selection gradient . . . . .	37
2.9.2	Frequency-encoding gradient . . . . .	39
2.9.3	Phase-encoding gradient . . . . .	41
2.10	Tomographic Image Generation & Fourier Transforms . . . . .	46
2.11	Spin-Echo Pulse Sequence . . . . .	48
<b>3</b>	<b>Advanced Non-Fourier MR Imaging</b>	<b>51</b>
3.1	Slice Selection in Light of Spatially-Selective Excitation . . . . .	52
3.2	Linearity of the MR Process in the Low Flip-Angle Regime . . . . .	53

3.3	Matrix Representation of RF Encoding . . . . .	54
3.3.1	Fourier-encoded MRI by spatially-selective RF pulses . . . . .	58
3.4	Dynamic Adaptive MRI . . . . .	62
3.4.1	Invertible separable unitary transforms . . . . .	65
3.5	Optimal Spatially-Selective Excitation by SVD . . . . .	67
3.6	Excitation by Krylov Subspace Methods . . . . .	70
3.6.1	Aside: Excitation by row-space encoding . . . . .	71
3.7	Excitation by Non-Orthogonal Vector Sets . . . . .	72
<b>4</b>	<b>Pulse Sequence Design And Results</b>	<b>75</b>
4.1	PSD for GE 1.5T Signa LX 8.2.5 Scanner . . . . .	77
4.1.1	Host and TPS processes . . . . .	80
4.2	Spatially-Selective Fourier Encoding . . . . .	84
4.3	Random Orthogonalized Basis Encoding . . . . .	88
4.4	SVD Encoding . . . . .	88
4.5	Krylov Subspace Encoding . . . . .	94
<b>5</b>	<b>Real-Time System Design and Implementation Issues</b>	<b>97</b>
5.1	How to Tap Into the MRI Scanner Internals . . . . .	99
5.2	Computation Timeline . . . . .	104
5.3	Functional Units and Interconnections . . . . .	108
5.4	Near Real-Time System Performance . . . . .	109
<b>6</b>	<b>Conclusions</b>	<b>111</b>
6.1	Further Research . . . . .	112
<b>B</b>	<b>Bibliography</b>	<b>114</b>

## SECTION 1

---

# *Introduction*

Magnetic Resonance Imaging (MRI) is based on the nuclear magnetic resonance principle (NMR), which refers to enhanced energy absorption by atomic nuclei, at some resonant radio frequency (RF), when they are under the influence of an external magnetic field. NMR yields a measurement of the nuclear magnetic moments of specific nuclei. MR Imaging is a manifestation of NMR, wherein the immediate chemical surroundings of the nuclei cause a measurable variation in those measured magnetic moments and the relaxation properties of the nuclei.

By the early 1980s, MR Imaging had become the preferred non-invasive medical tomographic imaging modality, partially due to its adjustable, tissue contrast, usage of hazard-free, non-ionizing, radiation, rendering it especially suited for clinical applications. MR images are thin slices of the body, produced by measuring the nuclear magnetic moments of the hydrogen atoms found in the water contained in tissues within the slice. NMR possesses great sensitivity in differentiating between normal and damaged tissues or external probes.

Despite its virtues, MR data acquisition is however, very slow, currently requiring between 20 seconds to several minutes per high resolution image, limiting its usage in dynamic diagnostic imaging, as well as interventional radiology, which require the generation of images within a few seconds from each other. This temporal limitation is due in part to the spatial encoding employed and the physics of nuclear spin relaxation.

In this thesis we will present well-known linear-algebraic methods, placed in the context and framework of MRI, which allow us to circumvent the physical

constraints of spin relaxation up-to a certain degree, and increase the efficacy of data acquisition by an order of magnitude. In fact, mathematically, our methods provide the *optimal* strategy for MR signal acquisition for a certain image based on previously acquired images, but only nearly-optimal in the MR framework, since changes and noise in the FOV can never be exactly anticipated. An implementation of our techniques on a parallel multicomputer yields two-dimensional images with minimal temporal overhead apart from the acquisition process itself. Our system can produce MR images in under one second.

The first step in understanding the well-foundedness of our methodology and its implementation is to look at the MRI process and the related physics. MRI is a *multi-parametric* technique; a number of parameters can be almost arbitrarily associated, leading to a vast array of different possible pulse sequences.[10]

In general, an *MRI pulse sequence* is defined as a series of RF pulses in conjunction with magnetic field gradient pulses that together generate a NMR signal which can be processed in a well-defined manner to produce an image. It is exactly the form and timing of the RF and gradient pulses which determines the information content of the image, based on relaxation properties of tissues, diffusion, fluid flow and so on. The magnetic field gradients are used in MRI to modify the precession frequencies of the atomic spins in a known fashion, allowing a linear relationship between the position of the nuclei and their frequency of precession, therefore producing signal modulation along a single arbitrary axis.

In the MRI literature, a spatially dependent complex-valued *signal density*  $S(x, y)$  function is often used to represent the image information. This is not the directly measured quantity in MRI, instead, the actual signal observed in an MR experiment is the integrated weighted (e.g., by a Fourier basis function) signal density over a volume. In order to reconstruct an image from the observed signal, multiple measurements must be made under different conditions so as to reconstruct an estimate of  $S(\vec{x})$ . These differing conditions can be a weighting of the signal by various encoding functions. If these encoding functions form an orthogonal basis set, the ensemble of the measurements yield enough acquired data to reconstruct  $S(\vec{x})$ , e.g. representing a function in terms of some chosen basis.

In the standard method of data acquisition, known as Fourier encoding, the encoding functions are simply Fourier basis functions. Let the  $n \times n$  complex-valued matrix  $\mathbf{S}$  be the discrete representation of the in-plane MR signal intensity

function  $S(x, y)$ . Let the 2D DFT matrix

$$\mathbf{W} = \frac{1}{\sqrt{n}} \begin{bmatrix} \mathcal{W}^0 & \mathcal{W}^0 & \mathcal{W}^0 & \dots & \mathcal{W}^0 \\ \mathcal{W}^0 & \mathcal{W}^1 & \mathcal{W}^2 & \dots & \mathcal{W}^{n-1} \\ \mathcal{W}^0 & \mathcal{W}^2 & \mathcal{W}^4 & \dots & \mathcal{W}^{2(n-1)} \\ \vdots & & & & \vdots \\ \mathcal{W}^0 & & \dots & & \mathcal{W}^{(n-1)^2} \end{bmatrix}, \quad (1.1)$$

where  $\mathcal{W} = e^{-2\pi i/n}$ . In this case, the Fourier-domain matrix  $\mathbf{F}$  is related to the image  $\mathbf{S}$  as

$$\mathbf{F} = \mathbf{W}\mathbf{S}\mathbf{W} \quad (1.2)$$

$$\mathbf{S} = \mathbf{W}^*\mathbf{F}\mathbf{W}^*, \quad (1.3)$$

where the superscript  $\star$  signifies the Hermitian conjugate of a matrix. In a Fourier-encoded MR experiment, the matrix  $\mathbf{F}$  is acquired, by using a magnetic field gradient which enforces frequency modulation of the MR signal in one dimension, and producing successively  $n$  lines of echos, wherein a small phase increment has been enforced on each echo. If the time between two consecutive sample acquisitions is  $TR$ , then in order to acquire the full matrix  $\mathbf{F}$  to yield an  $n \times n$  image, a total of  $n \times TR$  time is required.

A number of methods have been developed to try and overcome this exorbitant time limitation. For example, so-called *Echo-Planar Imaging* (EPI) pulse sequences, try to minimize acquisition time by using a single excitation to acquire more than one echos from the atoms, [52, 18] by “blipping” to cause a small amount of phase shift after each acquisition. The downside of EPI sequences is the requirement of special hardware, due to the extreme requirements placed on the gradient field switching times (e.g., gradient slew rates,) as well as inhomogeneous signal intensity across frequencies. Another example is so-called *Gradient Recalled Echo in the Steady State* (GRASS) pulse sequences, which acquire an echo shortly after excitation, without first refocusing it [18, 53]. The downside of GRASS-type sequences is that they are susceptible to magnetic field inhomogeneities caused by tissue etc (although in some cases, such as BOLD sequences, used in functional MRI, *fMRI*, such inhomogeneities are actually exploited).

An alternative to increasing the efficiency of MR signal acquisition stems from realizing that Fourier encoding is a special case of a change of representation of a 2D image, first identified by Andrews and Pratt.[2] A matrix  $\mathbf{S}$  can be mapped into a differently encoded representation,  $\mathbf{R}$  by applying a separable linear unitary transformation (e.g. that is invertible,) [39]

$$\mathbf{R} = \mathbf{A}_C^* \mathbf{S} \mathbf{A}_R \quad (1.4)$$

$$\mathbf{S} = \mathbf{A}_c \mathbf{R} \mathbf{A}_r^*. \quad (1.5)$$

The matrix  $\mathbf{R}$  is the new image representation, and  $\mathbf{A}_c, \mathbf{A}_r$  are unitary (orthogonal) matrices, spanning the column and row spaces of  $\mathbf{S}$  respectively, merely an alternative definition of the weighting functions. The column-matrix  $\mathbf{A}_c$  spans the column space of  $\mathbf{S}$ : each column of  $\mathbf{S}$  can be formed exactly by a linear combination of the column vectors  $\mathbf{u}_i, i = 1 \dots n$  comprising the matrix  $\mathbf{A}_c$  [17]. Similarly the column-matrix  $\mathbf{A}_r$  spans the row space of  $\mathbf{S}$ .

It has been shown [32] that for small flip-angle excitations<sup>1</sup> ( $\leq 30^\circ$ ), the MR imaging process is linear: if the row-matrix  $\mathbf{P}$  whose rows define unique RF pulses to be used for excitation of the volume, and if some time ( $TE$ ) after each excitation a set of echo samples is acquired, during which time a frequency encoding gradient is applied, then the response of the system is

$$\mathbf{R} = \mathbf{P}\mathbf{F}, \quad (1.6)$$

where the row-matrix  $\mathbf{R}$  contains the digitized echo response, each row  $i$  corresponding to the  $i$ 'th RF excitation defined by  $\mathbf{P}$ .

It is this linearity of the MR imaging process which gives us the freedom to use optimal encoding basis sets,[58] such as the well-known singular value decomposition (SVD), to acquire an MR image in a fraction of the time that is required to sample the Fourier domain using other standard methods.

We will first examine the physics behind the formation of an MR image in Section 2. We will turn to understanding the linear algebra-derived encoding methods and introduce encoding by non-orthogonal vector sets methods in Section 3. In Section 4, we will present results obtained on a standard clinical GE 1.5T scanner for each of the linear-algebraic methods described. Finally we will describe the hardware and implementation of our near real-time system in Section 5. Our concluding remarks and future extensions for near real-time, three-dimensional, very high resolution volume imaging are offered in Section 6.

---

<sup>1</sup>As will be explained in the next chapter, the flip angle refers to the angle between the magnetization vector quantity defined by the sum of nuclear moments in the excited volume and the static, external, magnetic field axis of the MR scanner.

## SECTION 2

---

# *MRI Basics*

An essential goal of this thesis is to explain the process by which an MR image is constructed in order to understand how to optimize the associated data acquisition. In this section, we present the well-known physics of NMR and how it is used to produce a two-dimensional image. Armed with this knowledge, we can then turn to optimizing the data acquisition process.

With this intent, this entire section is merely a review of the relevant physics, which can be found in many college-level electromagnetism and quantum mechanics books. The reader already comfortable with these concepts may wish to skip to Section 2.8. On the other hand, readers already familiar with the basics of MRI may find it useful to skip the entire section and immediately turn to our spatially selective encoding methods in Section 3. Finally, advanced readers already familiar with spatially-selective pulse sequences, such as wavelet-encoded MRI and SVD encoding may wish to turn to the last two sections, where we describe the main contribution of this thesis: the design and implementation of a complete, highly inhomogeneous, system which achieves an order of magnitude faster MR imaging over an equivalent Spin Echo pulse sequence.

### *2.1 Overview of the MRI process*

Standard MR images are the end-product of a number of steps, that can be briefly summarized as follows: [53]

1. First, a strong, uniform, magnetic field is used to align all the randomly



aligned nuclei in the tissue. This alignment produces an overall “magnetization” in the tissue.

2. This equilibrium magnetization is tipped off-axis by the magnetic component of an RF pulse at some resonance frequency. In effect, the nuclei respond by departing from the equilibrium by an amount proportional to the applied RF power.
3. Once the RF pulse is turned off, the nuclei begin to relax back to the equilibrium state. Due to their surrounding (sometimes time-varying) chemical environment, they relax at different rates.
4. The relaxation process produces an RF signal proportional to the amount of the initial alignment. The different relaxation rates of the nuclei give rise to tissue contrast.
5. The positions of the nuclei are localized by modulating the signal they produce, by using spatially varying magnetic fields, called “gradients”.
6. The modulated signal is acquired after a freely specified “echo time” ( $TE$ ) and its Fourier transformation yields back the originating positions of the signals, yielding an MR image.

We now turn to a more rigorous description of the basic physical principles that underlie the above process. We will turn to both the classical and quantum mechanical descriptions of magnetic resonance, as the first is most convenient in understanding the dynamic and transient effects, whereas the latter is what give rise to the resonance phenomena.

## 2.2 *NMR & Physics Essentials*

Hydrogen, the most abundant element in the Universe, and also the human body, is the main focus of clinical MR imaging, and hence the focus of our discussion. It is composed of a single proton nucleus, positively charged. Classically, it possesses a magnetic field, causing the nucleus to act as a tiny bar magnet. The proton can be thought of as spinning about the North-South axis, in analogy to the magnetic field produced by a moving charge.

The hydrogen nucleus is paramagnetic<sup>1</sup> and thus, under the influence of a strong, uniform, magnetic field tends to align itself with the external field – in

---

<sup>1</sup>A paramagnetic molecule is one with an unpaired electron.

a sense, it becomes “magnetized.” Although it is the classical description of the particles that provides almost all the laws needed to explain the formation of the MR signal, it is the quantum mechanical description which gives rise to nuclear magnetic resonance, therefore we explore the two in parallel.




Figure 2-1: A proton and the magnetic field that can be thought of as emanating from it.

### 2.2.1 Classical view of orbiting charged particles

Let us consider an electric charge  $q$  traveling with speed  $v$  in a circular orbit of radius  $r$ . Consider a single point along the trajectory of the charge. Clearly, the charge passes through that point with a period of  $T = 2\pi r/v$ , and is therefore equivalent to a current  $I = q/T = qv/2\pi r$ . [14] Classical electromagnetism [40] tells us that a circuit enclosing an area  $A$  with a current  $I$  flowing through it possesses an *orbital magnetic dipole moment*

$$\vec{M} = A\hat{r} \times \vec{I}, \quad (2.1)$$

where  $\hat{r}$  is a unit vector from the center of rotation to the particle,  $A = \pi r^2$  is the *area* traced by the particle motion, and  $\vec{I} = \frac{vq}{2\pi r}$  is the vector representation of the current.


Combining the above two, we have that the orbital magnetic moment is

$$\vec{M} = \frac{1}{2}q|r|\hat{r} \times \vec{v}. \quad (2.2)$$

Since the particle is moving in a circular orbit, it also possesses an *orbital angular momentum*,  $L$ , defined as

$$\vec{L} = m\vec{r} \times \vec{v} \equiv mr\hat{r} \times \vec{v}, \quad (2.3)$$

where  $m$  is the rest mass of the particle.



**Figure 2-2:** The flow of a current in a circular orbit tracing an area  $A$  gives rise to a magnetic dipole moment.

The ratio of the magnetic moment and the angular momentum is known as the *gyromagnetic ratio*,  $\gamma$ :

$$\gamma = \left| \frac{M}{L} \right| = \frac{q}{2m}. \quad (2.4)$$

### 2.3 Quantized view

In the quantum mechanical description of things, basic units are quantized in multiples of Planck's constant, the orbital angular momentum and magnetic dipole moment are not an exception. The orbital angular momentum of atoms is quantized in magnitude[14]

$$L^2 = l(l + 1)\hbar^2, \quad (2.5)$$

where  $l$  is a non-negative integer, the orbital angular momentum quantum number, and  $\hbar$  is Planck's constant over  $2\pi$ . It must also be quantized in orientation. The allowed orientations of the orbital angular momentum relative to a fixed axis, e.g., the  $z$  axis<sup>2</sup>, is given by[14]

$$L_z = m_l \hbar, \quad (2.6)$$

<sup>2</sup> $L$  is a vector and can thus be thought of as being aligned with an axis; choosing the  $z$  axis, arbitrarily defining the particle motion to lie on the  $x - y$  plane, the classical angular momentum lies on the  $z$  axis.

where  $m_l = 0, \pm 1, \dots, \pm l$ . Note that given some orbital angular momentum quantum number  $l$ , there are  $2l + 1$  compatible orientations.<sup>3</sup>

Similar to the orbital angular momentum, elementary particles possess an *intrinsic spin angular momentum*<sup>4</sup>[14]

$$\vec{M}_s = \gamma_s \vec{S} = g\gamma_s \vec{S}, \quad (2.7)$$

where  $g$  is the *spectroscopic splitting factor* and  $\vec{S}$  is the intrinsic spin, with a quantized magnitude

$$S^2 = s(s + 1)\hbar^2, \quad (2.8)$$

where  $s$ , the spin momentum quantum number, can be a *half integer* (thus explaining the even number of possible atom states<sup>5</sup>)  $\vec{S}$  is also, quantized in direction,[14]

$$S_z = m_s \hbar. \quad (2.9)$$

The electron and the proton have a spin of  $s = 1/2$ , and thus  $S^2 = 3/4\hbar^2$  and the two allowed orientations are  $m_s = \pm 1/2$ .<sup>6</sup>

Now, we are ready to express the important result the above physics lead us to: the energy states of particles under the influence of a magnetic field.

Under the influence of an external magnetic field  $B_0$ , an *isolated* proton (or electron) has two possible energy levels corresponding to each of the two possible spin orientations:

$$\begin{aligned} E_{\pm} &= -M_s B_0 \\ &= -\gamma_s S_z |B_0| \\ E_{\pm} &= \mp \frac{1}{2} \hbar \gamma_s |B_0|. \end{aligned} \quad (2.10)$$

The energetically favorable energy state is always the minimum energy state, which immediately reveals itself from (2.10) as the one wherein the spin is aligned with the direction of the imposed magnetic field ( $E_+$ ).


---

<sup>3</sup>The number of allowed orientations of  $\vec{L}$  ( $2l + 1$ ) is exactly the number of compatible eigenvalues of  $L_z$ . [14]

<sup>4</sup>When considering individual atoms, this is an essential addition to the quantized angular momentum, since, experimentally, some atoms exhibit an even number of possible orientations. In particular, this was the observation in the Stern-Gerlach experiment, wherein a beam of cesium atoms under the influence of a transverse magnetic field gradient was split into two beams. [44]

<sup>5</sup>The quantization of possible  $z$  orientations in half-integer multiples of  $\hbar$  was first proposed by Uhlenbeck and Goudsmit. [49]

<sup>6</sup>For a proton, the spin angular momentum is about 2.793 times as effective as the orbital angular momentum.



**Figure 2-3:** A spin  $s = 1/2$  elementary particle can only be in one of two allowed orientations with respect to a fixed axis. It can however exist at any angle on the perpendicular plane (tracing a cone).

The energy equation above, further reveals that MR can be performed using any nucleus so long as it possesses a spin (having an odd number of protons or neutrons). For example, sodium with a spin of  $3/2$  can be imaged, but helium which has no spin, can not. The reason for imaging spin  $1/2$  nuclei is that they are particularly easy to observe.[53] In particular, hydrogen is not only the most naturally abundant element, but it also has the highest gyromagnetic ratio of any element. From the moment on that each state has a unique energy associated with it, we are immediately given the ability to interact with the nuclei.

## 2.4 Quantum Mechanics & NMR

When an external magnetic field is applied, the random orientations of the particle spins are in one of the possible spin states, corresponding to the allowed orientations. The equilibrium of the system favors the least energetic state. Considering the special case of protons or electrons with  $s = 1/2$ , by equation (2.10) the minimum energy state is the one wherein the spin is aligned to the field. Therefore there will be an excess of particles which are aligned with the field. The other orientation, that opposed to the field, will also be present, simply due to the ambient thermal energy. The reason for this is that the energy difference between the two states is much smaller than the thermal energy. The energy difference, given by

$$\begin{aligned}
 \Delta E &= E_- - E_+ \\
 &= \hbar\gamma_s|B_0| \\
 &= \hbar\omega_L
 \end{aligned}$$

$$= h\nu_L \quad (2.11)$$

reveals that a photon of frequency  $\nu_L = \gamma_s|B_0|/2\pi$  can force a proton to shift from the aligned (parallel) state to the opposed (anti-parallel) state. This frequency is called the *Larmor* frequency[53] and is about 42.57MHz for a proton in an external field of strength  $|B_0| = 1T$  (Tesla). This frequency is in the RF range.

Because of this energy difference, the applied field causes a net excess of protons aligned with the field. The population difference is then[53]

$$\frac{n_+}{n_-} = e^{-\mu}, \left[ \mu = \frac{\Delta E}{k_B T} \right], \quad (2.12)$$

where  $k_B$  is the Boltzmann constant and  $T$  is the absolute temperature in degrees Kelvin.[53]  $k_B T$  is the average thermal energy of the system. At 1 Tesla, the population difference is about 3 parts per million. It is this population difference that gives rise to the signal observed in NMR.

If we consider the overall alignment of protons under the influence of the magnetic field, we can associate a classical vector  $\vec{\mu}$ , referred to as the *net magnetization*, as a representation of this overall alignment.


## 2.5 Observing $\vec{\mu}$

Once the net magnetization  $\vec{\mu}$  has been aligned with the external magnetic field, and has reached an equilibrium, so long as there is no change, we can not observe it.

We can however, disrupt the equilibrium magnetization by applying an RF signal at the Larmor frequency in order to observe it. An RF signal is an oscillating electro-magnetic field, or alternatively, photons of energy proportional to the frequency. If that frequency matches the Larmor frequency, its effect is to cause the body's net magnetization to oscillate from the parallel to the anti-parallel orientation since individual nuclei can absorb a photon and flip from the parallel to the anti-parallel state. This is the first step of the process by which an MR signal is generated.

Classically,  $\vec{\mu}$  can be thought of as being much like a spinning top which is momentarily flicked by a force  $F$  and begins to wobble from side to side.

This wobbling is called precession. The frequency of the precession is exactly the Larmor frequency. If the RF pulse is applied at intervals equal to the precession



**Figure 2-4:** Under an external magnetic field, proton spins are aligned parallel or anti-parallel to the field, with a small fraction of extra spins aligned parallel to the field. By absorbing a photon of energy  $h\nu_L$ , a proton can flip from the parallel to the anti-parallel state, while the reverse causes an emission of a photon.




Figure 2-5: Spin precession is equivalent to a spinning top.

frequency, the precession angle will increase at each interval. The same effect can be obtained by increasing the strength of the RF pulse. A  $90^\circ$  angle of precession (in the classical sense, of group behavior of protons) is said to be caused by a  $90^\circ$  pulse.

In the following sections we model the behavior of the magnetic moments, first at the level of individual nuclei, and then extend this to the group behavior of large numbers of atoms.

### 2.5.1 Classical description of precession of a single spin


Suppose that we consider a nucleus to be a large rotating sphere, with a magnetic dipole moment  $\vec{M}$ . If the charge of the sphere is positive then the magnetic moment vector is pointing in the same direction as the angular momentum vector. On the other hand, if the charge is negative, then it is pointing in the opposite direction. In the presence of a magnetic field  $\vec{B}_0$ , the dipole does not act as a compass needle which oscillates about the direction of the field until it eventually aligns itself with it due to friction. Instead, due to its rotational angular momentum, the sphere will *precess*, that is, the angle between  $\vec{M}$  and  $\vec{B}_0$  will remain *constant*. Let us see why this is the case.

We can think of the dipole as two equal and opposite magnetic “poles”  $\pm q_m$  separated by some distance  $l$ , so that  $M = lq_m$ . The magnetic field will impose forces on *both* poles, and therefore a *torque* is created which is perpendicular to the angular momentum vector. So long as the magnetic field is homogeneous, this condition creates *steady precession*. [11, 12]

The torque applied to the system is defined as the time rate of change of the total angular momentum: [11]

$$\frac{d\vec{L}}{dt} = \vec{\tau}, \quad (2.13)$$





**Figure 2-6:** The classical description of a nucleus in terms of a tiny magnetic dipole explains the steady-state precession of the spins.

where  $\vec{\tau}$  is the torque. During some short time  $\Delta t$ , the torque adds to the angular momentum of the particle a vector increment  $\Delta\vec{L}$  in the direction of the torque vector. In a uniform magnetic field  $\vec{B}_0$ , the torque tends to rotate the magnetic moment  $\vec{M}$  into the direction of the field, expressed by the following cross product:[14]

$$\vec{\tau} = \vec{M} \times \vec{B}_0. \quad (2.14)$$

Since  $\vec{M} = \gamma\vec{L}$ , it follows that for an isolated magnetic moment:

$$\frac{d\vec{L}}{dt} = \vec{M} \times \vec{B}_0 \Leftrightarrow \frac{d\vec{M}}{dt} = \gamma\vec{M} \times \vec{B}_0. \quad (2.15)$$


This last equation describes the behavior of a single spin in a magnetic field  $\vec{B}_0$  according to the laws of classical physics. In solving the differential equation, assume without loss of generality, that the magnetic field has only a positive-z component, i.e.,

$$\vec{B}_0 = |B_0|\mathbf{k}. \quad (2.16)$$

Under this condition, the solution for the magnetic dipole moment becomes

$$\vec{M} = M_t \cos(\omega_L t + \phi)\mathbf{i} - M_t \sin(\omega_L t + \phi)\mathbf{j} + M_{z0}\mathbf{k}, \quad (2.17)$$

wherein the constants  $M_t$ ,  $\phi$  and  $M_{z0}$  are determined by the initial conditions. The angular speed  $\omega_L$  is the, by now familiar, Larmor frequency  $\gamma B_0/2\pi$ . As we will see later, it is exactly the transverse magnetization (the  $\mathbf{i}$ ,  $\mathbf{j}$  components) which give rise to the observable MR signal.



**Figure 2-7:** Relationship between the vectors  $\vec{M} = \gamma\vec{L}$ ,  $\tau = \vec{M} \times \vec{B}_0$  and the fixed coordinate system for a single nucleus.

That solution describes a magnetic moment  $\vec{M}$  of the nucleus as simply a rotating vector about the z-axis counterclockwise (that is about the direction of  $-\vec{B}_0$ , at the Larmor frequency.<sup>7</sup> Note that by the above solution, we have that

$$\frac{d\phi}{dt} = -\omega_L, \quad (2.18)$$

$$\frac{d\theta}{dt} = 0 \quad (2.19)$$

e.g., the precession frequency is the Larmor frequency, regardless of the angle  $\theta$  between the magnetization and the field.

This is the classical physics of the precession of a single, isolated, nucleus under the influence of a homogeneous external magnetic field. The next step towards understanding MR is to explain the formation of the measured quantity, the MR signal. We must bear in mind that the MR signal comes from voxels composed of countless individual nuclei, and not individual nuclei. However, note that under the laws of classical physics, the net magnetization vector is merely the integral over the distribution of the individual nuclear magnetic moments.

<sup>7</sup>It is customary to define the angular velocity  $\vec{\omega}_L = -\gamma\vec{B}_0$ , called the *Larmor vector*.

### 2.5.2 Motion of the net magnetization vector $\vec{\mu}$

The net magnetization  $\vec{M}$  of a voxel in the imaged volume is simply made up out of all the individual magnetic moments  $\vec{\mu}_i$ . Then, the net magnetization, for  $n$  nuclei in a sample, is simply

$$\vec{M} = \vec{\mu}_0 + \vec{\mu}_1 + \dots + \vec{\mu}_{n-1} = \sum_{i=0}^{n-1} \vec{\mu}_i \quad (2.20)$$

The time dependency of the net magnetization vector exposed to an external magnetic field follows immediately

$$\frac{d\vec{\mu}}{dt} = \gamma(\vec{M} \times \vec{B}_0) \quad (2.21)$$

At equilibrium, this net magnetic moment  $\vec{M}$  will be aligned with  $\vec{B}_0$  and will thus have no measurable components (it is the  $x - y$  components that give rise to the signal as we shall shortly see).

### 2.5.3 Generation of the MR signal: Induction

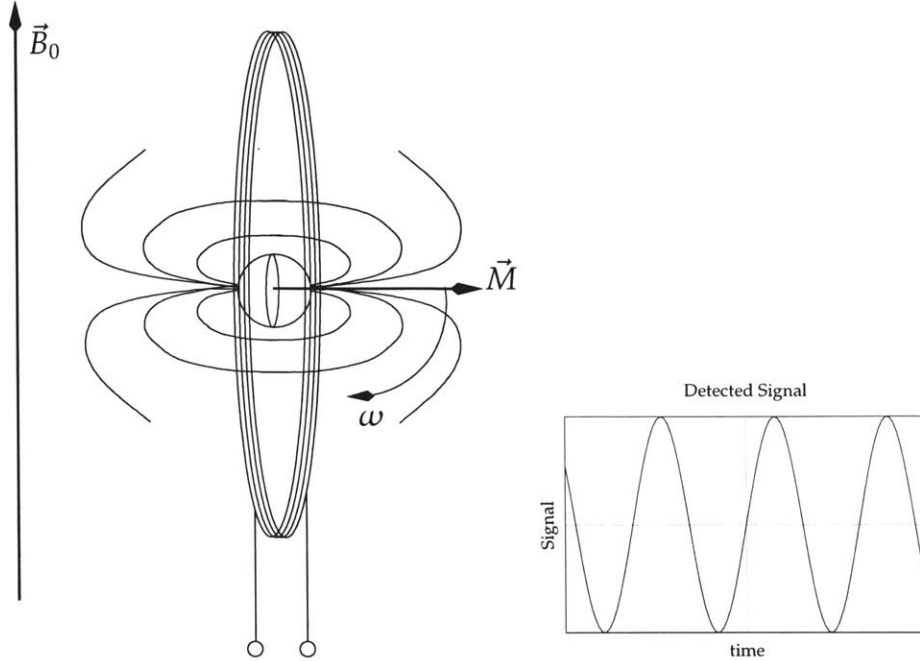
We have seen that if the magnetic moment of a nucleus has a component in the  $x - y$  plane, and is under the influence of a homogeneous magnetic field the particle will experience a torque that is proportional to the strength of the magnetic field. As a result, it will precess at the Larmor frequency. In this case, if the nucleus is surrounded by a receiver coil, then by Faraday's law[40], an alternating current will be induced on the coil, the frequency of the current being equal to the precession frequency of the particle.

If the angle ( $\theta$ ) of  $\vec{M}$  to the static magnetic field is 0, in which case the net magnetization does not have an  $x - y$  component,<sup>8</sup> there will be no MR signal, since the particle will not precess in the absence of torque. Unfortunately, as we have seen, the equilibrium state, i.e. the minimum energy state, of the net magnetization vector is that aligned with the external magnetic field. Therefore, merely placing a volume in a very strong magnetic field is alone insufficient to produce an MR signal.<sup>9</sup> Instead, it is essential to force the net magnetization into departing from the  $z$ -axis, by making  $d\theta/dt \neq 0$ . We already have the necessary physics to understand this, so we now turn on how to accomplish it.

---

<sup>8</sup>Since individual paramagnetic nuclei are always at an angle to the magnetic field, due to their quantum mechanical states, they are always precessing, each moment  $\vec{\mu}_i$  possessing an  $x = y$  component. However, these precessions are phase incoherent, and thus cancel each other out.

<sup>9</sup>The magnetic field of the earth is 0.5Gauss, which produces a very tiny overall alignment, too small to be measured.



**Figure 2-8:** The rotating precession of the magnetic moment of the proton, much like a bar magnet, induces a current on a receiver coil. In the absence of any dampening, or relaxation factors, the current detected is sinusoidal at the Larmor frequency.

## 2.6 $\vec{\mu}$ Under a Pulsed Magnetic Field

In MR imaging, the necessary rotation of the magnetization is accomplished by the magnetic field of an RF pulse. It is important to note that (2.15) holds for a *time-dependent* magnetic field, not only for a static one. This enables us to consider the effects produced by alternating magnetic fields under the scope of classical physics.

Assume that we apply a second magnetic flux density  $\vec{B}_1$  such that it is parallel to the  $x$ -axis and varies with time, by rotating about the  $x - y$  plane. That is, let

$$\vec{B}_1 = 2B_1 \cos(\omega_L t) \mathbf{i}. \quad (2.22)$$

The equation of motion of the dipole moment would then become

$$\frac{d\vec{M}}{dt} = \gamma \vec{M} \times (B_0 + B_1). \quad (2.23)$$

As will be demonstrated later, the time dependency of  $\vec{B}_1$  is needed in order to make  $d\theta/dt$  maximal, since the net magnetization is precessing at the Larmor

frequency during the application of this field. But first, we explore the caveats of generating such a field.


To make such a field is apparently not an easy task,[10, 45] instead, a plane polarized field is applied,<sup>10</sup> such that

$$\begin{aligned}\vec{B}_1 &= \{B_1\cos(\omega_L t)\mathbf{i} + B_1\sin(\omega_L t)\mathbf{j}\} + \{B_1\cos(\omega_L t)\mathbf{i} - B_1\sin(\omega_L t)\mathbf{j}\} \\ &\equiv 2B_1\cos(\omega_L t)\mathbf{i}.\end{aligned}\quad (2.24)$$

This field can be thought of as having two components, one rotating on the  $x - y$  plane clockwise ( $\vec{B}_1^c$ ) and the other rotating counterclockwise ( $\vec{B}_1^{cc}$ ):

$$\vec{B}_1^c = B_1(\cos(\omega_L t)\mathbf{i} + \sin(\omega_L t)\mathbf{j}), \quad (2.25)$$

$$\vec{B}_1^{cc} = B_1(\cos(\omega_L t)\mathbf{i} - \sin(\omega_L t)\mathbf{j}), \quad (2.26)$$



**Figure 2-9:** The decomposition of a plane-polarized field into two rotating components.

We can solve the equation of motion of the magnetic moment ((2.15)), obtaining[43]

$$\frac{d\theta}{dt} = \gamma B_1, \quad (2.27)$$


<sup>10</sup>Such a field can be produced by using two identical coils at right angles to each other, and with alternating currents which are 90 degrees out of phase.

therefore obtaining the maximal torque on the vector  $\vec{M}$  and subsequent rotation relative to the static field.

We must now turn into formulating and understanding the equation of motion for a spin under the effects of both the static and rotating magnetic fields.

### 2.6.1 Rotating frame of reference

It is now convenient to switch into the frame of reference that rotates along with  $\vec{B}_1^{cc}$ . The magnetic moments which are precessing about  $\vec{B}_0$ , by convention, the z-axis, at the Larmor frequency, will then maintain the same orientation as the counterclockwise component of the  $\vec{B}_1$  field,  $\vec{B}_1^{cc}$ . On the other hand, the clockwise component  $\vec{B}_1^c$  is rotating away from this frame of reference at twice the Larmor frequency, and thus has no effect on the magnetic moments.<sup>11</sup> However, because  $\vec{B}_1^{cc}$  is maintaining the same direction relative to the spins, its effect is maximal.



**Figure 2-10:** In a frame of reference rotating along with the net magnetization, the counterclockwise component of the  $\vec{B}_1$  field is always aligned with the magnetization, rendering its effect on the latter maximal.

Note that from the point of view of the magnetic moments, the field  $\vec{B}_1^{cc}$  is acting much like  $\vec{B}_0$  if we rotate our viewpoint. This means that the magnetic moments will also be precessing at some Larmor frequency  $\omega_1$  proportional to  $B_1$ , about

<sup>11</sup>It can be shown that near resonance the counter-rotating component of the field may be neglected.

the direction of  $\vec{B}_1^{cc}$ ! The magnetization vector, precesses on the  $y' - z'$  plane, in effect, flipping away from the equilibrium longitudinal magnetization along  $\vec{B}_0$ . If the field  $\vec{B}_1^{cc}$  is applied for an amount of time, the magnetization vector will rotate down to the  $y'$  axis by a certain amount, the so-called *flip angle*.

The complete motion of the sum of the magnetic moments, in the fixed frame, is the superposition of these two precessional motions: starting from the  $z$ -axis, it spirals down to the  $x - y$  plane.

Since the effect of the counter-rotating component can be neglected, the new equation of motion becomes

$$\frac{d\vec{M}}{dt} = \gamma\vec{M} \times \vec{B}_1^{cc} + \vec{B}_0. \quad (2.28)$$

By using a rotating frame of reference, which rotates counter-clockwise about the  $z$  axis at the Larmor frequency, we can eliminate the time dependency of  $\vec{B}_1$  and reformulate the equation in a simpler form. Before we embark on this, let us first understand the transformation by looking at the equation of motion under the static field  $\vec{B}_0$  which is aligned with the  $z$  axis.

Given a vector function of time in rectangular coordinates,

$$\vec{F}(t) = F_x(t)\mathbf{i} + F_y(t)\mathbf{j} + F_z(t)\mathbf{k}, \quad (2.29)$$

assume the elementary vectors retain a fixed length, that is we only allow rotations. Consider the case wherein the frame of reference defined by the elementary unit vectors is rotating with an angular velocity  $\vec{\Omega}$  such that

$$\frac{d\mathbf{i}}{dt} = \vec{\Omega} \times \mathbf{i}. \quad (2.30)$$

The time derivative of  $\vec{F}$  is then given by:[4, 43]

$$\left(\frac{d\vec{F}}{dt}\right)_{\text{lab}} = \left(\frac{\partial\vec{F}}{\partial t}\right)_{\text{rot}} + \vec{\Omega} \times \vec{F}, \quad (2.31)$$

where the subscripts “lab” and “rot” signify the laboratory and rotating frames of reference, respectively.

Given this result, we can rewrite the original equation of motion of the magnetization under the  $B_0$  field (equation (2.15)), as

$$\frac{\partial\vec{M}}{\partial t} + \vec{\Omega} \times \vec{M} = \gamma\vec{M} \times \vec{B}_0 \Leftrightarrow \quad (2.32)$$

$$\frac{\partial\vec{M}}{\partial t} = \vec{M} \times (\gamma\vec{B}_0 + \vec{\Omega}). \quad (2.33)$$

The last equation tells us that under the rotating frame,  $\vec{M}$  acts exactly the same as in the laboratory frame, *provided* the actual magnetic field  $\vec{B}_0$  is exchanged for an *effective* field  $\vec{B}_{eff}$  where

$$\vec{B}_{eff} \equiv \vec{B}_0 + \frac{\vec{\Omega}}{\gamma}. \quad (2.34)$$

This agrees with our previous analysis: consider the case wherein  $\vec{B}_{eff} = 0$ , that is, wherein  $\vec{\Omega} = -\gamma B_0 \mathbf{k}$ . This special case describes a frame of reference wherein the magnetization vector,  $\vec{M}$ , maintains its orientation with the elementary vectors  $\mathbf{i}, \mathbf{j}, \mathbf{k}$  of the reference frame, but the reference frame itself is rotating at the Larmor frequency with respect to the laboratory!

We are now ready to formulate equation (2.28) in the rotating frame: first, note that in this frame,  $\vec{B}_1^{cc}$  is static, and, since the axis of rotation coincides with the direction of  $\vec{B}_0$ , it will also be static. Then, (2.28) becomes:[43]

$$\left( \frac{d\vec{M}}{dt} \right)_{\text{lab}} = \vec{M} \times \gamma \vec{B}_1^{cc} + \vec{B}_0 \Leftrightarrow \quad (2.35)$$

$$\left( \frac{\partial \vec{M}}{\partial t} \right)_{\text{rot}} = \vec{M} \times [\gamma B_1^{cc} \mathbf{i} + (\omega_L + \gamma B_0) \mathbf{k}]. \quad (2.36)$$

This equation is equivalent to the original equation of motion (2.15), where we let[43]

$$\vec{B}_{eff} = B_1^{cc} \mathbf{i} + \left( \frac{\omega_L}{\gamma} + B_0 \right) \mathbf{k}. \quad (2.37)$$

Physically, we are simply stating that the moment acts as if it is experiencing a static magnetic field  $\vec{B}_{eff}$ . The moment precesses in a cone of fixed angle about the direction of  $\vec{B}_{eff}$  at an angular frequency  $\gamma |\vec{B}_{eff}|$ .


In this rotating frame of reference, we can see from (2.36) that the motion of the moment lies in a cone around  $\vec{B}_1^{cc}$ , and also, that it is periodic. Furthermore, if it was initially aligned to  $z$  then it will return to it in one period.

Finally, another thing the above equation is telling us in the case where we only turn  $\vec{B}_1$  on for a short amount of time is the amount by which the magnetization vector is tilted from the  $z$  axis assuming it was initially aligned with it, e.g., the flip angle.

## 2.6.2 Flip angle

The forced rotation of the magnetization vector from the  $z$ -axis down to the  $x - y$  plane in the classical sense is called the magnetization flipping. In the quantum





**Figure 2-11:** The precession of the net magnetization vector about the effective magnetic field traces out a cone whose tip is the center of the coordinate system and whose base is perpendicular to the effective field.

mechanical view, this is the equivalent of a forced shift in the ratio of nuclear spin states via individual spin “flips” after absorption of photons at the resonance frequency. The angle by which the magnetization flips (*FA*) from the  $\vec{B}_0$  axis, is dependent on the  $\vec{B}_1$  field which can itself be dependent on time (e.g., a pulse):

$$\omega_1 = \frac{d\alpha(t)}{dt} = \gamma \vec{B}_1^{cc}(t), \quad (2.38)$$

where  $\alpha(t)$  is the angle between the magnetization vector and the z-axis, the flip angle. We thus have that

$$FA = \alpha(t_e) - \alpha(t_s) = \gamma \int_{t_s}^{t_e} \vec{B}_1^{cc}(t) dt, \quad (2.39)$$

where  $t_s$  is the starting time of the field application, and  $t_e$  is the end time.

What we have seen so far suggests a way to observe magnetic resonance: we could simply put the material we wish to study, in a magnetic field  $\vec{B}_0$ , surround it with a coil whose axis is perpendicular to the  $\vec{B}_0$  field. In the thermal equilibrium there will be a net excess of protons aligned with  $\vec{B}_0$ . By applying an alternating current on the coil, thereby generating the  $\vec{B}_1$  field, and applying it for some time  $t_w$ , we could force the net magnetization to flip by 90 degrees. The magnetization would then be perpendicular to  $\vec{B}_0$  and would therefore precess at the Larmor frequency. This would produce a flux through the receiver coil, alternating as the spins precess. The induced current can then simply be observed.

An apparent problem with the above description of our MR experiment related to the magnetic dipole environment of each nucleus, is that according to the physics we have seen so far, the induced current on the coil would persist indefinitely! In reality, by several physical processes, the spins tend to return to the original equilibrium state as time progresses. The rate by which they “relax” is dependent on their chemical surroundings.

We finally turn our attention to understanding this relaxation process.

## 2.7 *Relaxation Processes*

In the absence of any relaxation or dampening factors, once the magnetization vector is flipped from the  $\vec{B}_0$  axis, and precesses, the signal induced on a nearby detector would be sinusoidal and persist indefinitely in time. In dense nuclear environments (e.g. molecules and assemblies of molecules,) the magnetization vector decays over time, via “relaxation” processes [41], a descriptive mechanism of the macroscopic effects of numerous individual nuclear magnetic dipole moments. There are two types of relaxation: longitudinal (also called  $T_1$  or spin-lattice) relaxation and transverse (also called  $T_2$  or spin-spin) relaxation. Each relaxation process affects a different attribute of the net magnetization, and each is a manifestation of a different interaction.


Once the net magnetization ( $\vec{M}$ ) has been tilted away from its alignment to the external magnetic field  $\vec{B}_0$ , it will begin to return to it, via an exponential process of characteristic time  $T_1$ . This relaxation is a manifestation of the exchange of energy between the protons and their molecular surroundings, the lattice being the thermal bath that the particles are in.

Furthermore, some exchange of energy in-between the nuclei also occurs, dipole-dipole coupling, causing phase decoherence. This is also an exponential process, with a characteristic time  $T_2$ . Intrinsically, the two are very different:  $T_1$  is related to spin flips between the allowed quantum states, and hence a return to the equilibrium state, whereas  $T_2$  is a dephasing phenomenon, wherein no spin-flips are involved, rather the presence of the spins cause local differences in the magnetic field experienced by other nuclei. The former type of relaxation ( $T_1$ ) yields information about the chemical surroundings and thus molecular structure, whereas, the latter ( $T_2$ ) yields information about the mobility of the protons.[53]

### 2.7.1 Spin-lattice ( $T_1$ ) relaxation

Protons in the higher energy state can make a transition to the lower energy state, returning to the equilibrium magnetization, by either a spontaneous or stimulated emission of a photon. Most longitudinal relaxation transitions are stimulated[5].

This stimulated relaxation is a result of other existing magnetic fields in the sample, caused by translational and rotational molecular motions in the sample. Spin-lattice relaxation occurs because these magnetic fields may be oscillating at the correct Larmor frequency dictated by the main  $\vec{B}_0$  field. The excited proton spins make a transition by releasing a photon in the form of thermal or kinetic energy, hence, an exchange of energy between the spin system and the lattice. One such example of stimulated relaxation is the result of magnetic dipole-dipole coupling. Consider a water molecule, made of two hydrogen atoms and an oxygen one. These molecules are moving randomly, creating magnetic fields by their dipole moments, which are fluctuating with time (due to the motion). In Figure 2-12, we consider the water molecule rotating about one of the hydrogen atoms, about the  $\vec{B}_0$  axis, precessing at the Larmor frequency, merely a choice of reference frame, then the other hydrogen atom experiences a sinusoidal magnetic field  $B$  acting on it, caused by the former atom.



**Figure 2-12:** Magnetic dipole-dipole coupling in a water molecule.

Returning to the quantum mechanical picture, the protons can only take a spin state that is parallel or anti-parallel to the external field, as the other fields are much weaker and can be ignored. Therefore, if magnetic fields due to the lattice cause protons to flip their spins, the net magnetization vector will be changed along the external field direction. That is, the net magnetization vector will only change along the direction of  $\vec{B}_0$ , by convention, the z-axis.


The characteristic time of this process is called  $T_1$ . If a  $90^\circ$  pulse is applied, then

the return of the longitudinal magnetization to the equilibrium is

$$\mu_z = \mu_0(1 - e^{-t/T_1}), \quad (2.40)$$

where  $\mu_0$  is the thermal equilibrium magnetization.

A number of things come out of a close examination of the spin-lattice relaxation process. First of all, in a sample of molecules, there is a large distribution of rotations. However, only rotations at the Larmor frequency can influence  $T_1$ . Since the Larmor frequency is proportional to the applied field  $B_0$ ,  $T_1$  will vary by the field strength. In general,  $T_1$  is inversely proportional to the density of molecular motions at the Larmor frequency. However, molecular rotation distribution in a sample is also proportional to the temperature of the sample. Another factor affecting  $T_1$  relaxation is the viscosity of the sample. Although the temperature of the human body does not vary enough to significantly affect  $T_1$ , the viscosity from tissue to tissue does vary significantly.




**Figure 2-13:** Dependence of the number of molecules rotating at the Larmor frequency in a sample, as a function of sample temperature or viscosity.

In the case where repeated  $90^\circ$  pulses are used, the effect of the pulse depends directly on the longitudinal magnetization left in the sample, and thus, tissues with short  $T_1$  relax quickly and produce an intense signal throughout the excitations. On the other hand, long  $T_1$ , with slow longitudinal recovery, translates to signal of lesser magnitude throughout the images or spectra.

### 2.7.2 Spin-spin ( $T_2$ ) relaxation

The second prevalent form of relaxation and signal loss over time in MRI is an effect of the tiny nuclear magnetic dipoles on each other, due to their very close proximity. A dipole, produces a small magnetic field, which falls off as  $1/r^3$ . Therefore,



**Figure 2-14:** Expected magnetic field experienced by a proton in a sample, under a completely homogeneous field  $B_0$ , giving rise to phase decoherence, or  $T_2$  relaxation.

nearby protons will be experiencing the external field  $\vec{B}_0$  with some small variation because of the other protons in the vicinity. For example, in water, a typical local magnetic field ranges over  $\pm 5 \times 10^{-4}$  Tesla. Even if the external field  $B_0$  is completely homogeneous, an incorrect assumption, different protons in a sample of water will experience a  $B'_0$  which is variable in that range.

The effect of this variation is of course that the Larmor frequency of protons in a sample varies by a small amount. Even if two individual protons begin to precess in phase, after some time, because their precession frequencies will be slightly different, they will be out of phase. In the rotating frame, this results in ablation of the  $x - y$  component of the magnetization. The characteristic time for this process is  $T_2 \leq T_1$ . It is inversely proportional to the number of molecular motions less than or equal to the Larmor frequency. The net magnetization in the  $x - y$  plane drops to zero due to this dephasing, while the longitudinal magnetization grows until it reaches  $\mu_0$  due to spin-lattice relaxation.

$$\mu_{xy} = \mu_{xy0} e^{-t/T_2}. \quad (2.41)$$

In equilibrium  $\mu_{xy} = 0$ . As was noted above, we can not assume that the external field  $\vec{B}_0$  is completely homogeneous. The effect of inhomogeneities in this field is modeled as part of the spin-spin relaxation time  $T_2^*$ :

$$\frac{1}{T_2^*} = \frac{1}{T_2} + \frac{\gamma |\Delta \vec{B}_0|}{2}, \quad (2.42)$$

where  $\Delta \vec{B}_0$  captures the inhomogeneity of the external magnetic field and  $T_2^*$  is the *observed* or *effective* spin-spin relaxation time.[53]

## 2.8 The Bloch Equations

The above sections fully explain the NMR principle and the generation, detection and decay of the MR signal. We finally put together all the equations and arrive at the full equation of motion for the net magnetization vector in the presence of a static external magnetic field  $B_0$ , an excitatory pulsed magnetic field  $B_1$  and relaxation due to both dipoles giving up radiation as well as due to the interactions in-between the dipoles.

Starting from the equation of motion of the magnetization vector, equation (2.28), we must now add the  $T_1$  and  $T_2$  dampening terms which leads to the so-called *Bloch* equations (2.45):[4]

$$\frac{d\mu_x}{dt} = \gamma(\vec{\mu} \times \vec{B})_x - \frac{\mu_x}{T_2} \quad (2.43)$$

$$\frac{d\mu_y}{dt} = \gamma(\vec{\mu} \times \vec{B})_y - \frac{\mu_y}{T_2} \quad (2.44)$$

$$\frac{d\mu_z}{dt} = \gamma(\vec{\mu} \times \vec{B})_z - \frac{\mu_z - \mu_0}{T_1}, \quad (2.45)$$

in the static or laboratory frame of reference.

When properly integrated, the above coupled differential equations will yield the  $x$ ,  $y$  and  $z$  components of the magnetization vector as a function of time. It is the Bloch equation that describes the MRI physics phenomenon.

### 2.8.1 The Bloch equations in a rotating frame

Before transferring the Bloch equations to the rotating frame of reference that was used above, we will add another component to the MRI field configuration. Consider adding a gradient to the field, aligned to the  $z$  axis,  $\vec{G}(z) = \vec{B}_g(z) = \partial B_0 / \partial z \neq 0$ ,<sup>12</sup> which will be described in further detail at a following section. This gradient is used to select a slice out of the volume under the influence of  $\vec{B}_0$ . Simply put, depending on the  $z$  position of a sample, the  $\vec{B}_0$  field is modified by the amount  $B_g(z)$ , hence modifying the Larmor frequency of the slice at that  $z$  position, compared to any other part of the sample. This allows selective excitation based on the  $z$  position. The field configuration is then

$$\vec{B}_0 = \vec{B}_0 + \vec{B}_1^{cc} + \vec{B}_g(z) \quad (2.46)$$

$$= B_1(t)\cos(\omega_L t)\mathbf{i} - B_1(t)\sin(\omega_L t)\mathbf{j} + (B_0 + B_g(z))\mathbf{k}, \quad (2.47)$$

---

<sup>12</sup>Ideally, we would also have  $\partial B_0 / \partial x = \partial B_0 / \partial y = 0$ .

wherein we approximate the  $\vec{B}_1$  field only by its counter-clockwise component as, it only, has any significant effect.

Now, in order to obtain the magnetization vector  $\vec{\mu}'$  in the rotating frame we again look at the time derivative in the rotating frame

$$\frac{d\vec{\mu}}{dt} = \frac{d\vec{\mu}'}{dt} + \vec{\omega}_L \times \vec{\mu}', \quad (2.48)$$

where  $\vec{\omega}_L = -\gamma\vec{B}_0$  (the negation comes from the fact that the precession and thus rotation frame is counter-clockwise).

The equation of motion for the magnetization vector then becomes

$$\frac{d\vec{\mu}'}{dt} = \gamma\vec{\mu}' \times \vec{B}' - \gamma\vec{\mu}' \times \vec{B}_0, \quad (2.49)$$

since the cross product  $\vec{\mu} \times \vec{B}$  is independent of the rotation of the reference frame – the two vectors maintain their mutual orientation. Clearly, the  $\vec{B}_0$  term drops out, and what we have left is

$$\frac{d\vec{\mu}'}{dt} = \gamma\vec{\mu}' \times (\vec{B}'_1 + \vec{B}'_g). \quad (2.50)$$

Then, note that in the new rotating frame the  $\vec{B}_1$  field only has an  $x$  component, as it is rotating along. That is,

$$\vec{B}'_1 = B_1(t)\mathbf{i}, \quad (2.51)$$

whereas the  $z$  gradient field remains unmodified, as the whole frame is rotating about the  $z$  axis.

We can finally rewrite equations (2.43)–(2.45) in the rotating frame, including the relaxation terms which remain unaffected:[43]

$$\frac{d\mu_{x'}}{dt} = \gamma\mu_{y'}B_g(t) - \frac{\mu_{x'}}{T_2} \quad (2.52)$$

$$\frac{d\mu_{y'}}{dt} = -\gamma\mu_{x'}B_g(t) + \gamma\mu_{z'}B_1(t) - \frac{\mu_{y'}}{T_2} \quad (2.53)$$

$$\frac{d\mu_{z \equiv z'}}{dt} = -\gamma\mu_{y'}B_1(t) - \frac{\mu_z - \mu_0}{T_1}. \quad (2.54)$$

This last set of equations completely describes the nuclear magnetic resonance principle and its solution yields the MR signal for a certain sample, whose relaxation properties are known. The importance of the Bloch equations is that in low  $\vec{B}_1$  fields, which do not tilt the net magnetization's precession angle by more than  $\sim 30^\circ$ , the signal produced by the excitation caused by  $B_1$  is linear to the input  $B_1$  excitatory field. This fact is central to the success of our non-Fourier encoding

methods: linearity of the MR process is the reason why spatially selective non-Fourier encoded MRI works.

Before we depart from the physical principles and turn our attention to our novel non-Fourier encoding methods for MRI, we must first understand how the signal acquired in an MR experiment can be made to contain spatial information and hence can be reconstructed into an image, based on the knowledge provided by the Bloch equations. The addition of the gradient magnetic field  $B_g$  in this section was a first step towards localizing the acquired signal: by using the gradient field during excitation of the sample, we can selectively excite some slice profile of the sample rather than the whole volume. However, we must now discuss how to localize the signal in the remaining two in-plane dimensions. In the next sections we will discuss slice MR imaging, with in-plane Fourier (spin-warp) encoding, attributed to [22, 9], although other approaches based on “true” 3D Fourier encoding, or 2D projection and reconstruction exist.

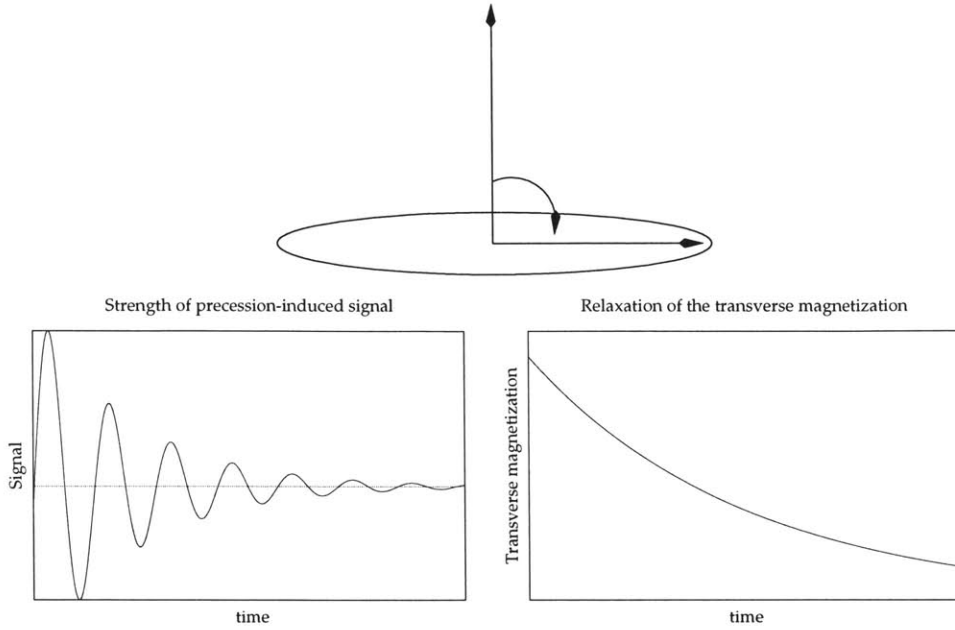
## ***2.9 Spatial Localization of the MR Signal***

So far, we have seen how each proton possesses an intrinsic spin, the ensemble of the spins giving rise to a net magnetization vector. Under the influence of an external magnetic field, if the net magnetization has an  $x - y$  component, then it precesses around the  $z$ , field, axis. This precession induces a current on a wire surrounding the sample, then, that is the observed MR signal. The frequency of the alternating current is exactly the Larmor frequency.

The problem however, lies in our inability to extract the position of the spins throughout the sample which gave rise to the signal. If we break up the sample volume into elementary cubes, called “voxels”, and consider each voxel to have its own net magnetization vector, then the problem is in finding out which voxels produce what amount of the observed signal. If we can not break up the volume into voxels, then we can not produce an image.

Clearly, all “voxels” in a sample will be precessing at the same Larmor frequency. As was suggested in the previous section, a gradient in the magnetic field can be used to modify that frequency as a function of position. Depending on when this gradient is applied, it is relatively simple to analyze its effect on the signal and hence find the location of the voxel whose magnetization produced a certain frequency or phase in the signal. The next few sections explain how a gradient can be used at different times through the MR imaging experiment, first, in conjunction with an RF excitation, to selectively excite a slice through a volume, and later without an RF excitation, modify the frequency and the phase of the signal produced





**Figure 2-15:** After a  $90^\circ$  pulse, the transverse magnetization induces a signal by precessing, called the *free induction decay*. The signal is dampened by the relaxation processes.

by an excited voxel, depending on its in-plane position.

### 2.9.1 Slice selection gradient

We begin by trying to pick out only a thin slice from the rest of the volume to be excited. If, when the RF pulse is issued, only the protons in the wanted slice match the frequency of the pulse (i.e., the Larmor frequency of these protons  $\omega_L$  matches the frequency  $\omega_p$  of the pulse) then only those protons from the entire sample will be excited and thus only they will afterwards possess a precessing  $x - y$  component and give rise to a signal. As noted above, we can apply a  $z$  dependent magnetic field  $\vec{B}_g$  of much lower strength than the main field  $\vec{B}_0$  in which case, the Larmor frequency of the protons then becomes

$$\omega_L = \gamma |\vec{B}_0 + \vec{B}_g(z)|, \quad (2.55)$$

and we can simply let  $\vec{B}_g(z) = \vec{B}_g z$ , that is, there is a linear dependence of the field and the  $z$  position. Fixing either the pulse frequency or the gradient field strength allows us to select a  $z$  position. Then, since the excitation slice will occur where[8]

$$\omega_p = \omega_L = \gamma |\vec{B}_0 + z\vec{B}_g|, \quad (2.56)$$

we can calculate the position of the slice

$$z = \frac{(\omega_p/\gamma) - |\vec{B}_0|}{|\vec{B}_g|}. \quad (2.57)$$

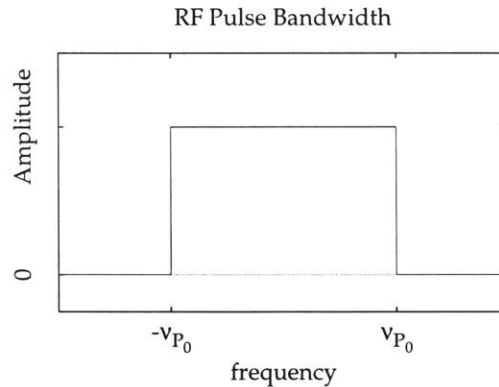
The gradient can also be dependent on time.

In reality, the RF pulse is composed of a number of frequencies covering a certain bandwidth, so as to allow a slice of finite thickness from the sample to be excited. As can be seen from the last equation, the ratio of the bandwidth of  $\omega_p$  frequencies, contained in the pulse, to the strength of the gradient field determine the thickness of the slice.[8]

$$\Delta z = \frac{\Delta\omega_p}{\gamma|\vec{B}_g|}. \quad (2.58)$$

When the strength of the gradient is increased, the same range of frequencies will excite a thinner slice, and likewise, a weaker gradient increases the slice thickness.

If we want to excite a well-defined slice, assuming that the magnetic fields are completely homogeneous, we need to apply an RF pulse covering a well-defined range of frequencies  $\Delta\omega_p = \pm\omega_{p_0}$ , with equal power throughout this range.



**Figure 2-16:** In the frequency domain, the RF pulse used to selectively excite a slice out of an entire volume, is a boxcar function, with equal coverage over a frequency range which depends on the slice thickness.

Therefore, the profile of the RF signal which will excite a slice of thickness  $\Delta z$  must have the form of a boxcar function in the frequency domain, covering the range  $[-\omega_p, \omega_p]$ . However, since the RF pulse will be applied as a time-dependent magnetic field, the actual time domain waveform to be applied is the Fourier transform of this frequency-space square pulse. The Fourier transform of the boxcar is a

$\text{sinc}(x)$  wave extending *infinitely* in time. Of course this will have to be truncated, as the RF pulse must have a finite temporal existence. This can be modeled as the convolution of the infinitely-long sinc wave with a square signal, which results in a ripple-like square-ish signal in frequency-space. For this reason, although rather accurate, slice selection can never be perfect. There is always a drop-off of magnetization at the edges of the excited slice.

Note that the slice-selection gradient is applied *simultaneously* with the RF pulse. This gradient is used in order for the pulse itself to excite a specific frequency-encoded part of the volume, therefore, the gradient must be effective precisely during the pulse play-out.

Clearly, the excited spins throughout the thickness of the slice, will be precessing at slightly different frequencies since they are experiencing a differing gradient field depending on their position within the slice. By the time the gradient is turned off, this will have induced a phase shift along the slice selection direction. After the gradient is turned off the spins begin to precess at the Larmor frequency of the external  $\vec{B}_0$  field, as it is the only contributor to the Larmor equation. This phase shift will be dependent on the position of the spin along the slice-selection gradient (for simplicity the  $z$  axis) and will be [8]

$$\phi_{r=z} = \gamma \int_0^t r \vec{B}_g dt = \gamma \int_0^t z \vec{B}_g dt = \gamma z |\vec{B}_g| t, \quad (2.59)$$

where  $r$  is the spin position along the slice selection direction, the  $z$  axis in this example. The limits of the integration are based on the simplifying assumption that halfway through the pulse, the magnetization is flipped instantaneously by the entire flip angle. In order to rephase the spins, it is then only necessary to apply a gradient of *opposite* sign, containing the same amount of power as the integral in the last equation.

So far, we have localized the induced signal to a single slice, wherein all spins are precessing at the same frequency and are in phase. Now, we need to localize the signal to a specific position on the in-plane tissue.

## 2.9.2 Frequency-encoding gradient

We can deduce one of the two remaining directions, again, by using a magnetic field gradient, in the  $\vec{B}_0$  magnetic field, to modify the Larmor frequency of precession. Arbitrarily, let this direction be the  $x$  axis. We again apply a linearly dependent magnetic field along the  $x$  axis, forcing the resonance frequency to be proportional to the position of the spin:

$$\omega_L = \gamma |\vec{B}_0 + \vec{B}_v(r)| \quad (2.60)$$

$$\begin{aligned}
&= \gamma |\vec{B}_0 + \vec{B}_x(x)| \\
&= \gamma |\vec{B}_0 + x\vec{B}_x|,
\end{aligned} \tag{2.61}$$

where  $r$  is the position of the spin in the frequency encoding direction,  $r \equiv x$ . Note that at the time when the  $x$  gradient turns on, both the slice-selection gradient as well as the excitation pulse have been turned off, otherwise equation (2.61) would not hold. Since excitation has already occurred, each spin will be precessing and will be generating an MR signal. The effect of this gradient then, equivalent to equation (2.57), is that the total signal received (that is, the sum of the signals from all the spins) will contain a number of different frequencies, wherein each frequency is generated by all the spins at the same position along the selected axis. We can now invert the problem, and clearly state that if the signal received contains a frequency  $\omega_s$ , then there must have been spins at the  $x$  position

$$\begin{aligned}
x &= \frac{\omega_s/\gamma - |\vec{B}_0|}{|\vec{B}_x|} \\
&= \frac{\omega_s - \omega_0}{\gamma |\vec{B}_x|},
\end{aligned} \tag{2.62}$$

where  $\omega_0$  is the resonance frequency corresponding to the contribution of  $\vec{B}_0$  only, and  $\omega_s$  is the precession frequency of some magnetization vector at the specific  $x$  position.

In order for this frequency encoding to be useful, we want the frequency shift, caused by the application of the gradient, to take effect during the detection of the signals. That is, we want the *detected* signal to contain multiple frequencies, wherein each frequency can be used to deduce the spatial position, along a single axis, of the spin that gave rise to that signal component. Therefore, this gradient is applied during the signal-detection period, and usually called the *readout-gradient*. Since we are free to choose which direction is frequency encoded, this direction is called the *frequency-encoded direction*.

The signal  $S(t)$  experienced at the receiving coil, assuming no relaxation, at time  $t$  after the gradient is switched on is[8]

$$S(t) = c \int \rho(x) e^{i[\gamma x \int_0^t \vec{B}_x(x) dt]} dx, \tag{2.63}$$

where  $\rho(x)$  is the spin density. Fourier transformation can of course be used to convert  $S(t)$  to  $\rho(x)$  and vice versa. In the last equation let

$$k(x) = \gamma \int_0^t \vec{B}_x(x) dt, \tag{2.64}$$

be the *wavenumber*,  $k \equiv 1/\lambda$  where  $\lambda$  is the wavelength. Of course in this case the frequency/wavelength/wavenumber depends on the position of the spin along the gradient direction, as equation (2.64) clearly shows. The wavenumber is simply the number of complete wavelengths per unit distance, and is very convenient to use when a number of different frequency signals are superposed.[13] The notion of wavenumber will become more convenient later on, when we can describe the Fourier-domain representation of the sample using the so-called “k-space”.[48]

Another important definition is the size of the *Field-Of-View*. This is the largest wavelength (smallest wavenumber) present:[8]

$$FOV = \frac{1}{\Delta k'}, \quad (2.65)$$

which, for a constant gradient field diminishes to

$$FOV = \frac{1}{\gamma |\vec{B}_x| \Delta t'}, \quad (2.66)$$

where  $\Delta t$  is time increment in-between sampling points.

We can also define the spatial resolution along the frequency-encoding direction,  $\Delta x$ , since the smallest wavelength (largest wavenumber) corresponds to the smallest feature encoded. In the case of  $n$  equidistant samples and coverage of  $\pm k_{max}$ , the resolution is simply[8]

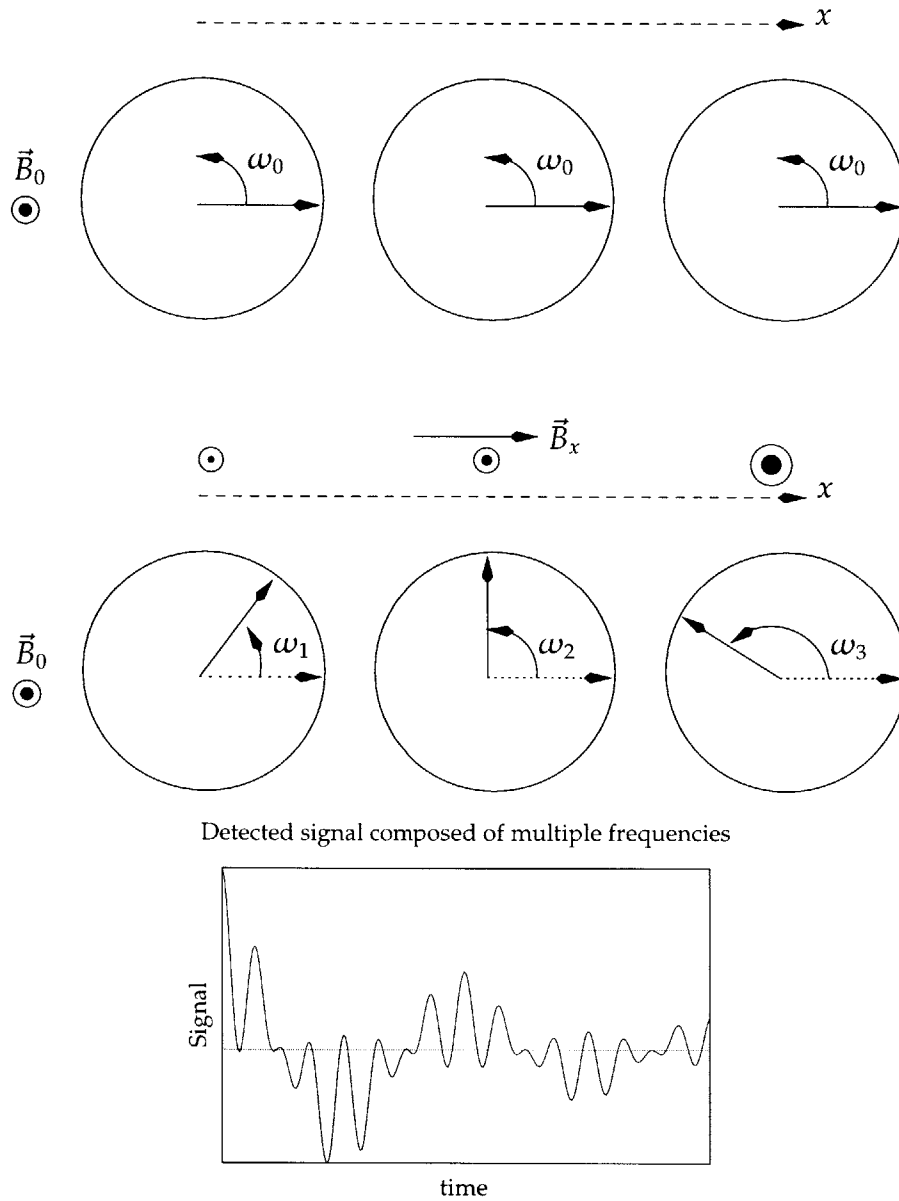
$$\Delta x = \frac{2}{n \Delta k} = \frac{2}{n \gamma |\vec{B}_x| \Delta t'}, \quad (2.67)$$

where  $\Delta k = 2 \times k_{max}$ .

### 2.9.3 Phase-encoding gradient

In order to localize the third and last orthogonal dimension, we need yet another gradient, the phase encoding gradient, in this case, the  $r \equiv y$  axis. Since, we can no longer use the frequency of the precession to encode this direction, we instead use the phase of the precession. This gradient is applied after the slice selection gradient and RF pulse, but before the frequency-encoding gradient (signal readout time). It is very similar to the frequency-encoding gradient, except that in this case, it is only the phase of the precessions that matters, after the end of the application of the gradient.

The idea is that by the end of the application of this gradient the phase of precession will be linearly related to the  $y$  location of the transverse magnetization. Consider some spins immediately after the application of a  $90^\circ$  RF pulse. That is,



**Figure 2-17:** Looking down the z-axis, before applying a frequency-encoding gradient, all spins are precessing at the same frequency. However, when the gradient is applied along the  $x$ -axis, the spins precess at different frequencies. If the signal is read out at the time when the gradient is applied, it will be composed of a number of frequencies, allowing localization of the signal amplitude along the  $x$ -axis. Note that since the signal is sampled, we need to obtain twice the samples as we want to have resolution along the  $x$ -axis, as dictated by the Nyquist limit  $\Delta\omega_{max} = 1/2T$ . Typical read-out, or, sampling time is 5-30ms.

local magnetization vectors have been rotated down to the  $x$  axis. Now, consider what happens if we apply a gradient in the same direction as the main  $\vec{B}_0$  field as always, but whose strength is dependent along the  $y$  direction. Each magnetization vector will precess about the main magnetic field, at a frequency given by the usual resonance equation

$$\omega_L = \gamma|\vec{B}_0 + \vec{B}_\phi(r)| \quad (2.68)$$

$$\begin{aligned} &= \gamma|\vec{B}_0 + \vec{B}_y(y)| \\ &= \gamma|\vec{B}_0 + y\vec{B}_y|, \end{aligned} \quad (2.69)$$

and the  $y$  position of a spin precessing at an angular velocity of  $\omega_s$  is given as usual by

$$y = \frac{\omega_s - \omega_0}{\gamma|\vec{B}_y|}. \quad (2.70)$$

While the phase-encoding gradient is on, each magnetization vector at a specific position along the  $y$  direction is precessing at a unique, different frequency, depending on its  $y$  position. So far, this is exactly the same as the frequency-encoding scheme described in the previous section.

Now, consider what happens if the gradient is turned off. Then, the field currently experienced by each spin is the same, and the magnetization vectors at each voxel will now speed up or slow down and will be precessing at the same frequency  $\omega_0$ . The phase  $\phi$ , however, of each vector will be different if we do not apply a rephaser gradient as after applying the slice selection gradient. The phase shift at the end of the phase-encoding gradient pulse will then be given by


$$\phi(r \equiv y) = \gamma \int_0^t \vec{B}_\phi(r) dt \quad (2.71)$$

$$\begin{aligned} &= \gamma|\vec{B}_y(y)|t \\ &= \gamma y|\vec{B}_y|t, \end{aligned} \quad (2.72)$$

where  $\vec{B}_y(y)$  is the linear  $y$ -axis magnetic field gradient, and  $t$  is the total time the gradient was applied for, assuming that the gradient is kept constant over time.

The question then arises as to how we can use this phase of the precessions to differentiate among the positions. The difference between frequency and phase encoding is that a phase-modulated signal resolves to one spatial frequency in the phase-encoding direction. That is, in order to resolve  $n$  spatial frequencies, a total of  $n$  different phase encodings need to be acquired.

This last part needs further clarification, which we now explain in detail. Let's consider two spins in different positions along the  $y$  axis, after we have applied




**Figure 2-18:** The phase encoding gradient after applied for some time  $t$  causes magnetization vectors to acquire some distinct phase  $\phi$  depending on their position along the phase-encoding direction. During the application of the gradient the vectors are precessing at different frequencies. After the gradient application, the vectors precess at the same frequency, but are out of phase. Note that, as is the case for all magnetic field gradients, the gradient is parallel to the  $\vec{B}_0$  field, in order to have the desired effect. The arrow in the image labeled  $\vec{B}_y$  shows the direction if increasing gradient strength (as is hinted at by the arrowheads shown underneath it).



a phase encoding gradient. Assume the two spins are the only ones present and they are situated at the same position along the  $x$  axis, that is during the read-out gradient they are precessing at the same Larmor frequency. For simplicity, also assume that there is no relaxation. Then, each of the two spins will induce a sinusoidal current in the receiver, but each of the two signals will have a different phase. However, the induced signal in the receiver coil will be the superposition of the two signals. The sum of two sinusoidal waves of different phase but same frequency, is still a sinusoidal wave, as can be seen by adding the complex exponential form of their rotating-vector oscillation description[13]

$$\begin{aligned} z_1 &= A_1 e^{i(\omega t + \phi_1)} \\ z_2 &= A_2 e^{i(\omega t + \phi_2)} \\ z_1 + z_2 &= \left( A_1 + A_2 e^{i(\phi_2 - \phi_1)} \right) e^{i(\omega t + \phi_1)}, \end{aligned}$$

where the term  $e^{i(\phi_2 - \phi_1)}$  is just a rotation through the angle  $\phi_2 - \phi_1$ , and  $A_1, A_2$  are the amplitudes.



**Figure 2-19:** The sum of two oscillations of the same frequency but different phase is still an oscillation of a single frequency and hence a single phase.

It is therefore not possible to deduce the position along the phase-encoding direction of the two spins after the application of a phase gradient, as it gives rise to a single sinusoidal wave with a single phase shift. In effect, we need one equation for each unknown. The solution in this example is to apply two different values of the phase-encoding gradient, in turn. Then, we acquire two signals, which are now

made up of two sinusoidal waves, with two different phase shifts, corresponding to the two different spin locations.

This is the idea behind phase-encoding. At each application of the phase-encoding gradient, the signal received from all the spins in the same frequency-encoding direction location, is made of a single frequency and is phase-shifted by the average contribution of all the phase-encoding direction spin locations.

By successively iterating the phase-encoding gradient strength  $|\vec{B}_\phi(r)|$  by a fixed amount  $\Delta|\vec{B}_\phi|$  and measure the signal after each iteration, The phase corresponding to each phase-encoding direction location  $r$  will in effect periodically change as a function of the gradient. This periodic change can of course be Fourier transformed to yield the position  $r$ . This is more clearly expressed in terms of the wavenumber  $k(r_\phi)$ . Multiplied by time, it gives a phase. By iteration over some increment  $\Delta k(r_\phi)$ , the subsequent measurement, of the phase at some position  $r_\phi$  undergoes periodic change as a function of  $k(r_\phi)$  much like thinking of frequency in terms of the phase at a certain time, is what happens along the frequency-encoding direction at some position  $r_\nu$  as a function of  $k(r_\nu)$ .

The number of equal increments  $\Delta k(r_\phi)$  that are used in iterating  $k(r_\phi)$  from its minimum value  $-k_{max}(r_\phi)$  to its maximum value  $+k_{max}(r_\phi)$  is the resolution along the phase encoding direction  $r_\phi$ . The time between each iteration of  $k(r_\phi)$  is the *repetition* or *recycle* time  $T_R$ , since for each phase iteration, excitation and readout must be repeated – for single echo Spin-Echo imaging, other types of sequence acquire more samples per single repetition time. In the next section we examine how to acquire the data and produce an image, by using all the gradients and RF excitations examined so far.


## 2.10 Tomographic Image Generation & Fourier Transforms

By assembling the pulse the gradients, RF excitations and readout into a “pulse sequence,” we have everything we need to create tomographic MR images. First, the slice selection gradient is turned on, while the sinc-shaped RF pulse is issued. A slice of hydrogen atoms is excited. The phase-encoding gradient is then applied causing a phase shift, followed by the readout (frequency) gradient and signal sampling during the gradient. As mentioned in the previous sections, rephasing gradients must be applied to avoid the dephasing caused by the gradients in the slice selection and frequency encoding directions.

At a certain time called  $T_E$  after the issue of the center of the RF pulse, the

coherent signal from the spins is sampled. As in the case of the readout gradient, the rephasing gradient is rather a dephasing one applied before the readout, which in effect causes rephasing of the spins.

The signal sampled at each readout consists of a single line of the  $k$ -space image matrix. The frequency encoding gradient is applied during readout, which causes



**Figure 2-20:** At each  $T_R$  a single line of the Fourier-domain of the image is sampled, each line containing all the frequency-encoded direction points, but only a single point along the phase-encoding direction.

the sampling to proceed from  $-k_v$  to  $+k_v$ , left to right in the previous image. Each subsequent phase encoding step employs a decreased-amplitude phase-encoding gradient allowing to span the  $k_\phi$  axis, the first step starting at  $+k_{max}$ . A 2D discrete Fourier transform along the rows and then columns produces the final output MR image.

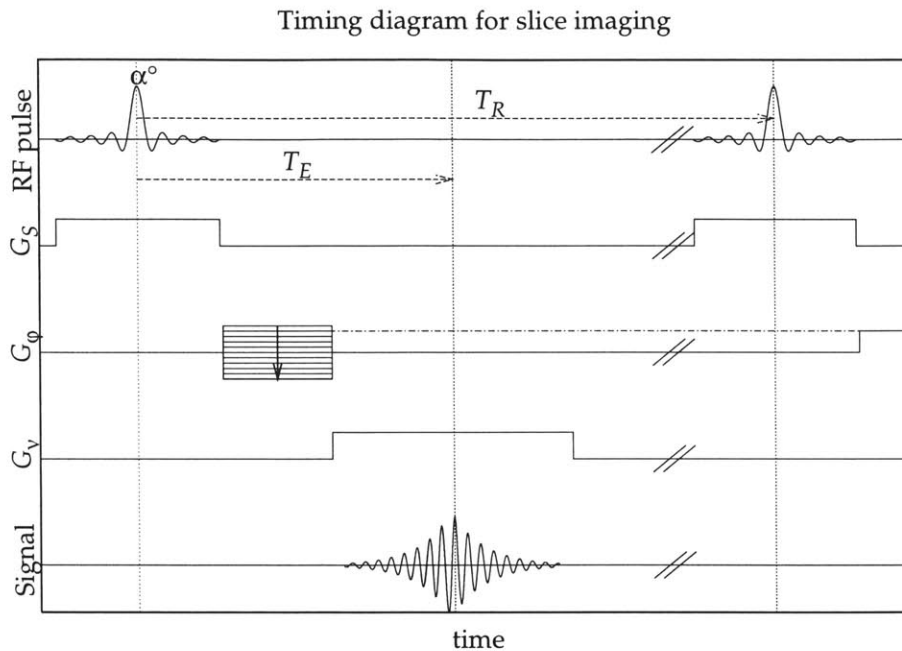
Figure 2-21 shows the (simplified<sup>13</sup>) pulse timing diagram with the application of the gradients and RF pulses of a so-called *Gradient Recalled Echo* pulse sequence.

One problem with the above timing diagram is the inability of the gradient power supplies to produce sharp square pulses. In fact all the square pulses in the above diagram will be trapezoidal, due to the finite gradient rise time<sup>14</sup>. This in turn means that  $T_E$  and  $T_R$  are increased.<sup>15</sup>

<sup>13</sup>The slice-select rephasing gradient and pre-frequency encoding dephaser lobes are not shown.

<sup>14</sup>This is referred to as the gradient "slew rate" and is the gradient strength over the rise time, in units of Tesla/m/sec.

<sup>15</sup>One way around this problem is so-called *spiral sampling* through  $k$ -space, instead of raster-like (line) sampling. Spiral sampling uses sinusoidal-like gradients of increasing amplitude, a much more natural use of the underlying electronics. Spiral sampling holds promise for faster imaging and is the focus of our continuing research.



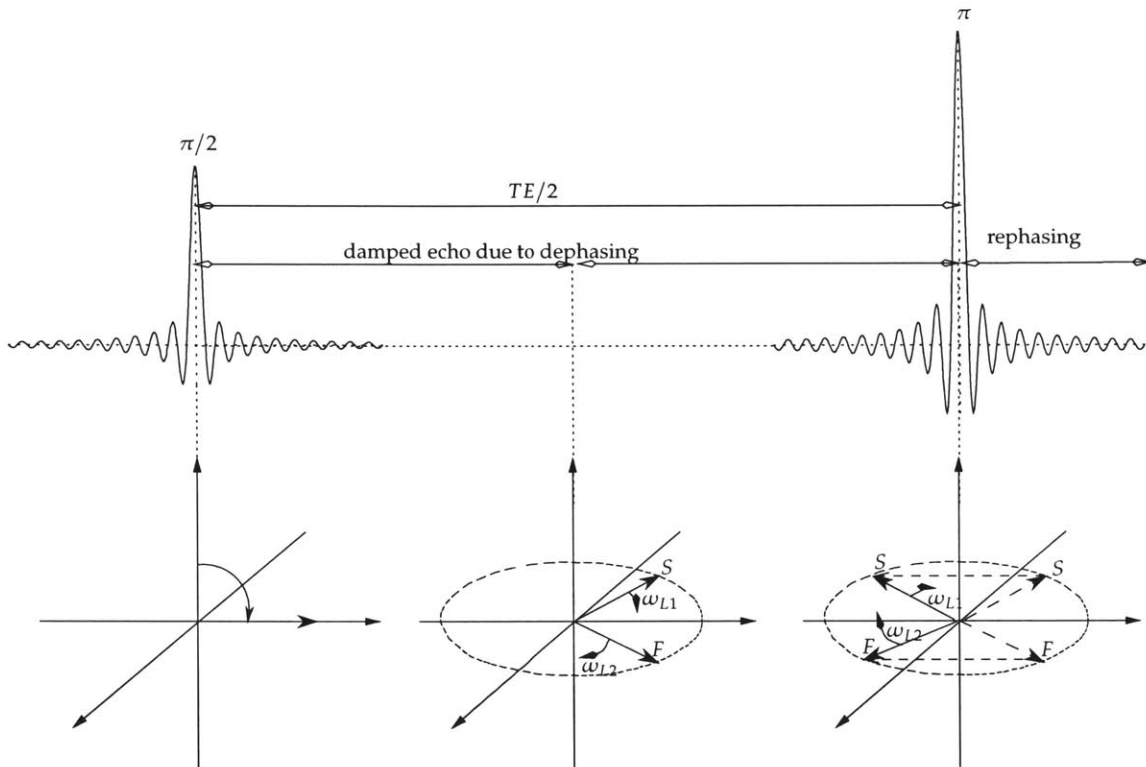
**Figure 2-21:** Gradient Recalled Echo pulse sequence timing, showing the applications of the gradients, the RF pulse play-out and sampling time. With the addition of the rephasing lobe for the slice-selection and dephaser lobe for the frequency-encoding gradient, the above pulse sequence when used on an MR scanner will produce a matrix containing the Fourier-domain representation of the 2D image.

## 2.11 Spin-Echo Pulse Sequence

Soon after the first NMR experiments, it was observed that if in-between the  $90^\circ$  excitatory pulse and the signal acquisition, a  $180^\circ$  excitatory RF pulse is used, then some time after it, a “spin echo” signal was acquired.[19] The reason for this increased signal is easily explained via the physics of spin-1/2 nuclei and is a manifestation of the  $T_2^*$  relaxation introduced above, wherein the interference of the magnetization vectors becomes constructive.

Consider the first RF pulse issued in this experiment, a  $\pi/2$  radian pulse. Almost immediately after this excitation, at time  $t_0$ , the longitudinal magnetization is rotated down to the  $x$  axis. This magnetization is a coherent sum of all the spins and induces a detectable signal. However, according to  $T_2$  effects, the spins will be precessing at slightly different Larmor frequencies due to local field inhomogeneities. Some spins will be precessing faster than others. Let us consider only two spins, a fast one ( $F$ ) and a slow one ( $S$ ).

The result of this difference in precession frequency is that some time later, at  $t_1 = \tau + t_0$ , the two spins  $F$  &  $S$  will have acquired different phases,  $\alpha$  and  $\beta$  respectively. Now, let us issue a  $\pi$  radian RF pulse. This *refocusing* pulse has the effect of “mirroring” the two spins about the  $x$  axis, since it represents a  $180^\circ$  rotation about the  $y$  axis.<sup>16</sup> In effect, after the refocusing pulse, the two spins have the same angular relation to the  $-y$  axis as the originally had to the  $+y$  axis. During the next  $\tau$  seconds after the refocusing pulse,  $F$  and  $S$  will acquire another  $\alpha$  and  $\beta$  phase respectively, and will thus become aligned along the  $-y$  axis, at time  $2\tau$  after initial excitation by the  $\pi/2$  RF pulse.[38]



**Figure 2-22:** The RF pulses used in a Spin Echo sequence. The second  $\pi$ -radian pulse has the effect of refocusing the rapidly dephasing spins, at time  $2\tau = TE$  after the initial excitation.

Since the above holds for all phases the spins may have acquired before the refocusing pulse, all the spins will align at the  $-y$  axis at that time, the strong

<sup>16</sup>The  $\pi$  pulse is exactly like the  $\pi/2$  pulse; the latter represents a rotation about the  $y$  axis bringing the magnetization from the  $z$  to the  $x$  axis. The former represents rotation about the  $y$  axis forcing the  $x$  component of the spins to end up in the  $-x$  axis.[10, 38]

transverse magnetization from the alignment will cause an “echo” induced on the receiver at that time. The Spin Echo sequence results in the refocusing of the dephasing effect of field inhomogeneities that are constant throughout the sequence time, but does not refocus fluctuating inhomogeneities. The clinical usefulness of the Spin-Echo sequence stems from its ability to produce very high signal-to-noise ratio images, and remove inhomogeneities due to the construction of the imperfect magnets used.

In our non-Fourier encoding methods, we will use a Spin-Echo sequence. Apart from the above two reasons, a Spin-Echo is desirable, as we will use the initial excitatory pulse to do spatial selection, and not slice-selection. In this case, we can still use the  $\pi$  pulse to perform slice-selection in our sequences.<sup>17</sup> It is now time to turn to the theory of spatially-selective non-Fourier basis RF encoding.

---

<sup>17</sup>Note however that other methods can be used to perform slice selection, however, we chose Spin-Echo as it is a well established sequence in the clinical world.

## SECTION 3

---

# *Advanced Non-Fourier MR Imaging*

The last section, covered standard MR imaging, as it is commonly used in everyday clinical settings. We have seen examples of both Fourier encoding, which was used to in-plane encode precessing spins, as well as spatially-selective encoding, which was used for slice-selection. It is the case that slice-selection is just one of an infinite space of spatially-selective excitations. This degree of freedom can be used as an advantage to encode much more than just a slice, in particular, it may be used to encode any arbitrary spatial profile.[21, 25, 36, 51, 55, 32, 34, 58] In this section we will follow the aforementioned references and describe spatially-selective excitations based on the linearity of the MR imaging process. We will then turn our attention to *efficient* encoding methods such as encoding by the *Singular Value Decomposition* (a special case of *Rank-Revealing Orthogonal Decompositions*) [55] and finally, introduce the usage of non-orthogonal vector sets such as a small number of overlapping, shifted, Gaussian distributions that can be used to acquire and reconstruct an MR image.[28, 58] Experimental results from the RROD methods described in this Section will be provided in Section 4.

### 3.1 Slice Selection in Light of Spatially-Selective Excitation

In the case of slice selection, with in-plane Fourier encoding and slice-selection along the  $z$  axis, the acquired signal equation (2.63) becomes in the time domain:[34]

$$S(k_x, k_y, z_0) = \int \int \int s(x, y, z) P(z - z_0) e^{-i(k_x x + k_y y)} dx dy dz, \quad (3.1)$$

where  $k_x = \gamma G_x$ ,  $k_y = \gamma G_y$ , and where the spatial excitation profile  $P(z - z_0)$  is ideally zero everywhere but a thin slab of thickness  $\pm \Delta z$ , centered about  $z_0$ , captured by the function  $P$ . The 2D Fourier transformation of (3.1) with respect to  $k_x$  and  $k_y$  yields an image after integration along the  $z$ -axis, this last step representing the projection of the contents of the slice onto the imaging plane:

$$s_{img}(x, y) = \int s(x, y, z) P(z - z_0) dz, \quad (3.2)$$

where  $s(x, y, z)$  is the Fourier-transformed acquired signal, hence in the spatial domain, and  $s_{img}(x, y)$  represents the reconstructed image.[34]

Roughly speaking, the result of the application of the RF pulse without any gradients playing out at the same time, is simply an excitation of all spins in the FOV volume, *iff* the Larmor frequency of the spins matches that of the applied RF. Once gradients are used, the form of the RF pulses can be manipulated to selectively excite spins in the volume, as the operator  $P$  acted in the last few equations. Intuitively, using a gradient along any one direction, in theory, any profile can be excited, along that direction, so long as the Fourier transform of the desired profile is used as the RF excitation pulse. This is equivalent to the sinc RF that was applied to select a slice profile. The desired profile was covering a certain frequency bandwidth, while the gradient applied was in the  $z$  direction. The axis of the gradient gave rise to encoding along that axis, where the sinc RF pulse used was the Fourier transform of the boxcar: the desired profile.

An important twist of (3.1) is that if a number of acquisitions are made, each with a different profile  $P_i$  along the spatially-selective dimension, it is possible to fully “encode” that dimension, if for example a full, orthogonal, set of profiles is used. Then, that entire dimension can be reconstructed upto a maximum resolution equal to the number of profiles used, simply by applying an appropriate inverse transform for the set. To this end we will shortly introduce a discretized matrix representation of this process which will clarify this.

As an example, let us assume that we wish to excite an arbitrary profile given by  $p(t)$ , along some arbitrary axis  $d$ . In our experiment we will use a gradient field



such that at position  $d_\alpha$  along  $d$ , the magnetic field will have strength  $|B_0 + Gd_\alpha|$  (this is called a gradient offset of  $G$  and is usually of the order of Gauss in units of Gauss/cm). The excitation will be performed using  $P(\omega)$ , the Fourier transform of  $p(t)$ . Under some assumptions we will shortly turn to, the resulting profile of this experiment will be  $P(d)$  where  $d = \omega/\gamma G$ . [34]

The above-mentioned results hold only when the response of the magnetization vector is linear with respect to the input RF field. It is well known in MR that this is the case [21, 20] in the *low flip-angle* regime, that is, when (2.39) leads to  $FA \leq 30^\circ$ .<sup>1</sup>

In the next section we will give a short version of the well-known, low flip-angle, linearity principle [21, 43, 36]. Furthermore, we will follow [34, 32] in using a matrix representation to describe the spatially-selective encoding process, rendering it amenable in a generalized framework for non-Fourier encoding, the basis for the work presented in this thesis.

## 3.2 Linearity of the MR Process in the Low Flip-Angle Regime

Let us reconsider the Bloch equations, in the absence of relaxation,<sup>2</sup> and low values of  $\vec{B}_1$ .<sup>3</sup> Then, in the rotating frame, equation (2.36) can be written as: [21]

$$\begin{pmatrix} \frac{dM_x}{dt} \\ \frac{dM_y}{dt} \\ \frac{dM_z}{dt} \end{pmatrix} = \gamma \begin{pmatrix} 0 & \vec{x}G & -B_{1,y} \\ -\vec{x}G & 0 & -B_{1,x} \\ B_{1,y} & -B_{1,x} & 0 \end{pmatrix} \times \begin{pmatrix} M_x \\ M_y \\ M_z \end{pmatrix}, \quad (3.3)$$

where we let the  $\vec{B}_g$  gradient field be linear, equal to some amplitude  $G$  times the position  $\vec{x}$ ,<sup>4</sup> and where the  $\vec{B}_1$  field is composed of the  $B_{1,x}$  and  $B_{1,y}$  components.

In the small flip-angle regime, the longitudinal magnetization is assumed to be constant, [21] e.g.

$$M_z \approx M_0 = \text{constant}. \quad (3.4)$$

---

<sup>1</sup>Note that methods are known to surpass this limitation; for example, an entire class of *inherently refocused pulses* exists, which may be concatenated to produce large flip angle excitations [37], or, alternatively, an iterative, error back-projection method, based on a numerical solution to the Bloch equations, may be used to “massage” arbitrary RF pulses so that they produce the expected excitation profile, only at a much higher flip angle. [23]

<sup>2</sup>This leads to the simple requirement that the RF pulse width be shorter than  $T_1$  and  $T_2$ .

<sup>3</sup>Low enough to avoid saturation.

<sup>4</sup>Recall that these equations are for the rotating frame of reference, wherein the magnetization vector is aligned to the  $x$  axis.

In order for this to be the case, we must have that  $\sin(FA) \approx FA$ ,  $FA$  in radians.

Given this, the last component of (3.3) can be decoupled from the others. First, define the transverse magnetization as a vector on the complex plane,

$$M_{xy} = M_x + iM_y, \quad (3.5)$$

and similarly, the applied RF field as

$$B_1 = B_{1,x} + iB_{1,y}. \quad (3.6)$$

Now, assume that the magnetic field gradient is constant along  $x$  during the RF excitation pulse and the transverse magnetization is governed by[36]

$$\frac{dM_{xy}}{dt} = -i\gamma GxM_{xy} + i\gamma B_1M_0. \quad (3.7)$$

The solution of this differential equation assuming an initial condition of  $M_x = M_y = 0, M_z = M_0$  at time  $T$  is

$$M_{xy}(x) = i\gamma M_0 \int_0^T B_1(t)e^{-ixk(t)} dt, \quad (3.8)$$

where  $k(t)$  defines a path through frequency space. This last equation, (3.8), appears like a familiar Fourier transform of  $B_1(t)$  with respect to time. Further analysis [36] shows that if we define a parametric path through space  $p(k)$  then the transverse magnetization becomes the Fourier transform of the  $k$ -space trajectory, weighted by  $B_1(t)/|\gamma G(t)|$ .

Since the transverse magnetization is linearly related to the input RF field, and since the induced MR signal is a manifestation of the transverse magnetization, it follows that the MR signal is also linearly related to the RF field.

Given the linearity of the MR signal in the low flip-angle regime, we will now explore a matrix representation of the MR imaging process,[34, 32] as this representation lends itself to the linear-algebraic manipulations which lie at the heart of efficient encoding methods.

### 3.3 Matrix Representation of RF Encoding

Based on the linear response of MR Imaging in the low flip-angle regime, we can reformulate the excitation and acquisition of MR as a simple matrix equation, as was shown by Panych et al. [34, 35] The value in this change of representation is that it allows us to ask whether any of the well-known linear-algebraic methods have any meaning in MRI, and most importantly, whether they are useful,

which translates to increased efficiency in the MR imaging process over the standard methods employed.

In reality, RF pulses, as used on typical MR scanner waveform generators, are not continuous as the last few sections may have misled one to believe, since the waveform generators produce digital waveforms, e.g. digital approximations. We will claim that discretization does not pose any problems to the above analyses, if the RF pulse is approximated using a train of *hard* pulses.[46] A hard pulse is a discrete approximation to  $p(t)$  wherein the power enveloped by  $p(t)$  in-between time  $t$  and  $t + \Delta t$  is delivered by a delta function (impulse) at time  $t$ . Such a hard pulse tips the magnetization off of the magnetic field axis at time  $t$ , *instantaneously*. With this approximation, the usual *superposition* principle of linear systems holds. That is, the result of some excitation profile  $p(t)$  approximated by a hard pulse train  $p_H(t)$  such that [34]

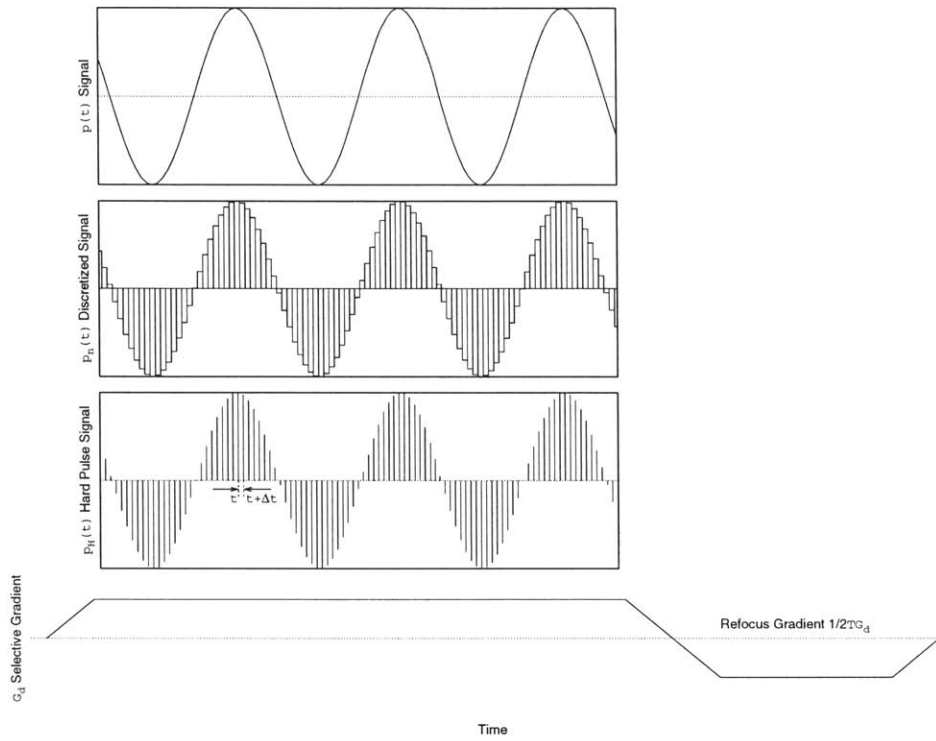
$$p_H(t) = \sum_n p_n \delta(t - n\Delta t), \quad (3.9)$$

wherein  $\delta$  is the Kronecker delta, will be the sum of the linear responses to each of the elements composing  $p_n$ , so long as the excitation remains in the low flip-angle regime.

The RF pulse is not the only continuous component of the system that must be approximated by a discretized digital signal. Once the magnetization has been flipped and is precessing, the induced signal on the receiver will also be continuous. The receiver electronics must digitize the signal so that it can be later processed to produce an image. Once the continuous signal is sampled, we have introduced a sampling or *point-spread function* (PSF). Imaging systems tend to have a Gaussian point-spread profile. However, as was pointed out by Panych [29, 30] due to the finite-length FOV of the MR imaging system, the effective PSF is best approximated by a sinc function.

Now, consider an ordered sequence of complex numbers (magnitude and phase)  $[c_1, \dots, c_k]$  being applied as hard pulses during a gradient offset. The  $i$ th magnitude and phase will be applied at time  $t_i$  as a hard pulse, which assuming that the length of the RF pulse is  $t_p$  would be  $i\Delta t_p$  where  $\Delta t_p = t_p/k$ . We can formulate this sequence as a function exactly as in equation (3.9), wherein the  $\delta$  function becomes an elementary unit hard pulse (of course any linear combination of the  $c_i$  will produce a predictable effect as well). This is merely a reformulation of  $p_n(t)$  as a row vector  $\mathbf{p}_n$  describing the RF pulse.

This RF excitation will produce some nuclear induction signal,  $f(t)$ , on the receiver which will have to be sampled during some frequency encoding gradient. This sampling can be modeled [34] as a function, in the temporal domain, that in-



**Figure 3-1:** An arbitrary spatially-selective bipolar RF pulse is approximated by impulses delivering the RF power while a gradient along the desired direction is activated. The spatial profile excited by  $p_H(t)$  along the direction of the gradient will be approximately the Fourier transform of  $p(t)$ .

tegrates the signal on the receiver coil in-between increments  $\Delta t$ . That is,  $f_i$ , the integral of the induced signal over a window of length  $\Delta t$  beginning at some time  $t_i$  becomes the  $i$ th MR sampled signal. more accurately [34]

$$f_i = \int f(t)\Phi(t_i - i\Delta t_f)dt \quad (3.10)$$

$$= \int f(t)\Phi_i(t)dt. \quad (3.11)$$

The function  $\Phi_i$  serves the purpose of first being a window of length  $\Delta t_f$  in-between time  $i\Delta t_f$  and  $(i + 1)\Delta t_f$  and absorbing the characteristics of the receiving coil and other such details.

Suppose then, that the MR signal density, that is, the representation of the spin precessions,<sup>5</sup> in the temporal (Fourier) domain, is captured by some discretized matrix  $\mathbf{F}_{(n,m)}$ . Each row of this matrix is the response of the MR system to a  $c_i = (\mathbf{e}_j)_i$   $1 \leq i \leq n$ , where  $\mathbf{e}_j$  is the  $i$ th row of the identity matrix  $\mathbf{I}$ . Then, the response to some arbitrary RF pulse  $p'_n$  was shown by Panych et al [34, 32] to be simply

$$f_i = \sum_j (\mathbf{p}'_n)_j \mathbf{F}_{(j,i)}. \quad (3.12)$$

Call the train of  $m$  samples acquired the an ordered sequence  $[f_1, \dots, f_m]$ . Then, equation (3.12) can be written as a vector-matrix multiplication leading to a row vector response  $\mathbf{f}$ :

$$\mathbf{f} = \mathbf{p}_n \mathbf{F}. \quad (3.13)$$

The final MR image is the 2D discrete Fourier transform of  $\mathbf{F}$ . To summarize, in the last equation, if  $\mathbf{p}_n$  is a train of RF hard pulses played out during a magnetic field gradient, then the acquired signal during readout the response  $\mathbf{f}$ . Obviously, it is  $\mathbf{F}$  that we need in order to produce the MR image, by applying a 2D discrete Fourier transform:

$$\mathbf{S} = \mathbf{W}^* \mathbf{F} \mathbf{W}^*, \quad (3.14)$$

where  $\mathbf{W}$  is the 2D DFT matrix and the superscript  $\star$  signifies Hermitian conjugation.

In order to produce the image, it is then necessary to first obtain enough samples  $\mathbf{f}_k$  each from a different RF pulse  $\mathbf{p}_k$  that is orthogonal to all other pulses, so

<sup>5</sup>Or more accurately, the *free induction decay*, since the signal decays by relaxation as was explained in Section 2.

that equation (3.13) can be inverted and yield  $\mathbf{F}$ . The number of excitations and acquisitions becomes the resolution of the image along this spatially-selective direction:

$$\mathbf{F}'_{k \times n} = \mathbf{P}_{k \times n} \mathbf{F}_{n \times k} \quad (3.15)$$

So long as  $\mathbf{P}$  is invertible, and thus  $k \equiv n$ , the acquired signal from the scanner  $\mathbf{F}'$  can be used to yield the Fourier domain image.[32, 34] Multiplying both sides by the inverse of  $\mathbf{P}$ :

$$\mathbf{P}^{-1} \mathbf{F}' = \mathbf{P}^{-1} \mathbf{P} \mathbf{F} \Leftrightarrow \quad (3.16)$$

$$\mathbf{F} = \mathbf{P}^{-1} \mathbf{F}'. \quad (3.17)$$

In fact, orthogonality and squareness of  $\mathbf{P}$  (so that it is invertible,) is not even a necessary condition; the least stringent condition is that  $\mathbf{P}$  has a left-side inverse, known as the Moore-Penrose *pseudo-inverse*  $\mathbf{P}^\dagger$ . Of course, different RF matrices  $\mathbf{P}$  will produce images with different desirable qualities, such as *Signal-to-Noise Ratios*, but we are not currently concerned with this.<sup>6</sup>

From equation (3.13) we can see that there is a vast space of RF pulse sequences that can be used to obtain enough information to yield an MR image, in fact, an infinite number of them. As a first step into the following sections, where we describe some of the most efficient ways to acquire the MR images as described by Zientara et al [58], let us first try and gain some insight by performing an alternative MR experiment which will yield a 2D Fourier-encoded image, but which uses spatially-selective RF excitations to do so. As most readers will have realized, this is accomplished by using the rows of the identity matrix  $\mathbf{I}$  as RF pulses, as was first explained in [34].

### 3.3.1 *Fourier-encoded MRI by spatially-selective RF pulses*

According to equation (3.13), if we use the rows of the  $n \times n$  identity matrix as RF hard pulses, each row being one square pulse, and each pulse played out at  $TR$  time after the previous one, with the signal acquired at  $TE < TR$  time after the pulse play-out, then the resulting matrix we will get should be the Fourier domain representation of the volume. This is the case since the square RF impulses that the identity matrix would represent as a waveform is the “best” approximation to the  $\delta$  functions in equation (3.9).

---

<sup>6</sup>An in-depth theoretical approach to SNR for spatially-selective RF pulses can be found in [29].

We can use a Spin-Echo sequence so that the  $\pi$  pulse performs the slice selection, in which case the matrix acquired should be the Fourier domain representation of the selected slice. The pulse sequence which achieves this result is shown in Figure 3-2.

Let us consider the effect of the  $i$ th excitatory pulse,  $e_i$ , which will be played out at time  $i \times TR$  from the starting time. Assume that the net magnetization is fully relaxed after one  $TR$ . The length of the excitatory RF pulse,  $t_p$  is partitioned in  $n$  equal-length segments,  $\Delta t_p$ , and the  $i$ th excitation involves a single hard pulse at  $i \times TR + i \times \Delta t_p$ . A spatially-selective gradient along the  $y$  axis ( $G_y$ ) is turned on during the entire RF pulse width, that is, starting at time  $i \times TR$  and until time  $i \times TR + n \times \Delta t_p$ . A rephase gradient equal to one half the area of the encoding gradient is applied immediately afterwards. Define the time  $t_0 = i \times TR$ .

Then, consider what happens during the  $i$ th repetition interval: the gradient is turned on at  $t_0$ , but the magnetization is not flipped until  $t_1 = t_0 + i \times \Delta t_p$ . Therefore, the gradient has no effect until  $t_1$ . At this time, the single hard pulse for this repetition interval is issued. The net magnetization immediately flips down to the  $x - y$  plane and begins to precess. Since the  $y$  gradient remains turned on until the end of the RF pulse width, at time  $t_2 = t_0 + n \times \Delta t_p$ , the precessing magnetization obtains a phase increment proportional to  $(n - i) \times \Delta t_p$  and, of course, also proportional to its position along the  $y$  axis. The phase shift incurred by a spin at position  $y_s$  after the  $i$ th excitation, will be

$$\phi_i(y) = \gamma \int_{t_0+i\Delta t_p}^{t_0+n\Delta t_p} G_y y dt \quad (3.18)$$


$$= \gamma(n - i)\Delta t_p G_y y. \quad (3.19)$$

Since the rephasing gradient is always applied for half the time, after the application of the rephasing gradient, the spin will have attained a total phase shift of

$$\phi'_i(y) = \gamma(n - i - \frac{n}{2})\Delta t_p G_y y. \quad (3.20)$$


Since the last equation is dependent on the variation of  $i$  only, from  $i = 0, \dots, n - 1$ , the result will be a periodic change of the phase over a fixed increment, starting at some maximum value  $+k_{max}(y)$  to some minimum value,  $-k_{max}(y)$  with increments  $\Delta k(y) = \gamma\Delta t_p G_y y$ . This is exactly equivalent to phase-encoding via a gradient after excitation. Starting from the top of the Fourier domain, each RF pulse  $e_i$  will produce the  $i$ th line:

Clearly, the matrix acquired after all  $n$  excitations will be  $\mathbf{S}$  which can be directly Fourier-transformed to yield the final MR image. Note that in the absence of the



**Figure 3-2:** A spatially-selective pulse sequence which produces a Fourier-encoded image. Each excitatory RF pulse is a row of the identity matrix, subsequent slice selection accomplished by a  $\pi$ -degree RF refocuser and slice selective gradient. The signal from that slice is then sampled. In this sequence, the  $x$  axis is frequency encoded, the  $y$  axis is phase encoded, and the  $z$  axis is slice-selection.





**Figure 3-3:** By using the rows of the identity matrix as spatially-selective RF pulses, a standard Fourier-encoded image can be acquired. In this case however, it is intrinsically the RF pulse which encodes the amount of phase offset for each line of  $k$ -space *during* excitation, and not a gradient applied *after* excitation, as is normally done.

$\pi$  refocusing pulse, we would not have been able to perform slice selection. We can only use the excitatory RF pulse in order to encode one dimension, albeit an arbitrary one, at a time.

One last bit of intuition comes from considering an RF pulse that is a linear combination of  $\mathbf{e}_i, i = 0, \dots, n - 1$ , when played out during a spatially-selective gradient in one of the two standard directions. First of all, any spatially selective RF pulse is such a linear combination. The effect of such a pulse, is to excite each line of  $k$ -space by a varying amount, in the direction of the gradient and hence during readout in the other direction, the equivalent linear combination of all the lines is frequency-encoded and obtained. This is readily incorporated in arbitrary orthogonal directions:  $k$ -space is no longer excited in the vertical and acquired in the horizontal direction, instead, if a linear combination of the gradients is used, the excitation happens along some direction in both Cartesian dimensions.

In that respect, a spatially-selective RF hard-pulse train is nothing more than merely a *weighting* function ( $w$ ). In the spatial domain, one can think of this as a weighted projection of the contents of the FOV ( $\Omega$ ):

$$Sample_x = \int_{\Omega} w(y) S(x, y, z) dy dz. \tag{3.21}$$

In the MRI experiment described above with the identity matrix as RF pulses, we simply used Fourier basis functions as the weighting functions.

### 3.4 Dynamic Adaptive MRI

Based on the above intuition, it was first proposed in [51] that spatially selective RF pulses can be used to obtain an incomplete image matrix  $\mathbf{S}_{app} \approx \mathbf{S}$ , by using Wavelet-shaped RF pulses. With some loss of information, the number of wavelet functions used could be tailored depending on the use of the images. For example, in a interventional MR surgical or therapy procedure setting, it is important that images be produced with very high temporal resolution, and even though spatial resolution can not be sacrificed, at least not inside the region of interest, some level of detail can be sacrificed. Spatial resolution is of importance in the context of MRI since most current imaging is already performed at extremely low resolution. For example, brain imaging is performed over a  $24^2\text{cm}^2$  FOV, at a resolution of  $256^2$ , and a slice thickness of 5mm, yielding a voxel size of approximately  $0.94 \times 0.94 \times 5\text{mm}^3$ . Wavelet MRI encoding was first implemented by Panych et al.[31]

The importance of the above is that in many cases it is more useful to produce images faster, rather than detailed ones capturing *exactly* all the information of the underlying sample. Note however, that there is no *error-free* MR image since usual SNR in MR images is between 20 and 80 due to thermal noise.

An alternative motivation is that in fact, MR images, like most images, disregarding noise, tend to be *rank-deficient*. The *rank* of a matrix is the dimension of its column or row space, that is, the dimensionality of the space spanned by the vectors that compose the columns or rows of the matrix.

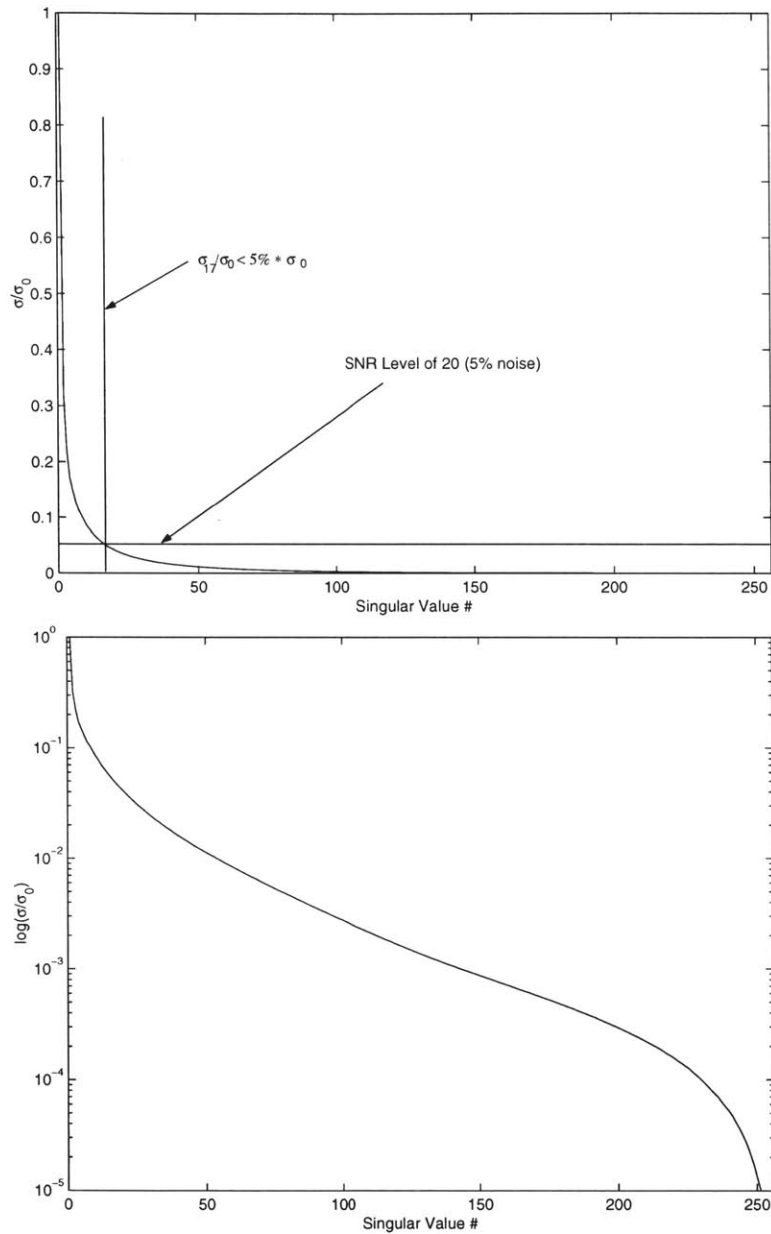
Figure 3-4 shows the *average* of the 256 singular values<sup>7</sup> computed from 84 images, each  $256 \times 256$ , that were randomly selected from the Brigham And Women's Hospital digital MR Image database. The images ranged in contents from torso and spine to brain, in both axial and sagittal planes, and from different pulse sequences and different (unknown) human subjects.

Note that for an SNR of 20 which is typical of MR images in the low flip-angle regime (SNR of 15 – 20), on the average, the 17th singular vector (out of 256) falls below the SNR significance level of 5%! What this effectively means is that any more of the 239 remaining rank-one images are insignificant compared to thermal noise in the system, and hence can not encode useful information.


What we have tried to convey thus far is that in some cases very high spatial resolution and fine detail is required (e.g. diagnostic imaging prior to surgical planning), which can currently be achieved, although at an exorbitant cost, time-wise. However, in many cases, very high temporal resolution is required, rather

---

<sup>7</sup>The singular values of a matrix  $\mathbf{F}$  are the square roots of the eigenvalues of  $\mathbf{F}^T\mathbf{F}$  or equivalently  $\mathbf{F}\mathbf{F}^T$ . The next section introduces the Singular Value Decomposition in some detail.



**Figure 3-4:** Eighty four  $256^2$  images were randomly chosen from the Brigham And Women's hospital MR Image database, their singular values computed, and the average of those is plotted against their relative order.



**Figure 3-5:** Four of the images used to generate the graph in Figure 3-4.

than detail, but spatial resolution should not be sacrificed. Currently, this can not be achieved for some of the better MR modalities such as Spin-Echo imaging.

It is in this setting that linear-algebraic techniques, rank-revealing orthogonal decompositions in particular, can be successfully applied to compute more efficient encoding basis than Fourier basis functions. We wish to produce a series of MR images of possibly time-varying FOV contents, as fast as possible, *without* reducing the spatial resolution, but possibly dropping some detailed information.

In the following discussion we will assume that the MR signal intensity function  $S(x, y, z)$  when in-plane localized and discretized is represented by  $\mathbf{S}$  a square,  $n \times n$  matrix. We must manipulate some non-square complex-valued RF matrix  $\mathbf{P}$  of size  $k \times n, k < n$ , composed of row-vector, spatially-selective, RF hard pulses in order to be able reconstruct a good estimate of  $\mathbf{S}$ . Each row of  $\mathbf{P}$  will be played out during one repetition time  $TR$  and during a gradient  $G_d$  along some direction  $\mathbf{j}$ . Readout will occur during a gradient  $G_v$  along a direction  $\mathbf{i}$  orthogonal to  $\mathbf{j}$ . We will assume a Spin-Echo sequence wherein the  $\pi$  pulse is used to perform slice selection along the third orthogonal direction  $\mathbf{k}$ . The result of the application of  $\mathbf{P}$ ,  $\mathbf{S}_{acq}$  will be the data acquired from the MR scanner. The total time for the excitation and acquisition of  $\mathbf{S}_{acq}$  will clearly be  $k \times TR < n \times TR$ . An explanatory block diagram of this method appears in Figure 3-6

Alternate representations of a matrix  $\mathbf{S}$  based on *invertible* transformations have been very successfully used in the image processing field since first introduced by Andrews and Pratt [2]. Such transforms are useful in e.g., bandwidth and dimensionality reductions, which can significantly reduce the cost of processing. The transformations we are mostly interested in are dimensionality reductions.

### 3.4.1 Invertible separable unitary transforms

A unitary matrix  $\mathbf{A} \in \mathbb{C}^{n \times n}$  is unitary if its inverse is its hermitian conjugate:[47]

$$\mathbf{A}^{-1} = \mathbf{A}^*. \quad (3.22)$$

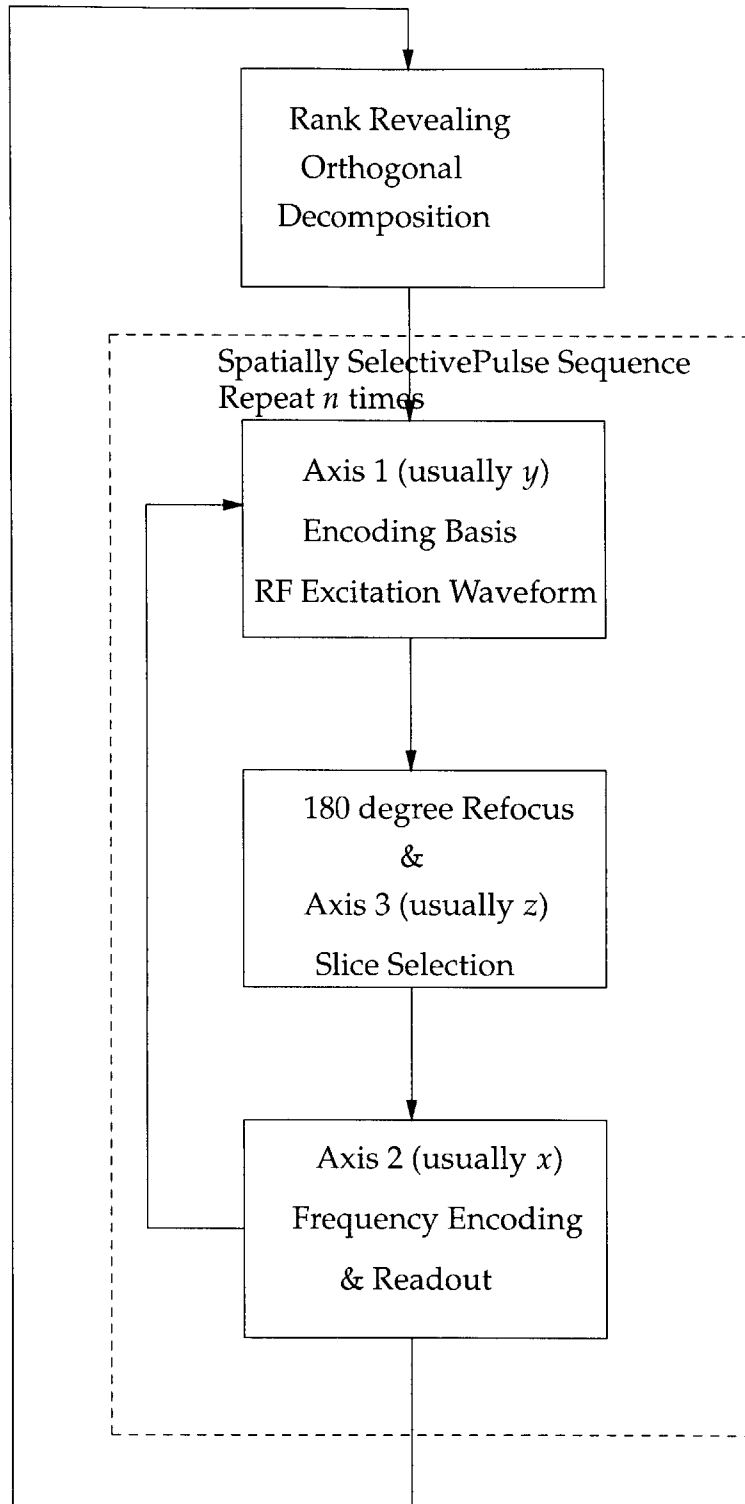
Note that by definition, the columns of  $\mathbf{A}$  form an orthonormal basis for  $\mathbb{C}^n$ . This follows from the fact that  $\mathbf{A}^* \mathbf{A} = \mathbf{I}$ .

Given two such unitary matrices,  $\mathbf{A}_C$  and  $\mathbf{A}_R$ , we can define an invertible transform for a matrix  $\mathbf{F}$ :[2]

$$\mathcal{F} = \mathbf{A}_C \mathbf{F} \mathbf{A}_R^* \Leftrightarrow \quad (3.23)$$

$$\mathbf{F} = \mathbf{A}_C^{-1} \mathcal{F} (\mathbf{A}_R^{-1})^* \quad (3.24)$$

$$\mathbf{F} = \mathbf{A}_C^* \mathcal{F} \mathbf{A}_R \quad (3.25)$$



**Figure 3-6:** The steps in acquiring a spatially-selective Spin-Echo RROD encoded image.

The above transformation can be written as a sum of rank-one matrices, resulting from the outer product of columns and rows of  $\mathbf{A}_C$  and  $\mathbf{A}_R$ . Let  $\mathbf{a}_C(i)$  be the  $i$ th column of  $\mathbf{A}_C$  and  $\mathbf{a}_R(j)$  be the  $j$ th column of  $\mathbf{A}_R$ . Then, the forward transform in equation (3.23) can be rewritten as:[2]

$$\mathcal{F} = \sum_{i=1}^n \sum_{j=1}^n \mathbf{F}_{(i,j)} \mathbf{a}_C(i) \mathbf{a}_R^*(j). \quad (3.26)$$

In the above equation, if  $\mathcal{F}$  is rank-deficient, e.g.,  $\text{rank}(\mathcal{F}) = m < n$ , then the transformation need only range over a basis for the  $m$ -dimensional space, since only  $m$  rank-one matrices are need be summed to exactly produce  $\mathcal{F}$ . The first obvious condition on the transform then is that it must be *rank-revealing*. Without this condition we would not be able to truncate the basis set during MR excitation and acquisition.

Now, consider the case wherein  $\mathcal{F}$  is not rank-deficient. When this is the case, equation (3.26) by itself is not useful in *reducing* the dimensionality of  $\mathcal{F}$ . If however the transform also reveals, apart from the rank, some *ordering* of the rank-one matrices, in the composition of the sum, so that some criterion may be used to limit the dimensionality by dropping the dimensions which are least important, then the usefulness of the transform becomes immediately apparent.

Importance is a very general notion here. In many computational fields, such as machine learning, the dimensions least important are those which exhibit the least amount of variance, one of the reasons *Principal Components Analysis* (PCA) has been so successfully employed.[3] In MRI, one such measure is the mean spatial frequency distribution of the  $\mathbf{a}_C(i)$  and  $\mathbf{a}_R(j)$  vectors which provides a useful ordering.[56]

In the field of MRI, based on the work of Panych et al [31] it was first realized by Zientara et al [55] that such alternate representations of the MRI signal matrix  $\mathbf{S}$ , may be used to increase the efficiency of MR signal acquisition.

### 3.5 *Optimal Spatially-Selective Excitation by SVD*

The full Singular Value Decomposition of a complex matrix  $\mathbf{S} \in \mathbb{C}^{m \times n}$  is defined as the decomposition:

$$\mathbf{S} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^*. \quad (3.27)$$

This decomposition is a *unique* (upto a magnitude  $\pm 1$  complex vector permutation) rank-revealing orthogonal decomposition wherein  $\mathbf{\Sigma}$  is a diagonal matrix of

non-negative real values, the singular values, arranged in decreasing order, and  $\mathbf{U}, \mathbf{V}$  contain the left and right singular vectors, orthonormal basis sets spanning the column and row space of  $\mathbf{S}$  respectively. The SVD is related to the Eigen-decomposition since:[47]

$$\mathbf{\Sigma} = \mathbf{\Lambda}^{1/2}, \quad (3.28)$$

where  $\mathbf{\Lambda}$  is the diagonal matrix containing the eigenvalues of  $\mathbf{S}$ . The number of non-zero singular values in  $\mathbf{\Sigma}$  is the rank  $m$  of  $\mathbf{S}$ .

Since  $\mathbf{\Sigma}$  is diagonal, equation (3.27) can be rewritten as a sum of rank-one matrices, in the spirit of (3.26)

$$\mathbf{S} = \sum_{i=1}^m \sigma_i \mathbf{u}_i \mathbf{v}_i^*, \quad (3.29)$$

where  $\sigma_i = \mathbf{\Sigma}_{(i,i)}$ .

The most important property of the SVD in the MRI setting is that given some  $k \leq m \leq n$ , the rank  $k$  reconstruction of the matrix  $\mathbf{S}$  produced by the first  $k$  singular vectors is the optimal, in a least-square error sense, reconstruction of  $\mathbf{S}$ . That is, the  $k$ -rank reconstruction captures as much of the energy of  $\mathbf{S}$  as possible, in a  $F$ -norm (Frobenius-norm<sup>8</sup>, but also holds for the 2-norm) definition of energy:[47] let the rank- $k$  reconstruction of  $\mathbf{S}_k = \sum_{i=1}^k \sigma_i \mathbf{u}_i \mathbf{v}_i^*$ . Then, we have that

$$\|\mathbf{S} - \mathbf{S}_k\|_F \leq \|\mathbf{S} - \mathbf{B}\|_F, \forall \mathbf{B} \in \mathbb{C}^{n \times n} \text{ s.t. } \text{rank}(\mathbf{B}) \leq k. \quad (3.30)$$

Given the SVD of an initial estimate of the contents of the FOV,  $\mathbf{S}$  and based on equation (3.29) we can see how to use the column space to acquire an image from the MR system [55]: if the row vectors  $\mathbf{u}_i^*$ ,  $1 \leq k \leq m$  are used as RF pulses, the response to the  $i$ 'th pulse will be

$$\mathbf{y}_i = \mathbf{u}_i^* \mathbf{S} \quad (3.31)$$

$$= \mathbf{u}_i^* \sum_{j=1}^n \mathbf{\Sigma}_{(j,j)} \mathbf{u}_j \mathbf{v}_j^* \quad (3.32)$$

$$= \sigma_i \mathbf{v}_i^*, \quad (3.33)$$

since  $\mathbf{u}_i^* \mathbf{u}_i = 1$ , and  $\mathbf{u}_i^* \mathbf{u}_j = 0, \forall j \neq i$ , by the orthonormality of the left singular vectors. The process can be reversed by multiplication with  $\mathbf{u}_i$  on the left to obtain the rank one image after the acquisition of  $\mathbf{y}_i$ .

---

<sup>8</sup>The Frobenius norm of a  $n \times m$  matrix  $\mathbf{A}$  is defined as  $\|\mathbf{A}\|_F = \left( \sum_{i=1}^n \sum_{j=1}^m A_{(i,j)}^2 \right)^{1/2}$ .



Clearly, at the beginning of an MR experiment, the estimate  $\mathbf{S}$  of the contents of the FOV is not available to compute the encoding basis set, however, in a fast dynamic MR imaging setting, the last acquired image can be used to near-optimally encode the next acquisition,[57] since changes in the FOV are expected to be very small.

Let us now walk through a dynamic MR experiment, in order to clarify the encoding method presented in this section.

1. The experiment begins by acquiring an  $n \times n$  Fourier-encoded MR image in the time domain,  $\mathbf{F}_0$ , e.g., by using the identity matrix as pulses in a spatially-selective Spin-Echo pulse sequence. This set-up acquisition step requires  $n \times TR$  scanning time.
2. The time-domain image matrix  $\mathbf{F}_0$  is then decomposed using the SVD, and its column-space  $\mathbf{U}$  is computed. The MR technician running the experiment selects an appropriate for the situation, number of singular vectors  $k$ . The  $k \times n$  matrix  $\mathbf{U}_{exc}^*$  is composed by taking the hermitian conjugate of the first  $k$  columns of the column-space matrix  $\mathbf{U}$ .
3. The rows of  $\mathbf{U}_{exc}^*$  are then used to excite the sample, and acquire the  $k \times n$  new response matrix  $\mathbf{F}_{acq} = \mathbf{U}_{exc}^* \mathbf{F}_0$ . This acquisition step requires  $k \times TR$  scanning time.
4. The new, time domain, representation of the new  $n \times n$  MR image is then reconstructed as:

$$\mathbf{F}_1 = \mathbf{U}_{exc}^* \mathbf{F}_{acq}. \quad (3.34)$$

5. Given the new time-domain image matrix  $\mathbf{F}_1$ , the process is repeated by going back to step 2.

The above process also explains the presence of the word “adaptive” at the heading of the previous section: the basis set used for acquiring an image at imaging step  $t_i$  is computed from the contents (or alternatively some processed contents) of the image at step  $t_{i-1}$ .

Quite literally, what will be acquired in the future is merely a projection onto the basis of the Fourier Transform of the current image estimate. A possible caveat here is that in a changing FOV it is possible for something to occur in the near future which has a zero projection onto the current encoding basis (e.g., is orthogonal to all the  $k$  basis vectors used). In a dynamic series, this is avoided by a two-tier method. First, the encoding basis is computed from the contents of the

current FOV, which are, unfortunately, noisy. Due to partially encoding the noise, it is very unlikely that changes in the FOV will have an exactly zero projection onto any of the  $k$  basis vectors, or that  $k$ -dimensional space. The second level of resistance to this occurrence is that, again, in a dynamic MR experiment, we can always re-acquire a full Fourier-encoded image every few minutes (usually every one hundred images in our near real-time system.)

Although using the SVD as above produces the optimal low-rank reconstruction, it is unattractive in a real-time setting due to its high computational complexity. Moreover, no efficient parallel algorithms for the SVD are currently known (i.e. higher operation counts than the serial algorithm are required in the parallelized algorithms). One way to avoid this problem is to use the same set of basis vectors for more than one acquisitions. Alternatively other near-optimal encoding methods such as fast Krylov subspace methods can be employed.

### 3.6 Excitation by Krylov Subspace Methods

One class of numerically stable Krylov subspace iterative techniques determines an approximation of the symmetric decomposition:

$$\mathbf{S}_{sym} = \mathbf{U}\mathbf{R}\mathbf{U}^* \quad (3.35)$$

The  $l$ -dimensional Krylov subspace  $\mathcal{K}(\mathbf{S}_{sym}, \mathbf{q}, l)$  of a symmetric matrix  $\mathbf{S}_{sym}$  is spanned by vectors formed by repeated multiplication of an initial random basis vector  $\mathbf{q}$  by the matrix:

$$\mathcal{K}(\mathbf{S}_{sym}, \mathbf{q}, l) = \text{span}(\mathbf{q}, \mathbf{S}_{sym}\mathbf{q}, \mathbf{S}_{sym}^2\mathbf{q}, \dots, \mathbf{S}_{sym}^{(l-1)}\mathbf{q}) \quad (3.36)$$

$$\equiv \{\mathbf{q}_0, \mathbf{q}_1, \dots, \mathbf{q}_{l-1}\} \quad (3.37)$$

where  $\mathbf{q}$  must have a finite projection onto the column space of  $\mathbf{S}_{sym}$ .

The orthonormal basis vectors  $\mathbf{q}_i$  do not reveal the rank of  $\mathbf{S}_{sym}$  and moreover there is no immediately apparent ordering that may be used to limit the dimensionality of  $\mathbf{S}_{sym}$ . There is however a relation in-between the so-called *Ritz* values and vectors, that is, the eigenvalues and eigenvectors of the projection of  $\mathbf{S}_{sym}$  onto  $\mathcal{K}(\mathbf{S}_{sym}, \mathbf{q}, l)$ , and the extremal eigenvalues and eigenvectors of  $\mathbf{S}_{sym}$  respectively. The Krylov subspace approximately spans the column space of  $\mathbf{S}_{sym}$ .

As far as an ordering is concerned, it has been realized [56, 58] that the mean spatial frequency of the vectors can be used, since phenomenologically, the eigenvalues in the spatial and Fourier domains are also related by the Fourier transform.[58]

The last interesting points to note about Krylov subspace methods is that first, the MR image matrix  $\mathbf{S}$  is not symmetric, however  $\mathbf{S}\mathbf{S}^*$  is. Using this matrix, the

vectors will approximately span the desired column space by using the decomposition

$$\mathbf{S}_{sym} = \mathbf{S}\mathbf{S}^* = \mathbf{U}\mathbf{T}_{tri}\mathbf{U}^*, \quad (3.38)$$

where  $\mathbf{T}_{tri}$  is a tridiagonal matrix.[17]

Computing the Krylov subspace is a numerically unstable problem, for example, due to the large values involved in exponentiating a matrix. This is overcome by iterative procedures, e.g., the Lanczos iteration, [17, 47] wherein the recurrence for the  $i$ th basis vector becomes:

$$\mathbf{S}_{sym}\mathbf{q}_i = \beta_{i-1}\mathbf{q}_{i-1} + \alpha_i\mathbf{q}_i + \beta_i\mathbf{q}_{i+1}, \quad (3.39)$$

where the scalars  $\alpha_i, \beta_i$  are the diagonal and off-diagonal elements of the tridiagonal matrix  $\mathbf{T}_{tri}$ , e.g.,[17, 47]

$$\mathbf{T}_{tri} = \begin{bmatrix} \alpha_1 & \beta_1 & & & & & \\ \beta_1 & \alpha_2 & \beta_2 & & & & \\ & & \ddots & \ddots & \ddots & & \\ & & & & & & \beta_{n-1} \\ & & & & \beta_{n-1} & & \alpha_n \end{bmatrix}. \quad (3.40)$$

Excitation and reconstruction are defined similar to that of the SVD encoding method. The advantage of this method is the intrinsic parallelizability of the algorithm, rendering an excellent candidate for real-time processing. The disadvantage is that it is not a rank revealing decomposition, and hence only an approximate ordering and truncation can be postulated.

### 3.6.1 *Aside: Excitation by row-space encoding*

So far it has been convenient to use the column spaces as the spatially-selective excitations to be input in the MR system. There is nothing intrinsically necessary that this be the case, except for the fact that it is most convenient since the rows of the acquired matrix line up with the  $x$  axis of standard scanners, and the phase encoding direction is then usually the  $y$  axis, with the  $z$  axis defined as the direction of the  $B_0$  magnetic field. So-called “vertical” encoding, e.g., using the column space of the FOV contents is the easiest, since the direction of the frequency encoding remains the  $x$  axis, and the spatially-selective RF excitation is played out during a gradient in the  $y$  axis, as would normally be done for a non spatially-selective pulse sequence.

In the two methods we have described above, the row space is equally useful in reconstructing a good approximation. In the case of the SVD, the row space is spanned by the columns of  $\mathbf{V}$  and thus the basis row vectors  $\mathbf{v}_i^*$  would be used to excite the sample. In the Krylov subspace, equivalent to equation (3.38), we would compute the decomposition of  $\mathbf{S}_{sym} = \mathbf{S}^*\mathbf{S}$ .

The only significant difference is that these basis vectors now describe a linear combination of the rows of the MR image matrix, not the columns. Therefore, even though the starting matrix for computing the decompositions would be acquired with  $x$  axis frequency encoding and  $y$  axis spatially-selective excitations, the subsequent excitations by the row space basis vectors would be acquired with a pulse sequence that simply swaps the order of the gradients, such that  $y$  axis becomes the frequency encoding direction and similarly, the  $x$  axis becomes the spatially-selective direction.


Although we have in few occasions demonstrated results from row-space excitations, the complication is that unfortunately, a number of “tricks” take place inside the MR scanners’ systems in order to boost performance. For example, each excitation on the scanner uses an alternating RF magnitude. The effect of this is that the “d.c.” component (e.g., a line of average brightness) after Fourier transformation appears at the edges of the image, rather than the center of the image. We have resolved most of these issues and have successfully acquired “horizontally” (row-space) encoded images, but for the sake of simplicity, we will present results from “vertical” (column-space) encoding.

### 3.7 Excitation by Non-Orthogonal Vector Sets

In some cases, a non-orthogonal set of encodes might be able to yield both efficient encoding, as well as some desirable properties. In such cases, when excitation is performed along the spatially-selective direction with non-orthogonal sets, an inverse operation must be performed in order to obtain the time-domain representation of the excited FOV. The linear systems description,[34]

$$\mathbf{Y} = \mathbf{P}\mathbf{F}, \tag{3.41}$$

where  $\mathbf{Y}$  is the MR system response, and  $\mathbf{F}$  is the time domain representation of the FOV contents, holds for any encoding matrix  $\mathbf{P}$ , regardless of orthogonality. Clearly, an RF matrix does not have to be orthogonal in order for us to be able to use its rows to just *excite* a sample with it. If however  $\mathbf{P}$  is not orthogonal, then the question is how to obtain the system response  $\mathbf{F}$ , the sought-after quantity in the



**Figure 3-7:** The standard Cartesian coordinate system orientation for many MR magnets.

MR experiment, given the result of the excitation  $\mathbf{Y}$ . The “obvious” solution is a Moore-Penrose pseudo-inverse, e.g. the left-side inverse based on the SVD:

$$\mathbf{P}^\dagger = \mathbf{V}\mathbf{\Sigma}^+\mathbf{U}^*, \quad (3.42)$$

where  $\mathbf{\Sigma}^+ = \text{diag}(1/\sigma_0, 1/\sigma_1, \dots, 1/\sigma_n)$ . The first problem stems from the fact that since the rows of  $\mathbf{P}$  are not orthogonal, the matrix is not of full rank, and thus some of its singular values will be 0, rendering the inversion of the diagonal matrix  $\mathbf{\Sigma}$  problematic. In fact, even if the smallest singular values are non-zero, but extremely closed to it, the matrix is said to be ill-conditioned, and reconstruction can yield severe numerical instabilities, rendering the result unusable.

This is because ill-conditioned matrices tend to greatly amplify noise. Unfortunately, noise is exactly one feature of MRI that we exploited above, in the SVD and Krylov subspace methods, in order to reduce the amount of data that needs to be gathered in order to accurately reconstruct  $\mathbf{F}$ . In effect, noise can be modeled in the linear system equation as an extra term:

$$\mathbf{Y} = \mathbf{P}\mathbf{F} + n. \quad (3.43)$$

Our early experiments with non-orthogonal vector excitations showed that the

presence of this noise term in MR imaging renders such inversion processes useless.

One solution to this, using the SVD pseudo-inverse, is to simply not include the contributions of the rank-one matrices whose corresponding singular values are below some tolerance level. In the context of MR excitation and reconstruction, this is inadequate, since the dimensionality of the reconstructed spatial domain is reduced even further than the rank of  $\mathbf{P}$ . Even if the RF matrix is of full rank (e.g. orthogonal) reconstruction via a matrix inversion is still hampered if the matrix has some very small singular values. Instead, an alternative method must be employed.

Equation (3.43), has been at the center of digital signal and image processing, and efficient methods are known to invert it.[2, 24]

We are currently focusing our research on two methods. The first, is an iterative error back-projection technique, in which the current estimate of  $\mathbf{F}^{(k)}$  is considered a “blurred” version of  $\mathbf{F}$ . When  $\mathbf{F}^{(k)}$  is used to simulate the MR signal acquisition process, the difference between the simulated and acquired samples can be used to modify the subsequent estimate. The problem with this method is that in a naive approach, an inverse operation must be performed to obtain the first estimate. Ill-conditioning of the matrix to be inverted is a major source of error in this inversion. Furthermore, another inversion must be performed in each step of the algorithm in order to project the error samples onto the Fourier domain in which the system response is described. A possible solution to ill-conditioning is to can treat singular values less than the SNR level as zero, in all pseudo-inverses.

Finally, it should be noted that although the above method might be adequate for the general non-orthogonal set case, in some cases better solutions can be formed. In particular, consider a case which we find interesting: a set of Gaussians in the spatial domain, equally spaced about the FOV. In this case, a *dual basis* can be constructed and an inverse metric computed, allowing an inverse transform onto a basis composed of Dirac basis vectors on a Cartesian grid.[54]

## SECTION 4

---

# *Pulse Sequence Design And Results*

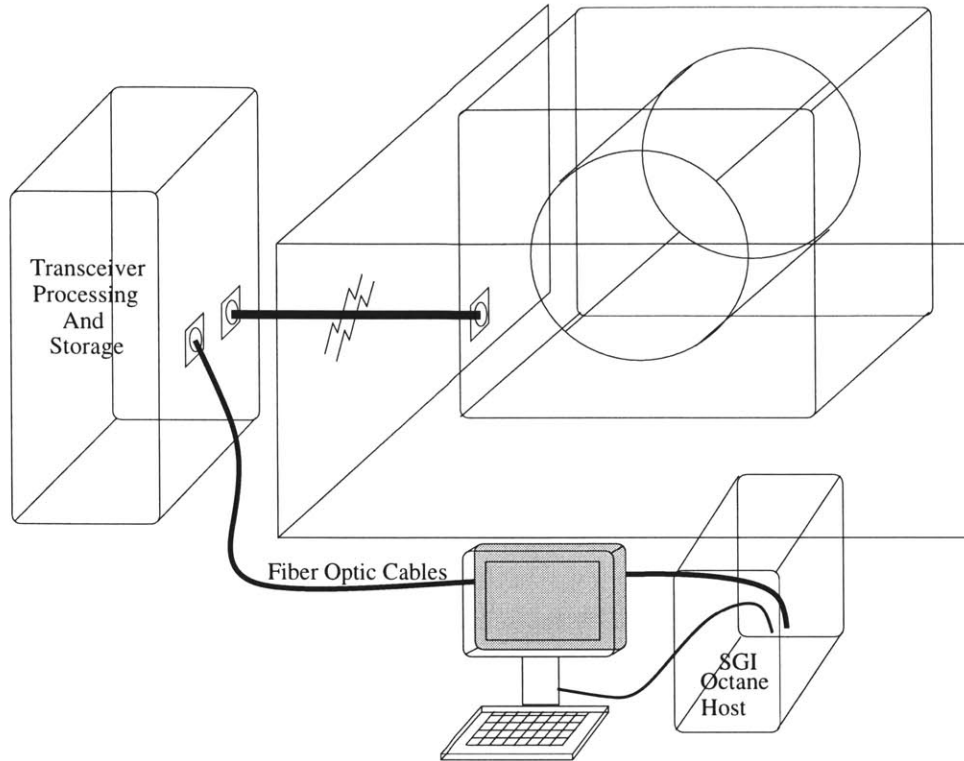
A commercial MRI scanner is composed of a number of complex proprietary parts, with little allowed user interaction and modification, as its primary purpose is to abide by FDA directives for clinical usage. In this section we will look at a specific scanner, manufactured by General Electric Medical Systems Corporation<sup>1</sup>. The GE Horizon LX scanner used in this study, installed at the Brigham And Women's Hospital<sup>2</sup>, employs a 1.5 Tesla magnetic field (achieved through a superconducting magnet), has a 60cm wide and 148cm long, cylindrical, horizontal configuration, magnet bore, with gradient coils capable of producing a magnetic field differential of 4Gauss/cm, with a slew rate of 200 mTesla/m/ms. The software suite loaded on the scanner is the GE Signa Horizon LX 8.2.5 suite. The main parts of this MRI scanner are:

- The main magnet, gradient, RF receiver and transmitter coils and all supporting hardware structure, placed in an RF-shielded "scan" room.
- The *Transceiver Processing And Storage* chassis (TPS), a VME backplane which contains all the necessary computer and electronics hardware to produce the RF excitation pulses and gradient events, acquire the resulting signals and reconstruct the images.
- The *host* computer, an SGI Octane computer, used to control the scanning sessions.

---

<sup>1</sup>GE Medical Systems, Milwaukee, WI 53201.

<sup>2</sup>Brigham And Women's Hospital, 75 Francis Street, Boston, MA 02115



**Figure 4-1:** A typical GE Signa LX MRI scanner setup.

A high-level overview of the system is shown in Figure 4-1.

The Octane workstation, operating on IRIX 6.5, acts as a front end to the MRI scanner. Its main functions are to set up the imaging parameters, load the pulse sequence program to be executed onto the TPS, and transmit the imaging parameters to it, and finally display the resulting reconstructed images after receiving them from the TPS. It also provides disk services in a global database of clinical exams as well as concurrent, FDA-required, magneto-optical storage archiving.

Once a pulse sequence has been loaded on the MR scanner host and all parameters have been set, the object code of the pulse sequence that arranges the RF and gradient pulses as well as sets up the acquisition loop, is dynamically loaded into the TPS chassis. The TPS is a VME backplane containing various systems to generate the actual RF signals, transfer them in analog form to the transmitter, operate the RF transmitter/receiver and gradient coils, and perform image reconstruction. The TPS CPU board is operating on vxWorks 5, a flavor of real-time Unix, which allows one to *dynamically* load and *link* object code into running programs. If the object code provides some predefined “entry-point” symbols, then the pre-running programs can jump into those symbols and execute code from the object file.



In the next section we will describe the computer hardware/software interaction in some small amount of detail and then discuss the implementation issues for a spatially-selective pulse sequence, including some of the caveats involved. In the remaining latter sections we will give results and make comparisons between the methods already described.

## 4.1 *PSD for GE 1.5T Signa LX 8.2.5 Scanner*

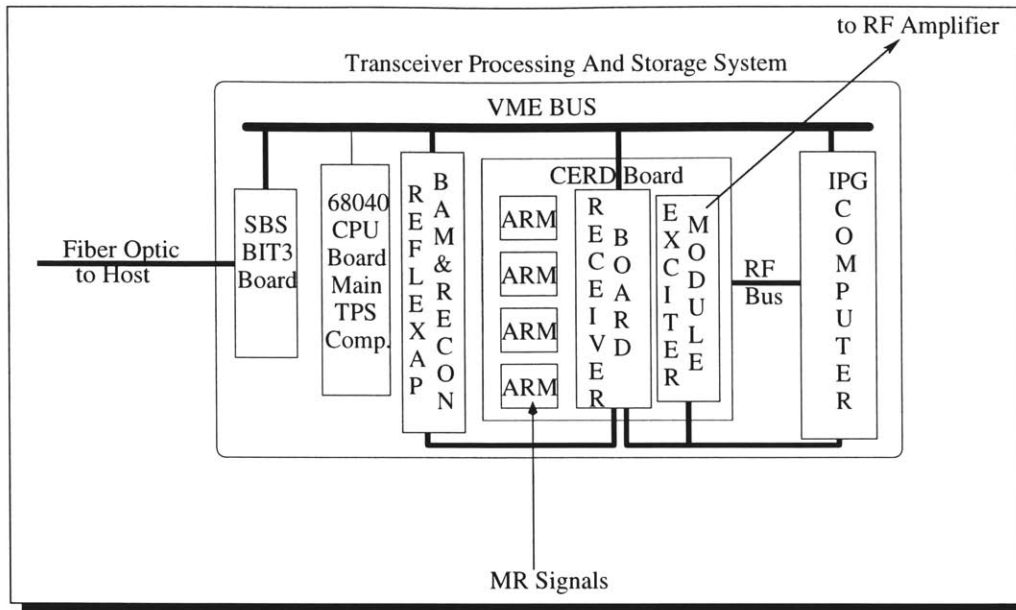
An overview of the TPS system is shown in Figure 4-2. The TPS chassis contains the Integrated Pulse Generator (IPG),[16] a system capable of transforming digital waveforms into analog RF pulse signals. The input to the IPG is a sequence of *short* (2-byte signed) integer pairs, a magnitude and a phase. The IPG provides 256K of SRAM wherein the RF pulse can be stored before the transformation takes place. This limitation is one of the barriers in our studies, since we found that 512 points of resolution for a 5.12ms RF pulse width provides optimal results. However, at this resolution, there is not enough SRAM to hold 256 RF pulses, which would be required for acquiring a “baseline” Fourier-encoded image at the beginning of a scan session. In the next section we will see how this may be overcome by directly tapping onto the VME bus, while knowing the VME base address of this SRAM, and continually updating it – a necessity for the Real-Time system.

The IPG itself is a set of sequencers, which multiply an “instruction” memory with a “waveform” (amplitude) memory to produce the actual outputs.[16] There are separate sequencers for each of the  $x$ ,  $y$  and  $z$  gradients, the magnitude  $\rho_1$  and the phase  $\theta$ .

The instruction memory is simply a “table” composed of entries that are a tuple consisting of:

- a waveform memory pointer (providing the amplitude for the entry),
- a period which when multiplied by the minimum resolution of the IPG ( $2\mu\text{sec}$  for the Signal Horizon LX system) gives the total time to play out the amplitude for,
- an “EOS” (End Of Sequence) bit, designating the end of the sequencer loop.

The waveform memory is a set of 16-bit words wherein the LSB bit is also an EOS bit, in particular used to halt the sequencer until the instruction memory is incremented, after the period has expired. If the EOS bit in a waveform entry is not set, then the waveform memory pointer is incremented and the same instruction from the instruction memory is re-executed using the new amplitude in the



**Figure 4-2:** The Transceiver, Processing And Storage System is a diverse set of VME boards that produce the RF signals, control the gradient coils, digitize the MR signal and reconstruct the images.

incremented waveform pointer entry. Figure 4-3 shows an example of the steps performed by any one of the sequencers given a certain instruction and waveform memory.[16]

Then, our strategy for forcing the sequencers to execute a series of hard pulse amplitudes for an equal amount of time per impulse, as needed for a spatially-selective RF pulse is clear: we form the waveform memory by the values of the magnitude and phase of the RF pulse we wish to play out (one for each of the  $\rho$  and  $\theta$  sequencers), sampled at as many resolution points as we want. Then, we fill the instruction memory with only two entries, the first instruction tuple contains the period for each RF resolution point, which for our purposes is simply  $\Delta t_p$ , the width of a single hard-pulse.<sup>3</sup> The waveform pointer in the first instruction tuple points to the first 16-bit entry of the waveform memory that holds our RF pulse amplitudes. We must take care to make the LSB bit of each waveform amplitude 0, except the last one! Finally, the second instruction memory tuple is useful only in the EOS bit to end the entire sequence.

In this section we will assume a single-image experiment, e.g. a non-dynamic

<sup>3</sup>E.g., for an RF pulse width of 5.12ms and minimum period of  $2\mu\text{sec}$  and 512 RF resolution points, we have  $\Delta t_p = \frac{5.12 \times 10^3}{2 \times 512} = 5$ , also called the "dwell-time".



MRI setting. This is done to avoid unnecessary complications while trying to understand the MR scanner part of the system. We will treat the issues for the dynamic MRI case in Section 5.

The waveforms can be generated however we wish, e.g., using MATLAB<sup>®</sup><sup>4</sup> on a workstation in the same network as the SGI host workstation. The only particular points to pay attention to is that the byte-endian order must be the same as that of the SGI workstation (which, like most RISC processors, is big-endian), and also that an appropriate checksum is appended and the file is prefixed with a vxWorks file descriptor. Then, the waveform file can be transferred via the network to the SGI workstation.

#### **4.1.1 Host and TPS processes**

GE Medical systems provides an environment for writing programs that can be compiled into a host and a TPS program. This environment, called *EPIC*, is a superset of the C programming language. The EPIC pre-processor converts the C source files with intermixed EPIC compiler directives and macros into a set of ANSI C source files that can be compiled into an IRIX “Host PSD Process” and another set that can be compiled into a vxWorks “IPG PSD Process”. A number of runtime shared and static libraries are linked into the object files resulting from compilation of the C sources, in order to produce the executables.

The EPIC pre-processor provides two different types of variable declarations useful to us: *Control Variable Definitions*, or CVs, which are initialized in the host process, and later on are transmitted to the IPG process. It is such variables which allow the MR technician running the MRI experiment to set the *TE* and *TR*, resolutions along the Fourier-encoded axes, the position and orientation of the slice and so on. These variables are extremely useful in that they allow the user to interactively set them via the Graphical User Interface on the host, and can later on be used to control all the aspects of the generation of the pulse sequence that will be executed on the scanner.

The second set of variable declarations we need is the *RSP* variables. These may also be initialized in the host process, but their usefulness lies in the fact that they can be updated in real-time while the IPG process is executing, during the actual play-out of the pulse sequence. These will be necessary in the later section, when we will need real-time *hooks* into the TPS to modify and extract its memory contents in real-time.

---

<sup>4</sup>MATLAB is a registered trademark of The MathWorks Inc., 24 Prime Park Way, Natick, MA 01760.

The EPIC source file needs to define a number of function entry-points that must perform a specific task in order for the experiment to occur. These functions are summarized below.

- CV Initialization, *cvinit()* function. This function is RCP-called by the host system, and can initialize any static variables, or set default values, such as pulse-widths, amplitudes, RF resolutions etc.
- CV evaluation, *cveval()* function. This function is called by the host system, every time a GUI event occurs on the host which changes the value of some CV. For example, when the *TR* or *TE* is modified, this function will get called, and the effect of it should be a re-definition of the appropriate sections of the pulse sequence.
- Predownload, *predownload()* function. This function is again called by the host system after the operator has completed the definition of the CVs. Its purpose is to perform the final preparations for the pulse sequence, such as amplitude scaling.
- Pulse Generation, *pulsegen()* function. This is the most important function for us. It is called on the IPG after all the control variables have been downloaded to it. Its purpose is to make the necessary EPIC macro calls to download the pulse definitions into the sequencer memory so that the pulse can be played-out by the sequencers by starting them. Effectively, for each sequencer the PSD needs to use, such as the RF magnitude and phase as well as gradients, this function must set up the instruction and waveform memories for each one, as was explained above. Note that this is the only function on the TPS side that can still read files. After this function, the ability to read and write files is lost. In our single-image PSD we use this function to read the waveforms from files on the host, whereas for the real-time system we have built and will explain in Section 5, we set up the sequencer memory with some default pulses.
- Real-time entry points, e.g., *scan()* function. The programmer is allowed a total of 20 real-time entry point functions, that will be executed on the IPG computer. The primary purpose of these entry points is to star the pulse sequence on the sequencers, and to modify the sequence as it is running, e.g., in a Fourier-encoded MR experiment, at each *TR*, the magnitude of the phase-encoding gradient must be modified by a small incremental amount. In order to implement spatially-selective pulses, we use the *scan()* function to change

the pointer to the instruction memory that will be played out by the RF magnitude and phase sequencers. An instruction memory and waveform table is set up in *pulsegen()* for each of the RF excitations, and in this entry-point we just shuffle through them, equivalently to modifying the phase encoding gradient.


Figure 4-4 presents the result of executing a spatially-selective Spin-Echo sequence, that the author implemented in EPIC, on the GE Horizon LX 8.2.5 system. All the results appearing in the next section were acquired with this pulse sequence. For each result below, in the single-image setting, the waveforms were generated using MATLAB, and transferred to the host, previous to starting the pulse sequence. As explained above, the *pulsegen()* section reads the waveform files and generates the instruction and waveform memories for the magnitude and phase RF sequencers. The instruction and waveform memories for each of the gradients are generated with standard EPIC macros. The amplitude of the gradients in this sequence will produce a *square* FOV. The resolution of the RF files can be set by a user CV so as to ensure enough waveform memory is available for both  $128^2$  and  $256^2$  imaging.

We are now ready to present and compare the results from various spatially-selective encoding methods. The experiment begins by acquiring a Fourier-encoded image in the low flip-angle regime, using the identity matrix method, as explained in Section 3.3.1. Given this image, we either use SVD or Krylov subspace decomposition, and use the resulting basis set to acquire the raw data, which are then used to reconstruct the image, with a varying number of encodes. In the non-orthogonal set cases, we would of course not need an initial estimate of the FOV contents. Since this last method is work in progress, we do not provide results herein.

The results in the following sections are from a human brain (attached to a person of course...) and a GE provided phantom, composed of plastic and NiCl-doped water. Results from the brain image are only available for some of the methods<sup>5</sup>. We will use two different approaches to comparing the encoding methods. The first, is geared towards the mathematician, and as such should be interpreted with care. This is the percent-difference between the *full* image, e.g. the image computed from an entire basis set, and the image resulting from fewer encodes. The reason why this should be carefully interpreted is because first of all, it does not address any of the important attributes of the image, as far as a medical doctor

---

<sup>5</sup>The author does not have permission to experiment with human subjects. Human cases were acquired by Dr. G.P. Zientara with full, informed consent of volunteer subjects (normal, healthy, adults).



**Figure 4-4:** The above diagrams show an actual spatially-selective Spin-Echo pulse sequence implemented in EPIC, as it executes on the scanner with a 25msec  $TE$  and 400msec  $TR$  (only one  $TR$  shown in the top, and a zoom onto the RF pulse is shown on the bottom). The  $x$ -axis is time in  $\mu\text{sec}$ , and the  $y$ -axis is short integers.

is concerned: Signal-to-Noise-Ratio, Contrast-to-Noise-Ratio, and the size of the minimal structures that can be resolved in the image.

SNR and CNR computations and comparisons are not addressed herein, (those were addressed to a limited extent in [28], the interested reader is referred to [29] for a unified approach,) since our current methods are low flip-angle. Suffice it to say that the SNR is higher for fewer encodes rather than more encodes used. The difference in-between Fourier-encoded and SVD-Encoded MRI in terms of SNR is minimal (and at least for the brain results, both Fourier-encoded and SVD-Encoded have an SNR of about 20.)

The second method of comparison is based on the GE phantom. We will look at the number of encodes required by each method to adequately resolve a structure on the phantom (which is the purpose of the existence of the structure!)

Without further ado, we turn to the results.


## 4.2 Spatially-Selective Fourier Encoding

The results in this section are obtained using the method presented in Section 3.3.1, using the identity matrix to obtain the full Fourier-encoded response matrix. It has been recently suggested [50] that it is not necessary to fully acquire the entire Fourier domain, instead, we can drop some amount of detail in order to produce an image faster. Such *keyhole* methods only acquire the few center lines of  $k$ -space and add high frequency lines acquired once from an earlier “baseline” image. This method is based on the empirical observation that typical MR images tend to have most power in the largest spatial frequencies, and drop off radially from the center frequencies. However, such keyhole methods are much less efficient in capturing information, compared to SVD or Krylov subspace methods.


Figures 4-5 and 4-6 show the reconstructed brain and phantom images for a varying number of phase-encoding acquisitions respectively. The number imprinted on the top left corner of each image thus represents the number of lines starting from the center of the Fourier -domain matrix and extending to both top and bottom, that were acquired. Table 4.2 presents the percent-error for the brain image reconstructions, and Table 4.1 similarly for the phantom.

As can be seen from Figure 4-7, with 64 lines, *half* the Fourier-encoded matrix, it is not possible to resolve the comb-like feature of the GE phantom. At 80 encodes, the spatial resolution is adequate. At 128 encodes, the maximal spatial resolution is obtained. The importance of this small test is that even though 64 encodes capture 89% of the image data, they are still not enough to resolve what looking at those images, one would be interested in.






**Figure 4-5:** Fourier-encoded acquisitions for a GE phantom composed of plastic, and NiCl-doped water. A varying number of phase encodes was acquired from the center of  $k$ -space and extending outward to the top and bottom.



**Figure 4-6:** Fourier-encoded acquisitions lead to the above reconstructed images of a human brain, for a varying number of phase encodes acquired from the center of  $k$ -space and extending outward to the top and bottom.



**Figure 4-7:** Detail from the images presented in Figure 4-5, showing the resolving capability of reduced basis sets. In Fourier imaging, at least 80 encodes are needed to resolve this standard “comb” feature of the GE phantom.

Number of Encodes Used	Percent Error
4	74%
8	62%
12	55%
20	45%
24	40%
28	37%
32	35%
40	30%
48	27%
64	21%
80	15%
128	0%

**Table 4.1:** Fourier-Encoded GE Phantom Percent Error ( $\|S - S_k\|_F / \|S\|_F$ ).

Number of Encodes Used	Percent Error
4	40%
8	25%
12	24%
20	21%
24	20%
28	19%
32	18%
40	16%
48	14%
64	12%
80	10%
128	0%

**Table 4.2:** Fourier-Encoded Human Brain Percent Error ( $\|S - S_k\|_F / \|S\|_F$ ).


### 4.3 *Random Orthogonalized Basis Encoding*

In order to better understand the implications of the linearity result presented above, consider taking a *random* matrix  $A$ . Then, we orthogonalize the matrix, via, e.g., Gram-Schmidt, and use the resulting basis set, to acquire an MR image. This basis set is completely random, and has no connection whatsoever with the contents of the FOV. Figure 4-8 shows the reconstructed image, simply the 2D DFT of the acquired data after being left-multiplied with the inverse of the random matrix.

### 4.4 *SVD Encoding*

We have seen how we can acquire a “baseline” image using either a fixed Fourier basis for encoding, or a random orthogonalized matrix as a basis set. Given the baseline image, we can use the matrix decomposition methods discussed in Section 3, and produce a truncated basis set, which when used to encode an image, theoretically should yield an image which contains more useful information. In this section we present results obtained by the SVD decomposition.

Figures 4-10 and 4-11 show the image reconstructions from the basis sets computed by the SVD of either the brain or the phantom images. The number in the




**Figure 4-8:** A random matrix was orthogonalized and used to acquire this image. Since the rows of the orthogonal matrix forms a basis set for  $\mathbb{C}^{128}$ , the image is fully reconstructed.

top left corner of each image represents the number of columns of  $\mathbf{U}$  starting from the first column that were used. Clearly, the results are rather pleasing, with between 16 and 20 encodes capturing most salient image information. In particular, Figure 4-9 shows that the comb structure in the GE phantom is fairly clear already with just 16 encodes. Compared to the 80 that the Fourier-encoded image required, this is a significant improvement. This was already alluded at in Section 3.4; since low flip-angle MR images have an SNR of about 20–25 (which is the case in the brain images presented herein,) only few of the first singular values encode information significant compared to the noise in the image. This method then presents a potential of  $> 50\%$  reduction in image acquisition time during dynamic imaging.

Table 4.3 summarizes the percent difference in-between the truncated and full images for both cases. Compared to the Fourier-encoded images, the errors are approximately 30% lower. This does not quite scale with the increased spatial frequency resolution ability, which as was mentioned, based on Figure 4-9, SVD appears to require 80% fewer acquired data to produce an equivalent resolving capability. One reason behind this can in fact be seen from the images in Figures 4-11 and 4-10: starting at about 32 encodes and increasing, it is mostly extra noise which gets encoded. In particular, note how there is minimal noise outside the left and right of the image content region (in these images, the spatially-selective excitation direction is left to right, since the  $x$  and  $y$  axes are flipped in the scanner).

We can furthermore claim<sup>6</sup> that it is possible to analyze the frequency distribution of the acquired samples as they are being acquired and use this information to determine whether more encodes would be acquiring mostly noise. In effect, an adaptive rank-reduction algorithm geared towards MRI appears well within

<sup>6</sup>Although we have no preliminary results on this yet.



**Figure 4-9:** Detail from the images presented in Figure 4-10, showing the resolving capability of reduced SVD basis sets; 20 encodes appear enough to resolve the standard “comb” feature of the GE phantom.

reach.

There are a couple more elements that need to be quantified, in convincing our audience as to the usefulness of SVD encoding in MRI.

First, since the RF waveforms are spatially-selective, composed of very short hard-pulses, it is important that the RF penetration of a sample in the FOV leads to uniform RF power deposition. Clearly, this is the case as can be seen in Figure 4-11, since the signal is consistent across the FOV. Secondly, both SVD encoding and Krylov subspace encoding are *adaptive* techniques: the current image is being encoded with a basis computed from a previous image. This is orthogonal to the fact that FOV contents are rarely static. Even in the next section, were we present our near real-time system which has the ability to produce one SVD-encoded image in under one second, movement in the FOV is to be expected in that time scale. What we must show is that encoding an MR image with SVD is *robust*. The experiment we performed involved using the *same* encoding basis as was used for Figure 4-10, but acquiring the image after first *rotating* and *translating* the phantom in the scanner by approximately 25 degrees counterclockwise and about one cm to the left. The results of this experiment are displayed in Figure 4-12. As expected, the result is not nearly as pleasing. The importance of these results however, is that even with 32 encodes, and such an unusually large change in the FOV, the image is still recognizable.

The obvious argument of robustness is that as

$$\lim \frac{\Delta t_{\text{image acq}}}{\Delta t_{\text{FOV motion}}} \rightarrow 0, \quad (4.1)$$

adaptive encoding becomes more efficient. Zientara et al [57] used the distribution of principle angles between the subspaces spanned by an MR images’ column space and white noise column space, as a limiting worst case. It was then shown






Figure 4-10: SVD-encoded acquisitions for the GE phantom. Same as for brain experiment.





**Figure 4-11:** SVD-encoded acquisitions for a human brain. A varying number of column-space singular vectors starting from the first column vector, corresponding to the largest singular value, were used to acquire the echo responses. Reconstruction by multiplication with those singular vectors yielded the image, after 2D DFT.





**Figure 4-12:** SVD-encoded images for the GE phantom using the same basis set computed for Figure 4-10, but using it to acquire a both rotated and translated phantom.

Brain		Phantom	
Encodes Used	Percent Error	Encodes Used	Percent Error
4	23%	4	50%
8	19%	8	41%
12	17%	12	36%
20	15%	20	28%
24	14%	24	25%
28	14%	28	23%
32	13%	32	22%
40	12%	40	20%
48	11%	48	18%
64	10%	64	16%
80	8%	80	13%
128	0%	128	0%


**Table 4.3:** SVD-Encoded Percent Error (computed as  $\|S - S_k\|_F / \|S\|_F$ ). Brain on the left, GE phantom on the right.

that in any case, the subspace spanned by an earlier MR image in a dynamic sequence, had a better projection onto the current FOV contents than did the worst limiting case of the subspace spanned by the columns of white noise.

## 4.5 Krylov Subspace Encoding

The second adaptive encoding method presented in Section 3, was based on Krylov subspaces, in particular the Lanczos method. Apart from Krylov and SVD, a number of RROD methods were proposed and simulated results produced for MR images, by Zientara et al, [58] e.g. SVD, Krylov, rank-revealing QR decompositions and ULV decompositions. Although sub-optimal (since the SVD decomposition is the optimal,) Lanczos encoding produces rather good results, as can be seen in Figure 4-14. Note that for Lanczos encoding, we currently only have phantom data. From Figure 4-13, we can see that approximately 24 to 28 acquisitions are required to differentiate the detail in the phantom. This is expected, over the 16 to 20 acquisitions that were needed by the SVD, since the SVD in the absence of FOV movement is optimal.

Again, RF power deposition is uniform across the FOV. Similarly, changes within



**Figure 4-13:** Detail from the images presented in Figure 4-14, showing the resolving capability of reduced Krylov-subspace basis sets; 28 encodes appear enough to resolve the standard “comb” feature of the GE phantom.

Number of Encodes Used	Percent Error
4	54%
8	43%
12	37%
20	27%
24	23%
28	20%
32	17%
40	12%
48	9%
64	0%

**Table 4.4:** Lanczos-Encoded GE Phantom Percent Error ( $\|S_{64} - S_k\|_F / \|S\|_F$ ).

the FOV, although having an effect on the performance proportional to the amount of change, as in the SVD case, can still produce useful images.




Figure 4-14: Krylov-subspace encoding for the GE phantom.

## SECTION 5

---

# *Real-Time System Design and Implementation Issues*

Our group has recently acquired, on loan, from Mercury Computer Systems<sup>1</sup> a VME-bus card cage, with two MCH6 cards, each containing four PowerPC 603 Compute Elements (CEs) with 16MB of local memory per CE, and a FORCE CPU-5V host Board with a microSparc processor and an accompanying SBUS daughter-card. Each PowerPC CE operates at 100MHz whereas the host processor, which has 32MB of physical RAM, operates at 110MHz, and acts as the bus-master. Extra VME cards can be added in each of two free VME bus slots, ranging in functionality from extra RAM to i860 DSPs, more PowerPCs, or other VME bus adapters such as RINTs. Each of the PowerPC CEs is operating on a proprietary embedded OS (MC/OS), whereas the FORCE board is operating on Sun's Solaris 2.5.

This multi-processor system is a distributed memory system which can appear to the programmer as an SMP.[26] The preferred mode of operation is through high-speed low-latency programmed burst transfers, using the facilities of the embedded OS, and utilizing the special DMA controllers attached to each CE. Broadcasting capabilities are also available. The interconnecting fabric of the system has a peak throughput of 840MB/s.

The implementation of our near real-time MRI system will utilize the MCH6 in order to perform the non-Fourier encoding computations, transferring the RF pulses to the waveform generators on the TPS, performing raw Fourier-domain

---

<sup>1</sup>Mercury Computer Systems, Inc., 199 Riverneck Road, Chelmsford, MA 01824-2820

reconstructions, subsequent image reconstructions and finally communicating the appropriate results to the display subsystems.

First, let us explore some of the problems that must be overcome in order for such a real-time system to operate.

- As was mentioned in Section 4.1.1, there is no provided way of communicating with the TPS after the pulse sequence object code has been downloaded to the TPS sequencers. The TPS begins execution of the *scan()* entry-point, and will not return control until scanning is completed, and the PSD IPG process has exited. Clearly, in a dynamic adaptive MR experiment we *must* have the ability to modify the RF waveforms and possibly gradient events after the initial download of the sequence to the TPS, and *while* the TPS is executing the real-time *scan()* entry-point. We clearly also need to know *where* the sequencers' memory is, in order to, in real-time, modify it, in-between iterations of the *scan()* loop.
- Once the *scan()* section has finished acquiring an image, we need some way to extract it from the TPS, in order to perform the reconstruction and displaying, as well as computing the next set of spatially-selective RF waveforms. We need to both know *where* to find the raw acquired data in the TPS, as well as *when* it is ready for us to extract.
- Synchronization in-between the external computing system and the TPS is necessary. As was just mentioned, we need to know when the acquired digitized data resulting from the adaptive RF waveforms is ready for us to extract. However, we also need to be able to halt the TPS once an iteration of the *scan()* loop is finished, until we have updated the sequencer memory with the new waveforms that were computed from the result of the last acquisition.

Unfortunately, exact answers to most of these problems were not available to us, and as a result we had to reverse-engineer a fair part of the system. In the following section we will provide exact answers to all these problems, in hope that future researchers will be able to quickly progress in such situations.<sup>2</sup>

Once the technical details are clear we will turn towards the general system design we have implemented in order to achieve minimal latency in-between acquisitions and reconstructions.

---

<sup>2</sup>Present-day MR scanners represent a decade of iterative improvements held together by chewing-gum and duct-tape. Due to the complexity of these systems, they are rarely, if ever, completely re-engineered. As a result, even the GE Medical Systems chief engineer gave us some very inaccurate information at times, as to what goes on in the TPS.

## 5.1 *How to Tap Into the MRI Scanner Internals*

We first note that the SGI host performs most, although not all, of the functions we need to perform from the add-on Mercury system: it can upload variables and waveforms to the TPS, which the *pulsegen()* section can use to build the sequencer memories; it also gets a copy of the standard Fourier-reconstructed image from the TPS once the *scan()* section is finished executing, and displays it.

The SGI host is connected via a *Bit3*<sup>3</sup> Fiber Channel adapter directly onto the VME bus backplane in the TPS. Since all the boards are VME cards in the TPS, this leads us to believe that most of the components and memories we need to access on the TPS, have a dedicated window on the VME bus. This does turn out to be the case.

Therefore, our communication means is via a fiber channel, in our case a VME to VME,[42] adapter. This adapter is capable of abstracting communications between two VME busses, in a transparent fashion: once control registers are set up on the local and remote cards, an address reference on the *local* VME bus in a special VME window, becomes a reference on the *remote* VME bus. The result of the reference is transferred by the adapter on an as-needed basis. The peak sustained transfer rate of a Bit3 adapter is 26MB/s, rendering it a perfect approach to our add-on real-time system.

This method provides us with direct access to the memory of the TPS, by mapping it into a fixed address space window on the VME bus of our client system, the MCH6. Of course extreme care must be taken in order to not disturb the VME bus on the TPS side, as our card is an add-on and absolutely no components of the TPS expect it to be there. The next problem to be resolved is how to obtain the addresses we need in the remote VME bus, in order to be able transfer the waveforms into the TPS IPG, and to be able to extract the acquired data from the Reflex Recon board.

First, we will return to the *RSPs* that were mentioned in Section 4.1.1. Suppose we set up an array of raw memory bytes in the RSP section of the PSD. These real-time variables are globally updated whenever they are changed by the PSD. If we have a pointer we want to communicate to the Mercury system from the TPS, we can place it in this RSP buffer, and transfer it over the Bit3 (since the address of the RSPs is a raw VME address, in order for them to be globally visible). Unfortunately, we still need to communicate the base address of this buffer to the Mercury, manually. This is the approach we take in order to perform two-way communications and synchronization variables from the Mercury system to the TPS: a buffer

---

<sup>3</sup>SBS Technologies, Inc., 1284 Corporate Center Drive, St. Paul, MN 55121-1245

is set up in *pulsegen()* on the TPS by our PSD. The buffer is declared in the EPIC RSP section. A pointer to the base of the buffer is also set up as an unsigned integer RSP. Once the user downloads the sequence onto the TPS, and just before beginning the scanning loop, the user can query the RSP pointer by name (from the RSP window on the SGI host GUI). The user then enters the decimal address onto the R/T system we have implemented on the Mercury.

From that point on, all communication is handled automatically via the RSP buffer. The R/T system executing on the Mercury system can halt the scanning loop, obtain pointers to IPG RAM, be informed of the scanning loops' progress (by having the scan-loop update a synchronization variable) etc.

Now that we have established a channel of communication in-between the TPS and the Mercury system, the last step is to obtain pointers to the right memory banks on the right boards of the TPS VME, to transfer waveforms to and extract acquired data from. Although apparently simple, this is a slightly more involved task.

The IPG computer, as was already mentioned, has 256KB of RAM. We need a pointer to a block of it in into which to transfer the waveforms. The *pulsegen()* section of the single-image acquisition PSD, described in the last section, already allocates some memory in order to read in waveforms from binary files from the host. Instead of deallocating this memory once we have produced the instruction memory and transferred it over to the exciter boards, as was done in that PSD, we instead retain the pointer to this block of RAM. The Mercury can request it from the RSP buffer and the *pulsegen()* section provides this pointer. The necessary condition is that the address visible in *pulsegen()* is a raw VME address. This is indeed the case; necessarily so, because the *pulsegen()* function is executing on the TPS 68040 processor, and this RAM is on the IPG embedded computer. Therefore, it must be that the vxWorks operating system can only handle this address as a VME address. We take immediate advantage of this situation.

Whenever the Mercury system computes new adaptive spatially-selective RF waveforms, and the scanning for the last acquisition has completed on the TPS, one of the Mercury CEs is designated to transfer the waveforms via the Bit3 to the address communicated to it from the *pulsegen()* function. The only unnecessary complication on this, is the fact that the proprietary MC/OS on which the Mercury CEs are operating on, does not provide a driver for the Bit3 adapter. While implementing a driver from scratch based on the respective manufacturer's manuals ([42, 27],) we found out through much labor that the MC/OS can not map a VME window smaller than its minimum *virtual memory page size* of 4KB. Unfortunately, the Bit3 adapter provides exactly 256 bytes worth of control registers (which need



to be set up individually for each transfer). Of course, as is the case with most software systems, failure to comply with the undocumented “features” results in no warnings and subsequent non-deterministic system crashes. Thanks to Mercury support<sup>4</sup> we were able to resolve this issue and implemented a (almost) functional device driver for the adapter.<sup>5</sup>

Given synchronization and the ability to reset the waveform memory during the scanning session, the final missing element is the ability to read the acquired data from the TPS memory. This is considerably more convoluted. During the *scan()* entry point, the TPS computer which is executing the function does not have access to the acquired data. Instead, a global variable is provided (called “view”). The scan loop is supposed to increment this variable every time it plays-out a read-out. The receiver board uses that variable to insert the acquired digitized data into the *Balk Access Memory* (BAM) on the Reflex board. Once the data has been inserted into the BAM by the receiver, the Reflex (which is a Mercury board as well) takes over and performs the reconstruction (which is limited to standard Fourier-encoded imaging). The image is then transferred to the SGI host.

The problem arising in this system is that the TPS computer executing the PSD code never has access to the BAM, which contains the acquired data resulting from our spatially-selective RF waveforms. In fact, it appears there are only two ways to get the VME address of the BAM and the offset wherein the acquired data is placed. One is to remotely login to the vxWorks OS on the TPS and execute a special C function call<sup>6</sup> which yields the address. Note that, also unfortunately, the address into the 256MB BAM where the acquired data is placed does vary with time. The other alternative to getting the desired address is by telnet’ing to the TPS computer at port 1071! The daemon running on the TPS returns a character string containing the hex address of the next BAM acquisition. Again, the unfortunate fact here is that the TPS computer is only connected to the SGI host via a reserved IP address.

The solution we chose, was to write a second daemon, on the SGI host which is connected to the Brigham And Women’s Hospital subnet. The Mercury computer, apart from the Sun fast Ethernet 10/100Mbps HME card, which is dedicated for imaging in our system, has a 10Mbps LE card which is connected to the hospital’s network. The R/T system on the Mercury connects via TCP to the daemon on the SGI which in turn connects via TCP to the TPS, and the desired address automatically is passed to the Mercury, by following the inverse path. Given that address,

---

<sup>4</sup>support@mc.com

<sup>5</sup>Some issues still remain unresolved, although we are making progress towards solutions.

<sup>6</sup>The call is `ed_show_all_passes()`.

the system executing on the Mercury can access the acquired data on the BAM via the Bit3 adapter! The final piece of information is the arrangement of the real and imaginary parts of the image in the BAM. Unfortunately, this is not equivalent to the separate raw data files the SGI host can be made to write out to disk. We have found experimentally that the contents of the BAM are organized as follows:

- Starting at the base BAM address, the first 1024 bytes are never used.
- After that, the real and imaginary parts are interleaved. That is, the signal at each point on the sampling grid is saved as a tuple  $\langle \text{real}, \text{imaginary} \rangle$ .
- The BAM always contains a  $256^2$  grid, in row-major order. When 128 points of resolution are selected in the frequency-encoding direction, each row contains useful information only in the first 512 bytes, whereas the remaining 512 bytes are noise. Regardless of the spatially-selective encoding direction, the first row of the acquired data begins at the 1024-byte offset.
- The real and imaginary values of each *odd* row (0-based) are negated. This is actually an effect of inverting the magnitude of every second RF waveform played out, in order to avoid the DC offset.
- The real and imaginary values of each *odd* column (again, 0-based) are negated. This is beyond our understanding.

For comparison, the data written out to disk by the SGI host, when requested, are two separate files, real and imaginary respectively, always  $256^2$ , where the values of each odd row are inverted. Unlike the BAM, in  $265^2$  imaging, the first line of the file does contain useful data, and also, unlike the BAM, when 128 points of frequency-encoding resolution and  $\leq 128$  points of spatially-selective encoding are requested, the first useful row appears at the 65th line of the  $256^2$  matrix.

Having resolved these issues, we will shortly turn to the software system design for the Mercury multicomputer. First, a high-level overview of the entire system is in order. Figure 5-1 depicts the entire system once fully deployed.

As can be seen, the dedicated 100Mbps Ethernet connection on the Sun SBUS is used to transfer the reconstructed image sequence, as it is being generated, to a fast laptop that can display the image, and accept user input that is returned and is used to modify the actions of the MCH6, e.g. the type of decomposition to be used, the number of encodes, or the imaging plane. This networking choice was made as the MCH6 cage needs to be very close ( $< 20\text{ft}$ ) to the TPS, in a machine room, unless repeaters are used for the fiber cables. Moreover, the former was chosen




Figure 5-1: The ensemble of R/T subsystems and their interconnections.

also because it allows extra post processing (such as interpolation and scaling,) to be performed effectively at the client, specifically for each client.

The possible extensions which we are currently implementing, include a wireless network both inside the RF-shielded scan room<sup>7</sup> as well as in its vicinities.<sup>8</sup> Furthermore, provided a fast modem, it is possible to transfer images to clients in remote areas, via a modem dial-up server attached to the system, providing some non-interactive type of tele-medicine.

## 5.2 Computation Timeline

Given the above-mentioned hardware elements and interconnections, we now turn to the task of defining an adequate model of computation. We modeled the system after a simple repetitive pipeline which closely follows the inherently repetitive MR imaging process: as data is acquired, a new encoding basis must be computed which can then be used to excite and acquire the new set of data. The notion of a pipeline also naturally lends itself to stalling the computation while the encoding basis is being computed (although if so desired, the user can define the encoding as an asynchronous operation, limited by a user-controllable maximum number of acquisitions per encoding).

The first image acquired by the system is a full  $n \times n$  baseline image, either Fourier-encoded or resulting from excitation by a *full* basis set for  $\mathbb{C}^n$  e.g., an orthogonalized random matrix.

Given this initial image estimate, we perform the following steps:

1. First, the Fourier-domain data must be reconstructed as defined by the encoding method. For example, if SVD encoding was used to acquire the results (call this  $k \times n$  matrix  $\mathbf{F}'$ ) just transferred from the BAM to the Mercury, then the reconstruction would be performed as in equation (3.34):

$$\mathbf{F} = \mathbf{U}_k \mathbf{F}', \quad (5.1)$$

where the first  $k$  columns of  $\mathbf{U}$  were used for excitation, and  $\mathbf{U}_k$  is the  $n \times k$  matrix containing those columns.

---

<sup>7</sup>Since wireless LANs operate at 900MHz or 2.4GHz, and since most clinical MR imaging is based on hydrogen which at 1.5T has a resonance frequency of  $\sim 62.5$ MHz, there is not expected to be any destructive interference.

<sup>8</sup>This is useful since in MR-guided therapy, such as interventional surgery, a number of people need access to the images. The currently implemented solution is piling in-front of the SGI host computer.

Once this has been completed, we can both begin the computation of the next encoding basis, as well as begin the 2D FFT operation on the Fourier-domain  $\mathbf{F}$  data to yield the spatial domain image. These two processes are independent and proceed without any further synchronization.

2. We allow the possibility to begin the next acquisition immediately, even before the raw reconstruction  $\mathbf{F}$  is completed, if some efficiency in the encoding may be sacrificed for higher temporal resolution: the scanner is advised to complete a next acquisition using the currently available encoding basis. That is, the same waveforms used in the previous iteration are transferred to the TPS and the scanner begins acquisition. In this case, encoding and acquisition are asynchronous.

It is however guaranteed by our system, that the acquisition following the completion of the encoding basis computation will use the most recent basis. Of course, the encoding basis in the asynchronous model has been computed from an estimate preceding the last acquisition, unless, of course, the rate of the encoding computation is greater than or equal to the image acquisition time ( $\#encodes \times TR$ ).

Finally, in order to guarantee robustness of the encoding basis set, as was defined in Section 4.4, we allow a user-controllable limit on the number of images to be generated before a full basis set for  $\mathbb{C}^n$  is used to acquire a new baseline image. The basis set may be the full SVD (since it forms a full basis) or a Fourier basis. We also allow this interleaved baseline image to be acquired with a reduced non-orthogonal set, e.g. Gaussians. Since our goal is to ensure that some finite projection of the salient image information of the FOV is acquired, so long as the standard deviation of the Gaussians is large enough, some weighted amount of every column vector of the  $k$ -space will be acquired. Although the reconstructed image may not be of maximal rank for the number of encodes, the fixed basis can be chosen such that it produces both high and low frequency information.

3. Regardless of how the encoding RF waveforms were generated, (either by a new decomposition or using encodes computed for the last acquisition,) the waveforms are transferred by one of the CEs via the Bit3 to the TPS memory and from there, the PSD we have implemented builds the instruction and waveform memories for the sequencers. Some care needs to be taken since vxWorks employs a *watchdog* in order to avoid processes going into an infinite loop. If the turnaround time from the moment the PSD has finished

acquiring a data set to the moment when the instruction and waveform memories are ready for the next scanning loop, is more than one  $TR$ , the watchdog will kill the PSD. A solution to this is to acquire a nonsense line of data every one  $TR$  while the instruction and waveform memories are being built, or the system is idle (of course, this is not an easy task either, since there are not timing functions available on the TPS computer, and as a result, depending on the size of the RF matrix, we have “hand-tuned” an idle loop).

4. Once the scanner has begun the next scan loop, the CE that transferred the waveforms goes into a spinning loop, waking up every  $1/2TR$ , and requesting via the Bit3 the progress of the PSD. If the scanning loop has completed all  $k$  line acquisitions, the CE transfers the data from the BAM onto local memory, and forwards it to the reconstruction unit, jumping back to step 1.
5. Whenever during steps 1–4 the 2D FFT of  $F$  has been computed, the image is transferred, by the CE that performed the reconstruction, to the FORCE host board on the Mercury cage, and from the host processor there, via Solaris networking, it is transferred to the display clients via the Ethernet TCP connection.

Figure 5-2 provides a schematic of the R/T processes occurring through time as has been implemented.

Synchronization between the CEs is achieved using the only defined primitive on the MC/OS architecture, a semaphore, although we have implemented more complex synchronization primitives (such as readers-writers locks and producer consumer locks) on-top of the semaphore abstraction. Transfers in-between the CEs take place with as few lock acquisitions as possible. We feel that some systems are more critical than others, and is therefore more appropriate to guarantee coherency only between these. As an example of this, if one client display is too slow, new image data will keep getting transferred to it, even if it is possibly dropping frames, rather than stalling the acquisition process until the client has consumed all of the data.

The few cases when locks are needed are, for example, the ability to asynchronously compute new encoding bases. We must retain consistency between which basis was used to acquire which resulting data. Furthermore, if the RF transfer is occurring from the CE to the TPS, it must not be overwritten by the new encoding basis that could have possibly just finished being computed.

In other cases, however, such as reconstructing the images via the 2D FFT, we do not want to stall transferring data from the raw reconstruction unit to the 2D FFT unit, if the latter is running too slow. If we were to stall, the entire pipeline

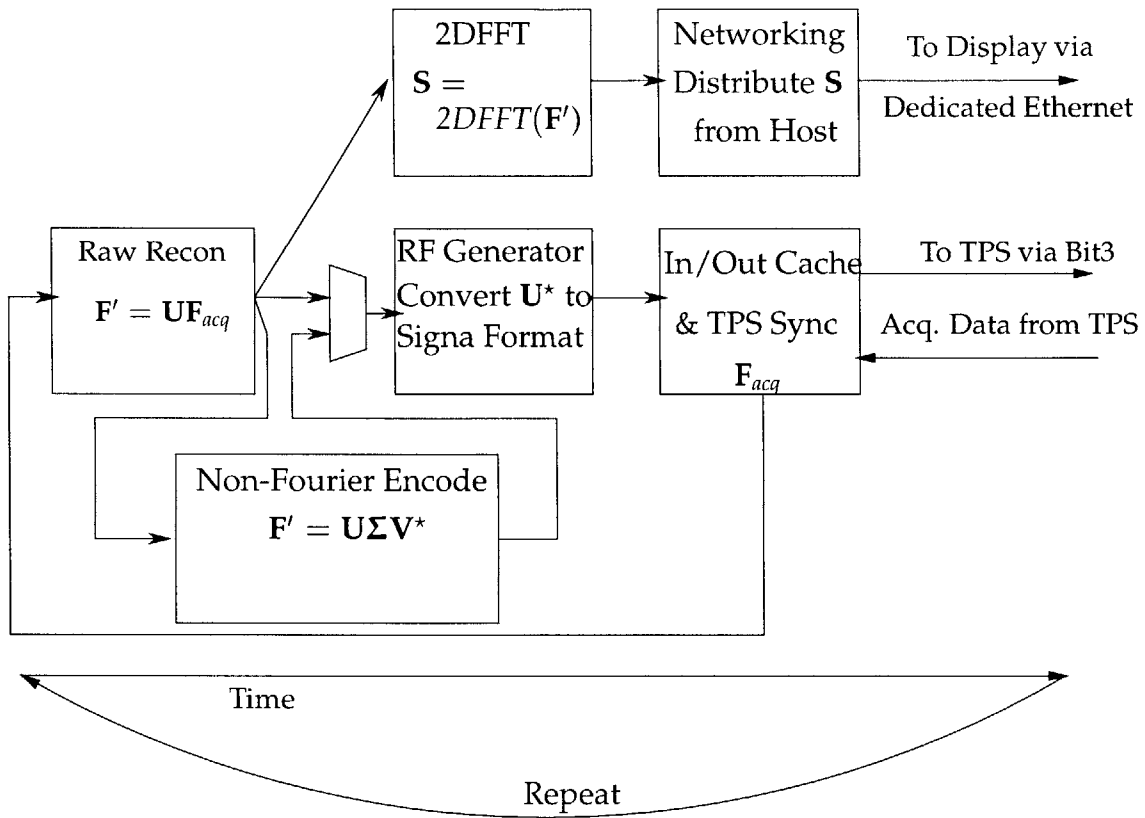


Figure 5-2: The “pipeline” stages of the R/T computation.

would stall, as Figure 5-2 clearly shows. Instead, we immediately make the transfer to the 2D FFT unit. In the worst case scenario, the 2D FFT is occurring at the same time as the new transfer into its local memory is taking place. In order to severely limit this from occurring, in each iteration of the 2D FFT unit, the input buffer is copied in a single memcopy operation.

This in fact is our approach to limiting inconsistencies in the cases where it would be harmful to stall while waiting for a lock: in each iteration of each functional unit, the first instruction sequence is performing a memory copy from the input buffer to a temporary buffer.

### 5.3 *Functional Units and Interconnections*

Some further explanation of this transferring and temporary storing of data is in order. Each of the PowerPC CEs on the Mercury boards is connected to the interconnecting fabric via a DMA controller. The controller provides efficient block-transfers in the fully interconnected fabric (ILK) on each board. Clearly, in order to take full advantage of the underlying system, we chose to arrange processes that need to communicate often, on CEs on the same MCH6 board. Each process is given two buffers: an input and an output one. Once woken up, a thread in each process will read the data in its input buffer, process it and write the result in the output buffer. Then, the thread sets up a DMA transfer (broadcasting capabilities are included and used in the ILK fabrics) from its output buffer to the input buffers on all the consumer processes. Once the transfer is started and completed, the thread simply wakes up the consumers and then goes back to sleep on a semaphore (or other synchronization variable) until it is notified that it has to more data to consume, process and pass on. Of course, the advantage of arranging the in/out buffers in such a manner comes only when the buffers are physically located on the local RAM of the CE (although the ILK is intelligent enough to find a good path if the buffers are located on an adjacent CE – e.g., on the same board).

Figure 5-3 depicts the mapping of the processing tasks onto CEs, as well as the DMA block transfer paths of the data. The example of Krylov subspace encoding was used to label the computations. Perhaps this was not a good choice, since our implementation of the SVD algorithm based on [6].<sup>9</sup> only takes advantage of one of the four encoding processors. It appears that, at least for now, an exact SVD is somewhat slow for real-time MRI. Instead, newer methods, based on Monte-Carlo like averages of randomly sampled sub-matrices [15, 7] hold promise both in terms

---

<sup>9</sup>We modified the version of the Singular Value Decomposition code found in CLAPACK.[1]



of speed, as well as the fact that by not being exact basis sets for the underlying MR image, there is less of a chance of getting “trapped” into some lower dimensional vector space wherein changes in the FOV may have zero projections onto the subspace. In exact SVD calculations, always in the context of MR, it is the noise which allows this robustness. The current system, as was mentioned above, allows the user to control a limit of the number of adaptive acquisitions before a full baseline image is re-acquired, therefore guaranteeing robustness.

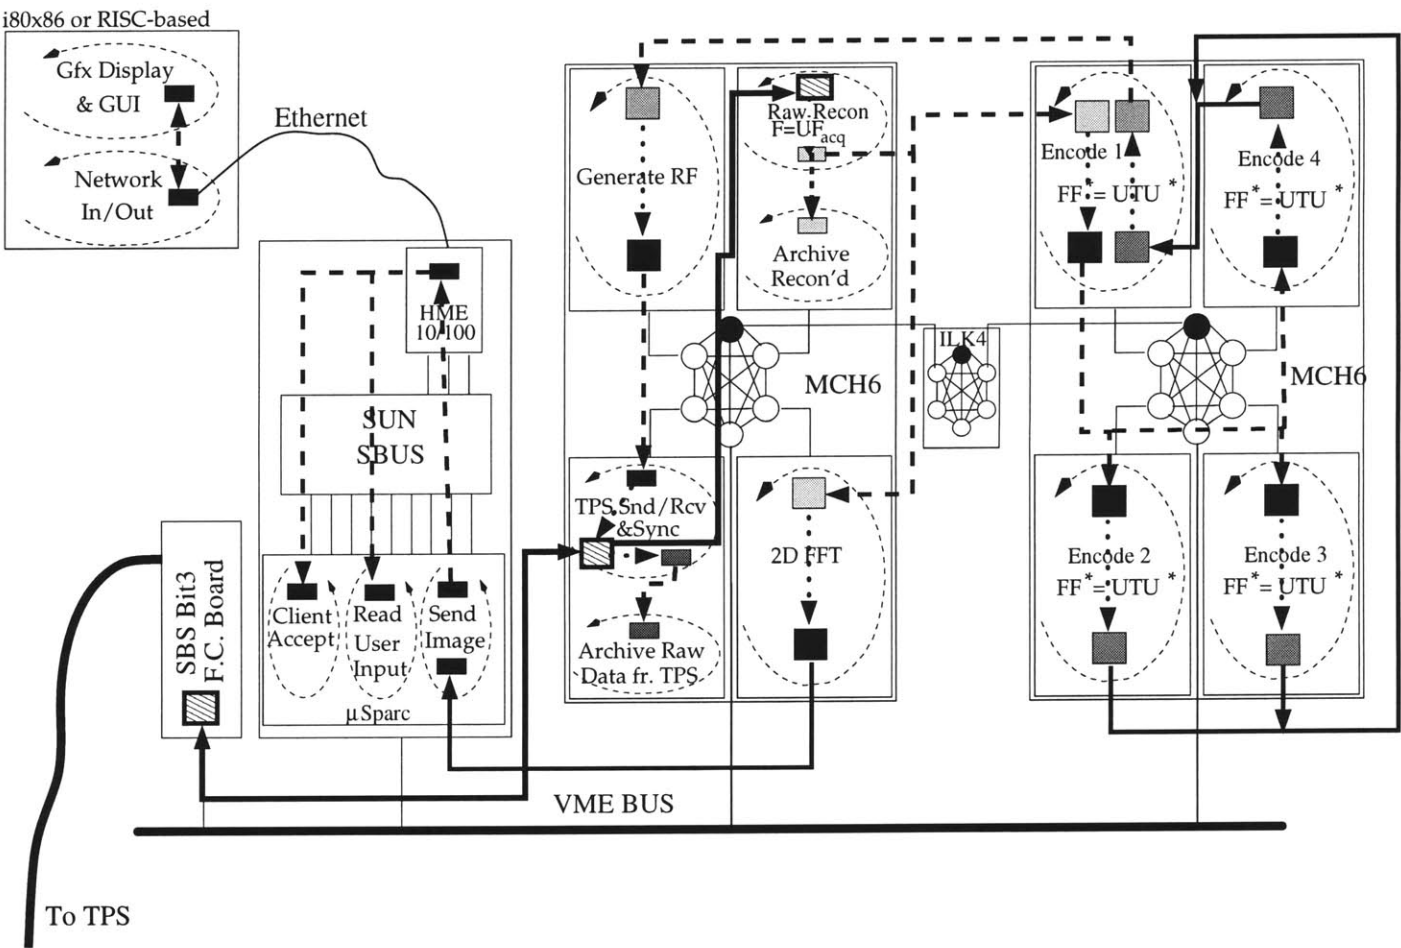
As can be seen from Figure 5-3, each of the steps in Section 5.2 has a direct mapping onto one process. We retained a one-to-one process to processor mapping since multi-threading using POSIX compliant threads was sufficient and more easily maintainable.

### 5.4 Near Real-Time System Performance

The following table summarizes the performance of the functional units under 8, 32 and 128 encodes, for 128<sup>2</sup> SVD imaging. Note that 2D FFT as well as networking and displaying is always 128<sup>2</sup>, thus the times are independent of the number of encodes. The time needed for SVD encoding weakly depends on the rank of the matrix, and is thus dependent on the number of encodes. Values are approximate averages. As of the writing of this thesis, the GE scanner used in this study was upgraded to LX 8.3, for which the EPIC environment and compiler has not been yet made available to us. Unfortunately, even minor revisions of the LX suite are not binary compatible. As a result, we are unable to provide more accurate results, since the PSD we have developed no longer executes on the TPS side.

SVD Encoding				
# Encodes	Raw Recon.	Encoding	2D FFT	Networking & Displaying
10	<25ms	~ 3s	≤ 100ms	< 50ms
32	<50ms	~ 3s		
128	< 200ms	~ 3.7s		

Given the performance of the units, we have been able to produce and display a sustained train of 128<sup>2</sup> Spin-Echo MR images, with a  $TE = 25\text{ms}$ ,  $TR = 70\text{ms}$  and 10–12 SVD encodes, in under one second ( $\sim .83\text{--}.96\text{sec}$ ), at a frame rate of 1.2 to 1 frames per second.



**Figure 5-3:** Data movements in-between the CEs. A filled box denotes a local memory buffer (some are color-coded for ease of understanding), dashed lines with arrowheads denote *non-blocking* data DMA transfers. Solid lines with arrowheads denote *blocking* DMA transfers. Dotted lines denote data copying/processing inside a thread. Threads are denoted by oval dashed arrows.

## SECTION 6

---

# *Conclusions*

MRI systems have profoundly impacted medical diagnosis and therapy in the last twenty years with the application of Fourier-encoded imaging in commercial MR scanners in hospitals and imaging center sites. Nevertheless, this thesis proves that further improvements are possible, without radical redesign of complex legacy code or machinery. Non-Fourier based approaches can, in a number of circumstances, naturally and easily replace Fourier encoding, not only on paper, but in practice, particularly for dynamic imaging.

In this thesis we have described the theory and implementation of a near real-time system for 2D MR imaging. An order of magnitude speedup is obtained through increased MR signal acquisition efficiency, stemming from our use of non-Fourier spatial encoding based on linear-algebraic techniques. Our fully functional system is currently part of a commercial clinical GE MR scanner installed at the Brigham And Women's Hospital.

Our near real-time system provides *proof-of-concept* for our methods. In the process of this implementation we have reverse-engineered certain undocumented features of the MR scanner. Furthermore, the system we have built is composed of an inhomogeneous ensemble of over 20 general-purpose as well as embedded processors, split among 4 physically separated machines, and executing on 5 different operating systems.

We have shown that the linear-algebraic techniques we used are effective for the following reasons:

- *Rank-deficiency of typical MR images.* Since typical MR images tend to be rank

deficient, the amount of data that needs to be gathered during acquisition can be minimized without loss of image information.

- *Average to low SNR.* Since typical MR images are 1.3%–5% noise, (SNR of 20–80), and since the singular values of typical MR images drop off nearly exponentially, data acquisition reduction can also be achieved by not acquiring data that is below the noise-level significance.
- *Temporal rather than spatial frequency resolution.* In MRI-guided therapy, e.g. surgery, laser and focused ultrasound interventions, it is more important to obtain images faster rather than images with high quality. As a result, further truncation of the basis set may be possible.

Some MR systems have evolved in such a flexible manner that non-Fourier based approaches can become a natural extension to the normal mode of operation. Our add-on system merely requires some software addition to the scanner, and a second Bit3 fiber-channel to be installed in the TPS VME backplane. Our entire system can potentially be deployed with minor modifications on *any* GE MR scanner, by the purchase of the external Mercury system and a laptop, adding only a small fraction of the total cost of these scanners. These additions can represent an upgrade that a MR site's engineer can perform. Thus, we feel that we have reached our goal of providing a useful, general, real-time system for 2D MRI.

## 6.1 Further Research

*Pulse Sequences.* Although our system uses a Spin-Echo pulse sequence, which requires a large  $TR$ , of the order of the relaxation time of the sample, any other pulse sequence can be hot-swapped into the system, so long as it is adapted for spatially-selective excitation.[32, 33] Combining the shorter imaging times of faster pulse sequences such as Echo-Planar sequences, and the near 10-fold reduction of the required number of repetition intervals to acquire enough data, sub-100ms high resolution non-Fourier encoded MR imaging is for the first time possible. It is notable that the Mercury multi-processor system which was used for this study, has become outdated, therefore a newer system should provide a significant boost in processing power. Newer advanced hardware undoubtedly will sustain imaging at higher temporal resolution, given the limits of tissue relaxation properties. We expect to have such a system deployed in the very near future.

*3D Imaging.* 2D imaging does not in fact take full advantage of spatially-selective RF excitations. Our methods can be extended to 3D imaging, if instead of a 1D

spatially-selective excitation, a 2D spatial excitation is applied, and the third dimension is fully encoded via frequency-encoding. In this case, reconstruction entails not a sum of rank-one matrices, but a sum of rank-one tensors, and excitation is no longer a linear combination of lines of  $k$ -space along the single encoding dimension, but instead a linear combination of two-dimensional planes. Over phase encoded 3D imaging, RROD encoding efficiency scales *quadratically*, and may potentially lead to approximately two orders of magnitude increased efficiency! Although considerably more complicated, a 3D real-time system based on our current 2D system is feasible.

*Higher spatial resolution.* Higher temporal resolution is not the only desired goal, instead, our research is also geared towards extremely high spatial resolution. Today, it is standard practice to image a  $24\text{cm}^2$  FOV with a surprisingly low spatial resolution of  $\sim 1 \times 1 \times 5\text{mm}^3$ . Imaging time constrained by the relaxation properties of the tissue is not the only constraint on the spatial resolution. Also SNR drop-off becomes dramatic at higher resolutions (inverse-square), since the observable nuclei per voxel decreases as the inverse square of the in-plane voxel dimension.

One approach to obtaining higher spatial resolution at no loss of SNR is to employ error-backprojection techniques, related to what was described in the non-orthogonal vector set reconstruction in Section 3.7. Based on our ability to produce a train of images very fast, it is possible to dynamically sample each voxel at varying, sub-pixel scale, alignments, e.g. by slightly translating and rotating the coordinate system used to acquire each image. It is possible then to reproduce the underlying contents at a much denser grid, based on the voxel information contained in all the sub-pixel shifted images in the dynamic series. Our preliminary experiments, have yielded  $2\text{K} \times 2\text{K}$  images (as opposed to the *maximum possible*  $512^2$  in clinical scanners) over the same FOV, with higher SNR! The most important feature is that since thermal noise is white, the “time-integration” of a train of images as described herein, naturally reduces noise as  $\sqrt{\# \text{ images}}$ .

Present-day Fourier-based MRI technology can be metaphorically described as analogous to a wave of light: it only knows how to travel on a *straight line*. What a photon *perceives* as a straight line, *can*, however, be bent by gravity. What we have presented, is in turn the analogous of a system which allows one to *manipulate the geodesics*.

# Bibliography

- [1] ANDERSON, E., BAI, Z., BISCHOF, C., BLACKFORD, S., DEMMEL, J., DONGARRA, J., DU CROZ, J., GREENBAUM, A., HAMMARLING, S., MCKENNEY, A., AND SORENSEN, D. *LAPACK Users' Guide*, third ed. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1999.
- [2] ANDREWS, H., AND PRATT, W. Image transforms. In *Computer Techniques in Image Processing*, H. Andrews, Ed. Academic Press, 1970.
- [3] BISHOP, C. *Neural Networks for Pattern Recognition*. Oxford University Press, 1995.
- [4] BLOCH, F. Nuclear induction. *Phys. Rev.* 70 (1946), 460.
- [5] BRONSKILL, M., AND GRAHAM, S. NMR characteristics of tissue. In *The Physics of MRI 1992 AAPM Summer School Proceedings* (Alberta, Canada, 1992), P. Sprawls and M. Bronskill, Eds., American Institute of Physics, pp. 32–55.
- [6] DEMMEL, J., AND KAHAN, W. Computing small singular values of bidiagonal matrices with guaranteed high relative accuracy. *SIAM J. Sci. Statist. Comput.*
- [7] DRINEAS, P., FRIEZE, A., KANNAN, R., VEMPALA, S., AND VINAY, V. Clustering in large graphs and matrices. In *Proceedings of the Symposium on Discrete Algorithms* (1999), SIAM.
- [8] EDELMAN, R. Pulse sequence design in MRI. In *Clinical Magnetic Resonance Imaging*, J. Hesselink, Ed. W.B. Saunders Co., 1990.
- [9] ERNST, R., BODENHAUSEN, G., AND WOKAUN, A. *Principles of Nuclear Magnetic Resonance in One and Two dimensions*. Clarendon, Oxford, 1987.
- [10] FINN, J., AND EDELMAN, R. Pulse sequence design in MRI. In *Clinical Magnetic Resonance Imaging*, R. Edelman, J. Hesselink, and M. Zlatkin, Eds., second ed. W.B. Saunders Co., 1996.

- [11] FOWLES, G., AND CASSIDAY, G. *Analytical Mechanics*, fifth ed. Saunders College Publishing, 1986.
- [12] FRENCH, A. *Newtonian Mechanics*. The M.I.T. Introductory Physics Series. W. W. Norton & Company Inc., 1971.
- [13] FRENCH, A. *Vibrations and Waves*. The M.I.T. Introductory Physics Series. W. W. Norton & Company Inc., 1971.
- [14] FRENCH, A., AND TAYLOR, E. F. *An Introduction to Quantum Physics*. The M.I.T. Introductory Physics Series. W. W. Norton & Company Inc., 1978.
- [15] FRIEZE, A., KANNAN, R., AND VEMPALA, S. Fast monte-carlo algorithms for finding low-rank approximations.
- [16] GE MEDICAL SYSTEMS. *Signa Horizon LX EPIC Software Reference Manual*. General Electric Company, 1997–1998.
- [17] GOLUB, G., AND LOAN, C. V. *Matrix Computations*, third ed. The Johns Hopkins University Press, 1996.
- [18] HAACKE, E., AND TKACH, J. Fast MR imaging: Techniques and clinical applications. *A. J. R.* 155 (1990), 951–954.
- [19] HAHN, E. L. Spin echoes. *Phys. Rev.* 80 (1950), 580–594.
- [20] HINSHAW, W., AND LENT, A. An introduction to NMR imaging: From the bloch equation to the imaging equation. *Proc. IEEE* 71 (1983), 338–350.
- [21] HOULT, D. The solution of the bloch equations the presence of a varying B1 field – an approach to selective pulse analysis. *J. Magn. Reson.* 35 (1979), 69–86.
- [22] KUMAR, A., WELTI, D., AND ERNST, R.
- [23] KYRIAKOS, W., PANYCH, L., ZIENTARA, G., AND JOLESZ, F. A method for non-linear RF pulse design. Unpublished, 1997.
- [24] MAEDA, J., AND MURATA, K. Digital restoration of scintigraphic images by a two-step procedure. *IEEE Trans. Med. Imaging* MI-6 (1987), 320–324.
- [25] MAUDSLEY, A. Multiple line-scanning spin density imaging. *J. Magn. Reson.* 41 (1980), 112–126.
- [26] MERCURY COMPUTER SYSTEMS, INC. *Developer's Guide*. Mercury Computer Systems, 1998.

- [27] MERCURY COMPUTER SYSTEMS, INC. *Tools, Libraries and Devices Manual*. Mercury Computer Systems, 1998.
- [28] MITSOURAS, D., MASURKAR, A., ZHAO, L., PANYCH, L., EDELMAN, A., JOLESZ, F., AND ZIENTARA, G. Implementation of hybrid basis non-fourier spatial encoding for dynamic adaptive MRI. In *Proceedings of the 8th meeting of the Society of Magnetic Resonance* (Boulder, Colorado, Apr. 2000), ISMRM.
- [29] PANYCH, L. Theoretical comparison of fourier and wavelet encoding in magnetic resonance imaging. *IEEE Trans. Med. Imag.* 15 (1996), 141–153.
- [30] PANYCH, L. Personal Communication, Dec. 1999.
- [31] PANYCH, L., JAKAB, P., AND JOLESZ, F. An implementation of wavelet encoded MRI. *J. Magn. Reson. Imag.* 3 (1993), 649–655.
- [32] PANYCH, L., OESTERLE, C., ZIENTARA, G., AND HENNIG, J. Implementation of a fast gradient-echo SVD encoding technique for dynamic imaging. *Magn. Reson. Med.* 35 (1996), 554–562.
- [33] PANYCH, L., AND OSHIO, K. Selection of high-definition 2d virtual profiles with multiple RF pulse excitations along interleaved echo-planar k-space trajectories. *Magn. Reson. Med.* 41 (1999), 224–229.
- [34] PANYCH, L., ZIENTARA, G., SAIVIROONPORN, P., AND JOLESZ, F. A linear systems analysis of methods that employ spatially selective RF excitation to encode MR images. Unpublished, 1996.
- [35] PANYCH, L., ZIENTARA, G., SAIVIROONPORN, P., YOO, S.-S., AND JOLESZ, F. Digital wavelet-encoded MRI: a new wavelet encoding methodology. *J. Magn. Reson. Imag.* 8 (1998), 1135–1144.
- [36] PAULY, J., NISHIMURA, D., AND MACOVSKI, A. A k-space analysis of small-tip-angle excitation. *J. Magn. Reson.* 81 (89), 43–56.
- [37] PAULY, J., NISHIMURA, D., AND MACOVSKI, A. A linear class of large-tip-angle selective excitation pulses. *J. Magn. Reson.* 82 (89), 571–587.
- [38] PLEWES, D., AND BISHOP, J. Spin-echo MR imaging. In *The Physics of MRI 1992 AAPM Summer School Proceedings* (Alberta, Canada, 1992), P. Sprawls and M. Bronskill, Eds., American Institute of Physics, pp. 166–187.
- [39] PRATT, W. *Digital Image Processing*. John Wiley & Sons, 1978.



- [40] PURCELL, E. *Electricity and Magnetism*, vol. 2. McGraw-Hill, 1965.
- [41] REDFIELD, A. On the theory of relaxation processes. *IBM J.* 1 (1957), 19–31.
- [42] SBS TECHNOLOGIES, TECHNICAL PUBLICATIONS DEPARTMENT. *Model 413-1 VMEbus to VMEbus Adapter Hardware Manual*. SBS Technologies, Inc., 1999.
- [43] SLICHTER, C. *Principles of Magnetic Resonance*, third ed. Springer-Verlag, 1990.
- [44] STERN, W., AND GERLACH, O. *Z. Phys.* 8 (1921), 110.
- [45] STETTER, E. Instrumentation. In *Clinical Magnetic Resonance Imaging*, R. Edelman, J. Hesselink, and M. Zlatkin, Eds., second ed. W.B. Saunders Co., 1996.
- [46] SUBRAMANIAN, V., ELEFF, S., REHN, S., AND LEIGH, J. An exact synthesis procedure for frequency selective pulses.
- [47] TREFETHEN, L., AND BAU, III, D. *Numerical Linear Algebra*. SIAM, 1997.
- [48] TWIEG, D. The k-trajectory formulation on the nmr imaging process with applications in analysis and synthesis of imaging methods. *Med. Phys.* 10 (1983), 610–621.
- [49] UHLENBECK, G., AND GOUDSMIT, S. *Nature* 117 (1926), 264.
- [50] VAN VAALS, J., TUITHOF, H., AND DIXON, W. Increased time resolution in dynamic imaging. *J. Magn. Reson.* 2 (1992), 44.
- [51] WEAVER, J., XU, Y., HEALY, D., AND DRISCOLL, J. Wavelet-encoded MR imaging. *Magn. Reson. Med.* 24 (1992), 275–287.
- [52] WEHRLI, F. Fast-scan magnetic resonance: Principles and applications. *Magn. Reson. Quarterly* 6 (1990), 165–236.
- [53] WEISSKOFF, R., AND EDELMAN, R. Basic principles of MRI. In *Clinical Magnetic Resonance Imaging*, R. Edelman, J. Hesselink, and M. Zlatkin, Eds., second ed. W.B. Saunders Co., 1996.
- [54] WESTIN, C. *A Tensor Framework for multidimensional Signal processing*. PhD thesis, Link 1994.
- [55] ZIENTARA, G., PANYCH, L., AND JOLESZ, F. Dynamically adaptive MRI with encoding by singular value decomposition. *Magn. Reson. Med.* 32 (1994), 268–274.

- [56] ZIENTARA, G., PANYCH, L., AND JOLESZ, F. Lanczos spatial encodings for dynamically adaptive MRI. In *Proceedings of the 3rd meeting of the Society of Magnetic Resonance* (Nice, France, Aug. 1995), ISMRM.
- [57] ZIENTARA, G., PANYCH, L., AND JOLESZ, F. Applicability and efficiency of near-optimal spatial encoding for dynamically adaptive MRI. *Magn. Reson. Med.* 39 (1998), 204–213.
- [58] ZIENTARA, G., PANYCH, L., AND JOLESZ, F. Near-optimal spatial encoding for dynamically adaptive MRI: Mathematical principles and computational methods. *Inc. Int. J. Imaging Syst. Technol.* 10 (1999), 151–165.