

OPTIMAL DESIGN OF MR IMAGE

ACQUISITION TECHNIQUES

by

BRIAN MARSHALL DALE

Submitted in partial fulfillment of the requirements

For the degree of Doctor of Philosophy

Dissertation Adviser: Dr. Jeffrey L. Duerk

Department of Biomedical Engineering

CASE WESTERN RESERVE UNIVERSITY

May, 2004

CASE WESTERN RESERVE UNIVERSITY
SCHOOL OF GRADUATE STUDIES

We hereby approve the dissertation of

candidate for the Ph.D. degree *.

(signed) _____

(chair of the committee)

(date) _____

*We also certify that written approval has been obtained for any proprietary material contained therein.

I grant to Case Western Reserve University the right to use this work, irrespective of any copyright, for the University's own purposes without cost to the University or to its students, agents and employees. I further agree that the University may reproduce and provide single copies of the work, in any format other than in or from microforms, to the public for the cost of reproduction.

Brian M. Dale

To Barbara Ann

for giving me love

for giving me encouragement

for giving me three great children

for giving me the six best years of my life

but most of all

for giving me the chance to do the same for her

Table of Contents

TABLE OF CONTENTS	1
LIST OF TABLES	7
LIST OF FIGURES	8
LIST OF ABBREVIATIONS	10
ABSTRACT.....	13
CHAPTER 1. INTRODUCTION	15
1.1 SIGNIFICANCE	15
<i>1.1.1 MRI protocols and pulse-sequences</i>	<i>16</i>
<i>1.1.2 K-space trajectories</i>	<i>17</i>
1.2 OPTIMAL TRAJECTORY DESIGN	18
1.3 DCE-MRI.....	19
1.4 FLUOROSCOPIC GRIDDING-RECONSTRUCTION	20
1.5 fMRI APPLICATIONS	21
1.6 OVERVIEW OF DISSERTATION.....	21
<i>1.6.1 To develop an optimally precise DCE-MRI protocol</i>	<i>22</i>
<i>1.6.2 To develop a fast, software-based gridding-reconstruction method</i>	<i>22</i>
<i>1.6.3 To develop methodologies for optimal k-space trajectory design</i>	<i>22</i>
<i>1.6.4 To explore the impact of trajectory optimization on image quality.....</i>	<i>23</i>
<i>1.6.5 Extras</i>	<i>23</i>
CHAPTER 2. OPTIMIZING THE PRECISION OF DCE-MRI	25
2.1 INTRODUCTION	25

2.2 MATERIALS AND METHODS	29
2.2.1 <i>DCE-MRI Protocol</i>	29
2.2.1.1 Measuring Longitudinal Relaxation	29
2.2.1.2 Dynamic Contrast Enhanced Images	30
2.2.1.3 Evaluation of Contrast Agent Concentration	31
2.2.2 <i>Theory</i>	34
2.2.2.1 Propagation of Errors	34
2.2.2.2 Application to DCE-MRI protocol	34
2.2.2.3 Propagation of errors through model fitting	37
2.2.2.4 Validation through Monte-Carlo simulation	40
2.3 RESULTS	40
2.4 DISCUSSION	48
CHAPTER 3. RAPID GRIDDING-RECONSTRUCTION USING TABLES ...	57
3.1 INTRODUCTION	57
3.2 METHODS	58
3.2.1 <i>Description of a Lookup Table DFT method</i>	59
3.2.2 <i>Description of a Look-up Table Method for Gridding</i>	61
3.2.3 <i>Programming Methods</i>	64
3.2.4 <i>Experimental Methods</i>	66
3.2.4.1 DFT Methods	66
3.2.4.2 Gridding Methods	68
3.3 RESULTS	70
3.4 DISCUSSION	77

3.5 CONCLUSIONS	86
CHAPTER 4. TIME-OPTIMAL K-SPACE TRAJECTORIES	88
4.1 INTRODUCTION	88
4.2 THEORY	89
4.2.1 <i>State-space representation</i>	90
4.2.2 <i>Control regions and admissible controls</i>	90
4.2.3 <i>Pontryagin Minimum Principle</i>	93
4.2.4 <i>Time-optimality for 1D problems</i>	94
4.2.5 <i>Time-optimality for 2D and 3D magnitude slew-rate constraints</i>	98
4.3 METHODS	101
4.3.1 <i>Time Optimal EPI reversal</i>	101
4.3.2 <i>Time Optimal radial EPI reversal</i>	102
4.3.3 <i>Time Optimal read-dephasing, phase-encoding, and slice-refocusing ...</i>	102
4.4 RESULTS	103
4.4.1 <i>Time Optimal EPI reversal</i>	103
4.4.2 <i>Time Optimal radial EPI reversal</i>	105
4.4.3 <i>Time Optimal slice-refocusing and phase-encoding</i>	106
4.5 DISCUSSION	107
4.6 CONCLUSIONS	111
CHAPTER 5. TRAJECTORY DESIGN WITH A MULTI-OBJECTIVE GA	113
5.1 INTRODUCTION	113
5.2 METHODS	118
5.2.1 <i>Trajectory parameterization</i>	118

5.2.2 Objectives and Simulations.....	119
5.2.3 NSGA-II.....	123
5.2.4 Constraints.....	124
5.2.5 Implementation	125
5.2.6 Comparison of Theoretical and Actual Objectives.....	126
5.3 RESULTS	126
5.3.1 Trajectory parameterization	126
5.3.2 Computational Time.....	129
5.3.3 Optimal trajectories.....	130
5.3.4 Constraint activity.....	134
5.3.5 Objective function sensitivities	134
5.3.6 Simulation accuracy.....	135
5.4 DISCUSSION	136
5.4.1 Algorithm performance.....	136
5.4.2 Objectives.....	138
5.4.3 Parameters.....	141
5.4.4 Improvements.....	143
5.4.5 Optimal design.....	143
5.4.6 Conclusions.....	144
CHAPTER 6. SUBJECTIVE IMAGE QUALITY OF GA TRAJECTORIES	146
6.1 INTRODUCTION	146
6.2 METHODS	147
6.2.1 Experimental design.....	147

6.2.2 <i>Statistical analysis</i>	150
6.3 RESULTS	151
6.4 DISCUSSION	155
6.4.1 <i>Confounding effects</i>	155
6.4.2 <i>Objective effects</i>	156
6.4.3 <i>Conclusion</i>	159
CHAPTER 7. OPTIMAL TRUE-FISP PULSE SEQUENCES.....	160
7.1 INTRODUCTION	160
7.2 METHODS	161
7.3 RESULTS	165
7.4 DISCUSSION	171
CHAPTER 8. SUMMARY.....	175
8.1 SIGNIFICANCE	175
8.1.1 <i>Current Results</i>	175
8.1.2 <i>Future Impact</i>	178
8.2 FUTURE DIRECTIONS.....	180
8.2.1 <i>DCE-MRI Precision</i>	181
8.2.2 <i>Genetic Algorithm Optimizations</i>	184
8.2.2.1 <i>Methods</i>	184
8.2.2.2 <i>Results</i>	184
8.2.2.3 <i>Constraints</i>	185
8.2.2.4 <i>Parameters</i>	186
8.2.2.5 <i>Objectives</i>	187

8.3 CONCLUSION.....	188
APPENDIX A. KETY EQUATION DERIVATION.....	190
APPENDIX B. TRAJECTORIES AND GRADIENT WAVEFORMS.....	193
APPENDIX C. LONG AXIS CARDIAC IMAGES.....	197
REFERENCES.....	201

List of Tables

Table 3.1 DFT-table Reconstruction Times.....	72
Table 3.2 Computation Times for Various Trajectories	73
Table 5.1 Comparison of Trajectory Encoding Methods.....	127
Table 6.1 Spearman's Rank-Order Correlations.....	152
Table 6.2 Summary of logistic regression results	154
Table 7.1 Summary of test and standard sequence parameters	169
Table A.1 Summary of variables used in Kety equation	190

List of Figures

Figure 2.1 DCE-MRI Processing Steps	33
Figure 2.2 Propagated Noise Errors in $T1_0$ Measurement	41
Figure 2.3 Propagated Noise Errors in Concentration	43
Figure 2.4 Contrast Agent Concentration Curves	44
Figure 2.5 Propagated Noise Errors in K^{trans}	45
Figure 2.6 Average Propagated K^{trans} Errors by Source	47
Figure 3.1 Structure of the Gridding Table	64
Figure 3.2 Accuracy Comparison of Complex Number Representations	71
Figure 3.3 Mathematical Resolution Phantom with Spiral Sampling	74
Figure 3.4 Mathematical Resolution Phantom Difference Images	75
Figure 3.5 Mathematical Resolution Phantom with Radial Sampling	76
Figure 3.6 Pig Neck Showing Inserted Needle with Radial Sampling	77
Figure 4.1 Three Possible Control Regions	91
Figure 4.2 Time-optimal EPI Reversal Gradient Waveforms	104
Figure 4.3 Time-optimal rEPI Reversal Gradient Waveforms	105
Figure 4.4 Time-optimal Slice, Read, and Phase Waveforms	106
Figure 4.5 Control Region for Magnitude Slew Constraints at Gradient Limit	108
Figure 5.1 Flowchart of Typical Genetic Algorithm	114
Figure 5.2 Multi-Objective Terminology	116
Figure 5.3 Simulation Phantoms	122

Figure 5.4 Trajectory Parameterizations for 6-interleaf Spiral	128
Figure 5.5 Archived Trajectory Population	130
Figure 5.6 Sample GA-designed Trajectories.....	131
Figure 5.7 Comparison of GA-designed optimal and standard trajectories.....	132
Figure 5.8 Comparison of GA-designed optimal and radial trajectories	133
Figure 5.9 Correlation Between Simulated and Experimental Objectives	135
Figure 6.1 Screen shot of subjective experiment GUI.....	148
Figure 6.2 Images showing significant improvement in subjective scores	153
Figure 7.1 Single- and dual-echo true-FISP pulse sequences.....	163
Figure 7.2 Timing parameters.....	164
Figure 7.3 Pareto-optimal pulse sequences.....	166
Figure 7.4 Selected Pareto-optimal pulse sequences	168
Figure 7.5 Axial abdominal images.....	170
Figure 7.6 Axial head images	171
Figure B.1 Trajectory and waveforms for acquisition time group.....	193
Figure B.2 Trajectory and waveforms for aliasing energy group	194
Figure B.3 Trajectory and waveforms for flow-artifact energy group	195
Figure B.4 Trajectory and waveforms for off-resonance artifact energy group	196
Figure C.1 Long axis cardiac images.....	197
Figure C.2 Long axis cardiac images with regional saturation pulses.....	199

List of Abbreviations

1D	one dimensional
2D	two dimensional
3D	three dimensional
4D	four dimensional
ADC	analog to digital converter
ADC	apparent diffusion coefficient
AIF	arterial input function
BW	bandwidth or bandwidth/pixel
CNR	contrast to noise ratio
CT	computed tomography
CV	coefficient of variation
DCE	dynamic contrast enhanced
DFT	discrete Fourier transform
DSCQS	double stimulus continuous quality scale
DTPA	diethylenetriaminepentaacetate
EES	extravascular extracellular space
EPI	echo planar imaging
FA	flip angle
FFT	fast Fourier transform
FISP	fast imaging with steady-state free precession
FLASH	fast low-angle shot
flops	floating-point operations per second

fMRI	functional magnetic resonance imaging
FOV	field of view
FT	Fourier transform
GA	genetic algorithm
Gd	Gadolinium
Hct	hematocrit
HOT	hardware optimized trapezoids
iMRI	interventional magnetic resonance imaging
MAC	multiply accumulate
MR	magnetic resonance
MRI	magnetic resonance imaging
MTT	mean transit time
ND	n dimensional
NSGA	non-dominated sorting genetic algorithm
PC	personal computer
PET	positron emission tomography
pixel	picture element
PMP	Pontryagin minimum principle
PNS	peripheral nerve stimulation
RAM	random access memory
rEPI	radial echo planar imaging
RF	radio frequency
ROC	receiver operator characteristics

ROI	region of interest
SAFE	stimulation approximation by filtering and evaluation
SAR	specific absorption rate
SNR	signal to noise ratio
SPIDER	steady-state projection imaging with dynamic echo-train readout
T1	longitudinal relaxation time
T1 ₀	pre-contrast longitudinal relaxation time
T2	transverse relaxation time
TE	echo time
TR	repetition time
voxel	volume element
WHIRL	winding hybrid interleaved radial lines

Optimal Design of MR Image Acquisition Techniques

Abstract

by

BRIAN MARSHALL DALE

MRI (magnetic resonance imaging) is an important modern imaging modality due, in part, to its variety of contrast mechanisms and control parameters. However, this same wealth of control and contrast mechanisms poses a difficult problem for the design of MRI acquisition strategies. To date, most sequence design has been accomplished using experience and heuristic techniques. Despite the advances made in the field using these strategies, they are inherently limited by the skill of the designer. Many individuals who have a sequence-design need lack the necessary experience, and even experts may develop sub-optimal protocols. In addition, some areas of MRI, such as k-space trajectories, seem to offer advantages, but are still poorly understood even by the experts. In order to overcome these limitations it would be beneficial to use modern optimization algorithms in the design of acquisition techniques. To date, such optimal design has been quite limited in both objectives and parameters, and has led to techniques with limited benefit and applicability. This work overcomes these limitations by investigating a variety of algorithms and applications that should prove important in MRI. First, an optimally precise protocol for DCE-MRI (dynamic contrast enhanced MRI) was

developed. Second, a software-based fluoroscopic gridding-reconstruction technique was developed. This makes a variety of pulse sequences, such as spiral imaging, available for iMRI (interventional magnetic resonance imaging) applications and is a necessary step for the remainder of this work. Third, time-optimal k-space trajectories were developed using the calculus of variations and multi-objective GAs (genetic algorithms) were used to develop optimal trajectories with respect to time, aliasing energy, flow-artifact energy, and off-resonance artifact energy. Fourth, these trajectories were tested for improvements in image quality using both objective image-quality measures for experimental phantom images and subjective image-quality measures for in-vivo images. Finally, the multi-objective algorithms were adapted to optimize images acquired using rectilinear sequences with respect to acquisition speed, resolution, and BW (bandwidth). The accomplishment of these goals resulted in improvements for several MRI acquisition techniques and also resulted in the development of techniques that should have wider and more general utility for obtaining similar improvements for other MRI techniques.

Chapter 1. Introduction

MRI is one of the most important modern medical imaging modalities with approximately 18 million procedures performed in 2001 and an annual growth rate of 15% (Goldburgh 2002). While other modalities may be preferred over MRI for specific applications, no other modality has the same breadth of applications as MRI. Part of this general utility may be due to the variety of physical contrast mechanisms available to MRI (Haacke et al. 2001). With inherent sensitivity to magnetic relaxation parameters, susceptibility, motion, chemical shift, magnetization transfer, and other physical quantities, the possible types of image contrast are nearly innumerable. In addition, MRI pulse sequences are able to control 3 separate gradient coils as well as the transmission and reception of energy through one or more RF (radio frequency) coils at each instant in time. This results in a great deal of flexibility in increasing or reducing the sensitivity of a particular acquisition technique to a particular physical source of contrast or artifact.

1.1 Significance

MRI is currently used in three main areas: diagnostic imaging, interventional imaging, and quantitative imaging. Most commonly, it is used clinically for a variety of diagnostic procedures (Goldburgh 2002). These diagnostic procedures usually require high spatial resolution, SNR (signal to noise ratio), and low artifact levels. Additionally, it is important to have strong contrast between normal and pathological tissues (Hendrick et al. 1987). Available contrast mechanisms include diffusion-weighted imaging for stroke, perfusion imaging for vascular tumors and infarction, T1 (longitudinal relaxation time) and T2 (transverse relaxation time) contrast, flow-sensitive imaging and MR (magnetic resonance) angiography for detecting a variety of vascular pathologies,

malformations and cardiac defects, and others (Hahn et al. 1990; Brady et al. 1991; Aronen et al. 1995; Brittain et al. 1995; Sorensen et al. 1997; Beache et al. 1998; McKenzie et al. 1999).

The use of MRI in interventional procedures is rapidly expanding (Duerk et al. 1996). Although the main magnetic field restricts the use of some equipment, the lack of ionizing radiation adds a measure of safety for patients, physicians, and technicians. In iMRI the SNR, artifact, and contrast requirements are somewhat relaxed, but spatial and temporal resolution are critical.

Quantitative MRI shows significant potential for use in the laboratory (Reeder et al. 1996; Belle et al. 1998; Buxton et al. 1998). For such sequences, the precision and accuracy of the generated measurements are of paramount importance and all other considerations are usually only important insofar as they degrade the quality of the measurement. MRI can be used to quantitatively measure perfusion, blood-flow velocity, myocardial strain, ADC (apparent diffusion coefficient), diffusion-tensor, chemical shift and other important physical quantities.

1.1.1 MRI protocols and pulse-sequences

While the utility of this wide variety of controls and contrast mechanisms is beyond dispute, this same variety poses a great challenge for those involved in the design of MRI acquisition techniques. Most current design techniques rely on the skill and expertise of the designer and a few simple heuristic rules derived from the collective years of experience available in the literature. Many clinicians, technicians, researchers and other personnel may have an imaging need, but lack the ability to wisely select or develop the proper imaging technique.

No single pulse sequence is useful or appropriate for all imaging tasks and therefore it is important to specify the particular goal as completely and accurately as possible in order to choose the pulse sequence that will best meet the specific requirements. Once these imaging goals are clearly specified then the quality of a given image can be evaluated as a function of how well the image satisfies the specific requirements. This intuitive, goal-driven evaluation of image quality has rarely been used; instead image quality is usually specified as a function of less meaningful, but simple to compute, metrics such as the SNR (Prato et al. 1986; Hendrick et al. 1987; Star-Lack 1999). Many of these goal-driven measurements of image quality require some amount of modeling and simulation and it is expected that their usefulness and applicability will improve as the underlying models and simulations improve (Petersson et al. 1993).

Despite this important need in the MR imaging community, there is no such general, goal-driven, MRI protocol and pulse-sequence design procedure available in the literature. This work seeks to develop such a general design procedure using a variety of optimization and evaluation procedures.

1.1.2 K-space trajectories

The raw data in MRI is acquired in k-space, which is the spatial frequency, or Fourier domain. Most imaging techniques utilize a rectilinear traversal through k-space and reconstruct the image via the FFT (fast Fourier transform). However it is also possible to acquire data along any arbitrary, non-rectilinear, trajectory through k-space (O'Sullivan 1985; Jackson et al. 1991; Moriguchi H. et al. 1999; Oesterle et al. 1999; Qian et al. 2002; Rosenfeld 2002; Moriguchi et al. 2003). Some of better-known advantages of such non-rectilinear k-space trajectories are that they have rapid acquisition times and good

flow properties. However, they are also known to typically have high sensitivity to off-resonance blurring and significant amounts of aliasing energy (image errors from violating the Nyquist criterion).

Most current design techniques essentially begin with a trajectory shape that is easy to visualize and realize and then examine the properties of the trajectory. By far, the two most common classes of trajectories are spiral and radial (Ahn et al. 1986; Meyer et al. 1992; Noll et al. 1995; Thedens et al. 1999). These two classes have several desirable properties including rapid acquisitions and benign artifact patterns; however, a general understanding of the impact of k-space trajectories on image quality is quite limited (Glover and Pauly 1992; Nishimura et al. 1995; Liao et al. 1997; Shankaranarayanan et al. 2001). This work seeks to utilize this k-space trajectory design problem as an example of the potential utility for general optimal design procedures in MRI.

1.2 Optimal trajectory design

By using formal optimization techniques, the best pulse sequence and/or trajectory for a given task could be chosen as the one that results in images that optimize the appropriate image-quality metrics. This would not only result in imaging techniques that were superior to current techniques, but would also eventually allow independence between the quality of imaging techniques and the sequence-development skills of the designer. Due to the large volume of MRI procedures performed annually, even small improvements in such general design techniques could have a dramatic impact.

Despite the availability of good techniques for modeling, simulating, and evaluating image quality, there has been little effort in integrating image quality into the pulse-sequence design process (Petersson et al. 1993; Salem K.A. et al. 2002). The attempts

that have been made to date have generally centered on the more simplistic metrics and have been severely restricted in the choice of optimization parameters (Prato et al. 1986; Wang et al. 1987; Star-Lack 1999). Other researchers have optimized one small part of a pulse-sequence without rigorously examining the impact of other aspects (Van Lom et al. 1991; Simonetti et al. 1993). Thus they have usually searched a miniscule fraction of the available pulse sequences for a sequence that optimally accomplishes a single, less meaningful, imaging goal. To overcome these limitations we will develop a more general optimization technique that will allow the use of more important objective functions and more general pulse-sequence parameterizations. This should permit better exploitation of the desirable properties and reduction of the impact of the various associated disadvantages of using k-space trajectories.

1.3 DCE-MRI

DCE-MRI is one of the most important of all quantitative imaging techniques for measuring perfusion with MRI, and has been shown to correlate well with the microsphere technique without requiring excision of the tissue (Lombardi et al. 1999; Luo et al. 2002). The standard Gd (gadolinium) contrast agent-based DCE-MRI studies assume a two-compartment pharmacokinetic model relating the rate of change of the tissue concentration to the transport of the contrast agent across the endothelium between the plasma and the EES (extravascular extracellular space) (Tofts 1997). All quantitative DCE-MRI protocols attempting a complete implementation of the Kety model have some common features (Evelhoch et al. 2000). Specifically, they include the following steps:

- 1) Acquisition of raw data for the purpose of calculating spatial maps of $T1_0$ (pre-contrast longitudinal relaxation time) across tissues of interest.

- 2) A fast T1-weighted acquisition sequence used to acquire a pre-contrast signal baseline with respect to which the relative increase in signal after contrast can be calculated.
- 3) Rapid acquisition of a dynamic series of post contrast images using the same sequence, from which the relative signal change (S_{rel}) due to changing contrast agent concentrations can be obtained, and subsequently the local tissue concentration over time can be calculated.

The precision of T1 maps has been analyzed, but this is only part of the precision of the whole perfusion measurement (Wang et al. 1987). An optimally precise DCE-MRI technique would be a useful benefit to research about vascular tumors, and cerebral and myocardial infarctions. In the clinic it would allow more exact determination and monitoring of appropriate therapies and in the laboratory it would allow greater statistical power, smaller sample sizes, or higher confidence for detecting perfusion effects. This research will extend previous work in order to develop such an optimally precise DCE-MRI technique.

1.4 Fluoroscopic gridding-reconstruction

Obviously, data acquired along a non-rectilinear trajectory cannot be reconstructed with the FFT. Several techniques exist for reconstructing such data; the most commonly used being convolution-based gridding-reconstruction (O'Sullivan 1985; Jackson et al. 1991). Unfortunately, even this technique introduces a large computational burden relative to the FFT. In many situations this computational burden introduces an unacceptable delay. A software-based fluoroscopic gridding-reconstruction technique would allow the use of k-space trajectories in such situations, particularly in iMRI

applications. This should, in turn, benefit patient care and safety in iMRI applications by giving the interventional radiologist more rapid feedback. However, current fluoroscopic gridding-reconstruction techniques rely exclusively on hardware solutions, which may not be available at most sites, are more expensive, more difficult to maintain and distribute, and often require special expertise (Meyer et al. 1992; Eggers and Proska 1999). A software-based method would allow the reconstruction of such data without requiring additional hardware and would thus enable the use of non-rectilinear acquisition techniques at most sites. This work will attempt to develop such a software-based fluoroscopic gridding-reconstruction technique.

1.5 iMRI applications

As previously mentioned, one important application for MRI is the guidance of minimally invasive interventional procedures, usually involving the insertion of a needle or catheter (Duerk et al. 1996). This is a particularly promising area for the use of k-space trajectories because one of their known advantages, high temporal resolution, is one of the most important requirements for iMRI. In addition, it is a good area in which to use optimal design techniques because there is somewhat less information available in the literature for designing iMRI pulse sequences than for designing diagnostic imaging pulse sequences.

1.6 Overview of dissertation

This work was initiated with four specific aims in mind. The organization of the remainder of the dissertation roughly parallels this structure.

1.6.1 To develop an optimally precise DCE-MRI protocol

Chapter 2 details the use of the propagation of errors theory to analyze the sensitivity of DCE-MRI measurements to errors in the source images and obtain an optimally precise DCE-MRI protocol. The results of this analysis and optimization suggest that some of the fundamental assumptions used in designing these protocols, such as the use of a very short TR (repetition time) for increased T1 weighting, may actually lead to sub-optimal protocols (Dale et al. 2003b).

1.6.2 To develop a fast, software-based gridding-reconstruction method

Chapter 3 covers the development of a look-up table method for reducing the computational burden and achieving a software-based fluoroscopic gridding-reconstruction. This method has already found widespread acceptance (Dale et al. 2001). Although it does not represent the application or development of an optimization method, the table-based method was essential for the reconstruction of images acquired with the optimal trajectories developed later.

1.6.3 To develop methodologies for optimal k -space trajectory design

Chapter 4 covers the use of the calculus of variations to design optimal trajectories with respect to minimizing gradient transfer time. Chapter 5 details the use of multi-objective genetic algorithms to design Pareto-optimal trajectories with respect to minimizing acquisition time, flow artifacts, off-resonance blurring, and aliasing energy. Several of the objectives and constraints used in the GA optimizations are based on simulated MRI acquisitions. This is intended to produce results that are more useful than those previously obtained with more simplistic measures of image quality. The multi-objective GA, in particular, is a very powerful and flexible tool for optimal design. It is

conceivable that it will eventually achieve widespread use in the design of MRI acquisition techniques.

1.6.4 To explore the impact of trajectory optimization on image quality

It is understood that, even with perfect simulations of modeled quantities, effects that are not modeled may lead to significant image degradation. In order to explore this possibility, these trajectories were implemented on a 1.5 T Siemens Sonata. The quality of the acquired images was evaluated numerically in phantom studies for the achievement of the stated design objectives as described in Chapter 5. Volunteer images were also acquired and evaluated subjectively by human observers as described in Chapter 6. Both methods of validation confirmed the effectiveness of the objectives and optimizations utilized in the previous aim.

1.6.5 Extras

Although not one of the original aims of this project, Chapter 7 details a direct and important new extension of the multi-objective GA work of chapter 5. Specifically, it describes the application of the same multi-objective GA to the optimization of a true-FISP (fast imaging with steady-state free precession) pulse sequence using a standard rectilinear k-space trajectory. Because of the much more widespread use of rectilinear trajectories, Chapter 7 may represent the more immediately applicable portion of the work for most situations. Both single- and dual-echo true-FISP pulse sequences were developed which are optimal with respect to image quality, acquisition time, and resolution.

A summary of this work is presented in Chapter 8, along with some speculations on the future importance of this work and the directions that may prove most fruitful for

further investigations. The appendix contains a brief derivation of the Kety equation, which is the fundamental equation for quantitative measurements of perfusion such as those obtained using DCE-MRI.

Chapter 2. Optimizing the Precision of DCE-MRI

2.1 Introduction

Any product required by cells for normal function and survival, which they do not produce themselves, is delivered via the blood stream. The rate of delivery of these products (as well as the removal of waste products) is governed by the local tissue perfusion rate and the permeability and surface area of the local capillary network. These properties can vary dramatically between normal tissue types. In the central nervous system the foundation of the blood-brain barrier is a greatly reduced capillary permeability compared to elsewhere in the body. In the liver, both perfusion and permeability are high to facilitate delivery of a wide range of blood-born nutrients, toxins, waste products etc. to the cellular machinery responsible for their metabolism. Variations in perfusion and permeability also occur as a result of pathological conditions. Neurological diseases such as multiple sclerosis are known to cause breakdown of the blood brain barrier (Gadian et al. 1985; Grossman et al. 1986; Tofts and Kermode 1991). Cerebrovascular accident by definition leads to regions of local ischemia in brain parenchyma.

Tumors require proliferating neovasculature for continued growth, with microvessel density and permeability dramatically different from surrounding tissue. Tissue necrosis can also influence these properties. These variations in tissue perfusion and capillary permeability are two of the mechanisms responsible for the changes in contrast between tissue types obtained from intravenous contrast agents in medical imaging. Thus, measuring the pharmacokinetic behavior of contrast agents by repeated imaging has

opened the way to development of techniques for measuring local perfusion and permeability *in vivo*.

DCE-MRI assessment of vascularity is based on the application of models of the tracer kinetics of injected T1-shortening contrast agents. These models relate the exchange of tracer between various fluid compartments of the body over time; they are also applicable to other modalities like PET (positron emission tomography) and CT (computed tomography). They can be used to interpret *in vivo* measurements of signal increase or tracer concentration vs. time in various tissues of interest. The standard Gd contrast agent-based DCE-MRI studies assume a compartmental pharmacokinetic model relating the rate of change of the tissue concentration to the bi-directional transport of the tracer across the endothelium between the plasma and the EES (Tofts 1997). This assumption holds for standard clinically approved Gd-based contrast agents (Yuh 1999) which are lipophobic and do not cross the lipid membranes of erythrocytes or tissue cells.

Tofts, et al. (Tofts et al. 1999) published the results of a consensus effort to standardize the form and terms of this compartmental modeling approach for DCE-MRI. They presented a modified Kety formula (Kety 1951), derived in Appendix A, relating the rate of change of the tissue tracer concentration C_T to the instantaneous difference between the arterial plasma concentration (C_p) and the interstitial fluid concentration within the EES (C_e).

$$\frac{dC_T}{dt} = K^{trans} [C_p - C_e] \quad [2.1]$$

By assuming the tracers do not cross the cellular lipid membrane into the cytoplasm, tissue concentration C_T varies proportionally with C_e , and inversely with the volume

fraction v_e of the EES, according to the relation $C_e = C_T/v_e$. Making this substitution and multiplying through the right hand side by K^{trans} , the equation can be rewritten as

$$\frac{dC_T}{dt} = K^{trans} C_p - k_{ep} C_T \quad [2.2a]$$

or alternately as

$$\frac{dC_T}{dt} = K^{trans} C_p - \frac{1}{v_e} C_T \quad [2.2b]$$

The equation can be modified for use with arterial concentration C_A instead of C_p by knowledge of the Hct (hematocrit) according to $C_A = (1 - \text{Hct})C_p$. The parameters k_{ep} and v_e have unambiguous relations to physiology. The EES volume fraction, v_e , is defined above, while k_{ep} is the reciprocal of the MTT (mean transit time) required for a differential volume of tracer in plasma to cross a unit tissue volume from the arterial inflow side to the venous effluent side. The relation of K^{trans} to physiology depends on the balance between capillary permeability vis-à-vis the contrast agent and blood flow in the tissue of interest (Tofts et al. 1999). If the tracer is freely diffusible, i.e. permeability surface area product PS is high compared to flow F , then interstitial fluid and effluent venous plasma concentration are in equilibrium, and flux across the endothelium is only limited by the volume flow, F , of blood per unit volume of tissue. If the permeability is very low compared to flow, then the difference between incoming arterial and outgoing venous concentration is negligible, and only the permeability-surface area product PS of the endothelium governs tracer flux. In the flow-limited case $K^{trans} = F\rho(1-\text{Hct})$. In the permeability-limited case $K^{trans} = \text{PS}\rho$. In both cases ρ is the tissue density. In the mixed case where neither flow nor permeability can dominate, $K^{trans} = \text{EF}\rho(1-\text{Hct})$, where

$E = 1 - e^{\frac{-PS}{F(1-Hct)}}$ is the fractional reduction in capillary tracer concentration as it crosses a unit volume of tissue (Renkin 1959; Crone 1963). Most clinically approved Gd contrast agents, including Gd-DTPA (Gd-diethylenetriaminepentaacetate), fit the mixed case, with a tendency for the influence of flow to outweigh that of permeability surface area product.

Many techniques and approaches to DCE-MRI have been published. However, a Consensus Recommendation for DCE-MRI (Evelhoch et al. 2000) outlined the common features of quantitative DCE-MRI. In particular, those that attempt a complete implementation of the Kety model include the following steps

- 1) Acquisition of raw data for the purpose of calculating spatial maps of $T1_0$ across tissues of interest.
- 2) A fast T1-weighted acquisition sequence used to acquire a pre-contrast signal baseline with respect to which the relative increase in signal after contrast can be calculated.
- 3) Rapid acquisition of a dynamic series of post contrast images using the same sequence, from which the relative signal change (S_{rel}) due to changing contrast agent concentrations can be obtained, and subsequently the local tissue concentration over time can be calculated.

The parameter of greatest interest in DCE-MRI based on the two-compartment model is K^{trans} since it is the parameter associated with microvascular physiology. For any implementation of DCE-MRI, the accuracy and precision of K^{trans} from two-compartment modeling is a function of the source errors that exist in the raw MRI data. However, the route from raw MRI signals to local estimates of K^{trans} involves a complicated series of

mathematical processing steps. Determining which sources of image error have the greatest impact on the error in K^{trans} by simple inspection of the algorithm employed is difficult. Part of the difficulty lies in the fact that the magnitude of the propagated errors may be a function both of the magnitude of the source variation and other key parameters in the algorithm. Thus, a fixed measurement error from one source image may propagate to become a significantly larger error in K^{trans} than the same measurement error from a different source image.

Knowledge of the relative sensitivity of K^{trans} error to independent sources of image error can reveal to investigators which improvements in imaging techniques will lead to the greatest improvement in precision and which will lead to little or no improvement. This study describes the application of the *propagation of errors* technique to DCE-MRI methodologies. The goal was to evaluate the expected error variance in calculations of local measurements of K^{trans} , and to determine the influence of various experimental and imaging parameters, including $T1_0$, relaxivity α , and overall tissue concentration, on propagated error. Accomplishing this goal will answer questions such as, “If there is a small error in the dynamic images, how much error will be introduced into the estimate of K^{trans} , and conversely, what if that same small error was in the T1 measurement image instead?”

2.2 Materials and Methods

2.2.1 DCE-MRI Protocol

2.2.1.1 Measuring Longitudinal Relaxation

Several techniques for calculation of $T1_0$ exist. For the purpose of this experiment we selected the following two. The first is the dual spin echo technique, in which two T1-

weighted spin-echo images that differ only in their TR are acquired. The ratio, R, between these two signals is given by the following expression derived from the spin-echo signal equation (Haacke et al. 2001)

$$R = \frac{S_1}{S_2} = \frac{k_{SE} \left(1 - e^{-TR1/T1} \right)}{k_{SE} \left(1 - e^{-TR2/T1} \right)} \quad [2.3]$$

where k_{SE} includes the T2 effects (assumed constant given fixed TE [echo time]) as well as spin-density and receiver chain properties. This expression can be solved for T1 (numerically for arbitrary TR's or analytically for rational TR1/TR2). The second technique uses two FLASH (fast low-angle shot) sequences with identical TR/TE and different tip angles (Wang et al. 1987). The two signals can be used to solve for T1₀ by the equation

$$T1_0 = \frac{TR}{\ln \left(\frac{S_{\theta_2} \cos \theta_2 \sin \theta_1 - S_{\theta_1} \cos \theta_1 \sin \theta_2}{S_{\theta_2} \sin \theta_1 - S_{\theta_1} \sin \theta_2} \right)} = f_{T1}(S_{\theta_1}, S_{\theta_2}) \quad [2.4]$$

where S_{θ_1} and S_{θ_2} are the signals obtained using flip angles of θ_1 and θ_2 respectively.

2.2.1.2 Dynamic Contrast Enhanced Images

The dynamic T1-weighted images we chose to examine in this study were from 2D-FLASH (two dimensional) acquisitions since they are commonly used in DCE-MRI. Three precontrast baseline images were averaged while post-contrast dynamic images were repeated every 7.14 seconds.

2.2.1.3 Evaluation of Contrast Agent Concentration

The expression for the signal in a FLASH experiment, accounting for a particular contrast agent concentration C , can be written (Haacke et al. 2001):

$$S_{FLASH} = k_{FLASH} \frac{1 - e^{-TR(\frac{1}{T1_0} + \alpha C)}}{1 - \cos(\theta) e^{-TR(\frac{1}{T1_0} + \alpha C)}} \quad [2.5]$$

where k_{FLASH} includes the $T2^*$ effects as well as spin-density and receiver chain properties, and α is the relaxivity of the contrast agent. From pre- and post contrast images the relative signal change can be calculated as

$$S_{rel} = \frac{S_{FLASH}}{S_{FLASH}|_{C=0}} \quad [2.6]$$

Equations [2.5] and [2.6] can be combined to define an analytical expression for S_{rel} . Under certain conditions the k_{FLASH} terms in numerator and denominator can be canceled, and the resultant expression solved for C . This yields a concentration function f_c in terms of S_{rel} , $T1$, and α .

$$C = \frac{\ln \left(\frac{e^{-2TR/T1} \cos(\theta)(1 - S_{rel}) + e^{-TR/T1}(S_{rel} \cos(\theta) - 1)}{e^{-TR/T1}(\cos(\theta) - S_{rel}) + S_{rel} - 1} \right)}{TR \alpha} = f_c(S_{rel}, T1, \alpha) \quad [2.7]$$

Two assumptions are made in this derivation. First, that the post contrast longitudinal magnetization recovery is monoexponential with time constant $T1$, where $T1$ obeys the relation $1/T1 = 1/T1_0 + \alpha C_T$. In this case C_T is the tissue Gd concentration that would be obtained if the amount of contrast agent in the voxel (volume element) were uniformly distributed rather than restricted to the plasma and EES compartments. For this assumption to hold, the water exchange between compartments with a voxel must be in

the fast-exchange regime where protons affected by the extracellular contrast agent can quickly distribute throughout both the intra- and extra-cellular spaces (Landis et al. 1999).

The second assumption is that, for the range of sequence parameters we examined, the FLASH signal attenuation due to Gd-mediated T2 shortening is negligible over the range of tissue concentrations likely to occur with standard clinical doses of Gd contrast agents (0.0-0.1 mM). Negligible T2 effects allow the cancellation of k_{FLASH} terms from final expression for concentration.

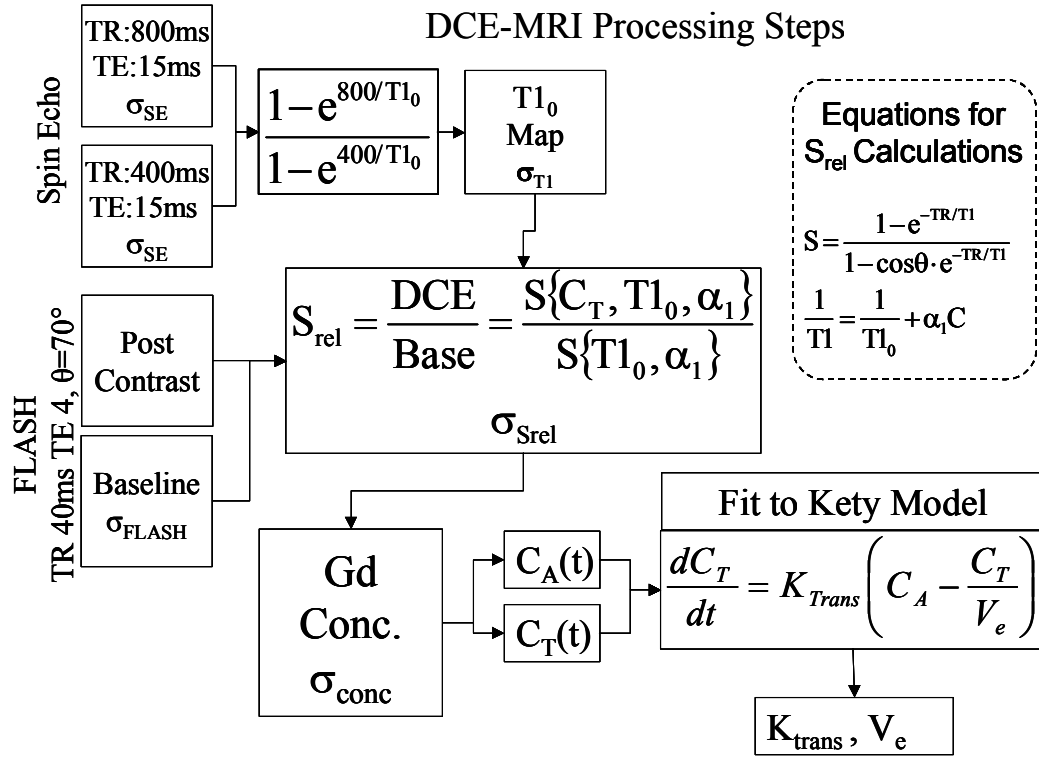


Figure 2.1 DCE-MRI processing steps for a typical protocol. $T1_0$ map is obtained using the dual spin-echo technique and dynamic images use the indicated parameters for a FLASH acquisition. $T1_0$ map, baseline, and dynamic images are combined as indicated to obtain concentration curves which are used to estimate K^{trans} . Each measured pixel has some error which contributes to errors in $T1_0$, S_{rel} , concentration and K^{trans} .

A flow chart outlining the processing steps required for DCE-MRI is given in Figure 2.1. It illustrates how raw images are first converted to T1 maps and relative signal curves, then combined to generate concentration curves. Once the concentration values versus time are calculated for every voxel, the figure shows how curves from tissue and arterial blood are fit to the modified Kety equation to obtain the principal perfusion parameter K^{trans} .

2.2.2 Theory

2.2.2.1 Propagation of Errors

Any measurement process introduces some error into the measured values. If a parameter is not directly measured, but instead is calculated as some function of measured values, then the calculated parameter will contain some error propagated through the function from the errors in the original measured values. The final error in the calculated parameter will depend both on the magnitudes of the various source errors and their modulation by the function. The theory of propagation of errors characterizes this modulation and the resulting errors. In particular, it states that for any parameter, f , calculated from measured values x , y , and z

$$\sigma_f^2 = \sigma_x^2 \left(\frac{\partial f}{\partial x} \right)^2 + \sigma_y^2 \left(\frac{\partial f}{\partial y} \right)^2 + \sigma_z^2 \left(\frac{\partial f}{\partial z} \right)^2 \quad [2.8]$$

where σ_x , σ_y , and σ_z are the respective standard deviations (Bevington 1969). The covariance terms are neglected in this expression, and therefore the errors must be uncorrelated. The propagation of errors does not require that the source errors be normally distributed, and even in the presence of bias, the partial differential terms are the sensitivity of the parameter to each source of error.

2.2.2.2 Application to DCE-MRI protocol

Only the propagation of errors from measured parameters was analyzed. User-defined parameters of the scanning sequence, such as TR and θ , were assumed to be exact quantities; they could not introduce errors themselves, but they could modulate errors introduced elsewhere.

Substituting equation [2.7] into equation [2.8], we obtain an expression for concentration variance in terms of S_{rel} , T1, and α :

$$\sigma_C^2 = \sigma_{S_{rel}}^2 \left(\frac{\partial f_C}{\partial S_{rel}} \right)^2 + \sigma_{T1}^2 \left(\frac{\partial f_C}{\partial T1} \right)^2 + \sigma_\alpha^2 \left(\frac{\partial f_C}{\partial \alpha_1} \right)^2 \quad [2.9a]$$

Practical DCE-MRI implementations typically use a published standard value for α , or at most use a lab standard measured once in a phantom study. Thus it is unlikely that any practical DCE-MRI implementation will have a reasonable estimate of σ_α^2 . For this reason we also neglect its contribution and set $\sigma_\alpha^2 = 0$ at this point to obtain [2.9b].

$$\sigma_C^2 = \sigma_{S_{rel}}^2 \left(\frac{\partial f_C}{\partial S_{rel}} \right)^2 + \sigma_{T1}^2 \left(\frac{\partial f_C}{\partial T1} \right)^2 \quad [2.9b]$$

Even if this is not a good assumption, it may prove to have very little practical effect because the errors in our concentration measurement due to propagated errors in relaxivity will be constant across all studies, patients, and measurements.

The values for T1 and S_{rel} are in turn subject to propagated errors from the source images. For example, the solution of [2.3] for the dual spin-echo experiment with TR1 = 800 ms and TR2 = 400 ms, is:

$$T1 = -\frac{400}{\ln\left(\frac{S_{800}}{S_{400}} - 1\right)} = f_{T1}(S_{800}, S_{400}) \quad [2.10]$$

where T1 is in ms, and S_{800} and S_{400} are the values of the signal intensity for corresponding voxels in the TR = 800 ms and TR = 400 ms images respectively. In a similar manner equation [2.6], the expression for the relative signal in the FLASH experiments, can be rewritten:

$$S_{rel} = \frac{S_t}{S_0} = f_{S_{rel}}(S_0, S_t) \quad [2.11]$$

where S_0 is the averaged baseline image (pre-contrast) and S_i is the post-contrast image for which the concentration is being calculated. By substituting [2.10] (or [2.4]) and [2.11] into [2.8], and using [2.3] and [2.5] to remove the S_i terms it is now possible to write

$$\sigma_{T1}^2 = \sigma_{S800}^2 \left(\frac{\partial f_{T1}}{\partial S_{800}} \right)^2 + \sigma_{S400}^2 \left(\frac{\partial f_{T1}}{\partial S_{400}} \right)^2 = \frac{\sigma_{SE}^2 T1^4 \left(2 + 2e^{-400/T1} + e^{-800/T1} \right)}{k_{SE}^2 160000 \left(e^{-400/T1} - e^{-800/T1} \right)} \quad [2.12a]$$

$$\sigma_{T1}^2 = \sigma_{\theta1}^2 \left(\frac{\partial f_{T1}}{\partial S_{\theta1}} \right)^2 + \sigma_{\theta2}^2 \left(\frac{\partial f_{T1}}{\partial S_{\theta2}} \right)^2 = \frac{\sigma_{\theta}^2 T1^4 \left(e^{-TR/T1} \cos(\theta_1) - 1 \right)^2}{k_{FLASH}^2} \cdot \frac{\left(e^{-TR/T1} \cos(\theta_2) - 1 \right)^2 \left(\csc(\theta_1) \left(e^{-TR/T1} \cos(\theta_1) - 1 \right)^2 + \csc(\theta_2) \left(e^{-TR/T1} \cos(\theta_2) - 1 \right)^2 \right)}{TR^2 e^{-2TR/T1} \left(e^{-TR/T1} - 1 \right)^2 (\cos(\theta_1) - \cos(\theta_2))^2} \quad [2.12b]$$

$$\sigma_{Srel}^2 = \sigma_{S0}^2 \left(\frac{\partial f_{Srel}}{\partial S_0} \right)^2 + \sigma_{St}^2 \left(\frac{\partial f_{Srel}}{\partial S_t} \right)^2 = \frac{\left(1 - e^{-TR/T1} \cos(\theta) \right)^2 \left(\sigma_{St}^2 + \sigma_{S0}^2 \frac{\left(1 - e^{-TR(1/T1 + C\alpha)} \right)^2 \left(1 - e^{-TR/T1} \cos(\theta) \right)^2}{\left(1 - e^{-TR/T1} \right)^2 \left(1 - e^{-TR(1/T1 + C\alpha)} \cos(\theta) \right)^2} \right)}{k_{FLASH}^2 \left(1 - e^{-TR/T1} \right)^2} \quad [2.13]$$

Finally, substituting [2.7], [2.12], and [2.13], into [2.9b] results in an expression for the variance in the concentration value in terms of tissue T1, contrast agent concentration, and SNR of the various source images:

$$\sigma_C^2 = f_{\sigma_C^2} \left(\frac{k_{SE}}{\sigma_{S800}}, \frac{k_{SE}}{\sigma_{S400}}, \frac{k_{FLASH}}{\sigma_{S0}}, \frac{k_{FLASH}}{\sigma_{St}}, T1, C \right) = f_{\sigma_C^2} (SNR_{800}, SNR_{400}, SNR_0, SNR_t, T1, C) \quad [2.14]$$

The full expression of equation [2.14] is omitted at this point for brevity. We now have expressions for σ_{T1}^2 , σ_{Srel}^2 and σ_C^2 in terms of the various imaging parameters, the tissue T1 and the contrast agent concentration.

2.2.2.3 Propagation of errors through model fitting

While T_1 and concentration are obtained by single point evaluations of analytical functions, K^{trans} is obtained by fitting discretely sampled time functions of arterial and tissue concentration to some representation of the Kety model. Several approaches to such model fitting have been employed. Some involve fitting a derived analytical expression for C_T under an assumed AIF (arterial input function). Others deconvolve measured AIF with measured tissue concentration curves to obtain the tissue residue function. Alternately one may numerically differentiate the tissue concentration curve to obtain dC_T/dt and solve the Kety equation directly using linear regression with the tissue concentration and AIF as independent linear predictors of the differential term. Other approaches may be conceived but all will have at least one element in common: the data used in each will be comprised of a time series of n sampled concentrations, each of which has a propagated error component which contributes uniquely to the overall error in the fit parameters. The error contribution of each sample is unique due to the fact that the propagated error in concentration varies with concentration itself (see [2.11]) and the concentration varies for each sample. Tracking the propagation of these contributions through a fitting a process to a final single parameter will now be described.

For directness, assumed the technique of numerical derivation of $dC_T(n)/dt$ from the sampled tissue curve $C_T(n)$ for purpose of direct fit to the Kety equation using linear regression with $C_T(n)$ and $C_A(n)$ as independent predictors of $dC_T(n)/dt$. Once these three curves are obtained, a least-squares fit to the Kety model gives K^{trans} and $1/V_e$. Analysis of the least squares fit shows that K^{trans} can be written:

$$K^{trans} = \frac{\sum(C_A dC_T) \sum(C_T^2) - \sum(C_A C_T) \sum(C_T dC_T)}{\sum(C_A^2) \sum(C_T^2) - (\sum(C_A C_T))^2} \quad [2.15]$$

It is not appropriate to use [2.8] directly on [2.15] because [2.8] is only valid if the errors in the arguments are uncorrelated (Bevington 1969). In this case, all of the concentration measurements depend on the T1 measurement and the errors should therefore be correlated. Instead, by first substituting [2.7], [2.10] (or [2.4]), and [2.11] into [2.15] we obtain an expression for K^{trans} in terms of all of the n+3 arterial samples S_{Ai} , and the n+3 tissue samples S_{Ti} as follows.

$$K^{trans} = f_{K^{trans}}(S_{800A}, S_{400A}, S_{0A}, S_{1A}, \dots, S_{800T}, S_{400T}, S_{0T}, S_{1T}, \dots) \quad [2.16]$$

We then make the assumption that the error in any of these images will be uncorrelated and so [2.8] may be used to obtain an expression for $\sigma_{K^{trans}}^2$.

$$\sigma_{K^{trans}}^2 = \sum_i \left(\sigma_{Si}^2 \left(\frac{\partial f_{K^{trans}}}{\partial S_i} \right)^2 \right) \quad [2.17]$$

As before [2.3] and [2.5] can be used to eliminate all of the S_i terms. The use of an analytical representative C_A curve results in an expression for the variance in K^{trans} in terms of the tissue and blood T1 and the SNR of the various source images.

$$\sigma_{K^{trans}}^2 = f_{\sigma_{K^{trans}}^2} \left(T1_A, \frac{k_{SE}}{\sigma_{S800A}}, \frac{k_{SE}}{\sigma_{S400A}}, \frac{k_{FLASH}}{\sigma_{S0A}}, \frac{k_{FLASH}}{\sigma_{S1A}}, \dots, T1_T, \frac{k_{SE}}{\sigma_{S800T}}, \frac{k_{SE}}{\sigma_{S400T}}, \frac{k_{FLASH}}{\sigma_{S0T}}, \frac{k_{FLASH}}{\sigma_{S1T}}, \dots \right) \quad [2.18]$$

The equations outlined above were used to obtain a closed-form expression for $f_{\sigma_C^2}$ and $f_{\sigma_{K^{trans}}^2}$ using a symbolic math program (Mathematica, Wolfram Research, Champaign Ill). The large numbers of terms, combined with the complexity of the MR signal

equations, prevent the closed form expression for σ_{Ktrans}^2 to fit within the space constraints of this dissertation.

Estimates for k_{FLASH} , σ_{S0}^2 , and σ_{St}^2 were obtained from the FLASH (TR=37ms, TE=4ms, FA [flip angle]=70°, slice thickness = 5mm, matrix=192x256, FOV [field of view]=320x240 mm) images obtained using the DCE-MRI protocol described above. Estimates for k_{SE} , σ_{S800}^2 , σ_{S400}^2 , were obtained from spin-echo sequences for T1 measurement (TR=400/800 ms, TE=15ms, FA=90°, slice thickness = 4mm, gap=1mm, matrix=192x256, FOV=320x240 mm, 1 acquisition). The FLASH sequence was repeated 120 times, once every 7.14 seconds. The FLASH sequence for DCE-MRI employed 1 acquisition, while that for the baseline measurement was averaged over 3 acquisitions. All images were obtained using a 1.5T Siemens (Erlangen, Germany) Symphony MRI system.

These estimates were substituted into the $f_{\sigma_C^2}$ expression, along with the appropriate imaging parameters, to obtain:

$$\sigma_C^2 = f_{\sigma_C^2}(T1, C) \quad [2.19]$$

for the imaging system and protocol under consideration.

Eq [2.20] was chosen as a representative AIF (Simpson et al. 1999). The Kety equation, [2.21] was solved using the AIF to obtain an expression for the tissue concentration as a function of time.

$$C_A(t) = a \frac{t}{t_p^2} e^{-t/t_p} + b \left(1 - e^{-ct/t_p} \right) \quad [2.20]$$

$$\frac{dC_T}{dt} = ef \rho C_A - \frac{ef \rho (1 - hct)}{ve} C_T \quad [2.21]$$

2.2.2.4 Validation through Monte-Carlo simulation

Monte-Carlo simulations were used to estimate variance and confidence limits in the fit of K^{trans} values due to errors propagated from source image noise. Specifically, several K^{trans} , $T1_0$ values, and AIFs were chosen and used to simulate the sampled signals described above. Zero-mean uncorrelated Gaussian errors were introduced into each signal. The resulting simulated noisy data was then used to obtain a least-squares fit to [2.2]. The process was repeated multiple times and the corresponding values for K^{trans} were analyzed to obtain confidence limits for K^{trans} . The Monte-Carlo simulations were used as an independent method to validate the results of the propagation of errors analysis.

Additionally, the propagation of errors results above were used to obtain an optimal DCE-MRI protocol. Using similar Monte-Carlo simulations as above, this protocol was compared to the standard protocol currently in use at this institution. The Levene test for homogeneity of variation was used to test the null hypothesis that the variances of the standard and optimal protocols were the same. The number of Monte-Carlo simulations was determined by statistical power considerations (power is the probability of detecting a true difference - the probability of getting a true positive). Specifically, it was chosen so that a true difference of $\frac{1}{2}$ of the theoretical difference could be detected with probability 0.80.

2.3 Results




Figure 2.2 The coefficient of variation in $T1_0$ as a function of true $T1_0$ for two different measurement techniques. Only errors due to typical levels of noise are considered. 16 dual tip-angle acquisitions can be performed in the same amount of time as one dual spin-echo acquisition.

Figure 2.2 is a plot of the CV (coefficient of variation = standard deviation/mean) of the $T1_0$ measurement as a function of $T1_0$ (see Eq [2.12]) for both measurement techniques. The FLASH technique is 16 times faster than the spin-echo technique and therefore the CV is plotted both for a single acquisition and for 16 acquisitions. The FLASH technique with 4 or more acquisitions was found to be more precise than the spin-echo technique for all $T1_0$. Due to its improved precision and speed, the dual tip-angle FLASH technique with 4 acquisitions was used in the remainder of this work. Although this analysis assumes that TR and tip-angle are exact quantities, it should be noted that tip-angle could be inexact or non-uniform causing a small distortion (increase

or decrease in slope of flat region) in the corresponding curve. Note that the most precise $T1_0$ measurements ($CV_{SE} = 1.6\%$, $CV_{FLASH1} = 2.5\%$, $CV_{FLASH16} = 0.6\%$) are obtained for low $T1_0$ (250 ms for SE, 380 ms for FLASH). The least precise $T1_0$ measurements ($CV_{SE} = 7.8\%$, $CV_{FLASH1} = 3.7\%$, $CV_{FLASH16} = 0.9\%$) occur for high $T1_0$ tissues such as blood ($T1_{blood} = 1200$ ms). The source error variances used in this plot were taken from the noise variance estimates described above.

Propagated Noise Errors in Concentration (C.V.)




Figure 2.3 The coefficient of variation in concentration (σ/μ) as a function of true $T1_0$ and concentration. Only errors due to typical levels of noise are considered. $T1_0$ is measured using the dual tip-angle (7° , 40°) technique with four acquisitions. The FLASH images are measured using $TR = 40$ ms and tip angle = 70° . The baseline is averaged over three FLASH acquisitions.

Figure 2.3 is a contour plot based on equation [2.14]. The contours are plotted along lines of constant CV for the concentration measurement. Note that, when the concentration is approximately 150 μM or greater, we expect that the measured concentration values will have less than a 10% CV over almost all values for $T1_0$ while concentrations less than 75 μM will have more than a 15% CV. For most tissues and

Contrast Agent Concentration Curves




Figure 2.4 Simulated concentration curves used in this study. These three arterial input functions were used during this study. Tissue curves correspond to the tissue concentration from the sharpest of the three input functions with three different perfusion values. These concentration curves were assumed in the evaluation of the precision of K^{trans} . Other functions could be assumed for situations where these do not provide a reasonable approximation to the observed range of concentration curves.

concentrations of interest the CV is nearly independent of $T1_0$. Note also, that the most accurate data (CV = 2.2%) will be obtained with relatively high contrast agent concentrations ($C = 1810 \mu\text{M}$) in low $T1$ tissues ($T1_0 = 310 \text{ ms}$).

Figure 2.4 is a plot of the simulated contrast-agent concentration curves used to evaluate the precision of K^{trans} measurements. The plots are based on equations [2.20] and [2.21] with the parameters chosen to give reasonable approximations to the curves observed in previous studies using the DCE-MRI protocol described above. The highest tissue concentration is $182 \mu\text{M}$ for highly perfused tissues ($K^{\text{trans}} = .0040$). The more typical tissue ($K^{\text{trans}} = .0025$) had a peak concentration of $130 \mu\text{M}$.




Figure 2.5 Coefficient of variation in K^{trans} as a function of $T1_0$ and as a function of TR and Tip Angle for several different values of $T1_0$. Note that the secondary effect of $T1_0$ and Tip Angle is greater than the main effect of $T1_0$ alone and that longer TR measurements are more precise.

The CV of K^{trans} was plotted against tissue $T1_0$ (Figure 2.5a), TR of the dynamic FLASH sequence (Figure 2.5b), and tip angle of the dynamic FLASH sequence (Figure 2.5c). For each plot, all other variables and parameters were held constant, (e.g. acquisition time, $T1_0$ measurement tip angles, etc.). Source image variances used in these computations were obtained from the DCE-MRI protocol images as described above. The variance for the arterial concentration was based on a 100 voxel ROI (region of interest), which improved the arterial SNR by a factor of 10. The CV displayed little

variation as a function of $T1_0$ though there is a minimum of approximately 5.0% for tissues with a $T1_0$ around 500 ms (CV = 5.5% for $T1_0 = 250$ ms). The CV was more sensitive to TR than $T1_0$, with longer TR yielding more precise measurements. Though no minimum was found over the range of TR studied, the CV varied from 10% for TR = 20 ms to 2.6% for TR = 100 ms. The CV dropped 3.3% as TR changed from 20 ms to 30 ms and only dropped 0.2% as TR changed from 90 ms to 100 ms. The CV was also found to be sensitive to the tip-angle of the dynamic FLASH sequence, and the sensitivity function was strongly modulated by changes in the tissue $T1_0$. For tissues with $T1_0 = 250$ ms the CV varied from a maximum of 7.3% at 30° to a minimum of 5.2% at 50° , while tissues with $T1_0 = 500$ ms varied from a maximum of 6.6% at 90° to a minimum of 3.7% at 40° .




Figure 2.6 The coefficient of variation in K^{trans} using the optimally precise protocol is averaged over $250 \text{ ms} < T1_0 < 1250 \text{ ms}$ and $.001 \text{ min}^{-1} < K^{\text{trans}} < .004 \text{ min}^{-1}$. Both errors due to typical levels of noise (reduced by averaging) and errors due to a small artifact (not reduced by averaging) are considered. The small artifact is equal to 1% of the signal from fat. Only one type of error from one source of error is considered for each bar.

Figure 2.6 shows the average error introduced into the assessment of K^{trans} from errors in the T1 maps, baseline images, and dynamic FLASH images. Errors from each source image are separated into arterial and tissue components for a total of six sources of error. Each average CV is generated by evaluating [2.18] for the imaging parameters of the standard protocol currently in use at this institution and averaging the resulting CV over a range of $T1_0$ from 250 ms to 1250 ms and a range of K^{trans} from 0.0010 min^{-1} to 0.0040 min^{-1} . In order to investigate the sensitivity both to noise and to other errors (artifacts) these average K^{trans} errors were generated both from the noise variance

estimates above and from a small artifact deviation equal in magnitude to 1% of the signal from fat. The noise error was reduced appropriately to account for signal averaging and ROI averaging, but the artifact error was not. The largest source of propagated error was an average CV = 7.67% when the small artifact occurred in the tissue portion of the baseline image. In contrast, noise in the baseline image caused only a 3.64% average CV. The smallest source of propagated error was an average CV = 0.04% due to noise in the arterial input portion of the dynamic images. Note also the much greater sensitivity to artifact than to noise of the arterial portions of both the T1 measurement (4.17% v. 0.17%) and baseline image (4.14% v. 0.20%).

The optimal protocol, for $T1 = 300$ ms and $K^{\text{trans}} = 0.0025 \text{ min}^{-1}$, was determined to be the same as the standard protocol described above except with a TR increased to 110 ms and contrast injected as quickly as possible ($t_p = 15$ s). The optimal protocol had approximately one fourth of the variance of the standard protocol or one half of the standard deviation (optimal CV = 4% v. standard CV = 8%). The number of Monte-Carlo simulations was set to 120, which is the number required to distinguish between CV = 5% and CV = 7% normally distributed random variables with probability 0.80 using the Levene test for homogeneity of variation. The null hypothesis that the variances were equal was rejected with $p < .001$ for the Monte-Carlo simulations, well below the 95% confidence limit.

2.4 Discussion

DCE-MRI is a powerful tool for the assessment of vascular physiology, particularly in the arena of tumor perfusion and angiogenesis. To solve for both K^{trans} and K_{ep} (or equivalently V_e) *in vivo*, dynamic concentration data from both tissue and AIF are

required. Difficulties in obtaining case-by-case measurement of the AIF have motivated the use of a standardized AIF. Standard signal processing techniques can then be used to extract the desired parameters. For example, the impulse response for the two-compartment model is given by $h(t) = K^{\text{trans}} \cdot \exp(-k_{\text{ep}}t)$. Since $C_T(t) = \text{conv}(C_p(t), h(t))$, $h(t)$ can be obtained by deconvolution of the tissue concentration curve and AIF to obtain $h(t)$, and non-linear regression used to fit the parameters K^{trans} and k_{ep} (Tofts et al. 1999). Another approach is to assume an analytical form for $C_p(t)$, develop an analytical solution for $C_T(t)$ and fit the result to the tissue data using a nonlinear fitting technique. Also, one can calculate dC_T/dt at each sample point numerically from tissue data, and, treating it as an independent variable, perform a least squares linear fit to the differential equation itself with $C_p(t)$ and $C_T(t)$ as independent predictor variables of dC_T/dt .

Several models exist, and the choice between them depends both on the *in vivo* behavior of the contrast agent employed and the assumptions the investigator is able to make regarding the characteristics of the tissue of interest. A simple plasma-compartment model is appropriate for blood pool tracers that do not enter the tissue (Axel 1980). Additional compartments are required for tracers that cross the capillary endothelium. Port et al. (Port et al. 1999) suggested that in mammary tumors, the tissue compartment may be comprised of from one to three sub-compartments, and developed a Bayesian approach for selecting the appropriate combination of sub-compartments in a particular tumor based on the fit of tissue enhancement curves to acquired data.

Assuming a typical two-compartment model is chosen and that the AIF is to be measured from the images, there are still a variety of considerations in obtaining optimally precise (minimum variance) measurements of K^{trans} . For example, the best

choice of $T1_0$ measurement technique depends on both the expected $T1$ of the tissue of interest (Wang et al. 1987) as well as the expected sources of errors as can be seen from Figure 2.2. For example, in muscle ($T1 = 900$ ms) even a single acquisition of the dual tip-angle technique is superior to the spin-echo technique. In more fatty tissues ($T1 = 300$ ms) the choice is not so clear. If the expected errors are mostly noise, then they may be reduced by acquiring multiple averages, making the dual tip-angle technique superior for a given amount of time. However, if the expected errors are due to artifacts then it is possible that they will not be reduced by multiple averages and the spin-echo technique could be superior.

In Figure 2.3 it is apparent that, at the concentrations typical in such studies (0-200 μM), the errors in concentration are almost independent of the tissue $T1_0$, which is consistent with the flatness of the dual-tip angle curve over most of the range of $T1_0$ values. For such typical concentration values the errors will never be less than a 7% CV for an individual measurement, and at contrast-agent concentrations around 100 μM , where the CV is at least 11% for all tissue $T1_0$, the 95% confidence interval is between 78 and 122 μM .

An interesting result can be noted by examining Figures 2.3, 2.4, and 2.5 together. Specifically, the simulated tissue ($K^{\text{trans}} = .0025$) concentration curve in Figure 2.4 peaks at 130 μM , which, from Figure 2.3, results in concentration measurements with at least an 8.5% CV. It would be reasonable to expect that K^{trans} measurements from such data could be no more precise than 8.5% CV, however Figure 2.5 clearly shows that the K^{trans} measurements can have CVs as low as 3% or 4%. The least squares fit can thus be shown to suppress some of the variability and result in K^{trans} measurements that are more

precise than the source concentration measurements. Whether such a result can be generalized to other techniques, such as deconvolution, remains to be seen.

Some of the results shown in Figure 2.5 are counter-intuitive and clearly demonstrate the need for rigorous analysis of the precision of all DCE-MRI techniques. Specifically, the precision of K^{trans} was found to improve with lower tip-angles and longer TRs. As both of these tend to reduce the T1 weighting, this result is contrary to popular wisdom in designing DCE-MRI protocols, however the Monte-Carlo simulations were specifically used to verify this result. This counter-intuitive result appears not to be a spurious artifact introduced at some stage of the propagation of errors analysis, but rather a real trend not previously noted in the literature.

Figure 2.5 also shows the presence of a strong secondary effect as well as some main effects. Main effects are those where changing a parameter results in changes to the CV. For example, 2.5a shows that the main effect of $T1_0$ is fairly minimal; while in Figure 2.5b the main effect of TR is noticeably greater. Secondary effects are those where changing one parameter alters the main effect of another parameter. For example, in Figure 2.5b the main effect of TR is relatively unaffected by variations in tissue $T1_0$ meaning there is little secondary effect of $T1_0$ on TR. Figure 2.5c, however, shows the strong secondary effect of $T1_0$ on tip-angle. This secondary effect seems to be stronger than the main effect of $T1_0$ and perhaps even as strong as the main effect of tip-angle. Therefore, it is important to match the tip-angle of a given DCE-MRI protocol to the expected T1 of the tissue of interest in order to obtain the most precise K^{trans} measurements possible. At the time that this matching is performed, the $T1_0$ measurement technique should probably also be matched to the same T1. Note that the

choice of tip-angle will also impact the precision of AIFs obtained from the image data. The technique used here includes this effect by propagating errors from both the tissue pixel (picture element) and the blood ROI, and can thus be used to select the optimal parameter in the face of such potentially competing factors. The use of an assumed or standardized AIF may lead to different results, both in terms of the optimal parameter values and the final precision of the measurement.

The increased TR of the optimal DCE-MRI protocol also coincides with the results of Figure 2.5, and invites the question, “Why has this trend not been previously discovered in the literature?” This is particularly surprising considering the magnitude of the improvement: a four-fold reduction in variance or a two-fold reduction in standard deviation. Such a large improvement would allow the same statistical power in scientific DCE-MRI studies with as little as one-fourth of the number of samples. Alternatively, this improvement could allow detection at the same statistical power of true differences of as little as one-half of the current detectable differences with the same number of samples. There are several reasons why such an impressive trend may not have been noticed previously. First, all methodological studies to date have focused on the accuracy, not the precision, of DCE-MRI. Although it is not impossible to accidentally discover precision effects during such studies, it is much less likely than with the direct approach used here. Second, statistical tests on measures of spread are, in general, much less powerful than statistical tests on measures of central tendency. Recall that $N = 120$ for the Monte-Carlo validation despite the size of the effect. Most scientific studies performed using DCE-MRI are designed for tests on the mean and thus have insufficient number of samples to statistically detect differences in precision despite the potential

benefit in statistical power that could be attained for the tests on the mean. Third, even normal tissues can exhibit fairly large natural variation in perfusion values. Because variances are additive, the natural variance of the tissue perfusion will be added to the error variance of the measurement technique to obtain the total variance of the measurements, which is the only observable variance. This will further reduce the statistical power of tests on the variance and require even greater N . Finally, Monte-Carlo methods, although widely accepted in the DCE-MRI literature, are inherently limited by the assumptions made in generating the random signal. Although care was taken to use the best and most complete methods possible for generating and processing the signals, there is always a danger that some important factor was overlooked. Despite this well-known weakness of Monte-Carlo simulations, they are commonly used and accepted primarily because of the lack of any perfusion phantoms and secondarily because of its success in correctly identifying and guiding the development of several important DCE-MRI techniques (Su and Nalcioglu 1993; Bahn 1995; Ostergaard et al. 1996a; Ostergaard et al. 1996b; Boxerman et al. 1997; Dennie et al. 1998; Karlsen et al. 1999; McMahon and Oldfield 1999).

The artifact level chosen in the calculation of the results in Figure 2.6 is slightly greater than the noise level and should be barely perceptible. The most important difference between such errors and the noise is not their magnitude, but rather that the artifact errors are not reduced through spatial or temporal averaging. This error is used to demonstrate the absolute sensitivity of K^{trans} precision to small errors in the various sources. In particular, consider the question posed in the introduction. If there are no image errors other than a small error that occurs in the arterial portion of the dynamic

images, then the error propagated into K^{trans} will be less than a 0.5% CV. Conversely, if the same error is instead introduced into the same portion of the $T1_0$ measurement image it will propagate into more than a 4% CV in K^{trans} .

Errors in the arterial portion of the dynamic images had the smallest impact on the precision of the K^{trans} measurement. This indicates that future developmental efforts would probably be more productively focused elsewhere (e.g. optimizing the dynamic sequence for the tissue of interest rather than for blood). The large impact of small artifacts in the tissue portion of the baseline image indicates that care should be taken in this portion of the DCE-MRI protocol. If *any* perceptible artifact exists in the baseline image, then it can be assumed that the K^{trans} of the corresponding tissue cannot be measured more accurately than within $\pm 15\%$ and is probably much worse (e.g., $\pm 20\%$ for noise in all images + small artifact only in baseline).

It should be emphasized that, as with any analysis, the results of this analysis were necessarily limited by the assumptions made throughout the propagation of errors and K^{trans} analysis. For example, water exchange was assumed to be fast, only two tissue compartments were considered, and off-resonance and flow effects were not modeled in the signal equations (Landis et al. 1999). While the specific results may need to be recalculated in situations where one or more of the assumptions do not hold, a careful examination of the propagated errors is important to consider in the development of any DCE-MRI protocol as evidenced by the aforementioned counter-intuitive results. It is possible that similar rigorous analyses of other quantitative measurement techniques, such as diffusion-tensor imaging, may also yield counter-intuitive results, and that failure

to perform such analyses could lead to sub-optimal precision in the quantity of interest for currently proposed protocols.

DCE-MRI perfusion assessment requires the measurement of multiple signals as inputs for a complex post-processing algorithm that yields a single index of blood flow. A standard index that may be obtained is K^{trans} which is proportional to blood flow and dependent upon capillary perfusion for clinically approved diffusible tracers such as Gd-DTPA. The accuracy of measured K^{trans} is dependent upon the accuracy of the source image data used in its calculation and also a function of how these errors are propagated through the signal equations and modified Kety two-compartment model. We have shown how K^{trans} error has different sensitivities to both sequence and tissue parameters. In T1 weighted source images, signal decreases with increasing $T1_0$. With constant noise, SNR varies inversely with $T1_0$, causing K^{trans} error to increase with $T1_0$. The results of our analysis (Figure 2.2) suggest that at high $T1_0$ the sensitivity to error may not be negligible. For example, with the spin-echo measurement technique, the T1 of tissues such as gray matter ($T1 = 950$ ms) cannot be measured more accurately than within $\pm 11\%$. The analysis outlined above provides a means for determining the magnitude of the increase, as well as a means for determining the SNR changes necessary to reduce K^{trans} error to acceptable levels when attempting voxel-resolution perfusion imaging of tissues with high $T1_0$. This will be particularly important when attempting to compare perfusion in tissues with different $T1_0$, or in the same tissue undergoing changes over time (e.g. apoptosis) that concurrently lead to changes in $T1_0$.

In conclusion, our findings suggest that estimates of K^{trans} in tissue at a voxel within an image are most sensitive to errors in the images used to estimate baseline T1. Thus,

errors in pre-contrast T1 estimates of the arterial blood and tissue leads to far greater error in estimates of perfusion than comparable error in images used to measure the relative signal change after contrast administration. A fixed error in post contrast arterial signal that propagates to a 0.4% error in K^{trans} will propagate to a much greater error when it occurs in pre-contrast arterial signal (5%), pre-contrast tissue signal (5-12%), post contrast tissue signal (4-10%), tissue T1₀ measurement signal (5-25%), or arterial T1 measurement signal (27%). However, in spite of these differences in sensitivity, the signal to the uncorrelated background noise (SNR) in all of our images did not propagate to large errors. On the other hand, larger errors from other sources such as breathing or other motion could contribute to large variation in K^{trans} estimates, especially when they occur during T1₀ measurement signal acquisition. In other words, SNR is not as important as minimizing artifact in acquisitions used to estimate T1₀. Finally, due to strong main and secondary effects, it is important to optimize a DCE-MRI protocol for the tissue of interest, particularly with regards to matching the technique to the anticipated T1₀ of the target tissue.

Chapter 3. Rapid Gridding-Reconstruction Using Tables

3.1 Introduction

The reconstruction of non-uniformly sampled Fourier data is a well-known problem in many fields including radio astronomy, synthetic aperture radar, and MRI (Hogg et al. 1969; Choi and Munson 1998). In MRI this reconstruction problem arises when an image is acquired using a non-rectilinear or non-Cartesian k-space trajectory. Examples include radial k-space acquisitions for (filtered) back projection reconstruction (PR), spiral trajectories, rosette trajectories, or any number of other proposed trajectories. Such trajectories may be desirable for a variety of reasons including rapid acquisition times or good artifact reduction properties (Meyer et al. 1992; Irarrazabal and Nishimura 1995; Zhou et al. 1998). For example, in an iMRI setting, many trajectories could allow simultaneously high temporal and spatial resolution while providing data collection throughout a targeted region of k-space, determined by resolution requirements. This is in contrast to the currently more common keyhole techniques, which increase temporal resolution by sampling only a portion of the targeted region of k-space (Duerk et al. 1996). However, to be useful in iMRI and other fluoroscopic applications, these non-rectilinear trajectories currently require access to numerically intensive special reconstruction code, sophisticated computer hardware or special reconstruction hardware, thereby limiting their use in some applications (Kerr et al. 1997; Eggers and Proska 1999).

The solutions to the problem of how to reconstruct non-uniformly sampled k-space data can be broadly grouped into two categories. The first group is solutions that rely on direct application of the DFT (discrete Fourier transform) from the sampled data points.

The second group is solutions which use a gridding algorithm to estimate an equivalent set of uniformly sampled Fourier data from the non-uniformly sampled acquired data, followed by reconstruction using the FFT (O'Sullivan 1985; Jackson et al. 1991; Rosenfeld 1998).

An efficient method for performing either the DFT reconstruction or data gridding operations would allow completely real-time reconstruction (i.e., limited by the acquisition time) of these non-rectilinear trajectories, which could open the door for their use in a wider variety of clinical situations. Moreover, if the method could allow the calculation of the result due to each individual data point immediately upon acquiring the data point, it would have an advantage even over current keyhole or sliding window reconstruction methods. The image could be updated after each individual data point, or after any given number of data points, to reflect the most recent information. This would be a necessary precursor to truly fluoroscopic MRI acquisitions.

The purpose of this work was to determine if the computational efficiency of pre-calculated look-up tables could be effectively exploited to achieve real time reconstruction, or real time gridding during non-rectilinear or non-Cartesian k-space sampling when using only modest conventional computational resources. The remainder of this paper will describe table-based computational methods that permit rapid reconstruction of non-uniformly sampled k-space data even on a typical, PC (personal computer) for the DFT and for some of the currently used gridding algorithms.

3.2 Methods

A linear look-up table arrangement (in time) and hence a simple memory allocation scheme allows very predictable access into a look-up table, which in turn suggests the

possibility for fast table-driven computation (Booth 1997). Proposed methods to achieve an efficient DFT and gridding table method are described below.

3.2.1 Description of a Lookup Table DFT method

A DFT table is a pre-calculated table of weights describing how each data point affects the entire image space. That is, the precalculated table consists of the DFT of a delta function at the sampled location in k-space. Consider the DFT transform pair for a 1D (one dimensional) signal sampled at N equidistant points in both time and frequency.

$$\begin{aligned} F(k) &= \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} f(n) \cdot e^{-j2\pi kn/N} \quad k = 0, 1, \dots, N-1 \\ f(n) &= \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} F(k) \cdot e^{j2\pi kn/N} \quad n = 0, 1, \dots, N-1 \end{aligned} \quad [3.1]$$

where $F(k)$ is the Fourier data and $f(n)$ is the reconstructed signal in the space (or time) domain. The synthesis equation may be rewritten (Eq. [3.2]) when the requirement for equidistant frequency samples is relaxed.

$$f(n) = \frac{1}{\sqrt{N}} \sum_{i=0}^{N_s-1} \rho^{-1}(k(i)) \cdot F(k(i)) \cdot e^{j2\pi k(i)n/N} \quad n = 0, 1, \dots, N-1 \quad [3.2]$$

where the $k(i)$ are the sampled locations in frequency, N_s is the number of frequency samples and $\rho^{-1}(k(i))$ is the density compensation function at the sample locations. The function $\rho^{-1}(k(i))$ is equal to 1.0 if the $k(i)$ are equidistant; otherwise it is large in the relatively undersampled regions and small in oversampled regions (Hoge et al. 1997; Pipe and Menon 1999).

Letting

$$T(i, n) = \frac{1}{\sqrt{N}} \rho^{-1}(k(i)) \cdot e^{j2\pi k(i)n/N} \quad n = 0, 1, \dots, N-1 \quad i = 0, 1, \dots, N_s-1 \quad [3.3]$$

then Eq. [3.2] may be rewritten as

$$f(n) = \sum_{i=0}^{Ns-1} T(i, n) \cdot F(k(i)) \quad n = 0, 1, \dots, N-1 \quad [3.4]$$

If the $k(i)$ are known in advance, as is the case in any MRI pulse sequence, then it is possible to pre-calculate and store all of the $T(i, n)$, and reload them only during the acquisition/reconstruction. This is the mathematical basis of the DFT table-enhanced method. Note that all of the transcendental functions and the density compensation functions are pre-computed, and the task, at run time, consists entirely of multiplication and addition. Note also that each of the $f(n)$ can be processed independently and simultaneously. Because of this, it is convenient to write the above equations in a vectorized format.

$$\mathbf{f} = \sum_{i=0}^{Ns-1} \mathbf{T}(i) \cdot F(k(i)) \text{ where}$$

$$\mathbf{f} = \begin{bmatrix} f(0) \\ \vdots \\ f(N-1) \end{bmatrix} \text{ and } \mathbf{T}(i) = \begin{bmatrix} T(i, 0) \\ \vdots \\ T(i, N-1) \end{bmatrix} \quad [3.5]$$

For a 2D reconstruction

$$T(i, x, y) = \frac{1}{N} \rho^{-1}(k_x(i), k_y(i)) \cdot e^{j2\pi(k_x(i)x + k_y(i)y)/N}$$

$$x = 0, 1, \dots, N-1 \quad y = 0, 1, \dots, N-1 \quad i = 0, 1, \dots, Ns-1$$

$$f(x, y) = \sum_{i=0}^{Ns-1} T(i, x, y) \cdot F(k_x(i), k_y(i)) \quad x = 0, 1, \dots, N-1 \quad y = 0, 1, \dots, N-1 \quad [3.6]$$

or, in vector/matrix format

$$\mathbf{f} = \sum_{i=0}^{Ns-1} \mathbf{T}(i) \cdot F(k_x(i), k_y(i)) \text{ where}$$

$$\mathbf{f} = \begin{bmatrix} f(0,0) & \cdots & f(N-1,0) \\ \vdots & \ddots & \vdots \\ f(0,N-1) & \cdots & f(N-1,N-1) \end{bmatrix} \quad \text{and}$$

$$\mathbf{T}(i) = \begin{bmatrix} T(i,0,0) & \cdots & T(i,N-1,0) \\ \vdots & \ddots & \vdots \\ T(i,0,N-1) & \cdots & T(i,N-1,N-1) \end{bmatrix} \quad [3.7]$$

In MRI, each of the $F(k_x(i), k_y(i))$ is one of the acquired data points. A table-based DFT reconstruction, therefore, consists of the following steps:

- 1) Compute and store all of the $\mathbf{T}(i)$. This file is the pre-calculated look-up table and may be used repeatedly for a given trajectory. That is, $\mathbf{T}(i)$ needs only to be calculated once since it can be used, for example, for any pulse sequence in which that k-space trajectory is used.
- 2) Load the look-up table into memory.
- 3) Set the current estimate of \mathbf{f} to zero.
- 4) Acquire the first data point, $F(k_x(0), k_y(0))$.
- 5) Calculate $\mathbf{T}(0) \cdot F(k_x(0), k_y(0))$.
- 6) Accumulate (add) the result into the current estimate of \mathbf{f} .
- 7) Repeat steps 4 – 6 for the rest of the data points.

3.2.2 Description of a Look-up Table Method for Gridding

A table for gridding operations used prior to 2D-FFT reconstruction of non-rectilinearly sampled MRI data would be a pre-calculated table of weights describing how each data point contributes to a small, rectilinearly sampled, k-space neighborhood. The gridding algorithm presented by O’Sullivan consists of several steps designed to compute, not the actual Fourier data at the grid points, but a set of gridded Fourier data

whose FFT is equal to a sampled version of the integral FT (Fourier transform) of the non-uniformly sampled function (O'Sullivan 1985). The basic steps are as follows.

- 1) Multiply the trajectory data, $F(k_x(i), k_y(i)) \cdot \delta(k_x - k_x(i), k_y - k_y(i))$, by an appropriate density compensation function, $\rho^{-1}(k_x(i), k_y(i))$.
- 2) Convolve the density compensated trajectory data by an appropriate function, $C(k_x, k_y)$. Ideally, a sinc function covering all of k-space would be used, but this would lead to a significant computational burden and thus finite windows, like Kaiser-Bessel functions, are chosen instead.
- 3) Sample the convolved function onto a rectilinear k-space grid by multiplying by the comb function, $comb(k_x, k_y) = \sum_{u=0}^{N-1} \sum_{v=0}^{N-1} \delta(k_x - u, k_y - v)$. [3.8]
- 4) Transform the gridded data, $G(k_x, k_y)$, from the previous step using the FFT.
- 5) Deconvolve the transformed image by dividing it by $c(x, y)$, the FT of the convolution function used in step 2. This step is unnecessary if a sinc function was used in the gridding convolution; for many convolution functions this correction produces minimal improvements in the central portion of the image.

The result from step 3 of the gridding algorithm is the gridded Fourier data that lies on a rectilinear grid (Jackson et al. 1991), expressed here as

$$G(k_x, k_y) = comb(k_x, k_y) \cdot \left[C(k_x, k_y) * \sum_{i=0}^{N_s-1} \left(\rho^{-1}(k_x(i), k_y(i)) \left(F(k_x(i), k_y(i)) \cdot \delta(k_x - k_x(i), k_y - k_y(i)) \right) \right) \right] \quad [3.9]$$

Letting

$$\begin{aligned}
Tg(i, k_x, k_y) &= comb(k_x, k_y) \cdot \\
&\quad \left(C(k_x, k_y) * \left(\rho^{-1}(k_x(i), k_y(i)) \cdot \delta(k_x - k_x(i), k_y - k_y(i)) \right) \right) \\
i &= 0, 1, \dots, Ns - 1
\end{aligned} \tag{3.10}$$

then

$$G(k_x, k_y) = \sum_{i=0}^{Ns-1} Tg(i, k_x, k_y) \cdot F(k_x(i), k_y(i)) \tag{3.11}$$

or, in matrix format,

$$\mathbf{G} = \sum_{i=0}^{Ns-1} \mathbf{Tg}(i) \cdot F(k_x(i), k_y(i)) \text{ where}$$

$$\mathbf{G} = \begin{bmatrix} G(0,0) & \cdots & G(N-1,0) \\ \vdots & \ddots & \vdots \\ G(0,N-1) & \cdots & G(N-1,N-1) \end{bmatrix} \tag{3.12}$$

and

$$\mathbf{Tg}(i) = \begin{bmatrix} Tg(i,0,0) & \cdots & Tg(i,N-1,0) \\ \vdots & \ddots & \vdots \\ Tg(i,0,N-1) & \cdots & Tg(i,N-1,N-1) \end{bmatrix}$$

Eq. [3.12] has the same form as Eq. [3.7]. Therefore, a similar table-based reconstruction process used above to calculate \mathbf{f} could also be used here to calculate the gridded data, \mathbf{G} . Once complete, the image can be reconstructed via the 2D (or ND [n dimensional]) FFT of \mathbf{G} . Further, any operation that can be expressed as a linear function on the sampled data points, e.g. $F(k_x(i), k_y(i))$, can be computed using the table-based reconstruction method.

Once the gridding table, \mathbf{Tg} , is created and stored it can be loaded at run time into memory and then used to greatly speed the processing. The table-enhanced gridding algorithm proceeds in the following steps.

- 1) Do a table-based reconstruction using \mathbf{Tg} to reconstruct \mathbf{G} (as opposed to using \mathbf{T} to reconstruct \mathbf{f}). This covers steps 1 – 3 in the gridding algorithm.
- 2) Transform the gridded data using the FFT.
- 3) Deconvolve the transformed image by dividing it by $c(x,y)$, the FT of the convolution function used to calculate the table.

3.2.3 Programming Methods

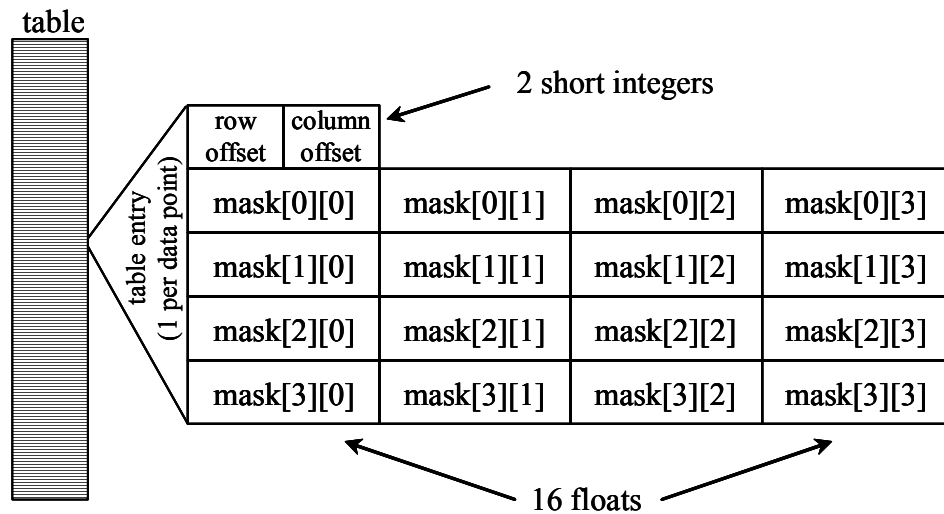


Figure 3.1 Structure of the gridding table. Table example based on a 4 by 4 convolution window. The row offset and column offset values accomplish the sparse matrix representation. The mask elements are the only non-zero elements of the sparse matrix.

For a small convolution window, the vast majority of the points in the $\mathbf{Tg}(i)$ are equal to zero. This allows use of a sparse matrix representation where the zeroes need neither be stored nor calculated at run time. Thus the tables are organized as shown in Figure 3.1. The i 'th data point acquired from the scanner corresponds to the i 'th entry in the table

that is, in turn, simply a sparse matrix representation of $\mathbf{Tg}(i)$. When a data point is acquired, the row and column offsets are read and used as indices into the gridded array, which is the current estimate of \mathbf{G} . Then all of the weights in the table entry are multiplied by the data point and accumulated into the indicated locations in the array. The next entry in the table is then loaded and the processor waits, if necessary, to acquire the next data point. If the processor has idle time while waiting for the next data point, the size of the convolution window may be increased or other computations may be performed without adding any time to the reconstruction. A larger convolution window generally results in less aliasing artifact from the gridding procedure itself (Jackson et al. 1991).

All of the gridding tables used here were calculated with the Kaiser-Bessel window Eq. [3.13] (O'Sullivan 1985; Jackson et al. 1991):

$$C(k_x, k_y) = C(k_x) \cdot C(k_y)$$

$$C(k) = \begin{cases} \frac{1}{W} I_0 \left[\beta \sqrt{1 - \left(\frac{2k}{W} \right)^2} \right], & |k| \leq \frac{W}{2} \\ 0, & \text{otherwise} \end{cases} \quad [3.13]$$

where I_0 is the zero-order modified Bessel function of the first kind, W is the width of the window, and β is a free parameter. It was set to the optimal value (5.7567 for $W=4$) presented by Jackson et al. for windows of width 5.0 or less or otherwise to the approximately optimal value (9.4248 for $W=6$ or 12.566 for $W=8$) from the simple equation presented by Wajer et al. (Jackson et al. 1991; Wajer et al. 1999).

Both radial and spiral trajectories were considered. The density compensation function was multiplication by the k-space radius of the data point for both cases (Hoge et al. 1997). The comb function used to produce the k-space grid had a sampling interval

of unity as shown above; it was not oversampled as has been suggested in order to increase the FOV and thereby move the aliasing energy away from the ROI (Jackson et al. 1991).

Three different complex number representations were investigated for the DFT table reconstruction. The first was the standard rectangular representation (i.e., $A + jB$). The second and third were polar representations ($Ae^{j\theta}$) with the phase discretized into 16 and 8 bits respectively and the magnitude being constant (equal to the value of $\rho^{-1}(k_x(i), k_y(i))$, representing the density compensation function) for each table entry. The Kaiser-Bessel window function is entirely real, and therefore the various complex number representations used for a DFT table are neither possible nor necessary for a gridding table.

All algorithms were implemented in either the Interactive Data Language (IDL: Research Systems, Inc., Boulder CO) or in C++ (Visual C++ v5.0: Microsoft Inc., Redmond, WA). For C++ programs, the OptiVec C/C++ complex math library (by Martin Sander, shareware v1.6), which allows for complex data and includes the FFT, was used. The functions in this library have been partially optimized by using assembly-level instructions and vectorized computational techniques. For example, the FFT algorithms were written in assembly level.

3.2.4 Experimental Methods

3.2.4.1 DFT Methods

The DFT table experiments were designed to determine differences in accuracy, computational speed, and table size of the various complex number representations. A mathematical phantom, similar to a CT resolution phantom for large data sets ($>100^2$ data

points) or a single square for small data sets, was calculated based on a series of 2D rectangular pulses of width (w_x, w_y) centered about the point (c_x, c_y) (u in Eq. [3.14]).

$$u_{w_x, w_y, c_x, c_y}(x, y) = u_{w_x, c_x}(x) \cdot u_{w_y, c_y}(y)$$

$$u_{w, c}(x) \xleftrightarrow{F} 4 \frac{e^{-jkc} \sin(\frac{kw}{2})}{k} \quad \text{where}$$

$$u_{w, c}(x) = \begin{cases} 1, & (|x - c| \leq \frac{w}{2}) \\ 0, & \text{otherwise} \end{cases} \quad [3.14]$$

Its analytical FT provided simulated k-space data that was uniformly sampled on a 32-by-32 square grid. Look-up tables were developed to reconstruct from the 32-by-32 k-space data onto a 32-by-32 image-matrix. Three different look-up tables were pre-calculated using IDL according to the three different representations of complex numbers: one using the standard rectangular representation and the other two using polar representations with the phase discretized to 8 and 16 bits respectively. The size of each of the look-up tables was noted. The table-based DFT reconstruction was performed following the aforementioned method and the time required was measured using the computer's internal clock and the standard time routines. After measuring the reconstruction time with these tables, the fastest one was re-coded in C++ and compiled using the previously mentioned compiler set to optimize for speed. The deviations from use of the FFT were calculated and compared. The speed results obtained with IDL were checked by using C++ to measure the average speed of one-million multiply-accumulate operations on random data for each of the various complex number formats. The experiment was not designed to assess or compare the computational time of the 2D-FFT versus a table-based reconstruction, but rather to explore only errors and reconstruction

times associated with different table formats that might lend themselves to more rapid reconstruction.

3.2.4.2 Gridding Methods

All of the gridding procedures described in Methods Section B and Eq. [3.8] – Eq. [3.12] were coded with the Microsoft C++ compiler. All reported reconstruction times were obtained using the standard C++ time routines with the code compiled to optimize for speed. They were performed on a Windows NT 4.0 single processor 600 MHz Pentium III PC with 512MB of RAM (random access memory). All of the gridding-table experiments used the same gridding code and the same 2D-FFT code for reconstruction; the only differences were in the code used to generate the pre-calculated look-up tables and the raw data.

The first gridding-table trial measured the time required to grid simulated data resulting from sampling along a spiral trajectory. The simulated spiral trajectory data was generated from evaluation of the analytic expression for the 2D-FT of the mathematical resolution phantom at each k-space sampling location. An Archimedian spiral trajectory was used:

$$(k_x(i), k_y(i)) = (re(\frac{i}{256} e^{j\pi i / 64}), im(\frac{i}{256} e^{j\pi i / 64})) \quad i = 0, 1, \dots, (128 \cdot 128) - 1 \quad [3.15]$$

where the real part of the number represents the distance in k-space along the k_x direction and the imaginary part represents the position along the k_y direction. This spiral has 128 turns and 128 sampled data points per turn with the radius of the last data point being 64. The data were gridded onto a $(128+W-1)^2$ point matrix where W is the width of the Kaiser-Bessel window as is dictated through the use of the convolution gridding method without oversampling. The gridded matrix was zero padded to 256^2 points as

required by the OptiVec FFT routine. A reference image was generated by applying the gridding reconstruction algorithm to data obtained by sampling the transform of the mathematical resolution phantom along a 128^2 Cartesian trajectory.

These images and times were also compared to a non-table-based gridding operation. The non-table-based operation used the same code as the table-based operation except in the innermost loop where the simple multiply-accumulate of the pre-calculated table value was replaced with the usual inline computation of the density compensation function and the Kaiser-Bessel window. This window utilized a fast Taylor-series expansion of the zero-order modified Bessel function of the first kind. Enough terms were used to ensure accuracy to one part per million. This is in contrast to the much slower function used in the table calculations, which proceeds to full machine precision. However, since this is pre-calculated, the increased time does not impact on the gridding/reconstruction time.

The second gridding-table test reconstructed simulated data sampled along a radial k-space trajectory, as would occur in projection reconstruction acquisitions, which are now regaining attention in a variety of MRI applications. The same mathematical resolution phantom was used to generate simulated data by evaluation of the analytic expression for its 2D-FT. The radial trajectory consisted of 180 views, each offset by 1° from the previous; 256 points were collected for each view with the 128^{th} data point located at the origin of k-space. The data were gridded onto a $(256+W-1)^2$ point matrix where W is the width of the Kaiser-Bessel window. The gridded data was generated, the gridding time was measured, and the resulting times and images were compared to those obtained through a non-table-based algorithm as above.

The third and final trial of the table-based gridding algorithm reconstructed in-vivo radial k-space data acquired in our lab. A pig's neck was scanned using a 180 view, 256 points per view radial k-space True-FISP acquisition during iMRI needle insertion. The pulse sequence was implemented on a Siemens 0.2T Magnetom Open MRI system (Siemens Medical Systems, Erlangen, Germany). The echo occurred at the 128th point in each view. The views were separated by approximately (but not exactly) 1° due to the MRI system's hardware truncating the gradient tables to a fixed precision. The data were again gridded onto a $(256+W-1)^2$ point matrix where W is the width of the Kaiser-Bessel window. The time required for table-based gridding was measured and compared with non-table-based gridding as with previous experiments.

3.3 Results




Figure 3.2 Accuracy comparison between the different complex number representations for the DFT table. For the polar representations, the number of bits indicates how many bits of precision were used to discretize the phase. Deviations are scaled to the range of bytes (0-255) and are measured relative to the FFT of the same data

Results from the DFT experiment showed that the unavoidable Fourier artifacts, such as Gibbs ringing, overwhelm the miniscule differences between the FFT and the table-enhanced DFT by five orders of magnitude as would be expected. Figure 3.2 shows the deviations, with respect to the FFT, generated by using the various table representations from the first DFT table experiment. The deviations in this image are calculated by taking the absolute value of the magnitude error. Before the errors were calculated, the images were scaled up to the range of bytes (gray levels) to give a reference for how visible those errors would be on a typical display device. The rectangular complex table generated the most compact error pattern with the lowest peak value. The two polar tables generated different error patterns with the errors from the 8-bit table being highest,

5 gray levels or approximately 2% peak error. Also note that the error for the 8-bit table was at or near its peak value (5 gray levels) for many more pixels than was the 16-bit table.

Table 3.1

DFT-table reconstruction time for various complex number representations

Representation	Reconstruction Time
Rectangular	422 ms
16-bit Polar	839 ms
8-bit Polar	731 ms
C++ (rectangular)	110 ms

Unfortunately, neither of the polar representations performed better than the standard rectangular table in speed as shown in Table 3.1. The general trend of Table 3.1 was also evident with the multiply-accumulate (MAC) measurements on random data in C++ (rectangular: 351 ns/MAC; 16-bit polar: 1041 ns/MAC; 8-bit polar: 972 ns/MAC). The fastest time, obtained for the optimized C++ implementation of the rectangular table, was 110 ms for a 32 x 32 reconstruction.

Table 3.2

Computation times for various k-space trajectories with lookup table gridding

Trajectory	W = 4 ^a	W = 6 ^a	W = 8 ^a	2D-FFT
Spiral^b (table-enhanced)	12.8 ms	29.3 ms	47.17 ms	44.3 ms ^d
Spiral^b (standard)	582.3 ms	1499.5 ms	2992.9 ms	44.3 ms ^d
Radial^b (table- enhanced)	36.2 ms	80.1 ms	131.5 ms	219.4 ms ^e
Radial^b (standard)	1639.2 ms	4218.1 ms	8415.3 ms	218.7 ms ^e
Radial^c (table- enhanced)	36.1 ms	80.0 ms	130.8 ms	218.4 ms ^e
Radial^c (standard)	1636.9 ms	4211.7 ms	8399.0 ms	218.8 ms ^e

a. W is the width of the Kaiser-Bessel convolution kernel

b. Simulated data

c. Actual data from pig

d. 256² point 2D-FFTe. 512² point 2D-FFT

In Table 3.2 the left-hand column indicates the trajectory and the top row indicates the size of the Kaiser-Bessel window used in the computations for gridded data trials. The size of the window determines the number of mask points in each table entry. These convolution kernels were generally larger than those typically used elsewhere (Irrazabal et al. 1993; Irrazabal and Nishimura 1995; Butts et al. 1997; Sedarat and Nishimura 1999). The images and results from the individual trials will be mentioned below, but it is important to note the ratio of the computation times of the table-based and the non-table-based trials. The table-enhanced gridding operation ranged from 45 to 64 times




Figure 3.3 Gridded and reconstructed image from the spiral trajectory sampling of the mathematical resolution phantom, gridded in 12.8 ms using the table-based reconstruction with a 4 by 4 Kaiser-Bessel convolution window. 256^2 point FFT took 44.3 ms on the same system.

faster than the standard gridding operation, and in most cases the table-enhanced gridding was faster than the corresponding 2D-FFT while the non-table-enhanced gridding was always slower.

The spiral trajectory, the first gridding-table experiment, produced the results shown in Figure 3.3 and Table 3.2. Figure 3.3 is the gridded and reconstructed data from a mathematical resolution phantom, produced using an analytical Fourier-domain representation of the rectangular regions.




Figure 3.4 Difference images with respect to a reference image, obtained by applying the gridding reconstruction algorithm to data from sampling the transform of the mathematical resolution phantom along a 128^2 Cartesian trajectory. Peak difference with the gridded/reconstructed spiral image was 93 gray levels in the outer regions and 29 in the central region. The spiral vs. reference difference image (a), which shows artifacts due to the spiral trajectory, was scaled by 2.7 to improve contrast visibility. Peak difference with the non-gridded FFT-reconstructed image was 5 gray levels. The non-gridded vs. reference image (b), which shows artifacts due only to the gridding, was scaled by 51 to improve contrast visibility. K-space data gridded using the table-based reconstruction with a 4 by 4 Kaiser-Bessel convolution window.

Figure 3.4a is the difference between the spiral image and the reference image, scaled up by 2.7 times to improve contrast visibility. The figure shows that the only artifacts in the original image are some streaking artifacts of 93 gray levels (peak difference) on the edge or 29 gray levels (peak difference) in the central region. This difference arises from the non-uniform sampling. Figure 3.4b is the difference between the non-gridded Cartesian data image and the reference image, scaled up by 51 times to improve visibility. The peak difference is 5 gray levels. The peak difference between the




Figure 3.5 Reconstructed image from the radial trajectory sampling of the mathematical resolution phantom, gridded in 36.2 ms using the table-based reconstruction with a 4 by 4 Kaiser-Bessel convolution window. 512^2 point FFT took 219.4 ms on the same system.

reconstructed image obtained with the table-based gridding and the one obtained with standard gridding is one gray level. These small errors occurred slightly more frequently in the corners of the image.

The second experiment, the simulated radial k-space acquisition, resulted in the image in Figure 3.5 and the computational speeds given in Table 3.2. It took nearly three times longer to reconstruct the PR sequence than the previous spiral trajectory, as expected since the radial k-space PR trajectory had nearly three times the number of data points. The artifacts in this reconstruction were similar to, but fainter (35 gray levels peak), than the artifacts in the spiral experiment.




Figure 3.6 Reconstructed magnitude image (a) from the *in-vivo* radial acquisition of a pig neck showing inserted needle, gridded in 36.1 ms using the table-based reconstruction with a 4 by 4 Kaiser-Bessel convolution window. 512^2 point FFT took 218.4 ms on the same system. Also shown for comparison is a filtered back projection reconstruction (b) of the same data. The 180-view radial data was acquired using a True-FISP sequence with TE = 9.8 ms and TR = 20.8 ms at 0.2 T.

The third and final gridding-table trial is summarized in Figure 3.6a and Table 3.2. Even using this actual MRI data, with non-zero imaginary parts, the computational speed is essentially the same as from the previous experiment, as was expected for two trajectories with the same number of data points. Figure 3.6b is a filtered back projection reconstruction of the data. Observe that most of the artifacts in the gridded image (with the notable exception of the rings in the corner) also occur in the back-projected image.

3.4 Discussion

Many of the points of this paper may seem fairly obvious to some readers, but a thorough investigation regarding the applicability of look-up tables to non-rectilinear

MRI reconstruction in general, and these algorithms in particular, is absent in the literature. Prior to this work, it was doubtful whether such large tables would be useful in accelerating any algorithm. The conventional wisdom regarding look-up tables is that they must be small enough to fit inside the cache to achieve computational efficiency, while the tables used in this manuscript are necessarily many times larger.

The DFT-table experiment, the comparison of rectangular vs. polar representations, indicates that the rectangular representation is preferable over either of the polar representations (rectangular vs. 16-bit polar vs. 8-bit polar representations had peak deviations of 0.000444 vs. 0.0224 vs. 5.42 gray levels and reconstruction times of 422 ms vs. 839 ms vs. 731 ms). The computational burden of the DFT reconstruction is proportional to the number of data points in the trajectory times the number of pixels in the final image (1024 times 1024 for the above experiments). Thus the fastest results from these experiments (110 ms) can be extrapolated to give anticipated best-case reconstruction times for larger trajectories or image sizes. For a 128 by 128 trajectory and final image (or any arbitrary trajectory with 128^2 data points) the extrapolated reconstruction time is 28.16 s; for a 256 by 256 trajectory and final image the time is 7.5 min. These reconstruction times are much greater than would be useful in any kind of time-sensitive application. The DFT-table method is still useful to consider because, due to its inherently parallel structure, specialized hardware and advances in computer technology may eventually make the table-based DFT even faster than the FFT (discussed in further detail below). Unfortunately, direct application of the DFT on the non-rectilinear data points is simply too computationally intensive, and thus currently

takes too long, to warrant further investigation for use in a typical situation with limited computer resources.

The first gridding experiment was the first one to use a non-Cartesian sampling scheme and thus the first one that needed a non-FFT reconstruction. The most interesting aspects of this experiment are the difference images between the reconstructed gridded data and the reference image and between the non-gridded 2D-FFT of the Cartesian data and the reference image. Though not necessary to reconstruct the image, the use of the gridding operation on the Cartesian data allows one to separately examine the effects of the trajectory and the effects of the gridding. The spiral vs. reference difference image is shown (peak difference: 93 outer region, 29 central region), scaled by 2.7 times, in Figure 3.4a. This image indicates that the artifacts in Figure 3.3 are largely due to the trajectory itself. The non-gridded vs. reference difference image is shown (peak difference: 5 gray levels), scaled by 51 times, in Figure 3.4b. The most evident features of the spiral vs. reference image are the radial streaks around the outer portion of the image caused by the spiral sampling of the Fourier data. However, note the distinct delineation of the central portion of the image where there is much less difference between the two images. This central portion is almost exactly 128 pixels in height and width, or it is almost exactly the central quarter of the image. With some values of β it is possible to get sharp peaks in the corners of the image. This problem could be avoided by setting the outer portion of the image to zero, keeping only the portion that is the most accurate reconstruction of the desired image (see the central portion of Figure 3.4). This also supports the use of over-sampling of the Fourier data, as has been suggested in the literature, to increase the FOV prior to this cropping (Jackson et al. 1991). The low peak difference (5 gray levels) in

the non-gridded vs. reference image indicates that most of the artifacts in the table-based gridding reconstructions are attributable to the trajectory itself, rather than to the reconstruction.

The simulated radial k-space data experiment showed similar, though much less intense, artifact patterns as the previous spiral sampling. This is not surprising since the spiral ultimately generates a sampling pattern in k-space that is very similar to the radial acquisition (Ahn et al. 1986). The salient point to note from these final two experiments is the computational speed. First, the radial trajectories have nearly three times the number of data points as the spiral trajectory used (46080 vs. 16384). Consequently, the gridding time for the radial acquisition was nearly three times as long as the gridding time for those other trajectories (36.2 ms vs. 12.8 ms). Thus, the number of computations in the table method, and hence the gridding time, is directly proportional to the number of data points. This is one reason for using gridding rather than the DFT tables, since with the DFT the reconstruction time should be roughly proportional to the square of the number of sampled data points (Proakis and Manolakis 1996).

The radial k-space acquisition in a living pig, the third gridding-table experiment, showed the utility of this method for actual image reconstruction. The radial pulse sequence used in this experiment requires 1300 msec to acquire all of the data while the gridding process, when performed on a modest PC, required 36.1-130.8 msec. This indicates that the entire gridding process could take place during the acquisition itself even using a single processor PC for the reconstruction.

If this were implemented on-line, the table-based gridding procedure could be accomplished in real time (limited by the acquisition) and the total lag time would be the

time it takes to do the 2D-FFT and display the data. This would allow a variety of pulse sequences to be used in a more real-time fashion as is important in fMRI (functional magnetic resonance imaging), perfusion, cardiac and iMRI applications where currently the selection is quite limited. In addition, this efficient method would free up computational resources for other desirable activities, especially for groups with access to specialized computational resources.

When examining all of these results it should be remembered that expertise and familiarity with a particular computer language might play a significant role in affecting the speed results of any such experiment. Undoubtedly results will vary according to the skill of the programmer. Thus, less attention should be paid to the absolute times found in the tables, but rather the relative differences between conventional and table-based methods. The reason for this is that differences in times for any task are expected with different computer systems. However, use of speed-optimized code allows comparison between relative time differences to be made. For example, if conventional regridding on a different computer requires 100msec (and 582.3 msec here for $W=4$: Table 3.2), then a table based method on this alternative hardware would likely require approximately 2.2 msec for table-based gridding since 12.8 msec was needed on our 600MHz Pentium.

The DFT has the advantage of much simpler mathematics and direct calculation of the final reconstructed image, when compared to the various gridding approaches. However, as shown here, it is computationally intensive and therefore usually considered too slow for widespread use, even when using these efficient table-based methods. On the other hand, all of the gridding approaches introduce some error during the change from non-uniformly sampled data to an estimate of a uniformly sampled Fourier data set.

The errors result from the need for interpolation, convolution, matrix inversion, or other such operations in the data estimation process (O'Sullivan 1985; Jackson et al. 1991; Rosenfeld 1998). Further, gridding methods create Fourier data, or k-space data, that must still be transformed into the reconstructed image by use of the FFT. The considerable numerical advantage of the FFT, however, can make the total gridding-reconstruction process less computationally intensive than the corresponding direct DFT reconstruction, at the expense of modest image errors and artifacts.

The speed of the reconstruction really became potentially useful in dynamic settings beginning with the gridding experiments (12.8 ms to 130.8 ms). As mentioned earlier, the primary advantage of the gridding algorithm is that most of the values in $\mathbf{Tg}(i)$ are equal to zero and, using a sparse matrix representation, they need to be neither stored nor calculated. Another advantage of the gridding table over the DFT is that the DFT table must be complex while the convolution function, and therefore the gridding table, is usually purely real (though the acquired or simulated data is still complex). This further reduces the computational burden by simplifying each complex multiply-accumulate from four multiplications and four additions in the DFT table to two multiplications and two additions in the gridding table.

The final two steps (i.e., the FFT and the deconvolution) of both the normal and the table-enhanced versions of the gridding algorithm are identical and so the anticipated savings in time are in the first step. The processor is able to prepare the data for the FFT in a minimal amount of time with the table-enhanced method by reducing all of the transcendental or other functions in the first three steps to a single multiply-accumulate. Even on a single processor this can mean that the gridding process can be done as fast as

the data is acquired. In that case, the entire reconstruction process would only be twice as long as the reconstruction of a normal acquisition in which the row FFT is evaluated for each line of data as it is acquired. Table 3.2 shows that non-uniform k-space sampling with a table-enhanced gridding algorithm permits gridding during the acquisition, leaving both the row and the column transformations for calculation upon completion of the acquisition.

These table-based methods are highly amenable to parallelization. With a sufficient number of processing elements the DFT-table methods could, theoretically, become even faster than the FFT. This is because the FFT must wait until the entire row is collected before beginning the row transformation. Then it must wait until all the rows are collected before beginning the column FFTs. The DFT, on the other hand, can process the result from each data point as it arrives from the scanner. All processing tasks could be equally distributed between up to as many different processing units as there are elements in the $\mathbf{T}(i)$ or $\mathbf{Tg}(i)$ matrices. Groups have already begun exploring multiprocessor and dedicated hardware systems for non-rectilinearly sampled k-space data (Kerr et al. 1997; Eggers and Proska 1999). If enough elements exist to process one data point before the next arrives, then the entire process is performed in real-time and limited only by the MRI data-sampling rate. The end result of using a table method in conjunction with sufficient computer power, as described above, would be the ability to reconstruct an image after the acquisition of each data point or after any arbitrary number of data points. This would even be possible for the gridding-table method, but would be somewhat more natural for the DFT-table method. Even without reconstructing a new image after every data point, the gridding process can be carried out concurrently with

the image collection and thus the 2D-FFT could be applied as often as desired to reflect the updated data. This would be important in a MR fluoroscopy application where it would be desirable to obtain updated images in sub-second increments. Certain acquisition strategies, such as spiral or rosette, could acquire sufficient information each second to warrant an updated image reconstruction several times per second; and these table-based methods could be a means to facilitate that high rate of image reconstruction, even without specialized computational equipment. In situations with specialized computational resources, this method would allow for more processing, such as motion correction, to occur on-line rather than expending all the computational resources on the gridding/reconstruction task.

The use of lookup tables to improve the computational speed of algorithms is a generally accepted and common practice. In fact, many compilers, when optimizing for speed, will automatically generate small tables in order to handle complicated branching structures such as switch statements (Booth 1997). Tables often allow much more complicated numerical or logical functions to be mapped into a simple memory access function. For example, the table saves tens or hundreds of processor clock cycles (dependent on the processor) by pre-calculating and storing the results from complicated functions, such as Bessel, exponential and other transcendental functions. In a similar fashion, the information needed to perform complicated operations, such as convolution, can be easily stored in the table for quick use at run time. The net effect is a significant reduction in the computational burden during the reconstruction. This mapping, in essence, trades bytes for flops (floating-point operations per second). This can be

advantageous since it is easier to increase the maximum number of MB of available memory than it is to increase the number of Mflops on most available computer systems.

Other methods have been proposed to reduce the number of Mflops needed prior to reconstruction for use in gridding algorithms. Some groups have used fixed-point calculations, rather than the more accurate floating-point, in order to use the faster integer portions of the processor (Liao 1999). The table lookup methods developed here for flops could also be extended to fixed point operations and thus gain all of the advantages (and disadvantages) of using integer arithmetic.

The two major limitations to the speed of a table look-up operation are the size of the table and the predictability of the access pattern. Both of these relate to the probability of a given table entry being in the data cache of the processor (Booth 1997). A strong computational time penalty is incurred using the table methods each time the algorithm must wait for data from main memory, or worse, from the hard drive. The computational advantage can therefore be lost if the tables are poorly organized. Thus the fastest tables are small and generally accessed in a straightforward manner (Booth 1997). Thus the speed advantages gained by the gridding tables seem to violate some of the known stereotypes of fast tables simply because of their sheer size. The reconstruction tables in this project are not small, but their access patterns are so predictable that they enhance the computational speed significantly, even without explicit concern over cache access. There are two principal requirements that must be satisfied to be able to use this table method. The first is that the sampled locations of k-space must be known beforehand. In MRI this is not a difficult requirement since the sample point locations are determined by the pulse sequence, which is created beforehand. The second more limiting, requirement

is that the operation replaced by the table must be a linear operation on the sampled data points. This allows a table implementation of the DFT algorithm or the commonly used convolution gridding algorithm. However, other gridding operations (e.g., interpolation) are generally not linear and therefore would not qualify for a table implementation as described here (O'Sullivan 1985; Rosenfeld 1998). Also, portions of the convolution gridding algorithm that are linear operations on the transformed data rather than on the original sampled data, such as the final deconvolution, cannot be implemented in this type of table. Fortunately, by combining all linear operations on the sampled data, operations like density compensation, convolution, and resampling can all be contained in a single table.

Future table-based methods could include other corrections or enhancements. For instance, Meyer et al. include a step for normalizing the energy in each grid point that could be included in the table with no run-time penalty (Meyer et al. 1992). Other enhancements could include corrections for eddy currents or any other influence that could be anticipated in advance and compensated for with a linear function on the original, non-uniformly sampled, k-space data.

3.5 Conclusions

In conclusion, a table-based gridding operation can be computed in less time than the acquisition, even when the computations are performed on a single processor PC. Using this method, the rate-limiting step in generation of the MR image for many applications is the MRI pulse sequence itself since the table based method proposed here can be run coincident with data collection. When implemented on-line, this method would allow the total reconstruction time to equal the acquisition time plus the time required for the 2D-

FFT. Further improvements in reconstruction time can be achieved with this method using specially designed hardware or multi-processor systems. The current implementation of the table-based method allows for larger convolution kernels than are usually used with other methods. This in turn leads to reduced sidelobes and therefore less aliasing energy in the image, all within a computation time (when performed on a 600MHz Pentium III processor with 512 MB RAM) less than many MRI acquisitions.

The lookup tables require a substantial amount of memory, but the predictable access into the table allows for rapid memory access despite their size. The inclusion of the first three steps of the typical gridding algorithm in a single step and the transformation of all complicated functions into simple MACs allow for rapid computation at run time. Together the quick access and the easy computation lead to a very rapid execution even with a single processor PC system. Look-up table methods for reconstruction should be considered as an alternative method for gridding and reconstruction in non-uniformly sampled k-space MRI acquisitions.

Chapter 4. Time-optimal K-space Trajectories

4.1 Introduction

In many MRI pulse sequences, a relatively large amount of time is spent neither transmitting nor receiving RF energy. This time is typically used to either allow magnetization relaxation to generate a desired contrast between tissues or to allow the gradients to transition to a desired gradient level and/or k -space location. Often this gradient waveform control results in dead time within a pulse sequence that should be minimized. For example, the time required to move from the end of one line in an EPI (echo planar imaging) sequence to the beginning of the next contributes to $T2^*$ decay related loss in resolution and off resonance artifacts; reducing this time would result in less sensitivity to these effects (Butts et al. 1997; Jesmanowicz et al. 1998; Reber et al. 1998). Unnecessary dead time can also result in a larger TR that can increase the severity of the banding artifact in true-FISP sequences (Duerk et al. 1998; Larson et al. 2002; Dale et al. 2003a). Minimization of sequence dead time is an important concern for steady-state EPI sequences such as the recently proposed SPIDER (steady-state projection imaging with dynamic echo-train readout) technique (Larson and Simonetti 2001).

A variety of methods have been introduced in order to minimize this dead time. Some of the most effective are the HOT (hardware optimized trapezoid) waveforms by Atalar and McVeigh and the corresponding higher order moment extension by Bolster and Atalar (Atalar and McVeigh 1994; Bolster, Jr. and Atalar 1999). One significant limitation of the HOT approach is that the rotation between the logical and hardware gradient axes (i.e., the amount of obliquity of the slice plane and in-plane rotation) must

be known at the time of the design of the gradient waveforms. A further result of this limitation is that the total duration of a waveform generated with the HOT method depends on the rotation in any sequence and on the azimuthal angle of each view in any radial acquisition.

It is desirable, in many circumstances, to use gradient waveforms that will not violate hardware limitations regardless of the rotation between the logical and hardware gradient axes. This allows design of the gradient waveforms prior to run-time and maintains timing consistency between acquisitions with different rotations or between different views in a radial sequence. This is often done by employing conservative methods such as reducing the maximum gradient strength by a factor of $\sqrt{2}$ or $\sqrt{3}$ depending on the number of active gradient axes (Atalar and McVeigh 1994). While this practice is common, it does not fully utilize the capability of the gradient hardware resulting in sub-optimal gradient waveform durations.

This work overcomes these problems by using modern control theory and the calculus of variations to derive the mathematical form of the time-optimal gradient waveforms for several possible design strategies. One of these strategies has both the property of being independent of the rotation and the property of fully utilizing the slew-rate capability of the gradient hardware. Each derived waveform is shown to be the minimum-time solution for the given design strategy and the differences between the various strategies are demonstrated through several examples.

4.2 Theory

The formulation of an optimal-control problem using the calculus of variations involves defining the state variables of the system, choosing the appropriate control

region and other constraints, and deriving the corresponding Hamiltonian and co-state equations for use with the PMP (Pontryagin minimum principle). The resulting systems of differential equations are then solved with the initial and final conditions to obtain a specific time-optimal gradient waveform. Each of these steps is described below. Together, a time-optimal control problem for design of MR gradient waveforms can be stated as: “Given the gradient amplitude and slew-rate constraints, how can the gradient system be controlled from its initial state to a desired final state in the shortest possible time?”

4.2.1 State-space representation

The solution of a control problem begins with a state-space representation of the system. For the MRI gradient system, the state equations are:

$$\begin{aligned} \mathbf{k}'(t) &= \gamma \mathbf{g}(t) \\ \mathbf{g}'(t) &= \mathbf{s}(t) \end{aligned} \tag{4.1}$$

where γ is the gyromagnetic constant, $\mathbf{k}(t)$ is the k -space location, $\mathbf{g}(t)$ is the gradient waveform, and $\mathbf{s}(t)$ is the slew rate. In this formulation, the state variables are $\mathbf{k}(t)$ and $\mathbf{g}(t)$, the control variable is $\mathbf{s}(t)$, and the system is linear time-invariant. This system is completely characterized for any time $0 < t < t_f$ given an initial state $\mathbf{k}(0) = \mathbf{k}_0$ and $\mathbf{g}(0) = \mathbf{g}_0$ and a control $\mathbf{s}(t)$ which is defined on all open subintervals of $0 < t < t_f$.

4.2.2 Control regions and admissible controls

Due to the hardware limitations described above, not all choices are reasonable for $\mathbf{s}(t)$. In particular, the slew-rate limitation can be expressed as a restriction on the range of $\mathbf{s}(t)$ and the gradient magnitude limitation can be expressed as a restriction on the




Figure 4.1 Three possible control regions for a 2D slew space. The vertical/horizontal axes represent the hardware coordinate system while the offset axes represent the logical coordinate system. The outer square is the entire set of slew rates that can be achieved by the hardware. The circle and inner squares are two sets of slew rates that can be achieved by the hardware regardless of the angle between the logical and hardware coordinate systems. The angle between these two sets of axes is not known until run-time due to the potential for oblique acquisitions and/or in-plane rotation.

range of $g(t)$. In the 1D case, the control region is straightforward: $|s(t)| \leq s_m$ where s_m is the maximum slew rate that can be achieved by the amplifier.

The 2D case is not as simple as the 1D case. Figure 4.1 shows a number of possible control regions. The outer square represents the total range of slew rates that could be achieved by two separate gradient channels. Expressed mathematically, $|s_i(t)| \leq s_m \quad i \in \{x, y\}$ where $\{x, y, z\}$ represent the hardware axes. The circular control

region can be expressed $\|s(t)\| \leq s_m$. This represents the largest range of slew rates that are guaranteed to not exceed the slew-rate constraint on any gradient axis regardless of the rotation between the logical and hardware axes. The inner square can be expressed $|s_i(t)| \leq \frac{s_m}{\sqrt{2}} \quad i \in \{r, p\}$ where $\{r, p, s\}$ represent the logical readout, phase-encode, and slice-select axes. This region, like the circular control region, is guaranteed to not exceed the slew-rate constraint regardless of the rotation, and it represents the largest such region where each logical gradient axis is independent of the other axes.

When a square (or rectangular) control region is chosen, the problem can be simplified into two independent 1D problems rather than a single 2D problem because no choice of $s_x(t)$ can have any influence on available choices for $s_y(t)$ or on the resulting $k_y(t)$ and $g_y(t)$. On the other hand, for the circular control region, a particular choice of $s_x(t)$ defines the limits for $s_y(t)$ and thereby indirectly affects the resulting $k_y(t)$ and $g_y(t)$ thus requiring the solution of a single 2D problem.

Let S denote the chosen control region (i.e., inner circle, inner square, etc.). A corresponding region can also be chosen for the gradient amplitude constraints and will be denoted G . A control function $s(t)$, defined on $0 < t < t_f$, that satisfies the control constraints $s(t) \in S$, and results in state functions $k(t)$ and $g(t)$ such that $g(t) \in G$, is called an admissible control. The corresponding pair of state functions $k(t)$ and $g(t)$ is called an admissible trajectory.

The time-optimal control is thus an admissible control $s^*(t)$ that transfers the system from the given initial state, $k^*(0) = k_0$ and $g^*(0) = g_0$, to the desired final state $k^*(t_f^*) = k_f$

and $g^*(t_f^*) = g_f$, and minimizes t_f^* . In other words, $t_f^* \leq t_f$ for any other admissible control $s(t)$ that transfers the system from $k(0) = k_0$ and $g(0) = g_0$ to $k(t_f) = k_f$ and $g(t_f) = g_f$. Thus, the minimal time is t_f^* , the optimal control is $s^*(t)$ and the optimal trajectory is $k^*(t)$ and $g^*(t)$.

4.2.3 Pontryagin Minimum Principle

In order to solve this optimization problem it is necessary to use the calculus of variations to examine the change in t_f due to infinitesimal variations in the function $s(t)$. The functionals used in variational problems are of the form (Kirk 1970; Hocking 1991)

$$F(f) = \int_0^{t_f} f(k, g, s, t) dt \quad [4.3]$$

For time optimality, we will let

$$f(k, g, s, t) = 1 \quad [4.4]$$

so that

$$F(f) = t_f$$

The PMP states that a necessary condition for a control function to minimize a cost functional is that the control function minimizes the Hamiltonian (considered as a function of the control variables) at each point in time. In order to obtain the Hamiltonian for a given system it is necessary to introduce the co-state equations. The Hamiltonian, the state equations, and the co-state equations are related as follows:

$$\begin{aligned} \mathbf{x}' &= \frac{\partial H}{\partial \mathbf{z}} \\ \mathbf{z}' &= -\frac{\partial H}{\partial \mathbf{x}} \end{aligned} \quad [4.5]$$

$$H = f + \mathbf{x}' \cdot \mathbf{z}$$

where \mathbf{x} is the vector of state variables, \mathbf{z} is the vector of co-state variables, H is the Hamiltonian and f is the integrand of the cost functional to be minimized. Differentiation of a scalar function with respect to a vector here refers to the vector formed by taking the partial derivative of the scalar function with respect to each vector component.

In variational problems, the optimal control at each point in time is either a local minimum of H or it is on the border of the admissible control region (Kirk 1970; Hocking 1991). Therefore, the solution to this optimization problem will be different for every different choice of control region. As mentioned previously, problems using the square control regions of Figure 4.1 can be separated into two 1D control problems. Therefore we will examine only 1D optimization problems and the 2D circular control region optimization problem.

4.2.4 Time-optimality for 1D problems

The slew-rate constraint for a 1D time-optimal gradient control problem is:

$$|\dot{s}(t)| \leq s_m$$

the Hamiltonian is:

$$H = 1 + \gamma z_1 g + z_2 s \quad [4.6]$$

and the co-state equations are:

$$\begin{aligned} z_1'(t) &= 0 \\ z_2'(t) &= -\gamma z_1(t) \end{aligned} \quad [4.7]$$

which have the solution:

$$\begin{aligned} z_1(t) &= -\frac{c_1}{\gamma} \\ z_2(t) &= c_2 + c_1 t \end{aligned} \quad [4.8]$$

The Hamiltonian is the equation of a straight line in the (s, H) plane with slope z_2 . Straight lines have no local minima and therefore can only be minimized by values on the boundary of the control region. Note that if $z_2 = 0$ there is no minimum even at the boundary values. This is not a difficulty for this problem since $z_2(t) = c_2 + c_1 t = 0$ only for one point in time. The value of the control at a single point cannot alter the state-space trajectory, and therefore we can choose $s(t) = 0$ at $z_2(t) = 0$ in order to have an easily defined control at each point in time:

$$s^*(t) = -s_m \operatorname{sgn}(c_2 + c_1 t) \quad [4.9]$$

This result leads to two possible forms for the optimal control.

$$s^*(t) = \begin{cases} s_m, & 0 < t < t_1 \\ 0, & t = t_1 \\ -s_m, & t_1 < t < t_f \end{cases} \quad \text{or} \quad s^*(t) = \begin{cases} -s_m, & 0 < t < t_1 \\ 0, & t = t_1 \\ s_m, & t_1 < t < t_f \end{cases} \quad [4.10]$$

where the switching time

$$t_1 = -c_2 / c_1$$

Controls of this form are known as bang-bang controls and represent a maximal control effort throughout the transfer. Note that there is, at most, one switch between control extremes.

Using the two possible forms for the optimal control, along with the initial conditions, the final conditions and the condition of continuity at the switching time we obtain the following expressions for the switching time, the corresponding k -space location and gradient amplitude, and the final time:

$$\begin{aligned}
t_1 &= \frac{-2g_0 + \sqrt{\frac{2\gamma(g_0^2 + g_f^2) + 4s_m(k_f - k_0)}{\gamma}}}{2s_m} && \text{or} \\
t_1 &= \frac{2g_0 + \sqrt{\frac{2\gamma(g_0^2 + g_f^2) - 4s_m(k_f - k_0)}{\gamma}}}{2s_m} \\
k_1 &= \frac{\gamma(g_f^2 - g_0^2) + 2s_m(k_0 + k_f)}{4s_m} \text{ or } k_1 = \frac{\gamma(g_0^2 - g_f^2) + 2s_m(k_0 + k_f)}{4s_m} \\
g_1 &= \sqrt{\frac{g_0^2 + g_f^2}{2} + \frac{s_m(k_f - k_0)}{\gamma}} \text{ or } g_1 = -\sqrt{\frac{g_0^2 + g_f^2}{2} - \frac{s_m(k_f - k_0)}{\gamma}} && [4.10] \\
t_1 &= \frac{g_0 + g_f - \sqrt{\frac{2\gamma(g_0^2 + g_f^2) + 4s_m(k_f - k_0)}{\gamma}}}{s_m} && \text{or} \\
t_1 &= \frac{g_0 + g_f + \sqrt{\frac{2\gamma(g_0^2 + g_f^2) - 4s_m(k_f - k_0)}{\gamma}}}{s_m}
\end{aligned}$$

The addition of the gradient amplitude constraints is fairly intuitive in the 1D case.

This is done by making the control space depend on the state of the system as follows:

$$\begin{aligned}
&|s(t)| \leq s_m, \quad |g(t)| < g_m \\
&-s_m \leq s(t) \leq 0, \quad g(t) = g_m \\
&0 \leq s(t) \leq s_m, \quad g(t) = -g_m
\end{aligned} \tag{4.11}$$

The Hamiltonian remains unchanged, as do the co-state equations and their solutions.

The only change is the value of the minimizing slew-rate when $g(t) = \pm g_m$. Thus, if $g^*(t)$ never reaches $\pm g_m$, then the optimal trajectory with the constraint is the same as the optimal trajectory without it. If $g^*(t)$ does reach $\pm g_m$, then the control region is halved

so the optimal control switches to a slew-rate of 0 until the time $t_2 = -c_2 / c_1$ is reached.

This leads to two more possible forms for the optimal control:

$$s^*(t) = \begin{cases} s_m, & 0 \leq t < t_1 \\ 0, & t_1 \leq t \leq t_2 \\ -s_m, & t_2 < t \leq t_f \end{cases} \quad \text{or} \quad s^*(t) = \begin{cases} -s_m, & 0 \leq t < t_1 \\ 0, & t_1 \leq t \leq t_2 \\ s_m, & t_2 < t \leq t_f \end{cases} \quad [4.12]$$

Using the two possible forms for the optimal control, along with the initial conditions, the final conditions and the condition of continuity at the switching time we obtain the following expressions for the switching time, the corresponding k -space location and gradient amplitude, and the final time:

Again, the two possible forms for the optimal control are used, along with the initial conditions, the final conditions and the condition of continuity at the switching time to obtain the following expressions for the two switching times, the corresponding k -space locations (the corresponding gradients are equal to $\pm g_m$), and the final time:

$$\begin{aligned} t_1 &= \frac{g_m - g_0}{s_m} \quad \text{or} \quad t_1 = \frac{g_m + g_0}{s_m} \\ k_1 &= \frac{\gamma(g_m^2 - g_0^2)}{2s_m} + k_0 \quad \text{or} \quad k_1 = \frac{\gamma(g_0^2 - g_m^2)}{2s_m} + k_0 \\ t_2 &= \frac{\gamma(g_0^2 - 2g_m g_0 + g_f^2) - 2k_0 s_m + 2k_f s_m}{2\gamma g_m s_m} \quad \text{or} \\ t_2 &= \frac{\gamma(g_0^2 + 2g_m g_0 + g_f^2) + 2k_0 s_m - 2k_f s_m}{2\gamma g_m s_m} \quad [4.13] \\ k_2 &= \frac{\gamma(g_f^2 - g_m^2)}{2s_m} + k_f \quad \text{or} \quad k_2 = \frac{\gamma(g_m^2 - g_f^2)}{2s_m} + k_f \end{aligned}$$

$$t_f = \frac{g_m - g_f}{s_m} + \frac{g_0^2 - 2g_m g_0 + g_f^2}{2g_m s_m} + \frac{k_f - k_0}{\gamma g_m} \quad \text{or}$$

$$t_f = \frac{g_m + g_f}{s_m} + \frac{g_0^2 + 2g_m g_0 + g_f^2}{2g_m s_m} + \frac{k_0 - k_f}{\gamma g_m}$$

These results still represent a maximal effort throughout the transfer, but that effort is limited by different considerations during different portions of the transfer.

4.2.5 Time-optimality for 2D and 3D magnitude slew-rate constraints

The magnitude slew-rate constraint, $\|\mathbf{s}(t)\| \leq s_m$, results in the circular control region of Figure 4.1 for the 2D case or a similar spherical control region for the 3D (three dimensional) case. As mentioned previously, this control region results in a ND problem that cannot be separated into a number of 1D problems. In vector form, the state equations for the ND gradient control problem are:

$$\mathbf{k}'(t) = \gamma \mathbf{g}(t)$$

$$\mathbf{g}'(t) = \mathbf{s}(t)$$

and the Hamiltonian is:

$$H = \gamma \mathbf{z}_1(t) \cdot \mathbf{g}(t) + \mathbf{z}_2(t) \cdot \mathbf{s}(t) + 1 \quad [4.18]$$

The co-state equations are:

$$\mathbf{z}_1'(t) = \mathbf{0}$$

$$\mathbf{z}_2'(t) = -\gamma \mathbf{z}_1(t) \quad [4.19]$$

which have the solution:

$$\mathbf{z}_1(t) = -\frac{\mathbf{c}_1}{\gamma}$$

$$\mathbf{z}_2(t) = t \mathbf{c}_1 + \mathbf{c}_0 \quad [4.20]$$

The Hamiltonian (eq. [4.18]) is the equation of a plane in the (\mathbf{s}, H) space. As before, this has no local minima and can therefore only be minimized by values on the boundary of the control region where $\|\mathbf{s}\| = s_m$. In particular, the control that minimizes H is the control that minimizes $\mathbf{z}_2 \cdot \mathbf{s}$, therefore, the optimal control vector $\mathbf{s}^* = (s_x^*, s_y^*)$ can be expressed:

$$\mathbf{s}^* = -s_m \frac{\mathbf{z}_2}{\|\mathbf{z}_2\|} = -\frac{(\mathbf{c}_1 t + \mathbf{c}_0) s_m}{\|\mathbf{c}_1 t + \mathbf{c}_0\|} \quad [4.21]$$

This expression for the optimal control can be substituted into the state equations along with the above expressions for \mathbf{z}_2 . The resulting form of the state equations has the following analytical solution, with the constants of integration expressed in terms of the initial conditions:

$$\begin{aligned} \mathbf{k}^*(t) = & (\gamma \mathbf{g}_0 t + \mathbf{k}_0) + \\ & \left(\frac{\gamma t \mathbf{c}_0 s_m}{\|\mathbf{c}_1\|} - \frac{3\gamma (\mathbf{c}_1 \cdot \mathbf{c}_0)^2 \mathbf{c}_1 s_m}{2\|\mathbf{c}_1\|^5} + \frac{\gamma (\|\mathbf{c}_0\|^2 \mathbf{c}_1 + 2\mathbf{c}_1 \cdot \mathbf{c}_0 (\mathbf{c}_0 - \mathbf{c}_1 t)) s_m}{2\|\mathbf{c}_1\|^3} \right) \\ & \left(\ln \left(\|\mathbf{c}_0\| + \frac{\mathbf{c}_1 \cdot \mathbf{c}_0}{\|\mathbf{c}_1\|} \right) - \ln \left(\|\mathbf{c}_1 t + \mathbf{c}_0\| + \|\mathbf{c}_1\| t + \frac{\mathbf{c}_1 \cdot \mathbf{c}_0}{\|\mathbf{c}_1\|} \right) \right) + \\ & \frac{\gamma s_m}{2\|\mathbf{c}_1\|^2} (-2\|\mathbf{c}_0\|(\mathbf{c}_0 - \mathbf{c}_1 t) + \|\mathbf{c}_1 t + \mathbf{c}_0\|(2\mathbf{c}_0 - \mathbf{c}_1 t)) + \\ & \frac{3\gamma \mathbf{c}_1 \cdot \mathbf{c}_0 \mathbf{c}_1 s_m}{2\|\mathbf{c}_1\|^4} (\|\mathbf{c}_0\| - \|\mathbf{c}_1 t + \mathbf{c}_0\|) \end{aligned}$$

$$\mathbf{g}^*(t) = \mathbf{g}_0 + \frac{\|\mathbf{c}_0\| \mathbf{c}_1 s_m}{\|\mathbf{c}_1\|^2} - \frac{\|\mathbf{c}_1 t + \mathbf{c}_0\| \mathbf{c}_1 s_m}{\|\mathbf{c}_1\|^2} - \frac{\ln\left(2\left(\|\mathbf{c}_0\| + \frac{\mathbf{c}_1 \cdot \mathbf{c}_0}{\|\mathbf{c}_1\|}\right)\right) \left((\mathbf{c}_1 \cdot \mathbf{c}_0) \mathbf{c}_1 - \|\mathbf{c}_1\|^2 \mathbf{c}_0\right) s_m}{\|\mathbf{c}_1\|^3} + \frac{\ln\left(2\left(\|\mathbf{c}_1 t + \mathbf{c}_0\| + \|\mathbf{c}_1\| t + \frac{\mathbf{c}_1 \cdot \mathbf{c}_0}{\|\mathbf{c}_1\|}\right)\right) \left((\mathbf{c}_1 \cdot \mathbf{c}_0) \mathbf{c}_1 - \|\mathbf{c}_1\|^2 \mathbf{c}_0\right) s_m}{\|\mathbf{c}_1\|^3}$$

Substituting these expressions into the expression for the Hamiltonian and simplifying gives the following constant expression for the Hamiltonian:

$$H = 1 - \mathbf{c}_1 \cdot \left(\mathbf{g}_0 + \frac{\|\mathbf{c}_0\| \mathbf{c}_1 s_m}{\|\mathbf{c}_1\|^2} - \frac{\ln\left(2\left(\|\mathbf{c}_0\| + \frac{\mathbf{c}_1 \cdot \mathbf{c}_0}{\|\mathbf{c}_1\|}\right)\right) \left((\mathbf{c}_1 \cdot \mathbf{c}_0) \mathbf{c}_1 - \|\mathbf{c}_1\|^2 \mathbf{c}_0\right) s_m}{\|\mathbf{c}_1\|^3} \right)$$

Using the final conditions in the solution to the state equations and substituting the above expressions results in a system of 2 vector equations in 2 vector unknowns and one scalar unknown. The final scalar equation is obtained by setting the Hamiltonian equal to zero as indicated by the PMP. This gives the following system of 3 equations in 3 unknowns.

$$\mathbf{k}^*(t_f) = \mathbf{k}_f$$

$$\mathbf{g}^*(t_f) = \mathbf{g}_f$$

$$1 = \mathbf{c}_1 \cdot \left(\mathbf{g}_0 + \frac{\|\mathbf{c}_0\| \mathbf{c}_1 s_m}{\|\mathbf{c}_1\|^2} - \frac{\ln\left(2\left(\|\mathbf{c}_0\| + \frac{\mathbf{c}_1 \cdot \mathbf{c}_0}{\|\mathbf{c}_1\|}\right)\right) \left((\mathbf{c}_1 \cdot \mathbf{c}_0) \mathbf{c}_1 - \|\mathbf{c}_1\|^2 \mathbf{c}_0\right) s_m}{\|\mathbf{c}_1\|^3} \right) \quad [4.22]$$

This system of equations cannot be solved analytically for the remaining unknowns, but can be solved numerically when given specific values for the initial conditions, the final conditions, and s_m . The resulting values for the vectors c_g and c_s , and the scalar t_f^* completely characterize the gradient waveforms and k -space trajectory which transfer the system from the initial state to the final state in the minimum amount of time possible for a circular control region.

4.3 Methods

All constraints and waveforms are based on a system with a 200 T/m/s maximum slew rate and a 40 mT/m maximum gradient strength. In addition, there is an assumed 30° in-plane rotation between the logical and hardware gradient axes. Each section below presents the optimal waveforms for a particular transfer for each of the three control regions in Figure 4.1. All initial and final conditions are presented in the logical reference frame.

4.3.1 Time Optimal EPI reversal

Transfer from the end of the first line to the beginning of the second in an EPI sequence is defined as follows:

$$\mathbf{k}_0 = \begin{pmatrix} 500 \\ 500 \end{pmatrix} \quad \mathbf{k}_f = \begin{pmatrix} 500 \\ 450 \end{pmatrix}$$

$$\mathbf{g}_0 = \begin{pmatrix} .020 \\ 0 \end{pmatrix} \quad \mathbf{g}_f = \begin{pmatrix} -.020 \\ 0 \end{pmatrix}$$

where the readout gradient strength is 20 mT/m, the maximum k -space location is 500 m^{-1} and the separation between lines is 50 m^{-1} .

The time-optimal transfer for the control region corresponding to the rotated square of Figure 4.1 is calculated by evaluating equation [4.10] for the above initial and final

conditions with slew and gradient constraints of $200/\sqrt{2}$ T/m/s and $40/\sqrt{2}$ mT/m on each axis. The time-optimal transfer for the circular control region is calculated by numerically solving equation [4.22] after substituting in the initial and final conditions and the magnitude slew constraint of 200 T/m/s. The resulting control parameters are then fed back into the state equations to generate the optimal waveforms. The time-optimal transfer for the outer square control region is found by first translating the initial and final conditions into the hardware coordinate system. Equation [4.10] is then solved with the full hardware constraints of 200 T/m/s and 40 mT/m on each axis. The resulting waveforms are translated back into the logical coordinate system.

4.3.2 Time Optimal radial EPI reversal

Transfer from the end of the first line to the beginning of the second in a rEPI (radial echo planar imaging) sequence is defined as follows:

$$\begin{aligned} \mathbf{k}_0 &= 500 \begin{pmatrix} \cos(0^\circ) \\ \sin(0^\circ) \end{pmatrix} & \mathbf{k}_f &= 500 \begin{pmatrix} \cos(10^\circ) \\ \sin(10^\circ) \end{pmatrix} \\ \mathbf{g}_0 &= .020 \begin{pmatrix} \cos(0^\circ) \\ \sin(0^\circ) \end{pmatrix} & \mathbf{g}_f &= -.020 \begin{pmatrix} \cos(10^\circ) \\ \sin(10^\circ) \end{pmatrix} \end{aligned}$$

where the readout gradient strength is 20 mT/m, the maximum k -space radius is 500 m^{-1} and the angular separation between lines is 10° .

The three time-optimal transfers corresponding to the three different control regions are calculated as described above.

4.3.3 Time Optimal read-dephasing, phase-encoding, and slice-refocusing

Transfer from the end of the slice-selection pulse to the beginning of the first readout in a sequentially ordered FLASH sequence is defined as follows:

$$\mathbf{k}_0 = \begin{pmatrix} 0 \\ 0 \\ 255 \end{pmatrix} \quad \mathbf{k}_f = \begin{pmatrix} -500 \\ 500 \\ 0 \end{pmatrix}$$

$$\mathbf{g}_0 = \begin{pmatrix} 0 \\ 0 \\ .012 \end{pmatrix} \quad \mathbf{g}_f = \begin{pmatrix} .020 \\ 0 \\ 0 \end{pmatrix}$$

where the readout gradient strength is 20 mT/m, the maximum k -space location is 500 m^{-1} , the slice-select gradient is 12 mT/m, and the slice-select duration is 1 ms (corresponding to a 10 mm slice thickness with a 5.11 kHz transmit BW).

The three time-optimal transfers corresponding to the three different control regions are calculated as above except for the changes needed to handle the 3D transfer of this problem. Specifically, the hardware constraints are $200/\sqrt{3}$ T/m/s and $40/\sqrt{3}$ mT/m on each axis for the small cube. Also, some of the waveforms reached the gradient amplitude limit and therefore required use of equation [4.13] instead of equation [4.10].

4.4 Results

All of the figures below represent gradient waveforms that are time-optimal with respect to one of the control regions. In the case of the outer square the waveforms are designed in the hardware reference frame and then rotated and plotted in the logical reference frame. Slew and k -space waveforms are not plotted as they contain little or no additional information.

4.4.1 Time Optimal EPI reversal

The time-optimal EPI reversal gradient waveforms for each of the three control regions are shown in Figure 4.2. The minimum times are 283 μs , 221 μs , and 204 μs respectively for the rotated square, circular, and large square control regions. For the rotated square control region the minimum time is limited by the 283 μs required to

transfer the readout gradient. Note that the speed of the transfer is roughly proportional to the area of the control region with the fastest transfer occurring for the largest control region.

The switching time for the waveform in Figure 4.2a is 141 μs while the switching times for the waveform in Figure 4.2c are 39 μs and 189 μs . The time-optimal parameters for Figure 4.2b are $c_1 = 0$, $c_2 = .00379$, $c_3 = -29.5$, and $c_4 = .00326$. These parameters result in the time-optimal waveforms $g_x(t) = 25.6 \sinh^{-1}(862 - .00780t)$ and




Figure 4.2 Time-optimal EPI reversal gradient waveforms. Minimum time of 283 μs for the small square control region (left). Minimum time of 221 μs for the circular control region (center). Minimum time of 204 μs for the large square control region (right).

$$g_y(t) = -33.8 + \sqrt{1150 - 8.84t + .0400t^2} \quad \text{with } g \text{ in mT/m and } t \text{ in } \mu\text{s}.$$

For the small square control region there are an infinite number of waveforms that accomplish the required phase-encoding transfer in 283 μs , all of which are time-optimal in this case. The displayed phase-encoding waveform was generated by reducing the slew-rate to 58.7 T/m/s in the phase-encode direction. Other potential waveforms include those with the full slew-rate and an appropriate leading and/or lagging delay period. A

similar situation occurs for the large square control region. In this case the displayed waveforms were generated by reducing the slew-rate to 158.1 T/m/s in the direction of the hardware y-axis.

4.4.2 Time Optimal radial EPI reversal

The time-optimal rEPI reversal gradient waveforms for each of the three control regions are shown in Figure 4.3. The minimum times are 292 μs , 264 μs , and 231 μs respectively for the inner square, circular, and large square control regions. To generate the displayed waveforms the slew-rate was reduced to 120.8 T/m/s in the phase-encode




Figure 4.3 Time-optimal rEPI reversal gradient waveforms. Minimum time of 292 μs for the inner square control region (left). Minimum time of 264 μs for the circular control region (center). Minimum time of 231 μs for the outer square control region (right).

direction and 194.0 T/m/s in the y-axis direction as described above. Note that the fastest transfer corresponds to the largest control region.

The switching times for the inner square control region are 131 μs and 286 μs while the switching times for the large square control region are 25 μs and 159 μs . The optimal




Figure 4.4 Time-optimal slice-refocusing, readout-dephasing and phase-encoding gradient waveforms. Minimum time of 957 μs for the small square control region (left). Minimum time of 702 μs for the circular control region (center). Minimum time of 657 μs for the large square control region (right).

parameters for the circular control region are $c_1 = -3.05$, $c_2 = .00301$, $c_3 = 34.9$, and $c_4 = -.00437$. These parameters result in the optimal waveforms

$$g_x(t) = -2.49 + \sqrt{6.98 - .0801t + .000304t^2} + 14.9 \sinh^{-1}(1.76 - .0134t) \quad \text{and}$$

$$g_y(t) = 28.5 - \sqrt{912 - 10.5t + .0397t^2} + 1.30 \sinh^{-1}(1.76 - .0134t)$$

4.4.3 Time Optimal slice-refocusing and phase-encoding

The time-optimal slice-refocusing, readout-dephasing and phase-encoding gradient waveforms for each of the three control regions are shown in Figure 4.4. The minimum times are 957 μs , 702 μs , and 657 μs respectively for the rotated square, circular, and large square control regions. To generate the displayed waveforms the slew-rate was reduced to 51.5 T/m/s in the phase-encode direction, 54.2 T/m/s in the slice-select direction, 99.5 T/m/s in the x-axis direction, and 95.5 T/m/s in the z-axis direction as described above.

The switching times for Figure 4.4a are 200 μs , 448 μs , 508 μs , 583 μs , and 589 μs . For Figure 4.4c the switching times are 144 μs , 200 μs , 242 μs , 391 μs , and 407 μs . For

the spherical control region in Figure 4.4b the optimal parameters are $c_1 = -13.0$, $c_2 = .00409$, $c_3 = 7.96$, $c_4 = -.00263$, $c_5 = -6.75$, and $c_6 = .00236$. These parameters result in the time-optimal waveforms $g_r(t) = -44.9 + \sqrt{2560 - 15.8t + .0243t^2} - 1.46 \sinh^{-1}(25.0 - .0770t)$

$$g_p(t) = 33.5 + \sqrt{960 - 5.92t + .00913t^2} - .643 \sinh^{-1}(25.0 - .0770t)$$

$$g_s(t) = -22.3 + \sqrt{690 - 4.25t + .00656t^2} + 2.05 \sinh^{-1}(25.0 - .0770t).$$

4.5 Discussion

Some general trends are evident in the results presented above. Specifically, the fastest transfer will always be accomplished using the largest control region. This can be seen empirically in the transfer times presented above and demonstrated mathematically by noting that a larger control region will result in a larger set of admissible controls, and therefore any optimal control for the small region is an admissible control for the large region. The optimization procedure selects the control that minimizes the transfer time from the set of all admissible controls. Also, the time-optimal controls always represent a maximal effort throughout the transfer. This makes intuitive sense and can be verified by noting that the time-optimal control is always on the border of the control region due to the linearity of the Hamiltonian. These results indicate that the choice of the control region is one of the most important steps in the optimization problem.

In the slice-rephasing, phase-encoding, and readout-dephasing example the gradient waveforms for the spherical control region violated the gradient-magnitude constraint because the only constraint used was the magnitude of the slew-rate. At the present time there is no good method for incorporating these constraints. The variational techniques that exist to handle similar situations all result in a system of differential equations with




Figure 4.5 Control region for magnitude slew-rate constraints when the gradient is at the gradient-magnitude limit. The control which minimizes the Hamiltonian, \mathbf{s}^* , depends on the gradient vector, \mathbf{g} , and the co-state vector, \mathbf{z} . The \mathbf{g} vector determines the semi-circular region of admissible controls while \mathbf{z} determines the optimal vector from within that region. The optimal control, \mathbf{s}^* , is the vector which minimizes $\mathbf{z} \cdot \mathbf{s}^*$. This is the vector directly opposite \mathbf{z} if $\mathbf{z} \cdot \mathbf{g} < 0$, for which $\mathbf{z} \cdot \mathbf{s}^* = -s_m |\mathbf{z}|$. Otherwise \mathbf{s}^* is the vector to the corner of the semi-circular control region most nearly opposite \mathbf{z} as shown, for which $-s_m |\mathbf{z}| < \mathbf{z} \cdot \mathbf{s}^* < 0$.

half of the boundary conditions specified at the initial time and the same half of the boundary conditions specified at an unknown final time. This system of differential equations must be solved numerically; however this is quite difficult due to the poor specification of the boundary conditions and the fact that the final time is unknown. With a sufficiently robust method it would be possible to solve this problem for a fixed final time and iteratively reduce the time until the method failed to converge to an admissible solution. Unfortunately, current numerical methods can fail to converge for a variety of reasons other than the non-existence of an admissible solution such as numerical instability, stiff problems, etc.

Despite the problem of incorporating gradient-magnitude constraints into the circular control region problem, certain characteristics of the solution can be ascertained by

examining the conditions that lead to minimization of the Hamiltonian when gradient constraints are included. If we follow the same line of reasoning that implemented the gradient constraints for the 1D problem we can express the admissible control region as a function of the gradient variables. In particular, if $\|\mathbf{g}\| < g_m$, then the control region is the circular control region of Figure 4.1. If $\|\mathbf{g}\| = g_m$, then the control region is the semi-circular region of the slew-space shown in Figure 4.5. This control region will allow the gradient vector to move around the edge of the gradient-magnitude-constraint region or to move back into the interior of the region, but will not permit it to move outside of the constraint region.

In this formulation none of the state or co-state equations are modified, nor is the Hamiltonian. The only modification is in the determination of the optimal control that minimizes the Hamiltonian. If $\mathbf{z} \cdot \mathbf{g} < 0$, then the Hamiltonian is minimized by one of the two corners of the semi-circle. Otherwise the Hamiltonian is minimized by the point on the arc of the control region that is in the direction opposite \mathbf{z} . Additionally, the solutions to the co-state equations indicate that the direction of \mathbf{z} changes smoothly. This means that the optimal gradient waveform must leave the $\|\mathbf{g}\| = g_m$ region tangentially even if it does not enter it tangentially. Finally, the paths from the initial point to the $\|\mathbf{g}\| = g_m$ entry point and from the $\|\mathbf{g}\| = g_m$ exit point to the final point will each be the time-optimal paths with the same control parameters as found above.

Recall that both square control regions in Figure 4.1 could be separated into two 1D problems. When the control problems for the two axes are separated, it is unlikely that they will both result in the same t_f^* . The larger of the separate t_f^* values is the minimum time for the overall transfer. On the other axis, any control that reaches the

final state at the proper time is optimal and therefore the time-optimal solution for a 2D problem with a square control region is not generally unique.

The following rules are proposed based on the results above. The fastest transfers will always be obtained by doing the gradient waveform design in the hardware reference frame after the initial and final conditions have been transformed from the logical reference frame. If it is not possible to know the rotation between the logical and hardware axes at design time, or if it is necessary to maintain timing consistency across acquisitions which may have different rotations, then the fastest transfers will be obtained using the circular control region. The only times when the rotated square control region should be used are when the circular control region results in a waveform that violates the gradient constraints or when a time-optimal transfer is not needed.

The variational approach appears to be a good technique except when it leads to results that cannot be solved using current methods as occurs when attempting to add the gradient-magnitude constraints or first-moment requirements to the circular control region problem. Other approaches, such as the waveform-building approach, can lead to sub-optimal waveforms that the variational approach avoids. The variational approach (despite its current limitations) also allows optimal controls to be generated for a wider variety of situations. In addition, once robust numerical methods are developed for the solution of such difficult boundary-value problems (half of the boundary conditions specified at an unknown final time), the variational approach will be easily extended beyond the current limitations.

Variational methods should also be useful for designing gradient transfers that have other constraints or that are optimal with respect to other functionals. For example, the

minimum-time transfer could potentially be extended to include gradient induced PNS (peripheral nerve stimulation) constraints, or an appropriate cost functional could be used for designing a slice-selection gradient waveform requiring the minimum RF energy deposition in a given amount of time. Optimizations utilizing such safety or biological-effect based cost functions and constraints might increase patient safety with current protocols or allow the use of imaging techniques that are currently unusable due to safety concerns.

4.6 Conclusions

Using constrained optimization in conjunction with variational methods it was shown that, for 1D transfers, the time-optimal gradient waveforms with gradient-amplitude and slew-rate constraints are either triangular with slopes equal to the maximum slew-rate or trapezoidal with slopes equal to the maximum slew-rate and the flattop at the maximum gradient amplitude. If first-moment requirements are included, the resulting time-optimal waveforms were shown to be bipolar with the same slope and maximum amplitude properties as above. Furthermore, any 2D transfer with a square control region (or any 3D transfer with a cubic control region) was shown to be separable into 2 (or 3) 1D transfers. The resulting time-optimal waveforms are not unique due to the separable nature of the problem. Finally, the form of the 2D and 3D transfers corresponding to constraints on the magnitude of the slew-rate was derived. The resulting waveforms are composed of the square root of a 2nd order polynomial and the inverse hyperbolic sine of a 1st order polynomial in time. These waveforms are not guaranteed to satisfy gradient-magnitude constraints, but they are guaranteed to satisfy the slew-rate constraints regardless of the rotation between the logical and hardware reference frames. The use of

these waveforms will result in gradient transfers that were shown to reduce the dead time to the minimum possible without violating hardware limits.

Chapter 5. Trajectory Design With a Multi-objective GA

5.1 Introduction

Spiral, radial, and other non-rectilinear k -space trajectories are an area of active research in the MRI community (Ahn et al. 1986; Glover and Pauly 1992; Noll et al. 1998; Bornert et al. 2000; Moriguchi et al. 2000; Larson and Simonetti 2001). Much of this interest has been generated because of some of the well-known advantages of these trajectories including rapid acquisition times and reduced artifact levels. However, despite this growing interest, there is little in the way of a general theory or method for k -space trajectory design that would allow easy derivation of new trajectories with specific desired properties.

This problem represents an excellent opportunity for formal optimal-design methods (Deb 2001). These methods would allow specification of the desired properties followed by derivation of the appropriate trajectory shape rather than the typical paradigm of specification of the shape followed by examination of the resulting properties. Such an approach would be beneficial in transforming trajectory design from an art to a science. In addition, an appropriate implementation could allow a search through all possible shapes rather than only trajectories with an easy analytical formulation.

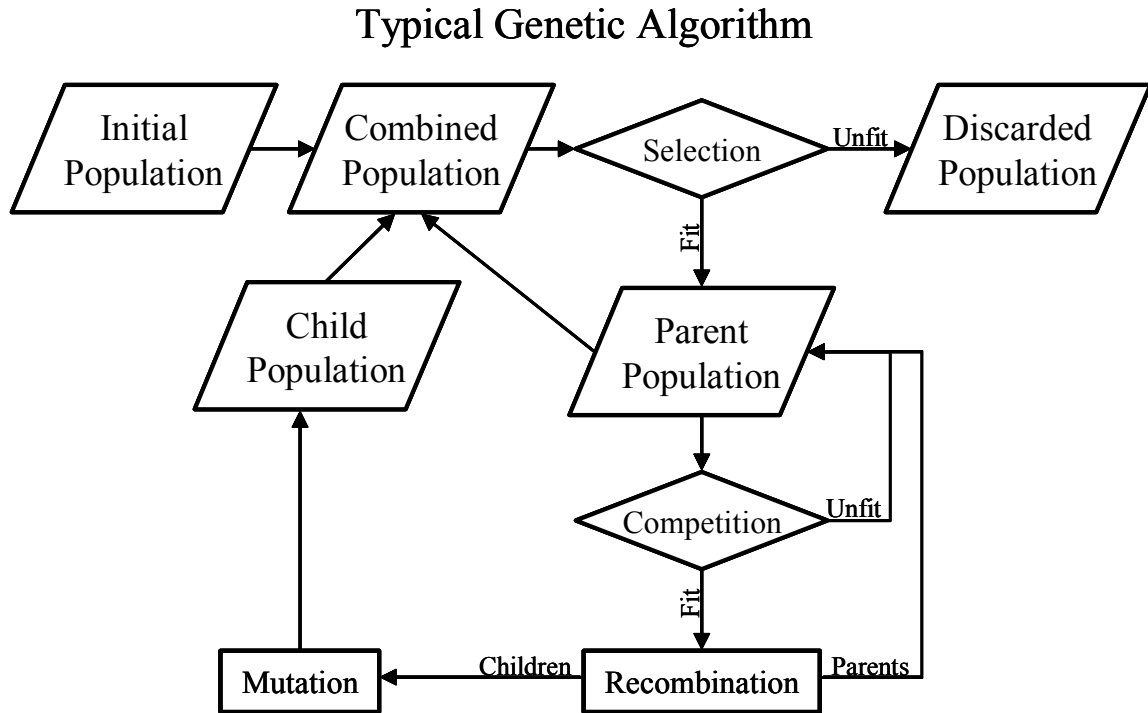


Figure 5.1 Flowchart of a typical genetic or evolutionary algorithm. Each individual is one potential solution to the problem (e.g. a specific trajectory). The initial solutions are evaluated according to some fitness criteria (e.g. acquisition time). The worst individuals are discarded and the fittest form the parent population. Parents are randomly selected and are evaluated according to some, potentially different, fitness criteria (e.g. image quality). New solutions (children) are generated by mixing the features (recombination) from two parent solutions. The child solutions have a small chance of undergoing a random change (mutation) to introduce new features. The parent and child populations are combined and the process is repeated for a number of iterations, or generations in order to gradually develop superior solutions.

There are many different classes of optimization algorithms, but one that has been gaining popularity is the class of evolutionary or GAs (GA's) (Deb 2001). As shown in Figure 5.1, these algorithms use a probabilistic biological metaphor rather than a deterministic hill-climbing metaphor. Parent solutions pass on characteristics to the next generation or results via their children; occasionally mutations occur which introduce

new characteristics to the population of potential solutions. Specifically, each individual solution in a parent population is evaluated with respect to the objectives to obtain a measure of reproductive fitness. Child solutions are obtained from the fittest individuals through crossover operations between multiple parents and/or mutation operations on a single individual. The child solutions join with the parent solutions and some solutions (potentially using a different measure of fitness for survival than for reproduction) are removed from the combined population. The surviving parent and child solutions combine to form the parent population for the next generation and the process is repeated.

Multi-Objective Terminology




Figure 5.2 An example of multi-objective terminology for trajectory optimization. Each trajectory is assigned a point according to its artifact severity and acquisition time, which are known as the objectives. Here, trajectories a) and b) are both clearly preferable to trajectory c), but without additional information, there is no way to choose between a) and b). Thus, out of these three trajectories a) and b) are considered non-dominated and c) is considered dominated. The feasible set is the set of all possible trajectories that do not violate the constraints. The Pareto-optimal set is the non-dominated subset of the feasible set; finding it is the goal of a multi-objective optimization.

Relative to more deterministic optimization methods, GA's are insensitive to the initial guesses and are more capable of handling discrete parameters (Deb 2001). More importantly, the GA's have the advantage of being able to find global optima in the face of non-linear, multi-modal, discontinuous, or otherwise ill-behaved objective and constraint functions (Deb 2001). In addition, multi-objective GA's are unique in their potential to sample the entire Pareto-optimal set in a single run, even when the Pareto-

optimal set is non-convex. In mathematics, Pareto-optimal means the best possible tradeoff between two or more competing objectives as shown in Figure 5.2 (e.g. the best image quality for a given acquisition speed), and non-dominated means that, in at least one objective, the solution is better than any other solution in the set without being worse in any other objective (e.g. a sequence with less artifact in the same time). However, GA's are generally quite computationally expensive and their convergence to the true optimum is usually not guaranteed.

One common use for GA's, particularly in the field of aerospace engineering, is for shape optimization (Sharatchandra et al. 1998; Deb 2001). Typically simulations are used to evaluate the properties of the shape, either because no analytical formulation exists or because of violations of the underlying assumptions in the analytical formulations that do exist. For MR image acquisitions this would be similar to optimization of the k -space trajectory shape using simulations to compute one or more measures of image quality.

Previous optimizations for MR image acquisition techniques have usually focused on objectives, such as SNR, that were chosen for their ease of computation rather than for their impact on image quality (Hendrick et al. 1987; Van Lom et al. 1991; Simonetti et al. 1993; Star-Lack 1999; Gao and Reeves 2000; Sabat and Irarrazaval 2002). Although some attempts have been made to use the more robust properties of single-objective GA's for MRI (typically for hardware design, but once for trajectory design), none have yet used multi-objective GA's to optimize trajectory shape using multiple simulated measures of image quality (Fisher et al. 1997; Williams et al. 1999; Sabat and Irarrazaval 2002). The closest to this type of optimization are Wager and Nichols who have recently

used a single-objective GA to attempt to simultaneously optimize multiple simulated objectives in fMRI experimental design (Wager and Nichols 2003). Their problem would have been a particularly appropriate application for a multi-objective GA.

This work describes the development and use of a multi-objective GA in designing k -space trajectories that Pareto-optimally achieve multiple simulated objectives. The objectives were chosen for their anticipated impact on image quality and not for computational convenience. Simulated images were acquired using the genetically designed trajectories in order to examine the effectiveness of the simulated measures of image quality. The new trajectories were translated into actual pulse sequences, and images were acquired to confirm predicted imaging improvements under real conditions.

5.2 Methods

5.2.1 Trajectory parameterization

Perhaps the most important choice in using a GA for trajectory optimization is the parameterization of the trajectory shape itself. This determines which trajectories are difficult or impossible to represent and therefore unlikely to be discovered. In addition, if a parameterization uses too many numbers to represent the shape, then convergence may be impeded. Several parameterizations were investigated to determine the amount of distortion introduced into the trajectories with a reasonable numbers of parameters. All parameterizations were designed to represent a single interleaf, with the remainder being obtained by successive rotations of the first. Due to the success of spirals and radials in the literature the first representation was a variable density spiral parameterization consisting of a sampling of the $d\theta(r)/dr$ function for the k -space trajectory (Tsai and Nishimura 2000). The second representation was a broadband parameterization where

the direction of travel was specified at each point in time and the trajectory went as far as possible in that direction subject to the hardware constraints (Hardy and Cline 1989). The third, fourth, and fifth parameterizations were direct representation of the slew, gradient, or trajectory waveforms respectively. Additionally, these representations were discretized using linear interpolation, Fourier decomposition, or Chebyshev decomposition (Sharatchandra et al. 1998). In other words, a continuous gradient waveform could be generated from a finite set of data by considering the data as samples of the waveform and interpolating, or by considering the data as coefficients in a finite Fourier or Chebyshev series and evaluating the series. Trajectory parameterizations were tested for accuracy of fit, determined by the maximum deviation from the desired trajectory, first on spiral and then on other trajectories including radial, WHIRL (winding hybrid interleaved radial lines), and SPIDER (Pipe 1999; Larson and Simonetti 2001). They were also tested for the introduction of gradient amplitude and slew-rate constraint violations.

5.2.2 Objectives and Simulations

Objectives were chosen for their potential impact on image quality and not for linearity or other purely mathematical considerations. Additionally, the objectives were required to be properties of the trajectories themselves. For example, contrast is a large factor in the quality of most clinical images; however it is not an inherent property of the trajectory but rather a property of the RF and gradient timings. Instead, artifact severity was considered to be the dominant image-quality property of a trajectory. For example, reduced flow-artifact is the most well-known advantage of spirals while severe off-resonance blurring is the most well-known disadvantage (Meyer et al. 1992; Bornert et al.

1999). Aliasing artifact is also an important property of any trajectory, especially for trajectories that violate the Nyquist criterion in order to reduce acquisition time (Peters et al. 2000; Tsai and Nishimura 2000). Therefore, the four objectives considered here were acquisition time, aliasing energy, flow artifact energy, and off-resonance artifact energy.

The dominant factor in the total acquisition time is the number of interleaves rather than the duration of each interleaf. Therefore, the acquisition time was calculated as: $N_{\text{leafs}} + t_{\text{leaf}}/100$, where N_{leafs} is the number of interleaves and t_{leaf} is the readout time for a single interleaf, in ms. The second term is intended to reduce the number of ties that would otherwise occur by choosing the sequence with the shorter interleaf when the total number of interleaves is the same.

Aliasing energy was calculated by evaluating a Lorentzian envelope: $E_A(k) = C_1/(C_2+|k|^2)$, where $|k|$ is the k -space radius and the constants are found by a least-squares fit to several actual spiral data sets. E_A is evaluated at each location that is not sampled on the k -space grid and represents the energy that will alias into the image. A grid point is considered not sampled if no trajectory point lies within its Voronoi polygon (the Voronoi polygons are unit squares for a rectilinear grid). E_A is then summed over the k -space grid. This measure of aliasing energy applies a greater penalty to a trajectory with a non-sampled grid point in the center of k -space than to a trajectory with a non-sampled grid point in the outer regions.

Measures of off-resonance and flow artifact energy depended on simulations. Analytical phantoms and their corresponding representations in k -space were developed using known properties of the FT, the velocity k -space formalism and an adaptation of the formalism for off-resonance effects (Nishimura et al. 1995). The 0th and 1st gradient

moments and sample time were evaluated for each ADC (analog to digital converter) sample in order to generate the simulated data. These data were gridded onto a 2x oversampled grid using the standard convolution-based gridding-reconstruction with Jackson's convolution-based density compensation function for a width 1 rectangular (for rapid computation of the density compensation function) convolution kernel (O'Sullivan 1985; Jackson et al. 1991).

This simulation process was repeated during the optimization to generate three magnitude images for each individual parent and child trajectory, one without flow or off-resonance effects, one corrupted by flow, and one corrupted by off-resonance. To obtain a measure of flow or off-resonance sensitivity, the uncorrupted image was subtracted from the corrupted image and the total energy of the resulting difference image was computed.

Simulation Phantoms




Figure 5.3 Example images for the three numerical phantoms (rectangular, spherical, and resolution from left to right) used in the simulations for the flow and off resonance objectives. Simulated data are sampled on a radial trajectory and reconstructed using standard convolution-based gridding-reconstruction with a 2x oversampled grid and a width 1 rectangular convolution kernel. The square root of the magnitude image is displayed in order to accentuate the artifacts.

Three different simulated phantoms, corresponding to three different physical phantoms, were used in order to investigate the sensitivity of the simulations to changes in the shape of the phantom (see Figure 5.3). The primary simulated phantom, the only one used during the actual optimization, was a long narrow rectangle in a vertical orientation; the second phantom was a uniform sphere; and the third phantom was a simulated resolution phantom (Bornert et al. 2000). The rectangular phantom was evaluated using $\text{sinc}(w_x k_x, w_y k_y) = \text{sinc}(w_x k_x) \text{sinc}(w_y k_y)$, and the circular phantoms were evaluated using $\text{jinc}(r k) = J_1(|r k|)/|r k|$, where $k = k_x + i k_y$, J_1 is the first order Bessel function of the first kind, w is the width of the rectangle, and r is the radius

of the circle. All three phantoms were also used to simulate a continuous range of off-resonance from -100 Hz at the bottom of the image to +100 Hz at the top. The rectangular phantom was also used to simulate a flow of 100 cm/s (Nishimura et al. 1995).

5.2.3 NSGA-II

Deb's elitist NSGA-II (non-dominated sorting genetic algorithm II) is a well-known multi-objective GA (Deb et al. 2002). It is capable of converging to the Pareto-optimal set while maintaining diversity along it using a crowded-comparison operator that incorporates both the non-domination rank and a crowding measure for each solution. The NSGA-II was implemented using Mathematica 4.0. In this implementation, the population size was 200. Such a large population was required to adequately sample the 3D Pareto-optimal set associated with four-objective problems (Deb 2001). For every child solution, two parent solutions were chosen through binary tournament selection. As previously mentioned, it is possible to use different selection criteria for survival and reproduction. This can continue to push the population towards the Pareto-optimal set even when the survival selection has generated a purely non-dominated population. In this case, the binary tournament selected the fastest trajectory if both trajectories met minimum image-quality standards (aliasing energy < 0.04 , flow and off-resonance artifact energy < 0.01 , as determined in prior trials using these measures of aliasing energy), otherwise it selected for the best image-quality. Child solutions were obtained through simulated binary crossover ($\mu_c = 3$) between the parents followed by lognormal mutation ($P_m = 0.01$, $\sigma_m = 0.5$) on the child (Deb 2001). Simulated binary crossover is a commonly used computational analog of genetic recombination, while lognormal

mutation perturbs a simulated gene by a small scale-factor. 40 children were produced for each of the 400 generations, resulting in a total of approximately 16,000 trajectories examined. To decrease the total computation time, the algorithm was executed in parallel on a local area network of PC workstations running Windows NT, 2000, and XP using Mathematica's Parallel Computing Toolkit. The Toolkit allows parallelization of a computation across a network of heterogeneous computers.

5.2.4 Constraints

The constraints for k -space trajectory design can generally be grouped into two categories: engineering constraints such as gradient amplitude and slew-rate limits and safety constraints such as gradient-induced PNS and SAR (specific absorption rate). SAR was not expected to be an active constraint for any of the tested k -space trajectories. These trajectories were to be implemented on a 1.5 T Sonata (Siemens Medical Solutions, Erlangen, Germany) and therefore the optimization used a slew-rate constraint of 200 T/m/s and a gradient amplitude constraint of 40 mT/m. The engineering constraints were enforced through a feasibility-preserving strategy where the trajectory is at one of the limits at every point in time for the variable-density spiral and broadband parameterizations. The engineering constraints were enforced through the constrained non-dominated sorting procedure for the other parameterizations (Deb et al. 2002). PNS constraints were evaluated using the SAFE (stimulation approximation by filtering and evaluation) model with the parameters provided by Siemens for the Sonata (Hebrank and Gebhardt 2000). The PNS safety constraint was enforced through the constrained non-dominated sorting procedure.

The FOV is also an important factor in designing k-space trajectories because it indirectly affects all of the above constraints through its impact on the gradient waveforms. There are at least two possible approaches for handling the FOV, the first is to consider it as both a genetic parameter and an objective, and the second is to simply use a fixed FOV. The latter approach was used in this work to avoid adding a 5th objective. In order to reduce the need for re-optimizing, a relatively small FOV of 128 mm was selected. The resulting trajectories may be used for any FOV ≥ 128 mm without violating any of the constraints.

5.2.5 Implementation

A total of 20 trajectories tested during the optimization were implemented on the 1.5 T Siemens Sonata mentioned previously using the standard Siemens pulse-sequence development tools. The trajectories to be implemented were selected in groups designed to vary, as much as possible, in only a single objective. Four groups were selected (see Appendix B), one for each objective (acquisition time, aliasing energy, flow artifact energy, and off-resonance energy). Each group contained a Pareto-optimal trajectory, a standard trajectory, and three sub-optimal trajectories.

Once implemented, the actual trajectory was measured and the objectives were recalculated based on the measured, versus the programmed, trajectory (Duyn et al. 1998). Image reconstruction was performed by a table-lookup method for gridding reconstruction (Dale et al. 2001). Gridding tables were calculated based on the measured trajectory acquired. Although contrast is not an inherent property of the trajectory, care was taken to maintain similar contrast within each group in order to aid comparison.

5.2.6 Comparison of Theoretical and Actual Objectives

The quality and optimality of the final trajectories was expected to depend largely on the accuracy and effectiveness of the simulations used to calculate the objectives. For each trajectory, images were acquired under conditions that recreate, as much as possible, the simulated images, with the number of averages set between 5 and 17 to both control for and reduce noise levels. To re-create the off-resonance objective, one image was acquired under conditions of careful shimming and a second after deliberately offsetting one of the first-order shim currents. To test the flow objective, a flow phantom was imaged, after careful shimming, under flowing and non-flowing conditions. The aliasing energy objective could not be recreated experimentally. Instead, after noting that aliasing causes energy from objects in one portion of the image to be displayed in other portions of the image, a point spread function was acquired by imaging a small object at the center of the FOV. The central portion of the image was masked out, and the total energy of the remainder was calculated. The time objective was measured by multiplying the TR by the number of interleaves. Linear regression was used to evaluate the effectiveness of the simulated objectives as independent linear predictors of the measured objectives.

5.3 Results

5.3.1 Trajectory parameterization

Table 5.1 Comparison of Trajectory Encoding Methods

Parameterization Method	Discretization Method	Max Deviation (grid points)	Max Gradient Overshoot (%)	Max Slew Overshoot (%)
High-bandwidth	Linear	13.9	0	0
	Fourier	141.	0	0
	Chebyshev	0.7	0	0
Slew	Linear	2565.	100	0
	Fourier	311.	13	13
	Chebyshev	14.9	1	6
Gradient	Linear	9.5	2	0
	Fourier	8.0	11	675
	Chebyshev	0.5	0	2
Trajectory	Linear	5.3	0	1740
	Fourier	31.7	1270	30,000
	Chebyshev	0.5	0	13

The variable-density spiral parameterization was found to be accurate at representing spiral and WHIRL trajectories with a relatively small number of data points, but was fundamentally incapable of representing any other class of trajectories and was therefore discarded from further consideration. As shown in Table 5.1, the slew waveform encoding resulted in severely distorted trajectories for all discretization methods while the trajectory waveform encoding resulted in spurious slew-rate constraint violations; these encodings were also dropped from consideration.

Trajectory Parameterizations for 6-interleaf Spiral

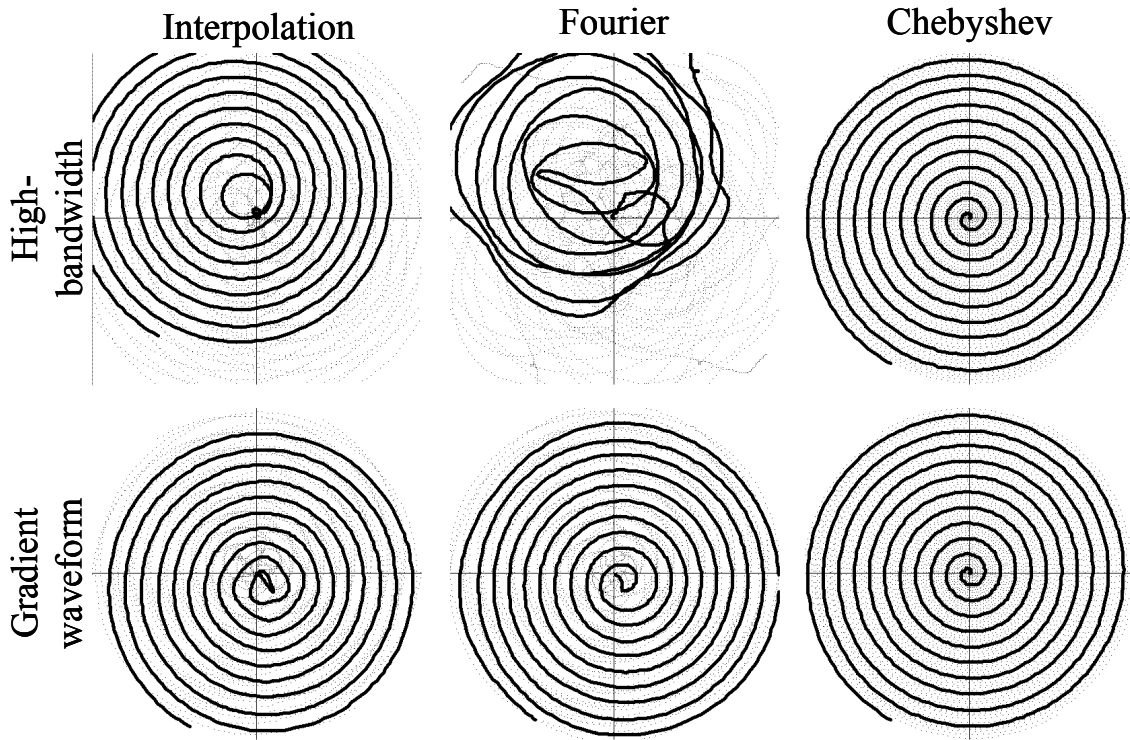


Figure 5.4 Approximations to a standard 6-interleaf spiral trajectory using the same number of terms for several different parameterization schemes. For each scheme a single interleaf of the trajectory is displayed by the bold line while the sampling patterns of the remaining interleaves are displayed with the points. The top row is a high-bandwidth parameterization where the direction of travel is encoded at each step in time and the trajectory proceeds as far as possible in that direction subject to hardware constraints. The bottom row is direct encoding of the required gradient waveforms. Each parameterization can be encoded using linear interpolation, Fourier series decomposition, or Chebyshev series decomposition.

The performance of the high-BW parameterization was found to depend strongly on the discretization chosen (see Figure 5.4) with the Chebyshev decomposition performing best. The gradient waveform encoding was generally more accurate than the high-BW representation (see Figure 5.4), with the Chebyshev discretization again being most

accurate. Comparisons of these two parameterizations with Chebyshev decomposition on a broad range of standard k -space trajectories (spiral, WHIRL, radial, and SPIDER) demonstrated that both could reproduce all classes of trajectories with less than 1.3 grid points of maximum deviation, in k -space, between the desired and realized trajectories. However, the high-BW encoding was generally both slightly less accurate and, by always driving the hardware at the maximum slew-rate or gradient-amplitude, it should tend to result in more gradient-induced PNS and trajectory distortions. Therefore, the remainder of this work utilizes the Chebyshev decomposition of the gradient waveform.

Four parameters were chosen to describe each trajectory during the optimization. The first was an integer for the number of interleaves, the second was an integer for the number of gradient waveform sample points, the third was a complex floating-point number for the initial k -space location, and the fourth was a variable-length list of complex floating-point numbers for the Chebyshev coefficients of the gradient waveform.

5.3.2 Computational Time

The evaluation of a single trajectory typically required between 1 and 96 seconds on a 3 GHz Pentium 4 running Windows XP with 1 GB RAM. Feasible trajectories with large numbers of data points or interleaves tended to take more time, while trajectories that violated one of the constraints tended to take less time. The non-dominated sorting procedure required approximately 15 seconds on the same computer system. Slower computers, or computers with less RAM required more time for the trajectory evaluation. Using parallelization across 7 to 9 heterogeneous workstations allowed a single generation to be completed every 3-5 minutes. The full 500 generations required 2 days




Figure 5.5 One projection through the four-dimensional objective space onto a two-dimensional data set. One dimension is the acquisition time objective while the other dimension is a linear combination of the three image quality objectives. Each point represents a k -space trajectory. Non-dominated trajectories (in the original objective space) are displayed in bold and other tested trajectories are displayed small.

to complete and were computed during a weekend, when the workstations were not otherwise occupied.

5.3.3 Optimal trajectories

With four objectives it is difficult to plot all of the trajectories on a single graph in objective space, but it is possible to use a linear combination of the three image quality objectives to project the 4D (four dimensional) data set down onto 2D. Figure 5.5 shows one such projection. Note the relatively dense sampling of the objective space in the region of the non-dominated set. Note also that some of the original non-dominated trajectories are not Pareto-optimal with respect to this particular weighting, but that all of

Sample GA-designed Trajectories




Figure 5.6 Three GA-designed non-dominated trajectories. All three trajectories begin with non-zero gradient amplitude at the k -space origin. They are slow trajectories with good off-resonance and flow properties. All three trajectories move outward, first in a straight line, and then with a fairly gentle curvature relative to a standard spiral.

the projection Pareto-optimal trajectories are from the original non-dominated set. Proceeding along the projection Pareto-optimal set from fast trajectories with low-quality images to slower trajectories with higher quality, we first encounter a region of standard spirals with progressively more interleaves. These give way to WHIRL trajectories followed by a group of new, genetically designed trajectories. This general trend was observed for most linear combinations investigated.

The non-dominated GA-designed trajectories generally start at the center of k -space with non-zero initial gradient amplitude and tend to move outward, first in a straight line, and then with a fairly gentle curvature relative to a standard spiral. Similar trajectories were also discovered in a previous optimization using slightly different flow and off-




Figure 5.7 Comparison of an optimal GA designed trajectory (middle trajectory in Figure 5.6) with a radial, spiral, and WHIRL trajectory with the same numbers of interleaves. The Voronoi area is reported in units of squared grid points so that values greater than 1.0 represent an undersampled region. Note that the new trajectory is more oversampled than any other trajectory. The radial trajectory samples each k -space radius at two points in time. Although the optimal trajectory begins with non-zero initial gradient amplitude, it generally uses a smaller gradient than the other trajectories.

resonance objective functions. Three of these trajectories are displayed in Figure 5.6 in order of increasing image quality and acquisition time.

Of those three trajectories, the middle one (in both image quality and acquisition time) is compared to 156-view radial, spiral, and WHIRL trajectories in Figure 5.7. Note that the optimal trajectory is more oversampled (smaller Voronoi areas) than the other




Figure 5.8 Comparison of a cardiac image acquired using a new GA designed trajectory (middle trajectory in Figure 5.6) with similar images acquired using two standard radial trajectories. The 157-view radial is matched to the GA trajectory in number of views, but requires 45% more averages in order to achieve the same SNR and therefore required a longer imaging time. The 109-view radial is approximately matched in terms of imaging time. Note the superior quality of the GA image, in particular note the reduced artifact near the chest wall, the more uniform signal in the aorta, and the improved definition of the coronary artery.

trajectories, and that only the radial trajectory ever becomes undersampled. Note also that, although the optimal trajectory begins with non-zero gradient amplitude, it generally uses lower gradients than the standard trajectories.

A cardiac image acquired using the same optimal trajectory is also compared to images from 109-view and 157-view radial trajectories in Figure 5.8. The 157-view radial is matched to the GA trajectory in terms of number of interleaves, but the radial trajectory requires approximately 45% more averages (9 for GA trajectory, 13 for radial trajectory) in order to obtain the same SNR. The 109-view radial is matched in terms of total acquisition time. Note the improvement in overall image quality, particularly in the regions of the chest wall, aorta, and coronary artery.

5.3.4 Constraint activity

Four thousand trajectories tested were considered non-feasible due to constraint violations. The gradient-amplitude constraint was active along 4% of the non-dominated set while the slew-rate constraint was active along 12%. The PNS constraint was active along 16% of the non-dominated set. No trajectory has triggered gradient hardware or PNS safety features even when run with minimum TE at a 128mm FOV on the 1.5 T Sonata. Although SAR constraints were not considered in the optimization, no trajectory has triggered the RF safety features.

5.3.5 Objective function sensitivities

After the optimization was completed, the 20 test trajectories were implemented and measured. For each group, the corresponding objective functions were re-calculated based on the measured trajectories, i.e. the flow artifact objective was re-calculated for trajectories in the flow group (see Methods: Objectives and Simulations above). There was a high degree of correlation ($R^2 > 0.98$) between the objective functions based on the designed and measured trajectories for all three image-quality objectives. In addition, the off-resonance objective was also recalculated for the two other simulated phantoms, the spherical phantom and the resolution phantom. The measures had the same trends as before, with the correlation parameter being positive with greater than 95% confidence in all cases. The exact details were object specific, but with R^2 values greater than 0.8 in all cases. No obvious indications of object-specific features (e.g. changes in curvature or density near the k-space radius corresponding to the width of the rectangle) were noticed in the genetically designed trajectories.

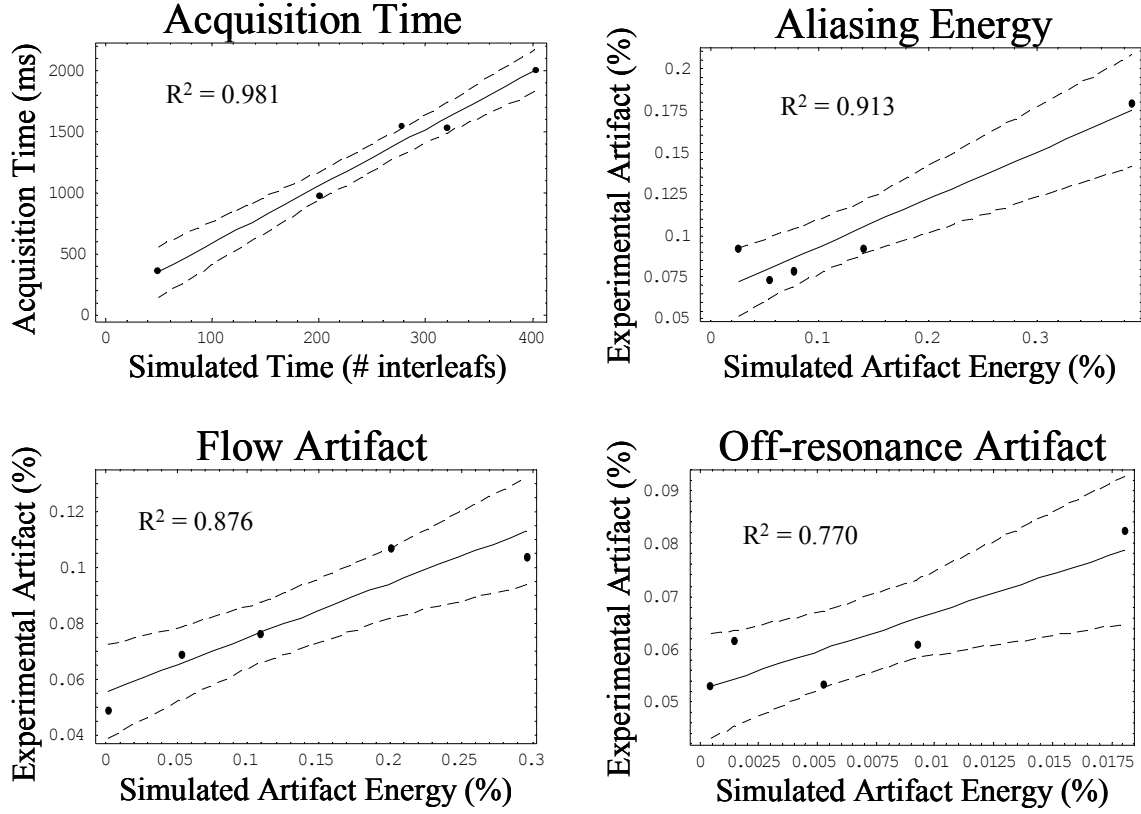


Figure 5.9 Plots of the correlation between the simulated objective functions and their experimentally measured analogs. Each point is a single simulated-experimental objective pair for a single trajectory while each plot is a single group. The solid line is the regression line. The dashed lines are formed by the mean prediction confidence interval at each location and demark a region 90% likely to contain the true regression line.

5.3.6 Simulation accuracy

Figure 5.9 displays the correlation between the simulated and experimental objective measurements for each objective. The simulated and experimental time objectives were strongly correlated ($R^2 = 0.981$) while the correlation for off-resonance objectives was not as strong ($R^2 = 0.770$). In all cases the slope of the regression was positive with 95%

confidence. All three of the image-quality objective regressions also had a constant term that was positive with 95% confidence.

5.4 Discussion

5.4.1 Algorithm performance

The algorithm appears to have converged well. The non-dominated set appears clearly defined and little advancement was observed in the last 50 generations. The population became entirely non-dominated within the first 100 generations. Subsequent progression towards the Pareto-optimal set can largely be attributed to the reproductive sorting function used in the binary tournament. The increased density towards the higher number of interleaves probably reflects the presence of the minimum-image-quality aspect of the binary tournament; faster trajectories failed to meet that minimum and were given lower priority. Alternatively or additionally, it may reflect the fact that the variable-length list of Chebyshev coefficients was usually shorter for high-interleaf trajectories, possibly leading to improved convergence towards that region of the Pareto-optimal set.

Consider the progression along the projection Pareto-optimal set of Figure 5.5. On the rapid (low-quality) side we first encounter spiral trajectories with progressively more interleaves, followed by WHIRL trajectories, and finally by the new trajectories shown in Figure 5.6. At each point along that progression, the Pareto-optimal trajectory represents the best possible image-quality for a given acquisition speed; the progression along the set represents an inherent trade-off between time and quality. This implies several things. First, although the quality of images acquired using the new trajectories is superior to images acquired with the rapid spiral trajectories, it is incorrect to say that the new

trajectories are superior to the rapid spirals. Instead, the preferred trajectory depends on the requirements of a particular application to determine if the sacrifice in speed is worth the improvement in quality. Second, the transition from spiral to WHIRL trajectories to the new trajectories indicates that at medium numbers of interleaves the WHIRL trajectories are superior to spiral trajectories with the same number of interleaves and high numbers of interleaves the new trajectories are superior to both WHIRL and spiral trajectories with the same number of interleaves. It appears that these new trajectories evolved the radial center of the WHIRL trajectories and improved on it by adding the non-zero initial amplitude component. Finally, because no weighting between the objectives was specified during the optimization, this optimization only needs to be performed once in order to derive the optimal trajectories for a variety of applications. In fact, the optimal trajectory for any application where the important factors are any function of these four objectives can be determined without repeating the optimization.

As shown in Figure 5.5, even with a projection through the data set, it is possible to see the dense sampling of the objective space in the vicinity of the non-dominated set. This increased sampling density is the key advantage of the GA over an exhaustive or random search approach; however, it also compounds the difficulty of selecting appropriate groups of trajectories by concentrating the tested trajectories within a relatively small region near the Pareto-optimal set.

The use of a true multi-objective optimization has some distinct advantages over a single objective optimization (Deb 2001). A single objective optimization would have to determine, in advance, how much image quality we would be willing to sacrifice for a given increase in acquisition speed. This would vary from application to application and

is not trivial to determine even given a particular application. In addition, the weighting coefficients themselves would be strange quantities, particularly when one or more of the objectives were in some sort of arbitrary unit (e.g. A.U.²/s). With the multi-objective optimization no such weighting coefficients are needed. Alternatively the time-quality trade off may be examined for any given weighting without needing to repeat the optimization, and for any given weighting the best trajectories will always be a subset of the original non-dominated trajectories. Because most MRI acquisition design problems represent a trade-off between multiple conflicting objectives, the multi-objective approach is generally preferred.

5.4.2 Objectives

The results of the “Objective Function Sensitivities” section indicate that the typical deviations from the intended k -space trajectory do not greatly disturb the results of the optimization, at least not for the objectives used here. However, there is a potential for the optimization to converge to a sub-optimal region simply because the simulated object is different from the experimentally imaged object. It is important to note that, even in the worst case, the slope of the regression was still positive with 95% confidence and that the simulated objects do not need to be exactly the same in order to have a very strong correlation. In addition, no obvious indications of features related to the simulated rectangle were noted in the optimized trajectories. These observations indicate that the optimization results may not be highly sensitive to the specific numerical phantom, but further investigations may still be warranted. Alternative phantoms include the Shepp-Logan phantom, point-spread functions, or circular phantoms, but the point-spread

function may favor uniform trajectories and rotationally-symmetric phantoms may cause problems related to low-signal rings in k-space.

One factor that may have limited the strength of the correlation between the simulated and experimental objectives is the fact that the various trajectories in each group are not exactly uniform in the other three objectives. Additionally, because it is not possible to obtain a perfect shim, the baseline images for the off-resonance cost function, and all images for the aliasing and flow cost functions, are unavoidably corrupted by some amount of off-resonance. This is especially true for the flow and point phantoms which were particularly difficult to shim. Despite the presence of such effects, the correlation between the simulated and experimental measures was positive with 95% confidence in all cases. This means that improvements in the simulated measures during the optimization generally resulted in measurably superior image quality. Because the GA only compares objective values to determine relative ranking, and not the magnitude of the difference, the actual slope of the regression is not as important as the fact that it is positive. With more data points it would be possible to get reasonable confidence limits on the value of R^2 . This might allow for convergence criteria based on stopping the optimization when we were confident that the improvements were smaller than some fraction of the unexplained variance for all objectives. The fact that the constant term was also positive with 95% confidence for all image-quality objectives is consistent with the presence of noise energy in all of the experimental measures, and the absence thereof in the simulations.

When evaluating these results, it is important to realize that the particular objectives chosen here impact the specific result of this optimization rather than the general

applicability of the optimization technique. In this sense, the fact that the algorithm was able to converge to a well-defined set of non-dominated trajectories may prove to be more important than the derivation of the particular trajectories presented in Figure 5.6. Specific applications may demand a different set of objectives or even a different set of parameters, but the convergence, robustness, and other properties of the GA itself will remain even though these specific results may not.

It would be useful to have a wider variety of standardized computational image-quality metrics in MRI. The current selection of metrics is largely limited to SNR and CNR (contrast to noise ratio), neither of which are designed to assess the impact of artifacts. ROC (receiver operator characteristic) curves, while capable of assessing the impact of artifacts, require a human expert to evaluate each image, which is possible but not practical for optimal design. One potential type of computational metric that has recently been introduced to the field of MRI and may overcome these limitations is the use of perceptual models (Salem K.A. et al. 2002). Other metrics might include measures of the precision or accuracy of quantitative techniques or computed analogs of detection thresholds for diagnostic techniques.

Despite the lack of good standardized objectives, the results obtained here with these objectives are encouraging. We have demonstrated convergence in a difficult shape-optimization problem with four simulated objectives each having positive correlation to the experimental analogs. Additionally, the similarity of the results in Figure 5.6 to trajectories obtained in previous optimizations with different cost functions is an indication that non-zero gradient amplitude at the k -space origin may indeed result in beneficial flow and off-resonance properties in conjunction with such curves.

5.4.3 Parameters

The Chebyshev gradient-waveform parameterization of the trajectory shape was found to provide an accurate representation of a variety of trajectories with a reasonable number of terms. Other researchers have also had success in using orthogonal series representations rather than interpolation in shape optimizations (Sharatchandra et al. 1998). One additional parameterization that may be useful is a rational polynomial parameterization. Often the minimax rational approximation to a function will require fewer terms for the same degree of accuracy as the minimax polynomial approximation (Press et al. 1992).

It should be noted that the variable-density spiral parameterization was capable of representing spiral trajectories with a much smaller number of parameters. Similarly, other classes of designed trajectories can be specified with a small number of parameters at the expense of generality. Due to the fact that all of the Pareto-optimal trajectories obtained using the Chebyshev parameterization (see Figure 5.6) could also have been described using a generalization (to allow non-zero initial gradient amplitude) of a variable-density parameterization, it may be reasonable for future optimizations to use a generalized variable-density parameterization without much concern about losing useful trajectories. An optimization with a smaller number of parameters may have better convergence properties; therefore there may be an opportunity to make a trade-off between generality and convergence. Although Figure 5.5 indicates that the optimization did converge, it is possible that it could have converged in fewer generations by using less general parameterizations.

In this case, the generality of the Chebyshev parameterization enabled the discovery of a new class of trajectories with apparently beneficial properties. These trajectories begin at the origin with non-zero gradient amplitude and move outward, first in a straight line, and then with a fairly gentle curvature relative to a standard spiral. The small amount of time required for the prephasing gradient pulse and ramp up was found to be a worthwhile investment even though both simulated flow and off-resonance effects accumulate during the application of the prephasing gradient lobes in the simulations. It is possible that there is some disadvantage to having the closely-spaced points near the origin that inevitably result from simply beginning at the center of k -space. This result can be applied in future trajectory design work by simply adding such prephasing gradient lobes to the beginning of spiral, WHIRL, or other standard k -space trajectories. This could be done even without utilizing a formal optimization procedure, although some amount of simulation may be necessary in order to determine the range of beneficial initial gradient amplitudes.

K -space trajectory design was selected because, like all shape optimizations, it is an inherently difficult problem due to the infinite number of degrees of freedom. The assumption was that, if an optimal-design technique could prove successful here, it should be powerful enough to handle most MR image acquisition design issues. In many applications, a standard rectilinear acquisition may be desired. For example, image-quality objective functions based on contrast could be utilized with parameters such as TE, TR, and the presence or absence of magnetization-preparation pulses. Such an optimization could still be used to find the fastest acquisition, subject to all of the engineering and safety constraints, for a given level of image-quality.

5.4.4 Improvements

In order for this technique to become generally accepted at least three specific improvements will be required. First, any increases in the ability to accurately simulate an MRI acquisition will lead to improvements in cost functions based on those simulations. For example, a standardized high-quality virtual heart, lung and rib cage would be useful for developing optimal cardiac sequences. Second, the current list of multi-objective GAs must improve. In particular, an ideal algorithm would be computationally inexpensive with few or no algorithmic parameters and guaranteed convergence. Third, it would be desirable to be able to simply insert particular cost functions onto pre-packaged, or at least standardized, parameterizations and optimization routines.

Some of the above-mentioned improvements will occur merely as a function of continuing increases in computer power. Others will require new innovations in the various fields. In particular, multi-objective GAs are still in their infancy and there is no reason to expect that their convergence properties will remain poorly understood for long. However, even without the benefit of these improvements, the optimization technique used here represents a powerful design tool.

5.4.5 Optimal design

This work represents, to our knowledge, the first attempt to use multi-objective GAs in the design of MR image acquisition techniques. It is hoped that such rational design procedures will transform the design process from an art to a science. As robust and effective design procedures are developed and utilized, the quality of future techniques can be rapidly improved relative to the normal progress and the quality will be less

dependent on the sequence-development skills of the designer. Such design procedures would be desirable in order to ensure that patients consistently receive the highest possible quality of care.

In addition, there are many researchers and clinicians who have a novel imaging need. Often, they can be very specific about the type of information that they want to obtain and the properties of the images that would provide that information. The technique developed here could prove quite useful for researchers and clinicians facing such situations by allowing them to obtain a pulse sequence that would optimally meet their specific requirements.

5.4.6 Conclusions

In conclusion, we have presented a technique for the optimal design of k -space trajectories. This technique allows the specification of the desired properties of the trajectory followed by a genetic-algorithm based search through all possible trajectories to find those that optimally accomplish the specified objectives. The algorithm was able to converge despite the inherent difficulty of shape optimization due to the large number of free parameters. Deviations from the designed trajectories that occur on modern MRI scanners were found to have relatively little impact on the objective functions. Changes in the simulated object had a relatively greater impact on the objective though similar objects such as the spherical and resolution phantoms were strongly correlated. Finally, reduction of the simulated artifact levels during the optimization was correlated with reduced artifact levels in the experimentally acquired images. No theoretical or practical limitation was encountered which would fundamentally exclude this technique from more general application to other MRI design problems.

We believe the method has significant potential for improving the utility and quality of non-rectilinear k -space trajectories. Feasibility of the method has been confirmed and results show predicted improvements. Novel trajectories with specific image-quality properties can now be derived without violating critical hardware or patient-safety constraints. The resulting improvements in image-quality and patient safety should hopefully lead to improved clinical outcomes. We also believe that extensions to this method will have similar potential for improving most MR image acquisition techniques.

Chapter 6. Subjective Image Quality of GA Trajectories

6.1 Introduction

Non-rectilinear k -space trajectories, including spirals and others, are an area of active research for MRI (Ahn et al. 1986; Glover and Pauly 1992; Liao et al. 1997; Noll et al. 1998; Bornert et al. 2000; Moriguchi et al. 2000; Cline et al. 2001; Larson and Simonetti 2001). Much of this interest has been generated because these trajectories typically have rapid acquisition times and benign flow or motion artifact patterns (Bornert et al. 1999). Many trajectories are known to have increased sensitivity to off-resonance effects, and some trajectories deliberately violate the Nyquist criterion in order to improve acquisition speed at the expense of increased artifact (Meyer et al. 1992; Peters et al. 2000; Tsai and Nishimura 2000).

Most trajectories have been designed using heuristic methods, but recently two groups have used GAs to generate trajectories that are optimal with respect to one or more measures of image quality (Sabat et al. 2003; Dale et al. 2004). In particular, a multi-objective GA has been used to generate trajectories that minimized simulated measures of aliasing energy, flow-artifact energy, and off-resonance energy (Dale et al. 2003c; Dale et al. 2004). The utility of such optimal trajectories depends, in large part, on the effectiveness of the objectives utilized during the optimization.

These simulated measures of image quality have been verified using numerical ratings of water phantom images acquired under carefully controlled experimental conditions (Dale et al. 2004). Clinical images, however, typically have signal from a variety of tissues and cannot be acquired under such carefully controlled conditions. Additionally, a given artifact may be very large according to these numerical ratings

while being visually insignificant. To overcome these limitations, it is important to study the relationship between the original simulated objectives and subjective image-quality ratings of *in-vivo* images as given by experienced medical imaging raters.

6.2 Methods

6.2.1 Experimental design

Fifteen trajectories were created as described in previous work (Dale et al. 2004). The trajectories were arranged in three groups (see Appendix B), each group corresponded to one of the three image quality objectives: aliasing energy, off-resonance artifact energy, flow artifact energy (the acquisition-time group of the original work is not considered here). Each group was selected to obtain a large variation in the corresponding objective and as little variation as possible in the other two objectives. Additionally, each group consisted of one optimal trajectory, one standard trajectory, and three sub-optimal trajectories.

A head image, a cardiac image, and a carotid image were acquired in a single volunteer for each of the fifteen trajectories. In addition, one simulated phantom image from the original optimization was also included for each trajectory as a fourth “location”. Care was taken to ensure similar contrast within groups and locations. The number of averages was set between 5 and 17 to both control for and reduce noise levels. Images were reconstructed with a table-based gridding-reconstruction using the measured trajectories (Duyn et al. 1998; Dale et al. 2001).




Figure 6.1 Screen shot of the graphic user interface used to obtain the subjective ratings. Each row contains three images from one group including one acquired using the standard trajectory. The TE, TR, and flip angle were held constant within each group in order to avoid non-artifact sources of contrast variation within each row. Each presentation had an associated ROI and question. The “Toggle ROI” button turned the ROI overlay on and off. During a training session, raters were instructed in the use of the GUI and in the appearance of various artifacts. They indicated their judgments using the sliders on the right, and clicked “OK” when they were satisfied with their ratings.

Image quality was assessed using a variation of the high-definition television industry standard method of DSCQS (double stimulus continuous quality scale) rating (International Telecommunication Union 2002). This method involves the presentation of pairs of images, obtaining separate ratings for each image using a pair of visual analog scales, and then averaging the scores across a relatively large number of observers to obtain a single quality score for each image. The only difference between the standard

DSCQS method and ours was the simultaneous presentation of 6 images as shown in Figure 6.1. Following standard practice, the test session was preceded by a brief training session and a rest period during which the subject who was evaluating the images could ask questions of the test administrator. The training session utilized a different series of images from the test session, but acquired using the same trajectories and also designed to span the full range of image quality. During the training session, specific types of artifacts, such as blurring or streaking, were indicated to the subject and specific instructions were given regarding scoring, the use of the GUI, and the criteria for judging images.

Each row of each presentation consisted of one image acquired using the standard trajectory and two using other trajectories within a single group and location. For a particular image quality issue, a ROI was marked using a graphical overlay and a particular question was asked. For example, to identify potential blurring, we would indicate a region of brain near the sinus cavities and ask, “How sharp are the structures in the brain?” Subjects recorded answers using the standard method of a visual analog scale ranging from “Excellent” to “Bad,” as shown in Figure 6.1. In all cases, subjects were asked to rate images based only on the severity of artifact within the ROI. In this experiment, we used 9 specific ROI with a maximum of three and a minimum of 1 ROI for each anatomical location.

To rate all 15 images (one from each trajectory) for each question required 3 presentations of 6 images each, with some repetition of the standard trajectory images as dictated by the DSCQS method. Therefore, excluding the training and stabilization images, each subject rated a total of 162 images from 27 presentations. Presentation

order was randomized according to standard practice. Images were evaluated in a darkened, quiet room under controlled conditions. Subjects consisted of 10 scientific MR imaging experts and 3 board-certified radiologists.

6.2.2 Statistical analysis

Data was analyzed according to the industry standard methods for the DSCQS (International Telecommunication Union 2002). Specifically, each slider rating was first converted to a numerical score. For each of the 162 images, a mean score was obtained across all observers along with an associated 95% confidence interval. The mean and confidence interval scores were then linearly transformed to a score between 0 and 1. The standard analyses included specific methods for testing normality, for rejection of raters, and for relating the subjective data to the original simulated measures of image quality.

For this work, the process of relating the subjective data to the original simulated image quality measure was particularly important. It is well known that these types of ratings often produce a sigmoidal curve (International Telecommunication Union 2002). Due to the sigmoidal shape, standard linear regression is not effective, but rank-order statistics such as the Chi-squared test and Spearman's rank-order correlation can be used instead. It is also possible to test the null hypothesis that the optimal trajectories did not have better subjective scores than the standard trajectories.

When a more detailed analysis is desired, the standard practice is to fit the data to a logistic curve of the form: $f(x) = (1/u + b_0 b_I^x)^{-1}$ where u is the upper limit, b_0 is the y-intercept, and b_I is the steepness of the logistic regression (International Telecommunication Union 2002). Such curve fitting is available on SPSS 11.5 (SPSS

Inc., Chicago, Illinois) and was applied separately for each of the three simulated objectives. The significance of each regression was determined by the use of the F statistic test, which is related to the confidence intervals on the regression parameters. Specifically, if the F statistic was significant at the 0.10 level and the slope of the curve was negative, then we can conclude with 95% confidence (one-tailed) that reductions in the simulated artifact objectives correlate with improvements in the subjective ratings. The proportion of variance accounted for, or R^2 , was computed as a measure of the relative importance of successful objectives.

6.3 Results

Before performing the analyses described above, the data was explored for evidence of confounding effects. The standard test for screening raters was applied and no raters were rejected. No significant correlations were observed between the simulated objective variables, although there was a slight negative correlation ($r = -0.061$) between aliasing energy and flow artifact energy and a slightly stronger positive correlation ($r = 0.172$) between aliasing energy and off-resonance artifact energy. There was no significant correlation of average rating with presentation order. No significant differences were found between image ratings across the three columns of images in each presentation. There was a small (5%), but statistically significant, decrease in the average rating of images in the bottom row, which coincided with an increased occurrence (11%) of the two worst trajectories (according to the simulated aliasing energy objective). A principal components analysis indicated that a single factor accounted for 72% of the total variance in the raw image ratings; this factor was always less than .05 standard deviations away from being a z-transformation of the mean scores.

Table 6.1**Spearman's Rank-Order Correlations**

Objective	Across Groups		Within Groups	
	Value	Sig.	Value	Sig.
Aliasing	.43	.00	.64	.00
Flow	.00	.96	.18	.00
Off-resonance	.24	.00	.26	.00

To obtain the subjective ranks for the Chi-square and Spearman's tests, the subjective scores from each presentation of six images (across two different groups of trajectories) were ranked from 1 to 6 and the scores from each row of three images in each presentation (within a single group of trajectories) were ranked from 1 to 3. Objective ranks were obtained similarly using each simulated objective score. All Chi-square tests were highly significant ($p < .001$). Table 6.1 presents the results of the Spearman's rank order correlation tests. Within groups (e.g. only considering images from the group of trajectories selected to vary, as much as possible, only in the simulated flow-artifact energy objective), all of the objective ranks had highly significant correlations with the subjective ranks, with the aliasing-artifact objective being the strongest and the flow-artifact objective being the weakest. A similar trend was observed across groups except that the flow-artifact objective was not significantly correlated.




Figure 6.2 Images showing significant improvements in the subjective scores for the optimal versus the standard trajectories. The X1 images were acquired using the optimal trajectories while the X2 images used the standard trajectories. The neck images are labeled A, simulated images are labeled B, and the cardiac images are labeled C and D. Note the decreased levels of background artifact in image A1 relative to image A2 and in image B1 relative to B2. The cardiac images show decreased background artifact (gray arrows) and aortic flow artifact (black arrows) in C1 and D1 relative to C2 and D2 respectively. Additionally, D1 does not show the slight increase in blurring seen in C1; D1 is optimal with respect to off-resonance while C1 is optimal with respect to flow artifact.

The data was combined across raters as described above for the remainder of the analysis. The null hypothesis that the optimal trajectories did not have better subjective scores than the standard trajectories was rejected for the aliasing objective in the simulated images, for the flow-artifact objective in the neck, heart, and simulated images, and for the off-resonance artifact objective in the heart images. Some of these images are shown in Figure 6.2. The null hypothesis was not rejected in the head images for any of the three simulated objectives. That is, the head images acquired using the optimal

Table 6.2Summary of logistic regression R^2 values by location

Case	Aliasing Energy	Flow Artifact Energy	Off-resonance Artifact Energy
All ROI ^b	.857	.053	.186
Simulated images ^b	.996	.987	.446 ^a
Background	.828	.005 ^a	.001 ^a
Brain	.845	.014 ^a	.000 ^a
Skull	.754	.158 ^a	.056 ^a
Carotid	.958	.032 ^a	.011 ^a
Coronary	.708	.003 ^a	.313
Aorta	.552	.000 ^a	.247

^a F statistic is not significant at 0.05 level (one-tailed)^b Regressions used only ratings from the corresponding trajectory group

trajectories were not considered better than the head images acquired using the standard trajectories.

Table 6.2 is a summary of the results from the logistic regressions. Across all ROI, we are 95% confident that improvements in all three of the simulated objectives were correlated with improvements in the subjective ratings, with the simulated off-resonance objective having the strongest correlation and the flow-artifact objective having the weakest (but still significant) correlation. Improvements in the simulated aliasing energy objective were significantly correlated with improvements of the subjective scores in all cases, and the effect was always at least twice as strong as the effect of any other simulated objective. Overall it accounted for 85.7% of the variance in the subjective image quality measure. The aliasing objective was the only objective with a significant correlation to the subjective scores in the head, in the neck, and when raters were instructed to consider only the background. The simulated flow-artifact energy objective

was only correlated with the subjective ratings across all ROI and for the simulated images. Improvements in the off-resonance artifact energy objective were correlated with improvements in the subjective scores both across all ROI and in the heart, accounting for 31.3% of the subjective-score variance for the coronary artery ROI and 24.7% of the variance for the aorta ROI.

6.4 Discussion

6.4.1 Confounding effects

One advantage of using the standard DSCQS method in evaluating MR images acquired using different trajectories, is that it is fairly robust and relatively well understood (International Telecommunication Union 2002). Despite the difficulties and complications possible in such subjective image-quality rating experiments, careful application of this method resulted in very few significant confounding effects. The one exception noted was a small, but statistically significant, decrease in the average score for images displayed on the bottom row of the GUI. Due to the very large number of ratings ($N = 2106$) even small differences are likely to be statistically significant. Additionally, the fact that there were more occurrences of the two worst trajectories on the bottom row indicates that the difference may be a real average difference in image quality rather than a measurement bias.

One other potential confounding effect common to these types of experiments is inter-rater variability. The training session is an important mechanism for reducing this variability and is required for the DSCQS method. This particular experiment also utilized ROI and associated questions to help direct the raters' attention towards the same part of each image, thereby reducing the inter-rater variability. Even with these features,

there remains a large amount of inter-rater variability. The DSCQS method compensates for this remaining variability by aggregating scores using the mean across all raters and by screening the raters. The principal component analysis of the raw ratings lends further support to the use of the mean scores as a method for normalizing or aggregating the scores from different raters. The only extracted factor was essentially the same as these mean scores obtained with the DSCQS method, and the lack of other significant factors indicates that the remaining intra-rater variance was essentially unstructured noise. Thus, the reduction and compensation procedures seem to have been effective in avoiding confounding effects due to inter-rater variability for this experiment.

6.4.2 Objective effects

The Chi-square and Spearman's rank-ordered correlation tests are attractive in this application because they require minimal processing of the data. Also, because the images were only ranked within the set of images visible on a single screen, these metrics would have been insensitive to any temporal effects that may have occurred. These tests are robust in the sense that they make fewer assumptions about the data (e.g. linearity of the regression), but by converting scale data to ordinal data they can under-estimate the true strength of a regression. Even so, they were able to clearly indicate two trends that were repeatedly uncovered in the remainder of the analysis. Specifically, all of the objective measures corresponded with the subjective ratings of image quality under certain circumstances, especially when images were compared only within each group of trajectories. Also, the aliasing objective was found to have the strongest and most consistent effect on image quality; while the flow objective was found to have the weakest and least consistent effect.

The only regressions that indicated the simultaneous success of all three objectives were the regressions using images from all ROI. These regressions only used ratings from images acquired with trajectories in the corresponding group (e.g. only images acquired with trajectories from the flow-artifact group were considered for the regression using the flow-artifact objective as a logistic predictor of the subjective scores). This is essentially a method for reducing the confounding influence of the other objectives; recall that the groups are designed to vary as much as possible in the corresponding objective while varying as little as possible in the remaining objectives. When considering the flow-artifact objective, for example, the effect of the aliasing and off-resonance objectives is to add non-correlated variance. This effect is also related to the difference between the significance of the Spearman's rank-order correlation "within groups" and "across groups" for the flow-artifact objective. This is particularly problematic for the flow and off-resonance objectives due to the strength of the aliasing objective effect on the subjective ratings. Although it is not possible to completely eliminate the variation in the other two objectives, by considering only the objectives within a group the uncorrelated variation is at least somewhat attenuated. With this attenuation it was possible to detect that, in general, improvements in all three objectives corresponded with improvements in the subjective measure of image quality.

The simulated images were included in this experiment in order to investigate whether or not humans would make the same choices as the computer when using the same image. This regression was very strong for both the aliasing and flow objectives, but was not significant for the off-resonance objective despite a reasonable proportion of variance explained (44.6%). This is likely due to both the small number of observations

(N = 6) for a single ROI considered within groups, and a restricted range problem. The subjective scores for simulated images acquired using the aliasing group of trajectories ranged from 0.06 to 0.93 (on a scale from 0.0 to 1.0), and the scores for the flow group ranged from 0.25 to 0.75, but the scores for the off-resonance group only ranged from 0.57 to 0.90 in the subjective scores. Considering the success of the off-resonance objective in other regressions and the R^2 of .446 in this regression, it is reasonable to hypothesize that the regression might become significant after resolving these problems.

In cardiac imaging fairly small and important structures, such as the coronary arteries, can be difficult to image due to susceptibility effects, movement, and the nearby presence of fat. The off-resonance objective, though not as universally important as the aliasing objective, was still quite important for cardiac imaging. It is somewhat surprising that the flow objective was not determined to be significant in the heart, but that could be due to the fact that the flow-artifact objective was computed using only in-plane simulated flow. Although motion certainly does occur in plane, it may be less than the through-plane motion, and the simulations would not detect such through-plane effects (see Appendix C). On the other hand, even with the aforementioned restricted range problem, small differences in the off-resonance artifact objective were correlated with blurring artifacts that could degrade the visibility of the coronary artery and the aortic root. For the coronary artery ROI, the off-resonance objective accounted for 31.3% of the total variance, the most for any objective besides aliasing when considered across groups.

The most important objective was the aliasing energy objective. It was highly significant and explained between 55.2% and 99.6% of the total variance in each case, 85.7% of the variance overall. When raters were asked to focus on the background, it

was the only objective that seemed to have any impact on the subjective scores. This is most likely due to the fact that the off-resonance and flow effects tend to generate more localized artifacts such as blurring or streaking while the aliasing typically causes extra signal to appear at a distance determined by the Nyquist criterion. The aliasing objective was also the only successful objective in the neck images. The carotid arteries have relatively rapid flow, but the direction is through plane with the same results as described for the heart. Off-resonance effects are also fairly limited in the neck due to its small size and more uniform susceptibility (relative to the head and chest). Thus it appears that the consistency of the aliasing effect is due to the fact that it is not dependent on the presence of a physical quantity (flow velocity or off-resonance frequency) in order to appear in an image. Due to the strength and consistency of this effect, it is important that all future optimization work include aliasing energy as either an objective or a constraint.

6.4.3 Conclusion

In conclusion, the simulated objectives utilized previously in the optimal design of k-space trajectories have been validated in subjective image-quality experiments, lending credence to both the method and the results of that optimization. The aliasing energy objective, in particular, has such a strong and consistent effect that it should be included in any future optimization work. Finally, in demanding applications with strong physical effects, such as cardiac imaging, it is important to include other suitable objectives to reduce the artifacts caused by these effects. The optimization of such objectives during the trajectory design process has been shown here to result in significant improvements in subjective ratings of image quality.

Chapter 7. Optimal True-FISP Pulse Sequences

7.1 Introduction

One common goal of active MR research is the design of pulse sequences with reduced acquisition times. Improving MR acquisition speeds results in reduced motion artifacts as well as increased utilization of costly MRI systems. Short acquisition times require rapid gradient switching which can lead to PNS (Irnich and Schmitt 1995; Price 1999). As a result, PNS effects, rather than gradient hardware, frequently limit the acquisition speed of pulse sequences. Because of its inherent complexity, few, if any, optimization studies have incorporated a PNS model into pulse sequence design (Simonetti et al. 1993; Reeder et al. 1999). Therefore, to date, reducing the acquisition time of a particular pulse sequence involves a tedious, iterative process involving adjustment of gradient lobe parameters until the PNS limits are approached but avoided overall.

Automating the pulse sequence optimization process requires a model capable of searching through many pulse sequence parameters with multiple objectives in a timely fashion. GAs are a class of optimization techniques that are capable of finding global minima in the face of highly non-linear or otherwise ill-behaved objective and constraint functions (Deb 2001). They use a biological metaphor where sequence parameters are encoded in a numerical “gene” and then sequences compete according to one or more objectives to pass that genetic information on to future generations. The population of sequences thus evolves to an optimal set after multiple generations.

In this study, a multi-objective GA incorporating the PNS model is used to optimize several variations of a True FISP pulse sequence (Duerk et al. 1998). Image quality,

resolution, and acquisition speed are three important considerations in most applications. This work sought to improve all three by simultaneously minimizing the TR, FOV, and the BW. Because FOV, TR, and BW are conflicting objectives, there will not be a single optimal sequence but rather a set of sequences, each representing the lowest possible TR for a given BW and FOV. Such “best trade-off” sequences are known as Pareto-optimal, and finding them is the goal of multi-objective optimizations (Deb 2001).

7.2 Methods

The GA utilized in this work was the NSGA-II first proposed by Deb for general work in optimization and later used by Dale for work specifically looking at MRI pulse sequence design (Deb et al. 2002; Dale et al. 2003a; Dale et al. 2003c; Dale et al. 2004). The NSGA-II is capable of simultaneously optimizing multiple objectives in the presence of ill-behaved objective or constraint functions. Here, the population size was set to 50 and the optimization was run for 500 generations resulting in a total of approximately 25,000 sequences tested.

Parent sequences were selected using a binary tournament selection, where two random sequences “compete” and the more fit of the two is then used in the simulated genetic operations. As in biology, the fitness for reproductive selection in a GA can be different from the fitness for survival. In this case, survival fitness is based on crowding and optimality, while the reproductive fitness was based on TR unless the BW was greater than 1000 Hz. The combination of these selection pressures is designed to push the population of sequences towards a set of the most rapid possible, best trade-off, sequences spanning the range of BW from the minimum possible to 1000 Hz. Two child sequences were produced from each pair of parent solutions using simulated binary

crossover ($\mu = 1$), which is a commonly used computational analog of genetic recombination. This was followed by Gaussian mutation ($\sigma = 4$), a common computational analog of genetic mutation, on an average of one parameter per child (Deb 2001).

In this study, the acquisition time of both a single-echo and a dual-echo True FISP pulse sequence were optimized for use on a Siemens Sonata 1.5 T scanner. Therefore the slew rate was constrained to be less than 200 T/m/s and the gradient magnitude on each axis was constrained to be less than 40 mT/m. In addition to the above hardware constraints, the sequence was also constrained to avoid causing PNS as determined by the SAFE model (Hebrank and Gebhardt 2000). The model uses two or more first-order digital low-pass filters to approximate the generation of action potentials within the nerve cells and the spread of the signal via synapses. The SAFE model takes all three gradient waveforms as input and returns a single stimulation waveform where values greater than 1.0 indicate that a PNS has occurred, and it is used on the Siemens Sonata to monitor the gradient waveforms and halt the sequence in the event of a predicted stimulation. All constraints were enforced using the standard constrained non-dominated sorting procedure (Deb et al. 2002).




Figure 7.1 Two true-FISP pulse sequences. All gradient axes are completely refocused at the end of each repetition. The single-echo sequence has eight gradient lobes while the dual-echo sequence has nine.

The single-echo True FISP sequence has eight gradient lobes while the dual-echo sequence has nine (Figure 7.1). The dual-echo sequence was designed to be a 2-Point Dixon acquisition for rapid fat suppression (Dixon 1984; Coombs et al. 1997). Note that the only free parameter of the 2nd readout gradient lobe was the ramp down time, and that the start time, duration, and ramp-up times were determined by the 2.2ms delay between the two same-BW gradient echoes required for fat suppression. The phase encode gradients of this multi-echo model could easily be adapted to cover different lines of k-space as in EPI sequences or 1-Point Dixon acquisitions (Flask et al. 2003). All gradient lobes were designed as trapezoidal gradient pulses. Nonlinear gradient waveform ramps could be included in the GA, but were not considered for this study (Dale et al. 2003c; Dale et al. 2004).




Figure 7.2 Timing parameters. Each trapezoidal gradient lobe can be completely described by the four indicated timing parameters. Once the timing is fixed, the gradient amplitude is determined by the FOV.

Each trapezoidal gradient lobe was specified using 4 timing parameters: Delay, Ramp Up Time, Flat Top Time, and Ramp Down Time (Figure 7.2). GAs are capable of accepting discrete parameters; therefore, the individual timing parameters were constrained to integer multiples of the hardware gradient raster time (10 μ s). The FOV was included as an additional parameter, resulting in a total of 33 parameters for the single-echo optimization and 34 parameters for the dual-echo optimization. The constant sequence values used during the optimization were a fixed matrix (256^2), slice orientation (axial), slice thickness (3 mm), and BW-time product of the RF pulse (1.6). Note, once the timing parameters are set, the amplitude of each gradient lobe can be uniquely determined from the desired FOV, slice thickness, RF pulse, and/or Readout BW and is therefore not a free parameter. All timing parameters were greater than or equal to zero and the FOV was constrained to between 50 mm and 300 mm.

7.3 Results

For three-objective optimizations, such as these optimizations, the Pareto-optimal set is usually a curved surface in the objective space. Recall that Pareto-optimal for this work means that each sequence represents the shortest TR possible (without causing stimulation) at a given combination of BW and FOV. These Pareto-optimal surfaces are shown in Figure 7.3 by plotting BW vs. TR curves for several different FOVs for each optimization. During both single echo and dual echo optimizations the sequence populations generally had more individuals with higher FOVs, and the final curves corresponding to higher FOVs also tended to be smoother.




Figure 7.3 Optimal pulse sequences. Each point represents a single pulse sequence plotted with the three objectives. All plotted sequences are optimal, the fastest sequence possible for a given BW and FOV. Note that the FOV curves are shown in steps of 100 mm for the single-echo and 50 mm for the dual-echo case.

Both optimizations generally exhibited a “diminishing returns” type of behavior where successive improvements in one objective came at progressively greater sacrifices

in another objective. This was observed by noting that the FOV curves are progressively more strongly shifted to the right, and that each curve generally becomes progressively steeper on the left than on the right. Note the relatively hard floor on the BW for the dual-echo optimization.

For the single-echo true-FISP the larger FOV sequences were almost universally slew-limited. The shorter TR, large FOV sequences also tended to be PNS limited. PNS limited sequences had stimulation curves that usually approached the safety limit at four time-points: before and after the excitation and before and after the readout. The importance of the slew constraints was not as significant for smaller FOV sequences. Instead, most small FOV sequences were PNS limited with the shorter TR, small FOV sequences also being gradient amplitude limited.

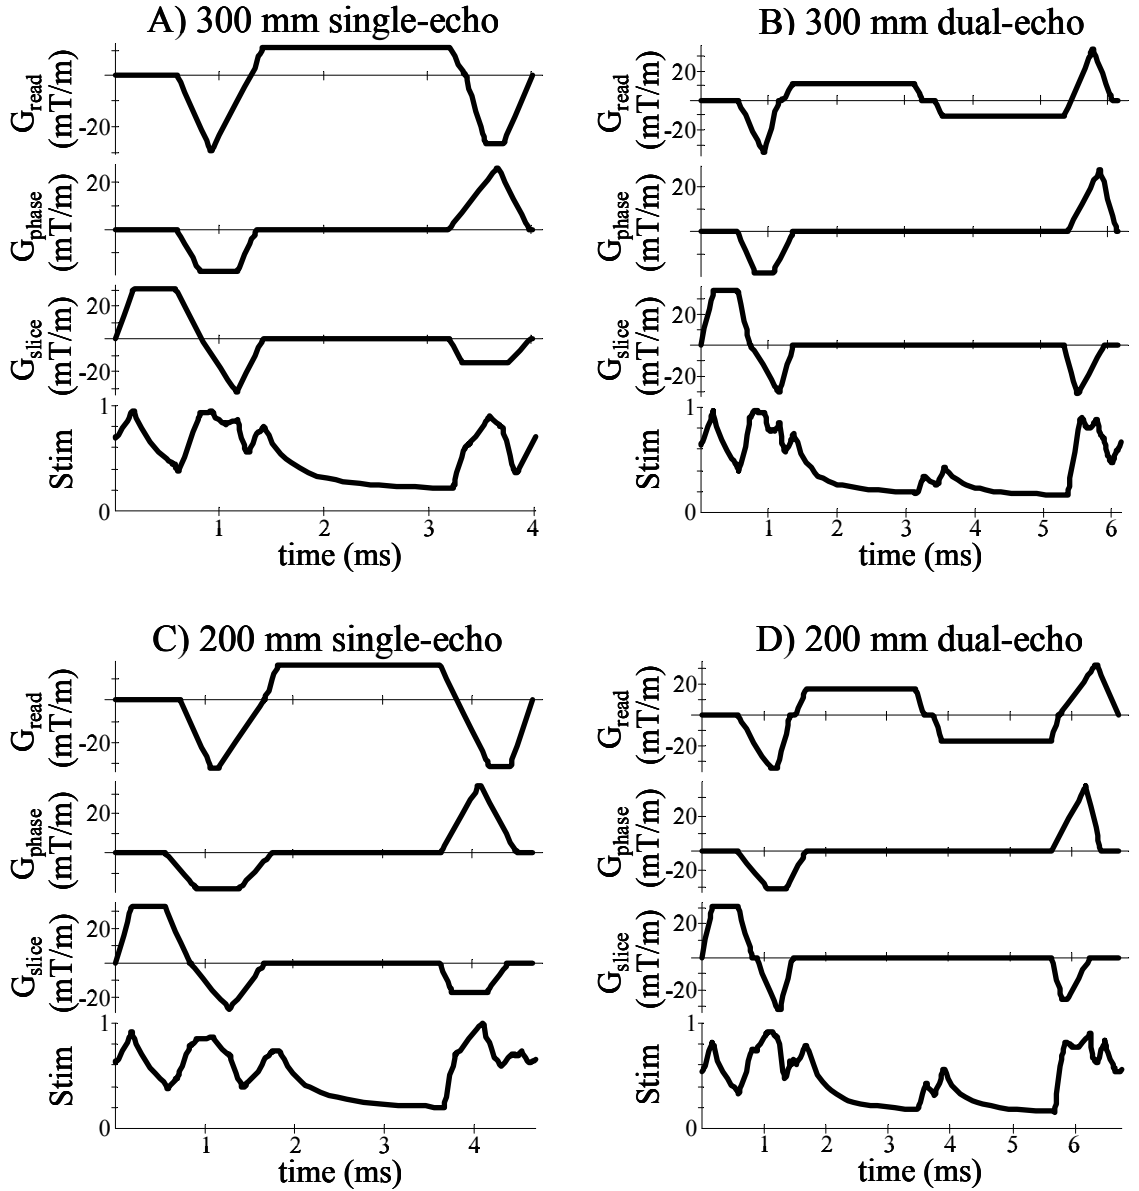


Figure 7.4 Selected optimal true-FISP pulse sequences. A) single echo, FOV = 300 mm, TR = 4.0 ms, BW = 560 Hz, stim = 96%. B) dual echo, FOV = 300 mm, TR = 6.1 ms, BW = 560 Hz, stim = 97%. C) single echo, FOV = 200 mm, TR = 4.7 ms, BW = 550 Hz, stim = 99%. D) dual echo, FOV = 200 mm, TR = 6.8 ms, BW = 560 Hz, stim = 92%

The behavior of the constraints was different for the dual-echo pulse sequences. The stimulation constraints were active across the entire population, regardless of FOV, TR,

and BW. Typically the stimulation approached the limit only between the two readouts for small FOV sequences, while the larger FOV sequences would approach the limit before and after the excitation and after the second readout. The slew-rate constraints were also active for both the largest and smallest FOV sequences across all TRs, while the gradient amplitude constraints were active for the shorter TR sequences across all FOVs.

Table 7.1 Summary of test and standard sequence parameters

Sequence	# Echoes	FOV (mm)	TR (ms)	TE (ms)	BW (Hz)	Max Stim
A	1	300	4	2.3	560	0.957
B	2	300	6.1	2.3/4.5	560	0.974
C	1	200	4.7	2.7	550	0.999
D	2	200	6.8	2.6/4.8	560	0.917
Standard	1	300	4.7	2.5	560	0.838

The selected Pareto-optimal pulse sequences are shown in Figure 7.4 and compared with a standard sequence in Table 7.1. Sequences A) and B) are 300 mm single- and dual-echo sequences respectively, and sequences C) and D) are the 200 mm single- and dual-echo sequences. Each plot displays the three gradient axes as well as the stimulation waveform generated by the SAFE model. Note the details of the pulse sequences and their impact on the stimulation waveforms. For example, in B) and D) note how the high slew-rate portion of each of the final refocusing pulses are somewhat staggered. In both cases, it appears that the stimulation waveform becomes truncated due to these





Figure 7.5 Axial abdominal images acquired with FOV = 300 mm and BW = 560 Hz. The standard sequence requires 18% more time than the optimal sequence, for the same BW and FOV. The two single-echo sequences have the same noise levels, while the dual-echo sequence requires 53% more time, has less noise, and more uniform fat suppression.

asymmetric gradient lobes. Note also that most of the pulse sequences have asymmetric waveforms on the PE axis with the first PE lobe being trapezoidal and the second being triangular.

Images of a volunteer's abdomen (FOV = 300 mm) were obtained for optimized single-echo and dual-echo true-FISP sequences as shown in Figure 7.5. Abdominal images were also acquired with a conventional, sub-optimal true-FISP sequence for comparison. The sub-optimal (Fig. 7.5a) and optimal (Fig. 7.5b) single-echo true-FISP images show comparable image quality and level of artifacts. Aside from the TR (4.7 ms), the parameters of the sub-optimal sequence were equivalent to the optimized single-echo pulse sequence. The optimized dual-echo image is also very similar to the single-echo image, but with uniform fat suppression throughout the abdominal region.



Optimal single-echo
4.7 ms, 200 mm, 550 Hz

Optimal dual-echo
6.8 ms, 200 mm, 560 Hz

Figure 7.6 Axial head images acquired with FOV = 200 mm. The 200 mm optimal single-echo sequence has the same TR as the 300 mm standard single-echo sequence used above. The dual-echo sequence requires 45% more time, has less noise, and more uniform fat suppression.

The optimized true-FISP sequences were also used to obtain images of volunteer's head/optic nerve (FOV = 200 mm) as shown in Figure 7.6. Multiple signal averages (NSA = 5) were obtained to improve the image quality. The single-echo images were reconstructed on-line. The dual-echo sequences were reconstructed offline incorporating an off-resonance correction algorithm resulting in uniform fat suppression (Coombs et al. 1997).

7.4 Discussion

The GA provided multiple, Pareto-optimal true FISP sequences that would have been very difficult to design by conventional means. Typical pulse sequence design involves an iterative process where the gradient pulse timing and shapes are varied until the TR is

at an acceptable level without violating the stimulation limits. Even when this process is successful, small changes in sequence parameters during imaging applications often lead to PNS violations. Prior to this optimization work, PNS violations caused by small parameter changes were particularly problematic during the development of the dual-echo sequences. The Pareto-optimal solution sets provided by the GA overcome this limitation and allow the sequence to remain time-optimal for a variety of specified imaging conditions.

The NSGA-II was able to converge to a reasonably well-defined Pareto-optimal set within 500 generations in both the true-FISP and the dual-echo true-FISP cases. The greater smoothness of the large-FOV curves suggests that these regions may have converged more completely than the small FOV regions. This is further substantiated by the observation that, during the optimization, the sequence populations tended to have a higher proportion of large FOV individuals. This may be a direct result of using a random initial population, because random small-FOV individuals are more likely to be severely penalized for hardware constraint violations. An alternative to using a random initial population would be to design an initial population based on product pulse sequences.

Another potential method to improve convergence would be to use a hybrid GA where both genetic optimization and more traditional “hill climbing” optimization routines are used together (Lahanas et al. 2003). Such algorithms have been shown to have faster and more complete convergence properties than pure GA’s, while being more able to escape from local minima than pure traditional optimization routines.

The relatively hard floor on BW for the dual-echo optimization is due to the fact that, even with an infinite slew rate, there is a fundamental minimum BW of 455 Hz for symmetric gradient echoes separated by exactly 2.2 ms. With limited slew rates, this BW floor is a function of the FOV as can be seen in Figure 7.4. Such features can be discovered automatically using a formal optimization without requiring specific encoding of the theoretical constraints.

The constraint activity has important implications for hardware design. Almost all of the Pareto-optimal pulse sequences were stimulation limited, with the exception of the low-BW/large-FOV single-echo sequences. This implies that improvements in hardware would only allow increases in acquisition speed if they were coupled with mechanisms to maintain or reduce stimulation levels. Without such mechanisms sequence designers will only be able to make improvements by sliding along the PNS constraint surface, and thus sequence-design methods like the one presented here will become more important.

Perhaps one of the most important advantages of a multi-objective GA is that it is possible to sample the entire Pareto-optimal set in a single run. This allows for a more complete understanding of the inherent trade-offs amongst the various objectives. The alternative is usually to do a single-objective optimization by assigning weighting coefficients to the various objectives. Such weighting coefficients are notoriously difficult to obtain with any degree of confidence and are generally different for each application. Instead, by using the multi-objective approach, the optimization can be performed once and the results may be used for any application without repeating the optimization. This is important because the 500 generations used here required approximately 12 hours of computation time (mostly for evaluating the SAFE model). It

is therefore neither feasible nor necessary to repeat the optimization on-line during the selection of sequence parameters, as would be required with the weighting methods.

There are several potential methods for obtaining the best sequence for a particular application from a table of all of the Pareto-optimal sequences without requiring the specification of weighting coefficients. The simplest involves plotting the trade-off curve for two objectives while fixing the remaining objectives and selecting the best trade-off for a particular application. This procedure can be repeated for other pairs of objectives until the best sequence is obtained. However, it becomes progressively more cumbersome with increased numbers of objectives. In such cases, a suitable alternative is to use a clustering algorithm to select a small number of representative sequences (Deb 2001). The best one for the application is selected and the corresponding cluster is repeatedly sub-divided into a similar number of sub-clusters until the single best sequence is obtained.

In conclusion, the NSGA-II was able to successfully converge and find minimum-TR, PNS-limited, true-FISP pulse sequences for a wide variety of combinations of BW and FOV without requiring selection of weighting coefficients. The multi-objective GA was able to converge despite the non-linear nature of the SAFE model and the discrete timing parameters used. This technique solved some of the challenges encountered when working with new pulse sequences such as the dual-echo true-FISP, and the inclusion of the FOV as an objective successfully overcame one of the fundamental limitations of previous work. We believe that techniques similar to those used here will prove useful for rapid refinement of a wide variety of new techniques developed in the future.

Chapter 8. Summary

This work has resulted in the development of a wide variety of techniques and methods for optimal design in the field of MR image acquisition. It has been motivated by the immense complexity and flexibility of MRI itself. Optimally precise DCE-MRI protocols, optimal k-space trajectories, and optimal true-FISP pulse sequences have been developed along with the means for fluoroscopic online gridding-reconstruction of non-rectilinear k-space data. This not only represents the accomplishment of all of the original specific aims of this project, but also the additional demonstration of the utility of the multi-objective GA optimization for more traditional types of sequences with rectilinear k-space trajectories.

8.1 Significance

8.1.1 Current Results

Several of the specific results are of immediate and practical application in the MR imaging community. For example, DCE-MRI is widely used as a non-invasive method for scientific studies that need to measure tissue perfusion *in-vivo*. Most such studies involve human clinical trials of new drugs, such as the testing of anti-angiogenic agents for safety and efficacy in starving cancer tumors. By improving the precision of the DCE-MRI measurements researchers can obtain the same information with a smaller number of samples. In this case, a smaller number implies that fewer patients are placed at risk by undergoing an experimental treatment. Even in studies where patient risk is non-existent, a smaller number always represents a smaller cost to gather the same information.

The time-optimal trajectories for k-space transfers have a wide potential for application in many high-speed applications, especially for any type of EPI application. For example, the SPIDER pulse sequence, an asymmetric radial true-FISP EPI pulse sequence, has shown some potential in cardiac imaging; however off-resonance and motion effects occasionally limit its effectiveness (Larson and Simonetti 2001). The application of the time-optimal trajectories to the SPIDER sequence can result in a 30% reduction in overall echo-train duration, or an increase in azimuthal separation from 1.4° to 18° . This should attenuate the off-resonance or motion effects and thereby improve the effectiveness of a promising technique. Because the time-optimal trajectory solution is general, for any initial and final gradient state, its application is not limited to the SPIDER sequence, but should provide similar reductions in sensitivity to off-resonance and motion effects for any multiple gradient echo technique. Such rapid sequences are particularly important in cardiac imaging, iMRI, and fMRI. Better, faster cardiac and iMRI images could allow for improvements in patient care, while better, faster fMRI images could help in a wide variety of scientific projects.

The new GA-designed Pareto-optimal trajectories represent a significant improvement in image quality over current standard trajectories at higher numbers of interleaves. Also, the unexpected presence of a consistent feature in the GA-designed optimal trajectories, the non-zero initial gradient amplitude, is a surprising discovery that may have further applicability. Specifically, it could be applied to current WHIRL and spiral trajectories, and may result in improved image quality. Such a potential general improvement may have never been discovered without the use of a formal optimization procedure such as that used here.

Although many techniques exist for correcting badly corrupted images, most of them require a great deal of post processing and are therefore not suitable for fluoroscopic applications such as iMRI. Rather than spending a long time in post processing, these trajectories allow for the improvement of image quality by investing a relatively small amount of time in the acquisition itself. The use of k-space trajectories in iMRI settings, such as device guidance, is particularly promising. The resulting improvement in temporal resolution and image quality over current techniques could reduce the total duration of most interventional procedures, which is often highly correlated with clinical outcome.

In addition to the new trajectories developed, one important result from the multi-objective trajectory optimization is the understanding of the relationship between spiral, WHIRL, and other standard trajectories. If an extremely high temporal resolution is required, and image quality is secondary, then spirals with low numbers of interleaves are the trajectory of choice. As the need for speed is reduced and the desire for image quality is increased, then spiral trajectories are no longer the best choice. While it has long been intuitively understood that there is an inherent trade-off between acquisition speed and image quality, the details of that trade-off have never been carefully described. This work now gives researchers a rationale for making specific choices in the use of particular trajectories for their specific application.

Finally, the optimization of the rectilinear true-FISP pulse sequences is expected to have much more widespread impact due to the more common use of rectilinear k-space acquisitions. In particular, the dual-echo pulse sequence optimization is particularly important. While true-FISP and Dixon techniques are both well-accepted techniques, the

typical implementation is to perform two separate single-echo acquisitions with different TEs and then combine the images. By placing both echoes in a single acquisition, the total imaging time is reduced by a factor of 2. This idea holds significant promise for a wide variety of rapid fat-suppressed applications, but was difficult to implement in a time-efficient manner without risking PNS. Early development of the dual-echo technique relied on a cumbersome, iterative, guess and test procedure for finding workable sequence parameters. Small changes in the sequence often resulted in unexpected violations of the PNS constraints. With the multi-objective GA optimization results, the parameters can easily be adjusted to meet the requirements of a particular application while both maintaining optimality and avoiding PNS. Essentially, the rectilinear dual-echo true-FISP technique has undergone a very rapid transition from a new to a mature technique.

8.1.2 Future Impact

Although these are each important improvements, the larger potential impact of this project lies, not in these specific results but in the continued application of the methods developed here. In particular, the multi-objective GA optimization shows immense promise for improving a wide range of MRI acquisition techniques. One major limitation of previous work in optimal design has been the usual single-objective approach. This approach is seldom appropriate in MRI due to the fact that there are almost always multiple important objectives and the objectives are usually conflicting. The multi-objective approach acknowledges that fact, and attempts to completely uncover the true nature of the inherent trade-offs that must be made between the various objectives.

One other limitation of previous work is the requirement of most optimization methods that the objective functions and constraints be well-behaved and even differentiable functions of the parameters. This often requires so many simplifying assumptions, and/or poor metrics of image quality, that the results can seem less important even after a well-done optimization. The use of a GA addresses this concern by using an algorithm that makes no requirements on the behavior of the objective or the continuity of the parameters and constraints. This allows discrete or otherwise discontinuous objectives and parameters, multi-modal objectives, disjoint feasible regions, and simulation-based objectives with arbitrarily complicated models. Although simplifying assumptions were made here, they were all related to reducing the computation time of the objective calculation, not to requirements of the optimization algorithm itself. This is an important distinction because the multi-objective GA method will be able to take immediate advantage of improvements in computer hardware to compute more realistic and useful objectives, instead of simply optimizing a poor objective in less time.

By addressing these two concerns and limitations of previous work, it is hoped that this technique will gain widespread acceptance and use in the design of MR image acquisitions. MRI is increasingly moving from being strictly a subject of research, to also becoming a tool of research. This means that there are an increasing number of people trying to use MRI to answer a question that has never been asked before. Usually these scientists are experts in the subject of their particular research question, not in MR physics and pulse sequence design. However, because the topic usually has one or more novel components, it is common for there to be no pulse sequences that specifically

address their particular requirements. Future extensions to this optimization work may help to move sequence design from the MR physics domain to the problem domain, thereby allowing such researchers to use their own specific expertise and general scientific skills to best advantage.

The other anticipated benefit from widespread use of this method would be the rapid maturation of new techniques. In other words, techniques that have been around for a long time, such as single-echo true-FISP, may see only a small, incremental improvement through the application of a formal optimization procedure. Such techniques are already mature in the sense that there is a large body of experience using them and most improvements that can be made will already have been discovered through the course of the many man-years spent in developing them. However, as our experience with the dual-echo true-FISP sequence indicates, even new techniques that seem quite similar to well-established sequences can have surprising difficulties and challenges. Instead of spending several man-years to refine these techniques, explore the obstacles, and discover the possible improvements through trial and error, the use of optimal design techniques allows the same improvements by spending a little time in careful preparation and a few computer-hours of exploration and refinement. In this sense, general optimization methods like the multi-objective GA allow the rapid maturation of new techniques, which in turn should help achieve an elevated rate of improvements in patient care.

8.2 Future Directions

In order to attain the desired future impact described above there are a number of necessary improvements, each relating to one or more specific facets of this project. This

penultimate section addresses some of the individual steps that should be taken next in order to realize the full potential of these methods.

8.2.1 DCE-MRI Precision

As mentioned in Chapter 2, one significant problem with the Monte-Carlo method is that it is inherently limited by the assumptions made in generating the random signal. The propagation of errors optimization therefore lacks some of the sense of validity achieved through the experimental validation of the trajectory optimization. However, in addition to the current lack of a perfusion phantom that motivates the use of Monte-Carlo methods in the DCE-MRI literature, the validation of the optimally precise protocol presents some extra experimental challenges. Specifically, because statistical tests on variance are so much less powerful than tests on the mean, it is absolutely essential that any attempt for experimental validation take care both to reduce the variability of the measured perfusion and to allow the collection of a large number of samples.

There are at least two methods that show some promise for the experimental validation of the optimal DCE-MRI protocol. The first would be to use a hollow-fiber hemodialysis filter, or dialyzer, as a perfusion phantom. The hollow-fiber dialyzer is the most similar manufactured product to the function of natural blood vessels. Gd-DTPA is known to pass from the blood to the dialysate during the course of normal hemodialysis just as it passes from the blood to the tissues during DCE-MRI (Niendorf et al. 1996; Niendorf et al. 1997; Ueda et al. 1998; Okada et al. 2001). In addition, because of the manufactured nature and the rigorous standards for such filters, the “natural” variance component should be quite low. Also, with the ability to completely and rapidly flush the dialyzer, it should be possible to repeat the experiment a large number of times in

relatively short succession. There are, however, some very important differences between a hollow-fiber dialyzer and normal tissues. First, the dialyzer may not be MR compatible. This should not be too difficult to remedy because the hollow-fiber component itself is manufactured from a cellophane resin, and it would only be the housing or other components that may need to be replaced. Second, the “blood” represents a much larger portion of the total volume in a dialyzer than in normal tissues. This could potentially be resolved through the removal or occlusion of a portion of the fibers. Third, the geometry of the dialyzer is such that there is likely to be significant concentration differences between the inlet and outlet sides of the fibers and therefore a significant component of diffusion along the fibers in the “tissue” compartment. A sufficiently rapid flow through the hollow fibers could reduce or eliminate this effect as well as potentially helping with the second concern. Finally, both the extracellular volume fraction and the clearance of Gd-DTPA (both directly related to K^{trans} and k_{ep}) are much higher in a dialyzer than in any kind of living tissue. These last two problems should not diminish the ability to use the hollow-fiber dialyzer as a perfusion phantom, but they will mean that the phantom will be quite different from any type of normal tissue. Because the results of the optimization are dependent on both the expected tissue and blood $T1_0$ and on the expected K^{trans} , the protocol described above will generally not be optimal for the dialyzer. Instead, it would probably be best to repeat the optimization for the dialyzer parameters and compare the dialyzer-optimal protocol to the standard protocol rather than comparing the tissue-optimal protocol to the standard protocol. If the tissue-optimal protocol is desired, then it will be essential to do the error propagation using the dialyzer parameters and the tissue-optimal protocol to determine the anticipated

change in precision. Even if the dialyzer experiments cannot be used to validate the specific optimal protocol, they can still be used to validate the propagation of errors method itself. If the propagation of errors is experimentally verified, then the optimization based on the propagation of errors will be at least partially validated.

The other method that has some promise in experimental validation of the optimal DCE-MRI protocol is the use of artificially perfused excised organs. The experimental setup would be similar to that of a Langendorff perfused heart. This setup would allow for the determination of both the arterial input and venous output components and therefore permit the determination of the overall perfusion with some, hopefully high, degree of accuracy and precision. By taking the difference between this measure and the MRI determined measure we should be able to obtain a fairly good measure of the variance of the DCE-MRI technique itself. The desired features of an organ for this experiment would be: a tough outer capsule to prevent changes in volume and the bulk accumulation of fluid, a single arterial input and venous output to simplify setup, a reasonably small size to perhaps permit multiple samples in a single run, and a single type of tissue with a high K^{trans} for increased precision and volume averaging. One potential candidate organ, although it contains two types of tissues, is the testicle. If it proves suitable for this type of preparation, it would have the additional advantage of being readily and cheaply available in large quantities due to common practices in the beef, pork, and lamb agricultural industries.

8.2.2 Genetic Algorithm Optimizations

8.2.2.1 Methods

This work clearly indicates that the multi-objective GA approach is a very powerful and general method for the optimal design of MR imaging techniques. The NSGA-II and the associated parameter settings used here, however, are not the only multi-objective GA or parameters available, and others should be investigated to determine convergence properties. In particular, a steady-state GA may prove useful. They avoid the use of generations and may have guaranteed convergence. Also, hybrid optimization routines, combining a GA and a deterministic hill-climbing optimization algorithm, are known to converge more rapidly and may even be able to converge more completely if care is taken to avoid becoming trapped in local minima.

8.2.2.2 Results

As mentioned above, there are at least two potential methods for obtaining the best sequence for a particular application from a table of all of the Pareto-optimal sequences without ever requiring the specification of weighting coefficients. The more practical, particularly for large numbers of objectives, is to use a clustering algorithm to select a small number of representative sequences. The best one for the application is selected and the corresponding cluster is repeatedly sub-divided into a similar number of sub-clusters until the single best sequence is obtained. This represents a very fast search. For example, there were almost 1200 Pareto-optimal trajectories developed, and this method could find the best after only five iterations with four clusters per iteration. For two objective sets, or by fixing the remaining objectives and considering only two at a time for higher dimensional sets, it may be reasonable to use the other method. In fact, the

two methods could be combined in order to easily allow small refinements of the selection determined by clustering, if experience indicates that the choice was not completely satisfactory. In either case, it would be quite useful to have one or both of these mechanisms available on the scanner for immediate investigation of any particular solution from the whole Pareto-optimal set, rather than the current system of selecting a few solutions at the PC for implementation on the scanner.

8.2.2.3 Constraints

The only major constraint not considered during this project is SAR. SAR should not be difficult to incorporate into any of the optimizations, and will become progressively more important as optimization and sequence development proceeds to 3 T and even higher field strengths. However, in addition to the SAR constraints, the current SAFE model for PNS constraints is not entirely satisfactory. It is quite cumbersome and computationally intense, and it also has some anisotropy that is not possible to remove simply by picking an inscribed sphere region similar to the approach of the time-optimal transfer method. This means that it is not possible to use the SAFE model, as currently implemented, to design sequences that are guaranteed to avoid stimulation in any orientation. It is probable that a simpler model, using fewer than the current 21 parameters used by the SAFE model and adaptable to such rotations, could fit the PNS data equally well. It would be desirable to have such a model for PNS, but its use may be limited by the fact that the SAFE model is likely to remain the model actually used by the Siemens MRI systems for the foreseeable future.

8.2.2.4 Parameters

The choice of parameters is a vital part of any optimization and can dramatically affect the results of any optimization routine by making certain types of solutions difficult or impossible to represent and discover. The Chebyshev series expansion of the gradient waveform has proved to be a very effective and general parameterization and it was absolutely critical that such a generality be utilized in the first experiences with trajectory optimization. The previous work done here using a variable-density spiral parameterization did not permit non-zero initial gradient amplitudes nor did it permit trajectories to begin at a non-zero k-space location (e.g. reversed spiral or radial). Optimizing using the general Chebyshev series expansion helped discover the importance of the initial gradient amplitude, but the general optimization also wasted some amount of time on testing sequences that did not sample the center of k-space; such sequences are generally not useful. The fact that the Pareto-optimal trajectories could all be represented with a slight generalization of the variable-density spiral parameterization is an important indication that future trajectory optimization work should probably utilize a generalized variable-density parameterization. Such a parameterization should have faster and more complete convergence properties, especially for low numbers of interleaves, without inadvertently discarding any useful solutions.

For the rectilinear and other pulse-sequence based optimizations, other specific additional parameters may prove useful. The gradient-timing parameterization seems to work well, but it is quite possible that non-linear ramp shapes may reduce PNS for the same change in gradient amplitude over the same amount of time. Such ramp-up and ramp-down shape parameters could be incorporated into the current optimization fairly

easily. In addition, there are many potential “binary” parameters in MRI, such as whether or not to perform fat-suppression, flow compensation, or magnetization inversion. These types of parameters may prove particularly important for some of the contrast-based objectives described in the next section.

8.2.2.5 Objectives

The objectives used here, particularly the simulation-based measures of artifact severity, were all validated using subjective rating experiments. However, there remains a lot of work that could be done to further refine and improve these objectives with the goal of designing objectives with the greatest possible impact on perceived image quality. This could be accomplished, for example, by using a perceptual difference model instead of the mean squared error as a measure of the degradation of image quality due to the influence of the physical effects. Also, using the perceptual difference model may allow the combination of multiple objectives, such as the aliasing and off-resonance objectives, into a single objective. This would be desirable due to the improvement in convergence that using fewer objectives usually provides. Additional artifact objectives that should be studied include through-plane flow and bulk motion between views.

The proper choice of objectives depends, in large measure, on the desired applications. For example, in scientific applications the multi-objective GA approach could be extended to other quantitative imaging applications, similar to DCE-MRI, where accuracy and precision would be a logical pair of objectives. Current quantitative techniques include ADC imaging, diffusion tensor imaging, flow quantification, myocardial stress/strain imaging, etc. If the new quantitative technique were not as well

established as DCE-MRI, then it may or may not be that accuracy and precision are conflicting objectives; either way the multi-objective approach will illuminate that fact.

Such objectives, however, would not be as useful for more traditional diagnostic or interventional imaging applications, where the usual objectives are acquisition time and one or more measures of image quality. This is a particularly important arena for future improvements. For example, contrast between blood and other tissue or between normal and pathological tissue is often a critical measure of image quality for many diagnostic and device-guidance applications. Specific kinds of contrast that could be optimized include dark-blood, bright-blood, fat-suppression, tumor-tissue contrast, vulnerable plaque, etc. Although the perceptual model may be more easily applied to artifact-based objectives, it could also be applied to such contrast objectives in order to maximize the desired perceptual difference between normal and pathological tissue. Image contrast was not optimized in either the k-space trajectory or rectilinear pulse-sequence optimizations, but probably represents the best and single most important future step for this work.

8.3 Conclusion

In conclusion, MRI is a very complicated and flexible imaging modality with an infinite variety of possible image acquisition techniques, and a corresponding variety in available information. The optimal design methods presented here are an essential step towards changing the design of MR image acquisition techniques from an art to a science. These initial optimizations have resulted in a four-fold improvement in the precision of quantitative DCE-MRI measurements of perfusion, as well as significant improvements in the quality of images acquired with non-rectilinear k-space trajectories.

Initial experiences with the rectilinear true-FISP optimizations indicate that the multi-objective GA, in particular, will be applicable to most possible imaging applications. It is anticipated that further developments and extensions to these methods will allow for improvements in established techniques and, perhaps more importantly, will permit the rapid maturation of new techniques.

Appendix A. Kety Equation Derivation

The Kety equation is used in MRI, PET, and other imaging modalities in order to quantify tissue perfusion rates. In general, it relates the rate of uptake of an inert substance into the tissues to the concentration of that inert substance in the arterial and venous blood. In MRI, the form of the Kety equation that is most commonly used is (all variables used in this section are defined in Table A.1):

$$\frac{dC_T}{dt} = E F (1 - Hct)(C_P - C_T / v_e) = K^{trans} C_P - k_{ep} C_T$$

To derive this equation, consider a differential element of plasma as it passes through the capillaries from the arterial to the venous side. If the tissue concentration is zero, then

a fraction $E = \frac{Q_A - Q_V}{Q_A} = \frac{C_A - C_V}{C_A}$ of the tracer will cross the capillary membrane from

Table A.1
Summary of variables used in Kety equation

Name	Description	Units
C_T	Concentration of tracer in tissue	mM
C_P	Concentration of tracer in plasma	mM
E	Initial Extraction Fraction	
F	Volume blood flow to tissue	min ⁻¹
Hct	Hematocrit	
v_e	EES volume per unit tissue volume	
K^{trans}	Transfer rate constant	min ⁻¹
k_{ep}	Back-transfer rate constant	min ⁻¹
C_A	Concentration of tracer in arterial blood	mM
C_V	Concentration of tracer in venous blood	mM
Q_A	Amount of tracer in arterial blood	mmol
Q_V	Amount of tracer in venous blood	mmol
C_E	Concentration of tracer in EES	mM

the plasma to the EES. This fraction, E , is related to the relative dominance of flow and vascular permeability. For example, if the vascular permeability is very high relative to the flow, then by the time the differential element passes through to the venous side the plasma concentration will be equal to the EES concentration and therefore E will be equal to 1 ($C_V = 0$ for $C_T = 0$). If permeability is not high relative to flow, then E will be some number less than 1.

The total volume of plasma that passes through a unit volume of tissue is equal to $F(1 - Hct)$, and therefore the total flux of tracer entering a unit volume of tissue can be written:

$$\frac{dC_T}{dt} = E F (1 - Hct) C_P$$

Assuming that at equilibrium $C_P = C_E$ (i.e. no active transport across the capillary so transport is equally likely in both directions), then E is also the fraction of the EES concentration that a differential element of plasma will have upon reaching the venous side if it enters with no tracer. Therefore the total flux of tracer leaving a unit volume of tissue can be written:

$$-\frac{dC_T}{dt} = E F (1 - Hct) C_E$$

These two expressions can be combined using the principle of superposition to obtain:

$$\frac{dC_T}{dt} = E F (1 - Hct) (C_P - C_E)$$

By making the substitution $C_T = v_e C_E$ we obtain:

$$\frac{dC_T}{dt} = E F (1 - Hct) (C_P - C_T / v_e) = K^{trans} C_P - k_{ep} C_T$$

which is the most commonly used form of the Kety equation for MRI.

The plasma concentration cannot be measured directly using MRI, but the arterial concentration can be measured if an appropriate artery is within the imaging plane. By making the substitution $C_A = (1 - Hct) C_p$ we obtain:

$$\frac{dC_T}{dt} = EF C_A - (EF(1 - Hct)/v_e) C_T = \frac{K^{trans}}{1 - Hct} C_A - k_{ep} C_T$$

which, although not commonly used, is an expression of the same form and meaning as the Kety equation, but in terms of variables that can be measured using MRI.

Appendix B. Trajectories and Gradient Waveforms

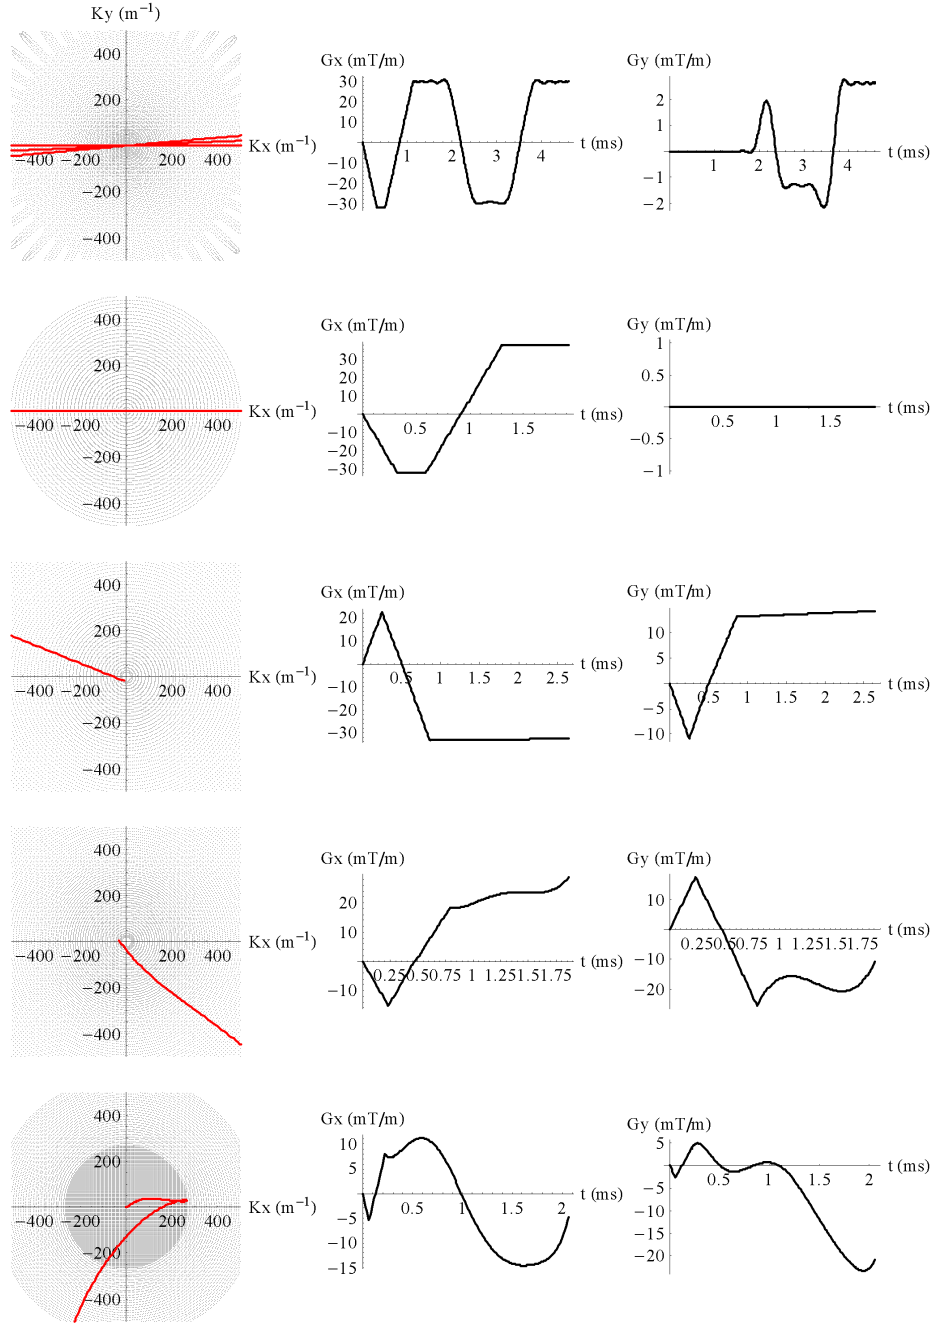


Figure B.1 Trajectories and gradient waveforms for the group of trajectories corresponding to the time objective. A single interleaf of each trajectory is displayed. Trajectories are displayed in order of acquisition time with the fastest trajectory on top and the slowest trajectory on bottom. The first trajectory is Pareto-optimal and the second is a standard 201-view radial trajectory.

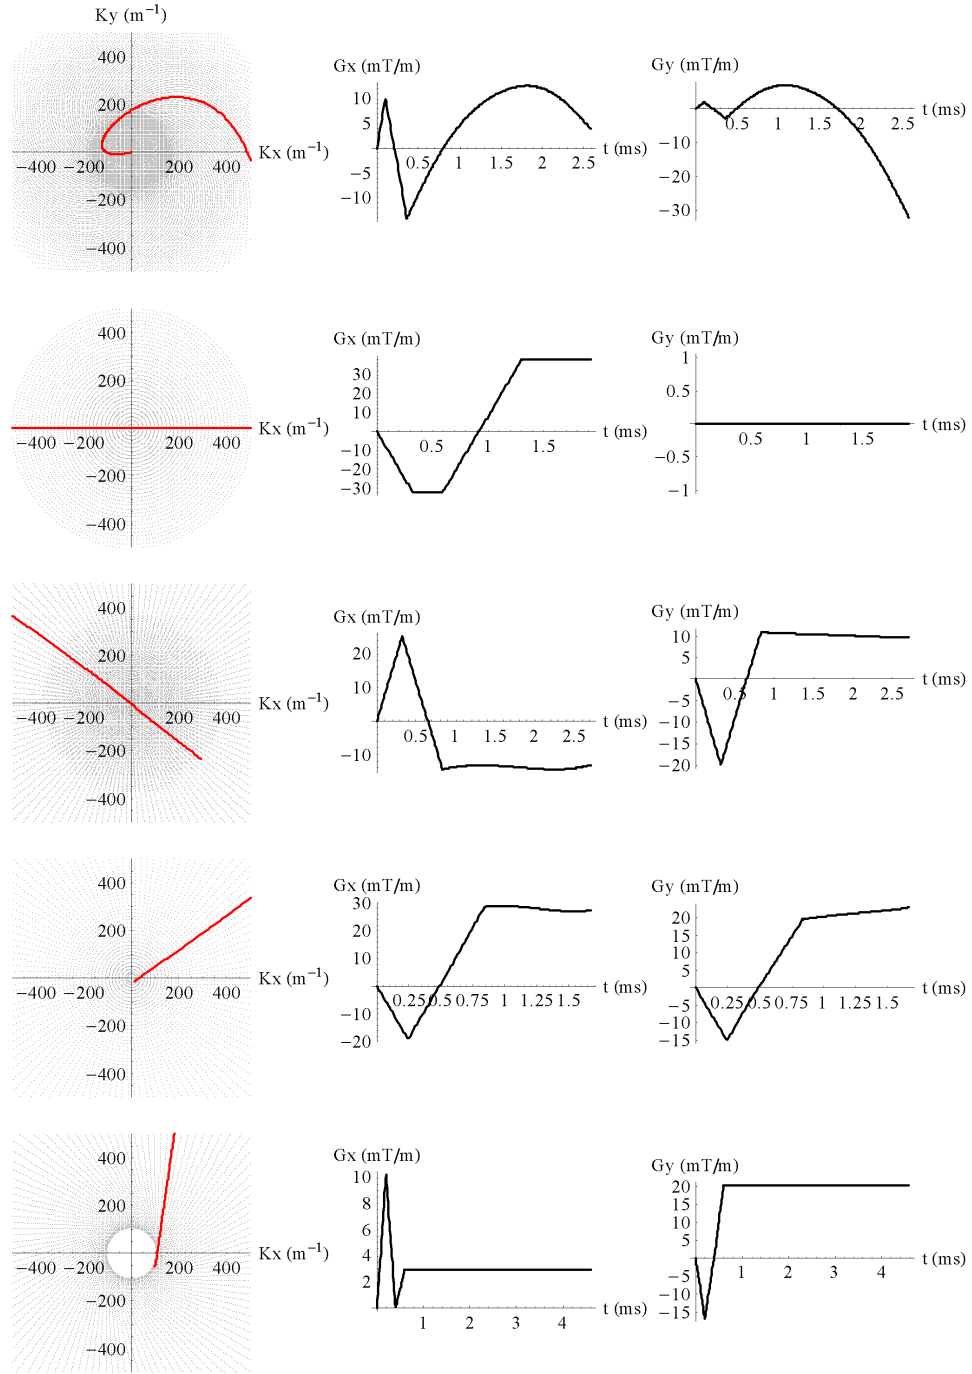


Figure B.2 Trajectories and gradient waveforms for the group of trajectories corresponding to the aliasing energy objective. A single interleaf of each trajectory is displayed. Trajectories are displayed in order of aliasing energy with the least aliasing energy on top and the most on bottom. The first trajectory is Pareto-optimal and the second is a standard 109-view radial trajectory.

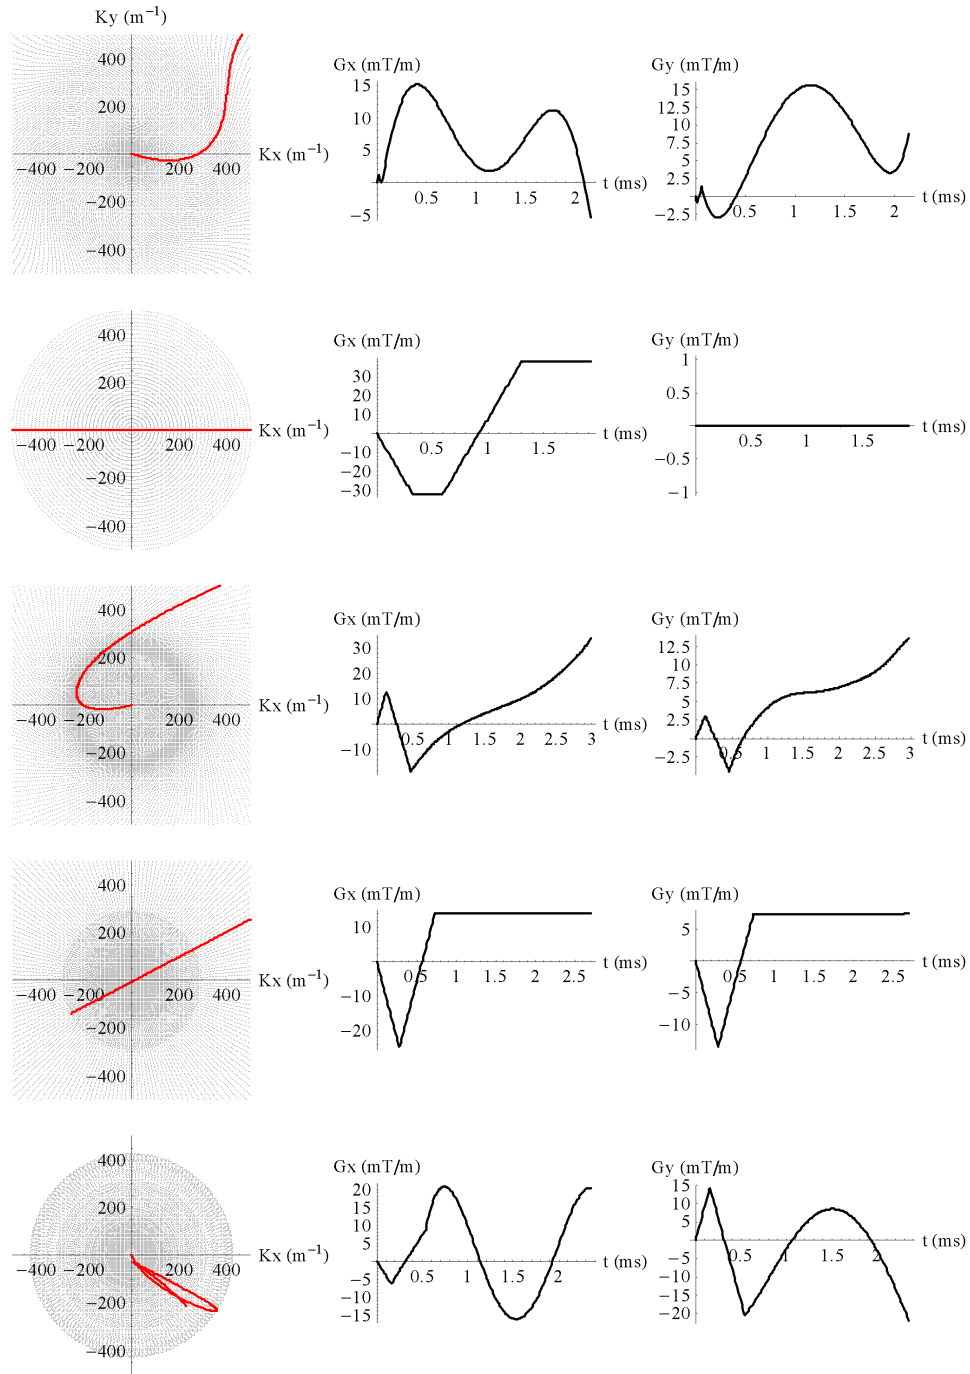


Figure B.3 Trajectories and gradient waveforms for the group of trajectories corresponding to the flow-artifact energy objective. A single interleaf of each trajectory is displayed. Trajectories are displayed in order of flow-artifact energy with the least flow artifact on top and the most on bottom. The first trajectory is Pareto-optimal and the last is a standard 157-view radial trajectory.

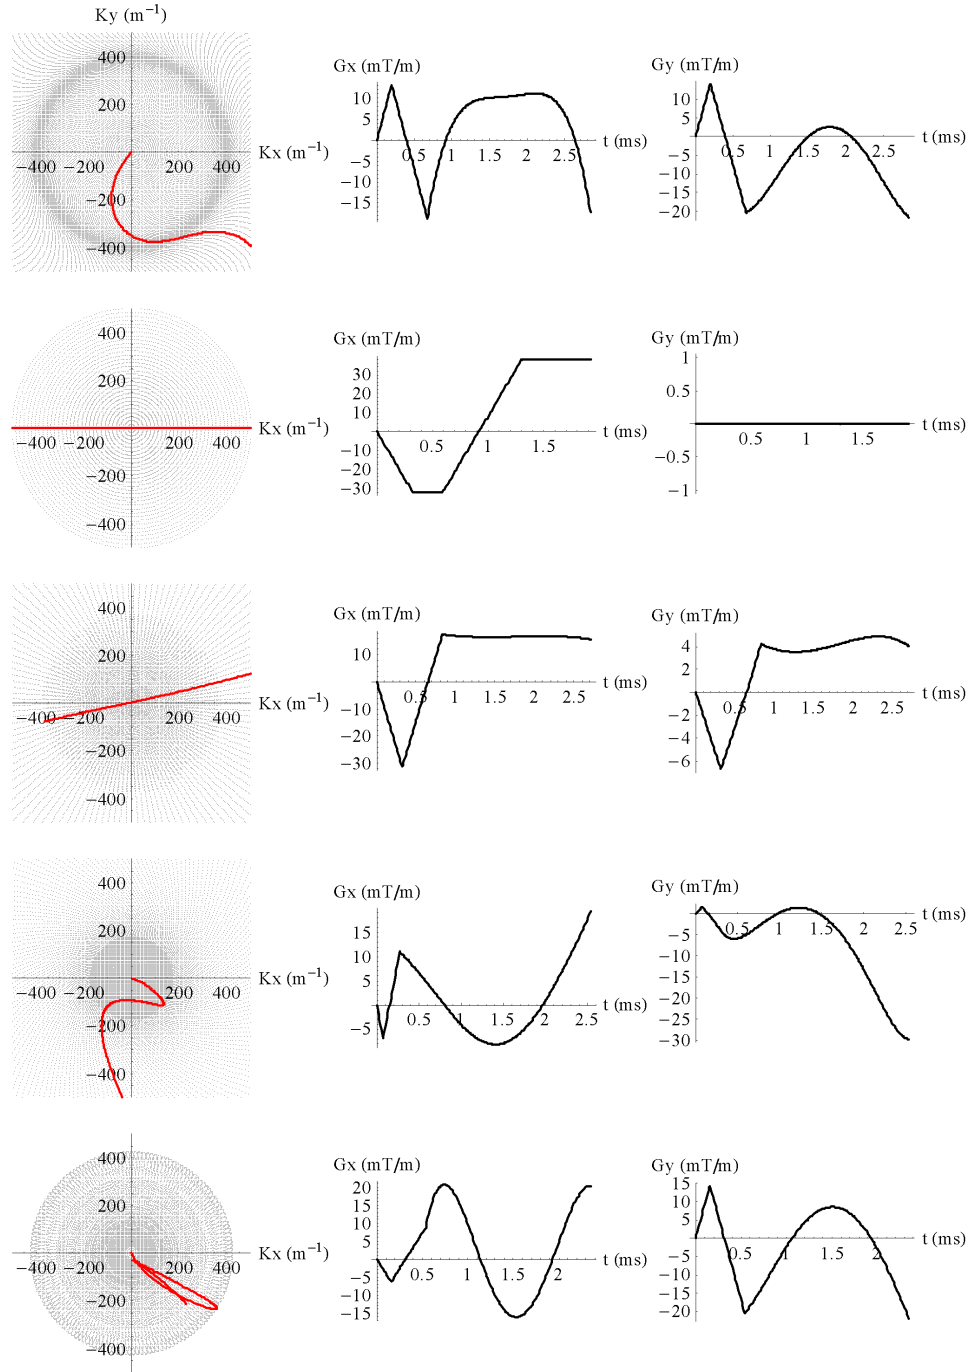


Figure B.4 Trajectories and gradient waveforms for the group of trajectories corresponding to the off-resonance artifact energy objective. A single interleaved of each trajectory is displayed. Trajectories are displayed in order of off-resonance artifact energy with the least off-resonance artifact on top and the most on bottom. The first trajectory is Pareto-optimal and the second is a standard 133-view radial trajectory.

Appendix C. Long Axis Cardiac Images

During the defense of this dissertation, one important comment from the committee was regarding the appropriateness of the axial slice orientation used to acquire the cardiac images for the subjective study. The simulations used for the flow-artifact energy objective computation simulated only in-plane flow, and the axial view has a large




Figure C.1 Long-axis cardiac images acquired using trajectories from the flow-artifact energy group. The trajectory for a) is Pareto-optimal while the trajectory for b) is a standard 157-view radial trajectory. The remaining trajectories, for images c), d), and e), are the sub-optimal trajectories and are arranged in order of increasing flow-artifact objective. Note that all images have a relatively large amount of background artifact. The position of the slice requires that a large amount of anatomy be located outside of the FOV. Using these trajectories there is no single “readout” direction along which to apply over-sampling and thus avoid aliasing. Due to these effects, the cardiac images for the subjective study were acquired using the axial slice positioning despite the greater amount of in-plane flow with the long axis view.

amount of through-plane motion. It was therefore suggested that a four-chamber or sagittal view of the heart, containing both ventricle and aorta, might have been preferable in order to increase the amount of in-plane flow and reduce the through plane motion. Figure C.1 shows such images acquired using each of the five trajectories in the flow-artifact objective group. Although there are visibly increased levels of artifact as a function of increasing flow-artifact objective (note the increased degradation on the bottom row from image c to image e), the overall quality of these images is much worse than the corresponding axial images.

This overall increased level of artifacts is most likely due to heightened amounts of aliasing energy. In rectilinear imaging, aliasing is only a concern for the phase-encoding direction because oversampling can be applied in the readout direction to increase the actual FOV to the required distance in that direction. However, with these non-rectilinear trajectories there is no single readout direction and no single phase-encoding direction. Even when readout oversampling is utilized, such trajectories are therefore still highly sensitive to signal from locations outside of the FOV. This would have a tendency to make the image quality more strongly related to the aliasing energy objective, and thus introduce even more variance for the logistic regressions related to the flow-artifact objective.




Figure C.2 Long-axis cardiac images acquired using trajectories from the flow-artifact energy group with cranial and caudal saturation pulses applied. The trajectory for a) is Pareto-optimal while the trajectory for b) is a standard 157-view radial trajectory. The remaining trajectories, for images c), d), and e), are the sub-optimal trajectories and are arranged in order of increasing flow-artifact objective. Note that all images have less background artifact than the corresponding images in Figure C.1. The aliasing energy due to signal from tissue outside of the FOV can be reduced using the saturation pulses at the expense of time and some changes in contrast due to the increased TR.

In order to reduce the amount of aliasing energy due to signal outside of the FOV it is possible to apply saturation pulses cranial and caudal to the heart and aortic arch. Such images are shown in Figure C.2. In this Figure it is possible to observe increasing levels of artifact with increasing flow artifact sensitivity levels, particularly for the three sub-optimal trajectories in Figure C.2a-c. It is also apparent that there are reduced levels of background artifact relative to Figure C.1.

Although the use of saturation pulses is shown here to be one potential method for improving the overall quality of these images, it was not utilized in the subjective image quality experiments for two principal reasons. First, all forms of magnetization preparation have been avoided in this work because any magnetization preparation necessarily increases the acquisition time and any associated improvements in image quality would be unrelated to the trajectory or the trajectory optimization. Second, by increasing the TR of the pulse sequence, the use of the saturation pulses would result in altered contrast, which was considered an important factor to control in the design of the perceptual experiments. However, it may be worthwhile for a follow-up study of the flow-artifact objective to utilize this technique in order to obtain a reasonable level of image quality in the sagittal orientation rather than using the transverse orientation without magnetization preparation.

References

1. Ahn CB, Kim JH, Cho ZH. 1986. High-speed spiral-scan echo planar NMR imaging - I. IEEE Trans Med Imaging MI-5:2-7.
2. Aronen HJ, Glass J, Pardo FS, Belliveau JW, Gruber ML, Buchbinder BR, Gazit IE, Linggood RM, Fischman AJ, Rosen BR. 1995. Echo-planar MR cerebral blood volume mapping of gliomas. Clinical utility. Acta Radiol 36(5):520-8.
3. Atalar E, McVeigh ER. 1994. Minimization of dead-periods in MRI pulse sequences for imaging oblique planes. Magn Reson Med 32(6):773-7.
4. Axel L. 1980. Cerebral blood flow determination by rapid-sequence computed tomography: theoretical analysis. Radiology 137(3):679-86.
5. Bahn MM. 1995. A single-step method for estimation of local cerebral blood volume from susceptibility contrast MRI images. Magn Reson Med 33(3):309-17.
6. Beache GM, Kulke SF, Kantor HL, Niemi P, Campbell TA, Chesler DA, Gewirtz H, Rosen BR, Brady TJ, Weisskoff RM. 1998. Imaging perfusion deficits in ischemic heart disease with susceptibility-enhanced T2-weighted MRI: preliminary human studies. Magn Reson Imaging 16(1):19-27.
7. Belle V, Kahler E, Waller C, Rommel E, Voll S, Hiller KH, Bauer WR, Haase A. 1998. In vivo quantitative mapping of cardiac perfusion in rats using a noninvasive MR spin-labeling method. J Magn Reson Imaging 8(6):1240-5.
8. Bevington PR. 1969. Data Reduction and Error Analysis for the Physical Sciences. New York: McGraw-Hill.
9. Bolster BD, Jr., Atalar E. 1999. Minimizing dead-periods in flow-encoded or -compensated pulse sequences while imaging in oblique planes. J Magn Reson Imaging 10(2):183-92.
10. Booth R. 1997. Inner Loops: A sourcebook for fast 32-bit software development. Massachusetts: Addison-Wesley Developers Press. -364 p.
11. Bornert P, Aldefeld B, Eggers H. 2000. Reversed spiral MR imaging. Magn Reson Med 44(3):479-84.
12. Bornert P, Schomberg H, Aldefeld B, Groen J. 1999. Improvements in spiral MR imaging. MAGMA 9(1-2):29-41.
13. Boxerman JL, Rosen BR, Weisskoff RM. 1997. Signal-to-noise analysis of cerebral blood volume maps from dynamic NMR imaging studies. J Magn Reson Imaging 7(3):528-37.

14. Brady TJ, Cohen MS, Weisskoff RM, Rosen BR. 1991. Equipment requirements to facilitate contrast-enhanced MR imaging. *Magn Reson Med* 22(2):273-9.
15. Brittain JH, Hu BS, Wright GA, Meyer CH, Macovski A, Nishimura DG. 1995. Coronary angiography with magnetization-prepared T2 contrast. *Magn Reson Med* 33(5):689-96.
16. Butts K, Pauly J, de Crespigny A, Moseley M. 1997. Isotropic diffusion-weighted and spiral-navigated interleaved EPI for routine imaging of acute stroke. *Magn Reson Med* 38(5):741-9.
17. Buxton RB, Frank LR, Wong EC, Siewert B, Warach S, Edelman RR. 1998. A general kinetic model for quantitative perfusion imaging with arterial spin labeling. *Magnetic Resonance In Medicine* 40(3):383-96.
18. Choi HK, Munson DC. 1998. Direct-Fourier reconstruction in tomography and synthetic aperture radar. *Intl J Imaging Sys Tech* 9:1-13.
19. Cline HE, Zong X, Gai N. 2001. Design of a logarithmic k-space spiral trajectory. *Magn Reson Med* 46(6):1130-5.
20. Coombs BD, Szumowski J, Coshov W. 1997. Two-point Dixon technique for water-fat signal decomposition with B0 inhomogeneity correction. *Magn Reson Med* 38(6):884-9.
21. Crone C. 1963. Permeability of capillaries in various organs as determined by the use of the 'indicator-diffusion' method. *Acta Physiol Scand* 58:292-305.
22. Dale B, Wendt M, Duerk JL. 2001. A rapid look-up table method for reconstructing MR images from arbitrary K-space trajectories. *IEEE Trans Med Imaging JID - 8310780* 20(3):207-17.
23. Dale, B. M., Flask, C. A., Lewin, J. S., and Duerk J.L. 2003a. Time-Optimal True-FISP Pulse Sequences with Gradient-Induced Peripheral Nerve Stimulation Constraints. Chicago, Illinois: In: Proc RSNA, 88th Annual Meeting, Chicago; 704 p.
24. Dale BM, Jesberger JA, Lewin JS, Hillenbrand CM, Duerk JL. 2003b. Determining and optimizing the precision of quantitative measurements of perfusion from dynamic contrast enhanced MRI. *J Magn Reson Imaging* 18(5):575-84.
25. Dale, B. M., Lewin, J. S., and Duerk J.L. 2003c. Genetic design of variable-density spiral trajectories to minimize off-resonance and flow effects. Toronto, Canada: In: Proc ISMRM, 11th Annual Meeting, Toronto, Canada; 1008 p.
26. Dale BM, Lewin JS, Duerk JL. Optimal Design of K-space Trajectories Using a Multi-objective Genetic Algorithm. *MRM Accepted for Publication*. 2004.

27. Deb K, Pratap A, Agarwal S, Meyarivan T. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans Evol Comp* 6(2):182-97.
28. Deb K. 2001. Multi-objective optimization using evolutionary algorithms. Chichester, N.Y.: John Wiley & Sons. xix-497 p.
29. Dennie J, Mandeville JB, Boxerman JL, Packard SD, Rosen BR, Weisskoff RM. 1998. NMR imaging of changes in vascular morphology due to tumor angiogenesis. *Magn Reson Med* 40(6):793-9.
30. Dixon WT. 1984. Simple proton spectroscopic imaging. *Radiology* 153(1):189-94.
31. Duerk JL, Lewin JS, Wendt M, Petersilge C. 1998. Remember true FISP? A high SNR, near 1-second imaging method for T2-like contrast in interventional MRI at .2 T. *J Magn Reson Imaging* 8(1):203-8.
32. Duerk JL, Lewin JS, Wu DH. 1996. Application of keyhole imaging to interventional MRI: a simulation study to predict sequence requirements. *J Magn Reson Imaging* 6(6):918-24.
33. Duyn JH, Yang Y, Frank JA, van der Veen JW. 1998. Simple correction method for k-space trajectory deviations in MRI. *J Magn Reson JID* - 9707935 132(1):150-3.
34. Eggers H and Proska R. 1999. Multiprocessor system for real-time convolution interpolation reconstruction [abstract]. In: Multiprocessor system for real-time convolution interpolation reconstruction.
35. Evelhoch J, Brown T, Chenevert T, Clarke L, Daniael B, Degani H, Hylton N, Knopp M, Koutcher J, Lee T-Y, Mayr N, Sullivan D, Taylor J, Toffts P, and Weisskoff R. 2000. Consensus Recommendation for Acquisition of Dynamic Contrast-Enhanced MRI Data in Oncology [abstract]. In: Consensus Recommendation for Acquisition of Dynamic Contrast-Enhanced MRI Data in Oncology.
36. Fisher BJ, Dillon N, Carpenter TA, Hall LD. 1997. Design of a biplanar gradient coil using a genetic algorithm. *Magn Reson Imaging* 15(3):369-76.
37. Flask CA, Dale B, Lewin JS, Duerk JL. 2003. Radial alternating TE sequence for faster fat suppression. *Magn Reson Med* 50(5):1095-9.
38. Gadian DG, Payne JA, Bryant DJ, Young IR, Carr DH, Bydder GM. 1985. Gadolinium-DTPA as a contrast agent in MR imaging--theoretical projections and practical observations. *J Comput Assist Tomogr* 9(2):242-51.
39. Gao Y, Reeves SJ. 2000. Optimal k-space sampling in MRSI for images with a limited region of support. *IEEE Trans Med Imaging* 19(12):1168-78.

40. Glover GH, Pauly JM. 1992. Projection reconstruction techniques for reduction of motion effects in MRI. *Magn Reson Med* 28(2):275-89.
41. Goldburgh M. 2001 MRI Census Market Summary Report. 2002. IMV Medical Information Division.
42. Grossman RI, Gonzalez-Scarano F, Atlas SW, Galetta S, Silberberg DH. 1986. Multiple sclerosis: gadolinium enhancement in MR imaging. *Radiology* 161(3):721-5.
43. Haacke EM, Brown RW, Thompson MR, Venkatesan R. 2001. *Magnetic Resonance Imaging - Physical Principles and Sequence Design*. New York: Wiley-Liss. 1 p.
44. Hahn PF, Stark DD, Weissleder R, Elizondo G, Saini S, Ferrucci JT. 1990. Clinical application of superparamagnetic iron oxide to MR imaging of tissue perfusion in vascular liver tumors. *Radiology* 174(2):361-6.
45. Hardy CJ, Cline HE. 1989. Broadband nuclear magnetic resonance pulses with two dimensional spatial selectivity. *J Appl Phys* 66(4):1513-6.
46. Hebrank, F. X. and Gebhardt, M. 2000. SAFE-Model - A New Method for Predicting Peripheral Nerve Stimulations in MRI. Denver, Colorado: 1764 p.
47. Hendrick RE, Kneeland JB, Stark DD. 1987. Maximizing signal-to-noise and contrast-to-noise ratios in FLASH imaging. *Magn Reson Imaging* 5(2):117-27.
48. Hocking LM. 1991. *Optimal control : an introduction to the theory with applications*. New York, N.Y.: Clarendon Press. xiv-254 p.
49. Hoge RD, Kwan RK, Pike GB. 1997. Density compensation functions for spiral MRI. *Magn Reson Med* 38(1):117-28.
50. Hogg DE, MacDonald GH, Conway RG, Wade CM. 1969. Synthesis of brightness distribution in radio sources. *Astronom J* 74:1206-13.
51. International Telecommunication Union. 2002. Methodology for the subjective assessment of the quality of television pictures. Report nr Rec. ITU-R BT.500-11. 1 p. Available from.
52. Irarrazabal P, Hu BS, Pauly JM, Nishimura DG. 1993. Spatially resolved and localized real-time velocity distribution. *Magn Reson Med* 30(2):207-12.
53. Irarrazabal P, Nishimura DG. 1995. Fast three dimensional magnetic resonance imaging. *Magn Reson Med* 33(5):656-62.
54. Irnich W, Schmitt F. 1995. Magnetostimulation in MRI. *Magn Reson Med* 33(5):619-23.

55. Jackson JI, Meyer CH, Nishimura DG, Macovski A. 1991. Selection of a convolution function for Fourier inversion using gridding. *IEEE Trans Med Imaging* 10:473-8.
56. Jesmanowicz A, Bandettini PA, Hyde JS. 1998. Single-shot half k-space high-resolution gradient-recalled EPI for fMRI at 3 Tesla. *Magn Reson Med* 40(5):754-62.
57. Karlens OT, Verhagen R, Bovee WM. 1999. Parameter estimation from Rician-distributed data sets using a maximum likelihood estimator: application to T1 and perfusion measurements. *Magn Reson Med* 41(3):614-23.
58. Kerr AB, Pauly JM, Hu BS, Li KC, Hardy CJ, Meyer CH, Macovski A, Nishimura DG. 1997. Real-time interactive MRI on a conventional scanner. *Magn Reson Med* 38(3):355-67.
59. Kety SS. 1951. The Theory and Application of the Exchange of Inert Gas at the Lungs and Tissues. *Pharmacological Reviews* 3:1-41.
60. Kirk DE. 1970. Optimal control theory; an introduction. Englewood Cliffs, N.J.: Prentice-Hall. ix-452 p.
61. Lahanas M, Baltas D, Zamboglou N. 2003. A hybrid evolutionary algorithm for multi-objective anatomy-based dose optimization in high-dose-rate brachytherapy. *Phys Med Biol* 48(3):399-415.
62. Landis CS, Li X, Telang FW, Molina PE, Palyka I, Vetek G, Springer CS, Jr. 1999. Equilibrium transcytolemmal water-exchange kinetics in skeletal muscle in vivo. *Magn Reson Med* 42(3):467-78.
63. Larson AC, Simonetti OP. 2001. Real-time cardiac cine imaging with SPIDER: steady-state projection imaging with dynamic echo-train readout. *Magn Reson Med* 46(6):1059-66.
64. Larson AC, Simonetti OP, Li D. 2002. Coronary MRA with 3D undersampled projection reconstruction TrueFISP. *Magn Reson Med* 48(4):594-601.
65. Liao JR. 1999. Real-time spiral image reconstruction using fixed-point calculations [abstract]. In: Real-time spiral image reconstruction using fixed-point calculations.
66. Liao JR, Pauly JM, Brosnan TJ, Pelc NJ. 1997. Reduction of motion artifacts in cine MRI using variable-density spiral trajectories. *Magn Reson Med* 37(4):569-75.
67. Lombardi M, Jones RA, Westby J, Torheim G, Southon TE, Haraldseth O, Michelassi C, Kvaerness J, Rinck PA, L'Abbate A. 1999. Use of the mean transit

- time of an intravascular contrast agent as arm exchange-insensitive index of myocardial perfusion. *Jmri-Journal Of Magnetic Resonance Imaging* 9(3):402-8.
68. Luo Y, Mohning KM, Hradil VP, Wessale JL, Segreti JA, Nuss ME, Wegner CD, Burke SE, Cox BF. 2002. Evaluation of tissue perfusion in a rat model of hind-limb muscle ischemia using dynamic contrast-enhanced magnetic resonance imaging. *J Magn Reson Imaging* 16(3):277-83.
 69. McKenzie CA, Pereira RS, Prato FS, Chen Z, Drost DJ. 1999. Improved contrast agent bolus tracking using T1 FARM. *Magn Reson Med* 41(3):429-35.
 70. McMahon MT, Oldfield E. 1999. Determination of order parameters and correlation times in proteins: a comparison between Bayesian, Monte Carlo and simple graphical methods. *J Biomol NMR* 13(2):133-7.
 71. Meyer CH, Hu BS, Nishimura DG, Macovski A. 1992. Fast spiral coronary artery imaging. *Magn Reson Med* 28(2):202-13.
 72. Moriguchi H., Wendt M., and Duerk JL. 1999. Image reconstruction method without using interpolation in nonuniform sampled k-space MR imaging. In: *Proc RSNA, 85th Annual meeting, Chicago*; 232 p.
 73. Moriguchi, H., Lewin, J. S., and Duerk, J. L. 2003. A New Approach for Optimal Reconstruction Using Rescaled Matrices from Non-uniformly Sampled K-space Data. Toronto: In: *Proceedings of the 11th annual meeting of ISMRM*; Submitted p.
 74. Moriguchi H, Wendt M, Duerk JL. 2000. Applying the uniform resampling (URS) algorithm to a lissajous trajectory: fast image reconstruction with optimal gridding. *Magn Reson Med* 44(5):766-81.
 75. Niendorf ER, Grist TM, Frayne R, Brazy PC, Santyr GE. 1997. Rapid measurement of Gd-DTPA extraction fraction in a dialysis system using echo-planar imaging. *Med Phys* 24(12):1907-13.
 76. Niendorf ER, Santyr GE, Brazy PC, Grist TM. 1996. Measurement of Gd-DTPA dialysis clearance rates by using a look-locker imaging technique. *Magn Reson Med* 36(4):571-8.
 77. Nishimura DG, Irarrazabal P, Meyer CH. 1995. A velocity k-space analysis of flow effects in echo-planar and spiral imaging. *Magn Reson Med* 33(4):549-56.
 78. Noll DC, Cohen JD, Meyer CH, Schneider W. 1995. Spiral K-space MR imaging of cortical activation. *J Magn Reson Imaging* 5(1):49-56.
 79. Noll DC, Peltier SJ, Boada FE. 1998. Simultaneous multislice acquisition using rosette trajectories (SMART): a new imaging method for functional MRI. *Magn Reson Med* 39(5):709-16.

80. O'Sullivan JD. 1985. A fast sinc function gridding algorithm for Fourier inversion in computer tomography. *IEEE Trans Med Imaging* MI-4:200-7.
81. Oesterle C, Markl M, Strecker R, Kraemer FM, Hennig J. 1999. Spiral reconstruction by regridding to a large rectilinear matrix: a practical solution for routine systems. *J Magn Reson Imaging* 10(1):84-92.
82. Okada S, Katagiri K, Kumazaki T, Yokoyama H. 2001. Safety of gadolinium contrast agent in hemodialysis patients. *Acta Radiol* 42(3):339-41.
83. Ostergaard L, Sorensen AG, Kwong KK, Weisskoff RM, Gyldensted C, Rosen BR. 1996a. High resolution measurement of cerebral blood flow using intravascular tracer bolus passages. Part II: Experimental comparison and preliminary results. *Magn Reson Med* 36(5):726-36.
84. Ostergaard L, Weisskoff RM, Chesler DA, Gyldensted C, Rosen BR. 1996b. High resolution measurement of cerebral blood flow using intravascular tracer bolus passages. Part I: Mathematical approach and statistical analysis. *Magn Reson Med* 36(5):715-25.
85. Peters DC, Korosec FR, Grist TM, Block WF, Holden JE, Vigen KK, Mistretta CA. 2000. Undersampled projection reconstruction applied to MR angiography. *Magn Reson Med* 43(1):91-101.
86. Petersson JS, Christoffersson JO, Golman K. 1993. MRI simulation using the k-space formalism. *Magn Reson Imaging* 11(4):557-68.
87. Pipe JG. 1999. An optimized center-out k-space trajectory for multishot MRI: comparison with spiral and projection reconstruction. *Magn Reson Med* 42(4):714-20.
88. Pipe JG, Menon P. 1999. Sampling density compensation in MRI: rationale and an iterative numerical solution. *Magn Reson Med* 41(1):179-86.
89. Port RE, Knopp MV, Hoffmann U, Milker-Zabel S, Brix G. 1999. Multicompartment analysis of gadolinium chelate kinetics: blood-tissue exchange in mammary tumors as monitored by dynamic MR imaging. *J Magn Reson Imaging* 10(3):233-41.
90. Prato FS, Drost DJ, Keys T, Laxon P, Comissiong B, Sestini E. 1986. Optimization of signal-to-noise ratio in calculated T1 images derived from two spin-echo images. *Magn Reson Med* 3(1):63-75.
91. Press WH, Teukolsky SA, Vetterling WT, Flannery BP. 1992. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge: Cambridge University Press. xi-994 p.

92. Price RR. 1999. The AAPM/RSNA physics tutorial for residents. MR imaging safety considerations. Radiological Society of North America. Radiographics 19(6):1641-51.
93. Proakis JG, Manolakis DG. 1996. Digital Signal Processing: Principles, algorithms, and applications. 3 ed. New Jersey: Prentice Hall. -968 p.
94. Qian Y, Lin J, Jin D. 2002. Reconstruction of MR images from data acquired on an arbitrary k-space trajectory using the same-image weight. Magn Reson Med 48(2):306-11.
95. Reber PJ, Wong EC, Buxton RB, Frank LR. 1998. Correction of off resonance-related distortion in echo-planar imaging using EPI-based field maps. Magn Reson Med 39(2):328-30.
96. Reeder SB, Atalar E, Faranesh AZ, McVeigh ER. 1999. Multi-echo segmented k-space imaging: an optimized hybrid sequence for ultrafast cardiac imaging. Magn Reson Med 41(2):375-85.
97. Reeder SB, Atalay MK, McVeigh ER, Zerhouni EA, Forder JR. 1996. Quantitative cardiac perfusion: a noninvasive spin-labeling method that exploits coronary vessel geometry. Radiology 200(1):177-84.
98. Renkin EM. 1959. Transport of Potassium-42 from blood to tissue in isolated mammalian skeletal muscles. Am J Physiol 197:1205-10.
99. Rosenfeld D. 1998. An optimal and efficient new gridding algorithm using singular value decomposition. Magn Reson Med 40(1):14-23.
100. Rosenfeld D. 2002. New approach to gridding using regularization and estimation theory. Magn Reson Med 48(1):193-202.
101. Sabat, S and Irarrazaval, P. 2002. Three Dimensional K-Space Trajectory Design Using Genetic Algorithms. Honolulu, Hawaii: 2332 p.
102. Sabat S, Mir R, Guarini M, Guesalaga A, Irarrazaval P. 2003. Three dimensional k-space trajectory design using genetic algorithms. Magn Reson Imaging 21(7):755-64.
103. Salem K.A., Lewin JS., Duerk J.L., Wilson D.L. 2002. Validation of a human vision model for image quality evaluation of fast interventional magnetic resonance imaging. J Electronic Imag 11(2):224-35.
104. Sedarat H and Nishimura DG. 1999. Gridding reconstruction using optimal, shift-variant interpolating kernels [abstract]. In: Gridding reconstruction using optimal, shift-variant interpolating kernels.

105. Shankaranarayanan A, Wendt M, Lewin JS, Duerk JL. 2001. Two-step navigatorless correction algorithm for radial k-space MRI acquisitions. *Magn Reson Med* 45(2):277-88.
106. Sharatchandra MC, Sen M, Gad-el-Hak M. 1998. New Approach to Constrained Shape Optimization Using Genetic Algorithms. *American Institute of Aeronautics and Astronautics Journal* 36(1):51-61.
107. Simonetti OP, Duerk JL, Chankong V. 1993. MRI gradient waveform design by numerical optimization. *Magn Reson Med* 29(4):498-504.
108. Simpson NE, He Z, Evelhoch JL. 1999. Deuterium NMR tissue perfusion measurements using the tracer uptake approach: I. Optimization of methods. *Magn Reson Med* 42(1):42-52.
109. Sorensen AG, Tievsky AL, Ostergaard L, Weisskoff RM, Rosen BR. 1997. Contrast agents in functional MR imaging. *J Magn Reson Imaging* 7(1):47-55.
110. Star-Lack JM. 1999. Optimal gradient waveform design for projection imaging and projection reconstruction echoplanar spectroscopic imaging. *Magn Reson Med* 41(4):664-75.
111. Su MY, Nalcioğlu O. 1993. Statistical description of microcirculatory flow as measured with an MR method. *J Magn Reson Imaging* 3(6):883-7.
112. Thedens DR, Irarrazaval P, Sachs TS, Meyer CH, Nishimura DG. 1999. Fast magnetic resonance coronary angiography with a three-dimensional stack of spirals trajectory. *Magn Reson Med* 41(6):1170-9.
113. Tofts PS. 1997. Modeling tracer kinetics in dynamic Gd-DTPA MR imaging. *J Magn Reson Imaging* 7(1):91-101.
114. Tofts PS, Brix G, Buckley DL, Evelhoch JL, Henderson E, Knopp MV, Larsson HB, Lee TY, Mayr NA, Parker GJ, Port RE, Taylor J, Weisskoff RM. 1999. Estimating kinetic parameters from dynamic contrast-enhanced T(1)-weighted MRI of a diffusable tracer: standardized quantities and symbols. *J Magn Reson Imaging* 10(3):223-32.
115. Tofts PS, Kermode AG. 1991. Measurement of the blood-brain barrier permeability and leakage space using dynamic MR imaging. 1. Fundamental concepts. *Magn Reson Med* 17(2):357-67.
116. Tsai CM, Nishimura DG. 2000. Reduced aliasing artifacts using variable-density k-space sampling trajectories. *Magn Reson Med* 43(3):452-8.
117. Ueda J, Furukawa T, Yamamoto T, Sakaguchi K, Araki Y. 1998. Permeability of gadolinium-DTPA through two types of hemodialysis membrane. *Invest Radiol* 33(10):734-7.

118. Van Lom KJ, Brown JJ, Perman WH, Sandstrom JC, Lee JK. 1991. Liver imaging at 1.5 tesla: pulse sequence optimization based on improved measurement of tissue relaxation times. *Magn Reson Imaging* 9(2):165-71.
119. Wager TD, Nichols TE. 2003. Optimization of experimental design in fMRI: a general framework using a genetic algorithm. *Neuroimage* 18(2):293-309.
120. Wajer FTAW, Woudenberg E, de Beer R, Fuderer M, Mehlkopf AF, and van Ormondt D. 1999. Simple equation for optimal gridding parameters [abstract]. In: Simple equation for optimal gridding parameters.
121. Wang HZ, Riederer SJ, Lee JN. 1987. Optimizing the Precision in T1 Relaxation Estimation Using Limited Flip Angles. *Magn Reson Med* 5:399-416.
122. Williams GB, Fisher BJ, Huang CL, Carpenter TA, Hall LD. 1999. Design of biplanar gradient coils for magnetic resonance imaging of the human torso and limbs. *Magn Reson Imaging* 17(5):739-54.
123. Yuh WT. 1999. An exciting and challenging role for the advanced contrast MR imaging. *J Magn Reson Imaging* 10(3):221-2.
124. Zhou X, Liang ZP, Gewalt SL, Cofer GP, Lauterbur PC, Johnson GA. 1998. A fast spin echo technique with circular sampling. *Magn Reson Med* 39(1):23-7.