EIGHT-CHANNEL HEAD ARRAY AND CONTROL SYSTEM FOR PARALLEL

TRANSMIT/RECEIVE MAGNETIC RESONANCE IMAGING

A Dissertation

by

KATHERINE LYNN MOODY

Submitted to the Office of Graduate and Professional Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Chair of Committee,	Mary P. McDougall
Committee Members,	Steven M. Wright
	Duncan J. Maitland
	Jeremy S. Wasser
Head of Department,	Gerard L. Cote

August 2014

Major Subject: Biomedical Engineering

ABSTRACT

Interest in magnetic resonance imaging (MRI) at high fields strengths (3 Tesla and above) is driven by the associated improvements in signal-to-noise ratio and spectral resolution. In practice, however, technical challenges prevent these benefits from being straightforwardly realized. High fields are associated with an increase in frequency and a decrease in the radiofrequency (RF) wavelength. The shortened wavelength causes potential inhomogeneity in the transmit field of the RF coil, resulting in non-uniform excitations. Susceptibility effects are also more pronounced at high field strengths, and can cause local distortions in the field and create areas of signal dropout.

Parallel transmit is one method in development to address these challenges at high fields. Parallel transmit involves using multiple independently driven channels with RF pulses varying in amplitude, phase, and pulse envelope to create desired transmit excitations. Parallel transmit has been implemented to create homogenous transmit patterns and compensate for magnetic susceptibility effects, but despite its proven usefulness, the technology has yet to receive widespread adoption. Few parallel transmit systems exist and little work has been done in the pre-clinical realm. Studies demonstrating the clinical benefits of parallel transmit constitute a gap in the current body of work. This works presents an approach to a parallel transmit array and control system that can be easily and safely integrated on a clinical whole-body scanner. The transmit array was designed for use with ultra-low output impedance amplifiers and demonstrates an array design with a simplified decoupling network augmented by

amplifier decoupling in both transmit and receive. The control system was programmed in LabVIEW using off-the-shelf hardware to manage pulse playback, correct transmit chain non-linearities, monitor on-coil waveforms, and drive the transmit hardware. The transmit array was shown with well-isolated patterns, and parallel transmit capability was demonstrated. This work progressed the translation of experimental parallel transmit technology to pre-clinical use.

DEDICATION

"To him who is able to do far more abundantly than all that we ask or think, according to power at work within us, to him be glory" *Ephesians 3:20-21*

ACKNOWLEDGEMENTS

Looking back on my graduate experience, I am overwhelmed with gratitude for those who have assisted, encouraged, loved, and prayed for me through this journey. I would like to begin by acknowledging my advisor, Dr. Mary McDougall. I will forever remember our first meeting and the amount of passion and depth she displayed for the field of magnetic resonance. She inspired me to enter the field and gave me a perspective that would carry me through the difficult times – knowing the effort we put forth is ultimately for those who may one day benefit. She has consistently challenged and pushed me to reach beyond what I thought myself capable and has molded me into the researcher I am today. She has truly been my strongest advocate even through the times I questioned my own abilities. She has provided technical expertise and mentorship at every stage, and I could not have made it to this point in my career without her continued support and guidance.

Dr. Steve Wright has continued to be a source of inspiration and encouragement. Thank you for providing your insight and expertise through innumerable conversations. Some of my favorite memories in the lab were our discussions on the decoupling network and measurements for the array coil. He has an excitement and passion for the field that is contagious, and many conversations lead to new thoughts and ideas. Thanks for pushing me to improve and iterate on the components of the system and for not allowing me to give up when times were difficult.

Dr. Duncan Maitland and Dr. Jeremy Wasser have provided a diverse expertise that has forced me to expand my knowledge and broaden my understanding of magnetic resonance and its applications.

This work has been a collaborative effort, and I would like to acknowledge Jon-Fredrik Nielsen, Feng Zhao, Hao Sun, and Dr. Douglas Noll for their support and efforts on this project. I would like to thank Will Grissom for using the parallel transmit system and sharing his expertise.

My labmates have been invaluable through my graduate school experience. Neal Hollingsworth, my partner on this work and friend, has been there at every stage – both successes and failures. He has kept me grounded, provided his insight, and been a willing listener. Without his perseverance and ability to build the modulators and amplifiers for the transmit system, the work described in the following pages would not have been possible. John Bosshard has always been a source of strength, peace, and a wealth of knowledge. Sam By has been a light in the lab and has always been a willing listener and friend. I acknowledge Joseph Rispoli, Wen-Yang Chiang, Eddie Eigenbrodt, and Jiaming Cui for their help, friendship, and insights.

I cannot begin to express the gratitude I have for my church family at Living Hope. They have consistently encouraged and prayed for me throughout this experience and reminded me of truth. Matt and Haeli Wey and Scott and Kristen Saunders have been our best friends and family. Haeli and Kristen provided meals during the long weeks when I was working hard to finish the array coil and have helped me maintain perspective on the truly important things in life.

TABLE OF CONTENTS

LIST OF FIGURES

xiii

xiv

LIST OF TABLES

CHAPTER I

INTRODUCTION

Magnetic resonance imaging (MRI) for in vivo use was pioneered in 1973 by Paul Lauterbur [1] and Peter Mansfield [2], and since its inception, the modality has become widely available in most clinical environments serving as the gold standard for neurological imaging. With advances in magnet technology, high field strength imaging has become available opening up a realm of new possibilities. The increase in signal-to-noise ratio (SNR) and spectral resolution that comes with high field (3 Tesla and above) magnetic resonance imaging (MRI) can be traded for improvements in spatial and temporal resolution [3], but in many cases, technical challenges prevent these benefits from being straightforwardly realized in practice. This work is dedicated to the exploration of an approach for overcoming the high field barriers through a technique known as parallel transmit in an effort to provide enhanced physiological information for future studies.

I.1 Motivation for High Field Imaging and Current Technical Challenges

One significant and well-known challenge at high field is potential inhomogeneity in the transmit B_1 field due to the higher frequencies (and shorter radiofrequency (RF) wavelengths) associated with higher magnetic field strengths [4-8]. In neuroimaging applications, this is often manifest as a central brightening artifact due to constructive interference in the center of the head preventing uniform tip angles [9-10].

Conventional systems use a single transmitter in combination with the RF coil to produce the transmit field. One approach to solving the B_1 inhomogeneity problem is to use multiple excitation channels and adjust the amplitude and phase of each channel to compensate for the non-uniformity, a technique known as B_1 or RF shimming. These "parallel excitation" systems involve a single transmitter with a network to divide the RF signal into multiple channels and apply fixed amplitude and phase adjustments on each channel. When these adjusted transmit pulses are implemented, the uniformity of the resulting excitation can be corrected/improved.

Another problem at high field strengths is magnetic susceptibility artifacts. These artifacts occur, for instance, near air/tissue interfaces in which air (slightly paramagnetic) distorts the local magnetic field and creates areas of signal loss [11]. One approach to solving this problem is to use complex multidimensional RF pulses, which can selectively emphasize or de-emphasize excitation areas. These pulses are typically impractical using a single channel transmitter, however, especially for in vivo applications, due to the duration of the RF pulses and limitations in the technique [12-13]. A single channel transmit configuration of B_1 shimming systems cannot provide the different pulse shapes needed in susceptibility-correcting pulse designs, and an alternative platform is needed to address this problem.

Parallel transmit is an emerging technology that uses multiple transmitters in combination with a transmit RF coil array to accelerate multidimensional RF pulse designs to create uniform excitations at these interfaces in a shorter excitation time [14-16]. The pulses on each channel of a parallel transmit system are separately modulated,

2

providing complete control over the pulse shape, amplitude, and phase of each channel. By using multiple excitation channels, the transmit patterns of the individual coils can be incorporated into the RF pulse design to divide the excitation space and reduce imaging times (transmit-SENSE) [16-17]. Parallel transmit systems can therefore also correct for B_1 inhomogeneity issues, and thus serve as a more general solution for the challenges of high field imaging.

I.2 Parallel Transmit Magnetic Resonance Imaging

Parallel transmission was recently introduced [14] and builds upon the concept of parallel receive technology which began to increase in popularity with the seminal paper by Roemer, "The NMR Phased Array" [18] and has since seen rapid developments enabling higher SNR through surface coil arrays and accelerated imaging. Parallel receive has been adopted by commercial vendors enabling parallel receive capabilities in routine clinical practice.

Parallel transmit technology has yet to receive widespread adoption, but as engineering challenges are overcome and benefits realized, the technology may quickly translate to the clinical environment. The MR experiment with parallel transmit technology begins with the hardware requirements of the RF coil array and multiple transmitters as described previously. The transmit array consists of independent resonant elements, each with a spatially-dependent transmit sensitivity pattern. By using the spatial degrees of freedom inherent in the array coil patterns, undersampled transmit patterns are simultaneously excited, and the resulting image produces the desired excitation [15].

I.3 Dissertation Objective and Organization

The primary objective of this work is to demonstrate an approach to a parallel transmit array system that can be easily and safely integrated on a clinical whole-body research scanner to facilitate the translation of experimental technology to pre-clinical use. Four specific aims were outlined to achieve this objective:

Aim 1: Development of a transmit RF coil array.

Aim 2: Development of a control system to drive the parallel transmission amplifier hardware.

Aim 3: Calibrate the transmit chain for improved accuracy.

Aim 4: Develop watchdog monitoring support software to protect against hardware anomalies.

Following an introduction to the subject matter, the aims translate clearly into chapter divisions. Chapter I of this dissertation has introduced the motivation for high field imaging and the associated engineering challenges and has presented a brief introduction to the concept of parallel transmit.

Chapter II aims to provide a brief background on the fundamental concepts of the MR experiment, introduce a parallel transmit pulse technique, and present an overview of the intended application.

Chapter III discusses the fulfillment of the first aim, the development of the transmit RF coil array. The evaluation process for choosing an appropriate design and the required measurements for verifying functionality are discussed.

Chapter IV covers the control system developed to drive the parallel transmission hardware and documents the completion of aim 2. The factors considered in specifying the system are discussed along with a detailed description of the programming implementation. The chapter concludes with parallel transmit results obtained on a 4.7 Tesla Varian animal scanner.

Chapter V discusses the algorithm used to correct the non-linearities in the transmit chain and represents the completion of aim 3. The method to extract the amplitude and phase data to characterize the non-linearites of the system is described in depth along with the routine used to perform the amplitude and phase pre-distortion.

Chapter VI describes the watchdog monitoring software and hardware used to protect against gross hardware anomalies. The monitoring philosophy is discussed with a brief overview of current monitoring techniques available for parallel transmission systems.

Chapter VII briefly summarizes the primary contribution of the work, details future work, and presents human data demonstrating the intended application.

CHAPTER II

BACKGROUND

II.1 MR Experiment

It is well-known that isotopes with an odd number of protons and/or neutrons have an intrinsic magnetic moment or "spin". In magnetic resonance imaging, hydrogen is most commonly interrogated. Its nucleus (the proton) has an odd spin moment that precesses about a magnetic field at the Larmor frequency, which in turn depends on the magnitude of the field. The Larmor equation is as follows:

$$\omega_0 = \gamma B_0 \text{ or equivalently, } f_0 = \frac{\gamma B_0}{2\pi} \qquad [2.1]$$

Where ω_0 is the angular frequency of precession, γ is the constant gyromagnetic ratio (for hydrogen, $\gamma = \frac{(2\pi)(42.57\ MHz)}{T}$), and B_0 is the static magnetic field assumed to be in the z-direction.

To obtain spatial information about the location of a spin, gradients are momentarily imposed which add to the static field, and the Larmor equation is adjusted as follows:

$$\omega(x, y, z) = \gamma\left(B_0 + G_x x + G_y y + G_z z\right) \qquad [2.2]$$

Where x, y, and z denote the spatial coordinates in the field and G_x, G_y, and G_z are the linear gradient fields in units of Tesla/meter.

With this in mind, the MR experiment begins with the sample being inserted into the static field. Without the presence of any additional gradients, all hydrogen nuclei

precess at the same frequency and align with or against the static field. A greater number of nuclei align with the field (the low energy state) than in the anti-parallel direction (the high energy state), and this gives rise to the magnetization vector defined as follows:

$$M_0 = \frac{p_0 \gamma^2 \hbar^2}{4kT} B_0 \qquad\qquad [2.3]$$

Where M_0 is the equilibrium magnetization vector, p_0 is the number of protons per unit volume, $\hbar$ is Plank's constant divided by 2π, k is Boltzmann's constant, and T is the absolute temperature. From the equation, it is evident that the equilibrium magnetization vector increases linearly with the static magnetic field giving rise to the higher SNR available at high field strengths as mentioned in Chapter I.

Consider the gradient echo sequence shown in Figure II.1.

With the magnetization vector aligned with the static field B_0 along the z-axis, a slice select gradient in combination with an RF pulse are applied. The slice select gradient adds a linear gradient in the z-direction, which adds to the static field and forces spins along the z-axis to precess at different frequencies corresponding to magnitude of the field at their spatial location as defined by the Larmor equation. Radiofrequency coils create a magnetic field in the transverse direction to the static field. A short RF pulse is applied with some finite bandwidth and center frequency that corresponds to a "slice" of spins spatially located along the z-axis. The transverse field created by the RF rotates the magnetization vector away from its alignment along the B_0 axis (the z-axis) at an angle defined as follows:

$$\alpha = \gamma \int_0^\tau B_1(t)\, dt \qquad\qquad [2.4]$$

Where α is the tip angle, B_1 is the transmit field created by the RF coil, and τ is the duration of the pulse. After the RF pulse is applied, the spins precess and relax back to align with the static field along the z-direction. The precession induces an electromotive force (EMF) in the RF coil (according to Faraday's law of induction), which gives the detected signal in MR, but before the signal is measured, localization must occur in the remaining two dimensions.
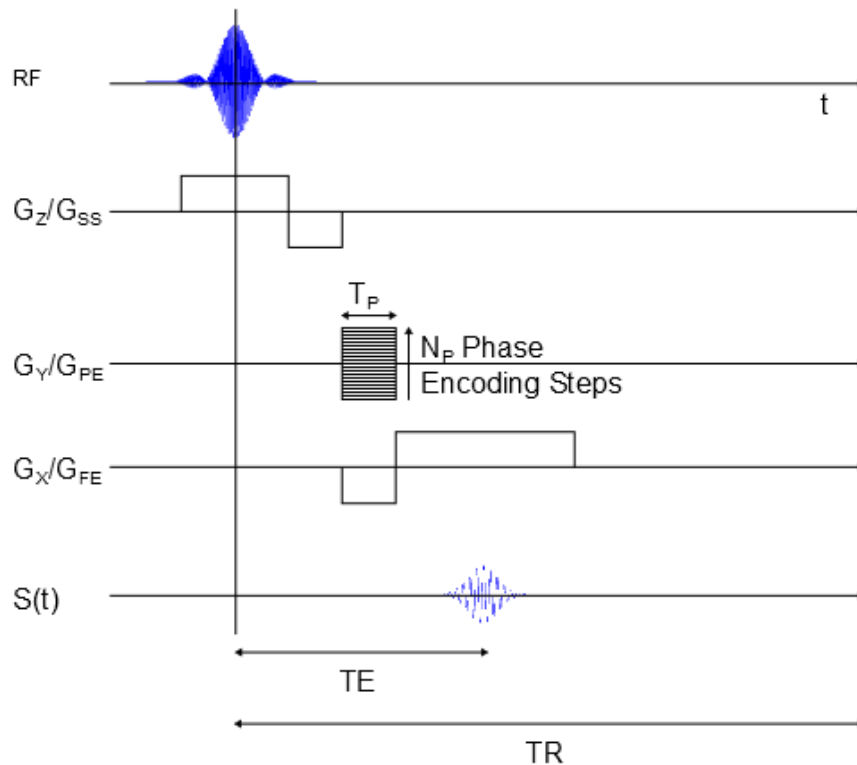
Figure II.1: Gradient echo pulse sequence.
RF is the pulse from the RF coil, G_{SS} is the slice select gradient (in the z direction for this discussion), G_{PE} is the phase encode gradient (in the y direction), and G_{FE} is the frequency encode gradient (in the x direction). S(t) is the acquired signal. TE is the echo time, and TR is the repetition time and is defined as the time between successive pulse sequences.

8

The slice select gradient localizes the spins in one dimension. To obtain a two-dimensional image, the spins need to be localized in the x and y directions. The phase encoding gradient creates a z-directed linear gradient in y that adds to the static field and imparts a phase to the spins allowing the spins to be differentiated spatially by their phase. The frequency encoding gradient is a z-directed x gradient that applies a spatially dependent frequency distribution across the sample. The frequency encoding gradient is on during acquisition, and a digitizer samples the resulting signal from the RF coil. The image information is contained in the sampled signal and is dependent on the transverse magnetization, the relaxation time due to spin-spin interactions, and the applied transverse transmit field, and for convenience, the image will be denoted as $I(x, y)$. Therefore, the signal detected by the RF coil is defined as:

$$S(t, T_p) = \iint I(x, y) e^{-j\gamma G_{FE} x t} e^{-j\gamma G_{PE} y T_p} dx dy \qquad [2.5]$$

Where S is the detected signal and T_p is the duration of the phase encoding gradient. The equation is simplified by substituting equation 2.6 into 2.5 as follows:

$$k_x = \gamma G_x t \text{ and } k_y = \gamma G_y T_p \qquad [2.6]$$

$$S(k_x, k_y) = \iint I(x, y) e^{-jk_x x} e^{-jk_y y} dx dy \qquad [2.7]$$

Where $S(k_x, k_y)$ is the detected signal from the RF coil in the frequency domain termed "k-space", and the integral is taken over the bounds of the slice. The gradient echo sequence is repeated multiple times with different magnitudes of the phase encoding gradient to acquire all points in k-space. A primary mechanism for reducing scan time in the MR experiment is to reduce the sampling requirements of k-space. It is evident from equation 2.7 that the detected k-space signal is simply the two-dimensional Fourier

transform of the image, and the image is obtained by simply applying a two-dimensional inverse Fourier transform to the acquired signal. Equations from this section were largely taken from [19] and [20].

II.2 Functional MRI and the BOLD Signal

Neural activity is closed linked to changes in blood flow and blood oxygenation in the brain [21]. The local susceptibility (the degree of magnetization) in the blood changes with blood oxygenation due to the state of the hemoglobin (red blood cells): oxyhemoglobin or deoxyhemoglobin. Oxyhemoglobin is diamagnetic and deoxyhemoglobin is paramagnetic. With increased neural activity, blood flow increases to the activated region, causing an increase in oxyhemoglobin and a decrease in deoxyhemoglobin. The susceptibility in the region of activation changes over the timecourse of the hemodynamic response leading to a detectable blood oxygenation level dependent (BOLD) signal that is the basis for functional MRI (fMRI).

The image intensity in MR images is determined by contrast mechanisms, and in fMRI experiments, T_2^*-weighted images are largely influenced by the presence of oxygenation in the blood [22]. T_2^* decay is caused by incoherence in the magnetization vectors of the spins due to interference between spins and field inhomogeneities (i.e. susceptibility artifacts). The diamagnetic properties of oxyhemoglobin have a small magnetic susceptibility effect, and therefore, the presence of oxyhemoglobin does not significantly alter the T_2^* properties of the region. Deoxyhemoglobin is paramagnetic and distorts the local field in the tissue region leading to a large observed T_2^* effect. The relative difference in concentrations of oxyhemoglobin and deoxyhemoglobin change

10

the observed T_2^* behavior over time, and therefore, a region of activation will change from having a short T_2^* decay to a long T_2^* decay, and an increase in the MR signal intensity will occur. It is this difference in the MR signal over time that allows regions of the brain to be "mapped" to correspond to functional tasks. Depending on the region being stimulated, the expected increase in the signal due to activation varies. In an early study on activation due to a visual stimulus, the change in the signal intensity was observed to be between 5-20% [22]. In general, the change in signal intensity is small. Another study on cognitive function in the prefrontal cortex reported signal intensity changes ranging from 2.2-5.7% depending on the region [23], indicating the need for high stability in a system performing fMRI experiments.

II.3 Imaging the Orbitofrontal Cortex

The orbitofrontal cortex (OFC) is found in the prefrontal cortex region. It is one of the least understood regions of the brain and is linked to the emotional response and risk/reward behavior [24]. The OFC is also known to be a player in obsessive compulsive disorder (OCD) [25]. Unfortunately, the OFC is often inaccessible for functional MRI (fMRI) studies due to magnetic susceptibility effects caused by the close proximity of the OFC to the air-filled frontal sinuses [26]. Air is paramagnetic and causes field distortions in the region of the brain near the sinuses. As mentioned in the previous section, this susceptibility artifact causes a decrease in T_2^* decay and confuses the interpretation of fMRI studies. With two sources of T_2^* decay, it is difficult to distinguish between magnetic susceptibility due to the air-filled sinuses and susceptibility due to deoxyhemoglobin. This severely limits the ability to adequately

assess activation and the orbitofrontal cortex, and it is the aim of this work to develop parallel transmission technology and enable collaborators with pulse design expertise to create robust methods for reliably imaging and determining activation in the orbitofrontal cortex.

II.4 Transmit SENSE

The targeted approach for addressing the signal loss issue in the orbitofrontal cortex is to use a multidimensional spatially tailored RF pulse method termed transmit SENSE. Sensitivity encoding (SENSE) for accelerating imaging on the receive side was introduced in 1999 by Pruessman et al. [27]. The theory behind SENSE involves the use of the sensitivity patterns of the individual receive coils to reduce the required k-space sampling in the Fourier domain. Instead of solely relying on the linear gradients to encode and spatially localize the spins in k-space, the coil patterns themselves are used to encode information.

With the arrival of parallel transmit technology, the priniciples of parallel imaging and the SENSE method were applied to RF transmission [14, 16]. In parallel imaging, the aliased artifact due to a constrained k-space trajectory is avoided by taking into account the coil sensitivity information during reconstruction [27-28]. Analogous to the parallel receive case, each transmit coil excites a specific magnetization pattern. Undersampling the excitation k-space produces artifacts, but by superimposing the information from multiple coils an artifact-free magnetization pattern is excited. In an N element array, with each element exhibiting a characteristic sensitivity profile $S_i(x)$, the desired excitation pattern $P_{des}(x)$ is obtained by determining which spatial patterns

$P_i(x)$ to excite by each of the transmit coils such that the superposition of the excitations yields the desired pattern.

$$P_{des}(x) = \sum_{i=1}^{N} S_i(x) P_i(x) \qquad [2.8]$$

The coil sensitivity patterns are determined by B_1 mapping techniques. The Bloch-Siegert method was used in the parallel transmit experiments described in this work to determine the individual coil B_1 patterns [29]. The pulse design method essentially weights the desired excitation pattern by the sensitivity pattern. To determine the waveforms required to excite the pattern, the Fourier transform of the desired B_1 magnetization pattern is computed along some trajectory. The first proposal of transmit-SENSE performed the computation in the Fourier domain [16] and was later adapted to the spatial domain which enables greater flexibility in adding additional requirements/constraints to the pulse design [17].

The transmit-SENSE algorithm enables multidimensional spatially tailored RF pulses to be accelerated due to the ability to undersample the excitation k-space, but in order to be implemented in a practical setting, parallel excitation systems are required with the ability to output RF pulses with varying timecourses (i.e. different amplitudes, phases, and pulse shapes) on all channels. The focus of this work is in developing the hardware to enable transmit-SENSE experiments.

CHAPTER III [*]

DESIGN OF THE EIGHT CHANNEL ARRAY

III.1   Introduction

A number of research groups have used multiple transmit channels to address the

challenges at high field strengths. Approaches range from relatively straightforward B_1

shimming [30-33] to more complex transmit SENSE techniques in which separate RF

excitation pulses are sent to each channel [14, 16-17]. With either approach, however,

the level of independence between transmit channels is a concern. Current applied to

one element will induce a voltage (more accurately, an electromotive force or EMF) in

other elements if they have any mutual impedance. This is generally referred to as

"coupling" between coils. If the induced EMF generates currents in the other elements,

this contaminates the desired pattern. Therefore, it is desirable to either eliminate the

EMF or eliminate the currents generated by the EMF, commonly referred to as

"decoupling". There are a number of possible ways to mitigate the effects of coupling,

such as geometrically overlapping the coils or constructing a lumped element network

on the coil to cancel the mutual inductance [34], introducing a high impedance across the

terminals of the coils as in the case of using isolating preamplifiers [18], or some combination of the two. In the first case, the current induced on the secondary coil is canceled; in the second case, current is prevented from flowing in the coil. To generalize, array coil design, specifically with respect to the degree of on-coil decoupling required, depends on and operates in concert with the preamplifier regime in the receive case and the amplifier regime in the transmit case [14].

Most transmit array coils are designed for use with standard RF power amplifiers, relying on decoupling of the coil elements themselves, just as with receive arrays operating with conventional preamplifiers. In this case, the elements are matched to 50Ω for optimal power transfer from the 50Ω amplifier. Some combination of geometric overlap and lumped element decoupling networks must be employed on these arrays to enable independent operation of the channels. Geometric overlap is limited in application to adjacent elements and imposes constraints on the array geometry. Lumped element networks require no overlap, but as the channel counts increase, the network required to fully decouple all elements increases in complexity due to the increasing number of decoupling capacitors required [34]. As an example, to decouple seven elements with a lumped element ladder network, three decoupling capacitors are required for each ladder stage [35], giving a total of 21 capacitors needed for full decoupling. In extending the approach to eight elements, four decoupling capacitors are required at each ladder stage for a total of 32 decoupling capacitors. The values for the decoupling capacitors are guided by closed-form equations that require iterative non-linear and numerical field solvers to compute [36], and further fine adjustments in

element tuning are done iteratively and experimentally. The amount of available decoupling is also sensitive to loading. As indicated in one study [35], the decoupling between opposing elements of a four-element array was found to increase from -30 dB in the unloaded case to -10 dB in the loaded case, indicating 10% of the power from one element being coupled into its opposing element. Due to these complications involved with the various on-coil decoupling strategies, stripline elements and shielded loop elements have emerged as the two most common designs for parallel transmit arrays in neuroimaging due to their inherent favorable coupling properties [31, 37-41]. Even still, an additional decoupling network is typically needed, adding complexity to the array design [31]. Recently, alternative approaches to amplifier design have been investigated that provide some degree of isolation between channels, decreasing or even eliminating the need for decoupling between the coil elements themselves.

One approach under investigation is current source amplification in which each series resonant element is driven with a prescribed current, insensitive to the effects of loading, including element-to-element coupling [37-41]. This allows for large amounts of flexibility in the design of the array coils used with current sources, as in principle no on-coil decoupling techniques are needed. While the current source approach provides high isolation, it is power-limited due to the high mismatch between the amplifiers and the series resonated coil elements [42]. Thus, the peak current produced is significantly lower than with a standard power amplifier. Recent work on current mode class-D (CMCD) amplifiers has shown to improve efficiency and may be able to produce higher

peak output levels than linear current source amplifiers [40-41, 43]; however, the CMCD design has other complications that must be addressed in the process [41].

Chu et al. introduced the "ultra-low output impedance (ULOI) amplifier", which provides decoupling in a manner analogous to preamplifier decoupling in parallel receive applications [44]. The array elements are matched to 50Ω input impedance using a network that forms a trap when connected to the ULOI amplifier or preamplifier. ULOI amplifiers are power matched and present a low impedance to the coil. The power match enables peak output levels comparable to standard power amplifiers while the low impedance provides isolation. The isolation obtained with ULOI amplifiers is substantially lower than that of current source amplifiers, but the peak current delivered to the coil is higher. In practice, the ULOI amplifiers represent a "middle ground" between current source and standard power amplifiers with respect to isolation and output power. We chose to implement a parallel transmit system with ultra-low output impedance amplifiers to achieve greater power output over previously built current sources [39] and to achieve decoupling benefits to simplify the on-coil decoupling network. This chapter discusses the design considerations and array coil progression, ultimately, concluding with the construction and characterization of the final eight-channel transmit/receive head array for use with an ultra-low output impedance parallel transmit system. The capacity of the ULOI amplifiers to eliminate the need for a complex lumped element network to decouple the eight-element array is verified. The characterization of the isolation provided by the stages of the system, from amplifier to coil, is discussed in detail. The following highlights the progression of the discussion:

17

- A brief discussion on series and parallel resonance

- Investigation into the appropriate element design and amplifier regime

- Element fabrication, array coil assembly, and phantom construction

- Evaluation of the isolation provided by the decoupling amplifier and the decision to implement a single-stage on-coil capacitive decoupling network

- Measurement considerations and decoupling network placement when using ultra-low output impedance amplifiers

- Testing procedure to validate array coil functionality on the bench

- Acquired coupling matrices for the array and the measured field patterns using the Bloch-Siegert mapping method

III.1.1 Series and Parallel Resonant Elements

Two modes of element tuning were implemented during the array coil investigation with different amplifier regimes: series and parallel resonance. Series resonance is required for operation with current source amplifiers and was accomplished by adding a capacitor in series with the current path to tune out the reactance of the element and maintain a resistance close to a short circuit to maximize on-coil current as shown in Table III-1.

To compute the value of the capacitor in the matching network, the coil impedance is measured without the matching network. A port extension is applied to compensate for the electrical length of the measurement probe. Without a proper port extension, the measured reactance of the coil is invalid. Once the coil impedance $R + j\omega L$ is known, the capacitor value is calculated using the knowledge that the

18

impedance looking into the coil after the matching network is $R + j0\Omega$ in the resonant condition as follows:

$$R + j0 = jX_{C_1} + R + j\omega L \qquad\qquad [3.1]$$

$$0 = Im\left\{\frac{1}{j\omega C_1} + R + j\omega L\right\} \qquad\qquad [3.2]$$

$$0 = \frac{-1}{\omega C_1} + \omega L \qquad\qquad [3.3]$$

Solving the above equation determines the value of C_1. If the rung impedance is $3 + j25\Omega$, the value of C_1 is 49.7pF.

Parallel resonance is the conventional resonance scheme used with standard 50Ω amplifiers and ULOI amplifiers (modified to include a trap circuit discussed in detail in the following sections). The basic matching network consists of a shunt and series capacitor at the element feed (shown in Table III-1). The shunt capacitor is referred to as the tuning capacitor and transforms the resistance of the element to 50Ω. The series capacitor is known as the match capacitor and tunes the remaining element reactance to zero to resonant the element. The capacitors in the parallel resonance matching network are calculated by first recognizing that the impedance looking into the coil after the shunt capacitor C_2 is $50 + jX\Omega$ in the resonant condition. Capacitor C_2 is in parallel to the rung, and the impedance after the capacitor C_2 is computed as follows:

$$X_{C_1}||R + j\omega L \qquad\qquad [3.4]$$

$$50 + jX = \frac{1}{\frac{1}{1/j\omega C_1} + \frac{1}{R + j\omega L}} \qquad\qquad [3.5]$$

$$50 + jX = \frac{1}{j\omega C_1 + \frac{1}{R + j\omega L}} \qquad\qquad [3.6]$$

19

$$50 + jX = \frac{R+j\omega L}{j\omega C_1 R - \omega^2 C_1 L + 1} \tag{3.7}$$

$$50 + jX = \frac{R+j\omega L}{j\omega C_1 R - \omega^2 C_1 L + 1}\left(\frac{-j\omega C_1 R - \omega^2 C_1 L + 1}{-j\omega C_1 R - \omega^2 C_1 L + 1}\right) \tag{3.8}$$

$$50 = Re\left\{\frac{R+j\omega L}{j\omega C_1 R - \omega^2 C_1 L + 1}\left(\frac{-j\omega C_1 R - \omega^2 C_1 L + 1}{-j\omega C_1 R - \omega^2 C_1 L + 1}\right)\right\} \tag{3.9}$$

$$50 = \frac{R}{\omega^2 C_1{}^2 R^2 + (-\omega^2 C_1 L + 1)^2} \tag{3.10}$$

This equation is solved for the value of capacitor C_1. To find the value of C_2, the reactance after capacitor C_1 is computed by taking the imaginary part of equation 3.8.

$$X = Im\left\{\frac{R+j\omega L}{j\omega C_1 R - \omega^2 C_1 L + 1}\left(\frac{-j\omega C_1 R - \omega^2 C_1 L + 1}{-j\omega C_1 R - \omega^2 C_1 L + 1}\right)\right\} \tag{3.11}$$

$$X = \frac{\omega L - \omega C_1 R^2 - \omega^3 C_1 L^2}{\omega^2 C_1{}^2 R^2 + (-\omega^2 C_1 L + 1)^2} \tag{3.12}$$

Capacitor C_2 is calculated from knowing that the impedance looking into the matching network is $50 + j0\Omega$. Capacitor C_2 is in series with the impedance $50 + jX\Omega$, and therefore, capacitor C_2 is determined as follows:

$$50 + 0j = jX_{C_2} + 50 + jX \tag{3.13}$$

$$0 = Im\left\{\frac{j}{\omega C_2} + 50 + j\left(\frac{\omega L - \omega C_1 R^2 - \omega^3 C_1 L^2}{\omega^2 C_1{}^2 R^2 + (-\omega^2 C_1 L + 1)^2}\right)\right\} \tag{3.14}$$

$$0 = \frac{1}{\omega C_2} + \frac{\omega L - \omega C_1 R^2 - \omega^3 C_1 L^2}{\omega^2 C_1{}^2 R^2 + (-\omega^2 C_1 L + 1)^2} \tag{3.15}$$

Since the value of C_1 was previously determined, C_2 is computed from equation 3.15. As an example, given a rung impedance of $3 + j25\Omega$, the values of C_1 and C_2 are determined from equations 3.10 and 3.15 to be 38.4pF and 13.8pF.

The coil impedance is adjusted by setting an appropriate value for the distributed capacitors to produce a coil reactance that results in reasonable matching network values.

	Series Resonance	Parallel Resonance
Schematic		
Coil Impedance	3+j25 Ω	3+j25 Ω
Computed Values	C_1=49.7 pF	C_1=38.4 pF C_2=13.8 pF
Smith Chart View		

Table III-1: A table comparing series and parallel resonance.
The circuit diagram for the matching network in each case is shown. The series resonant element appears as a short looking into the matching network whereas the parallel resonance circuit measures at 50Ω. The Smith Chart View shows how the computed capacitor values in the matching network transform the coil impedance to a short for series resonance and 50Ω for parallel resonance. The values were computed assuming a coil impedance of $3 + j25\Omega$ and following the equations in the discussion above.

III.2 Materials and Methods

III.2.1 Array Coil Design Investigation

In choosing an appropriate element design, unshielded loop, dual plane pair [45-46], and rung elements were considered. The array was ultimately intended to be used in a transmit/receive configuration with a 3 Tesla GE clinical research scanner, using the

add-on eight-channel prototype transmit system and the existing GE eight-channel receiver. The first prototype transmit system used current source amplifiers. In initial parallel transmit experiments, a transmit-only option was viewed as a straightforward first step with the GE body coil performing receive. Parallel transmit pulse design relies on accurate transmit field (B_1) maps, and with a uniform body coil for receive, reconstruction complexity was minimal and fields maps were easily extracted. A challenge in the transmit-only array/body coil receive implementation was the inability to easily detune the body coil during transmit. The transmitter required the output to the GE modulator, and in order to obtain the hard pulse, the scanner configuration was set to mimic a body coil transmit experiment. Therefore, for any transmit-only array/body coil receive option, coupling to the body coil during transmit factored in as an additional design consideration.

Conventional transmit/receive array configurations consist of an outside transmit coil and a smaller receive coil insert to maximize the signal-to-noise ratio. In our initial setup, the outside body coil was intended to be used for receive which would degrade SNR by increasing the noise level, and therefore, a transmit element with a high sensitivity pattern was desired to maximize the achievable signal level. Loop elements are RF transparent and have a high sensitivity profile. The B_1 field of both loop and dual plane pair elements were measured using a field probe to determine the expected sensitivity from both elements. Due to the higher sensitivity profile of the loop element, adjacent elements were more highly coupled, but element-to-element coupling would be minimized with the isolation provided by the current source amplifiers and was not

considered a constraint. The more pressing issue was the degree of coupling between the loop element and the body coil, which would cause some amount of detuning inside the bore. A standard S_{21} port measurement between the body coil and element was not easily managed due to the proprietary nature of the GE hardware connections. Instead, a dual plane pair element and loop element were resonated and matched to 50Ω and inserted inside the resonant body coil to evaluate the tuning shift. The loop element indicated significant coupling to the body coil with mode-splitting evident in the S_{11} measurement while the dual plane pair tuning shifted minimally as shown in Table III-2. Due to the high level of coupling to the body coil, the loop element was considered unsuitable.
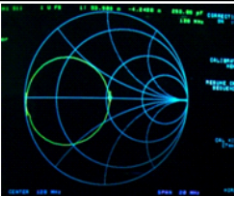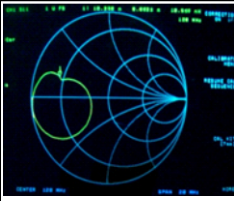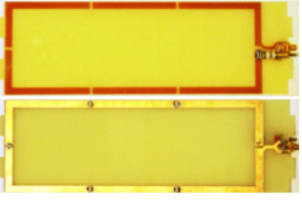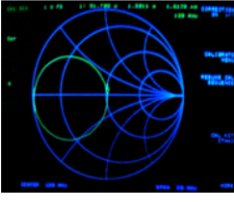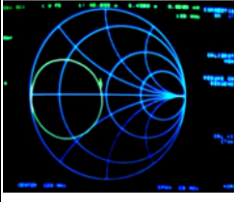
	Element	S_{11} Outside Bore	S_{11} Inside Bore
Dual Plane Pair			
Loop			

Table III-2: Table showing coupling to the body coil.
Coupling is evidenced by the S_{11} measurements made using the two element designs: dual plane pair and a loop. Outside the bore both element designs are well-tuned. Inside the bore the dual plane pair exhibits a minimal shift in tuning while the loop shows mode splitting, indicating a high degree of coupling.

Dual plane pair elements [45] (discussed in detail below) were implemented in a transmit-only eight-channel head array configuration for use with current source

amplifiers. Dual plane pair elements are RF transparent allowing the body coil to receive [46]. The dual plane pair design effectively decoupled the array from the uniform field of the body coil. The dual plane pair element consisted of a current path that followed the center conductor and split at the side opposite the feed to return along the two ground conductors. The element creates two current loops with opposing fields, effectively decoupling the element from the uniform external field of the body coil as shown in Figure III.1. During transmit element-to-element isolation was achieved using current source amplifiers. The array coil exhibited a high degree of isolation between elements because the amplifiers excited the elements with prescribed currents, minimizing coupling effects. The dual plane pair transmit-only array was built and tested on the GE 3 Tesla clinical scanner and was demonstrated to produce well-isolated field patterns as shown in Figure III.2.

The dual plane pair array with the current source amplifiers provided well-isolated sensitivity patterns, but the design was susceptible to noise due to the volume receive coil and power limitations caused by the impedance mismatch between the current source amplifiers and the series-resonant element. In order to achieve more power output per channel for more flexibility in pulse sequence design and applications, we moved to a recently investigated ultra-low output impedance amplifier configuration [44]. The decoupling amplifier operated as an analog on the transmit side to the low input impedance preamplifiers used for isolation during receive. As such, the coil element required a parallel resonant matching network scheme identical to those used on receive arrays for preamplifier decoupling. The 50Ω operation made a transmit/receive

configuration straightforward. In the transmit case, the ultra-low output impedance amplifiers provided isolation, and in the receive case, the low-input impedance preamplifiers provided isolation. This led to a fully shielded transmit/receive array design using a rung element and eliminated the need to use the body coil for receive.
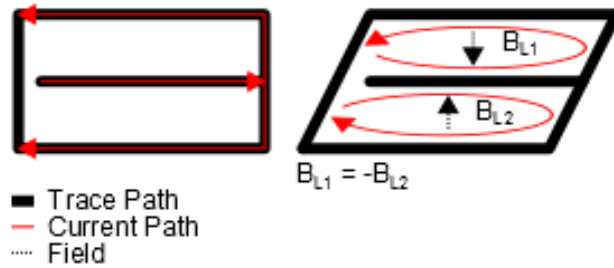
Figure III.1: Diagram of the dual plane pair element.
The center conductor is in a plane above the return ground conductors (not depicted). Current travels along the center conductor and splits at the opposite end of the coil feed before returning along the two ground traces. This creates two current loops that produce opposing fields that effectively decouple the element from a uniform external field.
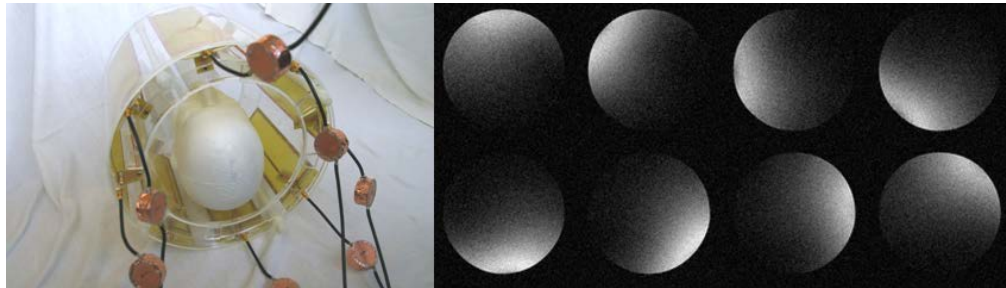
Figure III.2: Completed dual plane pair array (left) and single channel gradient echo images (right).
The array was used in a transmit-only configuration with current source amplifiers. The images show well-isolated patterns.

The first prototype of the eight-channel transmit/receive shielded rung array exhibited instabilities discussed in Figure III.3. The array was fabricated on an acrylic former with the supporting transmit/receive switches and preamplifier hardware housed on a separate acrylic piece. This created significant wear on the cables connecting the coil to the transmit/receive switches and caused instabilities in the form of intermittent connections. Additional instabilities occurred because the dc bias for the preamplifiers was supplied by an external power supply outside the bore. The cables ran the length of the bore and picked up gradient noise during the scan affecting the stability of the preamplifiers. Yet another issue was ineffective grounding that manifested as changes in element tuning when grounds at different stages of the transmit chain were connected at secondary locations.
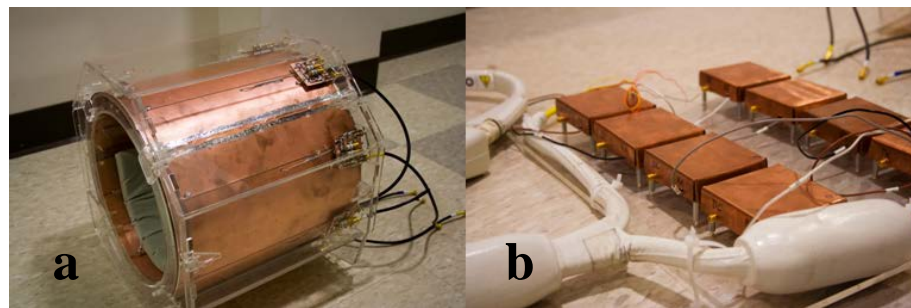
Figure III.3: Photos of the first prototyped fully shield array.
(a) Photo of the shielded array coil. (b) Photo of the transmit/receive switches and preamplifiers with the GE receive connector attached. The transmit/receive switches were separate from the array coil and handling caused quick degradation to the cables connecting the array to the switches resulting in instabilities.

A second iteration of the array coil was built to address the stability issues. A single structure was designed to house the array coil, transmit/receive switches, and

preamplifier hardware to minimize stress on cables and prevent intermittent connections. A solid ground plane was used to tie all coil, transmit/receive switches, and preamplifier grounds together. The cables providing the dc bias to the preamplifiers were eliminated, and instead, the dc bias was pulled from the GE receiver. The remaining sections discuss the details of this second iteration constituting the final eight-channel transmit/receive head array [47-48].

III.2.2 Array Coil Fabrication and Housing Design

The eight-channel head array [47] was designed with shielded rungs fabricated using copper sheet metal mounted to a ½ inch wide, 25 cm long acrylic piece with six breaks with 79 pF of capacitance at each break (Passive Plus, 1111C Series, Huntington, NY). The mounting piece for all the element hardware was a 12-inch outer diameter cylinder fabricated from white polycarbonate using a fusion deposition modeling (FDM) rapid prototyping machine.

The shield consisted of two layers of single-sided ½ ounce copper Pyralux (AC182500E, DuPont, Research Triangle Park, NC) mounted to the 12-inch cylinder. Each layer was slotted longitudinally to mitigate eddy currents and the two layers were oriented to alternate the position of the longitudinal slits. Rectangular slots were removed from the Pyralux shield and replaced with copper mesh to provide a view port through the coil for patient comfort and visual stimulus in functional imaging studies.

The elements were mounted on the inside of the cylinder with connections to the shield at one end and the matching network at the other. With the ½ inch thickness of the acrylic rung support and the ¼ inch thick cylinder, each rung was ¾ inch from the

shield, creating an effective 10.5 inch diameter coil array. Shielded current probes adjacent to each rung dropped down through the shield and connected separately on the match and tune board. Baluns for both the element and current probe were mounted to the match and tune boards.

A single-sided ¼ ounce Pyralux end piece (AC091200EV) connected the shield segments near the top of the head to prevent RF fields from interacting with the rest of the transmit and receive chains. The transmit/receive switches and low-input impedance (R=1.5Ω) preamplifiers (WMA3RA, WanTCom Inc., Chanhassen, MN), were mounted to an acrylic piece covered in a Pyralux ¼ ounce ground plane. All hardware components were contained within a modular polycarbonate housing manufactured in-house using a 3D printer. A diagram of the array coil hardware with photos of the array and housing are shown in Figure III.4.

Modularity was added as a design criteria in developing the array housing to enable quick access to the hardware beneath in the event of a failure or needed modification. In the full assembly, the coil array with the transmit/receive switches and preamplifiers consisted of one piece with mounted cable bundles exiting the back of the coil. Renderings of the modular construction designed in SolidWorks are shown in Figure III.5.
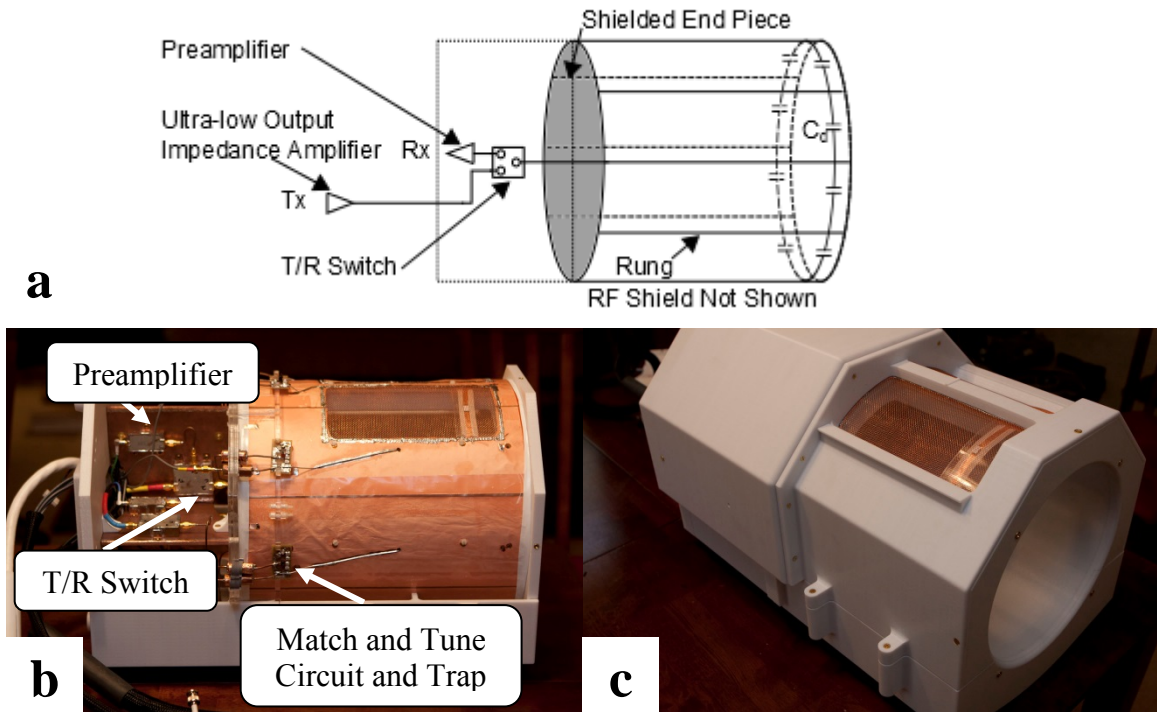
Figure III.4: Diagram and pictures of the array coil setup.
(a) Diagram of the array coil setup showing the rung elements with the decoupling capacitor at the end opposite the coil feed. The preamplifiers and t/r switches are mounted on the coil as indicated and shown in b. The elements are shielded and a shielded end piece (labeled and shaded) prevents stray RF from interacting with the hardware behind. (b) Outside view of the array coil showing the array hardware. (c) Fully assembled array coil shown in its completed housing.
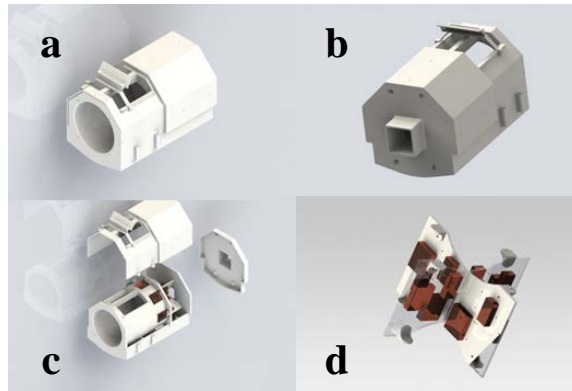
Figure III.5: Renderings of the array coil housing.
(a) Front view of the fully assembled housing. (b) Back view of the fully assembled housing. Cables bundles mount to the exit. (c) Exploded view of the housing showing some of the array coil hardware beneath (shield is not shown). The match and tune boards mount to the 12-inch cylinder, and the piece containing the transmit/receive switches and preamplifiers mounts to the back of the cylinder. (d): The piece containing the transmit/receive switches and preamplifiers. The copper foil ground plane (not shown) was adhered to the x-piece.

III.2.3 Transmit/Receive Switch Design and Construction

A PIN diode (UM9415, Microsemi Corporation) switch configuration with a quarter-wavelength pi section was used for the transmit/receive (t/r) switches as shown in Figure III.6. In the transmit configuration, the PIN diodes are forward-biased and act as short circuits. The transmit port is connected to the coil, and the short in parallel to the receive port is transformed by the quarter-wavelength section to a high impedance at the coil port effectively isolating the transmitter from the receiver. In the receive configuration, the diodes are reverse-biased and act as open circuits. The coil is connected to the receiver, and the transmitter is disconnected from the coil.

The t/r switches were designed to handle a kilowatt of power, requiring additional considerations to protect the preamplifier during transmit. While most of the

power is directed to the coil due to the isolation provided by the quarter-wavelength section, some power is passed to the receive port. The preamplifier could withstand powers up to 30 dBm, and the quarter-wavelength section provided approximately -20 dB of isolation. At a kilowatt of power (60 dB), the quarter-wavelength section alone would not sufficiently isolate the receive port. Therefore, fast-switching PIN diodes (UM9989, Microsemi Corporation) were placed back to back at the receive port behind a dc blocking capacitor to serve as passive protection for the preamplifiers. When the RF voltage exceeded the bias threshold for the diodes, the diode became forward-biased and acted as a short, shunting the power to ground.

PIN diode drivers for each channel were constructed and incorporated into the system at the power amplifier stage. The bias was input onto the RF line at the output of the amplifier.
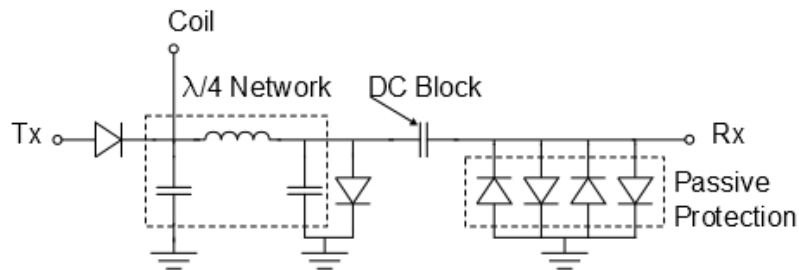
Figure III.6: Schematic for the transmit/receive switches.
A quarter-wavelength section isolates the receive port from the transmit port during transmit, and the fast-switching PIN diodes provide additional passive protection for the preamplifier.

The transmit/receive switches were redesigned in the final array to address instabilities. In the first array prototype, the t/r switch hardware and preamplifier were

31

combined on the same board, and while no known instabilities resulted directly from the board design, the final design split the transmit/receive switch and preamplifier onto two boards to physically separate the high and low power sections. Effort was made to reduce the RF trace lengths of both the t/r switch and preamplifier boards, and phase shifters were added at the transmit port of the t/r switch and the input of the preamplifier to provide another degree of freedom in adjusting the electrical lengths (described in detail below). The dc bias was pulled from the GE system via the receive RF cable, and the previous external dc supply and cables were eliminated. To address the observed grounding instabilities, all t/r switch and preamplifier hardware were mounted to a solid ground plane. Photos of the hardware noting the differences between the first array prototype and final version are shown in Figure III.7.

III.2.4 Phantom Construction

A phantom that simulated the loading conditions of a head was required prior to tuning the array coil elements. The literature describes head phantoms that use nickel sulfate, magnesium dichloride, sodium chloride, and distilled water [49]. The phantom described in [49] uses a solution of distilled water, magnesium chloride, and sodium chloride to simulate the loading properties of the head. A homogenous phantom containing these materials was manufactured, but the uniform solution did not mimic the asymmetric loading characteristics of the head. For instance, the back of the head loads the "bottom" elements more heavily than the "top" elements above the face. Instead, a series of six 12 ounce and six two ounce leak-tight plastic bottles were purchased and filled with distilled water with various quantities of monosodium glutamate and salt in

32

an effort to match the loading of a head. Elements were tuned with the actual head load of a volunteer, and then, monosodium glutamate and salt were added asymmetrically to the vials to load each element appropriately. The phantom was verified by comparing the tuning of the elements on both the phantom and actual head loads of volunteers. Small and large volunteer heads were used to evaluate the tuning as compared across a range of subjects. In all cases, the tuning of the elements remained within a standing-wave-ratio (SWR) of 2:1, indicating the phantom effectively simulated the loading characteristics of the head.
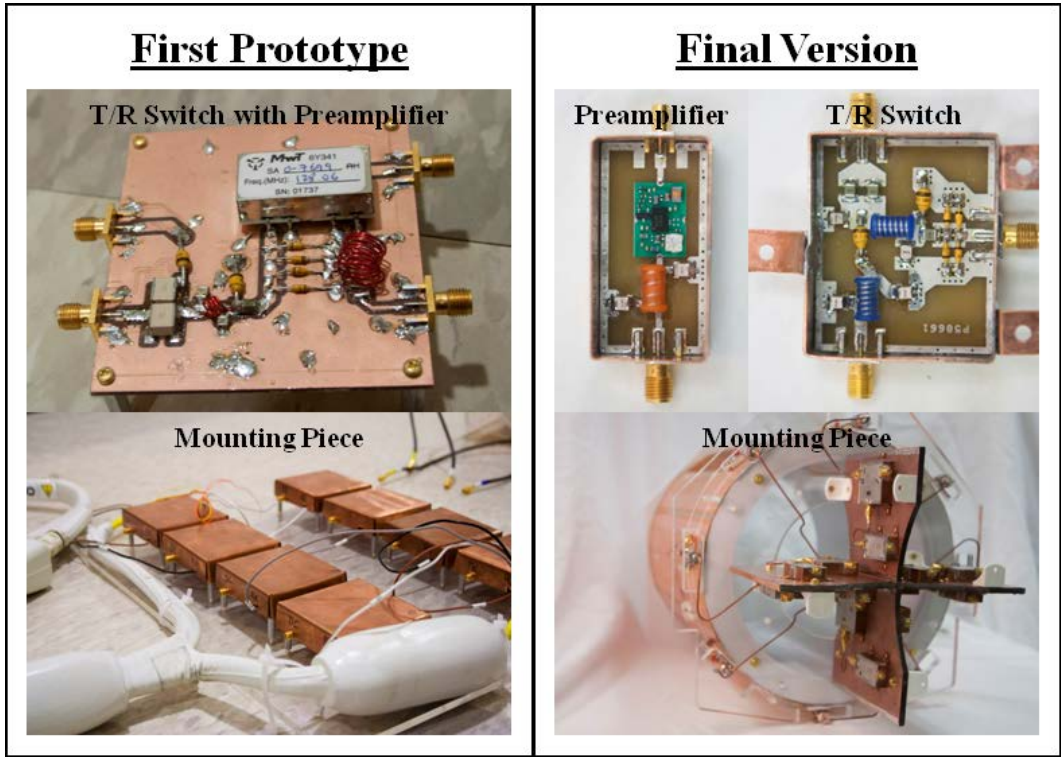
Figure III.7: Comparison of the transmit/receive switch and preamplifier hardware for the first prototype and final version.
Improvements made in the final version eliminated the instabilities.

III.2.5 Coil Matching Network and Transmit/Receive Configuration

The matching network used to match and tune the coil elements to 50Ω is diagrammed in Figure III.8 [50]. Because the ULOI amplifier operates as the analog on the transmit side to the low-input impedance preamplifier on the receive side with respect to providing isolation, the matching network scheme is conveniently the same in both transmit and receive modes. The values for the matching network components were theoretically determined by the equations in [50] and then tuned to account for variations in practice. The equations are detailed below.

$$A = X_L Z_0 + R_L X_{Amp} \tag{3.16}$$

$$B = \sqrt{R_L Z_0 \left(X_{Amp}^{\;2} + Z_0^{\;2}\right)} \tag{3.17}$$

$$C = R_L Z_0 - X_L X_{Amp} \tag{3.18}$$

$$X_{C_2} = \frac{XA}{XZ_0 - B} \tag{3.19}$$

$$X_{C_3} = X \frac{A - B}{B} \tag{3.20}$$

$$X_{L_4} = X \frac{A - B}{A} + Z_0 \frac{C}{A} \tag{3.21}$$

Where X_L is the rung reactance, R_L is the rung resistance, and X_{Amp} is the reactance of the amplifier. The quantity X_{Amp} is useful if trying to compensate for the required $\lambda/2$ electrical delay between the coil feed and the amplifier. For all elements, the X_{Amp} parameter was set to zero. The total electrical length from the coil port to the preamplifier and amplifier accounted for a phase delay of a multiple of λ/2 to ensure the low impedance from the amplifier/preamplifier was presented to the coil port. The phase

shifters in the transmit and receive chain paths enabled adjustments in the electrical lengths.

For each element, the capacitor C_x was selected to result in reasonable values for the remaining components that would not result in high voltages (small capacitance) or large currents (large capacitance) through the capacitors in the matching network. Adjustments in capacitors C_1 and C_2 were made to fine tune the coil. Tuning the trap was accomplished by short-circuiting the coil port (located on the match and tune board at the coil feed) to mimic the presence of the ULOI amplifier or preamplifier, placing the inductor L and the equivalent capacitance of the matching network in parallel to form a trap circuit. Slight adjustments to L were made to minimize S_{21} between a pair of two decoupled loop probes positioned near the rung, corresponding to minimum power being coupled into the coil – that is, when induced currents were suppressed and the amplifier or preamplifier was providing decoupling in the transmit and receive cases respectively. As seen in Figure III.8, the amplifier and preamplifier for each channel are connected to the rung by a PIN diode transmit/receive switch (UM9415, Microsemi, Aliso Viejo, CA) with typical insertion loss of -.2 dB and electrical length of .34λ in the transmit configuration.
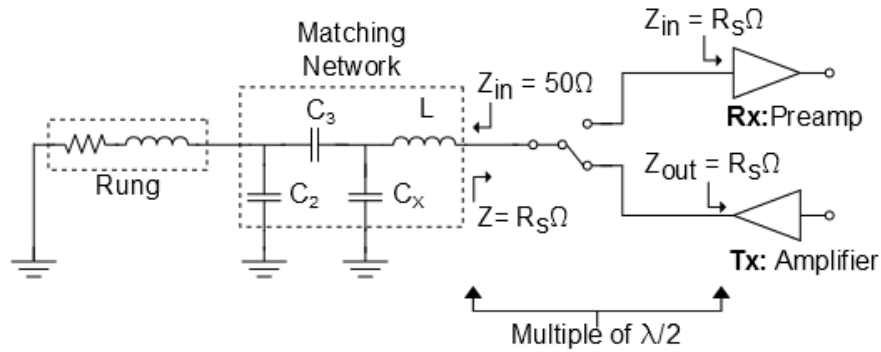
Figure III.8: Block diagram of the transmit/receive configuration of the array and matching network.
During transmit and receive, the coil is presented with a low input impedance that places the inductor L in parallel with the equivalent capacitance of the matching network, creating a trap and suppressing induced currents. The electrical lengths between the coil port and the amplifier/preamplifier are a multiple of a half-wavelength to prevent an impedance transformation. The matching network used was presented in [50].

III.2.6 Validating the On-Coil Probes

On-coil probes were needed to provide information about the relative current on each rung. Before manufacturing probes for all eight channels, two probes were built and tested to determine their dynamic range. The probe needed to sense the field produced by its corresponding rung while being sufficiently isolated from the field of the adjacent rung. The two test probes were secured to adjacent rungs, and power was input into one rung while the induced signal from both probes was monitored. The signal levels from the probes were measured at multiple power levels to verify the stability of the measurement. The coupling between a rung and its corresponding probe was -25.5 dB, and the coupling between the excited rung and the adjacent probe was -50 dB. The measurements indicate that the isolation between two adjacent probes is -24.5 dB.

36

Greater isolation is always preferred, but the -24.5 dB of isolation was acceptable. As a reference, the measurements were reacquired with the adjacent rung removed, and the results were the same, indicating the measurement from the adjacent probe corresponded to the field it sensed from the excited rung and was not due to coupling effects.

The -24.5 dB isolation value between adjacent probes imposed a fundamental limit on measurements made with the on-coil probes. Coupling measurements made with the probes with quantitative values of -24.5 dB or lower were suspect. In this case, the measurement would only qualitatively correspond to a high level of decoupling, and therefore, measurements made with the on-coil probes were capped at a minimum of -24.5 dB. For linearization and monitoring purposes, the probes are suitable due to the high isolation between probes, which corresponds to an adjacent probe sensing less than 0.35% of the power from a neighboring excited rung.

III.2.7 Measurements Overview

Because the decoupling benefits of ULOI amplifiers are realized by the impedance they present at the coil port, it is not always appropriate to characterize or optimize with standard 50Ω port network analyzer measurements. If a measurement needs to be acquired which involves the amplifiers being connected in an "operational" configuration, then this requires the use of a more direct measurement of current on the rungs of the array rather than a port measurement. The measurements described below that rely on the impedance the amplifiers present were acquired with current probes in one of two configurations. Disambiguation will be ensured by referring to them as "on-coil probes" or "crossed-probes", representing respectively, the single

37

current probe on each rung that is built into the coil (described above), or a pair of "free" crossed probes moveable to an individual rung. Three measurements are discussed in the sections that follow:

1) Characterizing the ULOI amplifiers: measuring the isolation per channel provided by the ULOI amplifiers, requiring a crossed-probe measurement on each channel

2) Characterizing the coil: measuring the on-coil decoupling added between elements, accomplished with a standard S_{21} port measurement under certain described conditions

3)   Characterizing the system: measuring the total channel-to-channel coupling with the amplifiers connected, measured with the on-coil probes

III.2.8 Isolation from the Power Amplifier

Because the amount of isolation provided by the power amplifier depends on its output impedance, the architecture is particularly sensitive to losses in the transmit cables and in-line transmit/receive switches. Ideally, the amplifiers might be placed at the coil and inside the bore [41], but to eliminate saturation and susceptibility effects from eight amplifiers with ferrite baluns, heat sinks, and MOSFET devices, as well as potential cooling problems, we opted to locate the power amplifiers outside the bore. This required a transmit cable to span the distance between the transmit/receive switch and the amplifier that maintained an electrical length of $n\lambda/2$ between the coil port and amplifier matching network. As expected, the losses associated with the added cable length and the transmit/receive switches reduce the achievable isolation from the amplifiers. Chu et al. discussed the effect of cable losses and observed that with a $7\lambda/2$

cable (insertion loss of 0.52 dB) at 3 Tesla, the isolation provided by the amplifier decreased to -9 dB as opposed to -14 dB obtained with the amplifier at the coil element [44]. For our particular setup, a $5\lambda/2$ transmit cable was required for each channel to place the amplifiers outside the bore.

Before measuring the isolation provided by the ULOI amplifier transmit system, the trap tuning on the coil was verified by measuring the ideal case with a short circuit at the coil port. An S_{21} measurement between a pair of crossed-probes positioned along the rung was used (see Figure III.9), with one probe exciting the rung, and the other sensing the field produced by the current on the rung. The isolation was calculated as the difference between S_{21} measurements collected with the coil port terminated in a 50Ω load and terminated in a short, in an analogous fashion to measuring the decoupling provided by an isolation preamplifier. With a short at the coil port, the isolation provided by the traps was observed to be -30 dB or better in all cases. The isolation provided by each amplifier was then measured using the same method. The isolation provided by each amplifier was calculated as the difference in the S_{21} measurements between a 50Ω load termination and termination with the amplifier connected through the transmit/receive switch and transmit cable. Based on these measurements, the average isolation provided by a well-tuned amplifier was found to be -11.5 dB with the $5\lambda/2$ transmit cable and transmit/receive switch in place.

To evaluate if the isolation provided by the amplifier independently provided sufficient decoupling between channels, the elements of the array coil were tuned individually with all other elements open-circuited as described above. The tuning of

one element then was observed using an S_{11} measurement while the other elements were connected to the ULOI amplifiers. With the ULOI amplifiers, the element tuning shifted outside a standing-wave-ratio (SWR) of 2:1, indicating insufficient decoupling. While there are other possible solutions, we decided to add a simple capacitive decoupling network on the coil to augment the isolation provided by the amplifiers.
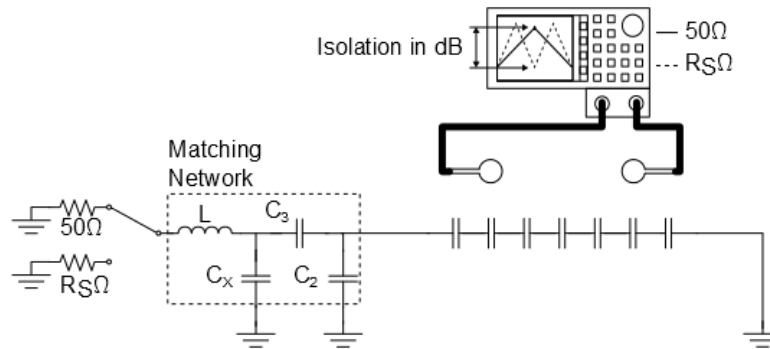
Figure III.9: Diagram of the isolation measurement.
Two decoupled probes are positioned above the rung ("crossed-probes" were additionally used but not shown and consist of two overlapped probes with the overlap set to maximize decoupling). An S_{21} measurement between the probes is acquired with the matching network terminated in 50Ω and a low impedance. The difference in the two S_{21} measurements corresponds to the isolation.

III.2.9 Providing Additional Decoupling Between the Elements

It is important to emphasize that the use of a conventional S_{21} port measurement to quantify coupling between elements is not always appropriate when using ULOI amplifiers, depending on the layout of any electrical connections between rungs. Specifically, the current distribution between two connected rungs is not always the

40

same when a rung is terminated in the 50Ω network analyzer (as with a conventional port measurement) and when a rung is terminated in the high impedance provided by the trap circuit when the rung is connected to an ULOI amplifier (or preamplifier).

For clarification, Figure III.10 illustrates the effect of the different placement of the decoupling capacitor between elements of this particular array. Importantly, placing a decoupling capacitor between rungs at the feed end of the coil and performing a conventional S_{21} port measurement with a network analyzer could provide an inaccurate indication of the coupling between those rungs that would occur if they were fed by ULOI amplifiers. This can be understood more clearly by comparing the current paths in Figure III.10a-b, which illustrate the difference in resulting rung currents (and thus coupling) between termination in the 50Ω network analyzer and termination in an infinite impedance. As shown in Figure III.10c, placing the decoupling capacitor between rungs at the opposite end of the coil in front of the last distributed capacitor avoided this difference between the two termination cases and allowed for use of a direct S_{21} port measurement for setting the decoupling capacitor value between rungs. Placing a capacitor between the rungs changed the rung impedance slightly, requiring a second iteration to adjust matching network and trap tuning, as described above. It is important to emphasize the simplicity, however, of the single capacitor decoupling network that was enabled by the isolation provided by the ULOI amplifiers as compared to conventional on-coil decoupling networks.

41

Current Paths for 50Ω Termination   Current Paths for AmplifierTermination   Current Paths with Decoupling Capacitor Moved

MN: Matching Network   NA: Network Analyzer R_S: Short C_C: Decoupling Capacitor
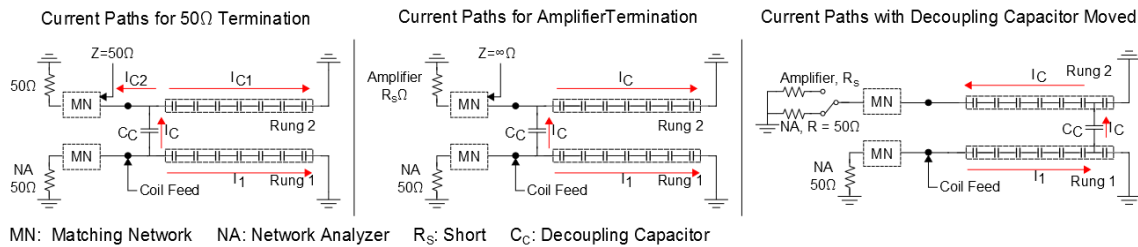
Figure III.10: Diagram showing the current paths for different decoupling capacitor placements to demonstrate appropriate coupling measurements when using ULOI amplifiers.
(a) The decoupling capacitor is placed between the coil feeds and both coil ports are terminated in 50Ω (as is the case for a conventional S_{21} port measurement). The current path splits at the juncture between the rung and decoupling network and splits again at the connection to the adjacent rung due to the 50Ω termination. (b) Instead of a 50Ω load, the adjacent rung is terminated in the ULOI amplifier, labeled with theoretical infinite impedance. The current paths do not split at the adjacent rung, and the ratio of rung currents is different from the case shown in (a) indicating a difference in coupling. Therefore, if a conventional S_{21} measurement were used to choose the decoupling capacitor value, the coupling provided by the network would change in the operating case with the ULOI amplifier (or preamplifier). (c) The decoupling capacitor is placed in front of the last distributed capacitor to avoid the difference in current paths for the two terminations cases allowing for direct S_{21} port measurements to set the decoupling capacitor.

III.2.10 Total Channel-to-Channel Isolation

To then accurately measure total channel-to-channel isolation with the coil decoupling capacitors and amplifiers in place, the on-coil probes were used to measure the relative currents on the rungs. The probes first were calibrated to compensate for differences in probe sensitivity due to slight variations in loop size and orientation with respect to the rung. To do this, a test fixture was fabricated to position a second probe at a repeatable fixed height and orientation above each element. To collect a relative "sensitivity measurement" for each on-coil probe, all other elements were open-

circuited, and the second probe was positioned above the respective element. A network analyzer was used to collect two S_{21} measurements: one between the element and the on-coil probe and one between the element and the second probe positioned above the element. The difference between the two measurements corresponded to the on-coil probe sensitivity and provided a set of calibration measurements for the eight probes. The measurements of the total channel-to-channel isolation obtained using the probes then were adjusted by this measurement.

To measure the isolation, one rung at a time was connected to the network analyzer and S_{21} measurements were collected between it and each of the eight on-coil probes. The measurements correlated to the relative currents on the rungs and thus to the element-to-element decoupling. In this manner 8x8 decoupling matrices (64 individual S_{21} measurements) were collected for two cases: 1) with the coil ports terminated in 50Ω, indicating the decoupling provided on the coil alone and 2) with the coil ports connected to the transmit system, indicating the total channel-to-channel isolation.

To demonstrate the overall system, B_1 maps were acquired on a GE 3 Tesla Signa clinical research scanner from each channel of the transmit system. Mapping was done in the flood region of the American College of Radiology (ACR) phantom using the Bloch-Siegert method with a TR of 200 ms and a TE of 18 ms [29, 51-52].

III.3 Results and Discussion

III.3.1 Bench Measurements and Power Testing

The coil tuning was evaluated using standard Z_{11}, S_{11}, and SWR measurements summarized in Table III-3. The trap tuning was verified by measuring the isolation in the

ideal short-circuited case. Element isolation was verified to be -30 dB or better for all elements in this "ideal" case. Average isolation from the ultra-low output impedance amplifier was -11.5 dB. The average isolation with the preamplifier was -20 dB. The quality factor (Q) of each element was measured using the 7 dB return loss method on the S_{11} plot where Q is:

$$Q = \frac{f_0}{\Delta f_{7dB}}$$ [3.22]

Where Δf_{7dB} is defined as the bandwidth for which the reflected power is 7 dB below the incident power [53]. The loaded Q for all coils was better than 100 and varied around the array due to asymmetric loading.

Element	Impedance (Ω)	S11 (dB)	SWR	Q	Element Layout
1	48.9-6.7j	-23.2	1.15	196	
2	51.2-3.4j	-28.9	1.07	220	
3	50.1-3.4j	-29.4	1.07	167	
4	44.7+1.3j	-24.9	1.12	125	
5	53.3-11.9j	-18.5	1.3	108	
6	56.7-11.3j	-18.2	1.28	139	
7	53.4-9.1	-20.6	1.21	174	
8	50.7-7j	-23.2	1.15	193	

Table III-3: Bench measurements for each array element in the loaded case.
All elements are well-tuned as indicated by the S_{11} and SWR measurements. Q varies asymmetrically around the coil due to the asymmetric loading properties of the head phantom as indicated by the layout on the right.

After verifying array tuning, the elements were tested at power using a conventional 50Ω amplifier (AMT) and a Bird Watt-meter as shown in Figure III.11. The ultra-low output impedance amplifier was not used in power testing to protect the

amplifiers against high reflected power in the event that an element failed in testing. The

input attenuation to the AMT amplifier was adjusted until the on-coil power (forward -

reflected) was measured by the Bird Watt-meter to be 350 Watts.
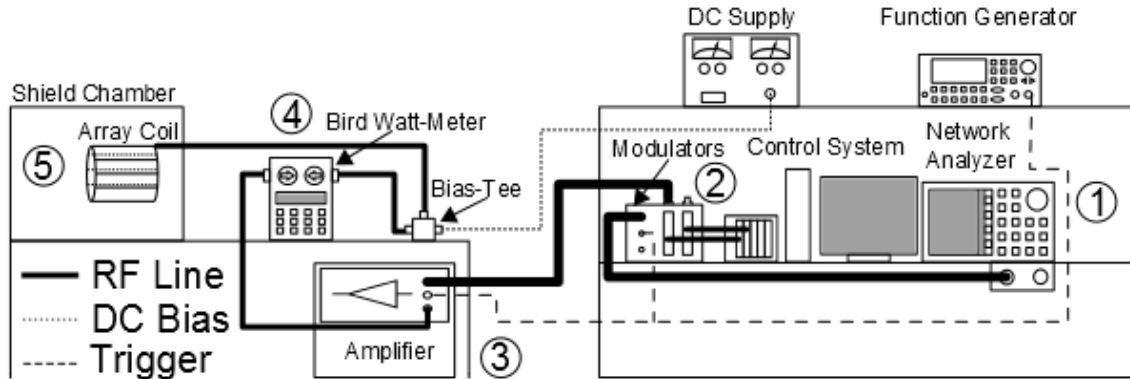
Figure III.11: High power setup for testing the array elements at 300 Watts.
The network analyzer (1) sourced continuous wave RF at the appropriate power level.
The control system delivered the baseband signals for modulation (2) and the output
waveform is input into a commercial kilowatt amplifier (3). The attenuation on the input
line to the amplifier was adjusted to obtain 350 Watts on the coil using the Bird Watt-
meter (4). The waveform was monitored on the coil (5) using the built-in probes.

The on-coil probe was monitored using a scope and the tuning of the element

was measured before and after applying power. If a capacitor failed, the on-coil

waveform would exhibit instabilities. The onset of a capacitor failure was detected by

observing any shift in resonance measured after applying power. The elements heat

during power testing, and heat causes the value of the distributed capacitance in the rung

to increase. An increase in the distributed capacitance causes the rung impedance to shift

to the right on the Smith chart indicating an increase in inductance. This shift to the right

requires less capacitance to resonate the element, and the resonant frequency of the

45

element increases. After some time, the element cools to room temperature and the resonant frequency shifts back to its initial value. The slight shift due to heating was expected. If an impending capacitor failure was detected, the behavior changed such that the distributed capacitance was reduced causing the resonant frequency of the element to decrease. These two tests detected any problems with the elements that were not detected at low power and ensured the elements were stable and ready for deployment.

III.3.2 Coupling Matrices

The 8x8 decoupling matrices acquired with the rungs terminated in 50Ω and in the ULOI amplifiers are shown in Figure III.12. When the elements were terminated in 50Ω, the single capacitor decoupling network on the coil provided an average of -11 dB (-17.4 dB max, -6.9 dB min) of decoupling between nearest neighboring elements, as seen in Figure III.12a. The decoupling matrix acquired with the rungs terminated in the ULOI amplifiers is shown in Figure III.12b, indicating the added isolation provided by the amplifiers. The total average decoupling between channels with the amplifiers in line was -23.2 dB (-30.8 dB max, -13.9 min). Further iterative tuning of the amplifiers and/or of the decoupling between elements can certainly improve this, but sufficient independence between channels was considered to have been achieved based on observing a minimal shift in element tuning (within an SWR of 2:1) for all elements as compared to the original tuning that was done with all other elements open-circuited.

In a few cases, the coupling between next-nearest neighbors is higher than the coupling between nearest neighbors. This is due to the addition of a current path between these elements via the decoupling capacitor that was tuned for nearest-neighbor

decoupling only. There are minor differences in the reciprocal measurements of the coupling matrices (i.e. $S_{ij} \neq S_{ji}$) due to the fact that all 64 measurements were acquired individually with the probes and therefore were sensitive to small differences in calibration.
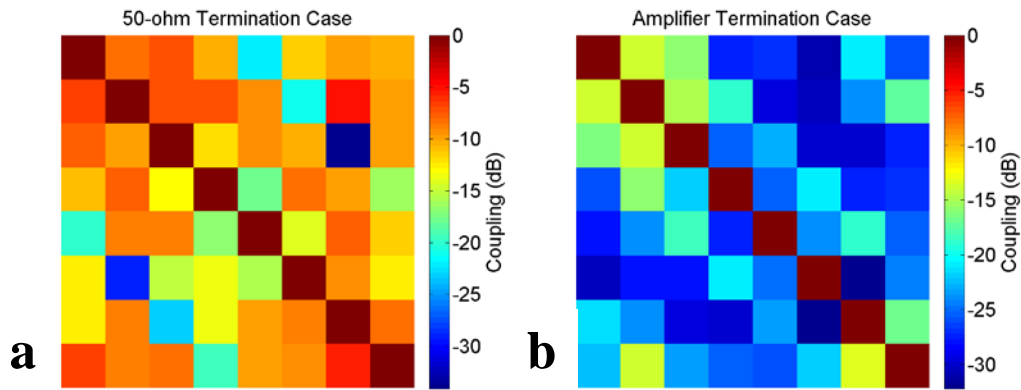
Figure III.12: Coupling matrices acquired using the on-coil current probes for the 50Ω and amplifier termination cases.
(a) The coupling matrix acquired with the coil elements terminated in 50Ω loads. The decoupling network adds an average of -11 dB of decoupling between nearest neighbors to augment the isolation provided by the ULOI amplifiers. (b) The coupling matrix acquired with the coil terminated in the ultra-low output impedance amplifiers. With the amplifiers in line, overall channel-to-channel coupling is an average of -23.6 dB. Minor differences in the reciprocal measurement in both matrices (i.e $S_{ij} \neq S_{ji}$) are due to the fact that all 64 measurements were acquired individually and were therefore subject to slight differences in probe sensitivities.

III.3.3 Scanner Setup and Transmit Field Patterns

All transmit array imaging experiments were performed on a 3 Tesla General Electric (GE) clinical research scanner at the University of Michigan. Two rooms housed the parallel transmitter: the equipment room and the scan room. The equipment room is

adjacent to the scan room, and connections passed through a filtered feed through panel. The control system, modulators, and driver amplifiers were in the equipment room, and the power amplifiers and array coil were in the scan room. For experiments, the array coil was positioned on the patient table and landmarked to center the array coil in the bore. The three cable bundles connected to the coil were accessed at the back of the magnet. The transmit cables connected to the output of the power amplifiers, and the current probe cables were connected to a scope to monitor the on-coil waveforms during the scan because the monitoring hardware was not yet setup. The receive bundle connected to the GE receiver via a custom connector obtained from GE. A diagram of the setup and connections are shown in Figure III.13.

The B_1 maps acquired from each element in the flood region of the ACR phantom are shown in Figure III.14. The patterns demonstrate an effective transmit system, generating well-isolated sensitivities and corroborating the expected behavior based on the coupling matrices.

III.4 Conclusions

In conclusion, we have discussed the design, construction, and characterization of the first eight-channel array coil for use with ultra-low impedance amplifiers. In particular, considerations when adding decoupling to the array coil with respect to appropriate measurements and placement of the decoupling capacitor were detailed. The head array working with ultra-low output impedance amplifiers provided effectively isolated channels and avoided the need for complex on-coil decoupling, which is typically load dependent. The approach made a transmit/receive array configuration

straightforward and demonstrates an option to consider for parallel transmission applications.
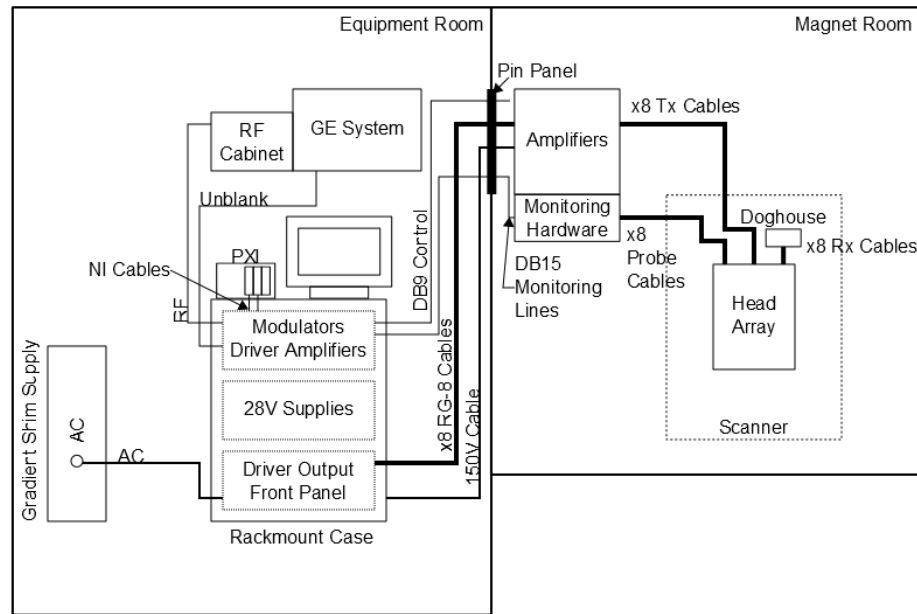
Figure III.13: Diagram of the parallel transmit system setup in the equipment room and scan room.
The control system, modulators, and driver amplifiers were housed in the equipment room. Connections passed through a filtered feed through panel between the two rooms. The power amplifiers, monitoring hardware (not yet setup), and head array were in the magnet room.
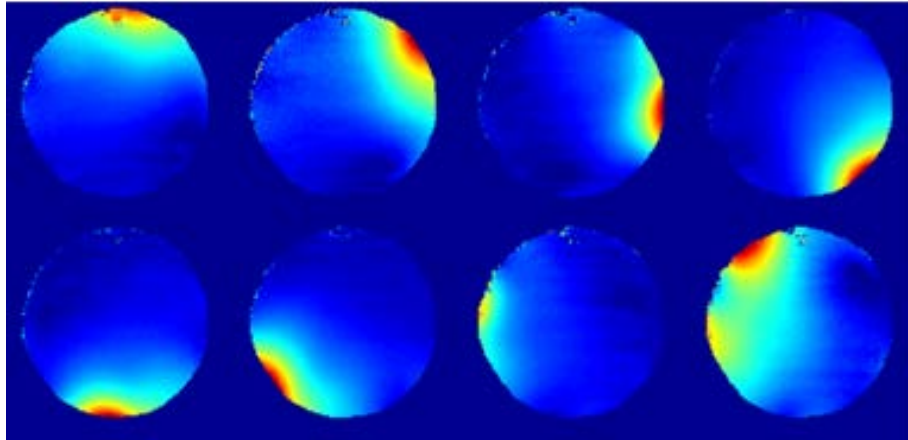
Figure III.14: B1 maps for the eight-channel array using ULOI amplifiers.
B_1 maps in the flood region of the ACR phantom using the Bloch-Siegert method acquired using the eight-channel array and ULOI amplifiers interfaced to a GE 3 Tesla clinical research scanner. The maps are scaled individually to better show the coil patterns.

CHAPTER IV [*]

CONTROL SUBSYSTEM TO DRIVE THE PARALLEL TRANSMIT HARDWARE

IV.1 Introduction

Current scanner hardware and software is proprietary and typically closed, making it difficult to adapt the existing technology for parallel transmit capability. The most straightforward approach to adding parallel transmit technology requires a solution that retrofits the existing scanner hardware. A separate control system dedicated to the parallel transmit system is needed to interface with the transmitter subsystems and manage waveform playback.

Previous work in developing custom control consoles has focused on receiver replacement [54-57], spectroscopy [58-59], and MR systems [60-62]. The MEDUSA console [63-64] has the capability of performing parallel transmit, but it is currently limited to an upper operational frequency of 100 MHz, making it unusable in high field imaging where the B_1 inhomogeneity and magnetic susceptibility problems are more pronounced and parallel transmit becomes necessary. The modularity of the MEDUSA console makes a transmit-only implementation possible, but the direct digital synthesizer (DDS) approach to generating the RF would complicate matching the transmitter frequency to the receiver frequency and make a retrofitting option difficult to implement.

We have developed a parallel transmit system using a vector modulation approach, requiring only two physical connections to the existing scanner hardware: the RF (set to a hard pulse) and the blanking digital control line.

Parallel transmit solutions with the capability to easily retrofit existing scanner hardware are needed to accelerate the translation of the technology from experimental research to routine clinical use. Since the MRI console software is closed, a separate parallel transmit control system capable of managing pulse playback and driving the transmitter subsystems is required. This chapter discusses the development of the control system for the vector modulation-based parallel transmit system. The chapter progresses as follows:

- Parallel transmit system overview

- Control system requirements

- Specifying the hardware for the control platform

- Software architecture

- Host-to-target communication options

- Programming implementation

- Developing the user interface

- Functionality testing

- Deployment and imaging experiments

IV.2 Materials and Methods

IV.2.1 Parallel Transmit System Overview

The eight-channel parallel transmit system was designed as a retrofit for a 3 Tesla GE clinical research scanner, requiring a rapid and transparent switchover from a standard single channel transmitter. Two inputs were required from the host GE scanner: the input to the RF amplifier and the master exciter unblank (RF gate) signal. A single hard pulse played out from the scanner is divided eight ways and then modulated by an in-house built vector modulator [44, 65]. The control system employs a PXI-7853R FPGA-based board with programs written in LabVIEW (National Instruments, Austin, TX) to drive the hardware and provide the baseband in-phase and quadrature signals to the vector modulator from the user-defined amplitude and phase information of each RF pulse [66]. The modulated waveforms pass through a first gain stage prior to the ultra-low output impedance amplifiers. The head array coil (discussed in Chapter III) is connected to the amplifiers through transmit/receive switches. Low input impedance preamplifiers on the array coil provide the first gain stage on the receive side prior to passing through to the scanner receiver chain. A block diagram of the parallel transmit system is shown in Figure IV.1.
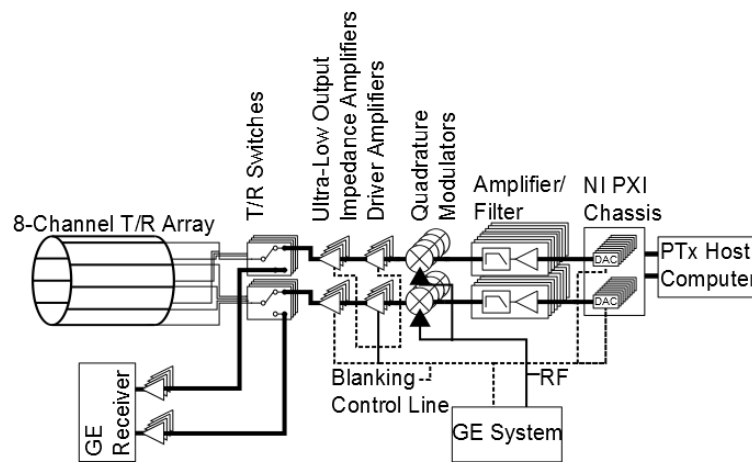
Figure IV.1: Block diagram of the parallel transmit system showing connections to the GE scanner.
Only two connections are required from the existing scanner hardware: the RF hard pulse from the output of the modulator and the master exciter unblank (trigger control line).

IV.2.2 Control System Requirements

The primary function of the control system was to provide the baseband in-phase and quadrature signals to the vector modulator. The vector modulator used these signals to shape and phase the RF hard pulse, and 16-bit resolution was necessary to preserve the pulse shape. Timing for pulse playback was critical as any instability or jitter in the system would cause inaccuracies in the output waveform leading to incorrect excitations in the MR experiment. Therefore, implementing the waveform playback on hardware independent from the operating system (such as a microcontroller or field-programmable gate array) referred to as a "target" device or streaming from a real-time operating system were pursued as viable solutions (discussed in the following sections).

54

The update rate was determined by examining current fast imaging techniques and their associated requirements. The system was designed for transmit SENSE applications. The first implementation of transmit-SENSE reported pulse dwell times of 25.6 microseconds but noted that this was considered long [16]. The price of hardware increases rapidly as the sampling rate increases, so a conservative requirement was set to allow for update rates of 1 MS/s corresponding to a 1 microsecond dwell time. The MEDUSA console operates with update rates up to 250 kS/s (4 microsecond dwell time) which is sufficient for modern fast imaging techniques [64].

The hardware for the digital-to-analog conversion (DAC) must simultaneously play out all the waveforms (16 waveforms total for eight channels). Some DAC hardware architectures employ multiplexing operations rather than using a dedicated DAC per channel. The multiplexing boards can create relative phase delays between channels resulting in inaccurate multi-channel excitations and were considered unsuitable.

Two implementation methods were considered for waveform memory management: storing the waveforms on the computer and streaming them to the target hardware between triggers, or storing them on the target hardware itself. The streaming method requires the waveforms be loaded in a buffer and subsequently read from the buffer on the target hardware. Interrupts between the host and target could facilitate handshaking, but minimum repetition times would be imposed due to bus rate limitations in the data transfer between the host and target. Storing the waveforms in the onboard memory of the target hardware is an elegant solution requiring only one initial transfer

of the data to the target before the start of an experiment. This approach eliminates the limit on the repetition time but requires a device with sufficient onboard memory.

In addition to waveform playback, the control system would perform functions requiring digital control lines: setting the attenuators on the vector modulator, adjusting the modulator potentiometers, controlling power sequencing, providing an amplifier gate protection line, and controlling switches on the modulator to blank the output when the amplitude level falls below the leakage level.

The RF transmit chain needed to be calibrated by linearizing the output (discussed in Chapter V below). To digitize the output waveform, a high-speed analog-to-digital converter was required with the capability of sampling a 128 MHz signal (the frequency at 3 Tesla). A hardware platform that could integrate the waveform playback hardware with the high-speed digitizer for linearization was added as an additional requirement.

The control system would also be responsible for detecting gross hardware anomalies or failures requiring analog-to-digital converters to provide feedback to the control system software. The monitoring lines did not have the stringent timing requirements of waveform playback, but they would require simultaneously sampling the waveform during waveform playback. The monitoring implementation details are discussed in detail in Chapter VI.

IV.2.3 Specifying the Control System Hardware Platform

With control system requirements established, multiple options were pursued and a decision was made to use off-the shelf hardware from National Instruments (Austin,

TX) to reduce development time. We purchased two PXI-7853R (8 AO 16-bit at 1 MS/s, 8 AI 16-bit at 750 kS/s, and 96 DIO) boards to execute RF pulse playback, monitoring, and digital control of the hardware. The boards contain a field programmable gate array (FPGA) with 384 MB of onboard memory allowing the RF waveform to be stored prior to the experiment. The FPGA approach ensured simultaneous output on all channels and precise timing such that a real-time operating system was not necessary. A PXI-chassis housed the hardware and a PC controlled the boards via the PXI interface bus. A PXI-5152 (2 AI 1GS/s 8-bit) high-speed digitizer was purchased to perform linearization. All programming was done in LabVIEW (National Instruments, Austin, TX) requiring no knowledge of a hardware description language.

IV.2.4 Software Architecture

The FPGA-based approach required a target program running on the PXI-7853R board controlled by a host program running on the desktop computer. A state machine implementation on the host and a case structure on the target facilitated task switching, allowing the host program to dictate which state executed on the FPGA. An event case running on the host program switches states based on user interaction with the front panel. The states with possible transitions are shown in Figure IV.2.
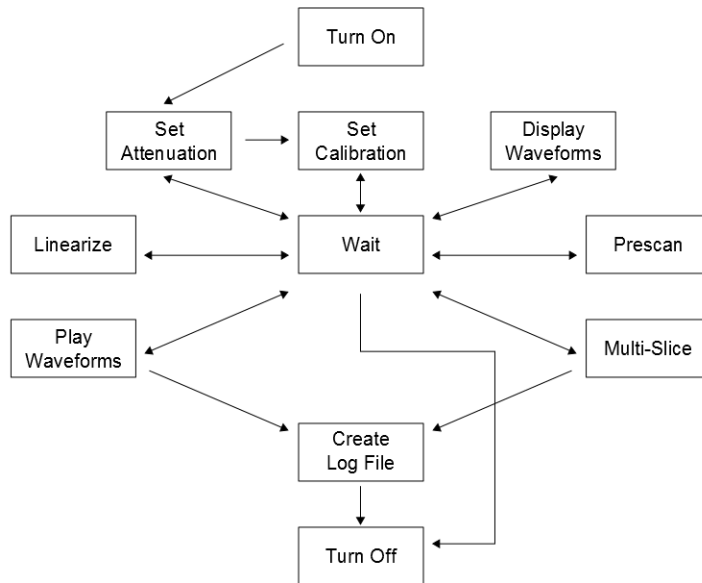
Figure IV.2: State diagram showing the possible states and transitions.

The "Turn-On" state initializes the hardware by setting potentiometers on the modulator board to a known state and switching relays to properly sequence on the transmitter power supplies. After the initial configuration, the control system transitions to the "Set Attenuation" state and checks to see if the user has a control to load the attenuation settings from the previous run. If the control is true, the program executes code to set the attenuation and potentiometer settings. The program transitions to the "Set Calibration" state and checks to see if the user has a control selected to load the previous calibration settings, and if so, loads the calibration constants from a file. The attenuation and calibration must be set before the user is permitted to play the waveforms. The "Wait" state contains the event case waits until the user interactaction triggers an event and transitions states appropriately. The "Display Waveforms" state

58

displays the magnitude and phase data from the user-defined waveform file. The "Linearize" state acquires the on-coil waveforms for the specified channel for a series of pulses and creates a look-up table that is subsequently applied to pre-distort the demand pulses and improve amplitude and phase linearity. The pre-distortion is not automatically applied to the pulses during the "Play Waveforms" and "Multi-Slice" states and requires the user to set a control to activate use of the pre-distortion. The "Prescan" state is required each day and plays out known waveforms and checks to ensure the acquired waveforms match the accepted waveforms acquired at install. The "Prescan" state detects any inconsistencies or gross errors in hardware performance over time. The "Play Waveforms" state reads the information from the waveform file, converts the data to in-phase and quadrature signals, loads the waveforms to the target FPGA, and plays out the waveforms upon receipt of a rising edge trigger. "Multi-Slice" duplicates the "Play Waveforms" state but plays out a different pulse on each trigger based on the information in the waveform file (the differences in file formats for single slice and multi-slice implementations are discussed below). If waveforms are played out, the state transitions to the "Create Log File" state when the user quits the program, which stores information about the experiment in a log file. The experiment ends when the control system is terminated after execution of the "Turn Off" state which resets the target FPGA and turns off the transmitter power supplies in reverse order from the "Turn On" state.

IV.2.5 Host to Target Communication Options

There are multiple methods for communicating between the host and target program, and the choice of implementation depends on the data being transferred and the timing requirements for the specific function. The simplest means of communication is a write/read operation: data is written to a control on the target via a "FPGA I/O" node prior to running the target program via a "FPGA Invoke Node". When the target program runs, the controls initialize with the values defined by the write operation on the host program. The target executes the function and stops. The host program has the option to either wait for the target program to complete execution or continue to implement other operations while the target program runs. Both mechanisms were implemented in the control system and are discussed in more depth in the sections below.

Once the target program is running, communication between host and target programs were facilitated by interrupts or case structures. An interrupt on the target program alerts the host of the target program status, and the host waits for an interrupt. When the host receives the interrupt, additional read/write operations may take place prior to releasing the interrupt and allowing the target program to continue running. Interrupts prevent race conditions, and in the event the host requires information from the target program, the interrupt ensures the data was written to the indicator before allowing the host to read the indicator value.

Transferring arrays of data between the host and the target were handled using either direct memory access buffers (DMA) or read/write operations to onboard memory

blocks. The number of configurable DMA buffers was limited to three per target. DMA buffers were configured as either host-to-target or target-to-host for a specific data representation. For the host-to-target buffer operation, the host configured the DMA by specifying the depth of the buffer and wrote the array of data to the buffer. After the write operation, the target program executed and read the data from the buffer. The target-to-host buffers work in a reverse manner though the host still configures the buffer depth and executes the call to run the target program. Interrupts on the target alert the host when data is loaded into the buffer and ready for transfer, and the host executes the read from buffer operation to extract the data. After data is read from the buffer, the process repeats if necessary to pass additional data. The DMA option is convenient for transferring different data sets between the host and target but is limited by the PXI bus speeds, requiring a finite amount of time for the transfer.

The other option for transferring arrays of data is to write to the onboard memory of the target. Memory blocks are configured on the target with the size and data representation specified. The host writes a data point, an address, and a memory block number to controls on the target program prior to running the target. The target executes and reads the data point from the control and writes it to the specified memory block at the address determined by the control and stops running. The write/read operations between the host and target are repeated until all data has been written to the onboard memory of the target. Once data has been stored in the onboard memory of the target, it can be repeatedly accessed without communication to the host. The amount of data stored is limited to the available onboard memory of 384 kB.

IV.2.6 Programming Implementation

IV.2.6.1 LabVIEW Terminology

In LabVIEW, the graphical user interface (GUI) is termed the front panel, and the backend code is contained in the block diagram. Inputs on the front panel are referred to as controls whereas outputs are called indicators. Each program written in LabVIEW is called a virtual instrument (VI), and subroutines are termed subVIs.

The control system requires communication between the host PC and the target PXI-7853R board in the PXI chassis. A VI runs on both the PC and the target, and to distinguish between the two programs, the host VI is referred to as the host program and the target VI is referred to as the target program.

IV.2.6.2 Attenuation to Adjust Gain Imbalances After Modulation

The attenuators (PE4302, Peregrine Semiconductor, San Diego, CA) on the modulators were used to adjust gain imbalances between the channels after modulation. The attenuators were adjusted using a serial protocol and were capable of setting attenuation values in 0.5 dB increments from 0–31.5 dB. The digital lines used to set the attenuation were on connector 0 of the PXI-7853R board and correlated with the control system as indicated in Table IV-1.

Control System Name	Connector 0 Name
Serial Data 1	DIO2
Serial Data 2	DIO3
Serial Data 3	DIO4
Serial Data 4	DIO5
Serial CLK	DIO6
Serial LE	DIO7

Table IV-1: Table correlating the attenuation controls in system with the output connector.
Both FPGA programs contain the same names, and the names on the second FPGA correspond to the attenuation data for channels 5-8.

Controls on the user interface contained the attenuation values input by the user. These controls were configured to force the value to an appropriate number from 0-31.5 dB in 0.5 dB increments. In the host program, the value was converted to an unsigned 8-bit integer and transferred to the FPGA along with a "Set Attenuation" state command. The target program converted the integer to a Boolean array, and a for loop was used to index the array and sequentially pass the Boolean value to the corresponding data output from the most significant bit (MSB) to the least significant bit (LSB) while transitioning a clock signal as shown in Figure IV.3.

After clocking the data, the latch enable was set high to set the attenuation. The target program alerted the host program that the attenuation state had finished executing, and the target program was set to idle until the next state was executed.
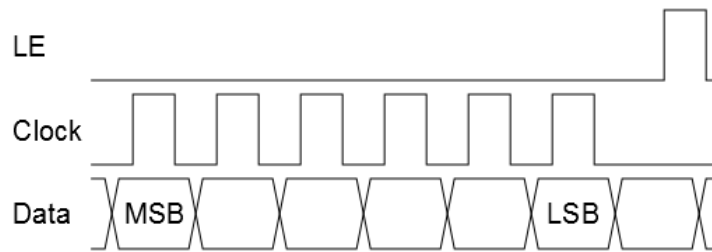
Figure IV.3: Serial timing diagram for setting the attenuation.
A Boolean array on the FPGA contained the appropriate state for each bit and was passed to the data output sequentially. Once all bits were clocked, the latch enable (LE) was set high to set the attenuation. All four channels were set simultaneously using different data outputs as indicated in Table IV-1.

On the front panel, the user is given the option to load the attenuation values from the previous run or set new values. In order to set the values, the user activates a control button. On the backend, the attenuation values are read from the controls, the data is pushed to the FPGA and the attenuators are set, and the values are saved to a file. The host program saves the file and the path to the file location. If the user opts to load the values from a previous run, the last saved file will be read and the attenuation controls are updated with the information from the file.

IV.2.6.3 Potentiometers to Adjust I and Q Imbalances

The potentiometers (MCP4651-103E/ST, Microchip, Chandler, AZ) adjust the common mode voltage point for the in-phase and quadrature signals specified to be 0.5 volts. The potentiometer wiper settings are controlled via the I^2C interface requiring two digital lines: clock and data. By adjusting the wiper setting the resistance of the potentiometer changes and varies the common mode voltage point setting. The commands implemented were the read, write wiper commands and the increment,

64

decrement wiper commands. To execute commands, the master referred to as the "target program" on the control system communicates with the slave containing two potentiometers termed "device".

All I^2C commands are executed across the I^2C bus in the following order: start bit, control byte, acknowledge or not acknowledge, command, acknowledge or not acknowledge, data bits, acknowledge or not acknowledge, and a stop bit. The read command requires an additional control byte before receipt of the data bits (discussed in detail below). All command sequences require interaction between the master target program and the slave device, and both the master and slave control the digital line at different time points. The start bit is defined as the data line being driven low while the clock line is high and initiates a command sequence. The acknowledge bit occurs when the device or target program drives the data line low while the clock is high. If this does not occur and the data line remains high, a not acknowledge bit occurs. The stop bit is defined as the data line being driven high while the clock line is high. Data is transferred on the I^2C bus using the start bit to initialize a transfer and subsequently passing data by changing the state of the data line (high or low) and holding the value while the clock transitions.

When the user runs the control system, the "Turn On" state executes and the potentiometers are set to present maximum resistance (lowest voltage) using the write command. This initializes the wipers and ensures the system starts up in a known state. The I^2C code to set the potentiometers is contained within the "Set Attenuation" state on both host and target programs, and in order to activate the target state, the host program

65

writes "Set Attenuation" to the target state control. Within the target "Set Attenuation" case are two parallel routines: one that sets the attenuation (described in the previous section) and one that sets the potentiometers. Boolean controls on the target program dictate which of the two routines are executed, and the host program writes the appropriate value to these controls before calling the target program to run. In the initialization state ("Turn On"), the potentiometer routine is executed but not the attenuation code. Once the controls on the target program are configured, the target program is called and begins running. The I^2C port is configured in the target program by setting the clock cycle to 100 kHz and setting appropriate downstream controls based on the clock timing. The host program waits for the target to acknowledge the configuration parameters have been set. Each PXI-7853R device controls four devices with two potentiometers in each device to correspond with the in-phase and quadrature wiper settings. A loop executes the write command eight times with different addresses to set all potentiometer wipers. The device addresses are shown in Table IV-2.

Channel	I^2C Address	I Wiper	Q Wiper
1	00101000	0000	0001
2	00101100	0000	0001
3	00101010	0000	0001
4	00101110	0000	0001

Table IV-2: Addresses for the I^2C interface.
The address contains both the fixed address (0010) and the specific address for the device (last 4 bits). The second byte in the I^2C sequence contains 4 bits to specify the I or Q wipers.

For the write command, the start bit is followed by a control byte that contains the fixed address (0010), and the variable address to specify the device as described in Table IV-2. The device responds with an acknowledge bit which drives the data line low. If a response is expected from the slave, the master control of the digital line is disabled to allow the slave to pull the digital line high or low. The control system responds by sending the command byte followed by the write data bits. The command and data bits are transferred in two bytes with the first byte containing the device memory address (see Table IV-2), followed by the write command (00), a do not care bit (0/1), and the first write data bit (D8). The slave responds after the first byte is sent with an acknowledge bit, and the control system sends the final byte of information containing D7-D0 write bits. The wiper values range from 0-255, so to write a value of 255 in the initialization stage, the data bits are all ones (111111111). The details of the write command are shown in Figure IV.4.

Once the wipers have been initialized in the "Turn On" state, they are not accessed again until the "Set Attenuation" state is executed on the host program. In the "Set Attenuation" state, the default behavior is to load the previous potentiometer wiper settings, which reads the potentiometer wiper settings from a file and converts the value to the appropriate control byte, command byte, and write data bits based on the channel number, I or Q wiper, and wiper value. The control system executes the write command as described above and shown in Figure IV.4.

The user has the option to adjust the potentiometer settings and launch a menu that executes commands to write, increment, and decrement the wiper setting based on

67

user interaction. An event case monitors the front panel and determines if the event to

launch the menu has occurred, and if so, the "Set Attenuation" state is executed and case

structures navigate the state code to open the potentiometer menu (shown in Figure IV.5)

and not reset the attenuators.
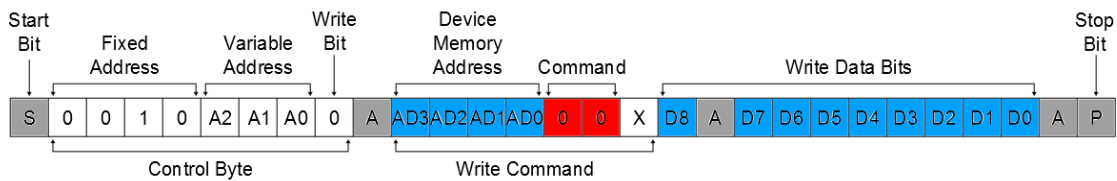
Figure IV.4: I²C write command sequence.
The command begins with a start bit followed by the control byte that specifies the
potentiometer. The potentiometer responds with an acknowledge bit, and the control
system sends the remaining two bytes of data containing the command byte to alert the
potentiometer of a write command with one data bit and the second byte containing the
remaining data bits. The potentiometer responds with an acknowledge bit, and if the
control system is done sending commands to the potentiometers, a stop bit is sent to end
the sequence.
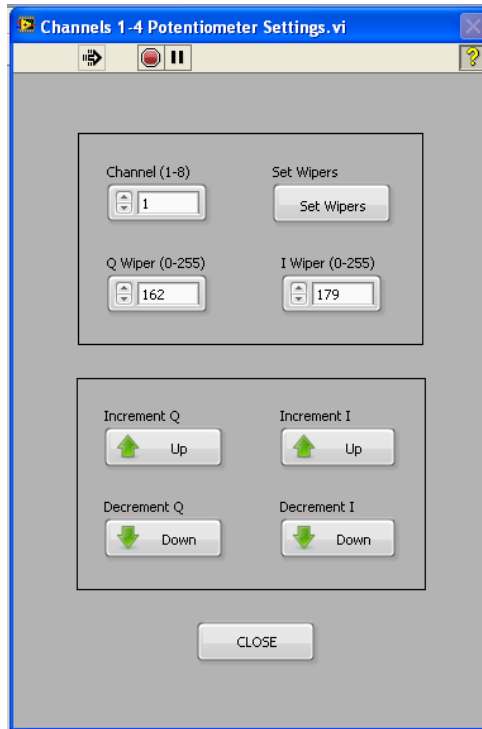
Figure IV.5: Potentiometer settings menu that allows the user to adjust the potentiometer wiper settings.
When the channel is changed, the Q and I wiper controls are updated to reflect the current value of the wiper. The controls can be adjusted, and the values are set by the "Set Wipers" control button. The wipers can additionally be adjusted by the increment and decrement commands which update the I and Q controls automatically. When all adjustments have been made, the menu is closed by the "Close" button and the current values saved to a file.

When the menu opens, the user can adjust the wiper settings for the potentiometers. The menu defaults to display the wiper settings for channel 1 and executes a read command to update the Q and I wiper controls. The read command is initiated with the start bit, followed by the control byte containing the fixed address, device address, and the write bit (0). The target program transfers data to alert the device to the appropriate address for the device/potentiometer combination for the upcoming

69

read operation by sending the read command (11). The device replies with an acknowledge bit, and the target program sends a repeated start bit to signal to the device the start of a new command. The target program sends the control byte with the read bit (1) and waits for device acknowledgement. The device responds by sending one byte containing half the read bits, the target program acknowledges receipt of the read bits, and the device transfers the remaining read bits. The target program responds with a not acknowledge bit to force the device to release the bus. The target program then executes a stop bit to signal the end of the command sequence. The read command sequence is shown in Figure IV.6. The read bytes are converted to a decimal in the host program, and the wiper controls are updated. To obtain the value for both the Q and I wipers, the read command executes twice with each device memory address (0000 for Q and 0001 for I).

Once the wiper controls are initialized, the user can execute the available commands to adjust the wiper setting. The "Set Wipers" command executes the write command previously discussed. The increment and decrement commands are identical to the write command with one exception being the command bits: 01 for increment and 10 for decrement. Once the user has completed the potentiometer adjustment, the menu is closed and the values stored in a text file.
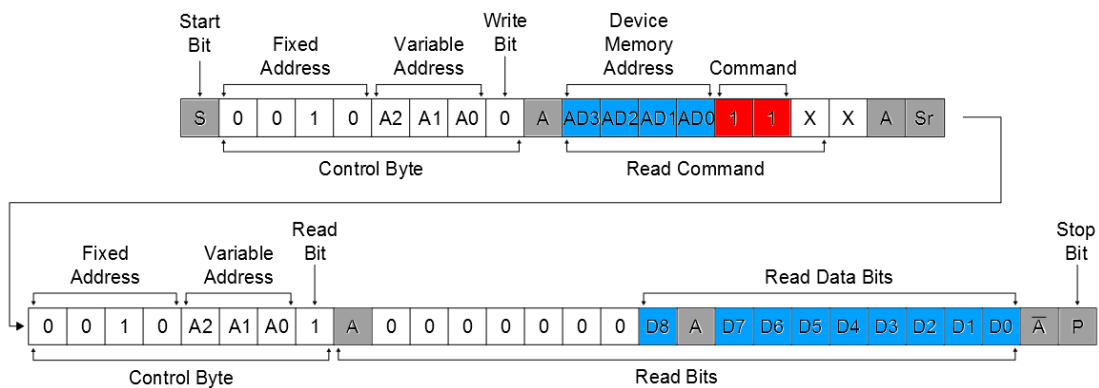
Figure IV.6: I²C read command sequence.
The sequence begins with a start bit followed by the control byte containing the write bit (0). The device acknowledges receipt of the control byte, and the target program sends the read command denoted by the 11 command. The device acknowledges the read command, and the target program sends a repeated start bit to signal the continued transfer of data across the bus. The target program sends the control byte with the read bit (1) and waits for the device to acknowledge receipt. The device responds with the first byte of read bits. The target program responds with an acknowledge bit, and the device sends the remaining read bits. The target program responds with a not acknowledge to signal to the device to release the bus and subsequently sends a stop bit to denote the end of the transfer.

The potentiometers were set when the hardware was installed and have not required a readjustment. The default operation of the control system is for the "Use Last Time's Attenuation" slider to be true. In this case, the control system loads the attenuation and potentiometer settings from a file and performs the necessary operations.

IV.2.6.4 Waveform Playback

Waveform data was read form the user-defined waveform file and converted into in-phase and quadrature signals using the following equations:

$$I(t) = A(t) \cos\big(\varphi(t)\big)$$ [4.1]

$$Q(t) = A(t) \sin\big(\varphi(t)\big)$$ [4.2]

71

Where A is the amplitude and φ is the phase. In addition to the waveform data, the file contains the update rate specified in "ticks" of a 40 MHz clock. The time step is determined from the number of ticks as follows:

$$dt = \frac{1}{40MHz} \cdot Update\ Rate \qquad\qquad [4.3]$$

Where dt is the time step. The format of the waveform file is shown in Figure IV.7. The waveforms play out after the user has pressed the "Play Waveforms" control on the front panel, and the waveforms are loaded to memory.
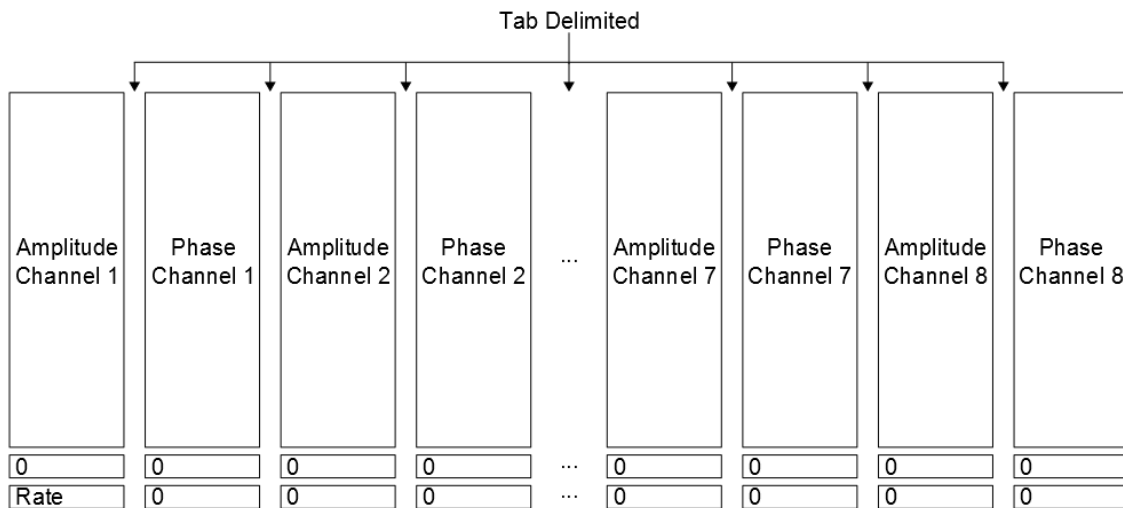
Figure IV.7: File format for the waveform file.
There is no header. Each column contains amplitude or phase data, and all channels are contained within the same file. The columns are delimited by tabs. A row of zeros below the last data point in each set is added to separate the waveform data from the update rate. The update rate is found at the end of the first column and is specified in ticks of a 40 MHz clock.

FPGA targets cannot handle double-precision; floating point numbers and all analog input and outputs are represented using signed 16-bit integers. After computing

72

all I and Q signals in double-precision, the data was converted to the signed interger-16 representation using the following formula:

$$\frac{x \cdot 10}{32767}$$ [4.4]

Where x is the I and Q data in double-precision format. The 10 in equation 4.4 arrives from the output range of the analog outputs of +/- 10 volts, and 32767 is one less than half the value of 2^{16} derived from the resolution of the analog-to-digital converter.

After converting the data to the appropriate datatype, the waveforms were transferred to the target. The direct memory access (DMA) option was initially implemented for waveform playback. The I and Q waveforms for four channels were packed into a single host-to-target buffer along with the size of the waveform and analog output channel. The format of the array consisted of the analog output channel, the size of the waveform, and the waveform data, repeated for all analog output channels. After the data is packed into the DMA buffer, an interrupt is thrown to prevent the host from overwriting the data in the buffer and to alert the target. The target program reads the data from the buffer one element at a time and writes the data to preconfigured local target-scooped buffers associated with the appropriate analog output channel. When the number of elements equals the size of the waveform, the target programmed transitioned to write data to the next local buffer until eight local buffers contained data: four channels each with I and Q signals. After sorting the data, the target acknowledged the interrupt to signal to the host that the DMA buffer was empty. The target then waited for a rising edge trigger, checking at each 40 MHz clock cycle. Upon receipt of a trigger, the data from the local buffers were read simultaneously point-by-point until all data was

73

output. The target waits for the host to throw the interrupt to signal the completion of the write operation to the DMA buffer from the host. When the interrupt is thrown, the process repeats, and the target reads the data from the buffer and writes it to the local buffers. The handshaking process between the host and target repeats until the user stops the program to signal the end of an experiment. The initial implementation with DMA was susceptible to finite time constraints due to the PXI bus transfer speeds, and it was found that a conservative estimate for the minimum time between triggers was:

$$TR_{Min} = \frac{1}{S_R}(N_P) + 0.05 \ s \qquad\qquad [3.5]$$

Where TR_{Min} is the minimum repetition time, S_R is the sample rate, and N_P is the number of points. In the MR experiment, 50 milliseconds is a significant limitation when attempting to accelerate the imaging experiment using parallel transmit, and ultimately, the DMA method for waveform playback was considered too slow and an alternative waveform playback mechanism was pursued. PXI express links have become more commonplace in the last few years and significantly increase bus speeds and may prove a more suitable option for future implementations of a DMA streaming approach.

Storing to the onboard memory of the target eliminates the minimum time between triggers by loading the waveforms to the memory prior to the start of an experiment. Memory blocks were configured on the target program for each analog output channel. On the host, the analog output channel, address, and one index of the waveform array were written to controls on the target. The target program was executed after the host write operation. The analog output channel determined the appropriate memory block, and the address specified the location in the memory block to write the

74

waveform data point. After completing the write operation to the onboard memory, the target program stops. The host program waits until the target program is finished running, and the process is repeated until all data points are written. A Boolean indicator on the host front panel alerted the user that the waveforms were loaded and an experiment could be started.

The user configures the scanner to output the hard pulse at the appropriate power level (20 mV peak-to-peak) and trigger. The target program checks for a trigger every cycle of a 40 MHz clock. When a rising edge is detected, the target program transitions to read the waveform data point-by-point from the memory blocks. All eight memory blocks are read in parallel to ensure simultaneous output. The index of a for loop provides the address for the memory block. A waveform data point is read from the memory at the specified location and output through the analog output channel. A wait function forces the target program to wait a specified number of clock cycles, termed "ticks", prior to releasing the for loop to execute the next iteration. The number of ticks is specified in the waveform file and defines the update rate of the waveform.

IV.2.6.5 Implementing Multi-Slice

A multi-slice implementation was required to enable the capability for complex pulse design methods. The previous section described waveform playback for a single slice. To enable multi-slice, two controls were added to the front panel: a Boolean multi-slice input and a control for the number of slices (see Figure IV.9). When the Boolean multi-slice control is set to true, the control system reads in a file and looks for the specified number of slices separated by # symbol and records the length of each slice

(see Figure IV.8). If the specified number of slices is not found, an error is thrown to alert the user. If the number of slices corresponds correctly to the file, the control system proceeds to load the waveforms to the onboard memory. An additional slice length memory block was created to store the size of each slice to be later used to determine the indexing of the waveform memory blocks during playback. There were multiple slice lengths corresponding to the number of slices, and the data was transferred to the target via the same read/write operations used to transfer the waveform data described in the previous section. Once the slice length data was loaded, the waveform data was transferred to the target as described in the previous section, and after completion of the target write operation, the target switched to a waiting state and the host alerted the user that the control system was experiment-ready through a Boolean indicator on the front panel termed "Waveforms Loaded" (see Figure IV.9).
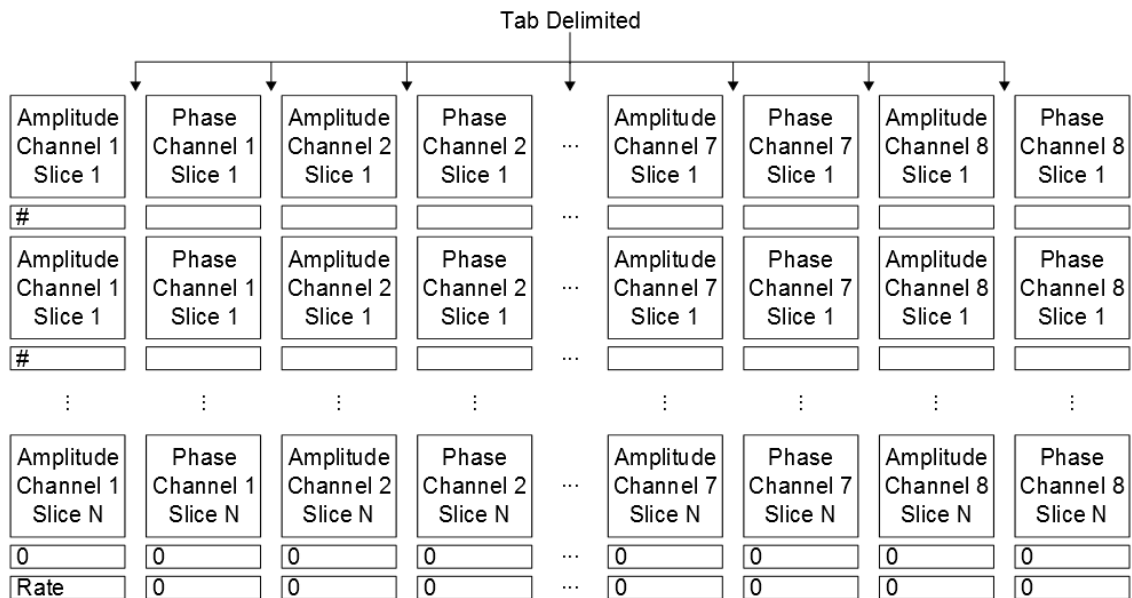
Figure IV.8: Multi-slice file format.
Each slice is separated by a # symbol. The update rate specified in ticks of a 40 MHz clock is located in the first column after the last slice and a line of zeroes.
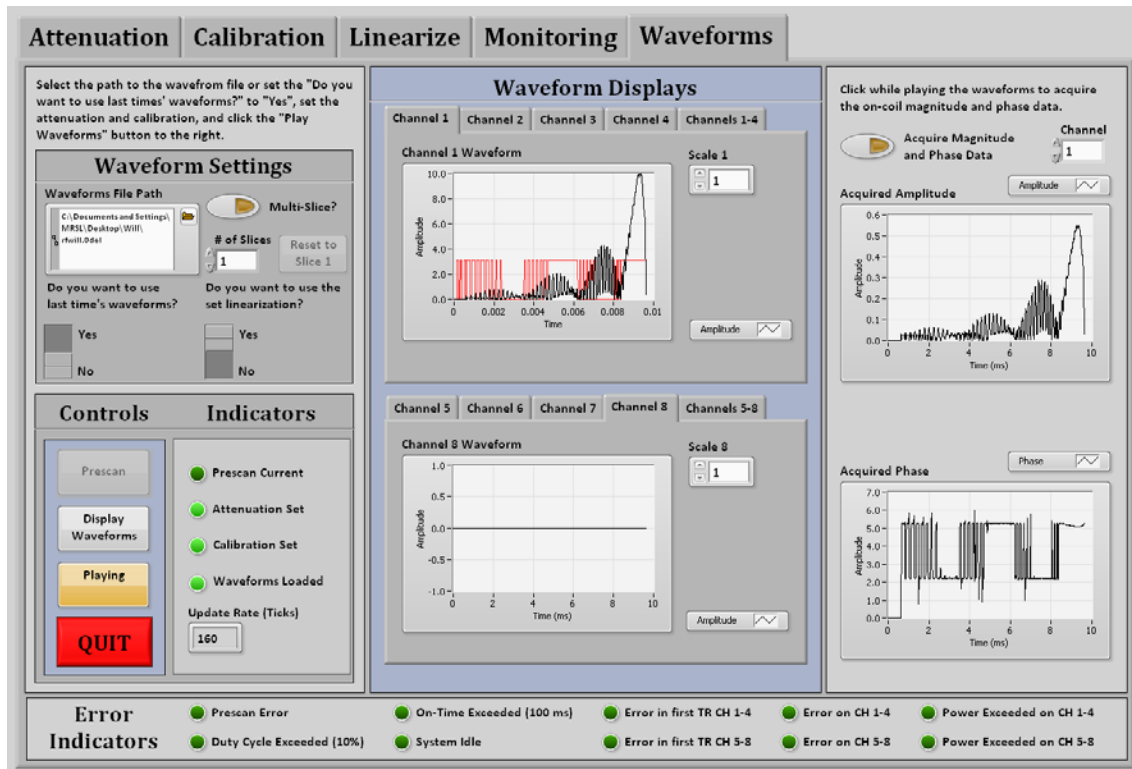
Figure IV.9: Front panel showing the "Waveforms" menu.
The waveform controls are in the top left corner under "Waveform Settings". The LED indicators below the waveform settings contain the "Waveforms Loaded" indicator that alerts the user when the waveforms have been transferred to the target.

The target program waited for a rising edge trigger using a timed single cycle while loop to check for a trigger each cycle of a 40 MHz clock. When a trigger was detected, the target read the first index of the slice length memory block corresponding to the size of slice one. A shift register counter termed "Address for Length" controlled the indexing of the slice length memory block. The target read from the waveform memory blocks until the iteration of the for loop equaled the size of slice one. A shift register stored the value of the length of slice one, and the target transitioned to wait for a second trigger and the process repeats. At playback of slice two, the slice length of

slice two was added to the value in the shift register to determine the final index of slice two, and the stored value in the shift register (slice one) determined the starting index of slice two. This process repeated until the last slice was output, and the shift register and counter reinitialized to zero to restart the process.

During experiments, the scanner was often configured to run a prescan (separate from the prescan the control system performs) prior to an experiment to ensure an accurate center frequency and receiver gain. This process was executed after the waveforms were loaded to the onboard memory and could result in the control system beginning execution at a slice other than slice one. To prevent this error and to continue to allow for the prescan operation, a control was added to the front panel named "Reset to Slice 1". This control became active after the waveforms were playing out, and when pressed, resulted in a return to slice one. The host program wrote to the "Reset to Slice 1" target control while the target VI ran. At each stage of waveform playback, the target program checks the value of the "Reset to Slice 1" control. If the value of the control was true, the shift register and counter were reinitialized to zero. An OR command inside the single cycle time loop that waits for the trigger killed the loop execution in the event the "Reset to Slice 1" control was true and the target program was waiting for a trigger. This enabled the prescan to execute, the scanner to stop sending a trigger, and the slice to be reset without requiring an additional trigger. After the appropriate counters and shift registers were reinitialized, the target program alerted the host program, and the "Reset to Slice 1" control returned to the false state to indicate to the user that the reset had occurred.

IV.2.6.6 Log File Creation

Log files were created if the user played out the waveforms, indicating an experiment occurred. The log files contained all pertinent information about the settings in the control system during the scan: attenuation and potentiometer settings, calibration constants, and waveforms. The format of the log file is noted in Figure IV.10. A waveform log number was recorded along with each file. A maximum of twelve files are kept to minimize memory usage, and a text file containing the log file number is updated at each execution of the "Create Log File" state. When twelve files are created, the next experiment overwrites the oldest log file.

IV.2.6.7 Shutting Down the Transmit System

The "Turn Off" state executes when the "Quit" button is pressed. If the control system is currently executing the "Play Waveforms" state, the target program is aborted and the log file is created before transitioning to the "Turn Off" state. Two primary functions occur during the "Turn Off" state: the power supplies are turned off in reverse order and the onboard memory of the target is cleared. There is an option to reset the target program, but this only reinitializes the controls and indicators to their default values. The onboard memory remains unaffected, and the waveforms data is still stored. The write/read operation described in Section IV.2.5.3 is executed in the "Turn Off" state to write zeroes to onboard memory of the target and reset the memory blocks.
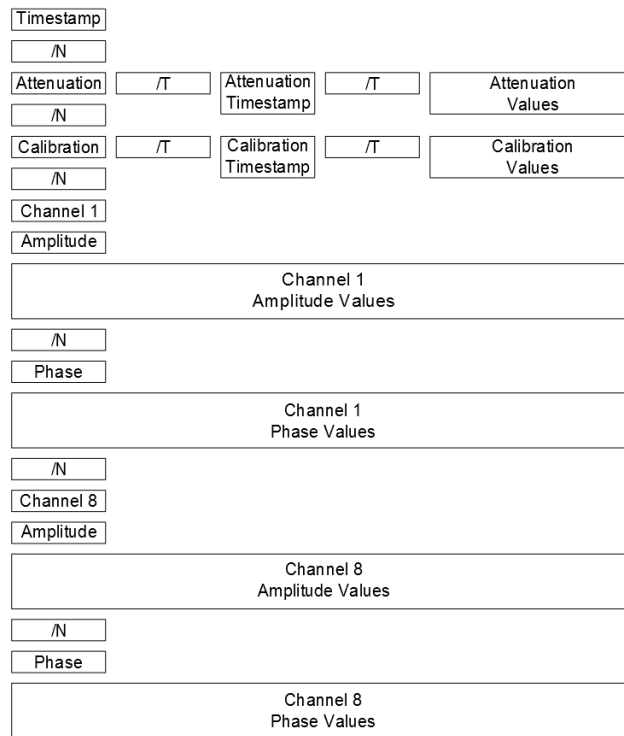
Figure IV.10: Log file format.
"/N" denotes a new line and "/T" indicates a tab. All values (i.e. attenuation, calibration, amplitude, etc.) are tab delimited.

IV.2.7 Developing the User Interface

The user interface was developed in LabVIEW and event handling was done in the "Wait" state on the backend. Each button on the user interface defined an event case that watched for the control to change values from low to high. When the event changed, the event case stopped execution and passed a state variable to transition to the appropriate state. The controls on the "Attenuation" menu that transition the program from the "Wait" state to the "Attenuation" state are shown in Figure IV.11.

81
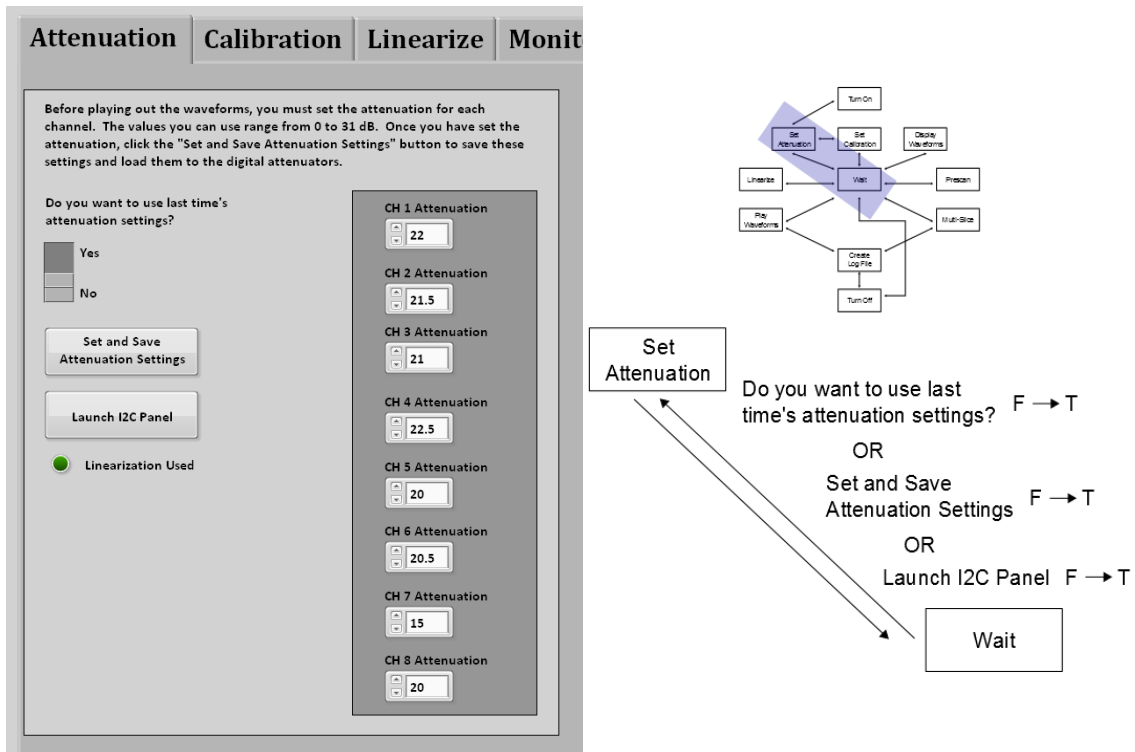
Figure IV.11: Attenuation menu and state diagram showing the controls that transition the host program from the "Wait" state to the "Attenuation" state.
The event case checks to ensure that the control values changed from a low state to a high state. The diagram at the top right is a reduced version of Figure IV.2 to show where the "Attenuation" state falls in the main state machine architecture.

The "Calibration" menu contains two controls that transitions the host program from the "Wait" state to the "Calibration" state shown in Figure IV.12. Both the "Attenuation" and "Calibration" states execute during the initial run of the host program as previously described. After the initialization phase, the "Attenuation" and "Calibration" states execute only if the user interacts with the described controls.
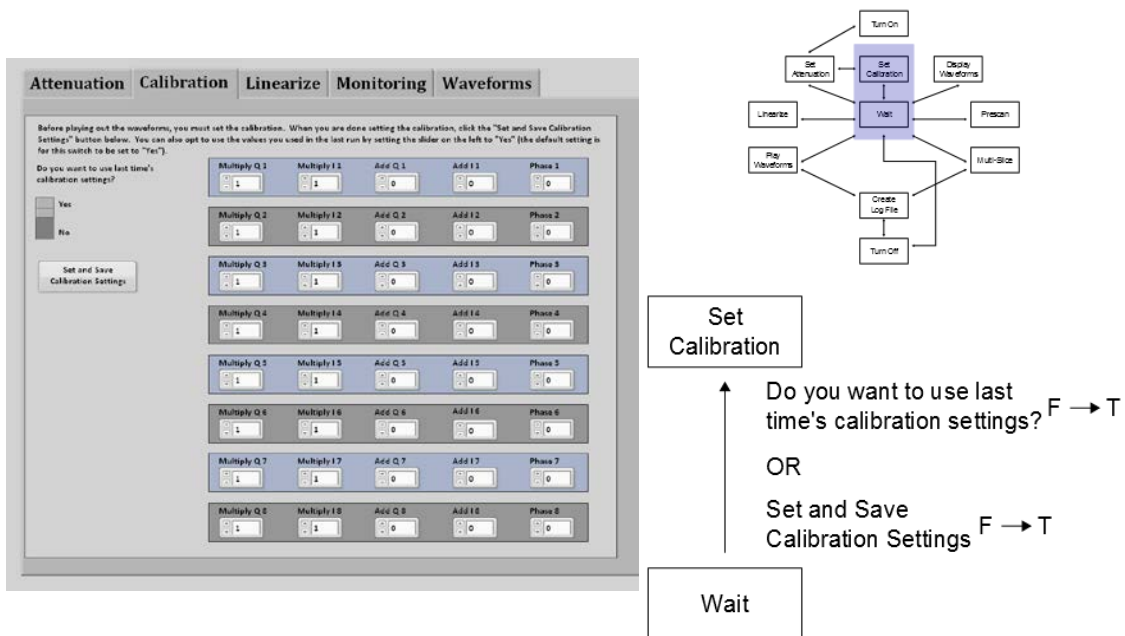
Figure IV.12: Calibration menu and state diagram showing the controls that transition the host program from the "Wait" state to the "Calibration" state. Each control is monitored by an event case in the "Wait" state.

The "Linearize" and "Monitoring" menus will be discussed in parallel to their implementation in Chapters V and VI below. The "Waveforms" menu contains several controls that are available for user interaction with some only available during execution of specific states. The "Reset to Slice 1" control is only available during the "Multi-Slice" state. It does not cause a state transition but simply resets the indexing on the target program to correspond with slice 1.

The "Wait" state monitors front panel activity using an event case and transitions states depending on the event as noted in Figure IV.13.
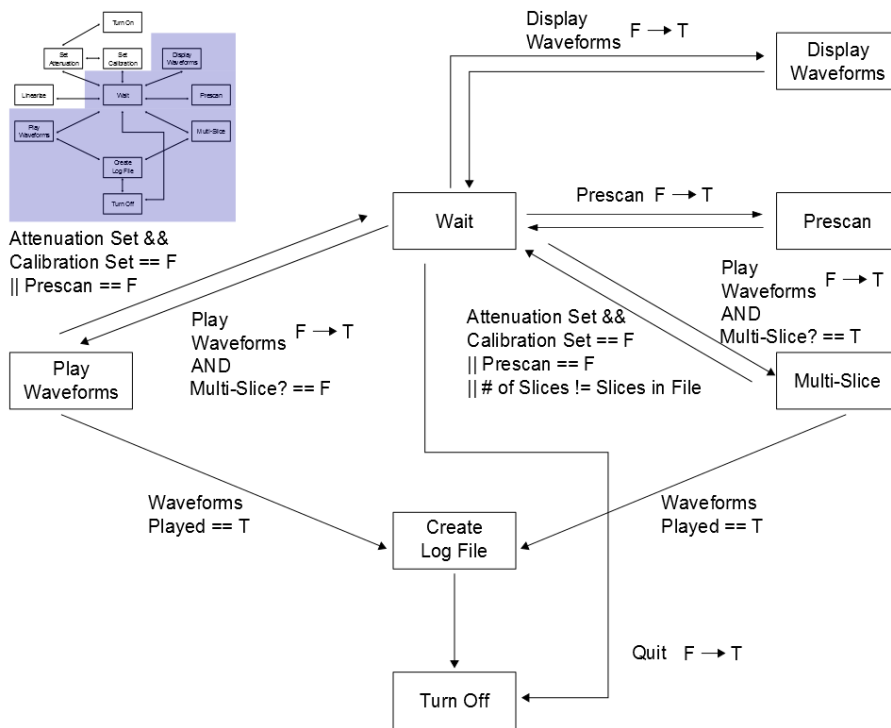
Figure IV.13: Diagram showing the possible state transitions available from the waveforms menu from the "Wait" state.
The waveforms menu was previously shown in Figure IV.9. The && symbol denotes an AND operation, || denotes and OR operation, and != indicates not equal.

IV.2.8 Functionality Testing

IV.2.8.1 Digital Control Testing

The attenuators and potentiometers were tested using a logic analyzer (LA1034, Intronix Test Instruments Inc., Phoenix, AZ) setup with appropriate interpreters. The logic analyzer has built-in interpreters for both the serial and I^{2}C interfaces simplifying testing. For the attenuation, the attenuation was set in the control system and verified to output the correct sequence of bits for the attenuation value. After testing the attenuation with the logic analyzer, an S_{21} measurement was made between the input RF waveform

and the modulator output using the network analyzer while playing out a hard pulse. The attenuation was set to zero and the resulting S_{21} measurement recorded to serve as a reference for remaining measurements. The attenuation was increased and the expected change in the S_{21} measurement was observed.

The I^2C code was tested using the logic analyzer in combination with a test board that contained a device with potentiometers. The master-slave communication between the control system and device necessitated the device be part of the test environment. The logic analyzer was connected and the I^2C interpreter configured. All stages of the code were tested and verified to produce expected results. As an additional test, the voltage level at the output of the potentiometer was measured with a handheld multimeter and verified to increase or decrease as the wiper setting was adjusted.

The digital lines to control the relays responsible for powering on the supplies in a specific order was tested and verified using the logic analyzer. Green LEDs in the supply chassis turn on with the supplies and corroborated the logic analyzer test.

IV.2.8.2 Verifying Waveform Playback

Waveform playback was tested using a high-speed scope (LeCroy 332A, Teledyne LeCroy, Chestnut Ridge, NY). Initial tests involved verifying the correct in-phase and quadrature signals from the target. A terminal block (SCB-68, National Instruments, Austin, TX) was used to probe the individual I and Q lines. A series of pulses with varying phase were played out, and each waveform was verified on the scope to have the correct voltage. The sampling rate was verified by measuring the duration of the pulse and ensuring it matched the specified rate in the waveform file.

After verifying the I and Q channels, the control system was connected to the vector modulator, and the output waveform was observed for a set of pulses and verified to have the correct pulse definition.

The multi-slice functionality was tested using the scope and looking at the output waveform from the modulator. The slice was verified to change at each trigger corresponding to the information in the waveform file. If the "# of Slices" control did not match the detected number of slices in the file, the multi-slice state was aborted and the user alerted to the error. The "Reset to Slice 1" error was verified by clicking the control and observing that the next slice played out corresponded to slice 1. The "Reset to Slice 1" behavior was also tested by turning off the trigger during waveform playback, and then, pressing the "Reset to Slice 1" control. When the trigger was output again, the slice played out corresponded to slice 1. This second test case corresponds most accurately to the scenario needed onsite since a prescan will occur giving a trigger to the control system, and then, the prescan will be stopped turning off the trigger. The control system will not be aborted between the prescan and the start of an experiment, and therefore, it was necessary to force the control system to reset to slice 1 without requiring a secondary trigger.

Imaging experiments were performed to validate the overall functionality of the control system. Single-channel images were acquired outputting a sinc pulse on each channel to verify excitation and synchronization with the scanner. Simultaneous output was demonstrated by generating pulses to produce a uniform excitation based on single-channel images. To demonstrate parallel transmit capability, the control system and

86

modulator subsystem were installed for use with a subset of four channels on the 64-channel transmit system [67].

IV.2.9 Deploying the Control System

After verifying functionality on the bench, the control system was deployed onsite. Identical control hardware was purchased for the onsite system, and the computer used contained a replica of the LabVIEW software installed in the lab. The system was installed and spot tested for base functionality. The vector modulators and remaining pieces of the transmit system were shipped to the site of the scanner and installed. The control system was designed for primary use on a 3 Tesla GE Signa clinical research scanner, but in addition, the control system and modulator subsystem have been deployed for use on a 4.7 Tesla Varian animal research scanner and a 7 Tesla Philips Achieva clinical research scanner.

IV.3 Results and Discussion

Single-channel gradient echo images shown in Figure IV.14 were acquired using the control and modulator subsystem from each channel of a 3 Tesla GE clinical research scanner (TR=250 ms). Commercially available power amplifiers were used in combination with the dual plane pair array discussed in Chapter III to acquire the images. The images demonstrate the ability of the control system to synchronize output with the scanner and generate appropriate excitations.
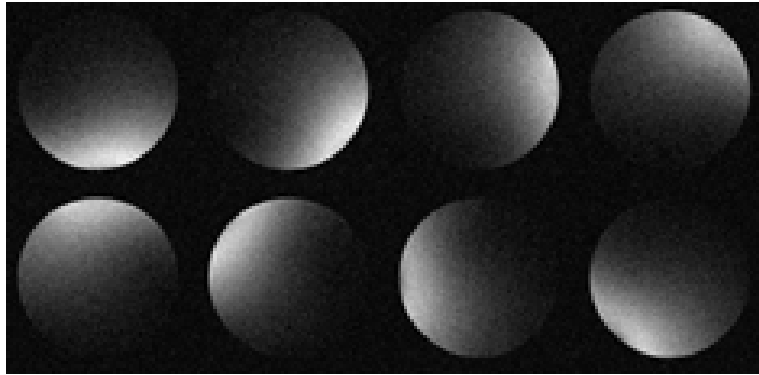
Figure IV.14: Single-channel gradient echo images acquired using the 3 Tesla GE clinical research scanner.
The excitations demonstrate synchronization with scanner hardware and output from the control system.

To demonstrate simultaneous output capability as well as amplitude and phase control, pulses were generated to create a uniform excitation using all eight channels and the image in Figure IV.15 produced. The uniform excitation is an example of B_1 shimming (discussed in Chapter I) where the same shaped pulse is played out on all channels but with relative differences in amplitude and phase. The image shows uniformity across the phantom and the ability of the control system to simultaneously output eight channels of data all synchronized to the scanner.
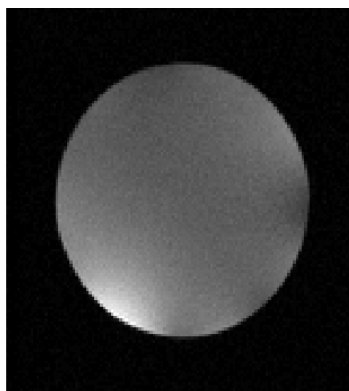
Figure IV.15: Uniform excitation image acquired playing out sinc pulses with various amplitudes and phases on all eight channels.

The control system was developed to perform parallel transmit applications, and to demonstrate parallel transmit capability, the control and modulator subsystem was setup for use with a subset of four channels of the 64-channel transmit system. A four-channel transmit-only decoupled loop coil with active detuning was used in combination with a quadrature receive birdcage coil to generate images. A collaborator generated the pulses used for the parallel transmit experiments. The images were acquired on a 4.7 Tesla Varian animal research scanner. We acquired B_1 maps from the individual coil elements and generated a fly-back echo-planar-imaging (EPI) sequence with an acceleration factor of two to excite a square pattern. A simulator developed by a collaborator was used to predict the excitation pattern given the RF pulse, gradient trajectories, and the B_1 maps for each channel. The predicted results are shown in Figure IV.16, and the measured results are shown in Figure IV.17. The normalized absolute error between the simulated and measured results is shown in Figure IV.18.
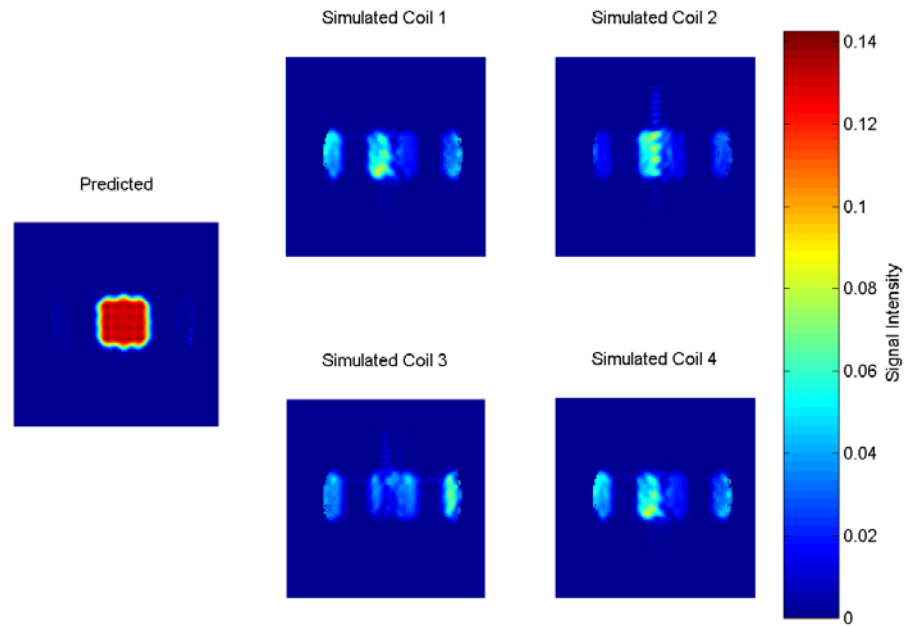
Figure IV.16: Predicted parallel transmit results for a flyback echo planar imaging (EPI) sequence.
The predicted (left) image shows the resulting excitation when transmitting on all four elements. The coil images (right) are the simulated individual excitation images for each channel. The image is aliased in the horizontal direction due to an acceleration factor of two in that direction.
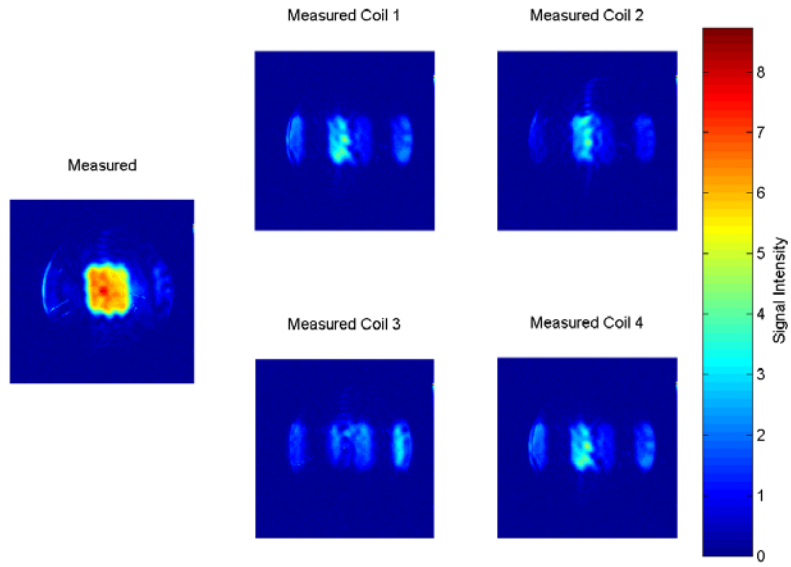
Figure IV.17: Measured parallel transmit results from a flyback echo planar imaging sequence.
The measured results show good agreement with the simulated results.
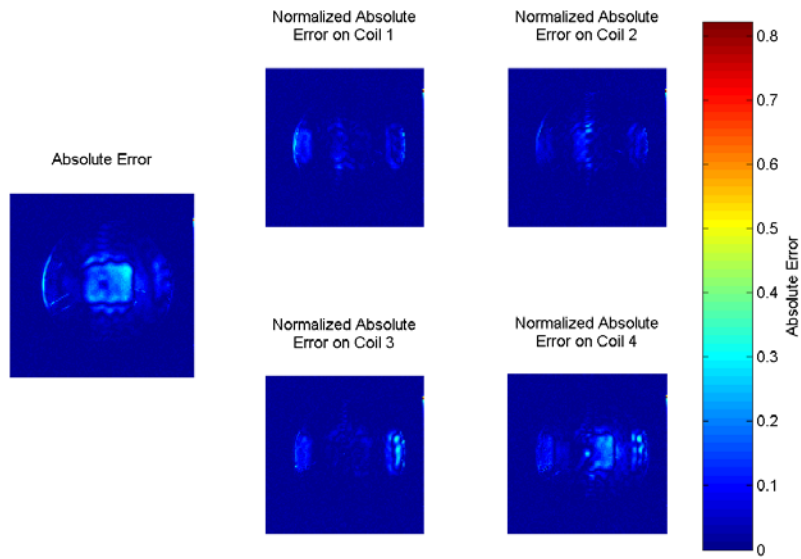
Figure IV.18: Normalized absolute error between the simulated and measured parallel transmit results.

These results indicate good agreement between the simulated and measured results. In the measured results, there is noticeable excitation outside the region of interest which can be partially explained by incomplete eddy current compensation causing inaccuracies in the actual gradient trajectory and remaining inaccuracies in the linearity of the transmit system.

In addition to the parallel transmit results, the calibration panel in combination with the scaling factor in the control system was used to manually adjust the phasing and signal level of the RF pulse on two channels in order to create quadrature (uniform excitation) and anti-quadrature (null in the center of the excitation) images. These images further demonstrate amplitude and phase control and are shown in Figure IV.19.
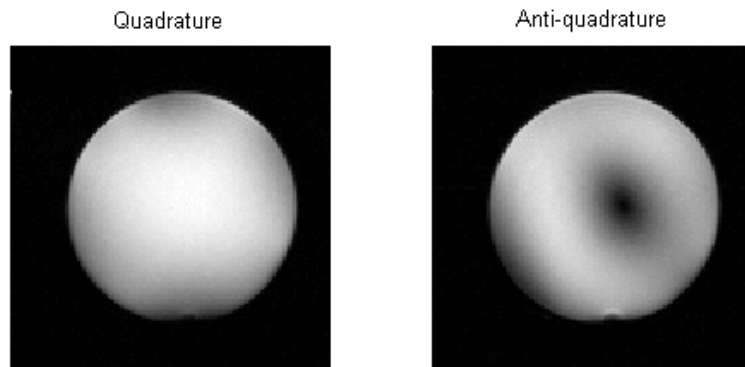
Figure IV.19: Quadrature and anti-quadrature images generated using two channels.
An amplitude and phase adjustment was applied using the scaling factor and calibration menu in the control system. The results further demonstrate amplitude and phase control.

IV.4 Conclusion

The control system effectively drives the transmitter hardware and provides a straightforward interface for the pulse designer. The system was demonstrated to adequately perform pulse playback for parallel transmit experiments, and amplitude and phase control through the system was verified. The control and modulator subsystem is flexible and able to easily retrofit existing scanner hardware. Results were acquired from both a 3 Tesla General Electric Signa clinical research scanner and a 4.7 Tesla Varian animal research scanner. The control and modulator subsystem was also used on a 7 Tesla Philips Achieva clinical research scanner though results were not shown. The variety of scanner hardware used demonstrates the flexibility of the system and the ease to which existing scanner hardware can be retrofitted for parallel transmit capabilities if multiple amplifier channels and a transmit array are available.

The system demonstrated the ability of a LabVIEW-based control system to perform magnetic resonance imaging. The use of LabVIEW for the programming environment enabled FPGA programming without the need to learn an additional hardware description language. With LabVIEW experience being more commonplace in academic circles, adapting the system for added functionality is more tangible as compared to other software architectures (i.e. microcontrollers or non-LabVIEW programmable FPGA-based boards) that require device-specific expertise and low-level programming experience. The control system described represents an approach to consider when developing parallel transmit hardware for use on existing scanners.

CHAPTER V

CALIBRATING THE TRANSMIT CHAIN FOR LINEARITY

V.1 Introduction

In magnetic resonance imaging, a slice-selective pulse in combination with a selective gradient excites a spatially restricted region. The Fourier transform of the pulse defines the excitation profile, and the bandwidth of the pulse in combination with the gradient determines the region of the sample excited. The gradient imposes a linear magnetic field across the sample, changing the spatial distribution of the field, and thus, the frequency of precession of the spins in the system. The radiofrequency pulse contains a set bandwidth and interacts with the spins corresponding to the frequencies contained in the pulse. In the simplest case, a sinc pulse excites a rectangular slice with a thickness corresponding to the pulse bandwidth.

Successful implementations of multi-slice imaging and spatially selective excitations depend on an accurate pulse. If the actual on-coil pulse deviates from the user-defined pulse, out-of-slice excitation occurs, degrading performance. Unfortunately, producing an accurate representation of the user-defined pulse is a nontrivial problem due to the nonlinearities in the transmitter chain, and adjustments in the pulse are needed to compensate for the system nonlinearities.

Several approaches for linearity correction exist with common implementations being Cartestian-feedback [68-70], prewarping [71], or vector iterative pre-distortion [72]. In a simplistic view, all approaches rely on monitoring the actual on-coil pulse or

amplifier output, extracting the amplitude and phase information from the pulse, and comparing the actual data to the demand data. A compensation is applied to correct for the nonlinearities in the system. The Cartesian feedback and vector iterative pre-distortion approaches perform a correction after each pulse, requiring a hardware and control subsystem with the capability of performing a downconversion and applying a correction within the time constraints of the pulse sequence. Prewarping requires a connection to the scanner receiver chain complicating setup. Simplified pre-distortion techniques rely on characterizing the transmit chain nonlinearities and performing a look-up at the start of an experiment to define a distorted pulse that given the system nonlinearities produces the desired demand pulse when output through the system.

A pre-distortion approach was implemented to correct for the system nonlinearities in both amplitude and phase. The built-in on-coil probes provided the feedback mechanism for collecting the amplitude and phase data required to create a look-up table. The remaining sections of this chapter will discuss the details of the implementation and pertinent results with the progression of discussion as follows:

- A description of the on-coil probes used to detect and characterize the system nonlinearities.

- An explanation of the linearization undersampling technique to extract the amplitude and phase data

- The programming implementation to process the data and create the look-up table

- A scattered interpolation routine to pre-distort the pulse

96

- Automating the linearization process

V.2 Materials and Methods

V.2.1 Terminology

In the following sections, the on-coil waveform prior to linearization is described as the "actual" waveform. The desired reference waveform is referred to as the "demand" or "accepted" waveform.

V.2.2 Constructing the On-Coil Probes

The on-coil probes were previously mentioned in Chapter III. The design of the probes consisted of a shielded loop design to minimize the effect of electric fields. The probes were positioned at the center of the rung between two sections of distributed capacitance. Flexible coaxial cable was molded into a loop by wrapping the cable around a 0.25" diameter tool. A small section at the center of the loop was scored using a razor and the shield removed to expose the inner conductor. At the base of the loop, the center conductor was exposed and connected to the shield. The probes were secured to the rung, and the cable dropped down to the array shield. The shield of the cable was electrically connected to the array shield for the run between the probe and the match and tune board. The cable was stripped and connected to the match and tune board where the signal was easily accessible via a SMA connector. Pictures of the current probe and the connection along the array shield are shown in Figure V.1.

V.2.3 Control System Additions to Acquire Linearization Data

The control system discussed in Chapter IV operates on a PXI platform with programs developed in LabVIEW. To add the capability to perform an automated

linearization routine, a high-speed digitizer capable of sampling the on-coil 128 MHz radiofrequency pulse was required. The amplitude and phase data was needed, requiring a digitizer with two inputs in order to sample the input RF hard pulse to acquire a phase reference for the on-coil data. A PXI-5152 high-speed digitizer (2 AI, 1 GS/s) was purchased and added to the existing PXI chassis.
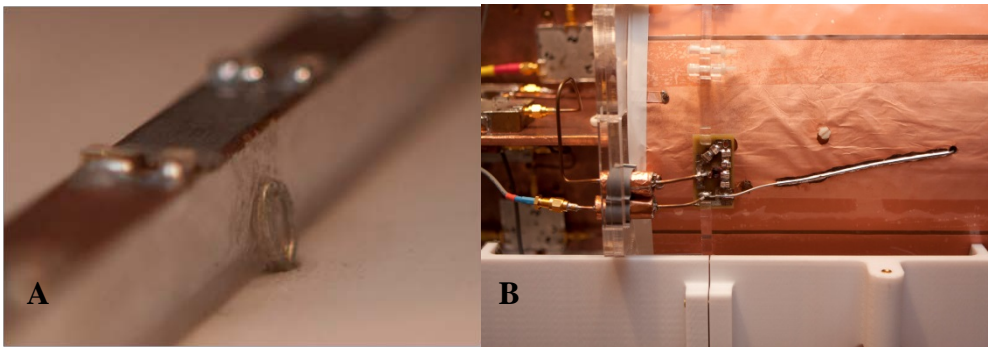
Figure V.1: Current probe details.
A: The shielded current probe positioned between two distributed capacitors along the rung. B: The current probe connection runs along the shield of the array and connects to the match and tune board.

V.2.4 Setup to Acquire the Linearization Data

The on-coil probe senses the output waveform from the transmit system and feeds back the signal to the high-speed digitizer. To linearize another channel, the current probe connection to the digitizer is physically changed. Switch hardware in development will enable full automation of the routine. In this configuration, the current probe signal is feedback into a chassis containing two single-pole-four-throw (SP4T) and one single-pole-double-throw (SPDT) absorptive switches. Each current probe connects to one of the inputs of the SP4T switch after the signal is attenuated and split via a power

98

splitter. The splitter is necessary to route the current probe signals to the monitoring and switch hardware simultaneously (discussed in Chapter VI). The outputs of the SP4T switch are connected to the inputs of the SPDT switch. Three digital control lines provided by the PXI-7853R control the switches and determine which probe is connected through the switch hardware to the high-speed digitizer as shown in Figure V.2.

The RF hard pulse is split at the input to the modulator chassis with one output connected to the modulator and the other used for the linearization routine. The RF is amplified to increase the 20 mV peak-to-peak signal to a voltage level with similar magnitude to the current probe signal in order to maximize the dynamic range of the digitizer. The unblank out signal from the modulator is connected to the PXI-5152 to trigger the acquisition.

V.2.5 Undersampling to Acquire the Magnitude and Phase Data

The linearization routine outputs and acquires a series of pulses to step through the possible amplitude and phase points using a series of triangle waveforms with constant phase as shown in Figure V.3.
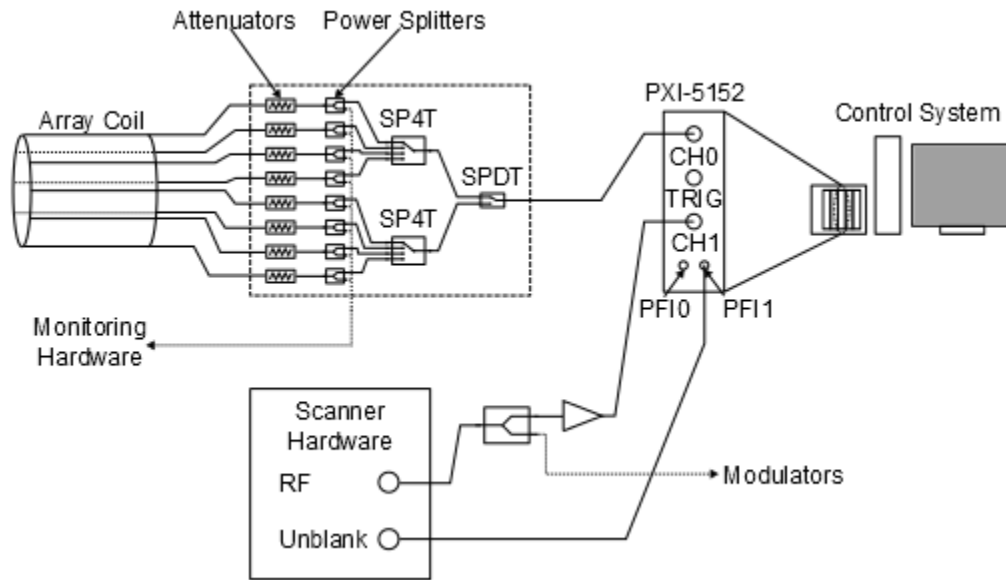
Figure V.2: Diagram depicting the setup to acquire the linearization data.
The signals from the on-coil probes are attenuated and fed through a power splitter. The
splitter divides the signals to provide both the monitoring and switch hardware with
access to the signals on the coil. The signals from the splitters are connected to the inputs
of two single-pole-four-throw (SP4T) switches, and the outputs of the SP4T switches are
connected to a single-pole-double-throw (SPDT) switch. The output of the SPDT is fed
to the input of the high-speed digitizer (PXI-5152). The RF hard pulse from the scanner
is split, amplified, and connected to the digitizer, and the master exciter unblank (trigger)
is connected to the digitizer. Three digital control lines from the PXI-7853R control the
switches for automation (connections not shown).
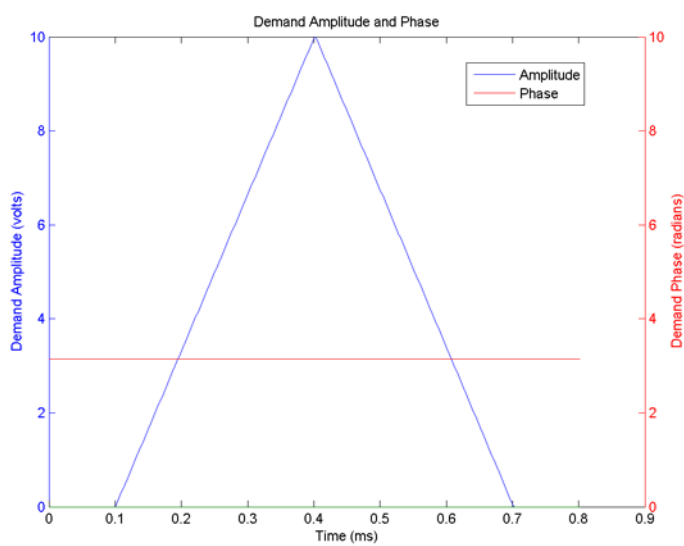
Figure V.3: Example demand amplitude and phase pulse.
All demand pulses for acquiring the linearization data consist of a triangle with a constant phase line. A series of 36 pulses are used with phase ranging from 0-350 degrees.

The magnitude and phase data are obtained from the on-coil waveform and the RF hard pulse. An undersampling technique was utilized to minimize the memory requirements of acquiring two fully sampled sets of 128 MHz data. Instead, the waveforms were undersampled at a rate of 40 MS/s. To extract the magnitude and phase information, the RF hard pulse was first characterized. The hard pulse operates at a single frequency and after undersampling, aliases to a new frequency associated with the original carrier frequency. The aliased frequency is determined by taking the Fourier transform of the RF hard pulse and finding the frequency with maximum content. The location of this frequency corresponds to the aliased carrier and defines the frequencies to extract from the Fourier transform of the more complex on-coil pulse. The Fourier transform of the on-coil waveform is computed, and a subset of points is extracted

101

around the aliased frequency for both the RF hard pulse and on-coil waveform Fourier transforms. The extracted data is shifted to the center of the Fourier transform corresponding to dc, and the remaining points in the transform are zeroed. The inverse Fourier transform is computed, and the downconverted magnitude and phase data of the pulse is extracted. The on-coil waveform phase is referenced to the RF hard pulse phase to provide meaningful phase information and ensure repeatability. A block diagram of the routine is shown in Figure V.4.
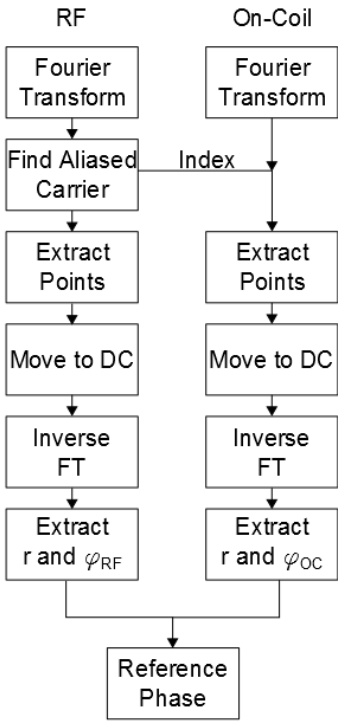
Figure V.4: Block diagram of the undersampling routine used to extract the magnitude and phase data from the on-coil waveform.

V.2.6 Processing the Magnitude and Phase Data

Some conditioning of the magnitude and phase data is required before extracting a subset of points for a look-up table. Both data sets are downsampled using a linear interpolation to match the time step of the original file. The phase data contains noise and follows the magnitude data in a near linear fashion. A combined approach using a linear fit and the real data was used to smooth the phase data.

A linear fit to the phase data is performed, and the error between the data and the fit computed. The mean and standard deviation of the error were calculated, and for each error point, error was subtracted from the mean and divided by the standard deviation as described in Equations 4.1 and 4.2.

$$e = Data - Fit \hspace{6cm} [4.1]$$

$$z_i = \left| \frac{\mu - e_i}{\sigma} \right| \hspace{6cm} [4.2]$$

Where e is the error, μ is the mean of the error, σ is the standard deviation of the error, and z is the computed z-score. With a data set and a linear fit, 99.7% of the data falls within 3 standard deviations of the mean (a z-score of 3), and with this in mind, the z-score was used as a tuning parameter to discard jumps in the phase data that were not real and due to measurement error. The z-score of the error was computed, and if less than 3, the corresponding data point was kept, but if greater than 3, the corresponding data point was discarded. A linear fit was recomputed on the data with the discarded points removed, and a second iteration was performed with a tighter tuning parameter of one, corresponding to 68% of the data. After cycling through all points a second time,

the linear fit was recomputed resulting in a fit minimally affected by jumps in the phase data.

Near the peak of the triangle, the phase data levels off in a non-linear fashion, and to fit this data with the linear fit alone would inject error into the pre-distortion rather than remove it. Thus, a routine was established that indexed the data from the center and computed the phase variation between adjacent points, which was observed to be minimal, and a threshold of 0.1 degrees was established based on the observed phase behavior. If the phase jump between points varied more than the defined threshold, the indexed data point was discarded and replaced with the linear fit. This allowed extraction of the real phase data without the noise spikes present in the original measured data as shown in Figure V.5. The spikes are not real but a result of the measurement, and they occur at indices corresponding to low amplitude values where phase is difficult to compute. Additionally, since the demand pulse consisted of a triangle ramp with constant phase, each ramp section in magnitude corresponded to duplicate phase data only reversed, and therefore, the two sets of duplicate data were averaged to further reduce error.
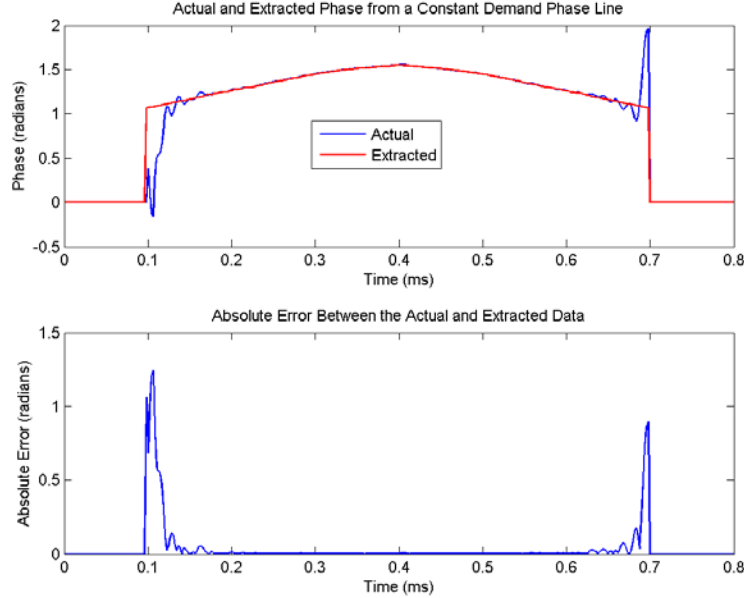
Figure V.5: Actual and extracted phase data.
The jumps in the phase data are evident near the beginning and end of the pulse, corresponding to low amplitude values.

The ramp section of the magnitude data was extracted, and both the magnitude and phase data were downsampled to 36 points. The routine steps through 36 phase steps corresponding to 10 degree step sizes, and the resampling created a 36x36 grid of points. The magnitude and phase arrays were inserted into the matrix at the appropriate index corresponding to the demand phase, and the process was repeated with a new demand pulse.

V.2.7 Conditioning the Data for the Pre-distortion

Prior to performing the pre-distortion, the data was pre-conditioned. The look-up table for the magnitude data after the previous operation was not normalized to the digital-to-analog converter (DAC) output range of 0 to 10 volts for the positive triangle.

The last column of the look-up table corresponded to the maximum point in the acquired waveform, the value corresponding to the demand waveform value of 10. Since the magnitude of the on-coil waveform varies significantly with phase due to the nonlinearities of the transmit chain, the minimum value in this column was located and the table normalized to force this value to 10.

The phase data required additional processing. The look-up table was unwrapped and forced to be within a range of 0 to 2π. The first row in the look-up table corresponded to data associated with the demand phase of zero, and the table was adjusted to force the mean value of the first row to a value of zero.

The amplitude and phase look-up table data was acquired with the control system before and after linearization and the resulting surfaces plotted to demonstrate the improvement in linearity. As a secondary demonstration, the amplitude and phase data was acquired for a grid of amplitude and phase points using the network analyzer. The control system was modified to automate the data collection by extracting the trace data from the network analyzer. The network analyzer was setup to perform a S_{21} power sweep with zero span at a power level corresponding to the required 20 mV peak-to-peak. The S_{21} measurement was taken between the input of the modulator and output of the transmit chain. The continuous wave frequency was set to 128.07 MHz, and the analyzer was setup with an external trigger. The number of points and intermediate frequency (IF) bandwidth were adjusted to correspond to the pulse length of the demand waveform. A series of hard pulses with constant phase were acquired and plotted before and after linearization.

106

V.2.8 Pre-distorting the Demand Pulse

A native LabVIEW VI was used to predistort the demand pulse. The VI required one-dimensional array inputs, so the look-up table data was restructured into the appropriate format. Arrays were created for the demand magnitude and phase data that corresponded to the look-up table array. The ask amplitude and phase data was imported from the user-defined file, and the data was checked to ensure it remained within the bounds of the look-up table to prevent extrapolation. All amplitude and phase data was converted to in-phase and quadrature data, and the interpolation was performed in I and Q to eliminate possible issues with a phase interpolation. The interpolation routine would not recognize phase values of 0 and 2π as corresponding to the same value and could create noise in the phase pre-distortion. A scattered interpolation routine was used to perform the interpolation based on computation of a convex hull containing the target point. After the pre-distortion in I and Q, the data was converted back to amplitude and phase data.

V.2.9 Automating the Linearization Routine

With the switch hardware, automating the routine to acquire all of the linearization data at once is straightforward. A for loop was added around the existing linearization code, and when the user selects "Acquire Linearization Data for All Channels", the loop executes eight times. The control lines are set after loading the waveform data to the onboard memory of the target, and the target program waits to ensure the control lines have switched prior to playing out the waveform and acquiring the data. The user has the option to acquire one channel at a time by selecting the

appropriate channel from the drop-down menu control, and selecting "Acquire Linearization Data for Single Channel". In this instance, the loop executes once and sets the control lines to correspond to the appropriate channel.

During execution the front panel display updates with the acquired amplitude and phase data and notes the current stage of the linearization routine. After acquiring all data for the look-up table, before and after linearization waveforms consisting of a triangle with a phase ramp were played out and displayed on additional front panel indicators shown in Figure V.6.

V.2.10 Experiments to Verify the Linearity Correction

During acquisition of the linearization data, two surfaces are acquired: an amplitude and phase surface. These surfaces provide the look-up table and characterize the system linearity. The surfaces were obtained in both cases with and without linearity correction to note the improvement.

Several pulses were acquired before and after linearization using the high-speed digitizer to demonstrate improvement with the routine and compute the error in both cases.

The amplitude and phase variation of the system with and without the linearity correction applied was acquired with a network analyzer to provide an additional check using a measurement system separate from the control system. The network analyzer was setup for an S_{21} power sweep with a 0 dB span and a RF at the appropriate 20 mV peak-to-peak power level to the modulators. The IF bandwidth and number of points were adjusted to correspond to a pulse length of approximately two milliseconds. The

network analyzer was setup to measure both amplitude and phase using the dual channel mode, and the analyzer was externally triggered by a function generator. A GPIB-to-USB controller was used to connect the network analyzer to the control system PC. An additional state termed "Collect Data" was added to the control system to loop through a series of hard pulses and extract the network analyzer trace data and write the extracted amplitude and phase data to a file. This data was then processed to further validate the linearization results.

Imaging experiments were performed on a 3 Tesla GE clinical research scanner to evaluate the slice profile before and after linearization was applied. Multiple slices were acquired through a slab orthogonal to the excited slice, and the signal intensity of a single pixel was plotted from each slice and the resulting slice profile plotted for both cases.

V.3 Results and Discussion

V.3.1 Graphical User Interface for Linearity Correction

The graphical user interface showing the linearization front panel is shown in Figure V.6. The user has the option to linearize a single channel using the "Linearize 1 Channel", "Channel" drop-down menu, and "Vertical Range" controls. The range control adjusts the dynamic range of the high-speed digitizer with available values of 0.1, 0.2, 0.4, 1, 2, 4, and 10 volts peak-to-peak. If the user desires to linearize all channels at once, the "Linearize All Channels" control is used, and the eight-way switch will automatically adjust to connect the appropriate on-coil probe to the high-speed

digitizer. The switch hardware is in development, but the control lines have been verified during the linearization operation.

While the linearization data is acquired, the acquired amplitude and phase data is plotted in the display on the right, and the "Linearization Status" indicator updates with the corresponding phase line grabbed. After the data is acquired for a channel, the "Linearization Done" indicator changes to a value of true (bright green LED), and the center displays update to show a waveform before and after linearization to demonstrate the correction. The demand waveform for the demonstration consists of a triangle with a linear phase ramp.

The remaining controls on the bottom right are used to collect data with the network analyzer to acquire amplitude and phase measurements.
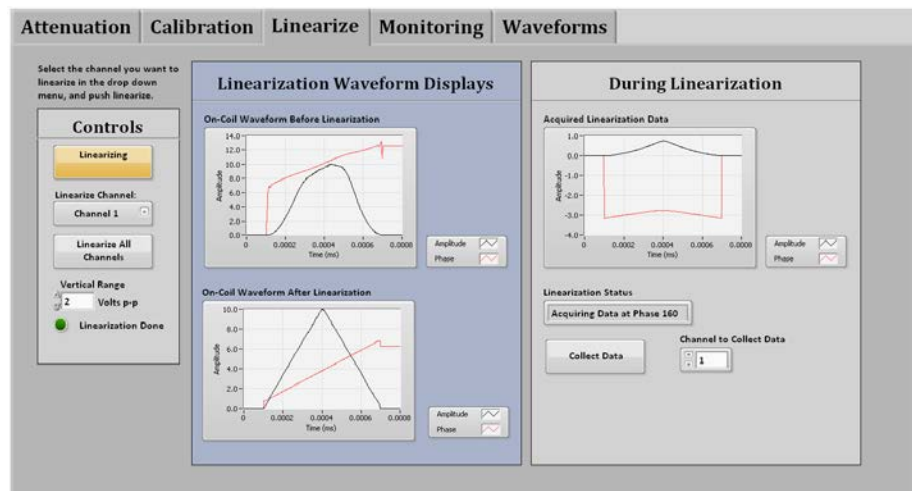
Figure V.6: Front panel for the linearization routine.
The user has the option to linearize one channel or all channels.

110

V.3.2 Amplitude and Phase Characterization

V.3.2.1 Data Acquired Using the Control System

The amplitude surfaces acquired using the control system with and without the linearity correction are shown in Figure V.7. The non-linearities in the system display significant variation in amplitude with changes in phase. For reference, the demand surface is shown containing the ideal amplitude surface with phase variation.
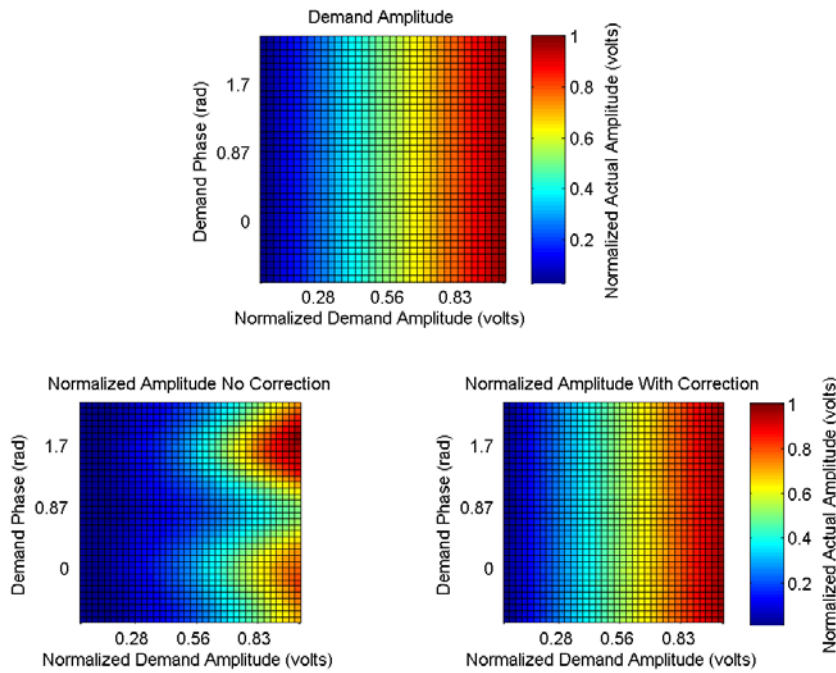
Figure V.7: Amplitude surfaces acquired using the control system without (bottom left) and with (bottom right) the linearity correction applied.
There are significant non-linearities in amplitude as the phase is varied. The demand surface (top) is shown as the ideal amplitude surface.

The absolute error between the demand surface and both the with and without linearization surfaces were computed and are shown in Figure V.8. The maximum

absolute error for the case with no correction is 0.5. With the linearity correction applied, the maximum absolute error is reduced to 0.03.

Figure V.8: Absolute error surfaces of the amplitude with and without the linearity correction applied.
The surfaces are identically scaled.

The phase surfaces with and without the linearity correction applied are shown in Figure V.9. The phase deviation with amplitude is less severe than the amplitude variations with phase, and the maximum discrepancy across a constant phase line with no correction is 0.404 radians (23 degrees) which reduces to .029 radians (1.6 degrees) with the linearity correction applied.
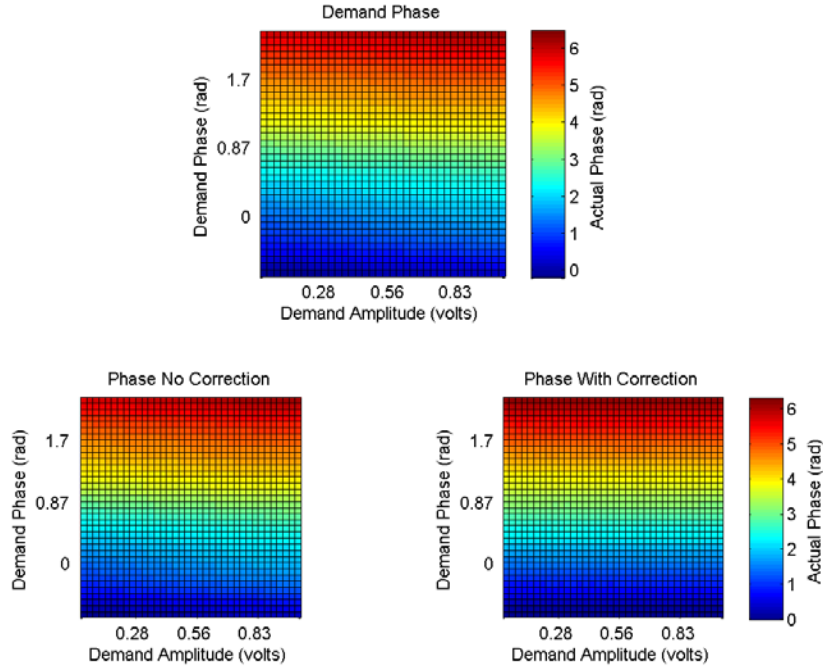
Figure V.9: The ideal phase surface (top) is shown for a reference.
The phase surfaces without (bottom left) and with (bottom right) linearity correction.
The variations in the phase surfaces are more subtle than the drastic changes seen in the
amplitude surfaces. For a better depiction of the variation in phase with amplitude, see
Figure V.10.

To depict the phase variation with amplitude more clearly, a single line

corresponding to a phase of π radians from each of the phase surfaces was plotted in

Figure V.10. The phase discrepancy is more evident in these plots with variations of 0.34

radians with no correction (19.4 degrees) and 0.01 radians with correction (0.8 degrees).

Recall the phase data is extracted using the control system routine described in the

previous section (the data is the acquired look-up tables for the two cases).

113
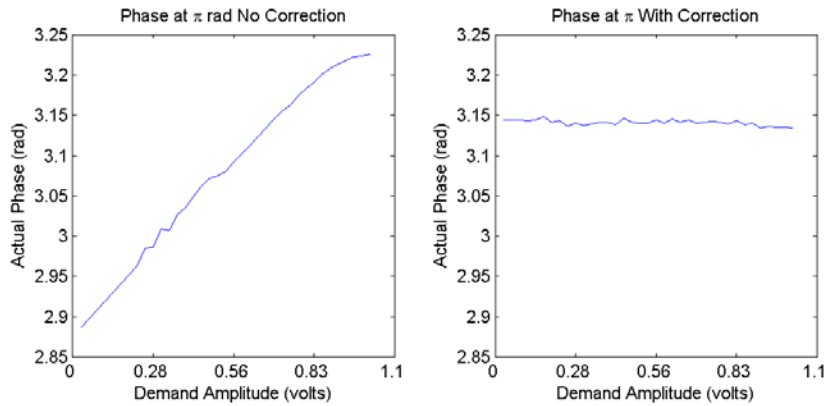
Figure V.10: The actual phase corresponding to a demand phase of π radians is plotted for both with (right) and without (left) linearity correction cases.

The absolute error for both phase surfaces was computed and is shown in Figure V.11. The maximum absolute error without linearization is 0.40 radians (22.8 degrees) which reduces to 0.04 radians (2.2 degrees) with the correction applied.
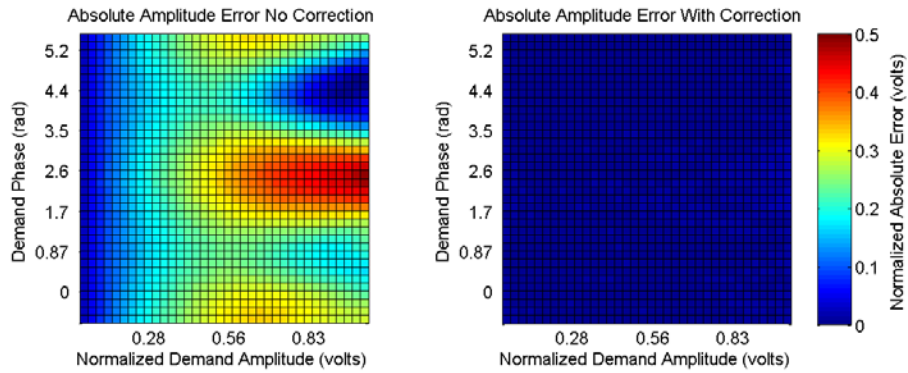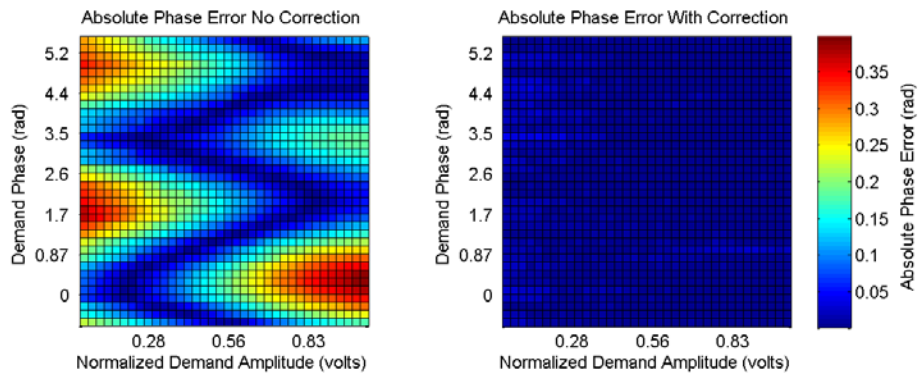

Figure V.11: Absolute phase error surfaces with and without the linearity correction applied.
Both surfaces are identically scaled.

V.3.2.2 Data Acquired with the Network Analyzer

The network analyzer was used to acquire amplitude and phase data to further validate the linearization routine. Data was collected at demand amplitude values of 1 – 10 volts and phase values of 0 – 324 degrees, creating a 10x10 grid. An arrow plot was generated to show the mapping between the demand data and the actual data. The plots with and without the linearity correction applied are shown in Figure V.12. The linearity of the system significantly improves with the correction as indicated by the alignment of the actual data to the demand data.

The absolute error between the demand and actual waveforms was computed from the same dataset and is plotted in Figure V.13 below. This data corroborates the results in the previous section acquired using the control system. The maximum absolute phase error was 20.4° in the uncorrected case and 2.8° in the case with the linearity correction applied. The amplitude error changed from an absolute error of 0.454 in the uncorrected case to 0.034 in the case with the linearity correction applied.

V.3.3 Amplitude and Phase Waveforms

A series of amplitude and phase waveforms were acquired with and without the linearity correction applied and the root-mean-square error calculated in each case. The acquired amplitude and phase waveforms are shown in the following figures for demand pulses consisting of: a triangle with a phase ramp (Figure V.14 and Figure V.15), a sinc pulse (Figure V.16 and Figure V.17), a hard pulse with a phase ramp (Figure V.18 and Figure V.19), a triangle pulse with constant phase (Figure V.20 and Figure V.21), a Bloch-Siegert pulse containing phase ramps to create the off-resonance in the Fermi

115

pulse (Figure V.22 and Figure V.23), and a flyback echo-planar pulse commonly used in

parallel transmit experiments and used to acquire the data in Chapter IV (Figure V.24

and Figure V.25).
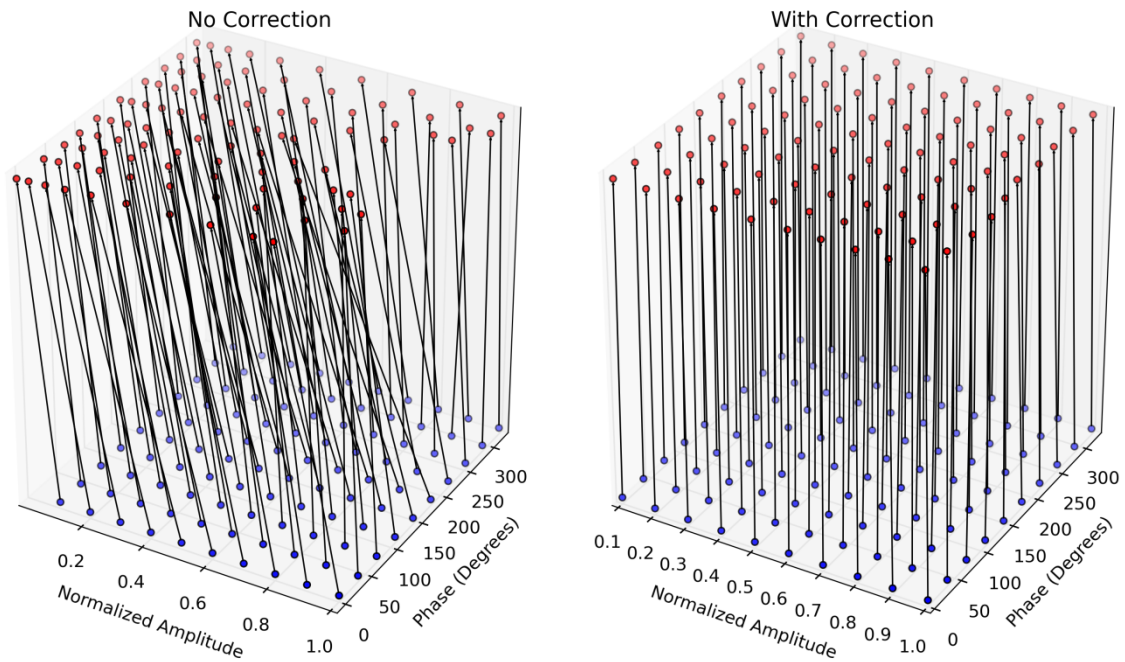
Figure V.12: Plots demonstrating the amplitude and phase mapping between the demand and actual values.
The blue circles correspond to the demand data, and the red circles correspond to the actual data. In the case with no correction, the lines connecting the blue and red circles cross, and the red circles do not align with the blue circles indicating poor linearity. In the case with the correction, the red circles align with the blue circles, and the red circle lies on a grid indicating linearity.
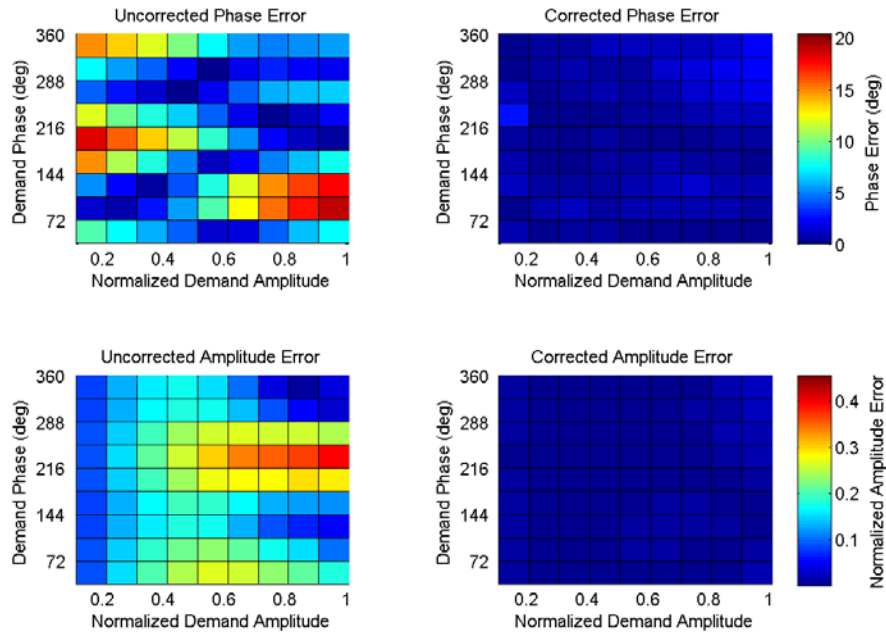
Figure V.13: Absolute error plots between the demand and actual waveforms.
The phase absolute error plots are shown in the first row (identically scaled), and the amplitude absolute error plots are shown in the bottom row (identically scaled). The left side displays the errors for the uncorrected case while the right side plots show the corrected case. The maximum phase error in the uncorrected case was 20.4° which reduced to 2.8° in the corrected case. The maximum amplitude error for the uncorrected case was 45.5% which reduced to 3.4% in the corrected case.
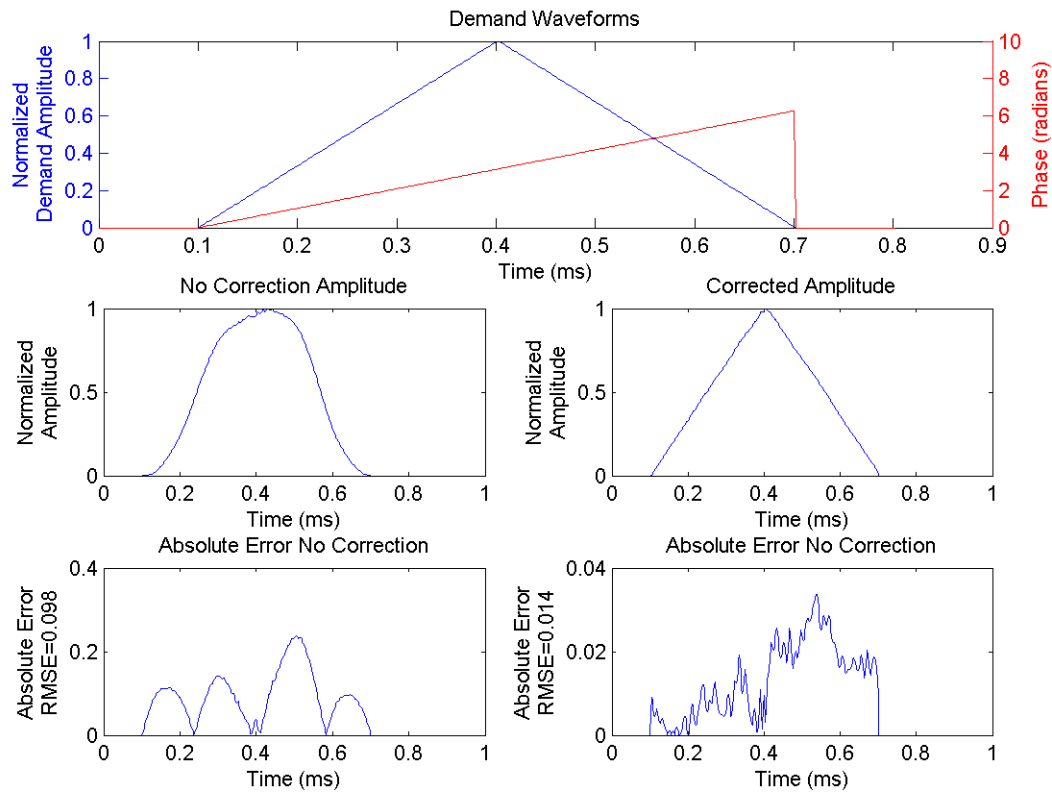
Figure V.14: Acquired amplitude waveforms for a demand pulse consisting of a triangle with a phase ramp for both the uncorrected (left) and corrected (right) cases.
The demand waveform (top) consists of a triangle with a constant phase ramp. The normalized amplitude and absolute error are plotted for both cases with the RMSE computed. In the uncorrected case, the RMSE was 0.098, and was reduced to 0.014 in the corrected case.

Figure V.15: Acquired phase waveforms for a demand pulse consisting of a triangle with a phase ramp for both the uncorrected (left) and corrected (right) cases. The demand phase waveform (top) is shown in red. The absolute error between the demand and actual waveforms for both cases was computed and plotted. The RMSE for each case was computed, and for the uncorrected case the RMSE was 0.236 which reduced to 0.011 in the corrected case.

Figure V.16: Acquired amplitude waveforms for both the uncorrected (left) and corrected (right) cases.
The absolute error between the demand and actual waveforms was plotted for both cases, and the root-mean-square error computed. The RMSE for the uncorrected case was 0.047 which reduced to 0.005 in the linearized case.

Figure V.17: Acquired phase waveforms for both the uncorrected (left) and corrected (right) cases.
The absolute error between the demand and actual waveforms was plotted for both cases, and the root-mean-square error was computed. The RMSE for the uncorrected case was 0.977 which reduced to 0.032 in the linearized case.
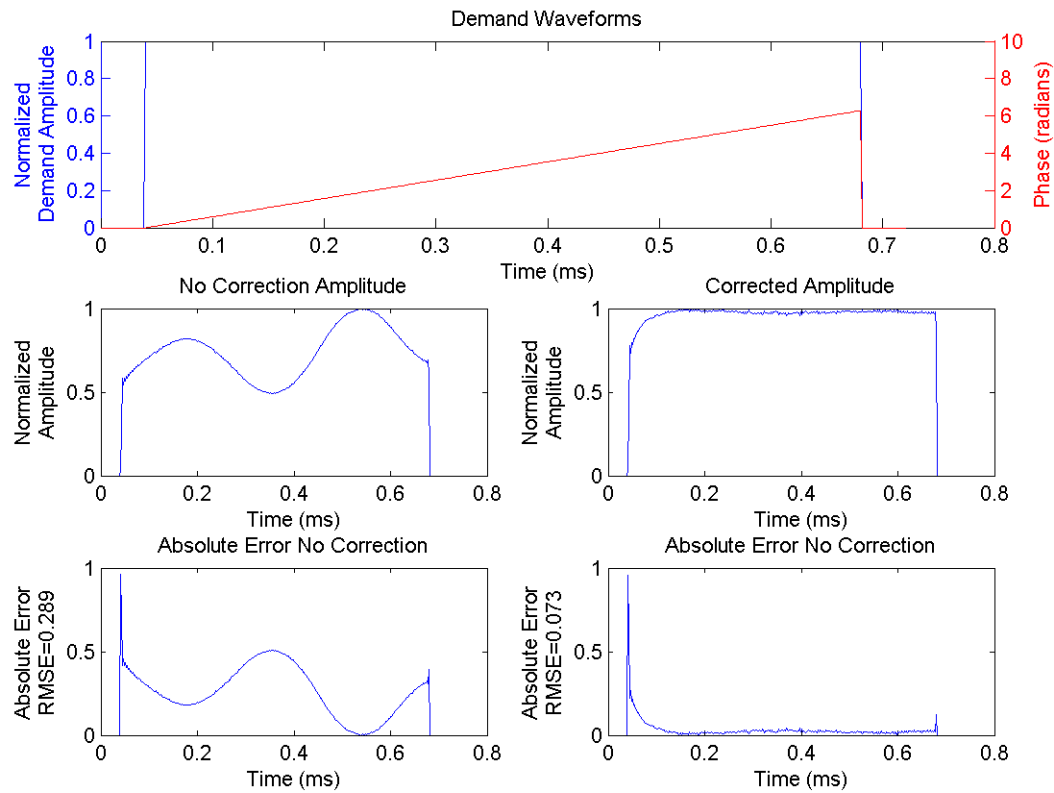
Figure V.18: Acquired amplitude waveforms for a demand waveform consisting of a constant amplitude and phase ramp for both the uncorrected (left) and corrected (right) cases.
The absolute error between the demand and actual waveforms was plotted for both cases, and the root-mean-square error was computed. The RMSE for the uncorrected case was 0.289 which reduced to 0.073 with the linearity correction applied. The corrected case shows some non-linearity in the beginning of the pulse due to the turn-on time for the amplifier.
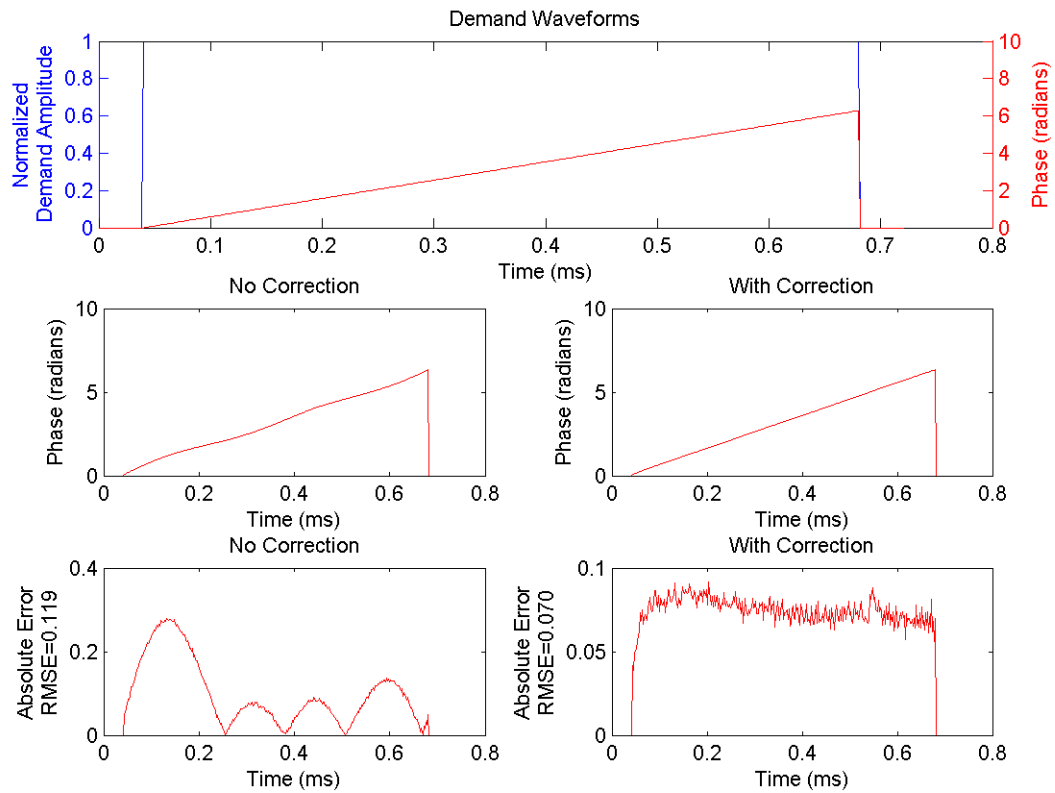
Figure V.19: Acquired phase waveforms for a demand waveform consisting of a constant amplitude and phase ramp for both the uncorrected (left) and corrected (right) cases.
The absolute error between the demand and actual waveforms was plotted, and the root-mean-square error was computed. The RMSE for the uncorrected case was 0.119 which reduced to 0.070 in the linearized case.
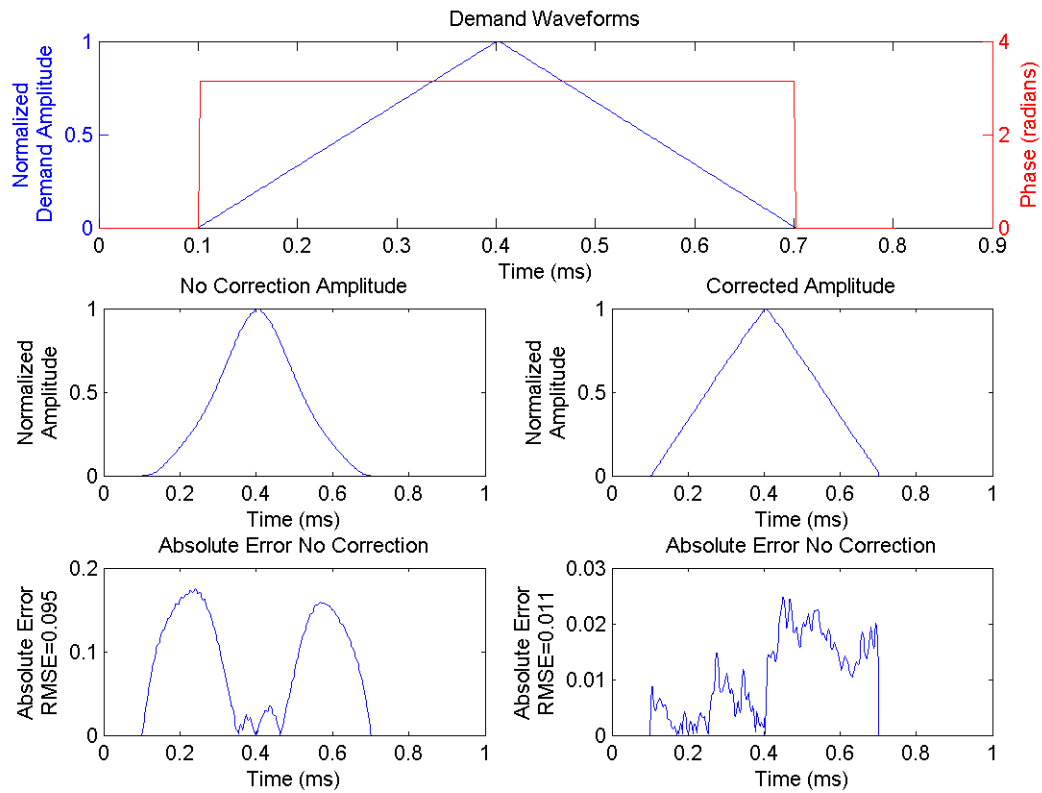
Figure V.20: Acquired amplitude waveforms for a triangle with a constant phase of π radians for both the uncorrected (left) and corrected (right) cases.
The absolute error was plotted, and the root-mean-square error for each case computed. The RMSE for the uncorrected case was 0.095 which reduced to 0.011 in the corrected case.

124
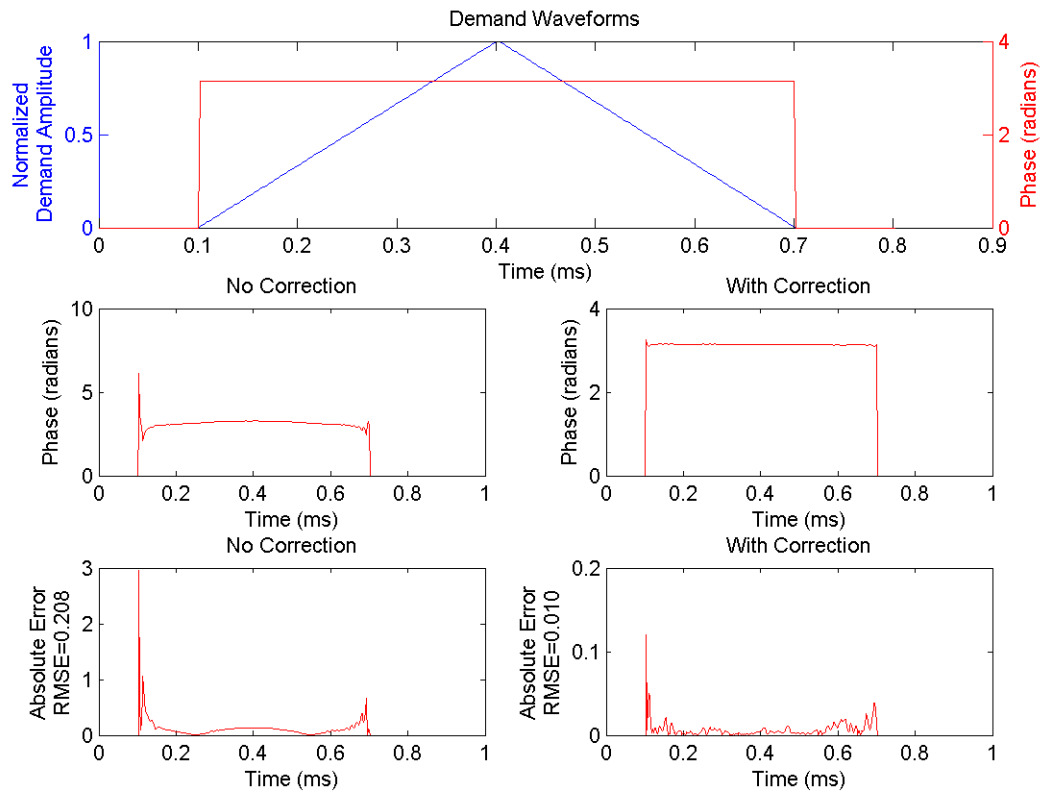
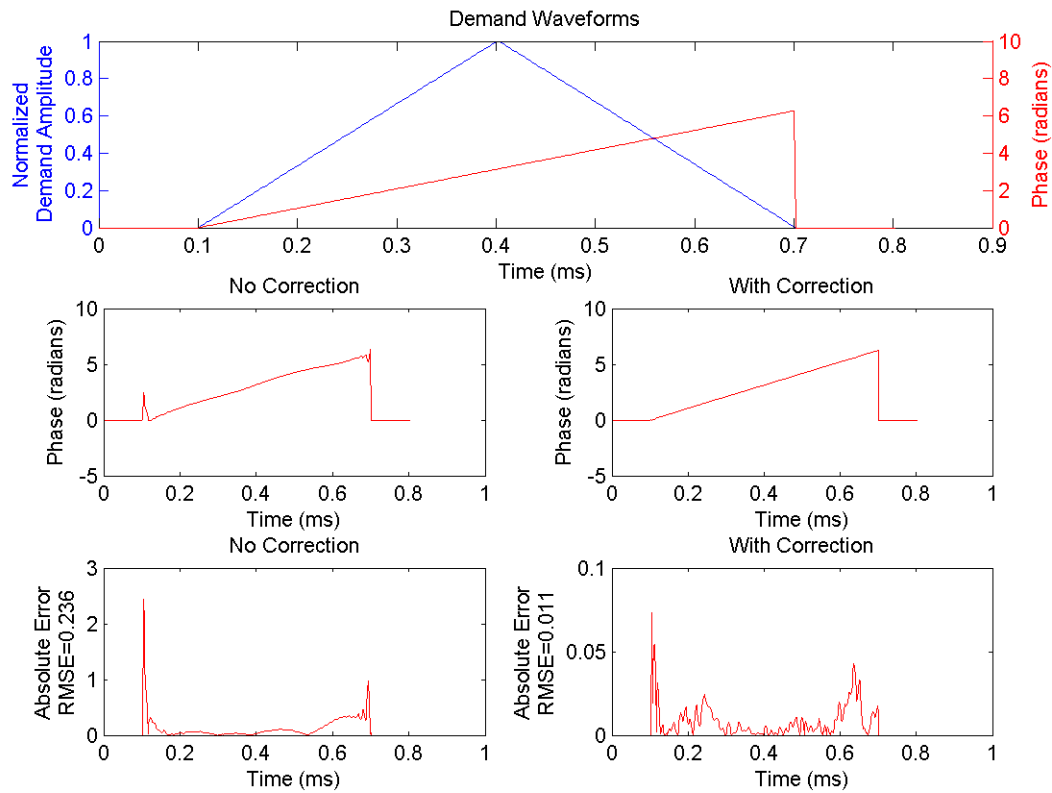**Figure V.21: Acquired phase waveforms for a demand triangle pulse with constant phase for both the uncorrected (left) and corrected (right) cases.**
The absolute error was plotted, and the root-mean-square error was computed. The RMSE for the uncorrected case was 0.206 which reduced to 0.010 in the corrected case.
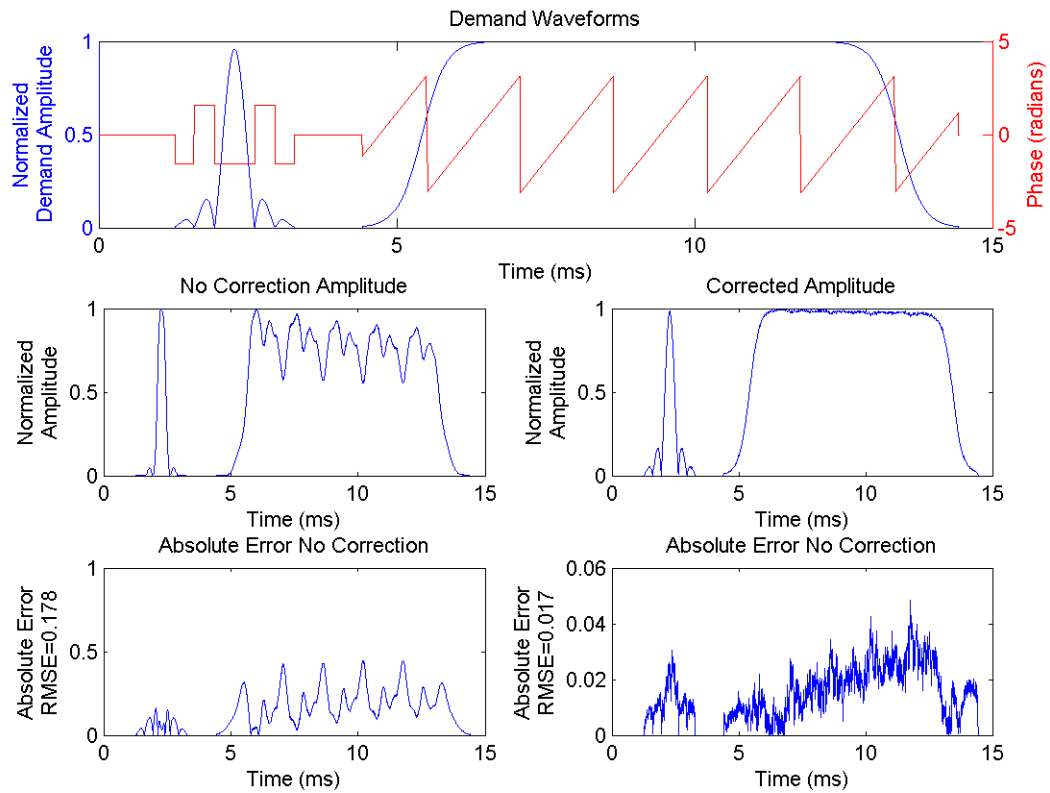
Figure V.22: Acquired amplitude waveforms for a Bloch-Seigert demand waveform for both the uncorrected (left) and corrected (right) cases.
The absolute error between the demand and actual waveforms was computed, and the root-mean-square error was computed. The RMSE for the uncorrected case was 0.178 which reduced to 0.017 in the corrected case.

126
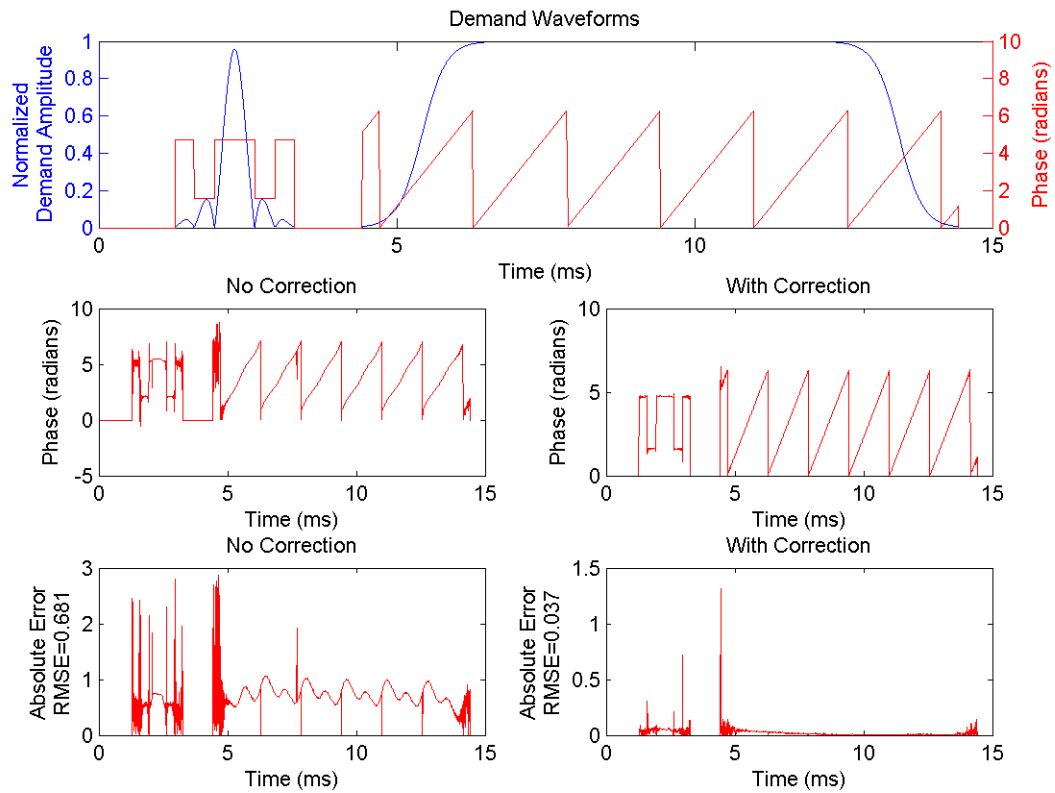
Figure V.23: Acquired phase waveforms for the Bloch-Siegert demand pulse for both the uncorrected (left) and corrected (right) cases.
The absolute error between the demand and actual waveforms was plotted, and the root-mean-square error was computed. The RMSE for the uncorrected case was 0.681 which reduced to 0.037 in the corrected case.
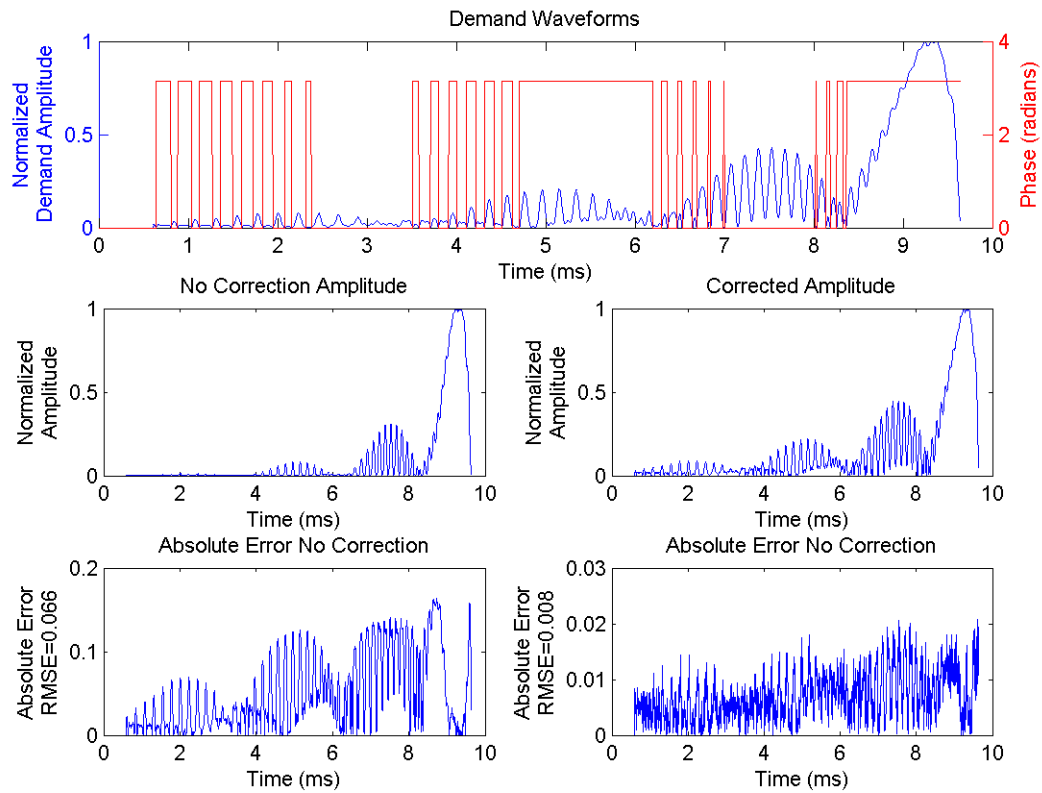
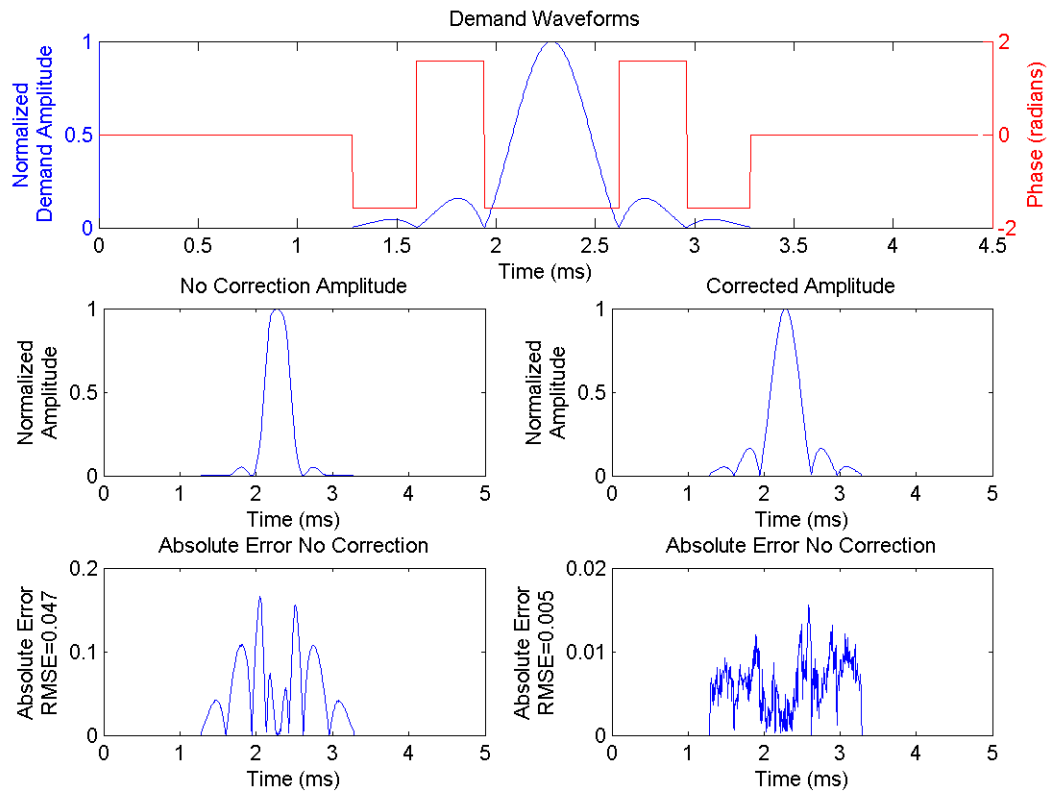**Figure V.24: Acquired amplitude waveforms for a echo-planar demand pulse for both the uncorrected (left) and corrected (right) cases.**
The absolute error between the demand and actual waveforms was plotted, and the root-mean-square error was computed. The RMSE for the uncorrected case was 0.066 which reduced to 0.008 in the corrected case.
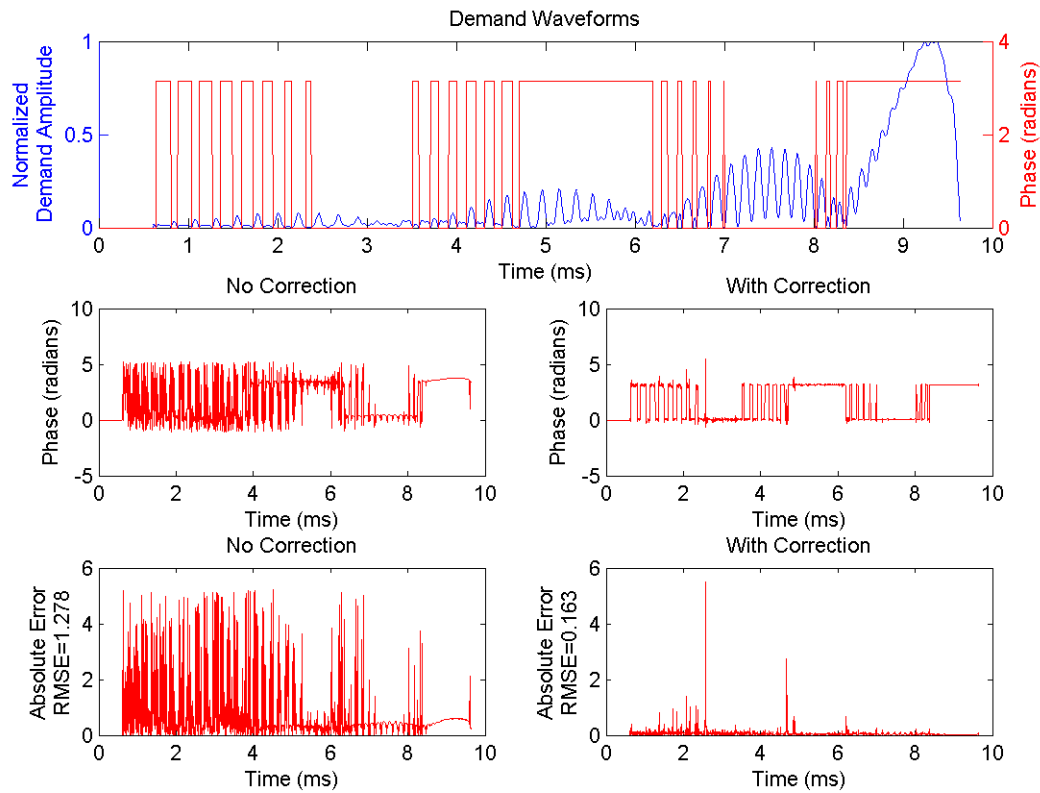
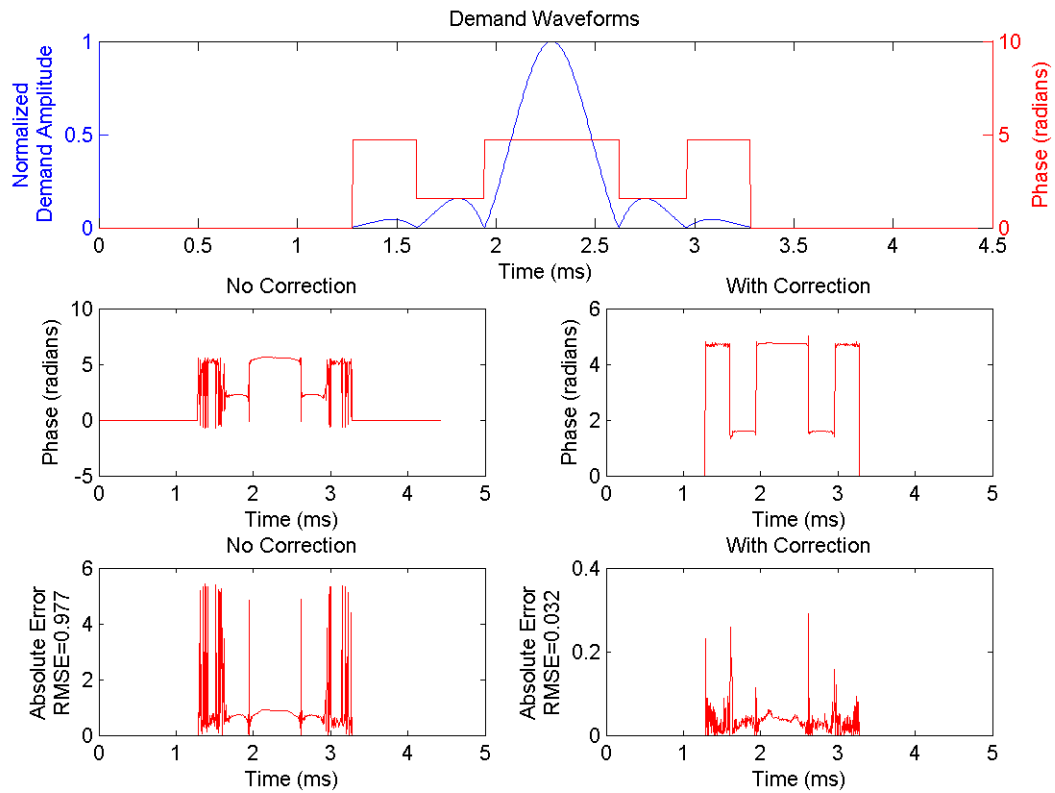**Figure V.25: Acquired phase waveforms for a echo-planar demand pulse for both the uncorrected (left) and corrected (right) cases.**
The absolute error between the demand and actual waveforms was plotted, and the root-mean-square error was computed. The RMSE for the uncorrected case was 1.278 which reduced to 0.163 in the corrected case.

In the examples shown, there are noticeable jumps in the phase data. The phase data is not processed to remove these jumps, and as a result, they are taken into account in the error computation. The phase jumps are due to measurement errors due to the difficulty in computing phase accurately at low amplitude values. The phase jumps are less noticeable in the corrected cases due partially to the improvement in the amplitude linearity, which increases the signal level at the low demand amplitude points.

129

Additionally, the system exhibited noticeable droop, which is evident in longer pulses with constant amplitude, particularly the corrected amplitude plot of the Bloch-Siegert pulse in Figure V.22. There is a slight decrease in the amplitude from left to right along the Fermi pulse (the region of constant amplitude with phase ramps) due to fluctuations in the supply voltages to the transmit system termed droop. This issue is more evident in looking at the histogram of the error (not the absolute error) between the demand and actual waveforms shown in Figure V.26. We would expect the routine to have a mean error at or near zero, but in the histogram, the mean of the error is to the right due to the system droop.

V.3.4 Slice Profile Results

To validate the linearization routine in the imaging environment, an experiment was performed to look at the slice profile. A 5 millimeter slice was excited with the transmit system, and multiple slices were acquired in a plane orthogonal to the excitation slice. A pixel was chosen in each of the orthogonal slices and the signal intensity plotted. The slice profile results with and without the linearization applied are shown in Figure V.27 and Figure V.28.
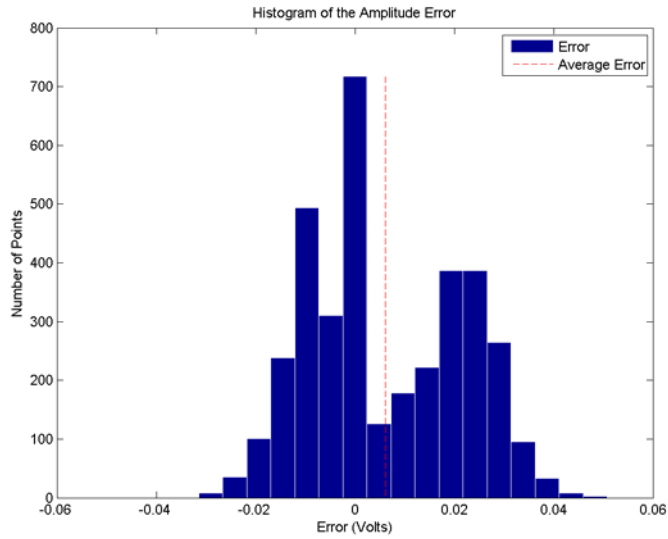
Figure V.26: Histogram of the amplitude error for the Bloch-Siegert pulse.
Most error values are located at or near zero, but the mean of the histogram (noted in red) is slightly off-centered to the right due to droop in supplies of the transmit system.
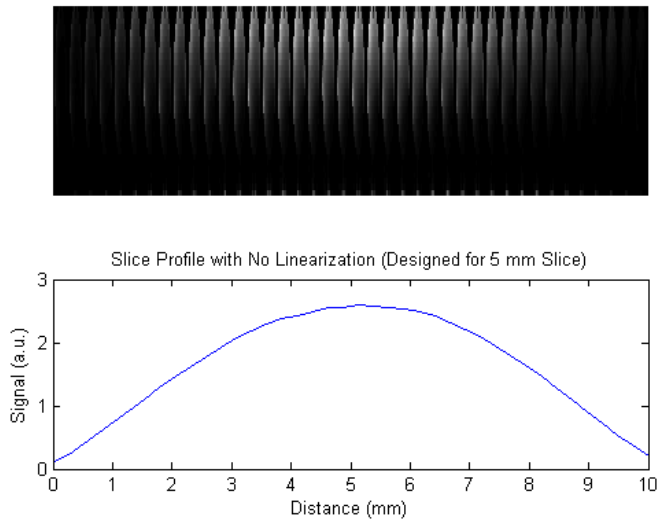
Figure V.27: Slice profile results with no linearity correction applied.
The series of images (top) are the slices orthogonal to the excitation slice used to create the plot (bottom). The plot shows a wide slice profile with significant excitation outside the slice, indicating poor linearity.
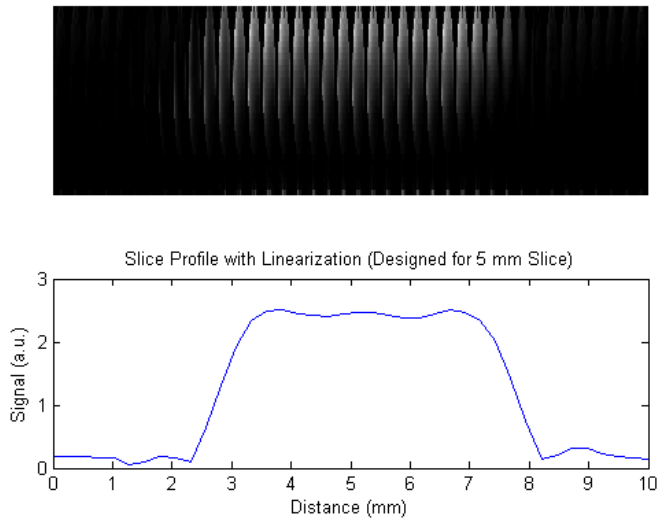
Figure V.28: Slice profile results obtained with linearity correction.
The series of images (top) are the slices orthogonal to the excitation slice used to create the plot (bottom). The plot shows a noticeable decrease in the signal level outside the slice, indicating improved linearity.

Since acquiring the slice profile data, the linearization routine has improved. The slice profile results are expected to improve as well but have not yet been reacquired.

V.4 Conclusion

The linearization routine effectively reduces the non-linearities of the system. If the system linearity proves inadequate for specific pulse designs, further improvements could be made by iterating the routine to further reduce the error, increasing the number of points in the look-up table to increase the resolution of the non-linearity characterization, or performing a pre-distortion directly on the user-defined pulse at the start of each scan. The maximum remaining amplitude error is 3.4% and the maximum remaining phase error is 2.8° according to the data acquired with the network analyzer.

The degree of amplitude and phase transmitter accuracy required for spatially selective pulses is not well-defined in the literature, but ultimately, the remaining question with the linearity asks whether the remaining error is sufficient. Commercial scanners operate under tighter error specifications but at greater costs and efforts. Other groups have reported vector modulators with phase errors of just above 1° and amplitude inaccuracy of 1% [73], but in their design, the board is not sufficiently fast to produce transmit-SENSE pulses. To answer the question, the system needs to be used exhaustively for multiple variations of complex pulse designs to determine the required linearity specifications. A cruder version of the linearization routine was used during acquisition of the parallel transmit results shown in Chapter IV, indicating that the linearity is sufficient for parallel transmit applications, which is the primary aim of the work.

CHAPTER VI

MONITORING FOR GROSS HARDWARE ANOMALIES

VI.1 Introduction

VI.1.1 Motivation for Monitoring

The conductivity of human tissue allows electric current to propagate, and the time-varying magnetic fields used for excitation in MRI induce currents in the tissue depositing power, quantified by the specific absorption rate equation (SAR):

$$SAR = \frac{\sigma E^2}{2\rho} \tag{[6.1]}$$

Where σ is defined as the conductivity in S/m, ρ is the density of the tissue in kg/m^3, and E is the peak value of the electric field in V/m. In neuroimaging applications, the extent of the heating is problematic due to the absence of pain receptors, which prevent the patient from noticing the heating effects. The International Electrotechnical Commission defines the standards for allowable SAR in the human head during an experiment with a global SAR limit of 3.2 W/kg and a local SAR limit of 10 W/kg or a 1°C temperature rise [74]. Global SAR is defined as the power transmitted through the RF coil divided by the mass of the exposed region whereas local SAR is calculated by equation 6.1 and is limited to any 10 grams of tissue. In the MR experiment, the local SAR is the limiting factor.

In parallel transmission, the spatially tailored RF pulses discussed in Chapter II are brought to practical durations due to the spatial degrees of freedom inherent in the parallel transmit system. Unfortunately, this benefit comes with a SAR penalty which

increases due to the shortened RF pulse length and corresponding increase in the RF voltage used for excitation, increasing SAR by the square of the applied voltage [75]. Multiple pulse design techniques in development add SAR as a constraint in the pulse design generation process in an effort to minimize SAR [76-79] or develop strategies to move the local hot-spots during the scan to distribute the local SAR load [80-83]. With the greater flexibility achieved with parallel transmit technology, SAR can be reduced as compared to a conventional single-channel transmit system.

There is no current standard for monitoring the superposition of E-fields in conductive tissue during an experiment. The current approach for mitigating SAR concerns is to use a full-wave simulation solution to compute the B and E fields associated with a given pulse and calculate global and local SAR values. Deviations in the pulse during a scan will change the SAR values, and minimum changes in amplitude and phase can cause large variations in the calculated SAR [84]. Since hardware anomalies can occur which effectively change the pulse amplitude and phase causing unexpected changes in SAR, monitoring the pulses during the scan is critical to ensure patient safety.

VI.1.2 Existing Monitoring Approaches

Multiple approaches exist for monitoring transmit pulses during a scan. Early monitoring approaches integrated a power monitoring unit to the transmit system that measured the power into a head coil over a period of time. If a predetermined power level were exceeded for the time interval, the system would shut down [85-87]. Unfortunately, this approach does not detect channel failures or deviations in the pulse

135

itself, and therefore, it does not adequately protect against local SAR increases [88-89].

Deviations in the amplitude and phase of the waveform on a channel will change the interference patterns of the produced magnetic and electric fields and change the SAR [90], and therefore, it is imperative to monitor the actual waveform rather than the average power. The most precise methods involve monitoring both amplitude and phase by sampling the pulse through pick-up-coils [89, 91] or directional couplers [92] but acquiring the phase data imposes additional complexities due to the need for direct digitization of the RF waveform. Commercial research approaches have used the scanner receiver chain to digitize the pulses during transmit [93], but for non-commercial research ventures, this approach adds complexity in integrating the monitoring subsystem with proprietary scanner software and hardware.

VI.1.3 Monitoring for Gross Anomalies

While an ideal approach would monitor both amplitude and phase, amplitude supervision is sufficient to detect gross hardware anomalies such as a channel failure or sudden change in the transmit system performance. It is important to note that this approach does not attempt to compute SAR in real-time, but instead, serves to ensure transmit system functionality and indirectly protects against changes in SAR.

The on-coil waveform envelopes were monitored in near real-time to detect channel failures or hardware anomalies, and if these errors are detected, the system will automatically shut down and alert the user. The shutdown consists of turning off the RF gate line to the amplifiers to prevent any additional power output from the transmit system and stopping the control system. The transmit chain was monitored at two stages:

the output of the driver amplifiers through directional couplers and the on-coil waveforms through pick-up-coils. The monitoring locations in relation to the parallel transmit system are shown in Figure VI.1.
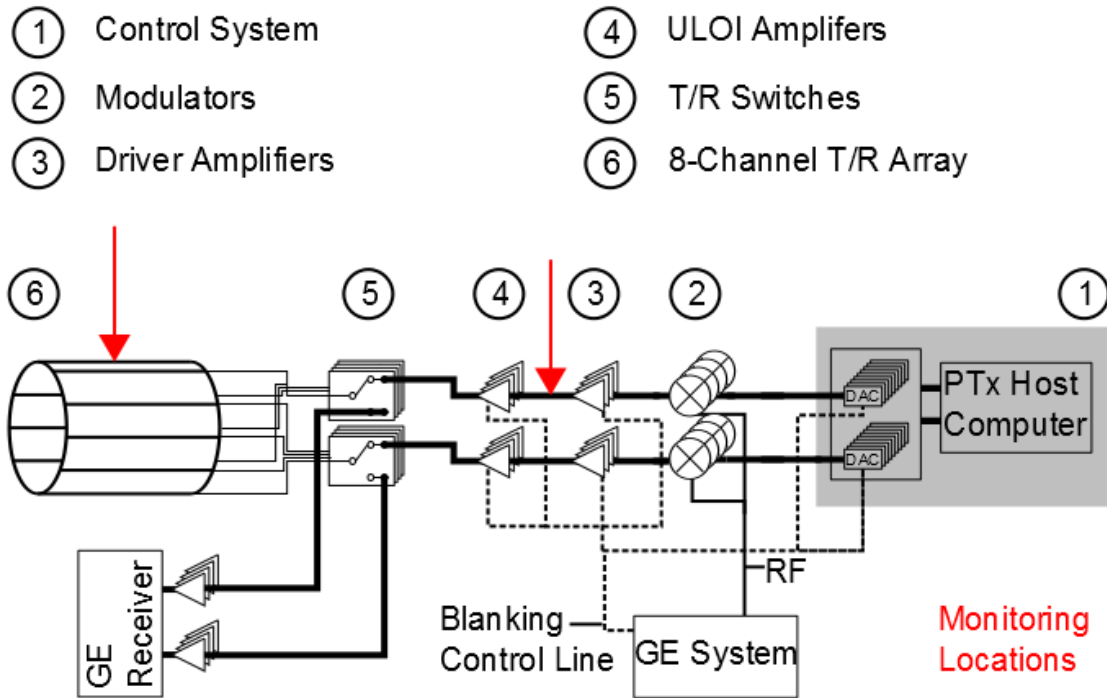
Figure VI.1: Block diagram of the parallel transmit system showing the monitoring locations after the driver amplifiers and on the coil.

VI.2 Materials and Methods

VI.2.1 System Level Checks to Prevent Operation Out-of-Specification

In addition to monitoring the transmit pulses, the monitoring software acts as a watch dog for user activity that causes the system to operate outside its specifications. The duty cycle and on-time of the RF pulse are monitored to prevent a duty cycle greater

than 10% during a 2 second time period, and the on-time is restricted to 100 milliseconds (a specification that if exceeded will damage the amplifiers).

For the transmit system to play out pulses, a digital RF gate line from the control system and the master exciter unblank (trigger) from the scanner must both be set to high. A simple AND gate in the vector modulator chassis accepts both signals and passes the resulting logic to the gate of the amplifiers. In the event where the pulse exceeds the set specifications, the control system immediately sets the RF gate line to low, and the amplifiers are no longer blanked even if the scanner trigger continues.

The 100 millisecond on-time limit is a more critical constraint than the duty cycle limit, and therefore, the on-time monitoring is all done on the target FPGA. A loop running in parallel to the waveform playback loop, measures the time between the rising and falling edges of the trigger line to compute the on-time. If the on-time exceeds the specification, the RF gate line from the control system is set low and the system shuts down.

The duty cycle limit is a less critical specification and is handled on both the host and target programs. The 10% duty cycle is the limit over a 2 second time period. The limit is adjustable on the monitoring menu of the user interface, but it defaults to 10%. To compute the duty cycle, the on-time is measured as described above and written to an indicator on the target program. It should be noted that the on-time monitors the trigger line and not the actual pulse to provide a conservative measure of the duty cycle (i.e. zeroes in the pulse occurring while the trigger is held high do not decrease the computed on-time). After the write operation, an interrupt is thrown to alert the host program that

138

the data is ready for a read operation. A loop on the host program waits for this interrupt from the target, and upon receipt, reads the on-time value and divides the value by the 2 second time period. The system checks to see if the calculated duty cycle exceeds the user specification, and if so, writes to a control on the target program to stop the system.

Another system level check involves the idle time of the control system measured on the target program. If the system does not receive a trigger after the specified system idle time has elapsed, the system shuts down and alerts the host program. The system also monitors for the power level on the coil. The power level cannot exceed the specified 300 Watts, so a threshold corresponding to a value of 300 Watts is written to a control on the target, and if during the waveform monitoring phase (described in detail below), a value is read that exceeds the limit, the system shuts down.

When a shutdown occurs, Boolean indicators on the host program light up to indicate which error occurred. The system level checks were verified for functionality by simulating the situation using the control and modulator subsystem (so as not to damage the amplifiers if the on-time were exceeded). A function generator was used to simulate the scanner trigger, and the on-time was exceeded by lengthening the duration of the trigger. When the error occurred, the appropriate indicator changed to alert the user to the on-time error. The duty cycle check was evaluated by setting up the function generator (pulse width of the trigger and frequency) to exceed the specified limit, and a shutdown occurred and the appropriate indicator changed. The system idle check was verified similarly by turning off the trigger and noting the system shutdown after the specified time with the corresponding Boolean indicator modified. The power level

139

check was evaluated by setting the threshold to a safe power level and observing a shut-down when it was exceeded.

VI.2.2 Monitoring Setup with the System

As mentioned previously, the waveforms were monitored at two stages in the parallel transmit system: the driver amplifiers and the array coil. Directional couplers were placed at the output of the driver amplifiers. The on-coil probes discussed in previous chapters (Chapter III and Chapter V) were used to sense the signal on the coil. A block diagram of the monitoring setup with the system is shown in Figure VI.2. Log amplifier power detectors (ZX47-60+, Mini-circuits, Brooklyn, NY) were used to extract the log of the envelope of the RF signal. Therefore, the phase information was discarded and only amplitude information was used to determine pulse accuracy. The envelope was digitized by the analog inputs on the PXI-7853R, and therefore, communication between the target and host programs facilitated data transfer to the host program for processing. On the host side, the data was converted to the actual envelope of the RF waveform and compared to the demand waveform or previously accepted waveforms for the current scan. Error between the actual and accepted waveforms was computed and a decision made based on the error threshold specification set by the user. If the error exceeded the specified threshold, the system shutdown by forcing the RF gate line from the control system low and turning off the control system.
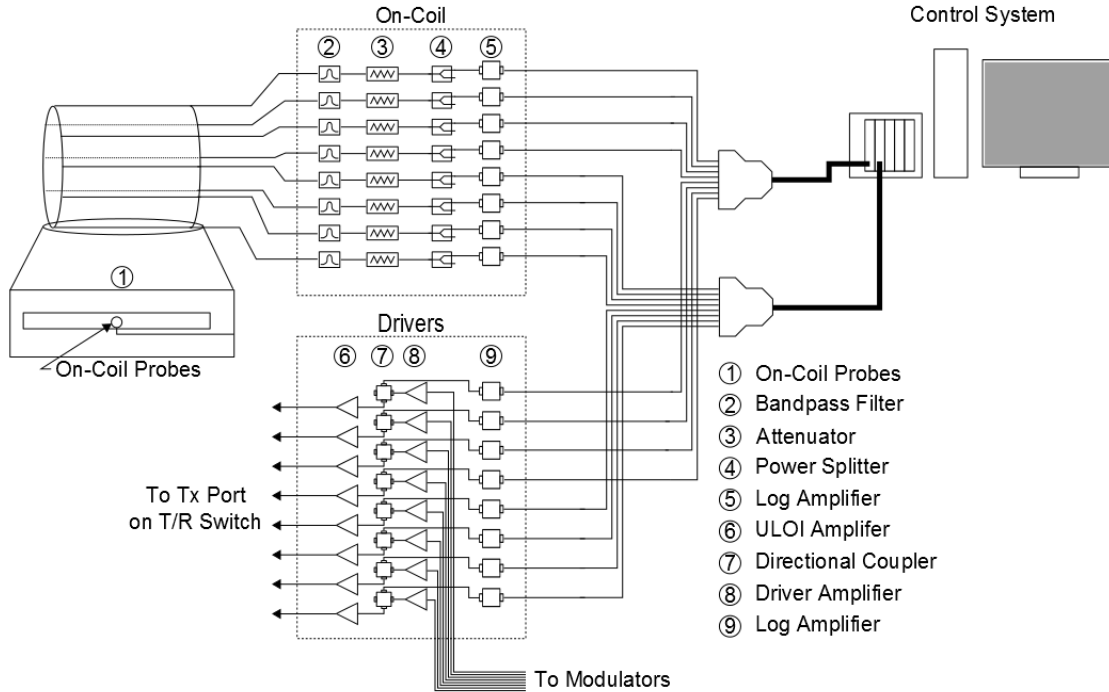
Figure VI.2: Block diagram of the monitoring setup.
Monitoring occurs at two locations: the on-coil probes at the array and the output of the directional couplers after the driver amplifiers. The on-coil probe signal are split in the switch hardware chassis as described in Chapter V prior to being input into a log amplifier. The coupled in port of the directional couplers is input into the log amplifier. The signals are bundled into cables and connected to the control system as shown.

VI.2.3 Monitoring Philosophy

The monitoring implementation was intentionally kept simple and straightforward to remove complexity. After verifying that the transmit system operated as expected, a series of triangle pulses with phase ramps were played out on each channel and the waveforms at both the driver amplifiers and coil were collected and stored as accepted waveforms. During acquisition, the waveforms are visually validated and stored as accepted waveforms.

141

In subsequent scans, a prescan option on the front panel forces the user to acquire waveforms using the same triangle pulse with a phase ramp to verify that the system has not changed substantially between the last known and validated "working" state. The user is forced to execute the prescan routine once a day. If the prescan is not current (determined by a timestamp stored in the software project directory), the user is not permitted to play out waveforms. The prescan compares the waveforms acquired to the previously accepted waveforms and computes the error between the two according to the following equation:

$$Percentage\ Error = \left(\frac{\sum_{i=1}^{N}(d_{accepted,i}-d_{acquired,i})^2}{\sum_{i=1}^{N}(d_{accepted,i})^2}\right)100 \qquad [6.1]$$

Where $d_{accepted}$ is the data of the accepted waveform and $d_{acquired}$ is the waveform acquired during the prescan. The resulting value is multiplied by 100 to obtain a percentage. The error threshold is specified by the user on the front panel, and if the error is exceeded, the prescan is not valid and the user alerted. The prescan saves the acquired waveforms enabling some remote diagnostics, but if the system deviates significantly from a previous operational state, it is indicative of a potential hardware failure and requires bench testing to further isolate the problem.

With a current validated prescan, the user is permitted to play out pulses and conduct experiments. During the scan, the waveforms at both stages are acquired during the first trigger. These waveforms are compared to the demand waveforms, and if there is a significant deviation between the demand and acquired waveforms as specified by a threshold (termed demand error threshold) set by the user, the system will shut down and

alert the user to the error. If the computed error is below the specification, the waveforms are saved as accepted waveforms for the current scan. Subsequent acquisitions of the waveforms are compared to these accepted waveforms. The error between the acquired and accepted waveforms is computed and validated to be within a secondary error specification (termed error threshold) defined by the user. If the error is outside the accepted limit, the system shuts down.

The error computation comparing the demand and actual waveforms is a more complex problem than comparing the acquired waveform to some previously accepted waveform due to the normalization required in the demand case. The demand pulses defined by the user range from amplitude values of 0-10 volts, but the pulses acquired through the system are on a smaller voltage range. Ultimately, during the first trigger, we are not concerned with the acquired waveform being at an improper power level (as long as the power level does not exceed the coil rating as discussed previously), and therefore, the problem was simplified by normalizing the acquired waveform to the demand waveform if the demand waveform was a non-zero pulse. Since the time step for the acquired waveform is the same as the time step specified in the file, the maximum value in the file is found, and the corresponding index in the acquired file is used to appropriately scale the waveform. This prevents noise spikes, which may occur at the low amplitude regions at the beginning or end of the pulse from interfering with the normalization (a hardware bandpass filter is placed on the RF line prior to being input into the log amplifiers to eliminate gradient noise). After normalization, a similar energy

error computation is used, but instead the energy of the error is referenced to the demand pulse with maximum energy as follows:

$$Percentage\ Error = \left(\frac{\sum_{i=1}^{N}(d_{demand,,i}-d_{acquired,i})^2}{\sum_{i=1}^{N}(d_{\max of\ demand,i})^2} \right) 100 \qquad [6.2]$$

Where d_{demand} is the demand waveform, $d_{acquired}$ is the acquired waveform, and $d_{\max of\ demand}$ is the maximum demand pulse. By referencing the error to the maximum demand pulse, errors occurring from a division by zero or large errors that may occur from referencing a low error to a near non-zero value are prevented.

For subsequent triggers, the error computation is straightforward and defined by the following equation:

$$Percentage\ Error = \left(\frac{\sum_{i=1}^{N}(d_{accepted,,i}-d_{acquired,i})^2}{\sum_{i=1}^{N}(d_{\max of\ accepted,i})^2} \right) 100 \qquad [6.3]$$

The error specifications were set experimentally but were guided by underlying principles. At the driver amplifier stage, the coupling between channels is minimal, and therefore, the monitoring waveforms are less affected by the currents on the other channels. With this in mind, the error tolerances at the driver amplifiers for both the first trigger (comparing the waveforms to the demand pulses) and subsequent triggers (comparing the waveforms to those acquired during the first trigger) can be set with tight margins of error. The waveforms sensed by the pick-up coils are sensitive to coupling effects, and therefore, the current on one element will induce a voltage (electromotive force or EMF) on all other elements with shared mutual impedance. As an example, consider two adjacent elements with some degree of coupling with one element transmitting a triangle pulse and the other a sinc pulse. Depending on the degree of

coupling between the channels, the triangle pulse will be superimposed on the sinc and vice versa. As discussed in Chapter III, array coil coupling was minimized but still exists and will affect the on-coil pulses. This prevents a tight error tolerance on the pick-up coils during the first trigger where the error is computed between the on-coil waveform and the demand waveform, but for subsequent triggers, the tolerance is tightened. A summary of the monitoring philosophy and implementation is shown in Figure VI.3.
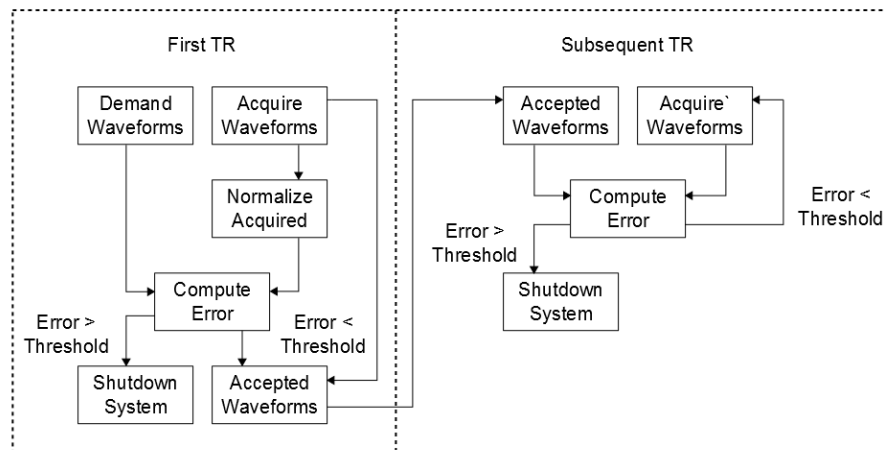
Figure VI.3: Block diagram of the monitoring setup during a scan.
During the first trigger (first repetition time), the acquired waveforms are compared against the demand waveforms. If the error between them is less than the specified threshold, the acquired waveforms are stored as the accepted waveforms for subsequent triggers. For subsequent triggers, the error between the accepted waveforms and acquired waveforms is computed. If the error is less than the threshold, the process repeats until the system is stopped or an error is detected.

VI.2.4 Waveform Extraction and Conversion

The log amplifier extracts the log of the envelope, and after digitization, the waveform is converted to the actual envelope prior to comparing to accepted or demand

waveforms and computing the error. The log amplifier has a transfer function of the following form:

$$V_{Out} = V_Y log_{10}\left(\frac{V_{In}}{V_X}\right)$$ [6.4]

Where V_{Out} is the output voltage in volts-rms, V_Y is the slope voltage expressed in volts-rms/decade, V_{In} is the input voltage in volts-rms, and V_X is the intercept rms voltage defined as the value where V_{Out} is zero. In converting the waveform to the input voltage, it is more beneficial to express the equation as follows:

$$V_{Out} = V_{Slope}(P_{In} - P_0)$$ [6.5]

Where V_{Slope} is the slope voltage expressed in volts/dBm, P_{In} is the input power in dBm, and P_0 is the intercept expressed in dBm. The logarithmic amplifier used has a slope voltage specification of -25 mV/dBm. Rearranging the equation, multiplying V_{Slope} by P_0 to obtain V_0 (the logarithmic intercept), converting the powers in dBm to watts, and extracting the input voltage, we obtain the following equation for the input waveform:

$$V_{In} = \sqrt{\left(10^{\frac{1}{10}\left(\frac{V_{Out}-V_0}{V_{Slope}}\right)}\right)(R)}$$ [6.6]

According to the datasheet, V_0 is approximately to 0.65 volts at 128 MHz.

The conversion was performed, and the error computed between the demand and converted waveforms for several pulses to validate the operation. The conversion is not perfect but represents a sufficiently accurate approximation for detecting gross changes in the transmit system functionality. The slope and intercept controls are available on the

monitoring front panel to enable future adjustments in the conversion if needed. The

monitoring front panel during normal operation is shown in Figure VI.4.
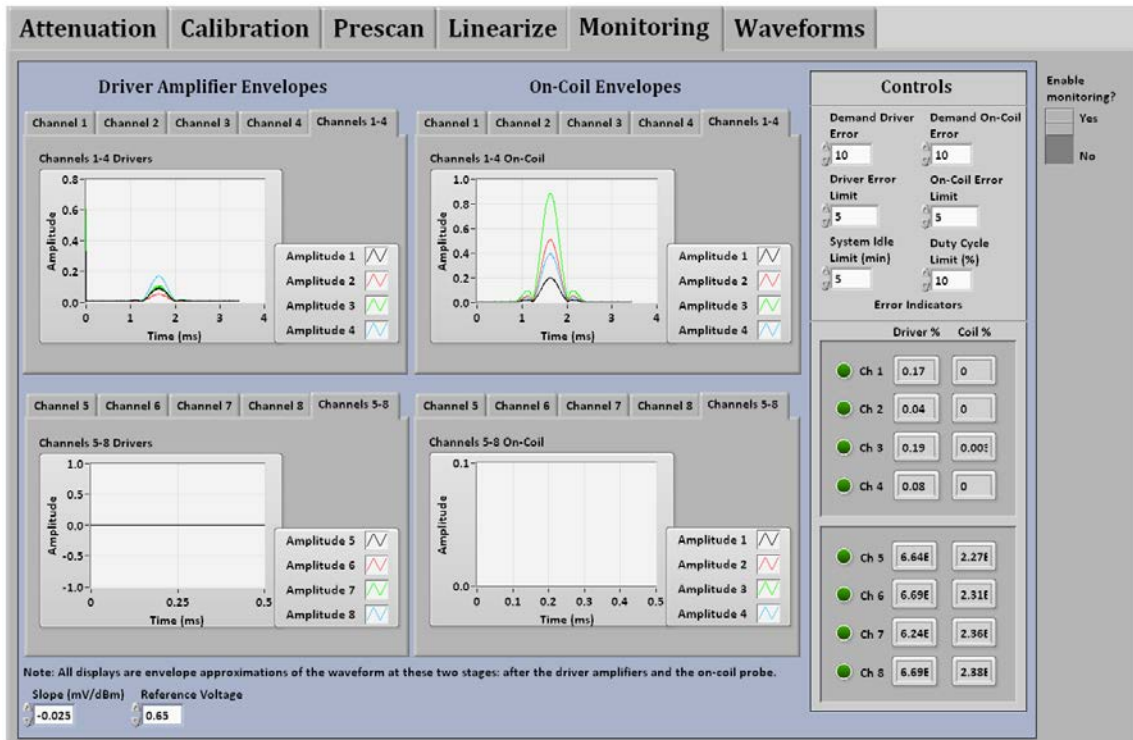
Figure VI.4: Monitoring front panel during normal operation.
Four channels for a uniform excitation are being played out and sensed at both the driver
and on-coil stages.

VI.2.5 Prescan Operation

VI.2.5.1 Acquiring the Accepted Waveforms for the Prescan

The accepted waveforms for the prescan operation are acquired immediately

after validating the transmit system functionality. A series of triangle pulses with a phase

ramp (the same pulse used for linearization) are played out one channel at a time, and the

waveforms on all channels are acquired. The waveforms are visually verified during the

acquisition to make sure a channel failure of some other anomaly does not occur, and the waveforms are manually accepted on the prescan front panel. These waveforms are saved as the accepted waveforms.

VI.2.5.2 Prescan to Validate Transmit System Functionality

The primary purpose of the prescan is to ensure that the transmit system operates as well as at some previously accepted state. The user is required to perform a prescan every day otherwise the control system will not allow waveform playback. The prescan repeats the experiment described above, but instead, the acquired waveforms are compared to the previously accepted waveforms, and the error between the waveforms is computed. In the event a channel fails the prescan test, the user is alerted. A channel can be disabled to allow the system to continue to operate the system despite a failed prescan on a single channel by changing the state of a Boolean disable control on the front panel.

The prescan saves the acquired waveforms. The waveforms are compared to the demand waveform as well in order to check for deviations in linearity. Small deviations are not cause to stop the scan, but the information is stored in a file if the information is needed later. An error threshold is available on the prescan panel to throw an error in the event the linearity varies outside the set threshold. Additionally, gross changes in linearity will be detected in the comparison to the previously accepted waveforms. The coupling information is computed and saved to a file for later access if needed. By pulling these files periodically, the drift in linearity and coupling can be monitored.

To validate the prescan functionality, a four-channel bench test was setup using four channels of the vector modulators, a subset of channels from the 64-channel

transmit system [67], and a four-channel transmit-only array with shielded probes. The waveforms were acquired, and the report containing the linearity and coupling information pulled for functionality verification. The prescan front panel is shown in Figure VI.5.
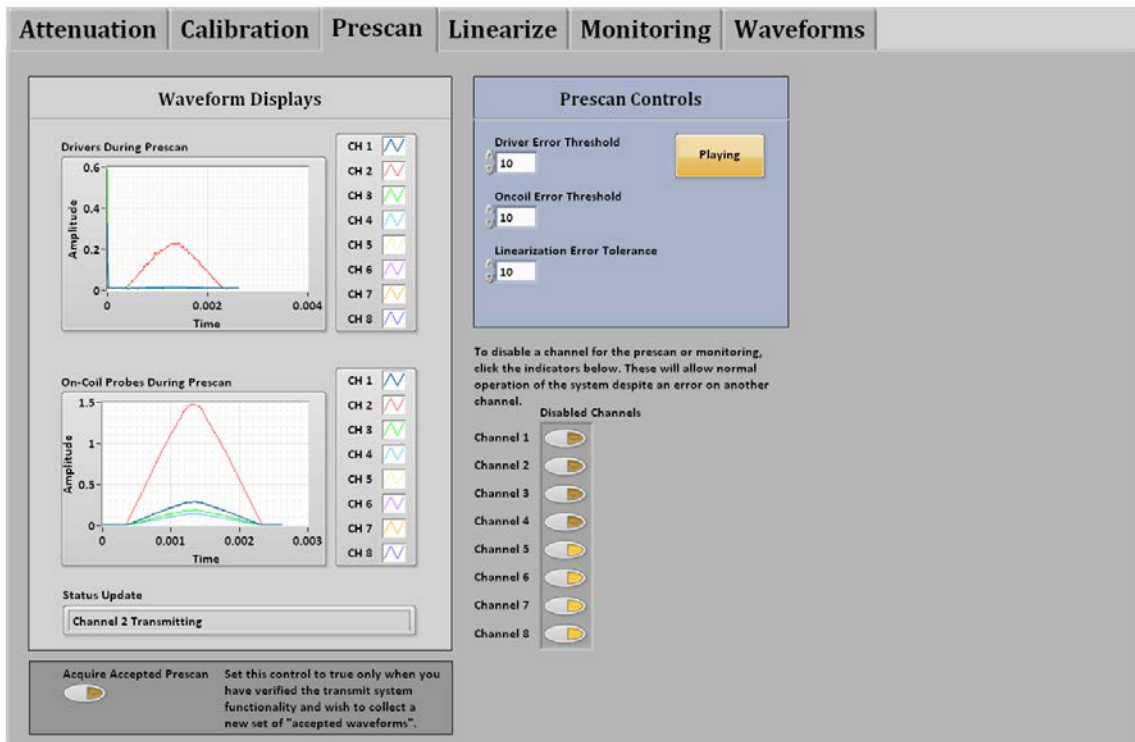
Figure VI.5: Prescan front panel.
The waveforms from all channels are displayed during the scan. For this demonstration of the prescan, channels 5-8 were disabled because the system was being used with a four-channel test setup.

VI.2.6 Monitoring During an Experiment

VI.2.6.1 Communication Between the Target and Host Programs

The available analog inputs on the PXI-7853R are used to acquire the monitoring signals. Each PXI-7853R enables four channels of monitoring at both the driver amplifiers and coil stages (eight analog inputs). A primary consideration when developing the monitoring software was to ensure that the monitoring functionality did not affect waveform playback, so all the monitoring software on the target program runs in parallel to waveform playback and communicates between the host and target using direct memory access (DMA) buffers.

The monitoring section of the target program waits for a rising edge trigger, and when the trigger is detected, data is acquired at a rate of 150 kS/s. The rate is controlled by using a sequence structure within a for loop that acquires data and waits the specified amount before allowing the loop to iterate. The number of iterations is determined by the sampling rate and size of the demand pulse. A single loop collects the data from all eight channels, and packs the data into a target-to-host DMA FIFO. The data for all channels is interleaved and packed into the buffer at each acquisition point. After acquiring the specified number of points, an interrupt is thrown to alert the host program that the buffer is filled.

The host program waits for the interrupt, and upon receipt, reads the data from the buffer, unpacks the array, and converts the data to the input waveform. A queuing and producer/consumer software architecture was setup to handle the monitoring operation and pass data from inside and outside a while loop. The while loop is in place

from the waveform playback implementation. The loop iterates, allowing the target program to continue to run, and waits until the user tells the control system to stop. Since the results of the monitoring error computations effect the decision to allow the target program to continue to function in addition to the "Quit" control on the front panel controlled by the user, the code to compute the errors is nested inside the same while loop. The converted monitoring data is packed into the queue, and an occurrence is thrown (a host-side only interrupt) to alert the code running inside the loop that data is available for processing.

When the occurrence is detected, the monitoring code runs. Otherwise, the function waiting on the occurrence times out, and the loop iterates. The data from the queue is unpacked and used to update the monitoring waveform displays. Error calculations are computed and the process repeats.

VI.2.7 Graphical User Interface and Alerting the User to Errors

There are three panels used for the monitoring software: the prescan panel, the monitoring panel, and the waveforms front panel. The waveforms panel is the primary menu used during the scan and contains the error indicators. When an error is detected, the system shuts down, and the appropriate error indicator on the front panel lights up (the error indicators mimic LEDs). The waveforms menu with the error indicators is shown in Figure VI.6.
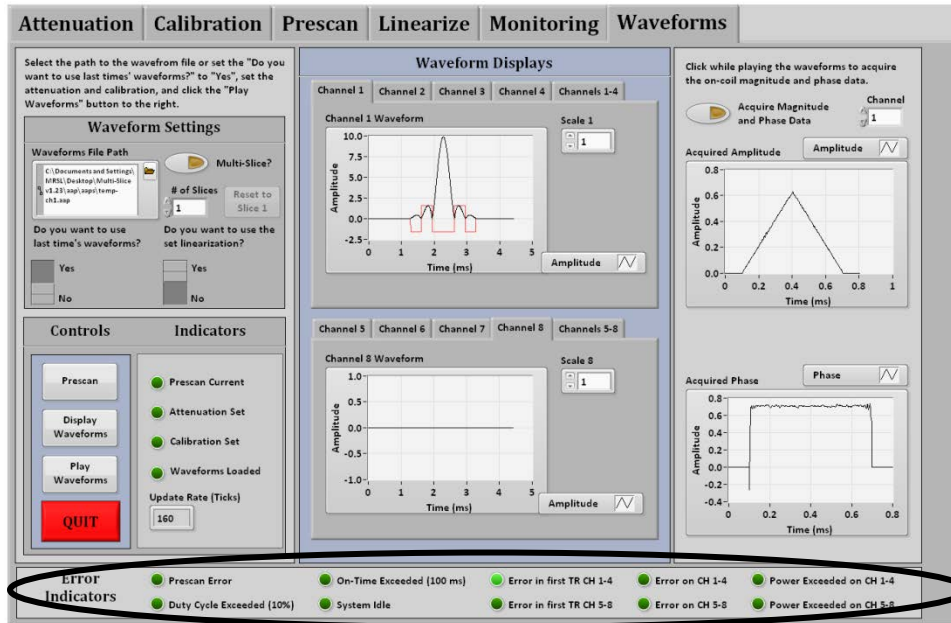
Figure VI.6: Waveforms panel with the error indicators circled at the bottom.
The error indicators mimic LEDs and light up when an error has occurred. Here an error
occurred during the first TR, the system shut down, and the indicator lit up to alert the
user to the error.

VI.3 Results and Discussion

VI.3.1 Converted Waveform Error

The log amplifier data after conversion was acquired in the control system, and

the RF waveform going into the log amplifier was acquired using the undersampling

technique to extract the magnitude described in Chapter V. The waveforms were

normalized, and the root-mean-square error computed between the two waveforms. An

external scope was not used due to memory depth limitations with longer pulses. The

results with the computed error are shown for a series of waveforms below in Figure

VI.7, Figure VI.8, Figure VI.9, Figure VI.10, and Figure VI.11.
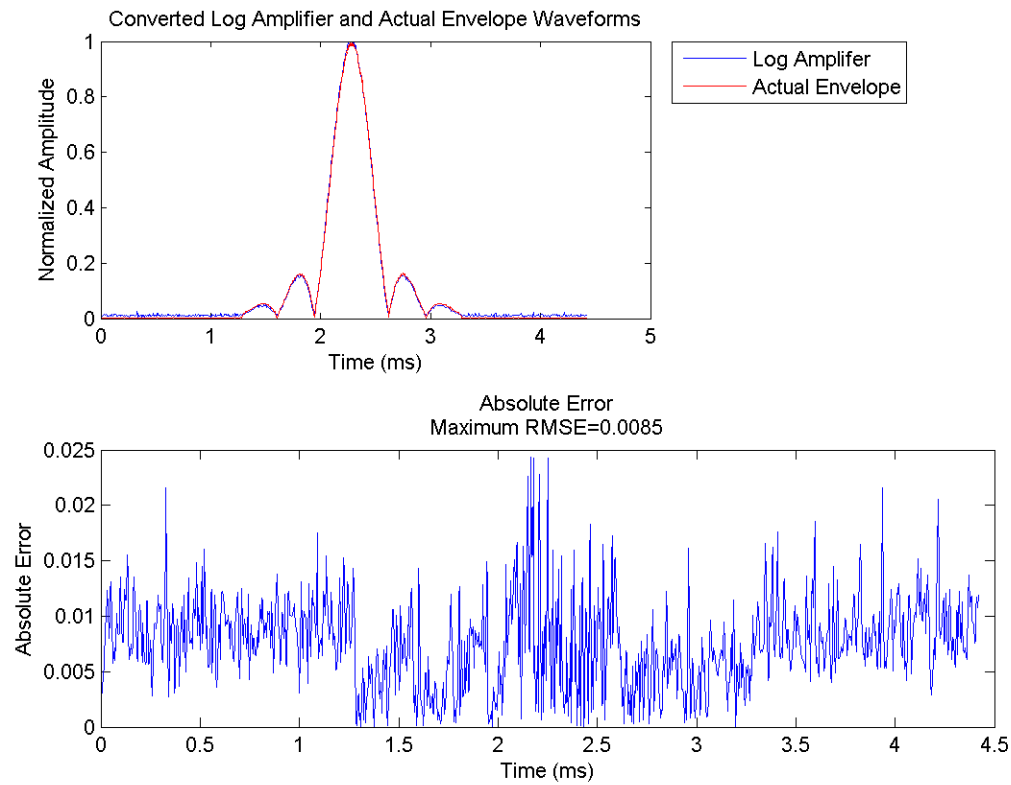
Figure VI.7: Converted log amplifier envelope and actual envelope for a sinc pulse.
Computed root-mean-square error is 0.0085. The two normalized waveforms are plotted
on the top, and the absolute error is plotted on the bottom.
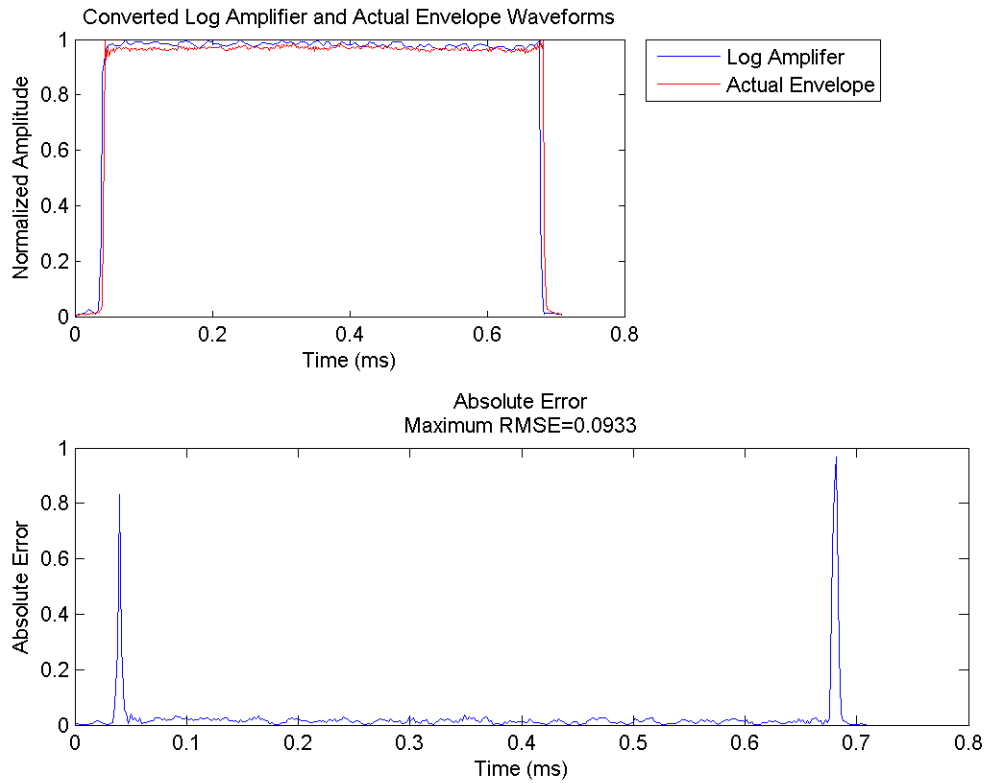
Figure VI.8: Converted log amplifier envelope and actual envelope for a hard pulse.
Computed root-mean-square error is 0.0935. The two normalized waveforms are plotted on the top, and the absolute error is plotted on the bottom.
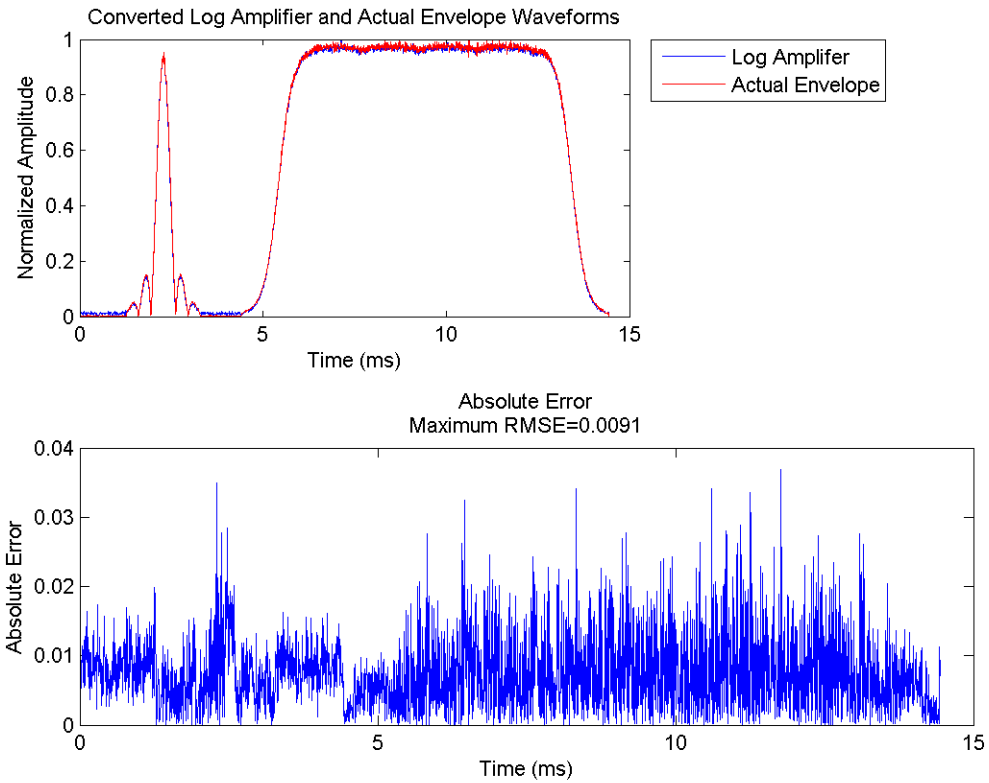
Figure VI.9: Converted log amplifier envelope and actual envelope for a Bloch-Siegert pulse.
Computed root-mean-square error is 0.0091. The two normalized waveforms are plotted on the top, and the absolute error is plotted on the bottom.
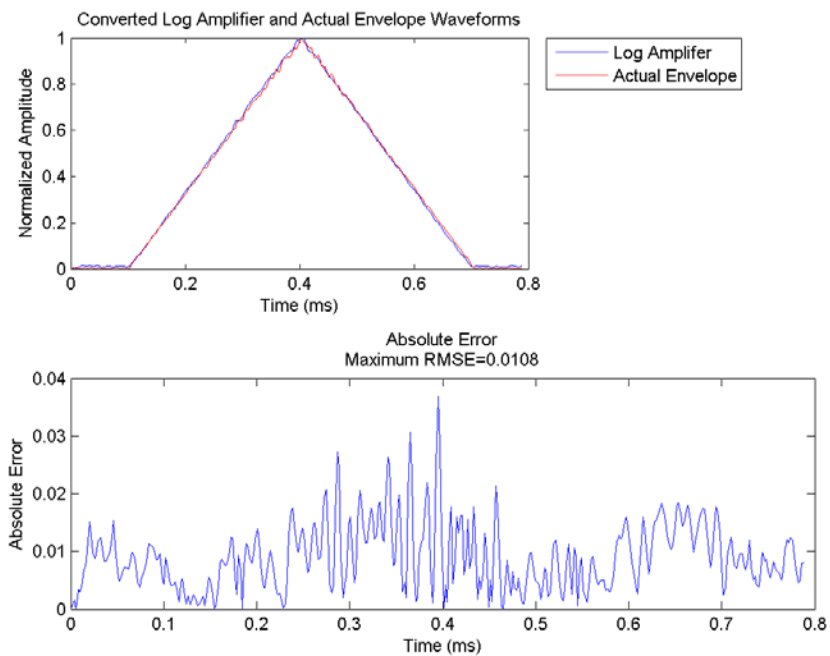
Figure VI.10: Converted log amplifier envelope and actual envelope for a triangle pulse.
Computed root-mean-square error is 0.0108. The two normalized waveforms are plotted on the top, and the absolute error is plotted on the bottom.

156
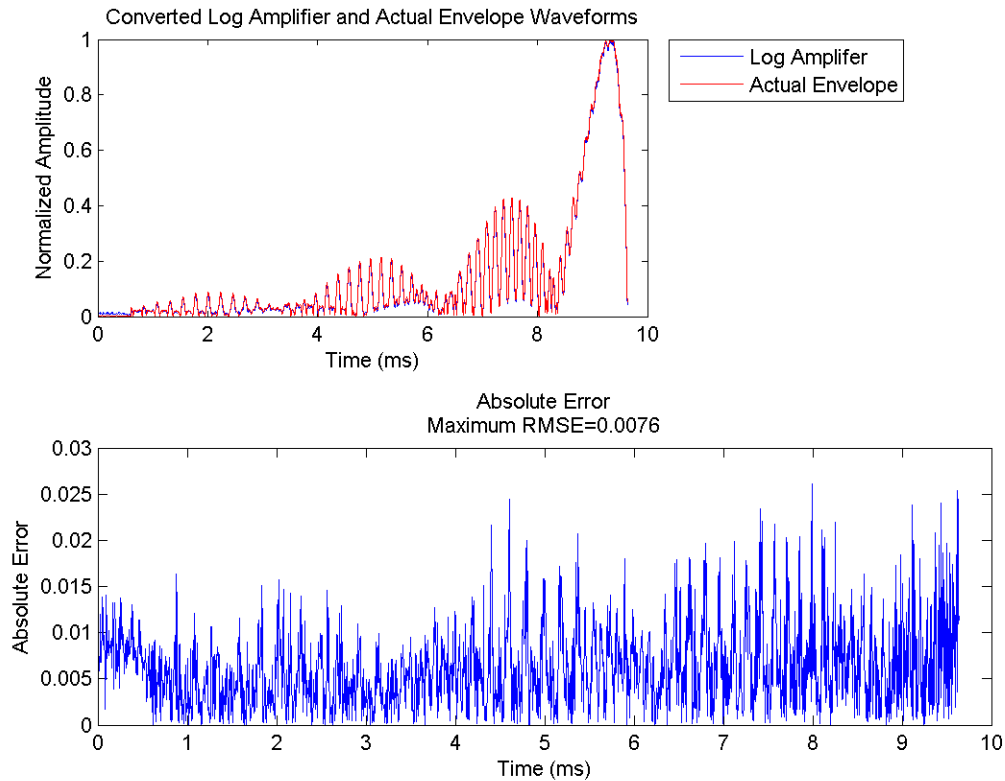
Figure VI.11: Converted log amplifier envelope and actual envelope for a echo-planar imaging pulse.
Computed root-mean-square error is 0.0076. The two normalized waveforms are plotted on the top, and the absolute error is plotted on the bottom.

The waveform is normalized during the first TR to the demand pulse, and during subsequent TRs, the waveforms are not normalized and simply compared to the previously acquired waveform. The log amplifier conversion was verified to scale with the actual amplitude envelope

VI.3.2 Error Computations for Different Scenarios

Several failures were safely simulated to determine if the system properly shutdown when specified error thresholds were exceeded. For demonstration, three error

cases are shown here. For the demonstration, the demand pulses shown in Figure VI.12 were played out on channels 1-4.
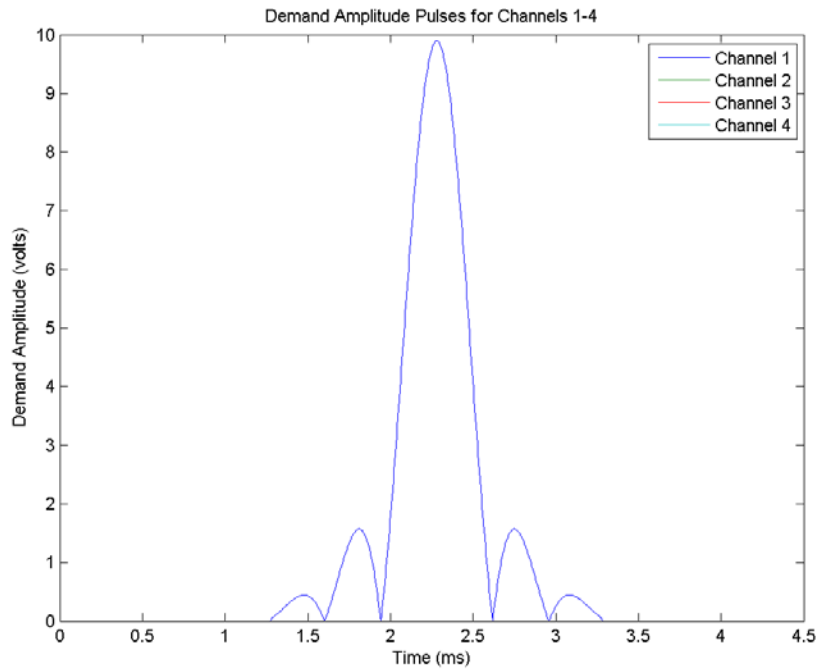
Figure VI.12: Demand pulses for channels 1-4.
A sinc is played out on channel 1 and all other channels are playing out zeroes.

As a reference, the error during the operational case when the waveforms are accurately being played out is shown in Figure VI.13. The error thresholds were set to 20% for both the driver and on-coil stages during the first TR (demand error threshold) and 10% for both the driver and on-coil stages during subsequent TRs. In the following screen captures of the software system, channel 1 has a lower amplitude than channel 2 because the current probe cable on channel 1 was passed through a power splitter in order to monitor the waveform on both an external scope and the software system. The

external scope was used to visually check to ensure the system shutdown correctly. Additionally, it is evident that there is substantially higher coupling at the on-coil stage than the driver amplifier stage as indicated by the sinc pulse being sensed on all probes to some degree. The computed error on all channels was less than 1% during the case with no errors.
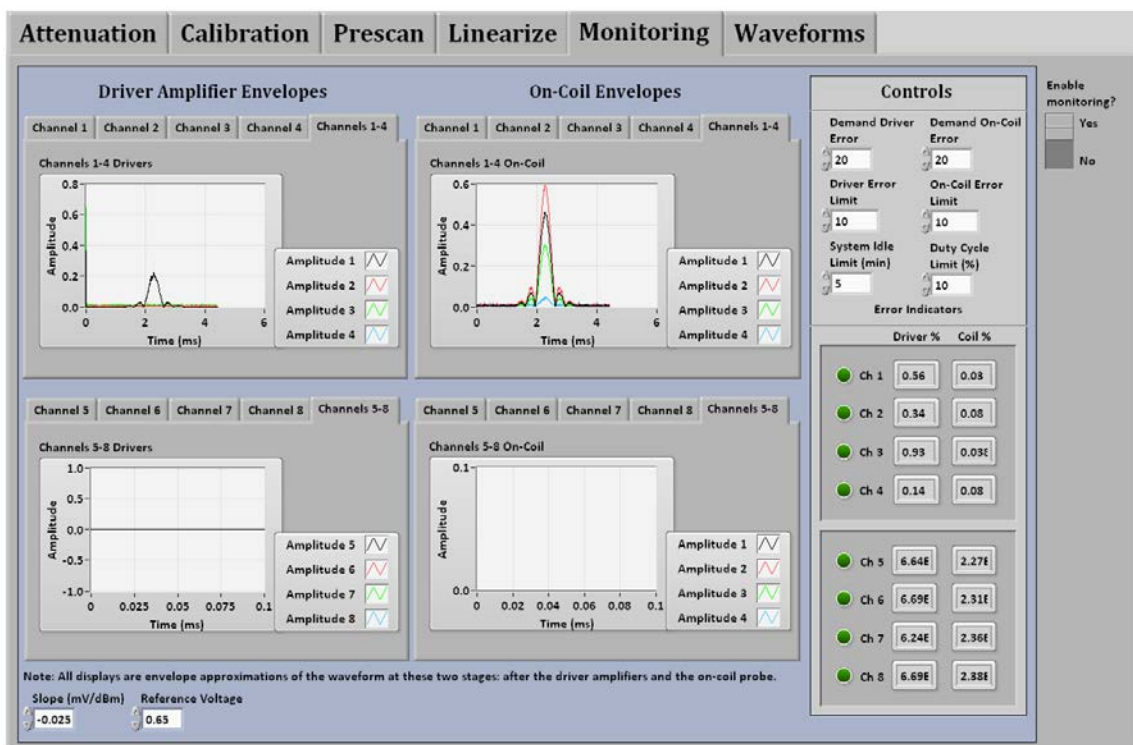
Figure VI.13: Monitoring front panel with no errors occurring.

The first error case simulated a channel failure on the channel being excited. The error was simulated by unplugging the current probe cable from the monitoring hardware (the connection through the splitter and to the external scope was still connected to verify shutdown). The resulting error was 64% on channel 1, and the

159

system shut down. The resulting screen capture with the error displayed is shown in Figure VI.14.
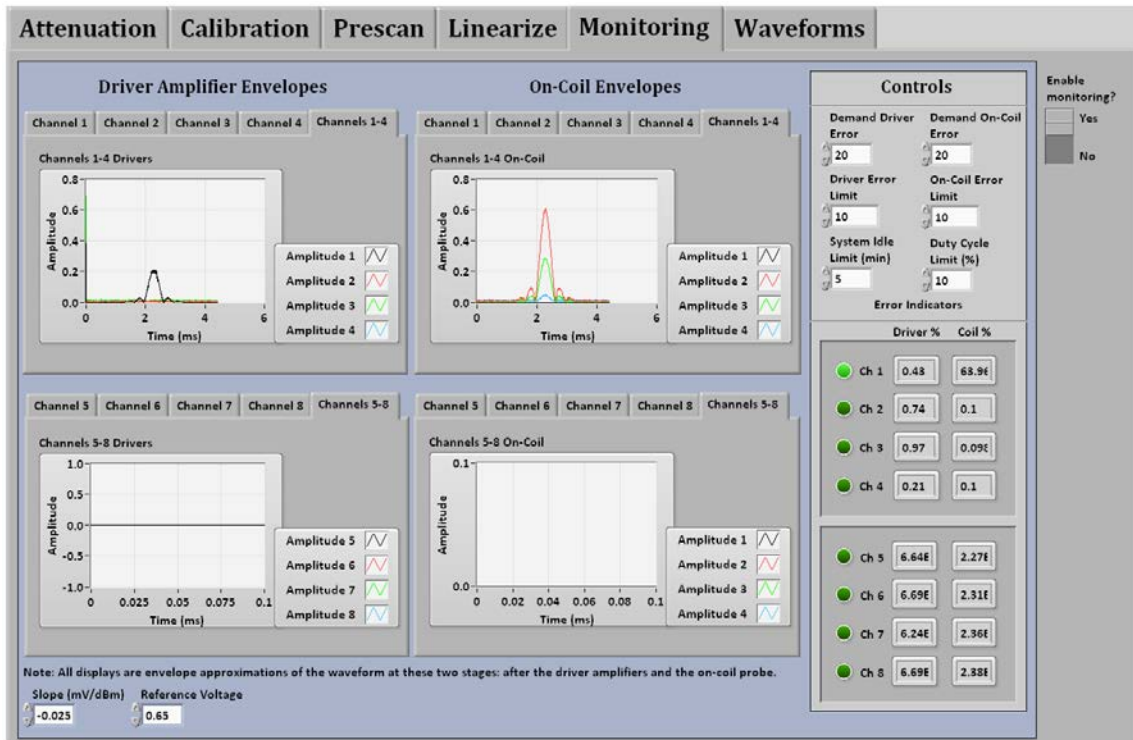
Figure VI.14: Monitoring front panel with a channel failure simulated on channel 1 by unplugging the cable from the monitoring hardware.
The error increased from less than 1% to 64%, and the system shut down.

The next error case simulated was a channel failure on a channel that was not being excited, and in this case, channel 3 was chosen. The error was simulated by unplugging the on-coil probe cable from the monitoring hardware, and the resulting computed error was 25%, and the system shut down. If coupling were minimal and a channel failure occurred, an error may not be detected if the sensed waveform was close

to zero. However, this case demonstrates that even if the demand waveform is zero, any

substantial deviation from one trigger to the next will cause a system shutdown.
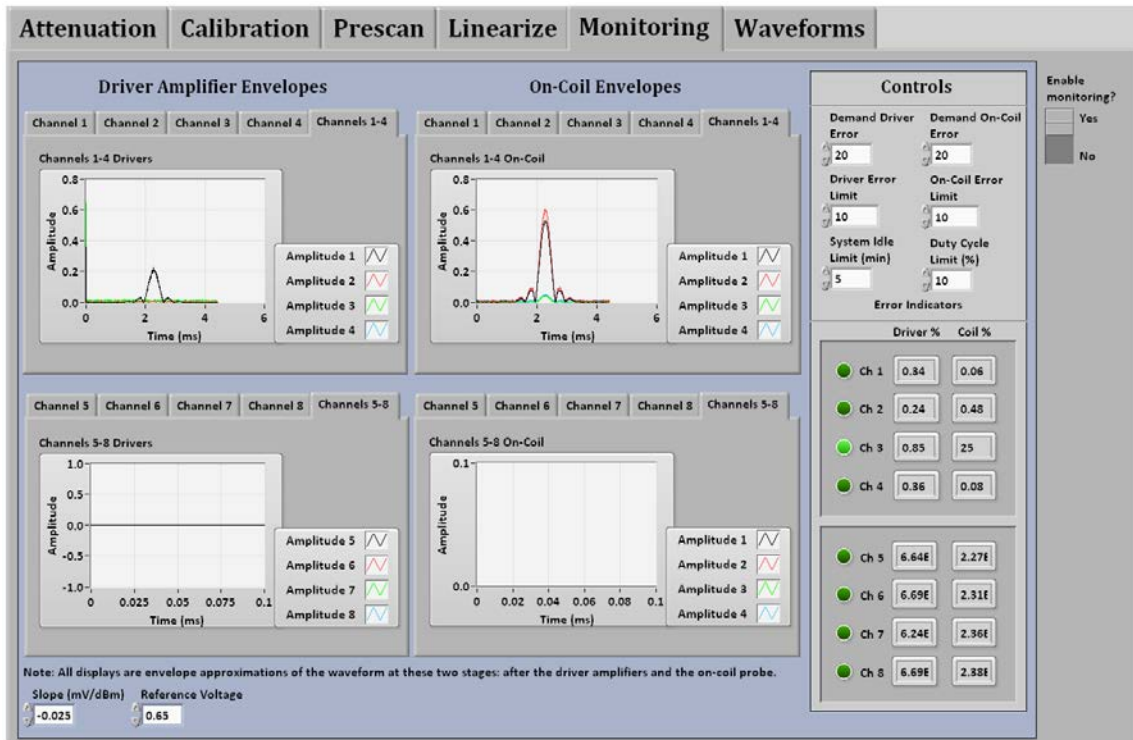
Figure VI.15: Monitoring front panel with a channel failure simulated on channel 3 by unplugging the on-coil probe cable from the monitoring hardware.
The error increased from less than 1% to 25%, and the system shut down.

The next error case simulated was a coil failure, which was accomplished by

turning off the forward bias to the coil. The coil used is a transmit-only coil and contains

PIN diodes for active detuning. During the transmit case, the diodes are forward biased,

and the coil is resonant. In the receive case, the diodes are reverse biased, and the coil is

detuned. By turning off the bias voltage, a mismatch occurs between the amplifier and

coil such that less power is transferred to the coil, and as a result, the on-coil probe

senses a reduced signal. The power going into the coil was verified to be low so as not to damage the amplifiers. When the bias was turned off, an error was detected on channels 1-3. The errors on channel 1-3 increased from less than 1% to 75%, 95% and 24%, respectively. An error was not detected on channel 4 because the signal on channel 4 was near zero initially.
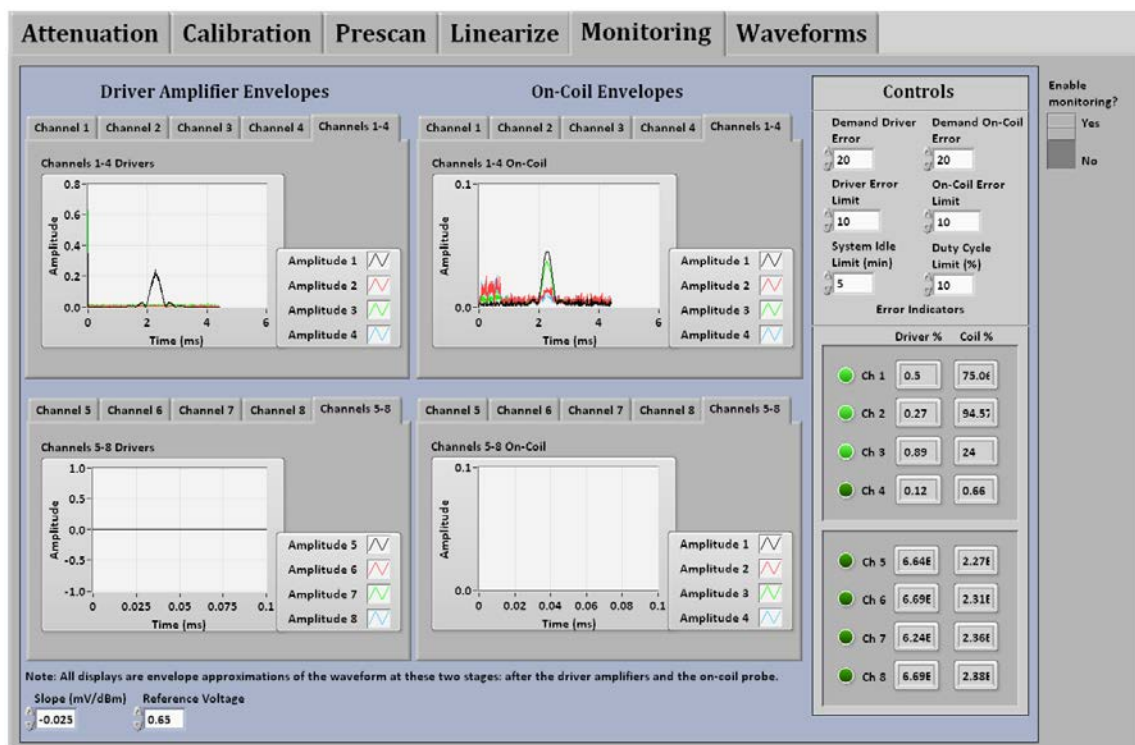
Figure VI.16: Monitoring front panel displaying the result from a coil failure accomplished by turning off the forward bias to the coil.
Errors were detected on channels 1-3, and the errors increased from less than 1% on all channels to 75%, 95%, and 24%, respectively.

As a final demonstration, the error detection at the driver stage was verified by unplugging the coupled in port from the directional coupler. The error increased from

less than 1% to 99%, and the system shut down. In the front panel, the waveform sensed by the on-coil probe on channel 1 is still present because the failure at the driver stage was only simulated by unplugging a cable. The actual transmit chain path to the coil was unaffected.
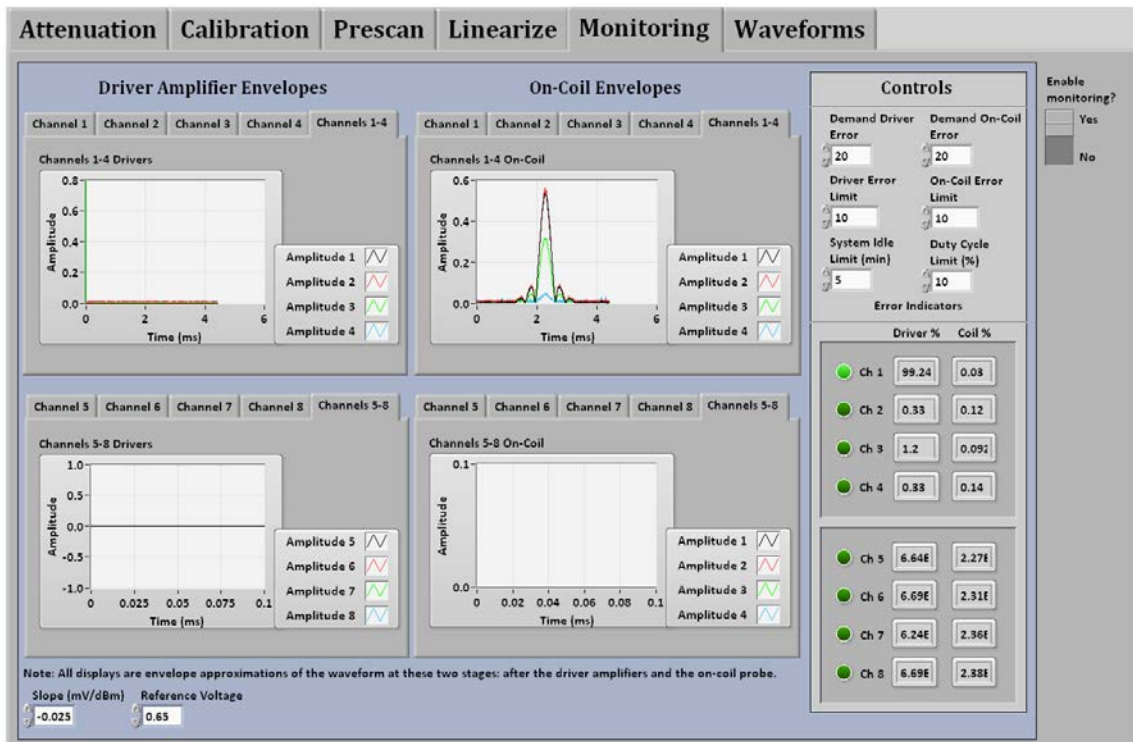
Figure VI.17: Monitoring front panel displaying the result from a simulated driver amplifier failure accomplished by unplugging the coupled in port from the monitoring hardware.
The error at the driver stage on channel 1 increased from less than 1% to 99%.

Multiple error cases were simulated for different pulses and scenarios and the software behaved as expected. In a few cases the effectiveness of the software prevented accidental improper operation of the system. The power-on sequence for the test setup

consisted of first switching the lever on the wall to power the circuit, turning on the modulators for the eight-channel system, powering on the modulators in the 64-channel system, powering on the amplifiers, and then, enabling the gate bias line to the amplifiers (an additional protection feature for the 64-channel system). In some instances, flipping the switch to enable the gate bias line was forgotten, and the system shutdown and indicated that an error had occurred during the first TR.

VI.3.3 Future Work

The monitoring software was tested and verified for functionality on a 4-channel test setup on the bench to verify all channels, but it has yet to be fully implemented on the 3 Tesla clinical research scanner as designed. The software installation and testing on the final system will be supported remotely to ensure expected operation.

VI.4 Conclusion

The monitoring software represents a straightforward approach to detecting gross hardware anomalies. The system effectively shuts down when an error exceeds the threshold and considerations were made such that the error computation can be easily understood with percentages, allowing the error thresholds to be set easily.

CHAPTER VII

FINAL RESULTS AND CONCLUSIONS

This chapter progresses with an initial experiment comparing the stability of the parallel transmit system and the GE system with a discussion on the importance of stability as it relates to functional MRI studies. The chapter concludes with images obtained by our collaborators and a discussion of future work. With the system installed, our collaborators used the system to demonstrate the signal recovery application in the prefrontal cortex discussed in Chapter II. The images and simulations shown are their results and demonstrate the capability enabled by the parallel transmit system discussed in this work.

VII.1 System Stability for Future Functional MRI Studies

Since fMRI experiments rely on small differences between an activated and unactivated state (refer to Chapter II), it is a critical system parameter to ensure the system instabilities are well controlled. To avoid correlating the stability measurements directly to the fMRI experiment, the parallel transmit system was compared to the commercial GE system, which routinely performs fMRI experiments. If the parallel transmit system instabilities were comparable to the GE system instabilities, the parallel transmit system would provide enough stability to accurately perform fMRI experiments.

An experiment was designed to acquire the same line of k-space 128 times by turning the phase encoding gradient off. Each line was compared to a reference line

chosen to be line 1 (the first line acquired in the experiment), and the repetition time between acquisitions was changed from a TR of 200 milliseconds to 2000 milliseconds to determine if there were any differences in the short term and long term stability of the system. When the parallel transmit system was used, the pulses played out were those corresponding to a uniform excitation. When the GE system was used, the birdcage coil was driven to excite a uniform excitation, and the GE 8-channel receive-only coil was used during reception. The results for the short-term stability test are shown in Figure VII.1 and Figure VII.2.

The excitation levels for the two systems are different as expected, but the percent deviation between each acquisition is the critical parameter, so the percent error was calculated for the magnitude data. The phase data does not require a reference, and therefore, the phase was reported as the maximum absolute error.

The results for the long-term stability test are shown in Figure VII.3 and Figure VII.4.

Figure VII.1: Stability results from the parallel transmit system using TR=200 milliseconds.
The resulting image from the experiment is displayed at the top. The magnitude absolute error between a reference line (line 1) and all other lines was computed and is displayed in the bottom left image. The maximum percent error was 2.03%. The phase error along the region with maximum signal is plotted on the bottom right, and the maximum absolute error was found to be 2.04 degrees.
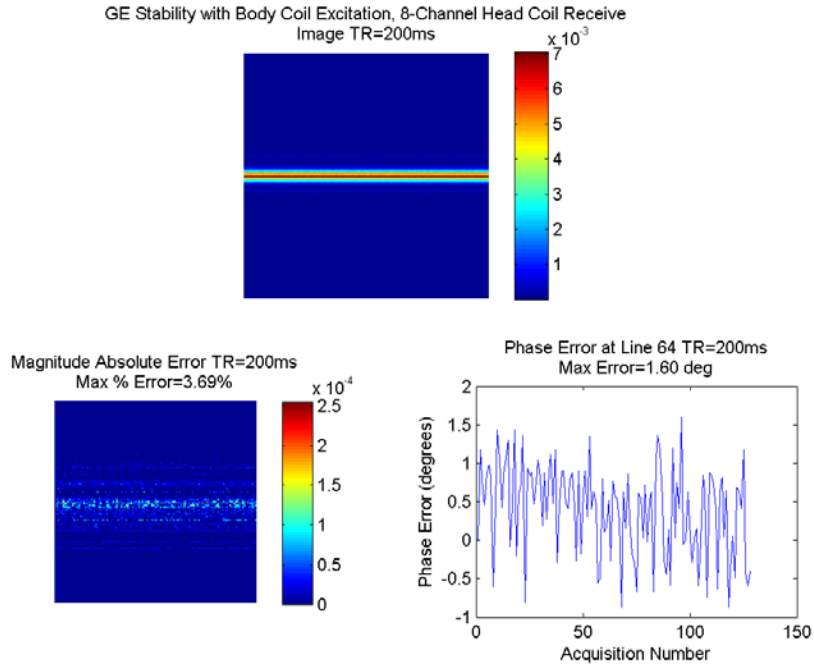
Figure VII.2: Stability results from the GE system using TR=200 milliseconds.
The resulting image from the experiment is displayed at the top. The absolute magnitude error between a reference line (line 1) and all others lines was computed and is displayed in the bottom left image. The maximum percent error was 3.69%. The phase error along the region with maximum signal is plotted on the bottom right, and the maximum absolute phase error was found to be 1.60 degrees.
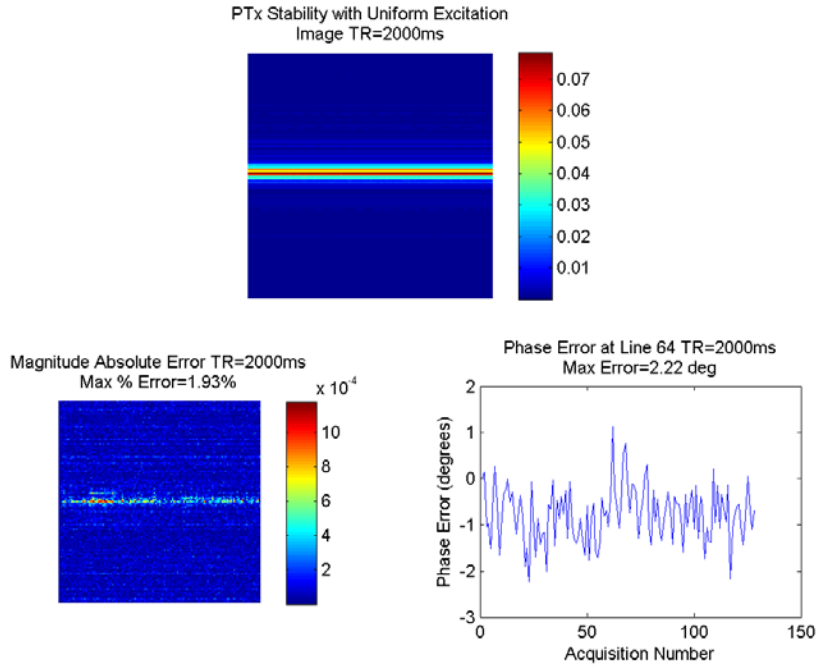
Figure VII.3: Stability results from the parallel transmit system using TR=2000 milliseconds.
The resulting image from the experiment is displayed at the top. The magnitude absolute error between a reference line (line 1) and all other lines was computed and is displayed in the bottom left image. The maximum percent error was 1.93%. The phase error along the region with maximum signal is plotted on the bottom right, and the maximum absolute error was found to be 2.22 degrees.
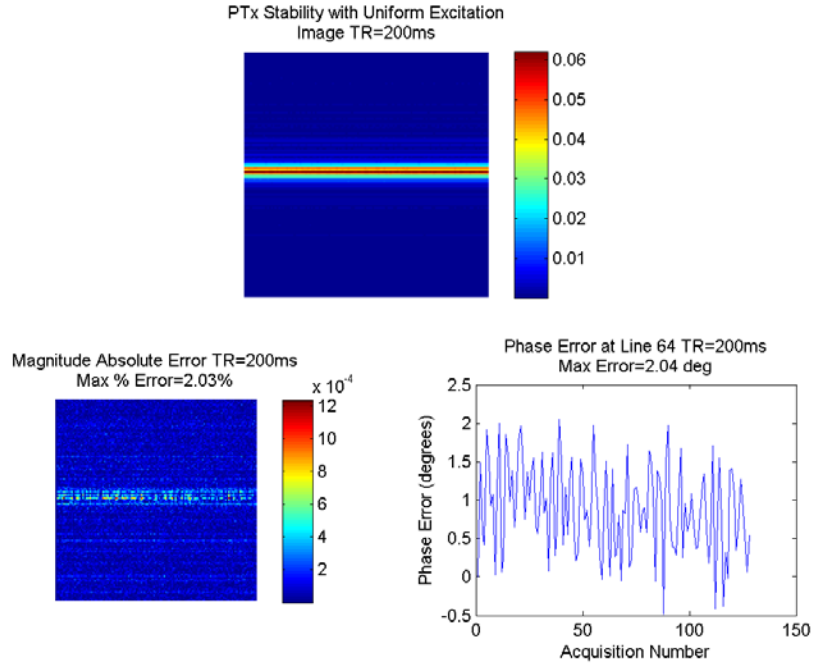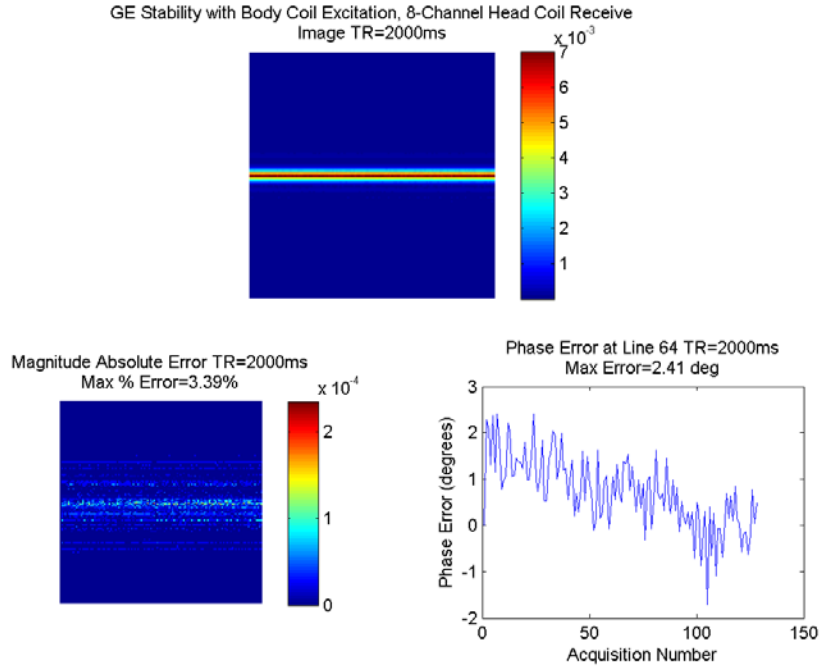
Figure VII.4: Stability results from the GE system using TR=2000 milliseconds.
The resulting image from the experiment is displayed at the top. The magnitude absolute error between a reference line (line 1) and all other lines was computed and is displayed in the bottom left image. The maximum percent error was 3.39%. The phase error along the region with maximum signal is plotted on the bottom right, and the maximum absolute error was found to be 2.41 degrees.

In summary, the short-term stability tests noted magnitude percent errors of 2.03% and 3.69% and phase errors of 2.04 degrees and 1.60 degrees for the parallel transmit system and GE system respectively. The long-term stability tests demonstrated magnitude percent errors of 1.93% and 3.39%, and phase errors of 2.22 degrees and 2.41 degrees for the parallel transmit system and GE system respectively. The parallel transmit system exhibits slightly less deviation in magnitude for both short and long repetition times than the GE system, and the phase error between the systems for both short and long-term cases is comparable. The GE system does exhibit an obvious

increase in the phase error between the short and long-term stability tests, indicating increasing instability with longer repetitions times whereas the parallel transmit system exhibits similar behavior for both the short and long-term tests.

VII.2 Demonstration of Signal Recovery in the Orbitofrontal Cortex

While parallel transmit technology opens a realm of possible applications, the primary one of interest for this work was to use spatially tailored RF pulses to recover signal in the orbitofrontal cortex region of the brain. Initial work was done on a phantom containing an air-filled region to simulate the effects of magnetic susceptibility and attempt to recover the signal from the region around the air pocket. Our collaborators acquired B_1 maps using the Bloch-Siegert method to obtain the transmit sensitivity patterns (similar to those show in Figure III.14). These maps were plugged into the pulse design to create a tailored RF pulse using a spokes trajectory [94]. The simulated results were based on the Bloch equations (the governing equations to define relaxation) are shown in Figure VII.5.

The simulated results correlate well with the imaging results shown in Figure VII.6. With signal recovery demonstrated on a phantom, initial results were acquired on a human volunteer to show signal recovery in the orbitofrontal cortex region of the brain. The simulated results are shown in Figure VII.7.
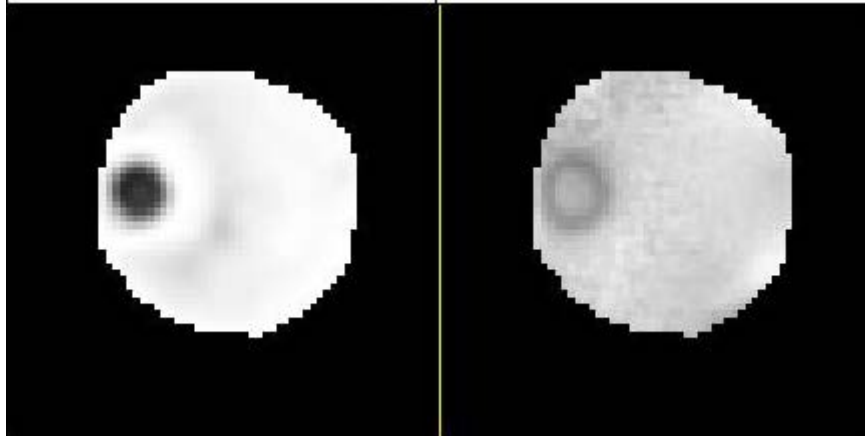
Figure VII.5: Simulated results for both an uncorrected pulse and a signal recovery spatially tailored RF pulse in a phantom containing an air pocket.
The left image shows the simulated result for the uncorrected (no signal recovery) case and demonstrates a null in the region of the air pocket. The right image is the simulated result for the corrected case and shows signal recovery in the region of interest.
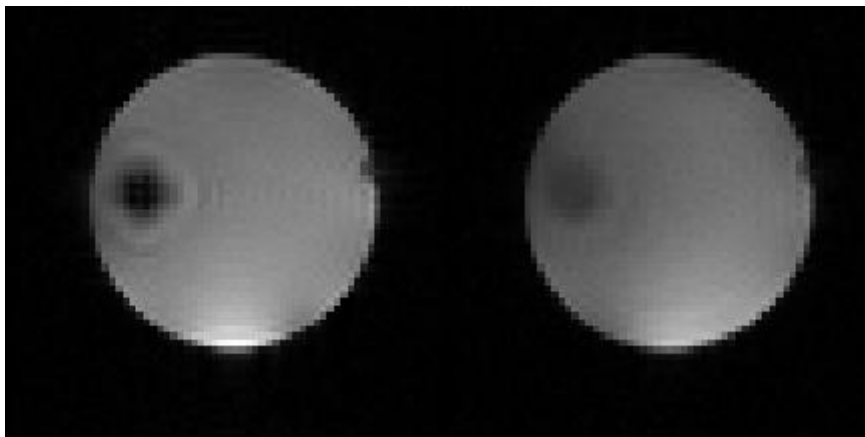
Figure VII.6: Imaging results obtained on the phantom containing an air pocket.
The left image shows the uncorrected result using a standard sinc pulse for excitation, and the right image shows signal recovery in the region near the susceptibility artifact using the spatially tailored RF pulse.
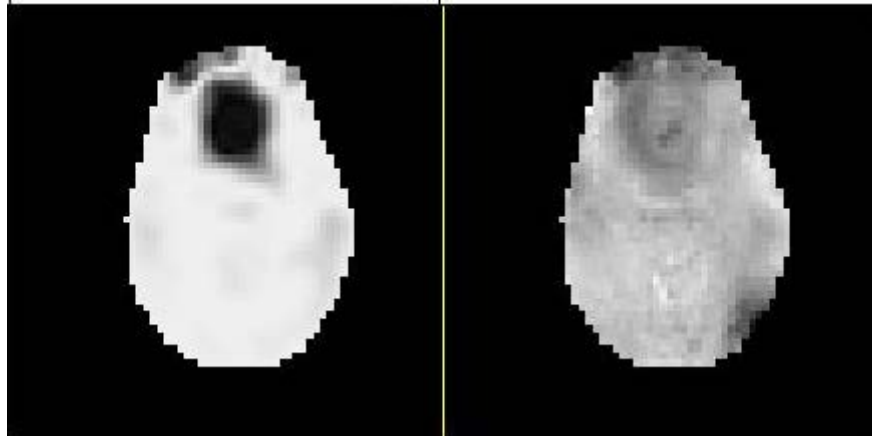
Figure VII.7: Simulated results for both an uncorrected pulse and a signal recovery pulse in the orbitofrontal cortex of the brain.
The left image shows the uncorrected pulse, and the right image shows the result for the signal recovery pulse.
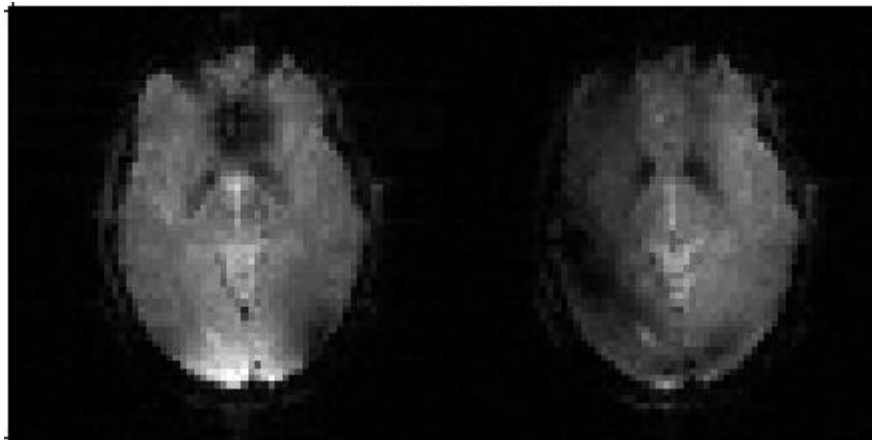
Figure VII.8: Imaging results obtained on a human volunteer demonstrating signal recovery in the orbitofrontal cortex region of the brain.
The left image shows the results using a conventional sinc pulse for excitation, and the null near the top of the image behind the sinuses is clearly visible. The right image shows the same slice and region of the brain but with the spatially tailored RF pulse applied. Signal recovery in the orbitofrontal cortex is evident.

173

Both the uncorrected and corrected pulses were played out in the experiment to obtain a reference image for the signal null in the orbitofrontal cortex region of the brain and a corrected image showing signal recovery in the same region due to the spatially tailored RF pulse. The results are shown in Figure VII.8.

There are still some non-uniformities in the right excitation image in Figure VII.8 particularly near the bottom left, and this may be due to incorrect B_1 maps. The B_1 maps obtained and used for the phantom were the same maps used to design the correction pulse in the human volunteer even though the transmit patterns will have changed slightly due to differences in loading behavior between the phantom and brain.

Since the acquisition of these images, improvements were made in the calibration routine performing the linearity correction and multi-slice capability was enabled. The added functionality and improved accuracy will further improve the quality of the system and allow for further improvements in the excitations.

These images were also acquired using only 7 channels of the parallel transmit system due to a failed gate bias line on channel one. With this problem solved, additional flexibility will be available to the pulse designer.

VII.3 Conclusions and Future Work

This work has described an eight-channel transmit/receive head array and control system for use in parallel transmit MRI. The array coil is the first head array designed for use with ultra-low output impedance amplifiers and demonstrates an unconventional approach to array coil design with the benefit of straightforward element decoupling. The control system with linearization and monitoring capabilities demonstrates a flexible

and non-proprietary solution for operating parallel transmit hardware. The control and modulator subsystem were used on three different systems (3 Tesla GE clinical research scanner, 4.7 Tesla Varian animal research scanner, and a 7 Tesla Philips clinical research scanner), demonstrating the flexibility in the control and modulator subsystem design.

The functionality of the eight-channel transmit/receive head array and control system were demonstrated, and in combination with the remaining pieces of the transmit chain, parallel transmit capability was achieved. In particular, initial results from a signal recovery experiment in the orbitofrontal cortex region of the brain were obtained, the primary application of this work. However, the system has yet to be exhaustively used in combination with pulse design experimentation, and the realm of benefits and possible applications attainable with the added technology are only just now being explored.

Components of the system remain to be implemented in the environment of an MR scan, but all functionality was tested. Future work will focus on ensuring the added monitoring and multi-slice software works well in the scan environment through remote support. The future will bring clinical studies to the project, and with increased use of the system, collaboration and support will continue and open new doors for investigation. The greatest contribution of this work is the future science it enables - providing a flexible approach to parallel transmit technology that in the hands of pulse designers, has the potential to accomplish much.

REFERENCES

[1]   P. C. Lauterbur, "Image formation by induced local interactions: examples employing nuclear magnetic resonance," *Nature,* vol. 242, pp. 190-191, 1973.

[2]   P. Mansfield and P. K. Grannell, "NMR'diffraction'in solids?," *Journal of Physics C: Solid State Physics,* vol. 6, p. L422, 1973.

[3]   W. A. Edelstein, G. H. Glover, C. J. Hardy, and R. W. Redington, "The intrinsic signal-to-noise ratio in NMR imaging," *Magnetic Resonance in Medicine,* vol. 3, pp. 604-618, 2005.

[4]   T. S. Ibrahim, R. Lee, B. A. Baertlein, A. M. Abduljalil, H. Zhu*, et al.*, "Effect of RF coil excitation on field inhomogeneity at ultra high fields: a field optimized TEM resonator," *Magnetic Resonance Imaging,* vol. 19, pp. 1339-1347, 2001.

[5]   D. I. Hoult, "Sensitivity and power deposition in a high-field imaging experiment," *Journal of Magnetic Resonance Imaging,* vol. 12, pp. 46-67, 2000.

[6]   R. Abraham and T. S. Ibrahim, "Proposed radiofrequency phased-array excitation scheme for homogenous and localized 7-Tesla whole-body imaging based on full-wave numerical simulations," *Magnetic Resonance in Medicine,* vol. 57, pp. 235-242, 2007.

[7]   T. S. Ibrahim, R. Lee, A. M. Abduljalil, B. A. Baertlein, and P.-M. L. Robitaille, "Dielectric resonances and B_1 field inhomogeneity in UHFMRI: computational analysis and experimental findings," *Magnetic Resonance Imaging,* vol. 19, pp. 219-226, 2001.

[8]   A. Kangarlu, B. A. Baertlein, R. Lee, T. Ibrahim, L. Yang*, et al.*, "Dielectric resonance phenomena in ultra high field MRI," *Journal of Computer Assisted Tomography,* vol. 23, pp. 821-831, 1999.

[9]   Q. X. Yang, J. H. Wang, X. L. Zhang, C. M. Collins, M. B. Smith*, et al.*, "Analysis of wave behavior in lossy dielectric samples at high field," *Magnetic Resonance in Medicine,* vol. 47, pp. 982-989, May 2002.

[10]   C. M. Collins, W. Z. Liu, W. Schreiber, Q. X. Yang, and M. B. Smith, "Central brightening due to constructive interference with, without, and despite dielectric resonance," *Journal of Magnetic Resonance Imaging,* vol. 21, pp. 192-196, Feb 2005.

[11]   C. G. Yang, W. R. Deng, V. Alagappan, L. L. Wald, and V. A. Stenger, "Four-dimensional spectral-spatial RF pulses for simultaneous correction of B_1+ inhomogeneity and susceptibility artifacts in T_2*-weighted MRI," *Magnetic Resonance in Medicine,* vol. 64, pp. 1-8, Jul 2010.

[12]   V. A. Stenger, F. E. Boada, and D. C. Noll, "Three-dimensional tailored RF pulses for the reduction of susceptibility artifacts in T^{*2}-weighted functional MRI," *Magnetic Resonance in Medicine,* vol. 44, pp. 525-531, Oct 2000.

[13]   C. Y. Yip, J. A. Fessler, and D. C. Noll, "Advanced three-dimensional tailored RF pulse for signal recovery in T_2*-weighted functional magnetic resonance imaging," *Magnetic Resonance in Medicine,* vol. 56, pp. 1050-1059, Nov 2006.

[14]   Y. D. Zhu, "Parallel excitation with an array of transmit coils," *Magnetic Resonance in Medicine,* vol. 51, pp. 775-784, Apr 2004.

[15]   U. Katscher and P. Bornert, "Parallel RF transmission in MRI," *NMR in Biomedicine,* vol. 19, pp. 393-400, May 2006.

[16]   U. Katscher, P. Bornert, C. Leussler, and J. S. van den Brink, "Transmit SENSE," *Magnetic Resonance in Medicine,* vol. 49, pp. 144-150, Jan 2003.

[17]   W. Grissom, C. Y. Yip, Z. H. Zhang, V. A. Stenger, J. A. Fessler*, et al.*, "Spatial domain method for the design of RF pulses in multicoil parallel excitation," *Magnetic Resonance in Medicine,* vol. 56, pp. 620-629, Sep 2006.

[18]   P. Roemer, W. Edelstein, C. Hayes, S. Souza, and O. Mueller, "The NMR phased array," *Magnetic Resonance in Medicine,* vol. 16, pp. 192-225, 1990.

[19]   J. D. Enderle and J. D. Bronzino, *Introduction to biomedical engineering.* Burlington, Massachussetts, United States: Academic Press, 2012.

[20]   E. M. Haacke, R. W. Brown, M. R. Thompson, and R. Venkatesan, *Magnetic resonance imaging: physical principles and sequence design* vol. 1: Wiley-Liss New York:, 1999.

[21]   S. A. Huettel, A. W. Song, and G. McCarthy, *Functional magnetic resonance imaging* vol. 1. Sunderland, Massachussetts, United States: Sinauer Associates, 2004.

[22]   S. Ogawa, D. W. Tank, R. Menon, J. M. Ellermann, S.-G. Kim*, et al.*, "Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging," *Proceedings of the National Academy of Sciences,* vol. 89, pp. 5951-5955, 1992.

[23]   B. Casey, J. D. Cohen, P. Jezzard, R. Turner, D. C. Noll*, et al.*, "Activation of prefrontal cortex in children during a nonspatial working memory task with functional MRI," *Neuroimage,* vol. 2, pp. 221-229, 1995.

[24]   M. L. Kringelbach, "The human orbitofrontal cortex: linking reward to hedonic experience," *Nature Reviews Neuroscience,* vol. 6, pp. 691-702, 2005.

[25]   S. Saxena, A. L. Brody, J. M. Schwartz, and L. R. Baxter, "Neuroimaging and frontal-subcortical circuitry in obsessive-compulsive disorder," *The British Journal of Psychiatry,* vol. 173, pp. 26-37, Aug 1998.

[26]   V. A. Stenger, "Technical considerations for BOLD fMRI of the orbitofrontal cortex," *The Orbitofrontal Cortex,* pp. 423-46, 2006.

[27]   K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, "SENSE: sensitivity encoding for fast MRI," *Magnetic Resonance in Medicine,* vol. 42, pp. 952-962, 1999.

[28]   D. K. Sodickson and W. J. Manning, "Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays," *Magnetic Resonance in Medicine,* vol. 38, pp. 591-603, 1997.

[29]   L. I. Sacolick, F. Wiesinger, I. Hancu, and M. W. Vogel, "B1 mapping by Bloch-Siegert shift," *Magnetic Resonance in Medicine,* vol. 63, pp. 1315-1322, 2010.

[30] F. Seifert and H. Rinneberg, "Adaptive coil control: SNR optimization of a TR volume coil for single voxel MRS at 3 T," in *Proceedings of the 10th Annual Meeting of International Society for Magnetic Resonance in Medicine*, Honolulu, United States, 2002, p. 162.

[31] G. Adriany, V. de Moortele, F. Wiesinger, S. Moeller, J. P. Strupp, *et al.*, "Transmit and receive transmission line arrays for 7 Tesla parallel imaging," *Magnetic Resonance in Medicine,* vol. 53, pp. 434-445, 2005.

[32] Y. Zhu and R. Giaquinto, "Improving flip angle uniformity with parallel excitation," in *Proceedings of the 13th Annual Meeting of the International Society of Magnetic Resonance in Medicine*, Miami, United States, 2005, p. 2752.

[33] P. Ullmann, S. Junge, M. Wick, F. Seifert, W. Ruhm, *et al.*, "Experimental analysis of parallel excitation using dedicated coil setups and simultaneous RF transmission on multiple channels," *Magnetic Resonance in Medicine,* vol. 54, pp. 994-1001, 2005.

[34] R. F. Lee, R. O. Giaquinto, and C. J. Hardy, "Coupling and decoupling theory and its application to the MRI phased array," *Magnetic Resonance in Medicine,* vol. 48, pp. 203-213, Jul 2002.

[35] J. Jevtic, "Ladder networks for capacitive decoupling in phased-array coils," in *Proceedings of the 9th Annual Meeting of International Society of Magnetic Resonance in Medicine*, Glasgow, Scotland, 2001.

[36] J. Jevtic, V. Pikelja, A. Menon, D. Seeber, N. Tatum, *et al.*, "Design guidelines for the capacitive decoupling networks," in *Proceedings of the 11th Annual Meeting of the International Society of Magnetic Resonance in Medicine*, Toronto, Canada, 2003, p. 428.

[37] K. Kurpad, E. Boskamp, and S. Wright, "Implementation of coil integrated RF power MOSFET as voltage controlled current source in a transmit phased array coil," in *Proceedings of the 12th Annual Meeting of International Society of Magnetic Resonance in Medicine*, Kyoto, Japan, 2004.

[38]   K. Kurpad, E. Boskamp, and S. Wright, "A parallel transmit volume coil with independent control of currents on the array elements," in *Proceedings of the 13th Annual Meeting of the International Society for Magnetic Resonance in Medicine*, Miami, United States, 2005, p. 16.

[39]   K. N. Kurpad, S. M. Wright, and E. B. Boskamp, "RF current element design for independent control of current amplitude and phase in transmit phased arrays," *Concepts in Magnetic Resonance Part B-Magnetic Resonance Engineering,* vol. 29B, pp. 75-83, Apr 2006.

[40]   N. Gudino, J. Heilman, M. Riffe, C. Flask, and M. Griswold, "7T current-mode class-D (CMCD) RF power amplifier," in *Proceedings of the 17th Annual Meeting of International Society for Magnetic Resonance in Medicine*, Honolulu, United States, 2009, p. 398.

[41]   N. Gudino, J. A. Heilman, M. J. Riffe, O. Heid, M. Vester, *et al.*, "On-coil multiple channel transmit system based on class-D amplification and pre-amplification with current amplitude feedback," *Magnetic Resonance in Medicine,* pp. 276-289, 2012.

[42]   K. N. Kurpad, "Transmit field pattern control for high field magnetic resonance imaging with integrated RF current sources," Texas A&M University, 2012.

[43]   J. A. Heilman, N. Gudino, M. J. Riffe, M. Vester, and M. A. Griswold, "Preamp-like decoupling and amplitude modulation in CMCD amplifiers for transmit arrays," in *Proceedings of the 16th Annual Meeting of the International Society for Magnetic Resonance in Medicine*, Toronto, Canada, 2008.

[44]   X. Chu, X. Yang, Y. F. Liu, J. Sabate, and Y. D. Zhu, "Ultra-low output impedance RF power amplifier for parallel excitation," *Magnetic Resonance in Medicine,* vol. 61, pp. 952-961, Apr 2009.

[45]   C. W. Chang, K. L. Moody, and M. P. Mcdougall, "An improved element design for 64-channel planar imaging," *Concepts in Magnetic Resonance Part B-Magnetic Resonance Engineering,* vol. 39B, pp. 159-165, Aug 2011.

[46]   K. L. Moody, N. A. Hollingsworth, J.-F. Nielsen, D. Noll, S. M. Wright, *et al.*, "RF transparent array for testing multi-channel transmit systems," in *Proceedings*

of the 19th Annual Meeting of the International Society for Magnetic Resonance in Medicine, Stockholm, Sweden, 2010.

[47]   K. L. Moody, N. A. Hollingsworth, F. Zhao, J.-F. Nielsen, D. C. Noll*, et al.*, "An eight-channel T/R head coil for parallel transmit MRI at 3T using ultra-low output impedance amplifiers," *Journal of Magnetic Resonance,* vol. 246, pp. 62-68, 2014.

[48]   K. L. Moody, N. A. Hollingsworth, J. F. Nielsen, D. Noll, M. P. McDougall*, et al.*, "Eight-channel transmit/receive head array for use with ultra-low output impedance amplifiers," in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, 2013, pp. 950-953.

[49]   M. J. Firbank, R. M. Harrison, E. D. Williams, and A. Coulthard, "Quality assurance for MRI: practical experience," *British Journal of Radiology,* vol. 73, pp. 376-383, 2000.

[50]   A. Reykowski, S. M. Wright, and J. R. Porter, "Design of matching networks for low noise preamplifiers," *Magnetic Resonance in Medicine,* vol. 33, pp. 848-852, 1995.

[51]   F. Zhao, J. A. Fessler, S. M. Wright, J. V. Rispoli, and D. C. Noll, "Optimized linear combinations of channels for complex multiple-coil B1 field estimation with Bloch-Siegert B1 mapping in MRI," in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, San Franscisco, United States, 2013.

[52]   F. Zhao, J. Fessler, J.-F. Nielsen, and D. Noll, "Regularized estimation of magnitude and phase of multiple-coil B1 field via Bloch-Siegert B1 mapping," in *Proceedings of the 20th Annual Meeting of the International Society for Magnetic Resonance in Medicine*, Melbourne, Australia, 2012, p. 2512.

[53]   F. Doty, T. Connick, X. Ni, and M. Clingan, "Noise in high-power, high-frequency double-tuned probes," *Journal of Magnetic Resonance (1969),* vol. 77, pp. 536-549, 1988.

[54]   C. A. Michal, K. Broughton, and E. Hansen, "A high performance digital receiver for home-built nuclear magnetic resonance spectrometers," *Review of Scientific Instruments,* vol. 73, pp. 453-458, 2002.

[55]   J. Bodurka, P. J. Ledden, P. van Gelderen, R. Chu, J. A. de Zwart, *et al.*, "Scalable multichannel MRI data acquisition system," *Magnetic Resonance in Medicine,* vol. 51, pp. 165-171, 2004.

[56]   M. P. McDougall and S. M. Wright, "64-channel array coil for single echo acquisition magnetic resonance imaging," *Magnetic Resonance in Medicine,* vol. 54, pp. 386-392, 2005.

[57]   S. Wright, M. McDougall, and D. Brown, "Single echo acquistion of MR images using RF coil arrays," in *Proceedings of the 24th Annual Conference and the Annual Fall Meeting of the Biomedical Engineering Society/Engineering in Medicine and Biology Society/Biomedical Engineering Society Conference*, 2002, pp. 1181-1182.

[58]   K. Takeda, "A highly integrated FPGA-based nuclear magnetic resonance spectrometer," *Review of Scientific Instruments,* vol. 78, pp. 033103-033103-6, 2007.

[59]   W. Mao, Q. Bao, L. Yang, Y. Chen, C. Liu, *et al.*, "A modularized pulse programmer for NMR spectroscopy," *Measurement Science and Technology,* vol. 22, p. 025901, 2011.

[60]   S. M. Wright, D. G. Brown, J. R. Porter, D. C. Spence, E. Esparza, *et al.*, "A desktop magnetic resonance imaging system," *Magnetic Resonance Materials in Physics, Biology and Medicine,* vol. 13, pp. 177-185, 2001.

[61]   K. Kose, T. Haishi, A. Nakanishi, S. Okada, and T. Tsuzaki, "Development of a desktop MR microscopy using a small permanent magnet," in *Proceedings of the 8th Annual Meeting of the International Society for Magnetic Resonance in Medicine*, Denver, United States, 2000, p. 1380.

[62]   M. Rokitta, E. Rommel, U. Zimmermann, and A. Haase, "Portable nuclear magnetic resonance imaging system," *Review of Scientific Instruments,* vol. 71, pp. 4257-4262, 2000.

[63]   P. Stang, S. Conolly, J. Pauly, and G. Scott, "MEDUSA: A scalable MR console for parallel imaging," in *Proceedings of the 15th Annual Meeting of the*

International Society for Magnetic Resonance in Medicine, Berlin, Germany, 2007, p. 925.

[64]   P. P. Stang, S. M. Conolly, J. M. Santos, J. M. Pauly, and G. C. Scott, "Medusa: a scalable MR console using USB," *IEEE Transactions on Biomedical Engineering,* vol. 31, pp. 370-379, 2012.

[65]   N. Hollingsworth, K. Moody, J.-F. Nielsen, D. Noll, M. McDougall, *et al.*, "An eight channel transmit system for Transmit SENSE at 3T," in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, Chicago, United States, 2011, pp. 775-778.

[66]   K. L. Moody, N. A. Hollingsworth, J. Nielsen, D. Noll, S. Wright, *et al.*, "A LabVIEW-based operating system for parallel transmit systems," in *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, Chicago, United States, 2011, pp. 771-774.

[67]   K. Feng, N. A. Hollingsworth, M. P. McDougall, and S. M. Wright, "A 64-channel transmitter for investigating parallel transmit MRI," *IEEE Transactions on Biomedical Engineering,* vol. 59, pp. 2152-2160, Aug 2012.

[68]   D. Hoult, G. Kolansky, and D. Kripiakevich, "A 'Hi-Fi'Cartesian feedback spectrometer for precise quantitation and superior performance," *Journal of Magnetic Resonance,* vol. 171, pp. 57-63, 2004.

[69]   M. Zanchi, J. Pauly, and G. Scott, "Frequency offset Cartesian feedback control system for MRI power amplifier," in *Proceedings of the 17th Annual Meeting of the International Society for Magnetic Resonance in Medicine*, Honolulu, United States, 2009, p. 398.

[70]   M. G. Zanchi, P. Stang, A. Kerr, J. M. Pauly, and G. C. Scott, "Frequency-offset Cartesian feedback for MRI power amplifier linearization," *IEEE Transactions on Medical Imaging,* vol. 30, pp. 512-522, Feb 2011.

[71]   F. Chan, J. Pauly, and A. Macovski, "Effects of RF amplifier distortion on selective excitation and their correction by prewarping," *Magnetic Resonance in Medicine,* vol. 23, pp. 224-238, 1992.

[72]   P. Stang, A. Kerr, W. Grissom, J. Pauly, and G. Scott, "Vector iterative pre-distortion: An auto-calibration method for transmit arrays," in *Proceedings of the 17th Annual Meeting of the International Society for Magnetic Resonance in Medicine*, Honolulu, United States, 2009, p. 395.

[73]   P. Yazdanbakhsh, S. Held, I. Brote, A. Bitz, S. Orzada, *et al.*, "16-bit vector modulator for B_1 shimming in 7T MRI," in *Proceedings of the 17th Annual Meeting of the International Society for Magnetic Resonance in Medicine*, Honolulu, United States, 2009, p. 4768.

[74]   International Electrotechnical Commission, 60601-2-33, "Medical electrical equipment and particular requirements for the safety of magnetic resonance equipments for medical diagnosis", *in Protection Against Hazardous Output*, 2002.

[75]   L. L. Wald and E. Adalsteinsson. (2010, Jan) Specific absorption rate (SAR) in parallel transmission (pTx). *MAGNETOM Flash*. Boston, United States, Siemens Healthcare.

[76]   A. C. Zelinski, L. L. Wald, K. Setsompop, V. Alagappan, B. A. Gagoski, *et al.*, "Comparison of three algorithms for solving linearized systems of parallel excitation RF waveform design equations: Experiments on an eight-channel system at 3 tesla," *Concepts in Magnetic Resonance Part B-Magnetic Resonance Engineering,* vol. 31B, pp. 176-190, Aug 2007.

[77]   I. Graesslin, F. Schweser, B. Annighoefer, S. Biederer, U. Katscher, *et al.*, "A minimum SAR RF pulse design approach for parallel Tx with local hot spot suppression and exact fidelity constraint," in *Proceedings of the 16th Annual Meeting of the International Society for Magnetic Resonance in Medicine*, Toronto, Canada, 2008.

[78]   M. Gebhardt, D. Diehl, E. Adalsteinsson, L. Wald, and G. Eichfelder, "Evaluation of maximum local SAR for parallel transmission (pTx) pulses based on pre-calculated field data using a selected subset of "Virtual Observation Points"," in *Proceedings of the 18th Annual Meeting of the International Society of Magnetic Resonance in Medicine*, Stockholm, Sweden, 2010, p. 1441.

[79]   J. Lee, M. Gebhardt, L. Wald, and E. Adalsteinsson, "Parallel transmit RF design with local SAR constraints," in *Proceedings of the 18th Annual Meeting of the*

International Society of Magnetic Resonance in Medicine, Stockholm, Sweden, 2010, p. 105.

[80]   A. C. Zelinski, K. Setsompop, E. Adalsteinsson, and V. Goyal, Massachusetts Institute of Technology, "Method for reducing maximum local specific absorption rate in magnetic resonance imaging," US8148985 B2, 2012.

[81]   I. Graesslin, S. Biederer, F. Schweser, K. Zimmermann, U. Katscher, *et al.*, "SAR reduction for parallel transmission using VERSE and k-space filtering," in *Proceedings of the 15th Annual Meeting of the International Society for Magnetic Resonance in Medicine*, Berlin, Germany, 2007, p. 674.

[82]   I. Graesslin, J. Weller, F. Schweser, B. Annighoefer, S. Biederer, *et al.*, "SAR hotspot reduction by temporal averaging in parallel transmission," in *Proceedings of the 17th Annual Meeting of International Society of Magnetic Resonance in Medicine*, Honolulu, United States, 2009, p. 176.

[83]   A. C. Zelinski, "Improvements in magnetic resonance imaging excitation pulse design," Massachusetts Institute of Technology, 2008.

[84]   I. Graesslin, S. Biederer, B. Annighoefer, H. Homann, H. Stahl, *et al.*, "Real-time global and local SAR monitoring for parallel transmission systems," in *Proceedings of the 17th Annual Meeting of the International Society of Magnetic Resonance in Medicine*, Honolulu, United States, 2009, p. 302.

[85]   M. A. Bernstein, Mayo Foundation for Medical Education and Research, "MRI RF power monitor," United States, United States Patent and Trademark Office, US6426623 B1, 2002.

[86]   A. Bitz, I. Brote, S. Orzada, O. Kraff, S. Maderwald, *et al.*, "An 8-channel add-on RF shimming system for whole-body 7 Tesla MRI including real-time SAR monitoring," in *Proceedings of the 17th Annual Meeting of International Society for Magnetic Resonance in Medicine*, Honolulu, United States, 2009, pp. 18-24.

[87]   I. Brote, S. Orzada, O. Kraff, S. Maderwald, H. Quick, *et al.*, "A multi-channel SAR prediction and online monitoring system at 7 T," in *Proceedings of the 17th Annual Meeting of International Society for Magnetic Resonance in Medicine*, Honolulu, United States, 2009, p. 4788.

[88]   I. Graesslin, K. Falaggis, P. Vernickel, P. Röschmann, C. Leussler, *et al.*, "Safety considerations concerning SAR during RF amplifier malfunctions in parallel transmission," in *Proceedings of the 14th Annual Meeting of International Society of Magnetic Resonance in Medicine*, Seattle, United States, 2006.

[89]   P. Vernickel, P. Röschmann, C. Findeklee, K. M. Lüdeke, C. Leussler, *et al.*, "Eight-channel transmit/receive body MRI coil at 3T," *Magnetic Resonance in Medicine,* vol. 58, pp. 381-389, 2007.

[90]   P. Ullmann and M. Bruker BioSpin, "Multiple-channel RF pulses and SAR control," in *Proceedings of the 16th Annual Meeting of the International Society of Magnetic Resonance in Medicine*, Toronto, Canada, 2008.

[91]   I. Graesslin, K. Falaggis, P. Vernickel, P. Röschmann, C. Leussler, *et al.*, "Ensuring patient safety in case of RF channel failure in parallel transmission," *Magnetic Resonance in Materials in Physics, Biology and Medicine,* vol. 19, p. 261, 2006.

[92]   Y. Zhu, L. Alon, C. M. Deniz, R. Brown, and D. K. Sodickson, "System and SAR characterization in parallel RF transmission," *Magnetic Resonance in Medicine,* vol. 67, pp. 1367-1378, 2012.

[93]   I. Graesslin, P. Vernickel, J. Schmidt, C. Findeklee, P. Röschmann, *et al.*, "Whole body 3T MRI system with eight parallel RF transmission channels," in *Proceedings of the 14th Annual Meeting of the International Society of Magnetic Resonance in Medicine*, Seattle, United States, 2006.

[94]   D. Yoon, J. A. Fessler, A. C. Gilbert, and D. C. Noll, "Fast joint design method for parallel excitation radiofrequency pulse and gradient waveforms considering off-resonance," *Magnetic Resonance in Medicine,* vol. 68, pp. 278-285, Jul 2012.