INSTRUMENTATION FOR PARALLEL MAGNETIC RESONANCE IMAGING

A Dissertation

by

DAVID GERALD BROWN

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

DOCTOR OF PHILOSOPHY

December 2005

Major Subject: Electrical Engineering

INSTRUMENTATION FOR PARALLEL MAGNETIC RESONANCE IMAGING

A Dissertation

by

DAVID GERALD BROWN

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

DOCTOR OF PHILOSOPHY

Approved by:

Chair of Committee,	Steven M. Wright
Committee Members,	Robert D. Nevels
	Andrew K. Chan
	Gerard L. Cote
Head of Department,	Costas N. Georghiades

December 2005

Major Subject: Electrical Engineering

ABSTRACT

Instrumentation for Parallel Magnetic Resonance Imaging. (December 2005)

David Gerald Brown, B.S., Texas A&M University

Chair of Advisory Committee: Dr. Steven M. Wright

Parallel magnetic resonance (MR) imaging may be used to increase either the throughput or the speed of the MR imaging experiment. As such, parallel imaging may be accomplished either through a "parallelization" of the MR experiment, or by the use of arrays of sensors. In parallelization, multiple MR scanners (or multiple sensors) are used to collect images from different samples simultaneously. This allows for an increase in the throughput, not the inherent speed, of the MR experiment. Parallel imaging with arrays of sensor coils, on the other hand, makes use of the spatial localization properties of the sensors in an imaging array to allow a reduction in the number of phase encodes required in acquiring an image. This reduced phase-encoding requirement permits an increase in the overall imaging speed by a factor up to the number of sensors in the imaging array. The focus of this dissertation has been the development of cost-effective instrumentation that would enable advances in the state of the art of parallel MR imaging.

First, a low-cost desktop MR scanner was developed (< $13,000) for imaging small samples (2.54 cm fields-of view) at low magnetic field strengths (< 0.25 T). The performance of the prototype was verified through bench-top measurements and phantom imaging. The prototype transceiver has demonstrated an SNR (signal-to-noise

ratio) comparable to that of a commercial MR system. This scanner could make parallelization of the MR experiment a practical reality, at least in the areas of small animal research and education.

A 64-channel receiver for parallel MR imaging with arrays of sensors was also developed. The receiver prototype was characterized through both bench-top tests and phantom imaging. The parallel receiver is capable of simultaneous reception of up to sixty-four, 1 MHz bandwidth MR signals, at imaging frequencies from 63 to 200 MHz, with an SNR performance (on each channel) comparable to that of a single-channel commercial MR receiver. The prototype should enable investigation into the speed increases obtainable from imaging with large arrays of sensors and has already been used to develop a new parallel imaging technique known as single echo acquisition (SEA) imaging.

DEDICATION

This work is dedicated to my loving wife, Laura, and our son, Thomas.

ACKNOWLEDGMENTS

The author wishes to thank the many persons and agencies who have helped to make this work possible. Among these are faculty advisors, colleagues, corporations, funding agencies, and family. There is too little room to mention everyone's contributions in detail here, but the most significant have been highlighted in the paragraphs below.

First among these contributors is Dr. Steven M. Wright of the Department of Electrical Engineering at Texas A&M University. Dr. Wright served as the chair of the author's doctoral committee and was his principal advisor. Without his advice, support, and the equipment/ facilities which he provided at the Magnetic Resonance Systems Lab (MRSL), it would have been impossible to conduct the research described in these pages.

Another invaluable resource has been Dr. Jay Porter of the Department of Engineering Technology at Texas A&M University. Dr. Porter provided his expertise in LabVIEW programming and a great deal of his own time to develop the control software for the desktop MR scanner. In addition, Dr. Porter generously provided the services of several of his students in the fabrication of the active shim set that was used on the desktop imaging system's magnet.

The author also wishes to thank his many colleagues at the MRSL who have assisted him during the course of this research. In particular, the author wishes to acknowledge David Spence for all of his help in the initial development of the desktop MR imaging system. Mr. Spence was responsible for developing the gradient amplifier system and was always available as a sounding board in the initial design process. Likewise, the

author wishes to thank Jason Payne, Hyokwon Nam, Dan Spence, Krishna Kurpad, Ed Jeter, Emilio Esparza-Coss, and Jorge Rosas, who helped in the repetitive (and often tedious) process of assembling the prototype 64-channel parallel receiver system.

IGC, Inc. is also gratefully acknowledged for its generous donation of a 0.16 T, whole body, permanent magnet. The IGC magnet was used as a testbed for the development of RF coils, gradient coils, and shims sets during the desktop MR imaging project.

Next, the author wishes to thank the State of Texas Higher Education Coordinating Board (grant no. 000512-0189-1997) for its support. From the spring of 1998 to the spring of 2001, this grant provided much of the funding required to develop the low-cost desktop MR imaging system.

The author also wishes to thank the Department of the Army and its National Medical Technology Testbed agency (Dept. of the Army DAMD17-97-2-7016) for its support. From the spring of 2001 to the winter of 2002, this grant provided funding for the development of the 64-channel parallel receiver prototype.

Last, but not least, the author wishes to thank his wife and family for their unending support, extensive patience, and unwavering belief in his ability to complete the work undertaken in this dissertation.

TABLE OF CONTENTS

CHAPTER Page

CHAPTER Page

CHAPTER Page

LIST OF TABLES

LIST OF FIGURES

FIGURE Page

FIGURE Page

FIGURE Page

FIGURE Page

CHAPTER I

INTRODUCTION

Reduction of imaging time has long been an important concern in magnetic resonance imaging (MRI). Shorter scan times have many benefits, including increased patient comfort, reduction of motion artifacts, and lower cost. Until recently, the primary method for increasing imaging speed in MRI has been to improve gradient-based imaging technology. This has been done to increase both the power and speed with which magnetic field gradient pulses can be applied to a sample. However, high speed gradient imaging methods are restricted by guidelines imposed by the FDA which limit the application of magnetic field gradients and radio-frequency (RF) energy to the patient in order to prevent bio-effects [1]. It turns out that the high switching speeds and large amplitude gradient pulses now achievable on most commercial MR scanners can cause unwanted peripheral nerve stimulation [1-3]. Additionally, the increase in imaging speed has increased the duty cycle of the RF energy applied to the patient during a scan. This higher RF duty cycle can lead to patient heating [1]. These limitations have led to the inception of alternative approaches for increasing imaging speed, such as parallel MR imaging.

I.1 Problem Statement

Parallel imaging is one method for increasing imaging speed without using rapidly

---

This dissertation follows the style and format of the *Institute of Electrical and Electronic Engineers (IEEE) Transactions on Biomedical Engineering*.

switched, high strength gradient fields. In MRI, parallel imaging may be performed in two different ways. The first is to use multiple MR scanners or multiple sample RF coils to "parallelize" the MR experiment. This method enables different samples to be scanned at the same time [4]. While this does not reduce actual imaging time for a given sample, it does increase sample throughput and thereby reduces the total imaging time required to scan multiple samples. The second parallel imaging method involves the use of arrays of receive coils to reduce the number of phase encoding steps required to scan a sample. The elements of a receive array have B-field patterns which are sensitive primarily to MR signals from regions of a sample near the coil. Hence, the signals received by a particular element of the array are known, *a priori*, to be within a specified region near that array element. This information may be used to decrease the amount of time required to scan a particular sample by eliminating or reducing phase encode steps.

I.1.1 Parallelization of the MR Experiment

In small animal imaging, large quantities of research animals (50 to 100 or more) must be imaged at regular time intervals to monitor the progression of new disease treatments. Using a clinical, whole body scanner to perform this task is both prohibitively expensive and time-consuming. One way to get around this obstacle is through the parallelization of the MR experiment by using either multiple, independent MR scanners or multiple sample RF coils within a single whole body MR scanner [4]. This would allow multiple animals (or other samples) to be imaged simultaneously. As a result, the overall scan time for large numbers of samples is reduced by increasing the throughput of the MR imaging experiment.

I.1.1.1 Multiple MR Scanners

Parallel imaging via multiple MR scanners would be useful for increasing the speed (i.e., throughput) of small animal imaging studies, were it not for the high cost and complexity of MR system hardware. Recently, however, advances in the consumer electronics industry have led to the availability of low-cost RF integrated circuits with remarkable capability. These range from single chip receivers to direct digital synthesis (DDS) chipsets. Such components have the ability to dramatically reduce the cost of the MR transceiver instrumentation. When combined with commonly available PC-based data acquisition cards and compact, low-field imaging magnets, these new RF circuits have the potential to form a new class of inexpensive desktop MR instruments. These low-cost MR scanners could be used to make parallel imaging with multiple MR scanners a practical reality.

I.1.1.2 Multiple Sample MR Coils

A second approach to the parallelization of the MR imaging experiment involves the use of RF coils capable of simultaneously imaging multiple samples. The idea of using a multiple sample RF coil for imaging laboratory mice has been the subject of some interest in the MRI community. In 2000, a group at the University of Toronto began developing a multiple-sample RF coil for simultaneously scanning up to sixteen mice at 7 T [5, 6]. At about the same time, a four sample RF coil was developed by a colleague in the Magnetic Resonance Systems Laboratory (MRSL) here at Texas A&M University for imaging mice at 1.5 T in the Department of Diagnostic Radiology at the University of Texas M. D. Anderson Cancer Center in Houston [7, 8]. This coil consisted of four,

lowpass, trombone birdcage coils mounted together inside a cross-shaped RF shield assembly. Both of these groups have shown that multiple coil RF coils have the potential for increasing the throughput of the MR imaging experiment via parallelization.

I.1.2 Array Imaging Techniques

MR imaging is used to obtain an image of a sample that has been placed within a static magnetic field. This image is formed by the application of a series of RF and static magnetic field gradient pulses (a pulse sequence) which interact with the nuclear magnetic dipoles, or spins, contained within the sample. Pulse sequences are used to scan the spatial frequency space, or k-space, of a sample's spins. With traditional gradient based imaging techniques, to acquire an $N \times M$ pixel image, N scans of k-space, called phase encode lines, must be acquired sequentially. Here, M is the number of readout points which is determined by the strength of the gradient pulse applied during the signal acquisition portion of each line of k-space. Since it takes a finite amount of time, TR, to perform each scan, it takes a total time of $N \cdot TR$ to acquire an $N \times M$ pixel image. The goal of parallel imaging with arrays of receive coils is to reduce the number of phase encode lines required to form a given $N \times M$ pixel image.

The idea of using an array of multiple, independent receive coils to reduce MR imaging time was first introduced in 1987 by two different groups [9, 10]. Indeed, Hutchinson and Raff went so far as to claim that a full $N \times M$ image could be acquired entirely without phase encoding by placing N closely spaced, small receive coils around a sample [10]. The main problem with this completely parallel imaging method was that

receiving an image from *N* receive coils required a receiver with at least *N* channels. By the mid-1990's no MR system manufacturer had produced a receiver with more than four channels; hence, the method proposed by Hutchinson and Raff was regarded as impractical and largely abandoned in favor of fast gradient based imaging.

By the late 1990's, fast gradient based imaging techniques were rapidly approaching the FDA imposed safety limits. This led to a revival of interest in parallel imaging with arrays; however, MR system manufacturers were still providing receivers with only a limited number of channels. For this reason, several partially parallel imaging (PPI) techniques came about which reduced imaging time not by eliminating phase encoding, but by reducing the number of phase encode steps required to take an image. In 1989, the idea of using MR arrays to perform partially parallel k-space scans was first proposed by Kelton, *et al.* [11]. In 1997, Sodickson released the SMASH technique for rapid MR imaging with arrays [12]. Roughly a year later, in 1999, Preussman published the SENSE method [13] which is essentially a refined version of the algorithm first proposed by Kelton. Since then, other PPI methods have been introduced, such as PILS and GRAPPA [14, 15]. All of these methods use arrays of receive coils to reduce the number of phase encode steps needed to acquire an image.

To date, SMASH (or SiMultaneous Acquisition of Spatial Harmonics) and SENSE (or SENSitivity Encoding) remain the most common partially parallel imaging techniques. Both of these techniques increase imaging speed by reducing the number of phase encoding steps needed to encode an image. In SMASH reconstruction the coil sensitivity patterns for each element in an array and the reduced k-space scan data are

used to synthesize the phase encode lines which were "skipped" in the initial scan. Using the SMASH technique, total scan time may be reduced at most by a factor equal to the number of coils in the array. However, it must be noted that at high acceleration factors (i.e., as the acceleration factor approaches the number of coils in the array), the SNR in the reconstructed image degrades significantly. In SENSE, a set of *N* aliased images are acquired from an *N*-element receive array. The sensitivity patterns for each element in the array are used to form a set of linear equations which relate the aliased images from each coil to the unaliased full field-of-view image. Solving this set of equations performs an "unwrapping" of the data in the image domain to yield the desired image. Like SMASH, the SENSE technique can be used to reduce scan time by at most a factor equal to the number of coils in the array. So, the amount by which these techniques can reduce imaging time is limited by the number of receiver channels available on a particular MR scanner.

It is clear that to increase imaging speed any further with partially parallel techniques, larger arrays of receive coils must be used. However, these larger arrays require a parallel receiver with enough channels to support them. When the work for this dissertation was begun in early 1998, the largest MR receiver reported in the literature had only sixteen channels [16], and no commercial MR system could support more than eight receiver channels. Thus, for scan times to be reduced any further by partially parallel imaging methods, it was determined that a dedicated MR receiver for parallel imaging was needed. Such a receiver would support a minimum of 64-channels,

enabling both fully parallel imaging and partially parallel imaging with high acceleration factors.

I.2 Research Objectives

The primary objective of this work is to identify cost-effective methods for implementing parallel MR imaging. To accomplish this goal, four specific aims have been defined. These aims encompass two methods for parallel MR imaging: parallel imaging via multiple low-cost MR scanners and parallel imaging with arrays of receive coils. (The third method, parallelization by multiple sample MR coils, has been excluded from this dissertation as it has been undertaken by a colleague within the MRSL.) The first aim investigates the former method for parallel MR imaging, while the last three aims deal with the latter. Accomplishment of the four aims will enable the overall research objective to be met.

I.2.1 Aim 1: Build a Low-Cost Desktop MR Scanner

In order to develop an understanding of the intricacies of the MR receiver, a prototype desktop MR scanner will be constructed. The scanner will be capable of imaging small volumes with high resolution on dedicated, low-field MR magnets. The scanner will be based upon a conventional analog transceiver design and will be built from off-the-shelf sub-systems to lower system cost. The prototype will demonstrate the feasibility of inexpensive, low-field MR scanners as a new-class of general use laboratory instrument for non-clinical applications.

I.2.2 Aim 2: Design a Digital MR Receiver

A cost-effective, single-channel MR receiver will be designed to serve as a building block for a highly parallel MR receiver system. The receiver will be suitable for imaging small volumes with surface coils at 4.7 T. The performance of the receiver will be comparable to that of commercial analog receivers. Finally, the single-channel receiver design must be easily scaled up to form a 64-channel receiver.

I.2.3 Aim 3: Implement a 64-Channel Parallel MR Receiver

A prototype receiver for parallel imaging will be implemented based upon the single-channel design of Aim 2. The parallel imaging receiver will be low in cost, have 64-channels total, and will easily interface to any commercial MR scanner.

I.2.4 Aim 4: Characterize the Performance of the 64-Channel Receiver

The 64-channel system will be fully characterized to establish how well it meets the requirements for a cost-effective parallel imaging MR receiver. The bandwidth, operating frequency range, noise performance, dynamic range, interchannel isolation, gain variation between channels, and system cost will all be determined for the receiver. As part of the noise performance characterization, the signal-to-noise ratio of the receiver will be compared to that of a single-channel commercial scanner.

I.3 Dissertation Organization

This dissertation has been organized to follow the completion of the research objectives detailed in section I.2. As such the dissertation has been divided into eight chapters. The first two chapters supply the motivation and required background theory for this work. Next, Chapters III through VII describe how each of the four main

research objectives was satisfied. Finally, in Chapter VIII, conclusions are drawn from the research and suggestions for future work are made. The contents of each chapter will be described in the next paragraphs.

Chapter I presents the problem statement, research objectives, and an outline of the dissertation. The outline has been provided to assist the reader in identifying the location of relevant information within the dissertation. Chapter II conveys the principles necessary for understanding the methods used to design cost-effective instrumentation for parallel MR imaging. These include a brief description of NMR theory, the basics of MR image formation, the received MR signal, the design of the MR receiver, and performance metrics for MR receivers.

Chapter III details the process by which the low-cost, desktop MR scanner was developed to satisfy the first research objective. The desktop scanner prototype was developed during a three year period from spring of 1998 to spring of 2001 as part of a grant from the State of Texas Higher Education Coordinating Board (no. 000512-0189-1997). The chapter begins with the initial design requirements which drove the design of the prototype desktop system. Then, the overall system design for the desktop scanner is presented. Finally, an evaluation of the desktop scanner is given which includes a performance summary, an evaluation of the desktop system cost, an SNR comparison to a commercial scanner, and a demonstration of system performance during 2D and 3D MR imaging.

Chapters IV through VI together present the design, implementation, and characterization of a dedicated, 64-channel receiver for parallel MR imaging. The

receiver was developed to fulfill research aims 2 through 4. Actual work on the 64-channel receiver prototype was carried out during an 18 month period from the spring of 2001 to winter of 2002 as part of a one year grant from the National Medical Technology Testbed (Dept. of the Army DAMD17-97-2-7016). Chapter IV presents the design of a scalable, single-channel, digital MR receiver. In Chapter V, the implementation of the 64-channel parallel MR receiver system is described. Then, in chapter VI, the characterization of the 64-channel prototype is presented.

Chapter VII presents methods for correcting the phase jitter problem that is commonly encountered in MR receivers. As such, it first describes the nature of phase jitter. Then, four techniques are presented for eliminating the phase jitter problem. Each of the four techniques has been used to accomplish the four research aims detailed in section I.2.

Finally, in Chapter VII, conclusions are drawn from the results of the primary research objectives. This is followed by suggestions for future work on both the desktop MR scanner and the 64-channel receiver.

CHAPTER II

BACKGROUND THEORY

In this chapter, the essential theoretical background for the work undertaken within this dissertation will be presented. As such, the key elements of NMR theory will first be briefly outlined to provide a basis for a discussion of MR imaging. Next, the basic operating principles underlying MR image formation will be discussed. Particular emphasis will be placed upon the equation of the MR signal that is received in a typical imaging experiment. This MR signal equation will lead to a description of the operation and design of the MR receiver. Finally, the methods used to characterize the performance of an MR receiver will be discussed. All of these topics should provide the reader with an understanding of the principles required to design instrumentation for parallel MR imaging. (The reader should note that it is not the author's intention to present anything more than a brief sketch of the fundamental principles which serve as a basis for the instrumentation required to perform parallel MR imaging with arrays of sensor coils. The theory of NMR and MR imaging has been widely discussed in great detail elsewhere in the literature.)

II.1 NMR Theory

The phenomenon of nuclear magnetic resonance (NMR) in bulk materials was first reported by Bloch in 1946 [17]. The theory of NMR has several key components which include the properties of nuclear spin and magnetic dipole moment, the Larmor frequency, bulk magnetization of nuclei within a static magnetic field, and the emission

of the free induction decay signal. This section will provide a brief introduction to the NMR phenomenon and these principle components.

II.1.1 Nuclear Spin and Magnetic Dipole Moment

Atomic nuclei which possess either an odd number of protons, neutrons, or both have certain magnetic properties which include both spin and a magnetic dipole moment. There are several nuclei in nature which meet this qualification, including ^{1}H, ^{13}C, and ^{31}P. (The hydrogen nucleus (or proton), ^{1}H, is of primary importance to MRI due to its natural abundance in biological tissues; hence, the hydrogen proton will be assumed in all discussions of NMR theory contained within this dissertation unless otherwise specified.) When any of these nuclei are placed within a static magnetic field they assume one of $(2I+1)$ possible energy states, known as spin states. Each of these spin states has an associated magnetic dipole moment. (The variable I is known as the spin number and represents the total angular momentum of a nucleus. In general, I is always some positive integer multiple of ½.)

II.1.2 Larmor Frequency

In the presence of a static magnetic field, B_0, the magnetic dipole moment, μ, of each nucleus will experience an aligning torque, τ, due to the static field (2.1). (For the purpose of this dissertation, the direction of the static magnetic field will always be assumed to be in the positive z-direction.) Under the influence of this torque, the magnetic dipole moment of the nucleus will precess, in a left-handed fashion, about the main magnetic field. This, in turn, will cause the z-component of each nuclear magnetic dipole moment vector to align itself either parallel or anti-parallel to the applied

magnetic field. The precession frequency of the nuclear magnetic dipole, or spin, is given by the Larmor relation (2.2) [18, 19]. In (2.2), ω_0 [rad/s] is known as the Larmor (or resonant) frequency of the nucleus, B_0 [T] is the strength of the static magnetic field flux, and γ [rad·T^{-1}·s^{-1}] is the gyromagnetic ratio. For the hydrogen proton, the gyromagnetic ratio, γ, is known to be $2\pi \cdot 42.578 \cdot 10^6$ rad·T^{-1}·s^{-1}.

$$\vec{\tau} = \vec{\mu} \times \vec{B}_0 \tag{2.1}$$

$$\omega_0 = \gamma \cdot B_0 \tag{2.2}$$

II.1.3 Bulk Magnetization

Since the hydrogen nucleus has a spin number of ½, it has only two possible spin states. When a sample containing a collection of hydrogen protons is not in the presence of a static magnetic field (and is in thermal equilibrium), there is no magnetic potential energy associated with either spin state. In this case, the population of spins is equally distributed between the two spin states. However, when a static magnetic field is applied to this same sample, each spin state has an associated magnetic potential energy. For the hydrogen proton, the energy for each spin state, E_m [J], is given by (2.3) [18]. In (2.3), m is the energy level and $\hbar = h/2\pi$, where h is Planck's constant (6.63×10^{-34} [J·s]). The two spin states are additionally differentiated by the orientation of the z-component of their respective magnetic dipole moments. The nuclear magnetic dipoles in the spin ½ state (i.e., m is ½) have a potential energy $E_{½}$ and align themselves parallel to the direction of the static magnetic field. The spins in the -½ state, on the other hand, have an energy of $E_{-½}$ and tend to align themselves anti-parallel to the static magnetic field. A slightly higher proportion of spins in the sample will occupy the spin ½ state since it

has a lower magnetic potential energy. Thus, if there are total of N spins contained within a sample, summation of the magnetic dipole moment vectors, $\vec{\mu}_i$, for all the spins (2.4) will result in a net magnetization vector, $\vec{M}_0$, that lies in the direction of the static field [18]. The magnitude of this net magnetization vector, M_0 [J/T/m^3], for a sample consisting of N hydrogen protons is given in (2.5), where k is the Boltzmann constant (i.e., 1.3807×10^{-23} [J/K]), T is the sample temperature [K], and the remaining variables have been previously defined [18].

$$E_m = -m\hbar\gamma B_0 \qquad m = \pm I = \pm \tfrac{1}{2} \tag{2.3}$$

$$\vec{M}_0 = \sum_{i=1}^{N} \vec{\mu}_i \tag{2.4}$$

$$M_0 = \frac{N\gamma^2\hbar^2 B_0}{4kT} \tag{2.5}$$

II.1.4 Free Induction Decay Signal

Once the magnetic dipole moment vectors of all the nuclei in a sample have aligned with the static field, an equilibrium condition is reached and the net magnetization vector ($\vec{M}_0$) will be aligned with the static field ($\vec{B}_0$). The net magnetization vector (and the spins associated with it) may then be tipped away from alignment with the static magnetic field (and into a plane transverse to the static field) by the application of an external time-varying magnetic field, B_1. This B_1-field must be orthogonal to the static field and have a frequency content that spans the Larmor frequency for the spins in the sample.

The angle that forms between the static magnetic field (B_0) and the net magnetization vector (M_0) during the application of the B_1 field is known as the "tip" angle. When the tip angle is 90°, the entire net magnetization vector is tipped into the transverse plane. Now, if the applied B_1 field is a square pulse (2.6) with duration, τ_p [s], then the tip angle, α [rad], is given by (2.7) [18]. In (2.7) the term B_{1eff} [T] is the effective B_1-field in the rotating frame. In the special case where the applied B_1 field is a linearly polarized square pulse, the magnitude of B_{1eff} may be specified by (2.8).

$$B_1(t) = \begin{cases} B_1 \cos(\omega_0 t) & 0 \le t \le \tau_p \\ 0 & elsewhere \end{cases} \tag{2.6}$$

$$\alpha = (\gamma B_{1eff})\tau_p \tag{2.7}$$

$$B_{1eff} = \frac{B_1}{2} \tag{2.8}$$

After the B_1 field has been applied, the net magnetization vector will again relax back to its equilibrium state while continuing to precess about the main B_0 field at the Larmor rate. The relaxation process, known as "spin-lattice" or "T_1" relaxation, may be described by the exponential relation of (2.9). In (2.9), the z-component of the net magnetization vector, M_z, is given as a function of time, t [s]. Here, the time constant of the relaxation process is T_1 [s]. Finally, M_z^0 is the net magnetization in the direction of the static field at time $t = 0$. (Note: If the sample is at its equilibrium state at time $t = 0$, then the term M_z^0 will be given by the net magnetization (M_0) of (2.5).)

$$M_z(t) = M_z^0(1 - e^{-t/T_1}) \tag{2.9}$$

The precession of the net magnetization vector in the transverse (or x-y) plane will induce an emf (at the Larmor frequency of the nuclei within the sample) inside any tuned RF pickup coil whose polarization vector is also orthogonal to the static magnetic field. This induced voltage is known as a free induction decay signal (or fid), and it serves as the basis for both NMR spectroscopy and MR imaging.

The voltage of the fid signal as a function of time, *v(t)*, has the form of a decaying sinusoid and may be described by (2.10) [18]. Here, t is time [s], j is simply the imaginary number ($\sqrt{-1}$), ω_0 is the Larmor frequency of the spins [rad/s], and ΔV is the volume of the sample [m^3]. The M_{xy}^0 term denotes the magnitude of the net magnetization vector in the transverse plane immediately after the RF excitation has been applied. The e^{-t/T_2} term describes the "spin-spin" or "T_2" relaxation process. T_2-relaxation is used to describe the loss of signal coherence due to a dephasing of the spins which compose the transverse plane magnetization vector. Finally, B_{1t} [T/A] is the effective sensitivity of the detector coil (2.11) [18]. In (2.11), the $\vec{B}_1$ term denotes the flux density vector that would be produced by the detector coil if it were operated in transmit mode, the I_1 term indicates the amount of current necessary to produce $\vec{B}_1$, and $\hat{\rho}_M$ is the unit polarization vector for the net magnetization of the spins in the transverse plane. The unit polarization vector (2.12) for the transverse plane magnetization is circularly polarized due to the fact that the net magnetization precesses about the z-directed static field (B_0). (The terms $\hat{a}_x$ and $\hat{a}_y$ are just unit vectors in the direction of the x- and y-axes, respectively.)

$$v(t) = j\sqrt{2}\omega_0 \Delta V B_{1t} M_{xy}^0 e^{-t/T_2} e^{j\omega_0 t} \tag{2.10}$$

$$B_{1t} = \left(\frac{\vec{B}_1}{I_1} \right) \bullet \hat{\rho}_M \tag{2.11}$$

$$\hat{\rho}_M = \frac{\hat{a}_x + \hat{a}_y}{\sqrt{2}} \tag{2.12}$$

II.2 MR Imaging Basics

The principle of using static magnetic field gradients to encode spatial information into the frequency content of the received NMR signal was first introduced by Lauterbur in 1973 [20]. This idea may be used to obtain an image of a sample that has been placed within a homogeneous, static magnetic field. An MR image is formed by the application of a series of RF and static magnetic field gradient pulses at specified time intervals (a pulse sequence) which interact with the nuclear magnetic dipoles, or spins, contained within the sample. This interaction causes a map or image of the spins within an object to be encoded in the NMR signals. In this section the method by which spatial information is encoded into the NMR signal will be presented. A typical MR pulse sequence will then be described to show how an MR image may be formed using standard 2D Fourier encoding techniques. This will be followed by a description of the MR signal that is received from the imaging sample during the pulse sequence. Afterwards, the MR image reconstruction process will be summarized. Finally, the resolution of the MR image will be presented and equations will be given for determining the signal-to-noise ratio of the received MR signal.

II.2.1 Encoding Spatial Information in the NMR Signal

The spins for a particular nucleus (such as the hydrogen proton) all have the same gyromagnetic ratio. When a sample is placed in a homogeneous, static magnetic field, B_0 [T], (which is assumed to be z-directed) the spins of all the hydrogen protons will have nearly the same precession frequency. (Note: There will be slight variations in this precession frequency throughout the sample due to varying amounts of shielding from local electrons within the molecules that the hydrogen protons inhabit.) If a linear gradient in the static magnetic field is then applied to the sample under investigation, the spins along that field gradient will precess about the main field at different frequencies. For example, if a z-directed magnetic field gradient, G_x [T/m], is applied along the x-axis of the sample, the spins along the x-axis will experience a linear variation in their Larmor frequencies governed by (2.13). It is easy to see that if additional, z-directed, linear, static magnetic field gradients (G_y [T/m] and G_z [T/m]) are imposed along the y- and z-axes over the region containing a sample, the position of each spin within the sample will be mapped to a unique frequency (2.13). In (2.13) and (2.14), position is denoted by x, y, and z [m] while ω_0 and γ are as given in section II.1. It turns out that by selectively applying the static magnetic field gradients (G_x, G_y, and G_z) at specified intervals during the time period in which a series of time-varying B_1 field pulses are applied to a sample (a series of events known as a pulse sequence), the spatial location of that sample's spins may be encoded in the received NMR signal. This encoding of the position to frequency is what is used in MRI to generate an image of the spins contained within a particular sample.

$$\omega(x) = \gamma(B_0 + G_x x) = \omega_0 + \gamma G_x x \tag{2.13}$$

$$\omega(x, y, z) = \gamma(B_0 + G_x x + G_y y + G_z z) = \omega_0 + \gamma G_x x + \gamma G_y y + \gamma G_z z \tag{2.14}$$

II.2.2 Pulse Sequences

Pulse sequences are used to collect the MR image data by scanning the spatial frequency space, or k-space, of a sample's spins. As an example, a simple, 2D Fourier encoded, gradient echo pulse sequence (Fig. 1a) and its corresponding traverse of k-space (Fig. 1b) has been shown. In order to acquire an $N \times M$ pixel image using this gradient echo sequence, N scans of k-space, called phase encode lines, must be acquired sequentially. A single line of k-space is acquired every time the pulse sequence is run from time $t = 0$ to $t = TR$ (Fig. 1a). The received NMR signal, a gradient echo in this case, is digitized in the time period from $t = t_3$ to $t = t_5$ which has a duration of T_{acq}. Note that the center of the received echo signal occurs at time TE (or the "echo time"). M (the number of readout points in the final image) is determined by the strength of the gradient pulse (G_x) applied during the signal acquisition, the length of the acquisition period, and the bandwidth of the pulse sequence. All of k-space is scanned by repeating the pulse sequence while changing the length of the step in k_y (Fig. 1b) that occurs between time $t = t_1$ and $t = t_2$. This change in the length of the k_y-step is accomplished by varying the strength of the y-gradient pulse, G_y (Fig. 1a), from ${}^{+}G_{y\ max}$ to ${}^{-}G_{y\ max}$ in N steps of size ΔG_y (called phase encoding). Since it takes a finite amount of time (TR or the "repetition time") to scan each line of k-space, a total time of ($N \cdot TR$) is required to acquire an $N \times M$ pixel image.

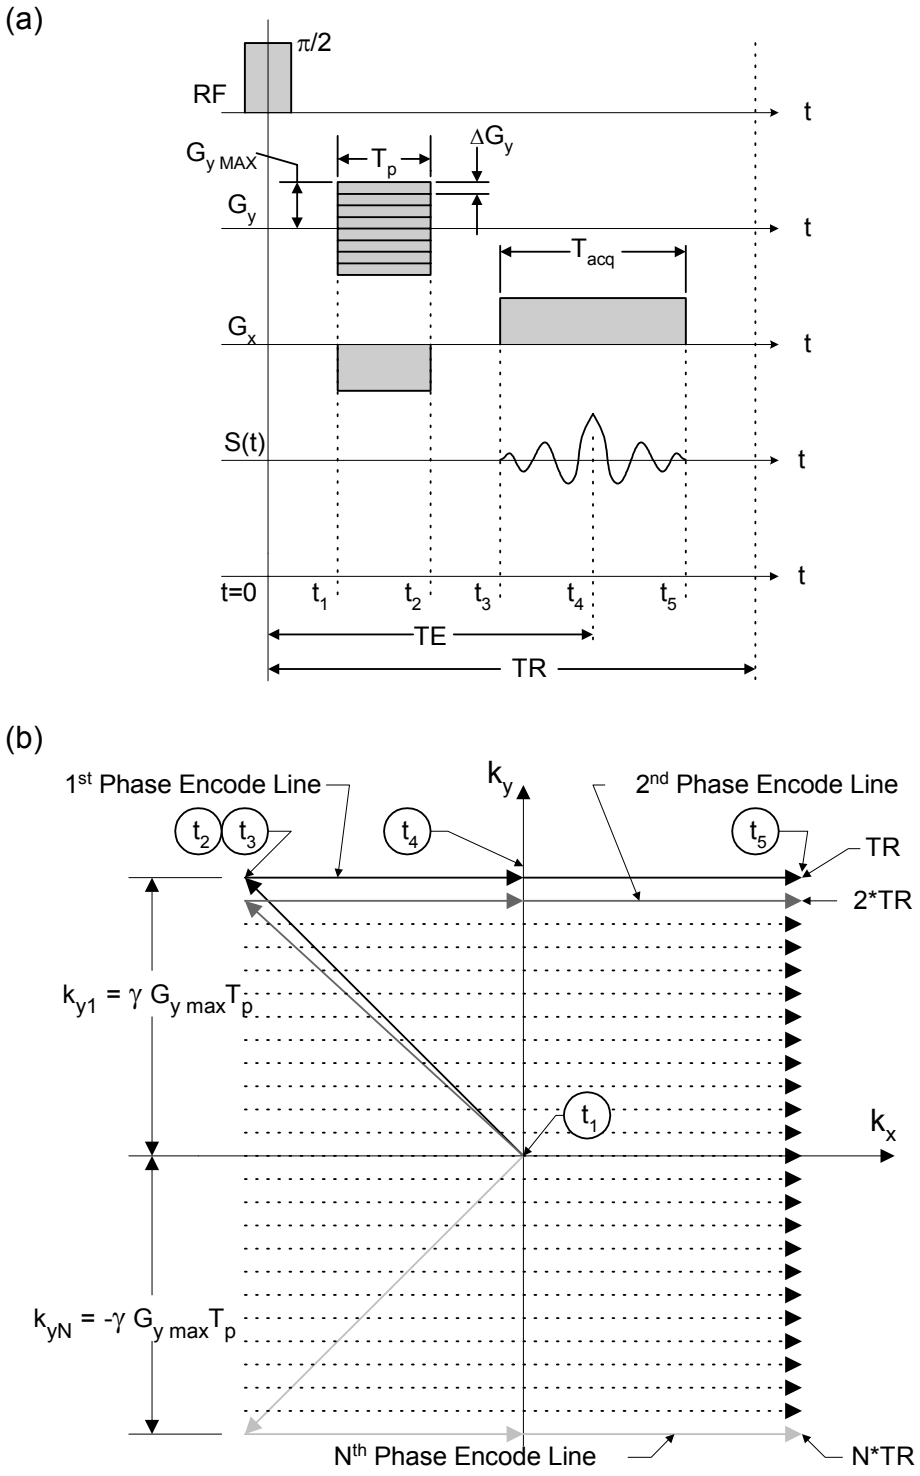

Fig. 1. An example of a 2D Fourier encoded pulse sequence for MR imaging. (a) A simple, 2D Fourier encoded gradient echo pulse sequence. (b) The k-space scan resulting from the pulse sequence in (a).

II.2.3 Received Signal

During a 2D Fourier encoded MR imaging experiment, an NMR signal is acquired for every line of k-space that is scanned. This signal is picked up as an induced emf in an RF detector coil that has been tuned to the resonant frequency of the spins within the imaging sample. The received MR signal itself may be represented by (2.15) assuming that the imaging sample consists of a thin slice of spins in the x-y plane with a constant slice thickness. (A thorough derivation of the signal received in a 2D Fourier encoded MR imaging experiment is widely available in the literature [14]-[15] and will not be given here.) In (2.15) the MR image is represented by the function *I(x,y)* which incorporates the effects of T_2 relaxation, transverse magnetization, and the sensitivity of the receive RF coil (which all vary with position over the imaging sample) as well as constant multiplicative factors due to the Larmor frequency and the slice thickness [18]. Additionally, T_{pe} is the duration of the phase encode gradient pulse while the remaining variables in (2.15) remain as previously defined in (2.2) and (2.14). It should be noted that by performing a change of variables according to (2.16) and (2.17), the expression of (2.15) can be rewritten in (2.18) as a function of the k-space coordinates k_x [rad/m] and k_y [rad/m]. Equation (2.18) shows that the received MR signal $S(k_x,k_y)$ consists of an information signal that has been modulated by a complex carrier with a frequency of ω_0. Finally, it should be noted that the information signal is actually just the 2D Fourier Transform of the image function *I(x,y)*.

$$S(t) = \left[\iint_{slice} I(x,y) e^{j\gamma G_x x t} e^{j\gamma G_y y T_{pe}} \partial x \partial y \right] \cdot e^{j\omega_0 t} \tag{2.15}$$

$$k_x = \frac{\omega_0 \cdot G_x \cdot t}{B_0} = \gamma \cdot G_x \cdot t \tag{2.16}$$

$$k_y = \frac{\omega_0 \cdot G_y \cdot T_{pe}}{B_0} = \gamma \cdot G_y \cdot T_{pe} \tag{2.17}$$

$$S(k_x, k_y) = \left[\iint_{slice} I(x,y) e^{jk_x x} e^{jk_y y} \partial x \partial y \right] \cdot e^{j\omega_0 t} \tag{2.18}$$

II.2.4 Image Reconstruction

The signals received during a 2D Fourier encoded MR imaging experiment contain the Fourier Transform of the MR image modulated by a complex carrier signal. In order to reconstruct the image from these signals, two operations must be performed. First, the signals must be demodulated in order to remove the complex carrier. Mathematically, this may be performed by simply multiplying the received signals by a complex exponential (2.19); however, the demodulation of the received signals is more involved in actual practice (this will be discussed in section II.4.2). Second, an inverse 2D Fourier Transform must be performed on the demodulated signals to reconstruct the MR image of the sample (2.20).

$$S_{demod}(k_x, k_y) = S(k_x, k_y) \cdot e^{-j\omega_0 t} = \iint_{slice} I(x,y) e^{jk_x x} e^{jk_y y} \partial x \partial y \tag{2.19}$$

$$I(x,y) = \iint S_{demod}(k_x, k_y) e^{-jk_x x} e^{-jk_y y} \partial k_x \partial k_y \tag{2.20}$$

II.2.5 Image Resolution

In a Fourier encoded MR imaging experiment each line of k-space is acquired and stored in a 2D matrix having *N* rows (i.e., one for each phase encode line) and *M*

columns (i.e., the signal is digitized and a total of M readout points are collected). The MR image is then reconstructed by performing an $N \times M$, 2D, inverse discrete Fourier Transform (2D-IDFT) on the k-space data set. For faster reconstruction, N and M are generally chosen to be powers of 2 so that fast Fourier Transforms (FFT's) may be used to reconstruct the final image. The resolution and field-of-view in the final MR image is given by (2.21) and (2.22). The terms FOV_x and FOV_y [m] refer to the fields-of-view along the x- and y-axes. Also, Δx and Δy [m] give the resolution in the x and y plane of the image space while Δk_x and Δk_y [rad/m] are the resolution in k-space. Finally, Δt [s] is the sample dwell time for the digitizer that acquires the MR signal and ΔG_y [T/m] is the phase encode gradient step size.

$$FOV_x = M \cdot \Delta x = \frac{2\pi}{\Delta k_x} = \frac{2\pi}{\gamma \cdot G_x \cdot \Delta t} \tag{2.21}$$

$$FOV_y = N \cdot \Delta y = \frac{2\pi}{\Delta k_y} = \frac{2\pi}{\gamma \cdot \Delta Gy \cdot T_{pe}} \tag{2.22}$$

II.2.6 Relaxation Processes and Steady-State Transverse Magnetization

During an MR imaging sequence, the received signal voltage level will vary from that of the fid signal (2.10) due to the presence of various relaxation processes which affect the behavior of the net magnetization vector. The two primary processes, T_1- and T_2-relaxation, have been previously discussed in section II.1.4. In addition to these two, there exist two other relaxation processes known as the T_2^*- and T_2^{**}-relaxations. These last two relaxations are actually just modifications of the T_2-relaxation process to include the effects of magnetic field inhomogeneities and magnetic field gradients. Hence, the T_2^*-relaxation process describes the dephasing of the transverse plane magnetization

(M_{xy}) due both to normal T_2-relaxation and to magnetic field inhomogeneities. Likewise, T_2^{**}-relaxation includes the effects of both T_2^{*}-relaxation and the dephasing from magnetic field gradients. As a general rule, the lengths of the time constants for the four relaxation processes have the following relationship: $T_2^{**} \leq T_2^{*} \leq T_2 \leq T_1$ [19].

Due to the actions of the various relaxation processes, the transverse plane magnetization of a sample during an MR pulse sequence will reach a steady-state value (usually within a few TR cycles) that is less than the net equilibrium magnetization (2.5) for that sample. Expressions for the steady-state transverse magnetization (M_{xy}) for different pulse sequences are tabulated in the literature [18]. Of particular interest for this work are the steady state transverse magnetizations for the spin echo and gradient echo imaging sequences which are given in (2.23) and (2.24), respectively. (The terms α and M_z^0 have been previously defined in section II.1.4. Likewise, the terms *TR* and *TE* have been defined in section II.2.2.) Note that the steady-state transverse magnetization for a gradient echo sequence has the same form as that of the spin echo at a tip angle of 90°.

$$M_{xy} = M_z^0 (1 - e^{-TR/T_1}) e^{-TE/T_2} \tag{2.23}$$

$$M_{xy} = M_z^0 \frac{(1 - e^{-TR/T_1}) e^{-TE/T_2^*} \sin\alpha}{(1 - \cos\alpha \cdot e^{-TR/T_1})} \tag{2.24}$$

II.2.7 Signal-to-Noise Ratio for the MR Signal

An important characteristic of the received MR signal is its voltage signal-to-noise ratio, SNR_v. This may be calculated by dividing the received signal voltage by the noise voltage. The received signal voltage level during an imaging pulse sequence (2.25) may

be found by taking the magnitude of the of the fid signal from (2.10). In this case, however, the value of the transverse magnetization immediately after the RF pulse, M_{xy}^{0}, is replaced by the steady-state transverse magnetization, M_{xy} (refer to section II.2.6). The primary source of noise in the NMR experiment is due to thermal resistance (or Johnson or Nyquist) noise. The r.m.s. level of thermal noise may be found from (2.26) [21], where k is the Boltzmann constant [J/K], T is the absolute temperature of the noise resistance [K], B is the noise bandwidth [Hz], and R is the resistance of the noise source [Ω]. Finally, the voltage SNR of the received MR signal may be computed (2.27).

$$V_{signal} = \sqrt{2}\omega_0 \Delta V M_{xy} B_{1t} \tag{2.25}$$

$$V_{noise} = \sqrt{4kTBR} \tag{2.26}$$

$$SNR_v = \frac{V_{signal}}{V_{noise}} = \frac{\sqrt{2}\omega_0 \Delta V M_{xy} B_{1t}}{\sqrt{4kTBR}} \tag{2.27}$$

II.3 MR Receiver Theory

The design of the receiver is key in MR imaging as it has a major influence on the quality of the final image. This section will explore several factors which play an important part in the theory and design of the MR receiver. These will include the Nyquist criterion for sampled data, MR signal detection, three methods for performing demodulation, a comparison of quadrature and single phase sensitive detectors, the four main single-channel receiver architectures for MR imaging, how to successfully undersample with a digital receiver, and the two primary receiver architectures for parallel MR imaging using arrays of sensors.

II.3.1 Nyquist Criterion for Sampled Signals

Nyquist first presented his criterion for sampling data without aliasing in 1928 [22]. This criterion is often improperly stated as "the minimum sample rate, f_{s_min}, to prevent aliasing in a sampled signal is twice the highest frequency, f_{max}, contained in the signal" (2.28). This version of the Nyquist criterion is intended for lowpass signals (i.e., signals which have frequency content that extends from dc to some upper frequency, f_U). Using the Nyquist criterion of (2.28) will always prevent aliasing in the sampled data stream, but in certain situations (i.e., if the signal to be sampled is bandpass in nature) it will lead to the use of a sample rate that is higher than necessary. A bandpass signal has frequency content which only extends from some lower frequency, f_L, to an upper frequency, f_U, and is zero elsewhere. The bandwidth, Δf, of a bandpass signal may be found by subtracting the lower frequency bound from the upper frequency bound (2.29). For bandpass signals the Nyquist criterion may be restated as "the minimum allowable sampling frequency is twice the bandwidth of the signal" (2.30). When (2.30) is used to determine the minimum allowed sampling rate, the sampling operation is often referred to as bandpass sampling, super-nyquist sampling, or undersampling. Undersampling will be discussed further in section II.3.3.3.

$$f_{s\_\min} \geq 2 \cdot f_{\max} \quad \text{(Lowpass Signal Formulation)} \tag{2.28}$$

$$\Delta f = f_U - f_L \tag{2.29}$$

$$f_{s\_\min} \geq 2 \cdot \Delta f \quad \text{(Bandpass Signal Formulation)} \tag{2.30}$$

II.3.2 MR Signal Detection

As noted in section II.2.3, the voltage produced in an MR receive coil is a complex signal at radio-frequency (RF). This signal consists of an information signal (i.e., the echo or fid) modulated by a complex carrier signal at the Larmor frequency. It is possible to detect the MR signal by direct digitization according to the Nyquist criterion (2.28), that is, by digitizing at a rate which is at least twice the highest frequency contained in the MR signal. However, for most MR imaging applications this is not practical for two reasons. First, lowpass Nyquist sampling would require high digitization rates, and large quantities of data would be generated. For instance, an MR signal with a 200 MHz carrier frequency would have to be digitized at a sample rate of 400 MHz or higher. If the MR signal lasted only 1.6 ms (a 160 kHz bandwidth echo), 640,000 samples would have to be acquired. Furthermore, if the samples were digitized with a 16-bit digitizer, then 1.28 giga-bytes (GB) of data would have to be stored for a single echo. Secondly, lowpass Nyquist sampling is impractical because the MR signal itself is a bandpass signal. In order to make digitization practical, either the carrier frequency of the MR signal must first be shifted to a lower frequency or the MR signal must be undersampled.

II.3.3 Demodulation Methods

The process of translating the carrier frequency of a signal to a lower frequency is known as demodulation or downconversion. Demodulation may be accomplished in three ways, depending upon whether a signal is analog or digital in nature. For analog signals, demodulation is performed by an analog mixer. Likewise, a digital signal may

be demodulated with a digital mixer. The third method of demodulation involves undersampling the analog signal with an analog-to-digital converter (ADC). In this case, the demodulation occurs as the analog signal is transformed into a digital signal.

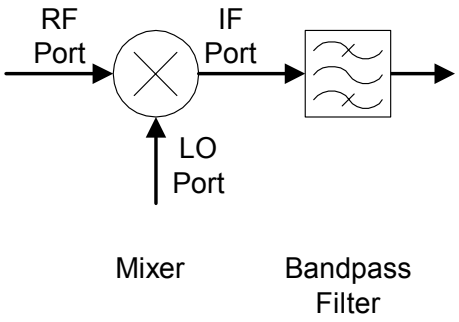

Fig. 2. Block diagram for an ideal analog mixer with an output bandpass filter.

II.3.3.1 Analog Mixer

An analog mixer is a nonlinear device used to perform frequency translation on an input signal. A block diagram for an ideal mixer with an output bandpass filter appears in Fig. 2. The operation of the mixer and filter is described by (2.31) through (2.35). The signals at the RF input port (2.31) and the LO (local oscillator) port (2.32) are multiplied together to produce a signal at the IF (intermediate frequency) output (2.33). Using standard trigonometric identities, (2.33) simplifies to (2.34). Thus, the IF output signal is the sum of two signals with different frequencies: the first (at $f_{RF} + f_{LO}$) is the upper sideband while the second (at $f_{RF} - f_{LO}$) is the lower sideband. Either sideband frequency may be selected by the use of an appropriate bandpass filter after the mixer. When the lower sideband (2.35) is selected, downconversion takes place. In practice,

the ideal analog mixer is not realizable. Real analog mixers do perform the same function as an ideal mixer, but they are plagued with problems such as unwanted harmonics, conversion loss, and distortion [21].

$$s_{RF}(t) = \cos(2\pi f_{RF} t) \tag{2.31}$$

$$s_{LO}(t) = \cos(2\pi f_{LO} t) \tag{2.32}$$

$$s_{IF}(t) = s_{RF}(t) \cdot s_{LO}(t) = \cos(2\pi f_{RF} t) \cdot \cos(2\pi f_{LO} t) \tag{2.33}$$

$$s_{IF}(t) = \frac{1}{2}\{\cos[2\pi(f_{RF} + f_{LO})t] + \cos[2\pi(f_{RF} - f_{LO})t]\} \tag{2.34}$$

$$s_{IF\_LOWER\_SIDEBAND}(t) = \frac{1}{2}\cos[2\pi(f_{RF} - f_{LO})t] \tag{2.35}$$

II.3.3.2 Digital Mixer

A digital signal is comprised of an array of numbers representing samples of an analog signal. To perform demodulation on such a signal, it is only necessary to multiply that signal by a digital sinusoid (i.e., another array of numbers representing samples of a sinusoid). This operation is just an element-by-element multiplication between two numerical arrays and may be thought of as a digital mixer. When the signal to be demodulated has a carrier frequency, f_{RF}, and the digital sinusoid has a frequency, f_{LO}, the operations described by (2.31) to (2.35) are carried out in the digital domain. To select the lower sideband from (2.34), the digital IF signal may be filtered using a digital bandpass filter to form the downconverted digital signal.

II.3.3.3 Undersampling

Undersampling an RF carrier signal with an analog-to-digital converter (ADC) may be used to perform downconversion as well [23, 24]. Consider what happens to the

spectrum of a bandpass signal when it is undersampled (i.e., sampled at frequency lower than the lowpass Nyquist rate of (2.28)) at a frequency, f_s (Fig. 3). From the diagram, it may be seen that when a bandpass signal (having a center frequency, f_o, and a finite bandwidth, Δf) is sampled at this f_s ($f_s << f_o$), aliases of the bandpass signal form which repeat every $f_s/2$ Hz. By selectively filtering out the alias which appears between dc and $f_s/2$ Hz, the original bandpass signal is effectively demodulated to an intermediate frequency, f_{IF}, that lies in the range $0 < f_{IF} < f_s$. Note, that in order for this operation to work, the sample rate must still satisfy the Nyquist criterion for bandpass signals (2.30).

II.3.4 Quadrature versus Single Phase Detection

Once the received MR signal has been downconverted to a reasonable IF frequency, it must be detected. During detection, the complex carrier is removed from the MR signal, leaving only the information signal. The detection process may occur in either the analog or the digital domain depending upon nature of the IF signal. In either case, signal detection may be performed using either a single phase sensitive or a quadrature phase sensitive detector [19]. These detection architectures are depicted in Fig. 4 and Fig. 5. A frequency domain comparison of the two detection schemes is presented in Fig. 6.

II.3.4.1 Single Phase Detection

In single phase detection (Fig. 4), the bandpass IF signal is downconverted to a lowpass signal via a single mixer and lowpass filter. The mixer uses a local oscillator frequency, f_{LO}, that is placed at the lower edge of the IF signal band that extends from f_L to f_U (i.e., $f_{LO} = f_L$) (Fig. 6a). The lowpass MR signal is then filtered with a lowpass filter

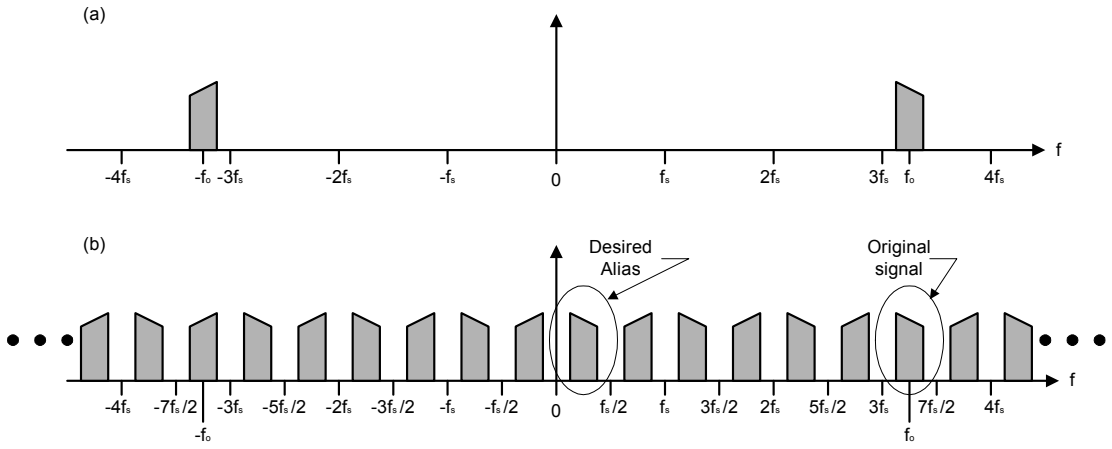

Fig. 3. Undersampling of a bandpass signal. (a) Spectrum of the original bandpass analog signal. Signal band is centered at fo. (b) Aliased spectrum of the sampled analog signal.

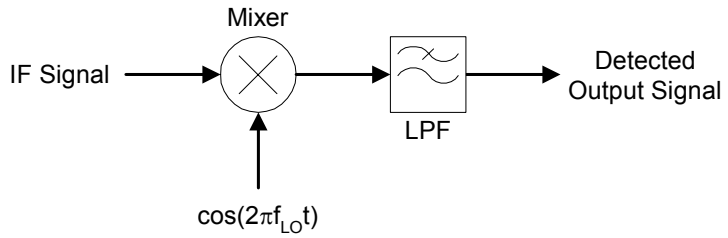

Fig. 4. Single phase sensitve detection architecture.

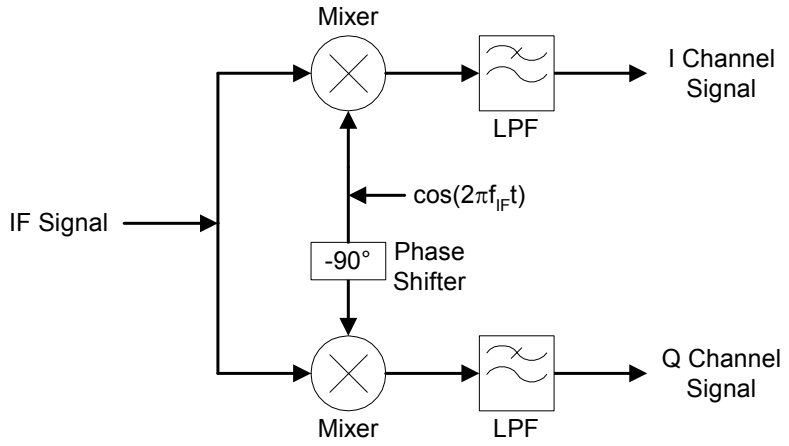

Fig. 5. Quadrature phase sensitive detector architecture.

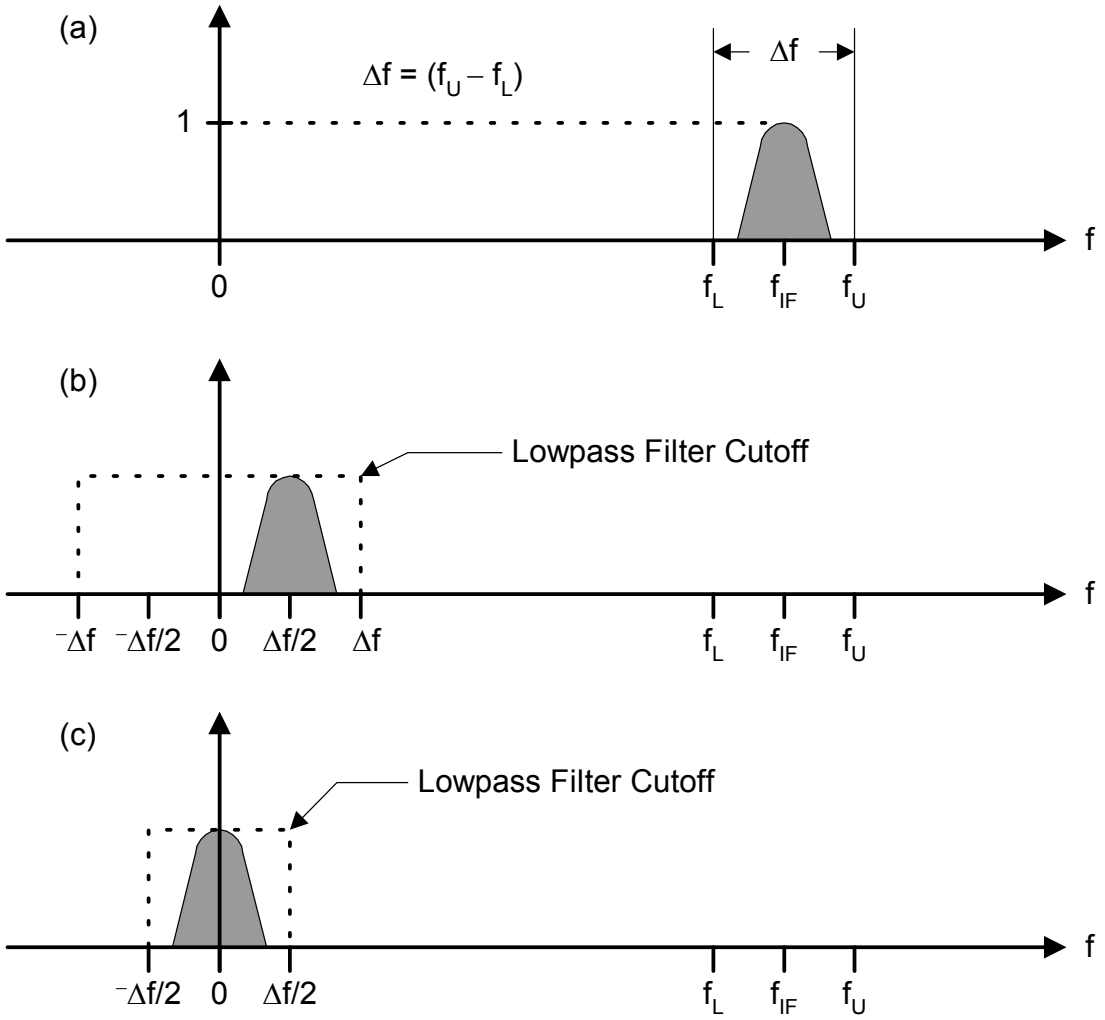

Fig. 6. Frequency domain comparison of single phase and quadrature phase sensitive detection of a hypothetical, IF frequency, MR signal. (a) Spectrum of the MR signal at IF frequency. (b) Spectrum of the MR signal after a single phase detection. (c) Spectrum of the MR signal after quadrature phase sensitive detection.

cutoff frequency of Δf (Fig. 6b). Since the lowpass filter (and, consequently, the single phase detection architecture) cannot distinguish between positive and negative frequencies, the noise bandwidth of the detected signal is actually $2 \cdot \Delta f$ [19].

II.3.4.2 Quadrature Phase Sensitive Detection

In quadrature phase sensitive detection (Fig. 5), the bandpass to lowpass signal conversion is accomplished by an RF splitter, two mixers, and two lowpass filters. The IF signal is first split into two halves and then mixed with two separate local oscillator waveforms: a cosine wave and a sine wave (i.e., a $-90°$ phase shifted version of the cosine wave). This allows the quadrature phase sensitive detector to distinguish between positive and negative frequencies in the downconverted MR signal. The frequency of both local oscillator signals is set to match the IF carrier frequency of the MR signal, causing the middle of the MR signal band to be downconverted to dc (Fig. 6c). Mixing with the cosine wave generates the in-phase (or I-channel) output (corresponding to positive frequencies) while the sine wave is used to create the in-phase quadrature (or Q-channel) output (corresponding to negative frequency content). Since the quadrature phase sensitive detector is able to distinguish between the positive and the negative frequencies in the downconverted signal, the detected signal outputs may be filtered to half the bandwidth of that needed in the single phase detection scheme [19]. This factor of two decrease in the noise bandwidth results in a $\sqrt{2}$ improvement in the detected signal's signal-to-noise ratio (SNR). Quadrature detection is generally preferred in MR imaging due to its lower bandwidth requirement and consequently higher SNR.

II.3.5 Single-Channel Receiver Architectures for MRI

There are several different ways to implement a receiver for MRI. This section will focus on four common receiver architectures. The architectures to be covered are the direct conversion analog receiver [25-27], the multiple conversion analog receiver [27], the single conversion digital receiver [23, 28], and the direct digitization digital receiver [23, 28]. The main difference between these architectures is the number of demodulation stages and the placement of the digitizers. The operation of each receiver and its corresponding advantages and disadvantages will be discussed in detail below.

II.3.5.1 Direct Conversion Analog Receiver

The direct conversion analog receiver (Fig. 7) has a single analog mixing stage and requires two A/D converters for signal detection [25, 26]. First, the MR signal (at the Larmor frequency, f_{RF}) from the low-noise preamp passes through a variable RF gain stage which is used to ensure the proper signal level at the receiver output. Next the signal is split in two so that one half of the MR signal may be mixed with a cosine and the other half with a sine wave (both at the Larmor frequency, f_{RF}). This will form the in-phase (I channel) and in-phase quadrature (Q channel) baseband signals. Note that the sine wave is generated by phase shifting the cosine wave by $^{-}90°$ and that the mixing signal must be phase locked to the frequency timebase of the MR scanner. Both the I and Q signals are then lowpass filtered to half the bandwidth of the MR signal (this is known as anti-alias filtering the signal to ensure that no frequency in the sampled signal is higher than half the sample rate). Finally, the I and Q signals are digitized by two separate, 16-bit ADC's at the sample rate, f_s, to generate the I and Q digital output data.

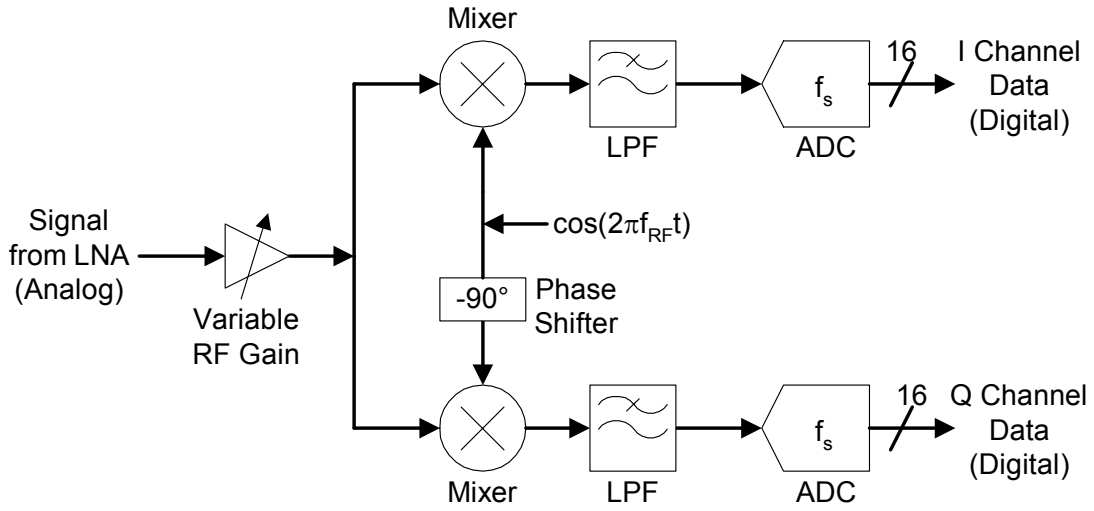

Fig. 7. Block diagram for a direct conversion analog receiver. Analog signal pathways are indicated with arrows. Digital signal paths are shown using arrows with a crossbar to denote a digital signal bus. Bit width of the bus is indicated by the number printed adjacent to each crossbar.

This output data forms the real (I) and imaginary (Q) parts of the received k-space data. The data is stored and later reconstructed to form an MR image.

II.3.5.1.1 Advantages

There are three main advantages to the direct conversion analog receiver architecture. First, this architecture is probably the simplest and least expensive of the four to implement. Second, this receiver presents a lower processing load to the control computer. Third, since the receiver is primarily analog, it can use a lower cost ADC chip. These three advantages will be discussed in the next paragraphs.

The primary advantage of the direct conversion analog architecture is its simplicity of implementation. It requires very few components; in fact, the portion of the receiver from the variable RF gain stage to the low-pass filters may be implemented on a single chip. It is this simplicity that makes the direct conversion analog receiver particularly inexpensive to implement.

The direct conversion architecture also requires less processing of the I and Q channel data for reconstruction of the MR image. The data is generally sampled at the appropriate rate for the bandwidth of the imaging sequence. In some receivers it may be oversampled to lessen the effects of trigger jitter, but the decimation of the data necessitated by this method does not place much of a load on the host processor.

A final advantage of the direct conversion receiver is that it requires less expensive ADC's than a comparable digital receiver. The lower f_s requirement (the signal being sampled is already at baseband) means that lower cost ADC's with lower available sample rates and input bandwidths may be used.

II.3.5.1.2 Disadvantages

The direct conversion analog receiver also has several disadvantages. Four of these stem from the architectures inherent simplicity while another two are due to the analog nature of the receiver. These former set of problems include dc offset error, I/Q mismatch, the necessity for two ADC chips, and even order distortion effects. The other two problems are due to the extensive use of analog components in this architecture. Each of these six disadvantages will be discussed.

Using the direct conversion scheme necessitates the generation of local oscillator signals (described above) at the Larmor frequency. It is possible for these signals to leak back up the receiver chain into the inputs of the preamp, the RF gain stages, and the analog mixer. It is also possible for the MR signal to leak through to the LO (local oscillator) port of the mixer. This phenomenon is known as LO leakage and will result in a DC offset at the mixer outputs which can corrupt the demodulated MR signal.

The problem of I/Q mismatch occurs when there is an amplitude mismatch between either the sine and cosine waves used to mix the input MR signal to baseband or in the two halves of the MR signal themselves. This will lead to the appearance of a quadrature artifact in the final MR image.

Another problem for the direct conversion architecture is that two ADC's are required. This is because I/Q demodulation of the received signal is performed entirely in the analog domain. Even though this architecture can use less expensive ADC circuits (see section II.3.5.1.1), the need for two ADC's can still drive up the cost of implementation.

The direct conversion architecture is also susceptible to problems with even-order distortion effects. This results when two strong interferers, having frequencies f_1 and f_2, are present close to the Larmor frequency. These interference signals will mix with the MR signal to produce a low frequency output at a frequency given by (f_1 - f_2). This undesired mixing product can then corrupt the demodulated MR signal at the receiver output.

The direct conversion receiver also suffers because of its intensive use of analog components. The receiver performs its demodulation and front-end processing functions in the analog domain. This means that more analog control signals and components are required, increasing the complexity of the receiver implementation. These analog components are susceptible to variations in performance due to thermal drift over time, causing them to have a lower degree of repeatability than their digital counterparts.

Finally, a conventional analog receiver (such as the direct conversion receiver) is less flexible than comparable digital implementations. For instance, to switch receive bandwidths on the lowpass filter sections, variable cutoff frequency filters must be used, which in turn increases the receiver complexity. In a digital receiver, the lowpass filters can be changed in software. This is a common problem among analog circuits when compared to a comparable digital circuit. The digital implementations tend to be more flexible as they can usually just be reprogrammed to change their operating characteristics. Similar changes for an analog circuit usually require modification fo the circuit hardware.

II.3.5.2 Multiple Conversion (Superheterodyne) Analog Receiver

The multiple conversion (or superheterodyne) receiver of Fig. 8 has multiple analog mixing stages and also needs two ADC's for signal detection [27]. The input MR signal first passes through an image reject bandpass filter followed by an RF gain stage. The image reject filter prevents unwanted noise at the at the image frequency of the mixer ($2 \cdot f_{LO} - f_{RF}$) from aliasing into the IF signal bandwidth. The signal is then mixed down to an intermediate frequency where it is again amplified and bandpass filtered. Similar to the direct conversion scheme, the signal is then split and I/Q demodulated to baseband where it is digitized, stored, and reconstructed to form an MR image. Also, the oscillator frequencies (f_{LO} and f_{IF}) used to mix the RF and IF signals must both be phase-locked to the timebase of the MR scanner.

II.3.5.2.1 Advantages

The multiple conversion analog receiver architecture has five main advantages. The first two are improvements on the design of the direct conversion receiver to eliminate some of its problems. The remaining three are similar to those of the direct conversion receiver. First, the multiple conversion architecture only has a minimal amount of dc offset error. The use of the IF bandpass filter in conjunction with the IF gain stages minimizes the problems of DC offset due to LO leakage in this receiver. Second, there are no problems with even-order distortion in the multiple conversion receiver. Like the direct conversion receiver, the multiple conversion architecture presents a low processing load on the control computer since I/Q demodulation is performed in hardware and the received signals are digitized at baseband frequencies. The multiple

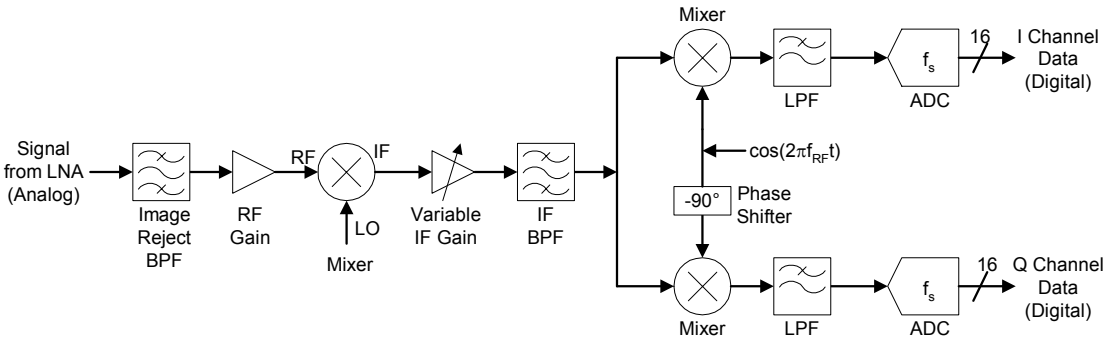

Fig. 8. Block diagram for a multiple conversion analog receiver. Analog signal pathways are indicated with arrows. Digital signal paths are shown using arrows with a crossbar to denote a digital signal bus. Bit width of the bus is indicated by the number printed adjacent to each crossbar.

conversion receiver is also relatively simple and inexpensive to implement. While more complicated than the direct conversion architecture, most of this receiver (from the first mixer to the I and Q outputs of the lowpass filters) can be implemented in a reliable fashion on a single chip (e.g., the Analog Devices AD607 single chip receiver [29]). Finally, like the direct conversion receiver, the multiple conversion architecture can use less expensive ADC circuits.

II.3.5.2.2 Disadvantages

The multiple conversion architecture suffers from many of the disadvantages that plague the direct conversion architecture. First, the multiple conversion receiver has an even higher analog component count than the direct conversion receiver. Thus, it suffers the same problems associated with analog components as the direct conversion receiver. Second, the receiver is also troubled by I/Q mismatch errors. The multiple conversion design also requires two ADC circuits. Finally, the receiver has a lower level of flexibility than either of the digital receiver architectures.

II.3.5.3 Single Conversion Digital Receiver (IF Digitization)

The single conversion digital receiver is similar to the conventional I/Q dual conversion receiver [23, 28]. The main difference is in the placement of the ADC (only one is required) and in the number of analog components required to implement this receiver (fewer for the digital receiver). This receiver has only one analog mixing stage, but may have multiple digital demodulation stages. As may be seen in Fig. 9, the input RF signal is image-reject filtered, amplified, mixed to an IF, amplified again, and then bandpass filtered. Next, the IF signal is digitized at f_s by a 16-bit ADC, after which the

sampled IF signal is I/Q demodulated to baseband and finally lowpass filtered. Note that the IF signal may be undersampled by the ADC, in which case the signal will be downconverted to a second IF frequency by the undersampling process. The single conversion digital receiver implements the downconversion of the IF signal to baseband in the digital domain (after the IF signal has been sampled) using numerical equivalents of the analog splitter, mixers, and lowpass filters found in the conventional analog receiver (described in the section II.3.2.2).

II.3.5.3.1 Advantages

The single conversion digital receiver has five distinct advantages over the other receiver architectures. First, this receiver requires a smaller number of analog components than either of the analog receivers, making it easier to implement. Secondly, the receiver needs only one local oscillator signal instead of two or more (as required by the analog receivers). Third, the single conversion digital receiver (unlike its analog counterparts) is not prone to I/Q mismatch since the sine and cosine functions used to mix the IF signal to baseband have identical amplitudes. In fact, the digital mixers themselves for both the I and Q channels are identical. This ensures that there will be no problems with quadrature artifacts in the final MR images. Fourth, this receiver requires only one ADC since I/Q demodulation is performed digitally. This further reduces the hardware complexity for the receiver. Finally, the single conversion digital receiver is capable of operating over a wide range of RF input frequencies. This may be accomplished by removing the front-end image reject filter. (The image reject filter on both of the analog receivers can be removed to achieve this as well, but the

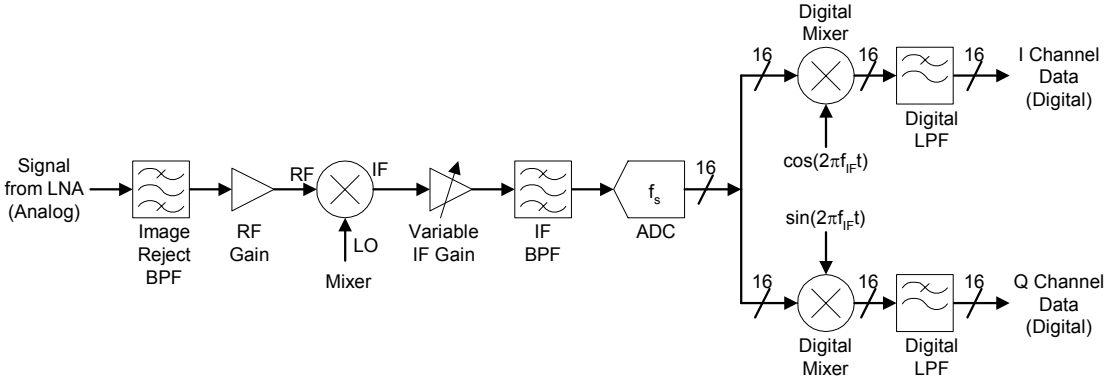

Fig. 9. Block diagram for a single conversion digital receiver. Analog signal pathways are indicated with arrows. Digital signal paths are shown using arrows with a crossbar to denote a digital signal bus. Bit width of the bus is indicated by the number printed adjacent to each crossbar.

same cannot be said for a direct digital receiver.) Unfortunately, removing the image reject filter will also decrease the available signal-to-noise ratio from the receiver by a factor of $\sqrt{2}$ as compared to a similar receiver with an image reject filter.

II.3.5.3.2 Disadvantages

Overall, the single conversion digital receiver has three main disadvantages. First, the receiver needs higher performance (and usually more expensive) components than the analog architectures. This is especially true if the IF frequency is to be undersampled as tighter constraints will be placed on both the ADC and the IF bandpass filter. Second, the receiver presents a heavier processing load to the host computer since I/Q demodulation is performed entirely in the digital domain. This means that there will be increased data storage requirements and longer data transfer times. Finally, synchronizing the digitizer to the MR scanner is a problem. If the sample clock on the ADC is not phase-locked to the reference frequency on the MR scanner, jitter will occur at the start of each digitization event. This will show up in the reconstructed MR image as phase jitter.

II.3.5.4 Direct Digital Receiver

The direct digital receiver (Fig. 10) needs no analog mixing stages and requires only one ADC [23, 28]. It bears many similarities to both the direct conversion analog receiver and to the single conversion digital receiver. The input RF signal is amplified (to raise the signal to a level which the ADC will detect), bandpass filtered, and then sampled by a 16-bit ADC at the rate f_s. Depending upon the frequency of the RF input signal and the capability of the ADC chip, the ADC may either directly sample or

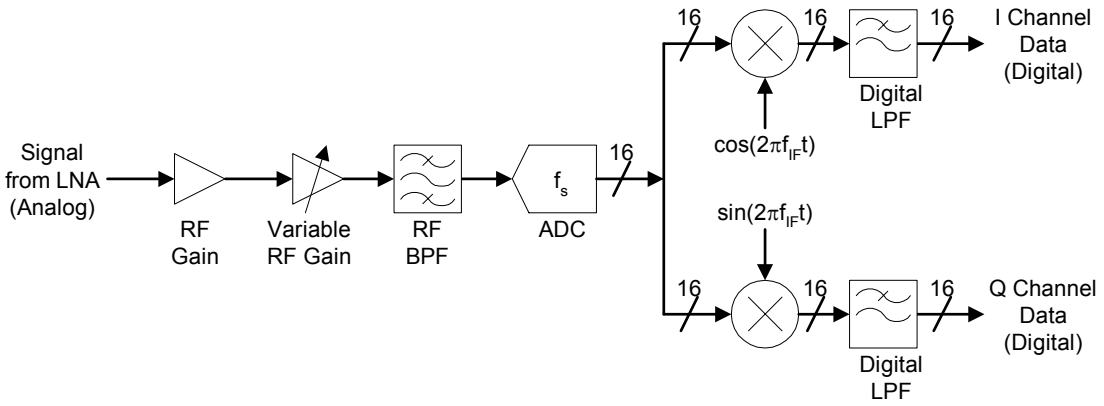

Fig. 10. Block diagram for a direct digital receiver. Analog signal pathways are indicated with arrows. Digital signal paths are shown using arrows with a crossbar to denote a digital signal bus. Bit width of the bus is indicated by the number printed adjacent to each crossbar.

undersample the RF input the signal. For lower frequency input signals ($f_{RF} \leq 10$ MHz), direct sampling of the carrier signal may be possible. For high frequency inputs ($f_{RF} >$ 10 MHz), however, the ADC must undersample the RF signal, thereby downconverting the RF signal to an IF frequency. This means that the RF bandpass filter and the ADC circuit must meet tighter constraints than those of the single conversion receiver. The final section of the direct digital receiver (which occurs after the ADC and includes the I/Q demodulation and lowpass filtering stages) is identical to that of the single conversion digital receiver.

II.3.5.4.1 Advantages

The advantages of the direct digital receiver are numerous. First, this architecture requires fewer analog components than any of the other three receivers. Thus, it is the easiest receiver to implement out of the four architectures considered. Also, since it is almost entirely digital, this receiver is more flexible than the two analog receivers. Finally, like the single conversion digital receiver, the direct digital receiver requires no LO signals, does not suffer from I/Q mismatch, and requires only a single ADC.

II.3.5.4.2 Disadvantages

While the direct digital receiver has many benefits, it is plagued by four problems. First the direct digital receiver requires higher performance components. Second, this architecture also presents a heavy processing load on the control computer. The direct digital receiver is also harder to synchronize to the timebase of the MR scanner. Finally, the direct digital architecture can only operate at a single fixed frequency.

The direct digital receiver requires a higher performance ADC than that used by the single conversion digital receiver. This is due to the fact that the receiver is directly undersampling the input RF signal. As a result, the track-and-hold network on the ADC must have an even higher input bandwidth and less aperture jitter than that required by the single conversion digital architecture. This also heightens performance requirements for the RF bandpass filter on the receiver front end in order to prevent unwanted out-of-band noise from aliasing into the undersampled signal.

Of the four receiver architectures, the direct digital receiver has the potential to present the heaviest load to the processor in the control computer. This is especially true if the ADC is used to directly sample a lower frequency RF carrier signal (whose frequency might be greater than the actual bandwidth of the signal). In such a case, higher sample rates would be required than if the input signal were just undersampled. Thus, the processing problems associated with the single conversion digital receiver would only worsen for the direct digital receiver. If the receiver's ADC is used in undersampling mode, then the load on the processor will be essentially the same as that of the single conversion digital receiver.

The problem of synchronizing the digitizer to the MR scanner on the direct digital receiver is similar to that of the single conversion digital receiver. In order to prevent phase jitter in the reconstructed MR image, the sample clock on the ADC must be phase-locked to the timebase on the MR scanner. Of course, the higher the sample rate on the ADC, the more difficult it is to properly phase-lock these two signals.

Finally, the direct conversion digital receiver can only operate at a single input frequency band. Unlike the single conversion receiver, the RF bandpass filter cannot be removed to increase the receiver operating frequency range. Thus, a change in the operating frequency requires that the RF bandpass filter be changed.

II.3.6 Undersampling Issues for the Digital Receiver

In a digital receiver, the A/D converter is used to implement the receiver's downconversion function. However, for a given ADC to work in this application, it must meet three requirements. First, the input bandwidth of the ADC track-and-hold circuit must be greater than the carrier frequency plus one half the bandwidth of the signal to be demodulated (i.e., it must meet the Nyquist criterion for lowpass signals (2.28)). Second, the ADC must be capable of sampling at the desired sample rate, f_s. Third, the sampling aperture error of the ADC must be low to prevent degradation of the signal being demodulated.

II.3.6.1 Track-and-Hold Input Bandwidth

The bandwidth of the ADC track-and-hold circuit is crucial to an ADC's success in an undersampling application. If the bandwidth specified for an ADC's track-and-hold circuit is lower than the highest frequency in the signal being undersampled, those frequency components above that bandwidth will be attenuated in the sampled signal. This can lead to significant degradation in the downconverted signal and can be problematic for MRI applications. This is because few 16-bit ADC's have front-end bandwidths higher than 10 MHz (whereas, most magnet frequencies are higher than this). To overcome this limitation, an external sample-and-hold (or track-and-hold) unit

may be used in front of the ADC. Several manufacturers produce track/sample-and-hold amplifiers with front-end bandwidths up to approximately 400 MHz [30-32]. These external track/sample-and-hold units allow ADC's with slower front-end track-and-hold circuits to be used in undersampling applications.

II.3.6.2 ADC Sample Rate Limitations

As mentioned above, for an ADC to perform in an undersampling application it must be able to sample at the desired sampling rate, f_s. There are, however, certain restrictions on the selection of f_s which must be followed for a signal to be undersampled properly. First, the minimum allowable sample rate, f_{s_min}, must satisfy the Nyquist criterion for bandpass signals (refer to (2.30) in section II.3.1). From the diagram in Fig. 3, it is also apparent that f_s must be carefully chosen so that the frequency band of the signal to be undersampled does not intersect with any of the aliased signal bands whose boundaries lie at (2.36):

$$\left(\frac{N \cdot f_s}{2}\right), \quad \text{where}: N = 0, \pm 1, \pm 2, \ldots . \tag{2.36}$$

II.3.6.3 Sampling Aperture Error

Finally, sampling aperture error is the phase noise in a sampled signal which results from jitter at the clock edge of the sampling clock and from the aperture error of the ADC track-and-hold circuit itself. The aperture error for a given ADC may be predicted by the approximation of (2.37):

$$t_a \cong \frac{1}{2^n \cdot \pi \cdot f_{\max}} \tag{2.37},$$

where n is the number of bits on the ADC, t_a is the aperture error [s], and f_{max} is the maximum frequency [Hz] to be sampled [23]. Thus, for less than one least significant bit of degradation to occur when sampling with an n-bit ADC, thetrack-and-hold aperture error for that ADC must be less than the value predicted by (2.37). Aperture errors are specified in the data sheets for all ADC's, allowing an appropriate ADC to be selected for a particular application.

II.3.7 Receiver Architectures for Parallel MR Imaging

Each of the four MR receivers presented so far (section II.3.4) has been a single-channel architecture. This means that only the MR signal from a single detector coil may be received at a time by the receiver. To perform parallel MR imaging (via an array of RF sensor coils as discussed in section I.1.2), however, a receiver must be capable of receiving the signals from multiple detector coils simultaneously. MR receivers able to perform this feat (i.e., receive the signals from multiple coils in parallel) are called parallel receivers. There are two main receiver architectures used to perform parallel MR imaging: the time domain multiplexed parallel receiver and the true parallel receiver [33].

II.3.7.1 Time Domain Multiplexed Parallel Receiver

The time domain multiplexed (TDM) receiver architecture relies upon the use of analog, RF switches to multiplex the signals from multiple sensor coils onto a single channel. The multiplexed signal is then demodulated using a standard, single-channel receiver and digitized. The signal from each sensor coil may then be separated out in post-processing. It should be noted that the TDM architecture is actually only quasi-

parallel as it cannot receive the signals from multiple coils simultaneously; however, this is not a serious limitation for MR imaging.

A major advantage of the TDM parallel receiver architecture is its relative simplicity and ease of implementation. In fact, the first ever 16-channel MR receiver was created in the Magnetic Resonance Systems Laboratory at Texas A&M University using the TDM architecture [16]. Unfortunately, the TDM architecture is limited by the wide bandwidth of a multiplexed MR signal. If N array coils are used to create a multiplexed signal, the bandwidth of that signal increases by $2N$ times the bandwidth of the MR imaging sequence [16]. Since most MR receivers have a limited operating bandwidth, this places a constraint on the bandwith of the MR sequences that may be used with a given N-channel TDM parallel reciever. This problem only worsens as the number of array coils (and thus the channel count) increases.

II.3.7.2 True Parallel Receiver

A true parallel MR receiver is formed by simply reproducing one of the single-channel receiver designs (see section II.3.4) to form a parallel receiver with as many channels as are needed for a particular application. For example, if a 64-channel true parallel receiver were to be formed based upon the single conversion digital receiver, sixty-four copies of that receiver (presented in section II.3.4.3) would need to be constructed. This would enable the signals from an array of up to sixty-four sensor elements to be received simultaneously. Each channel in the true parallel receiver would then possess all the advantages (and disadvantages) of the single-channel receiver architecture upon which the true parallel receiver was based (without the bandwidth

limitations that arise from time domain multiplexing). The main limitation of the true parallel architecture is that its implementation, especially for high channel count receivers, can be both costly and cumbersome [16].

II.4 MR Receiver Performance Metrics

There are five primary measures of performance that are of concern for the designer of an MR receiver. These metrics are noise figure, dynamic range, image signal-to-noise ratio (SNR), interchannel isolation, and channel-to-channel gain variation. The first three measures apply to both single-channel and parallel receiver architectures while the last two are only important for the parallel receiver. Simply stated, the noise figure for a receiver is a measure of how much the signal-to-noise ratio of an input signal is reduced by the components in that receiver. The dynamic range of a receiver, on the other hand, defines the region over which the output power of the received signal is linearly proportional to that of the input signal power. The image SNR is one measure of the quality of the reconstructed MR image that a given receiver produces. Interchannel isolation measures the amount of crosstalk that is possible between channels in a parallel receiver. Finally, channel-to-channel gain variation is used to quantify how much the performance of a parallel receiver varies across all of its channels. Each of these five performance metrics will be discussed in this section.

II.4.1 Noise Figure

The noise figure of a two port system is defined as the ratio of the input SNR (signal-to-noise ratio) to the output SNR of the system. A receiver is comprised of multiple two-port subsystems, or components, connected in cascade fashion, and each of these

components have an associated gain and noise figure (Fig. 11). The overall noise figure for the n-component receiver system of may be computed by (2.38):

$$F = F_1 + \frac{F_2 - 1}{G_1} + \frac{F_3 - 1}{G_1 \cdot G_2} + ... + \frac{F_n - 1}{G_1 \cdot G_2 \cdots G_{n-1}} \quad (2.38),$$

where F is the overall noise figure of the receiver, F_1 through F_n are the individual component noise figures, and G_1 through G_{n-1} are the individual component gains (all gains and noise figures must be in ratio form, not in dB) [21]. As is apparent in (2.38), the gain and noise figure of the first stage of the receiver system may be made to dominate the noise figure of the overall receiver if its gain is large and its noise figure is sufficiently small. In an MR system, the low-noise preamp is the first stage of the receiver. The noise figure of the preamp is typically less than 1.0 dB while its gain is usually between 20 and 40 dB. This causes the overall receiver noise figure to not be significantly higher than that of the preamplifier circuit.

II.4.2 Dynamic Range

The dynamic range defines the linear operating region (i.e., where the output power is linearly proportional to the input power) for a given receiver. In general, a high dynamic range is desired for a receiver so that it may operate over a large range of input signal levels [21]. The dynamic range itself is usually specified as the range of power levels, in dB, between the 1-dB compression point and the minimum detectable signal (MDS) level [21]. Hence, the 1-dB compression point and the MDS are key factors in determining the dynamic range of a receiver. In addition to 1-dB compression point and the MDS, processing gain also has a large influence on the dynamic range of a receiver,

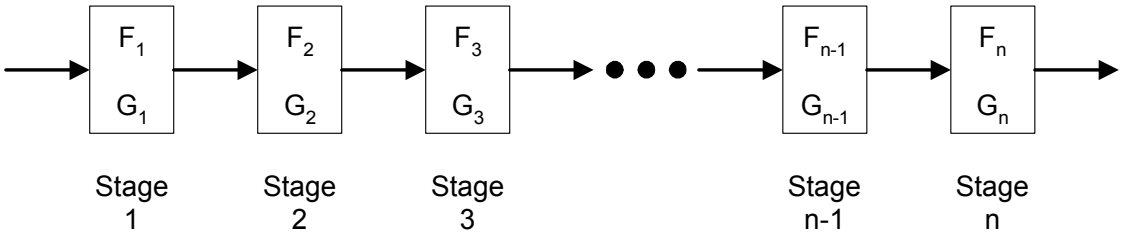

Fig. 11. Noise figure parameters for an n-stage cascaded system.

especially when the dynamic range of that receiver is measured. The effects of these three factors on dynamic range will be examined, and methods for estimating and measuring the dynamic range will be presented.

II.4.2.1 1-dB Compression Point

The 1-dB compression point is the input power level which causes the output power of the receiver to deviate by 1 dB from linearity. This is generally considered to be the maximum detectable signal level for a receiver. The 1-dB compression point for a receiver system is found by considering the gain and 1-dB compression point for each component within the receiver chain. (A 1-dB compression point is specified for each component by its manufacturer.)

II.4.2.2 Minimum Detectable Signal (MDS) Level

The MDS level may be found by summing the input noise level to the receiver (i.e., the noise floor), the receiver noise figure, and a detection threshold (2.39) [21]. The thermal noise floor (i.e., the r.m.s. power level of the ambient thermal noise) from a matched resistive load may be found using (2.40) [21]. In (2.40) the noise floor is computed in Watts, k is the Boltzmann constant, T is the ambient temperature [K], and B is the receiver bandwidth [Hz]. From (2.40), the noise floor for a 1 MHz bandwidth (in a matched, 50 Ω system) at 300 K is -114 dBm. The detection threshold may be arbitrarily set by the system designer depending on the application. For example, in radar, the detection threshold may be set as high as 10 to 16 dB in order to increase the probability of detection and lower false alarm rates [21]. For the MR receiver, however, the detection threshold may be set to a lower value, such as 3 dB.

$$MDS = Noise\ Floor + Noise\ Figure + Detection\ Threshold \tag{2.39}$$

$$Noise\ Floor = kTB \tag{2.40}$$

II.4.2.3 Processing Gain

In systems where an analog signal is converted to a digital signal by means of a digitizer, the dynamic range measurement can be increased by a factor known as processing gain. Processing gain is an improvement in the SNR of a signal due to digital signal processing techniques such as digital filtering, oversampling, coherent integration, or non-coherent integration [34]. This boost in SNR affects the overall dynamic range of a system by lowering the MDS level (i.e., processing gain effectively lowers the noise floor). Each of the MR receiver architectures presented so far are affected by processing gain. This section will explore the effects of coherent integration, non-coherent integration, and oversampling on the dynamic range of an MR receiver.

II.4.2.3.1 Coherent Integration

Coherent integration is a processing effect which occurs when the observation time of a coherent signal is increased. For example, the SNR for a CW (continuous wave) signal will increase by 3 dB when its observation time is doubled. If the coherent signal is digitized with n samples, then the effective coherent integration gain in dB is given by (2.41) [34]. In order for this coherent integration gain to occur, the observed signal must occupy only a small fraction of the total analysis bandwidth [34]. During normal MR imaging operations, however, the MR signal tends to occupy a significant portion of the analysis bandwidth (i.e., for a given imaging sequence, only the received signal

bandwidth is digitized). As a result, coherent integration only affects the dynamic range of an MR receiver when its single-tone dynamic range is measured.

$$Coherent\ Integration\ Gain = 10 \cdot \log_{10}(n) \tag{2.41}$$

II.4.2.3.2 Non-Coherent Integration

Non-coherent integration gain is obtained when several signal estimates are averaged together from successive time slices [34]. This effect is generally referred to as "signal averaging" in MR imaging. The result of averaging *N* MR signals acquired at different times is to increase the SNR by a factor of $\sqrt{N}$. Hence, signal averaging is commonly used in MR imaging to increase the dynamic range of a receiver.

To understand the processing gain incurred by signal averaging, consider a sequence of random variables ($\mathbf{x}_i$) formed by collecting samples of a signal in additive white noise at different times (i.e., the noise in each of the samples has a zero mean and is uncorrelated). Also let these samples have the same mean ($E\{\mathbf{x}_i\} = \eta$) and variance ($\sigma_i^2 = \sigma^2$). The variance of the average of the samples will then be given by (2.42) while the expected value of the average will be given by (2.43) [35]. Thus, the effect of averaging *N* samples of this type is to reduce the variance of the noise by a factor of *N*. Now, recall that thermal noise is the primary noise source in MR imaging and that thermal noise is a zero mean, white noise process. The variance of thermal noise (2.44) is given by the square of the r.m.s noise voltage from (2.26) [21]. Also, the average value of *N* samples of an MR signal (2.45) is just the steady-state voltage level for that signal (2.25). Thus, the total SNR obtained by averaging *N* signals is greater by a factor

of $\sqrt{N}$ over the SNR that might be obtained from any of the individual samples (SNR_i) (2.46).

$$\sigma_{\bar{\mathbf{x}}}^2 = \frac{1}{N^2}\sum_{i=1}^{N}\sigma_i^2 = \frac{\sigma^2}{N} \tag{2.42}$$

$$E\{\bar{\mathbf{x}}\} = E\{\mathbf{x}_i\} = \eta \tag{2.43}$$

$$\sigma^2 = V_{noise}^2 = 4kTBR \tag{2.44}$$

$$E\{\bar{\mathbf{x}}\} = \sqrt{2}\omega_0 \Delta V B_{1t} M_{xy} \tag{2.45}$$

$$SNR_{total} = \frac{E\{\bar{\mathbf{x}}\}}{\sqrt{\frac{\sigma^2}{N}}} = \sqrt{N}\left(\frac{\sqrt{2}\omega_0 \Delta V B_{it} M_{xy}}{\sqrt{4kTBR}}\right) = \sqrt{N}\cdot SNR_i \tag{2.46}$$

II.4.2.3.3 Oversampling Gain

Oversampling gain occurs when a received signal is first digitized at a rate higher than twice the actual signal bandwidth, then digitally filtered to a narrower bandwidth, and finally decimated down to a lower sample rate [36, 37]. The ratio of the oversampling rate to the decimated sample rate is called the oversampling ratio. The oversampling gain (in dB) is the just the logarithm of the oversampling ratio (2.47) [36, 37]. Thus, if a 1 MHz bandwidth signal is first oversampled at 20 MHz and then decimated to a sample rate of 2.5 MHz, the oversampling ratio would be 8 and the oversampling gain 9 dB. It would seem from (2.47) that dramatic increases in receiver dynamic range would be possible through oversampling. The relation of (2.47), however, only holds true if the received noise voltage level is on the order of the least significant bit (LSB) of the digitizer [36]. Unfortunately, for most MR imaging

experiments, the noise level of from the receiver coil is amplified to a level on the order of several LSB's of the digitizer [36, 38], thereby invalidating (2.47). As a result, oversampling does not have as much of an effect on the dynamic range of the MR receiver. Some improvement in dynamic range will occur (most likely due to digital filtering), but not to the extent predicted by (2.47) [36].

$$Oversampling\ Gain = 10 \cdot \log_{10}\left(Oversampling\ Ratio\right) \tag{2.47}$$

II.4.2.4 Dynamic Range Requirement Estimate

Determining the dynamic range requirement is a crucial part of the design process for an MR receiver. If the dynamic range requirement is underestimated, the resulting receiver may not be able to function properly in its intended imaging application. Hence, it is usually prudent to overestimate the dynamic range that will be required. Estimating the dynamic range involves assessing both the expected MDS and the expected maximum received signal levels for a given receiver application. This section will discuss the procedure for making these estimates.

II.4.2.4.1 MDS Estimate

The MDS power level is defined in (2.39) as the sum of the thermal noise floor, the receiver noise figure, and a detection threshold. For the determining the dynamic range requirement, however, the MDS level may be estimated by using only the thermal noise floor and the receiver noise figure or even by just the thermal noise floor (i.e., if the noise figure is unknown) (2.48). This simplifies the estimation process for the MDS level by allowing the estimate to be determined primarily by a knowledge of the receiver system bandwidth. The result is that the estimate will be lower than the true MDS level

(by about 3 to 5 dB), but that will only cause the dynamic range to be slightly overestimated.

$$MDS_{estimate} = Noise\ Floor + Noise\ Figure \approx kTB \tag{2.48}$$

II.4.2.4.2 Maximum Signal Level Estimate

The maximum signal level estimate, $MAXSIG_{estimate}$, may be found in two ways, both of which involve using an equation for the voltage that would be generated in a test coil by a sample imaging volume. The first method is to take the magnitude of the fid voltage (2.10) that would be produced in a test coil immediately after a 90° RF pulse has excited a test imaging volume, ΔV (2.49). In (2.49), the value for the initial transverse plane net magnetization (M_{xy}^{0}) is obtained by using the equation for the net equilibrium magnetization (M_0) from (2.5). The second method uses the signal voltage that would be received from a test sample during an imaging sequence (2.25) to determine the maximum signal level estimate (2.50). The main difference between the two methods is that the second uses the steady-state transverse magnetization (M_{xy}) to compute the estimate. (Refer to section II.2.6 to find equations for the steady-state transverse magnetization (M_{xy}).) As a result, the maximum signal estimate computed by the first (or fid voltage) method will always be higher (i.e., leading to a more stringent DR estimate) than one computed by the second (or steady-state) method. In many cases, however, (for instance, if it is known that the receiver in question will only be used for a particular type of imaging) the second estimate for the maximum signal level will be adequate. In order for the maximum signal level estimate to be accurate for either of these two methods, the B-field sensitivity of the test coil (B_{1t}) and the size of the test

volume (ΔV) must be similar to that of the detector coils and samples that will be imaged with the final receiver.

$$MAXSIG_{estimate} = \left|v(t=0)\right| = \sqrt{2}\omega_0 \Delta V B_{1t} M_{xy}^0 \qquad \text{(fid voltage method)} \qquad (2.49)$$

$$MAXSIG_{estimate} = \sqrt{2}\omega_0 \Delta V B_{1t} M_{xy} \qquad \text{(steady-state method)} \qquad (2.50)$$

II.4.2.4.3 Dynamic Range Estimate

To compute the dynamic range estimate, the value of maximum signal estimate must be converted from voltage to time average power (2.51). In (2.51), the resistance value, *R*, may be taken to be the characteristic impedance of the transmission line that will be carrying the received signal to the receiver input (i.e., usually 50 Ω). Then both the maximum signal and the MDS estimates need to be converted into decibels with respect to milliwatts (2.52). Finally, the estimated dynamic range requirement may be computed by subtracting the maximum signal estimate from the MDS estimate (2.53). The term, *SF*, in (2.53) is an optional safety factor that may be added in to prevent an underestimate of the dynamic range requirement.

$$P_{MSe} = \frac{(MAXSIG_{estimate})^2}{2 \cdot R} \qquad (2.51)$$

$$P_{MSe(dBm)} = 10 \cdot \log_{10}(P_{MSe} \cdot 1000) \quad P_{MDSe(dBm)} = 10 \cdot \log_{10}(MDS_{estimate} \cdot 1000) \qquad (2.52)$$

$$DR_{estimate} = P_{MSe(dBm)} - P_{MDSe(dBm)} + SF \qquad (2.53)$$

II.4.2.5 Single-Tone Dynamic Range Measurement

Measurement of the dynamic range for a receiver involves finding both the 1-dB compression point and the MDS level for that receiver. This may be accomplished by injecting a CW (continuous wave), single tone, RF test signal into the receiver and

observing the output response of the receiver [34]. (The output response of the receiver to the single-tone test signal may be analyzed by computing the power spectrum of the received signal.) The amplitude of the test signal is then varied until both the 1-dB compression point and the MDS level have been found. (The power level of the test signal may be controlled with a variable RF attenuator.) This is known as a measurement of the single-tone dynamic range for a receiver.

II.4.3 Image SNR

Another important performance metric for the MR receiver is the signal-to-noise ratio in the reconstructed MR image, or simply the image SNR. Image SNR is computed by taking a reconstructed MR image and calculating its average signal and the noise levels. These values for the average signal and noise are then divided to find the image SNR. The image SNR computation does not give an exact measure of the signal to noise ratio in the received signals, but it does give a feel for how well the receiver is performing. The image SNR calculation is most useful when the image is taken from a homogeneous phantom (e.g., a solid cylinder of either distilled water or 0.5 mmol $CuSO_4$(aq) solution). In this case, the signal values should be relatively constant over the imaging phantom, and a fairly accurate measure of the signal level in the image can be determined. There will, however, still be some variation in signal across the phantom due to the effects of the receiver coil sensitivity pattern (B_{1t}) and inhomogeneities within the static magnetic field (B_0) making an exact determination impossible.

To find the average signal value for a particular image, a square region of the image is first selected where a strong signal level is present. Often a 5×5 pixel (or larger)

region is chosen. The average signal value in this region is then found by summing the magnitude of the pixel values over this region and dividing that number by the total number of pixels within the region. The average noise level is determined in a similar manner. First, four or more square regions in the image are selected where it is known that no signal is present. These regions are generally each 25×25 pixels (or larger) to get a larger sample of the noise. Then, the average noise level is found by averaging the magnitude values of the pixels within the selected regions. An example of the image SNR computation is shown in Fig. 12. A standard spin echo image (TE/TR = 28/250 ms, 10 cm FOV, 256×128 pixels, 30 mm slice thickness, 1 average) was taken of a cylindrical imaging phantom filled with 0.5 mmol $CuSO_4$(aq) solution. The signal and noise levels were found by averaging a 25×25 pixel region and four, 25×25 pixel regions respectively. The signal and noise regions are indicated (Fig. 12). The image SNR was then computed (2.54) and found to be:

$$Image\ SNR = \frac{Average\ Signal}{Average\ Noise} = \frac{6.555}{0.0610} = 107.4\,. \tag{2.54}$$

II.4.4 Interchannel Isolation

An important consideration for the design of parallel receivers is the isolation between receiver channels. If there is insufficient interchannel isolation, the signal from one or more coils in a receive array may bleed through to other channels in the parallel receiver. This crosstalk would corrupt the signals on each of the bled-through channels and result in artifacts appearing in the reconstructed images from each of the corrupted channels. For this reason, it is standard practice to decouple the sensor coils in an NMR

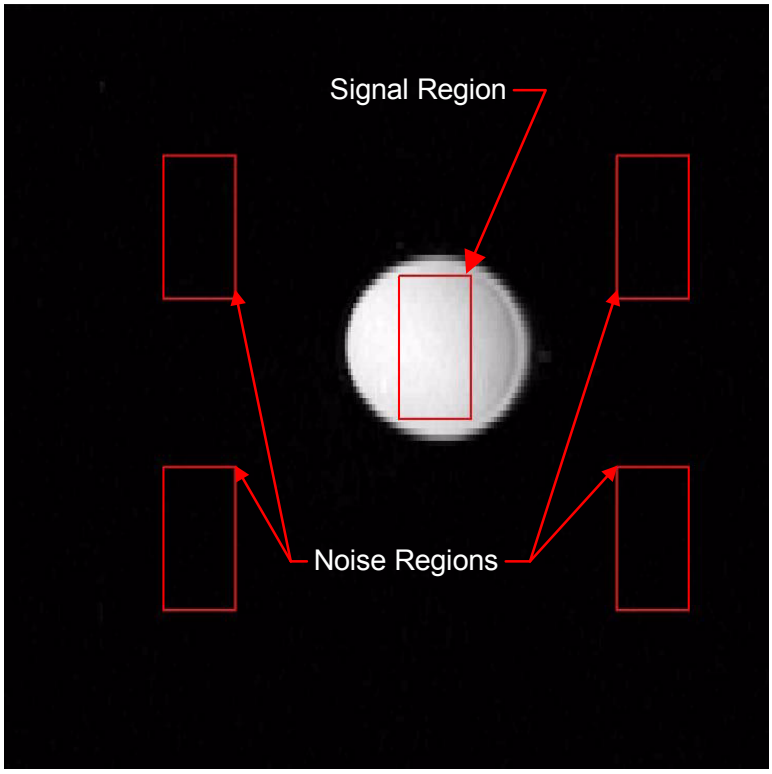

Fig. 12. Example of the location of signal and noise averaging regions for the calculation of image SNR on an image of a cylindrical test phantom. The signal region is 25 × 25 pixels large. Each of the four noise regions is also 25 × 25 pixels large.

phased array by at least 20 dB. (For details on standard methods of decoupling array elements in an MR phased array, refer to [33].) In order to ensure that the crosstalk due to interchannel coupling is negligible with respect to the crosstalk that results from coupling between coils in the receive array, the channel-to-channel isolation on a parallel receiver should be at least 35 dB [16].

II.4.5 Channel-to-Channel Gain Variation

In the ideal parallel receiver, the performance of each receiver channel, particularly the channel gain, would be identical to that of all the other channels. Thus, if the same RF signal were received by each channel in the receiver, the resulting baseband signal outputs from each channel would all have the same amplitude. This is because, in the case of the ideal receiver, the circuits for all the receiver channels would behave identically. Unfortunately, a parallel receiver must be constructed from real-world components, which means that the receiver gain will vary from channel to channel. For parallel imaging, the channel-to-channel gain variation should be kept as small as possible.

II.5 Summary

In this chapter, the fundamental ideas behind the design of instrumentation for parallel MR imaging have been presented. This included a brief introduction to the theory of NMR and MR imaging. Next, the details of the operation and design of the MR receiver were described. Finally, methods for the characterization of an MR receiver's performance were given. Taken together, these principles form the basis for the design of parallel MR imaging hardware.

CHAPTER III

A LOW-COST DESKTOP MR SCANNER*

In order to develop an understanding of the intricacies of the MR receiver, a prototype desktop MR scanner was constructed. The scanner is capable of imaging small volumes with high resolution on dedicated, low-field MR magnets. The scanner is based upon a conventional analog transceiver design and has been constructed from off-the-shelf sub-systems to lower the overall system cost. The prototype described here demonstrates the feasibility of inexpensive, low-field MR scanners as a new class of general use laboratory instrument for non-clinical applications.

III.1 System Implementation

In order to build a low-cost desktop MR system, the performance requirements for the desktop scanner system were first determined. Then the overall MR system was designed to meet these performance criteria. Finally, the performance of the prototype desktop MR system was evaluated. The details of this procedure will be discussed in the following sections.

III.1.1 System Requirements

The initial goal for the desktop MR scanner was to develop a compact MR imaging system capable of imaging 2.54 cm (1 in.) diameter phantoms at 0.16 and 0.21 T with a target resolution of 75 μm ×75 μm ×200 μm for a cost of under $10,000 (in addition to the magnet). This objective served as the starting point for the derivation of the system

*Part of this chapter is reprinted with permission from "A Desktop Magnetic Resonance Imaging System," by S. M. Wright, D. G. Brown, J. R. Porter, D. C. Spence, E. Esparza, D. Cole, and F. R. Husson, 2002, *Magnetic Resonance Materials in Physics, Biology, and Medicine,* vol. 13, pp. 177-185. Copyright 2002 by Elsevier Science B. V.

requirements for the desktop MR scanner. In addition to size and cost, the primary system specifications of concern for the desktop scanner were the operating frequency range, receiver noise figure, receiver bandwidth, receiver dynamic range, and the required RF transmit power. The requirements for each of these specifications are discussed in the following sections.

III.1.1.1 Operating Frequency Range

The operating frequency range of the MR scanner determines the field strengths at which the scanner will be able to image. For the desktop MR scanner, the only nuclear species of interest was the hydrogen nucleus (H^1). As a result, the required operating frequency range could be determined directly from the equation for the Larmor frequency (2.2) and the field strengths listed in the primary objective statement (i.e., 0.16 T and 0.21 T). The resonant frequencies for the hydrogen nucleus at these two field strengths are 6.81 MHz and 8.94 MHz, respectively. Thus, the desktop scanner needed an operating frequency range that would extend, at the very least, from 6.81 MHz to 8.94 MHz.

III.1.1.2 Noise Figure

The noise figure of a receiver measures how much that system will degrade the SNR of the signals which it detects. In general, the noise figure for an MR receiver should be as low as possible (in many cases commercial systems will have noise figures of 1 dB or less). This is accomplished by placing a high gain, low noise figure preamplifier stage at the front end of the receiver chain. For the desktop MR system, however, the noise figure requirement was made to be slightly less stringent due to cost concerns. The main

requirement was that the desktop receiver should be able to collect MR images with comparable SNR to those of a commercial system (see Chapter II section II.4.3 for a discussion of image SNR determination).

III.1.1.3 Bandwidth

The bandwidth required by an MR receiver is dependent upon the bandwidth of the MR signals which it will be used to receive. The bandwidth of the MR signal itself is determined by the pulse sequence used to create the MR image. In some cases, the signal bandwidth may be only a few kilohertz while in others as much as a few megahertz (for some fast gradient echo sequences). For most MR pulse sequences, however, a receiver bandwidth (through the digitizer) of 20 kHz is sufficient. Thus, the maximum bandwidth for the receiver in the desktop scanner was set to be at least 20 kHz.

III.1.1.4 Dynamic Range

The dynamic range requirement was estimated using the method presented in Chapter II, section II.4.2.4. As such, the MDS power level was first estimated using the equation for the thermal noise floor. Then, the maximum signal level was estimated via the equation for the magnitude of a fid voltage immediately after the application of a 90° RF pulse. From these two estimates, a final estimate for the dynamic range requirement of the desktop system receiver was formed.

III.1.1.4.1 MDS Estimate

The MDS estimate was formed directly from the equation for the thermal noise floor resulting from a matched resistive load (2.48), assuming a receiver bandwidth of 20 kHz

and an ambient temperature of 300 K. The result was an MDS estimate of 8.28×10^{-14} mW. Converting this figure into dBm yields an estimate of ¯130.8 dBm for the MDS.

III.1.1.4.2 Maximum Signal Level Estimate

The maximum signal level estimate was calculated based on the magnitude of the fid voltage (2.49) that would be induced in a test coil by a test sample volume immediately after the application of a 90° RF pulse in a static magnetic field of 0.21 T. This calculation required the determination of the following four parameters: the Larmor frequency of the hydrogen nuclei in the test sample (ω_0), the volume of the test sample (ΔV), the transverse plane magnetization of the test sample (M_{xy}), and the effective B-field sensitivity (B_{1t}) of the test RF coil. Determination of the Larmor frequency for the spins in the test sample was made using the upper limit of the operating frequency range requirement (see section III.1.1.1). This corresponded to the resonant frequency for hydrogen nuclei in a 0.21 T static magnetic field (i.e., $\omega_0 = 56.2\times10^6$ rad/s). The test sample consisted of a mouse-sized (i.e., 2.54 cm diameter, 2.54 cm long), cylindrical water phantom with a total volume of $\Delta V = 1.287\times10^{-5}$ m^3. The sample's transverse plane magnetization was found by using (2.5) from Chapter II, section II.1.3 to be 6.75×10^{-4} [J/T/m^3]. For a test coil, a 5-turn, 1.8 cm (0.7 in.) long, RF solenoid coil with a 3.35 cm (1.32 in.) diameter was used. The effective sensitivity (B_{1t}) of the solenoid detector coil was found by first using an equation for the on-axis B-field from a finite length, thin shell solenoid (3.1) oriented along the x-axis. Here, N is the number of turns, I is the excitation current [A], l is the length of the solenoid [m], a is the radius of the solenoid [m], and x [m] is the x-coordinate of the field observation point. For the

calculation of the maximum signal estimate, the field observation point, x, was taken to be at the center of the solenoid. (Using the center of the solenoid as the field observation point assumes that the flux density of the magnetic field generated by the solenoid, B_1, is constant over the entire sample volume. This assumption is not correct, and it will lead to a larger than necessary value for the maximum signal estimate. However, since this will not lead to an underestimation of the dynamic range requirement, the assumption does not adversely affect the maximum signal estimate.) The value of B_{1t} [T/A] was then found to be 1.171×10^{-4} T/A from the value for B_1 at the selected observation point by using (2.11) from section II.1.4 of Chapter II. Finally, the value of the maximum signal estimate was found to be 8.083×10^{-5} [V].

$$\vec{B}_1(0,x,0) = B_1(0,x,0)\cdot\hat{a}_x = \frac{\mu_0 NI}{2l}\left(\frac{\left(x+\frac{l}{2}\right)}{\sqrt{a^2+\left(x+\frac{l}{2}\right)^2}} - \frac{\left(x-\frac{l}{2}\right)}{\sqrt{a^2+\left(x-\frac{l}{2}\right)^2}}\right)\cdot\hat{a}_x \qquad (3.1)$$

III.1.1.4.3 Dynamic Range Estimate

To get the dynamic range estimate, the maximum signal estimate had to first be converted into a power level. This was accomplished using (2.48) and (2.49) from section II.4.2.4.3 of Chapter II. For the maximum signal estimate conversion to power, a characteristic impedance of 50 Ω was assumed which yielded a power level of -71.8 dBm. Using this value and that of the MDS estimate (from section III.1.1.4.1) in (2.53) (of Chapter II, section II.4.2.4.3), the final estimate for the dynamic range requirement

on the desktop receiver was found to be 59 dB. (This value was obtained using a safety factor (SF) value of 0dB in the calculation of the dynamic range estimate (2.49).)

III.1.1.5 RF Transmit Power

The amount of RF power available determines both the type of RF coils and the minimum 90° and 180° pulse widths that may be used for imaging on an MR transceiver. To determine the RF transmit power requirement, the amount of current required to tip all the spins within a test sample volume into the transverse plane in a given amount of time using an RF test coil must first be determined. The next step in the power estimation process is to determine the real resistance of the RF test coil. When these two parameters are known, the real time average power needed to produce the excitation current in the RF coil can be calculated. Each of these steps will be described in the following sections.

III.1.1.5.1 Excitation Current Requirement

In order to find the transmit excitation current requirement, the B-field required to form a 90° tip angle with a square RF pulse having a given pulse width must first be found. For the desktop scanner, the requisite pulse width was set at 10 μs and the 5-turn, 1.8 cm (0.7 in.) long, 3.35 cm (1.32 in.) diameter, RF solenoid coil of section III.1.1.4.2 was used as the test coil. The tip angle formula (2.7) (refer to Chapter II, section II.1.4) was then solved to find that the effective B_1 field in the rotating frame (B_{1eff}) which would produce a 90° tip in 10 μs was 5.87×10^{-4} T. This rotating frame B-field value was then related back to the field that would be produced by a square RF pulse from a

linearly polarized RF coil in the laboratory frame of reference via (2.7). Thus, the required filed in the lab frame (B_1) is 1.17×10^{-3} T.

Once the B-field required to create a 90° RF pulse is known, the current needed to produce that field strength with the test coil could be found. The equation for the on-axis B-field from a finite length, thin shell solenoid coil (3.1) from section III.1.1.4.2 was solved to find the excitation current (using the geometry parameters for the test coil and the B_1 value). Thus, the excitation current required to produce a 90° tip angle in 10 μs using the test coil was 7.09 A.

III.1.1.5.2 Test Coil Resistance

Finding the real resistance of the RF test coil used in section III.1.1.5.1 involved two steps. First, the inductance of the test coil had to be found. The inductance, L [H], of a thin shell, finite length solenoid is given by (3.2). Using (3.2) the test coil parameters, the inductance was calculated to be 6.73×10^{-7} H. Second, the real resistance, R [Ω], was determined from the Q of the loaded coil. For the test coil, a loaded Q of 100 was assumed. The Q of a coil is related to the coil's inductance and resistance by (3.3). Here ω is the Larmor frequency of the protons in the test sample (i.e., $\omega=\omega_0=56.2\times10^6$ rad/s at 0.21 T). The real resistance of the coil was then found to be 0.378 Ω.

$$L = \frac{\mu_0 N^2 \pi a^2}{l^2}\left(\sqrt{l^2 + a^2} - a\right) \tag{3.2}$$

$$Q = \frac{\omega L}{R} \tag{3.3}$$

III.1.1.5.3 RF Transmit Power Requirement

Finally, the RF power requirement was found by calculating the time average power that would be dissipated in the real resistance of the RF test coil during the10 μs, 90° pulse excitation. The time average power in this case was determined using (3.4), where P_{ave} is the time average power [W]. For the test coil considered above with a resistance of 0.378 Ω and an excitation current of 7.09 A, the time average power was 9.50 W. Thus, the desktop transceiver requires at least 9.5 W of time average power to function properly.

$$P_{ave} = \frac{I^2 R}{2} \tag{3.4}$$

III.1.2 System Design

Based on the design requirements, a desktop MR system was developed. The desktop scanner was designed to operate over a 2-11 MHz frequency range, have a receiver bandwidth of at least 20 kHz, greater than 59 dB of receiver dynamic range, and at least 20 Watts of RF transmit power. In addition, the desktop system was designed to collect images with SNR comparable to that of a commercial MR scanner. The desktop scanner was built from low-cost and off-the-shelf components wherever possible. A block diagram of the scanner appears in Fig. 13. From the figure, it may be seen that the desktop scanner is composed of three main elements: computer control, transceiver, and magnet subsystem. The function and design of each of these three sections will be discussed below.

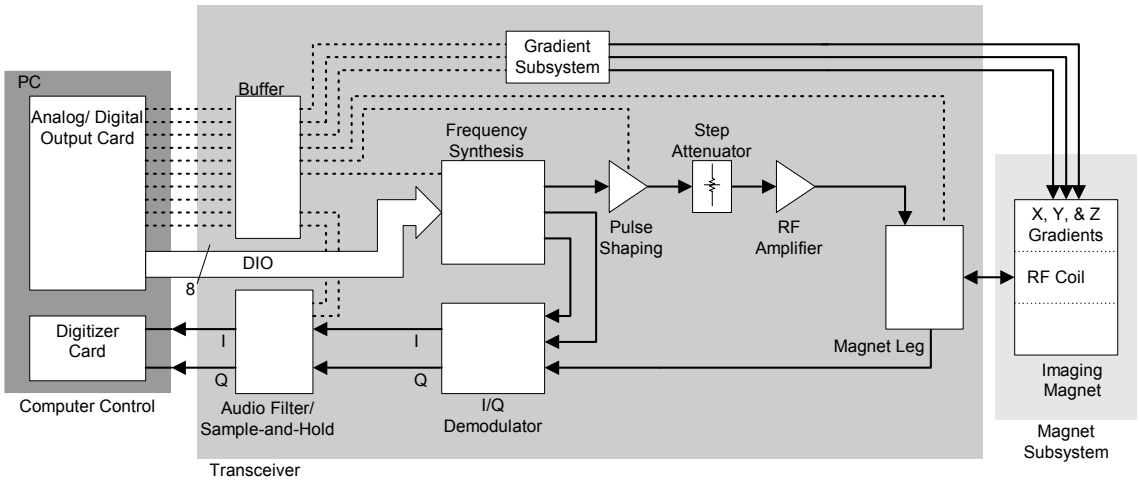

Fig. 13. Block diagram of the desktop MRI system. The three main components of the system are shown: computer control, transceiver, and the magnet subsystem.

III.1.2.1 Computer Control

The computer control portion of the scanner provides timing and control signals for the transceiver, performs data acquisition, processes the received image data, and provides a graphical user interface for pulse sequence programming and image analysis tasks. The computer control block is composed of four main sections. The first of these sections is the control computer itself. The second component is the analog and digital output unit which produces the timing signals for the rest of the MR system. Next is the analog input module which digitizes the received MR signal data from the transceiver. The last, and perhaps most important component, is the control software which enables the user to run the MR scanner.

The control computer itself is a 300 MHz, Pentium II machine. The computer has 384 MB of PC-133 SDRAM to enable the compilation of long pulse sequences. Additionally, the PC has a 15 GB hard drive for data storage as well as an on-board CD-ROM writer for backing up large data sets. These characteristics allow the control computer to run the user programmed pulse sequences, collect the resulting MR signal data, and perform all the necessary image processing tasks on the MR images which result.

The analog and digital output unit is made up of a PCI-based analog output card which resides in the control computer chassis. The card, a PCI-6713 (National Instruments, Austin, TX), has 8 simultaneous channels of 12-bit, 1 MSPS analog output. The PCI-6713 also has two, 20 MHz counter-timers and an additional 16 bits of slow

digital input/ output (DIO). This enables the analog and digital output module to provide all the timing and control signals required by the desktop scanner.

The analog input module is also comprised of a PCI card that is located within the control PC. This card is the PCI-4452 from National Instruments. This card is capable of simultaneous, 16-bit digitization of four analog input channels at 204.8 kSPS. Thus, the analog input module is able to acquire demodulated MR signals which have a bandwidth of up to 102.4 kHz. This is more than sufficient to digitize the demodulated analog signals from the receiver which will have bandwidths of up to 20 kHz.

The control software for the desktop MR scanner was developed entirely in the LabVIEW programming environment by Dr. Jay Porter and his students in the Department of Engineering Technology at Texas A&M University [39]. LabVIEW is graphical in nature, making it easy to generate highly modular code that is very user-friendly. These LabVIEW code modules may be rapidly chained together to form very complicated routines. In the case of the desktop MR scanner, LabVIEW was used to quickly develop a complete MR control console with a custom 'push-button' user interface.

The MR console control software consists of a number of user interface windows which each perform a specific function. The main console window appears in Fig. 14. Here the user may input such parameters as the number of frequency encoding points (Readout Points), the number of signal averages, and the repetition time (TR). Also provided by the user is the name of the pulse sequence file to be used in the imaging

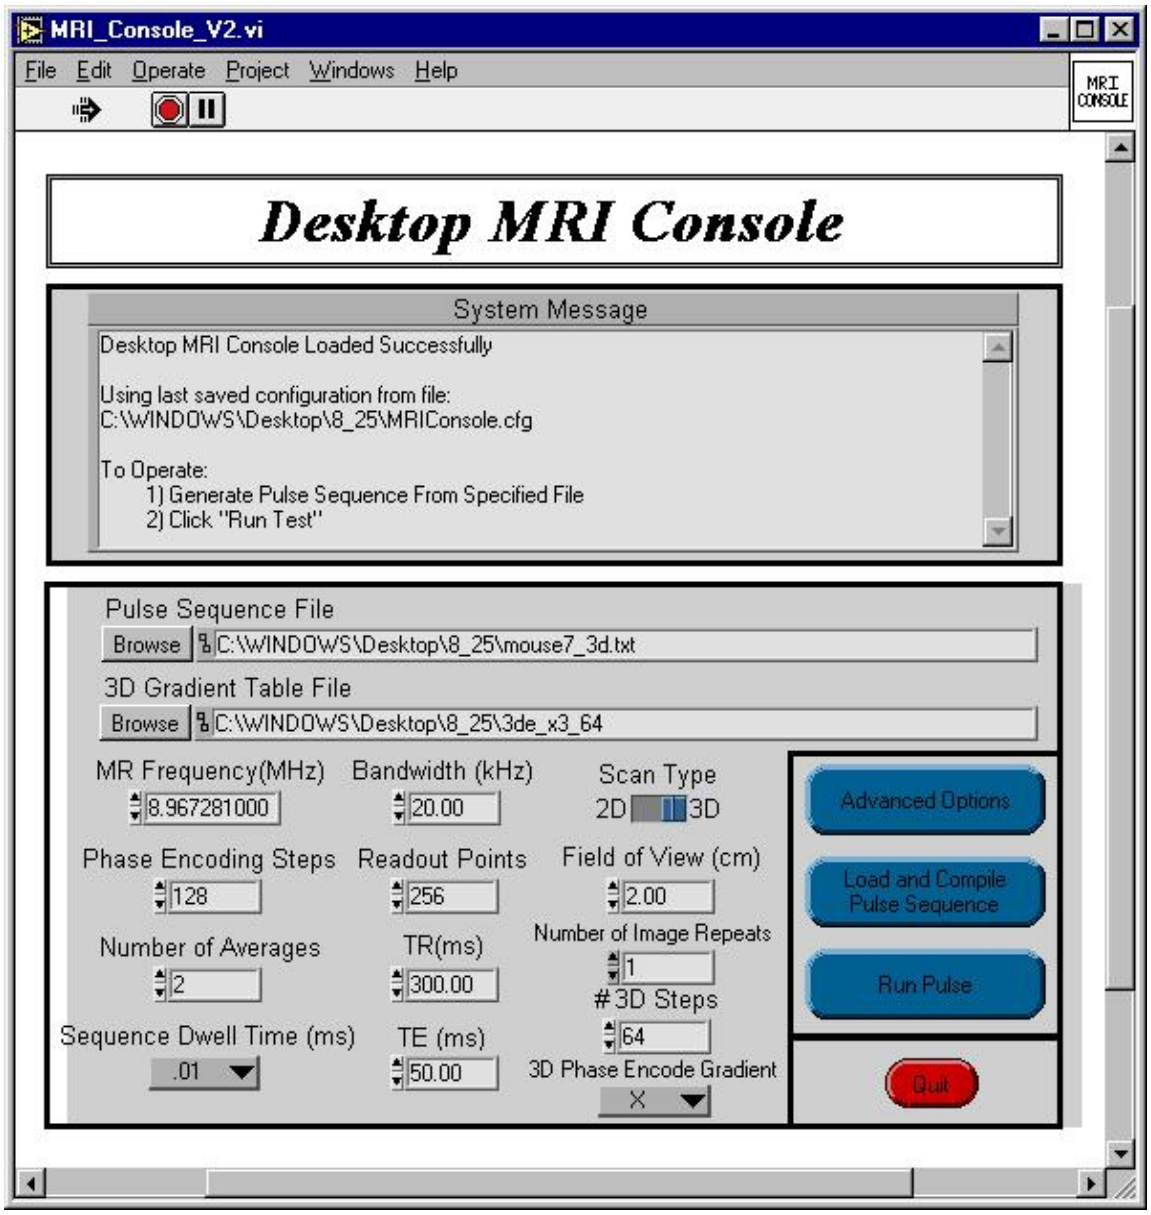

Fig. 14. Main MR control console window. Using this interface, the user is able to set the pulse sequence, imaging frequency, sequence bandwidth, number of phase encoding steps, the number of averages, and other important imaging parameters.

experiment. The information entered into the main console window is then used by the control software to compile the pulse sequence into RAM on the control computer.

Programming of the actual pulse sequence is performed by the user in a spreadsheet format. The pulse sequence data is entered in the form of rows for each control line and columns for each time event. The columns are each assigned a duration in milliseconds. For each control channel, the user must supply a text-based code for each time slot defined in the sequence. These codes determine the type of signals that appear on the eight control lines of the analog and digital output unit. Among the options available to the user are hard or soft RF pulse waveforms, fixed or ramped gradient amplitudes, and tables for programming phase encode values or RF phase rotation. A sample pulse sequence is shown in Fig. 15. In this example pulse sequence, the user has chosen to implement a spin echo pulse sequence with 9.46 ms TE time and a 24.87 ms TR time. Once the user has finished programming a pulse sequence, it is saved as a tab-delimited text file which the console control software can read. Since the actual format of the pulse sequence file is so simple, pulse sequence programming may be done offline by the user with software such as Notepad or ExCel (Microsoft, Redmond, WA).

After the console software reads in a pulse sequence file, it is compiled into memory on the control computer. The compiled version of the sequence is then streamed to the analog and digital output unit to provide the necessary timing signals for the transceiver during image acquisition. When the user acquires an image, the acquisition window (Fig. 16) appears. This window displays the pulse sequence control lines and is updated with each phase encode line. Additionally, the acquired echo is displayed for every

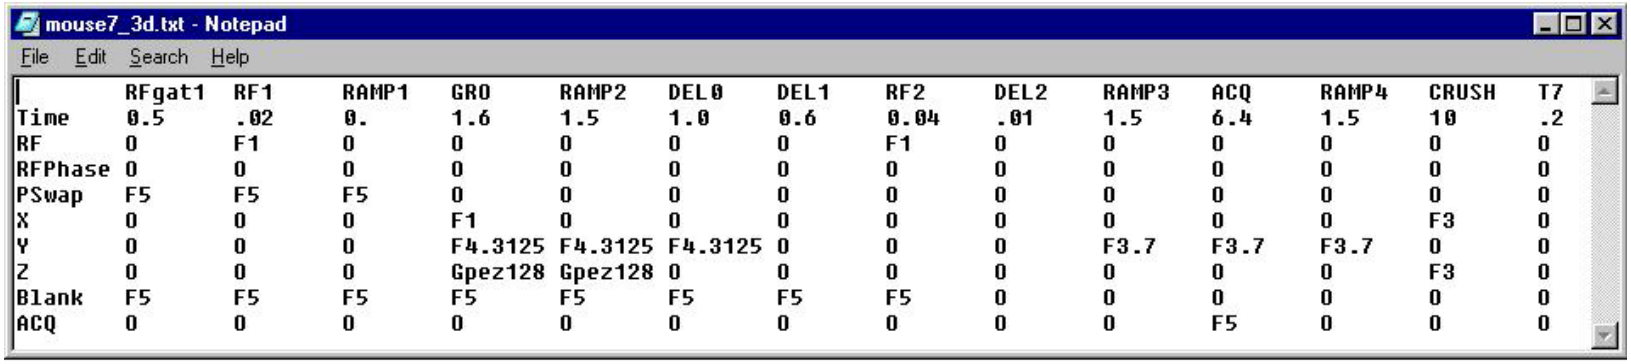

	RFgat1	RF1	RAMP1	GRO	RAMP2	DEL0	DEL1	RF2	DEL2	RAMP3	ACQ	RAMP4	CRUSH	T7
Time	0.5	.02	0.	1.6	1.5	1.0	0.6	0.04	.01	1.5	6.4	1.5	10	.2
RF	0	F1	0	0	0	0	0	F1	0	0	0	0	0	0
RFPhase	0	0	0	0	0	0	0	0	0	0	0	0	0	0
PSwap	F5	F5	F5	0	0	0	0	0	0	0	0	0	0	0
X	0	0	0	F1	0	0	0	0	0	0	0	0	F3	0
Y	0	0	0	F4.3125	F4.3125	F4.3125	0	0	0	F3.7	F3.7	F3.7	0	0
Z	0	0	0	Gpez128	Gpez128	0	0	0	0	0	0	0	F3	0
Blank	F5	F5	F5	F5	F5	F5	F5	F5	0	0	0	0	0	0
ACQ	0	0	0	0	0	0	0	0	0	0	F5	0	0	0

Fig. 15. Typical pulse sequence file. The time bins are labeled in the first row. The second row is used to assign a time length in milliseconds to each time bin. Rows three through ten correspond to the eight analog control signal outputs. Each of the rows (from two through ten) is labeled in the first column. In the pulse sequence depicted here, the total length of the pulse sequence file (TR) is 24.87 ms while the echo time (TE) is 9.46 ms.

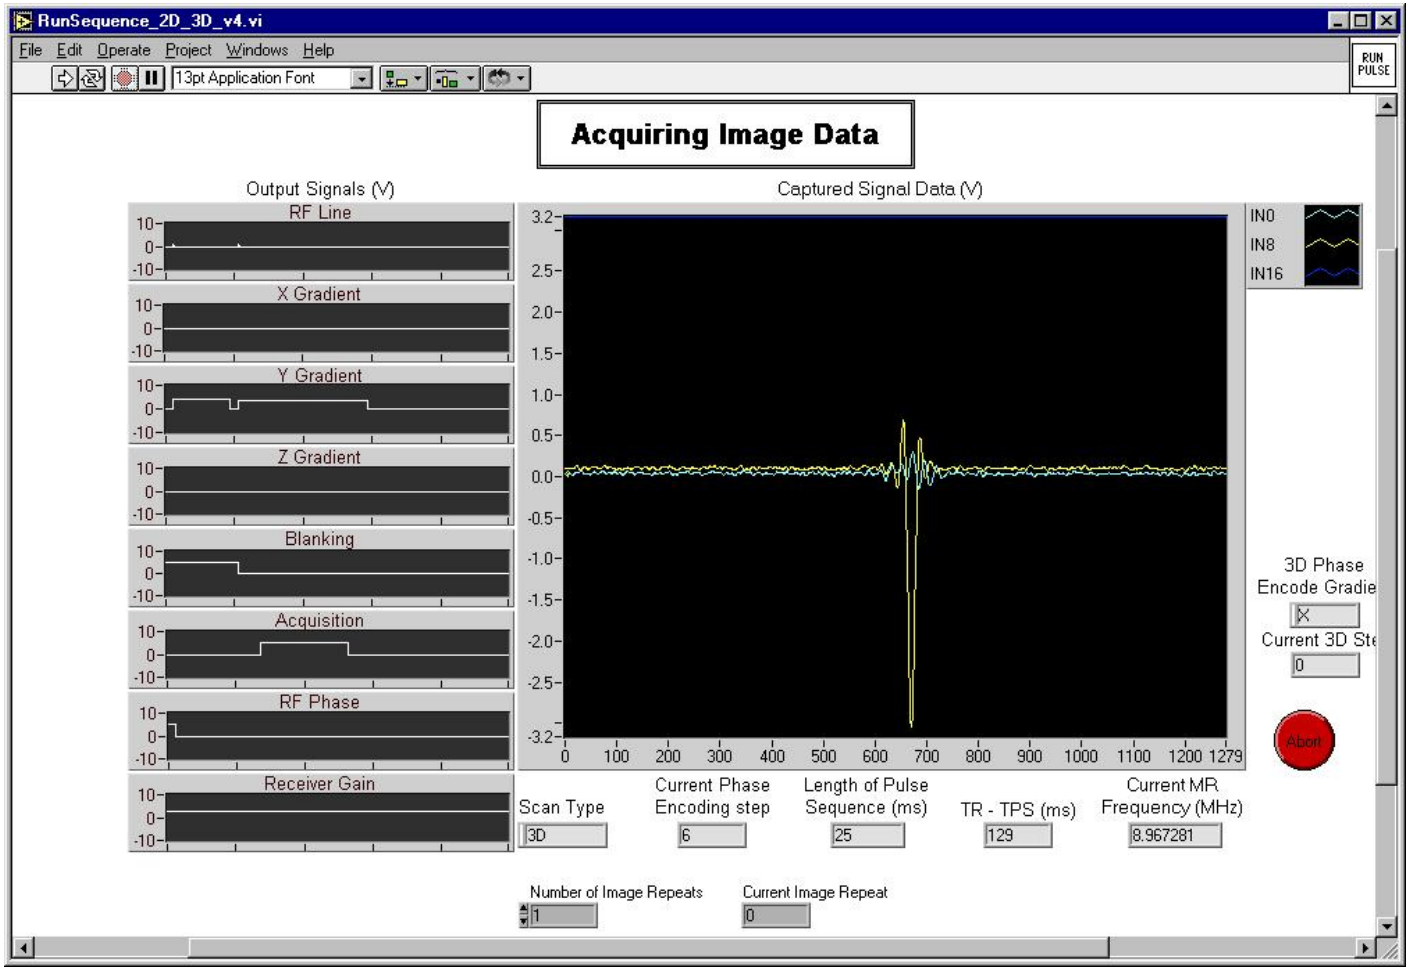

Fig. 16. Image acquisition window. This window provides the user with a display of the run time parameters for an imaging sequence. The waveforms for the eight analog control outputs appear on the left of the window. The acquired signals (both I and Q channels) are displayed in the middle of the window. The parameter information is updated after each signal acquisition.

phase encode step. After the pulse sequence is completed, the user is placed in an image analysis window. This last window allows the user to display both the image and k-space data, perform rudimentary statistics on the image to compute factors such as SNR, and save the acquired data to a file for later examination.

III.1.2.2 Transceiver

The transceiver section of the desktop MR system is primarily responsible for the modulation of the RF transmit waveforms at the MR imaging frequency and for the demodulation of the received MR signal data. From Fig. 13, it is apparent that the transceiver consists of several sub-modules. These modules include the buffer, frequency synthesis, pulse shaping, RF amplification, magnet leg, I/Q demodulation, anti-alias filtering and sample-and-hold, and gradient subsystems. The purpose and design of each of these subsystems will be presented in this section.

III.1.2.2.1 Buffer Module

The buffer module protects the output stages of the analog and digital output unit in the control computer from damage. The buffer stage is necessary due to the inability of most National Instruments cards to source more than 5 mA of output current. The buffer module (Fig. 17) consists of an 8-channel analog buffer and a 2-channel digital buffer. The analog buffer was constructed from eight, LM6221 high speed buffer chips from National Semiconductor (Santa Clara, CA). Each LM6221 has a 50 MHz operating bandwidth and is capable of supplying 300 mA of drive current. The digital buffer was made from a 74F244 fast TTL-level line driver chip from Texas Instruments (TI, Dallas, TX).

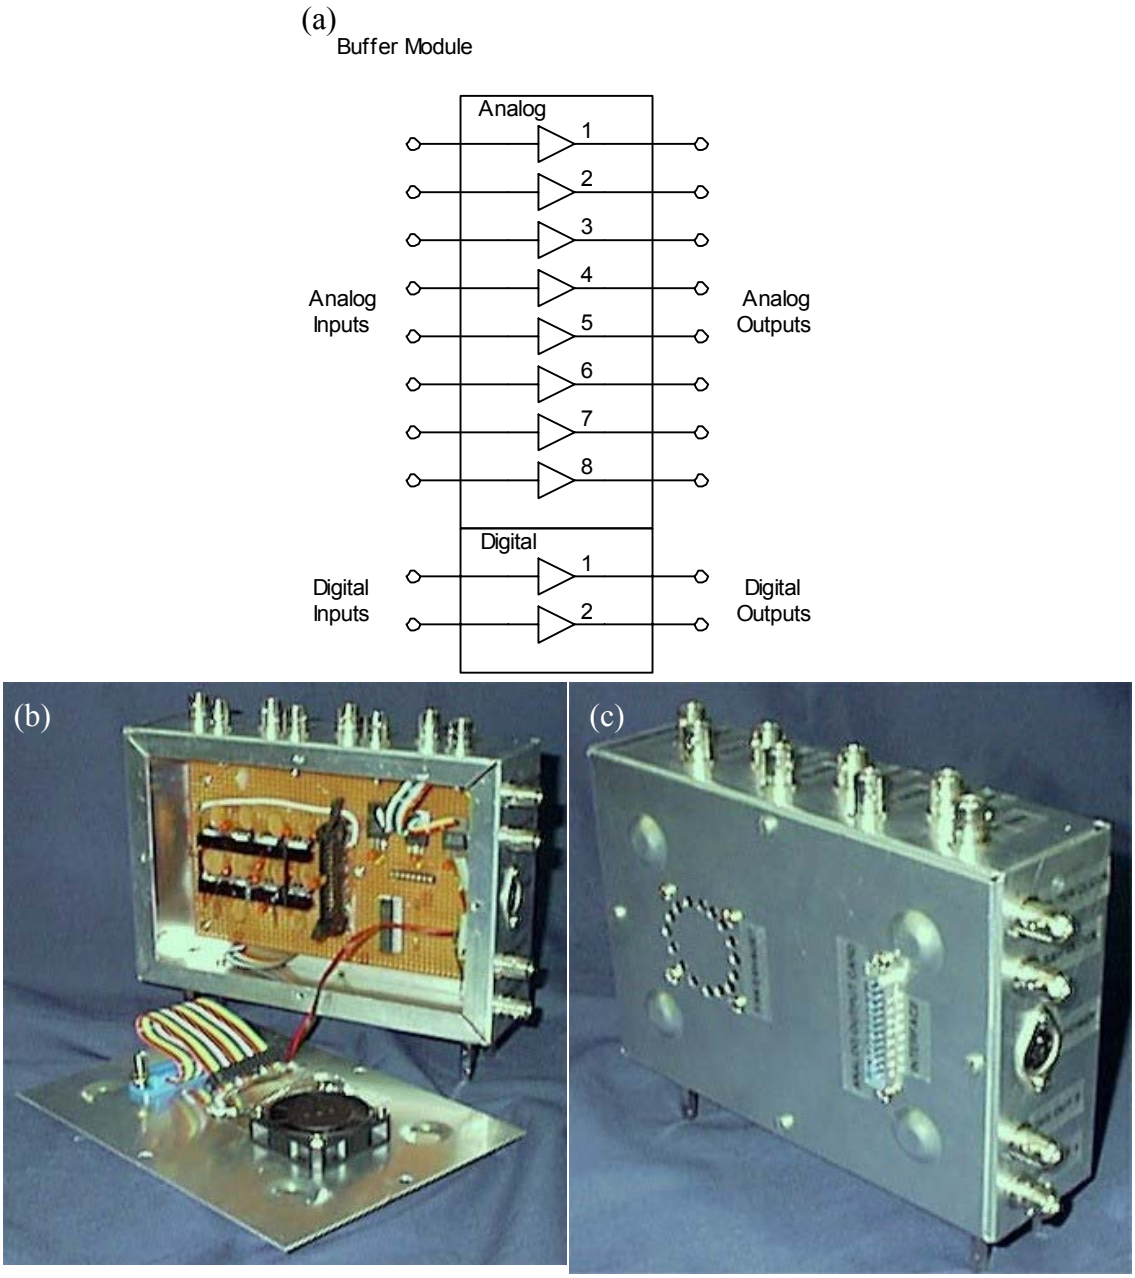

Fig. 17. Analog and digital buffer module. (a) Functional schematic of the buffer module. (b-c) Pictures of the buffer module showing the module contents (b) and the input, output, and power connections (c).

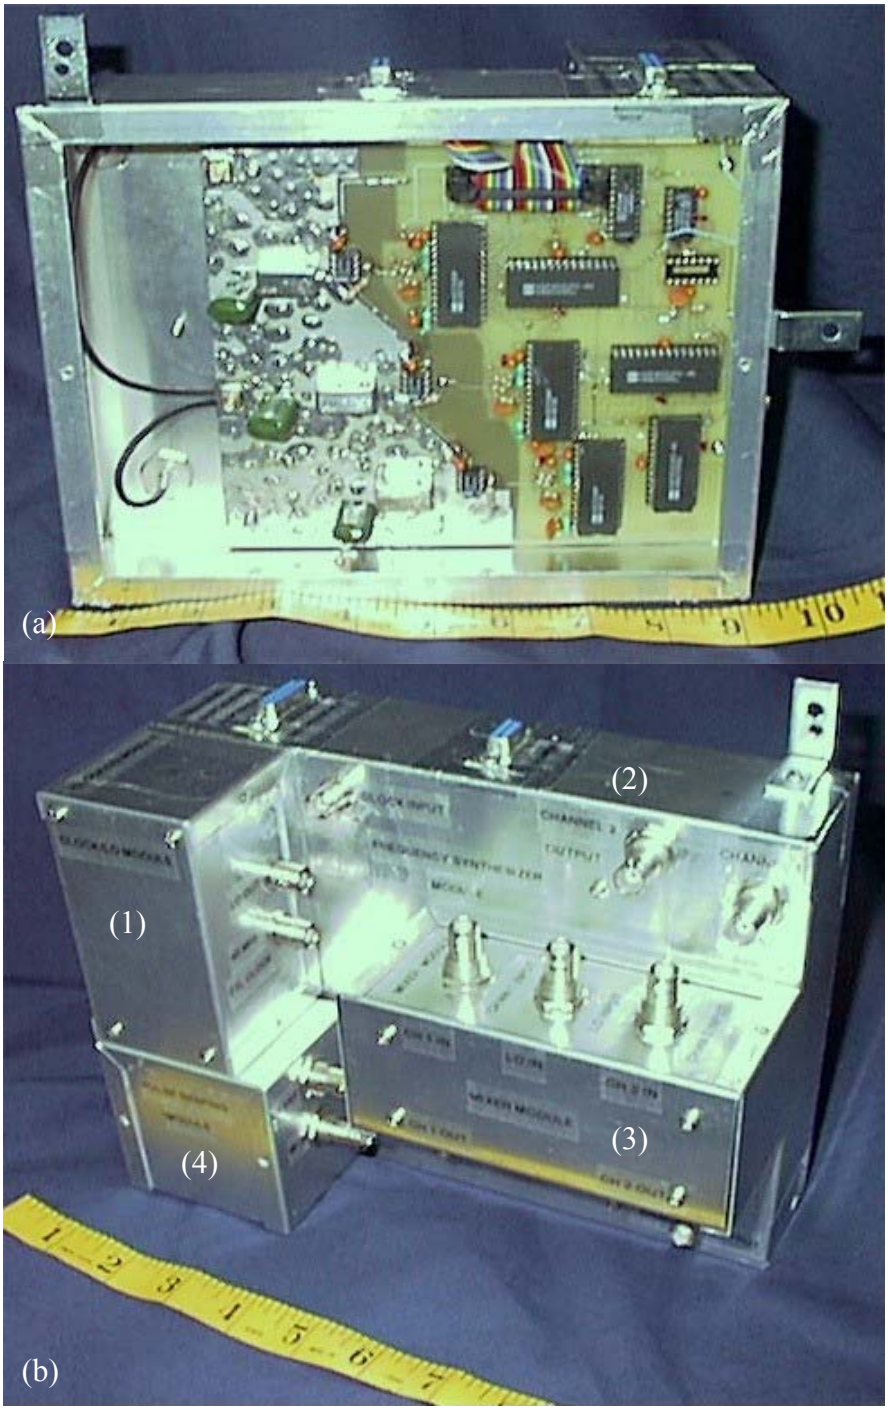

Fig. 18. Pictures of the frequency synthesis and RF pulse shaping modules. The measuring tape is in inches. (a) The direct digital synthesis (DDS) module PC board. (b) The complete frequency synthesis module with: (1) the clock/ LO module, (2) the DDS module, (3) the optional mixing stage module, and (4) the RF pulse shaping module.

Frequency Synthesis

The frequency synthesis module (Fig. 18) provides the three single frequency signals required by the transceiver. These signals are the RF transmit frequency (f_o), the first IF stage mixing frequency (f_0 + 10.7 MHz), and the second IF mixing frequency (10.7 MHz). The frequency synthesis module is composed of three main components. First there is a clock/ LO (local oscillator) card. Second is the direct digital synthesis (DDS) card. The third is an optional external mixing module. The clock/ LO card and the DDS module may be used by themselves to generate the three required frequencies (Fig. 19) or the clock/LO card and DDS board may be used with the external mixing module (Fig. 20). If the external mixing module is not used, the transceiver will only work at MR frequencies between 0 and 11 MHz; however, if the external mixing module is used, the transceiver can image at frequencies within the range of 0 to 50 MHz and 70 to 90 MHz.

The first component of the frequency synthesis module is the clock/ LO card (Fig. 19 and Fig. 20). This card provides both the master system clock for the transceiver and a user selectable LO signal for use with the external mixing module. The clock signal is generated by a TTL-level output, 80 MHz clock oscillator chip (Fox Electronics, Fort Myers, FL). To generate the LO signal, an SN74F161, fast TTL, modulo-16 binary counter chip is used as a frequency divider (Texas Instruments, Dallas, TX). The binary counter chip takes the input 80 MHz clock signal and outputs four, TTL-level frequencies at 40, 20, 10, and 5 MHz. Each of these four frequencies is phase-locked to the master system clock. The three highest frequencies (i.e., 10, 20 and 40 MHz) and the 80 MHz system clock are then buffered by an SN74F244 TTL-level buffer amplifier

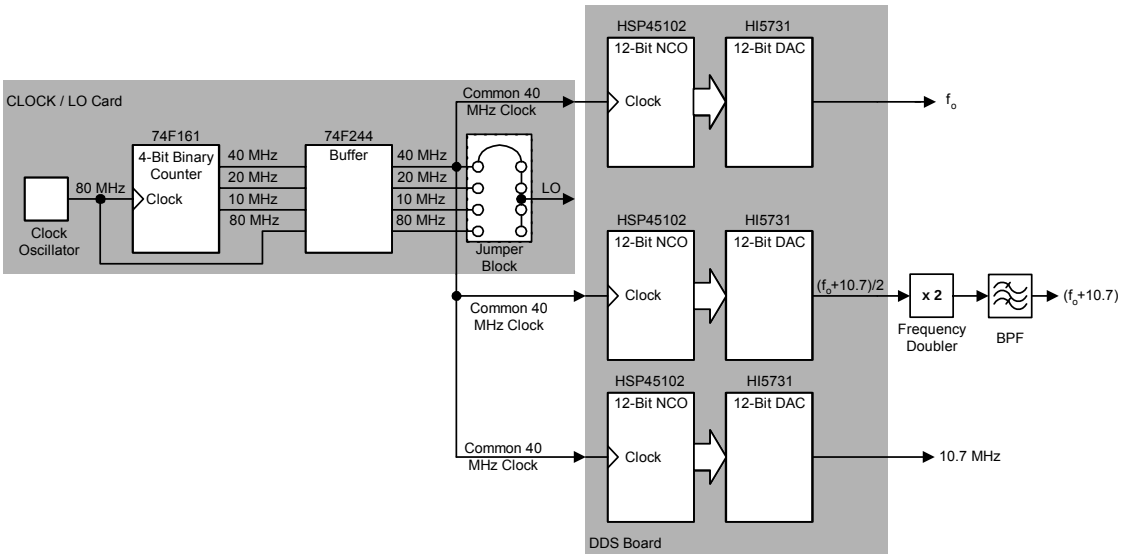

Fig. 19. Frequency synthesis module schematic. The module consists of a clock source card and a direct digital synthesis (DDS) board. The DDS card generates three frequencies: f_o, f_o+10.7, and 10.7 MHz. The first frequency, f_o, is the imaging frequency. The second frequency, f_o+10.7, is the first downconversion stage mixing frequency. Finally, 10.7 MHz is the second downconversion stage mixing frequency. The first frequency is used to generate the transmit excitation pulses while the last two frequencies are used to demodulate the received MR signal. This version of the frequency synthesis module allows the desktop system to image over the 0-11 MHz frequency range.

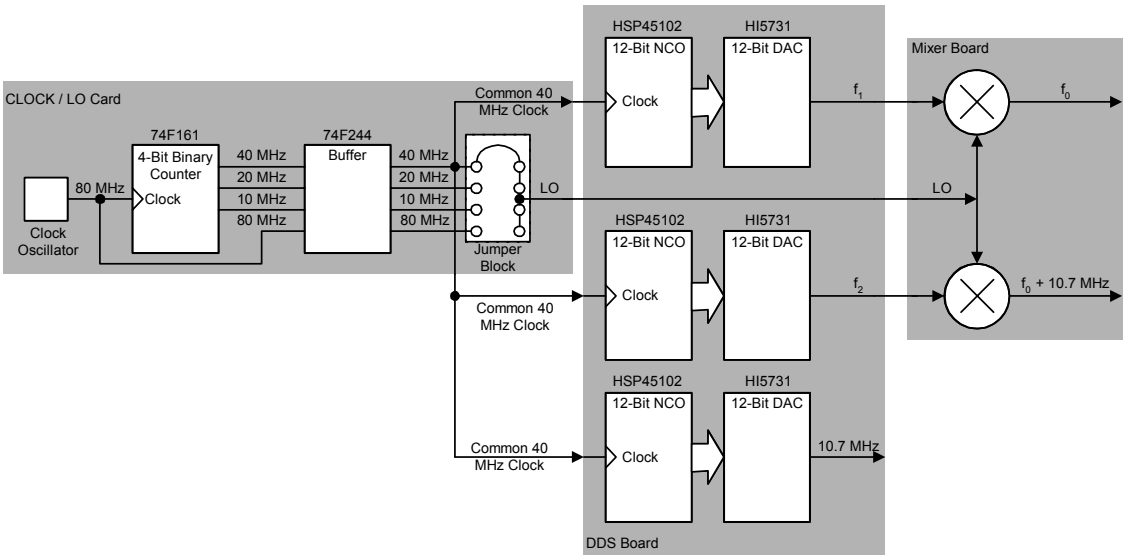

Fig. 20. Frequency synthesis module schematic with the external mixing module. The module consists of a clock source card, a direct digital synthesis (DDS) board, and an external mixer module. Like the version depicted in Fig. 6, the frequency synthesis card shown here generates three output frequencies: f_o, f_o+10.7, and 10.7 MHz. The main difference is the use of an external mixing card. In this implementation, the DDS card produces the following three frequencies: *f1*, *f2*, and 10.7 MHz. The *f1* frequency is equal to (f_o+10.7-*LO*) MHz; likewise, the *f2* frequency is given by (f_o-*LO*) MHz. The use of the mixing card extends the achievable imaging frequency ranges to cover the 0-50 and 70-90 MHz frequency bands.

(Texas Instruments). The buffered 40 MHz signal is sent to the DDS board where it serves as the primary clock for the DDS units on that card. Each of the four buffered signals also travels to a jumper block. The user is then able to set the frequency of the LO output to 10, 20, 40, or 80 MHz by setting the jumper appropriately.

The second element of the frequency synthesis block is the DDS module. The DDS module actually consists of three separate waveform synthesis circuits operating off the same 40 MHz clock signal. This allows the three DDS circuits to generate three independent, single frequency output signals that are phase-synchronized to each other. Each DDS circuit is based on the Harris HSP45102 and HI5735 integrated circuits (Intersil-Harris, Irvine, CA) [40-42]. The HSP45102 is a numerically controlled oscillator (NCO) which provides a 12-bit binary sine wave output. The HI5735 is a 12-bit, high speed digital-to-analog converter (DAC) which has a maximum throughput of 80 MSPS. Together, the two chips form a waveform synthesis circuit that is capable of generating 0-20 MHz sine waves with a 0.009 Hz digital tuning resolution. The output frequency of the three DDS circuits is programmed via an 8-bit wide serial interface. This programming interface is directly connected to the 8-bit, slow DIO channel on the PCI-6713 card that is resident within the control computer. The output phase of the generated waveforms is also controllable in 90° increments. However, in the transceiver DDS module, only the output phase of the NCO which generates the RF transmit frequency is externally programmable. The RF transmit NCO has been configured to allow phase jumps in increments of 180° to aid in the modulation of the transmit pulses. These phase jumps are directly controlled by one of the analog outputs from the

computer control module in order to ensure that precise timing is maintained during an imaging sequence.

The last part of the frequency synthesis block is the external mixing module. While the frequency synthesis module may be used without the external mixing stage (Fig. 19), it is limited to imaging at frequencies between 0 and 11 MHz. This is because the output stage buffer amplifiers on the DDS board become unstable at frequencies higher than 11 MHz. Using the mixing module allows the transceiver to operate over a broader range of frequencies (Fig. 20). The external mixing module was constructed from a PSC-2-1 power splitter and two SRA-3MH mixers from Mini-Circuits (Brooklyn, NY). In the mixing module, the outputs of the DDS board are mixed with the LO frequency from the clock/ LO card. This results in a frequency translation of the frequency outputs from the DDS board by 10, 20, 40, or 80 MHz, depending upon the jumper setting on the clock/ LO card. This means that with the external mixing module, the transceiver can operate over a frequency range that extends from 0 to 50 MHz and 70 to 90 MHz.

III.1.2.2.2 Pulse Shaping

The pulse shaping section of the transceiver performs modulation of the transmit frequency, f_0 (see Fig. 18), to produce 'hard' or 'soft' transmit pulses. In MRI terminology, a 'hard' pulse has a rectangular amplitude profile while a 'soft' pulse has a non-rectangular amplitude profile. The pulse shaping module consists of a variable gain amplifier (Fig. 18-b) used together with a push-button attenuator. The variable gain amplifier performs amplitude modulation while the attenuator allows further control over transmit pulse power levels. The variable gain amplifier is a CLC5523 chip from

National Semiconductor. The CLC5523 has linear-in-dB gain control applied via a 0 to 2 V modulation envelope signal and is capable of producing gain changes from 0 to -80 dB at rates of 4 dB/ns. The CLC5523 is available from the manufacturer with an evaluation PC board kit for only $4. The push-button attenuator is a model 50B-002 from JFW Industries (Indianapolis, IN) and is capable of providing 0 to 45.5 dB of attenuation. To achieve complex 'soft' pulse modulation, such as that required to for a sinc pulse, the phase of the transmit frequency may be hopped by 180° during the negative portions of the modulation envelope signal. An example of a two-lobe transmit sinc pulse produced by the pulse shaping module appears in Fig. 21. The phase of the transmit frequency was jumped by 180° during the two lobes adjacent to the main lobe of the sinc pulse.

III.1.2.2.3 RF Amplification

Final amplification of the RF transmit pulse occurs after the pulse shaping operation. To accomplish this task, there are a number of low-cost, linear amplifiers available which operate over the 1.6 to 30 MHz frequency band. Most of these amplifiers were originally intended for use by amateur radio enthusiasts, but with the simple addition of an external blanking switch, they can be made to work equally well for MR imaging applications. For the desktop transceiver, a model 20B (Henry Radio, Los Angeles, CA), 20 Watt, CW (continuous wave), RF amplifier was selected (Fig. 22). This amplifier sources ~43 dB of gain for an input signal of up to 1 mW (0 dBm). As such, the amplifier provides sufficient power for imaging the 2.54 cm (1 in.) fields-of-view for which the desktop scanner was designed. The amplifier sells for only $355 with a

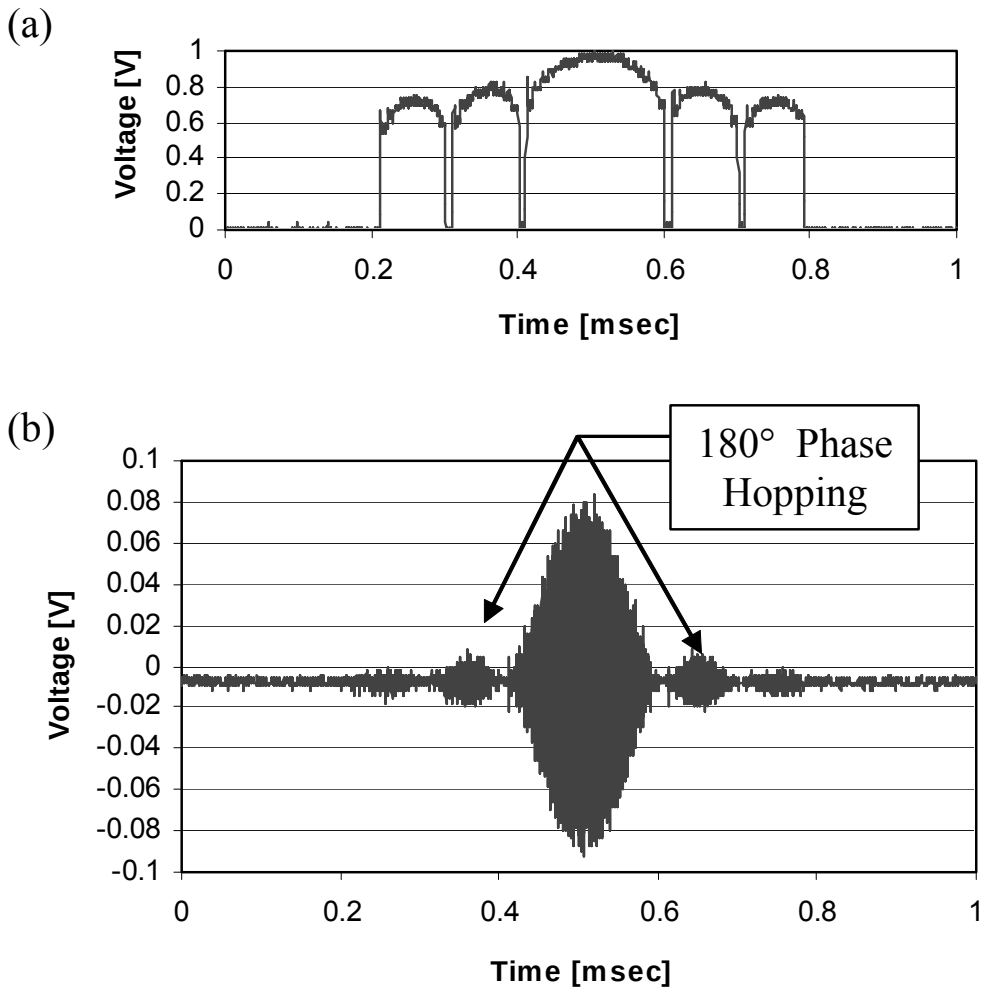

Fig. 21. Performance of the pulse shaping module as measured with a Tektronix TDS-220 digital oscilloscope. A sinc-modulated RF pulse with two side lobes is depicted here. (a) Amplitude modulation control signal. (b) Shaped pulse output. The side-lobes where the phase of the underlying RF transmit frequency has been hopped by 180° are indicated with arrows.

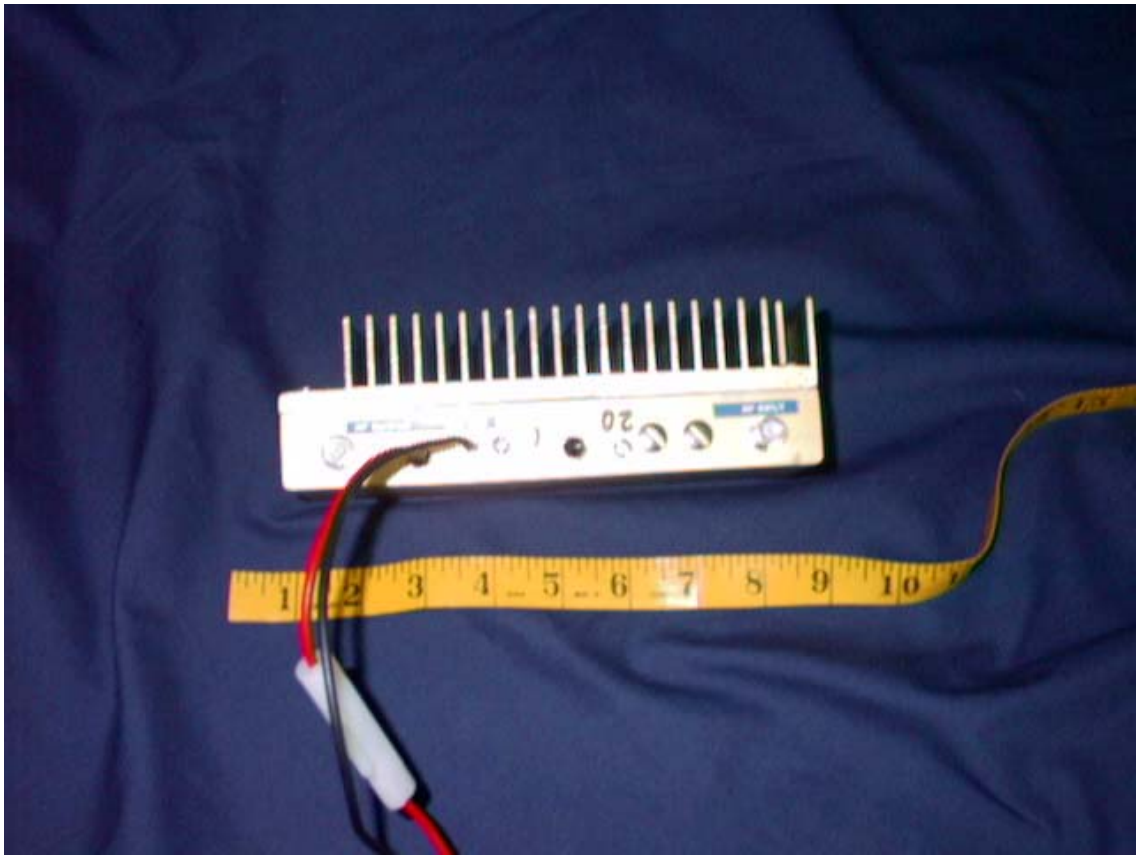

Fig. 22. RF power amplifier from Henry Radio. The amplifier provides 20 Watts of RF power over the 1.6 to 30 MHz frequency band. The dc power connections are made via the red (13.8 V) and black (GROUND) wires. The RF input and output connections are made via two front panel mount BNC connectors. The measuring tape indicates dimensions in inches.

(a)

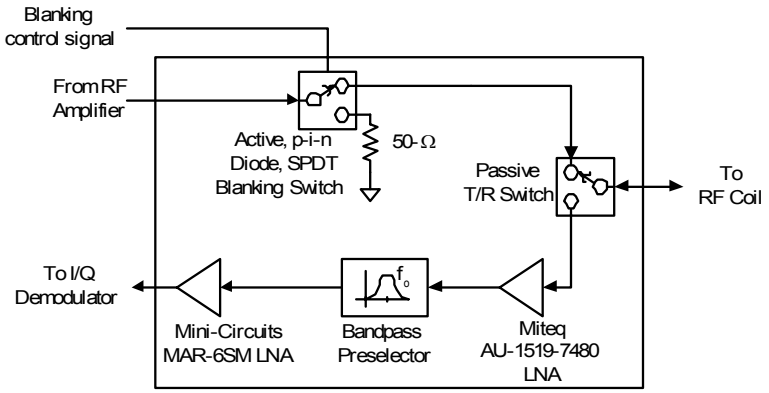

(b)

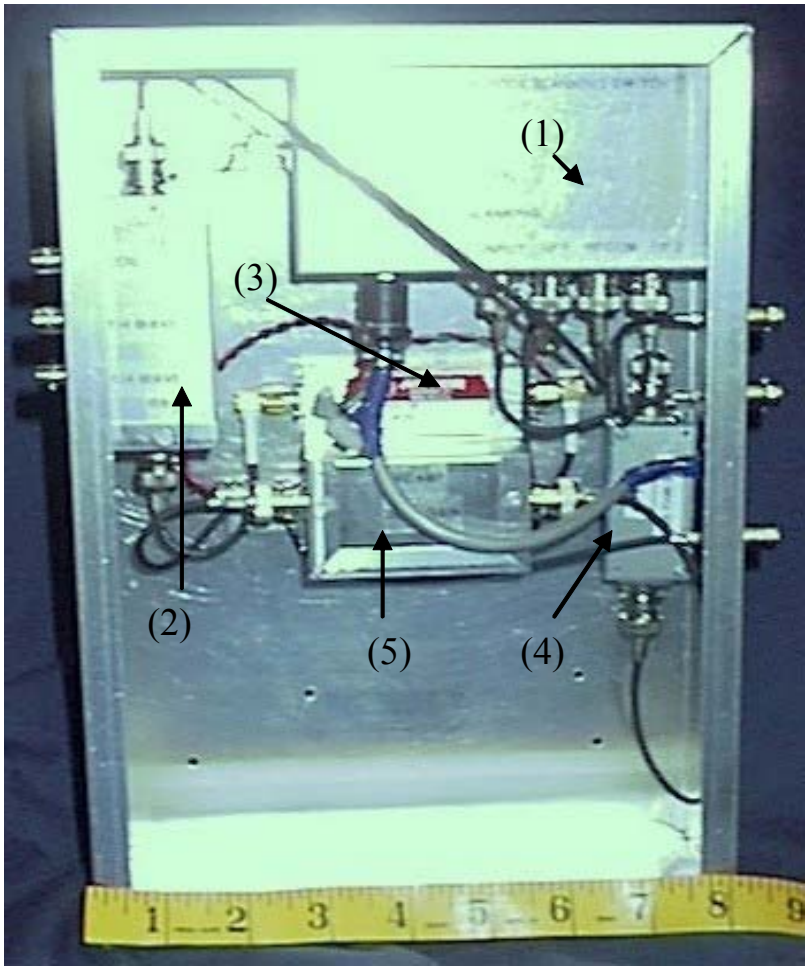

Fig. 23. Magnet leg module. (a) Functional schematic for the magnet leg module. (b) Photograph of the magnet leg module showing its five main components: 1) PIN diode blanking switch, 2) passive quarter-wavelength TR switch, 3) low noise preamplifier, 4) bandpass preselect filter, and 5) RF gain block.

13.8 V, dc power supply included. The external PIN diode blanking switch required by this amplifier is housed in the magnet leg module and will be described in the next section.

III.1.2.2.4 Magnet Leg Module

The magnet leg module is housed separately from the remainder of the transceiver in a box located near the imaging magnet (Fig. 23-b). The magnet leg module itself performs five main functions (Fig. 23-a). First, the module blanks output noise from the RF amplifier by means of a PIN diode blanking switch. Second, the module switches the MR imaging coil between transmit and receive sides of the transceiver via a passive, quarter-wavelength TR (transmit/ receive) switch. Third the module provides a high gain, low-noise amplification block as the first stage of the receiver. Next the module implements an input pre-selector stage with a fixed frequency bandpass filter. Finally, the module provides an additional RF frequency gain stage to further boost the received MR signal level.

The PIN diode blanking switch prevents unwanted transmit RF amplifier output noise from leaking through the TR switch to the receiver during image acquisition. The blanking switch is an electronically controlled SPDT (single pole, double throw) PIN diode switch that is based upon a commonly used commercial design [43]. The blanking switch has been optimized for operation at frequencies between 5 and 10 MHz (Fig. 24) and is capable of handling input RF power levels of up to 100 W (50 dBm). The switch consists of three main circuits: the RF switch circuit (Fig. 24-a), the bias control network (Fig. 24-b), and the voltage reference circuit (Fig. 24-c). Each of these

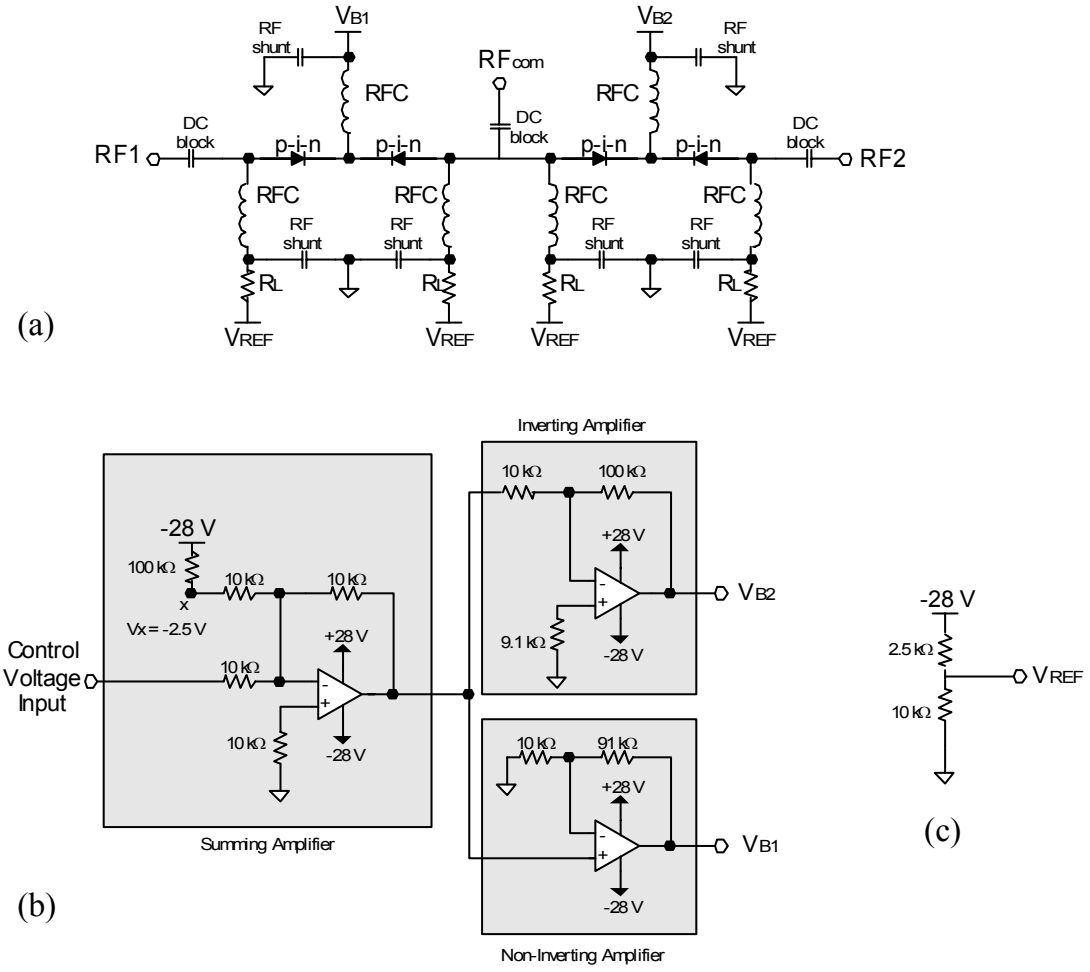

Fig. 24. PIN diode blanking switch circuit diagrams. (a) Circuit diagram for the SPDT PIN diode switch circuit used to blank the RF amplifier. (b) Circuit diagram for the bias control circuit for the switch. This circuit generates the two control voltages, V_{B1} and V_{B2}, which run the switch. (c) Voltage divider circuit for generating the reference voltage, V_{REF}, required by the switch.

circuits was constructed in-house, as a suitable blanking switch was not available commercially. Note: The op-amp chips used in the bias control are special high-voltage OPA445 operational amplifiers from Burr Brown-Texas Instruments (Dallas, TX). The bias control circuit takes a TTL-level input voltage and converts it to two bias control voltages, V_{B1} and V_{B2}, which are used to switch the PIN diodes in the RF switch circuit either 'on' or 'off.' The voltage reference circuit supplies the RF switch with V_{REF}, a $^{-}$ 22 V voltage level. When the applied control voltage is 5 V, the RFCOM port is isolated from the RF2 port by 60 dB, and signals at the switch inputs will travel only between the RFCOM and RF1 ports. Likewise, when the input control voltage is 0 V, ports RFCOM and RF1 are isolated by 60 dB, so incident RF signals at the switch terminals will only travel between the RF2 and RFCOM ports. To operate as a blanking switch for the RF amplifier, a 50 Ω load termination is connected to the RF2 port of the switch. A 5 V blanking control signal is then asserted during transmit to allow the RF pulse to pass and then unasserted during receive to switch the amplifier output noise into the termination resistor. This effectively blanks the output noise generated by the RF amplifier.

Table I
Active PIN Diode Blanking Switch Performance Summary at 6.75 MHz

Input Control Voltage [V]	S21 Measurement for the Indicated Signal Path [dB]	
	RF1↔RFCOM	RF1↔RFCOM
0	-4.02	-60
5	-60	-3.32

The completed blanking switch was tested at 6.75 and 9.15 MHz to quantify the isolation and transmission loss characteristics of the switch. The results appear in Tables

I and II. To measure S21 between the 2 ports of the switch, the indicated ports were connected to an HP4195A network/ spectrum analyzer (Hewlett-Packard, Palo Alto, CA) and the unused port was terminated in 50 Ω. When tested under actual operating conditions (as a blanking switch for the RF amplifier), the blanking switch was able to switch "on" in 132 μs. The turn off time was measured to be approximately 3.5 ms.

Table II
Active PIN Diode Blanking Switch Performance Summary at 9.15 MHz

Input Control Voltage [V]	S21 Measurement for the Indicated Signal Path [dB] RF1↔RFCOM	RF1↔RFCOM
0	-3.45	-61
5	-61	-2.78

The TR switch in the magnet leg module is a passive, quarter-wavelength switch (Fig. 25) which switches the RF coil between the transmit and receive portions of the transceiver during the collection of MR signals. This passive switch was used instead of the active PIN diode blanking switch described above because the switching times for the active switch were too slow for receiving fid signals. Also, using a second passive switch in addition to the active blanking switch provided an added layer of isolation between the transmitter and the receiver. When a transmit RF pulse is incident on the transmit port, the series crossed diode pair turns on allowing the RF pulse to propagate through to the coil port. At the same time a small dc voltage from the conducting series diode pair propagates to the shunt connected diodes and turns them "on". The short

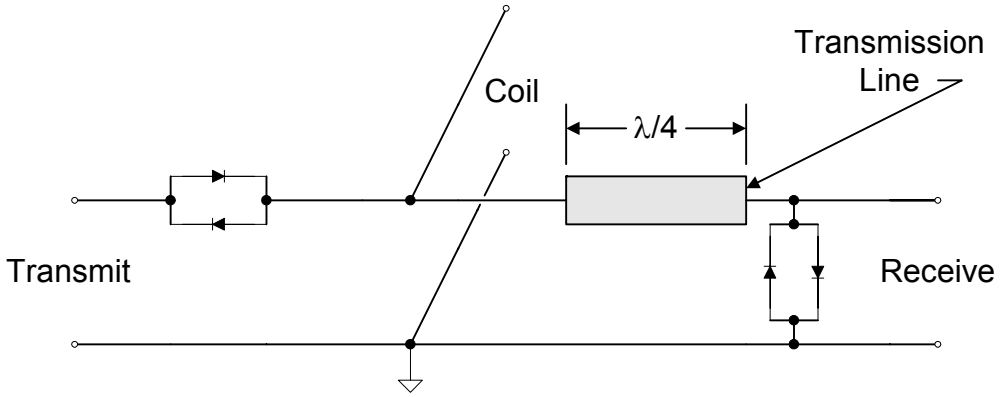

Fig. 25. Passive quarter-wavelength transmit/ receive (TR) switch. The transmission line is one quarter-wavelength long. The crossed diode pairs are MBD701 diodes from ON Semiconductor.

circuit formed by the conducting shunt diode pair is transformed by the quarter-wavelength transmission line to an open circuit back at the coil port terminal. Thus the impedance seen by the transmit pulse at the coil port appears to be nearly an open circuit, preventing the RF transmit pulse from leaking through to the receive port. During receive mode, the MR signal propagates from the coil port to the receive port. The voltage level of the MR signal, typically on the order of a few microvolts, is not high enough to turn on the series connected diode pair. Thus, there is no signal leakage from the coil port to the transmit port during receive mode. The switch was constructed from MBD701 diodes (ON Semiconductor, Phoenix, AZ) and a quarter-wavelength transmission line. The quarter-wavelength line is formed from a section of RG-58A/U coaxial cable cut to length for the imaging frequency.

The high-gain, low-noise amplification in the magnet leg module is provided by a low-noise preamp from Miteq (Hauppage, NY). The preamp is a Miteq model AU-1519-7480 low-noise amplifier. This preamp provides 60 dB of gain from 2-300 MHz with a noise figure of 1.2 dB. The preamp is available from the manufacturer for only $300.

The low-noise preamp in the magnet leg module is followed by a fixed frequency bandpass preselect filter. The preselect filter acts to reject the image frequency (refer to Chapter II section II.3.5.2) from entering into the IF stage of the receiver while limiting interference from coherent noise sources. The filter was constructed for a particular imaging frequency from low-cost inductors and capacitors using filter designs based on Butterworth, third-order, bandpass filter prototypes (Fig. 26). The performance of each

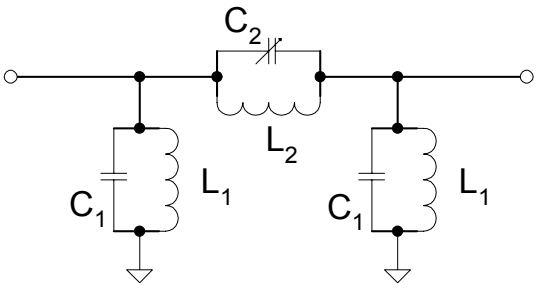

Fig. 26. Bandpass preselect filter schematic. The filter consists of three, LC tank circuits which are resonated at the center frequency of the filter passband. The center frequency of the filter may be adjusted by varying the capacitance of C2.

filter was simulated using *ARRL Radio Designer* (The American Radio Relay League, Newington, CT) software and measured using an HP 4195A network/ spectrum analyzer. Two filters were designed, one for 6.75 MHz (0.16 T) and one for 8.96 MHz (0.21 T). The 6.75 MHz filter has 14 dB of insertion loss and a 230 kHz, 3-dB bandwidth while the 8.96 MHz filter has a 5.6 dB insertion loss and a 750 kHz, 3-dB bandwidth. Since these filters have been packaged in their own cases, different filters may be easily interchanged to allow operation of the transceiver at different frequencies. The final section of the magnet leg module is the RF gain stage. This RF gain stage is comprised of an MAR-6SM single chip RF amplifier from Mini-Circuits (Brooklyn, NY). The MAR-6SM amplifier gives 20 dB of gain with a 3 dB noise figure over a frequency range from dc to 100 MHz. The MAR-6SM chip amplifier makes a highly modular gain block which can be easily added or removed from the receiver chain whenever additional gain is needed.

(a)

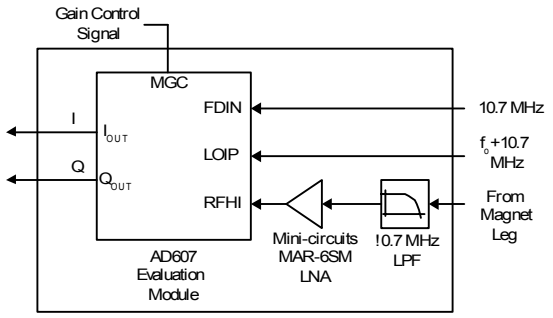

(b)

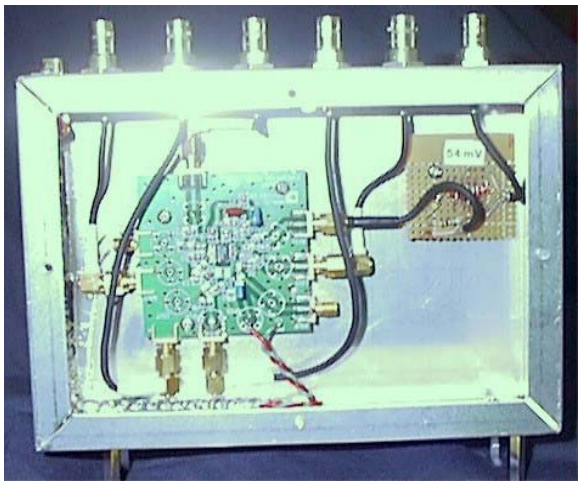

(c)

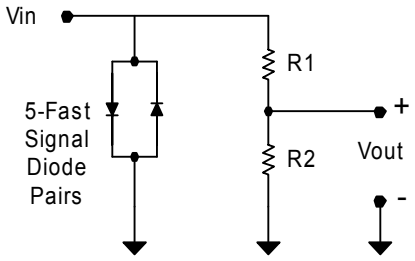

Vout	R1	R2
± 54 mV	47 Ω	1 Ω
± 400 mV	33 Ω	20 Ω

Fig. 27. I/Q demodulation block. (a) Functional schematic. (b) Picture of the I/Q demodulation module. Note the diode protection circuit in the upper right hand corner of the box. The populated demonstration module PC board for the AD607 single chip receiver appears in the center of the box. (c) Schematic for the diode protection circuits used to protect the AD607 board inputs. The LOIP and RFHI inputs require the ±54 mV output network; the FDIN input requires the ±400 mV output protection network. The protection network requires a ±0.7 V input signal and scales the output voltages via the voltage divider network. The input of the protection networks are matched to approximately 50 ohms.

III.1.2.2.5 I/Q Demodulation

The I/Q demodulation section of the transceiver (Fig. 27-a) is responsible for demodulating the MR signal at RF down to baseband I and Q channel output signals. This demodulation is performed using two separate downconversion stages. The first stage converts the MR signal to an IF frequency of 10.7 MHz while the second stage converts the IF frequency signal down to baseband. As such, there are three main components in the I/Q demodulation module: a low-pass filter, another RF gain stage, and a receiver. The operation of the demodulation module will be discussed in detail.

The filter and RF gain stages form the first part of the demodulation module. The filter is a BLP-10.7, 10.7 MHz low-pass filter from Mini-Circuits. The BLP-10.7 has less than 1 dB insertion loss from dc to 11 MHz. This filter is used to provide further rejection of unwanted high frequency signals outside the desired imaging band (since the simple preselect filter in the magnet leg does not have good rejection characteristics above ~70 MHz). The RF gain stage is simply another MAR-6SM amplifier module. This gain stage is only used to provide additional gain when imaging very small samples. The remainder of the time this secondary RF gain stage is simply bypassed.

The main section of the demodulation module is the receiver. The receiver is comprised of an AD607 single chip receiver from Analog Devices (Norwood, MA). The AD607 is a two-stage downconversion receiver on a single chip. It is capable of accepting RF input frequencies up to 500 MHz and demodulating them down to baseband I and Q channel outputs. The AD607 is available for $7 as a single chip or for $200 on a populated evaluation board module direct from the manufacturer. The

receiver portion of the desktop scanner makes use of an AD607 on an evaluation board (Fig. 27-b). For proper operation, the AD607 evaluation module requires only two local oscillator signals, an IF gain control voltage, and 3.3 to 5 volts of dc power (Fig. 27-a). The two local oscillator signals have the frequencies of 10.7 and (f_0+10.7) MHz and are provided by the frequency synthesis block of the transceiver. The IF gain control voltage is supplied by a divider circuit which has a 10-turn, precision potentiometer that is accessible to the system operator via the front panel of the desktop scanner. One limitation of the AD607 chip is that it accepts a maximum RF signal level of ¯15 dBm; however, this signal level was found to be sufficient for imaging small samples on low-field magnets. In order to protect the inputs on the AD607 from over-voltage, the diode protection network of Fig. 27-c is used. This allows larger signal levels than -15 dBm to be applied to the AD607 board without damage occurring. Each diode protection network is composed of five pairs of crossed, silicon, fast switching signal diodes in order to both allow the diode pairs to share the input RF current (in the event of a high power transient) and to provide redundancy to the protection network.

III.1.2.2.6 Anti-Alias Filtering and Sample-and-Hold Module

The anti-alias filter and sample-and-hold unit of the transceiver performs three functions in the processing of received MR signals (Fig. 28-a). The unit first removes the dc offset voltage from the output of the I/Q demodulation module. Then the unit filters both the I and Q channels to ensure that the signals will not be aliased by the digitizer. Finally, the signals are pre-sampled ahead of the digitizer to ensure synchronization of the digitizer to the transceiver master clock.

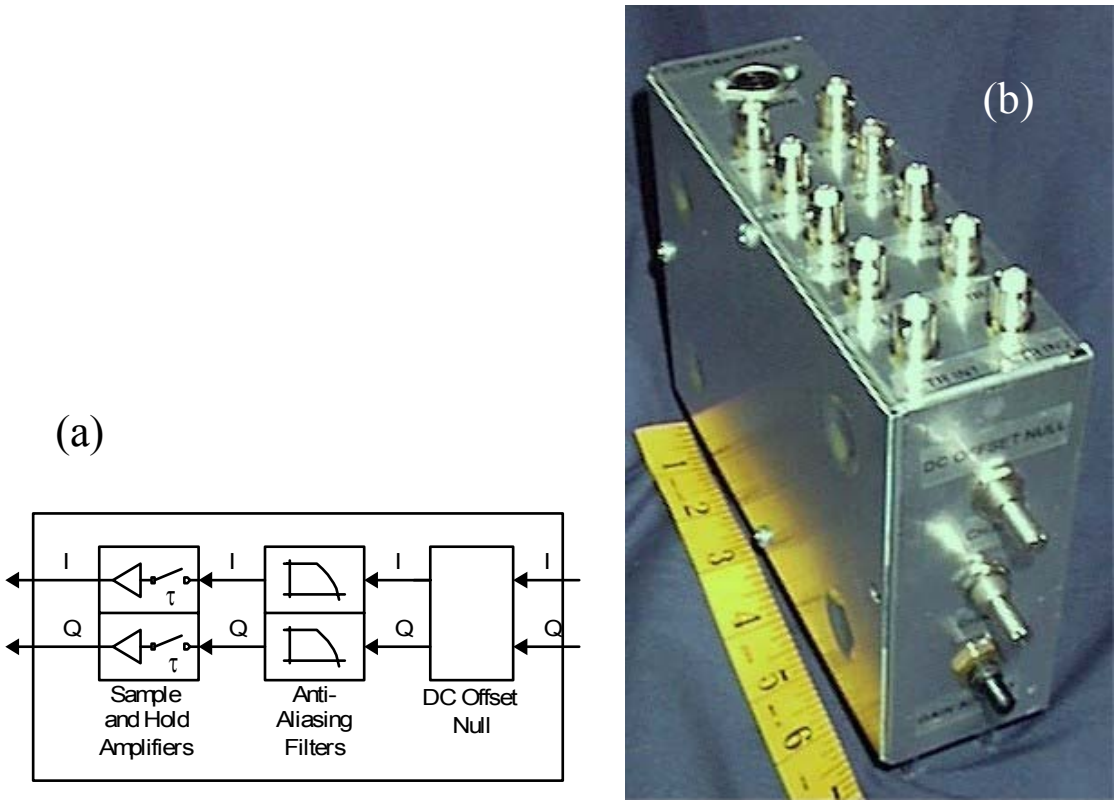

Fig. 28. Anti-alias filter and sample-and-hold unit. (a) Functional schematic. This module first strips the dc offset from the output of the I/Q demodulator. Then, the I and Q channel signals are filtered to prevent aliasing at the digitizer. Finally, the I and Q channel signals are simultaneously sampled. The sample-and-hold unit may be bypassed. (b) Picture of the actual anti-alias filter and sample-and-hold module.

One problem with using the AD607 chip as the receiver is that it automatically adds a variable dc offset voltage to the I and Q baseband output signals. If this dc offset voltage is not removed, a large dc artifact will show up in the reconstructed MR image. The dc offset voltage removal from I and Q channel signals is accomplished with a pair of op-amp summing circuits (Fig. 28-b). The I and Q output signals from the AD607 are each 'added' to a negative dc voltage level in order to null the dc offset voltage. The negative dc voltage level may be adjusted by a potentiometer which is mounted on the front panel of the transceiver. Each op-amp summing circuit was realized with a standard LM741 (National Semiconductor) op-amp and carbon composition, 0.25 Watt resistors.

The anti-alias filter block of the transceiver is crucial to proper digitization of the received MR signals. The anti-alias filters ensure that the Nyquist sampling theorem is satisfied by filtering out all signal content that is greater than half the sample rate of the digitizers. To implement the anti-alias filter, a model MSFS2P single chip, switched capacitor, 7 pole Butterworth lowpass filter from Mixed Signal Technology (San Jose, CA) was used. The MSFS2P has an electronically tunable cutoff frequency (1 Hz to 20 kHz) in addition to selectable on-chip gains (0, 10, or 20 dB). The cutoff frequency is controlled via the LabVIEW software while the gains are controlled by a switch mounted on the transceiver front panel (Fig. 28-b). Two MSFS2P filters were needed to filter both the I and the Q channels. The anti-alias filters limit the bandwidth of the desktop system receiver to a maximum of 20 kHz. (In the event that additional receiver bandwidth is needed, a model MSFHS2P filter from the same manufacturer (Mixed

Signal Technology, San Jose, CA) may be used. This chip is pin and function compatible with the MSFS2P chip filter, but has an electronically tunable cutoff frequency range from 10 kHz to 3 MHz.)

The sample-and-hold block is used to simultaneously pre-sample the received MR signals on the I and Q channels ahead of the digitizers. This allows the use of lower-cost digitizer cards, which often time multiplex one digitizer between two input channels and thereby introduce a timing skew between the I and Q channel signals. The sample-and-hold unit ensures both that the MR signals on the I and Q channels are sampled simultaneously and that the sampling is synchronous to the master system clock on the transceiver to prevent phase jitter. The sample-and-hold unit is comprised of two Analog Devices AD781 sample-and-hold amplifiers. Each AD781 is capable of 700 ns acquisition times and a full power sampling bandwidth of 1 MHz, which makes it well suited for this application.

III.1.2.2.7 Gradient Subsystem

The gradient subsystem (Fig. 29) of the transceiver is responsible for the amplification and filtering of the gradient waveforms generated by the computer control segment of the desktop scanner. To accomplish this, a 3-channel gradient amplifier and filter module was constructed. The gradient amplifiers are each based upon the LM12, a high power single chip op-amp from National Semiconductor which is able to drive a set of gradient coils at ±10 A. The gradient amplifier circuit is based on a current amplifier design [44], depicted in the LM12 datasheet [45], that was modified to include both dc offset and gain controls (Fig. 30). The gradient amplifiers were designed and

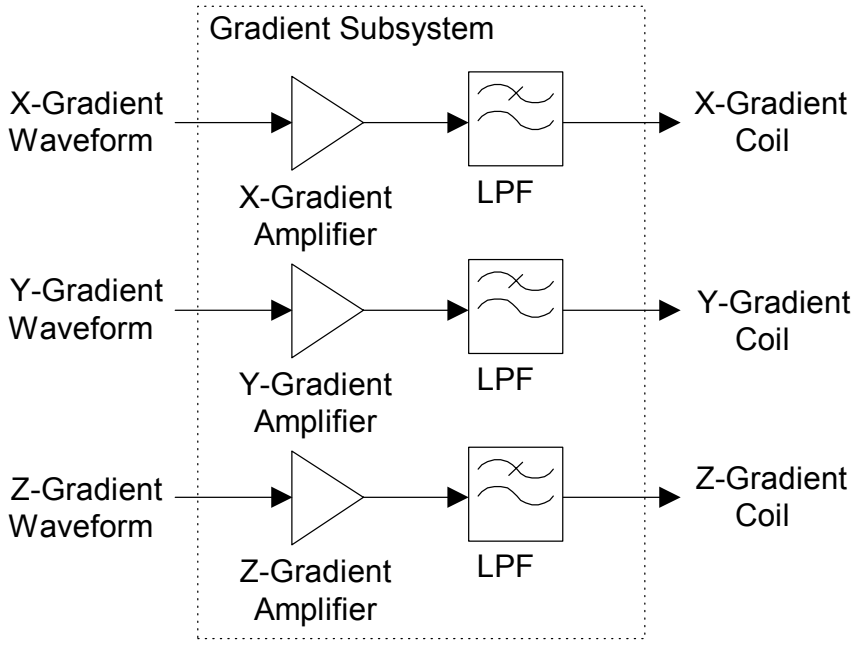

Fig. 29. Gradient subsystem block diagram. The gradient subsystem is composed of a 3-channel gradient amplifier followed by a 3-channel filter.

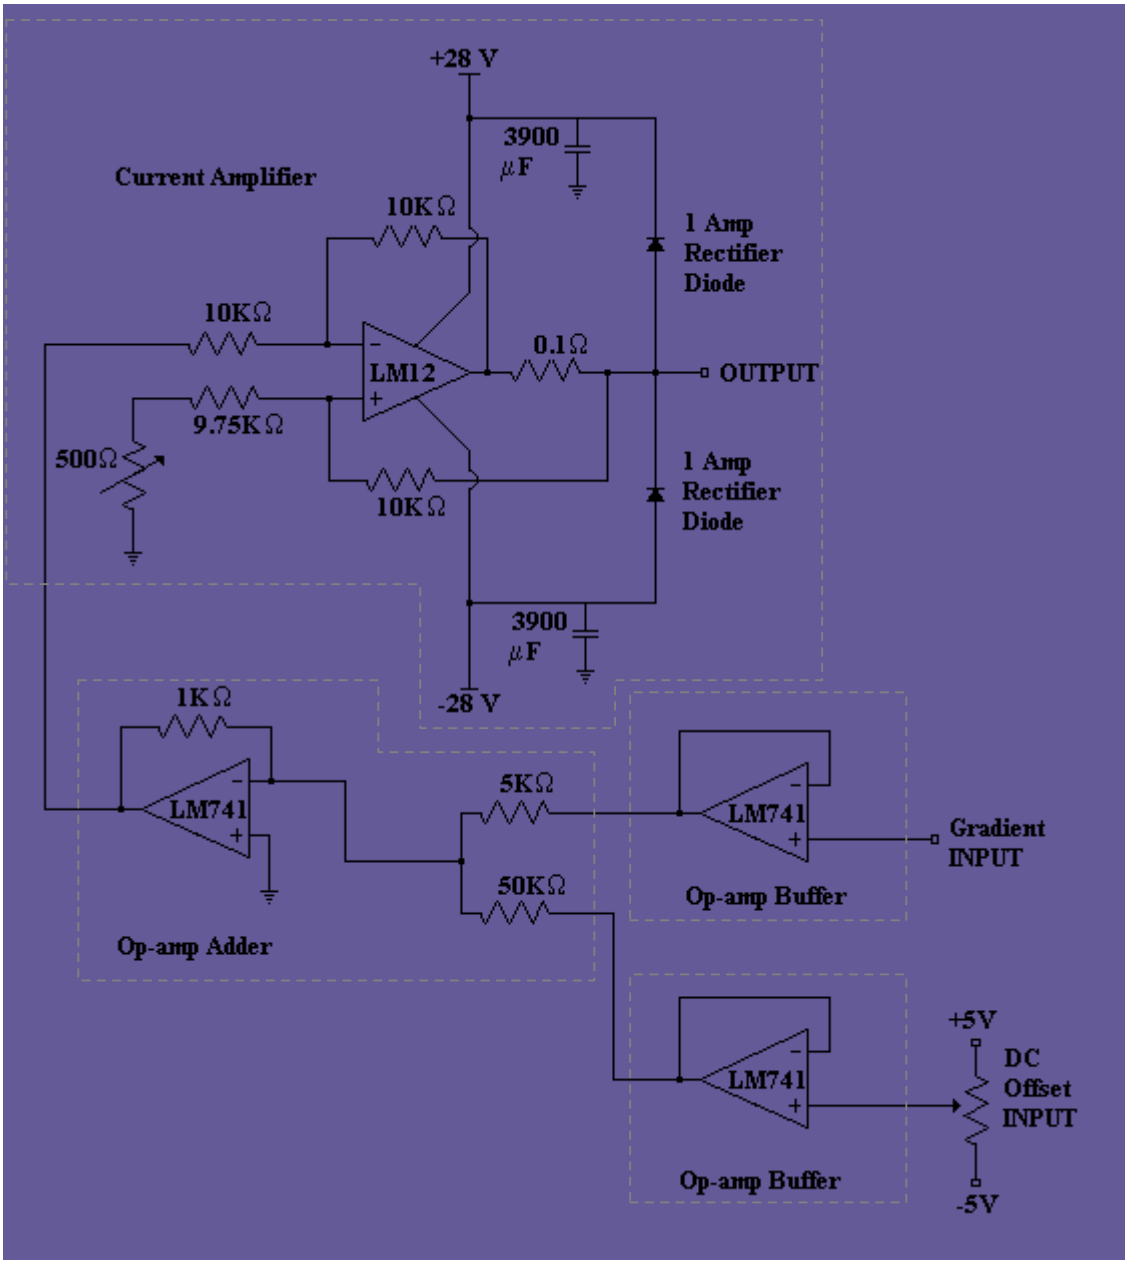

Fig. 30. Gradient amplifier schematic. The dashed lines enclose the four sub-circuits which compose the amplifier: two input buffers, an adder, and the current amplifier. The input buffers and adder circuit were added to the design of the current amplifier which was found in the data sheet for the LM12 op-amp. The input buffer and adder circuits allow a dc offset to be placed on the gradient waveforms for use as a static field shim. Note, also, that the LM741 op-amps are all biased with ±5 V.

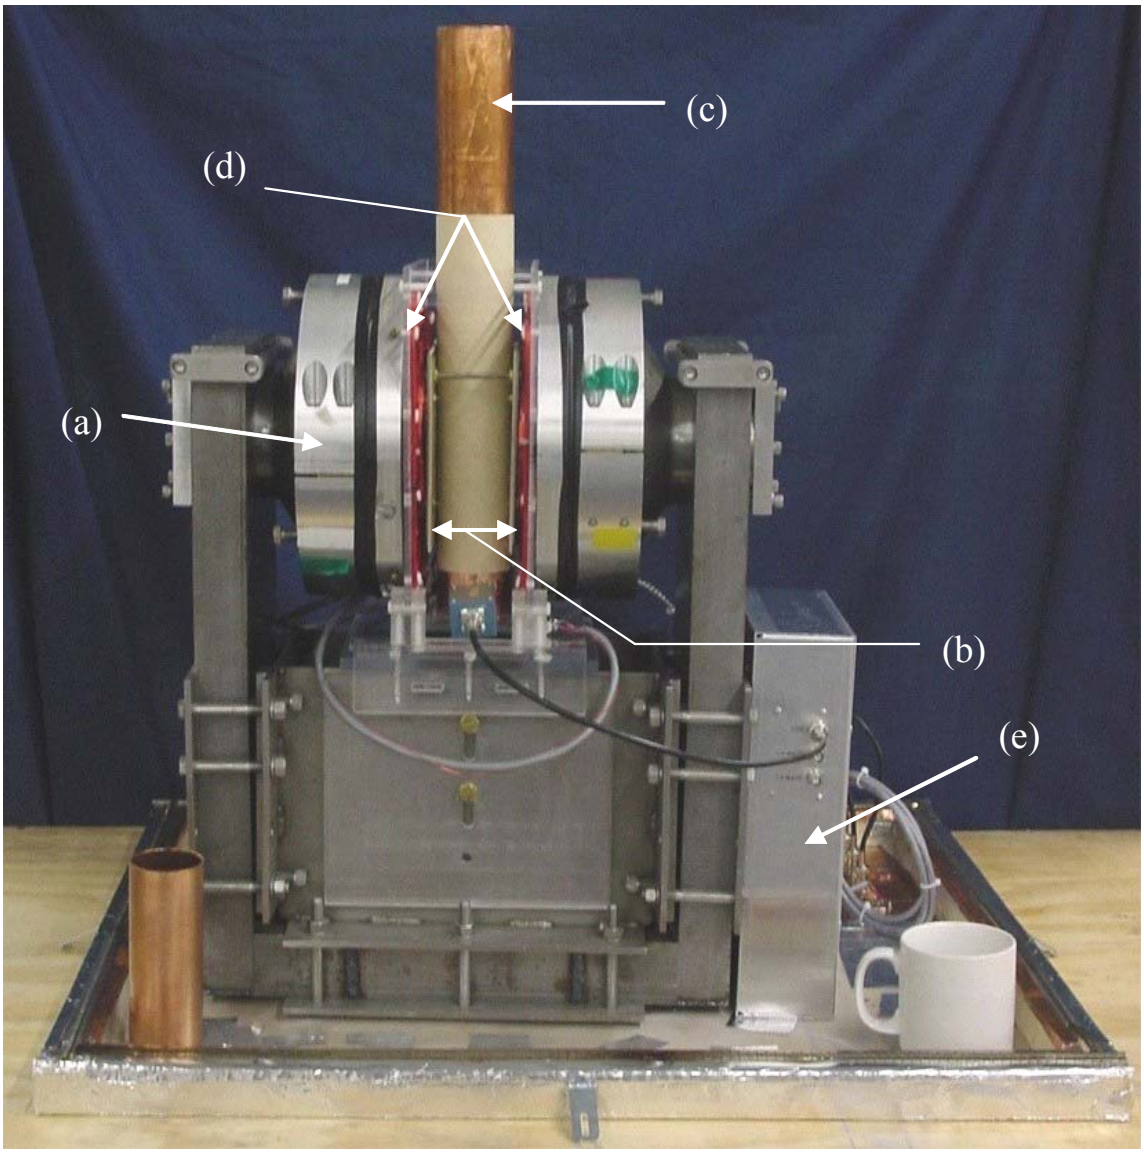

Fig. 31. Desktop MR scanner magnet subsystem. (a) Imaging magnet. (b) Gradient coils. (c) RF coil. (d) Electronic static field shims. (e) Magnet leg module from the transceiver.

constructed by David Spence, a Masters Thesis student in the Magnetic Resonance Systems Lab. The 3-channel gradient filter was implemented using three, 20 A, 60 Hz EMI power line filters (Tyco Electronics-Corcom, Exeter, NH).

III.1.2.3 Magnet Subsystem

The magnet subsystem (Fig. 31) in the desktop MR scanner has three main components. First and foremost is the imaging magnet (Fig. 31-a). The imaging magnet provides the homogeneous, static magnetic field which is required for MR imaging. The second component is the gradient coil set (Fig. 31-b) which is used to apply the static magnetic field gradients used to encode the spins in the sample being imaged. The final component of the magnet subsystem is the RF coil (Fig. 31-c) which both applies the RF magnetic fields used to tip sample spins during the imaging experiment and receives the MR signal from the sample.

III.1.2.3.1 Imaging Magnet

The imaging magnet (Fig. 31-a) used for the desktop MR scanner was previously constructed in the Magnetic Resonance Systems Lab at Texas A&M University as part of a collaboration with Professor Russ Huson of the Physics Department and his graduate students David M. Cole and Emilio Esparza-Coss. The design of the magnet has been fully described in other publications [46, 47]. The magnet is a 0.21 T, NdFeB permanent magnet having a C-shape design with a 10 cm gap and 17.5 cm pole face diameters. The magnet was constructed for ~ $3200 in material costs. A four element, static magnetic field shim coil (Fig. 31-d) and amplifier set were provided by Dr. Jay Porter and his students in the Department of Engineering Technology at Texas A&M

University. The dc currents on the shim coils must be adjusted manually using the current control potentiometers on the front panel of the shim amplifier set. After adjusting the electric shims, a half height linewidth of approximately 150 Hz at 8.96 MHz was achieved on the magnet (~20 ppm) using a cylindrical water phantom 1.9 cm (0.75 in.) long by 1.4 cm (0.55 in.) in diameter.

III.1.2.3.2 Gradient Coils

A custom gradient coil set (Fig. 32) was fabricated in-house for the desktop imaging magnet. The gradient coils were based on a standard bi-planar gradient coil design [48]. The gradient coil set was etched on six pieces of 15 cm (6 in.) × 15 cm (6 in.), 1.57 mm (0.062 in.) thick, double-sided, 8 oz., copper clad PC board. Together, the gradient coil set provides a 2.54 cm (1 in.) field-of-view with a sensitivity of 0.45 G/cm/A.

III.1.2.3.3 RF Coil

Several RF coils were constructed for imaging on the desktop scanner. Among these were surface coils and both solenoid and Helmholtz volume coils. The primary coil used for imaging on the desktop system is a solenoid coil (Fig. 33-a). This coil is a 5-turn, 1.8 cm (0.7 in.) long, shielded solenoid coil with a 3.35 cm (1.32 in.) inner diameter. The coil has an unloaded Q of 320 without a shield and 230 with a shield. When terminated in 50 Ω, the Q of the unloaded coil in the shield drops to 90, giving the coil a 3 dB bandwidth of 99.6 kHz at the 8.96 MHz operating frequency. The shield consists of a 6 cm (2.36 in.) diameter, 46 cm (18.25 in.) long circular waveguide shield (Fig. 33-b) that is closed on one end. The RF coil and sample holder (Fig. 33-c) were designed to image 2.54 cm (1 in.) diameter samples.

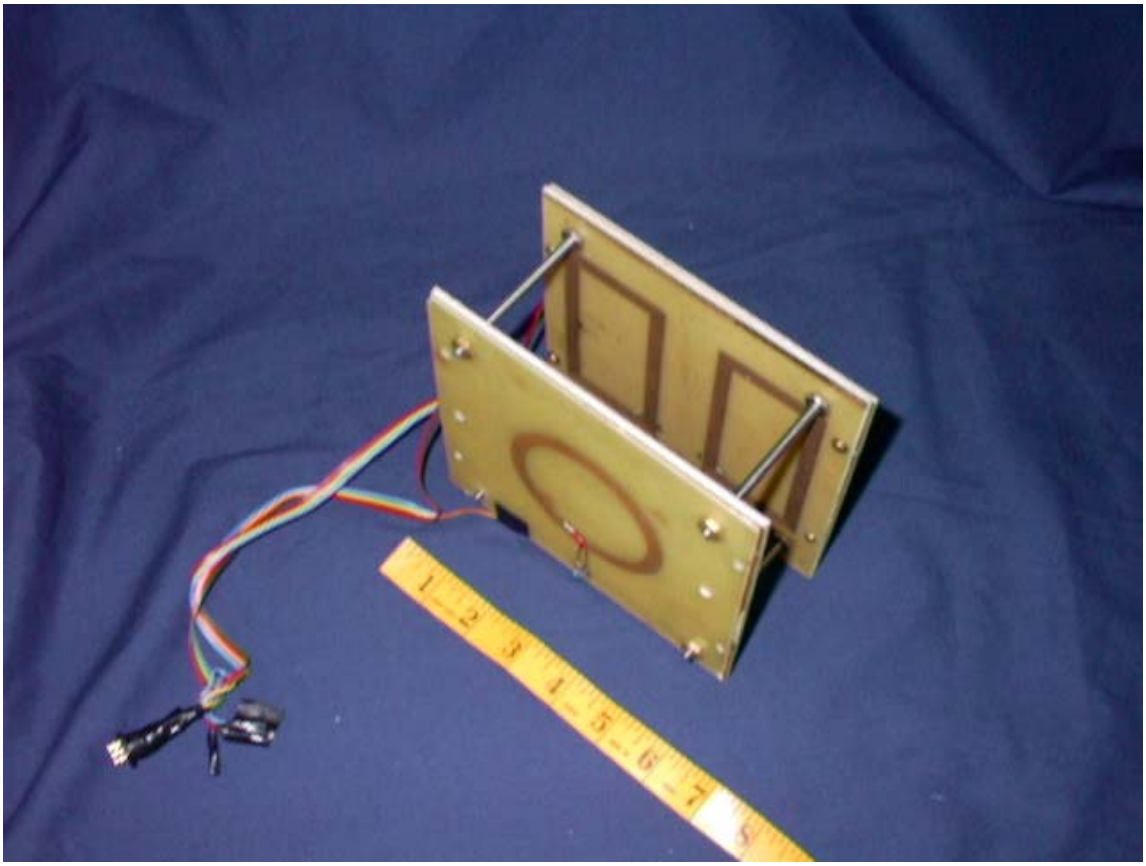

Fig. 32. Gradient coil set. The tape measure shown displays inches.

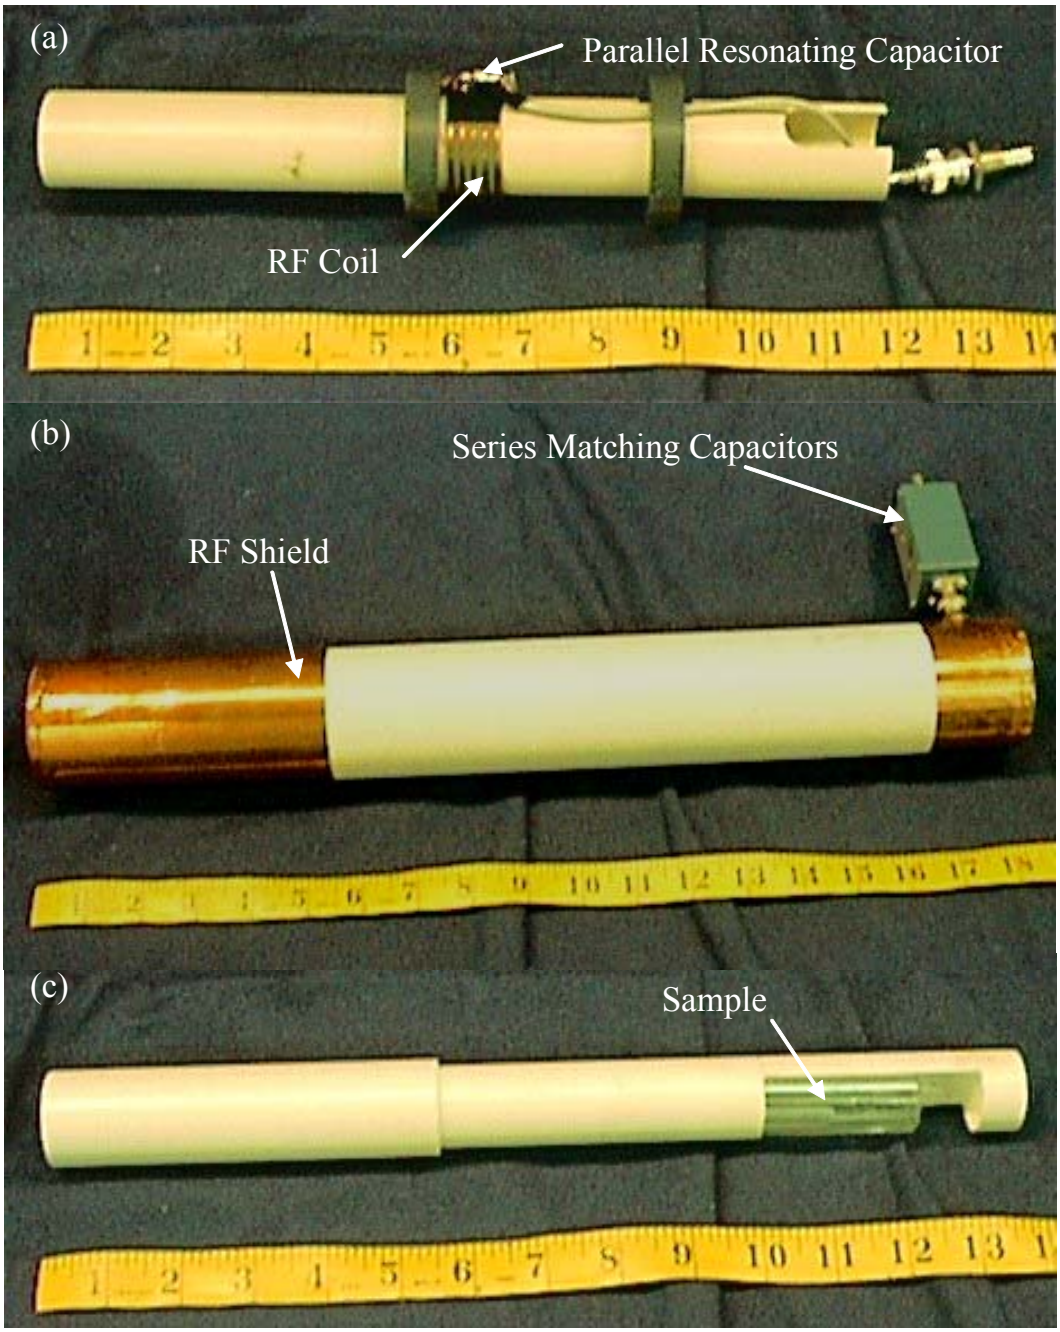

Fig. 33. RF imaging coil and shield assembly. The sample holder slides into the RF coil. Likewise, the RF coil fits inside the RF shield. The measuring tape shown indicates length in inches. (a) 5-turn RF solenoid volume coil with parallel resonating capacitors. (b) Circular waveguide RF shield and series matching capacitors. (c) Sample holder showing location of sample.

III.2 System Evaluation

The overall performance of the desktop MR system was evaluated. This consisted of benchtop tests, a cost analysis, and actual MR imaging trials to determine the operating characteristics of the system. The results of these evaluations are presented in the following sections.

III.2.1 Performance Summary

The performance of the desktop MR system was measured on the bench using an HP 4195A Network/ Spectrum Analyzer (Hewlett-Packard, Palo Alto, CA), PTS-160 frequency synthesizers (Programmed Test Sources, Littleton, MA), and a TDS-220 digital oscilloscope (Tektronix, Beaverton, OR). The transmitter is able to generate complex modulated RF pulses at power levels up to 20 W. The receiver has a dynamic range of 70 dB (measured using the single-tone dynamic range method discussed in Chapter II section II.4.2.5), and a maximum bandwidth of 20 kHz (limited at the anti-alias filter stage). By changing out the anti-alias filter chips (see section III.1.2.2.7), the receiver bandwidth could be expanded up to 102.4 kHz (the maximum bandwidth of the digitizer card in the control computer). The transceiver has an operating frequency range which extends from 0-11 MHz (without an external mixing module) or from 0-50 and 70-90 MHz (with the external mixing module). The desktop magnet system provides a 1.9 cm (0.75 in.) diameter homogenous field region at ~0.21. These performance characteristics have enabled the desktop MR system to image 1.9 cm (0.75 in.) fields-of-view at a field strength of 0.21 T, with a demonstrated resolution of 80 μm × 80 μm × 1000 μm.

Table III
Desktop MR System Cost Analysis

Module	Cost
Computer Control	$6000
Buffer	$160
Frequency Synthesis	$520
Pulse Shaping	$150
RF Amplification	$375
Magnet Leg	$675
I/Q Demodulation	$300
Anti-Alias Filter and Sample-and-Hold	$135
Gradient Subsystem	$550
Magnet Subsystem	$3200
Power Supplies, Cables, and Chassis	$700
Total:	$12,815

III.2.2 System Cost

The cost of the desktop MR system has been broken down in Table III. As may be seen in the table, the entire system was built for less than $13,000 in parts and materials. The cost of labor, engineering, and software development has not been included as all tasks were performed by graduate students working part time over a three year period. The most expensive portions of the scanner are the computer control ($6,000) and magnet subsystems ($3,200). The remainder of the desktop system was constructed from low-cost (and off-the-shelf when possible) circuits and systems. This is reflected in the cost of the remaining system components. In fact, neglecting the cost of the magnet, the MR system was built for less than $10,000. This price is significantly less than any MR scanner that is commercially available today.

III.2.3 SNR Comparison

In order to benchmark the SNR performance of the desktop scanner against a commercial system, image data sets were taken using both the desktop MR scanner and a Varian (Palo Alto, CA) SIS-85 MR system. The comparison image sets were obtained with each scanner using the same pulse sequences, gradient coils, RF coil, and test phantom (Fig. 34-a) on a 0.16 T whole body magnet (IGC-Magnet Business Group, Latham, NY). The 0.16 T whole body magnet mentioned here was used as a test platform during the initial development phase of the desktop MR system. A standard spin echo sequence (TE/ TR = 40 ms/ 1000 ms, 1 average, 2.5 kHz spectral width, 2 minute acquisition) was used to acquire the comparison image data. Representative images from the comparison data sets appear in Fig. 34. The SNR of the image from the desktop MR system (Fig. 34-b) is approximately 19.7 while the image from the Varian system has an SNR of about 19.4. (Measurement and calculation of the SNR in an MR image is discussed in Chapter II, section II.4.3.) Thus, the desktop MR scanner has demonstrated an SNR comparable to that of a commercial system transceiver operating at low-field and using the same MR magnet and gradient coil set .

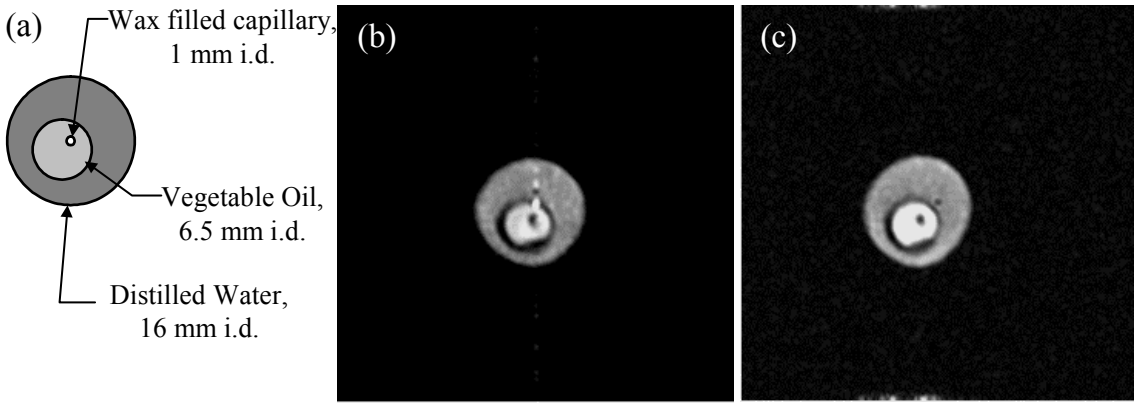

Fig. 34. SNR comparison images between the desktop MR system and a Varian SIS-85 commercial scanner. A standard spin echo sequence (TE/ TR = 40 ms/ 1000 ms, 1 average, 2.5 kHz spectral width, 2 minute acquisition) was used on each scanner. (a) Schematic of imaging phantom. (b) Image from desktop MR system. SNR~19.7. (c) Image from SIS-85 scanner. SNR~19.4.

III.2.4 2D Imaging

To further verify the performance of the desktop MR system, 2D slice select images of a test phantom (Fig. 35) were acquired. The test phantom (Fig. 35-b and Fig. 35-c) was constructed from a 3.5 cm (1.365 in.) long, 1.16 cm (0.46 in.) diameter, glass vial. A length of 2.2 mm (0.085 in.) diameter plastic tubing was coiled along the inside wall of the vial to form a slice thickness calibration phantom (Fig. 35-c). Then 3.3 cm (1.3 in.) long sections of both scrub brush bristles (0.38 mm (0.015 in.) diameter) and glass capillary tubes (1.5 mm (0.06 in.) outer diameter, 0.25 mm (0.01 in.) wall thickness) were placed inside the coils of tubing (Fig. 35-b). In all, seven sections of capillary tube (one filled with wax) were used. The entire phantom was then filled with 0.5 mmol $CuSO_4$(aq) solution which had been tinted with blue food coloring. Standard, axial, spin echo images were then obtained (TE/TR = 10 ms/ 500 ms, 120 averages, 1 mm slice thickness, 128 × 128 pixels) from the phantom. The image (Fig. 35-a) of the test phantom has a resolution of 80 μm × 80 μm × 1000 μm. Note that the image is distorted by an artifact which causes the outside edges of the vial to appear triangular rather than circular. This artifact is due to a couple of factors. First, the static field shim over the phantom was of a poor quality. The inadequate field shim could most likely be resolved with a better shim adjustment algorithm. A second factor is the fact that the gradient set is both unshielded and in very close proximity to the large, relatively flat metal pole faces of the imaging magnet. When short duration, high current gradient pulses (such as those used to perform slice selection for the 2D image) are applied to the unshielded gradient set, the currents on the gradients can cause eddy currents on the surface of the

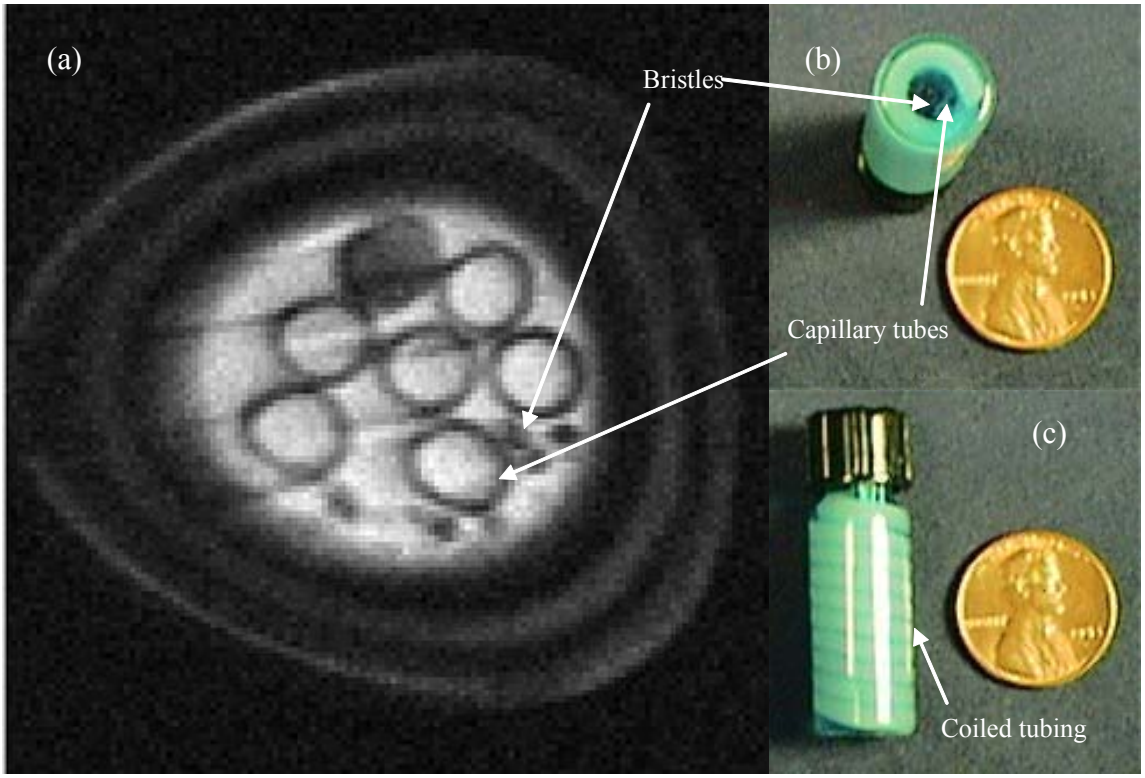

Fig. 35. 2D slice select image results from the desktop MR scanner. (a) Standard axial spin echo image of test phantom (TE/TR = 10 ms/ 500 ms, 120 averages, 1 mm slice thickness, 128 × 128 pixels). (b) End view of slice select imaging test phantom. The phantom was constructed from a 3.5 cm long, 1.2 cm diameter glass, screw top vial. A length of 2.2 mm diameter plastic tubing was coiled along the inside wall of the vial. Then 3.3 cm long sections of scrub brush bristles (0.38 mm diameter) and seven, glass capillary tubes (1.5 mm outer diameter, 1 mm inner diameter, six empty, one wax filled) were placed inside the coils of tubing. The entire phantom was then filled with 0.5 mmol $CuSO_4$(aq) solution which had been tinted with blue food coloring. (c) Side view of test phantom.

pole faces. These eddy currents can adversely influence the final MR image by affecting the linearity of the applied gradient fields during MR image formation. One way to resolve this problem is to use a 3D imaging sequence and thus eliminate the need for the slice select gradient pulses. Another method would be to replace the existing gradient coils with a shielded gradient set.

III.2.5 3D Imaging

3D MR image data sets were taken using a special test phantom (Fig. 36) constructed for use with the desktop system. The phantom was made from two pieces of machined, 0.64 cm (0.25 in.) thick, PVC plastic sheet (end plates) and nine glass hematocrit tubes. A section of insulation stripped from 24 AWG wire (0.94 mm diameter) was wrapped around four of the nine posts at a 30° pitch to aid in gauging slice thickness (Fig. 36-b). The assembly of hematocrit tubes and PVC end plates was then cemented inside a 16 mm o.d., glass test tube. The schematic of the machined PVC end plates (Fig. 36-a) depicts the phantom dimensions. The phantom has an overall length of 1.9 cm (0.75 in.) and a 1.4 cm (0.55 in.) diameter. The phantom was filled with 0.5 mmol $CuSO_4$(aq) solution. Standard 3D spin echo images (256×128×64 points, TE/TR = 250 ms/ 9.46 ms, 2.5 cm (1 in.) × 2.5 cm × 2.5 cm field-of-view, 32 averages, 18.2 h acquisition) were then taken of the phantom using the desktop system. Three central planes from the image data set appear in Fig. 37. The data set has been zero-padded in order to allow a displayed resolution of 256×256×64 points (100 μm × 100 μm × 400 μm). Thirty central slices out of the sixty-four acquired axial plane slices are also shown (Fig. 38). The spiral wound tubing is apparent. In Fig. 37 and Fig. 38, a slight image distortion is

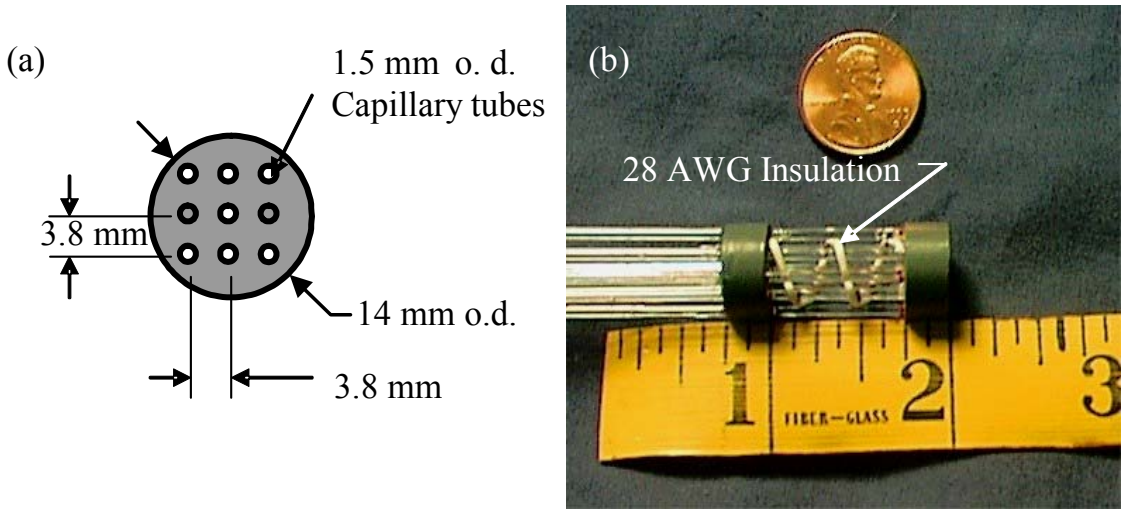

Fig. 36. 3D imaging phantom constructed for use on the desktop MR system. The phantom was made from two pieces of machined, 6.4 mm thick, PVC plastic sheet and nine glass hematocrit tubes that were cemented inside a 16 mm o.d., glass test tube. A section of insulation stripped from 28 AWG wire was wrapped around four of the nine posts at a 30° pitch to aid in gauging slice thickness. (a) Schematic diagram of the phantom end plate. The seven capillary tubes shown with a white center are filled with air while the two tubes shown with a gray center are filled with a relaxed water solution. (b) Photograph of the imaging phantom with the insulation indicated. Measuring tape shown is in inches.

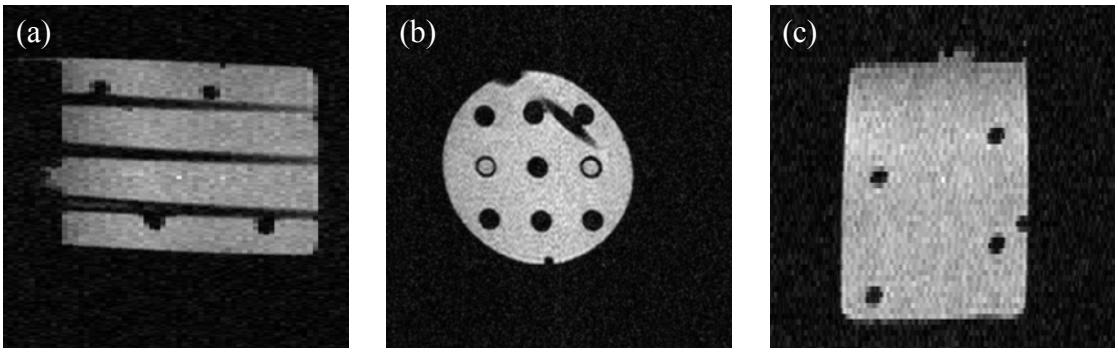

Fig. 37. 3D spin echo images of the test phantom. Images were taken on the desktop MR system with 256×128×64 points, TE/TR = 9.46 ms/ 250 ms, 2.5 cm (1 in.) × 2.5 cm × 2.5 cm field-of-view, and 32 averages for an 18.2 h acquisition. The center slices from 3 planes of the 3D data set are shown. (a) Sagittal plane image, 256×64 point display. (b) Axial plane image, 256×256 point display. (c) Coronal plane image, 64×256 point display.

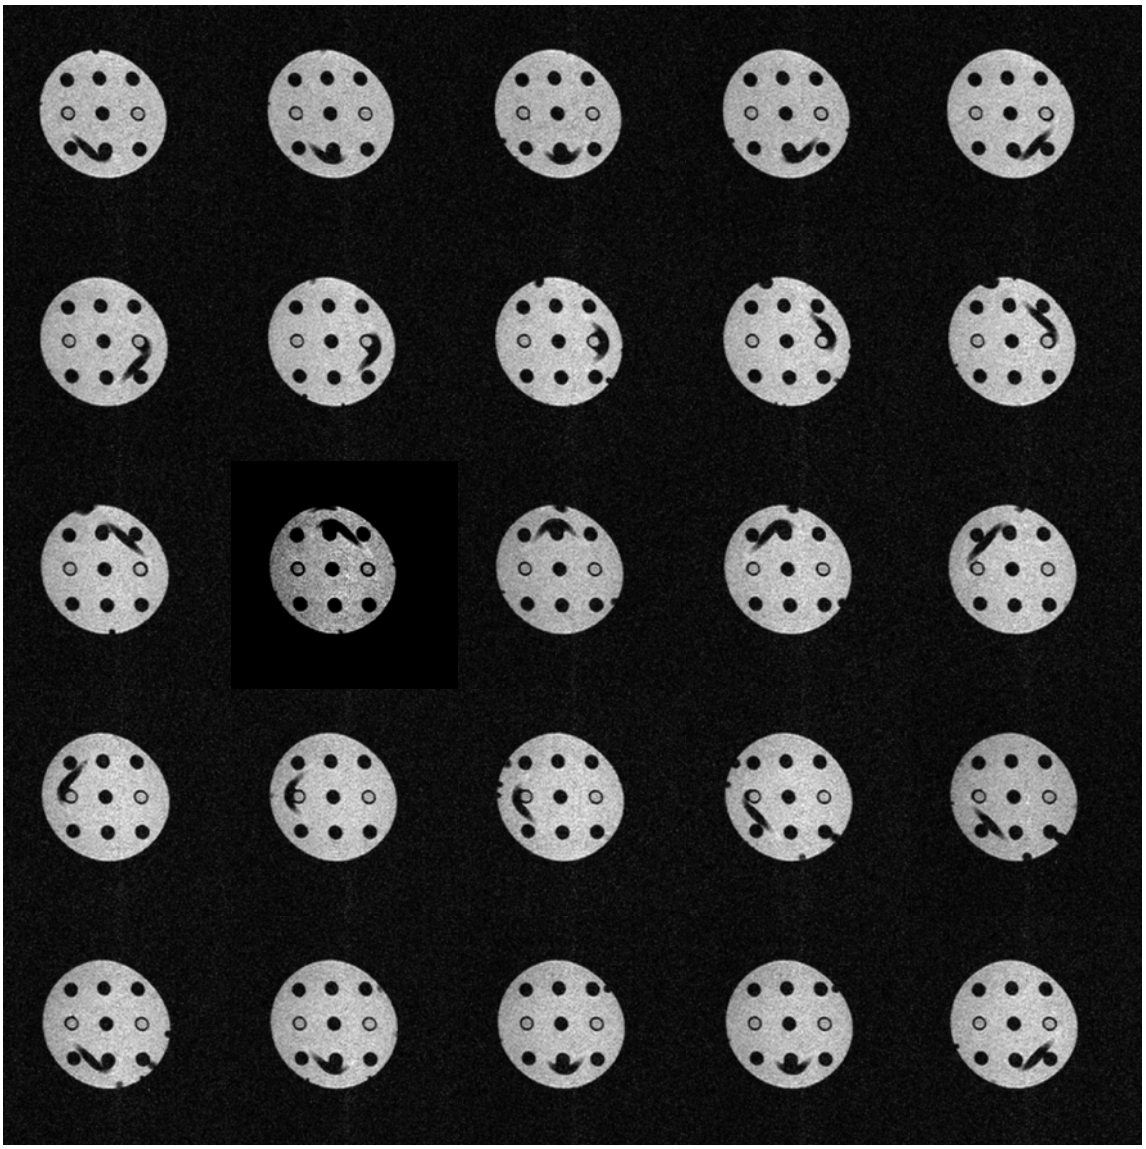

Fig. 38. Thirty central slices out of 64 axial slices from the 3D spin echo data set. Slice progression is right-to-left and top-to-bottom. The central slice (number 32) is the first image in row four. Field-of-view was 2.5 cm × 2.5 cm × 2.5 cm, with a 256×128×64 point acquisition, interpolated to 256×256×64 points for a displayed resolution of 100 μm × 100 μm × 400 μm. TE/TR = 9.46 ms/ 250 ms, and 32 averages were acquired for an 18.2 h acquisition time. The spiral wrapped tubing is evident in the images.

noticeable. This distortion causes the axial plane images to appear slightly out of round, and the outsides of the sagittal and coronal images to appear trapezoidal rather than rectangular. The distortion is due to a poor static field shim which the active shim set was unable to correct.

To evaluate the performance of the desktop MR scanner on a biological sample, 3D spin echo (TE/TR = 9.4 ms/ 300 ms) image sets were acquired of a mouse phantom (Fig. 39). The phantom was a formed from a juvenile mouse that was obtained "fresh frozen" from a local pet store where they are sold as food for snakes. The mouse carcass was 4 cm in length (without the tail). In all, three separate 3D image sets were taken (Fig. 40). The same axial and sagittal slices have been displayed from each of the three sets for comparison purposes. First, a thick slice image set (256×128×16 points, 20 averages, 3.4 h) was obtained (Fig. 40-a). This data set has relatively poor SNR and a lot of blurring in the sagittal view due to the slice thickness. Thus, resolution and SNR have been sacrificed for a lower imaging time. Next, a lower axial resolution set (128×128×32 points, 26 averages, 8.9 h) with thinner slices was acquired (Fig. 40-b). This data set has the best SNR of the three. There is also more detail present in the sagittal view, but less in the axial. Here a compromise has been reached between resolution, SNR, and imaging speed. Finally, a high axial resolution set with even thinner slices (256×128×64 points, 20 averages, 13.6 h) was taken. This last set has the highest resolution, but also the highest noise level and longest imaging time of the three.

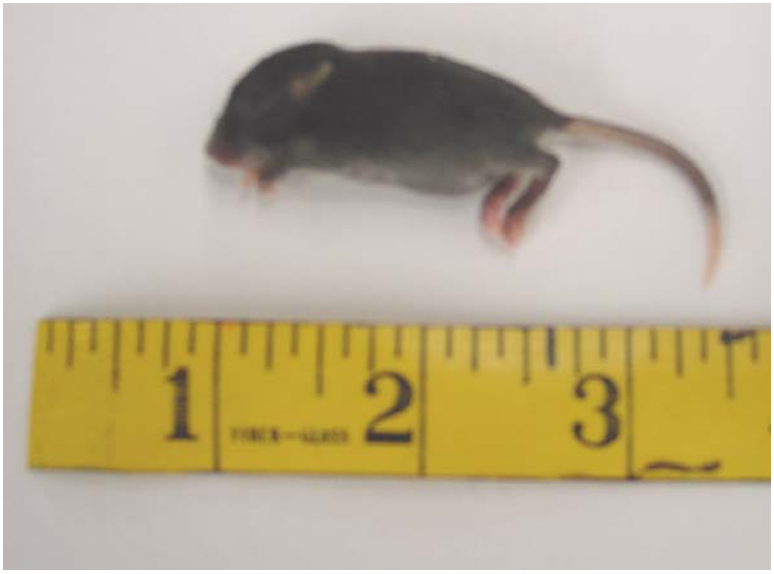

Fig. 39. A "fresh frozen" juvenile mouse phantom. Measuring tape shown is in inches.

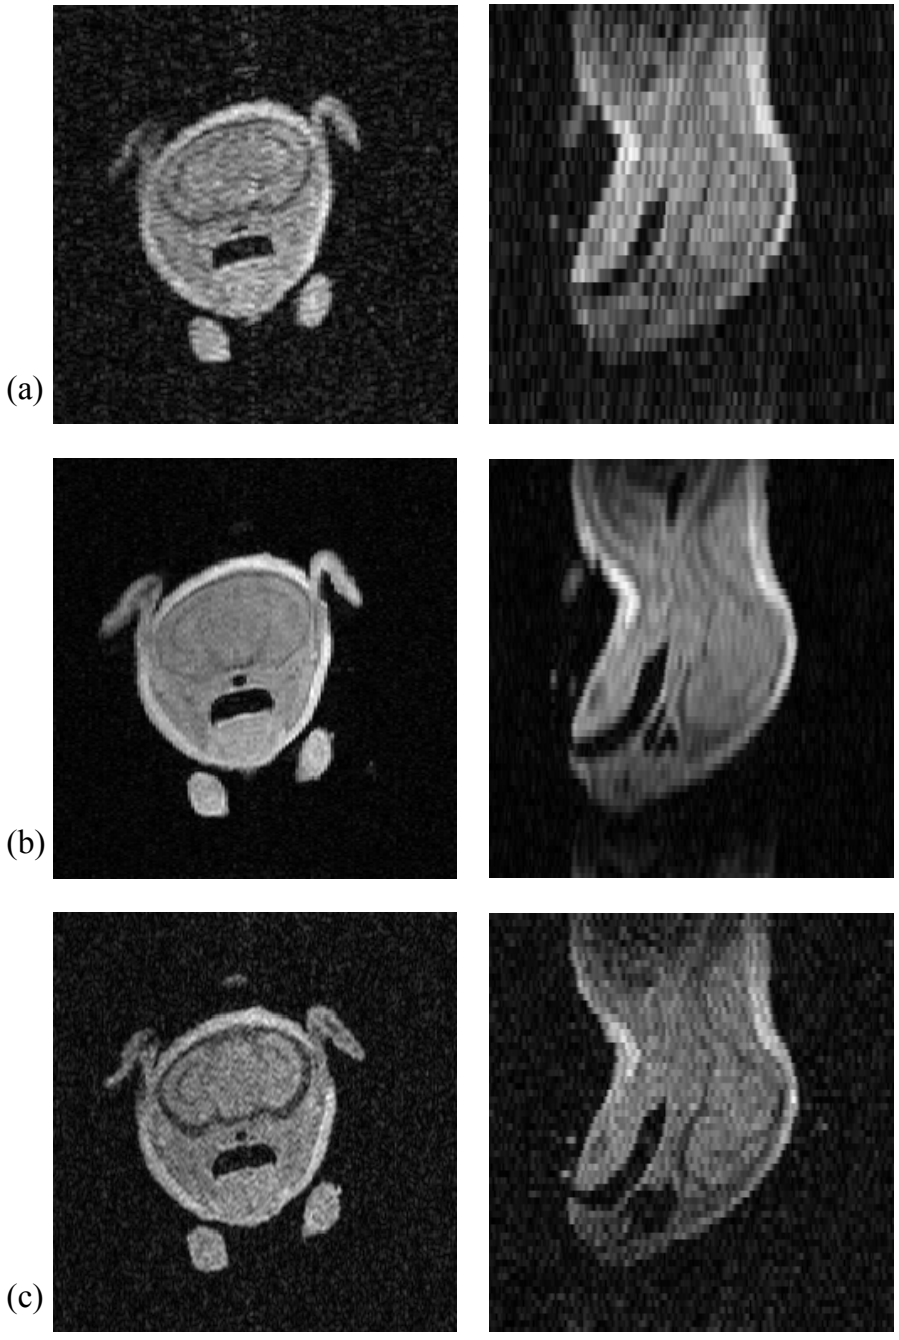

Fig. 40. 3D spin echo (TE/TR = 9.4 ms/ 300 ms) image data sets of a "fresh frozen" juvenile mouse acquired from the desktop MR system. Two images each (one axial and one sagittal slice) from three different 3D data sets are shown. (a) 256×128×16 points, 20 averages, 3.4 h. Displayed resolution is 100 μm × 200 μm × 1600 μm. (b) 128×128×32points, 26 averages, 8.9 h. Displayed resolution is 200 μm × 200 μm × 800 μm. (c) 256×128×64 points, 20 averages, 13.6 h. Displayed resolution is 100 μm × 200 μm × 400 μm.

III.3 Summary

A low-cost, desktop MR system has been built which is able to image small volumes at low field strengths. The design and construction of the desktop scanner have been described in detail. The system performance has been characterized through a series of bench measurements and imaging tests. These tests show that the desktop MR imaging system is well suited for imaging 1.9 cm (1 in.) fields-of-view at a field strength of 0.21 T, with a demonstrated resolution of 80 μm × 80 μm × 1000 μm. This prototype shows that low-field MR scanners, which have performance comparable to commercial systems, can be developed inexpensively. These systems may pave the way for a new type of low-cost MR scanner for use as a general laboratory instrument and other non-clinical applications.

CHAPTER IV

A SCALABLE, SINGLE-CHANNEL DIGITAL MR RECEIVER

In the first phase of the project to develop a parallel receiver for MR imaging, it was necessary to develop a cost-effective, single-channel, digital MR receiver. This receiver would then serve as a building block for a highly parallel MR receiver system. The receiver was designed specifically for the purpose of imaging small volume samples (i.e., samples with dimensions of 3 mm × 3 mm × 100 mm or less) with planar pair RF surface coils at 4.7 T. A principle goal of the design was to make the prototype receiver as good as possible given the constraint that it remain a low-cost system. As such, the performance of the receiver had to be comparable to that of a commercial analog receiver. A final prerequisite for the single-channel receiver design was that it could be easily scaled up to form a 64-channel parallel receiver. This chapter will present the methods used to meet these design criteria, the details of the final single-channel receiver design, and the steps taken to verify that the single-channel receiver prototype actually met the design requirements.

IV.1 System Design Methods

To design the single-channel receiver prototype three main steps were taken. First, the design specifications for the receiver were derived. These requirements were developed based upon the need for a cost-effective, digital receiver which could be easily scaled to create a 64-channel system. Second, the actual prototype design was developed. This development took place in a series of steps which led to the creation of

a detailed schematic for the receiver prototype. Finally, a prototype was constructed from the design schematic and tested to verify the integrity of the design.

IV.1.1 System Requirements

The specifications of primary concern for an MR receiver are bandwidth, operating frequency, noise figure, dynamic range, and the characteristics of the digitizer. The receiver bandwidth determines the types of pulse sequences with which the receiver will be able to operate. The operating frequency range of the receiver establishes the field strengths at which the receiver will operate. The noise figure is an indicator of how much the receiver will degrade the SNR of the detected signals. The dynamic range determines the largest and smallest signal levels that the receiver will be able to detect. Finally, the digitizer handles the final detection of the signal. The derivation for each of these requirements for the single-channel receiver will be detailed in the following sections.

IV.1.1.1 Bandwidth

The bandwidth required for an MR receiver can be as little as a few kilohertz. It should be noted, however, that the bandwidth of the received MR signal is determined by the pulse sequence used to form the MR image. While a receiver bandwidth of 20 kHz may be wide enough for many pulse sequences (such as those used for the desktop MR system) it is not wide enough for many of the faster imaging sequences. In order to overcome this limitation, the target bandwidth for the single-channel receiver was set at a maximum value of 1 MHz. This means that the single-channel receiver should be able to detect signals with bandwidths of up to 1 MHz.

IV.1.1.2 Operating Frequency Range

While the required bandwidth of the receiver is only 1 MHz, the desired operating frequency range is considerably higher. This enables imaging over a wide variety of field strengths. The principle operating frequency for the receiver design was set at 200 MHz to enable MR imaging on 4.7 T magnets. In order to allow imaging at field strengths from 1.5 to 4.7 T, the receiver was also designed to operate over the 60 to 210 MHz frequency range.

IV.1.1.3 Noise Figure

As discussed in Chapter II section II.4.1, the noise figure for a cascaded system may be set by using a high gain, low noise device in the first stage. In order to accomplish this, the first stage of an MR receiver system always consists of a low-noise preamp. Several manufacturers produce low-noise preamps for MR applications with varying bandwidths and noise figures.

Ar^2 Communication Products (Burlington, CT) provides narrowband, low-noise MR preamplifiers with noise figures of 0.5 dB and gains from 16 to 26 dB. These amplifiers operate at center frequencies ranging from 2.3 to 630 MHz, but generally have an operating bandwidths limited to about 7%. The cost per preamp unit varies from $137 to $147.

Another large MR preamplifier manufacturer is Miteq (Hauppauge, NY). Miteq offers a wide selection of wideband, fast-recovery, low-noise preamps that feature input crossed diode protection. These amplifiers are available with noise figures from 1.2 to 3.0 dB and gains ranging from 14 to 60 dB. The Miteq NMR amplifiers have operating

bandwidths of 400-500 MHz spanning the frequency range of 0.2-600 MHz. The cost per unit for these preamps starts at approximately $300.

A third manufacturer of low-noise MR preamps is Doty Scientific (Columbia, SC). Doty offers a line of four, wideband, low-noise preamps with noise figures of less than 1 dB over parts of their operating range and average gains of 31 dB. The operating bandwidth of Doty preamps ranges from 50-720 MHz over the 0.5-800 MHz operating band. The per unit cost for these amplifiers ranges from $1215 to $1290.

Since the noise figure determines the extent that a receiver will degrade the SNR of an input MR signal, the noise figure for a receiver system should be set as low as possible. Noise figure alone, however, was not the only factor used to determine the noise figure requirement for the single-channel receiver prototype. The cost of the preamp unit itself was also an important consideration (especially since the single-channel receiver was designed specifically to act as a building block for a sixty-four channel parallel receiver system). As a result, the MR preamps available from both Miteq and Doty Scientific, while providing a satisfactory noise figure, were too expensive to be used in the single-channel receiver prototype. Likewise, while the cost of the preamps from Ar^2 Communication Products was fairly reasonable, the bandwidth of the preamps was not wide enough to meet the receiver's operating frequency range requirement (see section IV.1.1.2). As a result, an alternative source had to be found for the low-noise preamplifiers to be used in the single-channel receiver. To allow the possible use of other, lower-cost, wideband preamp circuits, the design requirement for the noise figure of the receiver was set at 2 dB (or less).

IV.1.1.4 Dynamic Range

The dynamic range of a receiver is defined formally in Chapter II section II.4.2. This characteristic of the receiver is of particular interest as it determines the range of signal power levels which the receiver is able to detect. In order to determine the dynamic range requirement for a receiver, estimates for the minimum detectable signal (MDS) and maximum signal levels were made. The MDS level estimate was based upon the thermal noise floor resulting from the 1 MHz bandwidth requirement of the single-channel receiver (see section IV.1.1.1). The maximum signal estimate was made using previously measured signal levels from a prototype planar pair RF surface coil. These two estimates were then used to compute the minimum dynamic range requirement for the single-channel receiver.

IV.1.1.4.1 MDS Estimate

The MDS level may be estimated from the sum of the r.m.s. thermal noise floor and the system noise figure (2.48), as described in Chapter II section II.4.2.4.1. The thermal noise power may be determined from the ambient temperature, T [K], and the receiver bandwidth, B [Hz]. Since the required bandwidth for the receiver was set at 1 MHz (see section IV.1.1.1) and the noise figure was set at 2 dB (refer to section IV.1.1.3), the MDS estimate for the single-channel receiver (at a room temperature of T = 300 K) was found to be 6.31×10^{-12} mW, or -112 dBm.

IV.1.1.4.2 Maximum Signal Level Estimate

To obtain an estimate of the maximum signal level, a small, prototype planar pair RF surface coil was used to acquire spin echo signals on a GE Omega 4.7 T magnet system.

Small volume relaxed water samples were used (slightly larger than the target size of 3 mm × 3mm × 100 mm). The maximum received power levels from the acquired spin echos were measured to be on the order of ˜62.2 dBm.

IV.1.1.4.3 Dynamic Range Requirement

The dynamic range requirement for the single-channel receiver was computed by subtracting the MDS estimate from the maximum signal estimate (both in dBm). This calculation yields a required dynamic range of 49.8 dB. Since the single-channel receiver was designed primarily for imaging, the dynamic range requirement was set via the second estimate to be a minimum of 50 dB.

IV.1.1.5 Digitizer

The choice of the digitizer is crucial to the design of the MR receiver. In particular, a digitizer must be selected which can properly digitize the full bandwidth of the MR signal which is to be received by the receiver. This may be done using either direct sampling or undersampling techniques (see Chapter II, sections II.3.1. to II.3.4) for a discussion). Since the single channel receiver was designed in order to be easily scaled into a multiple channel receiver, the digitizer for the single channel prototype must also be chosen to easily accommodate multiple channels without sacrificing performance or adding undo cost to the receiver design. Additionally, the digitizer used in an MR receiver requires the dynamic range afforded by at least 16 bits of analog to digital conversion.

IV.1.2 System Design

A single channel receiver prototype was designed to meet the requirements for bandwidth, operating frequency, noise figure, dynamic range and the digitizer detailed in section IV.2.1. First, the prototype was designed to have an overall bandwidth of 1 MHz. To meet the wide operating frequency range requirements, the prototype was designed as a single conversion digital receiver without an image reject filter (see Chapter II, section II.3.5.3). The operating frequency for the prototype ranges from 2 to 500 MHz. Next, the noise figure was set to be less than 2 dB. The dynamic range for the prototype was designed to be at least 50 dB. Finally, to meet the digitizer requirements, an ICS-645 (Interactive Circuits and Systems, Ontario, Canada) was selected. The ICS-645 is a PCI card digitizer with thirty-two, independent, 16-bit, 2.5 MSPS (1.25 MHz bandwidth) input channels. This makes it an ideal choice for eventual expansion to a receiver system with sixty-four, 1 MHz bandwidth channels.

The final design schematic for the single channel MR receiver prototype appears in Fig. 41. In the receiver shown, the MR signal is first amplified at RF by a low-noise amplifier stage. The level of the RF signal from the output of a low-noise preamp is then adjusted by an RF frequency variable gain stage. The received signal is next downconverted to an IF frequency of 500 kHz. Afterwards the IF frequency signal is again amplified and then bandpass filtered. Finally, the received signal is digitized and processed on board the PC control computer. The design and function of each section of this receiver will be discussed in detail.

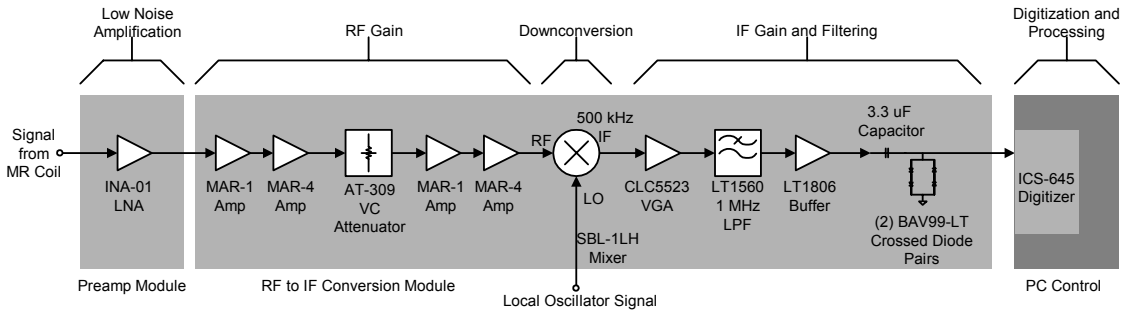

Fig. 41. Schematic for the single-channel receiver.

IV.1.2.1 Low Noise Amplification

Low-noise amplification of the received MR signal occurs in the preamp module which is located close to the MR coil. In order to reduce the overall cost of the preamp module, an inexpensive, single chip, 50 Ω, MMIC (monolithic microwave integrated circuit) amplifier was used to form the preamp circuit. The preamp circuit itself (Fig. 42) is composed of a shunt PIN diode switch followed by an Agilent (Palo Alto, CA) INA-01 single chip amplifier. The INA-01 chip delivers 32 dB of gain with a 1.7 dB noise figure over the frequency range of dc to 500 MHz. The PIN diode switch protects the amplifier chip from RF pulse breakthrough during the transmit portion of the MR imaging sequence. The switch turns 'on' during the application of the TTL-level RF blanking pulse from the MR scanner, preventing the transmit pulse from reaching the preamp chip. For normal operation, the switch remains in the 'off' state, allowing the MR signal to reach the low-noise preamp.

IV.1.2.2 RF Gain Subsystem

The RF gain section of the receiver consists of four amplification stages and one intervening voltage controlled attenuation stage. The variable attenuation stage allows adjustment of the total amount of RF gain for the receiver while increasing the receiver dynamic range. The RF gain and attenuation stages will be discussed below.

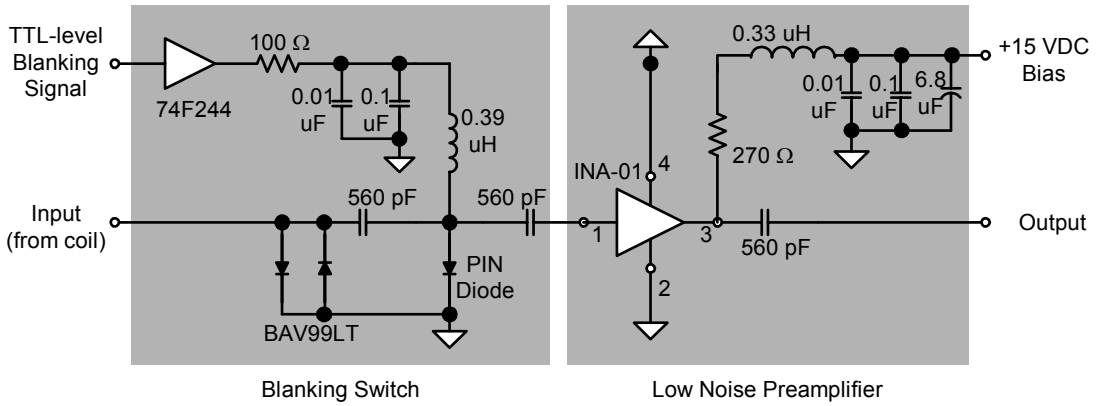

Fig. 42. Low-noise preamplifier circuit. The circuit is made up of a blanking switch and a low-noise preamplifier. When a 5 V (TTL-level) signal is applied to the blanking switch, the diode is turned 'on' forcing the input signal from the coil to be shorted to ground. For normal operation, the diode switch remains ‘off.’ The diode switch provides about 20 dB of isolation for the preamp when in the ‘on’ state. In the ‘off’ state, leakage is much less than 1 dB.

IV.1.2.2.1 Design with MMIC Amplifiers

Each amplifier stage in the RF gain section is based on a standard, single chip, MMIC amplifier (Fig. 43). A wide variety of these single chip MMIC amplifiers (each having different gains, noise figures, and power levels) is available from several different manufacturers. These amplifiers are very easy to use. First, the designer must select an amplifier based upon his requirements for gain, noise figure, power handling, and operating frequency. Next, the designer chooses a dc supply voltage level (V_{cc}). After this, the bias resistor, R_{bias}, may be calculated from (4.1), using the values for the device current, I_d, operating point voltage, V_d, and supply voltage. After determining the bias resistor value, the values for the RF choke inductor, L_{RFC}, and the dc block capacitors, C_{block}, are found from (4.2)-(4.4) using the desired frequency of operation, ω, in radians. Equations (4.3)-(4.4) place constraints on the reactance of the RF choke inductor (X_{RFC}) and the dc block capacitors (X_{Cblock}). The goal behind (4.3) is to make the combined RF impedance of the bias resistor and RF choke high enough at the operating frequency to block the passage RF and thus minimize RF leakage to the dc supply line. Likewise, (4.4) seeks to lower the RF impedance of the dc block capacitor to reduce signal degradation due to insertion loss. Note that (4.3) and (4.4) assume that the nominal impedance of the system is 50 Ω. All the MMIC amplifiers in the RF gain section of the receiver were designed for operation at a nominal frequency of 200 MHz using a 15 V supply voltage (V_{cc}). As such, each amplifier chip uses a 0.33 μH chip inductor for the RF choke and a 560 pF ceramic chip capacitor for the dc blocking components. RF frequency noise is filtered off the bias lines for each of the MMIC

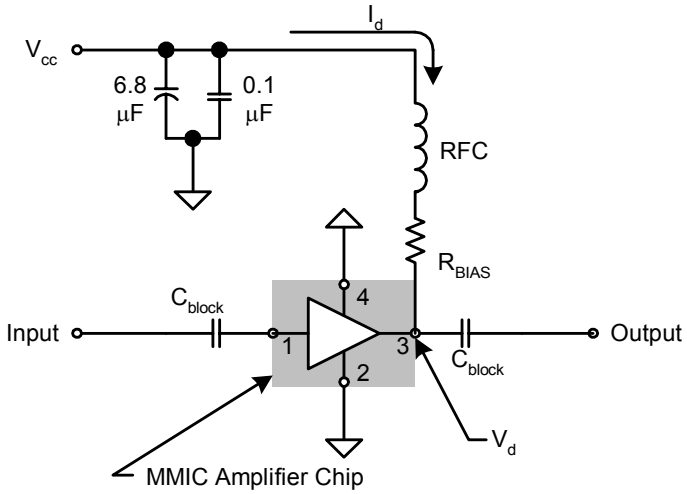

Fig. 43. Standard connection schematic for MMIC amplifier chips. V_{cc} is the dc supply voltage for the amplifier chip. The signal to be amplified is applied to the input terminal and travels through the dc block capacitor (C_{block}) to the amplifier and then to the output terminal. The bias resistor (R_{BIAS}) sets the amplifier bias point voltage (V_d) and operating current (I_d). The bias resistor and the RF choke inductor (*RFC*) act together to keep the RF signal from leaking to the dc supply line. The dc supply line is further filtered by the parallel combination of shunt capacitors to ground which act as a short circuit to RF frequencies.

amplifiers through the use of the RF chokes and a parallel combination of a 6.8 μF tantalum and a 0.1 μF ceramic disc capacitor. These capacitors are connected in shunt to ground.

$$R_{bias} = \frac{V_{cc} - V_d}{I_d} \tag{4.1}$$

$$\left|X_{RFC}\right| = \omega \cdot L_{RFC} \tag{4.2}$$

$$\left|X_{RFC} + R_{bias}\right| >> 50\,\Omega \tag{4.3}$$

$$X_{Cblock} = \frac{1}{\omega \cdot C_{block}} << 50\,\Omega \tag{4.4}$$

IV.1.2.2.2 Initial RF Gain Stages

The first two amplification stages in the RF gain section consist of Mini-Circuits, MAR series, single chip amplifiers. The first gain stage was formed from an MAR-1SM amplifier while the second is composed of an MAR-4SM amplifier. Together, the initial two RF gain stages are able to deliver a total of 26.8 dB of gain with a maximum linear power output of 12.5 dBm. Each of the amplifiers is described in detail below.

The MAR-1SM amplifier provides 18.5 dB of gain from dc to1GHz, an output 1 dB compression point of 1.5 dBm, and a noise figure of 5.5 dB. The chip draws 17 mA of dc current (I_d), and its operating point (V_d) is 5 V. The bias resistor (R_{bias}) is formed from the series connection of a 560 Ω and a 33 Ω, ½ Watt, carbon composition resistor.

The MAR-4SM amplifier has 8.3 dB of gain from dc to 1 GHz, an output 1 dB compression point of 12.5 dBm, and a noise figure of 7.0 dB. The amplifier draws 50

mA of dc current at its operating point of 5.5 V. The bias resistor consists of the series connection of two, 100 Ω, ½ Watt, carbon composition resistors.

IV.1.2.2.3 Variable Attenuation Stage

The attenuation stage follows the first two gain stages and acts as a gain control for the RF gain section of the receiver. The attenuation stage itself consists of the M/A-COM, model AT-309, single-chip, voltage controlled attenuator. The AT-309 provides a variable attenuation from 0 to 20 dB (with a maximum 1.2 dB of insertion loss). The attenuation level is governed by an external control voltage that must be between 0 and $^{-}$4 V. The attenuator operates over a 0 to 2 GHz frequency range. The 1 dB output compression point of the AT-309 varies between 18 and 21 dBm over this operating range, and the chip can handle a maximum RF input power level of 27 dBm. (Note: The 27 dBm maximum input level is well above the maximum allowable power output (12.5 dBm) of the intial two RF gain stages.) Thus, by varying the level of attenuation in this stage, the amount of RF gain may be adjusted by up to 20 dB.

IV.1.2.2.4 Final RF Gain Stages

The final two RF gain stages are identical to the initial RF gain stages. As such, the third gain stage consists of an MAR-1SM amplifier while the fourth gain stage is composed of an MAR-4SM amplifier. This allows the final two gain stages to deliver a total of 26.8 dB of gain with a maximum linear power output of 12.5 dBm.

IV.1.2.2.5 Performance Summary for the RF Gain Subsystem

The total RF gain of the receiver is 53.6 dB (if insertion loss of the attenuator is neglected) and 52.4 dB (if the maximum insertion loss for the attenuator is assumed).

The maximum linear power output of the RF gain section determined by the 1 dB compression point of the third gain stage (i.e., 1.5 dBm). Thus the maximum input signal levels to the RF gain subsystem, with only 1dB compression, are -22.6 dBm (with 20 dB of attenuation) and -42.6 dBm (with no attenuation). (The maximum insertion loss, or 1.2 dB, for the variable attenuation stage has been assumed for this calculation.)

IV.1.2.3 Downconversion

The downconversion section of the receiver is responsible for translating the carrier frequency of the received MR signal to the 500 kHz IF frequency. The downconversion section itself consists of a single, Mini-Circuits, model SBL-1LH mixer. The SBL-1LH mixer requires a 10 dBm local oscillator (LO) signal (this may be at $f_0 \pm 500$ kHz, where f_0 is the MR imaging frequency) and may take up to a 5 dBm RF input signal without saturation. The input 1 dB compression point of the SBL-1LH is 4 dBm. (Note: This input 1 dB compression point translates to an output 1 dB compression point of -2 dBm, assuming the nominal conversion loss of 6 dB.) The mixer has an RF/LO operating range from 2 to 500 MHz, an IF range from dc to 500 MHz, and a conversion loss of approximately 6 dB (8 dB maximum over its entire operating bandwidth). Since the maximum RF input power level to the mixer (with 1 dB compression) is 4 dBm, the maximum allowable power level into the RF to IF conversion stage of the receiver is further limited. Thus the maximum input signal level to the RF gain stage (without compression) becomes -28.4 dBm (with 20 dB of attenuation) and -48.4 dBm (with no attenuation).

IV.1.2.4 IF Gain and Filtering Subsystem

The IF gain and filtering section of the receiver performs four tasks. First, the 500 kHz IF signal is amplified by a variable gain amplifier. Second the received signal is filtered to a 1 MHz bandwidth by a bandpass filter. Then the IF signal is buffered through a high-speed op-amp. Finally, the IF signal is limited to an acceptable voltage level for the digitizer through a diode limiter circuit. The design of each of these four sub-circuits will be described below.

IV.1.2.4.1 Variable Gain Amplifier

The received signal level is scaled a final time for its detection in the digitizer in the variable gain amplifier (VGA). The VGA is a National Semiconductor, CLC5523 chip (which was previously used to perform pulse shaping in the desktop MR system). The amplifier chip is configured to provide gain levels between $^{-}$40 dB and 40 dB. The actual gain level is controlled by an external 0 to 2 V control signal. The amplifier is biased between $^{\pm}$5 V supplies and has a maximum output voltage swing of $^{\pm}$2.5 V. The input impedance of the chip is matched to 50 Ω by a 51.1 Ω chip resistor connected in shunt to ground. The supply voltages for the VGA chip are bypassed (i.e., noise filtered) with 10 μF tantalum and 0.1 μF ceramic chip capacitors at each supply pin.

IV.1.2.4.2 Bandpass Filter

The filtering section performs both anti-alias and dc offset voltage removal functions for the received IF signal prior to digitization. These tasks are accomplished with a bandpass filter. The bandpass filter consists of an active, 1 MHz, lowpass filter chip, the Linear Technology LTC1560, followed by a series connected 3.3 μF ceramic chip

capacitor. The LTC1560 is a fifth order, 1 MHz cutoff, elliptic lowpass filter in a single SO-8 chip package. The filter has 0.3 dB passband ripple, an average 0.14 dB passband gain, and better than 60 dB of stopband attenuation. The supply voltages for the LTC1560 chip ($^{\pm}$5 V) are bypassed with 10 μF tantalum, 0.01 μF ceramic chip, and 0.1 μF ceramic chip capacitors at each supply pin. Together, the series capacitor and lowpass filter form a 1.0 MHz bandwidth, bandpass filter that has a 500 kHz center frequency.

IV.1.2.4.3 Buffer Amplifier

The analog buffer amplifier section is required after the filter stage in order to drive the 500 Ω inputs of the digitizer. The buffer amplifier is a high speed op-amp (Linear Technology LT1806) wired in a unity gain follower configuration. The LT1806 chip is a 325 MHz gain bandwidth, rail-to-rail, low-distortion, low-noise op-amp housed in a SOT-23 package. It is biased with $^{\pm}$5 V and is able to support an output voltage swing of nearly $^{\pm}$5 V. The supply voltage for the chip is bypassed with 10 μF tantalum and 0.1 μF ceramic chip capacitors at each supply pin.

IV.1.2.4.4 Diode Limiter

Finally, the input to the digitizer is limited to $^{\pm}$2.5 V by a diode limiter circuit. This limiter circuit protects the digitizer inputs from damage due to over voltage. The limiter was formed from two crossed, high-speed, silicon, switching diode pairs (ON Semiconductor BAV99LT).

IV.1.2.4.5 IF Gain and Filtering Subsystem Performance Summary

In the IF gain and filtering section of the receiver, the IF signal is amplified by the VGA amplifier in order to produce a full-scale voltage level (i.e., $^{\pm}$1 V or 10 dBm in a 50 Ω system) for the digitizer. Once this voltage level is achieved, the buffer amplifier is then used to drive the 500 Ω input of the digitizer with the output voltage level from the VGA amplifier. This means that the minimum signal power level at the input to the RF gain section of the receiver (to give a full-scale digitizer voltage) is $^{-}$76.4 dBm (this assumes a 10 dBm output at the VGA, a 6 dB conversion loss in the mixer, and 1.2 dB insertion loss in the variable attenuator). The maximum input signal power level allowed into the receiver remains unchanged at $^{-}$28.4 dBm (since it is governed by the 1 dB compression point of the mixer).

IV.1.2.5 Digitization and Processing

In the PC control module, the IF signal output from the RF to IF conversion module is digitized, stored, and then processed. The control computer is a Pentium III, 1.0 GHz PC, with 1.0 GB of SDRAM, running under Windows XP Professional (Microsoft, Redmond, WA). The digitizer is an ICS-645 (Interactive Circuits and Systems, Ltd.) PCI-based card. The ICS-645 card supports up to 32 channels of simultaneous digitization at output rates of up to 2.5 MSPS. The card uses 32 separate ADC's which can be programmed to sample an input signal at rates of up to 20 MSPS (each ADC performs onboard decimation filtering, so the output rate for 32 channels is limited to 2.5 MSPS). For the single channel receiver, only one of these 32 channels is needed, but the ICS-645 is an ideal digitizer for later scaling the single channel receiver to a multiple

channel system. The ICS-645 has 1 Msample of on board memory (2^{20} samples), which means that in 32 channel mode, at most 2^{15} (32,768) samples may be stored per channel (before they must be offloaded to memory). At 2.5 MSPS, this translates to a maximum acquisition window length of 13.107 ms in 32 channel mode. The full-scale input voltage for each ADC channel on the card is $^{\pm}1$ V, with a maximum allowable input signal level of $^{\pm}2.5$ V. After digitization, the collected IF waveforms are stored, I/Q demodulated, decimated, and filtered using in-house software.

IV.1.3 Design Verification

Verification of the initial single channel receiver design took place in three steps. First, the noise figure and dynamic range of the receiver were calculated from the schematic to ensure that the design requirements were met. Next a prototype was constructed from the design schematic. Finally, the prototype was used to acquire spin echo images from a test phantom at 4.7 Tesla to verify its operation.

IV.1.3.1 Noise Figure and Dynamic Range

The noise figure for a cascaded system, such as the receiver, may be calculated from the gain and noise figure of each device in the cascade (Table IV-1 and Fig. 44), as initially discussed in Chapter II, section II.4.1. Using (2.38), the noise figure for the single channel receiver, from the preamp to the input of the digitizer, was found to range from a minimum of 1.705 dB to a maximum of 1.727 dB, during linear-mode receiver operation. The minimum noise figure occurs for input signal levels of less than $^{-}108.4$ dBm. For signal levels this low, the variable attenuator is turned to its minimum

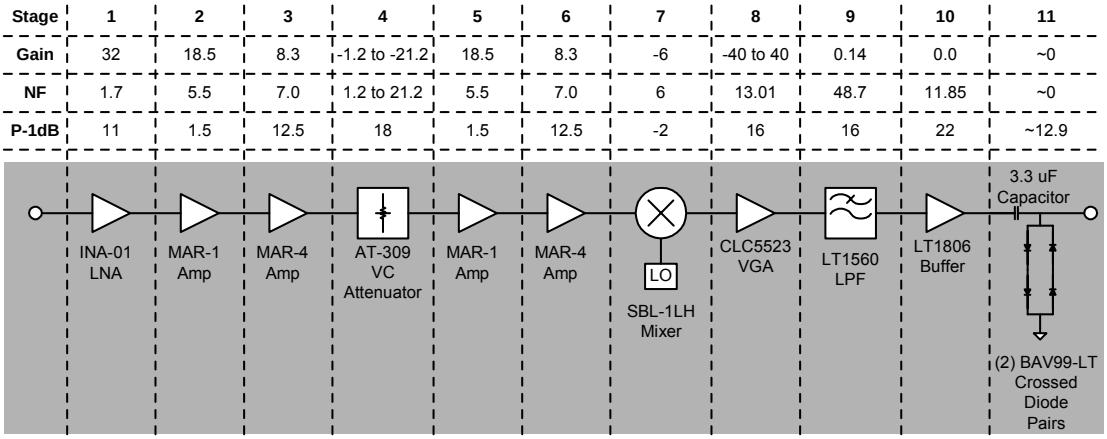

Fig. 44. Gain, noise figure, and output 1 dB compression points for each stage of the single-channel receiver prototype prior to the digitizer. The noise figure and gain values are given in dB, while the 1 dB compression points are given in dBm.

attenuation level of 1.2 dB, and the variable gain IF amplifier is set to its maximum gain of 40 dB. The maximum noise figure occurs when the maximum signal level (for linear mode receiver operation) of -60.4 dB is input to the receiver. In this case, the variable attenuator is set to its maximum level (21.2 dB) while the gain of the VGA amplifier is set to 12 dB in order to obtain a full-scale input voltage at the digitizer (i.e., 2 V_{pp} or 10 dBm). Since the maximum noise figure for the prototype receiver is less than 2.0 dB, the prototype receiver meets its noise figure requirement (refer to section IV.1.1.3).

Table IV
Gain, Noise Figure, and Output 1 dB Compression Point Values for the Single-Channel Receiver Components

Device	Gain [dB]	Noise Figure [dB]	1 dB Compression Point [dBm]
INA-01 Amplifier	32	1.7	11
MAR-1SM Amplifier	18.5	5.5	1.5
MAR-4SM Amplifier	8.3	7.0	12.5
AT-309 Attenuator	-1.2 to -21.2	1.2 to 21.2	18
SBL-1LH Mixer	-6	6	-2
CLC5523 VGA	-40 to 40	13.01	16
LTC1560 Low-pass Filter	0.14	48.7	16
LT1806 Buffer	0.0	11.85	22

The dynamic range of a receiver, as discussed in Chapter II, section II.4.2, is merely the power difference in dB between the largest and smallest signals detectable by that receiver. The largest detectable signal can be determined from the 1 dB compression points of the individual components in the receiver chain. In particular, the largest detectable signal is governed by the components in the receiver chain which have the lowest compression points. For the prototype, the maximum detectable signal is

determined by the mixer, which has an input 1 dB compression point of 4 dBm. Using this fact and the gains for the receiver components from Table IV, the maximum detectable signal level for the receiver was found to be -60.4 dBm. The minimum detectable signal level (see Chapter II, section II.4.2.2) is given by the sum of the thermal noise in a 1 MHz bandwidth (-114 dBm), the maximum noise figure of the receiver (1.727 dB), and a detection threshold. (Recall from section IV.1.1.1 that the bandwidth requirement for the single-channel receiver prototype was set at 1 MHz.) The detection threshold has been set to 1 dB for this receiver. The minimum detectable signal for the prototype receiver becomes -111.27 dBm. Thus, the dynamic range of the prototype receiver, not including the digitizer, is 50.87 dB, which just meets the minimum dynamic range specification for the receiver (i.e., 50 dB).

It should be noted that the dynamic range of the single-channel prototype may be increased by the addition of a fixed RF attenuator between the low-noise preamp and the first stage of the RF to IF conversion section. This will affect the noise figure of the receiver system to some extent. However, by using small levels of attenuation, the dynamic range may be increased by up to 10 dB without violating the system noise figure requirement. For example, by using a 7 dB fixed attenuator, the receiver dynamic range may be increased to 57.74 dB while only raising the maximum noise figure to 1.856 dB. Likewise, a 10 dB fixed attenuator may be used to increase the dynamic range to 60.6 dB while raising the system noise figure to its maximum acceptable level of 2 dB. These techniques may be used if it becomes necessary to collect fid signals from the

3 mm × 3mm × 100mm test volume used to determine the maximum signal level estimates in section IV.1.1.4.2.

Note that the above discussion of receiver dynamic range does not include the digitizer. The digitizer used with the prototype is the ICS-645 card which makes use of oversampling during signal detection. The ICS-645 card oversamples the 1 MHz receive bandwidth by a factor of 8. This oversampling ratio can translate to up to a 9 dB increase in the dynamic range of the receiver due to oversampling gain. (As discussed in Chapter II section II.4.2.3.3, oversampling gain (in dB) can be as high as $10 \cdot \log_{10}(\textit{Oversampling Ratio})$ (2.47).) In cases where the input MR signal bandwidth is less than 1 MHz and narrower bandwidth post acquisition digital filters are used to process the received signal, even larger oversampling ratios/ gains are possible. Thus dynamic range for the receiver prototype should always exceed the 50.75 dB calculated above. (Note that in the MR receiver, oversampling gain is not usually as high as (2.47) predicts since the input noise is usually amplified to a level that is on the order of several LSB's of the digitizer. There will, however, be an increase in the overall SNR and dynamic range of the MR receiver due to oversampling gain.)

IV.1.3.2 Single-Channel Prototype

The prototype single-channel receiver (Fig. 45) was constructed from the design schematic. The prototype consisted of a single preamplifier module (Fig. 45-b), an RF to IF conversion module (Fig. 45-a), and an ICS-645 digitizer housed in a control computer. Power for the preamp and RF to IF module was provided by four, 9 Volt batteries. The preamp was housed in a 5 cm (2 in.) × 10 cm (4 in.) × 15 cm (6 in.)

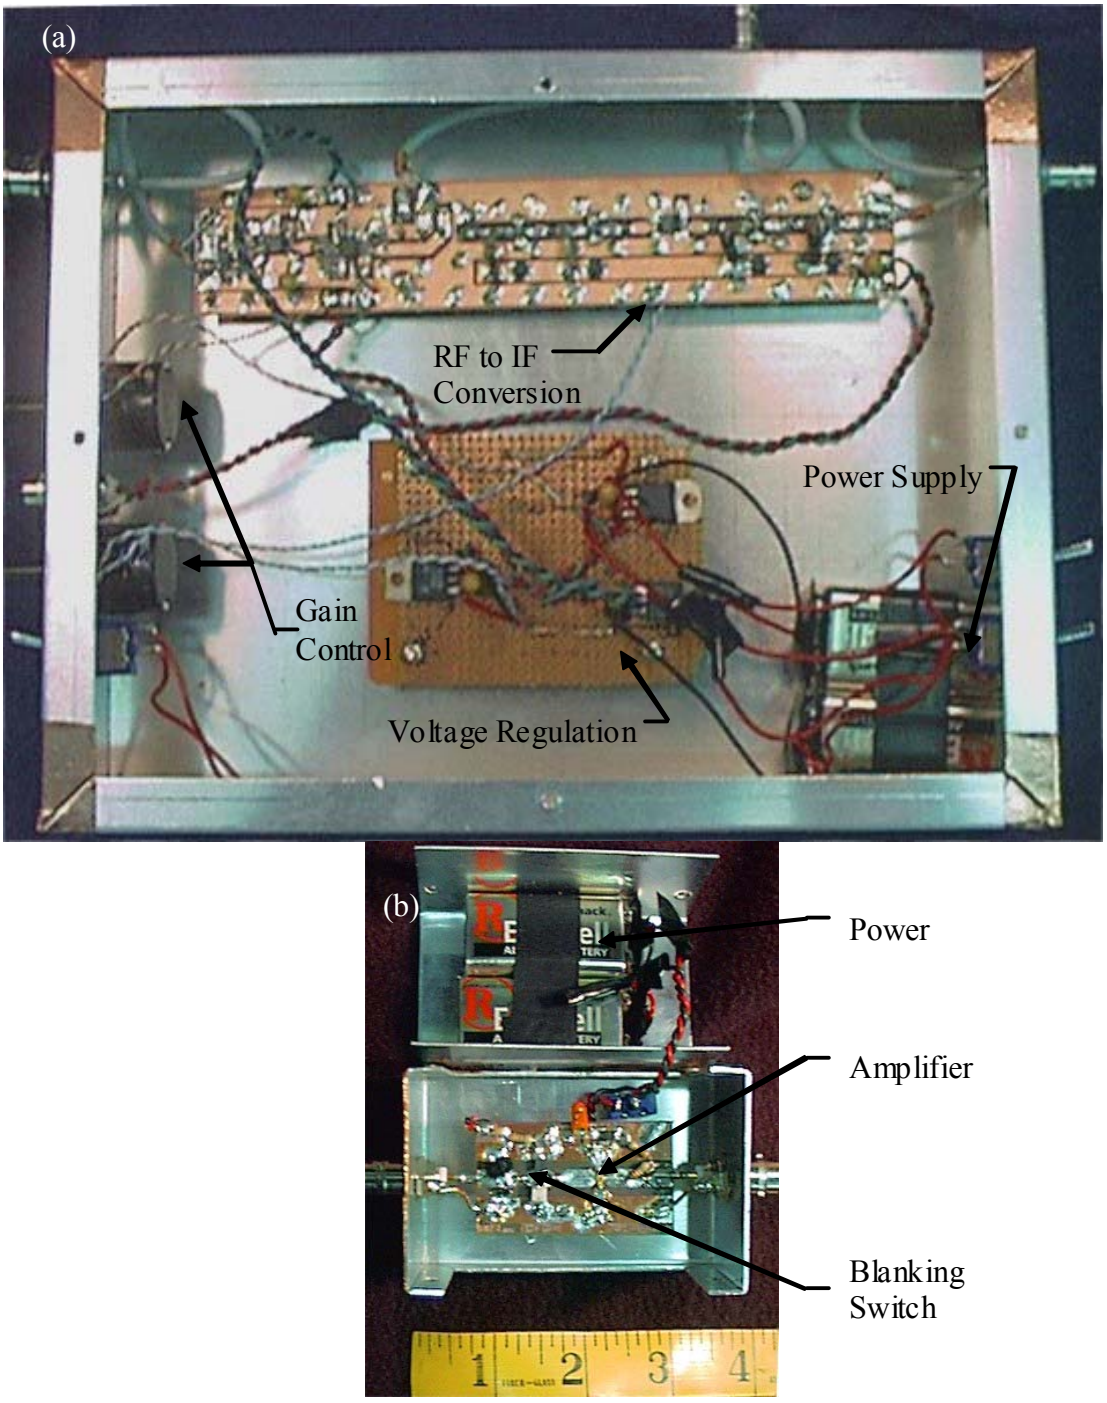

Fig. 45. Single channel receiver prototype. (a) The RF to IF conversion module. The power supply is composed of three 9 V batteries. The voltage regulation section provides 5 V, -5 V, and 15 V to power the receiver. Gain control is performed by two, 10-turn 100 kΩ, precision potentiometers which supply the gain control voltages to the VGA amplifier and variable attenuator chips. (b) Low-noise preamp module. Measuring tape shown is in inches.

aluminum box while the RF to IF conversion module was placed in a 5 cm (2 in.) × 18 cm (7 in.) × 23 cm (9 in.) aluminum box. Connections between boxes were made by shielded, 50 Ω, coaxial cables using BNC connectors. The local oscillator signal was provided by a PTS-250 (PTS, Littleton, MA) frequency synthesizer.

IV.1.3.3 Image Testing

As a final step in the verification of the single channel receiver design, the prototype receiver was used to acquire MR images. These images were acquired on a GE Omega 4.7 T (33 cm bore) magnet system to show that the receiver would function in a real imaging environment. A 10 cm (4 in.) inner diameter birdcage coil was used to both transmit the RF excitation and receive the MR signal from a test phantom. The test phantom consisted of a 2.2 cm (0.855 in.) inner diameter, 9.5 cm (3.75 in.) long cylinder with three, 1 mm (0.04 in.) inner diameter, 75 mm (2.95 in.) long capillary tubes spaced ~120º apart along its outer diameter (Fig. 46). Both the cylinder and capillary tubes were filled with 0.5 mmol $CuSO_4$(aq) solution. Initially, spin echoes alone (center line of k-space, TE/ TR = 28 ms/ 300 ms, 20 kHz spectral width, 75 mm FOV, 2 mm slice thickness, 256 point, 12.8 ms acquisition) were acquired (Fig. 47) to allow adjustment of the receiver gain stages to prevent saturation. After receiver adjustment, standard spin echo images (TE/TR = 28 ms/ 300 ms, 20 kHz spectral width, 75 mm FOV, 2 mm slice thickness, 256 × 128 point acquisition) were acquired with the prototype receiver (Fig. 48-a). Comparison images were then taken with the GE Omega receiver using the same pulse sequence (Fig. 48-b). The image from the prototype receiver has a pronounced phase jitter artifact and a significantly lower image SNR than the comparison image.

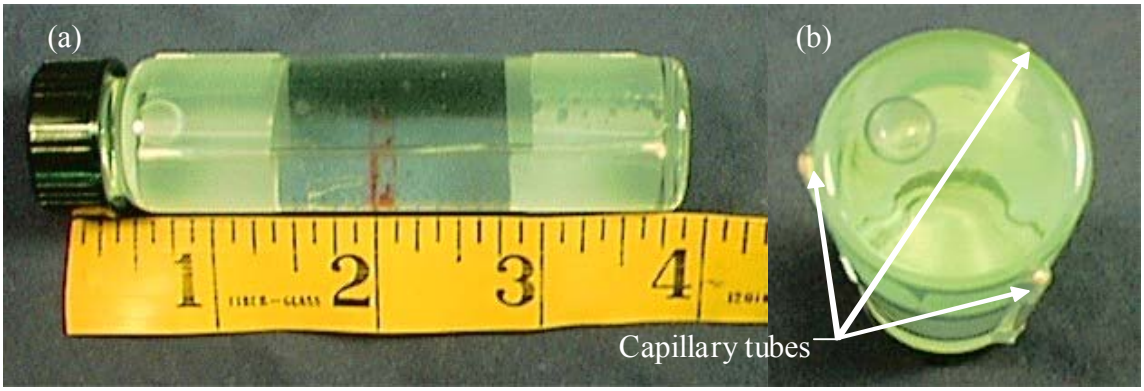

Fig. 46. Test phantom used with the prototype receiver to acquire spin echo images. (a) Side view of phantom. (b) End view of test phantom. Note the capillary tubes that have been taped to the outside wall of the phantom.

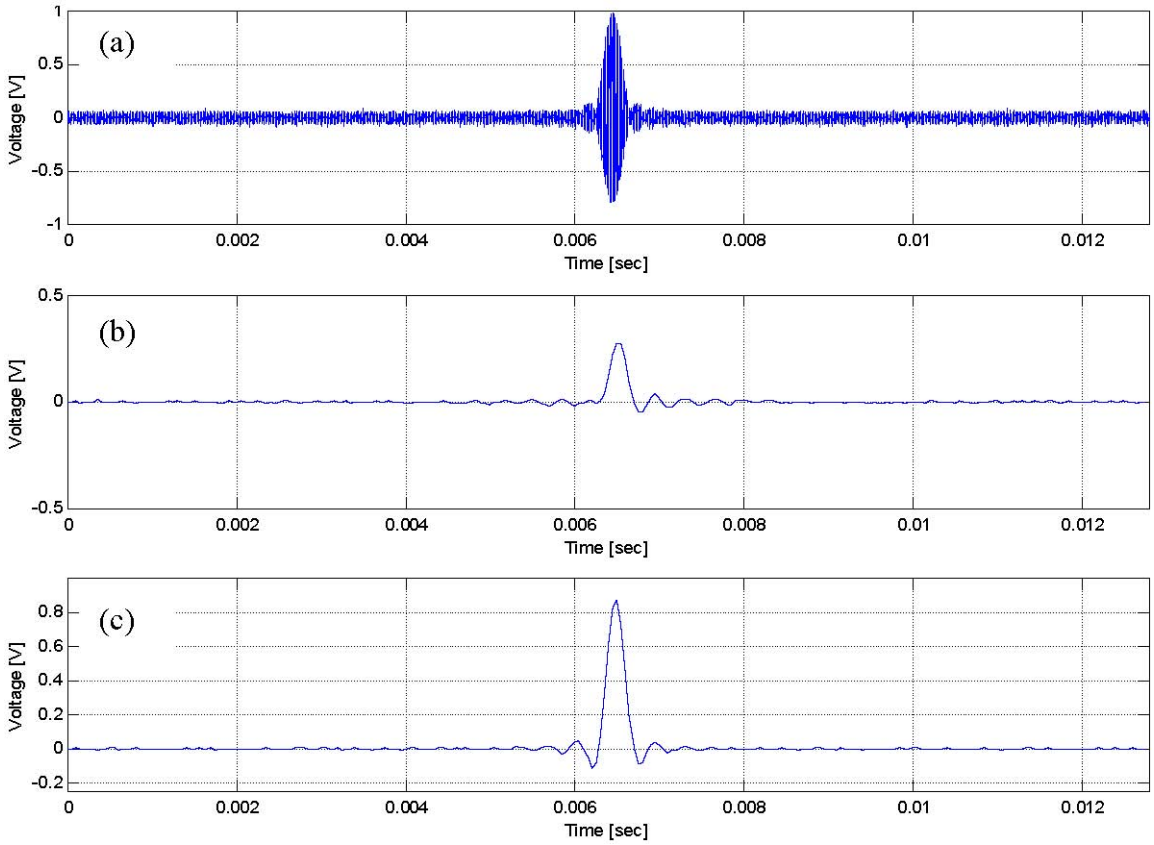

Fig. 47. A standard spin echo (center line of k-space, TE/ TR = 28 ms/ 300 ms, 20 kHz spectral width, 75 mm FOV, 256 point, 12.8 ms acquisition) received with the prototype receiver on a GE Omega 4.7 T/ 33 cm bore magnet system. (a) IF frequency signal. (b) In-phase part of the demodulated echo. (c) In-phase quadrature portion of the demodulated echo.

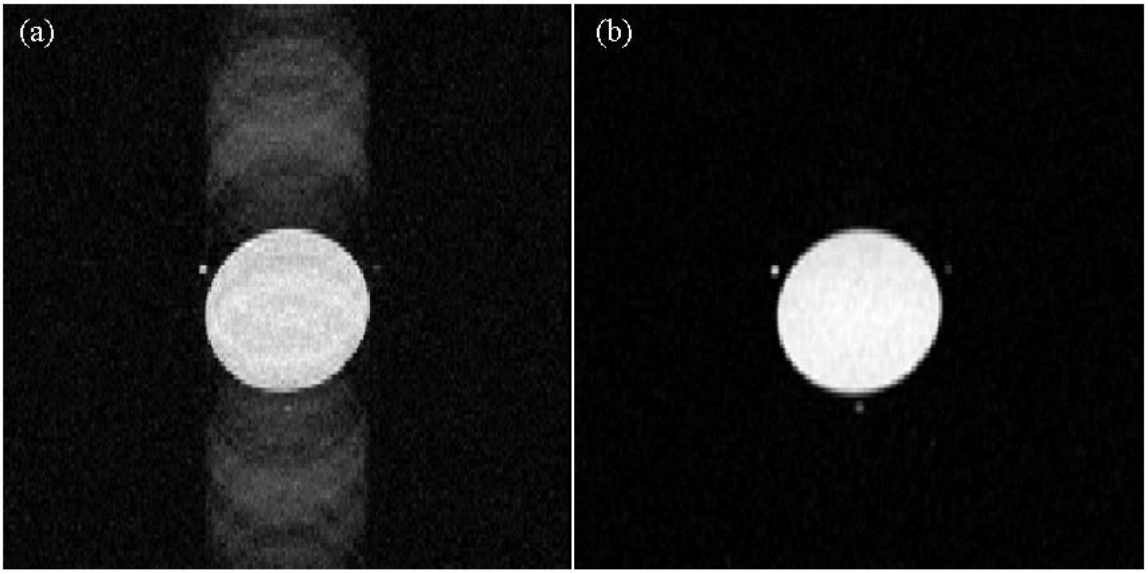

Fig. 48. Standard spin echo images (TE/ TR = 28 ms/ 300 ms, 20 kHz spectral width, 75 mm FOV, 2 mm slice thickness, 256 × 128 point acquisition) received with the prototype receiver and the GE Omega receiver on a 4.7 T/ 33 cm bore magnet system. (a) Image from the prototype receiver. SNR = 19.1. The phase jitter artifact has since been corrected (see Chapter VII for details). (b) Comparison image from the GE Omega system receiver. SNR=50.8.

The phase jitter artifact is due to a lack of synchronization between the 10 MHz master clock of the Omega system and the 40 MHz clock on the digitizer card (see Chapter VII). The low SNR in the prototype receiver image is likewise due to the phase jitter spreading the signal energy over the image. Synchronization of the digitizer clock to the master clock on the Omega system has since corrected this phase jitter problem (refer to Chapter VII for a complete discussion of the synchronization process). As a result, the phase jitter artifacts in the prototype receiver images have disappeared (see Chapter VII, section VII.4.2) and the image SNR has become comparable to that of the Omega system.

IV.2 Summary

A cost-effective, single channel, digital MR receiver was developed for imaging small volume samples at 4.7 T. The design process for the prototype has been discussed in detail. This involved the derivation of system requirements for the receiver, the development of a design for the single channel prototype, and verification of the prototype performance. Except for the problem of phase jitter in the initial test images (which has since been remedied), the performance of the prototype met all of the design specifications. Finally, this single channel design may be easily scaled to form a sixty-four channel receiver system.

CHAPTER V

DESIGN OF THE 64-CHANNEL PARALLEL MR RECEIVER

In phase two of the parallel receiver project, a prototype receiver for parallel imaging was developed based upon the single-channel receiver described in Chapter IV. This new receiver was designed to have a total of sixty-four channels, allowing the parallel imaging receiver to simultaneously receive MR signals from up to sixty-four element arrays of receive coils. Each channel in the parallel imaging receiver was designed to meet all the criteria developed for the single-channel prototype (see Chapter IV, section IV.1.1) and to easily interface to any commercial MR scanner. This chapter will present the overall design for the 64-channel parallel imaging MR receiver.

V.1 64-Channel Parallel Receiver System Requirements

The design specifications for each channel of the 64-channel parallel imaging receiver are the same as those initially developed for the single-channel receiver prototype (refer to Chapter IV, section IV.1.1). The specifications derived for the single-channel receiver included a channel bandwidth of at least 1 MHz, an operating frequency range from 60 to 210 MHz, a noise figure of less than 2 dB, and a dynamic range of at least 50 dB. An additional requirement on the design of the parallel receiver is that it must be able to easily interface with any commercial MR scanner. These requirements alone, however, do not completely specify the design requirements for a parallel receiver.

In designing a receiver for parallel imaging, two additional factors are of concern: the isolation between receiver channels and the channel-to-channel gain variation. First,

the isolation between channels in the parallel receiver must be at least 35 dB to prevent artifacts in the received images due to channel-to-channel bleed through [16]. The second factor is harder to quantify. In an ideal multi-channel receiver, the performance of each receiver channel, particularly the channel gain, would be identical to the performance of all the other channels. Thus, if the same RF signal were received by each channel in the receiver, the resulting baseband signal outputs from each channel would all have the same amplitude. This is because, in the case of the ideal receiver, the circuits for all the receiver channels would behave identically. However, the parallel receiver must be constructed from real-world components. This means that the gain of the channels in a real-world receiver will vary from channel to channel. For parallel imaging, the channel-to-channel gain variation should be as small as possible. Without prior knowledge of how much gain variation might creep into the 64-channel receiver due to manufacturing techniques and other factors, it was decided to not place fixed minimum specification on the allowable level of channel-to-channel gain variation.

V.2 64-Channel Parallel Receiver Design

A 64-channel receiver was designed to meet the specifications on the bandwidth, operating frequency range, noise figure, dynamic range, MR scanner interface, interchannel isolation, and channel-to-channel gain variation that were discussed in section V.1. The parallel receiver was created by scaling the design for the single-channel receiver prototype. This involved multiplying the design for the prototype by a factor of sixty-four and developing several new support subsystems for the distribution of power and control signals.

The 64-channel system, like the single-channel prototype, consists of three main parts: a preamplifier module, a main receiver module, and an acquisition control computer (Fig. 49). Additionally, the receiver is able to interface with any commercial MR scanner using only three control signals. The preamplifier module, which is housed in a magnet leg, amplifies the signals from up to sixty-four radio-frequency (RF) coils. The main receiver module then downconverts the MR signals from the magnet frequency (RF) to an intermediate frequency (IF) of 500 kHz. Both the preamp and main receiver modules were constructed in-house from low-cost, off-the-shelf components. In the control computer, the IF signals are digitized and stored for later reconstruction into an MR image. The main receiver module and control computer are housed within an 89 cm (35 in.) tall, 48 cm (19 in.) shock-mount rack enclosure for portability (Fig. 50). The design and operation of the MR scanner interface, preamp module, main receiver, and control computer will be described in the following sections.

V.2.1 MR Scanner Interface

The 64 channel receiver requires only three signals to interface with any commercial MR scanner: a 10 MHz frequency reference, an acquisition trigger signal, and an RF blanking signal. The 10 MHz frequency reference signal is used to phase lock both the local oscillator (LO) and clock synthesizers to the master clock of the MR scanner. The acquisition trigger signal is either a high- or low-true TTL-level signal which is asserted at the start of the MR scanner's digitization window for each acquired k-space line.

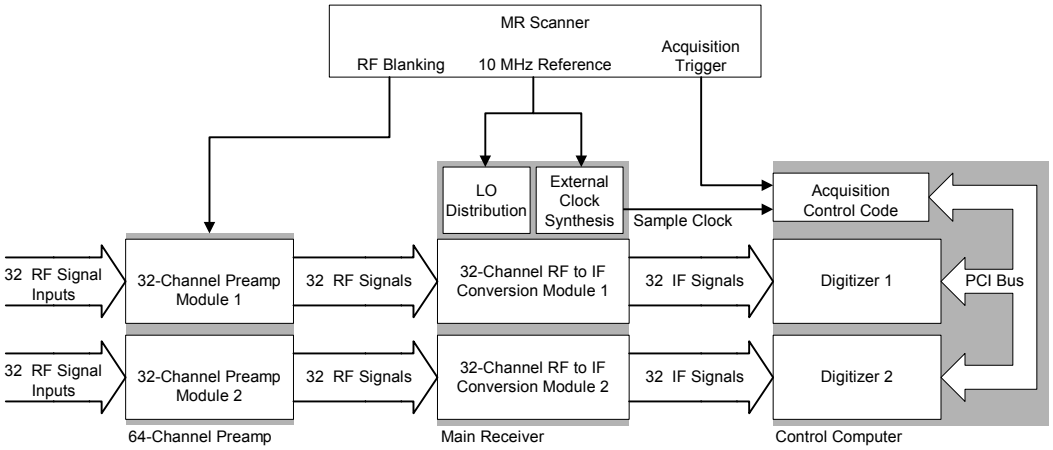

Fig. 49. Layout of the 64-channel receiver system. The receiver is composed of three main modules: the preamplifier, the main receiver, and the acquisition control computer.

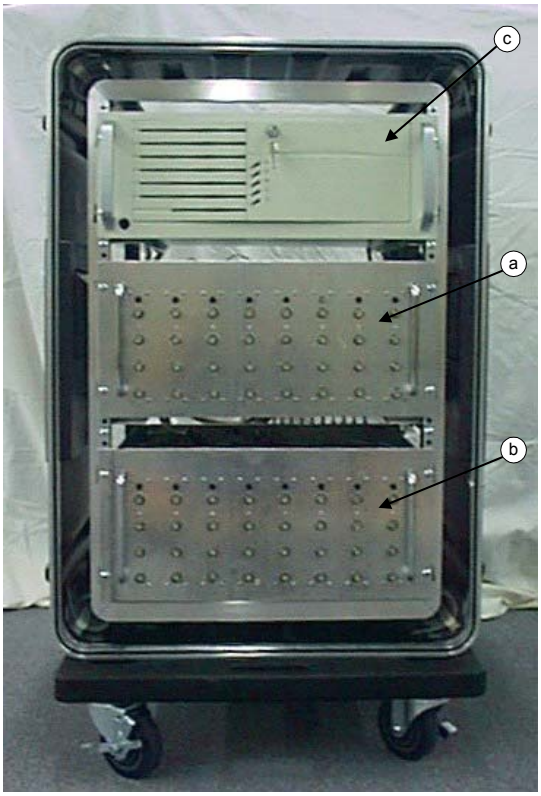

Fig. 50. The prototype 64-channel receiver system mounted in a 35-in. tall, 19-in. shock-mount rack. (a)-(b) Two, 32-channel RF to IF conversion modules. (c) Acquisition control computer.

Finally, the RF blanking signal is a TTL-level signal which is asserted during each RF transmit pulse. The blanking signal is used to protect the preamp modules from RF transmit pulse breakthrough.

V.2.2 64-Channel Preamplifier

The 64-channel preamplifier amplifies the MR signals from the receive coil array and establishes the overall noise figure for the receiver system (see section II.4.1 of Chapter II). The 64-channel preamp module consists of two, 32-channel, low-noise preamp modules. The design for each channel is based upon the preamplifier circuit developed for the single-channel receiver. Each channel in the preamplifier delivers 32 dB of gain with a 1.7 dB noise figure over the frequency range of dc to 500 MHz. In order to reduce the overall cost of the preamp module, inexpensive, single chip amplifiers were used to form the sixty-four individual preamp circuits. However, since each amplifier chip was designed for operation in a system with a characteristic impedance of 50 Ω, the preamp circuits cannot be used for decoupling the array coils [33]. Each amplifier circuit is protected from RF pulse breakthrough during the transmit portion of the MR imaging sequence by its own PIN diode switch. These switches turn on during the application of the TTL-level RF blanking pulse from the MR scanner (i.e., one of the three interface signals), preventing any transmit pulse breakthrough from reaching the preamp circuits.

V.2.3 Main Receiver

The primary function of the main receiver module is to perform downconversion and signal conditioning on the MR signals received from the 64-channel preamplifier. In

addition, the main receiver houses the support systems which provide power and control signals for all the channels in the receiver. The main receiver itself (Fig. 49) is comprised of a local oscillator (LO) distribution system, an external clock synthesis module, and two, 32-channel, RF to IF conversion modules. The two, 32-channel conversion modules downconvert the preamplifier signal outputs from RF to IF frequency signals. The LO distribution system provides a mixing signal to each RF to IF conversion channel within the main receiver. Finally, the external clock synthesis module generates a sampling clock for the digitizers located in the control computer.

V.2.3.1 Local Oscillator Distribution System

The local oscillator (LO) distribution system provides a mixing signal to all the RF to IF conversion channels within the 64-channel receiver. As such, the LO distribution system first generates a 10 dBm (in a 50 Ω system) local oscillator signal that is phase-locked to the 10 MHz master system clock on the MR scanner. This signal is then amplified and split to form sixty-four, 10 dBm copies of itself through a network of low-cost RF splitters. The sixty-four local oscillator signals are then delivered to each of the channels in the receiver.

V.2.3.2 External Clock Synthesis Module

The clock synthesis module generates an external sample clock for the digitizer cards used in the 64-channel receiver system. The generated clock signal is a 40 MHz, TTL-level square wave. The output signal produced by the clock module is phase synchronized to the 10 MHz master system clock supplied by the host MR scanner (one of the three scanner interface signals) by means of an on-board PLL (phase-locked loop)

circuit. The generated sampling clock is thus synchronous to the main system clock on the MR scanner. This clock synchronization eliminates the problem of phase jitter in the final MR image due to asynchronous sampling (refer to Chapter VII).

V.2.3.3 32-Channel RF to IF Conversion Modules

The two, 32-channel RF to IF conversion modules are responsible for downconversion and conditioning of the sixty-four MR signals from the preamplifier module. Each 32-channel module fits into a standard 19-in. rack chassis (Fig. 51). The 32-channel, RF to IF modules themselves are each comprised of eight, 4-channel, RF to IF conversion submodules (Fig. 52). Each channel in these conversion submodules is identical to the RF to IF conversion circuit first designed for the single-channel receiver prototype (see Chapter IV, section IV.1.2 for details). As in the prototype, each RF to IF conversion channel in a submodule first amplifies, then downconverts, and finally filters the received MR signal from the preamp before it travels on to the digitizer.

The 32-channel modules also house the power distribution and the gain control systems. The power distribution system supplies all the modules in the 64-channel receiver (except for the control computer) with dc power. The gain control subsystem on the 64-channel receiver allows the user to adjust the gain of the receiver channels.

V.2.4 Acquisition Control Computer

The control computer is a rack-mount, 1.0 GHz Pentium III machine. This computer performs three main functions for the receiver. First, it handles the digitization of the IF signals outputs from the main receiver module. Second, the control computer stores the

Fig. 51. A 19-in. chassis containing one of the two, 32-channel RF to IF conversion modules.

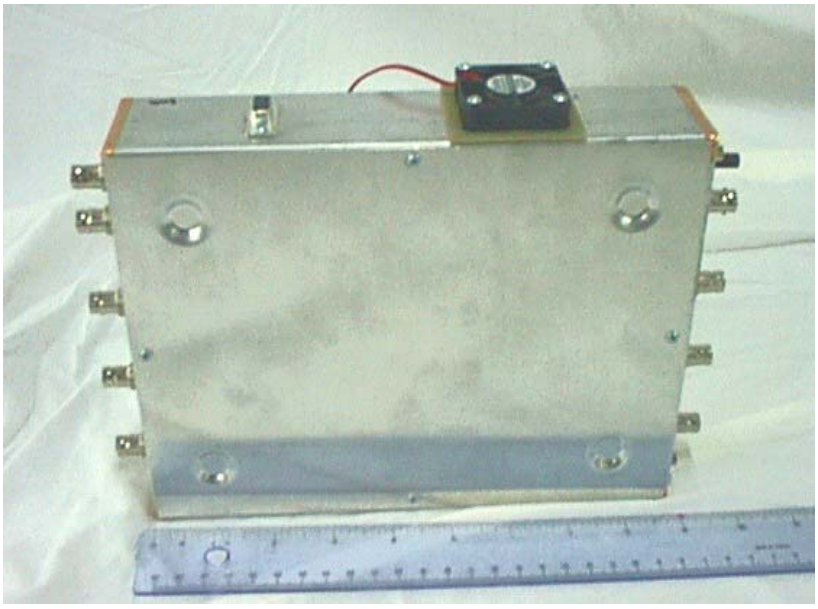

Fig. 52. One of the eight, 4-channel RF to IF conversion submodules used in each 32-channel RF to IF conversion module.

received data in memory for later processing using the acquisition control code. Finally, the acquisition computer performs the final signal processing necessary to reconstruct an image from the received data.

V.2.4.1 64-Channel Digitizer

Digitization of the IF signal data is performed using the 64-channel digitizer system. The 64-channel digitizer itself is based on the ICS-645, 32-channel digitizer card (Interactive Circuits and Systems, Ottawa, Canada). (One of these ICS-645 cards was used as the digitizer in the single-channel prototype.) The ICS-645 cards each support simultaneous, 16-bit digitization of up to thirty-two, 1.25 MHz bandwidth channels. Two cards reside in the control computer, enabling detection of up to sixty-four IF signals. Each card is able to digitize and store up to 13.1 ms of data (in 32-channel mode) before its on-board buffer is full. Once the on-board buffer fills, the digitized signal data must be transferred to the host computer memory over the system PCI bus. To facilitate this, each ICS-645 card supports sustained DMA transfer rates between 80 and 100 MB/sec. The two ICS-645 cards are synchronized to the host MR scanner by means of a 40 MHz sample clock (provided by the external clock synthesis module) and an acquisition trigger signal (from the MR scanner). The ICS-645 cards also come from the manufacturer equipped with a simple LabVIEW demonstration utility. The demonstration software allows the user to view, in real-time, the signals being digitized on any channel of the ICS-645 cards. This software provided a simple diagnostic tool for testing the operation of the digitizers.

V.2.4.2 Acquisition Control Code

The acquisition control code provides control for both the 64-channel digitizer system and the data storage process. For data storage, the acquisition control computer has 1.0 GB of SDRAM and a 60 GB hard disk drive. During MR image acquisition, the 64-channel digitizer acquires signal data from the receiver channels after each acquisition trigger signal is received from the main MR scanner. During the recycle time between successive triggers, the digitizers transfer the received signal data off-card to the host computer. To assist the speed of this data transfer, the signal data is first transferred over the PCI bus to a virtual buffer in SDRAM. After a complete image data set is acquired, the entire signal data set is finally written to files on the hard disk.

V.2.4.3 Signal Processing

In order to test the prototype receiver, software was developed in MATLAB (The MathWorks, Natick, MA) to perform the final processing of the digitized IF signal. The software first performs digital I/Q demodulation, filtering, and decimation on the IF signals to yield the baseband k-space data (Fig. 53). After digitization, each 16-bit IF signal is multiplied by both a 32-bit, digital sine and cosine wave (each having a frequency of 500 kHz) to demodulate the signal to baseband. This multiplication produces the in-phase (I channel) and in-phase quadrature (Q channel) signals. These 32-bit I and Q channel signals are then filtered to the spectral width of the imaging sequence by a 30-point FIR (finite impulse response) low-pass filter and decimated to the appropriate number of readout points based upon the spectral width of the imaging sequence. The I channel becomes the real part of the k-space data set while the Q

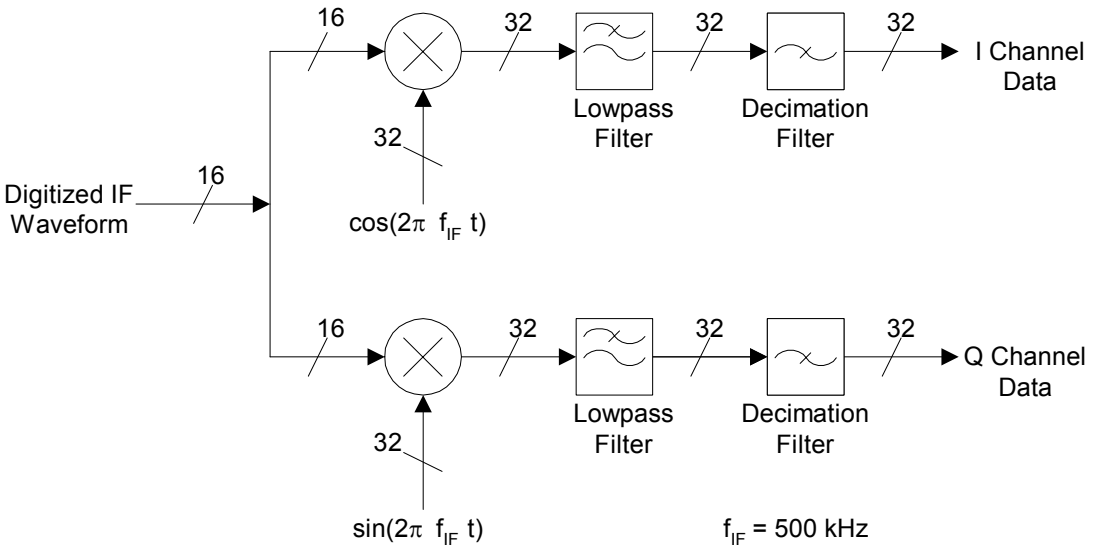

Fig. 53. A block diagram of the digital demodulation process for the digitized IF signal data. The 16-bit samples of the IF waveforms are I/Q demodulated, filtered, and decimated to yield the real and imaginary parts of the k-space data.

channel becomes the imaginary part. The k-space data is then reconstructed by Fourier Transform methods to form the final MR image.

V.3 Summary

The design for the single-channel receiver prototype was scaled to create a new receiver for parallel imaging. This receiver is capable of the simultaneous reception of MR signals from up to sixty-four receive coils. In addition, the parallel imaging receiver easily interfaces with any commercial MR scanner using only three control signals. A detailed summary of the design for the parallel imaging receiver has been presented.

CHAPTER VI

PERFORMANCE CHARACTERIZATION OF THE 64-CHANNEL RECEIVER

In the third and final phase of the parallel receiver project, the performance of the prototype parallel MR imaging receiver was fully characterized. The 64-channel receiver was designed to meet the requirements developed initially for the single-channel prototype receiver (Chapter IV, section IV.1.1) as well as the additional specifications discussed in Chapter V section V.1. These specifications included a channel bandwidth of 1 MHz, an operating frequency range from 60 to 210 MHz, a noise figure of less than 2 dB, a dynamic range of at least 50 dB, interchannel isolation of at least 35 dB, a minimum of channel-to-channel gain variation, and low system cost per channel. In order to establish how well the 64-channel prototype met its own design requirements, the bandwidth, operating frequency range, noise figure, dynamic range, interchannel isolation, gain variation between channels, and overall system cost have all been determined. The methodology and results of these characterizations will be presented in this chapter.

VI.1 Design Verification

The design specifications developed for the parallel receiver have all been discussed previously. This section will verify the performance of the 64-channel receiver design as compared to the specifications for bandwidth, operating frequency range, noise figure, dynamic range, interchannel isolation, gain variation between channels, and system cost. For each specification, the receiver was subjected to a test procedure to evaluate its

performance. These test procedures and the resulting receiver performance will be presented in the following sections.

VI.1.1 Channel Bandwidth

The bandwidth of the channels in the 64-channel receiver was measured to verify that the design specification of at least 1 MHz had been met. Neither the preamplifiers nor the digitizers were included in this measurement since the RF to IF conversion circuits are where bandwidth limiting occurs in each receiver channel. The preamplifier circuits are inherently wideband (each amplifier has a 3 dB operating bandwidth of 500 MHz) and the bandwidth of each digitizer channel is already known to be 1.25 MHz (as tested by the manufacturer). This leaves only the bandwidth of the RF to IF conversion circuits in question.

A test setup was devised for measuring the channel bandwidth (Fig. 54). An HP4195A network/ spectrum analyzer (Agilent, Palo Alto, CA) and a 50B-002 step attenuator (JFW Industries, Indianapolis, IN) were used to generate both a 10 MHz frequency reference signal and an RF test signal. For the test, an external, PTS-250 frequency synthesizer (Programmed Test Sources, Littleton, MA) generated a 199.9 MHz LO source signal for the LO distribution system in the 64-channel receiver (replacing the master LO source that is used for normal receiver operation). The outputs of the HP4195A and the PTS-250 were phase-locked together by the 10 MHz frequency reference. The RF test signal was a -70 dBm sinusoid that was injected separately into each of the sixty-four, RF inputs on the RF to IF conversion modules of the 64-channel receiver's main receiver module. (The dashed lines in Fig. 54 indicate the signal paths

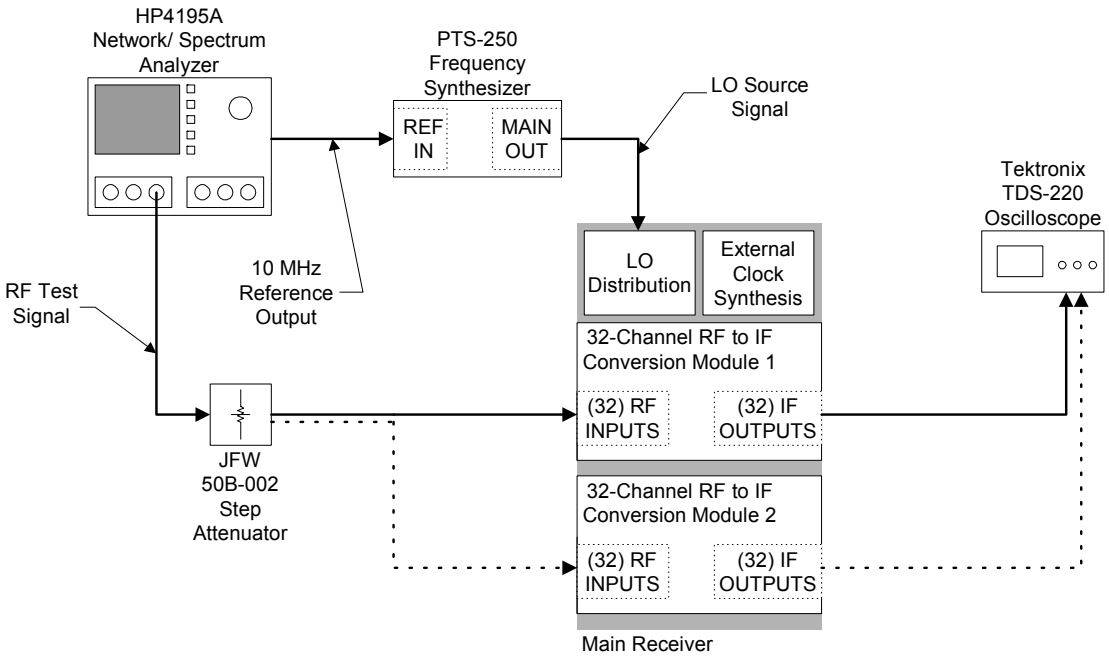

Fig. 54. Diagram of the apparatus used to test the bandwidth of the channels in the 64-channel receiver.

used to measure the bandwidth on the second 32-channel RF to IF conversion module). For each channel, the frequency of the RF test signal from the HP4195A was varied incrementally from 199.9 MHz to 201.175 MHz (a sweep range of 1.275 MHz). At each frequency increment, the voltage level of the RF to IF conversion channel's IF output response was measured using the TDS-220 (Tektronix, Beaverton, OR) digital oscilloscope and recorded. The IF output response voltages of the channels in the 64-channel receiver were then converted into dBm (assuming a 50 Ω system) and plotted.

The typical response curve for one of the sixty-four RF to IF conversion circuits is plotted in Fig. 55. A horizontal line depicting the ¯3 dB level has also been added. It is apparent from Fig. 55 that the 3 dB bandwidth of the receiver channels meets the design specification for a minimum of 1 MHz of channel bandwidth.

VI.1.2 Operating Frequency Range

The operating frequency range of the parallel receiver was not established by direct measurement. Instead, the extent of the operating range was first determined from the operating specifications for the parts used in the receiver. Then, the operation of the receiver was verified at a number of sample frequencies within the receiver's desired operating range. These procedures are explained in this section.

To find the operating frequency range of the receiver, the manufacturer's specifications for the parts used in the 64-channel system were examined. Since the design of each receiver channel is identical to that first developed for the single-channel receiver prototype, the operating frequency range for the sixty-four channel circuits in the parallel receiver should not differ from that of the single-channel prototype. In

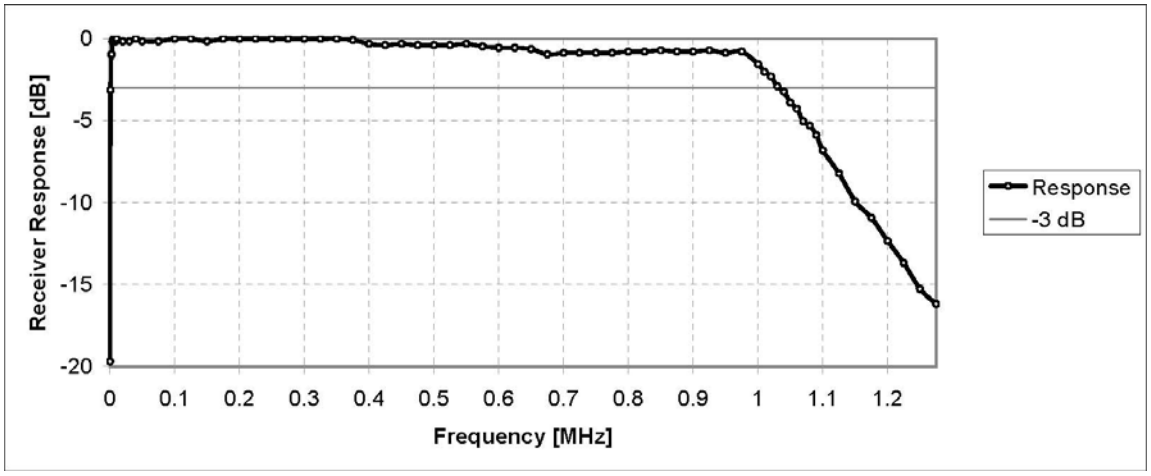

Fig. 55. Plot of typical bandwidth response of the channels in the 64-channel receiver.

Chapter IV, section IV.1.2, the operating frequency range for the single-channel receiver was found to extend from 2 to 500 MHz. In developing the parallel receiver, additional systems were created which further limited the receiver's operating range. Among these is the LO distribution system which has two components that limit the operating range of the parallel receiver. First is the MHL8115 (Motorola, Schaumberg, IL), a 1 W, RF amplifier that operates from 50 MHz to 1 GHz. Second is the SCP-4-1 (Mini-Circuits, Brooklyn, NY), a 4-way RF splitter which functions over the frequencies from 1 to 400 MHz. Thus, the LO distribution system limits the operating range of the 64-channel receiver to between 50 and 400 MHz.

To test the receiver's range of operation, the receiver was used to detect an RF test signal at the following frequencies: 63 MHz, 85 MHz, and 200 MHz (the approximate proton resonance frequencies for 1.5 T, 2 T, and 4.7 T magnets, respectively). The setup for this test appears in Fig. 56. The HP4195A and the step attenuator were used to generate both an RF test signal and a 10 MHz frequency reference signal. The LO source signal was provided by the PTS-250 frequency synthesizer which resides in the LO distribution system of the 64-channel receiver. The LO source signal frequency was set to 62.5, 84.5, and 199.5 MHz to allow the receiver to demodulate the three different RF test frequencies. The RF test signal, the LO source signal, and the 40 MHz sample clock were phase-locked together by the 10 MHz frequency reference signal. The RF test signal was a ¯70 dBm sinusoid that was injected into one of the RF to IF conversion inputs on the 64-channel receiver's main

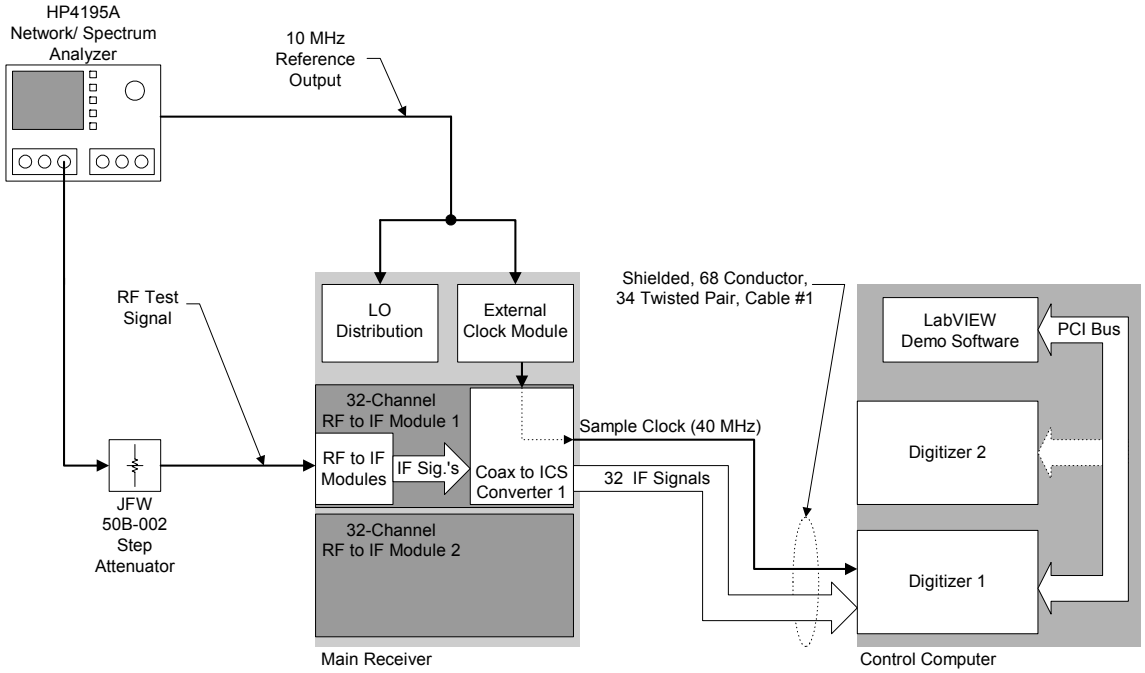

Fig. 56. Schematic diagram of the operating range test for the 64-channel receiver.

receiver module. Finally, an ICS-645 digitizer (denoted "Digitizer 1" in Fig. 56) and the LabVIEW demo software on the control computer were used to display the power spectrum of the detected test signals.

The parallel receiver was able to detect the test signal at each of the three test frequencies (i.e., 63, 85, and 200 MHz). Based on the ability of the 64-channel receiver to detect test signals at the three test frequencies and on the calculations of the receiver channel operating range, we concluded that the operating frequency range of the receiver meets the design requirements.

VI.1.3 Noise Figure

The noise figure of each of the receiver channels was not measured directly. Instead, the parallel receiver's noise figure was evaluated by comparing its image SNR to that of a commercial MR system. This was accomplished by conducting two different sets of imaging tests with the 64-channel receiver on a GE Omega 4.7 T (33 cm bore) magnet system (GE Medical Systems, Milwaukee, WI). First, the performance of the 64-channel preamplifier circuits was compared to that of a single-channel, commercial preamp with a known noise figure. Second, the image SNR performance of the rest of the parallel receiver system (composed of the main receiver and control computer modules) was compared to that of the Omega system single channel receiver. By combing the results from both sets of tests, the overall noise figure performance of the parallel receiver was determined.

VI.1.3.1 Preamplifier Noise Performance Evaluation

The 64-channel system preamplifier was used with the Omega receiver to collect single-channel images. Then, using the same imaging parameters, the Omega system alone was used to collect the same images using a Miteq (Hauppauge, NY) AU-1114 low-noise amplifier. Each of the 64-channel preamplifier circuits should be able to provide 32 dB of gain with a 1.7 dB NF from dc to 500 MHz; whereas, the Miteq AU-1114 is able to produce 30 dB of gain with a 1.4 dB noise figure from 10 to 500 MHz. This test effectively measured the performance of the parallel receiver's preamplifiers against that of a commercial preamp with a known noise figure.

The actual apparatus for the preamplifier noise performance evaluation test is diagrammed in Fig. 57. Single-channel, standard spin echo images (TE/TR= 28/300 ms, 20 kHz spectral width, 1 average, 100 mm FOV, 256×256 points, 2 mm slice thickness) were acquired using both the 64-channel receiver preamplifier circuits and the Miteq AU-1114 low-noise amplifier as the "preamplifier under test" mentioned in Fig. 57. These images were acquired at 4.7 T using a 10 cm (4 in.) inner diameter birdcage coil to transmit and a GE surface coil to receive the MR signals from a test phantom that consisted of an approximately 5.1 cm (2 in.) × 5.1 cm (2 in.) × 7.6 cm (3 in.) Rubbermaid (Wooster, Ohio) container filled with 0.5 mmol $CuSO_4$(aq) solution. The transmit birdcage and receive surface coil were matched and tuned to 50 Ω at 200.237 MHz (to -30 dB and -20 dB, respectively, using the return loss utility in S_{11} measurement mode on the HP4195A network analyzer). In addition, the GE surface coil was decoupled by 24 dB from the transmit birdcage (measured using S_{21} mode on the

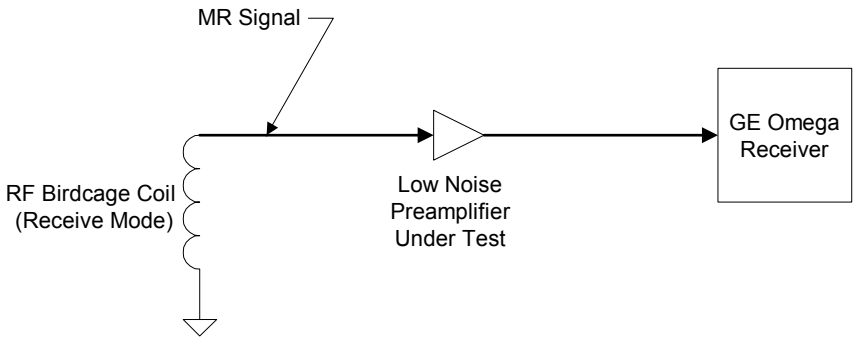

Fig. 57. Schematic of the test apparatus used to evaluate the noise performance of the 64-channel receiver preamplifier circuits using image SNR methods.

HP4195A). An image of the test phantom (using the birdcage coil for both transmit and receive and the same standard spin echo pulse sequence) depicts the topography of the test setup (Fig. 58). (Note that the RF coil and phantom have been rotated away from the horizontal in order to decouple the surface coil from the birdcage.) The image SNR from the resulting images was then calculated for the 64-channel preamplifier circuits and compared to that obtained using the commercial preamp.

Three representative images from the 64-channel receiver preamplifier and one reference image from the commercial preamp, showing the range of image SNR values, appear in Fig. 59. The average SNR in the reference images (using the commercial preamp) was approximately 172 while the SNR values for the images from the parallel receiver's preamplifier circuits ranged from 176.8 to 196.7. This indicates that the noise performance for the 64-channel preamp module circuits is comparable to, if not better than, that of the commercial preamplifier circuit.

VI.1.3.2 Main Receiver Noise Performance Evaluation

The noise performance of the remainder of the parallel receiver system (composed of the main receiver and control computer modules) was tested. In this test, each channel of the parallel receiver was used to separately acquire standard spin echo images (TE/TR= 28/250 ms, 50 kHz spectral width, 1 average, 75 mm FOV, 256×128 points, 20 mm slice thickness) from a test phantom on the GE Omega 4.7 T magnet system. A 10 cm (4 in.) inner diameter birdcage coil was used to both transmit the RF excitation and receive MR signals. The test phantom itself consisted of a 2.2 cm (0.855 in.) inner diameter, 9.5 cm (3.75 in.) long cylinder with three, 1 mm (0.04 in.) inner diameter, 75

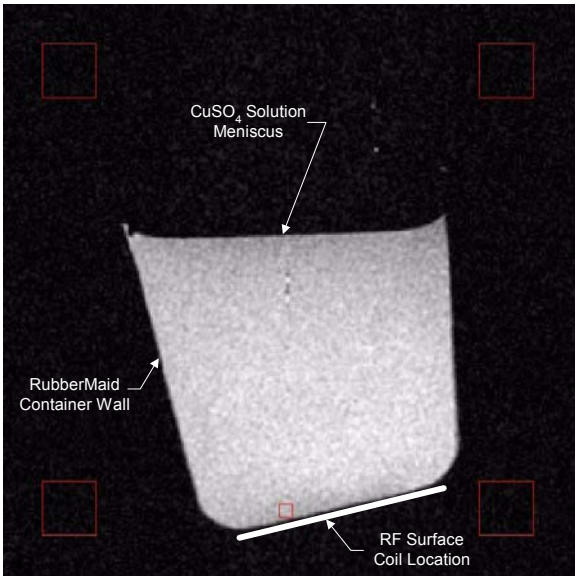

Fig. 58. A reference image of the preamplifier noise performance test phantom obtained using the birdcage coil for both transmit and receive, the Miteq preamp with the Omega receiver, and the same imaging sequence as the actual noise performance tests. The image shows a vertical cut through the plane of the test phantom. The meniscus of the solution inside the phantom and the walls of the Rubbermaid container are all labeled. The plane of the RF surface coil (i.e., perpendicular to that of the image) is shown as a thick white line.

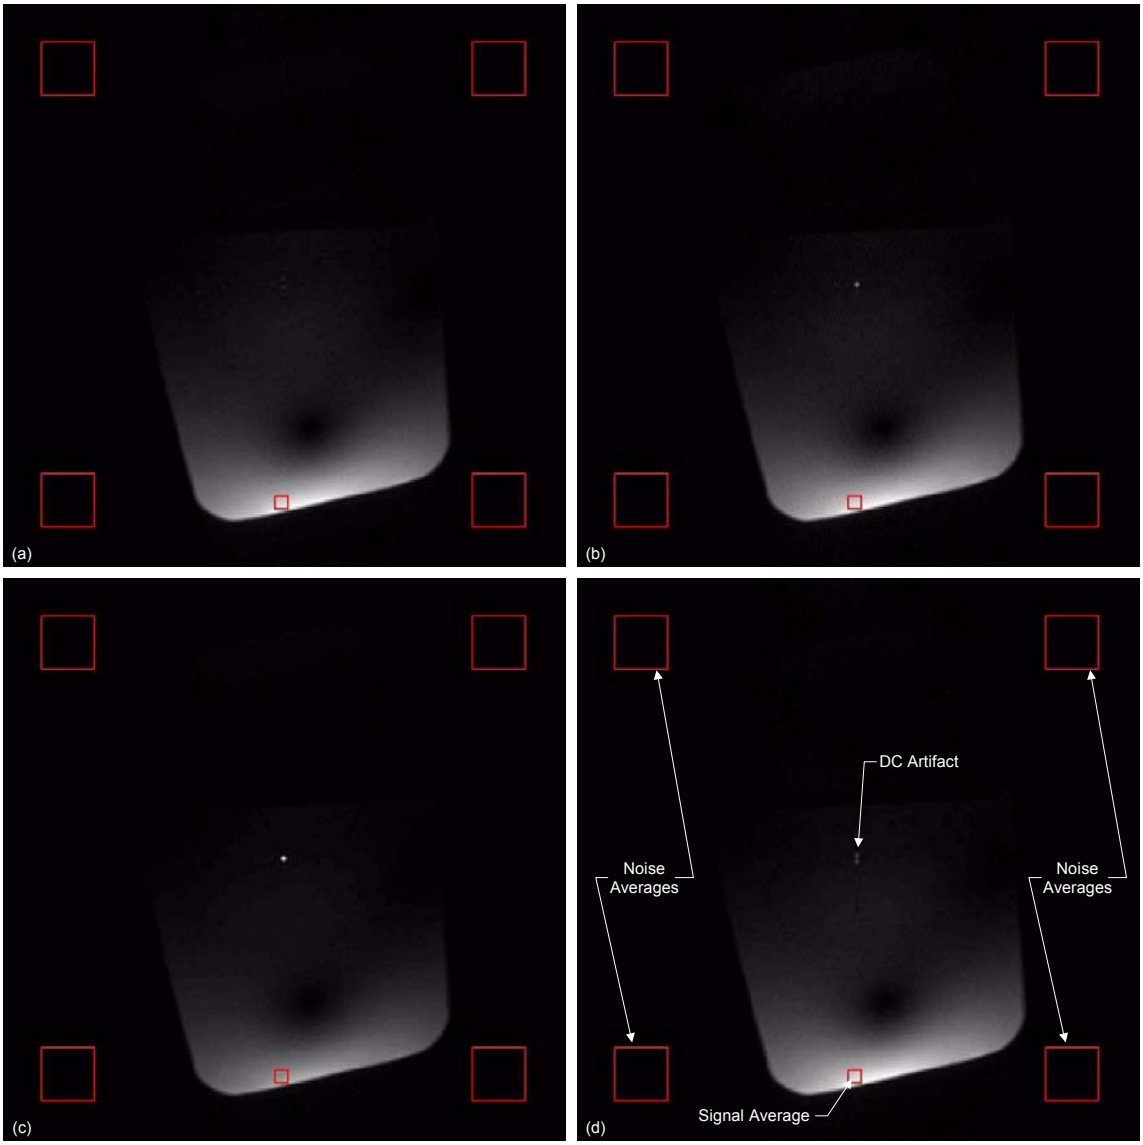

Fig. 59. Representative images from the preamplifier noise performance tests. Images (a)-(c) are from the 64-channel preamplifier module while (d) is one of the reference images from the Miteq preamp: (a) SNR=196.7, (b) SNR=176.83, (c) SNR=192.4, and (d) SNR=172.2. Note that (d) has been labeled to show the locations of the signal and noise sampling regions and the dc artifact (which are present on all four images).

mm (2.95 in.) long capillary tubes spaced ~120º apart along its outer diameter. Both the cylinder and capillary tubes were filled with 0.5 mmol $CuSO_4$(aq) solution. Reference images were collected simultaneously with the Omega receiver to serve as a basis for comparison with the parallel receiver. The image SNR values from each set of images were then calculated. This allowed the noise performance of the main receiver and control computer modules to be gauged against that of a working commercial scanner.

The simultaneous acquisition of single-channel MR image data from both the 64-channel and the Omega system receivers was accomplished by using the test arrangement of Fig. 60. The received signal from the test phantom was first amplified by the Miteq AU-1114 low-noise preamp before traveling to the 2-way RF splitter. The RF splitter was then used to divide the received MR signal equally between the Omega receiver and the parallel receiver channel that was under test. This allowed both receivers to acquire the MR image simultaneously.

The results of the image SNR comparison between the channels of the parallel receiver's main receiver module and the GE Omega receiver are plotted in Fig. 61. The noise performance for the main receiver module is comparable to, if not better than, that of the Omega system receiver on all channels.

VI.1.3.3 Parallel Receiver Noise Performance

Two sets of noise performance tests were carried out on the 64-channel receiver system. In the first test, the performance of the 64-channel preamplifier was evaluated against a single-channel commercial preamp using image SNR measurements. In the second test, the image SNR performance of the main receiver and computer control

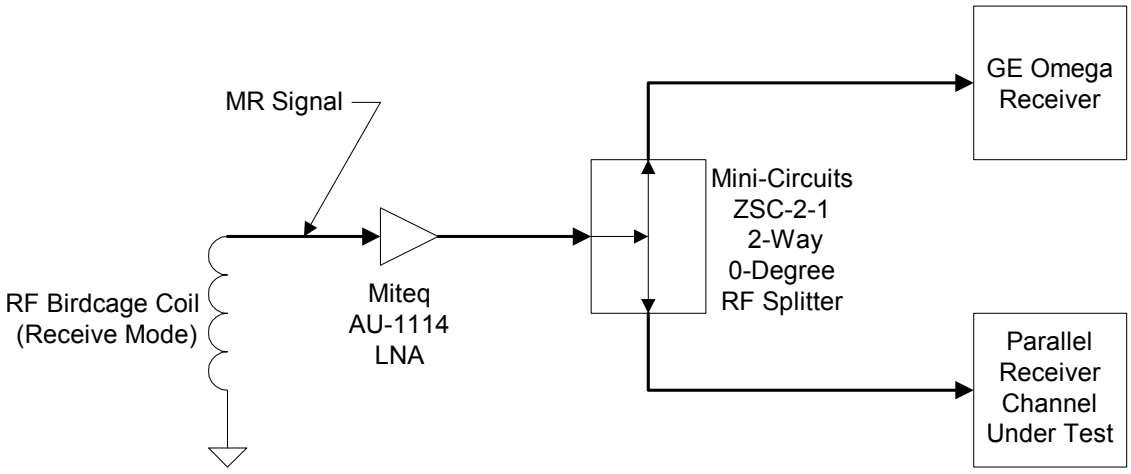

Fig. 60. Test arrangement for performing the image SNR comparison tests with a commercial receiver. The Miteq (Hauppauge, NY) preamp provides 30 dB of gain with a 1.4 dB noise figure from 10 to 500 MHz. The Mini-Circuits (Brooklyn, NY) RF splitter operates from 0.1 to 400 MHz with an insertion loss of approximately 3.4 dB and an amplitude inbalance of at most 0.2 dB at the test frequency.

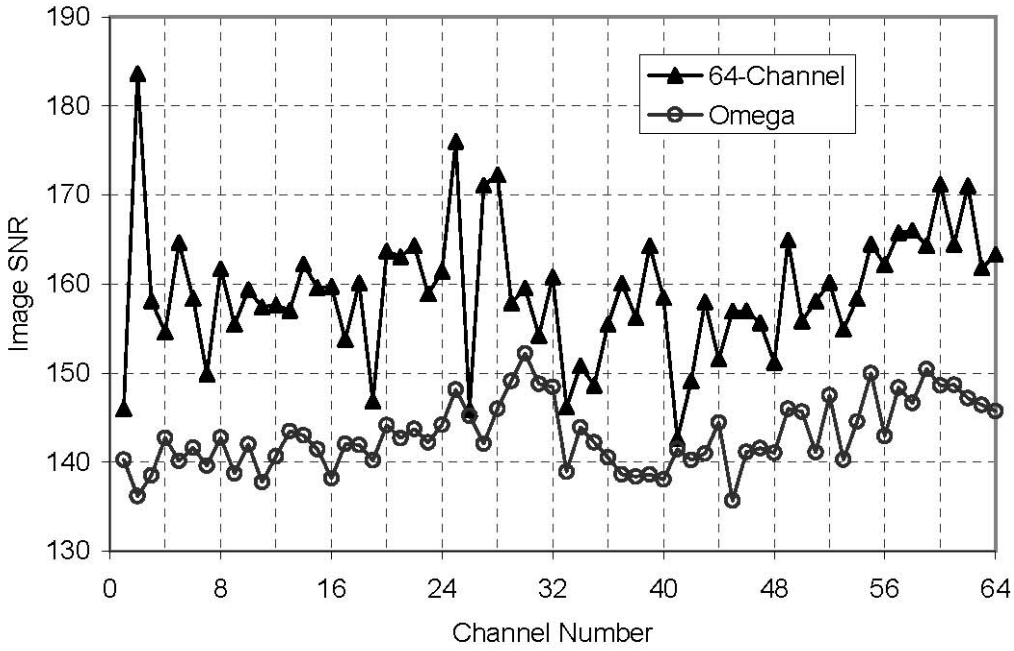

Fig. 61. Plot of the results of the image SNR comparison tests between the 64-channel receiver and a commercial MR system receiver.

modules was gauged against a single-channel commercial receiver system. In both cases, the components of the parallel receiver either out-performed or worked as well as the commercial components. This indicates that together, the overall 64-channel receiver exhibits noise performance that is comparable to, if not better than, that of a single-channel commercial system.

VI.1.4 Dynamic Range

The single-tone dynamic range of the parallel receiver system was measured using the test apparatus of Fig. 62 and Fig. 63 at an ambient temperature of 298 K (76.7° F). The HP4195A network analyzer was used to inject a test signal with a known power level into each channel of the parallel receiver system. The insertion loss of the cables and attenuator which carried the test signal to the preamplifier inputs of the receiver was measured prior to the test. A push-button attenuator was used in conjunction with the signal amplitude control on the network analyzer to vary the power level of the test signal (at the preamplifier inputs) from $^{-}57$ to $^{-}124.2$ dBm. The signal and noise response of each receiver channel was then observed using the LabVIEW demonstration utility on the ICS-645 digitizer card (denoted "Digitizer 1" in both Fig. 62 and Fig. 63). These signal and noise levels were recorded for each channel and later used to determine the maximum and minimum detectable signal levels for a computation of dynamic range.

The average minimum detectable signal level measured with the channels of the 64-channel receiver (taking into account a processing gain of 9 dB for oversampling and 33

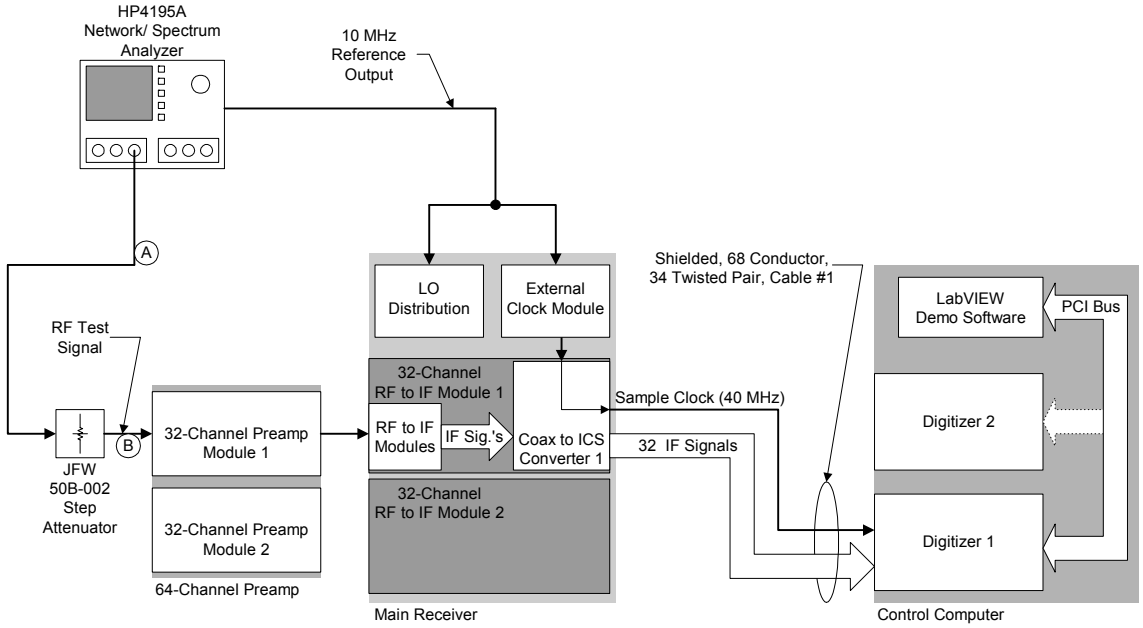

Fig. 62. Equipment setup for measuring the dynamic range of the first thiry-two channels in the parallel receiver. In the RF test signal path, the lines labeled "A" and "B" are two RG-58A/U coaxial cables each having insertion losses of 3.7 dB. The total insertion loss of the JFW 50B-002 step attenuator was measured to be 56.6 dB (at maximum attenuation).

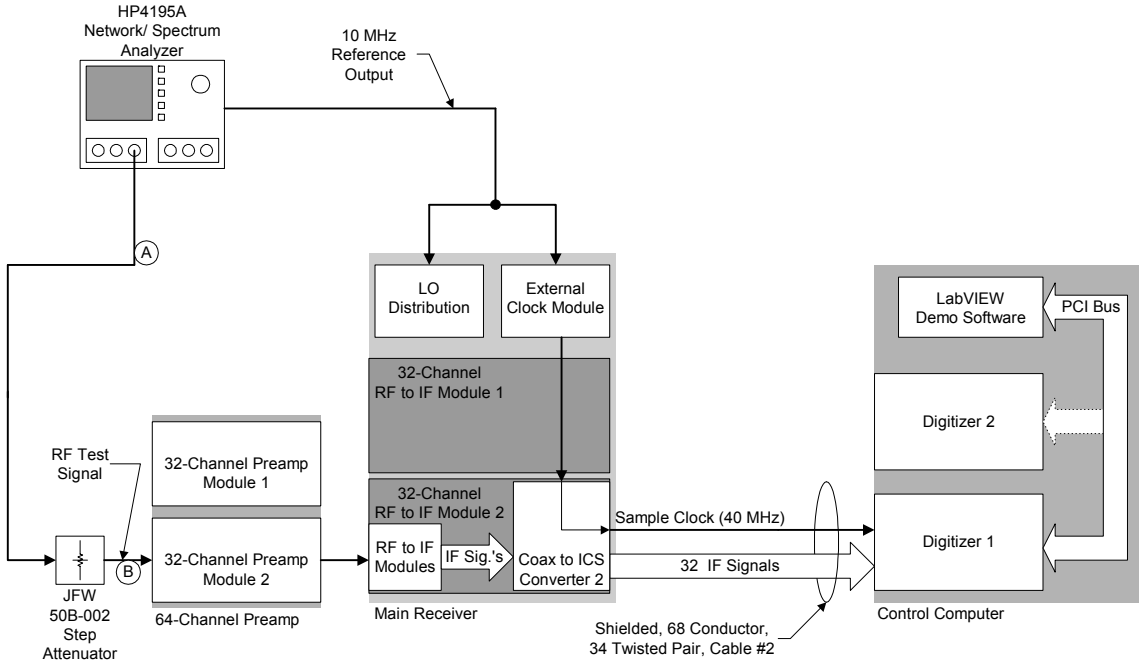

Fig. 63. Equipment setup for measuring the dynamic range of the second thiry-two channels in the parallel receiver. In the RF test signal path, the lines labeled "A" and "B" are two RG-58A/U coaxial cables each having insertion losses of 3.7 dB. The total insertion loss of the JFW 50B-002 step attenuator was measured to be 56.6 dB (at maximum attenuation).

dB for coherent averaging due to collecting 2048 point samples) was -111.2 dBm. (Refer to Chapter II, sections II.4.2.3.3 and II.4.2.3.1for descriptions of the effects of oversampling and coherent integration on dynamic range.) The average maximum signal detected (with a 1 dB output compression) was -59 dBm. This indicates an average single-tone dynamic range of 52.2 dB for the channels of the parallel receiver.

VI.1.5 Interchannel Isolation

The interchannel isolation on the 64-channel receiver was determined in two parts. First, the isolation between channels in each of the 32-channel preamp modules was measured to be greater than 20 dB. Then the interchannel isolation for the remainder of the receiver (i.e., the main receiver and the control computer modules) was measured. Finally, the isolation data from both measurements were combined to yield a net isolation figure for the parallel receiver.

The interchannel isolation for the main receiver and control computer modules was measured in two groups of thirty-two channels. This allowed each of the two, 32-channel RF to IF conversion modules to be measured separately. To perform the actual measurement, an RF test signal was injected, one channel at a time, into each channel of the 32-channel RF to IF conversion module. (During the test, the unused channel inputs were terminated with 50 Ω loads.) The IF signal output levels for all thirty-two channels in the module under test were then monitored to measure the amount of signal leakage between channels.

The actual test setup was similar to that described in section VI.2.2. An HP4195A network analyzer created the RF test signal and a 10 MHz phase reference signal. The

reference signal phase-locked the RF test signal, the receiver's LO source signal, and the 40 MHz sample clock. The gain controls for all of the receiver channels were set to the same positions (i.e., the VGA control network's fine adjustment potentiometers were set for maximum gain while the variable attenuator control network's control potentiometers were set for maximum attenuation). The RF test signal was a 199.9 MHz, -70 dBm, sine wave. (Fig. 64 shows the test setup and signal connections for measuring the isolation between the receiver channels in the first RF to IF conversion module; Fig. 65 depicts the connections for measuring the receiver channels in the second RF to IF conversion module.) The output IF signals from all thirty-two channels were then digitized by the ICS-645 digitizer (denoted "Digitizer 1" in both Fig. 64 and Fig. 65) in the control computer module. Finally, the power spectrum for each digitized IF signal was displayed by the LabVIEW demonstration software, and its power level was recorded. This allowed the amount of coupling between receiver channels to be determined.

The isolation between channels for the two, 32-channel RF to IF conversion modules was measured. The minimum isolation between two channels in each of the 32-channel modules was 37.9 dB; however, there were only five occurrences of interchannel isolation values lower than 40 dB. The average of the interchannel isolation values for all sixteen of the 4-channel, RF to IF submodules appears in Table V. Note that the average interchannel isolation within the 4-channel submodules is greater than 40 dB. Finally, the isolation between the channels of different 4-channel, RF to IF submodules was greater than 46 dB in all cases.

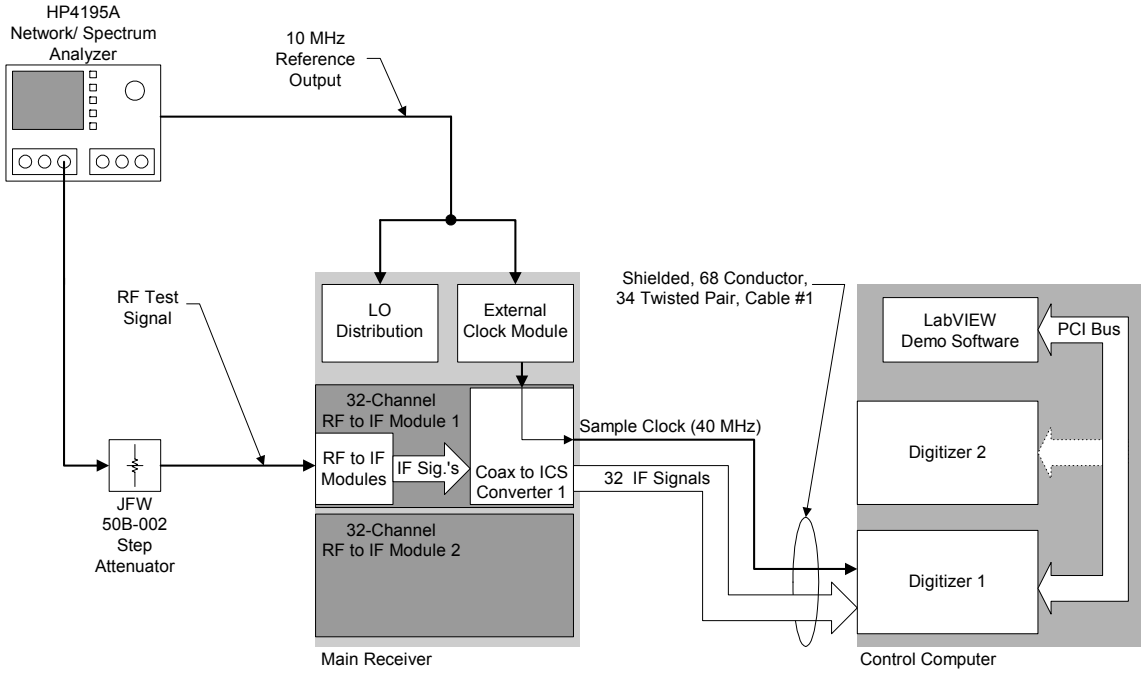

Fig. 64. Test schematic for the interchannel isolation and channel-to-channel gain variation tests for RF to IF Module 1.

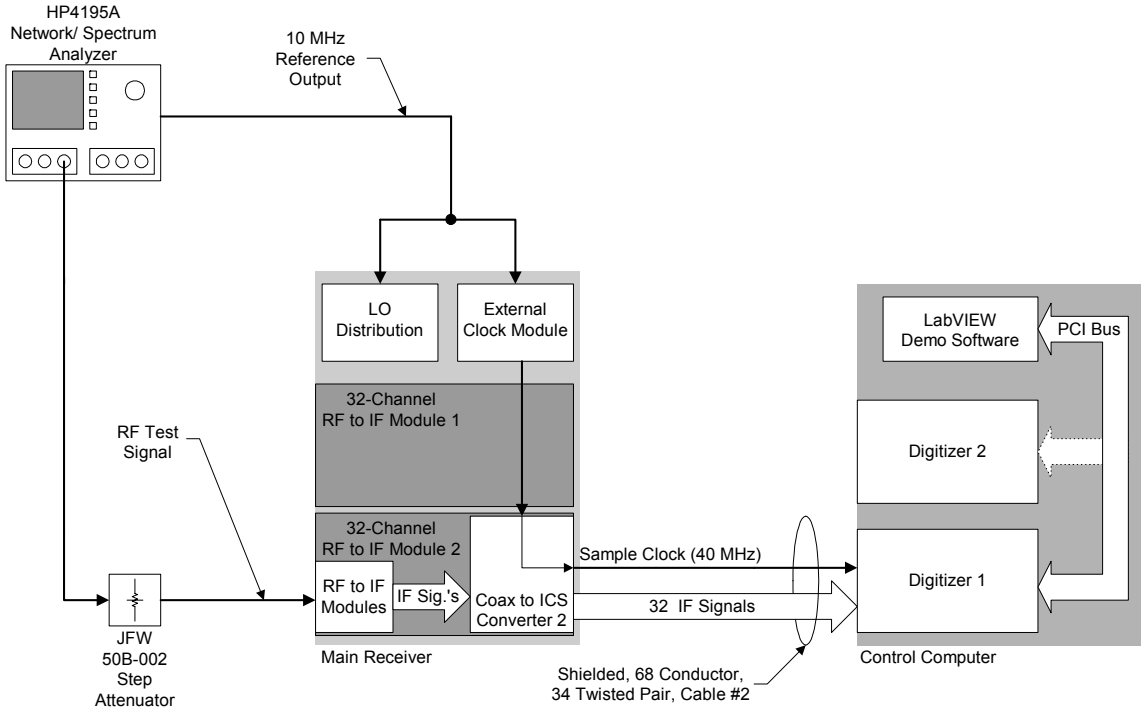

Fig. 65. Test schematic for the interchannel isolation and channel-to-channel gain variation tests for RF to IF Module 2.

In order to define an interchannel isolation level for the entire 64-channel receiver, the isolation data from both the preamplifier modules and that of the main receiver and control computer modules must be combined. The isolation between all channels in each of the two, 32-channel preamplifiers was found to be greater than 20 dB. Combining this result with the measurements for the remainder of the 64-channel receiver (i.e., at least 37 dB between channels) yields a net interchannel isolation level of 20 dB. (The interchannel isolation level in a cascaded multi-channel system cannot be higher than the lowest isolation value of its components.) This means that the interchannel isolation on the 64-channel receiver does not meet the initial design requirement of 35 dB of isolation between all channels.

Table V
Average Interchannel Isolation Data for the 4-Channel RF to IF Submodules

	Relative Output Power Level [dB]			
Input	Channel 1	Channel 2	Channel 3	Channel 4
Channel 1	0	47.8	53.2	54.6
Channel 2	48.9	0	45.7	52.4
Channel 3	53.1	45.0	0	48.6
Channel 4	51.5	51.1	43.3	0

VI.1.6 Channel- to-Channel Gain Variation

The total channel-to-channel gain variation for the parallel receiver, like the interchannel isolation, was found using a two part procedure. The 32-channel preamplifier modules were first measured to determine their gain variation. Then the

remainder of the receiver (i.e., the main receiver and the control computer modules) was measured to find its channel-to-channel gain variation. Finally, the data from both measurements were combined to yield a net channel-to-channel gain variation level for the parallel receiver.

The channel-to-channel gain variation was measured separately for each of the two, 32-channel RF to IF conversion modules. To perform the measurement, an RF test signal was inserted, one channel at a time, into each channel of the 32-channel RF to IF conversion modules. (During the test, the unused channel inputs were terminated with 50 Ω loads.) The output IF signal, for the channel into which the test signal had been inserted, was then monitored to measure its power level.

The test setup for measuring the channel-to-channel gain variation was identical to that used for the isolation measurement (refer to section VI.2.5, Fig. 64, and Fig. 65). The HP4195A network analyzer generated an RF test signal and a 10 MHz phase reference signal. The reference signal phase-locked the RF test signal, the receiver's LO source signal, and the 40 MHz sample clock. The gain controls for all of the receiver channels were set to the same positions (i.e., the VGA control network's fine adjustment potentiometers were set for maximum gain while the variable attenuator control network's control potentiometers were set for maximum attenuation). The RF test signal itself was a 199.9 MHz, ˉ70 dBm, sine wave. The output IF signal was digitized by the ICS-645 digitizer in the control computer module. Finally, the power spectrum for the digitized IF signal was displayed by the LabVIEW demo software, and its power level was recorded.

The gain variation between channels in each of the 32-channel RF to IF conversion modules is shown in Fig. 66. Note that the maximum variation in gain between receiver channels is 4.2 dB. This gain variation can be minimized to 3.3 dB through the use of the individual gain controls (these gain controls allow the gain to be adjusted in 4-channel groups) on each of the 4-channel modules (Fig. 67).

Finally, to define an the channel-to-channel gain variation level for the entire 64-channel receiver, the gain variation data from both the preamplifier modules and that of the main receiver and control computer modules is combined. The gain variation within each of the 32-channel preamp modules was found to be about 2 dB (i.e., an average gain of 32 dB ±1 dB). Combining this result with the measurements for the remainder of the 64-channel receiver yields a total interchannel gain variation level of 6.2 dB (without gain scaling) for the receiver. The total gain variation for the parallel receiver can of course be further reduced to approximately 5.3 dB by using the VGA network's fine adjust controls that are located on each of the 4-channel RF to IF conversion submodules.

VI.1.7 System Cost

The final version of the 64-channel parallel receiver prototype was built for approximately $33,500 in parts and materials. The cost of labor, engineering, and software development has not been included as all tasks were performed by graduate students working part time over an eighteen month period. Except for the control computer (see Chapter V, section V.2.4), the 64-channel digitizer (refer to Chapter V, section V.2.4.1), and the master LO source, the system was constructed in-house from

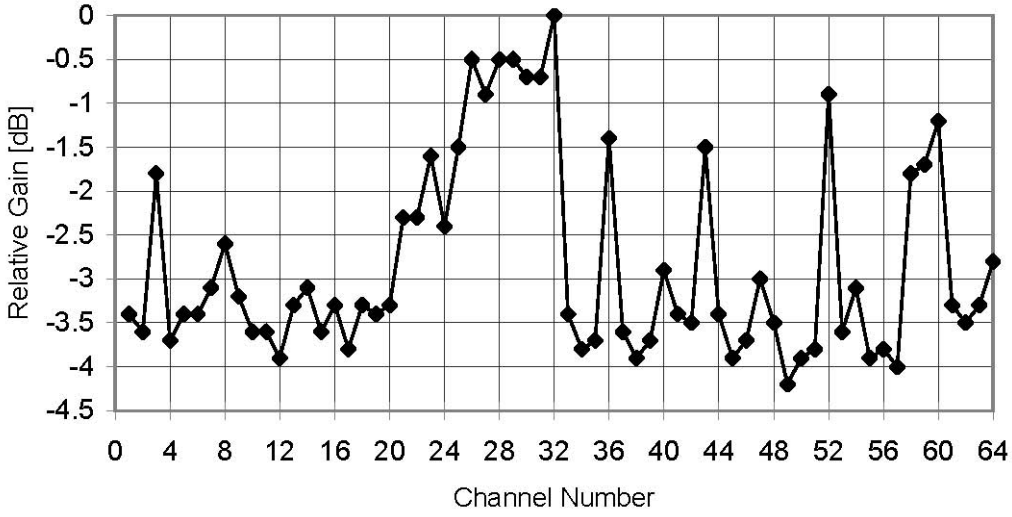

Fig. 66. Plot of the gain variation test results for each of the sixty-four receiver channels. The power output (in dB) from each channel relative to the maximum recorded response (designated as 0 dB) is shown.

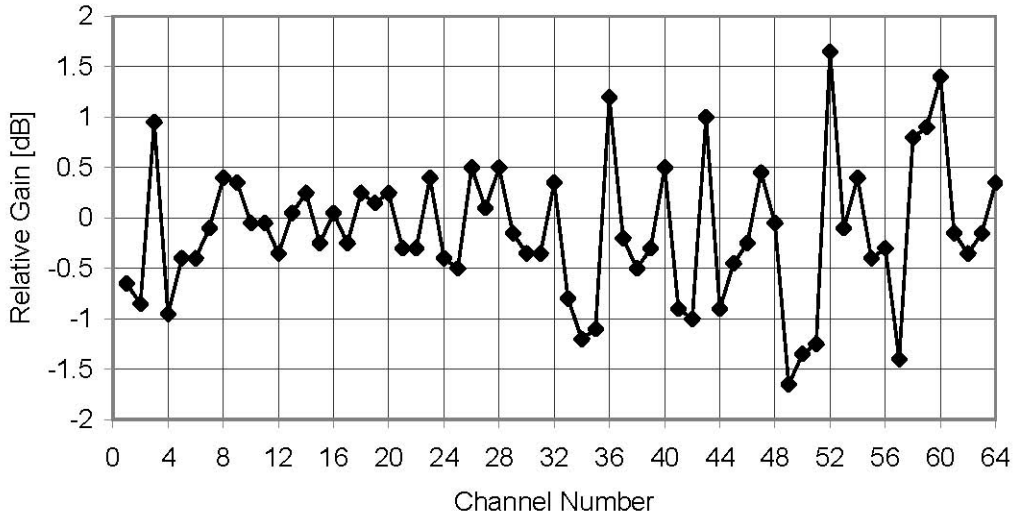

Fig. 67. Plot of the gain variation test results for each of the sixty-four receiver channels scaled by varying the VGA control network fine adjustment potentiometer on each 4-channel RF to IF submodule. The power output (in dB) from each channel relative to the midpoint in the recorded responses for each 4-channel submodule (designated as 0 dB) is shown.

low-cost components and systems. In all, the 64-channel prototype was built for less than $525 per channel.

VI.2 Summary

The performance of the 64-channel receiver system has been characterized to ascertain how well the system meets its own design criteria. To accomplish this evaluation, the channel bandwidth, operating frequency range, noise performance, dynamic range, interchannel isolation, and gain variation between channels on the 64-channel receiver have all been measured and the system cost was analyzed. The parallel receiver was found to have a 3 dB channel bandwidth exceeding 1 MHz. The receiver system's operating range was calculated to extend from 50 to 400 MHz, and its operation was verified at 63, 85, and 200 MHz. The actual noise figure for each of the channels in the parallel receiver was not measured directly; however, the noise performance of each channel was compared to that of a single-channel commercial MR scanner operating on a 4.7 T magnet system. The noise performance of the 64-channel receiver was found to be comparable, and in most cases better than, that of the commercial system. The average dynamic range of the parallel receiver was measured to be 57.2 dB. The receiver system's interchannel isolation was measured and found to be greater than 20 dB between all channels. Also, the channel-to-channel gain variation on the 64-channel receiver was less than 6.2 dB. Finally, the overall cost of the prototype was determined to be less than $525 per channel. In all, the performance of the 64-channel receiver prototype either met its design criteria in all areas except for that

of interchannel isolation. In this last area, the system failed to meet its design requirement due to a higher than expected level of coupling in the preamplifier module.

CHAPTER VII

DIGITIZER SYNCHRONIZATION TECHNIQUES FOR CORRECTION OF THE PHASE JITTER ARTIFACT IN MR RECEIVERS

When the sampling clock on the digitizer within an MR receiver is not properly synchronized to the pulse sequences of an MR system, phase encode jitter artifacts will appear in the MR images collected by that receiver. It is usually not possible to process these jitter artifacts out of the MR images. As a result, any MR image data collected by such a receiver is essentially useless. This situation may be remedied by synchronizing the digitizer to the host MR system in hardware. This synchronization may be accomplished by using one of four techniques. The first method involves oversampling the signal to be digitized. A second approach uses the Fourier time-shifting theorem to re-align the phase of the received data with that of the host MR system. A third method is to generate a synchronous sampling clock signal external to the digitizer card and use it to replace the internal sampling clock on that digitizer. The last technique is to pre-sample the waveform to be digitized with an external sample-and-hold circuit in front of the digitizer card. These techniques work by synchronizing the digitizer to either the acquisition trigger signal provided by the host MR scanner or to the master system clock on that MR scanner. The phase jitter problem as well as each of the four synchronization methods will be examined in the following sections.

VII.1 The Phase Jitter Problem

In order to understand how phase jitter arises from an asynchronously sampling digitizer, it is necessary to understand how sampling would occur on a digitizer that is

properly synchronized to a given pulse sequence. During an MR pulse sequence, the MR scanner signals the digitizer to begin acquiring data by sending an acquisition trigger signal. This signal indicates the start of a new acquisition period during which time the digitizer will acquire a preprogrammed number of samples at a preprogrammed sampling rate. The digitizer will then store the acquired samples (usually in memory on the MR scanner control computer). Finally, the digitizer will wait for the next acquisition trigger signal, at which time the whole data acquisition process will repeat itself. The time between the application of the acquisition trigger signal and the collection of the first data sample by the digitizer during each acquisition event is called the acquisition time delay. If the digitizer is properly synchronized to the pulse sequence, the acquisition time delay will remain constant for each signal acquisition (Fig. 68).

A common problem with many off-the-shelf digitizer cards is that they have a free-running, on-board, sampling clock. A sampling clock is used by the digitizer to control when an input analog signal is sampled. If this sampling clock is also free-running, the digitizer will collect samples continuously and store them to an on-board memory buffer. Each time the buffer fills, the data in that buffer is overwritten. In MR imaging, the digitizer is usually be operated in an externally triggered mode. In this mode, the digitizer can be made to store the sampled data to a more permanent memory location when an external trigger signal is received by the card. Upon reception of the external trigger signal, a preprogrammed number of samples (that are acquired after the trigger at a preprogrammed sample rate) will be stored in memory. The problem with this

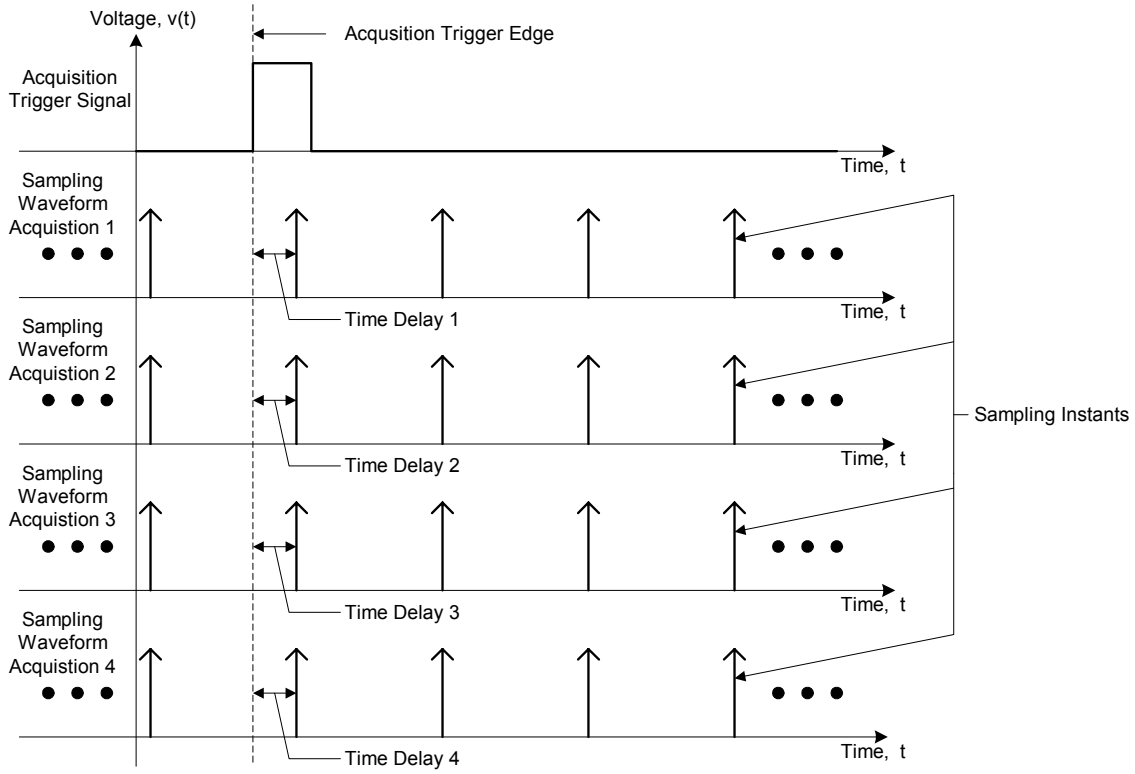

Fig. 68. MR Signal acquisition using a synchronous digitizer. The acquisition trigger signal (top waveform) and the portions of the sampling waveforms from four, successive MR signal acquisitions (lower four signals) are shown. In the lower four waveforms, each vertical arrow indicates a sampling point. Note that the first vertical arrow to the right of the acquisition trigger edge denotes the first sample of the MR signal that will be stored in each acquisition event. Also note that the acquisition time delay is constant for each acquisition.

situation is that the preprogrammed number of samples is not immediately acquired after the trigger signal; rather, the digitizer will wait until the next transition of its on-board clock signal to acquire the first data point. This on-board clock, however, is not guaranteed to be synchronized to the pulse sequence in any way. As a result, the acquisition time delay will vary from acquisition to acquisition within the pulse sequence. In Fig. 69, a digitizer with a free-running, asynchronous sample clock has been used to sample the MR signals received during an MR imaging experiment. The sampling clock is running at the Nyquist rate (i.e., at twice the bandwidth of the signal to be sampled) and is not synchronized to the acquisition trigger signal from the main MR system. On successive acquisition events, a different (and essentially random) time delay occurs between the trigger signal from the MR scanner and the onset of sampling. These random delays are responsible for the phase encode jitter artifact in the reconstructed MR image data.

The effect of a variable length acquisition delay time between the acquisitions in an MR imaging experiment is best understood by turning to the theory of Fourier transforms. There is a theorem of Fourier transforms which states that a delay in a time domain signal corresponds to a linear phase shift in the frequency domain content of that signal [49]. The Fourier transform pair which represents this theorem is given in (VII.1), where $x(t)$ is the time domain function, τ is the amount of time that $x(t)$ has been delayed, $X(\omega)$ is the Fourier transform of $x(t)$, and $\omega\tau$ is the phase shift imparted on the frequency content of $X(\omega)$. When a received MR signal experiences a time delay, a linear phase shift is introduced into the frequency content of that signal via (VII.1). If this time delay

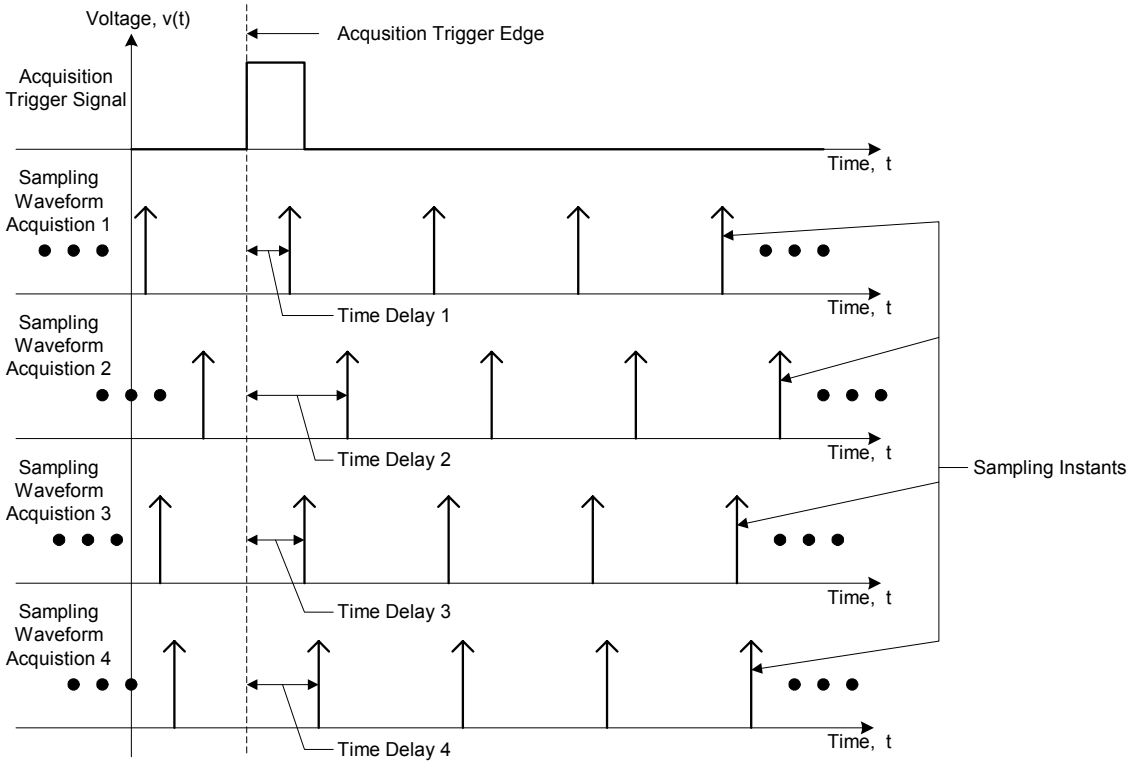

Fig. 69. MR Signal acquisition using a digitizer with an asynchronous sample clock. The acquisition trigger signal (top waveform) and the portions of the sampling waveforms from four, successive MR signal acquisitions (lower four signals) are shown. In the lower four waveforms, each vertical arrow indicates a sampling point. Note that the first vertical arrow to the right of the acquisition trigger edge denotes the first sample of the MR signal that will be stored in each acquisition event.

were constant for all the acquired signals in the MR imaging data set, the only artifact would be a linear phase ramp along the phase encode direction of the reconstructed image. However, if a different time delay occurs for each signal acquired in the MR imaging data set, a set of corresponding phase shifts will be imparted to the frequency content of each received line of k-space. These different phase shifts, in turn, result in a phase jitter artifact in the final MR image. Unfortunately, without prior knowledge of the duration of each of these time delays, the phase jitter artifact due to an asynchronous sampling digitizer cannot be corrected in post-processing.

$$x(t-\tau) \Leftrightarrow X(\omega)e^{-j\omega\tau} \qquad \text{(VII.1)}$$

VII.2 Oversampling

One way to synchronize a digitizer to an MR scanner is to oversample the received signals with that digitizer [16]. In this case, oversampling indicates sampling at a rate higher than the Nyquist rate for a given signal. So, for a received signal with a bandwidth of 20 kHz, oversampling would involve sampling at any rate higher than 40 kHz. Oversampling is the easiest synchronization method to implement and has been used successfully on the desktop MR system prototype; however, the oversampling method will not work for every digitizer. This section will present both the theory and the practical limitations of the oversampling technique.

VII.2.1 Oversampling Theory

Oversampling synchronizes a digitizer to the MR scanner by reducing the length, and consequently the variance, of the acquisition time delays which occur during the acquisition of an MR image. By increasing the sample rate on an asynchronous

sampling digitizer, the length of the acquisition time delays can be reduced considerably. This is shown in the example of Fig. 70, where the sample rate on the digitizer has been increased by a factor of five. By extension, if the sample rate is increased sufficiently, the acquisition time delays can be reduced to such an extent that the phase jitter artifact is eliminated from the MR image data. In this way the oversampling technique may be used to synchronize a digitizer to the acquisition trigger signal of an MR scanner.

The oversampling method for digitizer synchronization on an MR receiver may be implemented in three steps. First, the digitizer must be reprogrammed to acquire data at a faster rate (i.e., the oversampling rate). This, of course, requires a digitizer capable of collecting data at this higher sample rate. Next, the digitizer is used to collect MR signal data from the receiver at the oversampling rate. Finally, the oversampled data is decimated to the Nyquist sampling rate for the MR signal data before it is reconstructed into an MR image. If the oversampling rate is high enough, the length of the acquisition time delays will become negligible and the decimated signal will be synchronized to the acquisition trigger pulse.

The oversampling rate required to synchronize a given digitizer must be found qualitatively, as it can be difficult to measure the actual acquisition delay time. To do this, the user selects an oversampling rate, reprograms the digitizer to acquire data at the new rate, and then collects an MR image data set with the digitizer at the oversampling rate. If (after decimation back to the Nyquist rate) the resulting MR image data still has noticeable phase jitter artifacts, then a higher oversampling rate should be chosen. The programming and data collection process would then be repeated until an oversampling

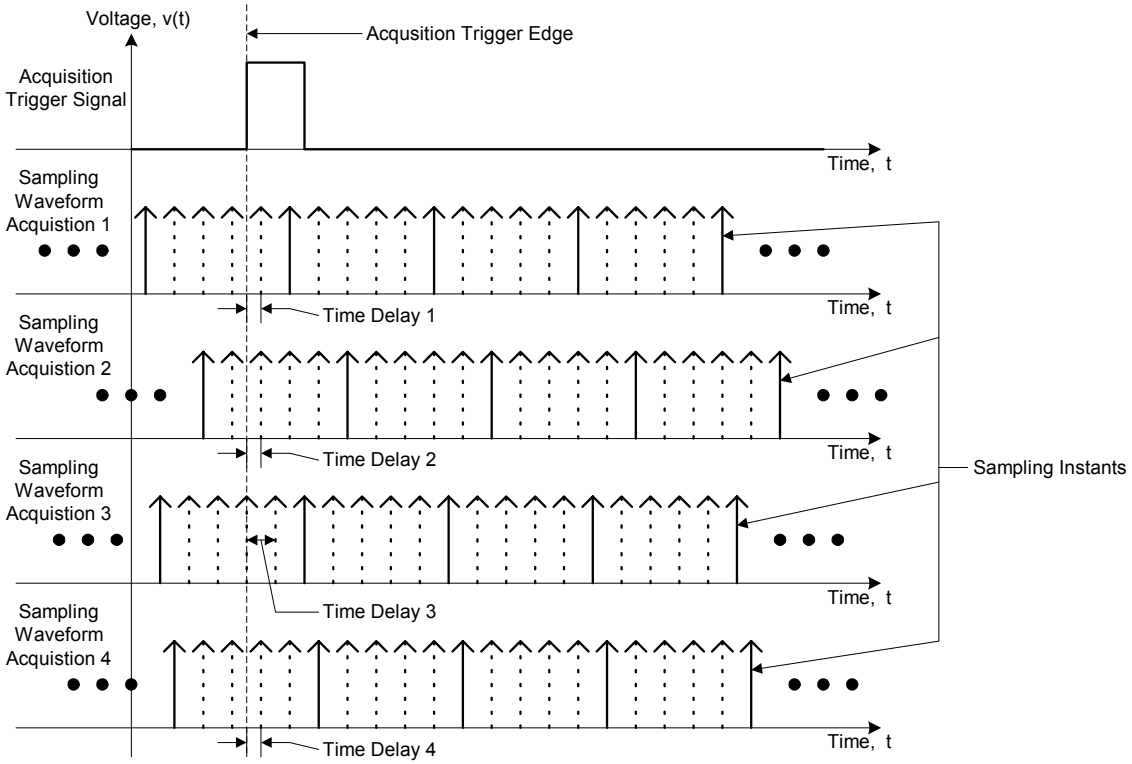

Fig. 70. Digitizer synchronization via oversampling. The acquisition trigger signal (top waveform) and the portions of the sampling waveforms from four, successive MR signal acquisitions (lower four signals) are shown. In the lower four waveforms, each vertical arrow indicates a sampling point. The solid vertical arrows indicate the samples that would be acquired if the sampling occurred at the Nyquist rate; whereas, the dashed vertical arrows and solid arrows together indicate oversampling at five times the Nyquist rate.

rate is found that diminishes the phase jitter artifact to an acceptable level. If an appropriate oversampling rate is not found (either because the digitizer cannot be programmed to sample at any higher rates or because no level of oversampling appears to solve the problem), then one of the remaining three digitizer synchronization methods must be used.

VII.2.2 Oversampling Limitations

Oversampling is by far the simplest digitizer synchronization method to implement; however, there are two cases in which the method tends to break down. If either of these situations arises, one of the other three digitizer synchronization methods must be used. The first case occurs when the Nyquist sampling rate for a signal is too high. When this happens, the digitizer in question may not be able to accommodate the oversampling rates required for synchronization. The second case occurs when the length of the acquisition time delay for a given digitizer varies by more than one sample period (at the Nyquist sampling rate) from acquisition to acquisition. In this case, even high oversampling rates can lead to large differences in the delay times between successive acquisitions. Both of these situations will be examined in the following paragraphs.

The oversampling technique is best suited to situations where the Nyquist sample rate for a received signal is relatively low (i.e., less than ~200 kHz). In this case, an acceptable oversampling rate is usually less than 1 MHz. If the bandwidth of the received signal is significantly higher than 100 kHz, then the required oversampling rates become difficult to achieve with readily available digitizer cards. (It is still difficult to find commercially available digitizer cards capable of 16-bit sampling at

rates higher than 10 MHz; however, inexpensive digitizers with 16-bit sample rates up to 1 MHz are common.) Thus, for signals with these higher Nyquist sample rates the oversampling method of synchronization often will either not work or will become too costly for practical implementation. This was the situation for the 64-channel receiver prototype where the received signal bandwidth for each channel was 1.25 MHz. The 64-channel system digitizers automatically incorporated a factor of eight oversampling ratio (i.e., the cards sampled the received signal at 20 MHz and decimated it down to 2.5 MHz). Not only was this oversampling rate not fast enough to eliminate the phase jitter problem, but replacing these digitizers with higher sample rate versions would have been impractical.

If the acquisition time delay varies by more than one sample period between successive acquisitions, the oversampling technique will not work. This situation usually indicates a digitizer card with a poor trigger response time. This poor response time may be the result of conflicts with the operating system under which the digitizer is running, an imaging sequence with a repetition time that is too short for the digitizer, or noise on the trigger signal. Whatever the cause, the oversampling method will not usually be able to synchronize digitizers which suffer from this condition. In general, synchronization will only be achieved once the underlying cause of the problem is addressed.

VII.3 Fourier Time-Shifting

In the Fourier time-shifting method of digitizer synchronization the acquisition time delay for each acquired MR signal is reset so that all the time delays have the same

value. This is accomplished by using the time-shifting property of the Fourier transform to realign the phase of the received MR signal data. In order for the Fourier time-shifting method to work, the length of each acquisition time delay must be measured and recorded. To do this the acquisition trigger signal from the MR scanner is integrated and acquired while the MR signal data is collected. The length of each acquisition time delay may then be determined mathematically from the integrated trigger signals. Once the acquisition time delays have been found, the phase of each MR signal may be realigned to the start of the signal acquisition period using a time-shifting algorithm. The Fourier time-shifting technique was used to successfully synchronize the digitizers in an early version of the 64-channel receiver prototype. This section will present the Fourier time-shifting method and its implementation on the 64-channel receiver.

VII.3.1 Fourier Time-Shifting Theory

Recall that when a digitizer with an asynchronous sampling clock is used to collect MR signal data, a series of random time delays (described in section VII.1) is introduced to the data collected for each line of k-space. These random time delays introduce a set of correspondingly random, linear, phase shifts into the frequency content of each sampled MR signal. These differing phase shifts, in turn, result in phase jitter in the final MR image. Digitizer synchronization via Fourier time-shifting seeks to time-shift the first sample in each collected line of k-space back to the same starting time (relative to the acquisition trigger signal) by nulling out each of these phase shifts. This involves first performing a Fourier transform on each received MR signal. Each transformed signal is then multiplied by a linear phase shift to cancel out the effect of the acquisition

time delay (i.e., for an MR signal delayed by a time, τ, the Fourier transform of that signal would be multiplied by $e^{j\omega\tau}$). Finally, each, phase-corrected, Fourier transformed MR signal must be inverse Fourier transformed. These operations effectively synchronize the MR data acquired from an asynchronous sampling digitizer to the acquisition trigger signal on the MR system.

In order to perform synchronization by Fourier time-shifting, some record of the length of each of the acquisition time delays must be kept during the collection of the MR signal data. One way to do this is to integrate the acquisition trigger signal and digitize it at the same time as the MR signal. (Digitization of the integrated trigger signal must be performed with the same digitizer as that used to acquire the MR signal data; otherwise, the recorded time delays would differ from those affecting the MR data.) Ideally, when the acquisition trigger signal (i.e., a square pulse) is integrated, a linear ramp function is generated (this is depicted in Fig. 71). Now, if the integrated trigger signal is digitized at the same time as each line of k-space in an MR imaging experiment (Fig. 72), the samples taken during the linear portion of the integrated trigger signal may then be used to determine the time delay between the acquisition trigger signal edge and the start of digitization. (Note: the integrated trigger signal in Fig. 72 has been inverted as the easiest integration circuit to build is an inverting op-amp integrator.)

Before the length of the acquisition time delay can be determined from the samples of the integrated trigger signal (Fig. 72), an equation for the ramp function within the

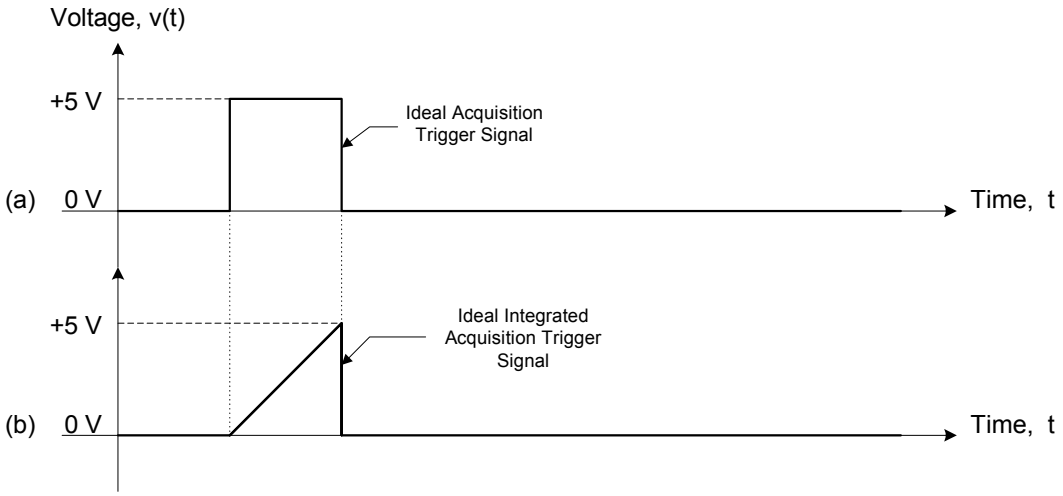

Fig. 71. Plots of (a) an ideal acquisition trigger signal and (b) an ideal, integrated acquisition trigger signal.

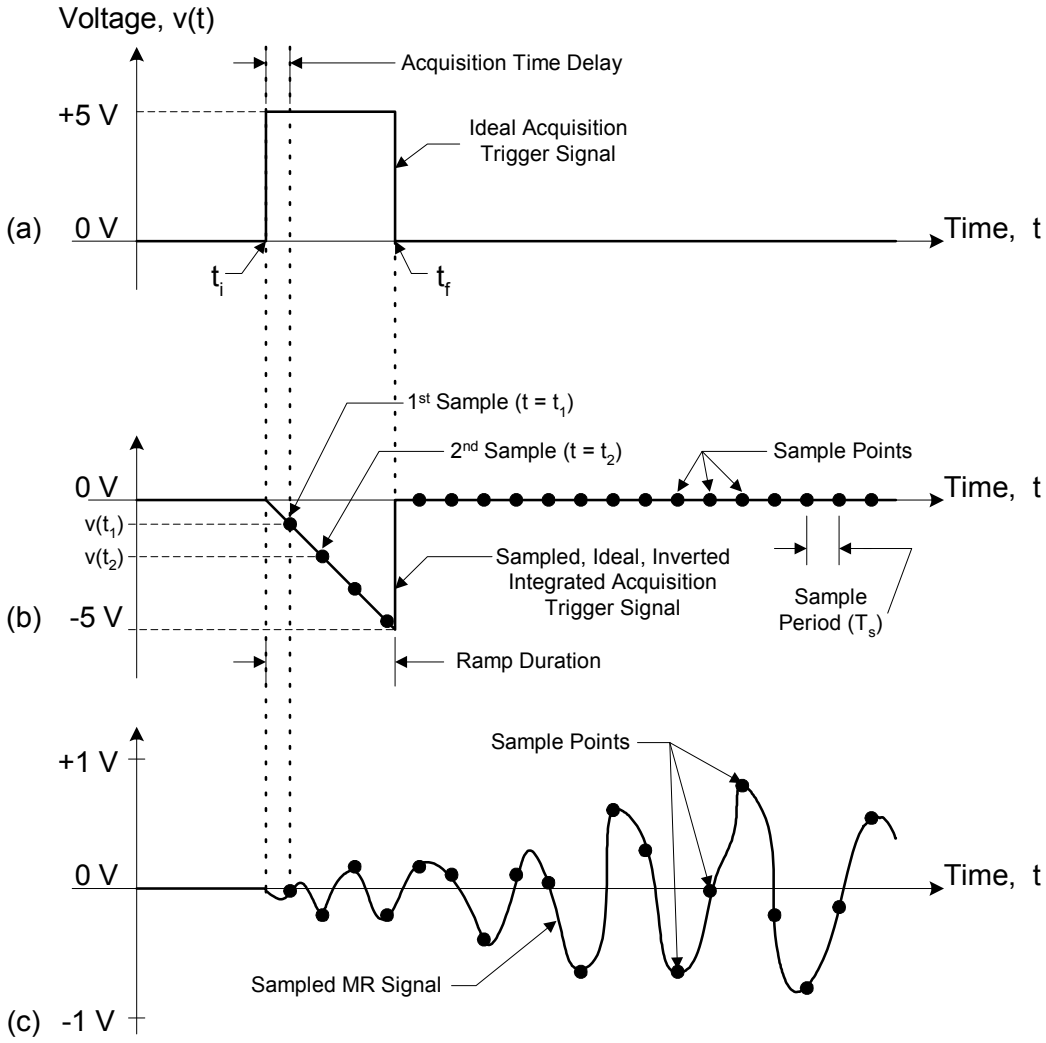

Fig. 72. The acquisition time delay for the simultaneous digitization of an integrated acquisition trigger signal and an MR signal. Plots of the (a) acquisition trigger signal, (b) an ideal, inverted, integrated trigger signal, and (c) a hypothetical, sampled MR signal are shown. The black dots on the integrated trigger and MR signals indicate the time points at which these two waveforms would be sampled.

integrated trigger signal must be found. The equation for the integrated trigger signal, itself, may be written in two parts (VII.2). The first part corresponds to the linear ramp portion of the signal (which occurs during the acquisition trigger in the time interval denoted by $t_i \leq t \leq t_f$ in Fig. 72) while the second part describes the integrated trigger signal at all other times. This first part of the function describes a line with slope, m, which varies as a function time, t. The term v_0 in (VII.2) is used to describe the presence of a possible a dc voltage offset in the integrated trigger signal. (This dc offset would be due to the output stage of the integrator circuit used to create the ramp signal; however, in Fig. 72 the v_0 term has a value of zero.)

Once the equation for the integrated trigger signal (VII.2) and the voltages of two points on the linear ramp are known, the acquisition time delay may be calculated. (The first two samples of the linear ramp in Fig. 72 are taken at times $t = t_1$ and $t = t_2$ to yield voltages $v(t_1)$ and $v(t_2)$, respectively. The exact values of the sample times, t_1 and t_2, are not known; however, the samples are spaced by an amount of time equal to the sampling period, T_s. Thus, the value of t_2 is known relative to t_1 (VII.3).) This calculation is a three step process. The first step in finding the acquisition time delay is to find the slope, m, of the ramp function. The value of the slope of the line between the first two sample points may be calculated from (VII.4). The second step is to determine the value of the dc offset term, v_0. In general, the value of the dc offset will not be zero (as in the case of the ideal integrated trigger waveform of Fig. 72), and must be found from the samples of the integrated trigger signal that are taken after the acquisition trigger signal has ended (i.e., when the value of the integrated trigger signal is equal to v_0). This

usually involves averaging a large number of these sample points in order to remove the effects of white noise. (There will be thermal noise present since the integrated trigger signal must be produced by an analog circuit composed of real components.) The final step is to solve (VII.2), at time $t=t_1$, for the value of the acquisition time delay (VII.5).

$$v(t) = \begin{cases} m \cdot (t - t_i) + v_0, & t_i \le t \le t_f \\ v_0, & \text{otherwise} \end{cases} \tag{VII.2}$$

$$t_2 = t_1 + T_s \tag{VII.3}$$

$$m = \frac{\Delta v(t)}{\Delta t} = \frac{v(t_2) - v(t_1)}{t_2 - t_1} = \frac{v(t_2) - v(t_1)}{T_s} \tag{VII.4}$$

$$Acquisition\ Time\ Delay = t_1 - t_i = \frac{v(t_1) - v_0}{m} \tag{VII.5}$$

VII.3.2 Acquisition Trigger Integration Circuit

The Fourier time-shifting synchronization method requires that the acquisition trigger signal from the MR system be integrated (see section VII.2.1). This may be accomplished through the use of an integrating amplifier circuit. The integrating amplifier circuit that was used in an early version of the 64-channel receiver prototype appears in Fig. 73. The circuit is actually an inverting op-amp integrator with a TTL-level buffer amplifier on its front-end. The TTL-level buffer, an SN74F244 fast digital buffer chip (Texas Instruments, Dallas, TX), prevents the integration circuit from overloading the output drive stage on the MR system that is responsible for generating the acquisition trigger pulse. The op-amp integrator circuit consists of an LT1797 fast op-amp chip (Linear Technology, Milpitas, CA) with a capacitor in its feed-back loop. There are two, 0-2 kΩ, potentiometer voltage dividers in the op-amp integrator

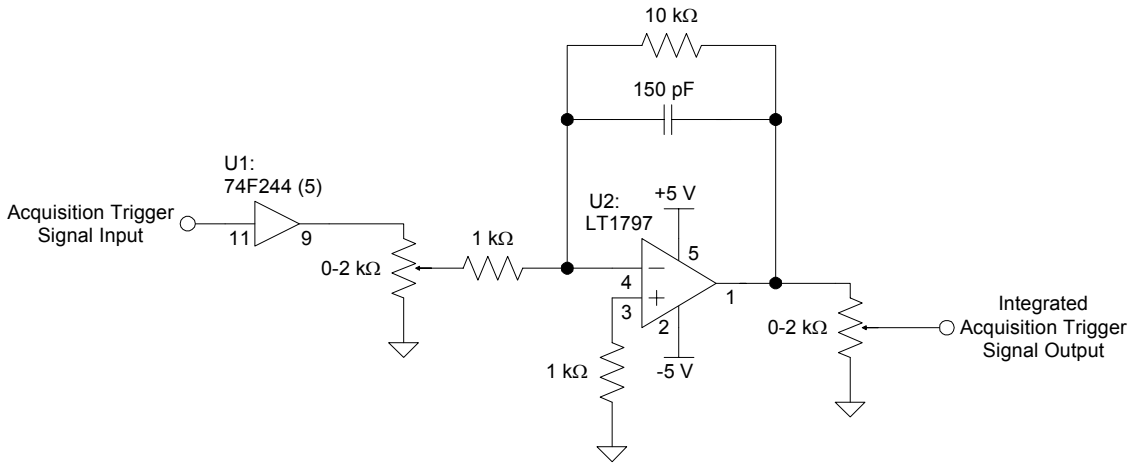

Fig. 73. Schematic diagram for the acquisition trigger integration circuit. The circuit consists of a TTL-level buffer followed by an inverting op-amp integrator. The voltage level of the output waveform may be adjusted by the two potentiometer voltage dividers in the op-amp subcircuit.

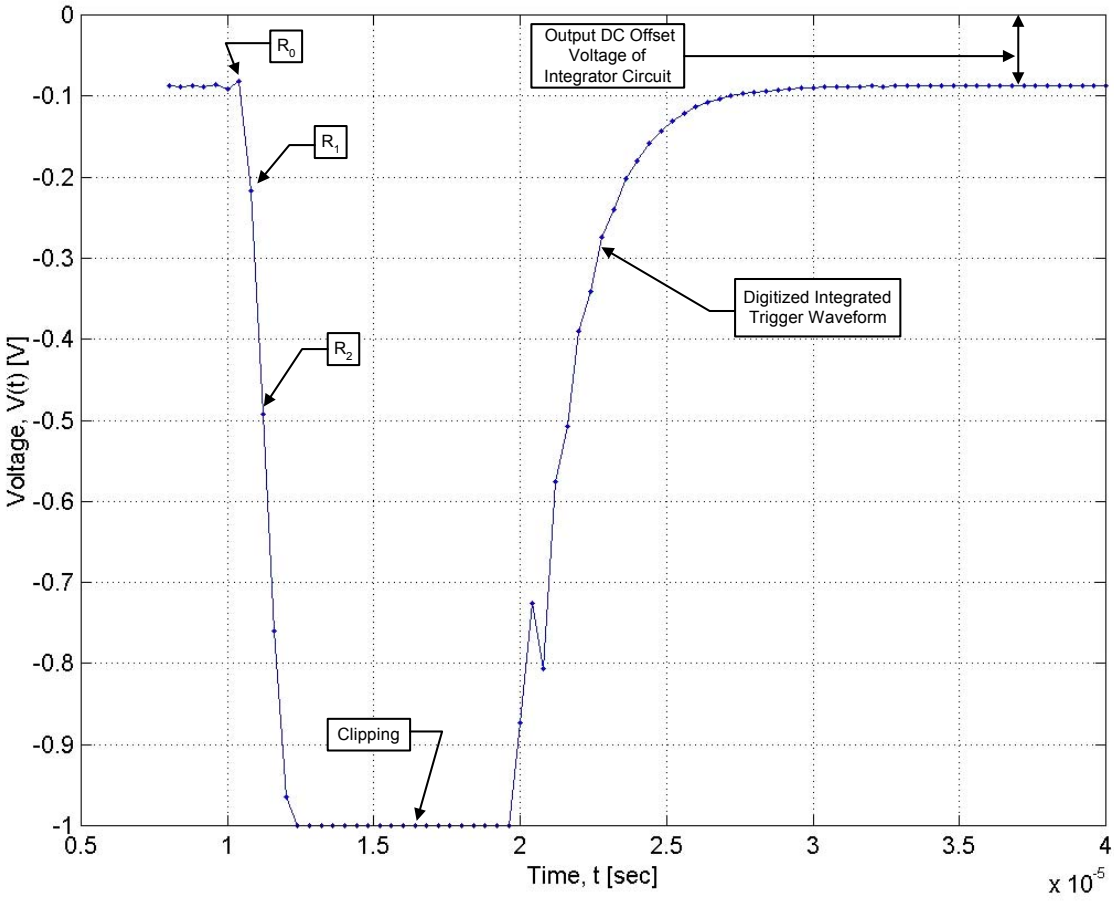

Fig. 74. Plot of one of the actual, digitized, integrated acquisition trigger signals used to implement digitizer synchronization by Fourier time-shifting on the 64-channel receiver. The R_0, R_1, and R_2 points (mentioned in section VII.2.3) are labeled, as are the DC offset of the integrator circuit and the clipping of the integrated waveform by the digitizer.

subcircuit which allow the voltage level of the integrated output waveform to be adjusted.

One of the acquisition trigger signals integrated by the circuit appears in Fig. 74. This signal was acquired using the digitizer on the 64-channel receiver and shows the result of integrating a 10 μsec, TTL-level, acquisition trigger pulse from a GE Omega 4.7 T MR scanner system (GE Medical Systems, Milwaukee, WI). The integrated trigger signal has several non-ideal characteristics. First, the dc offset of the integrator circuit (i.e., the v_0 term in (VII.2)) is non-zero. Second, the integrated signal voltage does not immediately return to the dc offset voltage level after the end of the acquisition trigger signal. This is due to a non-zero time decay on the parallel R-C network in the feedback loop of the op-amp integrator. Another non-ideal characteristic in the integrated trigger signal is the presence of noise. Finally, the digitized integrated trigger signal is also clipped to a minimum voltage of $^{-}1$ V. This clipping is present because the digitizers used on the 64-channel system allow a maximum input voltage swing of $^{\pm}1$ V while the output of the integrator circuit was adjusted to produce a voltage of $^{-}4$ V at the end of the acquisition trigger pulse.

VII.3.3 Fourier Time-Shifting Algorithm

When the Fourier time-shifting method was used to synchronize the digitizer on the 64-channel receiver, a slightly modified algorithm was needed. This was due to the non-ideal characteristics of the integrated trigger pulse (described in section VII.2.2). The actual integrated trigger signal had both a non-zero dc offset and contained thermal noise. In addition, the digitizer used on the 64-channel prototype always captured a

variable number of points prior to the acquisition trigger signal. The modified algorithm for performing synchronization via Fourier time-shifting accounts for these factors.

A hypothetical "real" integrated trigger pulse has been diagrammed (Fig. 75) to show its non-ideal characteristics. In the diagram, the integrated trigger pulse is shown as the time function, $R(t)$. Three sample points are indicated on the integrated trigger waveform at the voltage levels: R_0, R_1, and R_2. These points correspond to the last sample point acquired before and the first two points acquired immediately after the edge of the acquisition trigger signal. The time spacing between the sample points, T_s, and the duration of the acquisition trigger signal, T_D, are also shown. (Note: The acquisition trigger signal edge occurs at time $t = 0$ in Fig. 75. Also, the variables R_0, R_1, R_2, and T_s correspond to v_0, $v(t_1)$, $v(t_2)$, and T_s in (VII.2) through (VII.5).)

The algorithm for performing the Fourier time-shifting technique on the data from the 64-channel receiver is depicted in the flow chart of Fig. 76. The algorithm consists of three main parts: data input and parameter initialization, time delay determination, and phase correction. In the first portion of the algorithm, both the integrated acquisition trigger and the received MR signal data are read into memory. In the second part of the algorithm, the acquisition time delay for each acquired MR signal is determined. In this section, the parameters: R_0, R_1, R_2, and T_s are used to find the time delay for each received MR signal via (VII.2) through (VII.5). Due to the presence of thermal noise in the integrated trigger signal, however, the R_0 value must be determined by averaging the last half of each integrated trigger signal. Once the acquisition time delays are found, the phase of each received MR signal is corrected in the final part of the algorithm. This

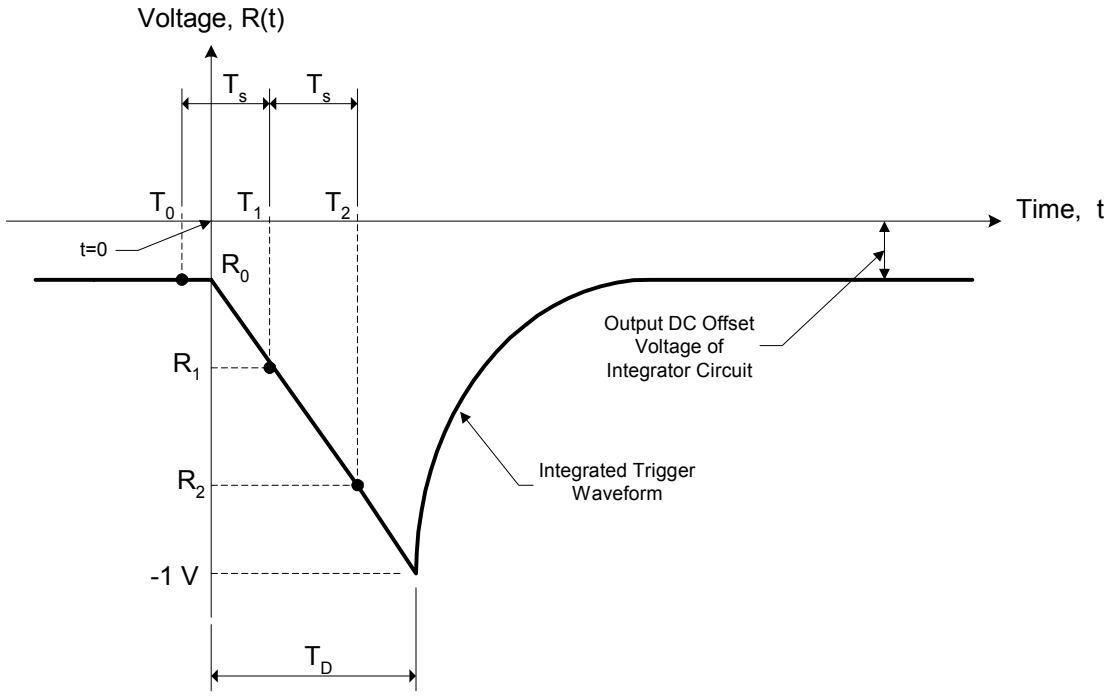

Fig. 75. Diagram of the important parameters for the Fourier time-shifting synchronization algorithm. The important points on the integrated trigger signal (i.e., R_0, R_1, and R_2) are shown on a hypothetical "real" integrated acqusition trigger waveform. The integrated trigger shown has a non-zero dc offset due to the integration circuit. The duration of the trigger pulse, T_D, is also indicated.

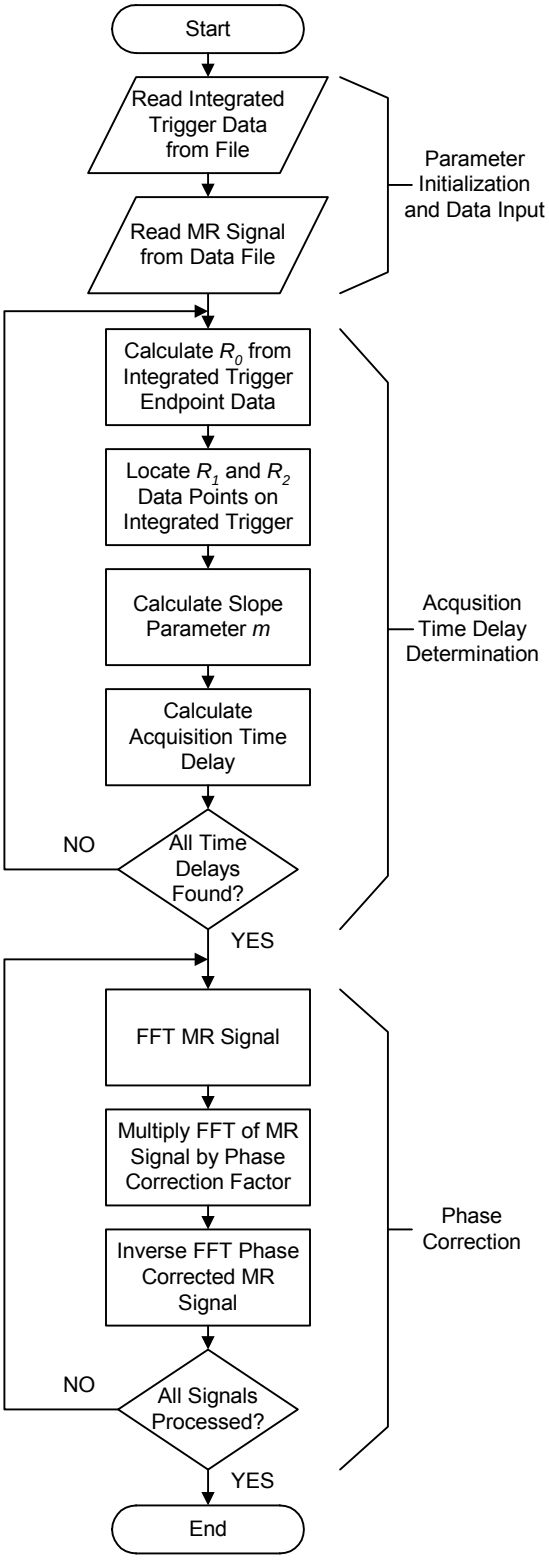

Fig. 76. Flow chart for the Fourier time-shifting synchronization algorithm.

involves performing a Fourier Transform on each MR signal, multiplying the Fourier transformed data by a phase correction vector, and then inverse Fourier transforming the phase corrected data. All of the Fourier transforms required in the algorithm are carried out via Fast Fourier Transform (FFT) techniques.

VII.3.4 Fourier Time-Shifting Synchronization on the 64-Channel Receiver

The Fourier time-shifting algorithm was used to synchronize the 64-channel receiver prototype with the main MR system. Standard spin echo images (TE / TR = 28 ms / 300 ms, 20 kHz spectral width, 1 average, 256 ×128 pixels, 10 cm FOV) were acquired of a cylindrical test phantom on a GE Omega 4.7 T MR system using a single channel of the 64-channel receiver prototype. The acquisition trigger signal from the GE Omega system was integrated (using the integration circuit of section VII.2.2) and then digitized on channel 1 of the 64-channel system's digitizer. The MR signal data was collected simultaneously on channel 2 of the receiver. The MR data set was then synchronized via the Fourier time-shifting technique (using the algorithm of section VII.2.3 and the code of section VII.2.4) and reconstructed to yield the final image. The results of this synchronization operation appear in Fig. 77. The MR image data is presented both before (Fig. 77a) and after (Fig. 77b) synchronization to show the phase jitter correction. The phase jitter was lowered significantly, but not entirely eliminated.

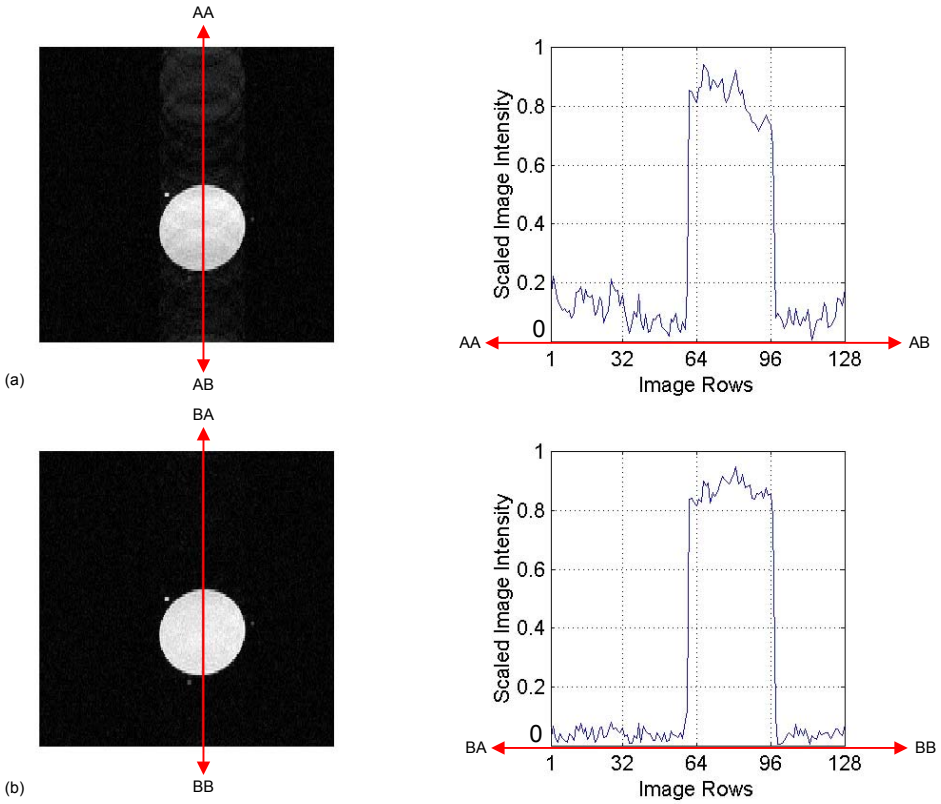

Fig. 77. Verification of the Fourier time-shifting synchronization technique. Integrated acquisition trigger (channel 1) and MR signals (channel 2) were collected simultaneously on the 64-channel receiver prototype. (a) An MR image collected with 64-channel receiver prior to synchronization. A profile of this image (along the line segment AA-AB) is shown in the plot on the right. Note the phase jitter along the line segment AA-AB. (b) The same MR image after synchronization. A profile through the image (along line BA-BB) is shown to the right of the image. Note the reduction in phase jitter.

VII.3.5 Limitations of the Fourier Time-Shifting Technique

The Fourier time-shifting synchronization technique has several limitations. First, the technique is limited by the performance of the integrator circuit. As may be seen in section VII.2.5, the Fourier time-shifting method will reduce (often significantly) the amount of phase jitter, but may not necessarily eliminate it. Noise and any other non-linear performance characteristics present in either the trigger integration circuit or the digitizer will affect the efficiency of the time-shifting synchronization method. Second, the Fourier time-shifting method requires an extra digitizer channel for recording the integrated trigger signal. In addition, the digitizer used must be capable of simultaneous sampling of multiple channels. In the case of the 64-channel receiver prototype, use of Fourier time-shifting limited the system to only sixty-three usable MR data channels (since channel 1 on the digitizer had to be dedicated to collecting the integrated trigger signal). A final limitation of the time-shifting technique is that the synchronization algorithm increases the data processing time. This leads to a decrease in the speed of image reconstruction.

VII.4 Synchronous Clock Synthesis

In this technique, an external, synchronous sampling clock signal is generated to replace the asynchronous sampling clock on-board a digitizer card. The new sampling clock signal must be synchronized to the master system clock on the main MR system. The synchronization operation is performed with a phase-locked loop (PLL) synthesizer circuit and by placing restraints on the sampling clock frequencies allowed on the digitizer. The synchronous clock synthesis technique was used in the final design of the

64-channel receiver prototype to align the sample clock on its digitizers with the main system clock on the MR scanner. The basic operating theory of the clock synthesis synchronization method, details of its use in the 64-channel receiver system, and a general discussion of the limitations of the method will be given in the following sections.

VII.4.1 Clock Synthesis Theory

The clock synthesis technique seeks to synchronize the digitizer not to the acquisition trigger signal but to the underlying master system clock on the main MR system. This is because the master system clock controls when every event in an MR pulse sequence occurs. On an MR scanner, the pulse sequence information is downloaded in digital form to a pulse programming module from which it is clocked out (at the master system clock frequency) when the sequence is run. As a result, all the events contained in a pulse sequence are synchronized to the master clock signal (i.e., usually to rising or falling edge transitions). If a digitizer's own sampling clock is synchronized to the master system clock on the MR system, a fixed length acquisition time delay (to within the jitter level of the external trigger circuit on the digitizer) can be guaranteed for all subsequent signal acquisitions by that digitizer.

Two components are necessary for performing synchronization by clock synthesis: a digitizer capable of accepting an external sampling clock signal and a synchronous clock signal. The first requirement is not met by every commercially available digitizer card. In order for this technique to work, the digitizer to be synchronized must allow its on-board sampling clock to be replaced with an externally applied signal (i.e., the

synchronous clock signal). The clock signal itself must be synchronized in both phase and frequency to the master clock on the MR scanner. This means that, in addition to being phase-locked to the master clock signal, the synthesized clock signal must have a frequency that is an integer multiple of the frequency of the master system clock. To generate this synchronous clock signal, a programmable clock synthesis circuit must be used.

The synchronous clock synthesis circuit is used to generate a sampling clock signal that is synchronous in both phase and frequency with the master clock signal on the MR scanner. The circuit itself can be formed by using an analog VCO (voltage-controlled oscillator) in combination with a programmable PLL (phase-locked loop) circuit. In this case, the PLL unit would be used to lock the phase of the clock signal generated by the VCO to that of the master system clock. Alternatively, the synchronous clock circuit may be created using a DDS (direct digital synthesis) synthesizer circuit, as long as the master clock from the MR scanner is used as the primary clock for the DDS circuit.

VII.4.2 Digitizer Synchronization on the 64-Channel Receiver via Clock Synthesis

The clock synthesis technique was used to synchronize the final version of the 64-channel receiver prototype with the main MR system. The synchronous clock signal was generated by the 64-channel receiver's external clock synthesis module. Standard spin echo images (TE / TR = 28 ms / 300 ms, 20 kHz spectral width, 1 average, 256 ×128 pixels, 10 cm FOV) were then acquired of a cylindrical test phantom on a GE Omega 4.7 T MR system using a single channel of the 64-channel receiver prototype. The results of the synchronization operation appear in Fig. 78. The

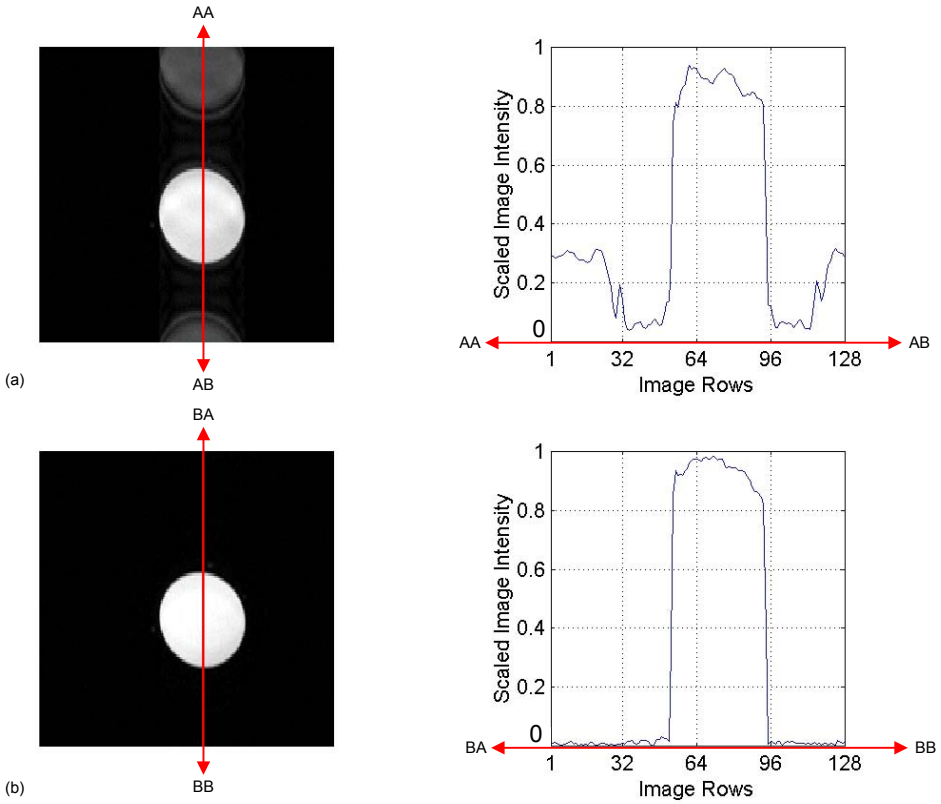

Fig. 78. Verification of the clock synthesis synchronization technique on the 64-channel receiver prototype. (a) An MR image collected with 64-channel receiver prior to synchronization. A profile of this image (along the line segment AA-AB) is shown in the plot on the right. Note the phase jitter along the line segment AA-AB. (b) The same MR image after synchronization. A profile through the image (along line BA-BB) is shown to the right of the image. Note the nearly complete reduction in phase jitter.

MR image data is presented both before (Fig. 78a) and after (Fig. 78b) synchronization to show the phase jitter correction. The phase jitter has been almost entirely eliminated from the image.

VII.4.3 Clock Synthesis Limitations

The clock synthesis synchronization method has two main limitations. First, the technique will only work on certain digitizer cards. This is due to the fact that the digitizer must allow its on-board, asynchronous, sampling clock to be replaced with an externally generated, synchronous, sampling clock signal (section VII.4.2). Many digitizers (especially the less expensive ones) do not have this feature. The second limitation is that the technique allows the use of only certain sampling clock frequencies on the digitizer. As mentioned above (section VII.4.2), the synthesized clock frequency must be an integer multiple of the master system clock frequency. For instance, if the master system clock has a frequency of 10 MHz (this is a common system clock frequency for commercial MR systems), then the only allowable synchronous sampling clock frequencies are integer multiples of 10 MHz (i.e., 10 MHz, 20 MHz, 30 MHz, 40 MHz, etc.). Thus, if other sampling frequencies are desired, they can only be created in post-processing by decimation of the originally acquired signal. Both of these factors limit the applicability of the clock synthesis synchronization technique.

VII.5 Pre-Sampling

Pre-sampling is both the most widely applicable and the most complex of the four synchronization methods. In this technique, the signal to be acquired is pre-sampled with an external sample-and-hold (or track-and-hold) amplifier circuit prior to being sampled by the digitizer. This allows the signal to be sampled synchronously before it is acquired with the asynchronous sampling digitizer. The pre-sampling technique has been used successfully in the desktop MR imaging system prototype (refer to Chapter III, section III.1.2.2.7). This section will cover the theory and limitations of the pre-sampling synchronization technique.

VII.5.1 Pre-Sampling Theory

Pre-sampling allows a signal to be sampled synchronously before it reaches a digitizer which has an asynchronous sampling clock. The pre-sampling operation may be explained using an ideal sample-and-hold circuit (Fig. 79). In this case, the analog waveform to be acquired would first be input to an ideal sample-and-hold circuit. A synchronous sampling clock signal would then be applied to the control input on the sample-and-hold circuit. This would cause the ideal sample-and-hold circuit to instantaneously sample the voltage of the analog waveform at each sampling instant and to then hold that voltage level until the next sample was taken. The result of this sample-and-hold operation is a "stairstep" waveform (Fig. 79) that has a constant value for the duration of each sampling period. The digitizer may then sample the "stairstep" waveform at any time within each sampling period without losing synchronization to the pulse sequence on the MR scanner. As long as the synchronous clock signal (which is

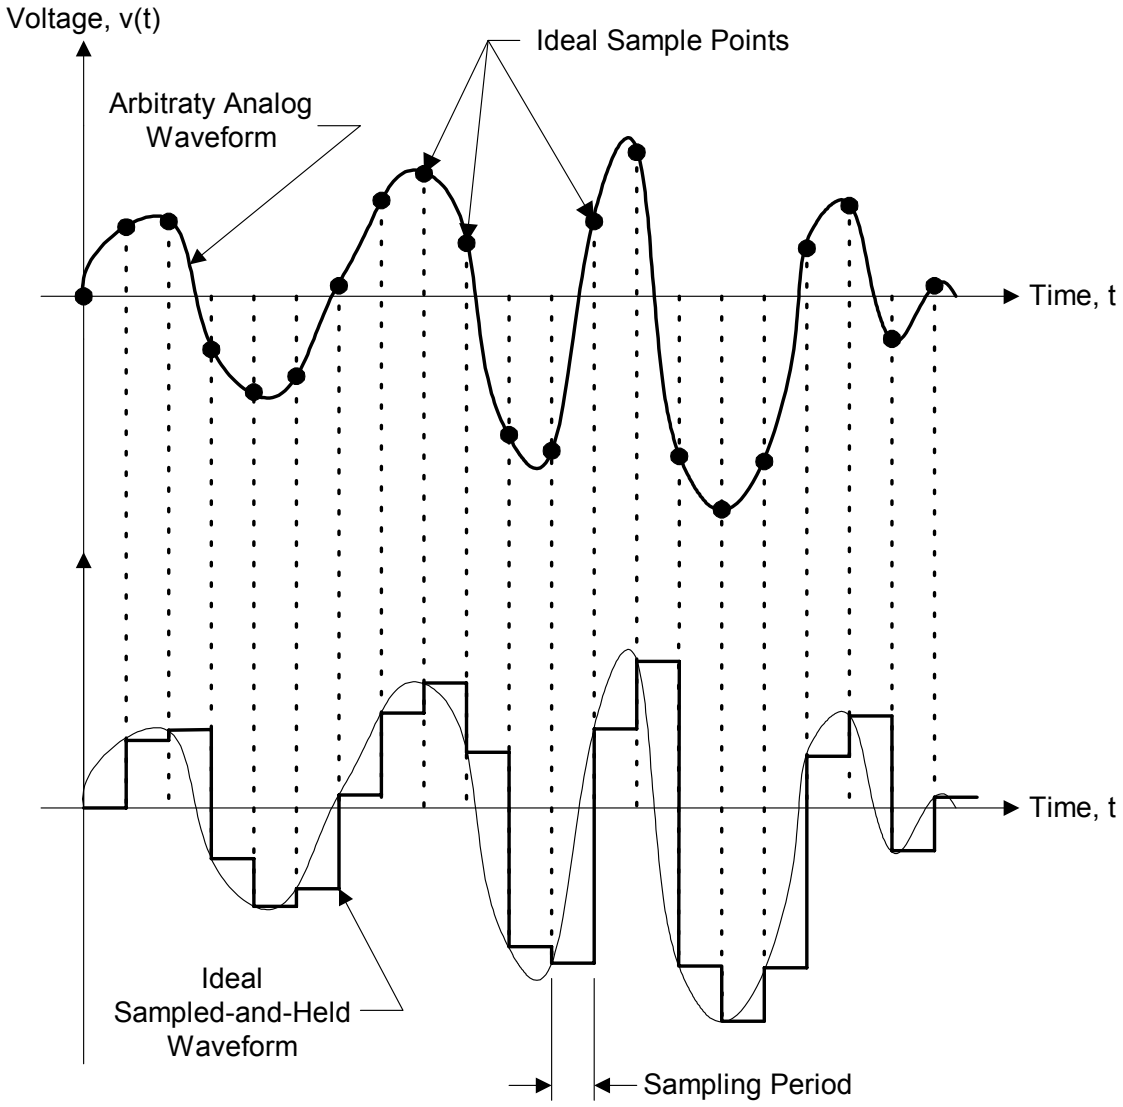

Fig. 79. The input and output waveforms for an ideal sample/ track and hold circuit. The top waveform depicts an arbitraty analog waveform to be sampled. The bottom waveform is the "stairstep" output waveform of an ideal sample and hold circuit. (The dots on the analog waveform indicate the time points at which the waveform was sampled.)

used to control the time at which each sample is taken by the ideal sample-and-hold circuit) always maintains a constant acquisition time delay, this pre-sampling operation will prevent an occurrence of the phase jitter artifact.

Real sample-and-hold circuits operate differently from the ideal sample-and-hold that was discussed above. A real sample-and-hold (or track-and-hold) amplifier circuit has two modes of operation: "sample" (or "track") mode and "hold" mode. In sample mode, the circuit acts as a non-inverting buffer amplifier or a unity gain follower circuit (i.e., the output voltage of the circuit "tracks" the input voltage). Thus, any signal applied to the input of the circuit during track mode will be copied to the circuit's output, as long as the slew rate of the amplifier is not exceeded. In hold mode, the circuit maintains at its output the last voltage level that occurred at the input of the device before that device entered hold mode. Now, in a real sample-and-hold circuit, the transition from hold to sample modes takes a finite amount of time, as does the transition from sample to hold modes. In addition, these sample-and-hold devices have a non-zero signal acquisition time (i.e., the time it takes for the device to switch from hold mode back to track mode and acquire the signal at the input of the device). This means that a real sample-and-hold circuit, unlike its ideal counterpart, cannot instantaneously acquire a sample from an analog input waveform. Consequently, there is length of time at the start of each sampling period during which the digitizer cannot be allowed to sample the output of the sample-and-hold network.

There are three components required to perform synchronization by pre-sampling: a sample-and-hold (or track-and-hold) circuit, a time delay network, and a sampling clock

signal. As described above, a real sample-and-hold (or track-and-hold) circuit is used to sample the analog waveform before it reaches the digitizer. The time delay network is used to ensure that the digitizer does not attempt to sample the output of the sample-and-hold network before the analog signal has been fully acquired. Finally, the sampling clock signal controls when the sample-and-hold (or track-and-hold) circuit samples the incoming analog waveform (i.e., the transitions between the sample and the hold modes). As such, this clock signal must be synchronized to the pulse sequence on the MR scanner. This clock synchronization may be either to the acquisition trigger signal from the host MR scanner or to the master clock on that MR scanner (as in clock synthesis technique in VII.4).

The implementation of the pre-sampling technique differs depending upon whether the digitizer is to be synchronized to the acquisition trigger signal or to the master system clock. If the digitizer is to be synchronized to the acquisition trigger signal, then the implementation of Fig. 80 will be used. In this case, the synchronous sample clock is a finite length pulse train generated by the clock synthesis module (i.e., for an acquisition requiring N samples with a sampling period of T_s, N pulses would be generated with a spacing of T_s between pulses). The clock synthesis module itself could be realized by a PC-based timing card. The clock synthesis block would be programmed to generate a pulse train every time it received the acquisition trigger signal in order to ensure a constant acquisition time delay. The acquisition trigger signal is delayed by the time delay network in order to prevent it from triggering the digitizer before the first sample of the analog waveform has been acquired by the sample-and-hold network. If,

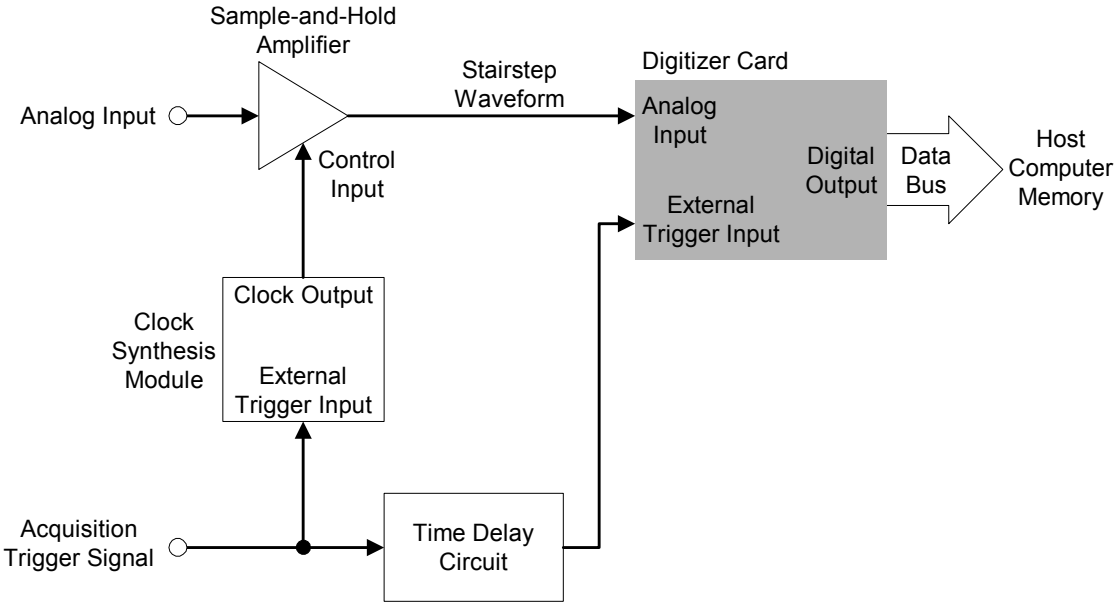

Fig. 80. Implementation of pre-sampling using synchronization to the acquisition trigger signal from the MR scanner. Note the presence of the sample-and-hold circuit, the synchronous sampling clock, and the time delay network.

however, the digitizer is synchronized to the master clock of the MR system, then the implementation shown in Fig. 81 would be used. For this implementation, the synchronous sampling clock signal must be generated by a separate clock synthesizer and be synchronized to the master system clock on the MR scanner in both phase and frequency (refer to section VII.4.1 for details). The synchronous clock is then delayed by the time delay circuit before it is used to replace the sampling clock on the digitizer card.

Finally, the pre-sampling technique may be used to provide an additional benefit aside from synchronization. The sample-and-hold (or track-and-hold) amplifier may be used to increase the operating bandwidth of the digitizer card. Few off-the-shelf digitizers have 16-bit ADC chips with front-end bandwidths higher than 10 MHz; however, an external sample-and-hold circuit with a high front-end bandwidth may be used in front of the digitizer to solve this problem. (This has been previously discussed in Chapter II, section II.3.6.1.)

VII.5.2 Pre-Sampling Limitations

The pre-sampling method (in particular synchronization to the acquisition trigger signal) is the most widely applicable synchronization technique. This method should be able to correct phase jitter due to variance in the acquisition delay time in nearly all instances. However, the pre-sampling method has one major drawback: the complexity of its implementation. Pre-sampling requires more circuitry than any of the other synchronization methods. In particular, pre-sampling requires a sample-and-hold circuit,

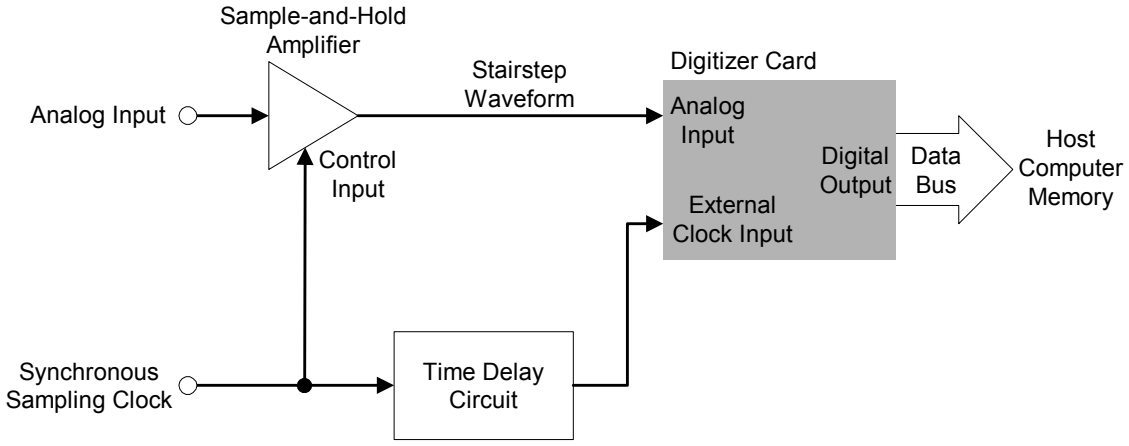

Fig. 81. Implementation of pre-sampling using synchronization to the master clock on the MR scanner. Note the presence of the sample-and-hold circuit, the synchronous sampling clock, and the time delay network.

a synchronous clock synthesizer, and a time delay network. An additional criticism pertains only to pre-sampling by synchronizing to the master clock on the MR scanner. This version of the pre-sampling technique suffers from all the weaknesses of the clock synthesis synchronization method (see section VII.4.3). The only possible reason for using this version of the pre-sampling technique would be to expand the operating bandwidth of a digitizer that has already been synchronized via the clock synthesis method.

CHAPTER VIII

CONCLUSION AND FUTURE WORK

This chapter will examine both the impact and the future direction of the work on instrumentation for parallel MR imaging undertaken within this dissertation. First, the significance of the prototype desktop MR scanner and 64-channel MR receiver will be stated. This will be followed by a discussion of future areas of improvement for both of these prototype systems.

VIII.1 Conclusions

Two main efforts for advancing the technology for parallel MR imaging have been undertaken within this dissertation. The first approach involved the parallelization of the MR experiment through the use of multiple, low-cost MR scanners. As such, a prototype desktop MR imaging system was developed. The second approach focused on the more traditional method of parallel MR imaging by using arrays of sensor coils to increase the speed of the MR experiment. This second effort resulted in the creation of a prototype 64-channel receiver system. Both of these efforts have yielded significant advances in the instrumentation available for parallel MR imaging.

VIII.1.1 Desktop MR Imaging System

The desktop MR scanner prototype shows that a compact MR system for imaging small sample volumes at low field strengths can be constructed inexpensively. This prototype scanner has performance comparable to that of a commercial MR system (operating at low field strengths on small volume samples), but costs thousands of dollars less. This brings about the possibility of expanding the use of MRI outside the

bounds of hospitals and dedicated laboratories to which it has thus far been limited due to the high cost and complexity of MRI hardware. In fact, the desktop scanner prototype may pave the way for a new class of MRI-based, general use laboratory instrument. Such an instrument would be ideal for non-clinical imaging of plant and animal samples in research laboratories and educational settings.

For example, several dedicated, inexpensive, desktop MR scanners could be placed within a small animal research colony, enabling the routine monitoring of diverse animal pathologies at reasonable cost. One such application would be in phenotyping the mouse genome [6]. Alternatively, using the technology developed in creating the desktop MR scanner, multiple independent MR scanner modules (complete with separate and independently controlled gradient and RF coils) could be placed inside a whole-body magnet to obtain scans of multiple samples simultaneously. Another group has already demonstrated this technology [4]. This is a form of parallel MR imaging in which the act of imaging is "parallelized" by allowing multiple samples to be scanned simultaneously via different MR scanners.

VIII.1.2 64-Channel Parallel MR Receiver System

The 64-channel parallel receiver prototype is the first ever receiver for parallel MR imaging capable of simultaneous reception of up to sixty-four, 1 MHz bandwidth signals. The receiver system was also designed to be scalable in 32-channel groups, facilitating expansion of the prototype system into an even higher channel count receiver. Prior to its reporting at the tenth annual scientific meeting of the International Society of Magnetic Resonance in Medicine (ISMRM) in 2002 [50], no receiver for

parallel MR imaging had ever been constructed with more than sixteen channels. In fact, the only 16-channel receiver up to that point was developed in the Magnetic Resonance Systems Laboratory (MRSL) at Texas A&M University by Dr. Wright and his students [16]. That receiver was a TDM parallel receiver with a channel bandwidth limitation of approximately 12.5 kHz per channel. The remaining commercial MR receivers at that time were all limited to no more than eight channels. Additionally, the cost of the prototype system was very low for a parallel MR receiver. The cost of the 64-channel prototype was approximately $33,500 (i.e., about $525 per channel) for parts and materials. Thus, the 64-channel receiver prototype represents a significant leap forward in parallel imaging receiver technology in terms of channel count, channel bandwidth, and overall system cost.

It should be noted that when the 64-channel receiver prototype was first reported in May of 2002, no commercial MR manufacturer had yet provided a system with more than eight receiver channels. Since that time, however, three of the major commercial MRI system manufacturers have either announced systems with thirty-two receiver channels or plans to produce them. At the eleventh annual meeting of the ISMRM in 2003, it was reported that GE Medical Systems had produced a special purpose, MR research system with a 32-channel receiver [51]. This special purpose system was formed by combining four, standard, 8-channel GE phased array receiver systems and occupied a small room (much larger than our 64-channel prototype system which fits inside a 89 cm (35 in.) tall, 19-inch shock-mount rack). Later, in November of 2003, Siemens Medical Solutions (Malvern, PA) described their new Total Imaging Matrix

(TIM) MR imaging system which featured up to thirty-two independent receiver channels [52, 53]. The first TIM capable system was installed in November of 2003 at the NYU Medical Center in New York with another slated for installation in Tubingen, Germany shortly thereafter [53, 54]. Also in November of 2003, Philips Medical Systems (Andover, MA) announced their new Achieva MR imaging platform which features Freewave, a scalable, 32-channel, 3 MHz bandwidth, direct digital sampling receiver architecture [55]. Initially, the Achieva was only available with eight to sixteen receiver channels, but the platform supports future upgrades to a total of thirty-two channels [56]. Thus, since the introduction of the 64-channel prototype, many of the commercial MR system manufacturers have moved to offer receivers with up to thirty-two channels, but MR systems with more than eight receiver channels remain scarce.

In addition to being a major advancement in parallel MR receiver technology, the 64-channel parallel MR receiver has also paved the way for a new parallel imaging modality. This modality is known as single echo acquisition (SEA) imaging and was developed by Dr. Wright and his students in the MRSL at Texas A&M University [57]. By means of SEA imaging, a custom, 64-element, planar RF coil array was used to collect a complete MR image (64 × 256 pixels) without gradient phase encoding. This image was collected using a single echo signal in approximately 300 ms while a comparable phase-encoded image would require an acquisition time of 19.2 s (this represents a factor of 64 in time savings). The SEA imaging technique has the potential of enabling dynamic MR imaging with extremely high frame rates [58]. In fact, frame rates as high as 125 frames per second have already been reported [59].

VIII.2 Future Work

Both the desktop MR system and the 64-channel parallel MR receiver prototypes represent significant advances in parallel MR imaging instrumentation technology. However, neither of these systems is ready for commercialization. In fact, each prototype system has several areas which need refinement before either system would make a viable commercial product. This section will explore the possible design improvements for both systems that would make a next generation version ready for debut as a commercial product.

VIII.2.1 Desktop MR System

The desktop MR system would profit from modifications in the many areas. These include the magnet subsystem, operating frequency range, receiver bandwidth, the control interface, and the system size, weight, and cost. Possible revisions to the design of the desktop MR system in each of these areas will be discussed in the next sections.

VIII.2.1.1 Magnet Subsystem

The area of the desktop system design that needs the most revision is the magnet and gradient subsystem. The desktop permanent magnet does not possess the homogeneity required to image 2.54 cm^3 (1 $in.^3$) volumes. The homogeneous region of the magnet is consists roughly of a spherical region with a diameter of 0.64 cm (0.25 in.). This was not discovered until work on the prototype system was well underway, so it was decided to continue the project using the existing magnet. As a result, the static field had to shimmed with a set of active shims (see Chapter III, section III.2.2.3). These shims generate heat, which warms the magnet pole faces, causing the metal to expand and the

magnet field strength to drift at a rate greater than that of the ambient temperature. The shims, which work to a degree, are still not able to homogenize the static field enough to allow imaging of 2.54 cm^3 (1 $in.^3$) volumes. As a result, the large gradient pulses must be used to in an attempt to overwhelm the remaining field inhomogeneities during an MR imaging experiment. However, even with the use of active shims and large gradient pulses, the desktop MR system is only able to image 1.9 cm^3 (0.75 $in.^3$) volumes without distortion.

In order to enable imaging of 2.54 cm^3 (1 $in.^3$) samples, the magnet subsystem needs to be redesigned. There are two ways to accomplish this redesign. The first (and best) is to use a different permanent magnet. Several companies construct custom, low-field, permanent magnets, so finding a supplier of a 0.2 T permanent magnet with a homogeneous region larger than 2.54 cm^3 (1 $in.^3$) should not be difficult. This would allow the active shims to either be eliminated or at least operated with much smaller currents (resulting in less pole face heating). In addition, the gradients, which were designed for a system with at least a 2.54 cm^3 (1 $in.^3$) homogeneous region, would not have to work as hard. The second option involves using the existing permanent magnet and redesigning the gradient and shim system. The gradient coils and amplifiers could be redesigned to make them stronger (this would, however, also increase their size and weight). The active shim system would also need modification. The current amplifiers which power the active shims would need to be redesigned to allow them to source more dc current. Finally, the shim coils would need to be remanufactured with smaller gauge

wire. This would both increase the current capacity of the shim coils and reduce the thermal output of the coils.

VIII.2.1.2 Operating Frequency Range

The desktop MR system prototype presently has an operating range which extends from 0-11 MHz (without the mixing module) and from 0-50 MHz and 70-90 MHz (with the external mixing module) (refer to Chapter III, section III.2.2.2). While the initial prototype transceiver was only designed to operate on low-field magnets (less than 0.25 T) at imaging frequencies between 0 and 11 MHz, a commercial version of the system would probably need to operate from at least 0 to 63 MHz. This expansion in the operating frequency range of the desktop system could be achieved by redesigning the frequency synthesis portion of the transceiver (please see Chapter III, section III.2.2.2). This redesign would involve replacing the DDS (direct digital synthesis) circuits currently used with DDS circuits that are able to generate frequencies up to at least 70 MHz. A number of these DDS chips are available from companies such as Analog Devices (Norwood, MA) and Intersil (Irvine, CA).

VIII.2.1.3 Receiver Bandwidth

The present bandwidth of the desktop receiver is limited to 20 kHz by the anti-alias filters which precede the digitizer. This is a rather narrow bandwidth for a commercial product; hence, the receiver bandwidth should be expanded to 1 MHz. This could be accomplished by making three changes in the design of the desktop receiver. The first would be to replace the existing anti-alias filter chips with model MSFHS2P filters from the same manufacturer (Mixed Signal Technology, San Jose, CA). These chips are pin

and function compatible with the existing anti-alias filter chips, but have an electronically tunable cutoff frequency range from 10 kHz to 3 MHz. The second step would be to replace the current digitizer (which has an input bandwidth of 102.4 kHz) with another digitizer capable of simultaneous, 16-bit sampling on two input channels at rates of 2 MSPS or higher. Finally, the bandpass preselect filter on the receiver would have to be replaced with a filter that has at least 1 MHz of bandwidth.

VIII.2.1.4 Slice Offset Imaging

The prototype desktop system has another flaw which would limit its desirability as a commercial product. This is its inability to perform offset slice imaging. At present, the desktop system can perform slice select imaging, but only if the slice of interest happens to be located at the center of the prototype gradient set. To get around this limitation, the prototype desktop system uses 3D MR imaging, but this lengthens the time required to acquire the image data from the slice in question. In order to perform slice offset imaging, the desktop transceiver must be able to apply a phase ramp (i.e., linear phase offset) to the RF transmit carrier waveform during the application of the shaped RF pulse; however, the current frequency synthesis circuitry is not able to perform this task.

There are two ways to enable slice offset imaging on future versions of the desktop system. The first is to use a phase modulation capable DDS circuit in the frequency synthesis portion of the transceiver. The second method is to add a voltage-controlled phase shifter circuit to the transmitter chain before the RF power amplifier. The phase modulator circuit could then be controlled by one of the analog output lines from the analog and digital output unit in the computer control module of the desktop transceiver

(see Chapter III, section III.2.2.1). Mini-Circuits (Brooklyn, NY) has recently come out with a voltage-controlled phase shifter (model JSPHS-12) which would work with the current prototype system.

VIII.2.1.5 System Size, Weight, and Cost

Finally, the desktop system would benefit from further reform in the areas of system size, weight, and cost. The desktop system is already much smaller, lighter, and less expensive than its commercial counterparts; however, a commercially viable version of the desktop MR system should probably be even smaller, lighter, and less expensive in order to make it more portable. To this end, the next two sections will explore ways to further reduce the size and weight of the desktop system.

VIII.2.1.5.1 System Size and Weight

The largest and heaviest item within the desktop prototype is the permanent magnet. Now, as has been pointed out above (section VIII.2.1.1), this is also the portion of the system in need of the most revision. It is highly unlikely that future editions of the permanent magnet (with a larger homogeneous region) will be much smaller or lighter than the one. Thus, the only part of the desktop system which can be revised to reduce the system size and weight is the transceiver system. The two primary methods for reducing the size and weight of the transceiver are by using more modern circuit board fabrication methods and by exchanging the existing control computer with a laptop PC. (Details for each of these methods may be found in sections VIII.2.2.4 and VIII.2.2.6.)

VIII.2.1.5.2 System Cost

The most expensive part of the desktop MR system prototype is the control computer subsystem (refer to Chapter III, section III.3.2). This is primarily due to the high cost of the digitizer and the analog/ digital output cards. Both of these are high performance, PCI-based cards from National Instruments (Austin, TX). With some effort, less expensive replacements for these cards should be able to be found from one of the many PC-card manufacturers, especially since these cards are now nearly seven years old.

VIII.2.2 64-Channel Parallel MR Receiver System

There are seven major areas in which the design of the 64-channel receiver system would benefit from revision. These include the operating frequency range, the noise figure, the dynamic range, interchannel isolation, the gain variation between channels, imaging speed, and the size and weight of the receiver. Possible future revisions in each of these areas will be detailed in the following sections.

VIII.2.2.1 Operating Frequency Range

Presently, the operating frequency range of the 64-channel receiver prototype extends from 50 to 400 MHz (see Chapter VI section VI.2.2). This range is acceptable for the prototype system, but a commercial version of the parallel receiver would benefit from a much wider operating range. The range of the prototype system is mainly limited by three components in the LO distribution network, the low-noise preamplifier circuits, and the mixers used in each 4-channel RF to IF conversion submodule. These limited bandwidth components are listed in Table VI. The remaining RF components in the 64-channel prototype all have operating ranges that span the range from dc to 1 GHz. The

operating range of the entire 64-channel system may be increased to extend from dc to 1 GHz as well by exchanging the components in Table VI with wider bandwidth versions. Suitable replacement parts are available from several manufacturers (such as Mini-Circuits), but these new parts will drive the cost of the prototype system higher.

Table VI
64-Channel Prototype Receiver Operating Frequency Limiting Components

Component	Description	Operating Range [MHz]	Location
PTS-250	Frequency Synthesizer	1-250	LO Distribution System
MHL-8115	RF Power Amplifier	50-1000	LO Distribution System
SCP-4-1	4-Way RF Power Splitter	1-400	LO Distribution System
INA-01	Low-noise amplifier	dc-500	64-Channel Preamplifier
SBL-1LH	Analog mixer	2-500	RF to IF Conversion Submodule

VIII.2.2.2 Noise Figure

The noise figure for each channel in the 64-channel system has not been exactly measured, but it has been calculated to be below 2 dB (see Chapter VI section VI.1.3 and Chapter IV section IV.1.3.1). While this noise figure is acceptable for the prototype system, a lower system noise figure would be better for the next generation of the 64-channel system. The noise figure for each channel in the system is set by the low-noise, high gain amplifier stage at the front end of the receiver. Presently, each LNA stage is implemented by a single-chip, 50 Ω, MMIC amplifier with 32 dB of gain and a 1.7 dB noise figure. This amplifier chip was chosen for its relatively high gain, stability over a broad frequency range, and low cost; however, its 1.7 dB noise figure is rather high for most MR imaging applications. In fact, narrowband preamp circuits with noise figures

of 0.3 to 0.5 dB are more commonly used in MR imaging. The easiest way to reduce the noise figure of the 64-channel prototype would be to replace the existing LNA's with narrower band, lower noise figure amplifier circuits. These circuits could be either designed in-house from GaAs-FET chips or obtained from a supplier (such as Advanced Receiver Research, Burlington, CT) for approximately $150 each. In either case, the redesign of the 64-channel preamplifier module would increase the cost of the next generation parallel receiver system.

VIII.2.2.3 Dynamic Range

The single-tone dynamic range of the prototype parallel receiver was measured to be 57.2 dB. For an MR receiver system, this dynamic range is still relatively low, as most commercial systems tend to at least 70 dB of dynamic range. There are two main approaches to increasing the dynamic range in the next generation of the 64-channel receiver. The first is to lower the level of the minimum detectable signal (MDS), and the second is to raise the 1-dB compression point for the receiver. Methods for accomplishing both of these tasks will be discussed in the following paragraphs.

VIII.2.2.3.1 Reduction of the MDS Level

The minimum detectable signal has been defined previously (refer to Chapter II section II.4.2.2) as the sum of the thermal noise floor, the noise figure of the receiver, and a detection threshold. Of these three factors, only the noise figure and the thermal noise floor are of much interest for lowering the MDS level (since the detection threshold is fairly arbitrary for MR imaging). Since reducing the noise figure has already been discussed (section VIII.2.2.2), the only factor left is the thermal noise floor.

Ordinarily, the thermal noise floor cannot be changed for a given bandwidth and temperature (assuming matched loads). However, in the prototype receiver, the thermal noise floor has been raised by approximately 3 dB (a factor of $\sqrt{2}$ in the r.m.s. noise voltage) due to the removal of the image reject filter in the front-end of each receiver channel. This means that the thermal noise floor can be lowered 3 dB by simply inserting an image reject bandpass filter into the receiver chain between the LNA stage and the RF to IF downconversion module on each channel. The image reject filter, however, must pass only the signal band and reject the image frequency band (which happens to lie very close to the signal band in the current implementation of the receiver).

VIII.2.2.3.2 Increasing the 1-dB Compression Point

At present, the 1-dB compression point is limited by the mixer, the four RF amplifier stages, and the one variable attenuator stage in the front-end of each receiver channel (refer to Chapter IV section IV.1.2 for details). The 1-dB compression point for the next generation of the receiver may be raised by changing out these components. Using mixer and RF amplifier components with higher RF power handling capabilities (such as the new ERA series of 50 Ω, MMIC amplifiers from Mini-Circuits) and a variable attenuator with a wider attenuation range will raise the 1-dB compression point of the receiver channels.

VIII.2.2.4 Interchannel Isolation

As mentioned in Chapter VI, the interchannel isolation on the prototype system did not meet its design requirement of 35 dB. This was due to a higher than expected level

of coupling within the 64-channel preamplifier subsystem. The isolation within this module can be increased in one of two ways. The first method would be to add better shielding between the preamp circuits contained within the two, 32-channel preamp modules which comprise the 64-channel preamplifier. A second, and probably better approach, would be to redesign the preamplifier module using commercially available (and individually shielded) preamps (as discussed in section VIII.2.2.2). This second approach would improve both interchannel isolation and the noise figure of the receiver system.

VIII.2.2.5 Channel-to-Channel Gain Variation

As reported in Chapter VI (section VI.2.6), the gain variation across all channels in the 64-channel prototype receiver has a maximum value of 4.2 dB. While this was acceptable in a prototype system, it is a rather large degree of variation for a commercial production parallel receiver. The majority of this gain variation is due to three main factors: manufacturing technique, variation in component characteristics, and limited gain control. Each of these factors and their possible remedies will be discussed in this section.

VIII.2.2.5.1 Manufacturing Technique

Each of the sixty-four, single channel receiver circuits which compose the parallel receiver prototype was constructed in house on double-sided, copper clad, G-10 circuit board material using surface mount components. The circuit boards were drawn up using Protel design software (Altium, Sydney, Australia) and then etched using a C-30 PC-board fabrication machine from LPKF Laser & Electronics (Garbsen, Germany).

Then, the components for these circuits were soldered in place by hand using ultra-fine tipped soldering irons. While the every attempt was made to follow standard high-speed PC board layout practices and care was taken during the assembly process to avoid ESD (electro-static discharge) damage, there were limitations to the quality and repeatability of this manufacturing process. There are several ways to address these problems in future versions of the receiver, such as multi-layer PC boards, automated component placement, and reflow soldering.

The use of multiple layer PC boards would allow a more efficient design of the power distribution circuitry on the 4-channel RF to IF conversion module circuit boards. Instead of having to bring the three dc supply voltages to one side of the PC board and distribute the voltages over jumper wires, power planes sandwiched within the circuit board could be utilized. This would both prevent resistive losses in the dc supply pipeline and help to further reduce RF noise on the dc supply traces. (At present, RF noise is kept off the supply lines using large numbers of shunt capacitors, RF chokes, and twisted pair jumper lines.)

Modern circuit board manufacturing often uses automated component placement. This ensures that the tiny, delicate RF, surface mount components are not damaged during circuit assembly either by physical mishandling or by ESD. The use of such equipment would help to ensure that each of the single-channel circuits in future versions of the parallel receiver would be assembled in a highly repeatable fashion.

Along with automated component placement, most modern circuit manufacturing also uses reflow soldering. Once the components have been fixed in place with adhesive

flux by the component placement apparatus, the entire circuit board is dipped in a solder bath. Solder connections are only made at the locations on the components where the adhesive paste was present. This method of soldering ensures that the components will not be damaged by overheating (as can often happen during hand soldering). Also, there will be no bad connections in the circuit due to "cold" solder joints (another problem common when components are soldered by inexperienced operators).

VIII.2.2.5.2 Component Characteristic Variation

Another problem which leads to gain variation between channels in the parallel receiver is variation in the performance characteristics of the circuit components. Now, when any electronic component is manufactured it will possess a difference in its operating characteristics between its actual performance and that specified for the part by the manufacturer (i.e., its nominal performance). For instance, the actual resistance of a discrete resistor will vary from its nominal resistance value by as much as ±1% to ±10%, depending on its manufacturing tolerance. This holds true for all active and passive electronic parts. Thus, even though each of the single-channel receiver circuits within the 64-channel receiver prototype is composed of components having the same part numbers (i.e., the mixer in each RF to IF conversion circuit has the same model number), the overall performance of each receiver channel will differ from its neighbors by varying amounts. The only practical way to remove channel-to-channel gain variation due to this type of component performance characteristic variation is to independently adjust the variable gain on each channel so that all channels have equal gain. This will be addressed in the next section.

VIII.2.2.5.3 Limited Gain Control

A final cause of the gain variation between channels in the parallel receiver prototype is limited gain control on the receiver channels. The 64-channel receiver's gain control subsystem does not provide independent control of the gain on each of the sixty-four receiver channels. This was due to the added cost and complexity that such a system would have required. Instead, the prototype receiver's gain is really only controllable in groups of four channels (i.e., the user may set the gain for each of the sixteen, 4-channel RF to IF conversion modules). Using the existing gain control system, the channel-to-channel gain variation may be reduced from 4.2 dB to 3.3 dB (see Chapter VI, section VI.2.6). In future versions of the parallel receiver, however, the gain variation between channels may be eliminated entirely if independent control of the gain for each of the sixty-four channels is provided. This would involve a complete redesign of the existing gain control subsystem. The new gain control system could even be automated by using a PC-based digital I/O (input/ output) card and a set of sixty-four or more digital potentiometers. This would allow the user to equalize the gain on the receiver by simply applying the same test signal to each receiver channel and individually adjusting that channel's gain.

VIII.2.2.6 Imaging Speed

In its current implementation, the 64-channel parallel receiver prototype supports only real-time acquisition of the MR signal data from up to 64 sensor elements. All reconstruction and processing of the received image data must be performed offline after the signal data has been stored. This results in a considerable time lag (in some cases

greater than 10 minutes) between data acquisition and actual image viewing. This time delay can limit the utility of the receiver system for monitoring dynamic processes. However, a real-time image reconstruction system for the 64-channel receiver could be developed using commercially available real-time DSP (digital signal processing) cards. This would eliminate the time lag and allow the user to view the parallel MR signal data as it is acquired. One possible incarnation of this real-time reconstruction interface will be described here.

Each of the two ICS-645 digitizers used in the current 64-channel receiver supports 32 channels of simultaneous, 16-bit digitization at 2.5 MSPS. This is a data accumulation rate of 160 MB/s per card (for 32 channels per card). The ICS-645 has an onboard FPDP (Front Panel Data Port) interface which may be used to stream received data off the digitizer at rates of 160 MB/s. Two, TS-P36N-4-B DSP cards (Transtech DSP, Ithaca, NY), with add-on FPDP interface modules could be used to perform real-time reconstruction of the MR signals from the 64-channel receiver. Each DSP card would receive digitized MR signal data from one of the ICS-645 cards in the prototype receiver at the FPDP transfer rate of 160 MB/s. The signal data would then be I/Q demodulated, filtered, and decimated in real-time by software resident on the DSP card. The processed MR signals would be transferred to the host computer for final image display.

Development of a real-time reconstruction system for the 64-channel receiver prototype is currently underway by other graduate students in the MRSL at Texas A&M University. This effort involves four main tasks. First, the digitizers and DSP cards must

be integrated with device drivers into the host computer. Second, the real-time processing software must be generated for the DSP cards. Third, the final image processing and display software must be developed for the host computer. Finally, the real-time processing system must be tested to evaluate its overall performance in terms of achievable frame rates, sustained imaging times for a given frame rate acquisition, and actual time delay between data acquisition and image display versus frame rate.

VIII.2.2.7 Size and Weight

Most of the prototype parallel receiver fits inside of an 89 cm (35 in.) tall, 19-inch shock-mount rack (i.e., the two, 32-channel RF to IF conversion modules and the acquisition control computer which are shown in Chapter V, section V.3). In addition to this, the prototype receiver also consists of a PTS-250 frequency synthesizer, the 64-channel preamplifier module, a 15-in. LCD (liquid crystal display) monitor, a standard mouse, and a standard keyboard which are not housed in the shock-mount rack. While the 64-channel receiver prototype is still considerably smaller than any of the commercial parallel receivers, reducing the size and weight of the receiver would make the system much more commercially attractive. This reduction may be accomplished by modifying the local oscillator distribution, the 32-channel RF to IF conversion modules, and the acquisition control computer systems in future versions of the parallel receiver.

VIII.2.2.7.1 Local Oscillator Distribution System

The local oscillator distribution network may be modified in one significant way to reduce the overall size and weight of the parallel receiver. This modification involves replacing the existing LO source with a smaller alternative. Currently, the LO source is

provided by a PTS-250 frequency synthesizer (Programmed Test Sources, Littleton. MA). The PTS-250 is a 19-in. rack mounted device that occupies 48.2 cm × 13.3 cm × 45.7 cm (19 in. × 5.25 in. × 18 in.) of space and weighs 15.9 kg (35 lbs.). In future versions of the parallel receiver, the PTS-250 could be replaced by a PC-based frequency synthesizer card which could be mounted inside of the acquisition control computer. Several of these cards, which allow their synthesized output signals to be phase-locked to an external reference clock (such as the CG400 or the DA4300 from the Chase Scientific Company in Langley, WA), are currently on the market from various manufacturers.

Another alternative for replacing the LO source in the local oscillator distribution network is constructing a frequency synthesizer in-house. This was actually attempted during the initial development phase of the prototype 64-channel receiver. A stand-alone frequency synthesizer system was built based upon the PCK-240 universal clock module from DST (Asaka City, Japan). The frequency synthesizer circuit took in a 10 MHz external reference clock and generated a phase-locked, output sinusoid with a programmable frequency in the range of 1-240 MHz. However, the PCK-240 clock source was not phase stable enough to act as a replacement LO source to the parallel receiver prototype. Due to time constraints, this part of the project was eventually abandoned.

VIII.2.2.7.2 32-Channel RF to IF Conversion Modules

The size and weight of the parallel receiver may also be reduced by making significant modifications to the design of the 32-channel RF to IF conversion modules.

These modifications involve redesigning both the 4-channel RF to IF submodules and the power distribution network within the 32-channel conversion modules. Some of the details of each redesign will be discussed in the following paragraphs.

Eight, 4-channel RF to IF conversion submodules are currently housed within each 32-channel conversion module in the parallel receiver prototype. Each of these 4-channel submodules is, in turn, housed in a 22.9 cm × 17.8 cm × 5 cm (9 in. × 7 in. × 2 in.) aluminum enclosure which helps both to shield the circuits from external RF noise as well as to reduce crosstalk between the submodules themselves. Unfortunately, this configuration wastes a large amount of space. More modern circuit design and fabrication techniques could be utilized (as discussed in section VIII.2.2.4), however, to reduce the size of the 4-channel RF to IF conversion circuit boards. If these fabrication methods were coupled with individual, integrated RF shields for each single-channel circuit on the 4-channel conversion circuit boards, the aluminum RF shielding enclosures could be eliminated completely. (Affordable, custom, PC board RF shields are available from manufacturers such as Fotofab in Chicago, IL.) This would allow the 4-channel submodule circuit boards to be placed closer together within the 32-channel conversion module's 19-in. rack chassis. With the aluminum enclosures removed, each of the 4-channel conversion submodule circuit boards could also be redesigned to plug into a custom power backplane. This would simplify the portion of the power distribution system which transports dc power to the 4-channel conversion circuits. By redesigning the 4-channel conversion submodule circuits in this way, it might even be possible to fit all sixty-four of the RF to IF conversion circuits into a single 19-in. rack

chassis. Doing so would allow the two, 32-channel RF to IF conversion modules in the current prototype to be replaced with a single, 64-channel conversion module.

The power distribution network inside each 32-channel conversion module currently consists of an ac power entry module, a dc power supply module, a main distribution module, and the cables which transport dc power to each of the modules in the receiver. Redesigning the 4-channel submodules (as discussed above) will necessitate a revision of both the main distribution module and the portion of the power cabling subsystem that is contained within each 32-channel module. The main distribution module must be modified to accommodate power connections to the custom power backplane. Use of a card backplane will also reduce the amount of transport cabling that is currently required to bring the dc supply voltages from the main distribution module to the RF to IF conversion submodules. As a result, the design of the power cabling subsystem will also have to be revised.

VIII.2.2.7.3 Acquisition Control Computer

A third method for decreasing the size and weight of future versions of the parallel receiver is to modify the architecture of the acquisition control computer. Currently, the control computer consists of a rack-mounted, 1.0 GHz Pentium III based CPU (which resides within the 19-in. shock rack), a 15 in. LCD monitor, and a standard PS/2 keyboard and mouse (which do not reside in the shock rack). Two, full-size PCI card digitizers also reside in the PCI chassis of the control computer's CPU. One way to reduce the size of the control computer would be to replace the rack-mounted CPU, the keyboard, mouse, and LCD monitor with a laptop PC. Of course, the digitizer cards

would not fit inside the laptop PC, but they could be installed within an auxiliary PCI chassis that would be mounted inside the 19-in. shock rack. The laptop could then communicate with the cards in the auxiliary PCI chassis via its on-board USB 2.0 interface (this would require a second USB interface card to be mounted in the PCI chassis and a USB cable between the PCI card chassis and the laptop). This would replace the currently bulky LCD monitor, keyboard, and mouse setup with a much more compact laptop computer.

VIII.3 Summary

Within this dissertation, two different methods for advancing the state of the art of instrumentation for parallel MR imaging have been studied. First, a prototype desktop MR imaging system was developed to explore parallelization of the MR experiment through the use of multiple, low-cost MR scanners (refer to Chapter III). In the second method, the first ever 64-channel receiver for parallel MR imaging was created (see Chapters IV-VI). This method focused on increasing the speed of the MR experiment through the use of arrays of sensor coils to perform parallel MR imaging. The significance of the work undertaken in each method was presented in section VIII.1. Following this, both a brief analysis of the limitations and suggestions for future improvements to the prototype systems were given in section VIII.2.

REFERENCES

[1] R. R. Price, "The AAPM/RSNA physics tutorial for residents: MR imaging safety considerations," *Radiographics*, vol. 19, pp. 1641-1651, 1999.

[2] B. A. Chronik and B. K. Rutt, "A comparison between human magnetostimulation thresholds in whole-body and head/neck gradient coils," *Magnetic Resonance in Medicine*, vol. 46, pp. 386–394, 2001.

[3] B. A. Chronik, "Gross physiological measurements do not correlate significantly to peripheral nerve stimulation thresholds exhibited by humans during magnetic resonance imaging," presented at Bioelectromagnetics Society 24th Annual Meeting, Quebec City, Canada, 2002, pp. 37-38.

[4] Y. Matsuda, S. Utsuzawa, T. Kurimoto, T. Haishi, Y. Yamazaki, K. Kose, I. Anno, and M. Marutani, "Super-parallel MR microscope," *Magnetic Resonance in Medicine*, vol. 50, pp. 183–189, 2003.

[5] N. A. Bock, N. B. Konyer, and R. M. Henkelman, "Multiple-mouse MRI," presented at ISMRM Ninth Scientific Meeting, Glasgow, Scotland, 2001, p. 1108.

[6] X. J. Chen, N. Bock, F. S. Foster, and R. M. Henkelman, "A mouse imaging centre for large-scale phenotyping," presented at Pathology of Genetically Engineered Mice II, Stockholm, Sweden, 2001.

[7] M. P. McDougall, D. G. Brown, D. Spence, and S. M. Wright, "A low-pass trombone birdcage coil with broad tuning range," presented at ISMRM Ninth Scientific Meeting, Glasgow, Scotland, 2001, p. 1092.

[8] S. M. Wright, M. P. McDougall, D. G. Brown, and J. Hazle, "Arrays of birdcage coils for imaging multiple samples," presented at ISMRM Ninth Scientific Meeting, Glasgow, Scotland, 2001, p. 18.

[9] J. W. Carlson, "An algorithm for NMR imaging reconstruction based on multiple RF receiver coils," *Journal of Magnetic Resonance*, vol. 74, pp. 376-380, 1987.

[10] M. Hutchinson and U. Raff, "Fast MRI data acquisition using multiple detectors," *Magnetic Resonance in Medicine*, vol. 6, pp. 87-91, 1988.

[11] J. R. Kelton, R. L. Magin, and S. M. Wright, "An algorithm for rapid image acquisition using multiple receiver coils," presented at Eighth Annual Meeting of

the Society of Magnetic Resonance in Medicine, Amsterdam, The Netherlands, 1989, p. 1172.

[12] 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.

[13] K. Preussman, "SENSE: Sensitivity Encoding for Fast MRI," *Magnetic Resonance in Medicine*, vol. 42, pp. 952-962, 1999.

[14] M. A. Griswold, P. M. Jakob, M. Nittka, and J. W. Goldfarb, "Partially parallel imaging with localized sensitivities (PILS)," *Magnetic Resonance in Medicine*, vol. 44, pp. 602-609, 2000.

[15] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase, "Generalized autocalibrating partially parallel acquisitions (GRAPPA)," *Magnetic Resonance in Medicine*, vol. 47, pp. 1202-1210, 2002.

[16] J. R. Porter and S. M. Wright, "A sixteen-channel multiplexing upgrade for single channel receivers," *Magnetic Resonance Imaging*, vol. 19, pp. 1009-1016, 2001.

[17] F. Bloch, "Nuclear induction," *Physical Review*, vol. 70, pp. 460-474, 1946.

[18] S. M. Wright, "Chapter 16: Nuclear magnetic resonance and magnetic resonance imaging," in *Introduction to Biomedical Engineering*, J. Enderle, S. Blanchard, and J. Bonzino, Eds. New York: Academic Press, 1999, pp. 783-842.

[19] Z. P. Cho, J. P. Jones, and M. Singh, *Foundations of Medical Imaging*. New York: Wiley, 1993.

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

[21] K. Chang, *Microwave Solid-State Circuits and Applications*. New York: Wiley, 1994.

[22] H. Nyquist, "Certain topics in telegraph transmission theory," *Transactions of the American Institute of Electrical Engineers*, vol. 47, pp. 617–644, 1928.

[23] "Application note AN9509.2: Digital IF subsampling using the HI5702, HSP45116 and HSP43220," Intersil Corp., Irvine, CA, 1998.

[24] W. Kester, "Practical analog design techniques: Section 5: Undersampling applications," Analog Devices Inc., Norwood, MA, 2003.

[25] B. Razavi, "Design considerations for direct-conversion receivers," *IEEE Transactions on Circuits and Systems—II: Analog and Digital Signal Processing*, vol. 44, pp. 428-435, 1997.

[26] B. Razavi, "Challenges in portable RF transceiver design," *IEEE Circuits and Devices Magazine*, pp. 12-25, 1996.

[27] *The ARRL Handbook for Radio Amateurs*, 75th ed. Newington, CT: The American Radio Relay League, 1998.

[28] B. Brannon, "Application note AN-502: Designing a superheterodyne receiver using an IF sampling diversity chipset," Analog Devices Inc., Norwood, MA, 1998.

[29] "AD607 data sheet: Low power mixer 3 V receiver IF subsystem," Analog Devices Inc., Norwood, MA, 2002.

[30] "AD9101 data sheet: 125 MSPS monolithic sampling amplifier," Analog Devices Inc., Norwood, MA, 1992.

[31] "SPT9101 data sheet: 125 MSPS sample-and-hold amplifier," Fairchild Semiconductor International, South Portland, ME, 1999.

[32] "SHC605 data sheet: High-speed operational track-and-hold amplifier," Texas Instruments Inc., Dallas, TX, 1993.

[33] P. B. Roemer, W. A. Edelstein, C. E. Hayes, S. P. Souza, and O. M. Mueller, "The NMR phased array," *Magnetic Resonance in Medicine*, vol. 16, pp. 192–225, 1990.

[34] B. Craig, "A scrapbook of digital signal processing," SkylondaWorks, Sky Londa, CA, 2004.

[35] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*, 3rd ed. New York: McGraw-Hill, 1991.

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

[37] "AD6640 data sheet: 12-Bit, 65 MSPS IF sampling A/D converter," Analog Devices Inc., Norwood, MA, 2003.

[38] R. A. Beckman and E. R. P. Zuiderweg, "Guidelines for the use of oversampling in protein NMR," *Journal of Magnetic Resonance: Series A*, vol. 113, pp. 223-231, 1995.

[39] J. R. Porter and S. Zha, "A labVIEW based magnetic resonance imaging console," presented at Engineering in Medicine and Biology, 1999: 21st Annual Conference and the 1999 Annual Fall Meeting of the Biomedical Engineering Society, Atlanta, GA, 1999, p. 1075.

[40] S. C. Hageman, "Build your own network analyzer--part 1," *QST*, vol. 82, pp. 39-45, 1998.

[41] "HSP45102 data sheet," Intersil Corp., Irvine, CA, 1999.

[42] "HI5735 data sheet," Intersil Corp., Irvine, CA, 2003.

[43] "Applications note AN-80-001: Using model PSW-1211 as a SPST switch," Mini-Circuits, Brooklyn, NY, 1994.

[44] T. Peck, "Personal Communication: Inexpensive Gradient Amplifiers." College Station, TX, 1997.

[45] "LM12CL data sheet," National Semiconductor Corp., Santa Clara, CA, 1999.

[46] D. M. Cole, E. Esparza, F. R. Huson, S. M. Wright, and A. Elekes, "Design improvements for permanent magnet-based MRI," presented at ISMRM Fifth Scientific Meeting and Exhibition, Vancouver, CA, 1997, p. 1079.

[47] E. Esparza and D. M. Cole, "A low-cost MRI permanent magnet prototype," presented at Second Mexican Symposium on Medical Physics, Mexico City, Mexico, 1998, pp. 119–129.

[48] W. Anderson, "Electrical current shims for correcting magnetic fields," *Review of Scientific Instruments*, vol. 32, pp. 241-250, 1961.

[49] J. W. Nilsson, *Electric Circuits*, 4th ed. Reading, MA: Addison-Wesley, 1993.

[50] D. G. Brown, M. P. McDougall, and S. M. Wright, "A receiver design for parallel imaging with large arrays," presented at ISMRM Tenth Scientific Meeting and Exhibition, Honolulu, Hawaii, 2002, p. 863.

[51] C. J. Hardy, R. D. Darrow, M. Saranathan, R. O. Giaquinto, K. W. Rohling, C. L. Dumoulin, Y. Zhu, and P. A. Bottomley, "Real-time large-FOV MRI with a massively parallel 32-channel MRI system and detector array," presented at ISMRM Eleventh Scientific Meeting and Exhibition, Toronto, Canada, 2003, p. 471.

[52] "Siemens develops revolutionary magnetic resonance technology," Siemens AG, Munich, Germany, 2003.

[53] C. Milano and C. Zindel, "A quantum leap in MR tomography: TIM [76x32]," *Medical Solutions Magazine*, pp. 10-15, 2003.

[54] "NYU medical center first to install revolutionary new MR technology: Siemens MAGNETOM Avanto incorporates TIM, the first seamless, whole body coil design," Siemens AG, Munich, Germany, 2003.

[55] "Philips unveils advanced digital imaging technologies at Radiological Society of North America annual meeting, showcases Ambient Experience," Philips Medical Systems, Andover, MA, 2003.

[56] "Achieva MR with freewave technology: Scalable 32-Channel architecture for today and tomorrow," Philips Medical Systems, Andover, MA, 2004.

[57] S. M. Wright, M. P. McDougall, and D. G. Brown, "Single echo acquisition of MR images using RF coil arrays," presented at 24th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Houston, TX, 2002, pp. 1181-1182.

[58] S. M. Wright, M. P. McDougall, K. Kurpad, and D. G. Brown, "Parallel imaging: System design and limitations," presented at 2nd IEEE International Symposium on Biomedical Imaging: From Nano to Macro, Arlington, VA, 2004, pp. 1200-1203.

[59] S. M. Wright and M. P. McDougall, "Single echo acquisition (SEA) imaging at 125 frames per second," presented at International Society for Magnetic Resonance in Medicine Twelfth Scientific Meeting and Exhibition, Kyoto, Japan, 2004, p. 533.

VITA

David Gerald Brown received his Bachelor of Science in electrical engineering (*cum laude*) from Texas A&M University in May 1996. From October 1997 to December 2002, the author worked as a research assistant in the Magnetic Resonance Systems Lab at Texas A&M University. His research focused on wavelet-based magnetic resonance imaging, low-cost MR scanners for low-field magnet systems, and highly parallel receiver hardware for magnetic resonance imaging sensor arrays. He received his Doctor of Philosophy degree in December 2005 and is currently working as an RF/ microwave systems engineer at a small consulting firm. His areas of specialty include electromagnetics, RF/ microwave design, receiver design, magnetic resonance imaging, MRI system design, and signal processing.

For further information, the author may be contacted at Brown Wolf Engineering, Inc., 5365 CR 1037, Celeste, TX 75423. His e-mail address is dgbrown@neo.tamu.edu.