Order Number 1347572

A digital-electronic video-rate reconstruction system for magnetic resonance imaging

Snell, Rodney James, M.S.

The University of Arizona, 1992

U·M·I
300 N. Zeeb Rd.
Ann Arbor, MI 48106

A DIGITAL-ELECTRONIC VIDEO-RATE RECONSTRUCTION SYSTEM FOR MAGNETIC RESONANCE IMAGING

by

Rodney James Snell

A Thesis Submitted to the Faculty of the

COMMITTEE ON OPTICAL SCIENCES (GRADUATE)

In Partial Fulfillment of the Requirements

For the Degree of

MASTER OF SCIENCE

In the Graduate College

THE UNIVERSITY OF ARIZONA

1 9 9 2

STATEMENT BY AUTHOR

This thesis has been submitted in partial fulfillment of requirements for an advanced degree at the University of Arizona and is deposited in the University Library to be made available to borrowers under rules of the library.

Brief quotations from this thesis are allowable without special permission, provided that accurate acknowledgement of source is made. Requests for permission for extended quotation from or reproduction of this manuscript in whole or in part may be granted by the head of the major department or the Dean of the Graduate College when in his or her judgment the proposed use of the material is in the interests of scholarship. In all other instances, however, permission must be obtained from the author.

SIGNED: Rodney Snell

APPROVAL BY THESIS DIRECTOR

This thesis has been approved on the date shown below:

Arthur F. Gmitro

A.F. Gmitro
Professor of Radiology and Optical Science

1/10/92

Date

Acknowledgements

I thank my advisor and friend, Arthur Gmitro, for his continual (and largely undeserved) assistance and inspiration, both in and out of the lab.

Special thanks to the "Crimetime after Primetime" crew: John Aarsvold, Andy Alexander, and Tim White. Sometimes the best psychiatrists are the craziest.

Much appreciation to Laura Lindberg for her patience and support.

I thank my brother, Reed Snell, for his encouragement and wisdom.

To my mother, Mary Snell; her constant support and sacrifice has made this possible.

Table of Contents

List of Illustrations

List Of Tables

ABSTRACT

A digital-electronic video-rate reconstruction system for Magnetic Resonance Imaging (MRI) has been designed to first order. The maturation of ultra-fast acquisition techniques in MRI has produced the need for a reconstruction system that will enable dynamic processes to be viewed on-line. Conventional reconstruction hardware is not configured for real-time reconstruction and previous developments are limited in accuracy and flexibility. The real-time reconstruction system presented here consists of three main subsystems. A digitizer interfaces with an MR scanner to digitize data matrices of resolutions up to 256 x 256 at arbitrary rates up to video rates. A Fourier processor performs either 2D Fourier transformation or projection filtering on the digitized data at video-rates. A backprojector performs the backprojection operation on filtered-projection data at video-rates. The complete system would be able to reconstruct data acquired from nearly any acquisition technique. True real-time MRI is then possible.

CHAPTER 1

INTRODUCTION

1.1 Motivation

Since its inception in 1973, Magnetic Resonance Imaging (MRI) has become one of the most versatile and useful imaging modalities available in diagnostic radiology. The main reasons for this are that MRI is non-invasive, does not use ionizing radiation, and provides excellent soft tissue contrast. In addition, MRI possesses a multitude of parameters that can be varied to easily extract pertinent diagnostic information. However, for all the benefits that MRI provides, it remains costly, is uncomfortable for the patient, and is useful mostly for imaging the central nervous system and the musculoskeletal system [1]. All of these drawbacks are due, in part, to the long times associated with conventional imaging sequences. One image can take many minutes to produce, with a total examination taking an hour or more. Patient throughput is slow which causes the cost per patient to be high. The patient must lay still in the imaging magnet for the length of an examination. This can be very uncomfortable, especially if the patient is in pain. Any object motion causes image degradation which limits the use of conventional MRI to mainly static anatomy.

A conventional MR image provides only static, time-averaged information about the object being imaged. If the object moves during data acquisition, severe artifacts may be seen in the final image. Patient motion or physiological

motion such as respiration, pulsation, and blood flow are enough to render an image unusable. Patient motion is nearly impossible to prevent, especially if the patient is in pain or is claustrophobic from the small clearance available in the bore of a conventional imaging magnet. It is likely that long imaging times only aggravate this type of motion. The effects from certain periodic physiological motions such as respiration and cardiac pulsation can be minimized by synchronizing the acquisition with the motion [2]. However, it is reasonable to expect that many patients will have irregular breath or heartbeat patterns, and so image degradation will occur anyway. In addition, most transient events go unseen with gating procedures.

Long acquisition times ultimately limit the utility of MRI. Fortunately, there are acquisition methods that can speed up the imaging process substantially. Some of these techniques reduce acquisition times to a few seconds [3]. This is fast enough to image the abdominal area in a breathhold but still not fast enough to capture dynamic processes without special synchronization schemes. With improved hardware, subsecond acquisition times can be achieved. Recently, developments in this area have produced methods with acquisition times as fast as 30 ms per image (video-rates) [4]. These ultra-fast techniques typically have stringent hardware requirements (high-power gradient drivers, shielded gradients, high field homogeneity), but have been shown to produce good, clinically useful information. It can be expected that these methods will only get better and that future MR scanners will come equipped with the appropriate hardware to implement them.

Ultra-fast acquisition methods acquire data in the range from 30 ms per image to 100 ms per image. Since these times effectively freeze physiological motion, motion effects are greatly reduced with ultra-fast techniques. Currently,

however, conventional MR reconstruction systems can not process at the rates required of these methods. Typically, a "snapshot" scan is performed where one frame of data is acquired very fast and then reconstructed off-line. The final image may not be seen for many seconds after the data are acquired; thus, the final image provides only retrospective information. To actually view dynamic events as they occur, ultra-fast acquisition must be accompanied by ultra-fast reconstruction; ie. true real-time MRI. In this thesis, real-time is defined somewhat loosely but it encompasses acquisition and reconstruction rates in the range from 30 ms per image to 100 ms per image.

Real-time MRI would have many uses. Dynamic processes could be viewed as they occur so that functional as well as anatomical information could be obtained. Flow dynamics and time-varying temperature mapping are examples of this kind of experiment [5]. MRI scan parameters such as pulse timing could be altered in real-time and any changes they introduce would immediately be seen [6]. For instance, the optimal contrast between a lesion and it's surroundings could be found quickly and easily. This would save time spent on superfluous scans. Pre-positioning and slice selection would be greatly aided by the ability to "glide" through a volume and immediately see the surrounding anatomy. Interventional techniques such as the positioning of catheters and biopsy needles could be performed under MR guidance. In addition, once real-time MRI is possible, new uses for the technology are inevitable.

Reconstruction of MR data in real-time places large demands on a reconstruction system. Common MR reconstruction algorithms such as 2D Fourier transformation and filtered backprojection are computationally intensive and so performing these algorithms at real-time speeds is a difficult task. A

real-time reconstruction system has the responsibility of performing any needed pre-reconstruction processing of the acquired data, which is required by nearly all the ultra-fast acquisition methods (chapter 3). Additionally, a real-time reconstruction system must concurrently input, process, and output data. Conventional reconstruction systems are not configured for these kinds of continuous operations; therefore, a dedicated approach must be used. Real-time hardware implementations of the reconstruction algorithms have previously been demonstrated, but most of these systems are not adequate for MR reconstruction because of limited accuracy and flexibility (chapter 2).

In the interest of realizing the full potential of real-time MRI, this thesis is concerned with the design of a video-rate digital-electronic MRI reconstruction system. The system discussed will be able to reconstruct data matrices up to a resolution of 256 x 256 by straightforward scaling of the hardware. The system is designed to continuously reconstruct data acquired from nearly all MR acquisition methods. This is accomplished by including hardware to implement both the 2D Fourier transform and filtered backprojection algorithms as well as designing the system to be as flexible as possible. Coupled with ultra-fast acquisition, this system would enable true real-time MRI.

1.2 Outline

In chapter 2, an overview of basic MRI principles is given. 2D Fourier and projection imaging are discussed in terms of their k-space trajectories along with the appropriate reconstruction algorithms, the 2D Fourier transform and filtered backprojection.

In chapter 3, the causes of long scan times are given and the ultra-fast techniques for reducing these times are briefly discussed. The additional tasks

that ultra-fast imaging techniques add to the reconstruction process are discussed. Sampling, interpolation, data reorganization, and special scan enhancements such as partial k-space scanning are among the considerations in designing a real-time reconstruction system.

In chapter 4, the computational rates required for video-rate reconstruction are discussed. Justification for the necessity of a dedicated digital-electronic reconstruction system is given in terms of the limitations of previous developments.

In chapter 5, a first-order design of a video-rate digital-electronic MR reconstruction system is presented. The system consists of a digitizer, Fourier processor, and backprojector. Considerations for each of these is given.

In chapter 6, the utility of the system for real-time MRI is discussed. Multiple acquisition scenarios are given, and the reconstruction processes for these are discussed.

CHAPTER 2

PRINCIPLES OF NMR AND MRI

Nuclear Magnetic Resonance (NMR) was introduced in the 1950's as a spectroscopic tool used in determining the chemical structure of complex molecules. Basic NMR spectroscopy utilizes the paramagnetic properties of certain atomic nuclei to identify the surrounding chemical environment. Magnetic Resonance Imaging (MRI) takes this one step further by enabling the NMR information to be dependent on its physical location within a sample. The basic concepts of MRI are presented in this chapter along with a description of the 2D Fourier and projection imaging techniques. Reconstruction of data acquired from these acquisition methods is also discussed. More detailed information of the basic principles of NMR and MRI can be found in references [7] and [8].

2.1 NMR Principles

Every magnetic nucleus possesses a quantity called spin. In MRI, the nucleus of most interest is the proton (^1H) because it is the most abundant and sensitive nucleus in biological systems. The classical model of spin depicts a nucleus (proton) rotating on it's own axis. Because the particle has charge and is rotating, it has both a magnetic dipole and angular momentum. When placed in a strong static magnetic field, B_0, the dipole aligns itself with the field. A collection of these spins produces a net magnetization which, under equilibrium conditions, lies in the same direction as the external field. If perturbed from

equilibrium, the net magnetization precesses around the static field with a frequency given by the Larmor relation,

$$\omega = \gamma B_0 \ . \qquad (2\text{-}1)$$

The precession frequency, ω, is called the Larmor frequency and is the resonance condition for the spin system. The magnetogyric ratio, γ, is unique for each magnetic isotope. For a given nuclear species (e.g. ^1H), γ is a constant, and thus the Larmor frequency is proportional to the strength of the magnetic field.

2.1.1 RF excitation

To perturb the net magnetization from it's equilibrium position, the sample is irradiated with radio-frequency (RF) energy applied perpendicular to B_0. This RF field must be oscillating at the Larmor frequency since this is the resonance frequency for the spins. The effect of this radiation is to rotate the magnetization down toward the transverse plane. This is shown in figure 2.1. The length and magnitude of the RF field (pulse) determines the angle of rotation, called the flip angle. The most common flip angles are 90° and 180°, but many fast-scan techniques in MRI use small flip angles to achieve a stronger signal under steady-state conditions (chapter 3).

A standard MRI system is configured so that the transverse magnetization is detected. There is no transverse magnetization in equilibrium but, after RF excitation, there can be a transverse precessing component to the magnetization. This component induces the signal detected in MRI. This signal oscillates at the Larmor frequency, and the magnitude of this signal is

proportional to the amount of transverse magnetization present, which, in turn, is proportional to the number of nuclei excited by the RF pulse.

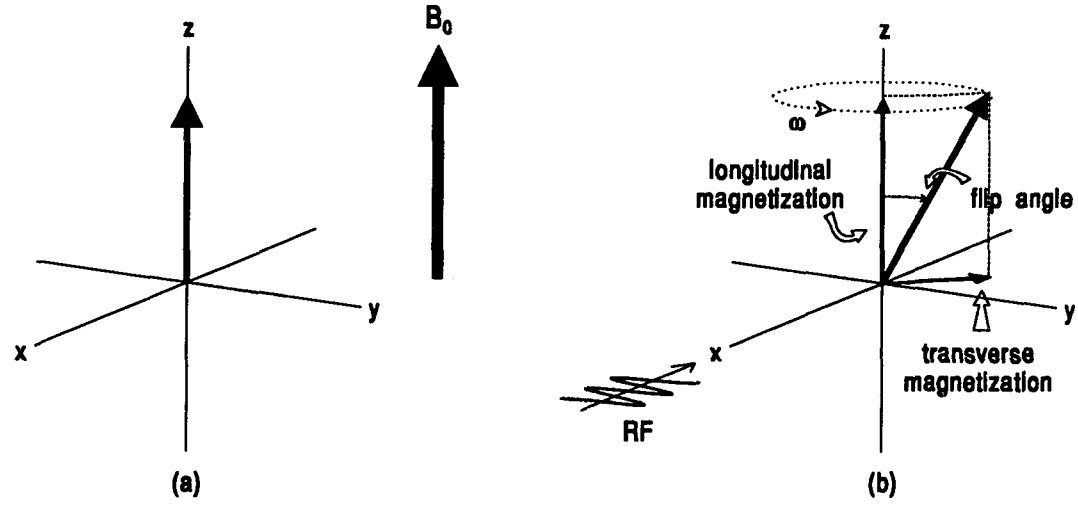

Figure 2.1 Spin magnetization; (a) in equilibrium the magnetization lies along the static field, (b) RF excitation causes the magnetization to rotate down toward the transverse plane and precess around the static field at the Larmor frequency.

2.1.2 Relaxation

Following RF excitation, the oscillating signal amplitude does not stay constant but rather decays to zero during a relaxation period. In one relaxation process, the excited spins slowly lose their energy to the surrounding environment. This causes the magnetization to lose it's transverse component and recover to it's equilibrium value (equilibrium longitudinal magnetization). This is known as spin-lattice relaxation and is often modelled as an exponential recovery of the longitudinal magnetization with a time constant T_1. This time is

sample dependent and varies from a few hundred milliseconds to about 2 seconds in biological tissues.

A second relaxation process affects the transverse component of the magnetization. Immediately after RF excitation, all individual nuclear moments in the sample are aligned (i.e. phase coherent) in the transverse plane. However, the moments gradually fan out with some precessing a little faster or slower than the resonant frequency depending on the actual field experienced by an individual moment. If the external static field was perfectly homogenous, this process would be due only to time-varying field fluctuations between interacting spins. This process is called spin-spin relaxation and is often modelled as an exponential decay of the transverse magnetization with time constant T_2. However, the static field has always some inhomogeneity and so spins at different spatial locations will experience different magnetic fields. Because of these additional field variations, the transverse magnetization decays with a time constant T_2^*, which is shorter than T_2. This process becomes very important in ultra-fast MRI where the data acquisition window is on the order of T_2^* (chapter 3). It is possible in some cases to reverse the dephasing process caused by static field inhomogeneities with a spin-echo sequence (see below). T_2 relaxation times are usually much shorter than T_1 relaxation times; tens to hundreds of milliseconds in biological samples.

2.1.3 Free induction decay

Following RF excitation of the spin system, the MR signal decays down to zero due to the relaxation processes described above. This signal is known as a free induction decay (FID). For a sample containing pure water, the spectrum of this FID (obtainable through 1D Fourier transformation of the FID) will consist of

one resonance peak at the Larmor frequency with a magnitude that is proportional to the number of spins in the sample. If there are nuclei in the sample in different chemical environments (e.g. protons of CH_2 or CH_3 groups), their resonance frequencies can be shifted. The spectrum of the FID from such a sample will consist of multiple peaks at the resonance frequencies of these nuclei. This is the basis for NMR spectroscopy.

2.2 MRI Acquisition

An FID signal contains information about the resonance frequencies and number of nuclei present in a sample. The signal does not, however, convey information about the spatial locations of these nuclei in the sample. MRI adds the ability to encode spatial information onto the NMR signal through the use of gradient magnetic fields. The magnitudes of these fields vary linearly in the three orthogonal directions but the fields are always parallel to the static field direction (denoted as $\hat{z}$). The total magnetic field becomes

$$\mathbf{B}(t) = [B_0 + xG_x(t) + yG_y(t) + zG_z(t)]\,\hat{z}\,, \qquad (2\text{-}2)$$

where B_0 is the static magnetic field and G_x, G_y, and G_z are the time-varying field gradients. In the MR tomographic experiments discussed here, G_z is used to select a slice to be imaged, while G_x and G_y are used to differentiate pixels in that slice.

Consider a sample that has been excited with an RF pulse. If a constant gradient is applied immediately after the RF pulse, individual spins corresponding to different planes orthogonal to the gradient direction will precess at different frequencies (equations (2-1) and (2-2)). The FID will contain all of these frequency components which can be related back to their spatial position

along the gradient direction. This is the basis for spatial encoding in MRI and will be discussed more fully in section 2.2.4. Note that because the gradient provides a field variation, the spins will fan-out or dephase as discussed above.

2.2.1 Spin-echo

An FID echo can be produced with a 180° RF pulse applied some time, τ, after an initial (usually 90°) RF pulse. This is shown in figure 2.2. The gradient G_x can be used to encode spatial information in the FID, but for purposes here, it is only important that the gradient causes additional dephasing of the individual spins. The initial RF pulse produces a net transverse magnetization which decreases rapidly because of the dephasing effects caused by the gradient, T_2 relaxation, field inhomogeneity, eddy currents, etc. The 180° RF pulse flips the partially dephased magnetization over in one direction. If the same gradient magnitude, G_x, is turned on after the 180° RF pulse, the faster precessing spins will catch up with the slow ones causing the dephased spins to refocus at a time 2τ after the initial RF pulse. This refocussed signal is called a spin-echo. Any dephasing of the spins due to the gradient, static field inhomogeneity, or chemical shift will be refocussed so that the only loss in signal will be that due to irreversible T_2 decay. Thus, one use of the spin-echo is to obtain true T_2 weighted images.

2.2.2 Gradient-echo

An alternative way to refocus the dephased spins is through gradient reversal after an RF pulse has been applied. This is shown in figure 2.3. Consider the x-gradient with magnitude G_x turned on immediately after a 90° RF pulse is terminated. From equations (2-2) and (2-1), the spins will be precessing at a frequency, $\omega_x = \gamma(B_0 + xG_x)$. If the gradient is reversed to magnitude $-G_x$ at

some time τ after the RF pulse, the spins will start precessing with a frequency $\omega_x = \gamma(B_0 - xG_x)$ in the same direction. The net magnetization will be refocussed at a time 2τ after the RF pulse, thereby forming a gradient-echo at this point. The main difference from a spin-echo is that dephasing of the spins due to static field inhomogeneity will not be refocussed. Gradient-echoes are used extensively in ultra-fast imaging because low flip-angle sequences can not use spin echoes.

2.2.3 Slice selection

MRI can be used in a tomographic mode through a process known as selective excitation. This technique is illustrated in figure 2.4a. A 90° RF pulse and a gradient, say G_z, are applied simultaneously to excite a slice of spins in a plane orthogonal to the gradient (z) axis. The gradient G_z forces each plane of nuclei orthogonal to the z axis to have a unique Larmor frequency depending on the plane's z coordinate. The finite width of a desired slice, Δz, corresponds to a particular frequency range, $\Delta\omega$ (figure 2.4b). The RF pulse must be shaped so that it uniformly excites only this frequency range and leaves the rest of the object unperturbed; i.e. the pulse contains a rectangular distribution of frequencies. As shown in figure 2.4c, a sinc-shaped pulse envelope produces the correct frequency excitation since the sinc and rect functions are a Fourier transform pair. The RF pulse selectively excites the frequency range corresponding to the slice selected by G_z. The exact location of the excited slice is determined by the RF carrier frequency. If this frequency is the on-resonance frequency, $\omega_0=\gamma B_0$, the selected slice will be at z=0. Any other slice can be excited by offsetting the RF carrier from ω_0. Because G_z is on while the spins are precessing down toward the transverse plane, dephasing effects due to the

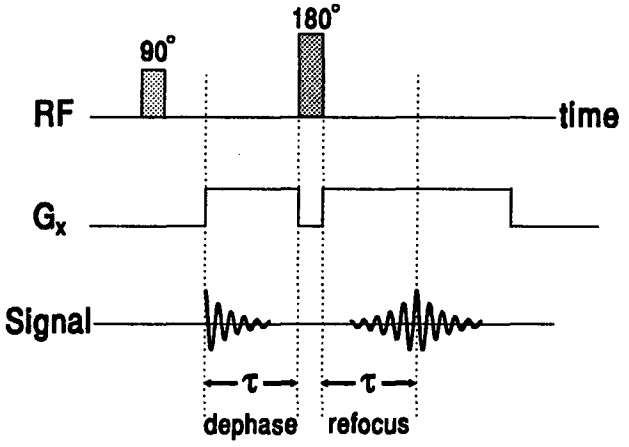

Figure 2.2 Spin-echo technique. The gradient causes spins located in different planes in the x-direction to precess at different frequencies. The net transverse magnetization is reduced because of this gradient dephasing, T_2^*, field inhomogeneity, etc. The 180°RF pulse refocuses the dephased spins to produce an echo signal.

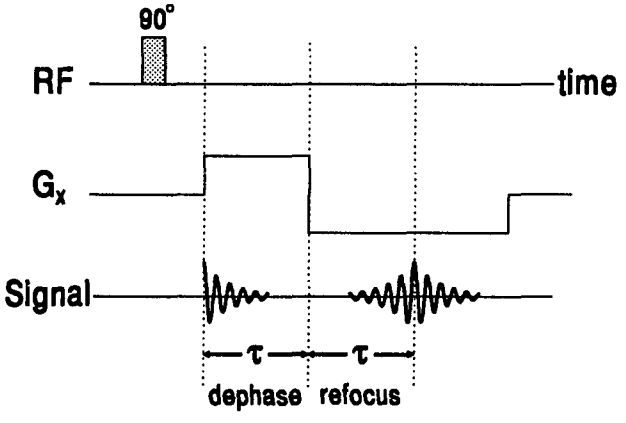

Figure 2.3 Gradient-echo technique. Gradient reversal refocuses dephased spins to produce an echo signal by appropriately altering the precession frequency of the individual spins.

gradient across the slice will cause the net transverse magnetization to be reduced. The spins can be refocussed immediately after the RF pulse by using the gradient-reversal technique described above.

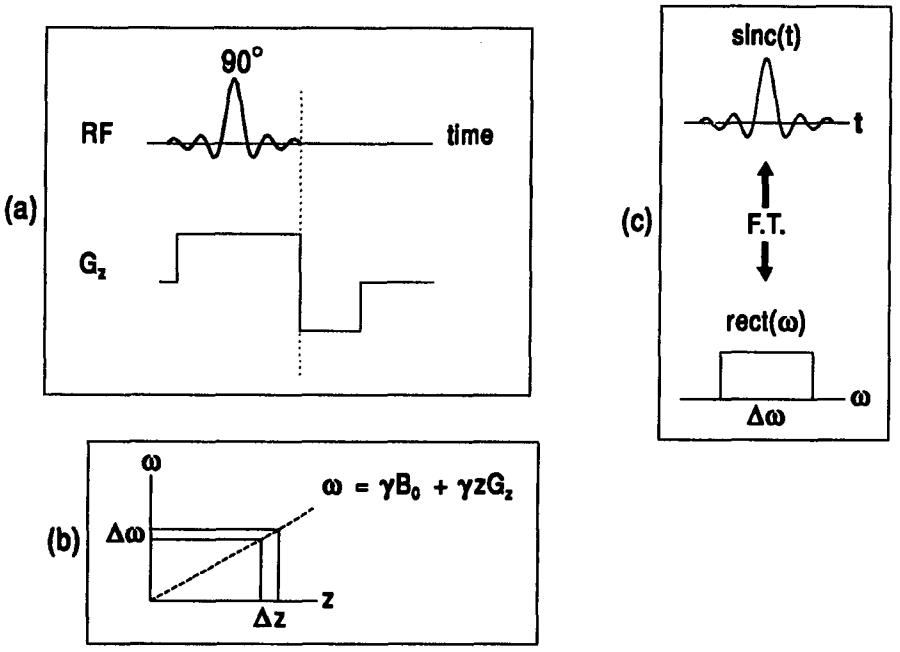

Figure 2.4 Slice selection. (a) pulse sequence used for selective excitation, (b) a desired slice corresponds to a particular frequency range, (c) uniform frequency excitation requires a sinc-shaped RF pulse.

2.2.4 Spatial encoding

Once a slice in the sample has been excited through the use of one of the gradients and a selective RF pulse, the detector receives a signal from the rotating transverse magnetization. The remaining two gradients are used to encode spatial information onto the signal. Assume that the coordinates in the slice are x and y and that the gradients used for spatial encoding are G_x and G_y. A pixel at a particular x and y generates a signal which is given by

$$dS(t) = \sigma(x,y)e^{j\phi(x,y,t)}\,dxdy\ , \tag{2-3}$$

where $\sigma(x,y)$ represents the signal strength which includes spin density and T_1 and T_2 effects. The phase, ϕ, depends upon pixel position and time and is given by

$$\phi(x,y,t) = \int_0^t \omega(x,y,t')\,dt'\ . \tag{2-4}$$

The frequency distribution follows the Larmor relation

$$\omega(x,y,t) = \gamma B_0 + \gamma x G_x(t) + \gamma y G_y(t)\ . \tag{2-5}$$

Integrating over the x,y plane gives the total detected signal,

$$S(t) = \iint \sigma(x,y)\, e^{j\phi(x,y,t)}\,dxdy\ . \tag{2-6}$$

The combination of equations (2-4), (2-5), and (2-6) yields

$$S(t) = e^{j\gamma B_0 t} \iint \sigma(x,y)\, e^{j\gamma x \int_0^t G_x(t')\,dt'}\, e^{j\gamma y \int_0^t G_y(t')\,dt'}\,dxdy\ . \tag{2-7}$$

The desired output is obtained by quadrature demodulation off of the leading exponential (i.e. demodulation of a complex signal from the Larmor frequency carrier) [9] which leaves,

$$S'(t) = \iint \sigma(x,y)\, e^{j2\pi k_x(t)\, x}\, e^{j2\pi k_y(t)\, y}\, dxdy \ , \tag{2-8}$$

where k_x and k_y are given by,

$$\begin{aligned} k_x(t) &= \frac{\gamma}{2\pi} \int_0^t G_x(t')dt' \\ k_y(t) &= \frac{\gamma}{2\pi} \int_0^t G_y(t')dt' \end{aligned} \tag{2-9}$$

The right side of equation (2-8) is recognized as the 2D inverse Fourier transform of $\sigma(x,y)$ with the Fourier conjugate variables k_x and k_y. The data collected in MRI, then, are in the spatial frequency space (k-space) of the object being imaged.

From equation (2-9), the gradients are responsible for the time evolution of the spatial frequencies $k_x(t)$ and $k_y(t)$. The gradients determine which data points in k-space are collected and the order in which they are collected. Therefore, the procedure used for reconstructing the object, $\sigma(x,y)$, will depend upon the k-space trajectory taken. Of interest here are two basic techniques: 2D Fourier imaging and projection imaging.

2.2.5 2D Fourier Imaging

The pulse sequence (RF and gradient waveforms as a function of time) shown in figure 2.5 is used in conventional MR imaging. As discussed previously, G_z is used with a frequency selective RF pulse to select the slice of

interest. G_x and G_y are then used to pre-position the scan at some k_x and k_y. The magnitude and length of these gradients determine where the k-space trajectory will start. The y gradient is then turned off and the x gradient reversed so that a horizontal line is traversed in k-space. Because of the G_x gradient reversal, a gradient-echo is formed during readout. If this sequence is repeated with a lower magnitude G_y, another horizontal line is traced below the first. By repeating the experiment with the appropriate magnitudes of G_y, the full 2D Fourier space is built up as a sequence of horizontal "raster" lines. This method is called 2D Fourier imaging.

2.2.6 Projection imaging

Another method of collecting data in k-space is shown in figure 2.6. Similar to the 2D Fourier method, a slice is selected and the x and y gradients pre-position the scan at some k_x and k_y. However, instead of turning G_y off during readout, it is reversed. The sequence traces a radial line whose slope depends on the ratio of the gradient strengths in the x and y directions. By adjusting these gradient magnitudes properly in subsequent experiments, the full 2D Fourier space can be built up as a sequence of radial lines. This method is called projection imaging.

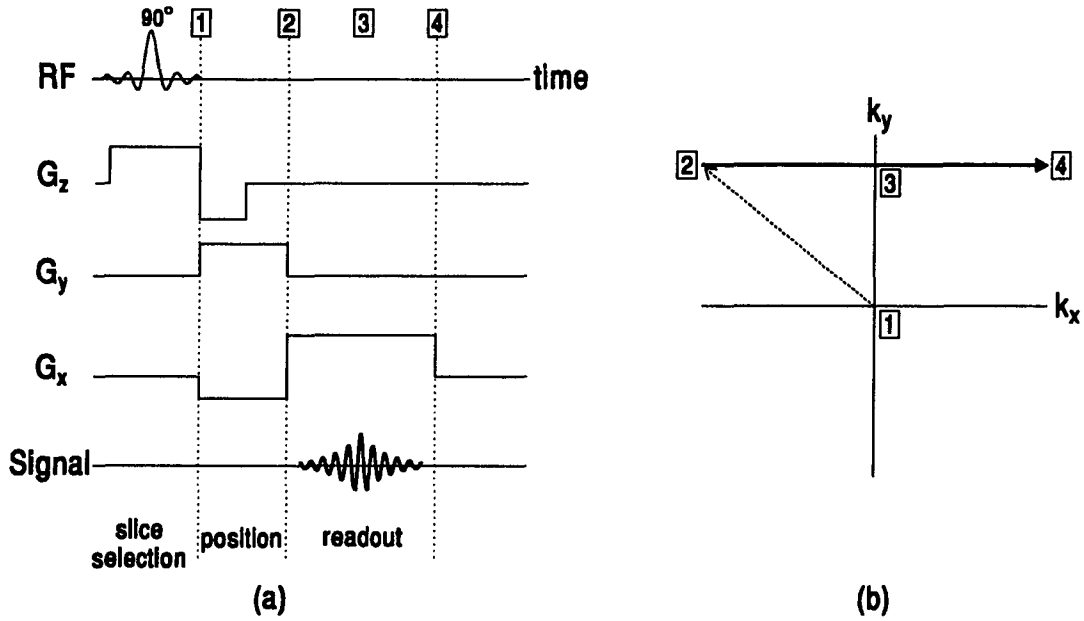

Figure 2.5 Raster Fourier imaging. (a) shows a typical pulse sequence and (b) shows the corresponding trajectory in k-space. The boxed numbers represent points in time and the corresponding locations in k-space. At [1], the scan is at the center of k-space. Pre-positioning places the scan at [2]. A horizontal line is traced from [2] to [4] with the echo becoming maximum at [3].

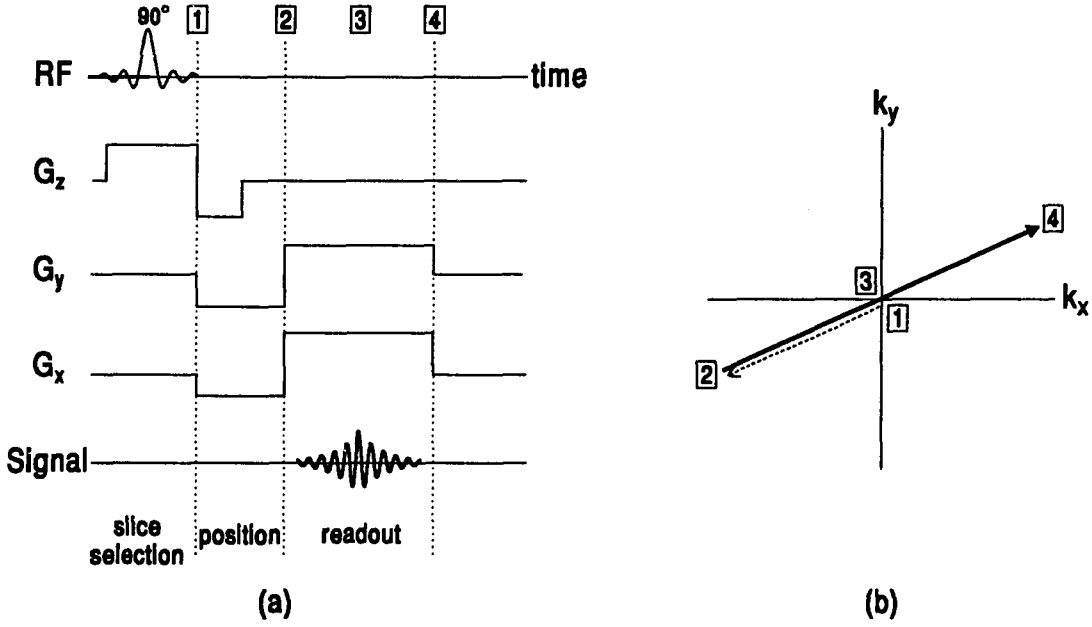

Figure 2.6 Projection imaging.(a) shows a typical pulse sequence and (b) shows the corresponding trajectory in k-space. The boxed numbers represent points in time and the corresponding locations in k-space. At [1], the scan is at the center of k-space. Pre-positioning places the scan at [2]. A radial line is traced from [2] to [4] with the echo becoming maximum at [3].

2.3 MRI Reconstruction

The algorithm needed to reconstruct the object depends on the scan-pattern in k-space. The scan patterns depicted above and the real-time scan patterns discussed in chapter 3 dictate reconstruction with either a 2D Fourier transform or filtered backprojection depending upon how the data are acquired. The basic flow of these algorithms follows.

2.3.1 2D Fourier imaging

Data collected from the 2D Fourier imaging method exist as raster lines in k-space. The combination of all lines, which are separate time signals, results in a data matrix that can be directly 2D Fourier transformed to reconstruct the object. This is given by the 2D Fourier transform of equation (2-8),

$$\sigma(x,y) = \iint \Sigma(k_x,k_y)\, e^{-j2\pi k_x x}\, e^{-j2\pi k_y y}\, dk_x dk_y \ , \qquad (2\text{-}10)$$

where $\Sigma(k_x,k_y)$ is the collection of data, S'(t), into the appropriate 2D array. The actual data is obtained by digitizing the analog time signals. Because the gradients are constant during acquisition, uniform sampling of S'(t) results in uniform sampling in k-space. Therefore, the sampled data points fall exactly on a rectilinear grid and a standard discrete 2D Fourier transform can be used to reconstruct the object. This is discussed further in chapter 4.

2.3.2 Projection imaging

Data are acquired as radial lines in projection imaging. Again, since the gradients are constant during readout, uniform sampling of S'(t) results in uniform sampling along a radial line. Therefore, the sampled data points fall exactly on a polar grid. Although these points could be interpolated onto a

rectilinear grid so that 2D Fourier transformation can be used to reconstruct the object, it is easier to use the filtered backprojection algorithm for reconstruction. Although a full description of the algorithm is beyond the scope of this thesis, the main steps for using the algorithm to reconstruct MR data will be outlined. A good discussion of the algorithm is given in Rosenfeld and Kak [10].

A projection, $P_\theta(p)$, of an object is the set of all possible ray integrals through the object at an angle θ where p is the distance from the origin to a ray. This is given by,

$$P_\theta(p) = \int_{ray} \sigma(x,y)\, ds = \iint \sigma(x,y)\delta(x\cos\theta+y\sin\theta-p)\, dxdy \ , \qquad (2\text{-}11)$$

where δ is the dirac delta function and $x\cos\theta + y\sin\theta = p$ describes a line at angle θ at a perpendicular distance, p, from the origin. Figure 2.7 shows a projection of an object. $P_\theta(p)$ is a 2D function of the variables θ and p and is known as the Radon transform of $\sigma(x,y)$. If a number of projections at different angles are taken, the object can be reconstructed with filtered backprojection. The data acquired in MR projection imaging are related to projection data through the Fourier slice theorem. This is shown in figure 2.8. This theorem states that the 1D Fourier transform of a projection at angle θ gives samples along a radial line at angle θ in the 2D Fourier space of the object. Data acquired in MR projection imaging are radial lines in the 2D Fourier space of the object (k-space) and so are related to the projections of the object by the Fourier slice theorem. This is given by,

$$S_\theta(w) = \int P_\theta(p) e^{-j2\pi wp}\, dp \ . \qquad (2\text{-}12)$$

Here, $S_\theta(w)$ is the MR data, S'(t), on radial lines in k-space where w and θ are the radial and angular spatial frequency variables, respectively (figure 2.8).

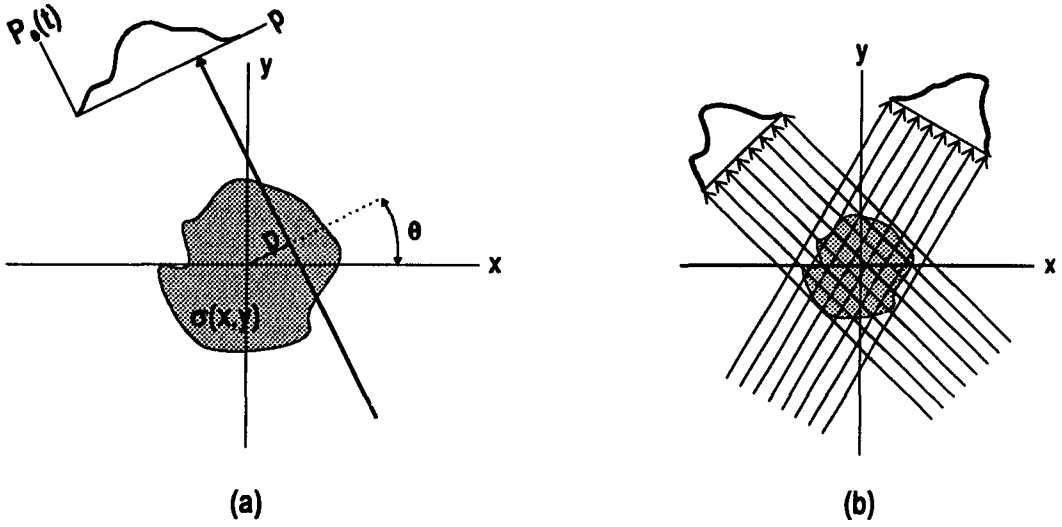

Figure 2.7 (a) A ray integral is a point in a projection. (b) projection data from multiple angles.

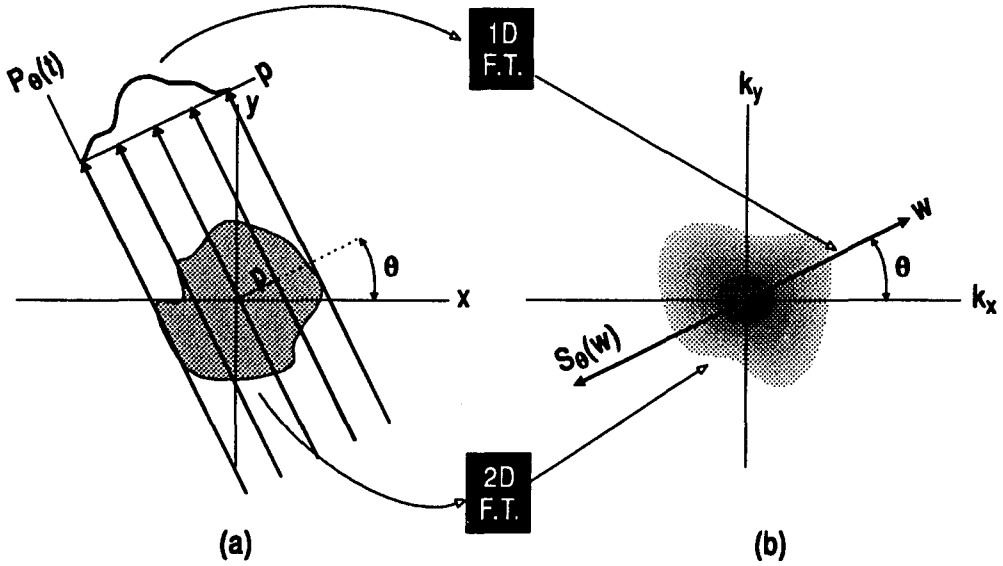

Figure 2.8 Fourier slice theorem. (a) object space, (b) k-space.

Because of the Fourier slice theorem, filtered backprojection can be used to reconstruct the data acquired with projection imaging. The continuous form of the filtered backprojection algorithm is given by,

$$\sigma(x,y) = \int_0^{\pi} Q_\theta(x\cos\theta + y\sin\theta)\, d\theta \,, \tag{2-13}$$

where

$$Q_\theta(p) = \int_{-\infty}^{\infty} S_\theta(w)|w|e^{j2\pi wp}\, dw \,. \tag{2-14}$$

These equations state that a filtered version of a projection, $Q_\theta(p)$, is calculated by multiplying the data on a 1D radial line, $S_\theta(w)$, by a 1D filter, $|w|$, and 1D Fourier transforming the result. These 1D filtered projections are then used by equation (2-13) to reconstruct the object, $\sigma(x,y)$. Equation (2-13) is called the backprojection step. Since all points along a ray through the object possess the same value of p ($p = x\cos\theta + y\sin\theta$), a sample in a filtered projection will make the same contribution to all points along the ray. In effect, all filtered projections are smeared back along the acquisition rays (backprojected) and summed to reconstruct the object. The discrete form of filtered backprojection is discussed in chapter 4.

CHAPTER 3

ULTRA-FAST ACQUISITION

Imaging times required for conventional 2D Fourier and projection methods pose obstacles to the ultimate utility of MRI. As shown in chapter 2, one line of k-space is acquired after an RF excitation and so multiple excitations are needed to fill in k-space. The time between successive excitations is known as the repetition time, T_R. Because T_1 constants are somewhat long, T_R is also long (500 to 3000 ms) to allow sufficient recovery of the magnetization. Therefore, a standard acquisition of 256 excitations to produce a final image of 256 x 256 can take from 2 to 15 minutes to complete. These long imaging times are objectionable when trying to image dynamic processes because the final image provides only static, time-averaged information. Physiological motion, such as respiration and flow, and patient motion can introduce severe artifacts into the final image. Efforts to eliminate these artifacts by drastically reducing imaging times have been ongoing since the inception of clinical MRI. Recently, some ultra-fast imaging methods have matured to the point of producing diagnostic quality images while creating the possibility of continuous real-time MRI. Although ultra-fast acquisition techniques produce good results, the reconstruction process can be more complicated than that in conventional MRI. This chapter presents a brief overview of some of these fast data acquisition methods and the complications they may add to the reconstruction process.

3.1 Multiple-shot methods

One way to reduce imaging time is to reduce T_R. However, simply reducing T_R in a conventional acquisition method also results in a reduction of the MR signal due to saturation of the magnetization. To avoid this problem, it is possible to use a small flip-angle RF pulse for excitation instead of a conventional 90° RF pulse. After a few RF excitations, the magnetization reaches a steady-state condition where the longitudinal magnetization is the same prior to each RF excitation. This magnetization will be stronger with small flip-angle pulses than that which would result from large (90°) flip angle pulses because less longitudinal magnetization is lost using smaller flip angles. Thus, with more longitudinal magnetization available in steady-state, the required recovery period is reduced and shorter repitition times can be used. Many such steady-state methods exist with names such as FLASH, FISP, GRASS, and FAST [11]. A simple FLASH pulse sequence is shown in figure 3.1. Note that an echo is obtained through gradient reversal rather than through a spin echo (chapter 2). Inverting the longitudinal magnetization with a 180° RF pulse (spin-echo) would lengthen rather than reduce the required recovery time and remove the advantages gained by using small flip angles. Because of the gradient reversal, these methods are also called gradient-echo methods. As in conventional MRI, each "shot" (one application of the pulse sequence) of a steady-state technique produces one line in k-space and multiple shots are needed to fill k-space. Imaging times for these steady-state techniques with short T_R's can be under one second, but are more

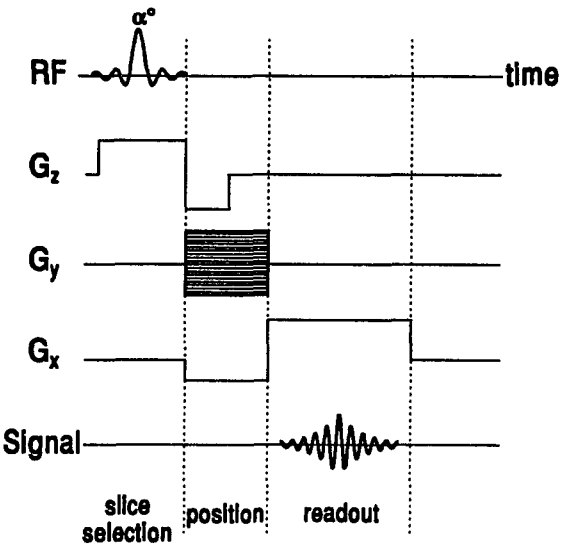

Figure 3.1 FLASH sequence. G_y is shown hatched to represent differing magnitudes for each excitation.

commonly in the range of a few seconds [12]. Although this is a bit slow to be considered real-time (at least as defined), these techniques are appropriate for many clinical applications. Reconstruction of data acquired from these techniques does not require special processing. The data acquired from a 2D Fourier FLASH method can be reconstructed directly with 2D Fourier transformation, and the data acquired from a projection FLASH method can be reconstructed directly with filtered backprojection.

3.2 Single-shot methods

There are a number of techniques to acquire all of k-space after a *single* RF excitation rather than after multiple RF excitations. These methods fall under the general heading of echo-planar [13], and a few of the variants are dubbed MBEST, InstaScan, blipped echo-planar, and spiral-scan. All share the ability to sample the entire k-space after a single RF excitation, and all produce echoes through rapid gradient reversal. The RF excitation can be in the form of a

conventional 90° RF pulse if a single frame is desired, or it can be in the form of a small flip-angle pulse if continuous acquisition is desired. The sequences differ mainly in the way k-space is traversed. These single-shot methods acquire data more rapidly than multiple-shot methods because all data collection must be done before the MR signal decays down to the noise floor. The decay is governed by T_2^* within a pixel and so typical acquisition times are in the range of 30 ms to 100 ms. Because these techniques can acquire data so fast, they place greater demands on the reconstruction system than do multiple-shot techniques.

To understand the demands that ultra-fast acquisition places on the reconstruction process, three single-shot methods will be examined: a raster-scan technique with a sinusoidal readout gradient and a pulsed gradient in the orthogonal direction (or blipped echo-planar), modified blipped echo-planar where the readout gradient approximates a square wave, and spiral-scan where both spatial encoding gradients are sinusoidal. The spatial encoding gradients and k-space trajectories for these are shown in figure 3.2. Not shown are the RF excitation, slice selection, and pre-encoding steps. These are the same as those in conventional or FLASH imaging. The first two methods scan k-space as a series of alternating left to right and right to left raster lines. The "blips" for the G_y gradient are shown as impulse functions in the figure, but in reality they must have a finite width. Therefore, the reversal points at extreme k_x will not be an immediate jump to the next line but rather will consist of some sort of smoother transition. The spiral-scan starts at the center of k-space and spirals outward. Because these sequences require strong rapidly switching gradients, they

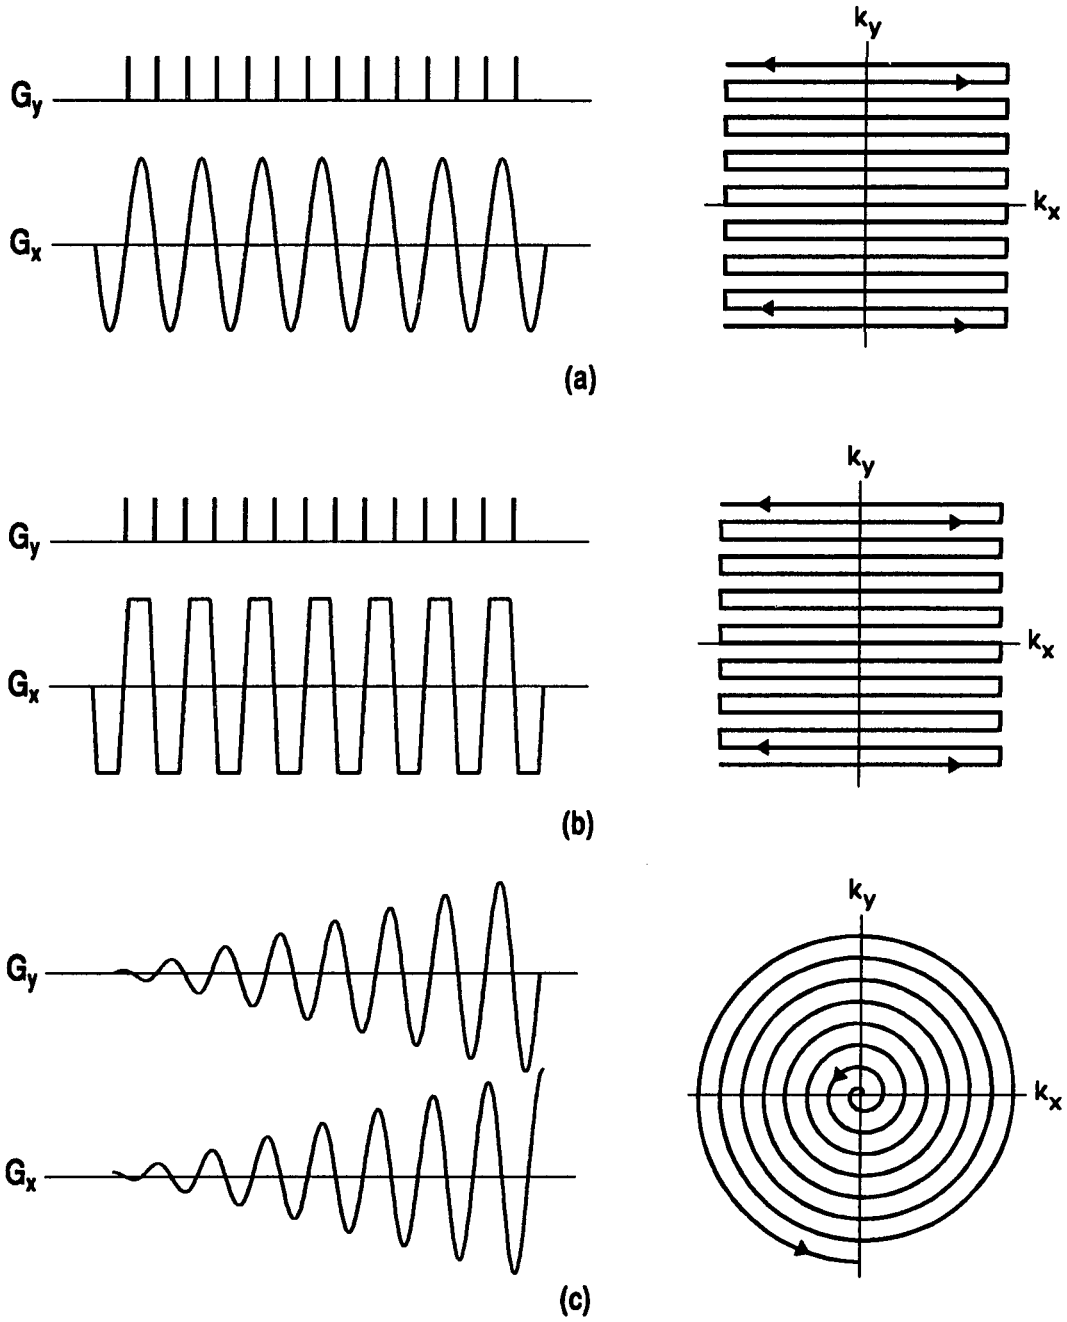

Figure 3.2 Gradient waveforms and associated k-space scan patterns for three single-shot methods. (a) blipped echo-planar, (b) modified blipped echo-planar, (c) spiral-scan.

have very stringent hardware requirements. The sinusoidal gradients require less gradient driving power than those approximating a square wave (trapezoidal gradients) if generated from a resonant tank circuit, but they can necessitate more complicated data sampling. This is discussed further, below. Barring reconstruction attributes, the relative merits of these scan techniques will not be presented here. Each of these methods have generated enough interest to justify their use in real-time MRI. Good discussions are available in references [14]-[15].

3.3 Factors affecting reconstruction of single-shot data

The blipped echo-planar and modified blipped echo-planar k-space trajectories produce data that lie on an approximately rectilinear grid and so 2D Fourier transformation can be used for reconstruction. The spiral-scan k-space trajectory produces data that lie on an approximately polar grid and so filtered backprojection can be used for reconstruction in this case. Figure 3.3 depicts the k-space sample locations for each of the scan methods when data sampling is performed uniformly in time. The resulting sample points fall on the ideal reconstruction grid for only the modified blipped echo-planar method. The samples for the blipped echo-planar scan do not exist on a rectilinear grid since they are further apart near the k_y axis than at the extremes. The sample points for the spiral-scan are close to lying on a polar grid, but the distance between the relative positions of the samples with respect to the origin changes with angle, and there is a shift in sample position when passing through the origin along a radial line. The simulations shown in figure 3.4 reveal the image errors that result when reconstruction proceeds as if the time-uniform samples actually existed on the ideal grids [16]. The image shown for the modified blipped echo-

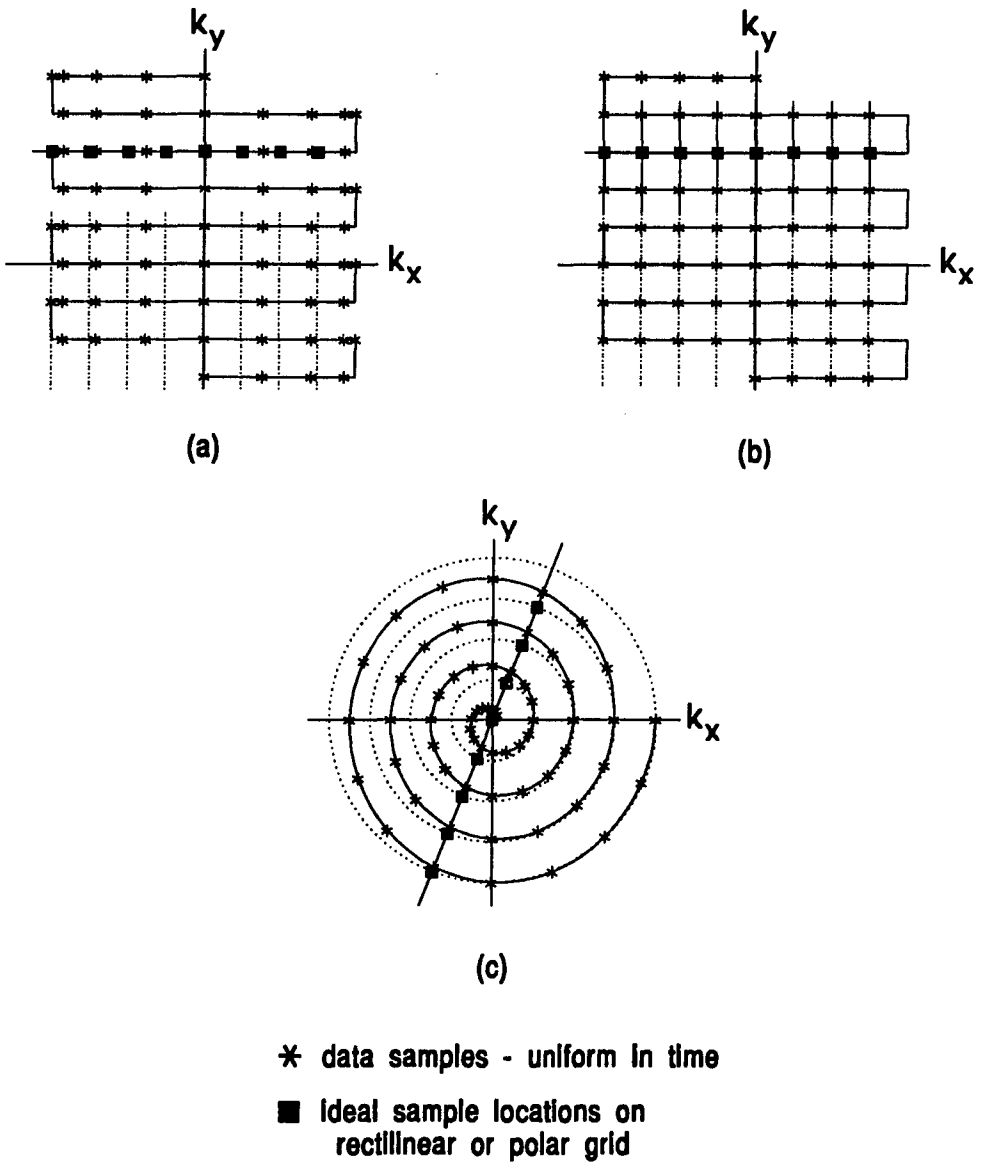

Figure 3.3 Uniform time sampling of the scan patterns (a) blipped echo-planar, (b) modified blipped echo-planar, (c) spiral-scan. Only part of the k-space trajectory is shown. One line with the correct sample locations is given in each case.

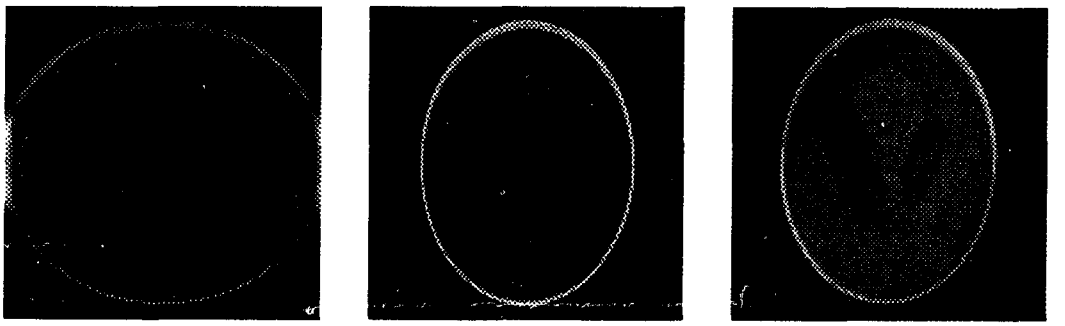

(a) (b) (c)

Figure 3.4 Simulation of scan-pattern error. (a) blipped echo-planar, (b) modified blipped echo-planar, (c) spiral-scan.

planar scan is perfect since in this simulation the data exist on the ideal rectilinear grid. The reconstructed blipped echo-planar image has unwanted ringing artifacts because of the sampling error. The reconstructed spiral-scan image shows some shading, but the quality is actually quite good. The artifacts present in the blipped echo-planar and spiral-scan images can be minimized by interpolating the collected data onto the correct grid before reconstruction. However, simulations show that a full 1D interpolation in the k_x direction is needed to generate an acceptable echo-planar image, and a full 1D interpolation in the radial direction is needed to improve the reconstructed spiral-scan image. Because this kind of interpolation is computationally expensive and would place a large burden on the reconstruction system, a better method is desired to obtain the correctly placed data . Non-linear time sampling could be used to produce the ideal data samples in the case of blipped echo-planar scanning. Implementation of this is straightforward (see chapter 5), but the receiver bandwidth must be increased to accommodate faster sampling near the k_x=0 frequencies. This results in lower signal-to-noise ratio (SNR). Also, the digitizer

speed must be increased. Still, this is probably the most practical technique to obtain the ideal data points for the blipped echo-planar scan. Non-linear time sampling does not help the spiral-scan, but, as evinced in figure 3.4, normal reconstruction produces good images that would be acceptable in most applications.

As data are collected successively in time, they must be stored in preparation for reconstruction. However, the single-shot techniques discussed here generate data samples that, if taken consecutively, are incorrectly ordered for reconstruction. For the raster-scan methods, alternating lines in k-space are scanned in opposite directions and so the data in every other line must be flipped before 2D Fourier transformation can occur. For the spiral-scan, time-successive data lie on different radial lines and neighboring samples on a radial line are separated in time by a full cycle of the spiral. The reorganization of this data into the correct reconstruction matrix is described in chapter 5.

3.4 Partial k-space scanning

It is possible to use the single-shot methods to scan only half of k-space in an excitation and then use conjugate synthesis to produce the rest of the k-space data. Additionally, k-space can be covered by multiple-shot tiling or interlacing. These place different demands on the reconstruction process and so are discussed separately.

3.4.1 Conjugate synthesis

The data collected in an MR experiment are in the Fourier domain (k-space) of the object being imaged. Since the object is real, it's Fourier transform is Hermitian [17]. Thus, the acquired k-space data should be Hermitian. A

common way to reduce the scan time in half is to acquire only half of k-space and synthesize the other half by taking the complex conjugate of the acquired data. This is known as conjugate synthesis [18]. Only raster-scan methods can be used in this type of acquisition since the spiral-scan can not acquire just a symmetric half of k-space. Note that the acquired half of k-space is arbitrary (top, bottom, left, or right) and that the ordering of this data into the correct reconstruction matrix is dependent on which half is acquired. The reconstruction system is responsible for synthesizing the rest of the data, reordering all data, and then performing the reconstruction on the full data matrix. This is a straightforward process on a conventional computer system, but it calls for more flexibility on the part of a dedicated real-time reconstruction system. The hardware discussed in chapter 5 is able to implement conjugate synthesis on the fly.

3.4.2 Multiple-shot tiling

A single-shot acquisition can be used to acquire only part of k-space during one excitation. Typically, two or four shots are used to "tile" k-space in some manner. For instance, each of two shots could be used to acquire separate halves of k-space. Although, these techniques increase the scan time, they are useful for increasing the spatial resolution of the final echo-planar image on systems with more conventional hardware since less gradient strength is required [19]. Again, the reconstruction system has the responsibility of properly ordering the data and then performing the reconstruction on the full data matrix.

3.4.3 Multiple-shot interlacing

Any of the single-shot techniques can also be used in an interlaced fashion where k-space is filled by interleaving multiple shots. As above, these

techniques can be used to implement higher resolution echo-planar imaging on more conventional systems. This is particularly useful for the spiral-scan where interleaving two shots doubles the final resolution in the radial direction [20]. The reconstruction requirements for interlacing are the same as above although the data ordering can become complicated.

CHAPTER 4

CONSIDERATIONS FOR REAL-TIME MRI RECONSTRUCTION SYSTEMS

Video-rate reconstruction of MR data is a difficult task. In addition to the intense arithmetic processing required for video-rate Fourier transformation or filtered backprojection, a complete real-time system must possess, at a minimum, the ability to continuously input and output the MR data stream. There have been many developments in performing the reconstruction algorithms at very high speed. Some analog-electronic or analog-optical systems can process at the needed rates but introduce error into the reconstruction and are limited in the accuracy they can achieve. Digital devices, on the other hand, introduce little error into the process, and can achieve any desired accuracy by simply increasing the data word length. Since conventional digital computers can not numerically process at the rates required for video-rate MR reconstruction, dedicated digital-electronic systems must be used. This chapter discusses considerations associated with video-rate Fourier transformation and filtered backprojection and previous developments in these areas. For comparison purposes, the acquisition of k-space data will be represented by a N x N matrix. For data existing on a rectilinear grid this means that k-space is acquired with N raster lines with N samples per line. For data existing on a polar grid this means that k-space is acquired with N radial lines with N samples per line. The values of N considered here are 64, 128, and 256.

4.1 Analog methods for Fourier transformation

When MR data exist on a rectilinear grid, a 2D Fourier transform is the appropriate reconstruction algorithm. Various optical systems can implement continuous Fourier transforms [21]. One optical approach for MR reconstruction is based on film recording of MR data [22]. Video-rate reconstruction, however, requires real-time input, so this system is not appropriate for the goals given here. A video-rate 2D Fourier transform system for reconstructing MR images using a spatial light modulator as the input device and a lens to perform the 2D transform has been demonstrated [23]. The quality of the reconstructions was poor, however, and so the system was deemed inadequate for use in MRI.

When MR data exist on a polar grid, filtered backprojection is the appropriate reconstruction algorithm. Recall that the filtering step is performed by multiplying the 1D data in a radial line with the 1D filter function, followed by 1D Fourier transformation. One system that accomplishes the 1D Fourier transform in real-time is based on the chirp Fourier transform algorithm and acousto-optic technology [24]. Although this approach works, the system is complicated, and, like most analog systems, is limited in the accuracy it can achieve. Other analog filtering systems have been studied for use in CT, but because CT data exist as real projections, the filtering operation is realized by convolution and is not appropriate for data acquired in MR.

4.2 Analog methods for backprojection

Assuming that filtered projection data have been obtained, the final step in the reconstruction of projection data is backprojection. Recall that this step is performed by finding the contributions of each filtered projection to each image pixel and then summing these contributions to obtain the final image. In effect,

each filtered projection is smeared back through the 2D image space, and the sum of these smearings is the reconstructed object. An analog-optical system has been constructed to perform this operation at video-rates [25]. The filtered projection data are written on a display as an intensity modulated line of light. A cylindrical lens smears this line of light perpendicular to the data direction and the resulting pattern is rotated by an image rotation prism and imaged onto a video camera. As a new filtered projection is written on the display, the prism is rotated to the corresponding projection angle, and the video camera integrates the resulting backprojection with all previous projection angles. The system is synchronous with the video frame rate; therefore, each new display frame is a new reconstructed image. This system was shown to work quite well and had a nominal resolution of approximately 128 x 128 in the image plane. When this system was designed, it was almost impossible to perform video-rate backprojection at this resolution with digital-electronic devices. However, it is now feasible to do so, and since the accuracy of the optical backprojector is limited, a digital-electronic system is preferred.

4.3 Limitations of conventional computer systems

There are two major factors that determine why a conventional serial computer can not be used in many real-time systems: raw computation speed and input/output (I/O) bandwidth. Both are vital for real-time MRI image reconstruction. Personal computers can "number crunch" at 1 million floating point operations per second (MFlops) or less, engineering workstations at 20 MFlops or less, and high-end minicomputers and supercomputers at tens to hundreds of MFlops. Because the computational requirements of both real-time Fourier transformation and filtered backprojection can reach hundreds or even

thousands of MFlops, only a few very high-powered computers possess the ability to compute at speeds fast enough for completing an image reconstruction in 33 ms; thus, there is a need for specialized processors and processor configurations in a real-time reconstruction system. Still, this does not mean that a system that is capable of numerically processing at these rates can actually *sustain* these rates. Real-time systems must have the ability to simultaneously receive a continuous data stream, process that stream, and send any resultant data to other systems for further processing or display. This requires not only fast processors but also a fast shared memory architecture and control logic. Since conventional computer systems are not architecturally configured to handle this type of concurrent I/O and arithmetic processing, a dedicated approach must be used.

4.4 Considerations for real-time 2D Fourier transformation

Digital computation of a Fourier transform requires the use of a discrete Fourier transform (DFT). Direct computation of the 2D DFT on a computer requires on the order of N^3 complex multiplications and additions for an N x N data matrix. A fast Fourier transform (FFT) algorithm can be used instead of direct calculation.[26] In general, a 2D FFT will reduce the number of complex multiplications and additions to $N^2\log_2 N^2$. This is a substantial savings that is crucial to the development of a digital-electronic real-time Fourier transformation system.

Digital implementations of the FFT are complicated by the nature of the algorithm. The algorithm causes naturally ordered data to become bit-reversed after the transformation. That is, the binary representation of the sample location in the output data stream is flipped over so that the ordering of the data

is not sequential. This type of addressing is not inconsequential when trying to execute the FFT at very high speeds.

Still, the main concern of implementing a 2D FFT is the amount of arithmetic processing that needs to be done. Given below are approximate computational rates (total number of multiplications and additions per second) needed for performing various complex 2D FFT's at video-rates.

Table 4.1

Video-rate Fourier transformation	
Matrix size	Computational rate (Millions of operations per second)
64 x 64	9
128 x 128	42
256 x 256	190

As discussed previously, most conventional computers can not numerically process at the rates given in the table above and so a dedicated approach must be used.

Currently, reconstruction of standard MR data (2D Fourier method) is done by performing the 2D FFT on an array processor. These processors typically complete a 256 x 256 complex FFT in the range from 150 ms to 2 seconds. There has been some work in configuring these types of processors to perform continuous real-time MRI reconstruction [27] but these systems fall short of the desired goal here of video-rate MRI reconstruction.

Early research into parallel processing of FFT's was limited by slow arithmetic processors and insufficient memory sizes and speeds [28]. The availability of multiprocessor supercomputers brought forth developments in efficient vectorized FFT algorithms [29], [30], and improvements in digital-

electronic processors and memory extended these ideas to dedicated hardware designs [31], [32]. Some of the most interesting developments came in the form of algorithms optimized for VLSI implementations [33], [34]. These types of algorithms have made possible an "FFT on a chip". One such integrated circuit (A41102, Austek Inc.) can perform a 1D FFT on a data stream of up to 256 complex samples without external memory or arithmetic processors. The IC takes care of all FFT addressing, processing, and storage on chip. Two of these IC's can be configured to perform video-rate FFT's at any of the above matrix sizes. In addition, the A41102 possesses other features that makes it ideal for use in a real-time MR reconstruction system. The Fourier processor presented in the next chapter is built around these IC's.

4.5 Considerations for real-time filtered backprojection

4.5.1 Filter

Digital filtering of spiral-scan data is straightforward since the data already exist in the spatial frequency domain (k-space). Each radial line need only be multiplied by the 1D filter function and then Fourier transformed to obtain the filtered projection. The Fourier transform step is performed with a 1D FFT. In general, a 1D FFT reduces the number of complex multiplications and additions needed in a 1D DFT to $N\log_2 N$. An acquisition with N projections and N samples per projection requires N 1D FFT's of length N to be performed. The total number of complex multiplications and additions for this situation is $N^2\log_2 N$. The operational rates required for filtering, then, are half those required for a full 2D Fourier reconstruction. These rates are still substantial and so the same considerations for the computation of video-rate 2D FFT's also apply here (table 4-1).

4.5.2 Backprojection

Recall that continuous backprojection is accomplished by using equation (2-13). With discrete data, however, the values of p ($p=x\cos\theta + y\sin\theta$) for which the filtered projection data are known may not be the values that equation (2-13) requires. Interpolation is then needed to obtain the correct filtered sample values to be backprojected. Linear interpolation is usually adequate and will be used as part of the backprojection analysis given below.

There are two common methods of implementing the backprojection operation in a discrete space. In the first, the filtered projection data (stored in memory) are indexed and the contribution of each data point to the image matrix is calculated. In the second, the image matrix is indexed and the contribution of the filtered projection data to each image pixel is calculated. Although both produce equivalent reconstructions, the latter is discussed here. Assume acquisition of N projection angles, θ_n, with N samples per filtered projection, $Q_n(m)$, where n indexes the projection angles and m indexes the samples within a filtered projection. The image matrix is N x N and is denoted by f(i,j) where i indexes the rows of the image matrix and j indexes the columns of the image matrix. Here, i,j,m,n run from one to N. The reconstruction sequence sums the appropriate contribution from each filtered projection to one image pixel, and then repeats the process for each image pixel. The basic steps are as follows:

(1) Find the sample position in a filtered projection.
for a given i,j, and n, calculate $p = i\cos\theta_n + j\sin\theta_n$.

(2) Interpolate to find the value of the sample.
calculate $\Delta f_{ij} = Q_n[\text{int}(p)]\cdot(1-(p-\text{int}(p))) + Q_n[\text{int}(p)+1]\cdot(p-\text{int}(p))$
where int(p) denotes the integer portion of p.

(3) Add the sample value to the image pixel.
sum Δf_{ij} with the previous value of f(i,j).

There are N^3 of these fundamental calculations required in a reconstruction. Because MR data are complex, the steps must be performed for both the real and imaginary parts of the data. However, values calculated in step 1 are the same for the real and imaginary parts and need to be calculated only once. Table 4.2 shows the approximate operational rates for video-rate backprojection of complex data.

Table 4.2

Video-rate backprojection	
Matrix size	Computational rate (Millions of operations per second)
64 x 64	111
128 x 128	890
256 x 256	7118

Clearly, serial computation cannot be used for calculation at these rates. Digital hardware for video-rate filtered backprojection has been investigated using massively parallel computers [35]. An interesting aspect of this research is the design of VLSI components tailored for backprojection. Dedicated parallel hardware for backprojection has also been investigated [36]. The digital-

electronic backprojector discussed in chapter 5 uses dedicated parallel processors to perform the backprojection step.

4.6 Reduction in backprojection processing

To ease the hardware implementation of video-rate backprojection, it is possible to reduce the rates given in table 4.2 by nearly an order of magnitude. This is done by precalculating the values in step 1 and reducing step 2 to choosing the nearest sample. These procedures are discussed below. No reductions can be made in the computation of step 3 since this is the actual backprojection step; the "smearing" of filtered projections into the 2D image space. Parallel processing will still be required to implement video-rate backprojection (at least for the two highest resolutions), but the hardware layout becomes much simpler with the computational reductions.

The linear interpolation in step 2 requires two multiplications and five additions to be performed in each iteration for both real and imaginary data. An interpolation could be constructed in hardware as a module. After the array indices are calculated, the two multiplications could be done in parallel and the products pipelined to an adder. A parallel processing system would need as many modules as the parallelism required. Hardware can be built to implement this arrangement, but the complexity of the system is greatly increased.

Simulations have shown that interpolation may not be necessary. That is, step 2 above can be reduced to simply picking the sample nearest to the desired one. Figure 4.1a shows a 128 x 128 image (Shepp-Logan phantom) reconstructed with filtered backprojection using linear interpolation. Figure 4.1b shows this same reconstruction with no interpolation (choosing the nearest neighbor). The differences between these two images are slight. Figure 4.1c

and 4.1d show the same reconstructions with gaussian noise added to the projection data to produce an image SNR of around 10:1. The images reveal that the reconstruction error due to no interpolation is far less than that inherent in the image.

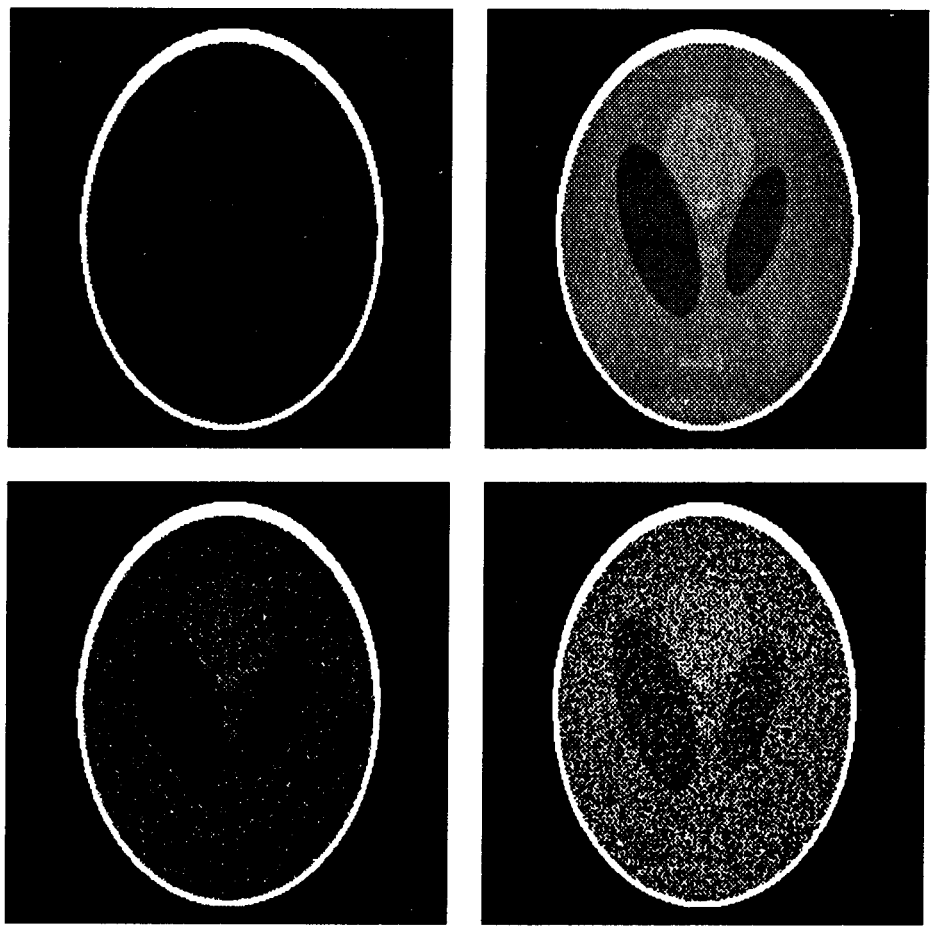

Figure 4.1 Simulation of interpolation error in backprojection. (a) linear interpolation, (b) nearest neighbor, (c) linear interpolation with noise, (d) nearest neighbor with noise.

Step 1 in the algorithm describes the calculation of the sample locations within a filtered projection. Two multiplications and one addition are required for this step (sample locations for real and imaginary data are the same). If no interpolation is performed, then this step must additionally take on the task of finding the nearest neighbor to the calculated p (either int(p) or int(p)+1). A hardware module to do this calculation could be built. However, if no interpolation is needed, it is possible to eliminate all computations by simply precalculating the appropriate filtered projection sample locations needed for each image pixel. These values represent locations in the filtered projection data and can be stored in a lookup table. The amount of memory required can be very large since N^3 addresses need to be stored. However, the computational savings and hardware simplification more than offset this inconvenience.

If the above reductions are taken into account, then the only computation needed for backprojection is that of step 3: N^3 complex additions must be completed in 33 ms. Table 4.3 shows the approximate operational rates (N^3/33 ms) and the times allotted for one serial operation (33 ms/N^3) to complete this step in 33 ms.

Table 4.3

Video-rate backprojection (accumulation only)		
Matrix size	Computational rate (millions of operations per second)	Accumulation rate (nanoseconds per operation)
64 x 64	16	64
128 x 128	127	8
256 x 256	1000	1

Because of the computational reductions, the operational rates have been lowered by nearly an order of magnitude. It is seen that serial computation will suffice for a matrix size of 64 x 64 because the time allotted per pixel for this resolution is an achievable operational speed with most adders. However, the time allotted for serial computation of the two larger matrices is prohibitively short. Some degree of parallelism must be used. The number and arrangement of parallel processors to perform video-rate backprojection is discussed in the next chapter.

CHAPTER 5

A REAL-TIME MRI RECONSTRUCTION SYSTEM

Conventional digital computers are not configured for video-rate reconstruction of MR data and dedicated systems have been shown to be lacking in either processing power or accuracy. This chapter presents a first-order design of a video-rate digital-electronic MRI reconstruction system that is adequate for reconstructing MR data acquired from nearly any technique. This system is configured so that any acquisition rate can be used, but reconstruction always occurs at video rates. Note that the terms "real-time" and "video-rate" are used interchangebly in this chapter. However, the reconstruction system always operates at video-rates. The acquisition rate may be slower than this and still be considered real-time. The system consists of a digitizer to interface with an MR scanner, a Fourier processor to perform either 2D Fourier transformation or 1D filtering, and a backprojector to perform the backprojection step on filtered data. The complete system could be constructed on a few PC boards for any bus architecture and controlled by a host microprocessor. Details of all timing, control, and support logic are not discussed here. All diagrams shown of the system contain a "timing and control" bar to represent these. The diagrams of the digitizer and Fourier processor depict the case of a 64 x 64 data matrix. It is straightforward to handle matrix sizes up to 256 x 256 and considerations for all common MR data matrix sizes are given. The memory sizes throughout this chapter will be denoted by *A*K x *B* where *A*K is the total number of memory locations and *B* is the bit depth for each location. For example, 4K x 16 denotes 4096 memory locations with 16 bits per location.

5.1 Digitizer

5.1.1 Data Conversion

Figure 5.1 shows the basic block diagram of the digitizer. MR data enters the digitizer as real and imaginary analog signals that are tapped off an MRI scanner. Each signal follows an identical conversion path so that the real and imaginary parts of a sample are kept in sync. A strobe from the scanner starts the conversion process. For data acquired with constant gradients (conventional imaging or modified blipped echo-planar), a time-uniform strobe is used to initiate conversion. However, recall from chapter 3 that non-linear time sampling may be needed for data acquired with sinusoidal gradients. If the external strobe signal is uniform in time, a non-uniform time signal must be generated. The appropriate non-linear time sequence can be calculated from the scan-pattern and the number of samples desired. A simple method for producing this clock is shown in figure 5.2. First, the normal time-uniform sample clock is divided (frequency multiplied) by some factor S. S dictates how well the sample clock approximates the exact non-linear pattern. The high frequency clock is then masked by the known non-linear pattern stored in an SN X 1 memory (N is the number of desired samples in one line of data), to provide clock pulses at the correct times. The bit sequence read out of this memory is the appropriate non-linear time clock. Only N bits in this sequence are non-zero with the rest used to mark increments of time. The larger S is, the smaller these increments of time are, and the closer the approximation is to the continuos non-linear time pattern. As an example, assume an echo-planar imaging experiment with a desired data matrix of 64 x 64 complex samples. The period of the time-uniform clock used to digitize this matrix in 33 ms (video rate) is about 8μs. If this clock

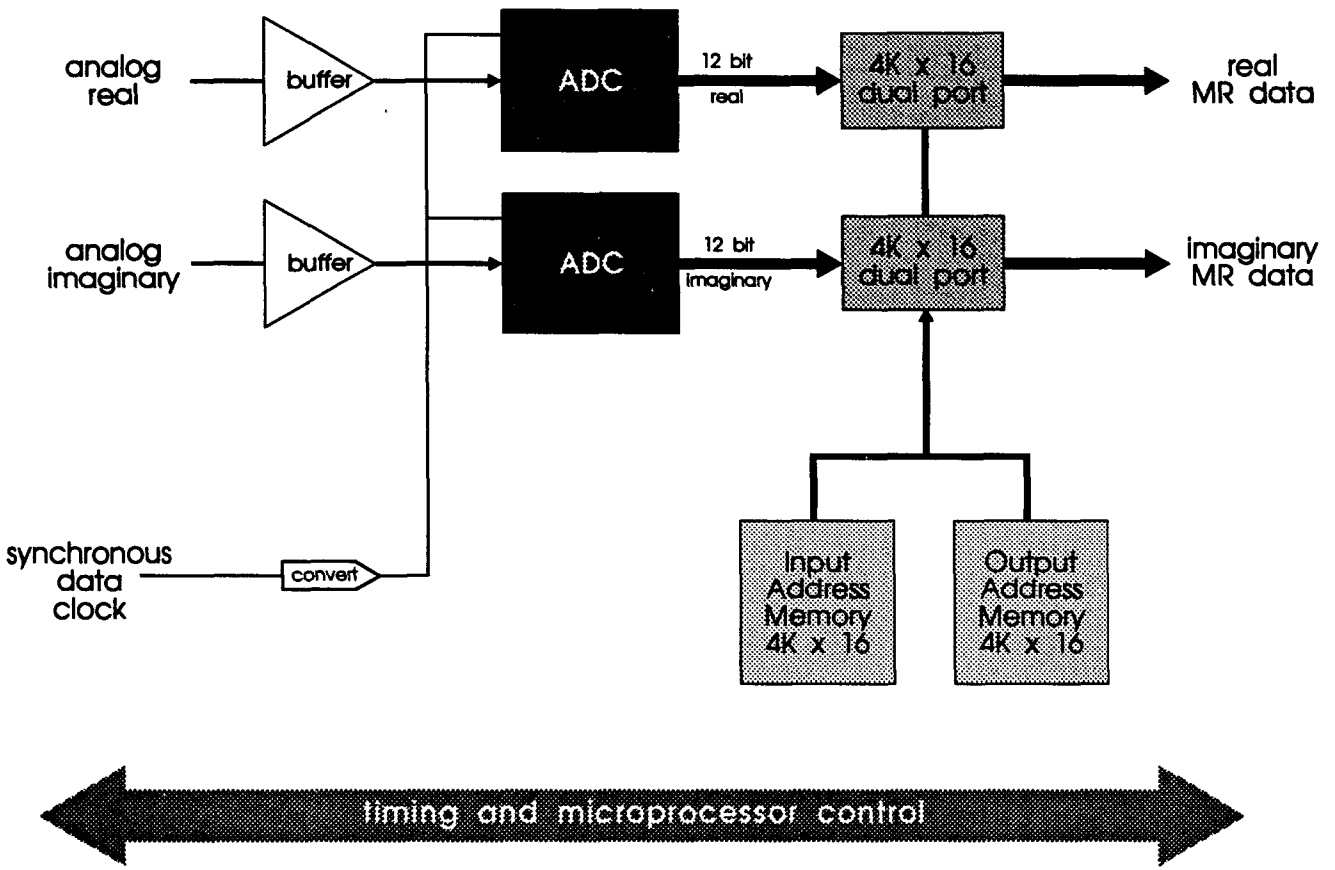

Figure 5.1 The digitizer used to convert the analog MR signal to digital form. Real and imaginary data are converted simultaneously and stored in dual-port memories for processing. The address memories provide the ability to easily reorder the data arbitrarily.

is frequency multiplied by a factor S=8, it's period becomes about 1μs. The data read out of the 512 x 1 memory is an approximation to the correct non-linear sequence. Only 64 of the memory cells have high bits while the rest are zero. The location of the high bits determines when a conversion should take place. The resulting digitized samples will be close to lying on a perfect rectangular grid.

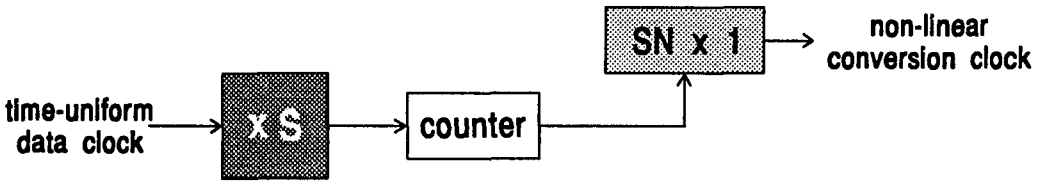

Figure 5.2 Generation of a non-linear time clock.

Each analog-to-digital converter (ADC) has 16 bit resolution and must be able to keep up with video-rate acquisition. Table 5.1 shows the maximum digitization times needed for the common data matrices.

Table 5.1

ADC sampling requirements	
Number of points acquired	maximum digitization time (microseconds)
64 x 64	8
128 x 128	2
256 x 256	0.50

The actual digitization times should be at least half those given in the above table to ensure the data will be converted properly and allowed to settle before the next conversion. These conversion times can be achieved with monolithic or hybrid ADC's that contain all the necessary sample-and-hold and reference

circuitry. Appropriate devices include Analog Devices HAS 1204 and AD872, and Signal Processing Technologies SPT7912.

5.1.2 Digitizer-interface buffer

After a sample is converted to a digital word, it is written to a memory buffer which provides the interface between the digitizer and the rest of the reconstruction system. The real and imaginary parts of the sample are written to separate memories. The digitizer-interface memory serves as an asynchronous buffer between the digitizer and the reconstruction system (Fourier processor and backprojector) so that each can be operated at different rates. Thus, the reconstruction system can be operated at video-rate speeds while the acquisition system can be operated at slower speeds if desired. The digitizer-interface buffer can be either dual-ported or double-buffered. Either method serves as an asynchronous buffer. The dual-port configuration allows data to be written to the digitizer-interface memory while data is being read out of the same memory for reconstruction. The double-buffered configuration allows data to be written to one bank of memory while data is being read out of another bank for reconstruction.

In the dual-port configuration (shown in figure 5.1), the reconstruction system processes at video rates whatever data is present in the memory. While new data is being written into the memory at an arbitrary acquisition rate, stored data is being read from the memory at video rates (a full image every 33 ms). If acquisition also occurs at video-rates, then data written to the memory is immediately read out so that new data is always being reconstructed. If, however, data are acquired more slowly than this, some of the data being reconstructed will be the same as the previous frame while some will be

different. This is referred to as a "sliding window" technique where the width of the reconstruction window equals the time to collect a full data frame but the window slides over 33 ms after each reconstructed frame. This type of approach has been implemented on one system with good results [37]. However, in that system, reconstruction is performed in about 120 ms (6 frames/second) for a 128 x 128 data matrix. Because the displayable output is updated once every 5 or 6 video frames, a perceptible "jumping" of the image may be seen with certain degrees of motion. Also, data acquisition at rates faster than the image reconstruction rate (6 frames/second) is not possible. The system proposed here differs in that it can reconstruct, in 33 ms, data matrices at resolutions up to 256 x 256. The displayable output will be updated at frame rates with some new data in each frame. This may result in a perceptibly smoother motion.

The double-buffering scheme uses two banks of single-port memory. Each bank has separate memory for the real and imaginary parts of the data. While one bank is being written with new data, the other is being read out for reconstruction. The reading and writing operations are alternated between banks every time a whole frame of data has been acquired. If the data are acquired at video-rates, this alternation occurs at video-rates so that each display frame will contain new data. However, unlike the sliding window technique, new data is not added *during* each image reconstruction, but *before* each image reconstruction. If the data are acquired slower than video speeds, the data in one bank is repeatedly read out for reconstruction and so not every display frame will contain new data. When acquired data fill the other bank, the reading and writing operations are alternated between banks. Note that at this point, the reconstruction of a frame may not be completed and the next sample read out for reconstruction could be anywhere in the full bank. Thus, there may

be perceptible "jumps" when the reconstruction bank is switched. This may result in a perceptibly less smooth display.

Table 5.2 shows the memory requirements for both dual porting and double buffering of the digitizer-interface memory. Because the memory sizes are small, SRAM can be used. This eliminates the need for refresh circuitry. Integrated SRAM circuits are available for all the specified memory sizes although combining smaller memories may be necessary for the dual-port buffer. Many dual-port devices contain arbitration logic to handle memory contention (trying to read and write the same memory cell simultaneously), so external arbitration logic should not be needed. The actual speed of the memories should be less than the values given in the table to ensure adequate read/write time. These memory access times are slow and therefore memory speed is not a limiting factor for this buffer. Appropriate devices are manufactured by many companies including Integrated Device Technologies (IDT7024, IDT7025, IDT8MP656, IDT8MP624) and Cypress Inc. (CY7C168, CY7C194).

Table 5.2

Digitizer-interface buffer memory requirements		
number of acquired points	required dual-port memory	required double-buffer memory
64 x 64	4K x 16 dual-port (2 @ 8 µs)	4K x 16 single-port (4 @ 8 µs)
128 x 128	16K x 16 dual-port (2 @ 2 µs)	16K x 16 single-port (4 @ 2 µs)
256 x 256	64K x 16 dual-port (2 @ 500 ns)	64K x 16 single-port (4 @ 500 ns)

It appears that dual-ported digitizer-interface memory will be superior to double-buffered memory. Each new video frame will have new data which will

give a smoother looking output for longer data acquisitions. Dual-porting also requires less memory and control logic. Note that with a sliding-window technique, the output is scan-pattern dependent. High spatial frequencies are collected first for a raster-scan method, and therefore the first to be included in a newly reconstructed image. In contrast, low spatial frequencies are collected first in the spiral-scan. This may cause a perceivable difference in the output. Also, any effects will be manifested differently for backprojection than for Fourier transformation. It will be interesting to see if these dependencies produce any undesirable effects.

5.1.3 Addressing of the digitizer-interface buffer

The interface memory needs to be addressed so that properly ordered data enter the reconstruction system. Most conventional acquisition sequences produce data that needs no reorganization. However, data from ultra-fast techniques usually require reordering. For example, both echo-planar and spiral-scan must be reordered prior to reconstruction, but the reordering sequence is different for the two. The natural order of digitized echo-planar data is reversed in every other line since alternating lines in k-space are scanned in opposite directions. The natural ordering of digitized spiral-scan data is more complex. The first point collected is at the center of k-space and is used as the center sample in all radial lines. Time-consecutive data samples lie on different radial lines, and neighboring samples on a radial line are separated in time by a full cycle of the spiral. Although a complicated counter arrangement could be used to generate addresses that reorder the data for either scan method, it is easier and more flexible to store the known reordered addresses in a dedicated address memory (lookup table). Data samples can be written to the buffer

memory in a certain order and read out in the proper order by stepping through the lookup table. Figure 5.1 shows both read and write addressing controlled by separate address memories. All reordering could be done by the output address memory or some combination of reordering could be shared between input and output address memories. Each address memory needs to be large enough to store addresses for all sample points. The size and speed of this memory is the same as given in table 5.3 although the bit-depth required for each matrix size is different (12, 14, and 16). It is possible that this memory could be ROM since the addresses do not change during acquisition. However, the ROM's would have to be removed and rewritten when the acquisition method changes. At these memory sizes and speeds, it is more desirable to use RAM instead since it can be loaded directly from the bus. The use of address memories makes for a flexible system that can be easily configured to handle different data acquisition methods.

5.2 Fourier processor

Figure 5.3 shows a system that can perform a 2D FFT on matrices up to 256 x 256 at video rates (for echo-planar reconstruction), or up to 256 1D FFT's at video rates (for spiral-scan filtering). This is called the Fourier processor. The processing hardware is based on a special FFT VLSI integrated circuit made by Austek Inc. (A41102). This fixed integer IC can perform a 1D FFT with up to 256 24-bit complex (24-bit real, 24-bit imaginary) samples without external arithmetic or memory components (the FFT sample size must be a power of 2). A functional diagram of this IC is shown in figure 5.4. Data enters the chip as a

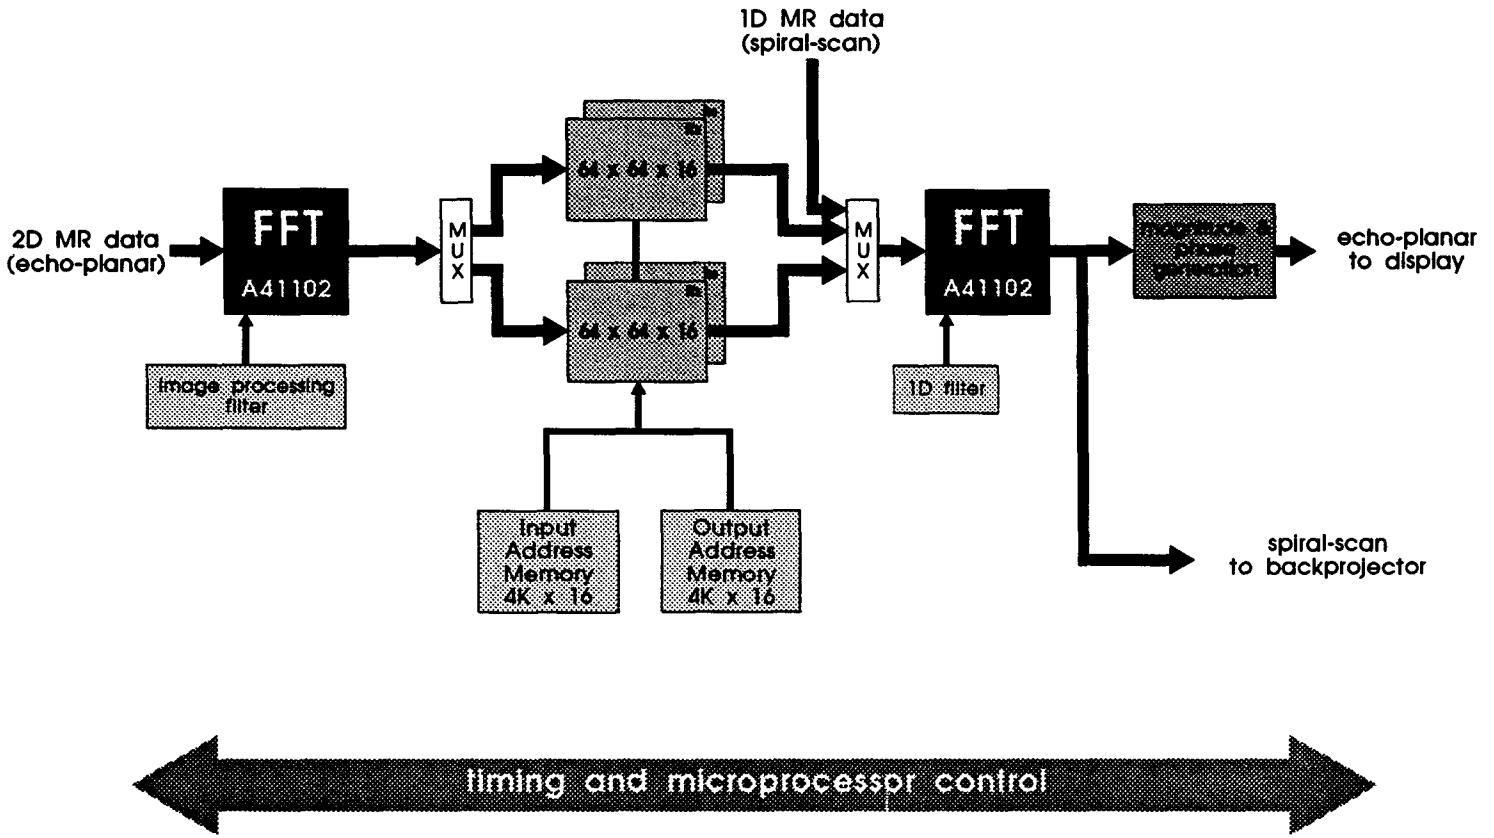

Figure 5.3 Fourier processor. Echo-planar data enters the processor, is transformed, and converted to magnitude and phase for display. Spiral-scan data enters the processor before the second FFT chip, is filtered, and sent to the backprojector.

serial stream and exits as a serial stream. The chip has three bi-directional ports; input, output, and one that enters an on-chip complex multiplier. The multiplier can be used at either the input or output, or it can be disabled. The chip also has conjugators at both the input and output. At it's maximum continuous speed, the A41102 can perform a 1D FFT on 256 complex data in 102.4 μs, on 128 complex data in 51.2 μs, or on 64 complex data in 25.6 μs. This is an effective operational rate of over 100 Mops. The word length for the transforms considered here is 16 bits but overflow from signal growth during the transform is prevented by enabling scalers on the chip. Recall that the FFT algorithm produces bit-reversed data from naturally ordered data. The A41102 also produces bit-reversed data in this manner but can be configured to produce naturally ordered data from bit-reversed input data. The A41102 can be configured to perform the proper length FFT, to enable or disable the multiplier or conjugators, and to set up overflow prevention through a microprocessor compatible port.

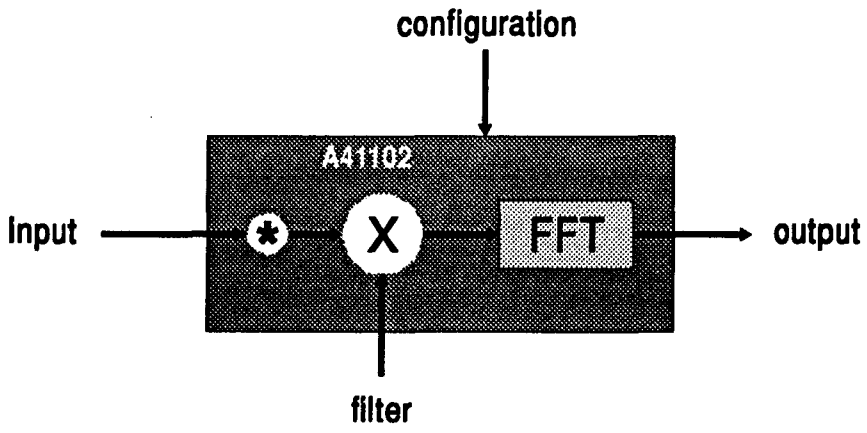

Figure 5.4 Functional diagram of the A41102 frequency domain processor. The data can bypass the conjugator and/or multiplier and feed directly into the FFT module.

5.2.1 2D Fourier reconstruction

To perform a 2D FFT, two A41102's are cascaded on either side of a double-buffered memory queue called the Fourier processor buffer. Data is read from the digitizer-interface memory (properly reordered) and enters the Fourier processor as a sequence of 1D rows of the data matrix. Each row is transformed by the first FFT chip and then written as rows to one of the memory banks of the Fourier processor buffer. Once transformed data fills this bank, it is read out as a sequence of 1D columns and Fourier transformed by the second FFT chip. Two banks of memory keep this process continuous. While one bank is being written to, the other is being read from. These functions alternate between banks every frame. This process of first transforming the rows and then transforming the columns (FFT rows, FFT columns) produces the 2D FFT.

Although a 2D FFT produces a reconstructed image from data existing on a rectilinear grid, the data from the second A41102 (appearing as a sequence of columns) is a bit-reversed version of the reconstructed image since both rows and columns have been bit-reversed by the FFT operations. To remedy this, bit-reversed addressing of the Fourier processor buffer can be used to reorder the bit-reversed data. The first A41102 is configured to perform a natural to bit-reversed transform and the second A41102 is configured to perform a bit-reversed to natural transform. Transformed rows from the first FFT can be written to the Fourier processor buffer in bit-reversed fashion thereby properly ordering the data in the rows of the matrix. Then, the columns of this memory can be read out in bit-reversed order and transformed to produce properly ordered data in the columns of the matrix. Thus, the data stream from the second A41102 is the correct reconstructed image appearing as consecutive columns of the image. (The data stream can also appear as consecutive rows of

the image if the first A41102 transforms columns and the second A41102 transforms rows.) Bit-reversed addressing can be implemented in a number of ways. The addresses can be hard wired for minimum cost and complexity, but it is difficult to handle multiple transform lengths in this way. Better solutions are programmable logic or lookup table memories since they can be configured easily for any transform length. Figure 5.3 shows dedicated address memories being used to look up the proper bit-reversed addresses. The size and speed of these address memories are the same as those with the digitizer-interface memory (table 5.2).

The reconstructed image appears as successive columns out of the Fourier processor. Since the data are still complex at this point, a magnitude (or phase) calculation can be performed on each sample prior to display. There are single chip implementations of this operation. One IC made by Plessey Inc. (PDSP 16330) converts 16 bit rectangular coordinates to 16 bit polar coordinates (complex to phasor) in 100 ns. This speed is more than adequate for the transform sizes considered here. The resulting magnitude or phase data can then be sent directly to a frame buffer on a dedicated display subsystem. There are many board-level display devices available to receive data and present it at video-rates (Datacube's MAXVideo, PsiTech's VME400).

5.2.2 Projection filtering

The Fourier processor can be used to filter data acquired from a projection imaging method or from a spiral-scan method, prior to backprojection. Recall that since the data exist as radial lines in k-space, the filtering operation is accomplished by multiplying one line of data with the 1D filter function, followed by 1D Fourier transformation. The data stored in the digitizer-interface memory

are read in a reordered and bit-reversed fashion directly into the second A41102 of the Fourier processor. The data enter the multiplier within the chip as a sequence of 1D radial lines. Each line of data is multiplied by the 1D filter function which is read out of a separate 1D memory in sync with the data. The filtered data are then 1D Fourier transformed to obtain filtered-projection data in the natural order. This data can then be sent to the digital-electronic backprojector.

5.3 Backprojector

Filtered-projection data leave the Fourier processor as a serial stream and are stored in another buffer memory for backprojection. As with the digitizer-interface memory, this filtered-projection memory is dual-ported. The real and imaginary parts of the filtered projections are stored in separate dual-port memories (see table 5.2 for specifications).

5.3.1 Serial backprojection

As discussed previously, the filtered backprojection algorithm can be reduced to choosing the appropriate samples from each filtered projection and summing them to generate one image pixel. This is done for each image pixel to create the final image. Because the data are complex, real and imaginary parts must be backprojected separately. The backprojection must incorporate some parallelism for the higher resolution data matrices but none is required for a resolution of 64 x 64. At 64 x 64, filtered-projection samples are read from the filtered-projection memory and accumulated in a single adder to form an image pixel (one for real, one for imaginary). Figure 5.5 shows the backprojector for a resolution of 64 x 64. Recall that the addresses of the samples appropriate for each image pixel are stored in an address memory. For a data matrix of

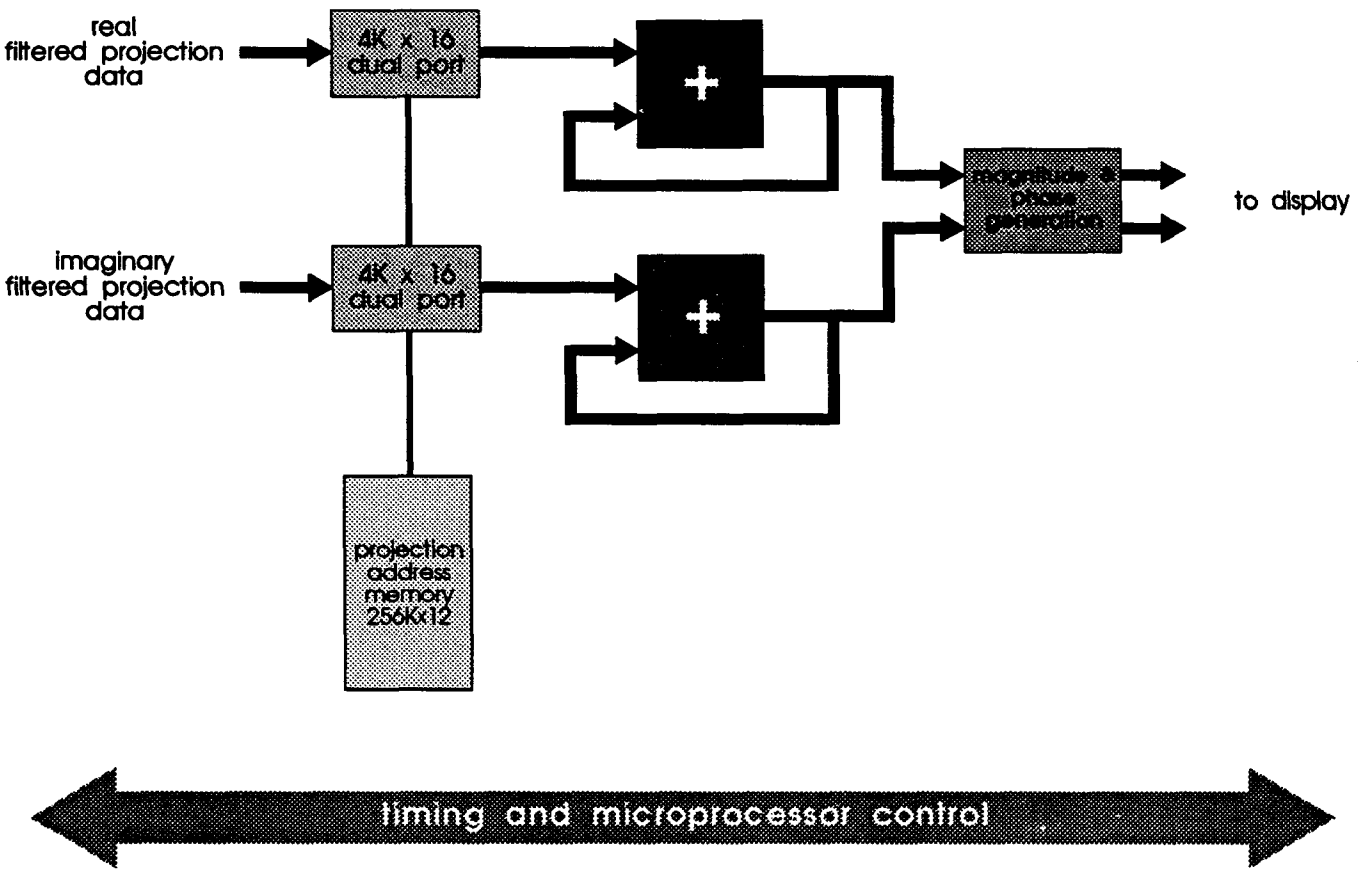

Figure 5.5 Backprojector. Each image pixel is formed through the summation of the appropriate projection data, is converted to magnitude and phase, and sent out for display. Separate channels are needed for backprojection of both real and imaginary data.

64 x 64, this memory must store 256K words at 12 bits per word and must deliver a new address every 126 ns, which is also the accumulation time. This memory is large; however, since the addresses are read sequentially, a high-density video RAM (VRAM) can be used for this memory. Appropriate VRAMS include Texas Instruments TMS44251 and Hitachi HM534251. The projection address memory is simply sequenced to deliver addresses of the filtered-projection memory. The final image pixel value is obtained after all the appropriate samples from the filtered-projection memory have been summed together by the accumulator.

The Fourier processor preserves data accuracy through the use of scalers on the FFT IC's. Here, to prevent loss of accuracy and overflow during the accumulation, a 22-bit sum must be maintained since the data are 16-bit words. Even more bit accuracy is needed for the larger data matrices. Most adders process 16-bit words so two must be cascaded to handle the larger word width. Nearly any fast accumulator or arithmetic logic unit (ALU) will do. One acceptable device is the Integrated Devices IDT7381, a 16-bit cascadable ALU with accumulation capability. An accumulated word (value of one image pixel) can be scaled before further processing by using only sixteen bits. The choice of which bits to use will be flexible so that optimal data accuracy is maintained. A magnitude (or phase) calculation can then be performed on the pixel before it is sent out for display. The considerations for this step were discussed in the section on the Fourier processor.

5.3.2 Parallel backprojection

The backprojection architecture for data matrices of 128 x 128 and 256 x 256 incorporates a parallel computational arrangement. If a parallelism of M is

used, then a serial processor (an adder) must be able to complete it's operation (a real addition) in M·A nanoseconds where A is the serial time allotted per pixel (table 4.3). Acceptable values of M are M=4 for a resolution of 128 x 128 and M=32 for a resolution of 256 x 256. In these cases, each serial processor has 64 ns to complete an addition. Here, the time per operation is half that of the 64 x 64 case but it is still slow enough for most processors and memories. If more time per operation is needed, more parallelism can be used. The parallel processors can be distributed in one of two ways; they can be dedicated to portions of the filtered-projection space or they can be dedicated to portions of the reconstruction space. The former approach turns out to be much more efficient because the parallel data selected from the filtered-projection memory can be summed in an adder network [38]. A network like this arranges the adders into a series of stages. The first stage contains half the total number of adders with each successive stage containing half the number of the stage to it's left. The final stage contains one accumulator to add the previously summed value with the new value. Figure 5.6 shows an example of an adder network for M=8.

Filtered-projection samples must be introduced in parallel to the adder network every 64 ns. The filtered-projection memory and the projection address memory, therefore, must be divided into parallel segments. Each segment of the filtered-projection memory will have it's own data and address lines. Similarly, each segment of the projection-address memory will correspond to a single segment of filtered-projection data memory. Table 5.3 summarizes the memory requirements for the backprojector.

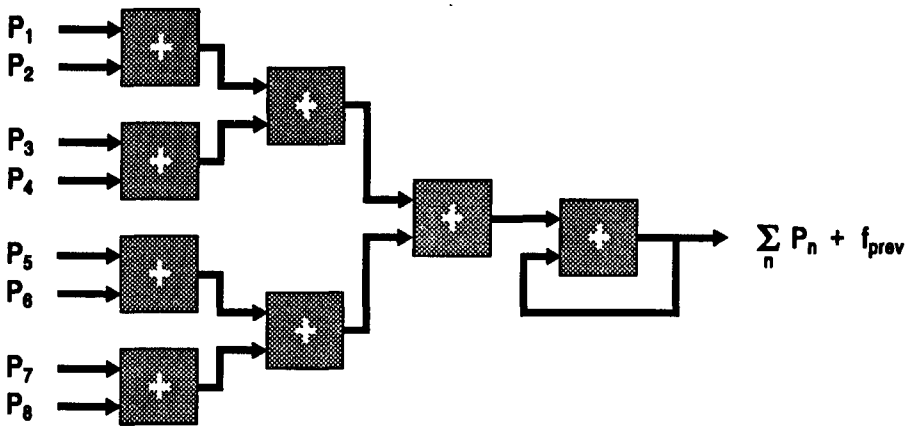

Figure 5.6 An adder network for a parallelism of 8.

Table 5.3

Backprojector memory requirements					
matrix size	total dual-port buffer memory	total address memory	M	parallel buffer memory segments	parallel address memory segments
64 x 64	4K x 16	256K x 12	1	4K x 16 (1 @ 126 ns)	256K x 12 (1 @ 126 ns)
128 x 128	16K x 16	2M x 14	4	4K x 16 (4 @ 63 ns)	512K x 12 (4 @ 63 ns)
256 x 256	64K x 16	16M x 16	32	4K x 16 (32 @ 63 ns)	512K x 12 (32 @ 63 ns)

The number of memories shown in this table are for one backprojector. This number must be doubled for an actual system since both real and imaginary parts are backprojected concurrently. The filtered-projection buffer must be much faster than the data memory in the digitizer or Fourier processor since N times more data must be read in a frame time. Fast SRAMS must be used for

these filtered-projection memories. As mentioned previously, VRAM can be used for the address memories since the addresses are read sequentially.

Figure 5.7 shows one backprojector (either real or imaginary) suitable for backprojection of a data matrix of 128 x 128. The parallelism of the system is 4 and so the adder network needed is in a 2-1-1 configuration. The filtered-projection data entering the backprojector get written to the segments of the filtered-projection memories through the use of a multiplexer. While this process is occurring, the segmented address memories are sequenced and the filtered-projection samples residing in these address locations are read into the adder network. The total communication rate from the full filtered-projection memory is a read every 16 ns, but the communication rate from the individual filtered-projection memories is 64 ns. Since four samples are fed to the adder network at once, an image pixel is produced with thirty-two passes through the network. As with the 64 x 64 backprojector, an image pixel can be converted to it's magnitude (or phase) and sent to a display buffer. The backprojector for a resolution of 256 x 256 is configured in the same way except that the adder network is in a 16-8-4-2-1-1 configuration and thirty-two parallel data and address memory segments are needed. Here, the total communication rate from the full filtered-projection memory is a read every 2 ns but the communication rate from the individual filtered projection memory segments is again 64 ns.

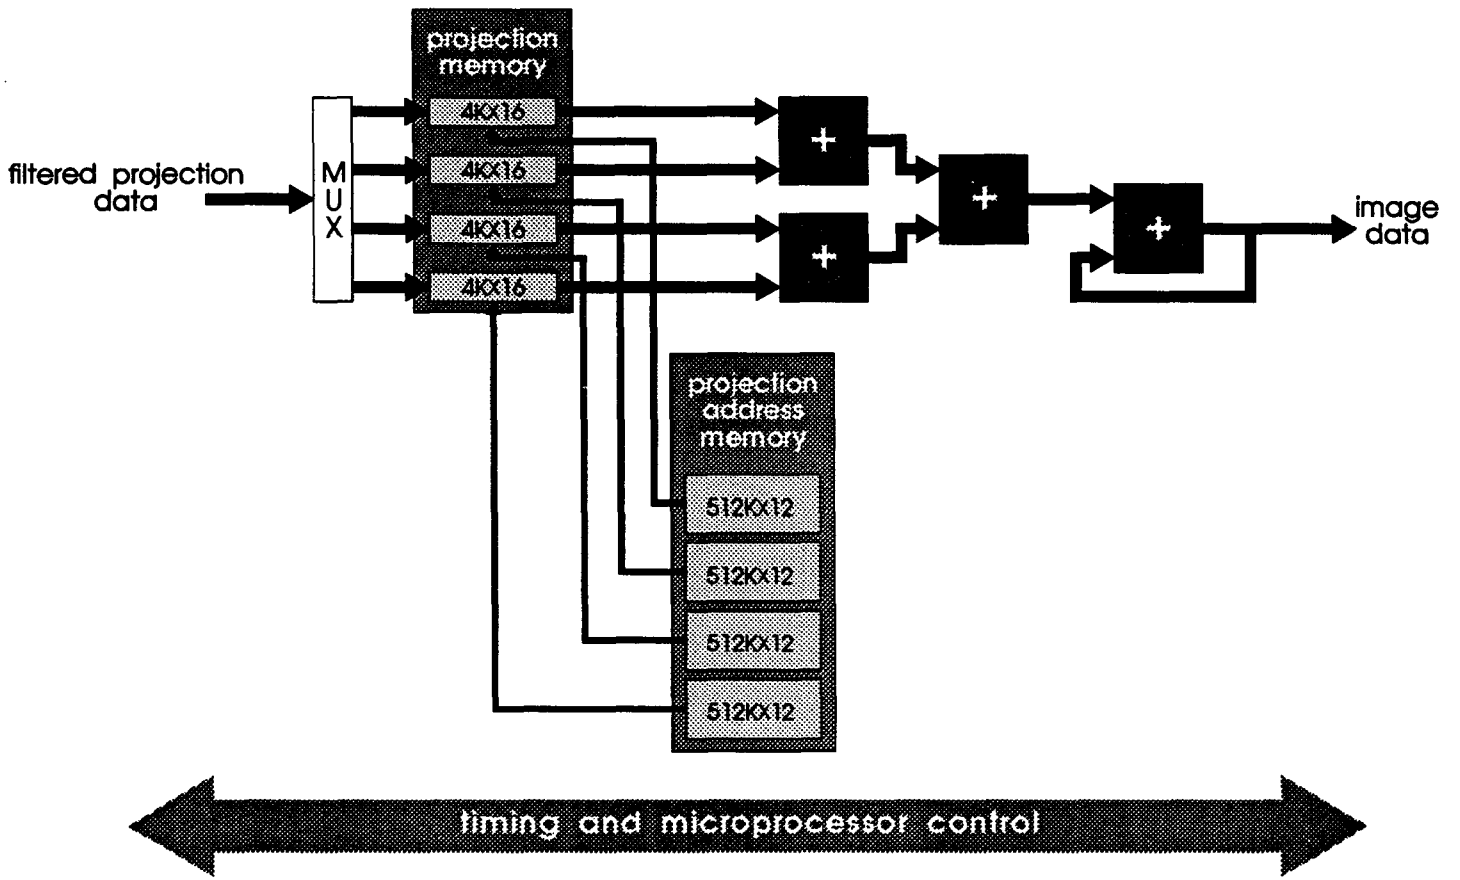

Figure 5.7 An example of the parallelism needed for backprojection of a 128 x 128 data matrix. The data and projection memories are divided into 4 separate banks, and the backprojection operation is performed by a network of 4 adders.

CHAPTER 6

UTILITY OF THE REAL-TIME RECONSTRUCTION SYSTEM

The real-time digital-electronic reconstruction system presented in the last chapter is designed to handle nearly any desired acquisition technique. This chapter discusses the utility of the reconstruction system for multiple acquisition scenarios.

6.1 Conventional Acquisition

Recall that conventional MR acquisition methods collect one line in k-space after an RF excitation. Typically, this takes 10 ms to 30 ms to complete and there is a rest period between excitations in the range of 100 ms - 3000 ms (chapter 3). Although real-time reconstruction is not needed for these techniques, there is no harm in it's use. In fact, the video-rate reconstruction system can enhance the utility of conventional acquisition methods. Any 2D Fourier or projection imaging experiment can be performed at any desired rate and the data will be reconstructed instantaneously. Video-rate reconstruction produces a display that immediately shows an image of the available data. (One image from a multi-slice experiment can be obtained by reconstructing data acquired from only one slice and ignoring all other data.) The image will be seen to fill in as more and more data samples are acquired. This may be useful in a clinical setting to save time wasted on a non-diagnostic scan such as that

caused by incorrect slice selection or image degradation due to motion. Normally, a full scan must be completed before reconstruction is done. With a system like the one here, the scan may be terminated once the image is seen to be incorrect or corrupted.

6.2 Ultra-fast Acquisition

The full potential of ultra-fast acquisition can only be realized when joined with real-time reconstruction. Currently, ultra-fast methods are used in a "snapshot" mode where the data is acquired very quickly, but reconstructed off-line. Cine-MRI movie loops of dynamic processes can be assembled off-line from multiple acquisitions but this is time-consuming and requires a large amount of fast data storage. The reconstruction system described here enables true real-time MRI, where a dynamic process is viewed as it occurs. Because of the video-rate sliding window configuration, new data is present in each video frame. The data do not have to be acquired at video-rates for a smooth motion to be visualized. Lower acquisition rates can be accommodated since the reconstruction system continuously reconstructs whatever data are fed to it.

6.3 Partial k-space scanning

Partial k-space scanning can be used with nearly any acquisition technique to reduce scan times. When used in conjunction with ultra-fast acquisition, an increase in spatial resolution can be achieved on systems with limited gradient hardware.

6.3.1 Conjugate synthesis

Recall from chapter 3 that because k-space data are Hermitian, only half of k-space need be acquired and the rest can be generated by conjugate

synthesis. The A41102 used in the Fourier processor has a conjugator that can be activated on the fly (within 100 ns). If only half of k-space is acquired, the system can synthesize the other half by turning on the conjugator and feeding the acquired data into the A41102 in the correct order. As an example, assume that the top half of k-space is continuously being acquired with a blipped echo-planar method and is stored properly reordered in the digitizer-interface buffer - the digitizer input address memory contains the correct reordered addresses in this case (chapter 5). The first half of the digitizer output address memory stores sequential addresses, and the bottom half stores the same addresses in reverse order. All the acquired data points (which constitute half of k-space) are read from the digitizer interface memory in sequential order and sent to the first A41102. The conjugator is off during this time. The conjugator is then turned on and the same data points are read in reverse order into the A41102. Thus, the A41102 synthesizes the rest of k-space from the acquired data. Again, because of the sliding-window configuration, new data points will be included in this process so that each display frame will contain some new data.

6.3.2 Tiling and interlacing

Both partial k-space tiling or interlacing can be handled by the reconstruction system. This is done by storing the proper addresses in the digitizer input and output address memories and performing the reconstruction on the full data matrix. The reconstruction system does not have to wait for the complete k-space to be acquired since it will reconstruct any new acquired data along with previously acquired data. The digitizer address memories make data reordering straightforward. As an example, consider an interlaced spiral-scan experiment where an image matrix of 128 x 128 is desired. This can be

accomplished by taking one shot with 128 x 64 sample points (128 radial lines with 64 points per line), and then taking a second interlaced shot with 128 x 64 sample points. If the digitizer input address memory contains sequential addresses so that the spiral data is written into consecutive memory locations in the digitizer-interface buffer, the digitizer output address memory is responsible for reordering the data prior to reconstruction. Reconstruction is performed on the full reordered data matrix (128 x 128).

6.4 Real-time image processing

The real-time reconstruction system can be easily configured to perform image processing on the data. This is because k-space data is directly accessible in the Fourier processor. Recall that the A41102 has an on-chip complex multiplier. The incoming data can be multiplied by an arbitrary filter function before Fourier transformation. Spiral-scan data is already manipulated in this way in the second A41102 by the 1D filter function. Echo-planar data can also be filtered as it enters the first A41102 by feeding the desired filter synchronously into the multiplier. Image processing of MRI images is not done very often in the clinical setting; processing functions such as edge enhancement do not seem to aid greatly in diagnosis. However, with real-time MRI, the utility of image processing may be increased. Real-time image processing may enhance the eye's natural ability to perceive motion. Note that because the filter is loaded into a memory from the data bus, the filter function can be modified in real-time. Thus, the form of the filter function can be altered interactively, and the effects on the image will be seen immediately.

6.5 Disadvantages

There are a couple of disadvantages to the system as described. Since the backprojector for a data resolution of 256 x 256 becomes hardware intensive, it will require a fairly large circuit layout. The cost of such a system may prevent it from being widely used. However, the backprojector at lower resolutions is not nearly as complex and would be useful in many cases. Note that the Fourier processor does not have this problem since the only hardware requirement to boost up the achievable resolution is a small increase in memory size and speed. Another disadvantage is that the system is a dedicated reconstructor, so general arithmetic operations commonly performed on array processors could not be performed with this system.

Even if data is not acquired at video-rates, the real-time reconstruction system described in this thesis is eminently suited to nearly any MR imaging experiment. There is no harm in reconstructing faster than the acquisition; in fact, there are advantages in doing so. A system like this will ultimately be needed if real-time MRI is to reach it's full potential.

References

[1] M. Stehling, R. Turner, and P. Mansfield. Echo-Planar Imaging: Magnetic Resonance Imaging in a Fraction of a Second. *Science*, 254, 43-44 (1991).

[2] G. Glover and N. Pelc. A rapid gated cine MRI technique. In Magnetic Resonance Annual 1988, H. Kressel, Ed. NewYork: Raven; 1988, pp.299-333.

[3] J. Frahm, A. Haase, and D. Matthaei. Rapid NMR Imaging of Dynamic Processes Using the FLASH Technique. *Magnetic Resonance in Medicine,* 3, 321-327 (1986).

[4] I. Pykett and R. Rzedzian. Instant images of the body by magnetic resonance. *Magnetic Resonance in Medicine,* 5:563-571 (1987).

[5] M. Stehling et. al. *ibid.* 254, 47 (1991).

[6] A. Holsinger, R. Wright, S. Riederer, F. Farzaneh, R. Grimm, and J. Maier. Real-Time Interactive Magnetic Resonance Imaging, *Magnetic Resonance in Medicine* 14,547-553 (1990).

[7] D. Bailes and D. Bryant. NMR Imaging, *Contemp. Phys.*, 25; 441,475 (1984).

[8] C. Partain, A. James, F. Rollo, R. Price, Nuclear Magnetic Resonance Imaging, W.B. Saunders (1983).

[9] W. Hinshaw and A. Lent, An Introduction to NMR Imaging: From the Bloch Equation to the Imaging Equation., *Proceedings of the IEEE*, 71; 333-350 (1983).

[10] A. Rosenfeld and A. Kak, Digital Picture Processing, Academic Press (1982).

[11] P. Van Der Meulen, J. Groen, A. Tinus, and G. Bruntink, Fast Field Echo Imaging: An Overview and Contrast Calculations, *Magnetic Resonance Imaging*, 6; 335-368 (1988).

[12] A. Haase, Snapshot FLASH MRI. Applications to T1, T2, and Chemical-Shift Imaging, *Magnetic Resonance in Medicine*, 13; 77-89 (1990).

[13] P. Mansfield, I. Pykett, Biological and medical imaging by NMR. *Journal of Magnetic Resonance*, 29; 355-373 (1978).

[14] F. Farzaneh, S. Riederer, N. Pelc, Analysis of T2 Limitations and Off-Resonance Effects on Spatial Resolution and Artifacts in Echo-Planar Imaging. *Magnetic Resonance in Medicine* 14; 123-139 (1990).

[15] C. Ahn, J.Kim, and Z. Cho, High-Speed Spiral-Scan Echo Planar NMR Imaging-I. *IEEE Transactions on Medical Imaging*, 5; 2-7 (1986).

[16] A. Gmitro and Y. Chen, Reconstruction of Fast-Scan MR Data, *Proceedings of the SPIE*, 1351; 331 (1990).

[17] J. Gaskill, Linear Systems, Fourier Transforms, and Optics, Wiley and Sons (1978).

[18] M. Cohen and R. Weisskoff, Ultra-Fast Imaging, *Magnetic Resonance Imaging*, 9; 1-37 (1991).

[19] M. Cohen and R. Weisskoff, *ibid.*

[20] C. Ahn, J. Kim, and Z. Cho, *ibid.*

[21] J. Goodman, Introduction to Fourier Optics, McGraw-Hill (1968).

[22] R. Turner, Optical Reconstruction of NMR Images, *J. Phys. E: Sci. Instrum.* 18; 875 (1985).

[23] A. Gmitro, V. Tresp, Y. Chen, R. Snell, and G. Gindi, Video-Rate Reconstruction of CT and MR Images, *Proceedings of the International Conference on Information Processing in Medical Imaginig (IPMI).* Berkeley CA (1989).

[24] A. Gmitro, V. Tresp, Y. Chen, R. Snell, and G. Gindi, *ibid.*

[25] V. Tresp, Electro-Optical System for the High-Speed Reconstruction of Computed Tomography Images, Ph.D. Dissertation, Yale University (1989).

[26] E. Brigham, The Fast Fourier Transform and it's Applications, Prentice-Hall (1988).

[27] R. Wright, S. Riederer, F. Farzaneh, P. Rossman, and Y. Liu, Real-Time MR Fluoroscopic Data Acquisition and Image Reconstruction, *Magnetic Resonance in Medicine* 12;407-415 (1989).

[28] A. Pomerleau, M. Fournier, and H. Buijs, On the Design of a Real Time Modular FFT Processor, *IEEE Transactions on Circuits and Systems*, October (1976).

[29] A. Nobile and V. Roberto, Efficient Implementation of Multidimensional Fast Fourier Transforms On A Cray X-MP. *Computer Physics Communications* 40; 189-201 (1986).

[30] W. Briggs, L Hart, R. Sweet, and A. O'Gallagher, Multiprocessor FFT Methods, *SIAM J. Sci. Stat. Comput.* 8; 27-42 (1987).

[31] N. Chowdary, and W. Steenaart, A High Speed Two-Dimensional FFT Processor, *Proceedings of the IEEE* 5; 1-4 (1984).

[32] V. Milutinovic, J Fortes, and L. Jamieson, A Multimicroprocessor Architecture for Real-Time Computation of a Class of DFT Algorithms, *IEEE Transactions on Acoustics, Speech, and Signal Processing* 34;1301-1309 (1986).

[33] A. Despain, Very Fast Fourier Transform Algorithms Hardware for Implementation, *IEEE Transactions on Computers*, 28; 333-341 (1979).

[34] E. Wold, and A. Despain, Pipeline annd Parallel-Pipeline FFT Processors for VLSI Implementations, *IEEE Transactions on Computers*, 33;414-425 (1984).

[35] S. Azevedo, A. DeGroot, D. Schneberk, J. Brase, H. Martz, A. Jain, K. Current, and P. Hurst, Tomographic Image Reconstruction Using Systolic Array Algorithms, *Review of Progress in Quantitative NDE Meeting*, La Jolla, CA (1988)

[36] B. Gilbert, S. Kenue, R. Robb, A. Chu, A. Lent, and E. Swartzlander, Rapid Execution of Fan Beam Image Reconstruction Algorithms Using Efficient Computational Techniques and Special-Purpose Processors, *IEEE Transactions on Biomedical Engineering* 28; 98 (1981).

[37] F. Farzaneh, S. Riederer, J. Lee, T. Tasciyan, R. Wright, and C. Spritzer, MR Fluoroscopy: Initial Clinical Studies, *Radiology* 171;545-549 (1989).

[38] B. Gilbert, S. Kenue, R. Robb, A. Chu, A. Lent, and E. Swartzlander, *ibid.*