Non-Uniform Fast Fourier Transformation of SPRITE MRI Data

James Rioux

CS 4997 Honours Thesis
Supervisors: E. Aubanel, B. Balcom

December 2003

Abstract

A new algorithm has been developed which permits extremely accurate and fast computation of the Discrete Fourier Transform for certain sets of non-uniformly sampled data. Specifically, this algorithm can be used to process the data acquired during a Magnetic Resonance Imaging experiment using the SPRITE technique with Multiple Point Acquisition, which samples data non-uniformly in a spatial frequency domain.

Non-uniform MRI data is generally processed with an interpolation or regridding algorithm which computes an approximate answer in $O(NlogN + Nlog(1/\epsilon))$ time; the accuracy depends on the parameter ϵ and is typically one part in 10^5. The algorithm developed herein uses a combination of Chirp Z-Transforms (CZT) to compute the Fourier Transform to within one part in 10^{15} (which is on the order of the machine epsilon) in $O(N_T NlogN)$ time, where N is the number of points reconstructed and N_T is the number of samples collected at each step of the SPRITE experiment.

When MRI data is processed by the CZT algorithm, the result is an improvement in the signal-to-noise ratio of the final image, by a factor of $\sqrt{N_T}$. Though a resolution increase is also observed in some images, it can be demonstrated that this effect is due entirely to the SNR improvement. This result allows us to use a modified version of the transform which operates in $O(NlogN)$ time for all practical purposes. This algorithm has immediate application in the rapid imaging of low-signal materials, and may also be developed for other uses, such as the numerical solution of differential equations.

Contents

Acknowledgements

The research which led to the algorithm described in this paper was the result of many different investigations coming together at the same time. I consider myself very fortunate to have been a part of the UNB MRI Centre at that time, such that I could take all of these related threads of inquiry and weave them into something new and worthwhile.

My supervisors for this thesis were Dr. Eric Aubanel of the Faculty of Computer Science, and Dr. Bruce Balcom, the director of the UNB MRI Centre. Their advice and guidance were invaluable at all stages of this project; Bruce and Eric are both excellent researchers and great friends, and it has been a pleasure working with them and learning from them.

Among the people at the MRI Centre who contributed to this research were Meghan Halse, who is responsible for bringing the Chirp Z-Transform algorithm to the MRI Centre in the first place; Dr. Bryce MacMillan, who helped me learn how to use the MRI instruments to acquire the test data; and Dr. Greg Stoch, who eagerly reviewed my initial derivations and provided some very helpful suggestions.

Finally, I extend my deepest thanks to Kim Brewer for all her love and support throughout this project. The problem of fully expressing my gratitude for all she's given me during these past eight months is one that may have no solution; I can only hope to find a suitable approximation.

1 Introduction

Magnetic Resonance Imaging (MRI) is a proven technique for medical imaging, and in recent years it has also found application in materials research. It allows non-invasive analysis of a wide variety of substances, including human tissue, foods, polymers and porous media. In addition to qualitative examination of the internal structure of materials, MRI also permits quantitative measurement of structural features in the sample being imaged. In some cases, dynamic processes such as fluid flow may also be quantified.

The MRI signal is generated by excitation of the sample with a radio-frequency pulse at an appropriate frequency. Using any of a variety of methods, the frequency spectrum is sampled in a reciprocal domain known as k-space, and converted into an image using a Fourier Transform. Because image reconstruction is performed on a computer, we require a fast but accurate algorithm to recover the desired image from the sampled k-space data. The algorithm of choice is traditionally the Fast Fourier Transform (FFT) which was introduced by Cooley and Tukey in 1965 [6], and which has an asymptotic complexity of $O(NlogN)$. In contrast, the direct Fourier Transform has a complexity of $O(N^2)$, which makes it too slow for processing datasets containing more than a few thousand samples.

The FFT is only suitable for processing uniformly sampled, or Cartesian, data; for this reason the majority of MRI techniques employ a rectilinear sampling scheme, enabling rapid processing using the FFT. However, this restriction is incompatible with some MRI methods. Contemporary clinical MRI methods acquire k-space data with spiral k-space trajectories [14]. These techniques feature a uniform angular separation, but a non-uniform linear separation. Single Point Imaging [8] typically samples k-space in a rectilinear way, but it is possible to acquire extra data during the SPI experiment at little time cost. This allows for a final image with higher signal quality and greater apparent resolution, but although this extra data will still exhibit some regularities, it will not fall onto the standard Cartesian grid.

The most common approach to processing non-uniformly sampled data is to use some form of interpolation, such as convolution regridding with an appropriate kernel [12], to restore the data to a Cartesian grid. The data can then be reconstructed with the FFT. Many of these interpolation schemes are based on (or equivalent to) the work of Dutt and Rokhlin [7], though other methods exist, such as that of Sha, Guo and Song [18]. However, interpolation is by definition an approximate method, and will always introduce some inaccuracy into the resulting image. It

would be preferable to implement a true Non-Uniform Fast Fourier Transform, allowing for the accurate reconstruction of an image in $O(NlogN)$ time.

This thesis will show that such an operation is possible for the specific case of extra data acquired during the course of an SPI experiment, as described above. For this particular case, it will be shown that the discrete Fourier Transform can be implemented with an algorithm known as the Chirp Z-Transform [16]. The Chirp Z-Transform is a more general version of the standard Fourier Transform which is implemented using a convolution, and therefore executes in $O(NlogN)$ time. This algorithm will be used to improve the quality of various MRI images, and will be compared to interpolation schemes to demonstrate its superior accuracy and speed.

2 Theory

In this section we shall outline some of the important mathematical formalisms upon which the FFT algorithm for SPI-type data will be based. This includes both the standard Fourier Transform and the Chirp Z-Transform. The physics of Magnetic Resonance Imaging will also be described, to demonstrate how the MRI signal is generated and collected, and how the Fourier Transform will be used to reconstruct the data. Finally we shall provide a description of Single Point Imaging and the extension of SPI known as SPRITE, which will be used in the experimental verification of the algorithm.

2.1 The Fourier Transform

It was shown in 1822 by the French physicist J. B. Fourier that any function $f(x)$ which is periodic on the interval $[-L, L]$ can be represented as an infinite series of sines and cosines:

$$f(x) = \sum_{k=0}^{\infty} a_k cos(k\pi x/L) + \sum_{k=0}^{\infty} b_k sin(k\pi x/L) \tag{2.1}$$

where the Fourier coefficients a_k and b_k can be calculated by exploiting the orthogonality of the sine and cosine functions. By making use of the Euler identity $e^{i\theta} = cos\theta + isin\theta$ the Fourier series can also be defined for complex functions periodic on $[-L/2, L/2]$ as

$$f(x) = \sum_{k=-\infty}^{\infty} c_k e^{2\pi ikx/L} \tag{2.2}$$

Transforming the complex coefficients c_k from discrete variables to a continuous function, and replacing the summation with an integral, we obtain the expression

$$f(x) = \int_{-\infty}^{\infty} F(k) e^{2\pi i k x} dk \tag{2.3}$$

which defines a function f(x) in terms of its *Fourier transform*,

$$F(x) = \int_{-\infty}^{\infty} f(x) e^{-2\pi i k x} dx \tag{2.4}$$

Equation 2.3 is typically referred to as the *inverse Fourier transform* since it allows us to recover an expression for $f(x)$ given knowledge of its Fourier transform $F(k)$.

$F(k)$ can be thought of as an alternative representation of the function $f(x)$ in a different set of basis vectors. In many cases the dependent variable is time and the Fourier representation is in terms of frequency. For this application, our variable is spatial position, and the corresponding reciprocal space is the spatial frequency k whose basis vectors are trigonometric functions of the form $e^{2\pi i k}$. In either case, the Fourier Transform is a useful tool in the solution of differential equations because the basis vectors are eigenfunctions of the differentiation operator.

It should be clear that the Fourier transform need not be confined to one dimension. In two dimensions, for example, we can express the Fourier transform as a double integral:

$$F(k_x, k_y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) e^{-2\pi i (k_x x + k_y y)} dx dy \tag{2.5}$$

It can be shown that this integral is separable, which means that we can execute this operation by applying the Fourier transform across the first dimension only, and then across the second. This will be of importance when it comes time to calculate the Fourier transform of a function using a computer.

In practical applications we generally do not have access to an analytic functional form for $F(k)$; we can only sample it at certain number (say, N) of discrete locations k_n in the frequency domain. To reconstruct $f(x)$ from this sampled data we set

$$F'(k) = F(k) \sum_{n=0}^{N-1} \delta(k - k_n) \tag{2.6}$$

to represent the sampling effect, such that equation 2.3 becomes

$$f(x_m) = \sum_{n=0}^{N-1} F(k_n) e^{2\pi i k_n x_m} \qquad m \,\epsilon\, \{0, 1, \ldots, M-1\} \tag{2.7}$$

This is known as the inverse *discrete* Fourier transform (DFT); notice that $f(x)$ is restricted to M discrete values x_m. There is also a forward form of the discrete Fourier transform for converting from $f(x_m)$ to $F(k_n)$.

In terms of computational complexity, the DFT is an $O(NM)$ operation, since evaluating the function $f(x)$ at one point x_m requires N complex multiplications and additions. For the special case when $M = N = 2^a$ for some integer a, and both the k_n and the x_m are uniformly spaced, Cooley and Tukey showed that it is possible to calculate the DFT in $O(NlogN)$ time. This is done by splitting the N-point DFT into two $N/2$-point DFT's, splitting those into four DFT's of length N/4, and so on until the trivial 1-point DFT is reached. This formulation is called the Fast Fourier Transform (FFT). Further revisions have since made it possible to compute transforms where N is not a power of two; it may even be a prime number.

In the case of data which do not fall onto a uniform rectangular (or Cartesian) grid, the FFT will not properly compute the Fourier Transform. To address the problem of computing the Fourier Transform of non-uniform data, a variety of interpolation schemes have been proposed; that of Dutt and Rokhlin [7] is one of the most common of these. These methods allow computation of the Fourier Transform to a specified accuracy, with the understanding that increasing the accuracy will involve a proportionate increase in computation time.

To conclude this section we introduce an important result called the *convolution theorem*, which states that the convolution $h(x)$ of two functions $f(x)$ and $g(x)$,

$$h(x) = f(x) \otimes g(x) = \int_{-\infty}^{\infty} f(t)g(x-t)dt \tag{2.8}$$

has a fourier Transform $H(k)$ equal to the product of the transforms $F(k)$ and $G(k)$ of the two functions being convolved; that is,

$$H(k) = F(k)G(k) \tag{2.9}$$

This follows directly from the definition of the Fourier transform and a change of variable. We are therefore presented with a method for calculating the convolution of two functions $f(x)$ and $g(x)$. After finding the Fourier transform of each function (using eqn. 2.4), multiplying them together and applying the inverse Fourier transform (using eqn. 2.3) will yield their convolution. This theorem will allow us to compute the convolution of two functions at high speed using the previously mentioned FFT algorithm.

2.2 The Chirp Z-Transform

The Fourier transform is a special case of a more general mathematical operation known as the Z-Transform. Another special case of the Z-Transform is the so-called Chirp Z-Transform which, though more general in nature than the Fourier transform, remains computable in $O(NlogN)$ time. In this section we introduce the Chirp Z-Transform and outline its computation. This derivation follows the example of Rabiner, Schafer and Rader [16] who are responsible for developing the Chirp Z-Transform for applications in signal processing during the late 1960's.

The most general Z-transform of a sequence of numbers x_n is defined as

$$X(z) = \sum_{n=-\infty}^{\infty} x_n z^{-n} \tag{2.10}$$

where z, x_n and $X(z)$ are all complex. For our purposes we will only concern ourselves with sequences having a finite number N of samples, and the z-transform will only be computed at a finite number M of points, z_k. That is to say,

$$X(z_k) = \sum_{n=0}^{N-1} x_n z_k^{-n} \qquad k \,\epsilon\, \{0, 1, \ldots, M-1\} \tag{2.11}$$

To make the interpretation of the variable z more clear, we shall select a contour in the z-plane of the form

$$z_k = AW^{-k}, \quad A = A_0 e^{i\theta_0}, \quad W = W_0 e^{i\phi_0} \tag{2.12}$$

This contour begins at $z = A$, a distance A_0 from the origin and at an angular displacement of θ_0 from the horizontal axis. The contour either circles the origin (if $W_0 = 1$), spirals inward (if $W_0 < 1$), or spirals outward (if $W_0 > 1$), with samples spaced at angular intervals of ϕ_0.

The Fourier Transform is a special case of the z-transform on the contour described by $A_0 = 1$, $W_0 = 1$, $\theta_0 = 0$, $\phi_0 = -2\pi/N$. This choice describes a set of equally spaced points on the unit circle, and our expression for the z-transform becomes

$$X(z_k) = \sum_{n=0}^{N-1} x_n e^{-2\pi ink/N} \tag{2.13}$$

which is the discrete Fourier Transform (cf. equation 2.7).

Rabiner et. al. [16] have shown that it is possible to compute the z-transform for the more general contour of equation 2.12,

$$X(z_k) = \sum_{n=0}^{N-1} x_n A^{-n} W^{nk} \qquad k \,\epsilon\, \{0, 1, \ldots, M-1\} \tag{2.14}$$

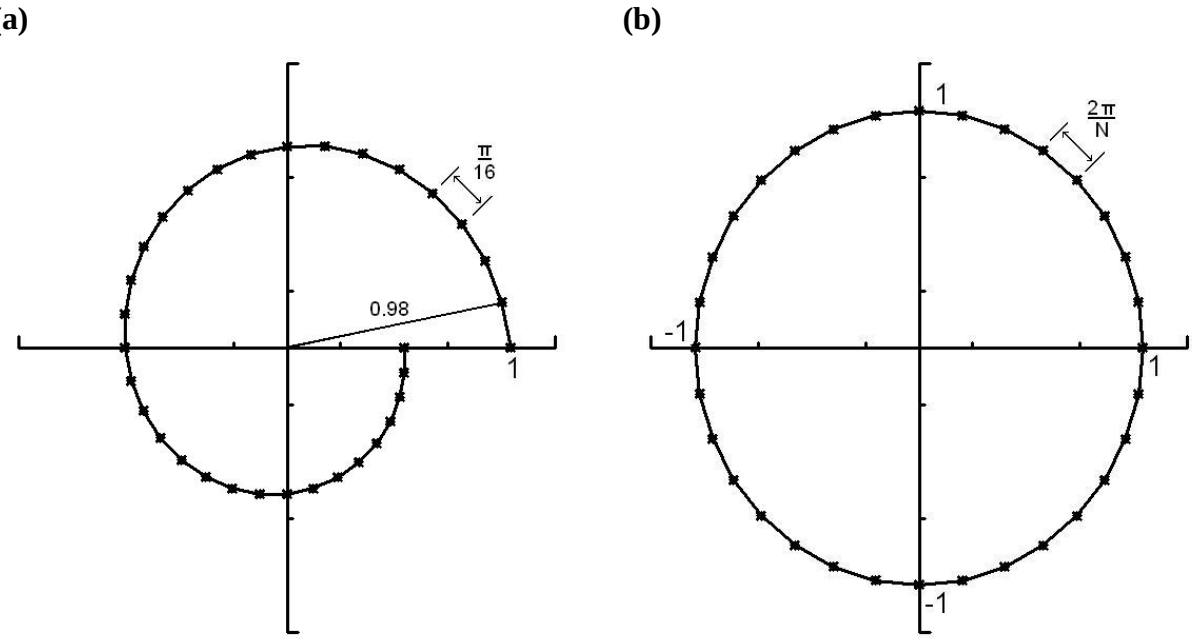

Figure 1: (a) A typical z-plane contour, $z = AW^{-k}$, $A = 1, W = 0.98e^{-i\pi/16}$. (b) The z-plane contour equivalent to the DFT, $A = 1, W = e^{-2\pi i/N}$.

by using the so-called Bluestein substitution, $nk = \frac{n^2+k^2-(k-n)^2}{2}$, to express the transform in terms of a discrete convolution:

$$X(z_k) = \sum_{n=0}^{N-1} x_n A^{-n} W^{n^2/2} W^{k^2/2} W^{-(k-n)^2/2} \quad (2.15)$$

$$= W^{k^2/2} \sum_{n=0}^{N-1} (x_n A^{-n} W^{n^2/2})(W^{-(k-n)^2/2}) \quad (2.16)$$

$$= W^{k^2/2}((x_n A^{-n} W^{n^2/2}) \otimes (W^{-n^2/2}))_k \quad (2.17)$$

As described at the end of the previous section, this convolution can be performed using a combination of Fast Fourier Transforms. For this reason, the z-transform can be calculated in $O(NlogN)$ time. According to Rabiner et. al. [16], the function $W^{n^2/2}$ found in the convolution is a sinusoid called a "chirp" in some radar systems, and they have therefore named this formulation of the algorithm the Chirp Z-Transform, or CZT.

The CZT has several properties not shared by the FFT. For our application, the two most important such properties are that the number of input points (N) need not equal the number of output points (M), and that the angular spacing of the points need not be $2\pi/N$ as with the FFT.

2.3 Nuclear Magnetic Resonance

The phenomenon of nuclear magnetic resonance (NMR) was discovered independently by Bloch [3] and Purcell [15] shortly after the Second World War; they were jointly awarded the Nobel Prize for this discovery in 1952. NMR arises in nuclei which have an intrinsic angular momentum, or "spin". Classically, an object (such as a proton) with a charge Q, a mass M, and an angular momentum $\vec{J}$ will also have a magnetic dipole moment

$$\vec{\mu} = \frac{Q}{2M}\vec{J} \tag{2.18}$$

which, when placed in a magnetic field $\vec{B}$, will provide the particle with a magnetic energy

$$E_\mu = \vec{\mu} \cdot \vec{B} = \gamma \vec{J} \cdot \vec{B} \tag{2.19}$$

Here we introduce the gyromagnetic ratio γ, a quantity which is characteristic of a given nucleus, and which depends on the charge and mass. The hydrogen nucleus, for example, has $\gamma = 42.2$ MHz/T.

In quantum mechanics the angular momentum of a particle is quantized along any axis $\hat{z}$ according to the relation

$$J_z = m\hbar \qquad m\,\epsilon\,\{-j, -j+1, \ldots, j-1, j\} \tag{2.20}$$

where $\hbar$ is Planck's constant and j is the particle's spin quantum number. There will be a different energy E_m for each value of the quantum number m. If we take $\vec{B} = B_0\,\hat{z}$, we find that this energy is

$$E_\mu = \hbar m \gamma B_0 \qquad m\,\epsilon\,\{-j, -j+1, \ldots, j-1, j\} \tag{2.21}$$

The energy difference between any two adjacent energy levels is simply

$$\Delta E = \hbar\gamma B_0 = \hbar\omega_0 \tag{2.22}$$

with $\omega_0 = \gamma B_0$ called the *Larmor frequency*.

Two states of a system with different energies will be populated such that more particles exist in the lower-energy state, and fewer in the higher-energy state. The difference in energy between the states of nuclei in a magnetic field is quite small, so there will only be a slight majority of particles in the lower-energy state, in which the dipole is oriented roughly along

the direction of the field. However, this small excess of about one particle in every 100,000 is sufficient to create a net sample magnetization in the direction of the external field $\vec{B}$.

If we now provide the system with energy at the Larmor frequency, this will cause some of the nuclei to make transitions from one state to another. Knowing γ, B_0 can be selected such that ω_0 is in the radio-frequency range. Applying an RF pulse at the appropriate frequency will excite the sample as desired. In the classical representation the application of an RF pulse perpendicular to $\vec{B}$ will introduce a torque that rotates the net magnetization vector downwards, into the transverse plane.

However, it can also be shown that a nucleus with a magnetic moment $\vec{\mu}$ which is not aligned exactly along the external field will experience a torque as the field tries to align the dipole's axis with the field. This torque will lead to rotation in the same way that a gyroscope rotates under the influence of gravity. For the nuclei in the magnetic field, this rotation takes place at the Larmor frequency, and it means that the net magnetization for the entire system of atoms will rotate in a similar fashion.

The laws of electrodynamics tell us that a rotating magnetic field will induce a voltage in a suitably oriented conductor. This voltage is collected and amplified, and becomes the NMR signal. This signal is typically known as the Free Induction Decay (or FID) since the voltage decays exponentially after the RF pulse terminates. The FID's time constant is called T_2^* and is characteristic of the material that is being imaged; this gives us a way to identify species by their NMR signal.

2.4 Magnetic Resonance Imaging

The method described above will excite all of the nuclei within the sample. While this is quite suitable for examining the bulk properties of the material, such as its chemical composition, in order to determine the sample's internal structure or how its constituents are distributed, we need to generate a spatially resolved image.

In order to distinguish various regions of the sample, we can apply a magnetic field gradient $\vec{G_x} = \frac{d\vec{B_0}}{dx}$ across the sample. This modifies the Larmor equation as follows:

$$\omega(x) = \gamma(B_0 + G_x x) \tag{2.23}$$

Now the frequency of precession will change linearly with spatial position, and different regions of the sample become distinguishable since they will induce voltages at different frequencies.

The total received signal from all of the nuclei in the sample can be written by integrating over all space, as

$$s(t) = \int_{-\infty}^{+\infty} \rho(x) e^{i\omega(x)t}\, dx \tag{2.24}$$

where $\rho(x)$ represents the density of nuclei at a given position x, and the complex exponential is used to illustrate that the nuclei at that location are rotating with a frequency $\omega(x)$. While the sample only occupies a finite region in space, since $\rho(x) = 0$ outside that region we can justifiably expand the limits of our integral to infinity. (Here we are treating only the one-dimensional case, but it can be shown that this will readily generalize to higher dimensions.)

Substituting our previous expression for $\omega(x)$ (eqn. 2.23), and noting that we can always demodulate the signal by a constant frequency of $B_0 t$ during the acquisition process, the MRI signal equation becomes

$$s(t) = \int_{-\infty}^{+\infty} \rho(x) e^{i\gamma x G_x t}\, dx \tag{2.25}$$

To clarify the significance of this signal equation, Mansfield [13] proposed the substitution $k = \gamma G t / 2\pi$, which changes the equation to

$$s(k) = \int_{-\infty}^{+\infty} \rho(x) e^{2\pi i k x}\, dx \tag{2.26}$$

By comparison with eqn. 2.3, this equation clearly describes a Fourier transform relationship between the actual proton density $\rho(x)$ and the intensity of the signal $S(k)$ in the reciprocal space, which is given the name k-space. This implies that, once we have determined the form of the signal in k-space, applying a Fourier transform will allow us to recover the desired image showing proton density as a function of position:

$$\rho(x) = \int_{-\infty}^{+\infty} s(k) e^{-2\pi i k x}\, dx \tag{2.27}$$

For the discrete case this becomes

$$\rho(x_m) = \sum_{n=0}^{N-1} s(k_n) e^{-2\pi i k_n x_m} \qquad m \,\epsilon\, \{0, 1, \ldots, M-1\} \tag{2.28}$$

The MRI experiment then consists of sampling k-space at certain discrete values, and using the sampled data to reconstruct an image using the Fourier relationship of equation 2.26. While it is theoretically possible to sample any desired values in k-space, generally the sampling is done such that all of the acquired data will fall onto a Cartesian grid. This will allow the data to be processed with the FFT.

2.5 Single Point Imaging and SPRITE

Magnetic resonance imaging is usually applied in the context of medical imaging, using techniques such as projection-reconstruction and spin-echo imaging [5]. All of these methods are based around the collection of a k-space signal that evolves in time, and while they are suitable for medical imaging, these techniques are not readily generalized to materials imaging because they exhibit several artifacts due to chemical shift, magnetic field inhomogeneity, and other time-dependent factors. The short signal lifetime of solid materials is also a problem, since these methods were designed for imaging of liquids which have relatively long signal lifetimes.

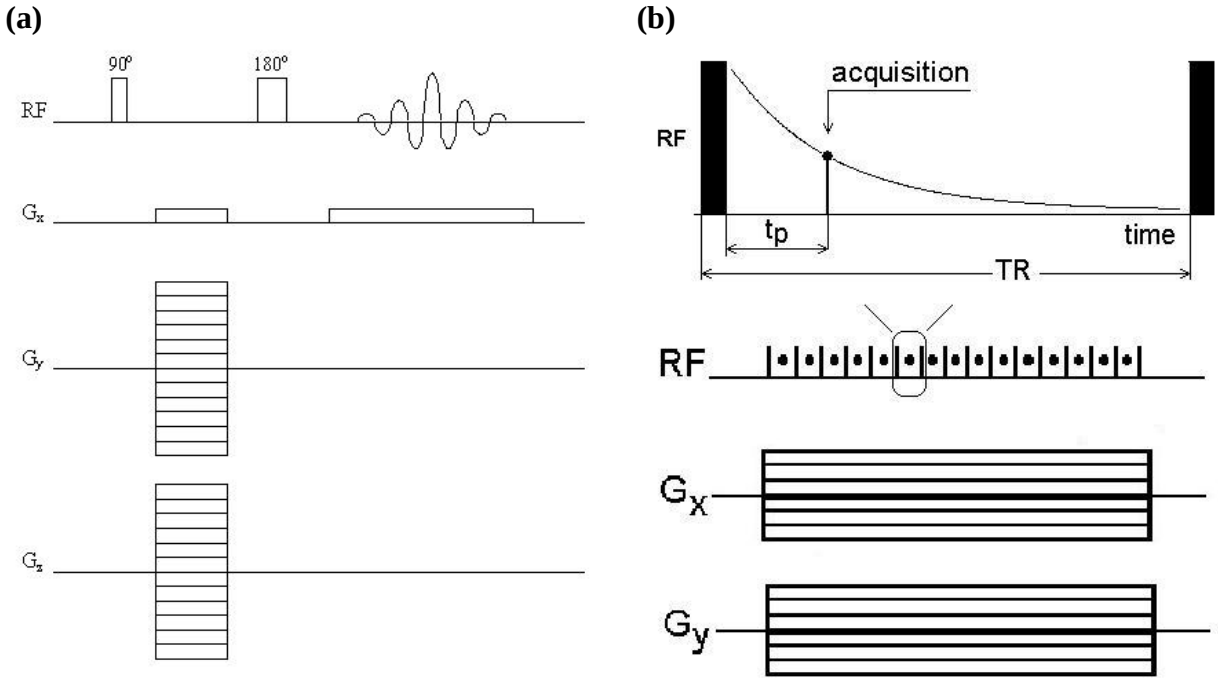

Figure 2: (a) Spin-echo pulse sequence. (b) SPI pulse sequence.

To address some of these problems, Emid and Creyghton [8] proposed an MRI method in 1985 that would become known as Single Point Imaging, or SPI. Instead of sampling data in k-space by varying the time t, a single encoding time t_p is selected and only the value of the magnetic field gradient is changed. For any given value of the gradient, an RF pulse is applied and the FID is sampled at $t = t_p$; the gradient is then changed to move to a new area of k-space. This allows imaging of short-T_2^* materials - that is, materials whose signal decays rapidly [2]. These materials are often difficult or impossible to image by other methods.

SPI in its original form is much slower than medical imaging techniques, and very demanding of the hardware (specifically, the gradient duty cycle is high). In 1996, Balcom et. al. [1] presented an extension of the standard SPI pulse sequence known as SPRITE, or Single Point Ramped Imaging with T_1 Enhancement. The major feature of the SPRITE sequence is a ramped gradient which decreases switching time and which makes it possible to rapidly acquire images of low-T_2^* materials. The standard SPRITE sequence is one-dimensional, though additional gradients can be added in orthogonal directions to allow imaging in up to three dimensions if desired. The sampling scheme is usually rectilinear, acquiring samples in k-space starting at one extremity and moving to the opposite extremity.

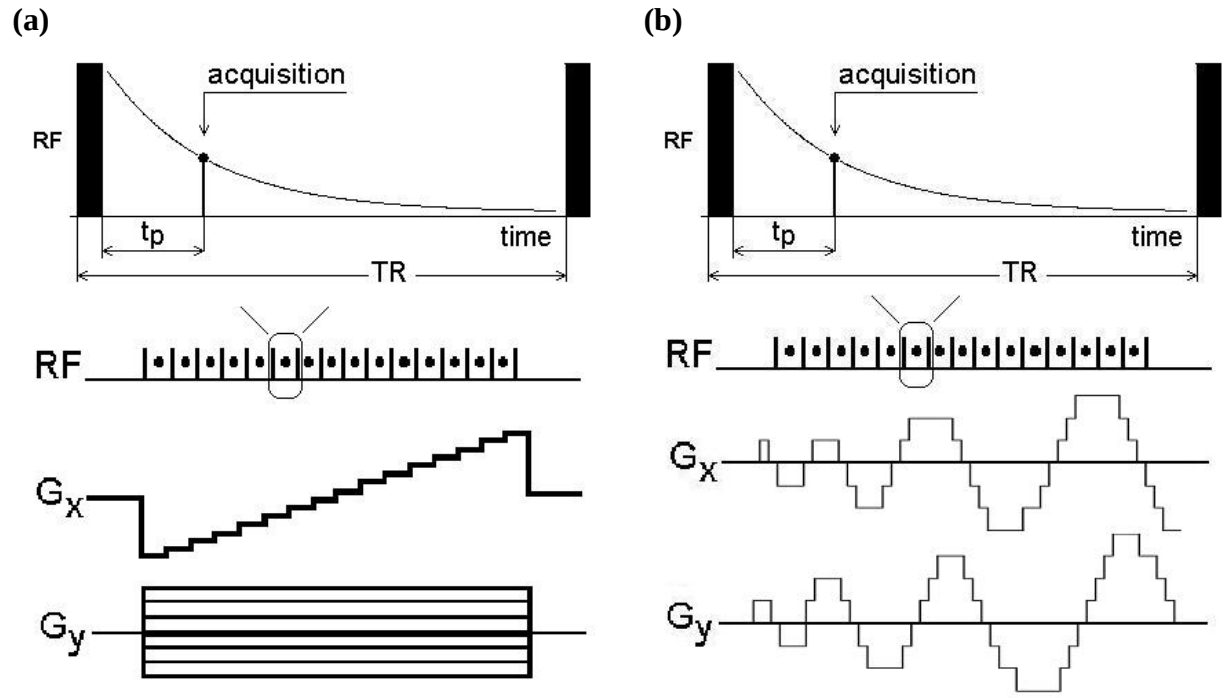

Figure 3: (a) SPRITE pulse sequence. (b) Spiral-SPRITE pulse sequence.

However, SPRITE has since led to a family of related techniques, and 2D imaging is more quickly accomplished by one of these variant techniques called Spiral-SPRITE [10]. Using Spiral-SPRITE, the magnetic field gradients in two orthogonal dimensions both have an initial value of zero, and are ramped sinusoidally such that the sampling trajectory follows a roughly spiral pattern, moving outward from the center of k-space. The gradient values are chosen such that the sampled data still fall on a Cartesian grid, and therefore the image can be reconstructed

with the FFT as usual. In this sense Spiral-SPRITE differs from the inherently nonuniform rounded spiral or rosette trajectory that is common in medical imaging [14].

Since not all of the k-space matrix is sampled, the Spiral-SPRITE method will be naturally faster than the corresponding standard SPRITE sequence. In addition, the centric-scan methodology has a number of advantages for quantitative imaging. A new method called Conical-SPRITE [10] provides a similar centric-scan technique for 3D imaging.

2.6 Multiple Point Acquisition

In all instances of the SPI and SPRITE techniques, the opportunity exists to sample more than one point along the FID at each gradient step. Additional points after the base encoding time t_p are easily acquired; for example, N_T points could be acquired at $t_p(j) = t_{p0} + j\Delta t_p, 0 \leq j < N_T$, with the spacing Δt_p typically on the order of a microsecond. (Figure 4) The additional data collection will not add to the duration of the acquisition, since we must wait for the signal excited by one RF pulse to decay sufficiently before applying another pulse. This process is referred to as Multiple Point Acquisition, or MPA.

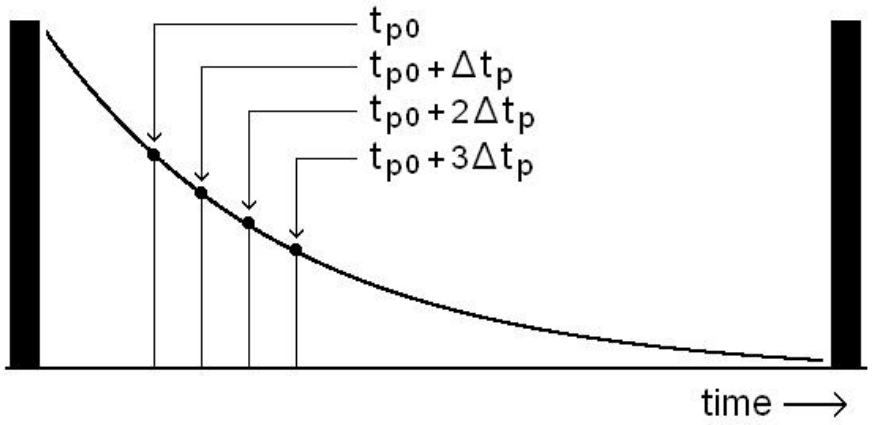

Figure 4: Illustration of Multiple Point Acquisition with $N_T = 4$.

In some instances it may be possible to acquire dozens of extra points before the FID decays completely, and this extra data could conceivably be used to provide more information about the sample. This will enhance the quality of the final image at no cost in terms of acquisition time. However, despite the many potential benefits it offers, MPA is not widely practiced, for two reasons.

(1) Although it is possible to separately process the data acquired at each encoding time, and reconstruct a number of separate images of the same sample, there is no straightforward way to combine these images, since they will not have the same field of view. The field of view for a given acquisition relates to the size of a step in k-space, given by the formula

$$FOV(j) = 1/\Delta k = 2\pi/\gamma \Delta G t_p(j) \tag{2.29}$$

The magnitude ΔG of a gradient step is constant for any given experiment, but as we collect more points along the FID, $t_p(j)$ will increase and the field of view will shrink. The same pixel in two separate images will no longer correspond to the same region of the sample. This means that we cannot, for example, average the images to improve the signal-to-noise ratio, unless we first correct the images so they all share a common field of view.

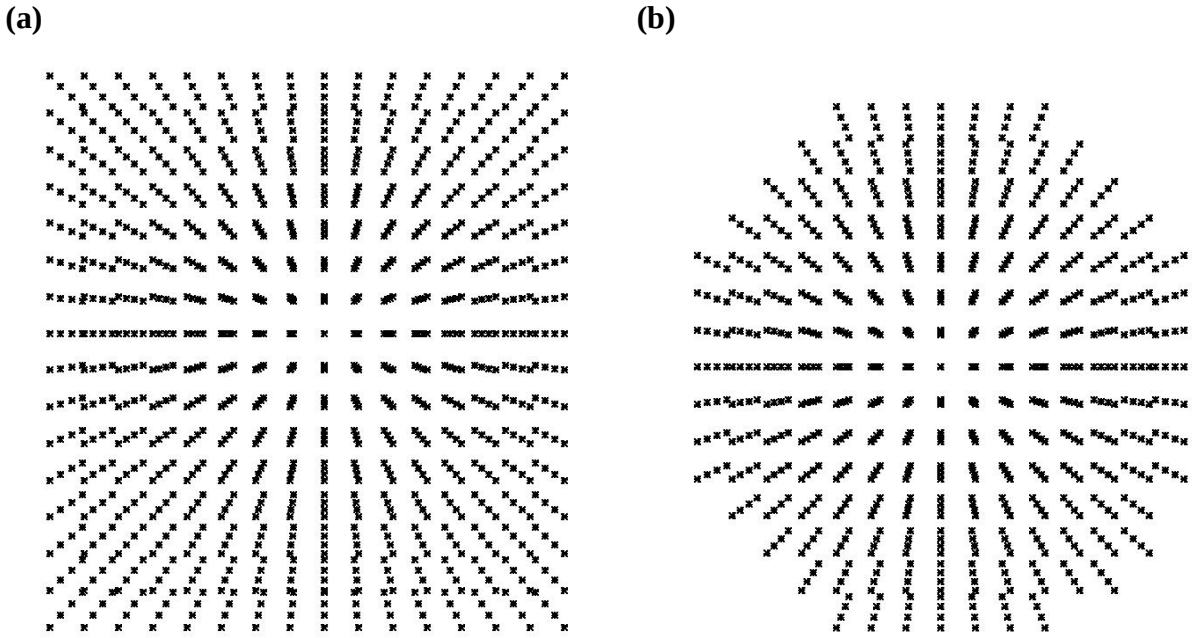

Figure 5: Examples of non-uniform data in k-space acquired during an MPA experiment. (a) Standard SPRITE acquisition, 16x16 gradient steps, 4 FID points. (b) Spiral-SPRITE acquisition, 187 gradient steps, 4 FID points.

(2) It is difficult process all of the acquired MPA data as a single dataset. Any extra data sampled after t_{p0} will not fall onto the Cartesian grid (see Figure 5), and the complete dataset is therefore no longer amenable to processing with an FFT. A possible solution is to use one

of the popular interpolation schemes described earlier for computing a non-uniform Fourier Transform, with the knowledge that this provides only an approximate answer. It would be beneficial, however, to compute the transform exactly but without an unreasonable amount of processing time.

Fortunately, both of these problems can be solved with the same data processing strategy. We have found that the FOV correction can be done using the algorithm described earlier, the Chirp Z-Transform [11]. Furthermore, we have also discovered that the CZT can also be used to exactly compute the Fourier transform for the entire dataset with the same $O(NlogN)$ complexity as the FFT.

3 Algorithms

In this section we derive the mathematical formalism necessary to implement the non-uniform discrete Fourier Transform using a combination of Chirp Z-Transforms. We shall then discuss the implementation of these algorithms in IDL (Interactive Data Language) 6.0, a language designed around matrix operations and which is therefore well suited for image processing applications.

3.1 The DFT for SPRITE MPA Data

We have already seen how, for sampled MRI data, the integral form of the MRI signal equation becomes a discrete summation, of the form

$$\rho(x_m) = \sum_{n=0}^{N_C-1} s(k_n) e^{-2\pi i k_n x_m} \qquad m \,\epsilon\, \{0, 1, \ldots, N_C - 1\} \tag{3.1}$$

where N_C is the total number of collected points. We will now reverse the variable substitution performed in equation 2.26 to restore explicit dependence on G and t:

$$\rho(x_m) = \sum_{j=0}^{N_T-1} \sum_{k=0}^{N_G-1} s(G_k t_p(j)) e^{-i\gamma x_m G_k t_j} \tag{3.2}$$

We have therefore split the single summation into a double summation, one for the N_T time points acquired after each gradient step, and one for the N_G gradient steps that comprise each SPRITE acquisition. The total number of acquired points in the experiment is $N_C = N_G N_T$, and the total number of calculated $\rho(x)$ values is also N_C.

While this expression could be implemented as shown, we will modify the expression in order to simplify the calculations and implementation. First, we wish to map the values of the complex exponential into the range $[0, 2\pi N_G]$, or equivalently, $[-\pi N_G, \pi N_G]$. This change yields the following:

$$\rho(x_m) = \sum_{j=0}^{N_T-1} \sum_{k=0}^{N_G-1} s(G_k t_p(j)) e^{-i\theta} \tag{3.3}$$

$$\theta = \pi N_G (x_m/x_{max})(G_k/G_{max})(t_p(j)/t_{max}) \tag{3.4}$$

where x_{max} represents the largest value of the position coordinate, G_{max} is the maximum value of the magnetic field gradient, and $t_{max} = t_{p0} + (NT-1)\Delta t_p$ is the longest encoding time,

With this choice of θ, both the position x and encoding time t have been scaled into the range $[0, 1]$, while the gradient value has been scaled into the range $[-1, 1]$ (since G_k varies from $-G_{max}$ to $+G_{max}$). θ therefore takes on values between $-\pi N_G$ and $+\pi N_G$ as desired. Notice that we have removed the dependence on the gyromagnetic ratio γ as this will only introduce a constant scaling factor.

In order to create a more generic algorithm we now seek to remove the dependence on the maximum gradient strength G_{max}, which may vary from one experiment to another. For a SPRITE experiment, we can state the following:

$$\begin{aligned} G_k &= -G_{max} + k\Delta G \\ \Delta G &= 2G_{max}/N_G \\ \therefore G_k/G_{max} &= 2k/N_G - 1 \end{aligned}$$

We also adjust the expression for x_m, such that $m = 0$ corresponds to $x = -\frac{1}{2}$ and $m = N_C$ would correspond to $x = +\frac{1}{2}$. This eliminates the need to shift the data by $N_C/2$ pixels when the DFT is complete, and this can be done by changing (x_m/x_{max}) to $(m/N_C - 1/2)$.

After these (mainly aesthetic) modifications we are left with our "computable" form of the Discrete Fourier Transform:

$$\rho(x_m) = \sum_{j=0}^{N_T-1} \sum_{k=0}^{N_G-1} s(G_k t_p(j)) e^{-i\theta} \qquad m \,\epsilon\, \{0, 1, \ldots, N_C - 1\} \tag{3.5}$$

$$\theta = 2\pi N_G (m/N_C - 1/2)(k/N_G - 1/2)(t_p(j)/t_{max}) \tag{3.6}$$

The question remains as to how the DFT given here can be executed on a computer in $O(N_C log N_C)$ time, since it appears to require $O(N_C * N_G * N_T) = O(N_C^2)$ operations.

3.2 Computation of DFT Using Chirp Z-Transform

It will now be shown that the form of the DFT given in eqn. 3.6 is equivalent to the following:

$$\rho(x_m) = \sum_{j=0}^{N_T-1} CZT(x_m, t_p(j))e^{i(\alpha n+\beta)} \tag{3.7}$$

$$CZT(x_m, t_p(j)) = \sum_{k=0}^{N_G-1} s(G_k t_p(j))A^{-k}W^{km} \qquad m\,\epsilon\,\{0,1,\ldots,N_C-1\} \tag{3.8}$$

for an appropriate choice of A, W, α, and β. A and W are simply the parameters discussed in section 2.2, describing a contour in the Z-plane corresponding to the desired image. The angles α and β are phase corrections which are necessary to bring the phase of the CZT and of our reformulated DFT into complete agreement. Without these corrections, the resulting sequences of complex data would not be exactly equal.

As has been mentioned, the fields of view for the datasets corresponding to different values of t_p have to be corrected, such that all of the datasets have the same field of view. The two most practical choices are to match all of the fields of view to either that of $t_p(j) = t_{p0}$ (the largest FOV) or that of $t_p(j) = t_{max}$ (the smallest FOV). Since the latter choice will provide an image with higher resolution, this is the one which we select.

If the entire field of view for a given dataset corresponds to the full unit circle in the z-plane, we are then interested in the region which spans an arc of length $2\pi t_p(j)/t_{max}$ for the image at time $t_p(j)$, and which is centered about the positive horizontal axis. (See Figure 6.)

We can describe this region of the z-plane by the following choice of the contour parameters A and W (see Figure 6):

$$A = e^{-\pi i T_j}, \quad W = e^{-2\pi i T_j/N_C} \tag{3.9}$$

where we have set $T_j = t_p(j)/t_{max}$. Substituting these into (3.8), and then into (3.7), yields

$$\rho(x_m) = \sum_{j=0}^{N_T-1}\sum_{k=0}^{N_G-1} s(G_k t_p(j))e^{\pi i k T_j}e^{-2\pi i k m T_j/N_C}e^{i(\alpha n+\beta)} \tag{3.10}$$

Comparing this with our earlier formulation of the DFT,

$$\rho(x_m) = \sum_{j=0}^{N_T-1}\sum_{k=0}^{N_G-1} s(G_k t_p(j))e^{-2\pi i N_G((m/N_C)-1/2)((k/N_G)-1/2)T_j}, \tag{3.11}$$

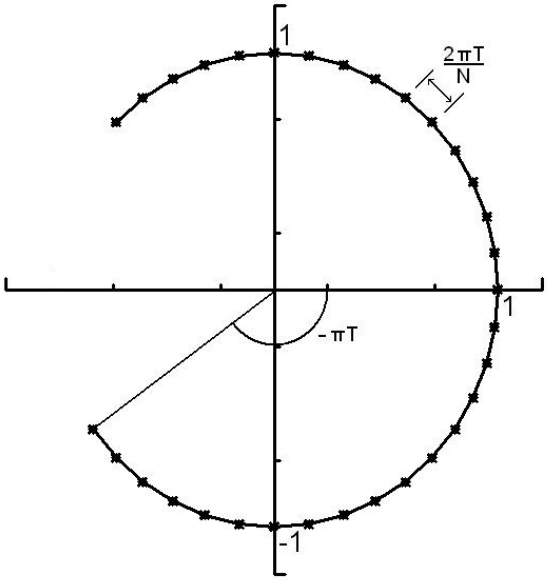

Figure 6: The z-plane contour for a transform that will shrink the field of view by a factor of T.

we see that for these to be equivalent, the arguments of the exponentials must match:

$$\begin{aligned}
\pi ikT_j - 2\pi ikmT_j/N_C + i\alpha m + i\beta &= -2\pi iN_G(m/N_C - 1/2)(k/N_G - 1/2)T_j \\
&= -2\pi iN_GT_j(km/N_CN_G - k/2N_G - m/2N_C + 1/4) \\
&= -2\pi iT_jkm/N_C + \pi ikT_j + \pi iT_jm/2N_T - \pi iN_GT_j/2
\end{aligned}$$

From this we identify the linear phase correction $\alpha = \pi T_j/2N_T$ and the global phase correction $\beta = -\pi N_G T_j/2$.

We have therefore shown that, for the appropriate choice of parameters, the CZT can be used to implement the non-uniform DFT of SPRITE MRI data. This is done by separately applying the CZT to the k-space data for each encoding time t_p to correct the field of view, performing the two phase corrections on the transformed data, and then summing over all N_T FOV- and phase-corrected datasets.

In algorithmic notation, we can describe the procedure as follows:

Input : An N_T by N_G matrix S containing N_T FID points and N_G gradient steps
An N_T-element vector times containing encoding times $t_p(j)$ for the FID points

Output : The N_C-point non-uniform DFT of the dataset S

$N_C \leftarrow N_G * N_T$
out $\leftarrow$ An empty N_C-element vector
sum $\leftarrow$ An empty N_C-element vector
for $i \leftarrow 0$ ***to*** $N_T - 1$ **do**
    $T_j \leftarrow$ times [i] / times [$N_T - 1$]
    out $\leftarrow$ `CalculateCZT` (S [i], N_C, T_j)
    for $j \leftarrow 0$ ***to*** $N_C - 1$ **do**
        $\alpha \leftarrow \pi T_j / 2N_T$
        $\beta \leftarrow -\pi T_j N_G / 2$
        out $\leftarrow$ out * $e^{i\alpha+\beta}$
    end
    sum $\leftarrow$ sum + out
end
return sum

Algorithm 1: `DFT_Using_CZT`

Here we have assumed that `CalculateCZT`(S, N_C, T) is a function that will correctly compute the N_C-point Chirp Z-Transform of the vector S, using the parameter T_j to generate A and W for the desired contour, as outlined in eqn. 3.9. The implementation of this algorithm is the topic of the next section.

The complexity of these algorithms will clearly depend on that of `CalculateCZT()`. As we shall demonstrate, `CalculateCZT()` operates in $O(N_C log N_C)$ time. The algorithm shown will therefore operate in $O(N_T N_C log N_C + 2N_T N_C)$ time, since it calls `CalculateCZT()` and performs two $O(N_C)$ phase corrections, a total of N_T times. The asymptotic complexity is therefore $O(N_T N_C log N_C)$.

It should be noted that the complexity is quadratic in N_T, since $N_C = N_G N_T$, and this will decrease performance significantly for large N_T. If N_T were limited to small values this might be acceptable, but this is not true in general. Indeed, Halse et. al. [11] have shown that the optimum final image is obtained when

$$N_T = \frac{N_G}{2}\left(\frac{1}{T_{max}} - 1\right) \tag{3.12}$$

where T_{max} denotes the maximum extent of FOV scaling ($0 < T_{max} \leq 1$). Although we will

generally attempt to keep T_{max} close to unity, to avoid artificial noise reduction and maintain good resolution, in some cases the SNR is the primary concern and we collect a larger number of points. For these situations we propose a modified version of the algorithm which will compute an image of size $N_C = N_G$. We shall refer to this as the "Non-Expanded" version of the algorithm, in contrast to the "Expanded" version which uses $N_C = N_G N_T$.

To construct the Non-Expanded algorithm, we simply replace N_C with N_G everywhere. This includes the equations of the previous section, and this modification will necessitate a slight change in the linear phase correction, to $\alpha' = \pi T_j/2$. The remainder of the data processing strategy is unaffected. This will provide a superior reconstruction methodology when signal intensity or processing time is the dominant concern, and it will be demonstrated that this method will not reduce the information content of the final image.

3.3 Implementation of the CZT

In this section we shall describe the implementation of the Chirp Z-Transform algorithm, following the example of Rabiner et. al. [16], but making some small modifications to suit our desired application.

Input : An N_G element vector X of SPRITE MPA data
N, an integer specifying the desired size of transform
T, a scaling factor ($0 < T \leq 1$)
Output : The N-point Chirp Z-Transform of the dataset X

$L \leftarrow 2 * N$
$W \leftarrow e^{2\pi T/N}$
$A \leftarrow e^{-\pi T}$
Y $\leftarrow$ an empty L-element vector
for $i \leftarrow 0$ ***to*** $N_G - 1$ **do** Y [i] $\leftarrow A^{-i} * W^{i^2/2} *$ X$[i]$
FY $\leftarrow$ FFT (Y)
V $\leftarrow$ an empty L-element vector
for $i \leftarrow 0$ ***to*** $N - 1$ **do** V [i] $\leftarrow W^{-i^2/2}$
for $i \leftarrow N + 1$ ***to*** $L - 1$ **do** V [i] $\leftarrow W^{-(L-i)^2/2}$
FV $\leftarrow$ FFT (V)
FG $\leftarrow$ FY * FV
G $\leftarrow$ InverseFFT (FG)
for $i \leftarrow 0$ ***to*** $N - 1$ **do** G [i] $\leftarrow$ G [i] * $W^{i^2/2}$
return first N elements of G

Algorithm 2: CalculateCZT

As mentioned, the reason why the Chirp Z-Transform can be calculated at high speed is that, along the contour of equation 2.12, the Z-Transform can be expressed as a convolution. The purpose of the algorithm is then to create the functions which must be convolved, and then carry out the convolution using the FFT.

It can be verified by inspection that Algorithm 2 will implement the convolution described in equation 2.17. The computationally significant steps of the algorithm are the three L-point FFT's used to perform this convolution, each of which will be executed in $O(L\,logL)$ time. However, since $L = 2N$ for this specific case, we can justifiably say that the complexity for this algorithm is $O(NlogN)$, where N is the number of points in the transformed image.

This algorithm has been implemented in IDL, based upon the czt() function in the Matlab signal processing toolbox, which has been optimized for speed. It should be noted that this implementation does not exactly match the explicit form of the Z-Transform given in equation 2.11. Though the magnitudes of the transformed data are identical to within the maximum precision available, there is a discrepancy in the phases of the complex data. At present it is not known why this phase difference is present; attempts to predict its value based on parameters supplied to the algorithm have failed. However, we must correct for this phase difference if the Chirp Z-Transform is to match the DFT precisely.

Fortunately, the phase difference between the CZT of equation 2.11 and the fast CZT as implemented for testing is always global; that is, the phase at each point differs from its true value by the same constant. Once we have applied the derived phase corrections α and β of the previous section, the leftover phase difference may be determined by calculating the phase of the true DFT at any point, and then comparing it with the phase of the CZT at that same point. The difference between these two values becomes our final phase correction, and applying this correction makes the transforms identical.

The logical point at which to calculate the phase of the DFT, for comparison to that of the CZT, is clearly that point which corresponds to x=0. At this point the exponential term in the DFT expression (equation 2.7) equals unity, and the expression for the DFT is simply

$$\rho(0) = \sum_{k=0}^{N_G-1} s(G_k) \tag{3.13}$$

which requires only N_G complex additions. This lets us determine and apply the right phase correction with no increase in the overall computational complexity.

3.4 Generalization to Higher Dimensions

Until this point, the algorithms we have discussed have all been designed to operate on one-dimensional data. The majority of MRI data for both materials and medical applications is two-dimensional, with three-dimensional imaging becoming increasingly common as the speed and sensitivity of MRI techniques improve. The algorithms discussed previously must therefore be generalized to higher dimensionality if they are to be of any practical value.

Since the one-dimensional Chirp Z-Transform, after the appropriate phase corrections, is identical to the one-dimensional DFT, we can exploit the separability of the Fourier Transform and use a succession of one-dimensional CZT operations to process data in multiple dimensions. For example, in two dimensions, we have the following:

Input : An N_T by N_G by N_G array S of SPRITE MPA data
An N_T element vector times containing encoding times for the FID points
Output : The N_C by N_C non-uniform DFT of the dataset S

matrInput ← An empty N_C by N_C matrix
matrStep1 ← An empty N_C by N_C matrix
matrStep2 ← An empty N_C by N_C matrix
matrSum ← An empty N_C by N_C matrix
for $i \leftarrow 0$ ***to*** $N_T - 1$ **do**
    $T \leftarrow$ times [i] / times [N_T-1]
    matrInput ← S [i]
    // CZT across first dimension
    for $j \leftarrow 0$ ***to*** $N_C - 1$ **do**
        matrStep1 [j] ← `CalculateCZT` (matrInput [j], N_C, T)
        apply phase corrections as before
    end
    matrStep1 ← transpose of matrStep1
    // CZT across second dimension
    for $j \leftarrow 0$ ***to*** $N_C - 1$ **do**
        matrStep2 [j] ← `CalculateCZT` (matrStep1 [j], N_C, T)
        apply phase corrections as before
    end
    matrSum ← matrSum + matrStep2
end
return matrSum

Algorithm 3: `2DDFTusingCZT`

For the Expanded algorithm, when we desire a larger image, we have $N_C = N_G\sqrt{N}_T$ as the total number of points along each side of the image. For this reason, this algorithm should only be used to process data with a perfect square number of FID points. In n dimensions, we require that the number of points must be some integer raised to the power n, such that each dimension can be expanded equally and by an integral value.

The Non-Expanded algorithm is not limited in this way and can process a dataset containing any number N_T of FID points, returning an image of size N_G by N_G pixels. As in the one-dimensional case, the Non-Expanded algorithm is derived from the Expanded algorithm by setting $N_C = N_G$ and modifying the linear phase correction as outlined earlier.

In terms of complexity, we already know that the computationally significant part of the algorithm will be the call to `CalculateCZT`, which executes in $O(N_C log N_C)$ time. In each of the two loops this function is called N_C times, and each loops executes N_T times, for an overall running time of $O(N_T N_C^2 log N_C)$ for an N_C by N_C image. As has been discussed earlier, for images with N_T large, the Non-Expanded algorithm will be preferable since its dependence on N_T is linear, whereas the dependence of the Expanded algorithm on N_T is quadratic.

4 Results

One must remember the primary motivation behind collecting non-uniform data during a SPRITE MPA acquisition. It is assumed that this additional data will provide new information about the sample being imaged, allowing us to somehow improve the image. For example, it may allow us to increase the resolution of the image, so we can discern smaller features that might not otherwise be perceptible. It may also increase the signal-to-noise ratio, lessening the effects of random background noise.

It should be noted that similar increases in quality would be observed if the non-uniform data were processed using one of the popular interpolation methods. Indeed, the results of the interpolated non-uniform transforms and the CZT, when compared, are often visually indistinguishable. However, the two algorithms may execute at different speeds, and the algorithm which runs fastest is clearly preferable for practical applications. There may be differences in the accuracy of the algorithms which, though not visible by eye, are significant when compared to the maximum precision available to the computer. Finally, since there is no need to select an interpolation factor, the CZT should generally be a more robust method.

The CZT algorithm as discussed in the previous section has been implemented and tested on a variety of MRI data. This will permit us to verify that its computational complexity is as expected, and that it properly reconstructs an image from SPRITE MPA data, improving the quality of the image without any cost in terms of acquisition time. The Expanded and the Non-Expanded forms of the algorithm will be compared to assess their differences. Finally, we will also show that the CZT is superior to interpolation techniques both in terms of accuracy and running time.

4.1 Running Time Analysis

To confirm that the running time for the CZT algorithm is as claimed, a number of datasets, containing a variety of matrix sizes and numbers of FID points, were reconstructed with the algorithm. Two independent tests were performed; one with N_T fixed and N_G changing, and another with N_G fixed and N_T changing, to determine the dependence of the running time on each parameter. The data were processed with both the Expanded and Non-Expanded versions of the 2D algorithm so running times of the two implementations can be compared.

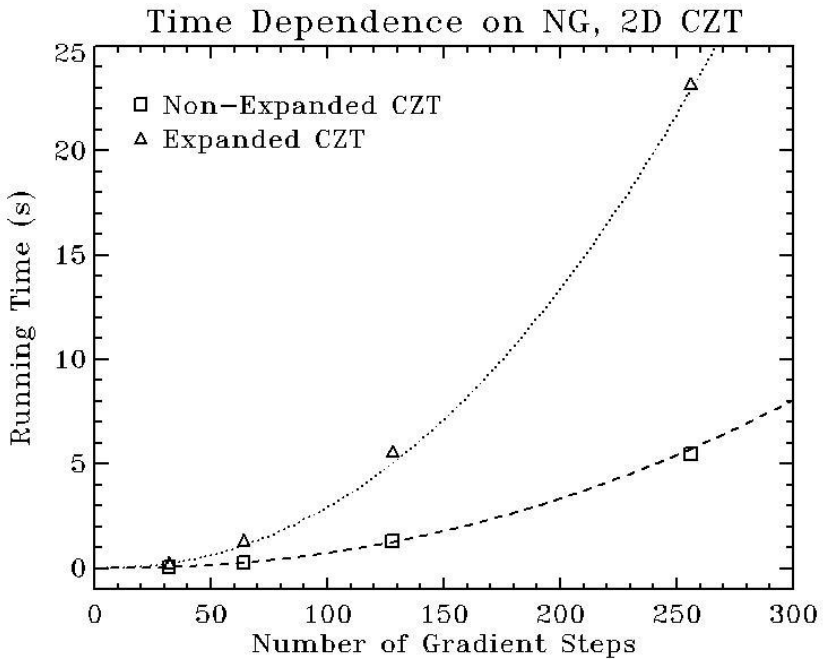

Figure 7: CZT performance for fixed N_T As expected, both Expanded and Non-Expanded forms demonstrate quadratic dependence on N_G in the two-dimensional case.

In Figure 7, we have set $N_T = 4$ and reconstructed datasets with 32, 64, 128 and 256 gradient steps along each side of the acquired k-space matrix. In both cases the data illustrate that the running time with N_T fixed is $O(N_G^2 log N_G)$. Also, we note that the running time for the Non-Expanded algorithm at any given N_G is the same as that of the Expanded case at $N_G/2$, but the curves are otherwise very similar.

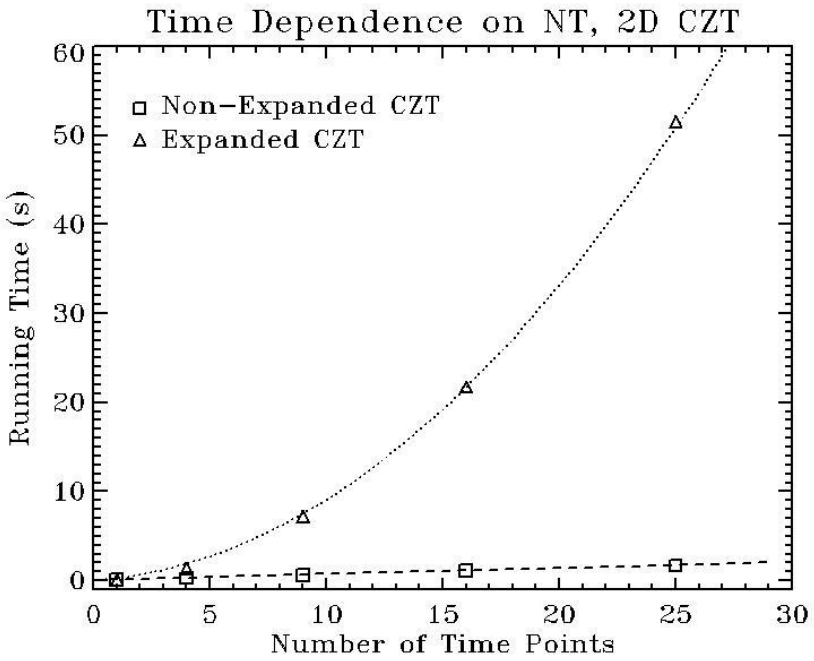

Figure 8: CZT performance for fixed N_G. Dependence of the Expanded form on N_T is quadratic, whereas the Non-Expanded form follows a linear dependence.

In Figure 8, the number of gradient steps N_G was fixed at 64, and the number of FID points varied between 1 and 25. Only perfect square values of N_T were chosen since the Expanded CZT can only reconstruct datasets of that size. Here the difference between the Expanded and Non-Expanded cases becomes evident, since the Non-Expanded case demonstrates an $O(N_T)$ complexity while the Expanded case is $O(N_T^2)$.

4.2 Processed MRI Data - Signal Enhancement

The first MRI image which has been enhanced with the CZT is an image of a fibre-reinforced polyester resin, a material which is of interest to the automobile industry for structural parts applications. Here we are imaging local water concentration through the hydrogen nuclei in each molecule. The overall water concentration of this sample is only 3-4%, hence the MRI

signal is very low. We will therefore attempt to use the CZT to enhance the signal-to-noise ratio of this image, and compare the data processed with the CZT to other data which was signal-averaged during the acquisition.

The experiment was performed on a 2.4T superconducting magnet having a horizontal bore of inner diameter 32 cm, and manufactured by Nalorac (Martinex, CA). The console was a Tecmag (Houston, TX) Apollo. A 32-strut quadrature birdcage coil (Morris Instruments, Ottawa), driven by a 2kW AMT (Brea, CA) 3455 amplifier, was used as the RF probe. The magnetic field gradient was created with a set of water cooled gradient coils having an inner diameter of 7.5 cm and a maximum gradient strength of 100 G/cm, powered by three Techron (Elkhart, IN) 8710 amplifiers.

The sample has a time constant of T_2^* = 650 μs and a spin-lattice relaxation time T_1^* = 65 ms. It is cylindrical with a height of 8 mm and a diameter of 38 mm. The imaging method was Spiral-SPRITE, and the spiral data from each time point were sorted into a 64x64 matrix for processing. A total of 125 FID points were collected at each gradient step; the field of view for the final FID point, t_{max} = 325 μs, is approximately 8 cm by 8 cm. All of the CZT-processed images were scaled to this field of view so their quality may be directly compared.

Three separate experiments were performed with this apparatus and sample, each with a different number of signal averages performed during the acquisition. This means that each k-space point is sampled multiple times and the results are averaged within the spectrometer; this increases signal quality but also increases acquisition time. The measurement with no extra averaging was performed in 6.5 seconds. The second and third measurements were undertaken with four and sixteen averages per point, respectively. The total acquisition time for these measurements was 25 and 101 seconds.

Data from all three measurements were processed using the CZT. Five different numbers of FID points were used to enhance the images; 1, 4, 9, 16 and 25 FID points. Using a single FID point corresponds to reconstruction with the FFT, and it has been verified that reconstructing one point with the CZT is equivalent to using the FFT. Nonetheless, the CZT was used throughout the experiment for consistency. The results are shown in Figure 9.

The signal-to-noise ratio for each image is the ratio of the average intensity of the central region, and the average intensity of a region in the background. The results of this calculation are given in Table 1, and plotted in Figure 10.

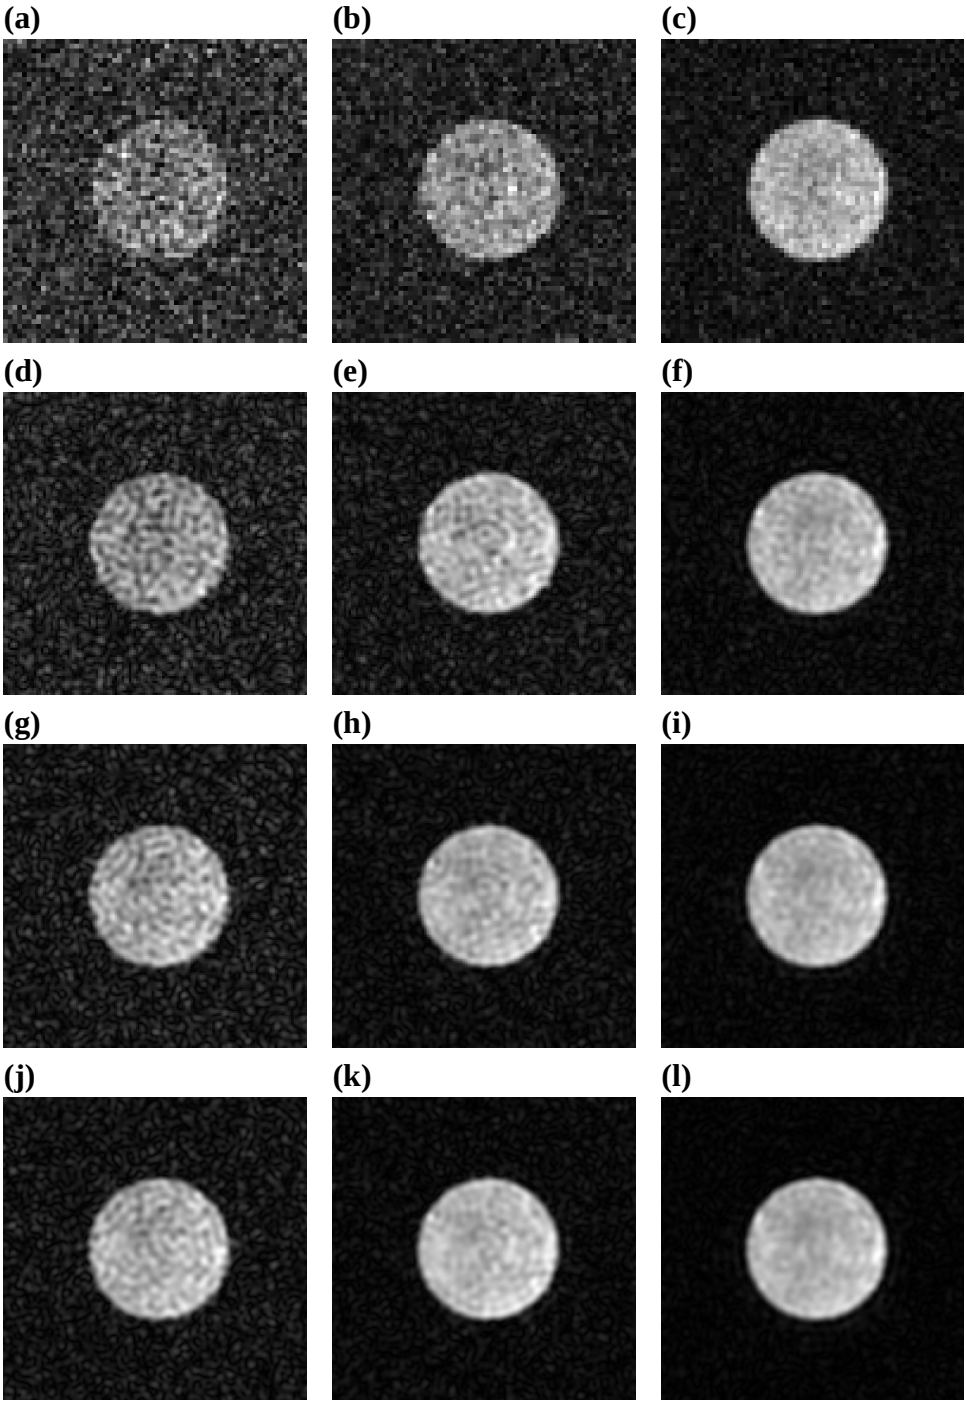

Figure 9: Signal-to-noise ratio improvement of an MRI image. Each column represents a different amount of signal averaging during the acquisition: 1, 4, and 16 averages for the first, second and third columns respectively. Each row represents a different amount of FID points used in CZT image reconstruction; 1, 4, 9 and 16 FID points respectively.

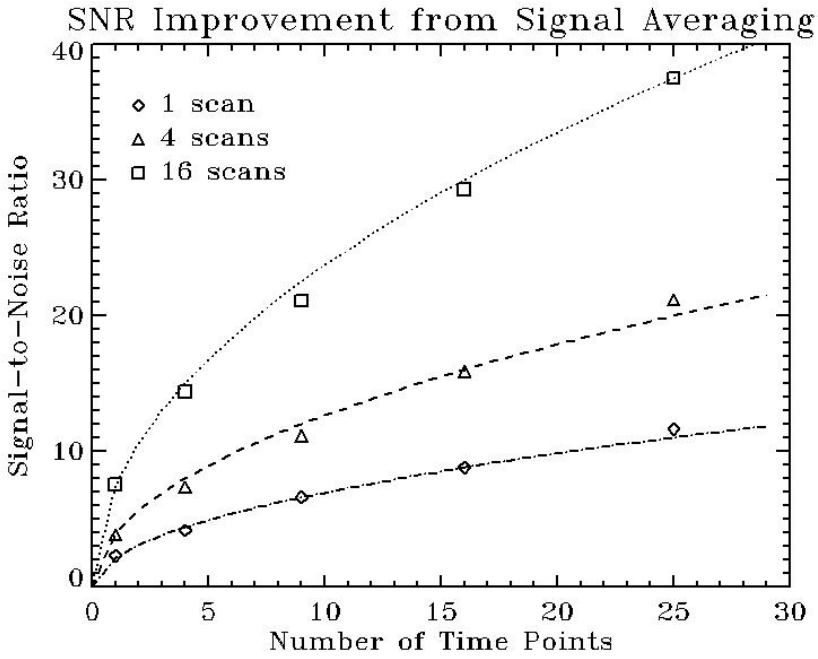

Figure 10: Signal-to noise ratio improvement due to signal averaging, both by performing extra scans during the acquisition and by processing extra FID points with the CZT.

These results are significant in two ways. First, the signal-to-noise ratio of an image increases proportionately with the square root of the number of averages, as expected from theory. Second, the signal-to-noise improvement obtained using the CZT is comparable to that obtained by extra signal averaging during the acquisition.

For example, the signal-to-noise ratios for the images with 16 averages (images **c**, **e**, **j**) are similar regardless of whether these averages were performed using traditional signal averaging (image **c**), MPA (image **j**), or some combination thereof (image **e**). However, the data used to create image **e** took four times longer to acquire than the data used to create image **j**, and the data

# FID Pts.	SNR, 1 scan	SNR, 4 scans	SNR, 16 scans
1	2.30622	3.84558	7.54888
4	4.14624	7.34596	14.3680
9	6.62382	11.1346	21.0802
16	8.79874	15.8625	29.2950
25	11.6291	21.1504	37.5141

Table 1: Raw signal-to-noise ratio values for various numbers of scans and FID points.

used to create image **c** took sixteen times longer to acquire. Though the CZT- enhanced images took correspondingly longer to process, this extra time investment is more than accounted for by reduced acquisition times, and in practical applications, processing times are not a significant issue unless they become extremely long.

4.3 Processed MRI Data - Resolution Enhancement

In a two-dimensional SPRITE acquisition with N_G gradient steps in each dimension, we collect enough information to create an $N_G \times N_G$ image, with a nominal resolution of FOV/N_G, where FOV is the field of view. To improve this resolution we must acquire a larger matrix in k-space, by increasing the number of gradient steps while maintaining a constant field of view. However, this will also increase the total acquisition time, since more data must be sampled. Doubling the number of gradient steps in a 2D image acquisition requires us to sample four times more data, for instance.

Since the Expanded version of the CZT returns an image whose size is $N_C^2 = N_G^2 N_T$, it seems reasonable to expect that we will observe some improvement in the resolution due to the larger matrix size. An experiment was performed to determine if such a resolution improvement was taking place.

For this experiment we require a sample that will allow us to quantitatively determine the resolution improvement yielded by the CZT, if any. We employ a *resolution phantom*, a sample with features of a known size that will allow us to determine the smallest size of feature which can be imaged by a given technique. Our phantom is a circular section of crosslinked polymer which contains several cuts and holes of various sizes. The dimensions of the phantom and its features are given in Figure 11.

The phantom was imaged using the same experimental setup as in the previous section. Its time constant T_2^* was $\sim 300\mu s$. 25 FID points were collected at each gradient step of the Spiral-SPRITE sequence, which collected a 64x64 k-space matrix. The maximum encoding time used was $t_{max} = 85\mu s$ and all CZT images were scaled to this value. Images were reconstructed with 1, 4, and 9 FID points; it was found that using more points does not noticeably alter the resolution of this particular image. The results are shown in Figure 12.

If we define the resolution of an image in terms of the smallest discernable feature, then there is clearly an apparent resolution increase in these images. In the original image, neither

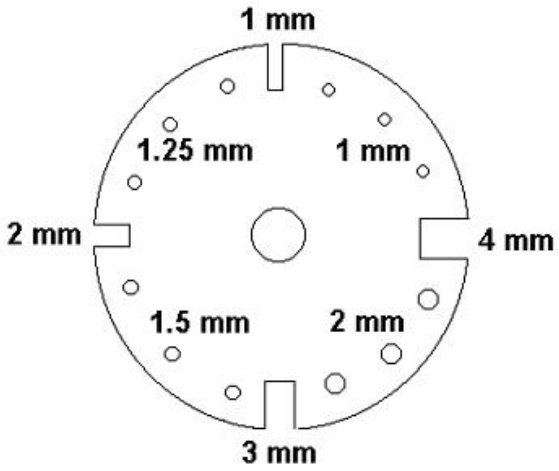

Figure 11: Dimensions of the polymer resolution phantom and its various features.

the 1.5mm holes nor the 1.0mm holes were visible. The image created with 4 FID points allows us to distinguish all of the smaller holes, and when 9 FID points are used in the reconstruction, many of the features are more sharply resolved.

However, using this definition of resolution, the resolution and signal-to-noise ratio of an image are always interrelated, with an increase in one leading to an increase in the other; an image with better SNR will have more easily distinguishable features. To determine whether the resolution enhancement observed in these images was greater than that which would be expected through SNR improvement alone, the same images were reconstructed with the Non-Expanded CZT, whose final size is that of the acquired matrix, and which should not exhibit any resolution enhancement beyond that provided by increasing SNR.

As illustrated in Figure 12(d), though an image created with the Non-Expanded CZT appear to be rougher and fills fewer pizels, all features that are distinguishable in an image created with the Expanded CZT can still be observed. This indicates that there is no resolution enhancement associated with the increased matrix size, and no information is lost if the image is processed with the Non-Expanded CZT instead of the Expanded CZT. This suggests that the Non-Expanded algorithm, with its lower complexity, is a better choice for most practical applications since it will execute faster but will reconstruct an image that nonetheless contains as much information as possible.

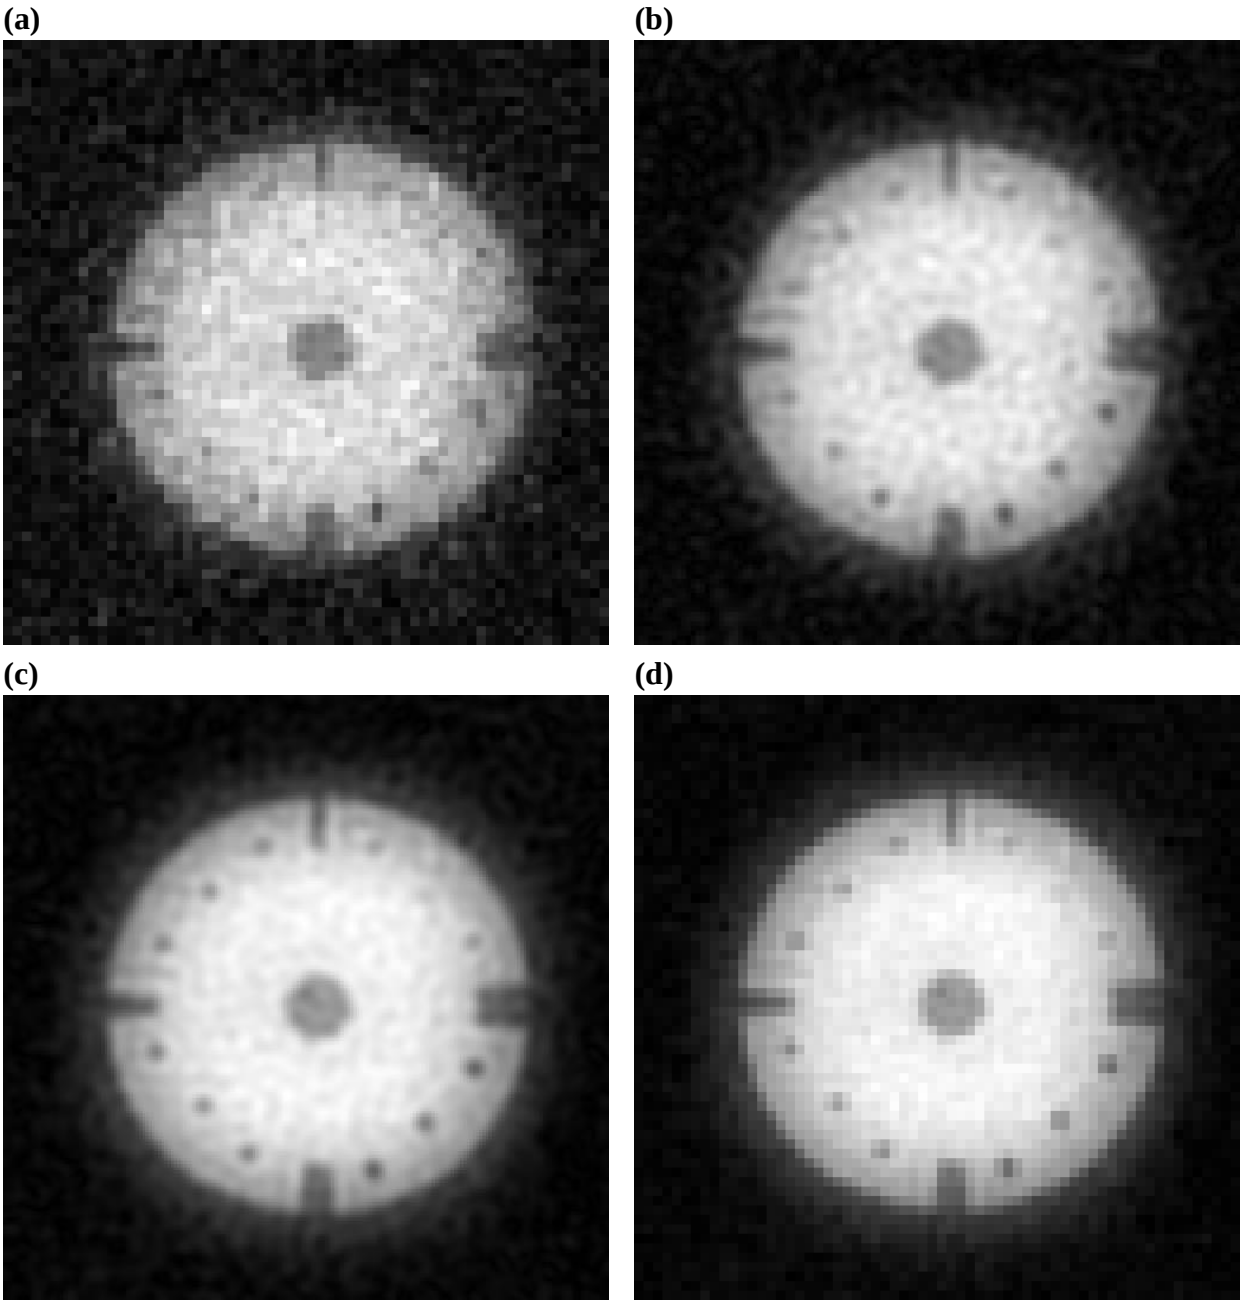

Figure 12: Reconstructed images of a polymer resolution phantom. (a) Original image, 1 FID point. (b) Image reconstructed with Expanded CZT and 4 FID points. (c) Image reconstructed with Expanded CZT and 9 FID points. (d) Image reconstructed with Non-Expanded CZT and 9 FID points; the resolution matches that of figure (c) despite smaller number of pixels.

4.4 Comparison with Existing Techniques

As mentioned earlier, many other non-uniform Fourier transform algorithms exist, and these often provide similar results to the CZT. However, these other algorithms may not execute in the same amount of time as the CZT, and they may not be as accurate in terms of their difference from the analytic DFT. We now proceed to an examination of one such algorithm, the interpolation scheme formulated by Dutt and Rokhlin [7], and generalized to two-dimensional MRI data by Sarty et. al. [17].

The Dutt-Rokhlin algorithm approximates the non-uniform Fourier transform by a combination of over-sampling and interpolation, which are controlled by two parameters m and q. The interpolation parameter q depends in turn on a value b according to the relation $q = \lceil 4b\pi \rceil$. Dutt and Rokhlin found that a reasonable choice of values for these parameters was $m = 2$, $b = 0.6$ and thus $q = 8$. These parameters are used to calculate an size-mN vector of uniformly spaced data from N non-uniform data points. After applying the FFT to the oversampled matrix, an exponential weighting is applied to recover the transformed data.

The running time for the algorithm is stated as $O(mNlogN+mNq)$ for the one-dimensional case treated by Dutt and Rokhlin in their original paper. Sarty provided the logical extension of the algorithm to two dimensions; this operates in $O(mN^2logN + mN^2q)$ time on a dataset containing N^2 points. This two-dimensional version has been implemented for comparison against the CZT in terms of speed and accuracy.

For this comparison we have used a different image of the previous section's resolution phantom (with slightly different imaging parameters), as well as a one-dimensional dataset which was acquired from a prototype Electron Paramagnetic Resonance Imaging unit [9]. Each was processed with the Expanded CZT, the Dutt-Rokhlin algorithm (using the Sarty implementation for the 2D case), and the Discrete Fourier Transform. The Non-Expanded CZT was not used in this comparison since the image it produces is smaller than those returned by the other methods; these images cannot be directly compared since they are of different sizes.

Speed testing was performed on a 1.07 GHz Intel Celeron processor. The accuracy of the algorithms was quantified by subtracting the images reconstructed by method being tested (CZT or Dutt-Rokhlin) from the DFT-reconstructed images, and dividing by the magnitude of the DFT-reconstructed image to obtain a relative error for each pixel. The mean values of this relative error then can be compared; the results are summarized in Table 2.

Algorithm	Mean Rel. Error, 1D	Mean Rel. Error, 2D
CZT	1.97 x 10^{-15}	1.23 x 10^{-13}
DRS, b=0.6	8.38 x 10^{-5}	4.12 x 10^{-5}
DRS, b=1.0	8.57 x 10^{-7}	1.08 x 10^{-6}
DRS, b=2.0	3.036 x 10^{-11}	6.77 x 10^{-11}
DRS, b=3.0	4.4 x 10^{-15}	6.54 x 10^{-14}

Table 2: Mean error (vs. discrete Fourier Transform) for various reconstruction methods.

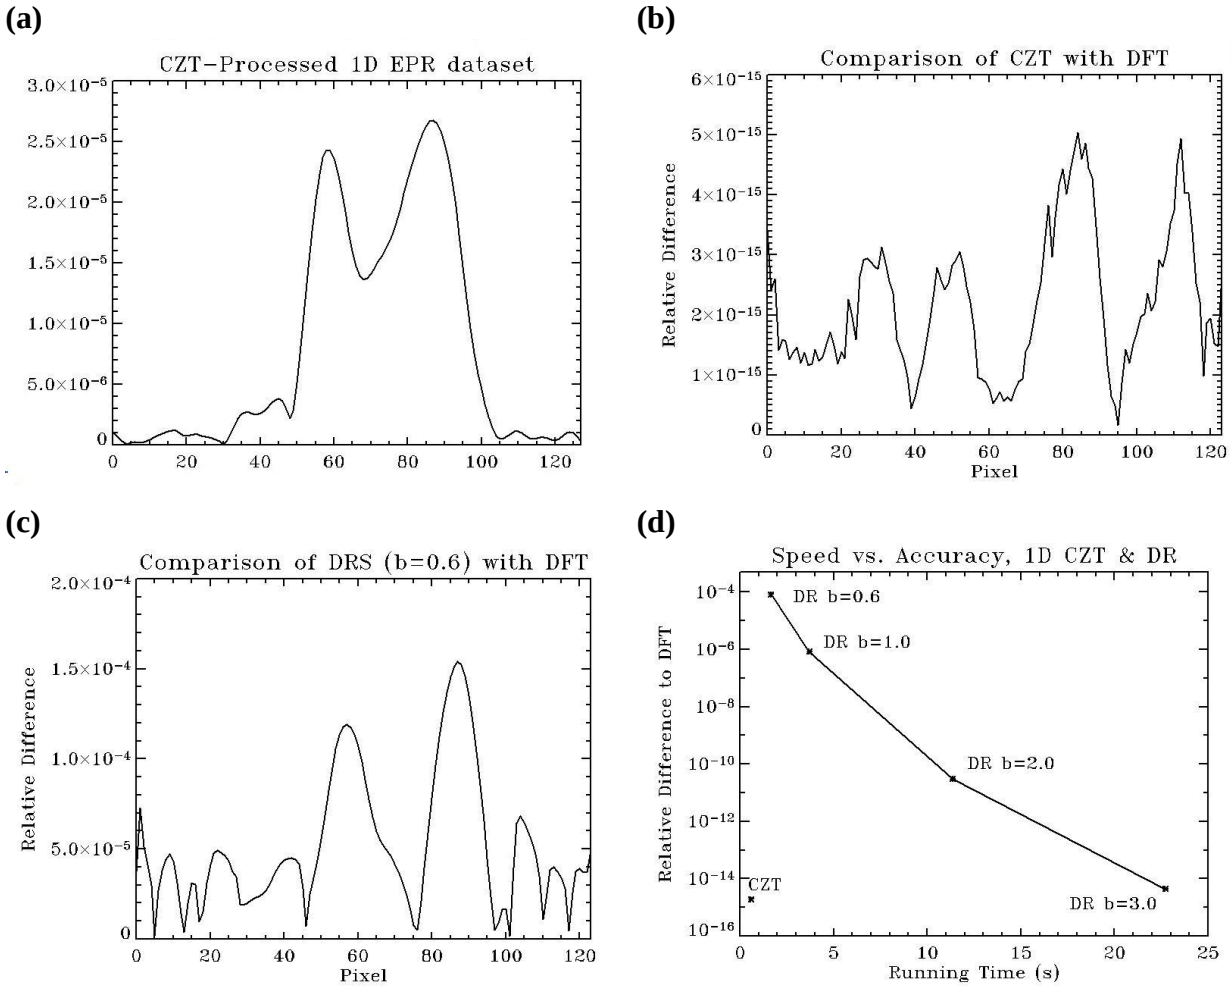

Figure 13: Comparison of Dutt-Rokhlin and CZT reconstructions for a 1D image, $N_G = 32$, $N_T = 4$. (a) MRI image as reconstructed by CZT. (b) Differences between the DFT and CZT. (c) Differences between the DFT and DRS. (d) DRS running times for various values of b.

Consider the one-dimensional images in Figure 13. The first important result is that there is a significant difference in the accuracy of the algorithms. As expected, the Chirp Z-Transform computes the non-uniform DFT to considerable accuracy. In this case, the accuracy of the CZT is generally on the order of the machine epsilon for double- precision floating point values, $\epsilon_{mach} \sim 10^{-16}$. However, with the suggested value of $b = 0.6$, the Dutt-Rokhlin algorithm has an accuracy of only 10^{-5}. Increasing b allows the accuracy to approach ϵ_{mach} but the running time will increase steeply.

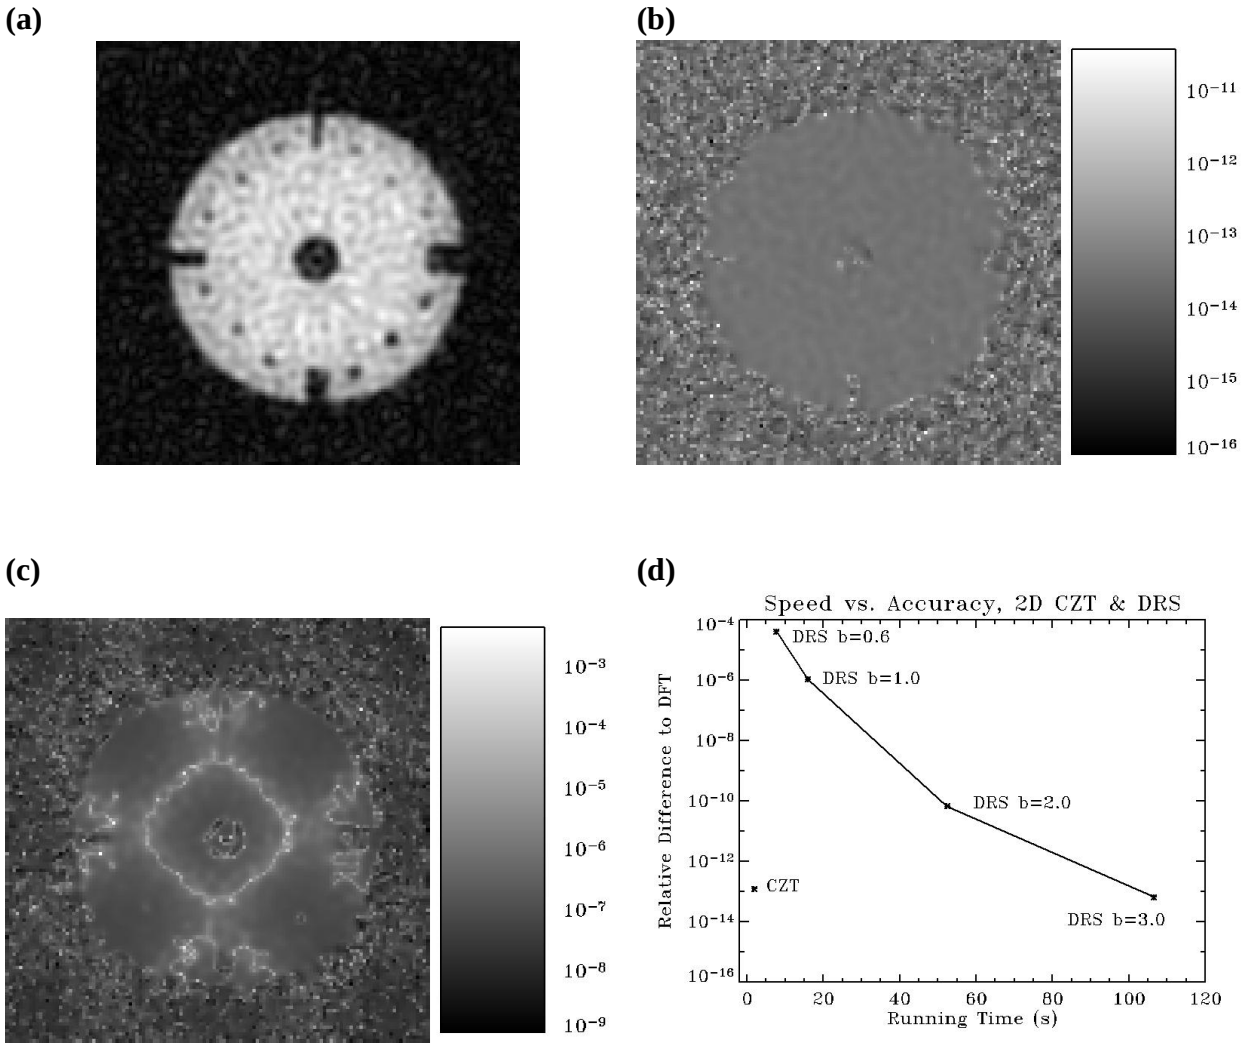

Figure 14: Comparison of the Dutt-Rokhlin-Sarty and CZT reconstructions for a 2D image, $N_G = 64$, $N_T = 4$. (a) MRI image as reconstructed by CZT. (b) Differences between the DFT and CZT. (c) Differences between the DFT and DRS with b=0.6. (d) DRS running times for various values of b.

The accuracy difference is similar in the two-dimensional case, as illustrated in Figure 14. Since the 2D CZT is a series of 1D CZT operations along orthogonal axes, errors accumulate throughout the transform. For the CZT, since the 1D error is on the order of 10^{-15}, we expect 2D errors in the range of 10^{-14} to 10^{-13}, which is confirmed in figure 14(b). The error in the Sarty reconstruction, however, is still on the order of 10^{-5} to 10^{-4} for the lowest value of b. Again, increasing b will increase accuracy, but a severe penalty is paid in terms of running time.

Figures 13(d) and 14(d) are plots of accuracy as a function of running time that summarize the results of this testing, and illustrate that the CZT is preferable to the Sarty reconstruction in all cases. The case where $b = 0.6$ yields similar running times to the CZT but only at the cost of greatly reduced accuracy. Comparable accuracy is achieved for $b = 3$ but running time has now been increased by a factor of almost 50 in both the 1D and 2D cases.

Finally, we will note once more that these comparisons have all been carried out using the Expanded CZT algorithm; as we have seen, the Non-Expanded version has a faster running time and would demonstrate even greater advantages over Dutt-Rokhlin for large N_T.

5 Conclusions

It has been known for some time that the standard algorithm used to process MRI data, the Fast Fourier Transform, is not well suited for processing the data acquired during an MPA SPRITE experiment. In the past, interpolation techniques have been used to reconstruct such non-uniform data with reasonable success. However, it has now been demonstrated that these data can be processed much more accurately, and much more rapidly, using an algorithm based upon the Chirp Z-Transform.

The Chirp Z-Transform is a more general formulation of the Fourier Transform, but which has the same $O(NlogN)$ asymptotic complexity. By a careful choice of the parameters which define the contour along which the Chirp Z-Transform is evaluated, and a proper correction to the phase of the transformed data, the Discrete Fourier Transform of the SPRITE data can be computed in $O(N_T NlogN)$ time for a dataset containing N_T FID points and whose reconstructed image contains N pixels. This is readily generalized to multiple dimensions.

Two versions of the algorithm exist; the Expanded CZT increases the size of the final image while the Non-Expanded CZT does not. Both methods reconstruct images with the same information content in terms of resolution and signal-to-noise ratio. Of the two, the Non-Expanded

CZT will be the algorithm of choice for practical applications because of its increased speed and because it has no restrictions concerning how many FID points can be used. In contrast, the number of FID points input to the Expanded CZT must be a perfect square.

When compared against the performance of existing interpolation methods, such as that of Dutt and Rokhlin, the CZT offers advantages both in terms of execution time and in terms of accuracy. For the one-dimensional case, accuracy is on the order of 10^{-15}, while accuracy for two-dimensional images is typically on the order of 10^{-13}. For applications in Magnetic Resonance Imaging, such high accuracy makes no perceptible difference to the image quality; our primary concern is the speed of the algorithm.

However, the reformulation of the Discrete Fourier Transform in terms of a sum of Chirp Z-Transforms has other potential applications beyond those considered in this work. Any kind of non-uniform mesh which contains certain regularities may be amenable to processing with the CZT. For example, since the Fourier transform can be used to solve some kinds of differential equations, the CZT might assist in the numerical solution of these equations. In such a scenario, the accuracy of the CZT would be an important factor, as would its speed.

In terms of further application to MRI, the CZT presents itself as a technique for enhancing the signal-to-noise ratio of images at virtually no time cost. This will be valuable in the study of materials whose MRI signal strength is naturally low. Instead of performing extra signal averaging scans, which may incur a time penalty, a Multiple Point Acquisition can be performed and a high-quality image reconstructed in a fraction of the time. The CZT could also be used to assist the imaging of dynamic systems by reducing acquisition time, allowing researchers to image rapid dynamic processes in materials.

At the moment, the application of the CZT in magnetic resonance imaging is limited to SPI-type measurements, such as SPRITE and the family of related techniques for 2D and 3D imaging. The non-uniform data often acquired in clinical spiral methods, for example, cannot be processed in this manner because the data lack the regularities in the k-space sampling scheme which allowed us to reformulate the Discrete Fourier Transform in terms of the Chirp Z-Transform. Within this restriction, however, we have demonstrated that the Chirp Z-Transform is a valuable tool for the enhancement of SPRITE MRI images, and future developments will surely permit us to extend the range of its capabilities.

Bibliography

[1] B.J. Balcom, R.P. MacGregor, S.D. Beyea, D.P. Green, R.L. Armstrong and T.W. Bremner, "Single Point Ramped Imaging with T_1 Enhancement (SPRITE)", *Journal of Magnetic Resonance* Series A, vol.123, pp.131-134, 1996.

[2] S.D. Beyea, B.J. Balcom, P.J. Prado, A.R. Cross, C.B. Kennedy, R.L. Armstrong and T.W. Bremner, "Relaxation Time Mapping of Short T_2^* Nuclei with Single-Point Imaging (SPI) Methods, *Journal of Magnetic Resonance* vol.135, pp.156-164, 1998.

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

[4] E. Oran Brigham, "The Fast Fourier Transform and its Applications", Prentice Hall, 1988.

[5] Paul T. Callaghan, "Principles of Nuclear Magnetic Resonance Microscopy", Oxford University Press, 1991.

[6] J.W. Cooley, J.W. Tukey, "An Algorithm for the Machine Calculation of Complex Fourier Series", *Mathematics of Computation*, vol.19 no.90 pp.297-301, 1965.

[7] A. Dutt and V. Rokhlin, "Fast Fourier Transforms for Nonequispaced Data", *SIAM Journal for Scientific Computation*, vol.14 no.6, pp. 1368-1393, 1993.

[8] S. Emid and J.H.N. Creyghton, "High Resolution NMR in Solids", *Physica B* vol.128, pp.81-83, 1985.

[9] Frederic G. Goora, "Electron Paramagnetic Resonance Imaging Using a Pure Phase Encode Technique" M.Sc.(Eng) Thesis, University of New Brunswick, 2003.

[10] M. Halse, D. Goodyear, B. MacMillan, P. Szomolanyi, D. Matheson and B.J. Balcom, "Centric Scan SPRITE Magnetic Resonance Imaging", *Journal of Magnetic Resonance*, in press 2003.

[11] M. Halse, S. Romanzetti, J. Kaffanke, B. MacMillan, J. Rioux, N.J. Shah and B.J. Balcom, “Centric Scan SPRITE Magnetic Resonance Imaging: Optimization of SNR, Resolution and Relaxation Time Mapping”, *Journal of Magnetic Resonance*, submitted 2003.

[12] J.I. Jackson, C.H. Meyer, D.G. Nishimura and A. Macovski, “Selection of a Convolution Function for Fourier Inversion using Gridding”, *IEEE Transactions on Medical Imaging* vol.10 no.3, 1991.

[13] P. Mansfield and P.K. Grannell, “NMR ‘diffraction’ in solids?”, *Journal of Physics* C vol.6 pp.L422-L426, 1973.

[14] C.H. Meyer, B.S. Hu, D.G. Nishimura and A. Macovski, “Fast spiral coronary artery imaging”, *Magnetic Resonance in Medicine* vol.28, pp.202-213, 1992.

[15] E.M. Purcell, H.C. Torrey and R.V. Pound, “Resonance Absorption by Nuclear Magnetic Moments in a Solid”, *Physical Review* vol.69 pp.37-38, 1946.

[16] L.R. Rabiner, R.W. Schafer and C.M. Rader, “The Chirp Z-Transform algorithm”, *IEEE Transactions on Audio and Electroacoustics*, vol.17 no.2 pp. 86-92, 1969.

[17] G.E. Sarty, R. Bennett and R.W. Cox, “Direct Reconstruction of Non-Cartesian k-Space Data Using a Nonuniform Fast Fourier Transform”, *Magnetic Resonance in Medicine*, vol.45 pp. 908-915, 2001.

[18] L. Sha, H. Guo and A.W. Song, “An Improved Gridding Method for Spiral MRI Using Nonuniform Fast Fourier Transform”, *Journal of Magnetic Resonance*, vol.162, pp.250-258, 2003.

Errata and Addenda

The following are errors in the manuscript which were not corrected until after submission.

On page 18, formula (3.10) should read

$$\rho(x_m) = \sum_{j=0}^{N_T-1} \sum_{k=0}^{N_G-1} s(G_k t_p(j)) e^{\pi i k T_j} e^{-2\pi i k m T_j / N_C} e^{i(\alpha m + \beta)}$$

On page 19, the linear phase correction should be $\alpha = \pi T_j / N_T$, and the last line of the series of equations should read

$$= -2\pi i T_j k m / N_C + \pi i k T_j + \pi i T_j m / N_T - \pi i N_G T_j / 2$$

This will also affect Algorithm 1 on page 20 and the correction to α on page 21.

Algorithm 1 also suffers from a poor choice of indices in the two **for** loops, and several small errors exist in the last few lines. The correct formulation of the algorithm is as follows:

for $j \leftarrow 0$ ***to*** $N_T - 1$ **do**
    $T_j \leftarrow$ times [j] / times [$N_T - 1$]
    out $\leftarrow$ `CalculateCZT` (S [j], N_C, T_j)
    $\alpha \leftarrow \pi T_j / N_T$
    $\beta \leftarrow -\pi T_j N_G / 2$
    for $k \leftarrow 0$ ***to*** $N_C - 1$ **do** out [k] $\leftarrow$ out [k] * $e^{i(k\alpha+\beta)}$
    sum $\leftarrow$ sum + out
end
return sum

On page 20, formula (3.12) is actually a formula for the number of time *intervals*, not the number of time points. (3.12) should read

$$N_T = \frac{N_G}{2}(\frac{1}{T_{max}} - 1) + 1$$

Finally, the discussion on page 22 regarding the difference in phase between the implemented CZT and the analytical form of the CZT has been rendered moot. A way has been found to modify the implementation of the algorithm such that this phase difference no longer manifests itself. This modification also perceptibly increases the accuracy of the algorithm. For the 1D dataset considered in table 2 and figure 13, for example, the mean relative error is decreased from 1.97 x 10^{-15} to 4.00 x 10^{-16}.