

MR Image Reconstruction from Pseudo-Hex Lattice Sampling Patterns Using Separable FFT

Jae-Ho Kim, Fred L. Fontaine *

Abstract— Common MRI sampling patterns in k -space, such as spiral trajectories, have nonuniform density and do not lie on a rectangular grid. We propose mapping the sampled data to a pseudo-hex lattice, taking advantage of its approximate isotropic nature in k -space and square nature in the reconstructed image space. The group structure of the lattice is exploited to implement the Fourier transform computations on the data using a separable FFT algorithm, which provides significant computational efficiency. We suggest this method can be generalized to multiresolution lattices, in which the signal is represented in different regions in k -space with varying sampling densities. The operations on index sets and mapping to separable FFT can be implemented efficiently in software or custom hardware (e.g., FPGA).

Keywords: magnetic resonance imaging, multidimensional signal processing, discrete Fourier transforms

1 Introduction

In magnetic resonance imaging (MRI), data is measured on a finite set of points in k -space (the spatial frequency domain), usually in two or three dimensions. Normally, the reconstructed image needs to be recovered on a finite set of points on a uniformly sampled square grid within a prescribed field of view (FOV). However, the measurements in k -space are usually in a nonuniform pattern, with varying sampling densities in different regions of k -space. The sampling pattern in k -space is usually determined by physical constraints such as the ability to change the excitation magnetic field precisely and quickly, as well as by the desire to have a higher sampling density near the origin of k -space, which generally contains more significant information. In fact, relative to standard results of sampling theory, the region near the origin of k -space is typically oversampled while the outer regions of k -space are undersampled.

The image reconstruction process usually involves regriding in k -space to a square lattice, and applying standard fast Fourier transform (FFT) techniques [1].

*Thanks to Weill Medical College of Cornell University. Manuscript submitted April 16, 2007. The authors are with The Cooper Union for the Advancement of Science and Art, Department of Electrical Engineering, New York, NY 10003 Tel/Fax: 1-212-353-4331/4341 Email: fred@cooper.edu

Sometimes these steps are combined, as in the case of the non-uniform FFT (NUFFT) algorithm [2]. Regriding is essentially an interpolation process that can involve approximating “ideal” interpolation functions, or can be achieved by solving a constrained (weighted) least-squares problem. The simplest form of regriding is mapping the data at an acquired point \vec{k}_i to the nearest point in a square lattice of a prescribed density; this is called nearest-neighbor regriding. In order to use a square grid that is not excessively dense, more precise interpolation techniques are generally required.

We propose, instead, to resample the data onto what we call a multi-resolution lattice (MRL): points that lie on a base lattice of maximal density, or sublattices (or cosets of sublattices) of the base lattice, for example, as shown in Figure 1. The local sampling density of the MRL reflects that of the original measurement process. Once an MRL sampling structure is obtained, a variety of techniques of multirate signal processing can be applied, including optimized decimation and interpolation filter banks and generalized Cooley-Tukey FFT computations. A key feature of a lattice is that, like a square grid, it has a group structure, which permits a Fourier transform to be defined on it. Specifically, if V is a nonsingular matrix, the lattice generated by V , denoted \mathcal{L}_V , is the set of all points of the form $V\vec{n}$ where \vec{n} is an integer vector. If U is a generator of a sublattice (i.e., $\mathcal{L}_U \subset \mathcal{L}_V$), then the elements of the quotient group $\mathcal{L}_V/\mathcal{L}_U$ can be associated with points in \mathcal{L}_V that lie within a unit cell for \mathcal{L}_U , and the “periodicity” that gives rise to a classical DFT here takes the form of periodicity with respect to the sublattice \mathcal{L}_U . A finite Fourier transform can be defined on this group, resulting in points in the target image space on the lattice $\mathcal{L}_{\hat{U}}$, where $\hat{U} = (U^{-1})^T$, within a unit cell for $\mathcal{L}_{\hat{V}}$, where $\hat{V} = (V^{-1})^T$. Note that the sampling density in k -space, $|\det V|$, and the area (or volume) of the FOV in the target image space, $|\det \hat{V}|$, are inversely proportional, as expected; similarly, the “bandwidth” in k -space, essentially the largest $|\vec{k}|$ at which samples are taken, determined by $|\det U|$, fixes the resolution (sampling density) in the target image space, essentially $|\det \hat{U}|$, and these quantities are inversely related.

By analogy, the ordinary DFT computes the Fourier transform on the quotient group $\mathbb{Z}/N\mathbb{Z}$, i.e., the cyclic group with N elements. The DFT is normally applied to the index set $\{0, 1, \dots, N - 1\}$. Now, identification of subgroups of $\mathbb{Z}/N\mathbb{Z}$ leads to decimation-in-time and decimation-in-frequency type operations, which are the basis for FFT. In the multidimensional case, when the sampling pattern is square (or, more generally, rectangular), the DFT operation is separable and separable FFT can be applied. The separability leads to computational advantages in addition to the Cooley-Tukey decompositions of the FFT. In the case of the DFT over a lattice quotient group, $\mathcal{L}_V/\mathcal{L}_U$, identification of subgroups also can give rise to decompositions similar to decimation-in-time and decimation-in-frequency operations. It turns out that the DFT over $\mathcal{L}_V/\mathcal{L}_U$ can also be decomposed in such a way as to be computed, at its core, by a separable FFT operation (this follows from the Smith form decomposition [3] of integer matrices).

When these more general sampling patterns are employed in k -space, many standard signal processing algorithms, such as multirate filtering (e.g., decimation and interpolation) can be applied. One sampling pattern that can provide certain advantages is a hexagonal lattice, which has a more isotropic nature than a rectangular lattice. Consider the task of nearest-neighbor regridding. For the same sampling density, the maximum distance from any point in k -space to a lattice point is about 13% less in a hexagonal lattice than in a square lattice, because of the more circular nature of the pattern. However, use of a hexagonal lattice in k -space leads to a hexagonal lattice in the target image space; this in turn would require regridding after image reconstruction to a square lattice, which is normally required for display and other post-processing operations. A pseudo-hex lattice, on the other hand, is a rational lattice that approximates a hexagonal one, and can lead to a rectangular lattice in the target image space.[4] That is, \mathcal{L}_V , the lattice of support in k -space, is pseudo-hexagonal, while $\mathcal{L}_{\hat{U}}$, the lattice of support in the target image space is square.

Our basic proposed image reconstruction process is as follows:

- Perform nearest-neighbor regridding from the original points in k -space to points on a pseudo-hex lattice \mathcal{L}_V , bounded within a Voronoi cell of a sublattice \mathcal{L}_U , which is chosen so that $\mathcal{L}_{\hat{U}}$ is square.
- Perform multirate operations, as necessary, to acquire data in a final form prepared for Fourier transformation. For example, the original lattice can have a high density so that the errors associated with nearest-neighbor regridding are negligible, and then the signal can be downsampled using optimized filters which can be determined in a more systematic way than optimized interpolation functions that are

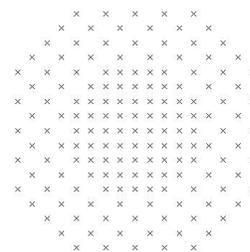


Figure 1: Points on a multi-resolution pseudo-hex lattice.

designed for original, nonuniform sampling patterns.

- Implement the computation of the generalized DFT for $\mathcal{L}_V/\mathcal{L}_U$ data set with a core separable FFT operation. This requires a remapping of the data, possibly both before and after the separable FFT operation, analogous to bit-reversed addressing that occurs in the one-dimensional case.

In this paper, we present efficient nearest-neighbor algorithms, which require minimal computational effort, and demonstrate the mapping of the general form of the DFT considered here to a separable FFT. A sample reconstructed image is shown.

2 Discrete Fourier Transform on Multi-Resolution Lattices

Although the pseudo-hex lattice described later in this paper is two-dimensional, in this section we will be more general. We associate vectors in \mathbb{R}^D with $D \times 1$ column matrices; $\langle \vec{x}, \vec{y} \rangle$ denotes the inner product of vectors \vec{x}, \vec{y} in \mathbb{R}^D ; for an invertible matrix A , $\hat{A} = (A^{-1})^T$; and $e(\alpha) = \exp(-j2\pi\alpha)$. For $\vec{x} \in \mathbb{R}^D$, $|\vec{x}|$ denotes its Euclidean length, and for a finite set \mathcal{S} , $|\mathcal{S}|$ denotes its cardinality. Our signals, corresponding to baseband representations of voltages detected by receiver coils, are complex scalar or vector valued (as arises in the case of multichannel MRI).

Given an invertible real matrix V , the lattice generated by V is $\mathcal{L}_V = \{\vec{x} : \vec{x} = V\vec{n}, \vec{n} \in \mathbb{Z}^D\}$. We denote an associated unit cell as \mathcal{U}_V . The reciprocal lattice is $\mathcal{L}_{\hat{V}}$, and a reciprocal unit cell is $\mathcal{U}_{\hat{V}}$. In general, $\mathcal{L}_U \subset \mathcal{L}_V$ iff $U = VM$ where M is an invertible integer matrix, and $\mathcal{L}_U = \mathcal{L}_V$ iff $U = VE$ where E is unimodular (an integer matrix with $|\det E| = 1$) [3]. Thus, given a lattice, the generator matrix is not uniquely determined, nor is the unit cell. However, the volume of the unit cell, $|\det V|$, and the sampling density, $1/|\det V|$, are uniquely determined. For $\mathcal{L}_U \subset \mathcal{L}_V$, the quotient group $\mathcal{L}_V/\mathcal{L}_U$ is com-

prised of the cosets of \mathcal{L}_U in \mathcal{L}_V ; $|\mathcal{L}_V/\mathcal{L}_U| = |\det V^{-1}U|$, does not depend on choice of generators. From $\hat{V} = \hat{U}M^T$, we have $\mathcal{L}_{\hat{V}} \subset \mathcal{L}_{\hat{U}}$, and $|\mathcal{L}_{\hat{U}}/\mathcal{L}_{\hat{V}}| = |\mathcal{L}_V/\mathcal{L}_U|$. An index set $\mathcal{I}(U, V)$ (or $\mathcal{I}(\mathcal{L}_U, \mathcal{L}_V)$) is a minimal complete set of coset representatives of \mathcal{L}_U in \mathcal{L}_V (i.e., contains exactly one point from each coset). Every unit cell \mathcal{U}_U of U generates an index set via $\mathcal{I}(U, V) = \mathcal{U}_U \cap \mathcal{L}_V$. A signal g is \mathcal{L}_V -periodic (or V -periodic) if, for all \vec{k} in the domain of g , $g(\vec{k} + \vec{\xi}) = g(\vec{k})$ for all $\vec{\xi} \in \mathcal{L}_V$. The set of \mathcal{L}_U -periodic signals with support on \mathcal{L}_V is denoted $\mathcal{P}(\mathcal{L}_U, \mathcal{L}_V)$, or $\mathcal{P}(U, V)$; every such signal is uniquely determined by its values on any index set $\mathcal{I}(\mathcal{L}_U, \mathcal{L}_V)$.

If $g : \mathcal{L}_V \rightarrow \mathbb{C}^L$ is a complex vector valued signal with support on \mathcal{L}_V , we define the *lexicographic form* of g as $g_V : \mathbb{Z}^D \rightarrow \mathbb{C}^L$ given by $g_V[\vec{n}] = g(V\vec{n})$. Note that the lexicographic form depends on the choice of generator matrix. Here, square brackets $[\cdot]$ denote lexicographic index vectors in \mathbb{Z}^D , and parentheses (\cdot) denote “physical” coordinates in \mathbb{R}^D . With this structure, we can express the finite Fourier transform and inverse transform formulas associated with the quotient group $\mathcal{L}_V/\mathcal{L}_U$ as:

$$\begin{aligned} G(\vec{r}) &= \sum_{\vec{k} \in \mathcal{I}(U, V)} g(\vec{k}) e(-\langle \vec{k}, \vec{r} \rangle) \\ g(\vec{k}) &= \frac{1}{|\mathcal{I}(U, V)|} \sum_{\vec{r} \in \mathcal{I}(\hat{V}, \hat{U})} G(\vec{r}) e(\langle \vec{k}, \vec{r} \rangle) \end{aligned} \quad (1)$$

For $g \in \mathcal{P}(U, V)$, we have $G \in \mathcal{P}(\hat{V}, \hat{U})$. We denote the formulas (1) as the *multiresolution lattice discrete Fourier transform (MRL-DFT)* and inverse transform formulas, respectively. Notice that these formulas do not depend on the choice of generators or index sets (by virtue of the signals’ periodicity). The MRL-DFT can be expressed in lexicographic form, with the result that $G_{\hat{U}}$ and g_V are related via the conventional multidimensional DFT [3] with respect to periodicity matrix $M = V^{-1}U$, as shown in equation (2) below. Note that the lexicographic formulation of the Fourier transform depends on the choice of generator matrices, and in fact the physical structure of the lattice is lost because it forces the base lattices of support in both domains to be \mathbb{Z}^D .

The MRL-DFT equations (1) can be rewritten in lexicographic form as follows. From $U = VM$, we have $\hat{V} = \hat{U}M^T$. With I the identity matrix, we denote $\mathcal{I}(M, I) = \mathcal{I}(M)$; that is, this is a complete set of integer vectors that are distinct modulo M . $\mathcal{I}(M^T)$ is defined similarly. Then:

$$\begin{aligned} G_{\hat{U}}[\vec{m}] &= \sum_{\vec{n} \in \mathcal{I}(M)} g_V[\vec{n}] e(-\langle \vec{m}, M^{-1}\vec{n} \rangle) \\ g_V[\vec{n}] &= \frac{1}{|\det M|} \sum_{\vec{m} \in \mathcal{I}(M^T)} G_{\hat{U}}[\vec{m}] e(\langle \vec{m}, M^{-1}\vec{n} \rangle) \end{aligned} \quad (2)$$

Now, every invertible integer matrix M can be expressed

in Smith form as:

$$M = E_1 M_0 E_2 \quad (3)$$

where E_1, E_2 are unimodular and $M_0 = \text{diag}\{\mu_1, \dots, \mu_D\}$ where μ_i are positive integers. From $U = VM$, if we take $U_0 = UE_2^{-1}$ and $V_0 = VE_1$, then we have $U_0 = V_0M_0$. Note that $\mathcal{L}_{U_0} = \mathcal{L}_U$ and $\mathcal{L}_{V_0} = \mathcal{L}_V$, and $M_0 = M_0^T$. Also:

$$\langle \vec{m}, M_0^{-1}\vec{n} \rangle = \sum_{i=1}^D \mu_i^{-1} m_i n_i \quad (4)$$

where $\{m_i\}_{i=1}^D, \{n_i\}_{i=1}^D$, are the components of \vec{m}, \vec{n} , respectively. If we take a “natural” choice for $\mathcal{I}(M_0) = \mathcal{I}(M_0^T)$ as $\mathcal{I}_0 = \{\vec{p} \in \mathbb{Z}^D : 0 \leq p_i \leq \mu_i - 1\}$, then, in light of (4), the DFT formulas (2) reduce to those of a D -dimensional separable DFT, with respective radices μ_1, \dots, μ_D in each dimension. Thus, a separable FFT can be applied.

Then what is the difference between the original equation (1) and the ordinary separable multidimensional DFT? We could select as generators U_0, V_0 at the outset, although as will be discussed later that may not always be possible. The problem is that the choice of \mathcal{I}_0 for both $\mathcal{I}(M)$ and $\mathcal{I}(M^T)$ corresponds to the index set $V_0\mathcal{I}_0$ in k -space, and $U_0\mathcal{I}_0$ in the target image space, and neither may be the desired set. Specifically, they will not be, in general, confined to the Voronoi regions for \mathcal{L}_U and $\mathcal{L}_{\hat{V}}$, respectively. For example, in the 1-D case, the natural index set for an $N = 8$ point DFT is $\{0, 1, \dots, 7\}$, but in many cases the centered set $\{-4, -3, \dots, 2, 3\}$ is more desirable. In the multidimensional case, this requires a mapping between each desired index set, $\mathcal{I}(U, V)$ and $\mathcal{I}(\hat{V}, \hat{U})$, and \mathcal{I}_0 . It may be difficult to characterize this mapping in a concise form (i.e., other than a look-up table). However, the procedure we outline in this paper does provide guidance for determining efficient representations of this mapping in many situations.

3 Nearest Neighbor Regridding for Pseudo-Hex Lattices

We consider MRI data originally sampled on an irregular pattern, specifically points that do not lie on a lattice and with a sampling density that varies through k -space. There is often a significantly higher density near the origin, and the sampling density in many cases is somewhat isotropic (i.e., does not vary significantly with direction). The first step is a nearest neighbor regridding onto a base lattice \mathcal{L}_V of maximal density. This is equivalent to a perturbation in k -space measurement. The Fourier transform relates the points in k -space to the physical coordinate vector \vec{r} in the target image space through the factor $e(-\langle \vec{k}, \vec{r} \rangle)$. Thus, the perturbation of the sample point in k -space introduces phase errors

on the measure data, with maximum phase error given by $\Delta\phi_{\max} \leq 2\pi R \left| \Delta\vec{k} \right|_{\max}$, where R is the maximum distance of any point in the FOV in image space from the origin, and $\left| \Delta\vec{k} \right|_{\max}$ is the radius of the circumscribing circle for the Voronoi cell of \mathcal{L}_V . Since there are other distortion effects, such as undersampling and bandlimited sampling in k -space (since the object of interest in image space has finite extent, its actual spectrum in k -space has infinite extent), and necessarily imperfect interpolation and noise, the value $\Delta\phi_{\max}$ does not have to be prohibitively small. For example, in standard regriding methods, to reconstruct a 256×256 image, k -space grid sizes from 256×256 up to 1024×1024 are typically used. Here, we employ pseudo-hex grids of comparable sizes.

A generator for a hex lattice is $V_{\text{hex}} = \begin{bmatrix} 2/\sqrt{3} & -1/\sqrt{3} \\ 0 & 1 \end{bmatrix}$. The Voronoi cell is a regular hexagon, but there is no square sublattice. A pseudo-hex lattice is obtained by using a rational approximation to V_{hex} , for example $V = \begin{bmatrix} 8/7 & -4/7 \\ 0 & 1 \end{bmatrix}$ [4]. In this case, $M = p \begin{bmatrix} 7 & 4 \\ 0 & 8 \end{bmatrix}$, for $p \in \mathbb{Z}^+$, yields a generator $U = VM$ for a square lattice, so that the MRL-DFT generates an image that has support on a square grid. For identical sampling density, $\left| \Delta\vec{k} \right|_{\max}$ is about 12.3% less in the pseudo-hex lattice than in a square lattice.

For an arbitrary point \vec{k} , we want to find a nearest lattice point onto which the data can be mapped. In other words, we must find an integer vector \vec{n} such that $\vec{e} = \vec{k} - V\vec{n}$ has minimum length, $|\vec{e}|$. In general, a search among several indices is necessary. In particular, $\hat{n} = \text{round}(V^{-1}\vec{k})$, where every component of $V^{-1}\vec{k}$ is rounded to the nearest integer, may not be the correct choice. This is because distance is measured in the physical coordinate space, \vec{k} , not the lexicographic space, \vec{n} . In fact, depending on the choice of V , the best choice for \hat{n} may not even be rounding each coordinate of $V^{-1}\vec{k}$ either up or down. The goal is to develop an algorithm that requires searching through a minimal possible set of \hat{n} vectors, since each test requires computing a distance.

In our approach, by expressing the column vectors in V in an appropriate orthogonal basis, we can express $|\vec{e}|^2$ as:

$$|\vec{e}|^2 = \frac{64}{65} (\rho_1 - n_1)^2 + \frac{65}{49} \left[\frac{-1008}{65} (\rho_1 - n_1) + (\rho_2 - n_2) \right]^2 \quad (5)$$

where $\vec{n} = [n_1 \ n_2]^T$ and $\vec{\rho} = [\rho_1 \ \rho_2]^T = V^{-1}\vec{k}$. Given n_1 , then n_2 must be chosen to minimize the second term. We should pick n_1 by rounding ρ_1 either up or down. However, for the pseudo-hex lattice we consider

here, the circumscribing radius of the Voronoi cell, which provides a bound on the maximum value of $|\vec{e}|$, is $65/98$; therefore, we can reject one of these choices outright if it would cause the first term to exceed (the square of) this bound. Thus, we obtain the following simplified nearest neighbor algorithm:

1. Compute $\vec{\rho} = [\rho_1 \ \rho_2]^T = V^{-1}\vec{k}$.
2. Compute $\varepsilon = |\rho_1 - \text{round}(\rho_1)|$.
3. If $\varepsilon \leq 1 - \frac{65\sqrt{65}}{8 \times 98} \approx 0.3316$, then $n_1 = \text{round}(\rho_1)$ and n_2 is given by:

$$n_2 = \text{round} \left(\rho_2 - \frac{1008}{65} (\rho_1 - n_1) \right) \quad (6)$$

4. Otherwise, compute the total error associated with the following two choices: $n_1 = \lfloor \rho_1 \rfloor$, and n_2 as above, or $n_1 = \lceil \rho_1 \rceil$, and n_2 as above. Select the (n_1, n_2) pair that minimizes $|\vec{e}|^2$.

Note that, since $|\rho_1 - \text{round}(\rho_1)| \leq 0.5$, on average we need to perform a comparison between two choices only about 33% of the time; otherwise, the nearest neighbor is computed directly without the need for trial-and-error comparisons.

4 Lexicographic Mapping

Here we discuss the algorithm for mapping the MRL-DFT over the pseudo-hex lattice to a form suitable for application of a separable FFT. The Smith form of the M matrix given above is $M = E_1 M_0 E_2$ where:

$$E_1 = \begin{bmatrix} 1 & 0 \\ 16 & -1 \end{bmatrix}, E_2 = \begin{bmatrix} 7 & 4 \\ 2 & 1 \end{bmatrix}, M_0 = \begin{bmatrix} 1 & 0 \\ 0 & 56 \end{bmatrix} p \quad (7)$$

with p is a positive integer determining the overall size of the data set. In standard MRI applications, the target is a 256×256 square grid, but sometimes a higher density lattice is used for pre-processing, say 1024×1024 . With the pseudo-hex lattice, the total number of points needs to be a multiple of 56 (as indicated by $|\det M| = 56p^2$). This suggests using 256×224 grid, with $p = 32$, or 1024×896 , with $p = 64$. The k -space indices in $\mathcal{I}(U, V)$ are typically scaled so that the coordinates lie in the range $-128 \leq k_i < 128$. Thus, the appropriate U_0 matrix (whose Voronoi cell lies in that region of k -space) is:

$$U_0 = \begin{bmatrix} -256 & 1024 \\ 512 & -1792 \end{bmatrix} \quad (8)$$

For a 256×224 grid, we use:

$$V_0 = \begin{bmatrix} -8 & 4/7 \\ 16 & -1 \end{bmatrix} \quad (9)$$

and for a higher density 1024×896 grid we use:

$$V_0 = \begin{bmatrix} -2 & 1/7 \\ 4 & -1/4 \end{bmatrix} \quad (10)$$

The previous section described an efficient nearest neighbor regridding algorithm for the first V_0 matrix; a similar algorithm can be derived for the second V_0 choice. Here we similarly describe a method for generating an index set $\mathcal{I}(U, V)$ comprised of points in \mathcal{L}_V inside a Voronoi cell of \mathcal{L}_U , and associating these points with the indices for the separable DFT associated with M_0 . Again, for illustration purposes, we consider the case $p = 32$ only.

Since $M_0 = \text{diag}\{32, 1792\}$, the core DFT operation is a 32-point DFT in one dimension and 1792-point DFT in the other dimension. The association between $\mathcal{I}(M_0)$, as required by the separable FFT operation, and $\mathcal{I}(U, V)$ is given as follows:

1. $\mathcal{I}(M_0) = \{0 \leq p_1 \leq 31\} \times \{0 \leq p_2 \leq 1791\}$.
2. Compute $\mathcal{J} = V_0 \mathcal{I}(M_0)$.
3. If $a \bmod b$ means the integer a' , $0 \leq a' \leq b - 1$ such that b divides $a - a'$, then:

$$a' = ((a + b/2) \bmod b) - b/2 \quad (11)$$

produces an index in the range $-b/2 \leq a < b/2$. We apply this formula componentwise to the points in \mathcal{J} as follows:

$$\mathcal{I}(U, V) = ((\mathcal{J} + 128) \bmod 256) - 128 \quad (12)$$

The resulting points are in the range $-128 \leq k_i < 128$.

4. The points in $\mathcal{I}(U, V)$ lie in \mathcal{L}_V . The corresponding lexicographic index vectors \vec{n} , such that $V_0 \vec{n} \in \mathcal{I}(U, V)$, are computed via:

$$\mathcal{I}'(M_0) = V^{-1} \mathcal{I}(U, V) \quad (13)$$

Thus, $\mathcal{I}'(M_0)$ is the set of integer vectors that produce points in the Voronoi cell of U_0 when multiplied by V_0 , and $\mathcal{I}(M_0)$ is the set of integer vectors that are used as indices for the separable FFT operation. A similar process can be used in reverse.

Once the FFT operation is applied, the indices $\mathcal{I}(M_0)$ associated with the standard DFT operation must be mapped to an index set $\mathcal{I}'(M_0)$ such that $\hat{U} \mathcal{I}'(M_0)$ corresponds to an index set of $\mathcal{L}_{\hat{U}}/\mathcal{L}_{\hat{V}}$ inside a Voronoi cell of $\mathcal{L}_{\hat{V}}$. This Voronoi cell has an approximately hexagonal shape. Because V is a rational matrix, this region can be described by a set of linear inequalities with rational coefficients, and thus the mapping of an arbitrary

point $\vec{r} \in \mathcal{L}_{\hat{U}}$ to a point in the Voronoi cell can be obtained by a sequence of integer modulo operations similar to (12). The formulation is a bit more complex because these operations cannot be applied componentwise, and the details are omitted here for brevity. However, it is important to note that a precise formulation can be developed, and it can be realized using fixed-point (i.e., integer) arithmetic. Thus it can be implemented in an efficient manner, both in software and in custom hardware (e.g., FPGA) implementations.

We consider one other situation where we must associate different sets of lexicographic indices. When we compute the $\mathcal{L}_V/\mathcal{L}_U$ MRL-DFT, the set of lexicographic indices representing points in k -space depends on the choice of generator V , and the lexicographic indices representing points in target image space depends on the choice of generator U . The physical coordinates in the two spaces do not change, since the Voronoi cells in each domain do not depend on the choice of generator. However, since any two choices for generators are related via a unimodular matrix, for example $V' = VE$, we can map the associated lexicographic indices readily through an integer arithmetic operation, namely if $V' \vec{n}' = V \vec{n}$, then:

$$\vec{n}' = E \vec{n} \quad (14)$$

5 Results

With the methods described above, an image is reconstructed from a set of 57344 lattice points via separable 32×1792 FFT. The computational complexity of the core FFT operation is $32 \times 1792 \times (\log 32 + \log 1792)$, compared with $(32 \times 1792)^2$ without a separable FFT approach. The nearest neighbor regridding operation requires computing and comparing the results of two possible choices for only about 33% of the points; in other cases, the nearest neighbor lattice point can be computed directly. We also have presented a systematic mapping between the lexicographic indices corresponding to the standard index set for the lattice (i.e., located in a Voronoi cell) and the indices used by the separable FFT.

In the case of the nearest neighbor regridding, when multiple points map to the same lattice point, their value is averaged. Where the sampling density is more sparse, typically for large $|\vec{k}|$, there are lattice points which are not the nearest neighbors to any of the original sample points, in which case the data at these lattice points is taken to be zero.

This result, when applied to data obtained from a phantom by a GE Signa 1.5T scanner, is shown in Figure 2.

We are not proposing that a simple nearest neighbor regridding is appropriate for general image reconstruction problems. Instead, we have demonstrated the feasibility of the core operation of image reconstruction built

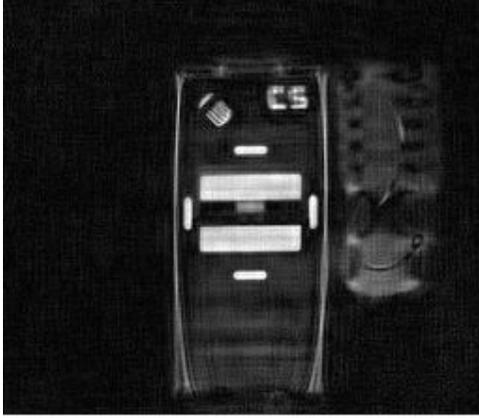


Figure 2: Reconstructed image.

around a non-rectangular grid. In the next section we discuss how these results can be utilized as the basis for more sophisticated image reconstruction operations to be developed.

6 The Multiresolution Lattice Framework

The proposed strategy is to first employ a nearest neighbor regridding to a lattice of sufficient density that the errors associated with the regridding are negligible compared to other effects, such as imprecision in the measurement process. In general, the original data is sampled at spatially varying densities. Thus, by applying appropriate multirate operations on the dense lattice (e.g., decimation filters), the signal can be represented in a multiresolution framework.

Suppose the base lattice of maximal sampling density is \mathcal{L}_V , and the data points are confined to the Voronoi cell associated with a coarse lattice $\mathcal{L}_U \subset \mathcal{L}_V$. The signal, in general, would be sampled on various sublattices (or their cosets) that lie between \mathcal{L}_U and \mathcal{L}_V . Thus, a multiresolution signal is associated with a structure similar to wavelet packets [5]. This structure allows optimized and flexible procedures, such as multirate filtering and decimation and interpolation to alternative lattice structures, as well as generalized Cooley-Tukey FFT computations, to be developed. For example, the passband of a decimation or interpolation filter is associated with a certain region in the FOV of the image space, and thus suppression of aliasing and imaging distortions, as well as amplitude and phase distortions, can be achieved in a prescribed region of the FOV [3]. In the regions of k -space that are undersampled, regridding to a lattice with uniform density creates a large number of samples that actually represent the same measurement. By contrast, in our approach, a sparser lattice is employed locally. Alternative methods, for example those which avoid regridding and

instead perform “direct” reconstruction, say from a least-squares approach or based on sophisticated interpolation kernels, are significantly more analytically difficult and therefore are harder to adapt or fine-tune to particular needs [1], [2].

A key process in reconstructing images from such a multiresolution structure is the ability to compute the DFT for an arbitrary quotient group $\mathcal{L}_{V'}/\mathcal{L}_{U'}$, where $\mathcal{L}_U \subset \mathcal{L}_{U'} \subset \mathcal{L}_{V'} \subset \mathcal{L}_V$. We have outlined a procedure for efficient mapping of such DFT operations to separable FFTs, specifically for the case of a pseudo-hex lattice. In particular, the process of index mapping can be expressed compactly using fixed-point arithmetic, and can be implemented efficiently in software or custom hardware (e.g., FPGA) when a look-up table implementation is not practical. Suppose we start with $\mathcal{L}_V/\mathcal{L}_U$ where the generator matrices are chosen so that $U = VM$ with M diagonal. If, say, we change the base lattice \mathcal{L}_V to a sublattice $\mathcal{L}_{V'} \subset \mathcal{L}_V$ (corresponding to decimation), then the matrix M' relating U and V' via $U = V'M'$ may no longer be diagonal. Thus, working with the MRL framework may require changing the generator matrix associated with a particular sublattice, but as indicated above, this change of generator can be achieved by integer arithmetic (i.e., multiplication by a unimodular matrix) on the lexicographic index set.

Our particular choice of a pseudo-hex lattice in k -space offers the advantages of direct reconstruction to a rectangular grid (avoiding the need for postprocessing regridding), with a hexagonal FOV (avoiding computation of points in the corners of image space, which are often unnecessary in MRI), and with smaller phase errors (caused by the distance in k -space from original measurements to lattice points) than those associated with square grids in k -space.

References

- [1] Sedarat, H., Nishimura, D. G., “On the Optimality of the Gridding Reconstruction Algorithm,” *IEEE Trans. Medical Imaging*, V19, N2, pp. 306-317, 4/00.
- [2] Sha, L., Guo, H. Song, A. W., “An Improved Gridding Method for Sprial MRI Using Nonuniform Fast Foruier Transform,” *J. Magnetic Resonance*, V162, pp. 250-258, 3/03.
- [3] Vaidyanathan, P. P., *Multirate Systems and Filter Banks*, Prentice-Hall, 1993.
- [4] Ehrhardt, J. C., “Hexagonal Fast Foruier Transform with Rectangular Output,” *IEEE Trans. Sig. Proc.*, V41, N3, pp. 1469-1472, 3/93.
- [5] Vetterli, M., Kovacevic, J., *Wavelets and Subband Coding*, Prentice-Hall, 1995.