Image reconstruction by deterministic compressed sensing with chirp matrices[†]

Kangyu Ni[a*], Prasun Mahanti[b], Somantika Datta[b], Svetlana Roudenko[a], and Douglas Cochran[b]

[a] Department of Mathematics, Arizona State University, Tempe, AZ, USA
[b] Department of Electrical Engineering, Arizona State University, Tempe, AZ, USA

ABSTRACT

A recently proposed approach for compressed sensing, or compressive sampling, with deterministic measurement matrices made of chirps is applied to images that possess varying degrees of sparsity in their wavelet representations. The "fast reconstruction" algorithm enabled by this deterministic sampling scheme as developed by Applebaum et al. [1] produces accurate results, but its speed is hampered when the degree of sparsity is not sufficiently high. This paper proposes an efficient reconstruction algorithm that utilizes discrete chirp-Fourier transform (DCFT) and updated linear least squares solutions and is suitable for medical images, which have good sparsity properties. Several experiments show the proposed algorithm is effective in both reconstruction fidelity and speed.

Keywords: compressed sensing, chirp, discrete chirp-Fourier transform, image reconstruction

1. INTRODUCTION

1.1 The compressed sensing problem

Consider a one dimensional signal x in $\mathbb{R}^N$. An image or any other higher dimensional array can be vectorized into a one-dimensional vector. If the set $\{\psi_i\}_{i=1}^N$, $\psi_i \in \mathbb{R}^N$, is a basis for $\mathbb{R}^N$, then x can be expressed as

$$x = \sum_{i=1}^{N} s_i \psi_i \quad \text{or,} \quad x = \Psi s, \tag{1}$$

where Ψ is the matrix $[\psi_1|\psi_2|\dots\psi_N]$ and $s \in \mathbb{R}^N$ is the vector of coefficients. If only k coordinates of s are non-zero, then the signal x is said to be k-sparse in the system Ψ. The goal in *compressed sensing*, or *compressive sampling*, is to be able to reconstruct a k-sparse signal x from only a small number of n linear measurements, where $k < n \ll N$ and n is much less than the rate suggested by the Shannon-Nyquist Sampling Theorem. The measured value of the signal x is denoted by $y = [y_1, y_2, \dots, y_n]^T \in \mathbb{R}^n$ and is obtained by computing n inner products between x and a collection of vectors $\{\phi_j\}_{j=1}^n$ as $y_j = \langle x, \phi_j \rangle$. If we define Φ to be the $n \times N$ matrix whose rows are the ϕ_j's, then

$$y = \Phi x = \Phi\Psi s = \Theta s. \tag{2}$$

The matrix Φ is called the *sensing* or *measurement* matrix and Ψ is called the *sparsifying basis*. Since x and s are equivalent representations of the same signal, recovering s is the same as recovering x. However, the matrix Θ is $n \times N$, making the system in the equation (2) ill-posed. This can be tackled by using the approach of compressed sensing which, in particular, includes a construction of good sensing matrices Φ which would permit recovery of the signal from only n measurements [2, 3].

[†] This work is partially supported by NSF-DMS grant # 0652833.
[*] Corresponding author, e-mail: kangyu@math.asu.edu

1.2 Mathematical Background

If we knew beforehand the exact location of the k non-zero entries of s, then we could solve an $n \times k$ over-determined system. Let Θ_{sub} denote the corresponding $n \times k$ matrix. It has been shown (e.g., in [4]) that a necessary and sufficient condition for this $n \times k$ system to be well-conditioned is that for any length-k vector v, the following inequality holds for some $\epsilon > 0$:

$$(1 - \epsilon)\|v\| \leqslant \|\Theta_{\text{sub}} v\| \leqslant (1 + \epsilon)\|v\|. \tag{3}$$

This is the so called *Restricted Isometry Property* (RIP) [4]. It says that the lengths of the vectors are nearly preserved under Θ_{sub}. In practice one does not know the exact location of the k non-zero entries which means that for a given Θ the RIP condition has to be verified for the $\binom{N}{k}$ possibilities of Θ_{sub}, a combinatorially complex procedure. So far this issue has been overcome by taking the sensing matrix Φ to be a random matrix for a fixed sparsifying basis Ψ. The entries of the matrix Φ are generated by an i.i.d. random variable such as a Gaussian or Bernoulli random variable. Both cases can be shown to give matrices Θ that satisfy the RIP with very high probability when $\Psi = \mathcal{I}$, the basis of delta spikes. For a Gaussian matrix, one needs $n > ck \log(N/k)$ measurements which is much smaller than N [2, 4, 5] . The obvious advantage of using a random matrix with i.i.d. entries is that if the RIP is satisfied for a particular selection of k columns, then it should hold with high probability for any other selection of k columns. The reconstruction in this case is by means of solving the following convex optimization problem:

$$\min \|\tilde{s}\|_1 \quad \text{such that} \quad y = \Theta\tilde{s}. \tag{4}$$

The optimization problem of (4) can be solved by a method called basis pursuit which has a computational complexity of $\mathcal{O}(N^3)$ and alternatives to basis pursuit such as the greedy matching pursuit also have computational complexities that depend on N.

1.3 Deterministic Compressed Sensing Matrices

In recent work, Howard, Applebaum, Searle, and Calderbank have proposed the use of deterministic compressive sensing measurement matrices comprised of "chirps" (i.e., certain families of frequency modulated sinusoids) [1] or second-order Reed-Muller sequences [6]. These matrices come with a very fast reconstruction algorithm whose complexity depends only on the number of measurements n and not on the signal length N. These papers give empirical evidence that compressed sensing matrices formed from these deterministic vectors share certain properties with Gaussian random matrices that make them suitable for compressed sensing. In particular, selecting k columns randomly (i.e., independently from a uniform distribution on all N columns) from the $n \times N$ sensing matrix Φ, yields $k \times k$ Gram matrices whose condition numbers are distributed with mean and variance essentially identical to those obtained by using the same procedure on Gaussian matrices. This observation appears to hold over a range of $k < n \ll N$, compatible with those discussed in the compressive sensing literature for random matrices (e.g., [3]). Furthermore, deterministic algorithms for reconstruction of a k-sparse signal vector $x \in \mathbb{C}^N$ from compressive measurements $y = \Phi x \in \mathbb{C}^n$ with $\mathcal{O}(kn \log n)$ computational complexity versus, for example, $\mathcal{O}(knN)$ for ℓ_1 optimization with matching pursuit [1].

In the context of image processing, the roles in which compressed sensing may ultimately prove most useful are still being defined. A key issue is whether various classes of images possess sufficient sparsity to make the use of compressed sensing efficacious. The "one-pixel" camera architecture developed at Rice University [7] is one specialized application that seeks to exploit image sparsity for compressive sampling. Lustig et al. [8] comment that certain medical images are sparse in several transform domains and make use of this sparsity to enable reconstruction of magnetic resonance (MRI) images from a relatively small number of samples of their Fourier transforms (i.e., in the natural sampling domain for MRI). Candès and Romberg [9] give an example of compressed sensing of a natural image with complex noiselet measurements.

This paper is primarily concerned with whether the deterministic compressed sensing matrices made of chirps as proposed by Applebaum et al. [1] and their associated reconstruction algorithms offer potential in the imaging regime.

2. CHIRP SENSING MATRICES

Fast decoders have been identified for two related classes of deterministic compressed sensing matrices. The first one, described in [1], is comprised of a family of "chirps" (i.e., frequency modulated sinusoids). The second, introduced in [6], is formed from second-order Reed-Muller sequences [10]. The work presented in the remainder of this paper focuses on the chirp matrices, which we briefly review in the following paragraph.

A chirp signal of length n with chirp rate r and base frequency m has the form

$$v_{r,m}(\ell) = e^{\frac{2\pi i r\ell^2}{n} + \frac{2\pi i m\ell}{n}}, \qquad r, m, l \in \mathbb{Z}_n. \tag{5}$$

For a fixed n, there are n^2 possible pairs (r, m). The full chirp sensing matrix Φ has size $n \times n^2$ and its j^{th} columns is

$$\Phi_{r,m}(\ell) = v_{r,m}(\ell), \qquad j = nr + m \in \mathbb{Z}_{n^2}. \tag{6}$$

Defining $N = n^2$, a k-sparse signal $x \in \mathbb{C}^N$ yields a measurement $y = \Phi x \in \mathbb{C}^n$, that is the superposition of k chirp signals

$$y(\ell) = s_1 e^{\frac{2\pi i r_1\ell^2}{n} + \frac{2\pi i m_1\ell}{n}} + s_2 e^{\frac{2\pi i r_2\ell^2}{n} + \frac{2\pi i m_2\ell}{n}} + \cdots + s_k e^{\frac{2\pi i r_k\ell^2}{n} + \frac{2\pi i m_k\ell}{n}}. \tag{7}$$

To recover x, Applebaum et al. use fast Fourier transform (FFT) to detect one by one the nonzero locations, (r_j, m_j) pairs, whose total computational complexity is $O(kn \log n)$. The magnitudes s_j of the nonzero locations r_j are found by solving the associated least squares problem. In detail, their algorithm repeats the following steps:

1. For a $T \neq 0$, $w =$ location of the peak of $\mathrm{FFT}\Big\{y_0(\ell)y_0(\ell+T)\Big\}$. Find r_j by solving $w = 2r_jT \pmod n$.
2. $m_j =$ location of the peak of $\mathrm{FFT}\Big\{y_0(\ell)e^{\frac{2\pi i r_l\ell^2}{n}}\Big\}$.
3. Determine s_l by minimizing $\|y_0(\ell) - \sum_{l=1}^{j} s_l e^{\frac{2\pi i r_l\ell^2}{n} + \frac{2\pi i m_l\ell}{n}}\|^2$
4. $y_0(\ell) = y_0(\ell) - \sum_{l=1}^{j} s_l e^{\frac{2\pi i r_l\ell^2}{n} + \frac{2\pi i m_l\ell}{n}}$. Terminate when $\|y_0\|$ is sufficiently small.

Note that for step 1 to work properly, n has to be a prime number to uniquely determine r_j. For reconstructing signals that are sufficiently sparse, this algorithm is more efficient than $\ell 1$ minimization with random matrices, which computational complexity is $O(knN)$, in terms of reconstruction speed and reconstruction fidelity.

3. IMAGE RECONSTRUCTION

Despite the success for sparse one-dimensional signals, we observed for real images, this algorithm becomes impractical even for 128×128 images. This is because real images usually are not sufficiently sparse in any transform domain compared to sparse one-dimensional signals. A modified algorithm by Applebaum et al. uses all $T \neq 0 \in \mathbb{Z}_n$ instead of only one $T \neq 0$ in step 1 above. This is more robust to noise and also can be performed with better image reconstruction fidelity, compared to the original version above. However, the computational complexity of the modified step is $O(kn^2 \log n)$, and therefore, the speed is hampered. For example, a 128×128 image with 10% sparsity in some known domain has 1638 nonzero coefficients. With their reconstruction algorithm, it requires about 1638 iterations and the least squares problems become very large. Moreover, by the rule of thumb, at least approximately 4000 measurements are needed for correct reconstructions. The number of measurements needed is about 25% of the signal length, which implies that only four chirp rates (denoted by r_1, r_2, r_3, and r_4 from now on) are needed in the sensing matrix. Therefore, the efficiency of finding nonzero locations is not utilized in the imaging regime.

To this, we propose to use discrete chirp-Fourier transform (DCFT) [11] to find the nonzero locations of and an updated least squares solutions to find the magnitude of these nonzero locations. The former allows detection of several nonzero locations instead of only one nonzero location in each step. The later utilizes calculations of the previous least squares solutions and requires much less memory storage. Our algorithm repeats the following steps:

1. Find multiple (r_j, m_j) using DCFT.

2. Determine s_j using updated linear least squares solutions.

3. $y_0(\ell) = y_0(\ell) - \sum_{l=1}^{j} s_l e^{\frac{2\pi i r_l \ell^2}{n} + \frac{2\pi i m_l \ell}{n}}$. Terminate when $\|y_0\|$ is sufficiently small.

The third step subtracts residual y_0 by the linear (their respective magnitudes) sum of the found chirp columns to obtain the new residual y_0. We now explain step 1 and step 2 in Sections 3.1 and 3.2, respectively.

3.1 Discrete chirp-Fourier transform

The DCFT defined in [11] of an n-point signal x can be written as

$$X_c(r, m) = \frac{1}{\sqrt{n}} \sum_{\ell=0}^{n-1} x(\ell) W_n^{r\ell^2 + m\ell}, \tag{8}$$

where $W_n = e^{-2\pi i/n}$. In the above definition r and m are the chirp rate-frequency pairs. For a fixed r_j, we have

$$X_c(r_j, m) = \frac{1}{\sqrt{n}} \sum_{\ell=0}^{n-1} x(\ell) W_n^{r_j\ell^2 + m\ell}. \tag{9}$$

Now defining

$$x_r(\ell) = x(\ell) W_n^{r_j \ell^2}, \tag{10}$$

we then have

$$X_c(r_j, m) = \underset{n}{\mathrm{DFT}} \left\{ x_{r_j}(\ell) \right\}, \qquad j = 1, 2, 3, 4. \tag{11}$$

It may be noted that only four n-point DFT's need to be evaluated as only 4 chirp sensing matrices were used. Of the $4n$ DCFT coefficients computed, the ones with largest absolute values are chosen. The corresponding (r, m) pairs from the DCFT plane are used to recover the chirp sensing matrix columns. In our setting, applying DCFT to detect nonzero locations is to de-chirps the measurement residual with all four chirp rates and then followed by applying DFT

$$w(r_j, \ell) = \underset{n}{\mathrm{DFT}} \left\{ y_0(\ell) \overline{\Phi_{r_j,0}(\ell)} \right\}, \qquad j = 1, 2, 3, 4. \tag{12}$$

The pairs (r_j, ℓ) with the largest $|w(r_j, \ell)|$ are then detected.

3.2 Updated linear least squares solutions

The second step finds s_l by solving $\min \|y_0(\ell) - \sum_{l=1}^{j} s_l e^{\frac{2\pi i r_l \ell^2}{n} + \frac{2\pi i m_l \ell}{n}}\|^2$. Let $z = [s_1, s_2, ..., s_j]^T$ and A be formed by the chirp matrix columns with (r_l, m_l), with $l = 1, 2, ..., j$. Then, we can rewrite the problem in this form, $\min_z \|Az - y_0\|^2$. By the nature of the algorithm, the matrix in the current iteration is expanded from the matrix in the previous iteration by adding the newly found chirp columns. To solve these least squares problems without treating each problem (iteration) independently, we use an updated pseudo inverse solution method whose calculations are based on the previous calculations. The pseudo inverse solution of $\min_z ||Az - y_0||^2$ is

$$z_{sol} = (A^* A)^{-1} A^* y_0, \tag{13}$$

where * indicates conjugate transpose. Since the current matrix A is obtained by concatenating the newly found columns c with the previous matrix, $A = [\widetilde{A} \quad c]$, finding the inverse of

$$A^* A = \begin{bmatrix} \widetilde{A}^* \widetilde{A} & \widetilde{A}^* c \\ c^* \widetilde{A} & c^* c \end{bmatrix} \tag{14}$$

can be made efficiently by the Schur-Banachiewicz blockwise inversion formula (e.g. see [12]),

$$\begin{bmatrix} D & E \\ F & G \end{bmatrix}^{-1} = \begin{bmatrix} D^{-1} + D^{-1}E(G - D^{-1}E)^{-1}FD^{-1} & -D^{-1}E(G - D^{-1}E)^{-1} \\ -(G - D^{-1}E)^{-1}FD^{-1} & (G - D^{-1}E)^{-1}, \end{bmatrix}, \tag{15}$$

because D^{-1} is known from the previous step and the size of $(G - D^{-1}E)^{-1}$ is small. The calculation of A^*y_0 can also be updated by

$$A^*y_0 = \begin{bmatrix} \widetilde{A}^*y_0 \\ c^*y_0 \end{bmatrix}, \tag{16}$$

where the size of c is much smaller than the size of A.

3.3 Reconstructed error v.s. time

Applying DCFT in the algorithm gives the ability to detect multiple chirp-frequency pairs in one iteration. In our experiments, selecting the largest 100 magnitudes gives good reconstruction. Therefore, the overall computational complexity of finding nonzero locations with DCFT is $O(\frac{4}{100}kn \log n) = O(\frac{1}{25}kn \log n)$. This is much smaller than $O(kn^2 \log n)$, by Applebaum et al.'s method with all shifts T, which is applicable for images. However, note that since n has to be prime, the efficiency of FFT cannot be fully utilized for the n-point DFT.

We next show the performance of the above proposed algorithm that combines DCFT and updated least squares solutions. Figure 1 shows the same experiment with various numbers of chirp-frequency pairs selected per iteration, (2, 5, 10, 20, 30, 40, 50, 100). The left plot shows that the reconstructed error does not increase (in fact, remains the same) when the number of chirp-frequency pairs detected by DCFT in each iteration increases. The reconstruction error is defined as

$$\text{Error(dB)} = 10 \log_{10} \left[\frac{||x_{\text{actual}} - x_{\text{reconstructed}}||^2}{||x_{\text{actual}}||^2} \right]. \tag{17}$$

The right plot shows that when the number of chirp-frequency pairs detected by DCFT in each iteration increases, the reconstructed time decreases rapidly. Therefore, the proposed algorithm is able to have good reconstruction fidelity and efficiency at the same time.

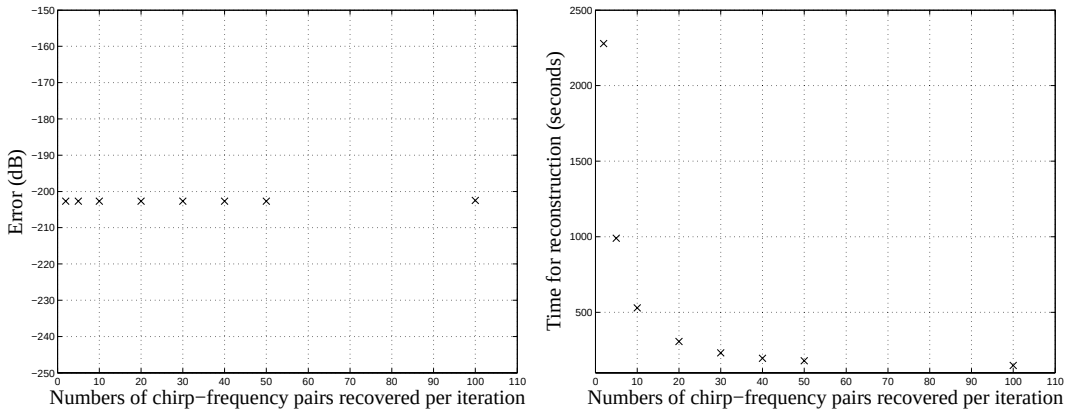

Figure 1. Left: the reconstructed error does not increase when the number of chirp rate-frequency pairs in each iteration increases. Right: Moreover, when the number of chirp rate-frequency pairs in each iteration increases, the reconstructed time decreases rapidly.

4. RESULTS

For the experiments, each original image was sparsified by computing its Haar wavelet transform and then retaining pre-determined fraction (e.g., 15% or 10%) of its wavelet coefficients, keeping the largest coefficients and setting the rest to zero. The image data were sampled (measured) with both chirp and noiselet matrices of the

same size and reconstructed images were obtained from the measurements. Figure 2 shows results reconstructed for a 256×256 pixel natural image. On the left is the image we wish to recover, whose sparsity is 15% in the Haar wavelet domain. The middle is the reconstructed image by our implementation of Candès and Romberg's method [9], $\ell 1$ minimization with randomly chosen noiselet measurements. The right is the reconstructed image by our algorithm with the chirp sensing matrix. Both chirp and noiselet took 16411 measurements ($\approx 25\%$) for reconstruction. The chirp reconstruction outperformed the noiselet one in terms of reconstruction error. The error for chirp was -41 dB, which is much better than noiselet (-24 dB).

Figure 2. Left: 256×256 pixel natural image with sparsity = 15% in the Haar wavelet domain. Middle: reconstructed image using about 25% data with noiselet. The error is -24 dB. Right: reconstructed image using chirp. The error is -41 dB, much better than that of noiselet.

In Figure 3, the same experiment is shown for a 256×256 angiography image. The sparsity of the image is 10% and 16411 measurements ($\approx 25\%$) were used for reconstruction. The chirp reconstruction in this case also outperformed the noiselet reconstruction. The error for chirp was -45 dB, whereas the error for noiselet was -17 dB.

Figure 4 shows experiment with the knee image with a 320×320 pixel resolution. The sparsity of the image is 10% and 25601 measurements ($\approx 25\%$) were used for reconstruction. The chirp reconstruction performed well and the error is small, -44 dB. Note that the discrete noiselet transform is defined only for signals of length that is a power of 2, and therefore, it is not possible to obtain a reconstruction by noiselet (unless introducing, for example, a zero padding, and thus, increasing the error level). This is a straight forward advantage of using the chirp compressed sensing.

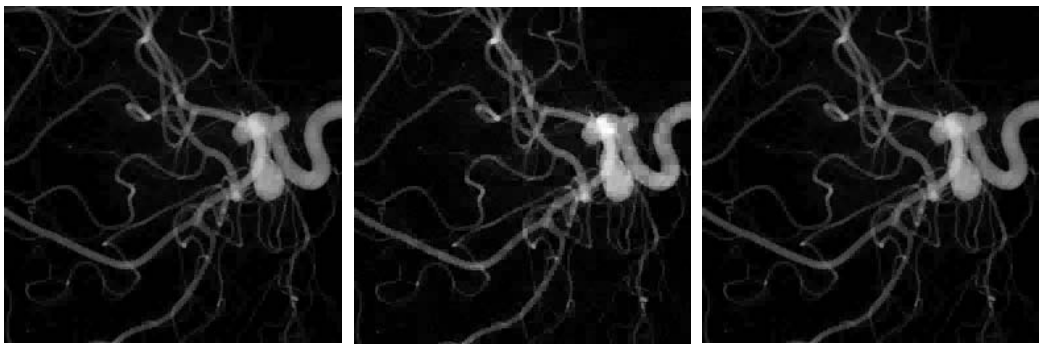

Figure 3. Left: 256×256 pixel MRI image with sparsity = 10% in the Haar wavelet domain. Middle: reconstructed image using about 25% data with noiselet. The error is -17 dB. Right: reconstructed image using chirp. The error is -45 dB, much better than that of noiselet.

Table 1 shows the reconstruction error for the examples in Figures 2–4. Although $\ell 2$ error is of limited value in assessing perceptual image quality, reconstruction error values in this instance are consistent with perceptual quality. The chirp reconstructions for images of the sizes and sparsities tested with the sample sized used are generally better in fidelity than reconstructions obtained using noiselet compressed sensing. In the above

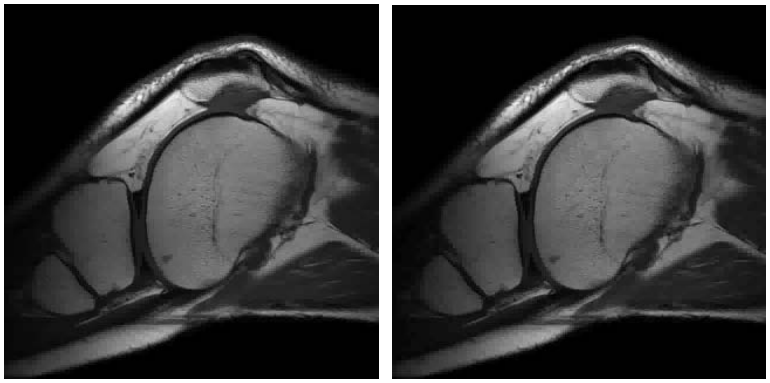

Figure 4. Left: 320×320 pixel MRI image with sparsity = 10% in the Haar wavelet domain. Right: reconstructed image using about 25% data with chirp. The reconstructed error = -44 dB is very small.

Fig	Size	Sparse	Method	Error
#2	256×256	15%	Noiselet	-24
			Chirp	-41
#3	256×256	10%	Noiselet	-17
			Chirp	-45
#4	320×320	10%	Noiselet	N/A
			Chirp	-51

Table 1. Reconstruction error for the examples in Figures 2–4. Error figures are in dB.

experiment, we see that the chirp sensing matrices are effective in compressed sensing when used in the image regime. In addition, as explained in Section 3.3, the computational complexity of finding nonzero locations with DCFT is $O(\frac{1}{25}kn \log n)$, much smaller than $O(kn^2 \log n)$, by Applebaum et al.'s method.

5. CONCLUSIONS

Deterministic compressed sensing matrices made of chirps are good sensing matrices for images with adequate sparsities. The reconstruction algorithms by Applebaum et al. are, however, not efficient for image reconstruction, even though great for reconstructing sparse one-dimensional signals. The proposed algorithm described in this paper provides much better image reconstruction in terms of both reconstruction errors and computational complexity. The former is compared to Candès and Romberg's compressed sensing method, random noiselet measurement with $\ell 1$ minimization for reconstruction, and the latter is compared to Applebaum et al.'s algorithm for chirp measurements.

6. ACKNOWLEDGMENT

The authors are grateful for the assistance and advice of Stephen Howard in connection with the work presented in this paper.

REFERENCES

[1] Applebaum, L., Howard, S., Searle, S., and Calderbank, R., "Chirp sensing codes: Deterministic compressed sensing measurements for fast recovery," *Applied and Computational Harmonic Analysis* **26**, 283 – 290 (2009).

[2] Baraniuk, R., "Compressive sensing," *IEEE Signal Processing Magazine* **24**, 118 – 121 (July 2007).

[3] Candès, E., "Compressive sampling," in [*International Congress of Mathematicians*], **III**, 1433–1452, European Mathematical Society, Zürich (2006).

[4] Candès, E., Romberg, J., and Tao, T., "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory* **52**, 489 – 509 (February 2006).

[5] Candès, E. and Tao, T., "Near optimal signal recovery from random projections: Universal encoding strategies?," *IEEE Transactions on Information Theory* **52**, 5406 – 5425 (December 2006).

[6] Howard, S., Calderbank, A., and Searle, S., "A fast reconstruction algorithm for deterministic compressive sensing using second order Reed-Muller codes," in [*Proceedings of the Conference on Information Sciences and Systems*], 11 – 15 (March 2008).

[7] Takhar, D., Laska, J., Wakin, M., Duarte, M., Baron, D., Sarvotham, S., Kelly, K., and Baraniuk, R., "A new compressive imaging camera architecture using optical-domain compression," in [*Computational Imaging IV at SPIE Electronic Imaging*], 43 – 52 (January 2006).

[8] Lustig, M., Donoho, D., and Pauly, J. M., "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magnetic Resonance in Medicine* **58**, 1182 – 1195 (December 2007).

[9] Candès, E. and Romberg, J., "Sparsity and incoherence in compressive sampling," *Inverse Problems* **23**(3), 969 – 985 (2007).

[10] MacWilliams, F. J. and Sloane, N. J. A., [*The theory of error correcting codes*], Elsevier (1976).

[11] Xia, X.-G., "Discrete chirp-Fourier transform and its application to chirp rate estimation," *IEEE Transactions on Signal Processing* **48**, 3122 – 3133 (November 2000).

[12] Björck, A., [*Numerical methods for least squares problems*], SIAM (1996).

[13] Candès et al., "ℓ_1-magic home page." http://www.acm.caltech.edu/l1magic/.

[14] DeVore, R., "Deterministic constructions of compressed sensing matrices," *Journal of Complexity* **23**, 918 – 925 (August 2007).

[15] Donoho, D., "Compressed sensing," *IEEE Transactions on Information Theory* **52**, 1289 – 1306 (April 2006).