Chirp Sensing Codes: Deterministic Compressed Sensing Measurements for Fast Recovery

Lorne Applebaum
Electrical Engineering
Princeton University
NJ 08544, USA

Stephen Howard
DSTO
PO Box 1500
Edinburgh 5111, Australia

Stephen Searle
Electronic & Electrical Engineering
The University of Melbourne
Victoria 3010, Australia

Robert Calderbank
Electrical Engineering
Princeton University
NJ 08544, USA

***Abstract*—Compressed sensing is a novel technique to acquire sparse signals with few measurements. Normally, compressed sensing uses random projections as measurements. Here we design deterministic measurements and an algorithm to accomplish signal recovery with computational efficiently. A measurement matrix is designed with chirp sequences forming the columns. Chirps are used since an efficient method using FFTs can recover the parameters of a small superposition. We show empirically that this type of matrix is valid as compressed sensing measurements. This is done by a comparison with random projections and a modified reduced isometry property. Further, by implementing our algorithm, simulations show successful recovery of signals with sparsity levels similar to those possible by Matching Pursuit with random measurements. For sufficiently sparse signals, our algorithm recovers the signal with computational complexity $O(K \log K)$ for K measurements. This is a significant improvement over existing algorithms.**

I. Introduction

The sparsity of signals is a fact often exploited in signal processing. In particular, the common way to compress a signal is to transform it to the basis in which it is sparse and subsequently store only the locations and values of the few non-zero elements. Recently, it has been discovered that, in addition to storage, signal sparsity can be leveraged to reduce the number of measurements for signal acquisition and detection. It has been shown that, if a signal is sufficiently sparse, a small number projections onto random vectors is enough to recover the signal [1], [2]. This method has been called *Compressed Sensing*.

In compressed sensing, the use of randomly generated projections to make measurements has the useful consequence of sidestepping the computationally difficult task of checking whether the measurements allow for signal recovery. By considering recovery stocastically, it has been shown that measurements generated from Gaussian or Bernoulli random variables allow for signal recovery with high probability. In some ways, the use of random measurements may be viewed as an analogy to random codes used by Shannon to prove theorems in channel coding. Though useful in proofs, purely random channel codes are never used in practice because encoding and decoding would be far too computationally intensive. Instead, practical channel codes are developed with an efficient coding and decoding scheme in mind. We have a similar situation in compressed sensing. Though ℓ_1 minimization has been shown to recover the signal from random projections [1], it is computationally expensive. The question arises as to whether we can design projections to facilitate the rapid recovery of the signal. This is an issue of practical consequence. If compressed sensing is to be used in real-time systems, we must have a method which, in addition to reducing the number of measurements, is able to recover the signal quickly. Here we present a proof of concept scheme which accomplishes this.

A number of decoding schemes have been proposed that improve upon the ℓ_1 minimization signal recovery technique (also known as *Basis Pursuit*). However, most schemes presuppose random measurements. Examples include Orthogonal Matching Pursuit [3] and its refinements [4]. In contrast, the scheme presented here exploits structure in deterministically designed measurements to make recovery much faster. There exists a small number of other schemes with less structurally random measurements [5]–[7]. The scheme presented here has lower recovery complexity.

The remainder of the paper is organized as follows. In Section II we provide necessary background and notation and in Section III we introduce our encoding scheme and the corresponding decoding algorithm. Section IV provides analysis of our encoding matrix in terms of the reduced isometry property commonly employed in compressed sensing. In Section V we consider our scheme in the special case of Fourier signals and present a modification to improve the scheme's robustness. In Sections VI and VII we examine our algorithm in terms of computational complexity, signal recovery and robustness to noise.

II. Compressed Sensing Background and Notation

We consider discrete signals of finite length. Let x be a length N signal which we would like to sense and recover. We assume that x is sparse in some orthonormal basis. Thus, we can write x as

$$x = \Psi s \quad (1)$$

where s is a length N vector with fewer than M non-zero elements. We measure x with $K < N$ projections which have results given in the vector y. The vectors projected upon are set as the rows of the $K \times N$ matrix Φ which gives

$$\begin{aligned} y &= \Phi\Psi s \\ &= \Theta s \end{aligned} \quad (2)$$

where the second equality is by definition of Θ. We are free to design Φ and thus Θ. Though, if we design Θ we should remain aware that actual sensing of the signal is done with Φ.

Since Θ is a wide matrix, solving for s given y is ill posed. However, using non-linear methods, we can leverage the fact that s has at most M non-zero elements. It has been shown in [8] that if Θ satisfies the *M-restricted isometry property* (M-RIP), s can be recovered perfectly using an ℓ_1 minimization. Several results exist showing that, when M satisfies

$$M < cK/\log(N/K), \quad (3)$$

with a known constant c, randomly generated matrices of various types satisfy M-RIP with high probability [9]. Thus, if a signal's sparsity is bounded by (3), then it can be recovered from K random measurements with high probability.

We summarize M-RIP as follows. Let $\mathcal{M}$ be a subset of $\{1, \ldots, N\}$ of cardinality M and $\Theta_{\mathcal{M}}$ be the matrix formed by concatenating the columns of Θ indexed by $\mathcal{M}$. Θ is M-RIP when the eigenvalues of $\Theta_{\mathcal{M}}^T\Theta_{\mathcal{M}}$ are bounded near 1 for every $\mathcal{M}$. The bound is dependent on M and its exact definition can be found in [8]. In this paper, the statistics of eigenvalues of the Gram matrices $\Theta_{\mathcal{M}}^T\Theta_{\mathcal{M}}$ are of concern rather than a formulation of their bound.

III. Chirped Sensing Codes

We approach the recovery problem by noting that finding s is equivalent to discovering which small linear combinations of the columns of Θ form y. We will design Θ to facilitate this. In particular, we will look at a Θ designed with chirp signals forming the columns.

A length K chirp signal has the form

$$v_{m,r}(l) = \alpha \cdot e^{\frac{j2\pi ml}{K} + \frac{j2\pi rl^2}{K}} \quad m, r \in \mathbb{Z}_K \quad (4)$$

where m is the base frequency and r is the chirp rate. For a length K signal, there are K^2 possible pairs (m, r). We will form a $K \times K^2$ sized Θ which has columns filled with all K^2 uni-modular chirp signals (setting $\alpha = 1$ for notational convenience, though in Section IV $\alpha = \frac{1}{\sqrt{K}}$ is used).

Consider a vector y, indexed by l, formed from the linear combination of some chirp signals

$$y(l) = s_1 e^{\frac{j2\pi m_1 l}{K} + \frac{j2\pi r_1 l^2}{K}} + s_2 e^{\frac{j2\pi m_2 l}{K} + \frac{j2\pi r_2 l^2}{K}} + \cdots \quad (5)$$

which have base frequencies defined by m_i and chirp rates defined by r_i. The chirp rates can be recovered from y by looking at $\bar{y}(l)y(l+d)$, where the index $l+d$ is taken $\mod K$. This gives

$$\begin{aligned} f(l) = \bar{y}(l)y(l+T) &= |s_1|^2 e^{\frac{j2\pi}{K}(m_1T + r_1T^2)} e^{\frac{j2\pi(2r_1lT)}{K}} \\ &+ |s_2|^2 e^{\frac{j2\pi}{K}(m_2T + r_2T^2)} e^{\frac{j2\pi(2r_2lT)}{K}} \\ &+ \cdots + \text{cross terms} \end{aligned} \quad (6)$$

where the cross terms are of the form

$$s_p\bar{s}_q e^{\frac{j2\pi}{K}(m_pT + r_pT^2)} e^{\frac{j2\pi}{K}l(m_p - m_q + 2Tr_p) + \frac{j2\pi}{K}l^2(r_p - r_q)} \quad (7)$$

and are therefore chirps. We see that $f(l)$ is a signal that has sinusoids at the discrete frequencies $2r_iT \mod K$. If K is prime, this is a bijection from chirp rates to FFT bins. Furthermore, the remainder of the signal consists of the cross terms. Since the cross terms are chirps, their energy is spread across all FFT bins.

As long as y consists of sufficiently few chirps (x is sparse), taking a FFT of $f(l)$ results in a spectrum with significant peaks at locations corresponding to $2r_iT \mod K$ from which we can glean chirp rates.

Upon discovering the chirp rate r_i we can "dechirp" the signal $y(l)$ by multiplying by $e^{-j2\pi r_i l^2/K}$. This converts only the chirps with rate r_i to sinusoids. Performing an FFT on the resulting signal can be used to retrieve the corresponding value(s) for m_i and s_i.

Setting the elements of Θ as

$$[\Theta]_{l,k} = e^{\frac{j2\pi rl^2}{K}} e^{\frac{j2\pi ml}{K}} \quad \text{with} \quad k = Kr + m \in \mathbb{Z}_{K^2} \quad (8)$$

we see that $y = \Theta s$ will have the form (5). Given y formed using Θ, we summarize the algorithm described.

1) Choose a $T \in \mathbb{Z}_K, T \neq 0$, and a stopping energy ϵ.
2) Form $f(l) = \bar{y}(l)y(l+T)$ and take length K FFT.
3) Find location of the peak in the FFT as $2r_iT \mod K$ and record the unique r_i corresponding to the location.
4) Multiply $y(l)$ by $e^{-j2\pi r_i l^2/K}$ and take length K FFT.
5) Find the location of the peak and record as m_i use the value to recover s_i.
6) Replace y with $y - s_i e^{\frac{j2\pi(m_i l + r_i l^2)}{K}}$
7) Repeat steps 2-6 until $\|y\|_2^2 < \epsilon$ or have iterated M times.

The recovery of s_i gives the value of an element in s while the pair (m_i, r_i) gives its locations in s as the index $Kr_i + m_i$.

IV. RIP Analysis of Θ

The standard compressed sensing formulation described in Section II considers the signal s as fixed and examines recovery using of random measurements. The probability of recovery is inherited from the randomness in the measurements. However, here we have developed deterministic measurements Θ. In general, it is not easy to show a non-random Θ satisfies RIP. Thus, we have no guarantees of recovery for all M-sparse s. However, we conjecture that, considering the signal rather than the measurements stocastically, a randomly generated sparse s can be recovered from Θ with high probability. Empirical results presented here support this claim.

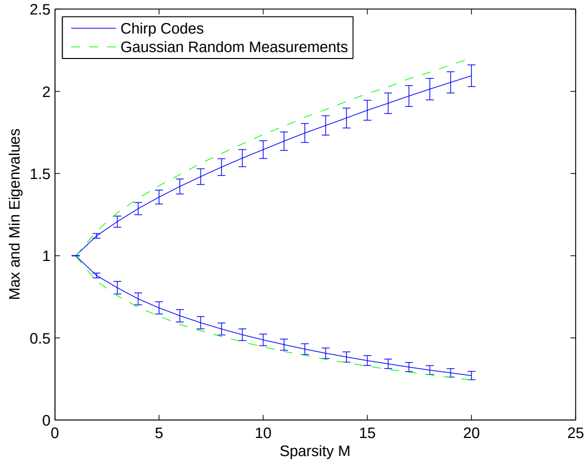

Fig. 1. Eigenvalue statistics of Gram matrices $\Theta_{\mathcal{M}}^T\Theta_{\mathcal{M}}$ of varying sparsity M for Chirp Codes and Gaussian Random Measurements

Recall that Θ satisfies M-RIP when $\Theta_{\mathcal{M}}^T\Theta_{\mathcal{M}}$ has eigenvalues sufficiently close to 1 for every $\mathcal{M}$. The sets $\mathcal{M}$ correspond to the possible supports of an M-sparse s. For a deterministic Θ, this, in general, requires the computationally difficult problem of checking all $\binom{N}{M}$ possible $\mathcal{M}$. Instead, we consider the probability that a random, though sparse, s has a support $\mathcal{M}$ with the property that $\Theta_{\mathcal{M}}^T\Theta_{\mathcal{M}}$ has eigenvalues appropriately close to 1.

After scaling Θ so that its columns have unit norm, we compare the statistics of eigenvalues of its Gram matrices to those of a matrix with Gaussian entries of zero mean and variance $\frac{1}{K}$. From well-known compressed sensing results, this Gaussian matrix is known to satisfy RIP with high probability when (3) is satisfied [8].

Figure 1 shows the sample means and standard deviations of the maximum and minimum eigenvalues of $\Theta_{\mathcal{M}}^T\Theta_{\mathcal{M}}$ for varying M. For every value M, sets $\mathcal{M}$ are generated uniformly random over all sets and the statistics are accumulated from 10,000 samples. A value of $K = 67$ was used for the simulation. For comparison, the sample means of the maximum and minimum eigenvalues of the Gram matrices of the Gaussian measurements are also shown.

From Figure 1 we see that the eigenvalues of $\Theta_{\mathcal{M}}^T\Theta_{\mathcal{M}}$ are, on average, closer to 1 by more than a standard deviation compared to the corresponding eigenvalues of Gaussian measurements. Thus, if Gaussian measurements satisfy M-RIP, then our Θ will also be able to recover a random M-sparse signal with high probability. Results are similar for other values of K.

It is important to note that here we have analyzed the measurement matrix Θ in isolation of the algorithm presented in Section III. In this section, we have shown that Θ is suitable as compressed sensing measurements in general. In Section VII, we examine how the measurements perform jointly with our corresponding algorithm.

V. Specializations

Here we consider the use of Chirp codes in the special case of sparse Fourier signals as well as a modification to the algorithm to mitigate cross-term interference and noise.

A. Φ for Sparse Fourier Signals

Though we are interested in being able to determine the combination of columns of Θ, measurements are taken upon x and thus are made with the matrix $\Phi = \Theta\Psi^{-1}$. We are therefore concerned with Φ for implementation. When Ψ is the Fourier matrix (if x is a sparse superposition of sinusoids), we can find the structure of Φ directly.

As described above, we set Θ to have the l,k entry of the form

$$[\Theta]_{l,k} = e^{\frac{j2\pi rl^2}{K}} e^{\frac{j2\pi ml}{K}} \quad \text{with} \quad k = Kr + m \in \mathbb{Z}_{K^2}$$

This construction of Θ groups the columns of chirps in blocks of chirp rates.

Since Ψ is the Fourier matrix, $\Phi = \Theta\Psi^{-1}$ is a matrix with rows formed by the K^2 length Fourier transform of the rows of Θ. The Fourier transform of the l^{th} row of Θ is given by

$$\begin{aligned}\sum_{k=0}^{K^2-1}[\Theta]_{l,k}e^{\frac{-j2\pi k\omega}{K^2}} &= \sum_{r=0}^{K-1}\sum_{m=0}^{K-1} e^{\frac{j2\pi rl^2}{K}} e^{\frac{j2\pi ml}{K}} e^{\frac{-j2\pi Kr\omega}{K^2}} e^{\frac{-j2\pi m\omega}{K^2}} \\ &= \left[\sum_{r=0}^{K-1} e^{\frac{j2\pi r}{K}(l^2-\omega)}\right]\left[\sum_{m=0}^{K-1} e^{\frac{j2\pi m}{K^2}(Kl-\omega)}\right] \\ &= \delta_{(l^2-\omega)} e^{j\pi(Kl-\omega)(1/K-1/K^2)} \\ &\quad \times \frac{\sin(\pi(Kl-\omega)/K)}{\sin(\pi(Kl-\omega)/K^2)}\end{aligned} \tag{9}$$

where δ_i is the Kronecker delta with i taken $\bmod K$. Thus, the rows of Φ are periodic trains of delta functions modulated by a sinc function. This means that the measurements y can be formed simply as a weighted sum of a sparse number of samples of x. Thus, the encoding of y has a relatively low computational cost.

This formulation of Φ along with the scenario of sparse signals in the Fourier domain is used in the simulations illustrated later.

B. Interference and Noise Mitigation

At the expense of more computation, we can improve the performance of the algorithm by exploiting the availability of $f_T(l)$ for different delays T. By adding the FFT bins of each chirp rate r with those of the FFTs formed from the other delays, we can mitigate the effect of noise and any significant values from cross terms in (6). A bin with a chirp at r is correlated across different T while with a white noise approximation of the cross term interference (or simply white noise), other bins are not correlated and have zero mean. In the extreme case, we take a FFT for all $K-1$ possible shifts T.

The modified algorithm is summarized here.

1) Choose a stopping energy ϵ.
2) Form $f_T(l) = \bar{y}(l)y(l+T)$ for every $T \in \mathbb{Z}_K, T \neq 0$ and take length K FFT of each.
3) Using $2r_iT \mod K$, reorganize the output of each FFT such that the bins are in order of increasing r_i.
4) Sum the absolute value of the reorganized FFTs and record the peak r_i
5) Multiply $y(l)$ by $e^{-j2\pi r_i l^2/K}$ and take length K FFT.
6) Find the location of the peak and record as m_i use the value to recover s_i.
7) Replace y with $y - s_i e^{\frac{j2\pi(m_i l + r_i l^2)}{K}}$
8) Repeat steps 2-7 until $\|y\|_2^2 < \epsilon$ or have iterated M times.

We compare the performance of original algorithm with this modified algorithm in Sections VII-A and VII-B.

VI. Computational Complexity

Reconstruction in compressed sensing is normally done by solving a linear program minimizing $\|s\|_1$. As remarked earlier, this method is computationally intensive and has complexity $O(N^3)$. An alternative scheme is the greedy Matching Pursuit algorithm with complexity $O(KMN)$, which we use in simulations for comparison. Here, we consider the complexity of the Chirp Code algorithm.

By using the chirp decoding algorithm we leverage the efficiency of the FFT. Similar to Matching Pursuit, the algorithm iteratively pulls out the strongest signals. In this algorithm, each "peel" requires the computation of two FFTs of length K: a first, to extract the coded chirp rate, and a second to extract the coded frequency. Since, for a M-sparse signal approximately M peels are required, the complexity of the computation is

$$O(MK \log K) \tag{10}$$

As noted in Section V-B we can trade additional computation for improved performance by using multiple delays T to form $f(l)$. In the most computationally intensive case, we perform K length K FFTs for each peel which gives an overall complexity of

$$O(MK^2 \log K) \tag{11}$$

In terms of computation, the algorithm is a significant improvement upon ℓ_1 minimization and Matching Pursuit. For comparison, a table of various algorithms and their complexities can be found in [6].

VII. Performance

Here we present some simulation results characterizing the performance of the algorithm. The simulations were produced using a measurement matrix formed as described in Section V-A acting upon a signal sparse in the Fourier domain. We compare its performance against the matching pursuit algorithm using a Gaussian random Φ.

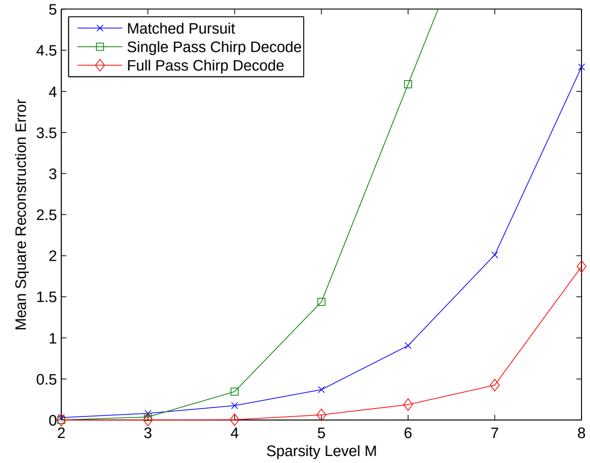

Fig. 2. Sparsity requirements of the algorithm for $K = 41$

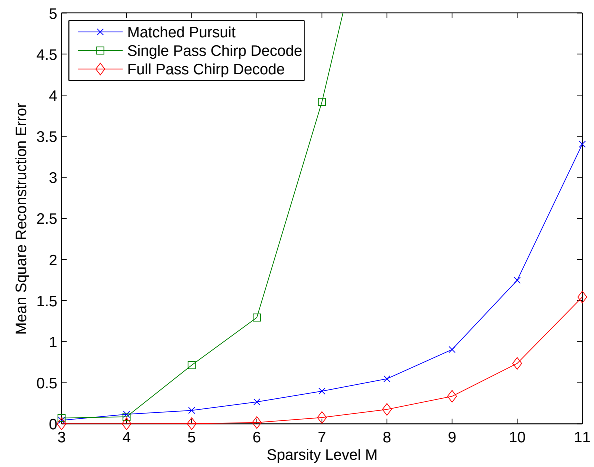

Fig. 3. Sparsity requirements of the algorithm for $K = 67$

A. Sparsity Requirements

An important examination is whether the algorithm's improved computation complexity degrades the sparsity level at which signals can be recovered. We look at the sparsity requirements for various signal lengths while using $K = \sqrt{N}$ measurements.

Figures 2 and 3 compare the reconstruction error of the chirp sensing code algorithm with that of matching pursuit for signal lengths $N = 41^2$ and $N = 67^2$. A signal comprised of a small number of sinusoids was measured and reconstructed. We include two chirp sensing code algorithms: using a single shift as well as using all possible shifts. We see that when all shifts are utilized, the chirp sensing algorithm is able to outperform matching pursuit, successfully reconstructing signals containing more sinusoids.

Results for other values of N show that the sparsity levels M required by the algorithm for signal recovery follow (3)

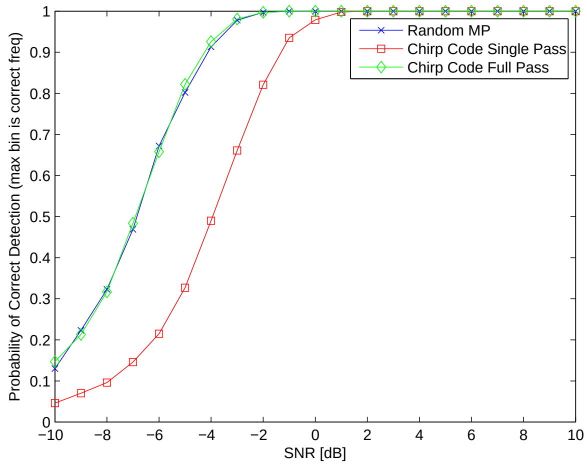

Fig. 4. Performance of algorithms in the presence of noise

with different values of c for single delays and multiple delays.

B. Detection in Noise

Strictly speaking, s is sparse if it has very few non-zero elements. However, this is not a good model of practical signals. Practical signals will have small values in all elements of s either due to noise or components that can be discarded. A practical recovery algorithm must be able to work under these circumstances. Further, it is important to know at which noise levels the algorithm can operate.

Figure 4 compares the performance of the chirped sensing code algorithm with matching pursuit. In the figure, we examine the detection of a single sinusoid in noise. We consider a correct detection if the first "peel" of the signal corresponds to the sinusoid. The figure was generated by simulation using $K = 41$ measurements and length $N = 41^2$ signals. Probabilities were estimated using 1000 samples. Both the single shift $O(MK \log K)$ algorithm and the $O(MK^2 \log K)$ algorithm using all shifts are included in the comparison.

We see that the algorithm that uses all the shifts achieves the performance of matching pursuit. These results can be compared to those in [10].

VIII. Conclusions and Extensions

The chirped sensing codes we introduced here are an illustration of how, by particularly selecting measurements in Φ, we can utilize a computationally efficient reconstruction scheme. The choice of the measurements in Φ were made such that from Θ, its form in the sparse basis, we can recover small linear combinations of columns. In particular, we designed a Θ filled with columns of chirps since we have an efficient method to recover chirp rates and frequencies from a small superposition. Further, with this design of Θ, Φ has a convenient form in the case of sparse Fourier signals.

Unlike most compressed sensing literature, we used deterministic measurements. As a result, verifying that the designed Θ satisfies the restricted isometry property and therefore can guarantee signal recovery is difficult. Instead, we considered a modified version of the RIP which regards the signal, rather than the measurements, stocastically. Empirical evidence showed that the majority of Gram matrices of Θ have eigenvalues closer to 1 than correspondingly sized Gram matrices of random Gaussian measurements. This, in turn showed that a signal recoverable from the random measurements is very likely recoverable from measurements made with Θ.

Signal recovery from our measurements was also shown by the implementation of our decoding algorithm. The recovery exploited the efficiency of the FFT in each of two steps: the first to recover the chirp rates and second to recover the chirp frequency. By identifying the chirp rates and frequencies, the superimposed columns of Θ are determined. In simulation the algorithm was shown to equal the performance of matching pursuit in noise resilience and exceed matching pursuit's performance in signal sparsity requirements.

A limitation of the chirp code algorithm is the restriction $K \geq \sqrt{N}$. This derives from the size of the family of length K chirps which necessitates Θ be $K \times K^2$ or narrower. As a result, this limits the algorithm's abilities in situations where K must be small. Stemming from this work, a similar algorithm based on second order Reed-Muller codes found in [11] addresses this. Second order Reed-Muller codes can be viewed analogously to chirps and decoding can be done using the Fast Hadamard Transform in place of the FFTs. The class of length K second order Reed-Muller codes is very large, essentially removing this lower bound on the number of measurements. This paper can serve as a preliminary read to [11] for those less familiar with Reed-Muller codes.

Regardless of the above limitation, chirp sensing codes excel as method for fast signal recovery with compressed sensing.

References

[1] E. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *Information Theory, IEEE Transactions on*, vol. 52, no. 2, pp. 489–509, Feb. 2006.

[2] D. Donoho, "Compressed sensing," *Information Theory, IEEE Transactions on*, vol. 52, no. 4, pp. 1289–1306, April 2006.

[3] J. A. Tropp and A. C. Gilbert, "Signal recovery from partial information via orthogonal matching pursuit," *Preprint*, Apr. 2005.

[4] D. L. Donoho, Y. Tsaig, I. Drori, and J.-L. Starck, "Sparse solution of underdetermined linear equations by stagewise orthogonal matching pursuit," *Preprint*, March 2006.

[5] S. Sarvotham, D. Baron, and R. G. Baraniuk, "Sudocodes – fast measurement and reconstruction of sparse signals," in *Proc. IEEE International Symposium on Information Theory (ISIT2006)*, Seattle, WA, Jul. 2006.

[6] S. Sarvotham, D. Baron, and R. Baraniuk, "Compressed sensing reconstruction via belief propagation," Electrical and Computer Engineering Department, Rice University, Tech. Rep., Jul. 2006.

[7] G. Cormode and S. Muthukrishnan, "Towards an algorithmic theory of compressed sensing," DIMACS, Tech. Rep. 25, Sep. 2005.

[8] E. Candes and T. Tao, "Decoding by linear programming," *Information Theory, IEEE Transactions on*, vol. 51, no. 12, pp. 4203–4215, Dec. 2005.

[9] E. J. Candès, "Compressive sampling," in *Proc. International Congress of Mathematiciains*, vol. 3, Madrid, Spain, Aug. 2006.

[10] M. Duarte, M. Davenport, M. Wakin and R. Baraniuk, "Sparse signal detection from incoherent projections," in *Proc. IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, Toulouse, France, May 2006.

[11] S. Howard, R. Calderbank, and S. Searle, “A fast reconstruction algorithm for deterministic compressive sensing using second order reed-muller codes,” in *Proc. Conference on Information Sciences and Systems (CISS2008)*, Princeton, NJ, Mar. 2008.