SENSITIVITY OF THE RANDOM DEMODULATION FRAMEWORK TO FILTER TOLERANCES

Pawel J. Pankiewicz, Thomas Arildsen, Torben Larsen

Aalborg University
Faculty of Engineering and Science
Department of Electronic Systems
DK-9220 Aalborg, Denmark
email: {pjp, tha, tl}@es.aau.dk
www.sparsesampling.com

ABSTRACT

The aim of the present paper is to demonstrate the impact of low-pass filter non-idealities on compressed sensing signal reconstruction in the random demodulator (RD) architecture. The random demodulator is a compressed sensing (CS) acquisition scheme capable of acquiring signals in continuous time. One of the main advantages of the system is the possibility to use off-the-shelf components to implement this sub-Nyquist framework. Low-pass filtering plays an important role in the RD analog acquisition process, which needs to be modeled carefully in the digital part of the compressive sensing reconstruction. Having a complete model of the analog front-end, CS algorithms conduct almost perfect reconstruction taking far less samples than for traditional Nyquist-rate sampling. This paper investigates reconstruction sensitivity to distortion in the impulse response of the low-pass filter caused by passive component value fluctuations. The authors simulate common CS recovery algorithms and show that the worst-case performance degradation due to filter component tolerances can be substantial, which requires special attention when designing reconstruction algorithms for RD.

1. INTRODUCTION

For more than half a century, traditional signal acquisition schemes have relied on Shannon-Nyquist sampling theory, which dictates sampling at a rate that is higher than twice the highest frequency of the sampled signal. In many cases, this constraint makes analog-to-digital converters (ADCs) power hungry devices. The recently emerged theory of CS [1, 2] states that this lower bound on the sampling frequency can be significantly reduced when the processed signal is sparse or compressible within a certain basis. CS theory exploits a priori knowledge of the signal structure, where signal sparsity is modeled by expressing the signal as the linear combination of a few elements taken in a particular dictionary (orthogonal or redundant) [3].

Compressed sensing is still mainly a mathematical concept, and only few successful attempts to sample continuous-time signals and hardware implementations have been reported [4, 5, 6]. One of the sampling architectures that exploits the signal sparsity is called random demodulator, the block diagram of which is shown on Fig 1. RD modulates the signal by multiplying it with a high-rate pseudo-random alternating sign sequence, "smearing" the signal frequency

The work of Pawel J. Pankiewicz and Thomas Arildsen is financed by The Danish Council for Strategic Research under grant number 09-067056.

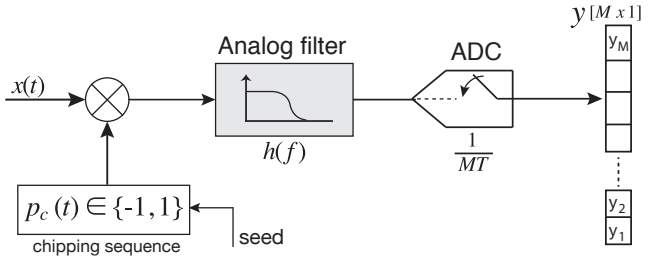

Figure 1: Random demodulator structure [4]

content across the entire spectrum. A low-pass filter assures no aliasing is introduced by the following low-rate ADC capturing compressed samples. The main advantage of this design is that it is easy to implement using robust, low-power and commonly available off-the-shelf components.

Thanks to its design simplicity, the RD hardware CS system is less exposed to the effects of non-ideal components compared to other CS hardware realizations [7]. Considering the fact that the random demodulator operates at relatively lower sampling rates, component non-idealities such as clock jitter of an ADC or mixer distortion can be neglected. One important aspect of the RD system is the measurement matrix defined within a reconstruction algorithm. It represents the analog sampling process where the pseudo-random sequence and impulse response of the filter need to be modeled precisely.

In reality, no physical device can be modeled perfectly. Due to filter component tolerances, the impulse response of the hardware low-pass filter will generally differ from the one modeled in the reconstruction algorithm. In this paper, we study the sensitivity of compressed sensing to this mismatch between the ideal and the actual analog front-end represented by the measurement matrix. For that purpose we demonstrate a MATLAB simulation framework and discuss results of our analysis.

2. METHODOLOGY

2.1 The Random Demodulator

When dealing with compressed sensing as it was initially defined [1, 3], we can distinguish two main operations that can be carried out separately. Namely, the data acquisition process and the subsequent recovery of the signal of interest by means of a reconstruction algorithm. In order to apply com-

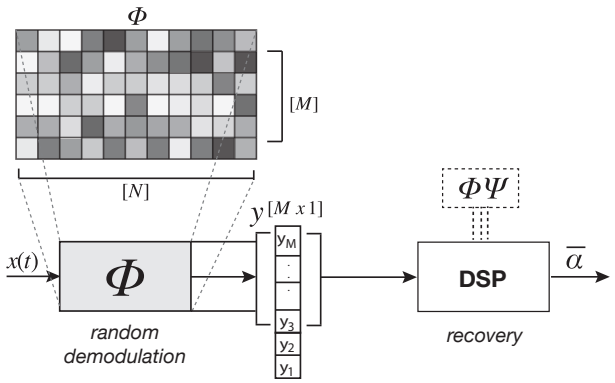

Figure 2: Process of compressed sensing acquisition.

pressed sensing to continuous signals, one needs an analog sampling front-end that will provide a signal representation that is compatible with the CS framework. The random demodulator is one of the analog acquisition schemes, which has been successfully employed in the compressed sampling field [4, 8, 7]. This data acquisition process provides non-adaptive linear projections of the analog input signal that are structured to directly conform to the CS approach:

$$\mathbf{y} = \mathbf{\Phi}\mathbf{\Psi}\alpha, \tag{1}$$

where: $\mathbf{y} \in \mathbb{R}^{M\times 1}$ represents compressed measurements, $\mathbf{\Phi} \in \mathbb{R}^{M\times N}$, $(M \ll N)$ describes system measurement matrix, $\mathbf{\Psi} \in \mathbb{C}^{N\times N}$ is signal sparsifying basis and $\alpha \in \mathbb{C}^{N\times 1}$ is a sparse vector containing coefficients, most of which are zero. Fig. 2 represents the general model of the CS system with RD framework included.

2.2 System design

The analog input signal model in our analysis is assumed to be a multi-tone signal that is band-limited and sparse in the frequency domain. Assuming finite energy in the signal it can be modeled as:

$$x(t) = \sum_{n=1}^{N} \alpha_n \psi_n(t), \quad t \in [0,T), \tag{2}$$

where α_n denotes Fourier series coefficients, T represents processed interval and

$$\psi_n(t) = \exp\left[j\frac{2\pi}{T}nT\right]$$

The random demodulator acquisition scheme consists of 3 operations: A) demodulation; B) low-pass filtering; and C) sampling. The operations performed by the analog front-end have to be modeled in the measurement matrix Φ. Properly defined Φ enables the reconstruction algorithm to successfully recover the sparse representation of the signal from the compressed measurements. Demodulation is carried out by multiplication by a randomly alternating zero-mean ± 1 sequence, often called chipping sequence, which is produced by a random number generator. The chipping sequence needs to be generated at minimum Nyquist frequency of the input signal. Continuous-time demodulation can be described by

$$p_c(t) = \sum_{n=-\infty}^{+\infty} P_c(n)\exp\left[\frac{2\pi}{T}t\right], \quad t \in [0,T), \tag{3}$$

where $P_c(n)$ represents Fourier series coefficients of the chipping sequence $p_c(t)$. The continuous-time multi-band signal is then multiplied with the $p_c(t)$ sequence, which is carried out by a mixer:

$$y(t) = x(t)\cdot p_c(t), \quad t \in [0,T). \tag{4}$$

This corresponds to frequency smearing, which is equal to convolution within the frequency domain [7]. The original sparse signal contains few tones, so it is sufficient to examine a small portion of the spectrum in order to extrapolate them. For that reason, we perform lowpass filtering to prevent aliasing, and we sample with a low-rate ADC at rate T_s. This operation produces compressed samples that are encoded with the representation of the original sparse signal [7]. The discrete measurement vector y can be characterized as a linear transformation of the discrete coefficient vector α. In CS it is expressed by the transform matrix $\mathbf{A} = \Phi\Psi$. The compressed measurement vector y can be modeled as follows [8]:

$$\begin{aligned} y[m] &= \int_{-\infty}^{+\infty} x(\tau)\, p_c(\tau)\, h(mt-\tau)\, d\tau \,|_{t=MT} \\ &= \sum_{n=1}^{N} \alpha_n \int_{-\infty}^{+\infty} \psi_n(\tau)\, p_c(\tau)\, h(mMT-\tau)\, d\tau \end{aligned} \tag{5}$$

In the compressed sensing framework, (5) is approximated by the measurement (sensing) matrix Φ. The measurement matrix, modeling the basic principle of the random demodulator, contains $\frac{N}{M}\cdot C$ pseudo-random ± 1s per row, where N defines the signal length, M is the number of measurements, and C is a scaling constant depending on the impulse response representation.

3. ANALOG FRONT-END AND MEASUREMENT MATRIX

The measurement matrix represents a model of the operations undergone by the signal during acquisition. It enables the recovery to find the sparsest solution to the system. In addition to the chipping sequence, that needs to be reflected in the $\mathbf{\Phi}$ matrix, a precise model of the filter needs to be included. Because the sampling system acquires samples in the time domain, the filtering operation is modeled through its impulse response $h(t)$.

In order to visualize how the matrix $\mathbf{\Phi}$ is constructed, one might consider two matrix factors representing the chipping sequence and impulse response of the filter, $\mathbf{P}$ and $\mathbf{H}$, respectively

$$\mathbf{\Phi} = \mathbf{HP}. \tag{6}$$

The demodulation process multiplies each sequence value $p_{c,n}$ with the corresponding signal input x_n within specified period

$$\mathbf{P} = \mathrm{diag}\{p_{c,1},\ldots,p_{c,N}\}.$$

The structure of the impulse response matrix is a little bit more complicated. The size of the matrix depends on how densely the filter impulse response is discretized. Each row of the matrix corresponds to one processed sample. The offset for each row corresponds to the decimation, which is

directly proportional to the sub-sampling factor

$$\mathbf{H} = \begin{bmatrix} h_n & h_{n-1} & \dots & h_{n-\tau} & \dots & 0 & 0 & 0 & 0 \\ 0 & 0 & h_n & & & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & & \ddots & \ddots & \ddots & 0 & 0 & 0 \\ \vdots & \vdots & & \dots & & h_3 & h_2 & h_1 & 0 \\ 0 & 0 & \dots & 0 & h_n & \dots & h_3 & h_2 & h_1 \end{bmatrix}.$$

The above system of linear dependencies is an approximation of the infinite dimensional analog architecture. Given the above model in (6) and the acquired data, which is described by the CS model in (1), we can try to solve the underdetermined system of equations and reconstruct the sparse vector α corresponding to the sampled input signal, linearly transformed by the sparsifying matrix. In order to obtain the sparsest solution, in principle we would have to solve the combinatorial search implied by:

$$\begin{array}{rl} \min_\alpha & ||\alpha||_0 \\ \text{subject to} & y = \Phi\Psi\alpha \end{array} \tag{7}$$

Solving (7) is an NP-hard problem. Fortunately, the problem can be relaxed to a convex problem, based on ℓ_1-norm minimization, and the sparse vector can be recovered.

$$\begin{array}{rl} \min_\alpha & ||\alpha||_1 \\ \text{subject to} & y = \Phi\Psi\alpha \end{array} \tag{8}$$

This particular approach is called *Basis Pursuit* [9] and it is one of the convex optimization methods used to recover signals within the CS framework. Apart from convex optimization approaches, there is a group of methods called *Greedy Pursuit* where the sparse solution is computed one step at a time by adding new signal components that yield least recovery approximation error. In our work, both types of methods have been employed to investigate recovery sensitivity to measurement matrix deviations due to filter impulse response variations. In the description of the performed simulations, we specify the exact algorithms used.

3.1 Modeling non-ideal effects

Following the methodology of RD, briefly described in this paper, we need a discrete model of the analog front-end in order to perform compressed sensing. The system cannot discard the information that is carried by the signal of interest. Since the acquisition process takes place in the time domain, we aim to design a discrete model of the analog sampler that closely approximates the time-domain system.

As mentioned, we represent the low-pass filtering operation by a discretized model of the filter impulse response. It is worth noting that our modeling approach entails a trade-off between computational complexity in the signal reconstruction and density of the discretized filter impulse response. We assume known filter architecture and parameters from which the filter transfer function $H(s)$ can be calculated, which can be further represented in the discrete-time domain, where one can obtain an approximation of the impulse response.

In order to investigate CS reconstruction sensitivity to changes in the filter response we have developed a rather simple, but for the purpose appropriate, MATLAB framework. We have employed a 4th-order *Butterworth* low-pass filter, with inductive passive network architecture depicted in Fig. 3.

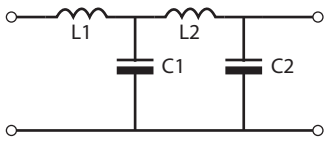

Figure 3: 4th-order low-pass filter

Assuming that our input impedance is represented by R (50 Ω) we can derive a transfer function[1], $H(s)$ for this filter:

$$H(s) = \frac{1}{A\mathbf{s}^4 + B\mathbf{s}^3 + C\mathbf{s}^2 + D\mathbf{s} + 1}, \tag{9}$$

where

$$\begin{array}{ll} A = C_1C_2L_1L_2 & B = C_1C_2L_2R \\ C = L_1C_2 + L_2C_2 + L_1C_1 & D = RC_2 + RC_1 \end{array}$$

Using bilinear transform (a.k.a Tustin's method), one can transform the model to discrete-time (Z) domain, obtain impulse response and create the $\mathbf{H}$ matrix representing the filter. In reality, the passive component values, due to the manufacturing process, deviate from their ideal technical specification. In particular, each component is produced under a process which cannot be controlled with sufficient precision. The process can be modeled with a Gaussian distribution where the mean μ is the target value. Quality assessment of the components (inductors, capacitors, resistors) selects them according to standard guidelines (e.g. IEC60063[2]) for choosing exact product dimensions within a given set of constraints. Depending on the design budget and component dimensions we end up with fabricated passive components with values different by up to, for example 1% or 10% from what we aimed for. Hence, the $\mathbf{H}$ matrix designed according to ideal values will be an inaccurate representation of the filter device. When possible, one might try to calibrate the algorithm according to performed measurements on the particular device, since component values due to production tolerances are static. Fig. 4 illustrates 4th-order Butterworth filter impulse response deviations from the expected response. Two worst-case scenarios are depicted, where all the capacitors deviate by 5% and inductors by 10% from their expected values.

Another problem with the elements in the hardware realization is that nominal values of the components might fluctuate dynamically due to external factors like e.g., ambient temperature variations. Dynamic component fluctuations directly introduce change to the filter response, which cannot be calibrated real time. We would like to point out that the problem of filter non-idealities, as well as RD hardware implementation issues is more complex and some of the other possible problems have been discussed in [7, 10], where some compensation schemes have been proposed. The idea of calibrating the device has been pointed out by [7].

[1] Assuming high load impedance we have neglected it for the purpose of simplicity

[2] *IEC, Preferred number series for resistors and capacitors, Standard IEC 60063, International Electrotechnical Commission, 1963.*

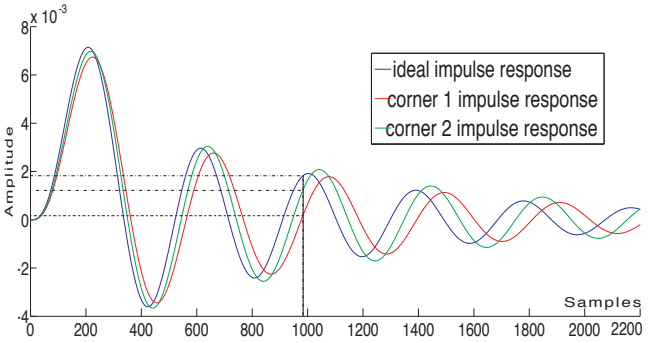

Figure 4: Impulse responses of the ideal component value low-pass filter and two worst case component deviations.

This calibration could help to decrease the impact of passive element fluctuations, but only in limited cases.

3.2 Simulation scheme

A MATLAB framework has been designed to simulate component deviations and examine the reconstruction error performance of different CS reconstruction algorithms. Algorithm 1 briefly describes the simulation methodology.

In our experiments, we use a multi-tone signal with five tones distributed within a 100 Hz to 2.1 kHz range, where 2.1 kHz represents the maximal signal frequency. The input x is represented by $N = 9000$ samples; oversampling imitates the analog front-end behavior. Additionally, the input signal is corrupted by additive white Gaussian noise (AWGN) and its signal-to-noise ratio (SNR) is set to 50 dB. The low-pass filter was synthesized with a 500 Hz cut-off frequency (f_{cut}), where the following values from (9) were calculated: $A = 2.053E{-}14$, $B = 8.428E{-}11$, $C = 3.459E{-}07$ and $D = 8.318E{-}4$. Sub-sampling operates at rate 1 kHz (twice the f_{cut}). The analog front-end for each input vector x outputs a 1000-samples compressed vector $\mathbf{y}$, also known as an observation vector ($M = 1000$). The trivial initial conditions have been chosen to assure almost perfect signal reconstruction, with high probability [7]. Additionally the low frequencies that were chosen for the input signal were dictated by computational limitations of our simulation framework.

The RD measurement matrix, which is our central object of interest, is generated from Tustin approximation of an analog filter input response and pre-generated chipping sequence. The horizontal size of the measurement matrix N corresponds to one processed period $[0, T]$ of an oversampled input signal. This dense measurement system representation enables accurate impulse response representation. Modeling the low-pass filter behavior with emphasis on component deviations enables the following benchmarking strategy for recovery sensitivity: the reconstruction system is generated according to the ideal filter specification, while the actual analog sampling simulation uses a filter with non-ideal component values. As one may guess from the system representation in (9) or Fig. 4, reconstruction without appropriate knowledge of the actual deviations in component values will increase the recovery error.

In order to investigate the performance degradation of the CS reconstruction, three common recovery methods have been utilized: Basis Pursuit, Basis Pursuit De-Nosing and Orthogonal Matching Pursuit [9, 11]. Orthogonal Matching Pursuit (OMP) is the canonical greedy algorithm for the sparse approximation. In order to solve an underdetermined system of equations using the mentioned algorithms, the following MATLAB toolboxes have been utilized: *CVX*[3], SPGL1[4] and *Sparsify*[5].

Algorithm 1 RD acquisition and reconstruction

Input: test signal x (multitone, frequency sparse)
Generate chipping sequence P_c, create $\mathbf{P}$ matrix
Declare num. of measurements M, $(M < N) \leftarrow F_s$
Set cut-off frequency $F_{cut} = \frac{F_s}{2} \leftarrow$ anti-aliasing
Synthesize *LP* filter and obtain $\mathbf{H} \in \mathbb{R}^{M \times N} \leftarrow h(n)$
Create sparsity basis **DFT** matrix $\leftarrow \Psi^{N \times N}$
for component K_1 **to** K_{nth} **do**
  Change nominal value according to tolerance **Tol**
  Perform *random demodulation* on x
  Reconstruct sparse vector $\hat{\alpha}$, using *BP*, *BPDN*, *OMP*
  Recover input x: $\hat{x} = \Psi\hat{\alpha}$
  Obtain recovery error vector: $x - \hat{x}$
  Calculate quality of the reconstruction $\leftarrow$ **SNR**
$$\mathrm{SNR} = 20 \log_{10}\left(\frac{||x||_2}{||x-\hat{x}||_2}\right)$$
  Compare performances of RD with deviated filter values $\mathbf{SNR_{tol}}$ to the ideal case $\mathbf{SNR_{ideal}}$
end for

4. RESULTS

We perform two types of experiments. First, we simulate the nominal values case and the 16 worst-case scenarios, i.e. the worst-case combinations of deviating values of components C_1, C_2, L_1, and L_2 from (9). We consider two capacitors and two inductors in the low-pass filter (Fig. 3), deviating from expected values by $\pm 5\%$ and $\pm 10\%$, respectively. Performing reconstruction with the changed values, we were able to estimate recovery performance boundaries for worst-case scenarios. Fig. 5 illustrates signal recovery performance ($\mathrm{SNR_{tol}}$ values) of all three algorithms solving underdetermined system of equations low-pass filtered with worst-case component tolerances. Each separate "corner" has been labeled as $(c1, \ldots c16)$. The "ideal" figure of merit in the plot indicates recovery performance in case of the ideal match between measurement matrix and actual filter component values. Ideal case reconstruction returns an almost perfectly recovered signal ($\simeq 48.9$ dB) for convex optimization algorithms *BP, BPDN* and $\simeq 46$ dB using the greedy approach (OMP). Since the amount of noise added to the input is very small, it is not a surprise that the convex methods perform alike during benchmark (maximum difference 0.6 dB). OMP appeared to react slightly different in each test case but the reconstruction degradation for all of the methods is substantial.

Results of the worst-case simulations indicate a drop in SNR down to 40 dB. Although the extreme cases do not occur with high probability, when designing RD signal acquisition systems for commercial purposes one cannot neglect this effect. Since we are focusing on practical issues related to the implementation, we have also performed Monte Carlo

[3] CVX: *http://cvxr.com/cvx/*
[4] SGPL1: *http://www.cs.ubc.ca/labs/scl/spgl1/*
[5] SPARSIFY:*www.personal.soton.ac.uk/tb1m08/sparsify/sparsify.html*

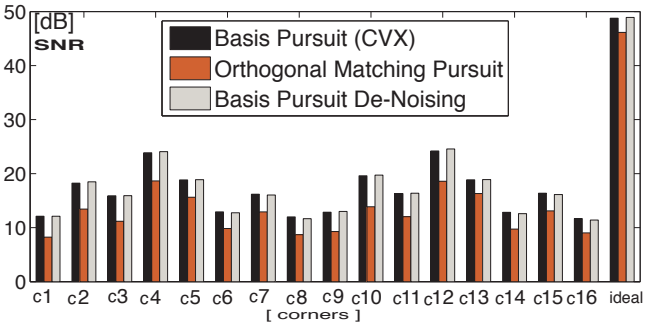

Figure 5: Performance of 16 worst-case combinations under *BP, BPDN* and *OMP* recovery methods.

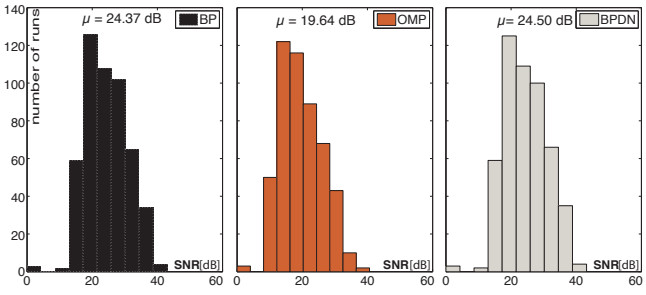

Figure 6: Monte Carlo simulation results, 500 runs (3σ truncated Gaussian distribution) with multitone input signal with AWGN at 50 dB SNR. Histograms represent number of simulation runs to SNR of the reconstruction.

simulations with 500 runs, using the same initial conditions for the input signal. Component values for each examination iteration were randomly chosen from a truncated *Gaussian distribution*. The truncation factor has been set according to the same component tolerance values as for the worst-case scenario tests (5%, 10%). Fig. 6 shows the simulation results. The observed average reconstruction quality drop for BP and BPDN is $\simeq 24$ dB, and $\simeq 29$ dB in the case of OMP.

5. CONCLUSIONS

In this paper we have presented an investigation of the common compressed sensing recovery methods, utilized with the random demodulator setup. The research elaborates on systems suffering from component deviations of a filter located in the analog front-end. Special attention has been paid to the representation mismatch in the measurement matrix between the ideal and actual analog front-end model.

The problem analysis emphasizes the implementation aspects. When we deal with a real system, sensitivity of the reconstruction framework plays an important role. Ongoing system fluctuations due to temperature or power supply changes, as well as static component deviations are difficult to avoid. A simple, but rather representative filter model was chosen to be tested under random demodulator setup, with common compressed sensing recovery techniques applied. Worst-case scenario simulations reveal substantial recovery degradation. Reconstruction error increase of up to 40 dB has been observed. Through Monte Carlo experiments we extend the scope of our methodology towards physical realizations, indicating mean recovery error increase of 24 dB for convex optimization methods and 29 dB for the greedy approach. Performed simulations clearly show that common CS recovery schemes are highly sensitive to the low-pass filter component tolerances. This problem deserves special attention when compressed sensing is used to sample moderate frequency analog signals using the RD framework. Moreover, utilizing low cost off-the-shelf parts without making additional improvements to the recovery scheme will result in poor system performance.

Future work of this research will aim at improving the reconstruction robustness to the measurement matrix inaccuracy in modeling the analog front-end. The compressed sensing method of sampling has great potential, but in order to make it economically reasonable for the industry to adopt, improved recovery methods need to be developed.

REFERENCES

[1] E. J. Candès, "Compressive sampling," in *Proceedings of the International Congress of Mathematicians*, vol. 3, pp. 1433–1452, 2006.

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

[3] E. J. Candès and M. Wakin, "An introduction to compressive sampling," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 21–30, 2008.

[4] S. Kirolos, J. Laska, M. Wakin, M. Duarte, D. Baron, T. Ragheb, Y. Massoud, and R. Baraniuk, "Analog-to-information conversion via random demodulation," in *IEEE Dallas/CAS Workshop on Design, Applications, Integration and Software, 2006*, pp. 71–74, 2006.

[5] M. Mishali, Y. Eldar, O. Dounaevsky, and E. Shoshan, "Xampling: Analog to digital at sub-Nyquist rates," *Circuits, Devices & Systems, IET 2011*, vol. 5, no. 1, pp. 8–20, 2011.

[6] T. Ragheb, J. Laska, H. Nejati, S. Kirolos, R. Baraniuk, and Y. Massoud, "A prototype hardware for random demodulation based compressive analog-to-digital conversion," in *51st Midwest Symposium on Circuits and Systems, MWSCAS 2008.*, pp. 37–40, 2008.

[7] J. Tropp, J. Laska, M. Duarte, J. Romberg, and R. Baraniuk, "Beyond Nyquist: Efficient sampling of sparse bandlimited signals," *IEEE Trans. Inf. Theory, 2010*, vol. 56, no. 1, pp. 520–544, 2010.

[8] J. Laska, S. Kirolos, M. Duarte, T. Ragheb, R. Baraniuk, and Y. Massoud, "Theory and implementation of an analog-to-information converter using random demodulation," in *IEEE International Symposium on Circuits and Systems, ISCAS 2007*, pp. 1959–1962, 2007.

[9] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM Review*, vol. 43, no. 1, pp. 129–159, 2001.

[10] Y. Chen, M. Mishali, Y. C. Eldar, and A. O. Hero, "Modulated wideband converter with non-ideal lowpass filters," in *IEEE International Conference on Acoustics, Speech and Signal Processing, 2010. ICASSP*, pp. 3630–3633, 2010.

[11] T. Blumensath and M. Davies, "Gradient pursuits," *IEEE Trans. Signal Process.*, vol. 56, no. 6, pp. 2370–2382, 2008.

The simulation framework described in this article as well as relevant documentation are available online at:

http://www.sparsesampling.com/eusipco2011