

# A 100MHz-2GHz 12.5x sub-Nyquist Rate Receiver in 90nm CMOS

Juhwan Yoo\*, Stephen Becker<sup>†</sup>, Matthew Loh\*, Manuel Monge\*, Emmanuel Candès<sup>‡</sup> and Azita Emami-Neyestanak\*

\*Department of Electrical Engineering, California Institute of Technology, Pasadena, CA 91125

<sup>‡</sup>Department of Statistics, Stanford University, Stanford, CA 94305

<sup>†</sup>Laboratoire Jacques-Louis Lions, Paris 6 University, Paris, France

**Abstract**—A fully-integrated, high-speed, wideband receiver called the random modulation pre-integrator is realized in IBM 90nm digital CMOS. It achieves an effective instantaneous bandwidth of 2GHz, with >54dB dynamic range. Most notably, the aggregate digitization rate is  $f_s = 320\text{MSPS}$ , 12.5 $\times$  below the Nyquist rate. Signal recovery can be accomplished for any signal with a concise representation. The system is validated using radar-pulses and tones as the input and recovering the time-domain waveforms.

**Index Terms**—Compressed Sensing, Random-Modulation Pre-Integrator, Sub-Nyquist Sampling.

## I. INTRODUCTION

The past 50 years have seen tremendous developments in electronics due to the rise and rapid development of IC-fabrication technology. In addition to making available cheap and abundant computing resources, tremendous effort has been devoted to developing wireless technologies. While the central focus of wireless research has been the mobile communication sector, an emerging area concerns the development of RF sensing and spectral applications over bandwidths exceeding multiple GHz. Such systems have many applications, ranging from scientific instrumentation to electronic intelligence. Although some solutions already exist, their large size, weight and power consumption make more efficient solutions desirable.

At present, the primary bottleneck in designing such systems is the power consumption of the back-end ADCs at the required *digitization* rate. ADCs are a dominant source of power consumption; it is also often the case that ADC block specifications are used to determine parameters for the rest of the signal chain such as the RF front-end and DSP-core which processes the digitized samples [1]. Conventionally, increases in system bandwidth have come from developing ADCs with superior performance.

In this work, an alternative method to reduce system power consumption via minimizing the required *digitization* rate is presented. The approach utilizes results from the field of compressed sensing (CS) [2,3]. Key to this approach is the use of random-sensing which is known to support undersampling. In §II we present a prototype wide-band receiver that implements random-sensing, called the random modulator pre-integrator (RMPI), which is an architecture specialized for the recovery of signals with a known concise structured representation (e.g. *compressible* signals).<sup>1</sup> The system captures an

<sup>1</sup>To be specific, we assume the signals of interest  $x(t)$  can be efficiently represented in a dictionary  $\Psi$ , with atoms  $\psi_\ell(t)$ ,  $\ell \in \{1, \dots, d\}$ , meaning that in the expansion  $x(t) = \sum_\ell a_\ell \psi_\ell(t)$ , only a few of the coefficients are significant.

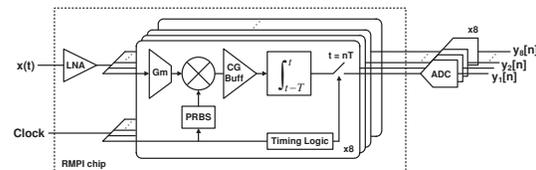


Fig. 1: Simplified block diagram of the RMPI

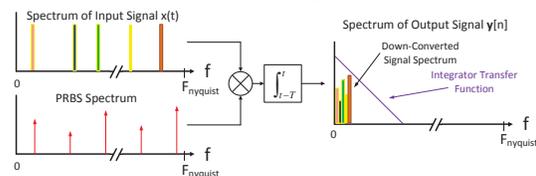


Fig. 2: Operation of RMPI depicted from the frequency domain

effective-instantaneous-bandwidth (EIBW) spanning 100MHz–2GHz with dynamic range > 54dB while digitizing samples at an aggregate rate of 320MSPS—a factor 12.5 $\times$  lower than the Nyquist rate.

## II. RMPI ARCHITECTURE

A block diagram of the RMPI is shown in Fig. 1. It consists of 8 parallel channels with a common input node driven by the input LNA. Each channel modulates the input signal  $x(t)$  with a distinct PRBS sequence  $p_i(t)$ ,  $i \in \{1, \dots, 8\}$ , where  $i$  is a channel index, toggling at the Nyquist rate. The output of the mixer  $r(t) = x(t)p_i(t)$  is then integrated over a fixed time-interval  $T$  and digitized at a rate  $f_{\text{ADC}} = 1/T \ll f_{\text{Nyq}}$ . Thus we write the digitized output samples  $y_i[n]$ ,  $i \in \{1, \dots, 8\}$  as  $y_i[n] = \int_{t-T}^t x(t)p_i(t)dt$ ,  $t = nT$ . In our system, we set  $T = 100/f_{\text{Nyq}}$ , where  $f_{\text{Nyq}} = 4\text{GHz}$ , for a digitization rate of  $f_{\text{ADC}} = 40\text{MSPS}$ , leading to an overall system back-end sampling-rate of  $f_s = 320\text{MSPS}$ .

At first glance, the idea that a signal lying anywhere in some bandwidth  $B$  (or even occupying the full band  $B$ ) can be recovered by samples collected at a sub-Nyquist rate  $f_s \ll 2B$  seems to violate the Shannon-Nyquist sampling theorem. We point out, however, that the Shannon-Nyquist theorem is a statement about an arbitrary signal with bandwidth  $< B$ ; the RMPI targets a more specific class of signals that possess structure beyond being band-limited. To give an idea of how the RMPI works, we show the frequency domain point of view in Fig. 2. A PRBS with period  $T_{\text{rep}}$  will have power concentrated at harmonics of the repetition rate  $f_{\text{rep}} = 1/T_{\text{rep}}$  (denoted by vertical arrows in Fig. 2). By mixing input  $x(t)$  with a PRBS, a shifted copy of the entire input spectrum is made by each harmonic of the PRBS; consequently, signal energy from the

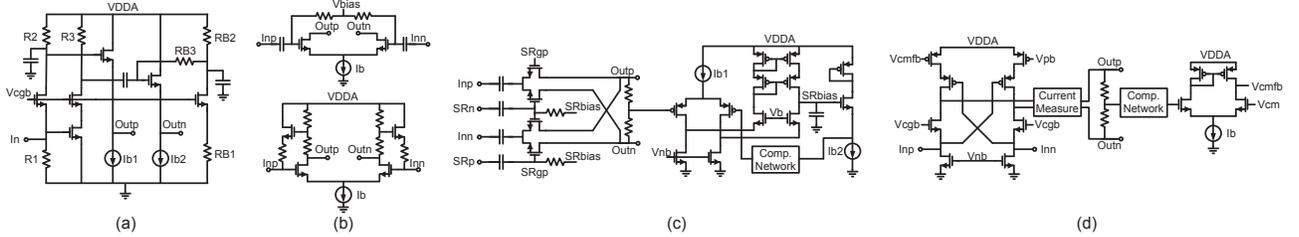


Fig. 3: Simplified schematics of key circuit blocks. (a) LNA, (b) clock distribution, (c) mixer, (d) current buffer.

entire spectrum is down-converted into the passband of the back-end digitizers. Thus, signal energy from the entire input spectrum is captured by the back-end ADCs.

In general, the information converted to baseband is insufficient to uniquely determine the original spectrum of an arbitrary band-limited signal. However, because we assume extra prior information, namely sparsity, CS techniques allow recovery; see §IV for details.

Successful CS recovery relies upon knowing a basis or dictionary in which the signal can be concisely represented which, is often the case in many applications. The basis need not be the frequency basis; in this paper, we reconstruct radar-pulses which are sparse in the *time-frequency* plane, hence we use a multi-scale Gabor dictionary [4].

### III. CIRCUIT IMPLEMENTATION

*a) Design Strategy and Implementation Description:* The input to the RMPI is a common-gate/common-source LNA shown in Fig. 3a. The LNA has a broadband  $50\Omega$  match set by the transconductance of the common-gate path and resistor R1 that sets the bias current. The gain of the LNA is 18dB and has a  $f_{3dB} \approx 3\text{GHz}$ . The  $s_{11} < -15\text{dB}$  from 10MHz–5GHz.

*b) Correlator Channel Circuits:* The output of the LNA drives 8 parallel correlator channels. Each channel is implemented as a modified direct-conversion receiver. In order to maximize dynamic range performance, the channels were implemented using a current domain approach [5,6]. The large voltage-amplitude output of the LNA is converted into a large current-domain signal with a source-degenerated transconductor. This current is then mixed by a passive-mixer with the LO-port driven by a programmable 128b shift register programmed with a PRBS.<sup>2</sup>

A schematic of the mixer is shown in Fig. 3c. The size of the passive-mixer switches and the  $Z_{in}$  of the current-buffer is set low to minimize the voltage amplitude at the mixer-node. This enables mixing of large amplitude current signals with minimum voltage amplitude mitigating nonlinearity from the switches in the mixer.

Fig. 3d shows the current-buffer topology. The buffer design is a hybrid pseudo-differential common-gate amplifier in parallel with a cross-coupled PMOS common-source stage, which improves current-gain [6]. The buffer also lowers the effective noise-contribution of the succeeding cascade of op-amps.

The cascade of class-A op-amp based TIA RC-integrator and unity-gain buffer serves both as the integrator in the

<sup>2</sup>A programmable shift-register was used to provide a flexible testing platform, but considerable power savings are possible by replacing it with a suitable LFSR.

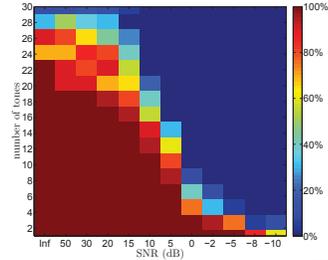


Fig. 4: Recovery rate as function of sparsity and noise. The rate is the mean number of success over 20 random trials; a “success” is recorded when the frequency estimates are all accurate within 0.5MHz. Recovery is by OMP [9].

correlator and the off-chip ADC driver. The output of the unity-gain buffer is digitized and exported to a PC where numerical optimization algorithms are used to recover the input signal §IV.

*c) PRBS:* As the RMPI reconstructs windows of finite length, a repeating PRBS is an admissible substitute for a random binary wave of infinite length [7,8]. The use of a PRBS imposes certain design constraints which are explained in [8].

*d) Clock Distribution:* Channel-to-channel timing accuracy as well as minimum duty-cycle distortion is crucial in producing compressed samples that allow high-fidelity reconstruction. Empirical simulations reveal that time-domain signal reconstruction is sensitive to the  $\sigma_{\text{jitter}}$  of the Nyquist-rate clock distributed to each channel. With all other system parameters fixed at nominal levels (thermal noise, non-linearity, etc.), the Nyquist-rate must have  $\sigma_{\text{jitter}} < 0.5\text{ps}$  in order to achieve 60dB dynamic range for radar-pulse reconstruction [8].

Thus, a current-mode analog clock distribution with a symmetric binary tree topology was selected to distribute the clock to the 4 pairs of channels. The analog clock-distribution is highlighted on the die photo of Fig. 10a. The current-mode distribution consists of 3 basic sub-blocks: an open-drain driver shown in Fig. 3b, which drives a  $100\Omega$  differential transmission-line, which is then terminated by a  $100\Omega$  differential input-impedance TIA shown in Fig. 3b. These 3 components are cascaded as necessary; the final clock receiver drives a CML-CMOS stage whose output is driven through a cascade of DCD reducing coupled-inverters. The output of the inverters drive the input to the clock distribution of the SRs. The input to the overall clock-distribution is an open-drain driver whose inputs have been terminated with  $50\Omega$  resistors.

### IV. SIGNAL RECOVERY ALGORITHMS

The RMPI *time-domain* signal recovery procedure consists of determining a vector  $\mathbf{x} \in \mathbb{R}^N$ . This finite-length discrete vector represents the desired input-signal sampled at

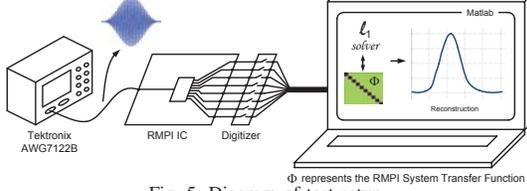


Fig. 5: Diagram of test setup.

the Nyquist-rate. Each channel simultaneously outputs samples every  $T = 100T_{\text{Nyq}}$ . For example, if the RMPI is reconstructing a  $N = 1000$  point vector, the matrix  $\Phi$  representing the sampling operation is block-diagonal, each block having 8 rows and length  $T_{\text{Nyq}}/T = 100$  Nyquist bits. Denoting the vector of samples obtained from the system by  $\mathbf{y}$ , the RMPI mode of acquisition can be modeled as  $\mathbf{y} = \Phi\mathbf{x}$  where  $\Phi \in \mathbb{R}^{80 \times 1000}$ . The rows of  $\Phi$  contain  $\pm 1$  entries over the integration window and 0 elsewhere.

Signal recovery consists of finding a solution,  $\mathbf{x} \in \mathbb{R}^N$ , to the under-determined linear system. Without prior knowledge, the problem is ill-posed since there are infinitely many solutions. CS theory provides rigorous proofs that if  $\mathbf{y} = \Phi\mathbf{x}$  and  $\mathbf{x}$  is sufficiently *sparse* (say, only  $k$  nonzeros), then  $\mathbf{x}$  can be recovered by searching for the feasible solution which has minimum  $\ell_1$ -norm as long as  $\Phi$  has  $\mathcal{O}(k \log N)$  rows.<sup>3</sup> Finding the minimum  $\ell_1$ -norm solution is a linear program known as *basis pursuit* and can be solved efficiently. Details of CS theory are in [2,3], and details pertaining specifically to the RMPI, including algorithms and variants, are in [4,7,8]. The exact amount of sparsity required can be determined by simulations. Fig. 4 shows a synthetic data simulation with signals composed of multiple equi-amplitude tones (frequencies and phases chosen randomly). The noise level, in dB, is synthetic noise and does not represent the real system. For noiseless input, the system can always recover 22 tones, and sometimes recover up to 30 tones. As the noise level increases, the system cannot handle as many tones until eventually, at extreme noise levels, it fails with a single tone.

## V. MEASUREMENT RESULTS

A diagram of the test setup used to validate the RMPI is shown in Fig. 5. Test stimuli consisted of tones and pulses of multiple widths and frequencies synthesized via an arbitrary waveform generator. Outputs of the RMPI IC are digitized off-chip and exported to a PC. The digitized samples are then used to obtain the input signal via a numerical optimization procedure [4]. We now present several case studies that demonstrate the potential of the RMPI. Fig. 6 shows reconstructions of 400ns pulses with carrier frequencies of about 87MHz and 1947MHz. This demonstrates the EIBW of the RMPI. There was no change in operating conditions (e.g. tuning of the LO) in order to recover these signals. The RMPI is capable of capturing signals from the entirety of the input bandwidth and not just the IF bandwidth of the baseband filters used about a single frequency.

<sup>3</sup>This isn't true for any matrix  $\Phi$ , but it is true if  $\Phi$  is sufficiently similar to a matrix with random  $\pm 1$  entries, which is why the RMPI uses a PRBS for mixing.

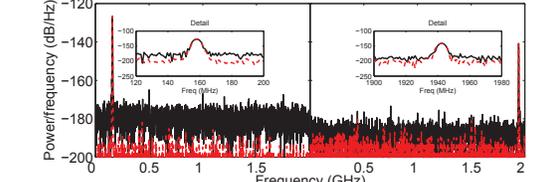
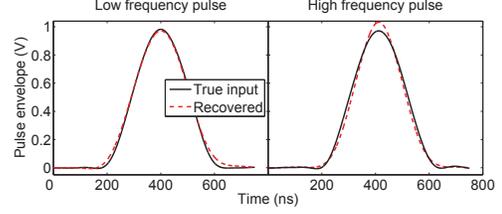


Fig. 6: Recoveries of pulses at both low and high carrier frequencies.

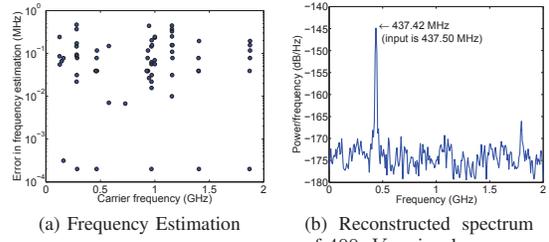


Fig. 7: Carrier frequency estimation vs.  $f_{in}$  and low-amplitude reconstruction.

Fig. 7a shows the error in carrier frequency estimation of baseband pulse envelope reconstructions with MSE  $< 10\%$ . The median frequency estimation error was  $< 69\text{kHz}$ , demonstrating the RMPI's potential in spectral-sensing applications.

The single-tone dynamic range was tested by sending in and reconstructing low-amplitude tones. Fig. 7b shows the reconstruction of a tone of  $400\mu V_{pp}$ , which is 54dB below the full-scale input set by the  $P_{1\text{dB}}$  compression point of the receiver.<sup>4</sup> A signal that is difficult for even standard Nyquist-rate receivers to handle is that of two overlapping pulses at different frequencies. Fig. 8 shows two pulses with 200ns overlap. The MSE of the pulse envelope reconstruction is less than 10%. Fig. 9 shows reconstructions of both a 50ns and 75ns pulse. Although the pulse-envelope reconstructions are of low-quality, what is notable is that the accurate frequency estimation of the pulses is possible from 16 and 24 RMPI samples respectively

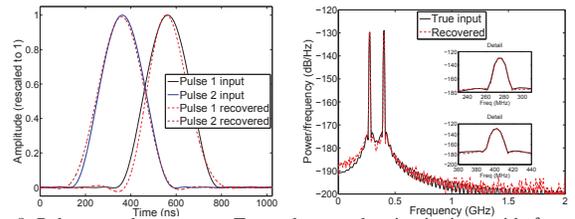


Fig. 8: Pulse-on-pulse recovery. Two pulses overlapping in time, with  $f_{\text{carrier}} = 275\text{MHz}/401\text{MHz}$ , are recovered from hardware data. The  $f_{\text{carrier}}$  of both pulses is estimated to within  $.234\text{MHz}$ . Left: time-domain, right: frequency-domain.

<sup>4</sup>The  $P_{1\text{dB}}$  of the receiver is not necessarily the upper-bound of the reconstruction capabilities of the receiver. Due to the variation in concentration of power at different harmonics for different frequencies, larger amplitude signals can still be recovered.

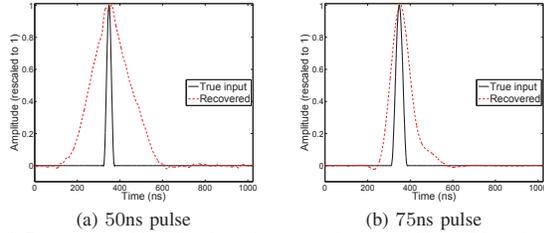


Fig. 9: Reconstructed baseband windows of pulses of width (a) 50ns, (b) 75ns.

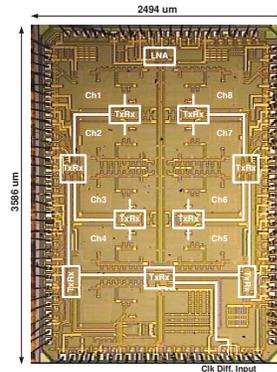
versus the  $\geq 200$  required by Nyquist. These results are not isolated incidents and demonstrate data-compressing aspects of the RMPI. A summary of RMPI performance when operated as a direct down-converter is shown in Fig. 10b.

## VI. DISCUSSION

While a fair comparison between a CS system and a Nyquist-rate system is a difficult proposition at best, we make a brief comparison between the RMPI and a high-speed ADC to provide a context for interpreting the results. Examination of state-of-the-art ADCs reported in ISSCC from 2009-2011 gives several possible points of comparison. For example, a 10b interleaved ADC implemented in 65nm CMOS with  $f_{ADC} = 2.6\text{GSPS}$  consumes 480mW and  $5.1\text{mm}^2$  [10]. Implementation of digital filtering necessitates oversampling the Nyquist-rate by a factor  $\geq 2$ . This means that the 2.6GSPS ADC would be used for a bandwidth of  $\leq 600\text{MHz}$ . In addition, two of these ADCs would be required to perform coherent detection via I/Q demodulation consuming 2W excluding any RF front-end that was required. Assuming a roughly linear scaling between sampling-rate and power consumption, this would equate to  $> 6\text{W}$  of power consumption to realize an EIBW  $\approx 2\text{GHz}$ .

In contrast, the prototype RMPI consumes only 506.4mW of power. We do not include the power consumed by the computational platform needed for signal recovery. CS-based signal reconstruction is a non-linear function of the acquired samples and more computationally expensive than Nyquist-rate reconstruction. As a result, real-time *time-domain* reconstruction is not currently practical in portable (low-power) applications and is the subject of extensive on-going research. To give a rough idea of the computational costs, a typical CS-recovery algorithm requires about 20-1000 FFTs.

This cost only applies when a complete time-domain reconstruction is needed. Often—if not almost always—the desired information is not the complete time-domain waveform but rather a small set of parameters (e.g. pulse width, carrier frequency, initial phase, etc.) which are extracted from the acquired waveform. It is possible to estimate parameters of the desired signal directly without first reconstructing the time-domain [4]. This procedure is less computationally expensive and potentially competitive with traditional Nyquist-rate approaches. The exact cost depends on the complexity of the signal model; for example, if the signal contains exactly one frequency, then frequency and phase information can be recovered at the cost of only one FFT. Furthermore, in many applications such as UAV, recovery is already done offline. In such applications, the RMPI can reduce the strain on the downlink budget in addition to reducing power consumption.



(a) RMPI IC

Implementation	
Technology	IBM 90nm CMOS9SF
Die Area	$8.85\text{mm}^2$
Supply Voltage	1.5V Analog/ 1.2V Digital/ 2.5V Digital I/O
PRBS Clk frequency	4 GHz
Rx Chain System Performance (as a Down-Converter)	
Gain	41 dB
NF	7 dB
$S_{11}$	$< -15\text{dB}$ , 10MHz-5GHz
p1dBm (CS mode)	-10 dBm
Rx Chain BW	2.7 GHz
Bandwidth	100MHz - 2GHz
Dynamic Range (CS Recovery)	54 dB
"Undersampling"	12.5X
$\sigma_{\text{system jitter}}$ (rms)	$< 300\text{fs}$
Channels (units)	8
Power Consumption (w/o out. Buff.)	506.4 mW

(b) Performance Table

Fig. 10: Performance Summary

The RMPI benefits greatly from parallelization. This is due to the fact that high-fidelity CS recovery relies upon knowledge of the transfer function of the overall system, as opposed to performance with rigorous specification. Thus, compared to other channelized system architectures, the need for well-matched blocks is lower, potentially allowing greater scalability.

## VII. CONCLUSION

To the best of the authors' knowledge, the first fully-integrated, high-speed, CS-based receiver capable of recovering structured signals at a sub-Nyquist rate has been demonstrated in 90nm CMOS. It can recover signals from an EIBW of 100MHz-2GHz with  $> 54\text{dB}$  dynamic range while undersampling the Nyquist-rate by  $12.5\times$ . The chip occupies  $8.85\text{mm}^2$  and consumes 506.4mW; a die photo is shown in Fig. 10a.

## VIII. ACKNOWLEDGEMENTS

This work was funded under DARPA grant FA8650-08-C-7853; in particular, the authors are indebted to Dr. Denis Healy for his foresight and encouragement. Professors J. Romberg & M. Wakin, and Drs. M. Grant & E. Keehr also contributed significantly.

## REFERENCES

- [1] B. Murmann, "Trends in low-power, digitally assisted A/D conversion," *IEICE Trans. Electron.*, vol. E93-C, no. 6, pp. 718-727, 2010.
- [2] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inform. Theory*, vol. 52, no. 2, pp. 489-509, 2006.
- [3] E. J. Candès and M. Wakin, "An introduction to compressive sampling," *IEEE Sig. Proc. Mag.*, vol. 25, no. 2, pp. 21-30, 2008.
- [4] S. Becker, *Practical Compressed Sensing: modern data acquisition and signal processing*, Ph.D. thesis, California Institute of Technology, 2011.
- [5] R. Bagheri et al., "An 800MHz-to-5GHz software-defined radio receiver in 90nm CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2006, pp. 1932-1941.
- [6] E. A. Keehr and A. Hajimiri, "A rail-to-rail input receiver employing successive regeneration and adaptive cancellation of intermodulation products," in *IEEE Radio Frequency Integrated Circuits Symposium*, 2010, pp. 47-50.
- [7] J. A. Tropp et al., "Beyond Nyquist: Efficient sampling of sparse bandlimited signals," *IEEE Trans. Inform. Theory*, vol. 56, no. 1, pp. 520-544, 2010.
- [8] J. Yoo et al., "Design and implementation of a fully integrated compressed-sensing signal acquisition system," *ICASSP, accepted*, 2012.
- [9] J. Mairal, *SPAMS: SParse Modeling Software*, v2.1, 2011, Available at <http://www.di.ens.fr/willow/SPAMS/>.
- [10] K. Doris et al., "A 480mW 2.6GS/s 10b 65nm CMOS time-interleaved ADC with 48.5dB SNDR up to Nyquist," in *ISSCC Dig. Tech. Papers*, Feb. 2011, pp. 180-182.