Theory and Implementation of an Analog-to-Information Converter using Random Demodulation

Jason N. Laska, Sami Kirolos, Marco F. Duarte
Tamer S. Ragheb, Richard G. Baraniuk, Yehia Massoud
Department of Electrical and Computer Engineering
Rice University, Houston, Texas

***Abstract*— The new theory of compressive sensing enables direct *analog-to-information conversion* of compressible signals at sub-Nyquist acquisition rates. We develop new theory, algorithms, performance bounds, and a prototype implementation for an analog-to-information converter based on random demodulation. The architecture is particularly apropos for wideband signals that are sparse in the time-frequency plane. End-to-end simulations of a complete transistor-level implementation prove the concept under the effect of circuit nonidealities.**

I. Introduction

The prevalence of digital signal processing in sensing and communication applications has popularized the use the analog-to-digital converters (ADC). The ADC process is based on the Nyquist sampling theorem, which guarantees the reconstruction of a band-limited signal when it is uniformly sampled with a rate of at least twice its bandwidth. Emerging applications like radar detection and ultra-wideband communication are pushing the performance of ADCs – and sampling systems in general – toward their physical limits.

In many cases of interest the signals have additional structure than band-limitedness alone; for example, some signals are *compressible* in some transform domain. Over the past two years, a new theory of *compressive sensing* (CS) has emerged, which exploits this knowledge to achieve signal acquisition using fewer measurements than the number prescribed by the Nyquist theorem for certain classes of signals. In particular, CS allows reconstruction of signals which are *compressible* by some transform (such as Fourier, wavelet, etc.). By leveraging the CS theory, an analog-to-information converter (AIC) can be designed to acquire samples at a lower rate while successfully recovering the compressible signal of interest.

In this paper, we present a complete transistor-level realization for a prototype AIC. We present an analysis of the reconstruction quality that arises in a hardware implementation. Our design is based on an extension of the CS framework to analog signals, first described in [1]. We find that for signals that feature considerable structure – and are thus compressible – the performance of the AIC system is comparable to that of the standard ADC system in terms of Signal to Noise Ratio (SNR) or Effective Number of Bits (ENOB), while the AIC sampling features a rate that is much lower than the Nyquist rate used by ADCs. We also develop a circuit implementation as a proof-of-concept, where we consider the non-idealities inherent in the transistor-level implementation for the different blocks in the system.

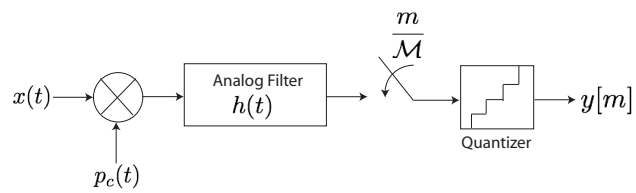

Fig. 1. Pseudo-random demodulation scheme for AIC.

This paper is organized as follows. First, we present a brief background of the CS theory and extend the mathematical framework to AICs in Section II. Second, we we provide a theory for non-ideal sampling systems to predict the SNR performance of real-world implementations in Section III. Third, we support the theory and mathematical concepts through an end-to-end simulation of a proof of concept implementation of the framework: we provide details of the transistor-level implementation in Section IV and present results in Section V. Finally we conclude in Section VI.

II. Compressive Sensing for AIC Systems

A. Compressive sensing background

Compressive Sensing (CS) provides a framework for acquisition of an $N \times 1$ discrete-time signal vector $\mathbf{x} = \mathbf{\Psi}\boldsymbol{\alpha}$ that is *compressible* in some *sparsity basis or frame* matrix $\mathbf{\Psi}$ (where each column is a basis or frame vector ψ_i). By compressible we mean that the entries of $\boldsymbol{\alpha} = [\alpha_1, \alpha_2, \ldots, \alpha_N]$, when sorted from largest to smallest, decay rapidly to zero; such a signal is well approximated using a K-term representation, consisting of the terms of $\boldsymbol{\alpha}$ with the K largest magnitudes while setting all the other terms to zero. Note that, by definition, signals that have only a few nonzero coefficients are compressible as well.

The CS framework [2], [3], demonstrates that a signal that is compressible in one basis $\mathbf{\Psi}$ can be recovered to a quality similar to that of a K-term approximation from $M = O(K \log \frac{N}{K})$ *nonadaptive* linear projections onto a second basis $\mathbf{\Phi}$ that is *incoherent* with the first. By incoherent, we mean that the rows ϕ_j of the matrix $\mathbf{\Phi}$ cannot sparsely represent the elements of the sparsity-inducing basis ψ_i, and vice versa. Thus, rather than measuring the N-point signal $\mathbf{x}$ directly, we acquire the $M \ll N$ linear projections which are then quantized such that $\mathbf{y} = Q(\mathbf{\Phi}\mathbf{x} + \mathbf{n})$. The effect of quantization may be modeled

1-4244-0921-7/07 $25.00 © 2007 IEEE.

as additive noise and thus we view the measurements as $\mathbf{y} = \mathbf{\Phi x} + \mathbf{n} = \mathbf{\Phi\Psi}\boldsymbol{\alpha} + \mathbf{n}$, where $\mathbf{n}$ represents the combination of the quantization effect and the noise inherent to the measurement process. For brevity, we define the $M \times N$ matrix $\mathbf{\Theta} = \mathbf{\Phi\Psi}$.

Since $M < N$, recovery of the signal $\mathbf{x}$ from the measurements $\mathbf{y}$ is ill-posed in general; however, the additional assumption of signal *compressibility* in the basis $\mathbf{\Psi}$ makes recovery both feasible and practical. The recovery of the set of transform coefficients $\boldsymbol{\alpha}$ can be achieved through *optimization* [4] by searching for the $\boldsymbol{\alpha}$ with the smallest ℓ_1 norm that agrees with the M observed measurements in $\mathbf{y}$. The margin of error is given by the magnitude of the noise $\epsilon \geq \|\mathbf{n}\|_2$:

$$\widehat{\boldsymbol{\alpha}} = \arg\min \|\boldsymbol{\alpha}\|_1 \quad \text{such that } \|\mathbf{y} - \mathbf{\Theta}\boldsymbol{\alpha}\|_2 \leq \epsilon \tag{1}$$

This optimization problem, also known as *Basis Pursuit with Denoising* (BPDN) [5] can be solved with traditional convex programming techniques whose computational complexities are polynomial in N. At the expense of slightly more measurements, iterative greedy algorithms like Orthogonal Matching Pursuit (OMP) [6] can also be applied to the recovery problem.

B. Real-Time CS

1) Analog signal: Suppose our analog signal has finite information rate $\mathcal{K}$ i.e., the signal can be represented using $\mathcal{K}$ parameters per unit time in some continuous basis. More concretely, let the analog signal $x(t)$ be composed of a *discrete, finite* number of weighted continuous basis or dictionary components

$$x(t) = \sum_{n=1}^{N} \alpha_n \, \psi_n(t), \tag{2}$$

with $t, \alpha_n \in \mathbb{R}$. In cases where there are a small number of nonzero entries in $\boldsymbol{\alpha}$, we may again say that the signal x is sparse. Although each of the dictionary elements ψ_n may have high bandwidth, the signal itself has few degrees of freedom.

2) Analog processing: Our signal acquisition system consists of three main components; demodulation, filtering, and uniform sampling. As seen in Figure 1, the signal is modulated by a psuedo-random maximal-length PN sequence of ± 1's. We call this the *chipping sequence* $p_c(t)$, and it must alternate between values at or faster than the Nyquist frequency of the input signal. The purpose of the demodulation is to spread the frequency content of the signal so that it is not destroyed by the second stage of the system, a low-pass filter with impulse response $h(t)$. Finally, the signal is sampled at rate $\mathcal{M}$ using a traditional ADC.

3) Analog system as a CS matrix: Although our system involves the sampling of continuous-time signals, the discrete measurement vector y can be characterized as a linear transformation of the discrete coefficient vector $\boldsymbol{\alpha}$. As in the discrete CS framework, we can express this transformation as an $M \times N$ matrix $\mathbf{\Theta}$ that combines two operators: $\mathbf{\Psi}$, which maps the discrete coefficient vector $\boldsymbol{\alpha}$ to an analog signal x, and $\mathbf{\Phi}$, which maps the analog signal x to the discrete set of measurements y.

To find the matrix $\mathbf{\Theta}$ we start by looking at the output $y[m]$, which is a result of convolution and demodulation followed by sampling at rate $\mathcal{M}$. Since our analog input signal (2) is composed of a finite and discrete number of components of $\mathbf{\Psi}$, we can write

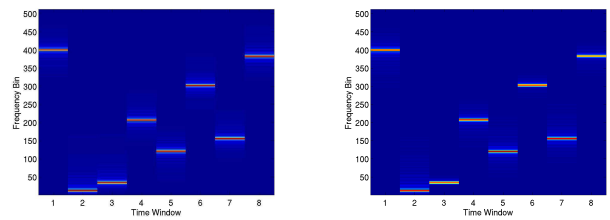

(a) (b)

Fig. 2. Comparison of Spectrograms obtained from full and CS compressed versions of a frequency hopping signal. The signal is a single side-band AM signal, whose carrier frequency changes periodically over time. (a) Spectrogram from original signal. (b) Spectrogram from CS reconstruction with measurement rate equal to 25% of Nyquist rate.

$$y[m] = \sum_{n=1}^{N} \alpha_n \int_{-\infty}^{\infty} \psi_n(\tau)\, p_c(\tau)\, h(m\mathcal{M} - \tau)\, d\tau. \tag{3}$$

It is now clear that we can separate out an expression for each element $\theta_{m,n} \in \mathbf{\Theta}$ for row m and column n

$$\theta_{m,n} = \int_{-\infty}^{\infty} \psi_n(\tau)\, p_c(\tau)\, h(m\mathcal{M} - \tau)\, d\tau. \tag{4}$$

C. Reconstruction for Analog Time-Frequency Sparse Signals

We consider the case of wideband signals that are time-frequency sparse in the sense that at each point in time they are well-approximated by a few local sinusoids of constant frequency. As a practical example, consider sampling a frequency-hopping communications signal that consists of a sequence of windowed sinusoids with frequencies distributed between f_1 and f_2 Hz. The bandwidth of this signal is $f_2 - f_1$ Hz, which dictates sampling above the Nyquist rate of $2(f_2 - f_1)$ Hz to avoid aliasing. We are interested in the case where $f_2 - f_1$ is very large and the signal is compressible, since the AIC will achieve much better performance than an ADC.

It is well known that signals that are localized in the time-frequency domain have near-sparse representation under the Gabor transform, which is defined as

$$\widehat{x}(\tau, f) = \langle x(t), \psi_{\tau,f}(t) \rangle,$$

i.e. the coefficient measures the inner product of the signal with the Gabor atoms

$$\psi_{\tau,f}(t) = g(t - \tau) e^{\pm j 2\pi f t}$$

where g is a window function with $\|g\|_2 = 1$ [7]. We will use a dictionary of Gabor atoms during the reconstruction of the signal to obtain a representation directly in the time-frequency domain, without performing reconstruction of the time signal. This representation is immediately useful because the Gabor transform is a uniform sampling of the coefficients of the signal under the short-time Fourier transform (STFT). Thus, a *spectrogram*, the conventional analysis tool for this class of signals, can be quickly generated by taking the squared magnitudes of the reconstructed coefficients.

An example is shown in Figure 2(a) where the spectrogram of a single sideband amplitude modulated (SSB-AM) frequency hopping signal is displayed. We see that for small ranges of time, the signal is well identified by its carrier frequency, but when

we consider the whole signal length there are many carriers to isolate. The spectrogram pictured in Figure 2(b) shows reconstruction of the signal from AIC measurements using a Gabor dictionary with a boxcar window. The signal was measured through the random demodulation system at 25% of the Nyquist rate from the original signal and then reconstructed using OMP. Since the SSB-AM signal is not represented with exactly K terms in the Gabor dictionary, we use the same constraint employed in BPDN(1) on an OMP solver to denoise the reconstruction.

III. Analog-to-Information System Performance

We aim to characterize the SNR of the AIC system using known analysis of CS performance. We present a theorem for K-sparse signals, which gives insight into the SNR behavior of the AIC system. The following definition is used in the theorem.

Definition 1: A $M \times N$ matrix $\mathbf{\Phi}$ has the K-Restricted Isometry Property (K-RIP) with constant δ_K if for all $\mathbf{x} \in \mathbb{R}^N$ with $\|\mathbf{x}\|_0 = K$,

$$(1-\delta_K)\|\mathbf{x}\|_2 \leq \|\mathbf{\Phi}\mathbf{x}\|_2 \leq (1+\delta_K)\|\mathbf{x}\|_2.$$

where the norm $\|\mathbf{x}\|_0$ is defined as the number of nonzero elements in $\mathbf{x}$.

Theorem 1: Let $\mathbf{x}$ be an K-sparse signal, i.e. $\|\mathbf{x}\|_0 = K$, and let $\mathbf{y} = \mathbf{\Phi}\mathbf{x}$ represent an AIC measurement setup, where we label reconstruction from the measurements $\mathbf{y}$ as $\mathbf{x}^\sharp$ with AIC reconstruction using BPDN. If Φ holds the K-Restricted Isometry Property (RIP) with constant δ_K and if $\delta_{3K}+3\delta_{4K} < 2$, then the SNR of the AIC system obeys the lower bound

$$\begin{aligned} \text{SNR}_{\text{AIC}} &= 20\log\left(\frac{\|\mathbf{x}\|_2}{\|\mathbf{x}^\sharp - \mathbf{x}\|_2}\right) \\ &\geq \text{SNR}_{\text{system}} - 20\log((1+\delta_K)C_{1,K}) \end{aligned}$$

where $\text{SNR}_{\text{system}}$ is the SNR of the sampling subsystem and $C_{1,K}$ is a constant depending only on K.

The condition on the RIP constants holds for random Gaussian matrices when the number of rows $M = O(K\log\frac{N}{K})$. Since Gaussian matrices are order optimal for RIP, the use of general matrices, including that posed in (4), incurs a loss in performance.. The theorem is proven in [8]. This bound on the performance decay will depend on the compressibility of the signal and the class of matrix $\mathbf{\Phi}$ applied. As an example, if a Gaussian random matrix $\mathbf{\Phi}$ is used with a large enough row-to-column ratio, and the signal has a sparsity $K = N/10$, the loss in performance is approximately 23dB. If the increase in SNR due to the reduction on the sampling rate is larger than the performance decay, the AIC performs better than an ADC.

IV. Transistor-Level Implementation

As a proof of concept, we have designed a transistor-level circuit implementation of the AIC described in the previous section. The design of each component is explained except for the back-end ADC which is assumed to be an off-the-shelf ADC capable of sampling at certain desired rates. In the hardware implementation of the AIC, the chipping sequence $p_c(t)$ and the mixer have the most demanding design specifications. Both components are required to operate at the fastest speeds in the system, and so they must be precise at such speeds. To make this possible, the two components are custom designed to efficiently operate together.

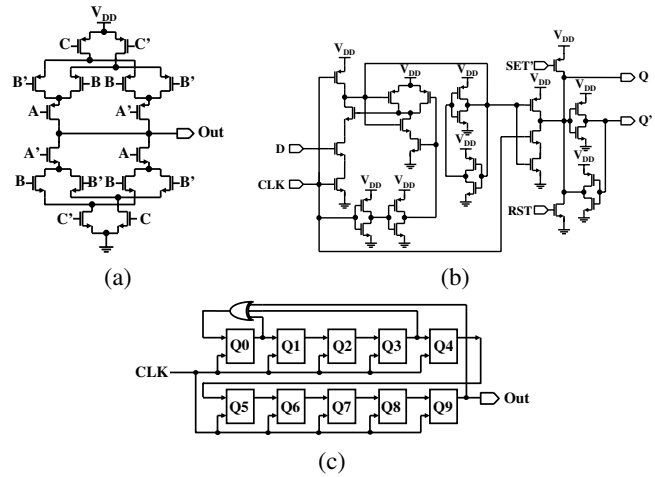

Fig. 3. (a) XOR compact circuit design. (b) The transistor level design for Semi Dynamic Flip Flop (SDFF). (c) The block diagram for 10 bit LFSR.

The chipping sequence is implemented with a maximal length - linear feedback register (ML-LFSR), for which the internal components are shown in Figure 3. This component produces values of -1 and 1 in a pseudo-random sequence. The period of the sequence is $2^n - 1$, with n being the number of registers in the LFSR, thus the sequence length can be tweaked to application requirements.

The primary bottleneck in the ML-LFSR is the delay flip flop (DFF) and so to provide the highest speed possible, we use a semi-dynamic flip flop (SDFF) [9]. Furthermore, in simulations we use the BSIM3v3 device models of a commercial 0.13μm CMOS technology. Other optimizations include using minimum gate length to reduce gate capacitance, which both reduces power dissipation and maximizes speed, and also using a compact and fast balanced XOR circuit design. This design is shown functioning properly at a 2 GHz clock frequency in the HSPICE simulation presented in Figure 5(b).

When operating at high rates, the most common artifact of the mixer, which computes $x(t) \times p_c(t)$, is non-linearity of the output. Passive mixers, such as diode and passive field effect transistor (FET) mixers, have good linearity, good noise performance, and can operate at frequencies up to 5GHz [10]. However, active mixers are preferred for low-power integrated circuits since they provide conversion gain, require less power at the local input port, and have a broader design space. Such mixers are generally based on the classic Gilbert cell [11]. Several modifications can be made to this circuit to increase its frequency response, decrease its nonlinearity, and limit its noise level, such as the double balanced Gilbert cell design, which is characterized by its rejection of common mode signals and its ability to operate as a four quadrant multiplier. Our AIC implementation uses a highly linear mixer operating at 2GHz using standard CMOS 0.13μm fabrication technology. The schematic diagram of the transistor-level implementation of this mixer is shown in Figure 4(a). For the AIC low pass filter, we utilize an integrator design. More specifically, a differential-input differential-output

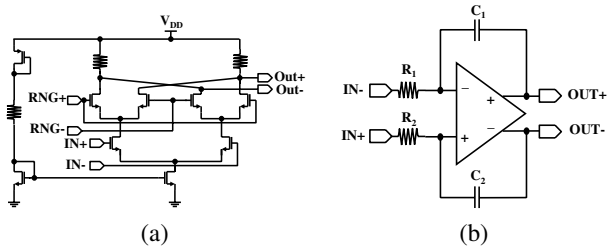

Fig. 4. (a) Modulator circuit. (b) Integrator circuit.

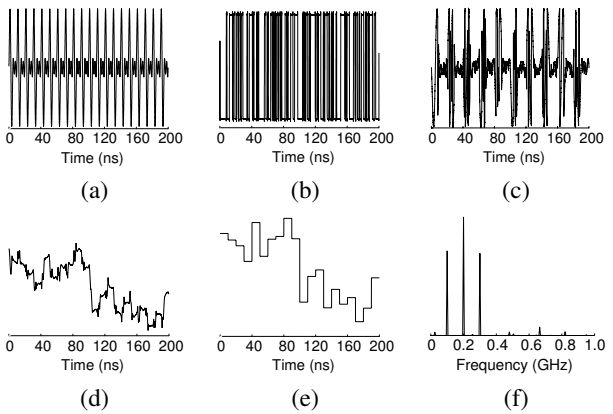

Fig. 5. HSPICE simulation results for: (a)The input analog compressible signal (b)The pseudo-random square wave generated from the random number generator (c)The multiplication result of both the analog signal and the pseudo-random square wave (d)The integration result of the randomized version from the analog signal (e)The result of the quantization (f)The reconstructed signal in DFT domain.

RC-active integrator is used and is shown in Figure 4(b). Such integrators have a finite time-constant due to limited amplifier gain, and therefore act as low pass filters with the time-constant determining cutoff frequency. In simulations, the cutoff frequency of the amplifier was adjusted at 100Hz with a linear slope of 20dB/decade that extends to 20GHz to cover the entire range of the modulated signal.

V. Results and Discussions

To measure the quality of our design, the optimized transistor-level implementation for each AIC component is simulated with HSPICE. The end-to-end results are shown in Figure 5, and demonstrate a successful recovery of the original input signal at 6x sub-Nyquist sampling rate.

More specifically, Figure 5(a) shows the input AM signal which is composed of a 100 MHz signal modulated with a 200 MHz carrier. Figure 5(b) presents the output of the ML-LFSR running at 2GHz. The modulator output in Figure 5(c) is the result of the multiplication of the input AM wave by the pseudo-random pattern. As explained earlier, we see in Figure 5(d) that the integrator acts as a low pass filter, smoothing the output of the modulator. The signal is then sampled at sub-Nyquist rate by the back-end ADC, which is shown in Figure 5(e). Finally it should be noted that in this example, the sampling frequency is 100 MSample/s, a reduction of one-sixth of the minimum Nyquist required rate for this signal with a conventional ADC. The spectrum of the reconstructed AM signal is shown in Figure 5(f).

Some artifacts appear in the spectrum in Figure 5(f) because the SNR of the CS reconstruction is sensitive to the non-ideal behaviors present in a circuit implementation. The most significant sources of non-idealities are: the clock jitter of the random number generator, the linearity and intermodulation distortion of the mixer, and the quantization error of the back-end ADC. The sensitivity of reconstruction is partially due to the fact that these behaviors reduce the feasibility of the system in the CS sense, but also because during reconstruction we cannot produce a matrix which is *exactly* tuned to the non-idealities, some of which which are stochastic in nature.

VI. Conclusion

A new design for analog-to-information converters (AIC) has been presented and with it, an expanded theory for compressive sensing on analog signals, bounds on performance, and practical methods for recovery of time-frequency signals. Furthermore, we have demonstrated an end-to-end transistor-level implementation of the AIC, which accurately recovers when sampled at a $6\times$ sub-Nyquist rate.[1]

References

[1] 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 *Proc. of the IEEE Dallas Circuits and Systems Workshop (DCAS)*, 2006.

[2] E. J. Candès and J. Romberg, "Quantitative robust uncertainty principles and optimally sparse decompositions," *Foundations of Computational Mathematics*, vol. 6, no. 2, pp. 227–254, April 2006.

[3] D. Donoho, "Compressed sensing," *IEEE Trans. Info. Theory*, vol. 52, no. 4, pp. 1289–1306, September 2006.

[4] E. Candès, J. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Communications on Pure and Applied Mathematics*, vol. 59, no. 8, pp. 1207–1223, Aug. 2006.

[5] S. Chen, D. Donoho, and M. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. on Sci. Comp.*, vol. 20, no. 1, pp. 33–61, 1998.

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

[7] S. Mallat, *A Wavelet Tour of Signal Processing*, 2nd ed. San Diego: Academic Press, 1999.

[8] M. F. Duarte, J. Laska, and R. G. Baraniuk, "Theoretical bounds for signal-to-noise ratio in analog to information conversion systems," Rice University ECE Department, Houston, TX, Tech. Rep. TREE-0608, Sept. 2006.

[9] V. Stojanovic and V. G. Oklobdzija, "Comparative Analysis of Master-Slave Latches and Flip-Flops for High-Performance and Low-Power Systems," *IEEE JOURNAL OF SOLID-STATE CIRCUITS*, 1999.

[10] S. Zhou and M. Chang, "A CMOS passive mixer with low flicker noise for low-power direct-conversion receiver," *IEEE JOURNAL OF SOLID-STATE CIRCUITS*, pp. 1084 – 1093, May 2005.

[11] M. Tsai and H. Wang, "A 0.3-25-GHz Ultra-Wideband Mixer Using Commercial 0.18 um CMOS Technology," *IEEE Microwave and Wireless Components Letters*, pp. 522 – 524, 2004.

[1]This work was supported by the grants DARPA/ONR N66001-06-1-2011 and N00014-06-1-0610, NSF CCF-0431150, ONR N00014-06-1-0769 and N00014-06-1-0829, AFOSR FA9550-04-1-0148, NSF Career 0448558, and the Texas Instruments Leadership University Program.