Generic Sensing Hardware and Real-Time Reconstruction for Structured Analog Signals

Moshe Mishali, Rolf Hilgendorf, Eli Shoshan, Ina Rivkin and Yonina C. Eldar
Department of Electrical Engineering
Technion—Israel Institute of Technology
Haifa, Israel 32000
Email: {moshiko@tx, rhilgen@ee, elis@ee, inna@ee, yonina@ee}.technion.ac.il

***Abstract*—Generic acquisition hardware promotes a unified treatment of various signal classes. In modern applications involving wide input bandwidths, uniform sampling, the common practice for generic digitization, leads to prohibitively large sampling and processing rates due to the wide Nyquist bandwidth of the input. In this paper, we present the X-ADC system which narrows down the input bandwidth by analog preprocessing prior to sampling at rates substantially lower than Nyquist. As we show, the X-ADC strategy is generic in the sense that multitude radio and medical imaging applications with structured signals can utilize the same architecture for lowrate acquisition. A recently published sub-Nyquist hardware, which was designed according to the proposed X-ADC scheme, provides a concrete reference for the present study. We complement the generic acquisition by reporting on a real-time embedded implementation of a sub-Nyquist reconstruction algorithm. The embedded design enables, for example, fast spectrum sensing which is essential to real-time cognitive radio applications.**

I. Introduction

Signal processing involves hardware and software design. The front-end amplifies and preprocesses the continuous input before analog to digital conversion (ADC) takes place. Digital signal processing (DSP) algorithms manipulate the incoming stream of numbers to achieve a desired effect, *e.g.,* denoising, source separation, estimation etc. The modern trend is to shift as much processing operations as possible from analog to digital. In the trend of digitization, the keys for a successful processing system are two: a generic hardware platform, so that a single hardware front-end supports a diverse range of applications, and real-time DSP algorithms with computationally-light digital complexities.

Compressed sensing (CS) is an emerging paradigm in signal processing, named after the works of Donoho [2] and Candès, Romberg and Tao [3]. The goal in CS is to sample signals at rates lower than what is traditionally required by the Shannon-Nyquist theorem. To accomplish this goal, the underlying structure of the input, usually sparsity in some transform domain, is exploited. Mainstream works in CS study measurement systems in discrete and finite settings. More specifically, the focus is on recovery of vectors with only a few nonzero entries from an underdetermined set of linear equations. Several frameworks have been proposed for extending CS to analog signals, including Xampling [4]–[6], finite rate of innovations (FRI) [7] and random demodulation (RD) [8].

In this paper, we propose a generic sensing architecture which can reduce the sampling rate of structured analog signal classes. The scheme, termed X-ADC, is based on mixing the input with a set of periodic waveforms prior to sampling. The approach is based on the modulated wideband converter (MWC) which is presented hereafter as one application of the X-ADC scheme. In Section II, we present the X-ADC structure in more detail. The X prefix hints at the reduced sampling rate. We also describe a recent board-level prototype of the MWC [1], which serves as a reference X-ADC circuit.

The first contribution of this work, Section III, describes how to employ the X-ADC scheme for reduced-rate sampling in multitude radio and medical imaging applications involving structured inputs, including sampling of multiband communications [4], spectrum sensing for cognitive radio (CR) [9], [10], ultrasonic imaging [11], [12] and RD [8]. The fact that only slight hardware modifications are required, on the level of assembling or disassembling components in the existing layout of [1], supports the generality of X-ADC. Once the samples are obtained by the generic system, application-dependent algorithms use different processing methods in order to extract the relevant information.

Our second contribution, presented in Section IV, is a real-time embedded implementation of the MWC recovery algorithms. The real-time implementation enables accurate frequency support detection in a time duration as low as several micro-seconds. Real-time support detection can, for example, alleviate one of the main bottlenecks in widespread deployment of CR technology [10].

II. X-ADC Hardware

In this section, we describe the X-ADC architecture, which reduces the input bandwidth so that it can be digitized using existing lowrate commercial ADCs, while preserving the information of interest for subsequent digital algorithms.

A. Architecture

A block diagram of an X-ADC is depicted in Fig. 1(a). The continuous input $x(t)$ is passed to m analog preprocessing units. Each unit multiplies the signal by a pre-defined periodic waveform $p_i(t)$ and applies a filter $H_i(f)$ on the product $\tilde{x}_i(t)$. The purpose of $H_i(f)$ is to shape the frequency response of the product and in particular to narrow down the bandwidth. Finally, the filtered output $y_i(t)$ is sampled every T_s seconds, resulting in a digital stream of numbers $y_i[n]$. The reduced bandwidth of $y_i(t)$ permits the use of commercial ADC devices with low sampling rate and low analog input bandwidth. This is in contrast with the premium devices needed to acquire $x(t)$ at its high Nyquist rate.

The second ingredient of X-ADC is a periodic generator unit that provides the waveforms $p_i(t)$. Since $p_i(t)$ is periodic, it has a Fourier-series expansion

$$p_i(t) = \sum_{\ell=-\infty}^{\infty} c_{i\ell}\, e^{j2\pi\ell t/T_p}, \tag{1}$$

where T_p is the period duration and $c_{i\ell}$ are a set of coefficients. The consequence of periodic mixing is deliberate aliasing of the spectrum of $x(t)$, so that the product $\tilde{x}_i(t)$ contains shifted copies of the input spectrum $X(f)$ at equally-spaced frequency spacing

$$\tilde{X}_i(f) = \sum_{\ell} c_{i\ell} X(f - \ell/T_p). \tag{2}$$

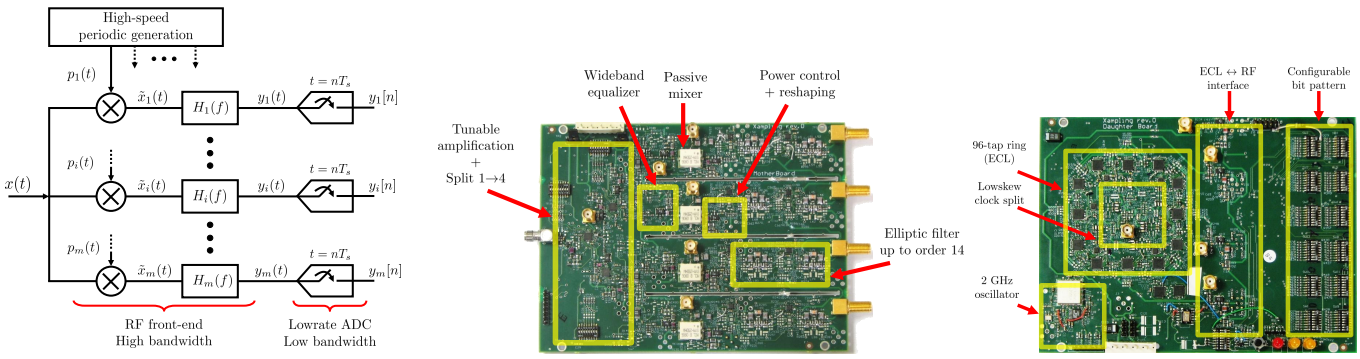

(a) (b) (c)

Fig. 1. Generic sensing architecture using periodic mixing (a). A hardware prototype with (b) four-channel analog front-end and a 2 GHz periodic generator (c) was reported in [1].

B. Circuit work

A hardware design that complies with the X-ADC scheme was reported in [1]. Fig. 1 depict photos of two hardware boards. One performs the RF mixing and filtering stage while the other provides four high-speed periodic functions $p_i(t)$. The markers in panels (b) and (c) of Fig. 1 highlight the interesting blocks in our design.

The RF board splits $x(t)$ into $m = 4$ analog processing branches. Tunable amplification gains along the path ensure a signal to noise-and-distortion ratio of 15 dB at the outputs $y_i(t)$, which was verified experimentally over a dynamic range of 50 dB input power [1]. In contrast to standard RF mixing with a single sinusoid, X-ADC requires periodic mixing, that is simultaneous multiplication with the multiple sinusoids (1) comprising $p_i(t)$. To support the nonordinary mixing additional circuitries were inserted: wideband equalization, passive mixing and adjustable power control with frequency reshaping option [1]. An elliptic filter with up to 14 stages allows a flexible choice for $H_i(f)$ in a wide range of pass and stop band combinations.

The periodic generator produces four sign alternating waveforms

$$p_i(t) = \alpha_{ik}, \quad k\frac{T_p}{M} \le t \le (k+1)\frac{T_p}{M}, \quad 0 \le k \le M-1, \qquad (3)$$

with programmable signs $\alpha_{ik} \in \{+1, -1\}$, derived from four taps of a single shift-register (SR) of length $M = 108$ and clock rate of 2 GHz. Other choices for $p_i(t)$ are possible, since in principle only the periodicity is a design constraint.

The X-ADC is a modular platform; parallelizing two four-channel X-ADC boards results in a structure with $m = 8$ branches and so forth. In particular, with sufficiently large m, the technology can scale up to the Nyquist rate.

III. Generic Sensing

We now explain how certain choices of $p_i(t)$ in conjunction with specific filtering $H_i(f)$ exploit the aliasing phenomenon to achieve innovative radio and medical applications, thereby supporting the generality of the method.

A. Sub-Nyquist sampling of multiband signals

The X-ADC concept of mixing with periodic waveforms is based on the MWC system [4]. The MWC application enables the design of a communication receiver which intercepts N narrowband transmissions, but is not provided with knowledge of their carrier frequencies. In this scenario, standard RF demodulation cannot be used. The X-ADC setup in this case is [4]

$$m \ge 2N, \quad \frac{1}{T_p} \ge B, \quad T_s = T_p, \qquad (4)$$

with identical lowpass filters $H_i(f) = H(f)$ having $1/2T_s$ cutoffs, where B is the maximal expected bandwidth of an individual transmission. The X-ADC aliases equally-spaced spectrum slices, each of width $1/T_p$ to the origin, as illustrated in Fig. 2. It follows from (4) that at most $2N$ spectrum slices intersect with the signal frequency support. The X-ADC circuit of Figs. 1(b)-(c) also supports a more advanced configuration of the MWC in which the number of channels is collapsed by a factor of q by taking $T_s = T_p/q$ and extending the filter cutoffs accordingly. In the prototype, $q = 3$ is used to provide effectively $mq = 12$ sampling branches from only $m = 4$ physical channels. This setup enables recovery of multiband inputs with $N = 6$ bands, or equivalently $N/2 = 3$ narrowband transmissions. Real-time reconstruction of the multiband input $x(t)$ from the lowrate sequences $y_i[n]$ is discussed in the next section.

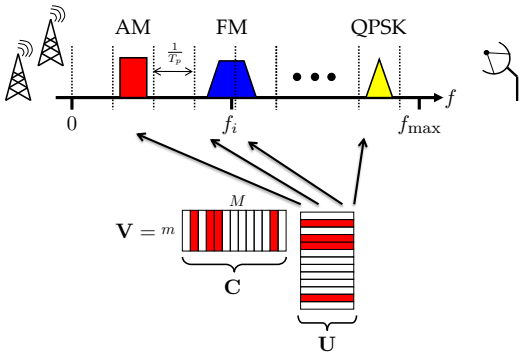

Fig. 2. Several narrowband transmissions result in an input signal with multiband frequency support.

B. Spectrum sensing for cognitive radios

A CR device exploits temporarily unused spectrum intervals, or "holes", for transmission on an opportunistic basis, as visualized in Fig. 3(a). Real-time spectrum sensing of the available holes is therefore essential for this radio technology. Clearly, the X-ADC platform together with the MWC reconstruction algorithm can provide spectrum sensing functionality. Nonetheless, in the CR setting, the hardware can be modified to better optimize the design for holes detection rather than multiband reconstruction. One option is setting up the system with $1/T_p < B$. The aliasing spacing becomes small and consequently the resolution of support/hole detection improves. For multiband sampling the setting $1/T_p < B$ results in splitting

the energy of an individual transmission between several adjacent spectrum slices, thereby increasing the reconstruction complexity. Since CR needs no reconstruction functionality, a large period T_p can be used by concatenating more registers in the chain of the periodic generator. Alternatively, a linear feedback shift-register (LFSR) can be constructed with XOR gates. For example, the LFSR in Fig. 3(b) outputs a periodic sign pattern with length $2^9 - 1 = 511$. The frequency resolution is five times higher than that of the $M = 108$ standard SR. The number of registers is reduced by a factor of 12. A theoretical study [13] indicated that popular LFSR sequences are suitable choices for support detection.

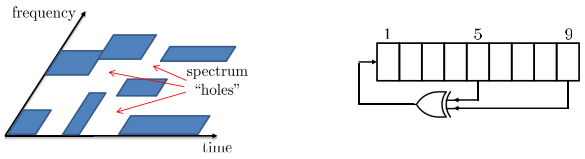

(a) (b)

Fig. 3. A map of available spectrum holes (a). The shaded areas represent frequency intervals occupied by licensed users. An optional maximal-length LFSR (b) for spectrum sensing using the X-ADC platform.

C. Ultrasonic imaging

An ultrasonic imaging system transmits an acoustic wave $g(t)$ towards a scanned tissue, which reflects several attenuated and delayed echoes

$$x(t) = \sum_{\ell=1}^{L} a_\ell g(t - t_\ell), \quad t \in [0, T], \tag{5}$$

where T is the duration of a single transmission-probing cycle, and L is the maximal expected number of echoes. The unknowns are the time delays t_ℓ and amplitudes a_ℓ. Fig. 4(a) illustrates a typical reflection pattern from a single acoustic path.

Typically, $g(t)$ has short time duration and therefore large Nyquist range. Recently, [11], [12] developed a stable and low complexity algorithm for recovery of an ultrasonic pulse stream, such as (5), from about $2L$ samples per interval T, effectively removing the dependency of the sampling rate on the bandwidth of $g(t)$. The approach relies on the fact that $x(t)$ is confined to a finite duration T. Therefore, in principle, $2L$ coefficients from its Fourier expansion

$$m_i = \frac{1}{T} \int_0^T x(t)\, e^{j2\pi k_i t/T} \mathrm{d}t, \tag{6}$$

with $\{k_i\}$ denoting a set of $2L$ consecutive integers, determine the unknowns t_ℓ, a_ℓ [11], [12]. Instead of generating $2L$ frequency sources to directly measure m_i, the generic X-ADC platform can be utilized to obtain a linear mixture of m_i by proper filtering of $p_i(t)$, which zeros out all but $2L$ Fourier coefficients $c_{i\ell}$. The reshaping circuit that is assembled on the local-oscillator port in the RF front-end provides this feature; see Fig. 1(b). Mixing with filtered periodic functions boils down to a linear relation $\mathbf{v} = \mathbf{C}\mathbf{m}$, with a square invertible matrix $\mathbf{C}$ containing only those coefficients $c_{i\ell}$ that passed the filtering. Inverting $\mathbf{C}$ returns the measurement vector $\mathbf{m} = \{m_1, \cdots, m_{2L}\}$, and the reconstruction flow of [11], [12] can proceed accordingly.

D. Random demodulation

RD is a strategy for reduced-rate sampling of multitone signals, consisting of a finite number of harmonic sinusoids. A single channel

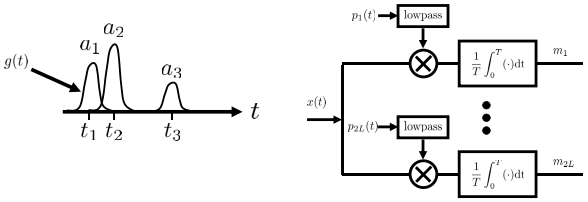

(a) (b)

Fig. 4. A typical ultrasonic signal (a) consists of reflected echoes of the known transmitted pulse shape $g(t)$. Lowrate measurements are realized in (b) by an X-ADC device that filters the periodic waveforms $p_i(t)$ prior to mixing.

mixes the input with a pseudo-random bit sequence (PRBS), integrates the product and dumps the output at a low rate. The PRBS needs to be sufficiently long to approximate a true random sequence. The X-ADC prototype can provide PRBS lengths up to $2^r - 1$, $r \leq 108$. The integrator is a first-order filter, thus requires assembling only a degenerated version of the elliptic lowpass filter. We point out that while realizing the RD method is possible using the circuit of Figs. 1(b)-(c), the method is computationally limited to signals with low Nyquist rate due to the complexity of the recovery algorithm which scales with the Nyquist rate of the input [5]. The approach also requires accurate time-domain properties, *e.g.,* sign transitions sharply aligned to Nyquist intervals and a rectangular integrator response, which are difficult to obtain at high rates. In contrast, the previous applications (MWC, cognitive spectrum sensing and imaging) require only the periodicity of $p_i(t)$, which can be achieved at high rates. Further details and comparisons appear in [5].

IV. Real-time Reconstruction

Reducing the sampling rate is one important aspect of analog sensing. The computational complexity in the digital domain is another important factor in practice. In this section, we describe a real-time fixed-point implementation of the MWC reconstruction algorithm [4] on an Altera Stratix III field-programmable gate array (FPGA) mounted on a Gidel PROCStar-III development board. Figure 5 depicts a high-level block diagram of the real-time design. Clock rates of the various modules are marked.

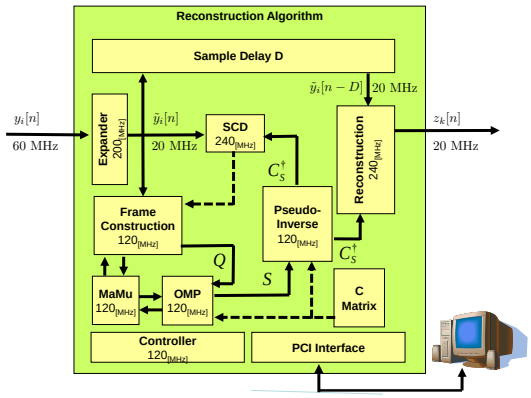

Fig. 5. Block diagram of a real-time design for sub-Nyquist reconstruction.

A. Frame construction

As described in [4], the reconstruction procedure begins by constructing a frame (or a basis) of the measurements and decomposition as follows

$$\mathbf{y}[n] \quad \rightarrow \quad \mathbf{Q} = \sum_n \mathbf{y}[n]\mathbf{y}^H[n] \quad \rightarrow \quad \mathbf{Q} = \mathbf{V}\mathbf{V}^H, \qquad (7)$$

where $\mathbf{y}[n] = [y_1[n], \ldots, y_m[n]]^T$ and $\mathbf{V}^H$ is the conjugate transpose of $\mathbf{V}$. The decomposition step to $\mathbf{V}$ enables removal of the noise space, but demands a substantial amount of multiplications and FPGA resources. In the FPGA design, we decided to avoid the decomposition, and instead increased the number N_{samples} of $\mathbf{y}[n]$ snapshots that are aggregated to $\mathbf{Q}$. Based on recovery results of Monte-Carlo simulations, Fig. 6, we chose $N_{\text{samples}} = 70$.

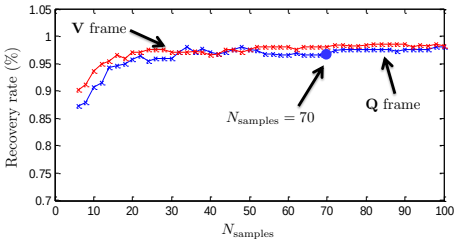

Fig. 6. Successful support recovery vs. number of samples used for constructing the frame matrix $\mathbf{Q}$.

B. Support detection

Given the frame $\mathbf{Q}$ (or $\mathbf{V}$) an underdetermined system

$$\mathbf{Q} = \mathbf{C}\mathbf{U}, \qquad (8)$$

is solved for $\mathbf{U}$ with minimal number of nonidentically zero rows, where $\mathbf{C}$ containing the Fourier coefficients $c_{i\ell}$. The sparsest $\mathbf{U}$ is proved in [4] to indicate the support of $x(t)$ at $1/T_p$ frequency resolution, as visualized in Fig. 2.

We chose the multi orthogonal matching pursuit (M-OMP) [14] algorithm for solving (8), since it involves simple algebraic computations. In M-OMP, every iteration reveals a new nonzero row of $\mathbf{U}$ and a corresponding column of $\mathbf{C}$. A pseudo-inverse matrix $\mathbf{C}_S^\dagger$, computed for the column subset $\mathbf{C}_S$ that is indicated by the current nonzero index set S, updates a residual $\tilde{\mathbf{Q}}$ for the next iteration. To reduce computations, we used a Gram-Schmidt procedure proposed in [14] to update $\mathbf{C}_S^\dagger$ sequentially. In addition, taking into account that since $x(t)$ is real-valued, so that S is symmetric, we modified the M-OMP matching rule to find two symmetric indices at each iteration. Consequently, S is recovered in N, rather than $2N$, cycles.

C. Parallel execution and gate-level simulation

The major advantage of FPGA programming is the ability to parallelize computations in time. We constructed a matrix multiplication unit (MaMu) of 144 complex multipliers, which is utilized for both frame construction and M-OMP computations. Once the index set S is detected, the contents $z_k[n]$ of the relevant spectral intervals are reconstructed at real-time by applying $\mathbf{C}_S^\dagger$ on the incoming measurements $\mathbf{y}[n]$.

To conclude, we present a result of gate-level simulation affirming real-time performance. Three signal generators were combined to the X-ADC input terminal: amplitude-modulated (AM) with carrier 871 MHz, frequency-modulated (FM) with carrier 631 MHz and pure sine at 981 MHz. Frame construction and the modified M-OMP algorithm

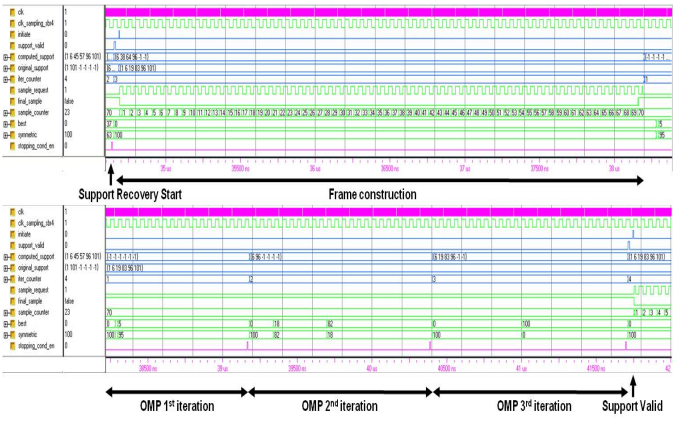

Fig. 7. Real-time support recovery.

detected the correct active slices. The entire procedure lasted less than 7μsecs.

ACKNOWLEDGEMENTS

The authors would like to thank Daniel Barsky, Natalie Pistunovich, Yoni Smolin, Daniel Primor, Omer Kiselov, Amir Bishara, Morad Awad, Tzvika Shirazi, Eli Sorin, Dima Kichin and Oleg Greenberg from the HS-DSL laboratory at the Technion for implementing and simulating the real-time recovery algorithm.

REFERENCES

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

[2] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, April 2006.

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

[4] M. Mishali and Y. C. Eldar, "From theory to practice: Sub-Nyquist sampling of sparse wideband analog signals," *IEEE J. Sel. Topics Signal Process.*, vol. 4, no. 2, pp. 375–391, Apr. 2010.

[5] M. Mishali, Y. C. Eldar, and A. Elron, "Xampling: Signal acquisition and processing in union of subspaces," *CCIT Report no. 747, EE Dept., Technion; [Online] arXiv.org 0911.0519*, Oct. 2009.

[6] Y. C. Eldar, "Compressed sensing of analog signals in shift-invariant spaces," *IEEE Trans. Signal Process.*, vol. 57, no. 8, pp. 2986–2997, Aug. 2009.

[7] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE Trans. Signal Process.*, vol. 50, no. 6, pp. 1417–1428, 2002.

[8] J. A. Tropp, J. N. Laska, M. F. Duarte, J. K. Romberg, and R. G. Baraniuk, "Beyond Nyquist: Efficient sampling of sparse bandlimited signals," *IEEE Trans. Inf. Theory*, vol. 56, pp. 520–544, Jan. 2010.

[9] Mitola III, J., "Cognitive radio for flexible mobile multimedia communications," *Mobile Networks and Applications*, vol. 6, pp. 435–441, 2001.

[10] M. Mishali and Y. C. Eldar, "Wideband spectrum sensing at sub-Nyquist rates," *to appear in IEEE Signal Process. Mag.; [Online] arXiv.org 1009.1305*, Sep. 2010.

[11] R. Tur, Y. C. Eldar, and Z. Friedman, "Innovation rate sampling of pulse streams with application to ultrasound imaging," *to appear in IEEE Trans. Signal Process.; [Online] arXiv.org 1003.2822*, Mar. 2010.

[12] K. Gedalyahu, R. Tur, and Y. C. Eldar, "Multichannel sampling of pulse streams at the rate of innovation," *to appear in IEEE Trans. Signal Process.; [Online] arXiv.org 1004.5070*, Apr. 2010.

[13] M. Mishali and Y. C. Eldar, "Expected-RIP: Conditioning of the modulated wideband converter," in *Information Theory Workshop, 2009. ITW 2009. IEEE*, Oct. 2009, pp. 343–347.

[14] S. F. Cotter, B. D. Rao, K. Engan, and K. Kreutz-Delgado, "Sparse solutions to linear inverse problems with multiple measurement vectors," *IEEE Trans. Signal Process.*, vol. 53, no. 7, pp. 2477–2488, July 2005.