

VLSI Design of a Monolithic Compressive-Sensing Wideband Analog-to-Information Converter

David Bellasi, *Student Member, IEEE*, Luca Bettini, *Student Member, IEEE*, Christian Benkeser, *Member, IEEE*, Thomas Burger, *Member, IEEE*, Qiuting Huang, *Fellow, IEEE*, and Christoph Studer, *Member, IEEE*

Abstract—One of the key tasks in cognitive radio and communications intelligence is to detect active bands in the radio-frequency (RF) spectrum. In order to perform spectral activity detection in wideband RF signals, expensive and energy-inefficient high-rate analog-to-digital converters (ADCs) in combination with sophisticated digital detection circuitry are typically used. In many practical situations, however, the RF spectrum is sparsely populated, i.e., only a few frequency bands are active at a time. This property enables the design of so-called *analog-to-information* (A2I) converters, which are capable of acquiring and directly extracting the spectral activity information at low cost and low power by means of compressive sensing (CS).

In this paper, we present the first VLSI design of a monolithic wideband CS-based A2I converter that includes a signal acquisition stage capable of acquiring RF signals having large bandwidths and a high-throughput spectral activity detection unit. Low-cost wideband signal acquisition is obtained via CS-based randomized temporal subsampling in combination with a 4-bit flash ADC. High-throughput spectrum activity detection from the coarsely quantized and compressive measurements is achieved by means of a massively-parallel VLSI design of a novel accelerated sparse spectrum dequantization (ASSD) algorithm. The resulting monolithic A2I converter is designed in 28 nm CMOS, acquires RF signals up to 6 GS/s, and the on-chip ASSD unit detects the active RF bands at a rate $30\times$ below real-time.

Index Terms—Analog-to-information (A2I) conversion, cognitive radio, compressive sensing, flash analog-to-digital converter (ADC), randomized subsampling, sparse signal dequantization, wideband spectrum sensing, very-large-scale integration (VLSI).

I. INTRODUCTION

SENSING the active frequency bands in the radio frequency (RF) spectrum finds use in a large number of practical applications. Cognitive radio [1], for example, aims at sensing (or detecting) unused frequency bands in order to re-use them opportunistically with the goal of improving the spectral utilization. Since bandwidth is a scarce and expensive resource, spectrum sensing is believed to play an important role in meeting the ever-growing demand for higher data rates in future

wireless communication systems [2]. Indeed, IEEE 802.22 [3] envisions to re-use unoccupied frequency bands in the television spectrum for private wireless networks. Communications (or signals) intelligence is another application that aims at sensing the active frequency bands in order to detect, identify, and localize certain wireless transmitters, such as radar systems or (narrow-band) communication transceivers [4]. In both of these applications, one typically requires expensive and often energy-inefficient integrated sensing circuits that are capable of acquiring and detecting active frequency bands over very large portions of the RF spectrum.

While conventional wideband analog-to-digital converters (ADCs) provide a straightforward solution for acquiring wideband signals in the GS/s regime, they are typically energy-inefficient and expensive [5]. These drawbacks prohibit their deployment in battery-powered devices. Furthermore, sampling of signals at Nyquist frequency additionally results in excessive data rates (on the order of tens of Gb/s) and requires cost-effective ways of processing the acquired data at very high throughput. Consequently, practical solutions for spectrum sensing applications in need of acquiring large bandwidths, while being able to extract information about the spectral occupancy, demand sophisticated integrated circuit solutions featuring low silicon complexity and low power consumption.

A. Analog-to-Information (A2I) Converters

In recent years, a considerable number of spectrum occupancy surveys observed an under-utilization of the available spectrum; in fact, only a small number of (typically narrow) spectrum bands are heavily used, while the utilization of the remaining spectrum is only a few percent in many practical scenarios [6], [7]. Hence, sampling several GHz of the RF spectrum at Nyquist rate, while only a few frequency bands are active at a given time, seems to be an ineffective way of extracting the low-dimensional spectral occupancy information.

Compressive sensing [8] is a recently introduced sampling paradigm that enables one to acquire sparse signals (i.e., signals having only a few non-zero coefficients in a specific transform domain) at sub-Nyquist rates, while enabling their stable reconstruction via sophisticated sparse signal recovery algorithms. As a result, CS allows for the development of so-called *analog-to-information* converters, which sample sparse signals using inexpensive analog circuitry consuming only little power, while sophisticated algorithms extract the information of interest, such as the spectral activity [9], [10].

Due to the high computational complexity associated with sparse signal recovery from compressive measurements, virtu-

Manuscript received May 31, 2013; revised Aug. 30, 2013; final version Sep. 30, 2013.

D. Bellasi, L. Bettini, T. Burger, and Q. Huang are with the Dept. of Information Technology and Electrical Engineering (D-ITET), ETH Zürich, 8092 Zürich, Switzerland (e-mail: {bellasi, bettini, burger, huang}@iis.ee.ethz.ch).

C. Benkeser is with D-ITET, ETH Zürich and is now with RUAG Space, 8052 Zürich, Switzerland (email: christian.benkeser@ruag.com).

C. Studer is with the Dept. of Electrical and Computer Engineering, Rice University, Houston, 77004 TX, USA (e-mail: studer@rice.edu).

The authors would like to acknowledge the help of B. Sporrer during the analog design. C. Studer also thanks R. G. Baraniuk, Mr. Lan, J. N. Laska, and D. E. Waters for fruitful discussions on sparse signal dequantization. We furthermore thank the anonymous reviewers for their valuable comments which helped us to improve the exposition of our results.

ally all existing CS-based A2I designs perform signal recovery off-line on CPUs, GPUs, or DSPs [11]–[13]. However, off-line processing results in excessive I/O data-rates and storage requirements. More importantly, it prohibits timely (or real-time) decisions based on the recovered information and prevents the use of adaptive sensing strategies. In contrast, high throughput on-chip sparse signal recovery has the potential to avoid all these drawbacks, but necessitates efficient algorithms and corresponding high-performance digital very-large scale integration (VLSI) circuits (see, e.g., [14] for more details).

B. Contributions

In this paper, we present—to the best of our knowledge—the first *monolithic* A2I converter design for wideband spectrum sensing. In particular, we detail the design of a single-chip mixed-signal VLSI circuit in 28 nm CMOS that includes a low-complexity, energy-efficient sub-Nyquist ADC and a high-performance digital spectrum recovery stage capable of detecting the active frequency components at high rates. Our main contributions are summarized as follows:

1) We introduce a novel A2I conversion framework for RF spectrum sensing. In addition to leveraging CS via randomized sub-Nyquist sampling [8], we acquire coarsely quantized measurements, inspired by recent results in 1-bit CS [15], [16]. This combined approach of randomized subsampling and coarse quantization substantially reduces the complexity of the ADC stage (compared to conventional sub-Nyquist sampling) and leads to a substantial reduction in output data rates (compared to high-precision ADCs), while still being able to accurately detect the active frequency components [17].

2) We present a novel computationally efficient first-order algorithm based on the FISTA framework [18] that is able to recover the sparse RF spectrum at high throughput from coarsely quantized and compressive measurements. The proposed method, referred to as the *accelerated sparse signal dequantization* (ASSD) algorithm, efficiently solves a convex sparse signal dequantization problem [19], and enables the design of high-throughput VLSI designs.

3) We detail an analog sub-Nyquist sampling and quantization front-end in 28 nm CMOS. The front-end consists of a 4-bit flash ADC and a high-speed digital standard-cell-based pseudo-random non-uniform clock generator unit with programmable undersampling rate. The front-end acquires wideband signals non-uniformly with aggregated sampling rates ranging from 0.3 GS/s to 1.5 GS/s, while CS extends the effective reconstruction bandwidth of the ADC to 3 GHz.

4) We present a high-throughput digital VLSI design of the ASSD algorithm. To this end, we deploy a variety of approximations on algorithm level that enable its efficient implementation in VLSI. We furthermore, develop a massively-parallel 2^{15} -point radix-32 (forward and inverse) fast Fourier transform (I/FFT) unit that enables our design to detect the active frequency bands at a rate of more than 1340 RF spectrum reconstructions per second.

5) We provide extensive system-level simulations with synthetic and real measured data to characterize the performance and limitations of the proposed A2I converter design for wideband spectrum sensing applications.

C. Existing A2I Converter Architectures

In recent years, a variety of CS-based A2I converters that avoid Nyquist sampling have been described in the literature. The most prominent architectures are summarized next.

1) *Random demodulator (RD)*: The RD performs mixing of the time-domain signal with a pseudo-random spreading sequence followed by integration over a block of samples. The integrated signal is then sampled uniformly and quantized at a sub-Nyquist rate (depending on the integration length). While the RD reduces the ADC sampling rate and features low implementation complexity, the reported systems only support the reconstruction of discrete multi-tone signals [20]–[23].

2) *Modulated wideband converter (MWC)*: The MWC builds upon the Xampling framework [24] and mixes an RF input signal to baseband in multiple channels with a specific set of periodic waveforms; each baseband signal is then low-pass filtered and sampled uniformly at sub-Nyquist rate. Corresponding implementations require one low-rate ADC per channel and trade the number of channels for the sampling rate in each channel [25], [26]. The MWC is suitable for static multiband spectrum environments as the active bands must be identified using a time-consuming procedure each time the spectral activity pattern changes.

3) *Random modulation pre-integrator (RMPI)*: The RMPI resembles the MWC and integrates the mixed signals over a certain time period instead of filtering each channel. The corresponding A2I converter designs are, for example, suitable for the acquisition of radar pulses [13], [27].

4) *Non-uniform sampler (NUS)*: The NUS samples the incoming signal at irregularly spaced time intervals by taking only a subset of the samples of a conventional Nyquist converter. The corresponding implementations only consist of a sample-and-hold stage and an ADC operating at a sampling rate corresponding to the shortest sampling period used by the NUS. Existing NUS implementations mainly differ in the used clocking scheme: (i) periodic non-uniform sampling (PNUS) relies on a sequence of non-uniform sampling periods that are repeated [11]; (ii) randomized non-uniform sampling (RNUS) deploys a sampling sequence that is composed of randomly chosen periods from a set of time intervals [12], [28]; (iii) level-triggered non-uniform sampling (LTNUS) samples the signal crossings with a given waveform [29], [30].

The A2I converter proposed in this paper relies on a novel RNUS approach employing coarse quantization (in contrast to using high-precision ADCs) and includes all the necessary components on a single chip, i.e., a non-uniform clock generation circuit, the analog signal acquisition front-end, and the digital sparse spectrum recovery stage. In contrast, existing A2I converters perform sparse signal recovery either off-chip, off-line, or at very low rates; in addition, the analog front-end of existing solutions only include certain blocks, e.g., required for sampling and/or quantization, and rely on high-precision but energy-inefficient ADCs [11], [12].

D. Notation

Lowercase and uppercase boldface letters stand for column vectors and matrices, respectively. The Hermitian transpose

of a complex-valued matrix \mathbf{A} is designated by \mathbf{A}^H . The i th entry of a vector \mathbf{x} is denoted by x_i or $[\mathbf{x}]_i$. The Euclidean (or ℓ_2) norm of a vector \mathbf{x} is denoted by $\|\mathbf{x}\|_2$ and the ℓ_1 -norm is defined as $\|\mathbf{x}\|_1 = \sum_i |x_i|$. The real and imaginary part of a scalar $x \in \mathbb{C}$ is denoted by $\Re\{x\}$ and $\Im\{x\}$, respectively.

E. Paper Outline

The remainder of the paper is organized as follows. Section II introduces the A2I framework and details the accelerated sparse signal dequantization (ASSD) algorithm. Section III details the non-uniform clock generation unit and the 4-bit flash ADC. Section IV presents the high-throughput digital spectrum recovery unit. Section V presents numerical simulation results and conclusions are drawn in Section VI.

II. QUANTIZED COMPRESSIVE SENSING-BASED A2I CONVERSION

A conventional way of detecting active frequency bands in RF signals is to deploy high-precision ADCs that sample the entire spectrum at Nyquist rate followed by peak detection in the frequency domain. For wideband signals, however, such an approach necessitates costly and energy-inefficient ADCs [5]. Compressive sensing [8], a recently introduced sampling paradigm, enables one to sample signals at their “information rate” rather than at their Nyquist rate. For spectrum sensing applications where only a few frequencies are active at a given time, this sampling paradigm enables the design of cost-effective, energy-efficient A2I converters that acquire the underlying frequency activity information.

In this section, we summarize the principles of CS and introduce a way of further reducing the cost of practical A2I converter implementations by means of coarse quantization. We then present a novel, low-complexity algorithm that reconstructs sparse signals from coarsely quantized measurements.

A. Compressive Sensing Basics

CS is concerned with the sampling and reconstruction of signal vectors $\mathbf{y} \in \mathbb{R}^N$ using fewer measurements than the Nyquist rate suggests. More specifically, CS acquires M linear measurements of the signal vector \mathbf{y} as follows [31], [32]:

$$\mathbf{z} = \Phi \mathbf{y} + \mathbf{n}, \quad (1)$$

where $\Phi \in \mathbb{R}^{M \times N}$ is a sensing matrix with (often substantially) fewer rows than columns (i.e., $M < N$) and $\mathbf{n} \in \mathbb{C}^M$ represents additive measurement noise. Recovering the signal vector \mathbf{y} from the noiseless measurements $\mathbf{z} = \Phi \mathbf{y}$ is, in general, an ill-posed problem. Nevertheless, many man-made or natural signals have a sparse representation \mathbf{x} in a given orthonormal basis Ψ , i.e., $\mathbf{y} = \Psi \mathbf{x}$, where only a few entries $K \ll N$ of \mathbf{x} carry most of the vectors energy; we say that such signals are approximately K -sparse. This sparsity property enables CS to obtain accurate estimates of the signal vector \mathbf{y} if the effective matrix $\mathbf{D} = \Phi \Psi$ satisfies certain mathematical conditions [32].

Relevant for our application is the case of Φ being a randomized subsampling operator¹ and Ψ the discrete Fourier transform (DFT) matrix. In this case, it was shown in [33] that

$$M \sim K \log^4(N) \quad (2)$$

compressive measurements are sufficient to guarantee the stable recovery of the (approximately-)sparse signal representation \mathbf{x} from \mathbf{z} with overwhelming probability. It is important to realize that (2) implies that the number of measurements, M , to be taken only scales in the number of non-zero entries K (apart from a logarithmic penalty in the dimension N). In other words, instead of sampling a spectrally sparse signal \mathbf{x} at Nyquist rate (i.e., by taking N samples), randomized temporal sub-sampling at sub-Nyquist rates guarantees the stable recovery of a sufficiently sparse RF spectrum. As a consequence, CS enables one to acquire sparse wideband RF signals by taking fewer samples than the Nyquist rate dictates.

Spectrum reconstruction from compressive measurements is typically carried out by a sparse signal recovery algorithm (see [34], [35] for algorithm surveys). One of the most prominent recovery methods, known as basis pursuit de-noising (BPDN), corresponds to solving the following convex problem [36]:

$$\text{(BPDN)} \quad \underset{\tilde{\mathbf{x}}}{\text{minimize}} \quad \lambda \|\tilde{\mathbf{x}}\|_1 + \frac{1}{2} \|\mathbf{z} - \mathbf{D}\tilde{\mathbf{x}}\|_2^2,$$

which delivers an estimate $\hat{\mathbf{x}}$ of the sparse spectrum [35]. The real-valued regularization parameter $\lambda > 0$ is used to trade sparsity in $\hat{\mathbf{x}}$ for consistency to the measurements \mathbf{z} .

B. Quantized Compressive Sensing

In virtually all practical systems, ADCs are used to sample and quantize the compressive measurements. Hence, instead of acquiring real-valued measurements as in the simplistic model (1), quantized measurements are acquired in practice. The effect of such a quantizer can be modeled as

$$\mathbf{q} = \mathcal{Q}(\mathbf{z}) = \mathcal{Q}(\mathbf{D}\mathbf{x} + \mathbf{n}), \quad (3)$$

where $\mathcal{Q}(\cdot): \mathbb{R} \rightarrow \mathcal{O}$ is a scalar quantizer (applied element-wise to a vector), which maps a real number x into $Q = |\mathcal{O}|$ ordered labels according to $\mathcal{Q}(x) = q$ if $b_{q-1} < x \leq b_q$, $q \in \mathcal{O}$, where the quantization-bin boundaries satisfy $-\infty = b_0 < \dots < b_Q = +\infty$. In the following, $B = \log_2(Q)$ defines the number of output bits of the quantizer.

Quantized CS deals with the reconstruction of the sparse vector \mathbf{x} from the quantized measurements collected in \mathbf{q} . Since quantized CS takes into account the effects of finite-precision ADCs, the model (3) is particularly relevant for systems employing coarse quantization. The extreme case of 1-bit compressive measurements has recently gained significant attention in the literature [15]–[17]. In particular, [16] has established that a stable reconstruction of the sparse vector \mathbf{x} via efficient algorithms from $M \sim K \log(N/K)$ measurements is possible, if the entries of the effective matrix \mathbf{D} are i.i.d. Gaussian distributed. Unfortunately, not much is known about

¹A sparse 0/1-matrix with $M < N$ and a single 1 per row and 0 otherwise. The location of each 1 entry per row is chosen at random and defines which sample from the N dimensional input vector is taken.

the general case, i.e., $Q > 2$ with arbitrary sensing matrices. Nevertheless, empirical studies carried out in, e.g., [17], [19], with several algorithms have shown that accurate sparse signal recovery from quantized measurements is possible with as few as 4-bit and for a variety of (non-Gaussian) sensing matrices.

A practical consequence of quantized CS is the fact that it further reduces the dimensionality of the measurements to be acquired (in addition to temporal subsampling). From a hardware perspective, coarse quantization (e.g., 4-bit or even less) enables the use of low-area, low-power ADC architectures. In the A2I converter design proposed in this paper, we take particular advantage of quantized CS and deploy a low-complexity, wideband 4-bit flash ADC (see Section III).

C. Basis Pursuit De-Quantization (BPDQ)

A number of sparse signal recovery algorithms for 1-bit measurements has been proposed in the literature [37]. However, for the non-binary (but quantized) case, we are only aware of the approach proposed in [19]; this approach assumes that the noise vector \mathbf{n} in (3) is i.i.d. zero-mean Gaussian distributed with variance σ^2 per complex entry—a reasonable assumption because most practical systems are subject to thermal noise. With this assumption, one can compute the likelihood of each measurement q_i as

$$p(q_i | \mathbf{d}_i^H \mathbf{x}) = \int_{\ell_i}^{u_i} \frac{1}{\sqrt{2\pi}\sigma^2} \exp\left(-\frac{|\nu - \mathbf{d}_i^H \mathbf{x}|^2}{2\sigma^2}\right) d\nu, \quad (4)$$

where $u_i = b_{q_i}$ and $\ell_i = b_{q_i-1}$ are the upper and lower bin boundary positions associated to $q_i \in \mathcal{O}$, respectively, and \mathbf{d}_i^H corresponds to the i^{th} row of the effective matrix $\mathbf{D} = \Phi\Psi$. We emphasize that (4) explicitly models the statistical input-output behavior of the quantizer $\mathcal{Q}(\cdot)$, avoiding the commonly used model of additive quantization noise (see, e.g., [38]).

Similarly to (BPDN), the main idea is to minimize the negative log-likelihood of (4) over all measurements, together with an ℓ_1 -norm penalty that induces sparsity of the solution $\tilde{\mathbf{x}}$. We refer to the resulting *convex* optimization problem as basis pursuit de-quantization [19]

$$\text{(BPDQ)} \quad \underset{\tilde{\mathbf{x}}}{\text{minimize}} \quad \lambda \|\tilde{\mathbf{x}}\|_1 - \sum_{i=1}^M \log p(q_i | \mathbf{d}_i^H \tilde{\mathbf{x}}),$$

where the parameter $\lambda > 0$ trades sparsity of the solution $\tilde{\mathbf{x}}$ for consistency to the quantized measurements in \mathbf{q} .

D. Accelerated Sparse Signal Dequantization (ASSD)

We now propose a novel computationally more effective alternative to the algorithm presented in [19], referred to as *accelerated sparse signal dequantization* (ASSD). The ASSD algorithm is capable of recovering sparse vectors from quantized measurements (3), while (i) requiring lower computational complexity than the method in [19] and (ii) being suitable for the efficient integration in VLSI (see Section IV).

A common way of solving convex optimization problems like (BPDN) or (BPDQ) is to use interior point methods [39]. However, such methods typically exhibit high computational complexity for large-dimensional problems and require considerable numerical precision, which prohibits their efficient

Algorithm 1 Accelerated sparse signal dequantization (ASSD)

```

1:  $\mathbf{x}_1 = \mathbf{y}_0 = \mathbf{0}_{N \times 1}$  and  $t_1 = 1$ 
2: while  $k = 1, \dots, K_{\max}$  do
3:    $\mathbf{y}_k \leftarrow \text{shrink}(\mathbf{x}_k + \frac{1}{L} \mathbf{D}^H \nabla f(\mathbf{D}\mathbf{x}_k))$ 
4:    $t_{k+1} \leftarrow \frac{1}{2}(1 + \sqrt{1 + 4t_k^2})$ 
5:    $\mathbf{x}_{k+1} \leftarrow \mathbf{y}_k + \left(\frac{t_k-1}{t_{k+1}}\right)(\mathbf{y}_k - \mathbf{y}_{k-1})$ 
6: end while

```

implementation in VLSI. As demonstrated in [40], first-order methods are the preferred choice for high-throughput sparse signal recovery in VLSI for approximately sparse signals having a large number of non-zero coefficients. We therefore build the ASSD algorithm on the FISTA (short for fast iterative shrinkage thresholding algorithm) framework [18], which enables the design of accelerated first-order methods for minimization problems of the form $f(\cdot) + g(\cdot)$, where $f(\cdot)$ is convex and continuously differentiable, and $g(\cdot)$ convex, but potentially non-smooth. By associating $f(\cdot)$ with the negative log-likelihood function of (BPDQ), and $g(\cdot)$ with the ℓ_1 -norm penalty, we obtain Algorithm 1. The ASSD algorithm performs the following three steps until a maximum number of iterations, K_{\max} , has been reached:

1) *Gradient step*: The gradient step enforces consistency to the quantized measurements \mathbf{q} . To arrive at an explicit gradient formulation, we set $w_i = \mathbf{d}_i^H \mathbf{x}$, and rewrite (4) as

$$p(q_i | w_i) = \Phi(\sigma^{-1}(u_i - w_i)) - \Phi(\sigma^{-1}(\ell_i - w_i)),$$

where $\Phi(a) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^a \exp(-\frac{1}{2}\nu^2) d\nu$ is the cumulative distribution function of a normal random variable. With the definition $f(\mathbf{w}) = -\sum_{i=1}^M \log p(q_i | w_i)$, the i^{th} entry of the gradient $\nabla f(\mathbf{w})$ is given by [19]

$$[\nabla f(\mathbf{w})]_i = \frac{\exp\left(-\frac{|u_i - w_i|^2}{2\sigma^2}\right) - \exp\left(-\frac{|\ell_i - w_i|^2}{2\sigma^2}\right)}{\sqrt{2\pi}\sigma^2 (\Phi(\frac{u_i - w_i}{\sigma}) - \Phi(\frac{\ell_i - w_i}{\sigma}))}. \quad (5)$$

To ensure convergence of Algorithm 1, we employ a constant step size that is determined by the Lipschitz constant L (see line 3). This constant is given by $L = \lambda_{\max}^2(\mathbf{D})/\sigma^2$, where $\lambda_{\max}(\mathbf{D})$ corresponds to the largest singular value of \mathbf{D} .

2) *Shrinkage step*: This step takes into account the ℓ_1 -norm in (BPDQ) and enforces sparsity on the vector \mathbf{x} . From the definition $g(\mathbf{x}) = \lambda \|\mathbf{x}\|_1$, the *complex-valued* shrinkage step in Algorithm 1 (line 3) is given by [18]

$$\text{shrink}(x) = \begin{cases} \frac{x}{|x|} \max\{|x| - \lambda/L, 0\} & \text{if } x \neq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

3) *Prediction step*: The prediction step in Algorithm 1 (lines 4 and 5) is used to obtain a new estimate of the sparse vector \mathbf{x}_{k+1} . As detailed in [18], the particular update of the ASSD algorithm yields accelerated convergence rates, which is key for achieving low computational complexity.

Note Algorithm 1 is suitable for general matrices $\mathbf{D} = \Phi\Psi$ and for processing on general-purpose processors. To arrive at an efficient implementation in VLSI, we derive a dedicated version for sparse RF spectrum recovery in Section IV.

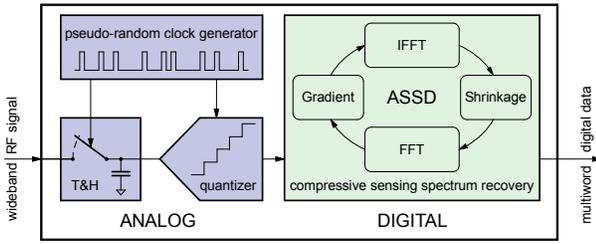


Fig. 1. Overview of the monolithic A2I converter. The design consists of a pseudo-random non-uniform clock generator, a low-bit flash ADC, and a high-throughput digital spectrum recovery stage.

III. ANALOG FRONT-END: SUB-NYQUIST FLASH ADC

We now detail the analog front-end, consisting of a pseudo-random non-uniform clock generator and a 4-bit flash ADC. The resulting design is illustrated in Figure 1. To the best of our knowledge, the only existing hardware implementation of a NUS-based A2I converter targeting a signal bandwidth of several GHz was reported in [11]; this design is implemented using an expensive InP HBT technology for the sample-and-hold (S&H) stage and a commercial, high-resolution, and off-chip ADC for the signal conversion stage. In contrast, the design proposed here builds on a standard 28 nm CMOS technology for the sampling, conversion, *and* recovery stage, and all three units are integrated into the same design.

A. Pseudo-Random Non-Uniform Clock Generator Unit

Random subsampling-based CS for wideband signal acquisition requires a (pseudo)-random non-uniform sampling clock at high rates. In existing A2I converter designs, such as the ones presented in [11], [12], the non-uniform sampling clock is generated off-chip and/or stored in large on-chip register arrays. Apart from requiring a considerable number of flip-flops to hold the sampling pattern, a major shortcoming of such an approach is the inability to adjust the undersampling rate—allowing to tune the A2I converter to changes in the sparsity level of the input spectrum—at run-time. To avoid the use of expensive and rather bulky external equipment, we propose a low-area and configurable standard cell-based clock-generation unit that derives a high-rate non-uniform sampling clock on-chip from an external Nyquist clock with period T_{clk} ; the architecture is depicted in Figure 2 and detailed next.

1) *Architecture*: The pseudo-random non-uniform clock generator unit consists of a linear feedback shift register (LFSR), a pipelined multiplexer tree, and a shift register (SR). The multiplexer tree and SR together form a circular shift register with a variable length according to the selection bits B'_0 – B'_3 of the multiplexer. A single logical 1, initially set by the RST signal in the first flip-flop (FF) of the SR, is propagated by the input clock signal through this circular shift register. The state of the multiplexer output represents the non-uniform sampling clock signal ϕ_{nus} . The pseudo-random sampling period T_{rnd} is an integer multiple of the uniform input clock period and is given by $T_{\text{rnd}} = T_{\text{min}} + N_{\text{rnd}} T_{\text{clk}}$, where the minimum sampling period T_{min} is equal to the time it takes for the logical 1 to propagate through the FFs in the SR whose

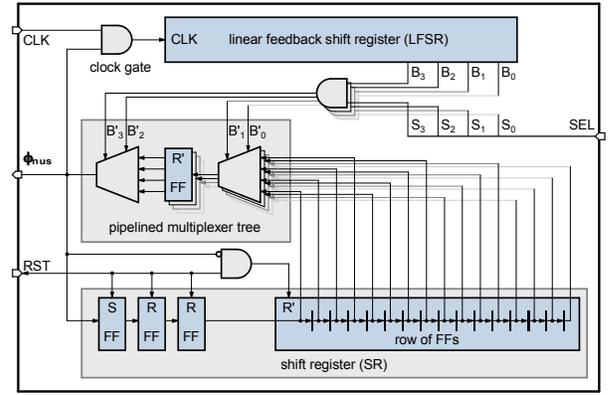


Fig. 2. High-rate pseudo-random non-uniform clock generator unit with configurable undersampling factor. The clock generation unit is configurable to generate the following undersampling factors: 4, 4.5, 5.5, 7.5, and 11.5.

output is not connected to the multiplexer tree. The logical 1 has to propagate through those FFs first, before it can reach any of the multiplexer inputs. N_{rnd} is the pseudo-random binary number coded by B'_0 – B'_3 , that are the 4 least significant bits of the LFSR state masked by the bits of the undersampling configuration signal SEL. In the present design, we chose the minimum sampling period to be $4 T_{\text{clk}}$ as the ADC operates in four phases (see Section III-B). The maximum sampling period T_{max} can be configured via the SEL signal by restricting the length of the circular shift register either to 19, 11, 7, 5, or 4 FFs. This enables one to tune the undersampling factor at run-time to one of the following values: 11.5, 7.5, 5.5, 4.5, and 4. Each setting corresponds to a set of sampling periods from which a new period is selected pseudo-randomly every time the logical 1 reaches the multiplexer output, as this event triggers the LFSR to generate a new (pseudo)-random number. The LFSR length is 11 and is chosen such that all generated sampling periods occur approximately equally often.

2) *Design*: The entire clock generation circuit is built from 33 flip-flops of the fastest type available in the used 28 nm CMOS standard-cell library. By inserting pipeline registers into the multiplexer tree, the critical path is reduced to one standard multiplexer, thereby allowing this design to achieve a maximum clock frequency of 6 GHz. The estimated power consumption at the maximum speed is 0.5 mW.

B. 4-bit Flash Analog-to-Digital Converter

The second component of the analog front-end is the ADC that samples and coarsely quantizes the compressive measurements. In the proposed A2I converter design, we use a high-rate ADC that acquires 4-bit samples in the GS/s regime.² High-speed ADCs with sampling rates of several GS/s and resolutions up to 6-bit often deploy expensive SiGe, InP, or GaAs technologies [41]. Alternatively, wideband CMOS ADCs targeting similar bandwidth and precision either rely on extensive time-interleaving [42], or use for instance integrated inductors to overcome the limits imposed by the technology [5]. However, both approaches lead to a considerable

²As it will be shown in Section V-B, 4-bit precision represents a good trade-off between circuit complexity and spectrum reconstruction accuracy.

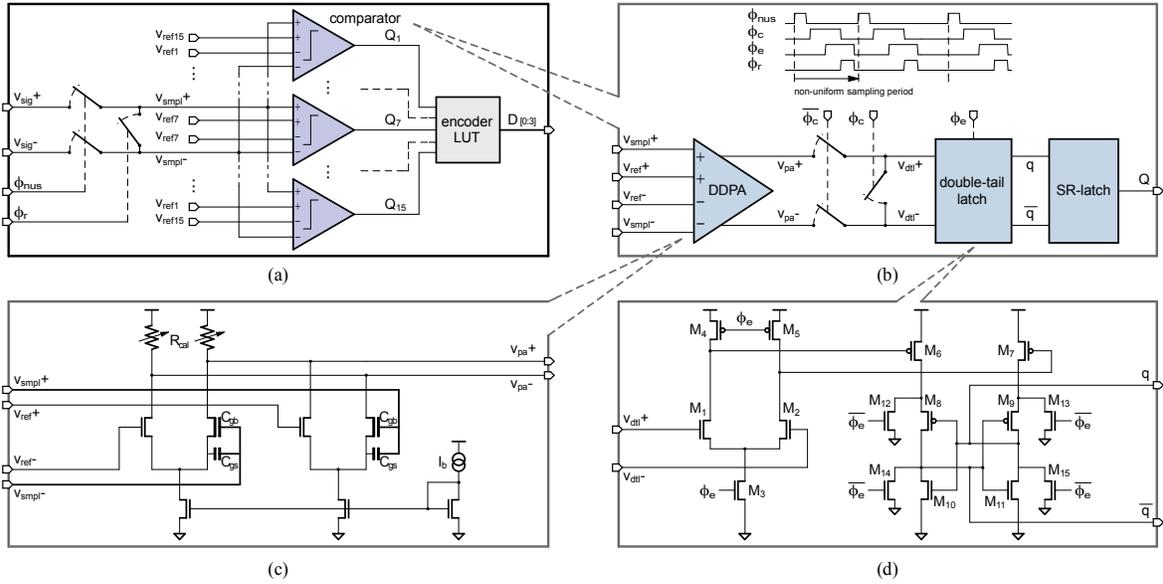


Fig. 3. Circuit diagram of the 4-bit flash ADC: (a) overview; (b) comparator and the associated timing diagram; (c) preamplifier; (d) double-tail latch.

increase in circuit complexity and power consumption, which would be in contrast with our attempt to keep the level of sophistication of the analog front-end at a minimum. Note that targeting resolutions above 8-bit at several GS/s would still be feasible in deep sub-micron CMOS technologies [43], [44]. Such implementations, however, require expensive digital calibration circuits to compensate for circuit imperfections. Hence, to push our paradigm of shifting complexity from the analog to the digital domain to the limit, we decided to develop a 4-bit flash ADC in 28 nm CMOS, which enables sampling rates of several GS/s at minimum hardware overhead.

1) *Architecture*: Targeting a modest resolution of 4-bit, a flash ADC appeared to be the natural choice for our design. An overview of the 4-bit flash ADC is shown in Figure 3(a). The ADC consists of 15 identical comparators Q_1, \dots, Q_{15} preceded by a shared sampling switch, and of a resistor ladder for the generation of the voltage references. Each comparator is composed of a static differential-difference-preamplifier (DDPA), followed by a double-tail latch (DTL) and a standard SR-latch [45], as shown in the lower half of Figure 3(b). The SR-latch simply keeps the comparator output stable, allowing the encoder to perform the thermometric-to-binary conversion. The preamplifier serves the main function of track-and-hold (T&H) stage, eliminating propagation delay differences between the comparators, and additionally provides a modest amplification. To avoid the use of dedicated sampling capacitors, we use the total gate capacitance of the input transistors (approximately equal to 500 fF) to sample the input signal. To eliminate the signal-dependent modulation of the switch on-resistance, which would degrade linearity, the sampling switch is controlled by a bootstrapped clock phase named ϕ_{nus} [46]; this clock is derived from the non-uniform clock-generator. The four phases required to operate the ADC are schematically depicted in the upper part of Figure 3(b). The DDPA is built from two resistively loaded differential

pairs that compare the differential input with the differential voltage reference, as depicted in Figure 3(c).

The latching stage is detailed in Figure 3(d), and is realized by means of a double-tail voltage sense amplifier similar to the one reported in [47]. The latch is operated with a single clock phase ϕ_e which keeps the circuit complexity to a minimum. However, compared to the circuit topology proposed in [47], the drain of the input transistors M_1 and M_2 is directly tied to the gate of M_6 and M_7 in order to inject the differential signal into the latch and to trigger it. The four reset devices M_{12} – M_{15} on the other hand are driven by the phase $\bar{\phi}_e$. The high current required in the latching stage necessitates the use of large transistors for M_6 and M_7 , whereas M_{12} – M_{15} can be kept at the minimum size. It is therefore convenient to connect the large devices M_6 and M_7 to the output of the first stage to minimize their contribution to the input-referred offset.

Moreover, the small gain of approximately 2 provided by the DDPA relaxes the matching requirements of the DTL, allowing us to employ small and fast transistors in the latch for maximum speed. The kick-back noise generated by the switching activity of the latch is attenuated by the static preamplifier, which prevents it from propagating backwards toward the input or toward the reference ladder. Furthermore, during the latch phase ϕ_e , the DTL is completely disconnected from the DDPA by means of the pair of switches controlled by ϕ_e to suppress kick-back noise even further.

2) *Design*: The entire analog front-end was designed in a 28 nm 1P7M bulk CMOS technology with a single 1.0 V supply. The resulting flash ADC achieves a maximum sampling rate of 6 GS/s. However, the entire A2I converter design is constrained by the external Nyquist clock from which the 4-phased timing of the ADC is deduced. Given the maximum operating frequency of the non-uniform clock generator of 6 GHz, the maximum sampling rate of the entire analog front-end is limited to 1.5 GS/s. In contrast to conventional Nyquist

rate data acquisition, this fact is not limiting the reconstruction bandwidth of the overall system, as the ASSD algorithm is based on sub-sampling, and effectively extends the conversion bandwidth of the ADC to 3 GHz.

The performance of the 4-bit flash ADC has been extensively characterized in all process corners by means of Cadence Spectre simulations. In a simulated single-tone test with a 400 mV_{pp} full-scale sinusoidal input the ADC achieved 25.8 dB SNDR at maximum clock rate, which corresponds to an effective number of bits (ENOB) of 4.0 bit. Each of the 15 static preamplifiers absorbs 200 μ A from the 1.0 V supply, while the entire analog front-end consumes an estimated power of 4.5 mW at 1.5 GS/s; this is expected to be negligible compared to the power consumption of the digital part, and compares favorably to designs reported in literature targeting similar speed and resolution, such as the one in [5]. We finally note that the layout of the analog front-end is part of ongoing work. However, a comparison of the proposed ADC design with a structurally similar reference 3.5-bit flash ADC implemented in 130 nm CMOS [48], allows us to obtain a pessimistic area estimate of about 0.1 mm² in 28 nm CMOS.

IV. DIGITAL PART: HIGH-THROUGHPUT ASSD UNIT

Virtually all existing A2I converter designs delegate the task of signal reconstruction to off-line CPU or GPU processing [11]–[13]. High-throughput and energy-efficient sparse signal recovery, however, can only be achieved by dedicated VLSI implementations, because even the most efficient algorithms exhibit high computational complexity [40]. To arrive at high-throughput sparse spectrum recovery from coarsely quantized and compressive measurements, we next detail a variety of optimizations for the ASSD algorithm. We then develop a corresponding VLSI design in 28 nm CMOS, which directly interfaces with the analog front-end.

A. ASSD Algorithm Optimizations for Spectrum Recovery

In order to obtain a high-throughput ASSD design with finite-precision (fixed-point) arithmetics, we introduce a host of new algorithm-level optimizations and approximations that facilitate an efficient integration in VLSI.

1) *Precomputed Lipschitz constant*: For general matrices \mathbf{D} , the ASSD algorithm requires the calculation of the Lipschitz constant $L = \lambda_{\max}^2(\mathbf{D})/\sigma^2$. In the present sparse spectrum recovery application, \mathbf{D} corresponds to a randomly-subsampled DFT matrix; for this particular sensing matrix, the maximum singular value is given by 1 and hence, we have $L = 1/\sigma^2$; this parameter is stored in a configuration register to avoid computation of the Lipschitz constant in VLSI.

2) *Precomputed prediction weights*: Straightforward computation of the sequence t_k , required to accelerate the convergence of Algorithm 1 on lines 4 and 5, involves the execution of costly square root and division operations. Since the sequence t_k depends only on the iteration counter k , we can precompute the quantity $\tau_k = (t_k - 1)/t_{k+1}$ on line 5 of Algorithm 1 and store it in a look-up table (LUT). This trick allows us to avoid costly arithmetic circuitry at the cost of a 128-entry LUT for the values τ_k , since the final design is able to carry out a maximum of 128 iterations.

3) *Approximate gradient calculation*: The gradient step (5) involves the evaluation of transcendental functions, which cannot be implemented efficiently in VLSI using fixed-point arithmetics. Nevertheless, inspection of (5) reveals that the gradient can be approximated with the piece-wise linear function

$$[\nabla f(\mathbf{w})]_i \approx \begin{cases} \frac{u_i - w_i}{\sigma^2} & w_i > u_i \\ 0 & \ell_i \leq w_i \leq u_i \\ \frac{\ell_i - w_i}{\sigma^2} & w_i < \ell_i, \end{cases} \quad (7)$$

especially for small values of σ^2 . This approximation can be implemented efficiently in VLSI using basic arithmetic circuitry and comparison logic. As will be discussed in Section V-C, the approximation delivers comparable performance to an algorithm variant that computes the gradient exactly.

4) *Approximate shrinkage*: The complex-valued shrinkage operation (6) requires a division of $x = \Re\{x\} + i \cdot \Im\{x\}$ by $|x| = \sqrt{\Re\{x\}^2 + \Im\{x\}^2}$, which involves significant hardware overhead and is prone to numerical issues in fixed-point arithmetics. To avoid both issues, we perform shrinkage for the real and imaginary part of $x \in \mathbb{C}$ as follows:

$$\text{shrink}(x) \approx \eta \Re\{x\} + i \cdot \eta \Im\{x\}. \quad (8)$$

Here, $\eta(v) = \text{sign}(v) \max\{|v| - \lambda/L, 0\}$ corresponds to *real-valued shrinkage*, which can be implemented at minimum hardware cost. Our own simulations have shown that using the approximation (8) instead of (6) causes only a minor performance loss (see Section V-C for the details). Note that we precompute and store the quantity $\lambda/L = \lambda\sigma^2$ in a configuration register of the recovery unit.

B. High-Level Architecture of the ASSD Algorithm

With the algorithm optimizations summarized above, we can implement the ASSD algorithm efficiently in VLSI; Figure 4(a) details the corresponding high-level architecture.

1) *Input memories*: The analog front-end delivers a single-bit signal indicating the sampling instant and the corresponding 4-bit time-domain samples to the digital recovery stage. Both, the incoming sampling instants and the 4-bit samples are stored in on-chip SRAMs ω and s_q , respectively. Targeting the reconstruction of a 2¹⁵-dimensional RF spectrum, both memories contain 2¹⁵ entries, each corresponding to a possible sampling instant. As we perform non-uniform sub-Nyquist sampling, only a pseudo-random sub-set of the 2¹⁵ time-domain entries of s_q contains valid (observed) samples. A logical 1 in the ω SRAM identifies a valid sample in the s_q memory. The look-up table q (called q LUT) contains the digital values representing the upper and lower quantization bin positions u_i and ℓ_i , respectively. This LUT is built from flip-flops and enables us to compensate for mismatches in the 15 analog reference voltages and for the comparators' offsets. The 4-bit samples directly address the entries of the q LUT.

2) *Architecture overview*: The ASSD architecture shown in Figure 4(a) comprises three main units: The *approximate gradient* unit implementing the piece-wise linear gradient approximation (7), the *approximate shrinkage* unit realizing (8), and the 2¹⁵-point radix-32 I/FFT unit performing the forward

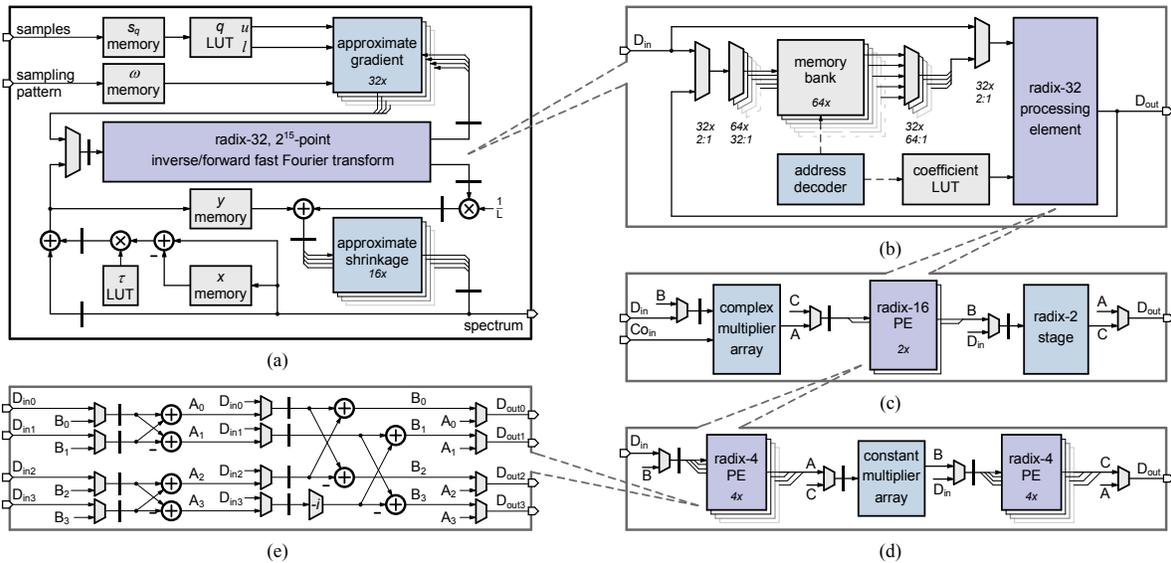


Fig. 4. ASSD recovery unit: (a) architecture overview; (b) radix-32 IFFT unit; (c) radix-32 processing element (PE); (d) radix-16 PE; (e) radix-4 PE.

and backward FFT. Spectrum recovery is achieved by alternately performing forward and inverse FFTs. The approximate gradient calculation is carried out during the last cycles of the inverse FFT. In this phase, data coming from the IFFT unit is processed by the approximate gradient unit and the result is directly fed back to the FFT memory, ready for the forward transform. Similarly, during the last cycles of the forward FFT operation, shrinkage and linear prediction are performed. The corresponding results are simultaneously written back to the FFT memory, ready for the inverse FFT required in the next iteration. The result of the shrinkage step is available at the output of the ASSD unit and, in the final iteration of the algorithm, this result corresponds to the RF spectrum estimate. The number of clock cycles required to carry out one ASSD iteration corresponds to the sum of the cycles required for one forward and one inverse FFT.

C. High-Throughput Parallel IFFT Unit

The input dimensionality and throughput of the FFT unit determine the spectral and temporal resolution of the A2I converter. While targeting a total reconstruction bandwidth of 3 GHz, sensing the activity within the narrow bands of today's communication standards needs an FFT of at least 2^{15} points corresponding to a resolution of 183 kHz per bin.³ This choice of the FFT size and bandwidth evidently results in spectral leakage, which can either be mitigated by adding an appropriate windowing filter or by acquiring and processing more samples at a given sampling rate. Increasing the FFT size can be done at compile time of the design, but results in longer processing time and larger memories, and therefore, would substantially increase the silicon area of the entire A2I converter. In addition, the proposed A2I converter architecture requires a high-throughput IFFT unit in order to achieve high temporal resolution, i.e., to detect fast changes in the spectral

activity. The highest throughput is achieved by fully parallel FFT architectures [49], which result in increased complexity compared to, e.g., cascade FFT architectures. However, for our case, this cost is balanced by the benefit of enabling the simultaneous processing of multiple data items by the gradient and shrinkage units, which further accelerates the ASSD implementation. To simultaneously achieve high throughput and high spectral resolution, we decided to develop a parallel memory-based 2^{15} -point IFFT unit detailed next.

1) *Choice of FFT architecture:* The number of clock cycles required to calculate an N -point FFT scales with $N/(mr) \log_r(N)$, where m is the number of parallel radix- r processing elements (PE). Hence, by choosing higher radix orders and/or instantiating multiple PEs, one can increase the throughput of the FFT unit. Our goal is to identify an FFT architecture that maximizes the throughput while not resulting in excessive silicon area. To this end, it is helpful to analyze the maximum number K of ASSD iterations that can be completed during N cycles which is the time it takes for the sampling phase to finish. Noting that the number of frequency bins of the FFT is also equal to N , we obtain $K = mr/(16 \log_r(N))$, where we assumed the digital logic to run $8\times$ slower than the Nyquist clock. The most hardware-efficient FFTs are obtained when the number of points is a power of the radix number. Thus, for a spectral resolution of 2^{15} points, one can choose between architectures based on radix-2, radix-8 and radix-32. By assuming a minimum of $K = 20$ iterations, all possible radix numbers clearly result in prohibitive silicon complexity under the real-time constraint. Consequently, as a good trade-off between silicon area and achievable recovery throughput, we decided to implement a 2^{15} -point inverse/forward fast FFT unit (Figure 4(b)) based on a single radix-32 PE. This particular choice is due to the observation that FFT architectures using higher radix orders are smaller in size for a given throughput; the number of instantiated complex-valued multipliers $m(r-1)$ is lower and

³The channel bandwidth of many established communication standards, such as GSM, is as low as 200 kHz.

the interconnect network between memory and PEs, which is a logarithmic function of the number of inputs $m \cdot r$ to the PEs, is less complex. The chosen FFT configuration allows ASSD-based spectrum recovery $30\times$ below real-time.

2) *VLSI design*: In each clock cycle, the radix-32 PE requires access to 32 data items from a memory and writing the corresponding 32 results back to the same memory locations. To achieve such a massive parallelism without causing access contentions, we partition the FFT memory into 64 independent banks (see Figure 4(b)). To minimize the area of the FFT memory, we use multiple single-port memories in combination with a specifically designed memory access scheme that ensures contention-free read and write access. Concretely, in each clock cycle, data for the radix-32 PE is retrieved from a specific set of 32 memories, whereas the FFT's output of a previous cycle is stored in the remaining 32 memories.

The radix-32 PE (see Figure 4(c)) is built from a combination of a radix-16 stage and a radix-2 stage in a split-radix fashion [50]. The radix-16 stage is performed by two identical radix-16 PEs (Figure 4(d)), each consisting of 8 multiplier-less radix-4 PEs (Figure 4(e)) and 31 complex-valued multipliers. The inverse FFT can be calculated using the same PE, with the aid of multiplexers in the data path. The data path of the forward FFT is inverted to conform to the data path of the inverse FFT, where the data items first pass through the radix-2 step and finally arrive at the complex multipliers fed with the complex conjugate coefficients. The same data-path inversion is used in the radix-16 PEs. To maximize the clock frequency and, hence, the recovery throughput, the radix-32 PE features a total of 11 pipelining stages, of which 2 stages are used in the complex-valued multipliers.

The resulting I/FFT unit is capable of computing a 2^{15} -point I/FFT in 3097 clock cycles, which is close to the theoretical minimum of 3072 cycles achievable with a single radix-32 unit. The extra cycles are due to pipelining and the contention-free memory access scheme. We finally note that this architecture is roughly $80\times$ faster than a conventional single PE radix-2 FFT design in the same technology. Post-synthesis timing results for the I/FFT unit in 28 nm CMOS show that we can achieve a maximum clock frequency of 830 MHz, which leads to a throughput of more than 8.7 GS/s. We note that a related VLSI implementation of a 2^{15} -point FFT was reported in [51]; this design contains 4 parallel radix-2 PEs and achieves 9 MS/s in 90 nm CMOS.

D. Fixed-Point Parameters

To minimize circuit area and power consumption, and to maximize the throughput, the entire ASSD architecture uses fixed-point arithmetic. All signal word-widths were established using extensive simulations of a Matlab golden model, to ensure an implementation loss well-below the quantization error of the ADC. The acquired time-domain signal is quantized with 4 bit precision. The quantization bin boundaries are programmable to any 14 bit value, which provides sufficient resolution to compensate for possible mismatches/offsets in the analog front end. The real and complex part of data in the radix-32 PE are represented by 24 bit, which also define the

word-length of the FFT memories, as well as the precision in the gradient and thresholding units. Thus, both the time- and frequency-domain signals are represented with 24 bit. The FFT twiddle-factors use 18 bit, while the τ LUT has 8 bit entries.

V. RESULTS AND DISCUSSION

We now characterize the performance and implementation complexity, as well as the limitations, of the proposed wide-band A2I converter. The front-end design of the digital part has been completed including register-transfer level (RTL) description and gate-level netlist compiled using the available 28 nm CMOS standard cell library, and will be discussed next.

A. Performance Measures and Algorithm Parameters

In order to evaluate the spectral activity detection capabilities of the proposed A2I converter, we conduct a series of experiments using the following performance metrics:

- *True positive detection rate*: The number of correctly detected active frequency bins divided by the total number of effectively active bins.
- *False positive detection rate*: The number of frequency bins falsely found to be active by the A2I converter divided by the total number of inactive bins.
- *Reconstruction signal-to-noise ratio (RSNR)*: The signal power in the active frequency bins (as detected by the A2I converter) divided by the remaining signal power.

The optimal detection threshold for identifying active frequency bands depends on algorithm parameters that are difficult to determine in practice (e.g., the ambient noise floor and the signal sparsity level). Therefore, we set the spectral activity threshold to $-6.02B - 1.76 - 10 \log_{10}(N) + 20$ [dB], i.e., 20 dB above the quantization noise floor, which performed best in our simulations. Here, $B = \log_2(Q)$ is the number of bits of the quantizer, and N the number of FFT points; the term $10 \log_{10}(N)$ takes the normalization constant of the FFT into account. We set the regularization parameter λ of the ASSD algorithm to the value which results in the highest true positive and smallest false positive detection rate separately for each resolution. As the quantization noise of the ADC is the predominant source of noise in the ASSD algorithm, σ^2 was set to the quantization noise power of the considered resolution and is given by $V_{\text{LSB}}^2/12$, with V_{LSB} being the voltage difference between the quantization levels.

B. Spectrum Sensing Performance

To characterize the impact of the signal sparsity and the noise sensitivity on the detection rate, we carried out simulations using a floating-point model of the A2I converter including the algorithmic approximations discussed in Section IV-A. Synthetic test data was used in which the percentage of active frequency bins was set according to the desired sparsity level, while the location and spectral magnitude were both chosen at random.⁴ To obtain the desired signal-to-noise ratio (SNR) in the test data, i.i.d. zero-mean Gaussian noise with appropriate

⁴The locations and non-zero entries were generated using an i.i.d. uniform and i.i.d. zero-mean Gaussian distribution with unit variance, respectively.

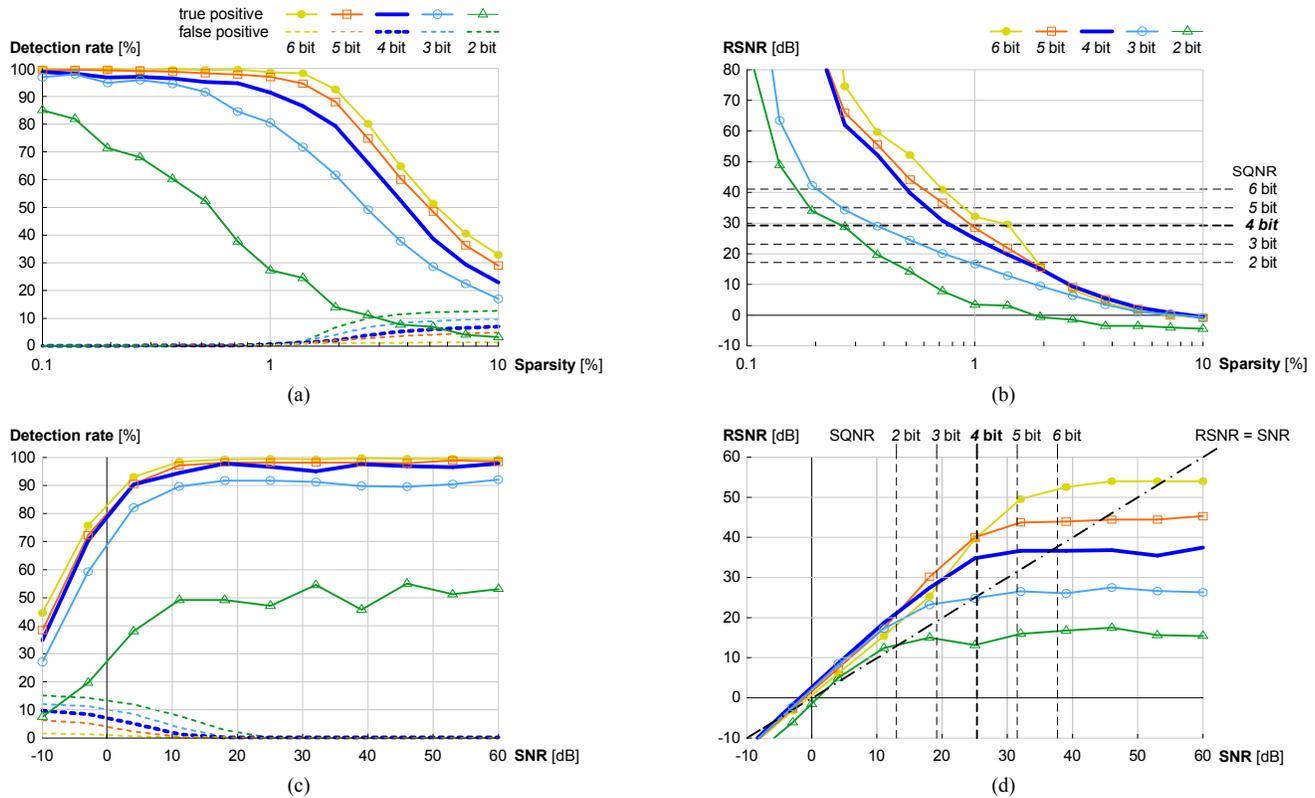


Fig. 5. Detection performance and RSNR of the ASSD algorithm for synthetic test data: (a) detection performance for varying signal sparsity levels; (b) RSNR for varying signal sparsity levels; (c) detection performance at different input signal SNRs; (d) RSNR at different input SNRs. Simulation parameters for $K_{\max} = 100$ ASSD iterations: SNR = 3 dB below quantization noise level (for sparsity trials), 0.5% sparsity (for SNR trials), undersampling factor is 11.5, λ values are 2.3, 3.7, 8.8, 55, and 165 for 2-to-6 bit resolution, respectively.

variance was added to the time-domain data. All results were averaged over 10 simulation trials, each running for 100 iterations of the ASSD algorithm.

1) *Impact of signal sparsity:* Figure 5(a) characterizes the impact of the signal sparsity level on the true positive and false positive detection rates for a different number of quantization bits B . The SNR of the input signal was set to exceed the corresponding signal-to-quantization-noise ratio by 3 dB. As it can be seen in Figure 5(a), reducing the signal sparsity level from 10% to 0.1% active bins improves the detection performance. The performance drop visible around a sparsity level of 1% active bins is a consequence of the chosen undersampling factor and the choice of λ . Both the parameters can be set at run-time in order to adapt the A2I converter to the signal's sparsity level. For the presented results, the undersampling factor was set to 11.5, while the appropriate λ values are 2.3 for 2 bit, 3.7 for 3 bit, 8.8 for 4 bit, 55 for 5 bit, and 165 for 6 bit resolution. While there are considerable performance differences between 2-bit and 4-bit quantization, 4-bit and higher achieve similar performance. Thus, we conclude that 4-bit quantization provides a reasonable trade-off between spectrum activity detection performance and ADC implementation complexity.

Figure 5(b) characterizes the impact of the signal sparsity level on the RSNR for a different number of quantization bits B . Interestingly, for each resolution, the RSNR exceeds the corresponding SQNR below a certain sparsity level; this

is because the ASSD algorithm is capable of mitigating the effects of thermal and quantization noise by imposing sparsity to the recovered spectrum. In other words, the ASSD algorithm effectively dequantizes the recovered sparse RF spectrum.

2) *Impact of thermal noise:* Figures 5(c) and 5(d) characterize the impact of the noise level on the detection rate and RSNR for input SNR levels ranging from -10 dB to 60 dB; the signal sparsity level was set to 0.5% for all trials, while the λ values are as in Section V-B1. We observe that the true positive detection rate quickly drops for input SNRs below 0 dB, whereas larger SNRs show good detection performance. Similarly, the RSNR starts to approach the input SNR at low SNR levels. In summary, for a sufficiently high input SNR, the proposed A2I converter is capable of achieving true and false positive detection rates close to 100% and 0%, respectively.

C. Impact of Approximations and Fixed-Point Arithmetic

The VLSI design of the A2I converter was facilitated by employing various algorithm-level approximations (see Section IV-A) and by means of fixed-point arithmetic. Both of these measures induce implementation-related non-idealities that evidently affect the detection performance. Figure 6 compares the detection rates using an ideal floating-point model, a floating-point model including the approximations, and the fixed-point golden model of the A2I converter. The parameter set was identical for all the models: The SNR was 3 dB below the quantization noise level, $\lambda = 2.0$, and the undersampling

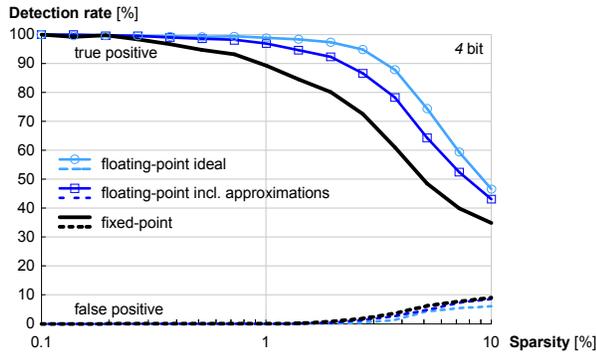


Fig. 6. Comparison of the detection rate for an ideal floating-point model, a floating-point model including the algorithm-level approximations detailed in Section IV-A, and the fixed-point golden model. 100 ASSD iterations were simulated using synthetic data with the noise floor 3 dB below the quantization noise level; λ was 2.0 and the undersampling factor was 7.5.

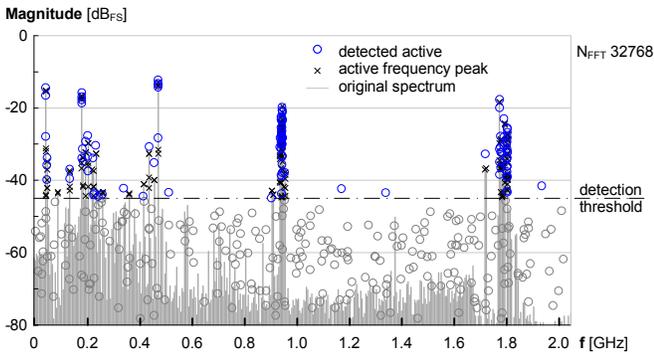


Fig. 7. Real-world RF spectrum recovery using a Matlab golden model of the A2I converter; blue circles correspond to the frequency bins detected by the ASSD algorithm. Simulation parameters for 100 ASSD iterations, $\lambda = 2.0$, undersampling factor 5.5, and a signal activity threshold of -45 dB.

factor was 7.5. The ASSD algorithm performed 100 iterations and the results were averaged over 10 Monte-Carlo trials.

As it can be seen from Figure 6, the used algorithm-level approximations only entail a small loss in detection performance; the performance of the fixed-point implementation is further reduced. However, we emphasize that careful parameter tuning recovers this performance loss by a large extent.

D. Simulations with Real-World Data

In order to assess the real-world performance of the proposed A2I converter, we carried out Matlab-based simulations using real-world signals acquired by a 2.1 GHz frequency analyzer during daytime at ETH Zurich, Switzerland. The input spectrum, shown in light gray in Figure 7, includes distinct channel aggregates with several MHz of bandwidth at 900 MHz and 1800 MHz, and several single and multi-tone signals. The RF spectrum was recovered by a Matlab fixed-point golden model of the proposed A2I converter (blue and dark gray circles); the results for this experiments are summarized in Table I and remain consistent for different signal trials. For this scenario the undersampling factor was set to 5.5, which corresponds to an average sampling rate of 1.09 MHz, i.e., the average sampling rate was $3.85\times$ below the

TABLE I
PERFORMANCE SUMMARY FOR REAL-WORLD TEST DATA

| | | | | | |
|-----------------|----------------------|-----|---------------------|------|-------|
| Quantization | [bit] | 4 | Active threshold | [dB] | -45 |
| T_{\min} | [T_{clk}] | 4 | Signal SNR | [dB] | 21.9 |
| T_{\max} | [T_{clk}] | 7 | True positive rate | [%] | 60 |
| Undersampling | | 5.5 | False positive rate | [%] | < 0.1 |
| ASSD iterations | | 100 | Reconstruction SNR | [dB] | 26.3 |

TABLE II
POST-SYNTHESIS RESULTS IN 28 nm CMOS

| Analog Front-End | | |
|--|--------------------|------------------------|
| Max. uniform sampling rate | [GS/s] | 1.5 |
| Non-uniform sampling rates | [GS/s] | 0.3–1.5 |
| Undersampling factors | | 4, 4.5, 5.5, 7.5, 11.5 |
| ADC max. power cons. ^a | [mW] | 4.5 |
| ADC idle power cons. | [mW] | 3.0 |
| Clock generator power cons. ^a | [mW] | 0.5 |
| Energy efficiency | [pJ/sample] | 2.9 |
| Area ^b | [mm ²] | 0.1 |
| Digital ASSD Unit | | |
| Max. clock frequency | [MHz] | 830 |
| Max. throughput ^c | [MS/s] | 220 |
| Standard cell area | [mm ²] | 1.1 |
| Standard cell based logic ^d | [MGE] | 1.7 |
| SRAM macro cell area | [mm ²] | 1.0 |
| Memory size | [MBit] | 3.3 |
| Power consumption ^e | [W] | 1.8 |

^aEstimated power consumption at max. frequency, $V_{dd} = 1$ V, and 300 K.

^bConservative estimate; layout is ongoing work.

^cAt $K_{\max} = 20$ ASSD algorithm iterations.

^d1 GE equals $0.4896 \mu\text{m}^2$ in the given 28 nm CMOS technology.

^eEstimated power consumption at 830 MHz, $V_{dd} = 1$ V, and 300 K.

Nyquist frequency of 4.2 GHz. The optimal value of $\lambda = 2$ has been determined via simulations; the detection threshold was set at -45 dB in order to obtain roughly 1% active bins.

From Figure 7, we can see that the dominant frequency peaks⁵ are detected with high accuracy. We note, however, that the true detection rate is reduced compared to the tests performed using synthetic data. This can be partly attributed to the fact that the input spectrum is only approximately sparse. In order to improve the detection performance for frequency bands of low SNR, one must resort to methods that either exploit specific features of the underlying communication standards or take multiple RF spectra into account [11].

E. Design Results of the A2I Converter

The design results for the analog, as well as the digital front-end, are summarized in Table II. The analog front-end is capable of acquiring wideband signals at an average sampling rate as low as 522 MS/s using sampling frequencies ranging from 0.3 GS/s to 1.5 GS/s. The non-uniform pseudo-random clock generator runs at a maximum clock frequency of 6 GHz, which results in a maximum sampling rate of 1.5 GS/s. We emphasize that the clock generator allows us to configure the undersampling factor up to 11.5, which renders the analog front-end power efficient in comparison to

⁵The two dominant aggregates at 900 MHz and 1800 MHz can be attributed to signals resulting from European GSM bands.

conventional Nyquist rate ADCs with the same reconstruction bandwidth. The digital ASSD unit includes memory macro-cells operating at low maximum clock frequency, i.e., in the range of 0.5 GHz to 1.5 GHz depending on the SRAM size and aspect ratio. Nevertheless, the limiting factor in terms of maximum operating frequency is the radix-32 PE of the I/FFT unit, which is capable of achieving a maximum of 830 MHz, which is roughly 1/7 of the 6 GHz Nyquist input clock. The designed A2I converter front-end delivers the non-uniform samples to the on-chip spectrum reconstruction stage, which in turn recovers a 2^{15} -bin spectrum with a spectral resolution of 183 kHz per bin. The effective recovered spectral bandwidth of the proposed A2I converter is 3 GHz. The power estimates of the analog front-end were obtained from transistor-level simulations using Cadence Virtuoso Spectre. The indicated area estimates for the digital part were reported by Synopsys Design Compiler after synthesis of the design. In addition, a coarse power estimate based on extracted switching activity from functional simulations was obtained from Mentor Graphics Modelsim and Synopsys Design Compiler.

F. Comparison to Existing NUS-based A2I Converters

We now compare the proposed monolithic A2I converter to other NUS-based A2I converter systems reported in literature.

Pfetsch *et al.* [52] propose a system based on off-the-shelf components. An FPGA is used to control an ADC via pre-calculated pseudo-random clock signals. Spectrum recovery is carried out on a DSP for bandwidths in the kHz regime.

Wakin *et al.* [11] report a system built around a custom, high-speed sample-and-hold stage (S&H) in a 0.45 μm InP HBT technology. The S&H is controlled by a pseudo-random clock signal generated off-chip. The quantization is performed by an off-the-shelf 14-bit 400 MS/s ADC. Spectrum recovery is realized off-line on a PC using a two-stage recovery method based on a 2^{16} -point FFT. The design achieves a 2.4 GHz effective bandwidth resulting in a spectral resolution of 73 kHz.

Trakimas *et al.* [12] demonstrate an integrated A2I sampling and quantization front-end in 90 nm CMOS. The design consists of a S&H stage and a 10-bit successive approximation register (SAR) ADC that is clocked asynchronously via a pseudo-random clock signal generated off-chip. Spectrum recovery up to 100 MHz of bandwidth is performed off-line.

With respect to the above implementations, our A2I converter encompasses all necessary components, i.e., a non-uniform sampling clock generator, an ADC stage, and a digital spectrum recovery unit. Moreover, our design is capable of acquiring RF signals in the GHz range using state-of-the-art CMOS technology, which is in stark contrast to the design in [11] that relies on expensive InP HBT technology.

We note that other A2I converter systems, such as the ones in [53], [54], remain at the conceptual stage of development, without actually providing circuit results. A corresponding performance analysis and comparison based on real circuit designs is certainly interesting and left for future work.

G. Limitations of the Proposed A2I Converter

The proposed A2I converter suffers from a variety of limitations, which are mainly a direct consequence of CS.

The considered spectrum recovery algorithm is based on the assumption that the RF spectrum is sparsely populated; hence, our approach naturally fails in situations where this condition is not satisfied. In situations where the spectrum is densely populated, one needs to resort to conventional high-rate and energy-inefficient ADCs. Nevertheless, as shown in Figure 5(a), the detection performance degrades *gracefully* for an increasing number of active frequency bands.

Virtually all CS recovery algorithms require a set of algorithm parameters, which are, in general, difficult to determine in practice. In our case, we deployed extensive simulations to determine the underlying parameter set. The development of a principled way of setting these parameters or even adapting them to the signal sparsity is left for future work.

Even though we deploy a highly optimized ASSD unit, our design recovers RF spectra at rates roughly $30\times$ below real-time. A corresponding real-time implementation would require a substantial investment in silicon area (e.g., by means of parallel recovery unit instances). However, such a brute-force solution would increase the circuit area and power consumption by about $30\times$ and, hence, would be rather unsuitable for applications targeting low cost, low area, and low power.

Finally, we emphasize that accurately detecting weak active frequency bands suffers from a fundamental dynamic range reduction problem [55] and thus, remains challenging. A possible way to overcome this limitation would be to aggregate a series of reconstructed RF spectra to improve the performance at low SNR, similarly to [11]; such an approach, however, sacrifices temporal resolution for dynamic range, and increases the storage requirements and processing latency.

VI. CONCLUSIONS

In this work, we have reported the design of a monolithic, wideband, CS-based analog-to-information converter for spectrum sensing applications. The proposed A2I converter is designed in 28 nm CMOS and contains a 3 GHz signal acquisition stage built from a 4-bit flash ADC that samples the time-domain signals at sub-Nyquist rates using a non-uniform sampling clock generated directly on-chip. The RF spectral activity is recovered on-chip from coarsely quantized and compressive measurements by means of a novel accelerated sparse signal dequantization (ASSD) algorithm. To achieve a high recovery throughput, we have developed a corresponding high-throughput VLSI architecture relying on a massively parallel radix-32 inverse/forward fast Fourier transform (I/FFT) unit. System simulations with synthetic and real-world data have shown that the proposed design is capable of accurately detecting sparse spectral activity information at low implementation cost.

The proposed monolithic A2I converter demonstrates a potential paradigm shift in the design of modern signal conversion circuits. In particular, CS with coarse quantization enables one to reduce the complexity and implementation effort in the analog front-end at the price of a more complex digital recovery circuit. Such an approach is of particular interest in advanced CMOS technologies, where the design of corresponding analog circuits becomes more and more

challenging; in contrast, standard-cell based digital design fully benefits from technology scaling, and digital logic will become even more inexpensive. In addition, the deployment of sophisticated digital signal processing algorithms enables one to compensate for quantization artifacts or mismatches/non-idealities of analog front-ends implemented in nanometer CMOS technologies.

We conclude by noting that alternatives to CS-based A2I converter designs for spectrum sensing have been proposed in the literature, such as energy detection [56], matched filter detection [57], cyclostationary feature detection [57], multi-resolution spectrum sensing [58], or Nyquist folding receivers [59]. All these methods exploit specific features of the underlying communication signals (e.g., periodicity) to detect the active frequency components. The consideration of such signal properties has the potential to further improve the sensitivity of spectrum sensing; a thorough investigation of such methods in combination with the proposed A2I converter design is an interesting open research topic.

REFERENCES

- [1] S. Haykin, "Cognitive radio: brain-empowered wireless communications," *IEEE J. Sel. Areas in Commun.*, vol. 23, no. 2, pp. 201–220, Feb. 2004.
- [2] S. Cherry, "Edholm's law of bandwidth," *IEEE Spectr.*, vol. 41, no. 7, pp. 58–60, Jul. 2004.
- [3] IEEE 802.22 Working Group on Wireless Regional Area Networks. [Online]. Available: <http://www.ieee802.org/22/>
- [4] D. Healy, "Analog-to-information (A-to-I)," *DARPA/MTO Broad Agency Announcement BAA*, pp. 5–35, Jul. 2005.
- [5] S. Park, Y. Palaskas, and M. P. Flynn, "A 4-GS/s 4-bit Flash ADC in 0.18- μm CMOS," *IEEE J. Solid-State Circuits*, vol. 42, no. 9, pp. 1865–1872, Sep. 2007.
- [6] K. Patil, K. Skouby, A. Chandra, and R. Prasad, "Spectrum occupancy statistics in the context of cognitive radio," in *Proc. Int. Symp. Wireless Pers. Multimedia Commun.*, Oct. 2011, pp. 1–5.
- [7] M. A. McHenry, P. A. Tenhula, D. McCloskey, D. A. Roberson, and C. S. Hood, "Chicago spectrum occupancy measurements & analysis and a long-term studies proposal," in *Proc. Int. Workshop Technol. Policy for Accessing Spectrum*, Aug. 2006.
- [8] E. Candès and M. Wakin, "An introduction to compressive sampling," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 21–30, Mar. 2008.
- [9] M. Davenport, J. Laska, J. Treichler, and R. Baraniuk, "The pros and cons of compressive sensing for wideband signal acquisition: Noise folding versus dynamic range," *IEEE Trans. Signal Process.*, vol. 60, no. 9, pp. 4628–4642, Sep. 2012.
- [10] M. Davenport, S. Schnelle, J. P. Slavinsky, R. Baraniuk, M. Wakin, and P. Boufounos, "A wideband compressive radio receiver," in *Proc. Military Commun. Conf.*, Oct. 2010, pp. 1193–1198.
- [11] M. Wakin, S. Becker, E. Nakamura, M. Grant, E. Sovero, D. Ching, J. Yoo, J. Romberg, A. Emami-Neyestanak, and E. Candès, "A nonuniform sampler for wideband spectrally-sparse environments," *IEEE J. Emerging Sel. Topics Circuits Syst.*, vol. 2, no. 3, pp. 516–529, Sep. 2012.
- [12] M. Trakimas, R. D'Angelo, S. Aeron, T. Hancock, and S. Sonkusale, "A compressed sensing analog-to-information converter with edge-triggered SAR ADC core," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 60, no. 5, pp. 1135–1148, May 2013.
- [13] J. Yoo, S. Becker, M. Monge, M. Loh, E. Candès, and A. Emami-Neyestanak, "Design and implementation of a fully integrated compressed-sensing signal acquisition system," in *IEEE Int. Conf. Acoust., Speech, Signal Process.*, Mar. 2012, pp. 5325–5328.
- [14] P. Maechler, "VLSI architectures for compressive sensing and sparse signal recovery," Ph.D. dissertation, ETH Zürich, Switzerland, 2013.
- [15] P. T. Boufounos and R. G. Baraniuk, "1-bit compressive sensing," in *Proc. Annu. Conf. Info. Science Syst.*, Mar. 2008, pp. 16–21.
- [16] Y. Plan and R. Vershynin, "Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach," *IEEE Trans. Inf. Theory*, to appear.
- [17] J. N. Laska and R. G. Baraniuk, "Regime change: Bit-depth versus measurement-rate in compressive sensing," *arXiv:1110.3450v1*, Oct. 2011.
- [18] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. Imaging Sciences*, vol. 2, no. 1, pp. 183–202, Jan. 2009.
- [19] A. Zymnis, S. Boyd, and E. Candès, "Compressed sensing with quantized measurements," *IEEE Signal Process. Lett.*, vol. 17, no. 2, pp. 149–152, Feb. 2010.
- [20] T. Ragheb, J. N. Laska, H. Nejati, S. Kirolos, R. G. Baraniuk, and Y. Massoud, "A prototype hardware for random demodulation based compressive analog-to-digital conversion," in *Proc. Midwest Symp. Circuits Syst.*, Aug. 2008, pp. 37–40.
- [21] J. N. Laska, S. Kirolos, M. F. Duarte, T. S. Ragheb, R. G. Baraniuk, and Y. Massoud, "Theory and implementation of an analog-to-information converter using random demodulation," in *IEEE Int. Symp. Circuits Syst.*, May 2007, pp. 1959–1962.
- [22] 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. IEEE Dallas/CAS Workshop on Design, Applications, Integration and Software*, Oct. 2006, pp. 71–74.
- [23] Y. Massoud, S. Smaili, and V. Singal, "Efficient realization of random demodulator-based analog to information converters," in *Proc. IEEE Biomed. Circuits Syst. Conf.*, Nov. 2011, pp. 133–136.
- [24] M. Mishali, Y. Eldar, O. Dounaevsky, and E. Shoshan, "Xampling: Analog to digital at sub-Nyquist rates," *IET Circuits, Devices & Syst.*, vol. 5, no. 1, pp. 8–20, Jan. 2011.
- [25] M. Mishali and Y. C. Eldar, "Sub-nyquist sampling," *IEEE Signal Process. Mag.*, vol. 28, no. 6, pp. 98–124, Nov. 2011.
- [26] M. Mishali and Y. 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, Mar. 2010.
- [27] S. R. Becker, "Practical compressed sensing: modern data acquisition and signal processing," Ph.D. dissertation, California Institute of Technology, 2011.
- [28] M. Ben-Romdhane, C. Rebai, A. Ghazel, P. Desgreys, and P. Loumeau, "Pseudorandom clock signal generation for data conversion in a multi-standard receiver," in *Proc. Int. Conf. Design Technol. Integrated Syst. Nanoscale Era*, Mar. 2008, pp. 1–4.
- [29] E. Allier, G. Sicard, L. Fesquet, and M. Renaudin, "Asynchronous level crossing analog to digital converters," *Measurement*, vol. 37, no. 4, pp. 296–309, Jun. 2005.
- [30] P. Maechler, N. Felber, and A. Burg, "Random sampling ADC for sparse spectrum sensing," in *Proc. Europ. Signal Process. Conf.*, Sep. 2011, pp. 1200–1204.
- [31] D. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [32] E. 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.
- [33] M. Rudelson and R. Vershynin, "On sparse reconstruction from Fourier and Gaussian measurements," *Commun. Pure Appl. Math.*, vol. 61, no. 8, pp. 1025–1045, Aug. 2008.
- [34] J. Tropp and S. Wright, "Computational methods for sparse solution of linear inverse problems," *Proc. IEEE*, vol. 98, no. 6, pp. 948–958, Jun. 2010.
- [35] A. Maleki and D. L. Donoho, "Optimally tuned iterative reconstruction algorithms for compressed sensing," *IEEE J. Sel. Topics Signal Process.*, vol. 4, no. 2, pp. 330–341, Apr. 2010.
- [36] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Sci. Comput.*, vol. 20, no. 1, pp. 33–61, Aug. 1998.
- [37] Rice University. 1-bit Compressive sensing. [Online]. Available: <http://dsp.rice.edu/1bitCS/>
- [38] G. Pope, C. Studer, and M. Baes, "Coherence-based recovery guarantees for generalized basis-pursuit de-quantizing," in *Proc. IEEE Conf. Acoust., Speech, Signal Process.*, May 2012, pp. 3669–3672.
- [39] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge University Press, 2004.
- [40] P. Maechler, C. Studer, D. E. Bellasi, A. Maleki, A. Burg, N. Felber, H. Kaeslin, and R. G. Baraniuk, "VLSI design of approximate message passing for signal restoration and compressive sensing," *IEEE J. Emerging Sel. Topics Circuits Syst.*, vol. 2, no. 3, pp. 579–590, Sep. 2012.
- [41] S. Shahramian, S. P. Voignescu, and A. C. Carusone, "A 35-GS/s, 4-Bit Flash ADC With Active Data and Clock Distribution Trees," *IEEE J. Solid-State Circuits*, vol. 44, no. 6, pp. 1709–1720, Jun. 2009.

- [42] S. Verma, A. Kasapi, L.-m. Lee, D. Liu, D. Loizos, S.-H. Paik, A. Varzaghani, S. Zogopoulos, and S. Sidiropoulos, "A 10.3GS/s 6b Flash ADC for 10G Ethernet Applications," in *Dig. Tech. Papers IEEE Int. Solid-State Circuits Conf.*, Feb. 2013, pp. 462–463.
- [43] L. Kull, T. Toifi, M. Schmatz, P. A. Francese, C. Menolfi, M. Braendli, M. Kossel, T. Morf, T. M. Andresen, and Y. Leblebici, "A 3.1mW 8b 1.2GS/s Single-Channel Asynchronous SAR ADC with Alternate Comparators for Enhanced Speed in 32nm Digital SOI CMOS," in *Dig. Tech. Papers IEEE Int. Solid-State Circuits Conf.*, Feb. 2013, pp. 468–469.
- [44] E. Janssen, K. Doris, A. Zanicopoulos, A. Murrioni, G. van der Weide, Y. Lin, L. Alvado, F. Darthenay, and Y. Fregeais, "An 11b 3.6GS/s Time-Interleaved SAR ADC in 65nm CMOS," in *Dig. Tech. Papers IEEE Int. Solid-State Circuits Conf.*, Feb. 2013, pp. 464–465.
- [45] M. Bolatkale, L. J. Breems, R. Rutten, and K. A. Makinwa, "A 4 GHz Continuous-Time $\Delta\Sigma$ ADC With 70 dB DR and 74 dBFS THD in 125 MHz BW," *IEEE J. Solid-State Circuits*, vol. 46, no. 12, pp. 2857–2868, Dec. 2011.
- [46] A. Abo and P. Gray, "A 1.5-V, 10-bit, 14.3-MS/s CMOS pipeline analog-to-digital converter," *IEEE J. Solid-State Circuits*, vol. 34, no. 5, pp. 599–606, May 1999.
- [47] D. Schinkel, E. Mensink, E. Kiumperink, E. van Tuijl, and B. Nauta, "A double-tail latch-type voltage sense amplifier with 18ps setup+ hold time," in *Dig. Tech. Papers IEEE Int. Solid-State Circuits Conf.*, Feb. 2007, pp. 314–315.
- [48] T. Christen and Q. Huang, "A 0.13 μm CMOS 0.1-20 MHz bandwidth 86-70 dB DR multi-mode DT $\Delta\Sigma$ ADC for IMT-advanced," in *Proc. IEEE Europ. Solid-State Circuits Conf.*, Sep. 2010, pp. 414–417.
- [49] S. Saponara, M. Rovini, L. Fanucci, A. Karachalios, G. Lentaris, and D. Reisis, "Design and comparison of FFT VLSI architectures for SoC telecom applications with different flexibility, speed and complexity trade-offs," *Circuits, Syst., Signal Process.*, vol. 31, no. 2, pp. 627–649, 2012.
- [50] E. O. Brigham, *The Fast Fourier Transform and its applications*. Prentice Hall, 1988.
- [51] S.-Y. Lin, C.-L. Wey, and M.-D. Shieh, "Low-cost FFT processor for DVB-T2 applications," *IEEE Trans. Consumer Electronics*, vol. 56, no. 4, pp. 2072–2079, Nov. 2010.
- [52] S. Pfetsch, T. Ragheb, J. Laska, H. Nejati, A. Gilbert, M. Strauss, R. Baraniuk, and Y. Massoud, "On the feasibility of hardware implementation of sub-nyquist random-sampling based analog-to-information conversion," in *Proc. IEEE Int. Symp. Circuits Syst.*, May 2008, pp. 1480–1483.
- [53] C. Luo and J. H. McClellan, "Compressive sampling with a successive approximation ADC architecture," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, may 2011, pp. 3920–3923.
- [54] J. Laska, S. Kirolos, Y. Massoud, R. Baraniuk, A. Gilbert, M. Iwen, and M. Strauss, "Random sampling for analog-to-information conversion of wideband signals," in *IEEE Dallas/CAS Workshop on Design, Applications, Integration and Software*, Oct. 2006, pp. 119–122.
- [55] D. Cabric, S. M. Mishra, and R. W. Brodersen, "Implementation issues in spectrum sensing for cognitive radios," in *Proc. IEEE Asilomar Conf. Signals, Syst. Comput.*, vol. 1, Nov. 2004, pp. 772–776.
- [56] Z. Quan, S. Cui, A. H. Sayed, and H. V. Poor, "Wideband spectrum sensing in cognitive radio networks," in *Proc. IEEE Int. Conf. Commun.*, May 2008, pp. 901–906.
- [57] D. Bhargavi and C. Murthy, "Performance comparison of energy, matched-filter and cyclostationarity-based spectrum sensing," in *Proc. IEEE Int. Workshop on Signal Process. Advances in Wireless Commun.*, Jun. 2010, pp. 1–5.
- [58] J. Park, T. Song, J. Hur, S. M. Lee, J. Choi, K. Kim, K. Lim, C.-H. Lee, H. Kim, and J. Laskar, "A Fully Integrated UHF-Band CMOS Receiver With Multi-Resolution Spectrum Sensing (MRSS) Functionality for IEEE 802.22 Cognitive Radio Applications," *IEEE J. Solid-State Circuits*, vol. 44, no. 1, pp. 258–268, Jan. 2009.
- [59] G. L. Fudge, R. E. Bland, M. A. Chivers, S. Ravindran, J. Haupt, and P. Pace, "A Nyquist folding analog-to-information receiver," in *Proc. IEEE Asilomar Conf. Signals, Syst. Comput.*, Oct. 2008, pp. 541–545.