Rabia Tugce Yazicigil, Tanbir Haque, Peter R. Kinget, and John Wright

Taking Compressive Sensing to the Hardware Level

Breaking fundamental radio-frequency hardware performance tradeoffs

Compressive sensing (CS) theory opens promising avenues toward building rapid and energy-efficient sensing systems in a wide range of applications that require inherently high temporal and/or spatial resolution while exhibiting a sparse signal structure [1]–[5]. The goal of this article is to review recent efforts to realize the benefits of CS in custom sensing hardware and the broad challenges that arise by investigating an example application in radio-frequency (RF) communications. We discuss in detail how using CS for the design of RF spectrum scanners can break through the fixed tradeoffs among scan time, hardware complexity, and energy consumption of traditional scanner architectures. Using the specific example of RF spectrum sensing [4], [6]–[11] we demonstrate how close collaborations between hardware and signal processing experts can yield new solutions that advance the state of the art in an important application domain. We stress the problems that arise when designing a custom hardware for CS [12]–[15] and address questions that often go beyond the currently available literature in CS, e.g., coping with the impairments of real hardware and avoiding catastrophic breakdown when the spectrum becomes nonsparse.

CS facilitates a sampling theory that allows for signal recovery using fewer measurements than required by the Shannon-Nyquist sampling theorem [16]–[21]. Previous hardware solutions based on the Shannon-Nyquist sampling theorem suffer from fixed tradeoffs between scan time and power consumption. Moreover, trading power for reduced scan time, which can significantly increase the complexity of the hardware, does not improve its overall energy consumption. In contrast, CS leverages the usage of signal structure, such as sparsity or compressibility in an appropriately chosen transform basis, to enable new tradeoffs with substantial improvements to both time and energy per scan.

Applications requiring a high-resolution acquisition of structured sparse signals are ideal candidates for CS. Some obvious application fields that could benefit from CS include medical diagnostics [1], [22], weather forecasting [23], [24], wireless communications [2]–[4], [6], [8], [10], [25], autonomous vehicles [26]–[28], and video systems [5].

Identifying essential application characteristics for CS utilization

CS has been widely used for RF communications [2], [3], including interferer detection [8], signal reception [4], [10], [25],

FIGURE 1. An example of a sparse signal in the frequency domain consisting of multitone and multiband signal models [11], [21] for RF signal reception and spectrum-sensing applications. Active tones and bands above the signal-level detection threshold are shown in green, while the inactive tones and bands are shown in blue.

spectrum sensing [4], [6], [10], [11], direction estimation [29], channel estimation, and symbol detection [2]. When compared to traditional RF communications, applying CS to RF communications can reduce the scan time for a desired resolution or increase the resolution while maintaining the scan time. Furthermore, CS introduces a compression in hardware via signal sparsity in contrast to traditional RF architectures with the same scan time, resulting in reduced hardware complexity.

For example, in RF spectrum sensing, CS-based approaches achieve scan times equivalent to multibranch architectures with many more parallel branches, e.g., an order of magnitude [4], [8]. The signal of interest is sparse in the frequency domain, which can be modeled with a few nonzero bands or tones in a wideband spectrum, as shown in Figure 1 [4], [8], [11], [25], [30]. We assume that the spectrum has been composed of two clusters, i.e., the lower and upper bands, with each divided into N_0-frequency bins. The signal of interest $X(f)$ is only supported on $K = 2K_0$ bands or tones out of $2N_0$ bins, and $X(f)$ is sparse in the frequency domain because $K_0 < N_0$. These K bands or tones, illustrated in green, exceed a predefined threshold level, while the inactive signals are illustrated in blue.

A second example of compression in hardware through CS pertains to reducing the number of antennas in a direction-of-arrival (DOA)-finding application. Current DoA-finding approaches have a fixed tradeoff between scan time and resolution with limited scaling in terms of the number of antennas and energy consumption. For this particular application, the signal of interest is sparse in the spatial domain across antenna elements. The power-versus-angle plot shown in Figure 2 represents the angle spectrum that corresponds to incident signal power at a particular angle θ. Applying CS to calculate the DoA of a spatially sparse signal has been demonstrated in [29], resulting in reduced energy consumption and scan time while reducing the number of antenna elements required.

To glean the benefits of CS, identifying a proper entry point in the application space is crucial. Consider two possible applications for using CS in the RF communications space: interferer detection, which only targets a few large signals that are above a detection threshold, and signal reception, which targets both weak and strong signals. Both applications typically have a sparse signal structure in the frequency domain, which positions them as potential candidates for CS. However, the limitations of CS become more prominent for the signal-reception application because, for example, noise folding [12], [13] limits the sensitivity of the system, thus resulting in signal-reception degradation. In contrast, the interferer-detection application is not severely impacted by noise folding because only strong signals are targeted. Furthermore, CS approaches for signal reception require obtaining the solution of very large optimization problems in the digital signal processing (DSP) domain. On the other hand, we will see that interferer detection leads to relatively small problems that can be solved rapidly. Taking further advantage of signal features specific to the use case, such as employing a bandpass CS approach for RF signal detection proposed in [8], can improve the sensitivity and greatly reduce the overall energy consumption by reducing the front-end power and also by compressing the system-matrix dimension for sparse recovery in contrast to a low-pass CS approach [4]. To illustrate how to exploit signal features specific to the interest scenario, we now briefly discuss mapping CS theory to hardware systems in general and highlight some of the common challenges associated with this process.

Bridging CS theory and practical hardware implementation

The goal of CS is to uniquely determine the sparse signal of interest, such as interferers or DoA from our examples, from an
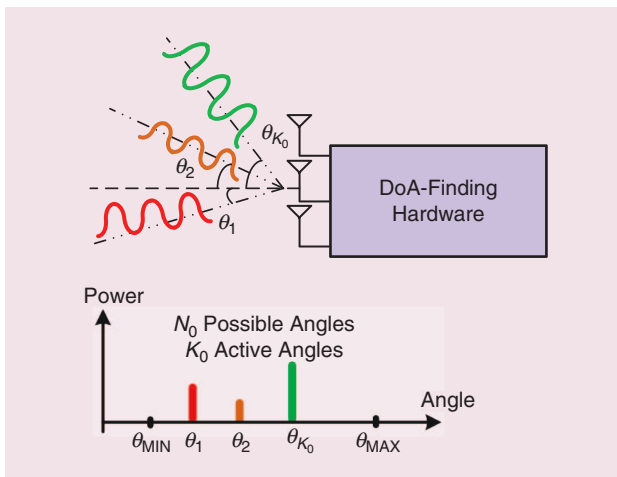

FIGURE 2. An example of a spatially sparse signal in a DoA-finding application.

undetermined system of linear equations. CS theory [16]–[21] solves a sensing problem in the form of $y = [A]z = [\phi][\psi]z$, where $y := y(m, 1)$ is the measurement vector, $[\phi] := \phi(m, N_0)$ is the sensing basis, $[\psi] := \psi(N_0, N_0)$ is the dictionary basis, $z := z(N_0, 1)$ is the unknown sparse coefficient vector with dimension m being the number of measured samples, and dimension N_0 being the number of Nyquist-rate samples. The matrix $[\psi]$ is used to define an input signal x as a sparse vector such that $x = [\psi]z$ contains only a few nonzero coefficients, i.e., K_0. Because the number of measured samples m is smaller than the number of unknowns N_0, there is an infinite number of admissible solutions. Fortunately, CS theory shows that if 1) the target solution z is sufficiently sparse and 2) the matrix $[A]$ is well-structured, then z (and hence, x) can be uniquely determined and efficiently recovered from the measurements, i.e., y.

The strongest available theoretical guarantees for CS pertain to random sensing bases ϕ (or random matrices $[A]$). For example, random matrices satisfy the restricted isometry property (RIP) as soon as the number of measurements m is sufficiently large. This implies that convex relaxations exactly recover z in the absence of noise, and stably estimate z in the presence of noise. Random matrices also provide guarantees for efficient greedy methods such as orthogonal matching pursuit (OMP) [31]. While many different matrix ensembles satisfy the RIP, one which is especially amenable to hardware implementation is the Rademacher ensemble, with independent ±1 entries [32].

Figure 3 illustrates how a CS sensing problem maps to a hardware system. A CS sensor is composed of an analog acquisition block and a digital back end. The analog acquisition block takes the input x from real-world scenarios such as RF signals and operates on x with $[\phi]$ by taking a linear random projection to generate y digital samples. Then, y samples are fed into the digital back end, which solves for z such that $x = [\psi]z$ by using an OMP [31], [33] or ℓ_1 minimization [34], [35].

One common challenge in taking CS to the hardware level is how to develop incoherent measurements that 1) are compatible with physical hardware and 2) reflect the sparsity structure encountered in signals of interest, i.e., the probable locations of nonzero entries. For example, in [36], the proposed signal-structure-aware CS-sampling scheme utilizes the Hadamard transform of the input signal, and hence, this scheme performs better in practice compared to the random Gaussian and Bernoulli schemes with a small circuit area. To discuss how to optimize the incoherent measurements in a CS system given the application scenario and key system performance metrics, we will use the specific example of RF communications for the remainder of this article [4], [8], [11], [13], [14], [25], [30]. Figure 4 shows an example of a CS RF hardware used to acquire an incoherent measurement for successfully detecting sparse signals in the frequency domain. Figure 4 also highlights the construction of the incoherent measurements through the dot product between the ith-row of the sensing matrix $[\phi]$, which has a spectral representation $P_i(f)$ shown as the black sinc-shaped curve, and the input vector x formed by the green and red multibands. For the specific example scenario of RF communications, we survey various implementations of the CS analog acquisition block known as the *analog-to-information converter* (*AIC*). We also briefly discuss the key components and energy consumption of the digital back end for those AIC implementations. Design considerations and challenges associated with the CS hardware implementation for RF communications are similar to the implementation challenges in a wide range of other CS applications as well.

RF spectrum sensors: An ideal application for CS

When comparing traditional hardware implementations for RF signal reception [10], [11], [25], [30], [37] and spectrum sensing [4], [7], [8], [11], [14] to CS-based implementations, we first introduce the signal model and spectrum scenario followed by key performance metrics used to evaluate such systems. Once established, we survey multiple traditional and CS implementations to demonstrate the potential of CS to fundamentally improve tradeoffs between key design considerations, including scan time, energy consumption, and hardware complexity [4], [8], [10], [25]. Finally, we discuss some difficulties and challenges associated with
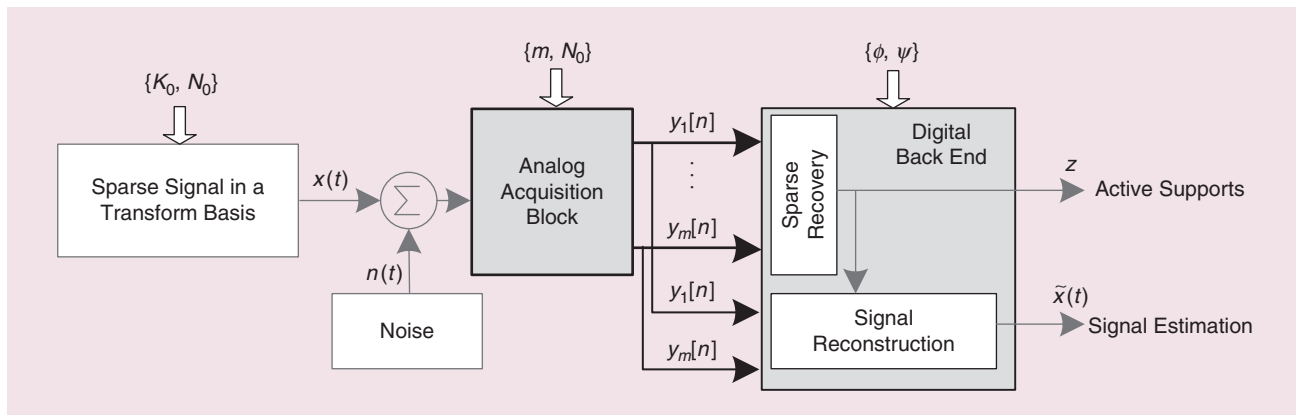

FIGURE 3. A CS sensing system consists of an analog acquisition block and a digital back end. The analog acquisition block performs a linear random projection of the K_0-sparse input signal with noise on the ϕ basis and outputs y-digital samples where the number of measurements, m, is significantly smaller than the number of unknowns N_0. The digital back end utilizes the random-sensing-basis ϕ and the dictionary-basis ψ to reconstruct the sparse signal.

CS implementations and how to circumvent issues such as implementing random (incoherent) measurements in a hardware- and energy-efficient manner, coping with the impairments of real hardware, and avoiding a catastrophic breakdown when the spectrum is not sparse but becomes densely occupied.

RF signal model and spectrum scenario

Signal structure and model

■ *Sparsity*: In many real-world applications, signals of interest are sparse or can be defined as a sparse vector in a well-chosen basis. A sparse vector with a dimension of N_0 has K_0 nonzero coefficients where $K_0 < N_0$. The sparsity level is defined as $S = K_0/N_0$ [14].

Sparsity is the prior information or assumption that can be leveraged to simplify the acquisition process by uniquely determining the original spectrum of a band-limited signal without requiring Nyquist-rate sampling of the instantaneous bandwidth (BW) [17].

■ *Multitone signal model*: Considers the following mathematical model for a class of discrete multitone signals [21], [38]:

1) The Fourier transform $X(f)$ of the continuous-time signal $x(t)$ has the highest frequency component below N_0, where N_0 is a positive integer.
2) There are only $K = 2K_0$ active tones, where K is substantially smaller than the instantaneous BW of $N = 2N_0$.

For each time interval normalized to a second, the signal model in [21] is given by

$$x(t) = \sum_{\omega \in \Omega} a_\omega e^{-2\pi i \omega t}, \qquad (1)$$

for $t \in [0, 1)$, where Ω is a set of K-sparse integer-valued frequencies that satisfies $\Omega \subset \{-N_0 + 1, ..., -1, 0, 1, ..., N_0\}$,
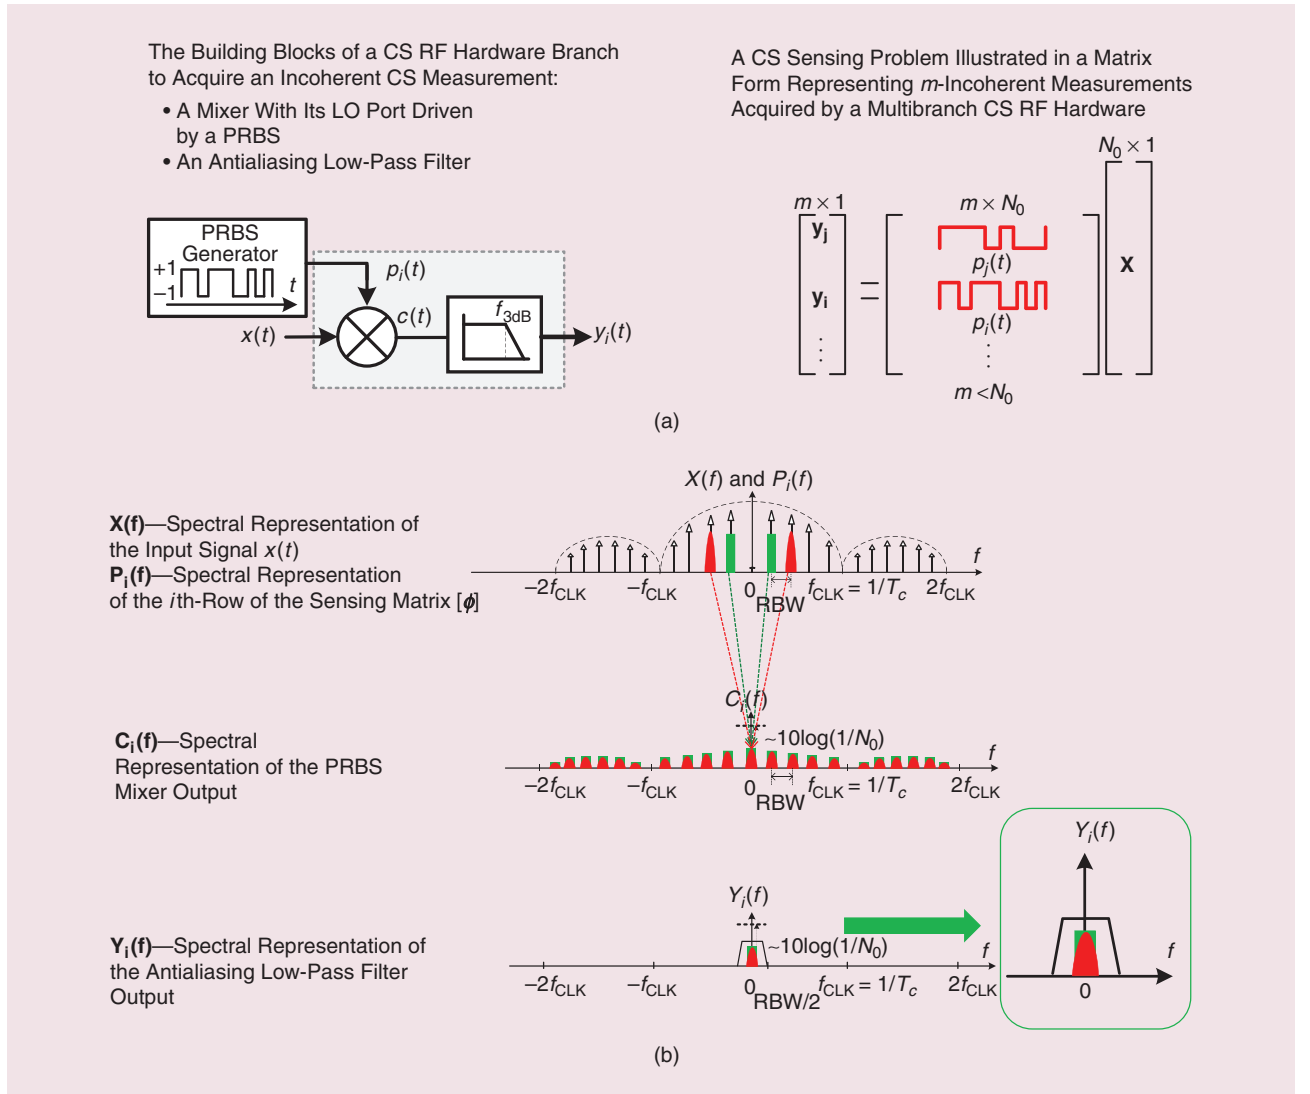

FIGURE 4. (a) The example hardware implementation of one branch of an RF analog-to-information converter providing one incoherent measurement $y_i[n]$. (b) An illustration using spectra shows how this incoherent measurement is obtained using a "well-conditioned" sensing matrix. LO: local oscillator; PRBS: pseudorandom binary sequence; RBW: resolution BW.

and $\{a_\omega : \omega \in \Omega\}$ is a set of complex-valued amplitudes, such as the Fourier series coefficients of $x(t)$. Figure 1 illustrates these assumptions.

- *Multiband signal model*: Considers sparse multiband signals [13], [38], [39] that are real-valued, square-integrable signals $x(t)$ satisfying two properties.
  1) The signal of interest has a valid Fourier transform $X(f)$ in the frequency range of $\mathcal{F} = (-f_{\text{MAX}}, -f_{\text{MIN}}] \bigcup [f_{\text{MIN}}, f_{\text{MAX}})$.
  2) The signal of interest is sparse in the frequency domain, meaning that the support of $X(f)$ is a relatively small subset of $\mathcal{F}$.

As shown in Figure 1, for a real-valued signal $x(t)$, $\mathcal{F}$ has been partitioned into $N = 2N_0$ disjoint bands with a resolution BW (RBW) of B. The sparse Fourier transform $X(f)$ is supported on only $K = 2K_0 < N = 2N_0$ of these bands. These $2K_0$ bands are referred to as *active bands* or *supports* when their power level is above a predefined or adaptive signal-level detection threshold.

If the spectral occupancy $S_0 = K_0/N_0$ is small, then the support of $X(f)$ has a Lebesgue measure [40] $\leq K_0 B$ Hz, which is much smaller than the Nyquist rate of the instantaneous BW $f_{Nyq} = 2f_{\text{MAX}}$ Hz [13].

Spectrum scenario

For the remainder of this article, we consider a case study of RF spectrum sensing or signal reception in a 1 GHz of interest spectrum band, e.g., the President's Council of Advisors on Science and Technology (PCAST) band ranging from 2.7 to 3.7 GHz with a 20-MHz RBW resulting in 50 spectrum bins, i.e., $N_0 = 50$ [8], [14]. In this scenario, there are three large signals, i.e., interferers that are above the predefined or adaptive signal-level detection threshold, i.e., $K_0 = 3$. The sparsity level for this scenario is $S = 3/50$, which indicates a spectrum occupancy of 6%. The $K_0 = 3$ large signals can be located in any of the $N_0 = 50$ spectrum bins. The goal is to efficiently acquire a one-dimensional spectrum image to locate those K_0 large signals, even if the spectral locations of the K_0 supports are not known in advance.

RF spectrum-sensing performance metrics

In this section, we provide an overview of key system performance metrics of the RF spectrum scanners and sensors that are designed to detect the K_0 large signals located in any of the N_0 spectrum bins. The key system performance metrics are the instantaneous BW, the scan time and energy consumption required to capture the information in that instantaneous BW, and the instantaneous dynamic range (DR), while simultaneously satisfying target detection and false alarm probabilities. An overview of the key system performance metrics discussed in this section is summarized in Table 1. These performance metrics also dictate the natural fit between a point in the application space and the choice of architecture.

- *Energy consumption*: Energy consumption for a scan is defined as the power consumption, P, times the scan time,

T_{scan}. The scan time is composed of two parts: front-end detector response time T_{resp} and DSP time T_{rec} [8].

$$E = P \cdot T_{\text{scan}} = P \cdot [T_{\text{resp}} + T_{\text{rec}}]. \qquad (2)$$

- *Scan time (T_{scan})*: This is defined as the combination of detector response time T_{resp} and DSP time T_{rec}. DSP time is proportional to N_s/f_s, where f_s is the analog-to-digital converter (ADC) sampling rate, and N_s is the number of samples collected. Front-end detector response time is proportional to the settling time of the low-pass antialiasing filters [8].
- *Detection and false alarm probabilities*: Detection probability P_D is the probability that a CS AIC correctly reports a signal in the RF spectrum as active. False alarm probability P_{FA} is the probability that a CS AIC reports a spectrum bin as occupied when there is no signal present in the RF spectrum [8].
- *Instantaneous BW*: The instantaneous BW of a spectrum sensor is defined as its *Span*, over which few signals can be successfully and rapidly detected, while meeting target detection and false alarm probabilities.
- *DR*: The instantaneous DR of a spectrum sensor is defined as its ability to successfully detect a weak signal in the presence of a strong signal or signals over a wide instantaneous BW, i.e., its *Span* [8], [13]. We note that this definition of DR is suitable for real-world applications when there are strong and weak signals present at the same time over a wide BW. One of the challenges of instantaneous DR is that the weakest signals do not enjoy maximum gain for a given full scale of the ADC. The reason for this challenge is the automatic gain control block prior to the ADC is typically set for the strongest signal that fits into the full scale of the ADC.

In addition to instantaneous DR, operational range is used when there are equal-power strong or weak signals present over a wide BW. Operational DR is the best-case DR performance [25]. In this case, the automatic gain control block maximizes the conversion gain to fit into the full scale of the ADC for both strong and weak signal scenarios.

- *Number of detectable signals*: The number of detectable signals is a valid performance metric for only CS spectrum sensors, and it depends heavily on the number of incoherent measurements collected from a CS AIC. The relation between the number of incoherent measurements, m, and the number of detectable signals, K_0, in a CS AIC is given by

$$m \approx \left\lceil C \cdot K_0 \cdot \log\left(\frac{N_0}{K_0}\right) \right\rceil, \qquad (3)$$

where K_0 is the number of detectable signals, N_0 is the number of spectrum bins in the range of interest, and C is a constant in the range of two to four [8], [13].

RF spectrum sensors: Traditional architectures

In this section, we survey the traditional spectrum sensor and scanner architectures [8]. Furthermore, we discuss their

advantages and limitations in terms of the key system performance metrics discussed in the "RF Spectrum-Sensing Performance Metrics" section.

Spectrum sensor and scanner architectures

The architectures can be organized into three main categories.

- *Sweeping spectrum scanners*: In a single-branch sweeping spectrum scanner, each bin is scanned sequentially by time-multiplexing the hardware; this is a slow approach, especially for fine-frequency resolutions. This approach results in large energy consumption, and there is also a risk of missing the changes in spectrum dynamics. Sweeping spectrum scanners offer a high DR for a small instantaneous BW equal to their RBW.
- *Nyquist-rate fast Fourier transform (FFT) spectrum sensors*: In a Nyquist-rate FFT spectrum sensor, the interest spectrum band is captured instantaneously by sampling at the Nyquist rate of the instantaneous BW equal to *Span*. The instantaneous BW, i.e., the *Span*, is subdivided into the desired RBW through FFT in the digital back end.

- *Multibranch spectrum sensors*: In a multibranch spectrum sensor, the interest spectrum band is captured instantaneously by deploying N_0 branches, which are equal to the number of spectrum bins to be observed.

In the following, we discuss the key design considerations and potential use cases of these three sensor/scanner types in more detail.

Sweeping spectrum scanners

A traditional sweeping spectrum scanner can be used to sequentially examine all $N_0 = Span/RBW$ bins in a time-multiplexed fashion to find the location of the K_0 signals [8]. Sweeping spectrum scanners have a fixed tradeoff between the resolution they offer and their scan-time performance. Because there is a risk of missing the changes in dynamic spectrum environments due to their long scan times, they are mostly preferred for stationary spectrum environments. A sweeping spectrum scanner architecture is also ideally suited to be used in high-quality measurement equipment since it offers a high dynamic range over a small instantaneous BW equal to its RBW. Cross-correlation spectrum analyzers [41]

Table 1. The best-known capabilities, limitations, and applications of traditional and CS RF spectrum scanners.

Type	Architecture	Application	Advantages	Limitations
Traditional RF spectrum scanners and sensors	Sweeping spectrum scanner	Spectrum analysis	High DR High sensitivity Low power Low hardware complexity Suitable for nonsparse signals	Small instantaneous BW Slow scan
	Nyquist-rate FFT spectrum sensor	Spectrum analysis	Large instantaneous BW Fast scan Suitable for nonsparse signals	High power High hardware complexity
	Multibranch spectrum sensor	Spectrum analysis	Large instantaneous BW Fast scan Suitable for nonsparse signals	High power High hardware complexity
CS spectrum sensors for a multitone signal model	RD [42], [43]	Sub-Nyquist receiver	Large instantaneous BW Fast scan Low power Low hardware complexity	Suitable for only sparse signals Moderate sensitivity Moderate DR
	RMPI [25]	Sub-Nyquist receiver	Large instantaneous BW Fast scan	Suitable for only sparse signals Moderate sensitivity Moderate DR High power
CS spectrum sensors for a multiband signal model	MWC [4], [10]	Sub-Nyquist receiver and spectrum sensing	Large instantaneous BW Fast scan	Suitable for only sparse signals Moderate sensitivity Moderate DR High power
	QAIC [8]	Sub-Nyquist interferer detection	Large instantaneous BW Fast scan Moderate power	Suitable for only sparse signals Moderate sensitivity Moderate DR Residual bands from I/Q linear impairments
	TS-QAIC [14]	Sub-Nyquist interferer detection	Large instantaneous BW Fast scan Moderate power Dynamic signal sparsity	Moderate sensitivity Moderate DR Residual bands from I/Q linear impairments
	DRF2IC [11]	Sub-Nyquist interferer detection, narrowband sensing, and high-sensitivity receiver	Large instantaneous BW Fast scan Moderate power High sensitivity	Moderate DR High hardware complexity

RD: random demodulator; RMPI: random modulation pre-integrator; MWC: modulated wideband converter; DRF2IC: direct RF-to-information converter; QAIC: quadrature analog-to-information converter; I/Q: in-phase/quadrature-phase; TS-QAIC: time-segmented QAIC.

improve noise performance, hence DR, at the expense of increased scan time.

The energy consumption for a sweeping spectrum scanner is given by

$$E_{\text{scanner}} = P_{\text{scanner}} \cdot [N_0 \cdot (C_{\text{settle}}/\text{BW}_{\text{filter}}) + N_0 \cdot (N_s/f_s)]$$
$$= P \cdot (N_0 \cdot T_{\text{scan}}), \tag{4}$$

where the power consumption of a single branch is P_{scanner}. Front-end response time $T_{\text{resp,scanner}}$ is proportional to the number of bins, N_0, and the antialiasing low-pass filter settling time, which is inversely proportional to its bandwidth $\text{BW}_{\text{filter}}$. For the desired power-reading accuracy, C_{settle} represents the number of filter time constants needed. For example, based on our experiments with the Hewlett-Packard 3585A spectrum analyzer, the estimated C_{settle} value is 4. DSP time, $T_{\text{rec,scanner}}$ is inversely proportional to the sampling rate, f_s and proportional to the number of bins, N_0 and the number of samples used for DSP, i.e., N_s [8].

Nyquist-rate FFT spectrum sensors

To sense signals over GHzs or wider instantaneous BWs, a Nyquist-rate FFT spectrum sensor would require a prohibitively high-aggregate-sampling rate, resulting in a power-hungry sensing approach [8]. Scan time is reduced at the expense of high power consumption, especially for a wide instantaneous BW.

DR for a Nyquist-rate FFT spectrum sensor is defined as capturing a weak signal successfully when there is at least one large signal present over its *Span*. The DR of the ADC is crucial for the DR of a Nyquist-rate FFT spectrum sensor and it is defined as the range between the noise floor and its specified maximum output level. As discussed in the "RF Spectrum-Sensing Performance Metrics" section, the automatic gain control block prior to the ADC is tuned to receive the strongest signal that fits into the full scale of the ADC, while there exists a weak signal over the *Span*. Quantization noise of the ADC must be kept below the lowest signal level that is desired for it to be detected. One way to design an ADC with a high DR is to increase the number of bits; however, this solution comes at the expense of increased energy consumption.

The energy consumption for a Nyquist-rate FFT spectrum sensor is given by

$$E_{\text{Nyquist}} = P_{\text{Nyquist}} \cdot [(C_{\text{settle}}/\text{BW}_{\text{filter}}) + N_0 \cdot (N_s/f_s)]$$
$$\approx (N_0 \cdot P) \cdot T_{\text{scan}}. \tag{5}$$

The response time for a Nyquist-rate FFT spectrum sensor $T_{\text{resp,Nyquist}}$ is proportional to the antialiasing low-pass filter settling time. This architecture subdivides the instantaneous BW into the desired RBW through the FFT. Therefore, DSP time $T_{\text{rec,Nyquist}}$ is proportional to the FFT size N_0. In addition to the FFT size, design considerations such as the number of samples used for DSP, N_s, and the ADC sampling rate f_s, impact the scan time. The energy consumption for a

scan stays within a fixed first-order envelop even though a Nyquist-rate FFT spectrum sensor offers a lower scan time with a compression factor of N_0 (due to an increase in sampling rate, f_s, by the same factor) in contrast to a sweeping spectrum scanner [8].

Multibranch spectrum sensors

A multibranch architecture consists of multiple narrowband scanners in parallel, equivalent to a single-branch scanner for each spectrum bin spaced closely with an operating frequency distance of RBW. As an alternative to a sweeping spectrum scanner, and, although the scan time is reduced through the parallelized branches, a multibranch spectrum sensor consumes higher power under a fixed first-order energy budget and has an impractical hardware complexity when compared to a sweeping spectrum scanner [8].

The energy consumption for a multibranch spectrum sensor is given by

$$E_{\text{mult}} = P_{\text{mult}} \cdot [T_{\text{resp,mult}} + T_{\text{rec,mult}}]$$
$$= P_{\text{mult}} \cdot [(C_{\text{settle}}/\text{BW}_{\text{filter}}) + (N_s/f_s)]$$
$$= (N_0 \cdot P) \cdot T_{\text{scan}}. \tag{6}$$

RF spectrum sensors: Moving CS to hardware

The use of CS [16], [18] in RF signal reception and spectrum sensing has the potential to fundamentally improve the fixed tradeoffs between key design considerations, including scan time, energy consumption, and hardware complexity found in traditional RF spectrum scanners and sensors [4], [8], [13]. We are interested in a use case where detecting a few active signals above a certain signal-level threshold in a wideband spectrum and the CS theory demonstrates the possibility of recovering sparse multiband or multitone signals from only a few linear measurements proportional to the number of active bands or tones K_0, in the signal [16], [18], [34]. Because the information BW is much smaller than the instantaneous signal BW under the sparsity condition, CS as a blind sub-Nyquist sampling approach allows for sampling at a rate defined by this information BW rather than the instantaneous BW [6], [16], [18], [21], [39].

In this section, we discuss the current CS architectures, i.e., AICs, which have been proposed for RF signal reception [10], [25], [30], [37] and spectrum sensing [4], [7], [9], [11], [14]. These CS architectures can be organized into the following categories: random demodulator (RD) [21], [42], [43]; random modulation preintegrator (RMPI) [17], [25]; modulated wideband converter (MWC) [4], [10]; and quadrature AIC (QAIC) [8]. We survey these architectures with a focus on signal structure and models, integrated hardware realizations, signal recovery, and reconstruction algorithms. The main focus of this section is to show how a combination of careful attention to the deployment scenario and RF architecture innovations make it possible to convert the abstract improvements promised by CS theory, such as fewer measurements, to concrete improvements in time and energy efficiency.

Integrated CS hardware realizations

CS spectrum sensors require a few incoherent measurements, as discussed in the "Bridging CS Theory and Practical Hardware Implementation" section. These incoherent measurements are achieved by mixing the wideband RF sparse input signal with independent unique pseudorandom binary sequences (PRBSs). PRBS mixing spreads and smears the spectrum while folding the wideband input signal onto narrowband baseband sampling channels. CS DSP techniques disentangle this folding mechanism to locate the few active signals, i.e., bands or tones, above the signal-level detection threshold by using the incoherent samples [8].

Some current CS architectures implement each of these few incoherent measurements as a corresponding physical hardware branch; however, once the number of hardware branches is fixed, integrated CS hardware realizations can only successfully detect up to a specific maximum number of active signals [14]. To understand how the number of hardware branches maps to the number of detectable signals, we refer to the required number of incoherent measurements given by (3), which show the scaling of m with the number of nonzero signals, K_0 that need to be detected [8], [13]. As discussed in [4], the number of branches m may be traded for the branch sampling rate by a static reduction factor q, where $q = 1, 3, 5, \ldots$. Furthermore, we discuss system-level techniques in the "Structured Sparsity and Adaptive Methods for Dynamic CS" section, which combine signal processing mainly sparsity estimation and adaptive thresholding with hardware adaptation to maintain reliable performance across all sparsity levels [14].

- *RD*: The RD [21], [42], [43] shown in Figure 5 can be used to acquire sparse band-limited multitone signals. This single-branch architecture mixes the signal with a Nyquist-rate pseudorandom sequence. The mixing operation with the high-rate sequence smears the tones across the entire spectrum [21]. The RD employs a low-pass antialiasing filter to limit the BW of this smeared signal across the entire spectrum. Thanks to CS, the RF front end avoids the need for a high-rate ADC and samples the sparse multitone signal with a sub-Nyquist rate ADC. As discussed in [21], successful identification of these sparse tones relies on the fact that the demodulation process with a pseudorandom
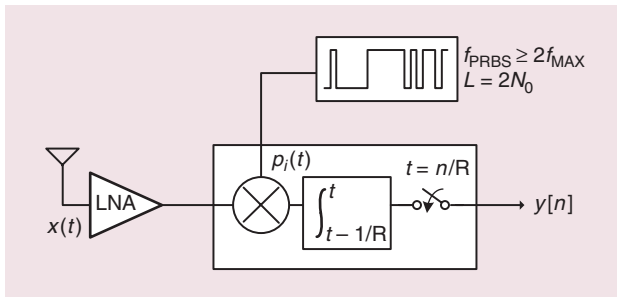
sequence provides a distinct signature within the passband of the low-pass filter for each tone.

This architecture is mostly suitable for signals with a finite set of harmonics chosen from a fixed uniform grid; however, real-world analog signals require a large number of harmonics to approximate them well within the discrete model. Therefore, the signal reconstruction becomes computationally intractable and performance is degraded through the sensitivity to the grid choice [4]. The analog acquisition hardware is simplified for the RD architecture through additional digital back-end complexity under the assumption of fast-paced advances in digital computing [21], [42], [43].

- *RMPI*: The RMPI [17], [25] shown in Figure 6 is a wideband receiver that implements random sensing for sparse multitone signals. The input signal $x(t)$ is mixed with a periodic distinct PRBS $p_i(t)$ in each channel to create a shifted copy of the entire spectrum by each harmonic of the PRBS. The output of the mixer is then integrated over a fixed-time interval T and digitized at a sub-Nyquist rate of $f_s = 1/T \ll f_{Nyq}$, where f_{Nyq} is the Nyquist rate of the input signal. The acquired sub-Nyquist samples from each channel represented as $y_i[n] = \int_{t-T}^{t} x(t) p_i(t) dt$ and $t = nT$ are used for CS support recovery and the spectrum of the output signal $y_i[n]$ is shown in Figure 6. Successful CS recovery relies upon knowing a basis or dictionary matrix $[\psi]$ to represent the signal, as discussed in the "Bridging CS Theory and Practical Hardware Implementation" section. In [25], the dictionary matrix is a multiscale Gabor dictionary to reconstruct radar pulses, which are sparse in the time–frequency plane.

Since the RMPI is a parallel-branch alternative of the RD architecture, for the remainder of this discussion, we compare the multiband architectures against the RMPI, which utilizes a low-pass CS approach. For the low-pass CS approach, the PRBSs have spectral content spanning from dc to f_{MAX}. Given that the spectral information below some f_{MIN} is not desired for RF applications such as in the example of PCAST spectrum scenario ($f_{MIN} = 2.7$ GHz), this only adds noise and suffers from the undesired blocking effect in the RF hardware when strong unwanted signals are present below f_{MIN}, as indicated by the red tone in the input spectrum shown in Figure 6(a). Furthermore, the frequency of the PRBS generator clock must be at least twice the maximum frequency ($f_{PRBS} \geq 2f_{MAX}$) of the input signal, which makes the PRBS generator the most power-hungry block of the RMPI architecture [8], [25].

The prototype wideband receiver, RMPI, demonstrated in the integrated performance table in Figure 6(c) [25], captures an effective instantaneous BW from 100 MHz to 2 GHz by eight parallel branches while digitizing samples at an aggregate rate of 320 megasamples/s (MS/s), i.e., 12.5-times lower than the Nyquist rate [25]. The RMPI prototype consumes 506.4 mW of power without including the power consumed by the computational platform needed for signal recovery.

- *MWC*: The MWC [4], [44] shown in Figure 7 is a multibranch sub-Nyquist sampling approach for sparse multiband signals.

FIGURE 5. A block diagram of the RD used for acquiring sparse band-limited multitone signals [42], [43]. This single-branch architecture mixes the multitone input signal with a Nyquist-rate pseudorandom sequence and samples the signal at a sub-Nyquist rate after low-pass filtering. LNA: low-noise amplifier.

The input signal $x(t)$ is mixed with periodic PRBSs at RF, low-pass filtered, and sampled at a low rate of $f_s \geq$ RBW.

Multiplication by a PRBS aliases the signal spectrum and spreads the information over the instantaneous BW as well by creating copies of the aliased signal spectrum in each bin. After low-pass filtering, only a single of copy of the folded signal spectrum remains for sampling at the sub-Nyquist rate. A suf-

ficiently large number of incoherent measurements from each parallel branch, represented by $y_1[n] \ldots y_m[n]$, allows for the successful recovery of the signal $x(t)$ if the m branches of the MWC are driven by m-unique, low cross-correlated PRBSs with a length of $L = 2N_0$ ($p_1(t), \ldots, p_i(t), \ldots, p_M(t)$) [11], [13].

Because the MWC also utilizes a low-pass CS approach as the RMPI, the frequency of the PRBS generator clock scales


FIGURE 6. (a) A block diagram of the RMPI, a parallel-branch alternative of the RD for sparse multitone signals, and spectra of various key signals including the input signal, the pseudorandom binary sequence, and the output signal. (b) The RMPI system parameters and its performance scaled for the example application of sensing 1 GHz in the 2.7–3.7-GHz range with a 20-MHz RBW ($q = 1$), and (c) the integrated RMPI [25] performance summary. LNA: low-noise amplifier; LO: local oscillator; Gm: transconductance.

up with the maximum frequency ($f_{PRBS} \geq 2f_{MAX}$) of the input signal. Note that the PRBSs will have frequency content spanning down to dc and the spectral information from dc to f_{MIN}, as indicated by the red band in the input spectrum shown in Figure 7(a) only adds noise and unwanted signals in the output spectrum for the RF spectrum scenario of interest [8], [11], [13].
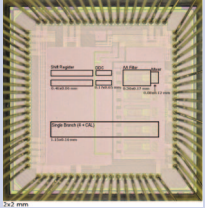
The MWC prototype discussed in the integrated performance table in Figure 7(c) [10] captures an effective instantaneous BW up to 900 MHz by five parallel branches, including one branch for calibration. The prototype MWC consumes 880.5 mW of total power from five parallel branches.

- *QAIC*: The QAIC [8], [13] illustrated in Figure 8 consists of an RF downconverter that limits the operation BW, *m*

MWC Architecture and Frequency-Domain Operation

(a)

	MWC System Parameters	Scaled for 2.7–3.7 GHz With a 20-MHz RBW
f_{LO}	—	—
$f_{3dB, IQ}$	—	—
L	$(2^r - 1) \geq 2N_0$	511
f_{PRBS}	$BL \geq 2f_{MAX}$	10.22 GHz
m	$\left\lceil 2K_0 C_{MWC} \log\left(\dfrac{L}{4K_0}\right)/q \right\rceil$	29 for $K_0 = 3$
f_s	$\geq qB$	20 MS/s

(b)

| Integrated MWC Realization [10]		Integrated MWC Performance [10]
	Application	Targeted Reception and Blocker Rejection
	Technology	65-nm CMOS
	Frequency Range	0–900 MHz
	m	5
	Blocker Rejection	62 dB
	Power Consumption	880.5 mW

(c)

FIGURE 7. (a) A block diagram of the MWC, a multibranch architecture for sparse multiband signals and spectra of various key signals including the input signal, the pseudorandom binary sequence, and the output signal. (b) The MWC system parameters [4] and its performance scaled for the example spectrum scenario of sensing 1 GHz in the 2.7–3.7-GHz range with a 20-MHz RBW ($q = 1$), and (c) the integrated MWC [10] performance summary. LNA: low-noise amplifier; LO: local oscillator.

in-phase (I) and quadrature-phase (Q) branches including PRBS mixers, antialiasing filters, and ADCs in each branch as well as a pairwise complex combiner. The band-pass filtered signal $x(t)$ is first downconverted to baseband, thereby enabling the use of shorter-length $L = N_0 < N$ and lower-frequency PRBSs $f_{PRBS} \geq f_{MAX} - f_{MIN}$.

In each I and Q branch, the downconverter output is multiplied by m-unique PRBSs $p_i(t)$ with low cross-correlation

(a)

	QAIC System Parameters	Scaled for 2.7–3.7 GHz With a 20-MHz RBW
f_{LO}	$(f_{MAX} + f_{MIN})/2$	3.2 GHz
$f_{3dB, IQ}$	$(f_{MAX} - f_{MIN})/2$	500 MHz
L	$(2^r - 1) \geq N_0$	63
f_{PRBS}	$BL \geq (f_{MAX} - f_{MIN})$	1.26 GHz
m	$\left\lceil K_0 C_{QAIC} \log \left(\frac{L}{K_0}\right)/q \right\rceil IQ$	8IQ (16) for $K_0 = 3$
f_s	$\geq qB$	20 MS/s

(b)

Integrated QAIC Realization [8]		Integrated QAIC Performance [8]
	Application	Interferer Detection
	Technology	65-nm CMOS
	Frequency Range	2.7–3.7 GHz
	m	8IQ (16) for $K_0 = 3$
	Scan Time	4.4 μs
	Power Consumption	81 mW

(c)

FIGURE 8. (a) A block diagram of the QAIC consists of an RF downconverter that limits the operation BW prior to the multibranch PRBS mixing of sparse multiband signals and spectra of various key signals including the input signal, the PRBS, and the output signal. By limiting the operation BW, the QAIC lowers the required rate and length of the PRBSs compared to low-pass CS architectures. (b) The QAIC system parameters [13] and its performance scaled for the example spectrum scenario of sensing 1 GHz in the 2.7–3.7-GHz range with a 20-MHz RBW ($q = 1$), and (c) the integrated QAIC [8] performance summary. LNA: low-noise amplifier; LO: local oscillator; IQ: in-phase/quadrature-phase.

between them, low-pass filtered, and sampled at a sub-Nyquist rate of the instantaneous BW. The output complex combiner allows for the selection of either the upper $[f_{\mathrm{MIN}}, f_{\mathrm{MAX}})$ or lower $(-f_{\mathrm{MAX}}, -f_{\mathrm{MIN}}]$ band cluster of the multiband input signal $x(t)$, as shown in Figure 1. The m-incoherent samples of the aliased spectrum, $y_1[n] \ldots y_m[n]$, are then used to recover the sparse multiband signal $x(t)$ [13].

The energy efficiency and sensitivity performance of the bandpass CS architectures are significantly improved when compared to the low-pass CS architectures, e.g., MWC for RF applications. Furthermore, the bandpass CS approach scales well to higher frequencies for the same instantaneous BW because it decouples the Nyquist-rate PRBS clock frequency from the maximum frequency of interest. To compare with the other CS AIC implementations, we use the performance numbers given in their corresponding system parameters tables in Figures 6(b) and 7(b), which are scaled to sense a 1-GHz span from 2.7 to 3.7-GHz with a 20-MHz RBW [8]. The number of hardware branches is scaled for three active signals, i.e., 29 for the low-pass CS architectures. Although the RMPI and MWC prototypes use more power-hungry shift-register PRBS generators, the PRBS generator power has been scaled assuming a linear feedback shift register (LFSR) with the required length of 511 and clock frequency of 10.22 GHz for a 20-MHz RBW. As reported in [8], the energy consumption per scan for the QAIC is five- to 10-times lower than the energy consumption of the low-pass CS implementations while maintaining a scan time of 4.4 μs when the performance numbers are normalized for the example RF spectrum scenario. For the example RF spectrum scenario, the QAIC requires a PRBS length of 63 and a PRBS generator clock frequency of 1.26 GHz for a 20-MHz RBW, as shown in the system parameters table in Figure 8(b). The measured power consumption of the QAIC implementation is 81 mW, as shown in the integrated performance table in Figure 8(c), without including the on-chip ADCs and phase-locked loops (PLLs) for the PRBS generator clock and quadrature downconverter fixed local oscillator (LO). When the power estimations of PLLs and ADCs are included for the QAIC, the total estimated power consumption is 115 mW [8]. As reported in [8], the scan time is 4.4 μs. Therefore, the front-end energy consumption per scan is 0.5 μJ [8].

System attributes and implementation parameters

The RMPI [25], MWC [4], and QAIC [8], [13] system parameters are defined in their corresponding tables in Figures 6–8. The RMPI and MWC sample a real signal $x(t)$ at RF. Because the RMPI and MWC utilize a low-pass CS approach, their frequency components extend from dc to f_{MAX}. The MWC input signal contains $N_0 = \lceil f_{\mathrm{MAX}} / B \rceil$ bands. The modulator bank employed by the MWC processes $2N_0$ total bands (i.e., N_0 positive and N_0 negative frequency bands) and $2K_0$ active bands (supports) [4]. Similar to the MWC with the only difference being the multitone signal model rather than a multiband signal model, the RMPI also processes $2N_0$ total tones and $2K_0$ active tones (see [21] and [25]). For the system attributes

and implementation parameters, we only compare within the architectures that employ multiband signal model.

In contrast to low-pass CS architectures, the QAIC samples a complex signal $I(t) \mp j \cdot Q(t)$ at a baseband intermediate frequency after quadrature downconversion. Its span extends from f_{MIN} to f_{MAX} and contains $N_0 = \lceil (f_{\mathrm{MAX}} - f_{\mathrm{MIN}}) / B \rceil$ total bands. The I and Q branches of the QAIC process N_0 total bands and K_0 active bands by selecting upper- or lower-band clusters. The number of bands $2N_0$ processed by the MWC is typically much larger than that which is processed by the QAIC when $f_{\mathrm{MIN}} >> 0$. The additional degree of freedom in adjusting the lower and upper boundaries of the QAIC *Span* in contrast to only adjusting the upper boundary of the MWC *Span*, provides significant improvements in energy consumption and sensitivity. The LO frequency (f_{LO}) and quadrature low-pass filter BW ($f_{\mathrm{3dB, IQ}}$) shown in Figure 8 are the system parameters for adjusting the QAIC frequency span. However, this scalability comes at the expense of unwanted residual bands due to quadrature linear impairments, e.g., phase and gain imbalance. Depending on the level of the impairment and compensation in the digital back end, this can increase false alarm probability.

When comparing a low-pass CS approach with a bandpass approach, it is assumed that both systems employ maximal length PRBSs generated with LFSR structures. The sequence length L is equal to $2^r - 1$, where $r \in \mathbb{Z}^+$ for a maximal length LFSR-type PRBS. The clock frequency of the PRBS generators f_{PRBS} employed by the MWC must be greater than twice the maximum frequency f_{MAX} of the input signal with a PRBS length of $L \geq 2N_0$ [4]. In contrast, f_{PRBS} for the QAIC must be greater than the $Span = f_{\mathrm{MAX}} - f_{\mathrm{MIN}}$ of the input signal with a PRBS length of $L \geq N_0$.

The number of branches or incoherent measurements m given in (3) required by the MWC and the QAIC for successful signal recovery, is proportional to the number of active bands multiplied by a logarithmic factor and constants $C_{\mathrm{MWC}} \geq 1, C_{\mathrm{QAIC}} \geq 1$. The MWC requires $m \propto 2K_0$ branches [4]. Because of its complex I/Q structure, the QAIC requires $m \propto K_0$ in-phase and quadrature-phase branches so the total number of physical hardware branches for the QAIC is $2m$ [13]. The sub-Nyquist sampling rate per branch for both systems is $f_s \geq B$ Hz. The aggregate sampling rate of the MWC is mf_s, while the aggregate sampling rate of the QAIC is $2mf_s$. To achieve a span of 1 GHz in the 2.7–3.7-GHz band with an RBW of 20 MHz, the number of branches m for the MWC is 29, and the number of I/Q branches m for the QAIC is eight, i.e., equivalent to 16 physical hardware branches. For the example RF spectrum scenario, the aggregate sampling rate of the MWC is 580 MS/s, while the aggregate sampling rate of the QAIC is 320 MS/s for a 20-MHz RBW [8].

Signal recovery and reconstruction algorithms

Support recovery and signal reconstruction are the main steps in a CS DSP [4]. K_0 supports (K_0 signals) of N_0 unknowns (N_0 spectrum bins) that are nonzero [4] are located using support recovery. Then, the time domain waveform of the input signal is estimated using signal reconstruction [4].

CS-based signal reconstruction is typically more computationally intensive than Nyquist-rate reconstruction; e.g., a typical CS algorithm requires roughly 20–1,000 FFTs for reconstruction due to its operation as a nonlinear function of the acquired samples [25]. Hence, a real-time time-domain signal reconstruction was not feasible for low-power mobile applications [25] until recently. In [45], a learning-based approach to subsampling by training signals and combinatorial optimization problems has been proposed as an opportunity for real-time and low-power signal reconstruction. The theory and recovery results for medical applications presented in [45] are based on efficient and scalable linear encoder and decoder pairs.

For applications concerned with only signal detection not estimation, e.g., RF interferer detection, support recovery is sufficient, and a complete time-domain reconstruction is not required. Support recovery extracts information such as carrier frequency, initial phase, and pulsewidth from the acquired waveform without reconstructing the time-domain signal. This is a less computationally expensive approach and is potentially cost-comparable to traditional Nyquist-rate approaches [25].

One simple greedy approach to sparse recovery is the OMP [31], [33]. A wide range of alternatives exist, including convex relaxations based on ℓ_1 minimization, nonconvex methods, more advanced greedy methods such as CoSamp, sparse Bayesian learning, and so on. Here, we focus on OMP because of its simplicity and tight connection to support recovery. OMP iteratively identifies the active supports by choosing elements that are maximally correlated with the measurement residual. The residual drops linearly as each signal support is recovered and removed from the measurements. Once all of significant signal supports above the signal-detection threshold are detected, there will be an abrupt change in the slope of the residual curve [14]. The adaptive thresholding concept discussed in the "Structured Sparsity and Adaptive Methods for Dynamic CS" section takes advantage of this iterative nature of the support-recovery algorithm by monitoring the OMP residual [14].

The CS system-measurement-matrix row dimension and the signal-level detection threshold are the two system parameters that impact the number of OMP iterations. The number of incoherent measurements sets the row dimension of the CS system measurement matrix. When the signal-level detection threshold is close to the system noise floor, the maximum number of OMP iterations performed in the digital back end is equal to the CS system-measurement-matrix row dimension. The choice of a threshold close to the noise floor might maximize the detection probability of a CS detector at the cost of an increased false alarm probability. If the signal-level detection threshold is reached before the maximum number of OMP iterations are performed, the CS DSP completes the recovery without further iterations [14]. We discuss the impact of the number of OMP iterations on the DSP energy consumption in the following section.

Evaluation of CS spectrum sensors: Key performance metrics

In this section, we discuss how to evaluate CS spectrum sensors and their key system performance metrics including detection and false alarm probabilities, scan time, front-end and DSP energy consumption, dynamic range, and scalability.

■ *Detection and false alarm probabilities*: CS signal detection is a statistical process and its key metrics are the detection probability P_D and the false alarm probability P_{FA}. P_D is defined as $\sum CD / N_E \cdot K_0$, where CD is the correct detections and N_E is the number of experiments. P_{FA} is defined as $\sum FA / N_E \cdot (L - K_0)$, where FA is the false alarms [8], [13], [14].

P_D and P_{FA} performance is typically used to demonstrate the sensitivity and DR of the CS signal detection. The P_{FA} depends on two key components: 1) the dimension of the system measurement matrix and 2) the signal-level detection threshold. The maximum P_{FA} is proportional to the dimension of the system measurement matrix when the threshold value is set close to the noise floor. This results in maximized P_D.

For example, Figure 9 demonstrates the measured sensitivity of a QAIC prototype through P_D and P_{FA} curves as a function of support power for a stationary sparse-spectrum scenario. The cluster of curves on the top correspond to P_D, and the cluster of curves on the bottom correspond to P_{FA}, with 1, 2, and 3 equal power signals above the signal-level detection threshold. The results are reported based on $N_E = 125$ experiments by using $N_s = 80$ samples per experiment. The measured sensitivity of this bandpass CS prototype is illustrated as the successful detection of three equal power signal bands as small as a -68 dBm/10 MHz band, while satisfying a target $P_D \geq 90\%$. For a single signal band, this sensitivity increases
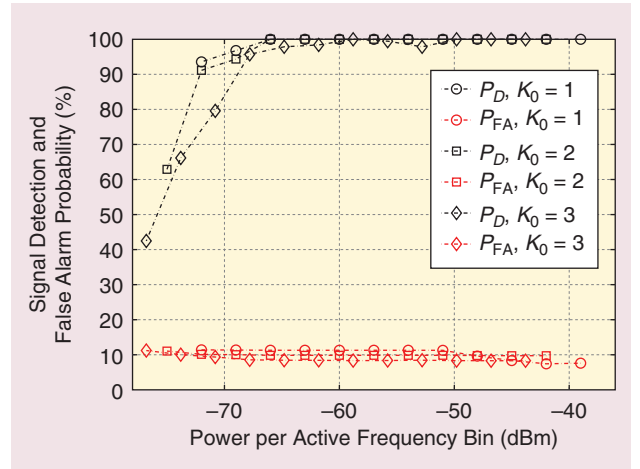

FIGURE 9. Characterization of the sensitivity of the CS AIC by measuring the detection probability, P_D, and the false alarm probability P_{FA} curves as a function of signal power for K_0 (=1, 2, and 3) equal power signals in a stationary spectrum scenario. The measured results from a QAIC prototype are reported based on $N_E = 125$ experiments each using $N_s = 80$ samples. The weakest signal level that the QAIC prototype can detect is -68 dBm/10 MHz band for the example spectrum scenario with $K_0 = 3$ while satisfying a target detection probability $\geq 90\%$ and false alarm probability $\leq 15\%$ [8].

to a –72 dBm/10 MHz band, while maintaining target-detection probability [8]. The measured P_{FA} remains below 15% for all of the experiments.

■ *Scan time*: Scan time has two main contributors: the detector response time and CS DSP time, as discussed in the "RF Spectrum-Sensing Performance Metrics" section. CS DSP time is N_s/f_s, where f_s is the ADC sampling rate and N_s is the number of samples collected. The front-end detector response time is very fast as a result of the PRBS mixing operation. Mixers driven by PRBSs spread the spectrum and allow users to capture the entire wide spectrum BW of interest instantaneously by creating folded copies of the spectrum information in all of the bins. The AIC front-end detector, therefore, is responsive to a change in the spectrum at the settling time of the low-pass filters [8].

■ *Front-end energy consumption*: Front-end energy consumption is defined as the power consumed by m physical hardware branches times the CS scan time [8]. For a CS spectrum sensor, front-end energy consumption is given by

$$E_{AIC} = P_{AIC} \cdot T_{scan}$$
$$= P_{AIC} \cdot [T_{resp,AIC} + T_{rec,AIC}]$$
$$= P_{AIC} \cdot \left[\left(\frac{C_{settle}}{BW_{filter}}\right) + \left(\frac{N_s}{f_s}\right)\right]. \qquad (7)$$

■ *CS DSP energy consumption*: For m the number of branches or incoherent measurements, L the length of the PRBSs employed by the AIC ($L = N = 2N_0$ for MWC and $L = N_0$ for QAIC) to a first order, $m \cdot L$ multiplications and additions are needed to complete a single iteration of the OMP for sparse support recovery. For a single set of measurements, K_0 iterations are required with an ADC sampling period of T_s, the time required to complete a single iteration is $T_{iteration} = T_s/K_0$. Therefore, the energy consumption of a single iteration is given by (8), where P_{adder} and P_{mult} are the power consumption of adders and multipliers, respectively.

$$E_{iteration} = m \cdot L \cdot (P_{adder} + P_{mult}) \cdot T_{iteration}. \qquad (8)$$

If N_s consecutive samples from the AIC are used for sparse support recovery to detect K_0 active signals, then the total number of multiplications and additions needed by the OMP is roughly $N_s \cdot K_0 \cdot m \cdot L$. To a first order, the total CS DSP energy consumption for sparse support recovery through OMP is given by (9).

$$E_{OMP} = m \cdot L \cdot (P_{adder} + P_{mult}) \cdot T_s \cdot N_s. \qquad (9)$$

Table 2 provides an estimation of DSP cost using the OMP algorithm for various CS signal-detector architectures [4], [8], [11]. It is assumed that the DSP power is dominated by the power dissipated in multipliers (P_{mult}).

Taking advantage of signal features, such as the bandpass signal model for RF applications, reduces the CS system-matrix dimension by reducing the length of the PRBSs L as well as the number of incoherent measurements, m, drastically. Since m and L play a significant role in DSP energy consumption, a smaller matrix dimension results in a lower DSP cost by performing the sparse recovery only for a desired signal band. For example, for the QAIC system that utilizes a bandpass CS approach, assuming $N_s = 80$, $m = 8$, $L = 63$, and $K_0 = 3$ for real multiplications and additions, the computational load of the OMP is estimated to be roughly $535 nJ$ [46]. Here, the estimated power consumption for a 16×16 multiplier and a 16-bit adder with a settling time less than 17 ns is 60 μW and 5 μW, respectively [47], [48].

■ *DR*: The instantaneous DR of an AIC is defined as the maximum power difference between $k < K_0$ weak signals and $(K_0 - k)$ strong signals that the AIC can successfully detect over GHz or wider instantaneous BW, while satisfying a desired P_D and P_{FA} given the linear and nonlinear impairments and noise [8], [13]. The "Open Challenges and Future Directions" section discusses how to model and compensate for some of these impairments as one of the open challenges in CS AICs design.

■ *Adaptive configurability and scalability*: Adaptive configurability and scalability are essential in real-world applications. CS AICs demonstrate scalability in terms of the number of hardware branches, samples per branch, and the number of bits employed in their ADCs. An example of this scalability is demonstrated in [8] and [14] through measured P_D and P_{FA} performance, respectively, in Figure 10.

The AICs can operate with varying ADC resolutions [49]. Reducing the number of bits will drastically consume less energy in DSP compared to its high-resolution counterpart

Architecture	DSP Algorithm	Estimated DSP Power	DSP Time
MWC [4]	OMP	$4NK_0^2 \log_2\left(\frac{N}{K_0}\right)P_{mult}$	$\frac{N_s}{qB}$
QAIC [8]	OMP	$2N_0 K_0^2 \log_2\left(\frac{N_0}{K_0}\right)P_{mult}$	$\frac{N_s}{B}$
TS-QAIC [14]	OMP	$2N_0 (n_{seg} K_0)^2 \log_2\left(\frac{N_0}{n_{seg} K_0}\right)P_{mult}$	$\frac{n_{seg} N_s}{B}$
DRF2IC [11]	OMP	$2N_0 K_0^2 \log_2\left(\frac{N_0}{K_0}\right)P_{mult}$	$\frac{N_s}{qB}$

Table 2. The estimated DSP cost of various CS-signal-detector architectures.

DRF2IC: direct RF-to-information converter; MWC: modulated wideband converter; QAIC: quadrature analog-to-information converter; TS-QAIC: time-segmented QAIC.

at the expense of a lower DR; this tradeoff between energy consumption and DR may be useful for many real-world applications. An example of the AIC operation under varying ADC resolutions is illustrated through the measured P_D and P_{FA} curves in [8]. As demonstrated in Figure 10(a), for $K_0 = 2$ active bands with a target $P_D > 90\%$ and $P_{FA} < 15\%$, a QAIC prototype with 8-bit ADCs can detect −72 dBm/10 MHz bands, while the same QAIC employing 1-bit ADCs can detect −68-dBm/10-MHz bands [8].

The AIC's operation under signal overload is also illustrated through an example in Figure 10(b). In this example, the CS system is designed to successfully detect $K_0 = 3$ signals by using a sufficient number of hardware branches for the required number of incoherent measurements; however, the system is tested with $K_0 = 6$ signals under the assumption of dynamic signal conditions. The CS system provides unreliable results with a maximum P_D of 50% since it commits to a fixed number of hardware branches under the assumption of static signal sparsity with a maximum bound of $K_0 = 3$ signals for this example [14]. A possible adaptive action in response to this system overload that handles more signals than the expected occupancy level is relaxing the threshold of recovery algorithms. The AIC becomes blind to lower power signals and effectively sparsifies the spectrum. Only signals above the new threshold are detected with high confidence. Another possible solution based on sparsity estimation and time segmentation (TS) to avoid system failure under the support overload is discussed in the "Structured Sparsity and Adaptive Methods for Dynamic CS" section [14]. As shown in Figure 10, the TS technique combined with adaptive thresholding improves the P_D performance back to ≥90% for reliable operation under signal overload at the expense of increased scan time [14].

RF spectrum sensors: Challenges of implementing CS in hardware

Taking CS to the hardware level has been a growing research interest both for RF communications and beyond with a specific focus on demonstrating how to overcome the traditional tradeoffs between scan time and resolution through integrated CS realizations for different signal structures. At the same time, CS AICs can suffer from nonidealities such as jitter noise, aperture [15], nonlinear and linear impairments [13], and noise folding [12] in addition to other performance limiting factors, e.g., static signal structure assumptions [14]. In this section, we discuss three major challenges of implementing CS in hardware for the RF spectrum-sensing example. However, these challenges are applicable to CS hardware used in other applications as well.

Jitter noise and aperture

Because the CS AICs sample the signal of interest with sub-Nyquist rate ADCs, they do not suffer from the jitter noise and aperture of the sampling stage. However, the improvement in the ADC performance comes at the expense of jitter noise and aperture of the PRBS mixing stage, which still operates at the Nyquist rate of the instantaneous signal BW. Since both of these nonidealities, i.e., jitter noise and aperture, depend on frequency, the Nyquist-rate PBRS mixing stage degrades the resolution and performance gain of CS AICs in high-BW applications [15]. Reducing the PRBS generator clock frequency by using a bandpass CS approach that limits the RF BW through downconversion makes the AICs less susceptible to jitter and aperture [13]. One interesting direction suggested in [15] is to investigate the effect of sensing matrix sparsity on the jitter noise and aperture in the PRBS mixing stage.
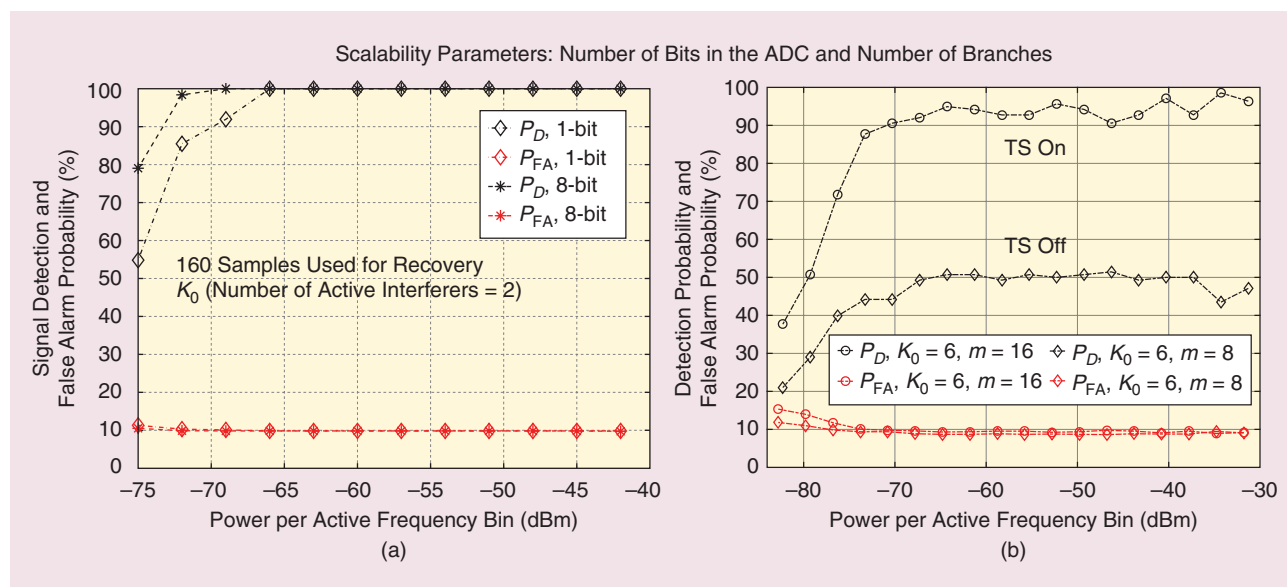

FIGURE 10. The scalability of the CS AIC is demonstrated in terms of ADC resolution and number of branches through measured detection probability, P_D, and false alarm probability P_{FA} curves [8], [14]. (a) The CS AICs operation under varying ADC resolutions is illustrated for an example stationary spectrum scenario of $K_0 = 2$ signals [8]. (b) Their operation under signal overload is illustrated for $K_0 = 6$ signals when the hardware uses $m = 8$-physical branches (i.e., TS OFF) and also when the hardware uses $m = 16$ (eight physical and eight virtual branches) (i.e., TS ON) [14].

Intrinsic tradeoff: Instantaneous DR and instantaneous BW

CS AICs have an intrinsic tradeoff between instantaneous DR and instantaneous BW. In the presence of a strong signal, the weakest signal that can be detected over the same instantaneous BW is mainly limited by noise. The effective sensitivity level is degraded in CS architectures due to the noise folding [12] effect of the PRBS mixing to instantaneously capture a wide instantaneous BW. Hence, the CS AICs offer a moderate instantaneous DR over a very wide instantaneous BW.

The CS AIC effective noise bandwith (NBWcs) that accounts for noise folding, tracks the instantaneous BW of the AICs in contrast to traditional spectrum scanners, where NBW is equal to the RBW. The effective noise bandwidth of CS AICs is defined by L.RBW, and it scales with the *Span*. For a given *Span*, if the RBW is halved, L needs to be doubled. Therefore, the NBW_{CS} remains the same. In essence, the impact of noise folding is independent of RBW and is fixed for a given *Span*.

Given the intrinsic tradeoff between instantaneous DR and instantaneous BW, CS AICs are more suitable for applications that do not require a very large DR, e.g., detecting a few large signals rapidly over a wide instantaneous BW rather than using the CS AIC as a high-sensitivity receiver. In [11], high-sensitivity signal reception and CS signal detection over a wide instantaneous BW are combined into a scalable architecture called *direct RF-to-information converter* (*DRF2IC*). This architecture couples the best features of CS wideband detection and traditional sweeping spectrum scanners for narrowband detection with high-sensitivity signal reception.

Structured sparsity and adaptive methods for dynamic CS

CS RF spectrum-sensing and signal-reception approaches assume that sparsity is fixed or bounded and rapidly break down when these bounds are violated. Fixed sparsity assumptions limit the applicability of sparse signal processing techniques in practical, dynamic wideband spectrum environments, in which the spectrum may occasionally be densely occupied.

Using signal information extracted from the sparse recovery algorithm, e.g., OMP, the CS detector system, e.g., TS-QAIC [14], can trade off time resolution versus system failure under dynamically varying signal conditions, as shown in Figure 11. For a dense spectrum, the TS technique demonstrated in [14] introduces a virtual extension of hardware to increase the signal-detection capability by creating unique measurements of the spectrum from a single physical branch through time-segmented independent-sensing waveforms. The virtualization technique is only applicable under the prior assumption of the spectrum being stationary across the time segments n_{seg}. However, the pseudostationary spectrum assumption is typically valid since CS systems offer very short scan times [50].

Sparsity estimation can be achieved by monitoring a normalized residual of the OMP in DSP. This residual information can be used for setting the signal-detection threshold adaptively. The OMP residual monitoring and adaptive thresholding do not require any additional steps in the OMP algorithm [14] even if it is implemented in real time. This adaptive system scaling is an alternative to always designing the CS hardware for the worst-case spectrum scenario. The worst-case design can be wasteful in terms of scan time and energy consumption [14]. An example system-operation protocol is shown in Figure 11 for $K_0'' > K_0' > K_0$ sparsity levels to illustrate the decision mechanisms for TS and adaptive thresholding in which the detection threshold values are set by system hardware design specifications, such as the compression point and the sensitivity for high- and low-threshold settings, respectively [14].

Discussion

In this article, we discussed the ongoing research in demonstrating the abstract improvements of CS theory in practical sensing system implementations by using an example application of RF signal reception and spectrum sensing. To summarize, Figure 12 shows the system performance plane of these RF spectrum scanners and sensors with their associated tradeoffs for the example RF spectrum scenario of scanning a 1-GHz-wide spectrum in the 2.7–3.7-GHz range with an RBW of 20 MHz. The reported energy consumption per scan and scan time for each architecture are normalized to the energy consumption and scan time of a CS QAIC detector, where E and T_{scan} denote the energy consumption and the scan time of a CS QAIC detector, respectively [8].

Traditional sweeping spectrum scanners and cross-correlation spectrum analyzers are located at the corner where the scan time is the longest and the energy consumption is the highest. They offer a high DR but at the expense of long scan time and small instantaneous BW equal to their RBW. For this example spectrum scenario, sweeping spectrum scanners require $50T_{scan} = 220\,\mu s$ scan time with a 20-MHz instantaneous BW in a 1-GHz span that results in high energy consumption and the risk of missing the changes in a dynamic spectrum. The cross-correlation spectrum analyzers shown in Figure 13 are the preferred method for increasing the DR by improving sensitivity at the expense of increased scan time [41]. It is demonstrated in [41] that reducing the NF by 1.5 dB doubles the measurement time. Both of these architectures are specialized for spectrum analyzers and instrumentation. Nyquist-rate FFT sensors offer 50-times faster sensing compared to that of sweeping spectrum scanners; however, the energy consumption remains the same as the first order, i.e., a power-hungry approach due to the high Nyquist-rate ADCs (e.g., 2 GS/s). We note that the CS architectures are suitable for the rapid detection of signals, e.g., interferers, over a wide BW with a moderate instantaneous DR, especially the bandpass CS approaches, such as the QAIC [8] the TS-QAIC [14], which are up to an-order-of-magnitude more energy efficient than low-pass CS approaches. The TS-QAIC discussed in [14] extends the signal-detection capabilities without any additional hardware resources compared to the QAIC, as shown in Figure 12.

Open challenges and future directions

In this section, we discuss some open challenges and future directions for CS in hardware. Specific open challenges include making recovery algorithms robust to clock jitter and nonlinearities in the signal path. Each of these nonidealities introduces

power and performance tradeoffs between digital and analog domains; more robust recovery algorithms enable better tradeoffs, and hence, more efficient sensors. Some of these future directions include adapting sparse recovery algorithms to hardware; and cooptimizing hardware, sensing waveforms, and recovery algorithms; optimizing CS for fixed system sizes and fixed computational budgets; and taking CS to the next level, i.e., functionally flexible unified architectures.

One of the open challenges is the modeling of the nonidealities limiting the instantaneous DR of AICs, such as noise folding, linear impairments, and nonlinearities. These nonideality models can be used for compensation or to make the sparse recovery performance robust against them. For example, the wideband I/Q downconverter employed by bandpass CS architectures such as QAIC, TS-QAIC comes at the expense of frequency-independent and frequency-dependent IQ gain and phase imbalance [13]. Image components due to the linear impairments can degrade the instantaneous DR. Digital-based calibration methods have been demonstrated to compensate IQ imbalance; however, the accurate calibration and modeling of these linear impairments are challenging for CS AICs since the PRBS mixing folds the wideband spectrum information onto a few narrow baseband sampling channels. One promising avenue is to use machine-learning techniques to compensate for nonidealities, e.g., by using neural network tools
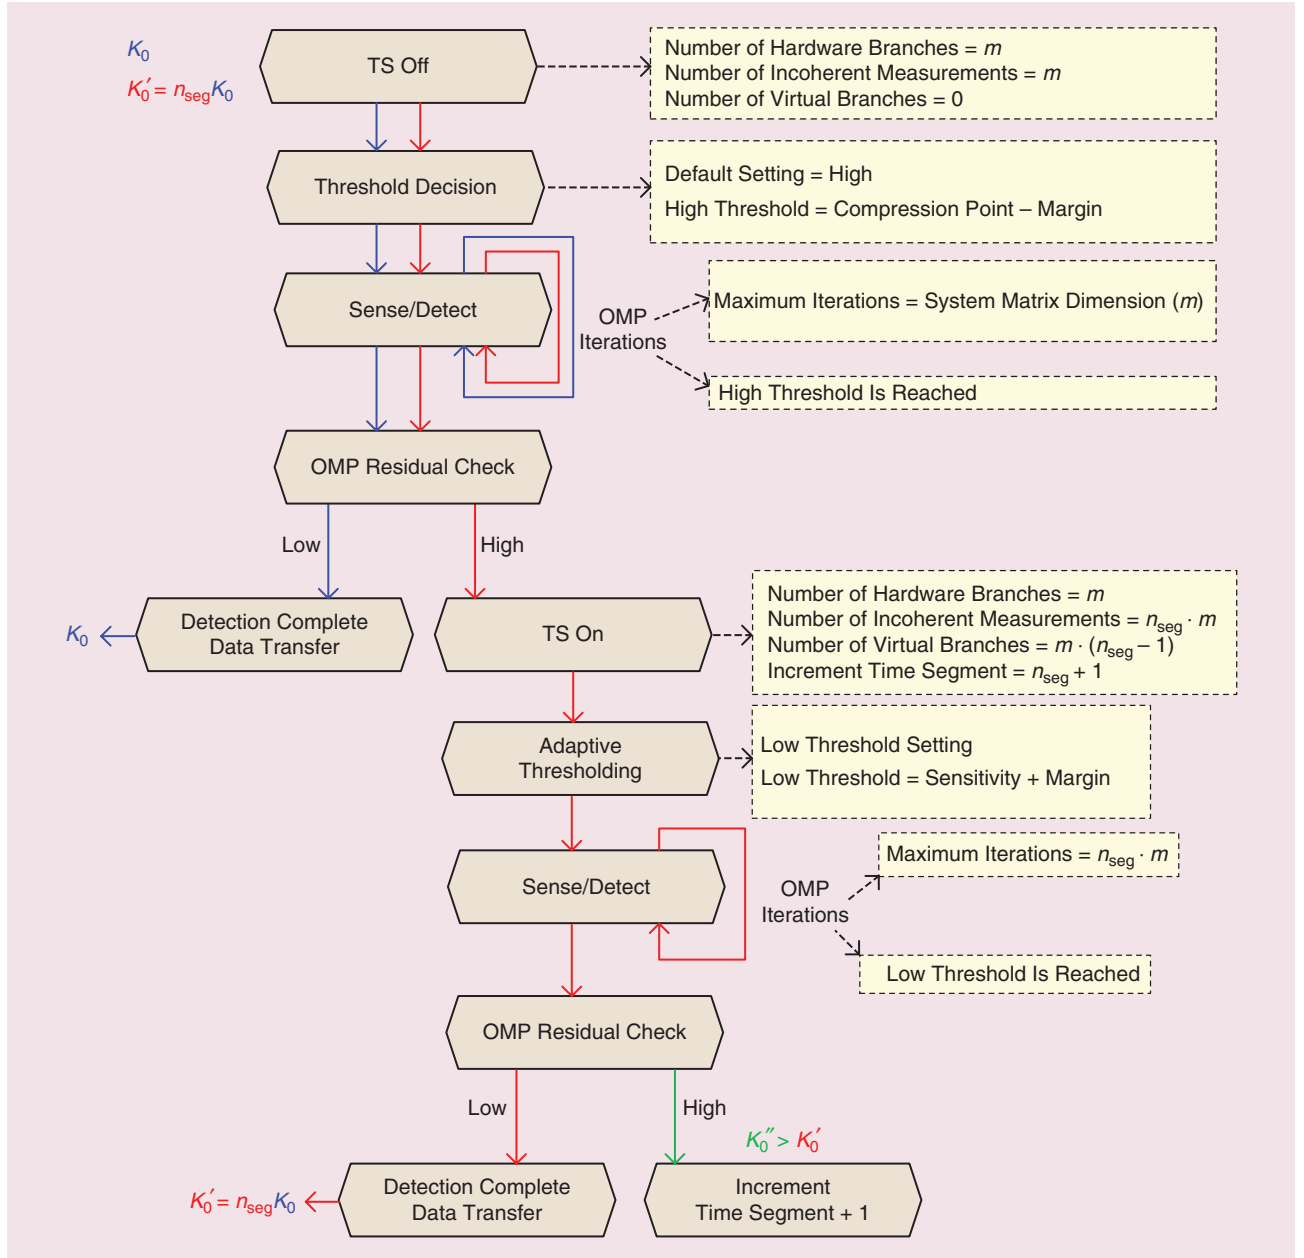

FIGURE 11. An example of an adaptive system scaling protocol for dynamic CS to perform reliably under changing signal conditions: the sparsity is estimated by monitoring a normalized residual of the OMP algorithm in DSP, and the system adapts through TS and adaptive thresholding. For each operation in the example protocol for $K_0'' > K_0' > K_0$ sparsity levels, critical design decisions, such as how to set the detection threshold value, the maximum number of OMP iterations, and the number of virtual branches, are in the dashed boxes [14].
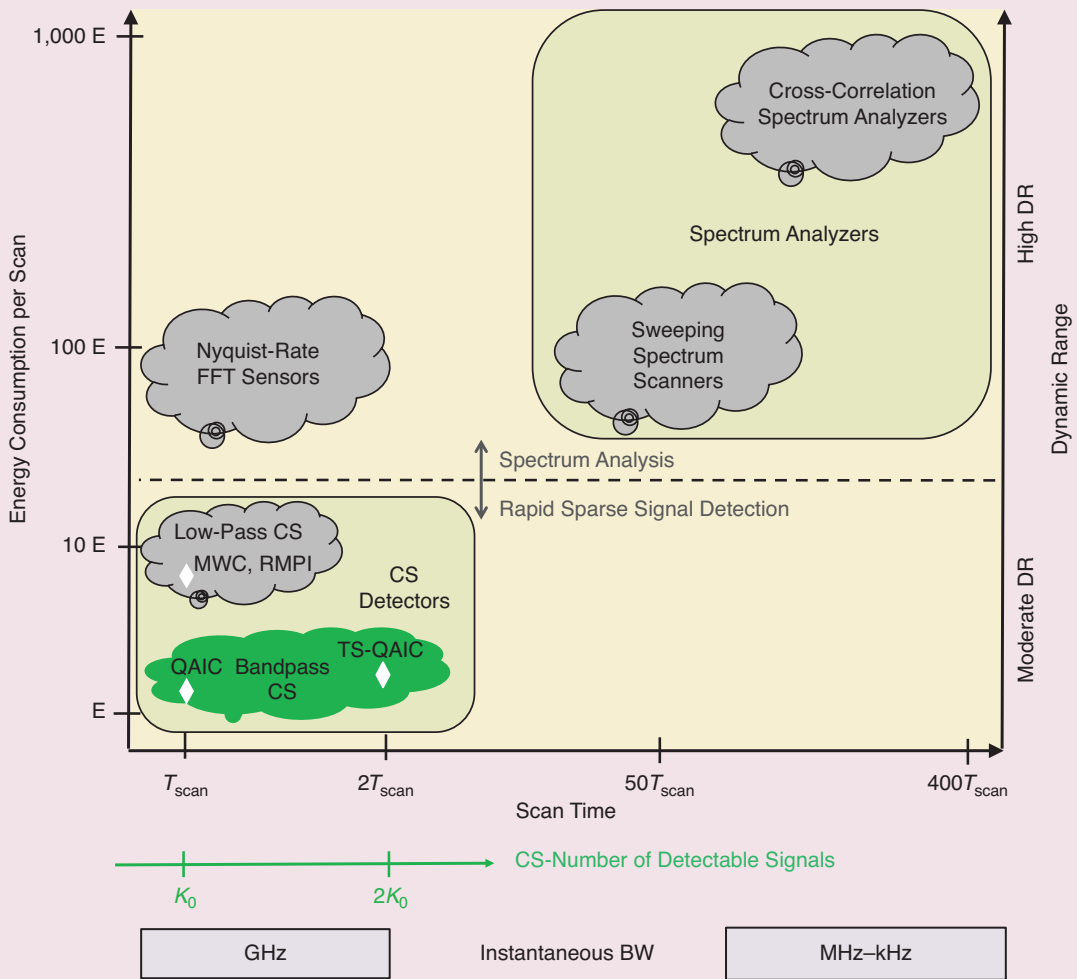
FIGURE 12. The system performance plane of RF spectrum scanners and sensors. The energy consumption and scan time of RF spectrum scanners and sensors are normalized to the energy consumption and scan time of the CS bandpass QAIC, E and T_{scan}, respectively.
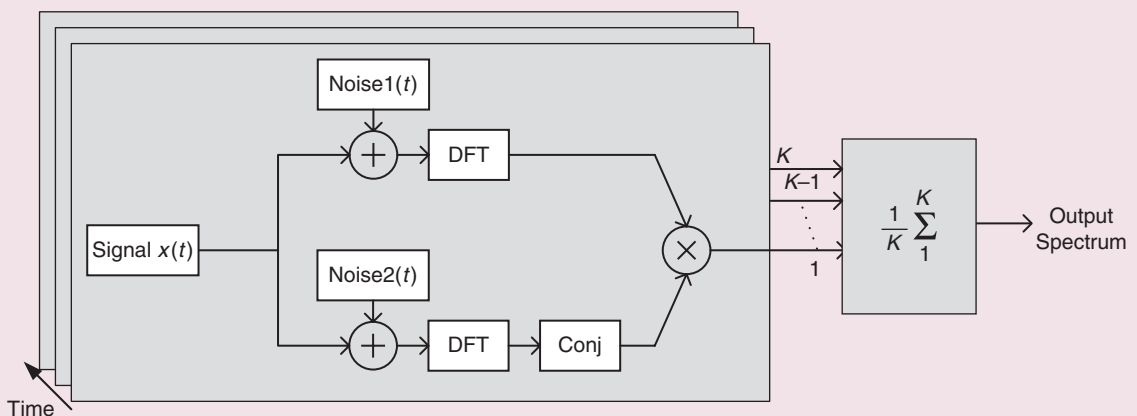

FIGURE 13. A mathematical block diagram of the cross-correlation spectrum analyzers discussed in [41]. Two independent signal paths add noise, i.e., Noise1(t) and Noise2(t), to the input signal $x(t)$. The cross-spectrum converges to the desired signal power for increasing integration time when the signal and noise sources are uncorrelated [41]. By utilizing cross-correlation, the system noise is lowered and the DR is improved at the expense of increased scan time. DFT: discrete Fourier transform; Conj: complex conjugate.

to directly learn sparse recovery procedures [51]. To cope with the deteriorated signal-to-noise ratio due to noise folding, cyclostationary detection has been recently proposed for spectrum sensing in algorithm-based research discussed in [52] and cyclic spectrum recovery from sub-Nyquist samples outperforms energy detection under noisy conditions in their presented simulations. One forward direction is to investigate architectural and circuit-level innovations based on these novel signal processing concepts through collaborations between hardware and signal processing experts.

Another interesting challenge is the effect of PRBS mixing on the nonlinearities in the signal path, and we briefly discuss two possible linearity scenarios while contrasting with traditional receivers. The intrabin spectrum-linearity scenario demonstrates the narrowband-input-to-narrowband output mapping of the PRBS mixing. On the other hand, the interbin spectrum-linearity scenario demonstrates PRBS mixing wideband-input-to-narrowband output mapping resulting from multi-LO spectral properties of PRBSs. This unique multi-LO feature of the PRBS mixer enables AICs to rapidly capture a wide instantaneous BW, while introducing a nonlinearity mechanism different from traditional receivers. A detailed analysis is needed for these nonlinearities in the signal path of CS AICs and making recovery algorithms robust to the nonlinearities through hardware–software codesign. To address this open challenge, iterative linearization for smaller nonlinearities and lifting to a semidefinite program for larger ones have been an ongoing interest to the signal processing research community [53].

In this article, we discuss taking CS to the hardware level for stand-alone CS applications. In addition to the open challenge of understanding nonidealities of a CS hardware, another important future need is unification through the combination of functionally flexible and rapidly reconfigurable architectures to achieve cost, size, functionality, and power targets for the mass deployment of CS in practical systems. While some recent work has started to address this issue by combining high-sensitivity signal reception, CS wideband detection, and sweeping spectrum analyzers into a compact architecture [11], this is just the beginning of a new generation of unified architectures combining novel signal processing concepts into a flexible hardware. Depending on the deployment scenario, these multifunction architectures are envisioned to quickly change between their different operation modes on the fly.

Acknowledgments

Authors

Rabia Tugce Yazicigil (rty@bu.edu) received her B.S. degree from Sabanci University, Istanbul, Turkey, her M.S. degree from École Polytechnique Fédérale de Lausanne, Switzerland, and her Ph.D. degree from Columbia University, New York, all in electrical engineering, in 2009, 2011, and 2016, respectively. She was a postdoctoral associate at the Massachusetts Institute of Technology, Cambridge, from 2016 to 2018. Currently she is an assistant professor in the Electrical and Computer Engineering Department at Boston University, Massachusetts. She was a recipient of the Columbia University EE Collaborative Research Award in 2016 and placed second at the Bell Labs Future X Days Student Research Competition in 2015. She is a Member of the IEEE.

Tanbir Haque (th424@columbia.edu) received his B.S. and M.S. degrees in electrical engineering from the State University of New York at Stony Brook and the New York University Tandon School of Engineering in 1994 and 1997, respectively. He received his M.S. degree in applied mathematics in 2010 from Columbia University, New York, where he is currently pursuing his Ph.D. degree in electrical engineering. He is a principal engineer with the Incubation and Strategy Group at InterDigital Communications, Melville, New York. His responsibilities include technology road map development, incubation, prototyping, and technology standardization. He has coauthored one book and holds 14 U.S. patents, with several applications under review. He is a Member of the IEEE.

Peter R. Kinget (peter.kinget@columbia.edu) received his B.S. degree in electrical and mechanical engineering in 1990 and his Ph.D. degree in electrical engineering in 1996, both from the Katholieke Universiteit Leuven, Belgium. From 1996 to 1999, he was a member of the technical staff in the Design Principles Department at Bell Laboratories, Lucent Technologies, New Jersey. From 1999 to 2002, he held technical and management positions in the field of integrated circuit design at Broadcom, Irving, California, CeLight, Iselin, New Jersey, and MultiLink, Somerset, New Jersey. In 2002, he joined the faculty of the Department of Electrical Engineering, Columbia University, New York, where he is now the Bernard J. Lechner Professor and the chair of the electrical engineering department. He is a Fellow of the IEEE.

John Wright (jw2966@columbia.edu) received his B.S. degree in computer engineering, his M.S. degree in electrical engineering, and his Ph.D. degree in electrical engineering from the University of Illinois at Urbana–Champaign (UIUC) in 2004, 2007, and 2009, respectively. From 2009 to 2011, he was with Microsoft Research Asia. He is currently an associate professor in the Electrical Engineering Department at Columbia University, New York. His has received a number of awards, including the 2012 Conference on Learning Theory Best Paper Award (with Dan Spielman and Huan Wang), the 2009 Lemelson-Illinois Prize for Innovation for his work on face recognition, and the 2009 UIUC Martin Award for Excellence in Graduate Research. His research interests include high-dimensional data analysis. He is a Member of the IEEE.

References

[1] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly, "Compressed sensing MRI," *IEEE Signal Process. Mag.*, vol. 25, pp. 72–82, Mar. 2008.

[2] J. W. Choi, B. Shim, Y. Ding, B. Rao, and D. I. Kim, "Compressed sensing for wireless communications: Useful tips and tricks," *IEEE Commun. Surveys Tut.*, vol. 19, no. 3, pp. 1527–1550, Feb. 2017.

[3] Z. Qin, J. Fan, Y. Liu, Y. Gao, and G. Y. Li, "Sparse representation for wireless communications: A compressive sensing approach," *IEEE Signal Process. Mag.*, vol. 35, no. 3, pp. 40–58, 2018.

[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, 2010.

[5] R. G. Baraniuk, T. Goldstein, A. C. Sankaranarayanan, C. Studer, A. Veeraraghavan, and M. B. Wakin, "Compressive video sensing: Algorithms, architectures, and applications," *IEEE Signal Process. Mag.*, vol. 34, no. 1, pp. 52–66, 2017.

[6] Z. Tian and G. B. Giannakis, "Compressed sensing for wideband cognitive radios," in *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing (ICASSP)*, 2007, pp. IV-1357–IV-1360.

[7] D. E. Bellasi, L. Bettini, C. Benkeser, T. Burger, Q. Huang, and C. Studer, "VLSI design of a monolithic compressive-sensing wideband analog-to-information converter," *IEEE Trans. Emerg. Sel. Topics Circuits Syst.*, vol. 3, no. 4, pp. 552–565, 2013.

[8] R. T. Yazicigil, T. Haque, M. R. Whalen, J. Yuan, J. Wright, and P. R. Kinget, "Wideband rapid interferer detector exploiting compressed sampling with a quadrature analog-to-information converter," *IEEE J. Solid-State Circuits*, vol. 50, no. 12, pp. 3047–3064, 2015.

[9] R. T. Yazicigil, T. Haque, J. Wright, and P. R. Kinget, "Band-pass compressive sampling as an enabling technology for rapid wideband RF spectrum sensing," in *Proc. 50th Asilomar Conf. Signals, Systems and Computers (ACSSC)*, 2016, pp. 1032–1036.

[10] D. Adams, Y. Eldar, and B. Murmann, "A mixer frontend for a four-channel modulated wideband converter with 62 dB blocker rejection," in *Proc. IEEE Radio Frequency Integrated Circuits Symp. (RFIC)*, 2016, pp. 286–289.

[11] T. Haque, M. Bajor, Y. Zhang, J. Zhu, Z. A. Jacobs, R. B. Kettlewell, J. Wright, and P. R. Kinget, "A reconfigurable architecture using a flexible LO modulator to unify high-sensitivity signal reception and compressed-sampling wideband signal detection," *IEEE J. Solid-State Circuits*, vol. 53, no. 6, pp. 1577–1591, 2018.

[12] E. Arias-Castro and Y. C. Eldar, "Noise folding in compressed sensing," *IEEE Signal Process. Lett.*, vol. 18, no. 8, pp. 478–481, 2011.

[13] T. Haque, R. T. Yazicigil, K. J.-L. Pan, J. Wright, and P. R. Kinget, "Theory and design of a quadrature analog-to-information converter for energy-efficient wideband spectrum sensing," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 62, no. 2, pp. 527–535, 2015.

[14] R. T. Yazicigil, T. Haque, M. Kumar, J. Yuan, J. Wright, and P. R. Kinget, "How to make analog-to-information converters work in dynamic spectrum environments with changing sparsity conditions," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 65, no. 6, pp. 1775–1784, 2018.

[15] O. Abari, F. Lim, F. Chen, and V. Stojanovic, "Why analog-to-information converters suffer in high-bandwidth sparse signal applications," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 60, no. 9, pp. 2273–2284, 2013.

[16] D. Donoho, "Compressive sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.

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

[18] E. Candés and T. Tao, "Robust uncertainty principles: Signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, 2006.

[19] R. G. Baraniuk, "Compressive sensing [Lecture Notes]," *IEEE Signal Process. Mag.*, vol. 24, no. 4, pp. 118–121, 2007.

[20] Y. C. Eldar and G. Kutyniok, *Compressed Sensing: Theory and Applications*. Cambridge, U.K.: Cambridge Univ. Press, 2012.

[21] 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, no. 1, pp. 520–544, 2010.

[22] B. Gozcu, R. K. Mahabadi, Y. H. Li, E. Ilicak, T. Cukur, J. Scarlett, and V. Cevher, "Learning-based compressive MRI," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1394–1406, 2018.

[23] K. V. Mishra, A. Kruger, and W. F. Krajewski, "Compressed sensing applied to weather radar," in *Proc. IEEE Geoscience and Remote Sensing Symp.*, 2014, pp. 1832–1835.

[24] B. M. Sanandaji, A. Tascikaraoglu, K. Poolla, and P. Varaiya, "Low-dimensional models in spatio-temporal wind speed forecasting," in *Proc. American Control Conf. (ACC)*, 2015, pp. 4485–4490.

[25] J. Yoo, S. Becker, M. Loh, M. Monge, E. Candes, and A. Emami-Neyestanak, "A 100-MHz–2GHz 12.5x sub-Nyquist rate receiver in 90-nm CMOS," in *Proc. IEEE Radio Frequency Integrated Circuits (RFIC) Symp.*, 2012, pp. 31–34.

[26] A. Kadambi and P. T. Boufounos, "Coded aperture compressive 3-D LIDAR," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 1166–1170.

[27] G. A. Howland, P. Zerom, R. W. Boyd, and J. C. Howell, "Compressive sensing LIDAR for 3D imaging," in *Proc. Conf. Lasers and Electro-Optics (CLEO)*, 2011, pp. 1–2.

[28] R. Hummel, S. Poduri, F. Hover, U. Mitra, and G. Sukhatme, "Mission design for compressive sensing with mobile robots," in *Proc. IEEE Int. Conf. Robotics and Automation*, 2011, pp. 2362–2367.

[29] M. Bajor, T. Haque, J. Wright, and P. R. Kinget, "Theory and design of a direct space-to-information converter for rapid detection of interferer DoA," in *Proc. 86th IEEE Vehicular Technology Conf. (VTC-Fall)*, 2017, pp. 1–5.

[30] X. Chen, E. A. Sobhy, Z. Yu, S. Hoyos, J. Silva-Martinez, S. Palermo, and B. M. Sadler, "A sub-Nyquist rate compressive sensing data acquisition front-end," *IEEE J. Emerg. Sel. Topics Circuits Syst.*, vol. 2, no. 3, pp. 542–571, 2012.

[31] T. Zhang, "Sparse recovery with orthogonal matching pursuit under RIP," *IEEE Trans. Inf. Theory*, vol. 57, no. 9, pp. 6215–6221, 2011.

[32] Y. D. C. S. Dallaporta, "Sparse recovery guarantees from extreme eigenvalues small deviations," *Statist. Theory, arXiv*, 2017. [Online]. Available: https://arxiv.org/abs/1604.01171v4

[33] J. A. Tropp and A. C. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Trans. Inf. Theory*, vol. 53, no. 12, pp. 4655–4666, 2007.

[34] D. L. Donoho, "For most large underdetermined systems of equations, the minimal ℓ_1-norm near-solution approximates the sparsest solution," *Commun. Pure Appl. Math.*, vol. 59, no. 7, pp. 907–934, 2006.

[35] E. J. Candés and T. Tao, "Decoding by linear programming," *IEEE Trans. Inf. Theory*, vol. 51, no. 12, pp. 4203–4215, 2005.

[36] C. Aprile, L. Baldassarre, V. Gupta, J. Yoo, M. Shoaran, Y. Leblebici, and V. Cevher, "Learning-based near-optimal area-power trade-offs in hardware design for neural signal acquisition," in *Proc. Int. Great Lakes Symp. VLSI (GLSVLSI)*, 2016, pp. 433–438.

[37] M. Mishali, Y. C. 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, 2011.

[38] M. A. Lexa, M. E. Davies, and J. S. Thompson, "Reconciling compressive sampling systems for spectrally sparse continuous-time signals," *IEEE Trans. Signal Process.*, vol. 60, no. 1, pp. 155–171, 2012.

[39] M. Mishali and Y. C. Eldar, "Blind multiband signal reconstruction: Compressed sensing for analog signals," *IEEE Trans. Signal Process.*, vol. 57, no. 3, pp. 993–1009, 2009.

[40] H. Lebesgue, *Intégrale, Longueur, Aire*. Milan, Italy: Bernandon de C. Rebeschini, 1902.

[41] M. S. O. Alink, E. A. M. Klumperink, M. C. M. Soer, A. B. J. Kokkeler, and B. Nauta, "A 50MHz-to-1.5GHz cross-correlation CMOS spectrum analyzer for cognitive radio with 89dB SFDR in 1MHz RBW," in *Proc. IEEE Symp. New Frontiers in Dynamic Spectrum*, 2010, pp. 1–6.

[42] 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 Design, Applications, Integration and Software*, 2006, pp. 71–74.

[43] 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 *Proc. IEEE Int. Symp. Circuits and Systems*, 2007, pp. 1959–1962.

[44] M. Mishali and Y. C. Eldar, "Wideband spectrum sensing at sub-Nyquist rates [Applications Corner]," *IEEE Signal Process. Mag.*, vol. 28, no. 4, pp. 102–135, 2011.

[45] L. Baldassarre, Y. H. Li, J. Scarlett, B. Gozcu, I. Bogunovic, and V. Cevher, "Learning-based compressive subsampling," *IEEE J. Sel. Topics Signal Process.*, vol. 10, no. 4, pp. 809–822, 2016.

[46] R. T. Yazicigil, *Compressive Sampling as an Enabling Solution for Energy-Efficient and Rapid Wideband RF Spectrum Sensing in Emerging Cognitive Radio Systems*. Columbia University Academic Commons, 2016. [Online]. Available: https://doi.org/10.7916/D8571BXM

[47] N. Rajput, M. Sethi, P. Dobriyal, K. Sharma, and G. Sharma, "A novel, high performance and power efficient implementation of 8×8 multiplier unit using MT-CMOS technique," in *Proc. 6th Int. Conf. Contemporary Computing (IC3)*, 2013, pp. 186–191.

[48] S. Goel, A. Kumar, and M. A. Bayoumi, "Design of robust, energy-efficient full adders for deep-submicrometer design using hybrid-CMOS logic style," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 14, no. 12, pp. 1309–1321, 2006.

[49] L. Jacques, J. N. Laska, P. T. Boufounos, and R. G. Baraniuk, "Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors," *IEEE Trans. Inf. Theory*, vol. 59, no. 4, pp. 2082–2102, 2013.

[50] J. Ziniel and P. Schniter, "Dynamic compressive sensing of time-varying signals via approximate message passing," *IEEE Trans. Signal Process.*, vol. 61, no. 21, pp. 5270–5284, 2013.

[51] S. Buchanan, T. Haque, P. Kinget, and J. Wright, "Efficient model-free learning to overcome hardware nonidealities in analog-to-information converters," in *Proc. Int. Conf. Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 3574–3578.

[52] D. Cohen and Y. C. Eldar, "Sub-Nyquist cyclostationary detection for cognitive radio," *IEEE Trans. Signal Process.*, vol. 65, no. 11, pp. 3004–3019, 2017.

[53] S. Ling and T. Strohmer, "Self-calibration and biconvex compressive sensing," *Inverse Problems*, vol. 31, no. 11, 2015. doi: 10.1088/0266-5611/31/11/115002.

SP