Deborah Cohen and Yonina C. Eldar

Sub-Nyquist Radar Systems

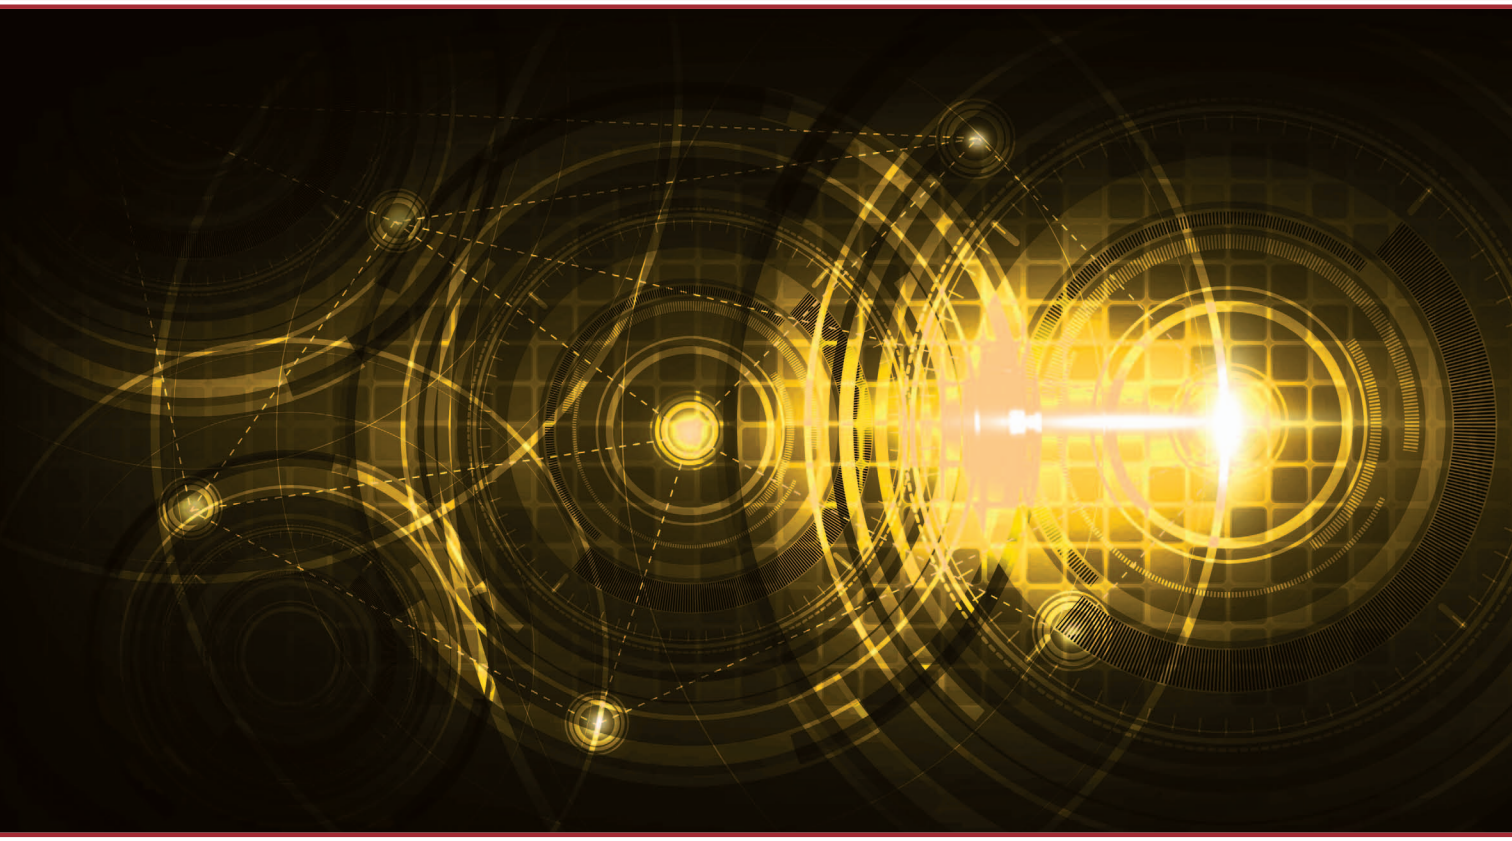

©ISTOCKPHOTO.COM/MRJUB

Temporal, spectral and spatial compression

Radar is an acronym for "radio detection and ranging." However, the functions of today's radar systems, both in civilian and military applications, go beyond simple target detection and localization; they extend to tracking, imaging, classification, and more and involve different types of radar systems, such as through-the-wall [1], ground-penetration [2], automotive [3], and weather [4]. Although radar technology has been well established for decades, a new line of compressed radars has recently emerged. These aim at reducing the complexity of classic radar systems by exploiting inherent prior information on the structure of the received signal from the targets. The goal of this article is to review these novel sub-Nyquist radars and their potential applications.

Conventional radar systems transmit electromagnetic waves of near-constant power in very short pulses toward the targets of interest. Between outgoing pulses, the radar measures the signal reflected from the targets to determine their presence, range, velocity, and other characteristics. Different systems use different radar waveforms and varying transmit strategies. One of the most popular methods is pulse-Doppler radar, which periodically transmits identical pulses. In contrast, stepped-frequency radars (SFRs) [5] vary the carrier frequency of each pulse. Some systems rely on simple traditional waveforms such as Gaussian pulses while others adopt more complex signals, such as chirps [6], [7]. Each configuration corresponds

Digital Object Identifier 10.1109/MSP.2018.2868137
Date of publication: 13 November 2018

to a certain choice in the complexity-performance tradeoff, between complex waveform and system designs and target detection and estimation.

State-of-the-art radar systems operate with large bandwidths, large coherent processing intervals (CPIs), and high number of antennas in multiple-input, multiple-output (MIMO) settings [8], [9], to achieve high-range velocity and azimuth resolution, respectively. This, in turn, generates large data sets to be sampled, stored, and processed, creating a bottleneck in terms of both analog system complexity, including high-rate analog-to-digital converters (ADCs), and subsequent digital processing [10].

Although the natural application of CS is typically the reduction of the required number of samples to perform a certain signal processing task, it was first used by the radar community to increase a target's parameter resolution.

In the past few years, novel approaches to radar signal processing have emerged that allow radar signal detection and parameter estimation using a much smaller number of measurements than that required by spatial and temporal Nyquist sampling. While temporal sampling refers to taking samples in time intervals determined by the sampling rate, spatial sampling extends this notion to placing transmit and receive antennas, whose locations are governed by the signal wavelength. These works capitalize on the fact that, in most radar applications, the reflectivity scene consists of a small number of strong targets. That is, the reflected signals by only a few targets have high enough power to be detected by the radar receiver. In pulse-Doppler radar, the target scene is often sparse in the joint time–frequency, or ambiguity, domain [5]. In synthetic aperture radar (SAR) [11], the scene is often sparse in the Fourier or wavelet domain, or even in the image domain.

Over the past decade, many works have exploited the inherent sparsity of the target scene to enhance radar-estimation capabilities. These rely on the compressed sensing (CS) [10], [12] framework, brought to the forefront by the works of Candes, Romberg, and Tao [13] and of Donoho [14]. Although the natural application of CS is typically the reduction of the required number of samples to perform a certain signal processing task, it was first used by the radar community to increase a target's parameter resolution [15]–[20]. It was later applied to reduce the number of samples to be processed [21]–[25] and finally to reduce the sampling rate [26], [27] and number of antennas [28] required in radar systems, performing time and spatial compression and alleviating the burden on both the analog and digital sides. In particular, the recently proposed Xampling, i.e., compressed sampling concept [10], [29], has been applied to radar [30]–[32] to break the link between bandwidth, CPI, and the number of antennas on the one hand, and range, Doppler, and azimuth resolution, respectively, on the other hand.

The reviews of compressive radar [22], [33]–[35] mostly deal with radar imaging. The works in [33] and [34] focus on SAR imaging and consider sparsity-based radar imagery using both greedy algorithms, which iteratively recover the sparse target scene, and convex relaxations of sparsity-inducing regularization. The special cases of interferometric, polarimetric, and circular SAR are presented in [33] for both two-dimensional (2-D) and three-dimensional (3-D) images. In [34], diverse SAR applications, such as wide-angle SAR imaging, joint imaging, and autofocusing from data with phase errors, moving targets, analysis, and design of SAR sensing missions, are reviewed. A survey of statistical sparsity-based techniques for radar imagery applications is presented in [35], including superresolution imaging, enhanced-target imaging, autofocusing, and moving-target imaging. The review of [22] presents three applications of CS radars: pulse compression, radar imaging, and airspace surveillance with array antennas. At the time it was written, there was a small number of publications addressing the application of CS to radar, as stated by the authors.

In this article, we focus on nonradar-imaging applications and survey many recent works that exploit CS in different radar systems to achieve various goals. We consider different transmit waveforms and processing approaches, while focusing on pulse-Doppler radar—one of the most popular systems—and its extension to MIMO configurations. Our goal is to review the main impacts of compressed radar on parameter resolution as well as digital and analog complexity. The survey includes fast time compression schemes, which reduce the number of acquired samples per pulse; slow time compression techniques, which decrease the number of pulses; and spatial compression approaches, in which the number of transmit and receive antenna elements is reduced. We show that, beyond a substantial rate reduction, compression may also enable communication and radar spectrum sharing [36]–[38], as elaborated on in [39]. Throughout this article, we consider both theoretical and practical aspects of compressed radar and present hardware prototype implementations [40]–[43] of the theoretical concepts, demonstrating the real-time target parameters' recovery from low-rate samples in pulse-Doppler and MIMO radars.

Radar systems

Radar systems aim to estimate targets' parameters to determine their location and motion. In its simplest form, the radar transmits a single pulse toward targets in one direction and recovers their range, i.e., distance to the radar, which is proportional to the received pulse delay. More elaborate systems are able to provide additional information on the targets. Pulse-Doppler radars transmit several pulses, enabling them to resolve both the targets' ranges and radial velocities, which are proportional to the Doppler frequency. Stepped-frequency-based approaches achieve highly effective bandwidths that increase range resolution, while allowing for narrow instantaneous bandwidth. MIMO radars use several elements both at the transmitter and at the receiver to illuminate the entire target scene and recover targets' azimuths in addition to their ranges and velocities. In this article, we consider the application of compression in

terms of the number of required samples, pulses, and antennas, as well as its impact on different aspects of the radar system, including parameter resolution and system complexity, for several types of radars.

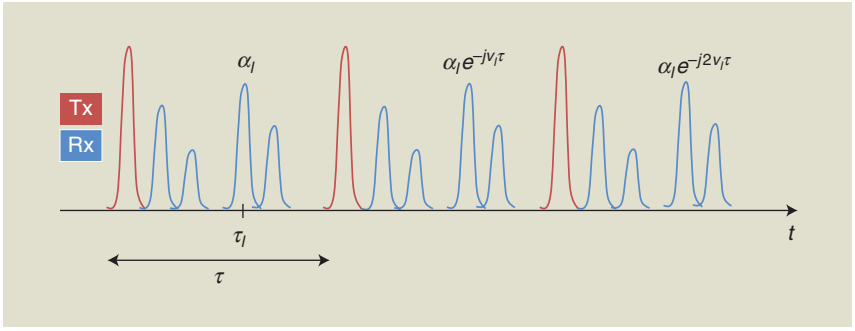

FIGURE 1. The pulse-Doppler radar transmitted and received pulse trains with $P = 3$ pulses and $L = 4$ targets Tx: transmitted; Rx: received.

Pulse-Doppler radar

A standard pulse-Doppler radar transceiver detects targets by transmitting a periodic stream of pulses and processing its reflections. The transmitted signal $x_T(t)$ consists of P equally spaced pulses $h(t)$ such that

$$x_T(t) = \sum_{p=0}^{P-1} h(t - p\tau), \quad 0 \le t \le P\tau. \tag{1}$$

The pulse-to-pulse delay τ is the pulse-repetition interval (PRI), and its reciprocal $1/\tau$ is the pulse-repetition frequency (PRF). The entire span of the signal in (1), i.e., $P\tau$, is the CPI. The pulse time support is denoted by T_p, with $0 < T_p < \tau$. The pulse $h(t)$ is typically a known time-limited baseband function with continuous-time Fourier transform (CTFT) $H(f) = \int_{-\infty}^{\infty} h(t) e^{-j2\pi ft}\, dt$ that has negligible energy at frequencies beyond $B_h/2$, where B_h is referred to as the bandwidth of $h(t)$. An example of a transmitted pulse train is illustrated in Figure 1.

It is typically assumed that the target scene is composed of L nonfluctuating point-targets, according to the Swerling-0 model [5]. This is one of the popular models in the radar signal processing literature since, by describing an idealized target, it allows simplifying the radar equations while constituting a fairly good approximation in many applications [6], [7]. Other models, such as Swerling-1, which applies to targets composed of many independent scatters, or fluctuating target models, are beyond the scope of this article. The pulses reflect off the L targets and propagate back to the transceiver. Each target l is defined by three parameters:

1) a time delay $\tau_l = 2r_l/c$, proportional to the target's distance to the radar or range r_l, where c is the speed of light
2) a Doppler-radial frequency $\nu_l = 2\dot{r}_l f_c/c$, proportional to the target's radial velocity to the radar, i.e., the target's velocity radial component $\dot{r}_l$, and the radar's carrier frequency f_c
3) a complex amplitude α_l, proportional to the target's radar cross section (RCS), dispersion attenuation, and other propagation factors.

The targets are defined in the radar radial coordinate system and are typically assumed to lie in the radar unambiguous time–frequency region: delays up to the PRI and Doppler frequencies up to the PRF. When this assumption does not hold, several processing techniques have been proposed that require the transmission of multiple pulse trains with different parameters, e.g., different PRFs. We review this setting in the "Range-Velocity Ambiguity Resolution" section.

Based on the three assumptions A1–A3 presented in "Targets' Assumptions," the received signal can be written as

$$x_R(t) = \sum_{p=0}^{P-1} \sum_{l=0}^{L-1} \alpha_l h(t - \tau_l - p\tau) e^{-j\nu_l p\tau}, \quad 0 \le t \le P\tau. \tag{2}$$

It will be convenient to express $x_R(t)$ as a sum of single frames

$$x_R(t) = \sum_{p=0}^{P-1} x_p(t), \tag{3}$$

where

$$x_p(t) = \sum_{l=0}^{L-1} \alpha_l h(t - \tau_l - p\tau) e^{-j\nu_l p\tau}, \quad 0 \le t \le P\tau. \tag{4}$$

An illustration of a received pulse train is shown in Figure 1 with $L = 4$ targets. In pulse-Doppler radar, the goal is to recover the three L parameters $\{\tau_l, \nu_l, \alpha_l\}$ for $0 \le l \le L-1$ from the received signal $x_R(t)$. In particular, estimating the time delays τ_l and Doppler frequencies ν_l enables an approximation of the targets' distances and radial velocities.

Stepped-radar waveforms

In classic pulse-Doppler radar, high-range resolution requires a large signal bandwidth. This technology bottleneck is partially overcome by stepped-frequency-based waveforms, in which the large bandwidth is obtained sequentially by stepping the frequency of each pulse, keeping the instantaneous bandwidth low. Two popular examples of such waveforms are SFRs and stepped chirps. An SFR [5] system transmits P-narrowband pulses, in which each pulse p has carrier frequency

$$f_p = f_0 + p\Delta_f, \tag{5}$$

for $0 \le p \le P-1$, with f_0 the initial frequency and Δf the frequency increment. The pth-transmitted pulse is a rectangular pulse modulated by its carrier f_p. The corresponding received signal is then of the form

Targets' Assumptions

To simplify the received signal model, the following assumptions of the targets' locations and motions are typically made [5]:

- *A1*: Far targets: The target-radar distance is large compared with the distance change during the coherent processing intervals (CPIs), which allows for constant α_l within the CPI:

$$\dot{r}_l P\tau \ll r_l \Rightarrow \nu_l \ll \frac{f_c \tau_l}{P\tau}. \tag{S1}$$

- *A2*: Slow targets: The constant-Doppler phase during pulse time,

$$\nu_l T_p \ll 1, \tag{S2}$$

and low target velocity allows for constant τ_l during the CPI. This condition holds when the baseband Doppler frequency is smaller than the frequency resolution:

$$\frac{2\dot{r}_l B_h}{c} \ll \frac{1}{P\tau} \Rightarrow \nu_l \ll \frac{f_c}{P\tau B_h}. \tag{S3}$$

- *A3*: Small acceleration: The target velocity remains approximately constant during the CPI allowing for constant ν_l. This condition is satisfied when the velocity change induced by acceleration is smaller than the velocity resolution:

$$\ddot{r}_l P\tau \ll \frac{c}{2f_c P\tau} \Rightarrow \ddot{r}_l \ll \frac{c}{2f_c (P\tau)^2}. \tag{S4}$$

Although these assumptions may seem hard to comply with, they all rely on slow enough relative motion between the radar and its targets. Radar systems tracking people, ground vehicles, and sea vessels usually comply quite easily [6].

In multiple-input, multiple-output settings, two additional assumptions are adopted on the array structure and transmitted waveforms:

- *A4*: Collocated array: The target radar cross sections α_l and θ_l are constant over the array [44].
- *A5*: Narrowband waveform: A small aperture allows τ_l to be constant over the channels:

$$\frac{2Z\lambda}{c} \ll \frac{1}{B_h}. \tag{S5}$$

$$x_p(t) = \sum_{l=0}^{L-1} \alpha_l \text{rect}(t-\tau_l) e^{-j2\pi f_p (t-\tau_l)} e^{j\nu_l p\tau}. \tag{6}$$

To process the received signal, the delay is neglected in the signal envelope because of the narrowband assumption. An SFR traditionally obtains one sample from each received pulse and computes the phase detector output sequence as

$$y_p = \sum_{l=0}^{L-1} \alpha_l e^{j2\pi f_p \tau_l} e^{j\nu_l p\tau}. \tag{7}$$

The phase detector signal y_p can be modeled as the product of the received signal (6) and the reference signal, followed by a low-pass filter (LPF). Conventional processing applies an inverse discrete Fourier transform (DFT) on the output to estimate the targets' time delays τ_l and Doppler frequencies ν_l. The range resolution achieved by SFR is $c/2P\Delta_f$, where $P\Delta_f$ is the total effective bandwidth of the signal over P pulses.

Another popular stepped waveform is the stepped chirp or multifrequency chirp signal. The corresponding transmitted signal is given by

$$x_T(t) = \sum_{p=0}^{P-1} e^{j\phi_p} \text{rect}\Big(\frac{t}{\tau}\Big) e^{j2\pi\left(f_p t + \frac{\gamma}{2}t^2\right)}, \tag{8}$$

where γ is the common chirp rate and f_p and ϕ_p are the frequency and complex phase of the pth subcarrier. The returned signal corresponding to the pth pulse, given by

$$x_p(t) = \sum_{l=0}^{L-1} \alpha_l e^{j\phi_p} \text{rect}\Big(\frac{t-\tau_l}{\tau}\Big) e^{j2\pi\left(f_p(t-\tau_l) + \frac{\gamma}{2}(t-\tau_l)^2\right)}, \tag{9}$$

is dechirped with a reference linear-frequency waveform of fixed frequency equal to the first carrier f_0:

$$m(t) = \text{rect}\Big(\frac{t-\tau_r}{\tau}\Big) e^{-j2\pi\left(f_0 t + \frac{\gamma}{2}t^2\right)}. \tag{10}$$

The receive window is $\tau_r = 2(r_{\max} + r_{\min})/c$, and the reference delay is $t_r = (r_{\max} + r_{\min})/c$, with $r_{\max}$ and $r_{\min}$ as the maximal and minimal ranges, respectively. The resulting dechirped received signal can be written as

$$x_p(t) = \sum_{l=0}^{L-1} \alpha_l e^{j(\phi_p - 2\pi f_p \tau_l)} \text{rect}\Big(\frac{t-\tau_l+t_r}{\tau}\Big) e^{j2\pi(f_p - f_0 - \gamma\tau_l)t}. \tag{11}$$

Classic processing of the received signal includes a DFT operation to recover the targets' delays τ_l.

MIMO pulse-Doppler radar

MIMO radar presents significant potential for advancing state-of-the-art modern radar in terms of flexibility and performance. This configuration [8] combines several antenna elements both at the transmitter and receiver. Unlike phased-array systems, each transmitter radiates a different waveform, which offers more degrees of freedom (DoF) [9]. There are two main configurations of MIMO radar, depending on the location of the transmitting and receiving elements; collocated MIMO

[46], in which the elements are close to each other relative to the signal wavelength, and multistatic MIMO [47], where they are widely separated. In this article, we focus on collocated pulse-Doppler MIMO systems.

Collocated MIMO radar systems exploit waveform diversity, based on mutual orthogonality of the transmitted signals [9]. Consequently, the performance of MIMO systems can be characterized by a virtual array constructed by the convolution of the locations of the transmit and receive antenna locations. In principle, with the same number of antenna elements, this virtual array may be much larger than the array of an equivalent traditional system [48].

The standard approach to collocated MIMO adopts a virtual uniform linear array (ULA) structure [49], where R receivers, spaced by $\lambda/2$ and T transmitters and spaced by $R(\lambda/2)$ (or vice versa), form two ULAs. Here, λ is the signal wavelength. Coherent processing of the resulting TR channels generates a virtual array equivalent to a phased array with TR $(\lambda/2)-$spaced receivers and normalized aperture $Z = TR/2$. Denote by $\{\xi_m\}_{m=0}^{T-1}$ and $\{\zeta_q\}_{q=0}^{R-1}$, the normalized transmitters' and receivers' locations, respectively. For the traditional virtual ULA structure, denote $\zeta_q = q/2$ and $\xi_m = Rm/2$. This standard-array structure and the corresponding virtual array are illustrated in Figure 2 for $R = 3$ and $T = 5$. The circles represent the receivers, and the squares are the transmitters.

Each transmit antenna sends P pulses, such that the mth-transmitted signal is given by

$$s_m(t) = \sum_{p=0}^{P-1} h_m(t - p\tau)e^{j2\pi f_c t}, \quad 0 \le t \le P\tau, \tag{12}$$

where $h_m(t), 0 \le m \le T-1$ are orthogonal pulses with bandwidth B_h and modulated with carrier frequency f_c. For convenience, it is typically assumed that $f_c\tau$ is an integer, so that the initial phase for every pulse $e^{-j2\pi f_c \tau p}$ is canceled in the modulation for $0 \le p \le P-1$ [6].

MIMO radar architectures impose several requirements on the transmitted waveform family. Besides traditional demands from radar waveforms such as low sidelobes, MIMO transmit antennas rely on orthogonal waveforms. In addition, to avoid cross talk between the T signals and form TR channels, the orthogonality condition should be invariant to time shifts, that is $\int_{-\infty}^{\infty} s_i(t)s_j^*(t-\tau_0)dt = \delta(i-j)$ for $i, j \in [0, T-1]$ and for all τ_0. The main waveform families typically considered are time-division multiple access (TDMA), frequency-division multiple access (FDMA), and code-division multiple access (CDMA), respectively. Time-invariant orthogonality is achieved by FDMA and TDMA and approximately achieved by CDMA, as the latter involves overlapping frequency bands [50].

Besides the traditional assumptions on the targets, MIMO systems present additional requirements on the radar array and waveforms with respect to the targets, as described in "Targets' Assumptions." In the MIMO configuration, the goal is to recover the targets' azimuth angles θ_l in addition to their delays τ_l and Doppler shifts ν_l from the received signals.

Current challenges

Standard radar processing samples and processes the received signal at its Nyquist rate B_h. For example, the pulse-Doppler classic radar processing, described in "Classic Pulse-Doppler and Multiple-Input, Multiple-Output Processing," first filters the sampled signal by a matched filter (MF). In modern systems, the MF operation is performed digitally and therefore requires an ADC capable of sampling at rate B_h. Other radar systems similarly require sampling the received signal at its Nyquist rate. The radar bandwidth B_h is inversely proportional to the system fast time, or range resolution, and can thus be hundreds of megahertz or even up to several gigahertz, requiring a high sampling rate and resulting in a large number of samples per pulse $N = \tau B_h$ to process.

The slow time (Doppler) resolution is inversely proportional to the CPI $P\tau$. The Doppler processing stage can be viewed as an MF in the pulse dimension, i.e., slow time domain, to a constant radial velocity target. As such, it increases the signal-to-noise ratio (SNR) by P compared to the SNR of a single pulse [7]. Since an MF is the linear time-invariant system that maximizes SNR, it follows that a factor P increase is optimal for P pulses. A large number of pulses increases resolution and SNR but leads to large time on target and a large total number of samples to process, given by PN.

The required computational power corresponds to P convolutions of a signal of length $N = \tau B_h$ and N-fast Fourier transforms (FFTs) of length P (see "Classic Pulse-Doppler and Multiple-Input, Multiple-Output Processing"). The growing demands for improved estimation accuracy and target separation dictate an ever-growing increase in the signal's bandwidth and CPI. This creates bottlenecks in sampling and processing rates in the fast time (intrapulse) domain and in time on target in the slow time (interpulse) dimension.

In MIMO radar, the additional spatial dimension increases the system's complexity, as may be seen in "Classic Pulse-Doppler and Multiple-Input, Multiple-Output Processing." In such systems, the array aperture determines the azimuth resolution. In a traditional virtual array configuration, the product between the number of transmit and receive antennas scales linearly with the aperture. Consequently, high resolution

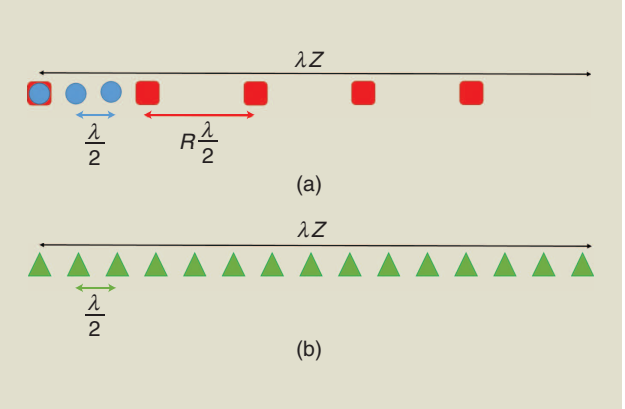

FIGURE 2. An illustration of MIMO arrays: (a) a standard array and (b) a corresponding receiver virtual array [32].

Classic Pulse-Doppler and Multiple-Input, Multiple-Output Processing

The classic methods for radar processing typically consist of the following stages [5], [45]:

1) *Sampling*: Sample each incoming frame $x_p(t)$ at its Nyquist rate B_h, equal to the double-sided bandwidth of $h(t)$, creating the samples $x_p[n], 0 \le n \le N-1$, where $N = \tau B_h$. We assume, for simplicity, that N is an integer.
2) *Matched filter* (*MF*): Apply a standard MF on each frame $x_p[n]$. This results in the outputs $y_p[n] = x_p[n] * h[-n]$, where $h[n]$ is the sampled version of the transmitted pulse $h(t)$ at its Nyquist rate and $*$ is the convolution operation. The time resolution attained in this step is $1/B_h$.
3) *Doppler processing*: For each discrete time n, perform a P-point discrete Fourier transform along the pulse dimension, i.e., $z_n[k] = \mathrm{DFT}_P\{y_p[n]\} = \Sigma_{p=0}^{P-1} y_p[n]\, e^{-i2\pi pk/P}$ for $0 \le k \le P$. The Doppler resolution is $1/\mathrm{P}\tau$.
4) *Delay-Doppler map*: Stacking the vectors $\mathbf{z}_n$ and taking absolute value, we obtain a delay-Doppler map $\mathbf{Z} = \mathrm{abs}[\mathbf{z}_0, \ldots, \mathbf{z}_{N-1}] \in \mathbb{R}^{P \times N}$.
5) *Peak detection*: A heuristic detection process, in which knowledge of the number of targets, targets' powers, clutter location, and so on, may help in discovering targets' positions. For example, if we know there are L targets, then we can choose the L-strongest points in the map. Alternatively, constant false alarm (FA) rate detectors determine the power threshold, above which a peak is considered to originate from a target so that a required probability of FA is achieved.

Classic collocated multiple-input, multiple-output radar processing traditionally includes the following stages:

1) *Sampling*: At each receiver $0 \le q \le R-1$, where R denotes the number of receivers, the signal $x_q(t)$ is sampled at its Nyquist rate Btot. In code-division multiple access and time-division multiple access, $B_{\mathrm{tot}} = B_h$ as all waveforms overlap in frequency, whereas in frequency-division multiple access, $B_{\mathrm{tot}} = TB_h$, where B_h denotes the bandwidth of a single waveform in both cases, and T is the number of transmitters.
2) *MF*: The sampled signal is convolved with a sampled version of $h_m(t)$, for $0 \le m \le T-1$. The time resolution attained in this step is $1/B_{\mathrm{tot}}$.
3) *Beamforming*: The correlations between the observation vectors from the previous step and the steering vectors corresponding to each azimuth on the grid defined by the array aperture are computed. The spatial resolution attained in this step is $2/TR$.
4) *Doppler detection*: The correlations between the resulting vectors and Doppler vectors, with Doppler frequencies lying on the grid defined by the number of pulses, are computed. The Doppler resolution is $1/\mathrm{P}\tau$.
5) *Peak detection*: This is similar to classic radar, but detection is performed on the three-dimensional range-azimuth-Doppler map.

requires a large number of antennas, thus increasing the system's complexity in terms of hardware and processing.

In the following sections, we review fast time-compressed radar systems that allow for low-rate sampling and processing of radar signals, regardless of their bandwidth, while retaining the same SNR scaling. We then demonstrate how compression can be extended to the slow time, thereby reducing time on target, and to the spatial dimension allowing one to achieve resolution similar to a filled array but with significantly fewer elements. In reality, the received signal $x_R(t)$ is further contaminated by additive noise and clutter. We will also demonstrate the impact of SNR and clutter on compressed radar system prototypes [30], [40], [54]. Finally, we show how compression and sub-Nyquist sampling may be used to address other challenges, such as communication and radar spectrum sharing.

Increased parameter resolution

In many radar applications, the reflectivity scene consists of a small number L of strong targets. Therefore, CS techniques (see "Compressed Sensing Recovery") are a natural processing tool for radar systems. Shortly after the idea of CS was brought forward by the works of Candes, Romberg, and Tao [13] and of Donoho [14] a decade ago, it was introduced to pulse-Doppler radar [15], [16], [55] and SFR [17].

While CS is typically applied to signal processing tasks to reduce the associated sampling rate [10], earlier papers that applied CS recovery to pulse-Doppler radar and SFR were aimed at increasing delay-Doppler resolution [15]–[17], [20] using Nyquist samples. More recent approaches use CS recovery techniques on low-rate, or sub-Nyquist samples, which enable sampling and processing rate reduction while achieving the same resolution as traditional Nyquist radars. Later in this section, we review radar recovery methods that increase delay-Doppler resolution using CS techniques on Nyquist samples. In the next sections, we consider the application of CS to reduce the fast time-sampling rate and the number of pulses and antennas, while preserving the resolution achieved by Nyquist systems.

In the works of [15]–[17] and [20], the signal is still sampled at its Nyquist rate B_h, but the delay and Doppler resolutions are determined by the CS grid containing $N > \tau B_h$ grid points, rather than the signal's bandwidth and CPI, respectively. The key idea in [15], which adopts a pulse-Doppler radar model, is that the received signal $x_R(t)$ defined in (2) is generally a sparse superposition of time- and frequency-shifted replicas of the transmitted waveforms. The time–frequency plane is discretized into an $N \times N$ grid in which each point represents a unique time–frequency shift $\mathbf{H}_i$, expressed as the product of

Compressed Sensing Recovery

Compressed sensing (CS) [10], [12] is a framework for simultaneous sensing and compression of finite-dimensional vectors, which relies on linear dimensionality reduction. In particular, the field of CS focuses on the recovery problem

$$\mathbf{z} = \mathbf{A}\mathbf{x}, \tag{S6}$$

where $\mathbf{x}$ is an $N \times 1$ sparse vector, i.e., with few nonzero entries, and $\mathbf{z}$ is a vector of measurements of size $M < N$. CS provides recovery conditions and algorithms to reconstruct $\mathbf{x}$ from the low-dimensional vector $\mathbf{z}$.

Two popular CS greedy recovery algorithms, orthogonal matching pursuit (OMP) and iterative hard thresholding (IHT), attempt to solve the optimization problem

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{x}\|_0 \quad \text{s.t. } \mathbf{z} = \mathbf{A}\mathbf{x}, \tag{S7}$$

where $\|\cdot\|_0$ denotes the ℓ_0-norm. OMP [51], [52] iteratively proceeds by finding the column of A most correlated to the signal residual $\mathbf{r}$,

$$i = \text{argmax} \mid \mathbf{A}^H \mathbf{r} \mid, \tag{S8}$$

where the absolute value is computed element-wise and $(\cdot)^H$ is the Hermitian operator. The residual is obtained by subtracting the contribution of a partial estimate $\hat{\mathbf{x}}_\ell$ of the signal at the ℓth iteration, from $\mathbf{z}$, as follows:

$$\mathbf{r} = \mathbf{z} - \mathbf{A}\hat{\mathbf{x}}_\ell. \tag{S9}$$

It is initialized by $\mathbf{r} = \mathbf{z}$. Once the support set is updated by adding the index i, the coefficients of $\hat{\mathbf{x}}_\ell$ over the support set are updated, so as to minimize the residual error.

Other greedy techniques include thresholding algorithms. We focus here on the IHT method proposed in [53]. Starting from an initial estimate $\hat{\mathbf{x}}_0 = 0$, the algorithm iterates a gradient-descent step with step size μ followed by hard thresholding, i.e.,

$$\hat{\mathbf{x}}_\ell = \mathcal{T}(\hat{\mathbf{x}}_{\ell-1} + \mu \mathbf{A}^H(\mathbf{z} - \mathbf{A}\hat{\mathbf{x}}_{\ell-1}), k), \tag{S10}$$

until a convergence criterion is met. Here, $\mathcal{T}(\mathbf{x}, k)$ denotes a thresholding operator on $\mathbf{x}$ that sets all but the k entries of $\mathbf{x}$ with the largest magnitudes to zero, and k is the sparsity level of $\mathbf{x}$ (assumed to be known).

Alternative approaches to greedy recovery are convex-relaxation-based methods using ℓ_1 regularization such as basis pursuit and least absolute shrinkage and selection operator, better known as *LASSO*. Further details on CS recovery conditions and techniques can be found in [10] and [12].

time-shift and frequency-modulation matrices, denoted by $\mathbf{T}^{(\cdot)}$ and $\mathbf{M}^{(\cdot)}$, respectively. In particular,

$$\mathbf{H}_i = \mathbf{M}^{i \bmod N} \mathbf{T}^{\lfloor i/N \rfloor}, \tag{13}$$

where

$$\mathbf{T} = \begin{pmatrix} 0 & 0 & & 1 \\ 1 & 0 & & 0 \\ & \ddots & \ddots & \\ 0 & & 1 & 0 \end{pmatrix}, \quad \mathbf{M} = \begin{pmatrix} 1 & & & 0 \\ & e^{j\frac{2\pi}{N}} & & \\ & & \ddots & \\ 0 & & & e^{j\frac{2\pi}{N}(N-1)} \end{pmatrix}. \tag{14}$$

Here, $\lfloor \cdot \rfloor$ and mod denote the floor and modulation functions, respectively.

The vector $\mathbf{y}$ that concatenates the Nyquist samples of a single pulse $x_p(t)$ can then be expressed as

$$\mathbf{y} = \mathbf{\Phi}\mathbf{s}, \tag{15}$$

where $\mathbf{s}$ is the L-sparse vector of size N^2 whose nonzero entries are the targets' RCS α_l with locations determined by the corresponding time–frequency shift. The ith column, i.e., atom of the $N \times N^2$ matrix $\mathbf{\Phi}$, is given by

$$\mathbf{\Phi}_i = \mathbf{H}_i \mathbf{f}, \tag{16}$$

where the vector $\mathbf{f}$ contains the Nyquist rate samples $h[n]$ of the transmitted signal $h(t)$. The latter is chosen so that the samples correspond to the Alltop sequence $h[n] = \left(1/\sqrt{N}\right) e^{2\pi j n^3/N}$ [56], for some prime $N \geq 5$. This yields a low-coherence matrix $\mathbf{\Phi}$, i.e., a matrix whose columns have small correlation.

The vector $\mathbf{s}$ is reconstructed from $\mathbf{y}$ using CS techniques, as described in "Compressed Sensing Recovery." The time–frequency shifts, determined by the targets' delays and Doppler frequencies, are thus recovered with a resolution of $1/N$.

The CS recovery in [15] is performed without an MF, which reduces performance in low-SNR regimes. Additionally, [15] considers only delay recovery. Alternatively, CS techniques can be performed after applying an analog MF [16] on the pulse-Doppler-received signal (2). The MF output of the pth pulse, sampled at the Nyquist rate $1/B_h$, is given by

$$w_p[k] = \sum_{l=0}^{L-1} \alpha_l e^{j\nu_l \tau_l} e^{j\nu_l p\tau} C_h[k - \tau_l/\tau], \tag{17}$$

where $C_h[k]$ is the discrete autocorrelation function of the transmitted waveform. For each sampling time k, the Nyquist samples have a sparse representation in the frequency (Doppler) domain using a Fourier matrix as a dictionary. A

two-step approach is therefore proposed to apply CS recovery for each k. However, the sidelobes of $C_h[k]$ lead to ambiguity. To avoid this, pairs of Golay complementary sequences x_1 and x_2 of length N, whose correlation functions satisfy

$$C_{x_1}[k] + C_{x_2}[k] = 2N\delta[k], \tag{18}$$

are transmitted alternatively by phased coding of the baseband waveform. This allows for unambiguous delay-Doppler recovery, provided that all of the Doppler coordinates are within the interval $[-\pi/2, \pi/2]$.

CS has also been applied to SFR to increase the range resolution [17]. As in pulse-Doppler radar, the target scene is discretized over an $N \times N$ delay-Doppler map [17]. The outputs of the phase detector (7) are then expressed, as in (15), where $\mathbf{y}$ is the vector of size P with the pth entry given by y_p, and $\mathbf{\Phi}$ is a DFT-based dictionary such that

$$\mathbf{\Phi}_{(p,(i-1)N+k)} = e^{j2\pi f_p \tau_i} e^{j\nu_k p\tau}. \tag{19}$$

The vector $\mathbf{s}$ is then recovered from $\mathbf{y}$ using CS techniques.

The approaches mentioned in this section may increase resolution by taking a large grid size N. However, bounds on N are not discussed, and it is not clear how large it can be. Denser grids reduce the sensitivity of the reconstruction to off-grid targets but increase the computational complexity by a square factor since the dictionaries contain N^2 atoms. More importantly, higher grid dimensions cause a significant increase to the coherence of the CS dictionary, which may degrade recovery performance.

The parameter space discretization, typically used in CS recovery techniques, assumes the targets' delays and Dopplers lie on the predefined grid. Several approaches have been proposed to solve off-grid issues, including grid refinement, which adjusts the detected delay-Doppler peak [32], parameter, perturbation-based, adaptive-sparse reconstruction techniques [21], and sensing matrix perturbation [57]. More references may be found in [58].

Fast time compression

In the works reviewed thus far, sampling and digital processing are still performed at the Nyquist rate. We next consider compressed radar that reduces sampling and processing rates.

Random sampling

Random sampling has been considered in SFR systems by selecting random measurements out of the Nyquist samples [21], [22]. The SFR approach of (1) is adopted in [22], with a random selection of M out of P pulses with different carriers. The sparse representation of the received signal used is a delay-Doppler shifted dictionary [21] similar to [15]. Consider the matrix $\mathbf{\Phi}$ whose ith column is given by

$$\mathbf{\Phi}_i = h(\mathbf{t} - \tau_i) \circ e^{j2\pi\nu_i \mathbf{t}}, \tag{20}$$

where $\mathbf{t}$ is the $N \times 1$ vector containing the sampling instants at the Nyquist rate, i.e., $t_i = i/B_h$, and $\circ$ is the Hadamard product operator. As in [15], the dictionary $\mathbf{\Phi}$ contains N^2 atoms. The Nyquist samples can then be expressed in the form (15), and the compressed samples $\mathbf{z}$ are given by

$$\mathbf{z} = \mathbf{A}\mathbf{y}, \tag{21}$$

where $\mathbf{A}$ is an $M \times N$ matrix, with $M < N$ constructed by randomly selecting M rows of the $N \times N$ identity matrix, which corresponds to the M-selected pulses.

In these approaches, processing is performed at a low rate; however, the random discarding of samples is difficult to implement in a sampling system for the purpose of effectively reducing the sampling rate. Furthermore, the large dictionary size discussed in the previous section remains an issue. Alternative practical radar systems using CS to reduce the sampling rate have been proposed and rely on two main techniques: uniform low-rate sampling using appropriate waveforms and analog preprocessing.

Uniform low-rate sampling

In [26], the authors consider SFR using multifrequency chirps, as described in (8). Low-rate samples are uniformly taken from the received signal (11) at rate $2\gamma\tau_r$, with $\tau_r = 2(r_{\max} - r_{\min})/c$, with γ being the common chirp rate. This results in the aliasing of the multiple sinusoids to baseband with random complex coefficients. Upon discretization of the target range, as denoted by $\mathbf{s}$, the low-rate samples may be modeled as

$$\mathbf{y} = \mathbf{A}\mathbf{s}. \tag{22}$$

Here, the kth column of the sensing matrix $\mathbf{A}$ is the FFT of the samples of (11) for a singular target at range bin k corresponding to a delay of $\tau_l = 2(r_{\min} - k\Delta)/c$, where Δ is the range-discretization step. The targets' delays are therefore recovered from the low-rate uniform sampling of the chirp waveforms.

Random demodulation

Many analog-to-information-conversion systems have been proposed to sample wideband signals at sub-Nyquist rates. Among them, the random demodulator (RD) [59], random-modulation preintegrator (RMPI) [60], and Xampling-based [29] systems have been used for radar applications. All three approaches consider pulse-Doppler radar.

The RD modulates the input signal using a high-rate sequence $p(t)$ created by a pseudorandom number generator, aliasing its frequency content. The random sequence used for demodulation is a square wave, which alternates between the levels ± 1 with equal probability. The mixed output is filtered by a bandpass filter $h_{bp}(t)$, with center frequency f_c and bandwidth $B_{CS} \ll B_h$, and sampled at a low rate, as shown in Figure 3(a).

The RD is adopted in [27] as the analog-mixing front end of a proposed quadrature-compressive-sampling (referred to as *QuadCS* by Liu et al.) system. The mixed and filtered output $y(t)$, shown in Figure 3, is given by

$$y(t) = \int_{-\infty}^{\infty} h_{bp}(\rho) p(t-\rho) x_R(t-\rho)\,\mathrm{d}\rho, \tag{23}$$

where $x_R(t)$ is defined as the real part of (7). The RD samples $y(t)$ at rate $f_s = 1/T_s = f_c/k$, with k an integer satisfying $k \le \lfloor f_c/2B_{CS} \rfloor$. The samples are fed to the quadrature-processing system [61], which extracts the baseband in-phase and quadrature (I and Q) components, respectively, of the radar echoes. As shown in [27], the complex samples of the RD output can be written as

$$\mathbf{y} = \mathbf{A}\mathbf{x}. \tag{24}$$

Here, $\mathbf{x}$ is a sparse vector that contains the complex amplitudes α_l at the corresponding delays τ_l, and the (m, p) element of the matrix $\mathbf{A}$ is given by

$$\mathbf{A}_{m,p} = \int_{-\infty}^{\infty} h_{bp}(\rho) e^{-j2\pi f_c \rho} p(mT_s - \rho) h(mT_s - p\tau - \rho) \mathrm{d}\rho. \tag{25}$$

The samples of P pulses are concatenated in a matrix $\mathbf{Y}$ such that each column corresponds to a pulse. The subsequent processing of the quadrature-compressive sampling, referred to as *compressive-sampling pulse Doppler*, is composed of a DFT step on the rows of $\mathbf{Y}$ that acts as an MF in slow time followed by a MF in each column, corresponding to the fast time.

The RMPI is a variant of the RD composed of a parallel set of RD channels driven by a common input, in which each RD uses a distinct pseudorandom binary sequence. A hardware RMPI-based prototype has been implemented in [43] that recovers radar pulses and estimates their amplitude, phase, and carrier frequency. In the next section, we discuss an alternative prototype with a different analog front end, which also recovers the targets' parameters from low-rate samples.

Note that the considerations behind waveform design for CS recovery in the approaches [15]–[17] presented in the previous section are similar to traditional radar requirements. The well-known ambiguity function (AF) impacts CS radar in a way that is similar to traditional radar systems. Indeed, the mutual coherence of the dictionary is linearly related to the highest sidelobe value of the AF [58], [62]. In contrast, we will see in the next section that the CS dictionary of the Xampling method is independent of the waveform, and MF is performed directly on the low-rate samples before parameter recovery.

Fast time Xampling

An alternative sub-Nyquist radar method is the Xampling-based system proposed in [30] and [40]. This approach, which may be used with any transmitted pulse shape, achieves the minimal sampling rate required for target detection, while providing optimal SNR.

The sub-Nyquist analog front end is composed of an ADC that filters the received pulse-Doppler signal (2) to predetermined frequencies before taking pointwise samples. These compressed samples, or "Xamples," contain the information needed to recover the desired signal parameters, i.e., the target' delay-Doppler map. To see this, note that the Fourier-series coefficients of the aligned frames $x_p(t + p\,\tau)$ are given by

$$c_p[k] = \frac{1}{\tau} H[k] \sum_{l=0}^{L-1} \alpha_l e^{-j2\pi k \tau_l/\tau} e^{-j\nu_l p\tau}, \quad 0 \le k \le N-1, \tag{26}$$

where $H[k]$ are the Fourier coefficients of the known transmitted pulse $h(t)$, and $N = B_h\tau$ is the number of Fourier samples. From (26), we see that the unknown parameters $\{\alpha_l, \tau_l, \nu_l\}_{l=0}^{L-1}$ are contained in the Fourier coefficients $c_p[k]$. We now show how the Fourier coefficients $c_p[k]$ may be obtained from low-rate samples of $x_p(t)$ and how the targets' parameters can then be recovered from $c_p[k]$ [30].

The received signals $x_p(t)$ exist in the time domain; thus, there is no direct access to $c_p[k]$. To obtain any arbitrary set of Fourier-series coefficients, the direct multichannel sampling scheme [63] illustrated in Figure 4 can be used. The analog input $x_p(t)$ is split into $k = |\kappa|$ channels, where, in each channel k_i with $i \in [0, K-1]$, it is mixed with the harmonic signal $e^{-j2\pi k_i t/\tau}$, integrated over the PRI duration, and then sampled. Xampling thus allows one to obtain an arbitrary set κ out of $N = \tau B_h$ frequency components from K pointwise samples of the received signal after appropriate analog preprocessing. An alternative Xampling method uses the sum-of-sincs filter described in [64]. This class of filters, which consists of a sum-of-sinc function

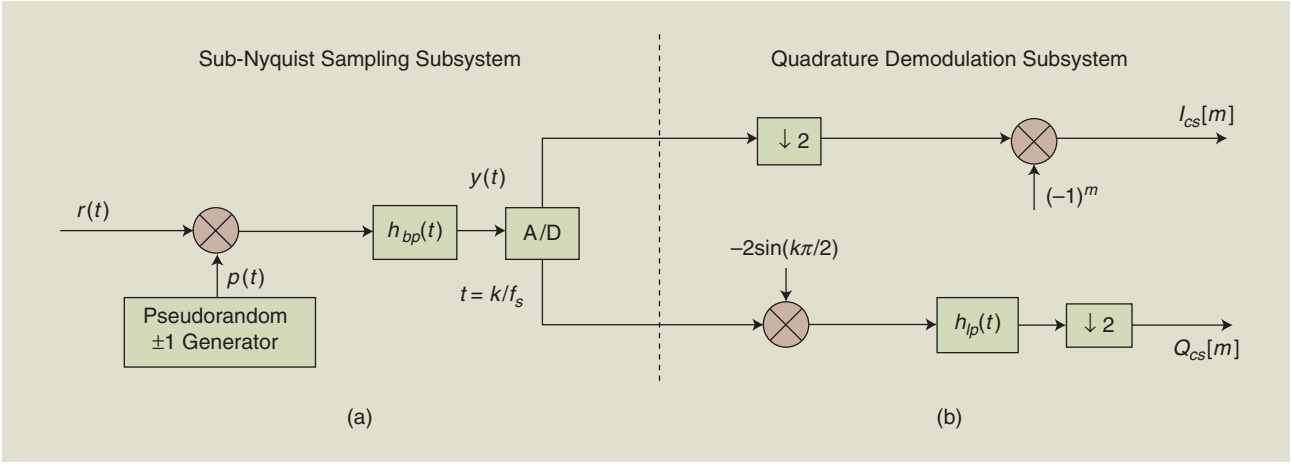

FIGURE 3. A quadrature-compressive-sampling implementation with (a) RD sampling followed by (b) quadrature demodulation [27].

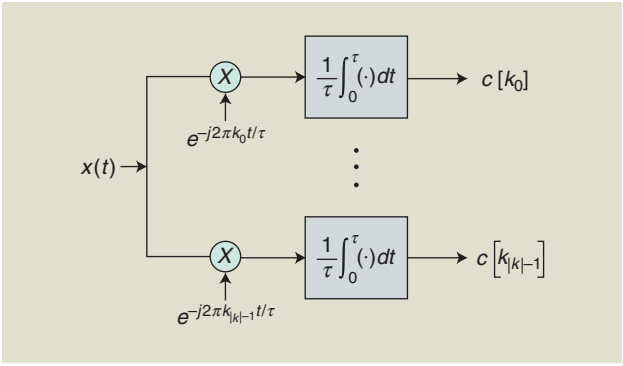

FIGURE 4. A multichannel direct sampling of the Fourier coefficients [63].

in the frequency domain, is a general sampling scheme for arbitrary pulse shapes.

A less expensive and more practical approach for the Fourier-series coefficients acquisition proposed in [40] is based on multiple bandpass filters and is adopted in the Xampling hardware radar prototype described in the next section. This system is composed of a few channels, with each sampling the content of a narrow frequency band of the received signal. Each channel thus yields a group of several consecutive Fourier coefficients. The multiple bandpass constellation has the advantage of acquiring the measurements over a wider frequency aperture. At the same time, it still allows for practical hardware implementation, as detailed in the next section. By widening the frequency aperture, a finer resolution grid may be employed during the recovery process. Moreover, empirical results show that highly distributed frequency samples provide better noise robustness [40]. However, widening the frequency aperture eventually requires increasing the number of samples K; otherwise, recovery performance may degrade. This tradeoff is observed in the experiments presented in [40].

Once a set of Fourier coefficients $c_p[k]$ has been acquired, the delays and Doppler frequencies can be recovered using different techniques. Doppler focusing [30], summarized in "Doppler Focusing," is one approach that has several advantages, as detailed next. This method uses target echoes from all of the pulses to generate a focused pulse at a specific Doppler frequency. It then jointly recovers the delay-Doppler map by reducing the detection problem to a one-dimensional, delay-only estimation. Performing the Doppler focusing operation in frequency results in computing the DFT of the coefficients $c_p[k]$ in the slow time domain:

$$\begin{aligned}\Psi_\nu[k] &= \sum_{p=0}^{P-1} c_p[k]e^{jp\nu\tau} \\ &= \frac{1}{\tau}H[k]\sum_{l=0}^{L-1}\alpha_l e^{-j2\pi k\tau_l/\tau}\sum_{p=0}^{P-1}e^{j(\nu-\nu_l)p\tau}.\end{aligned} \tag{27}$$

Note that $\Psi_\nu[k]$ is the Fourier series of $\Phi(t,\nu)$, defined in (S11), with respect to t. Following the same argument as in (S12), we have

$$\Psi_\nu[k] \approx \frac{P}{\tau}H[k]\sum_{l\in\Lambda(\nu)}\alpha_l e^{-j2\pi k\tau_l/\tau}. \tag{28}$$

The resulting equation (27) is a standard delay-estimation problem for each ν and may be solved using multiple techniques [10]. However, improved performance can be obtained by jointly processing the sequences $\{\Phi_\nu[k]\}$ for different values of ν. Thus, instead of searching separately for each of the delays $\tau_l, l\in\Lambda(\nu)$, the L delays are estimated by jointly processing overall Doppler frequencies.

A particularly convenient method in this case is to employ a matching pursuit-type approach in which the strongest overall peak ν, assuming a single delay, is first found:

$$(\hat{\tau}_l,\hat{\nu}_l) = \operatorname{argmax}_{\tau_l,\nu_l}\left|\sum_{k\in\kappa}\Psi_{\nu_l}[k]e^{j2\pi k\tau_l/\tau}\right|. \tag{29}$$

Once the optimal values $\hat{\tau}_l$ and $\hat{\nu}_l$ are determined, their influence is subtracted from the focused sub-Nyquist samples as

$$\Psi'_\nu[k] = \Psi_\nu[k] - \frac{1}{\tau}\hat{\alpha}_l e^{-j2\pi k\hat{\tau}_l/\tau}\sum_{p=0}^{P-1}e^{j(\nu-\hat{\nu}_l)p\tau}, \tag{30}$$

where

$$\hat{\alpha}_l = \frac{\tau}{P|\kappa|}\sum_{k\in\kappa}\Psi_{\hat{\nu}_l}[k]e^{j2\pi k\hat{\tau}_l/\tau}. \tag{31}$$

The same operations are performed iteratively to find all of the desired L peaks. This approach does not require discretization of the targets' parameters, and these are recovered over the continuous domain from a minimal number of samples.

In practice, the search for peaks can be limited to a grid, which enables all of the computations to be carried out using simple FFT operations. Suppose we limit ourselves to the Nyquist grid, i.e., the grid defined by the Nyquist resolution so that $\tau_l/\tau = s_l/N$, where s_l is an integer satisfying $0\le s_l\le N-1$. Then, (26) is approximately written in vector form as

$$\boldsymbol{\Psi}_\nu = P\mathbf{H}\mathbf{F}_N^K\mathbf{a}_\nu, \tag{32}$$

where $\boldsymbol{\Psi}_\nu = [\Psi_\nu[k_0]\ldots\Psi_\nu[k_{K-1}]], k_i\in\kappa$ for $0\le i\le K-1$, $\mathbf{H}$ is a diagonal matrix that contains the Fourier coefficients $H[k]$ of the transmitted waveforms, and $\mathbf{F}_N^K$ is the partial-Fourier matrix that contains the K rows of the $N\times N$ Fourier matrix indexed by k. The entries of the L-sparse vector $\mathbf{a}_\nu$ are the values α_l at the indices s_l for the Doppler frequencies ν_l in the "focus zone," i.e., $|\nu-\nu_l|<\pi/P\tau$. The P equations (32) are simultaneously solved using CS-based algorithms, which, during each iteration, the maximal projection of the observation vectors onto the measurement matrix is retained [30].

Some results comparing different configurations of low-rate sampling and processing are shown in Figure 5 [30]. The recovery performance of the classic processing applied to Nyquist samples is presented as a baseline. Sub-Nyquist approaches, performed at 1/10 of the Nyquist rate, include the same classic processing applied to sub-Nyquist samples, a two-stage CS

recovery method that performs delay and Doppler estimation in parallel, separately [30], and Doppler focusing. It is clearly seen that Doppler focusing applied to random Fourier coefficients, which are widely distributed with high probability leading to a wide aperture, outperforms other sub-Nyquist approaches. The use of consecutive coefficients yields small aperture and poor resolution.

The minimal number of samples required for the perfect recovery of $\{\alpha_l, \tau_l, \nu_l\}$ with L targets in a noiseless environment is $4L^2$, with at least $K \geq 2L$ samples per pulse and at least $P \geq 2L$ pulses.

The Xampling approach has several advantages. First, it recovers the targets' parameters directly from the low-rate samples without requiring sampling at the Nyquist rate. Second, previous CS-based methods typically impose constraints on the radar transmitter, which are not needed here. Indeed, as may be seen in (20) and (25), the CS dictionary depends on samples of the waveform $h(t)$, such that the mutual coherence of the dictionary is linearly related to the highest sidelobe value of the AF [58]. In contrast, the CS dictionary of the Xampling method is independent of the waveform, as shown in (32). Third, in the presence of additive white noise, Doppler focusing achieves an increase in SNR by a factor of P (a detailed analysis may be found in [30]). In addition, this approach can operate at the smallest possible sampling rate for recovering the targets' parameters, as derived in [30]. The minimal number of samples required for the perfect recovery of $\{\alpha_l, \tau_l, \nu_l\}$ with L targets in a noiseless environment is $4L^2$, with at least $K \geq 2L$ samples per pulse and at least $P \geq 2L$ pulses. The

Doppler Focusing

Doppler focusing is a processing technique, suggested in [30], which uses target echoes from different pulses to create a superimposed pulse focused at a particular Doppler frequency. This method allows for joint delay-Doppler recovery of all targets present in the illuminated scene. It results in an optimal signal-to-noise ratio (SNR) boost and may be carried out in the frequency domain, thus enabling sub-Nyquist sampling and processing with the same SNR increase as a matched filter.

The output of Doppler processing can be viewed as a discrete equivalent of the following time shift and modulation operation on the received signal:

$$\Phi(t,\nu) = \sum_{p=0}^{P-1} x_p(t+p\tau)\, e^{i\nu p\tau} = \sum_{l=0}^{L-1} \alpha_l h(t-\tau_l) \sum_{p=0}^{P-1} e^{i(\nu-\nu_l)p\tau}. \tag{S11}$$

Consider the sum $g(\nu \mid \nu_l) = \left| \Sigma_{p=0}^{P-1} e^{i(\nu-\nu_l)p\tau} \right|$. For any given ν, targets with Doppler frequencies ν_l in a bandwidth of $2\pi/P\tau$ around ν will achieve a coherent integration and an SNR increase of approximately P. On the other hand, since the sum of P equally spaced points covering the unit circle is generally close to zero, targets with ν_l not "in focus" will roughly cancel out. In summary, we have that

$$g(\nu \mid \nu_l) = \sum_{p=0}^{P-1} e^{i(\nu-\nu_l)p\tau} \approx \begin{cases} P & |\nu-\nu_l| < \pi/P\tau \\ 0 & |\nu-\nu_l| \geq \pi/P\tau, \end{cases} \tag{S12}$$

as shown in Figure S1.

We may therefore estimate the sum of exponents in (S11) as

$$\Phi(t,\nu) \approx P \sum_{l \in \Lambda(\nu)} \alpha_l h(t-\tau_l), \tag{S13}$$

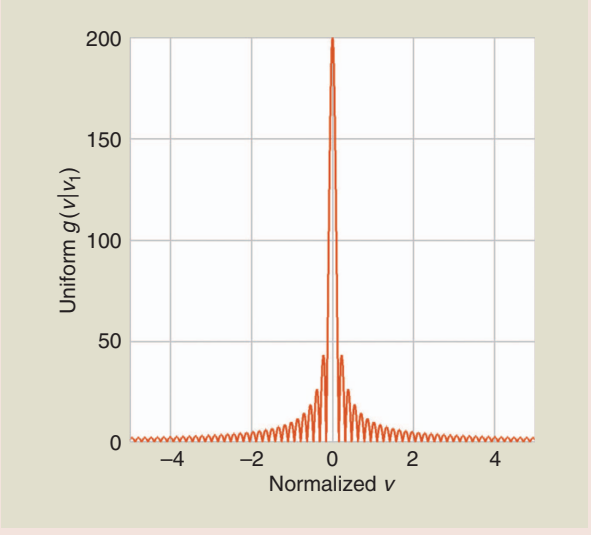

FIGURE S1. The sum of exponents $|g(\nu \mid \nu_l)|$ for $P = 200$, $\tau = 1$s, and $\nu_l = 0$.

where $\Lambda(\nu) = \{l : |\nu - \nu_l| < \pi/P\tau\}$. In other words, the sum is only over the targets whose Doppler shifts are in the interval $|\nu - \nu_l| < \pi/P\tau$.

For each Doppler frequency ν, $\Phi(t,\nu)$ represents a standard pulse-stream model in which the problem is to estimate the unknown delays. Thus, using Doppler focusing, the two-dimensional delay-Doppler recovery problem is reduced to delay-only estimation for a small range of Doppler frequencies, with increased SNR by a factor of P [10]. The Xampling radar of [30] performs Doppler focusing directly on the low-rate samples in the frequency domain, allowing for joint Doppler-delay recovery from the "Xamples."

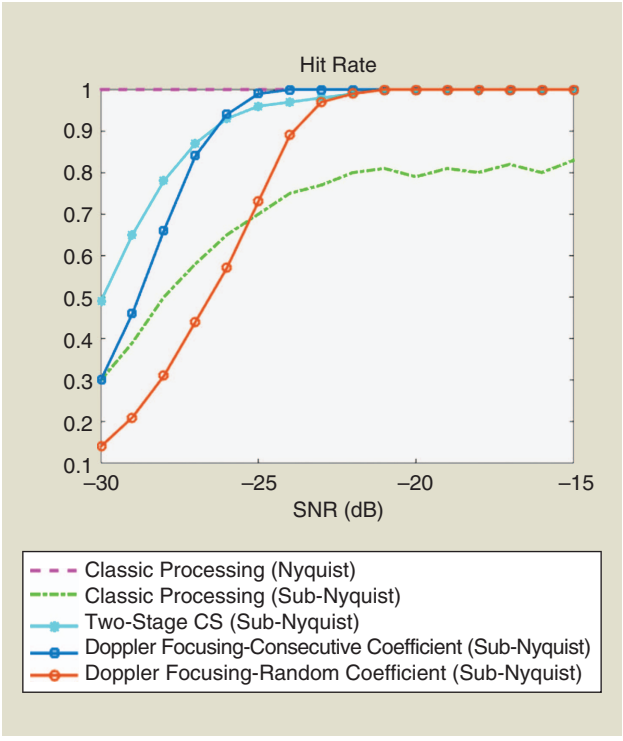

FIGURE 5. The hit rate of classic processing, two-stage CS recovery, and Doppler focusing for a fixed false alarm rate. A *hit* is defined as a delay-Doppler estimate circumscribed by an ellipse around the true target position in the time–frequency plane, with the axes equivalent to ± 3 times the time and frequency Nyquist bins. The two-stage CS recovery separates the delay and Doppler estimation, performing them in parallel [30]. The sub-Nyquist sampling rate was 1/10 of the Nyquist rate [30].

Doppler focusing approach achieves this minimal number of samples. Finally, Doppler focusing is able to deal with certain models of clutter and target dynamic range by adding a simple windowing operation in the sum (27) and by prewhitening in frequency [54].

The Xampling radar was implemented in hardware, as described in the next section, demonstrating real compressed radar capabilities. The hardware prototype is built from off-the-shelf components, which are bandpass filters and low-rate samplers, leading to low hardware complexity.

Hardware prototype

Xampling is used in combination with Doppler focusing in the sub-Nyquist prototype of [30], [40], which demonstrates radar reception at sub-Nyquist rates. The input signal simulates reflections from arbitrary targets and is corrupted by additive noise and clutter. The radar receiver implements the multichannel topology described in the previous section and samples a signal with Nyquist rate of 30 MHz with a compression factor of 30. Hardware experiments demonstrate the feasibility of detecting targets from the low-rate samples of an analog radar signal using standard radio-frequency (RF) hardware [30], [40]. Typical experiment results are shown in Figure 6, which depicts the input signal, the low-rate samples, and the original and recovered delay-Doppler maps, including close targets, both in terms of range and velocity.

At the heart of the receiver lies the Xampling-based ADC, which performs analog prefiltering of the signal before taking pointwise samples. A multiple bandpass-sampling approach

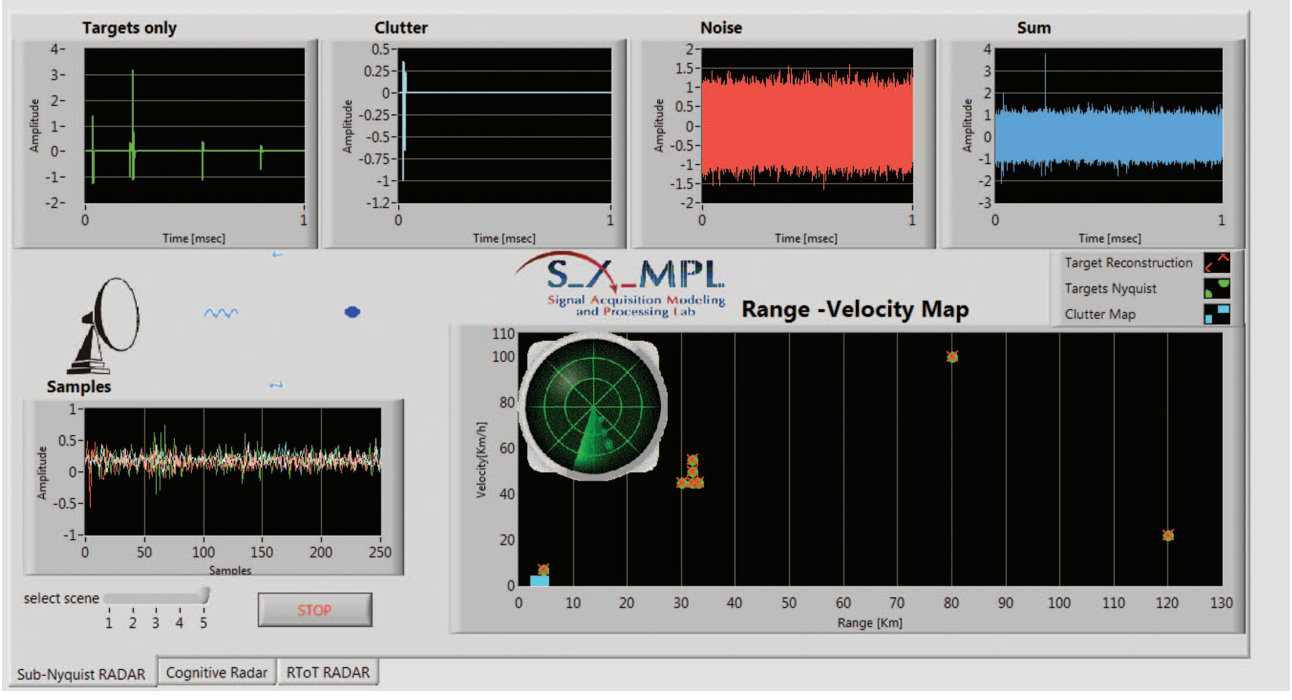

FIGURE 6. The Xampling radar LabView experimental interface. From left to right: at top is the received signal from targets only, then, the received signal from clutter, noise, and overall received signal $x_p(t)$. At the bottom are the sub-Nyquist samples of the four channels at 1/30 of the Nyquist rate, then, the true and recovered delay-Doppler maps. All of the targets (including close targets both in range and in velocity) are correctly detected.

with four channels is adopted. Each channel is composed of a crystal filter with a bandwidth of 80 KHz and extremely narrow transition bands and then sampled at a rate of 250 kHz. The front-end samples four distinct bands of radar–signal spectral content, yielding 320 Fourier coefficients after digital processing with a total sampling rate of 1 MHz. The samples are fed into the chassis controller, and a MATLAB function is launched that computes the 320 Fourier coefficients via FFT, composed of four groups of 80 consecutive Fourier coefficients. These are then used for digital recovery of the delay-Doppler map using the Doppler focusing reconstruction algorithm.

At the heart of the receiver lies the Xampling-based ADC, which performs analog prefiltering of the signal before taking pointwise samples.

The experimental setup is based on National Instrument (NI) PXI-series equipment that is used to synthesize a radar environment and ensure system synchronization. The entire component ensemble wrapped in the NI chassis as well as the analog receiver board are depicted in Figure 7. Additional information regarding the system's configuration and synchronization can be found in [40].

To demonstrate target detection from low-rate samples, the Applied Wave Research (AWR) software is used to simulate the radar scenario, including pulse transmission and accurate power loss due to wave propagation in a realistic medium. AWR software provides a computer-based environment for designing hardware for use with wireless and high-speed digital products. It is used for RF, microwave, and high-frequency analog circuits and system design. A large variety of scenarios, consisting of different targets' parameters, i.e., delays, Doppler frequencies, and amplitudes, are examined in [30] and [40]. An arbitrary waveform generator module produces an analog signal that is amplified and routed to the radar receiver board. The received radar waveform is contaminated with noise and clutter, showing the capabilities of the Xampling receiver to deal with these [30], [40], [54]. The Nyquist rate of the signal is 30 MHz, so that sampling at 1 MHz corresponds to a fast time compression factor of 30.

Slow time compression

Most works on CS radar focus on compression in the fast time domain, reducing the number of samples per pulse below the Nyquist rate. As we have seen, using appropriate CS techniques allows for preserving the range resolution while operating in low-rate regimes by breaking the link between bandwidth and sampling rate. This is illustrated in Figure 5, in which Doppler focusing is shown to achieve the same hit rate as classic processing above a certain SNR, and in Figure 6, where close targets are seen to be correctly recovered despite sampling at 3.3% of the Nyquist rate. We will now see that compression may be similarly performed in the slow time domain, as demonstrated in [65], where the number of transmitted pulses is reduced without decreasing Doppler resolution.

Nonuniform pulse Doppler

The resolution in Doppler frequency in standard processing is governed by the number of transmitted pulses P. More precisely, it is equal to $2\pi/P\tau$. However, a large P leads to large CPI and long time on target. Slow time compression breaks the relation between CPI and time on target. To that end, $M < P$ pulses are sent nonuniformly over the entire CPI $P\tau$, implementing nonuniform time steps between the pulses [65]. This way, the same CPI is kept, but a smaller number of pulses is transmitted, thereby reducing power consumption. In addition, the periods of time in which no pulse is transmitted in a certain direction can be exploited to send pulses in other directions. This allows the radar to scan several directions at the same time and obtain the corresponding delay-Doppler maps in a single CPI. However, at the same time, this reduces SNR because fewer pulses are transmitted in every direction.

Consider a nonuniform pulse-Doppler radar such that the pth pulse is sent at time $m_p\tau$, where $\{m_p\}_{p=0}^{M-1}$ is an ordered set of integers satisfying $m_p \geq p$. In this case, (1) becomes

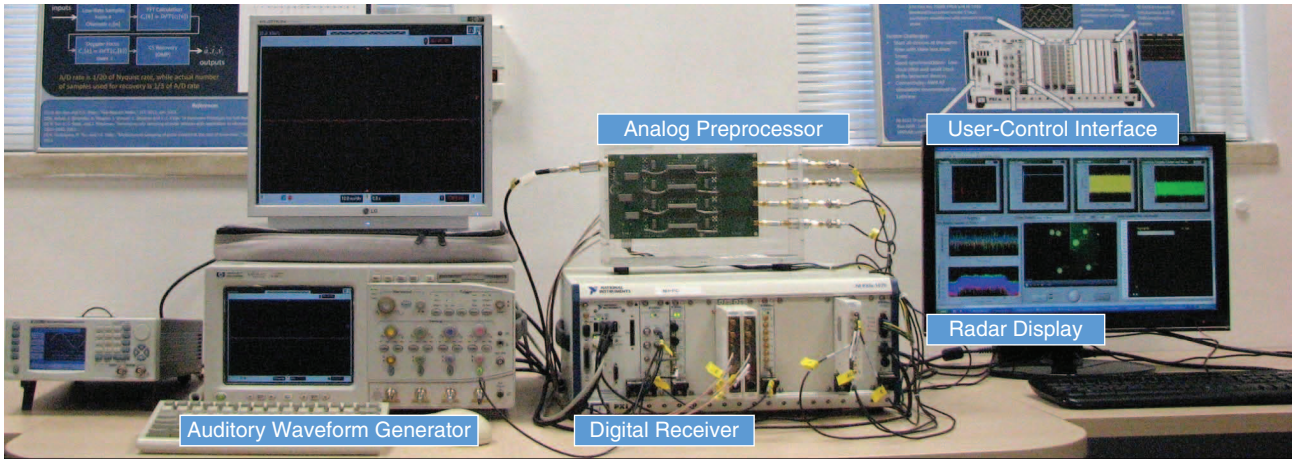

FIGURE 7. The Xampling radar prototype including an arbitrary waveform generator, receiver board, NI chassis, and display [40].

$$x_T(t) = \sum_{p=0}^{M-1} h(t - m_p\tau), \quad 0 \le t \le P\tau, \tag{33}$$

and the received frames (4) are written as

$$x_p(t) = \sum_{l=0}^{L-1} \alpha_l h(t - \tau_l - m_p\tau) e^{-j\nu_l m_p \tau}, \quad 0 \le t \le P\tau. \tag{34}$$

The same Xampling-based method in [30] is used to obtain the Fourier coefficients $c_p[k]$ of the received pulses. Suppose we limit ourselves to the Nyquist grid, as previously mentioned, so that $\tau_l/\tau = s_l/N$, where s_l is an integer satisfying $0 \le s_l \le N-1$, and $\nu_l\tau = 2\pi r_l/M$, where r_l is an integer in the range $0 \le r_l \le M-1$. Similar to the derivations in the previous section, we can write the Fourier coefficients $c_p[k]$in matrix form [65] as

$$\mathbf{X} = \mathbf{H}\mathbf{F}_N^K \mathbf{A} (\mathbf{F}_P^M)^T, \tag{35}$$

where $\mathbf{H}$ is a diagonal matrix that contains the Fourier coefficients $H[k]$. The partial-Fourier matrix $\mathbf{F}_M^P$ contains M rows from the $P \times P$ Fourier matrix, indexed by the values of the transmitted pulses m_p, $1 \le p \le M$; when sampling at the Nyquist rate, $K = N$ and $\mathbf{F}_N^K$ become the standard $N \times N$ Fourier matrix. Similarly, when considering uniformly spaced pulses, $M = P$ and $\mathbf{F}_P^M$ are the standard $P \times P$ matrix. The goal is to recover the sparse matrix $\mathbf{A}$ that contains the values α_l at the L indices $\{s_l, r_l\}$ from the Fourier coefficients matrix $\mathbf{X}$.

CS matrix-recovery algorithms are directly applicable to (35) by extending CS techniques presented in vector form, such as orthogonal matching pursuit or the fast iterative-shrinkage-thresholding algorithm [10], [12] to matrix settings [66]. Alternatively, instead of solving the matrix problem of (35), we can apply the Doppler focusing operation [30] described in "Doppler Focusing." As illustrated in Figure 8, the approximation from (S12) may still be applied in the nonuniform case, where $m_p \ge p$. Therefore, we can rewrite the Fourier coefficients from (27) by replacing p with m_p for the nonuniform case. These may then be approximately expressed in vector form as in (32) and recovered as previously described. It is shown in [65] that the minimal number of nonuniform pulses required to recover the Doppler frequencies of L targets is identical to the uniform case, that is, two L pulses.

Hardware simulation

The transmission of nonuniform pulses has been implemented in the Xampling prototype [40]. Recall that the received signal has a bandwidth of 30 MHz and is sampled at the rate of 1 MHz. To this fast time compression, we now add compression in the slow time domain. In the hardware simulation, $P = 50$ pulses over a CPI of $MP\tau = 2.5$ s are considered. Half of the pulses, i.e., $M = 25$, chosen at random, are sent in one direction, while the other half are sent in a second direction. Two delay-Doppler maps are then simultaneously recovered during a single CPI, as shown in Figure 9. Both of the maps are fully recovered, as previously mentioned, from compressed samples in both the fast and slow time domains.

Range-velocity ambiguity resolution

As presented thus far, targets are traditionally assumed to lie in the radar-unambiguous range-velocity region. For a given PRI τ, the maximum unambiguous range is $r_{\max} = c\tau/2$, and the maximum unambiguous velocity is $\dot{r}_{\max} = \lambda/(4\tau)$, where λ is the radar wavelength. When the target range and velocity intervals of interest are large, traditional pulse-Doppler radar

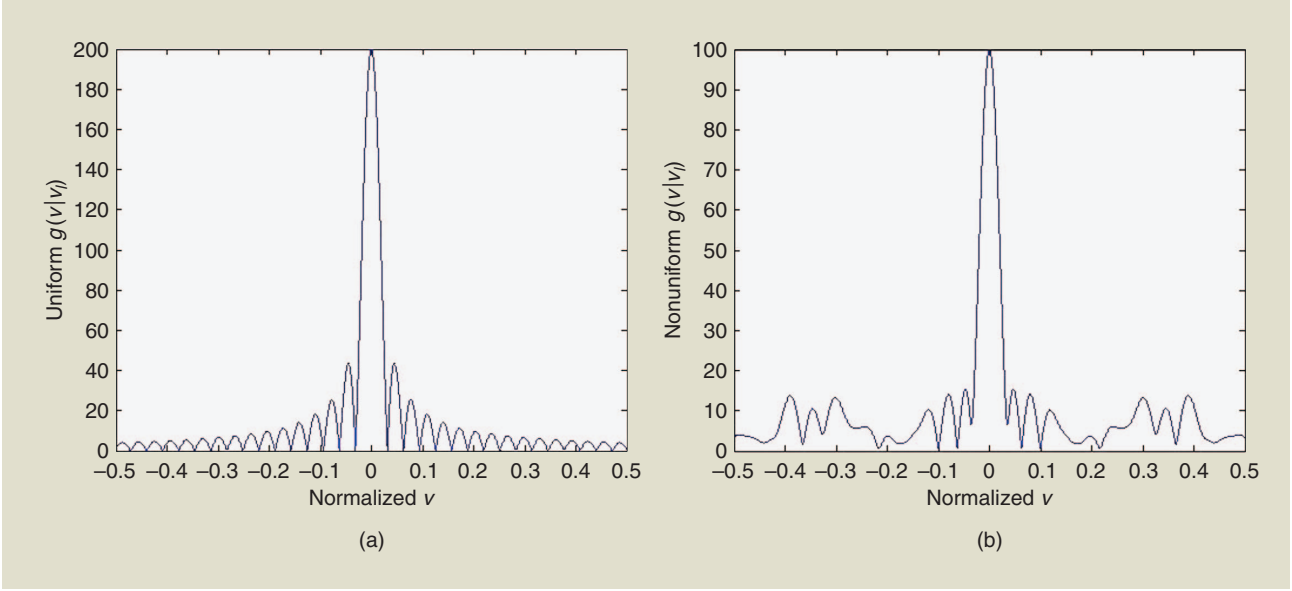

FIGURE 8. The sum of exponents $|g(\nu \mid \nu_l)|$ for $M = 200$, $\tau = 1$s, and $\nu_l = 0$ in the (a) uniform and (b) nonuniform cases. In the nonuniform case, $P = 100$ pulses are chosen uniformly at random [65].

systems suffer from the so-called Doppler dilemma [67], a tradeoff between range and velocity ambiguity, whose product is limited to $r_{\max}\dot{r}_{\max} = c\lambda/8$.

Several techniques have been proposed over the years to mitigate the range-velocity ambiguity by increasing either of these parameters. Two main PRF variation-based methods are staggered PRFs and multiple PRFs (MPRFs). Staggered PRFs are used to raise the first blind speed $\dot{r}_{\max}$ significantly without degrading the unambiguous range [7]. Pulse-to-pulse stagger varies the PRF from one pulse to the next, achieving increased Doppler coverage [68]. The main disadvantage of this approach is that the data correspond to a nonuniformly sampled sequence, making it more difficult to apply coherent Doppler filtering [7]. In addition, clutter cancellation becomes more challenging, and the sensitivity to noise increases [69]; therefore, MPRF techniques are typically preferred. We now review some of the MPRF-based methods and then present a Xampling approach that solves the delay-Doppler ambiguity using phased-coded-transmit pulses.

Ambiguity resolution is typically achieved by searching for coincidence between either unfolded Doppler or delay estimates for each PRF.

MPRF

The MPRF approach transmits several pulse trains, each with a different PRF. Ambiguity resolution is typically achieved by searching for coincidence between either unfolded Doppler or delay estimates for each PRF. A popular approach, adopted in [70], relies on the Chinese remainder theorem [5] and uses two PRFs, such that the numerator and denominator of the ratio between them are prime numbers. The ambiguous velocities are computed for each train i as

$$\hat{\dot{r}}_{i,k} = \hat{\dot{r}}_{i,0} + k\frac{\lambda}{2\tau}, \quad k \in \mathbb{Z}, \tag{36}$$

where $\hat{\dot{r}}_{i,0}$ is the velocity estimate within the unambiguous velocity interval $(-\dot{r}_{\max}, \dot{r}_{\max}]$. Congruence between these are found by an exhaustive search, so that all $\hat{\dot{r}}_{i,k}$ fall within a small, interval, or correlation bin. The resulting velocity estimate is computed by averaging overall $\hat{\dot{r}}_{i,k}$. Assuming $T = 2$ pulse trains with PRFs and ratio $\tau_1/\tau_2 = m/n$, where m and n are relative prime numbers, the expanded velocity interval is of size $m\lambda/2\tau_1 = n\lambda/2\tau_2$. However, in this approach, a small range error on a single PRF can cause a large error in the resolved range with no indication that this has happened [71].

A clustering algorithm proposed in [71] implements the search for a matching interval by computing average distances to cluster centers. The average squared error is defined as

$$C(k) = \sum_{i=1}^{T} |\hat{r}_{i,k} - \bar{r}_k|^2, \quad k = 0, \ldots, r_{\text{amb}}/r_{\max}, \tag{37}$$

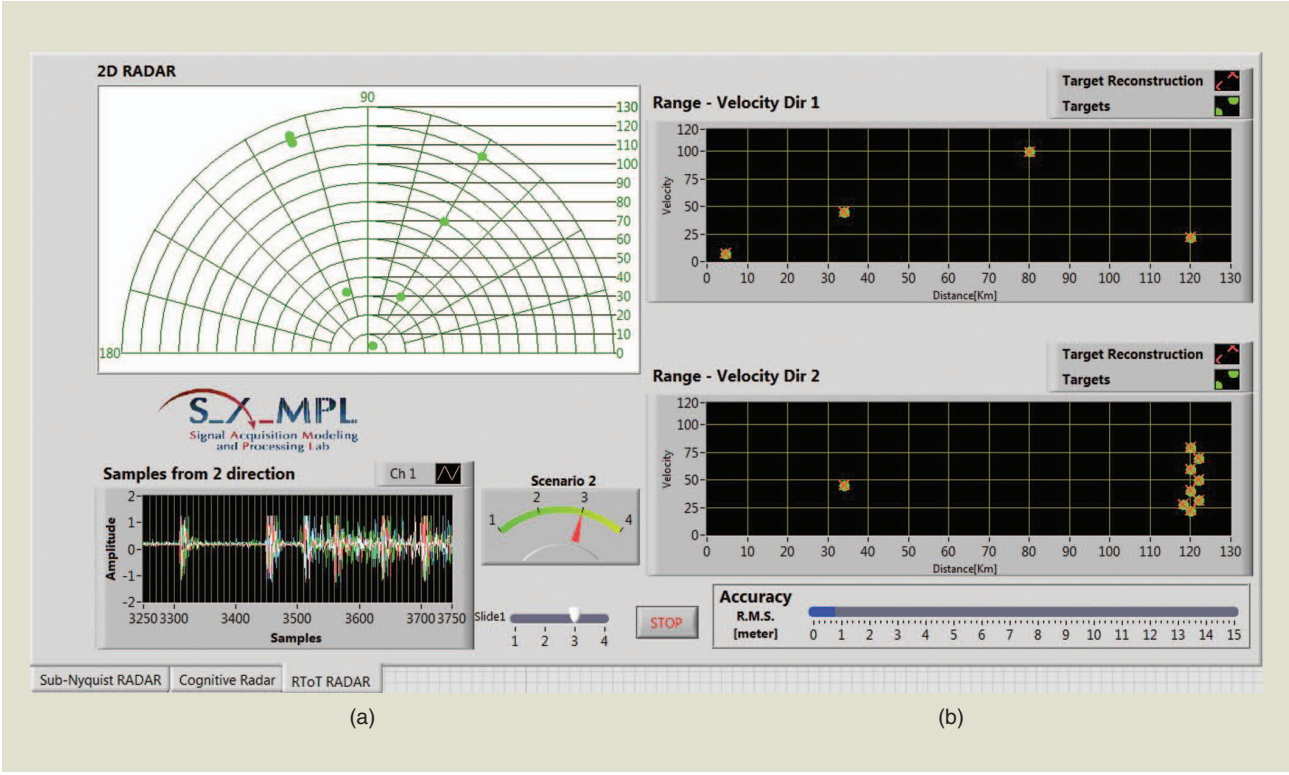

FIGURE 9. The experimental interface of the Xampling radar, with both fast and slow time compression. (a) The true targets' ranges in two directions (top) and superposed low-rate samples from both directions (bottom). (b) The range-velocity map of true and recovered targets in both directions [65].

where $\bar{r}_k$ is the median value of the T ranges with index k and r_{amb} is the maximal ambiguous range. The best cluster occurs at the value of k, where $C(k)$ is minimized. This happens when all of the ambiguous ranges are unfolded correctly, and, hence, all of the range estimates have nearly the same range. This technique still requires an exhaustive search over clusters and does not process the samples jointly, thereby decreasing the SNR.

Random PC has been adopted in polarimetric weather radars, which exploit the inherent random phase between pulses of the popular magnetron transmitters.

Phased-coded pulses

A random-pulse, phase-coding (PC) approach is adopted in [72] to increase the range-unambiguous region, while preserving that of the Doppler frequency and using a single PRF. A similar technique may be used to increase the Doppler-frequency-unambiguous region. Random PC has been adopted in polarimetric weather radars, which exploit the inherent random phase between pulses of the popular magnetron transmitters. In this context, PC mitigates out-of-trip echoes [73]. The approach of [72] introduces a random phase, which differs from pulse to pulse. The joint processing of received signals from all of the trains is the key to range ambiguity resolution.

The pulse-Doppler radar transceiver sequentially transmits one modulated pulse train, consisting of P equally spaced pulses. For $0 \le t \le P\tau$, the transmitted signal is given by

$$x_T(t) = \sum_{p=0}^{P-1} h(t - p\tau) e^{jc[p]}, \tag{38}$$

where $c[p]$ is uniformly distributed in the interval $[0, 2\pi)$ and represents the phase shift of the pth pulse.

As opposed to the common assumption in traditional radar, the targets' time delays $\tilde{\tau}_l$ are not assumed to lie in the unambiguous time region, i.e., less than the PRI τ, but rather in the ambiguous range $\tilde{\tau}_l \in [0, Q\tau)$, where $Q < P$ is the ambiguous factor defined by the targets' maximal range. For convenience, the delay $\tilde{\tau}_l$ is decomposed into its integer part (the ambiguity order) $q_l\tau$ and fractional part (the folded or reduced delay) τ_l such that

$$\tilde{\tau}_l = \tau_l + q_l\tau, \tag{39}$$

where $0 \le q_l \le Q - 1$ is an integer and $0 \le \tau_l < \tau$.

The received signal is then

$$x_R(t) = \sum_{p=0}^{P-1} \sum_{l=0}^{L-1} \alpha_l h(t - \tilde{\tau}_l - p\tau) e^{-j2\pi\nu_l(p+q_l)\tau} e^{jc[p]}, \tag{40}$$

for $0 \le t < (P + Q)\tau$. The main difference with traditional pulse-Doppler radar, aside from the coded phase, is that the PRI index in the Doppler shift term is $p + q_l$, rather than the pulse index p.

The Fourier series of the received signal (40) can be written in matrix form, similarly to (35), and recovered using matrix CS recovery techniques [72]. The minimal number of samples per pulse allowing recovery of $\mathbf{X}$ with high probability is found to be $K > 2L$, and the minimal number of pulses P is $2L + Q + 2$. This method resolves a maximum unambiguous range $r_{\max} = cQ\tau/2$, while preserving the maximum unambiguous velocity $\dot{r}_{\max} = \lambda/(4\tau)$, thereby increasing their product $r_{\max}\dot{r}_{\max}$ by a factor of Q, under the prior mentioned conditions on the number of samples and pulses.

This approach has three main advantages. First, it improves the delay estimation with respect to the MPRF methods since it preserves the resolution of traditional pulse-Doppler radar, i.e., $1/B_h$, while increasing the unambiguous delay region to $Q\tau$. Second, it increases the SNR by jointly processing the samples from all of the pulse trains, rather than matching the estimated parameters from each pulse processed separately. Finally, it provides a systematic delay-Doppler-recovery method that does not involve an exhaustive search. From a practical point of view, this approach does not require the use of several pulse trains with different PRFs, thus simplifying hardware implementation.

Cognitive radar and spectrum sharing

Recently, the concept of cognitive radar (CR) [74], inspired by the echo-location system of a bat, has been presented as a natural next step for traditional radar. The cognition property requires adaptive transmission and reception capabilities, i.e., both the transmitter and receiver are able to dynamically adjust to the environment conditions. Many interpretations of this idea have been proposed. We focus on one aspect of cognition, the dynamic and flexible adaptation to the spectral environment, which allows for spectrum sharing between communication and radar systems [36]–[38]. The interest in these spectrum sharing radars is largely due to electromagnetic spectrums being a scarce resource, with most services having a need for a greater access to it.

The spectrum sharing solution proposed in [39] capitalizes on the cognitive abilities of the radar system. It is shown how compressed radars may be adapted to allow for spectral coexistence between communication and radar signals and flexibility of the radar transmission. This demonstrates that, beyond increasing resolution and realizing compression in the time, frequency, and spatial domains, compressed radars have the potential to enable otherwise challenging technologies.

Spectral adaptive transmission

In previous works that implement fast time compression, e.g., Xampling radar [30], [40], the transmitter broadcasts a wideband signal, which reflects off the targets and propagates back to the receiver. The received signal is then filtered before sampling so that only the content of a few narrow bands is sampled and processed. These works only deal with the reception

side of the radar, providing sampling and processing techniques that can be used with any traditional radar transmitter. However, for broadband frequency occupation and power savings, only the narrow frequency bands that are to be sampled may be transmitted [31], [39]. This will not affect any aspect of the processing since the received signal is preserved in the bands of interest. In fact, since all the signal power is concentrated in the processed bands, the SNR increases, and the detection performance improves [75].

Let $\tilde{H}(f)$ be the CTFT of the new transmitted radar pulse,

$$\tilde{H}(f) = \begin{cases} H(f) & f \in \left[f_r^i - B_r^i/2,\, f_r^i + B_r^i/2 \right] \text{ for } 1 \le i \le N_b \\ 0 & \text{else,} \end{cases} \tag{41}$$

where N_b is the number of filtered bands and B_r^i and f_r^i are the bandwidth and center frequency of the ith band, respectively. Obviously, the computation of the relevant Fourier coefficients $c_p[k]$ from (25) will not change. Therefore, the recovery methods presented in the "Fast Time Xampling" section are applicable here as well.

The concept of transmitting only a few subbands that the receiver processes is one way to formulate a frequency-agile CR in terms of its ability to adapt to spectral demands. Complying with CR requirements, the support of the subbands varies with time to allow for dynamic and flexible adaptation. Such a system also enables the radar to disguise the transmitted signal as an electronic countermeasure or to cope with crowded spectrums by using a smaller, interference-free portion, as further discussed in the following section.

Application to spectrum sharing

The unhindered operation of a radar that shares its spectrum with communication systems has captured a great deal of attention within the operational radar community in recent years [36]–[39]. Recent research programs in spectrum sharing radars include the Enhancing Access to the Radio Spectrum project by the National Science Foundation [38] and the Shared Spectrum Access for Radar and Communication (SSPARC) program [37], [76], initiated by the Defense Advanced Research Projects Agency.

A variety of system architectures have been proposed for spectrum sharing radars, and most place an emphasis on optimizing the performance of either radar [77] or communication [78] while ignoring the performance of the other. In nearly all cases, the real-time exchange of information between radar and communication hardware has not yet been integrated into the system architectures. Exceptions to this are automotive solutions in which the same waveform is used for both target detection and communication [79].

In a similar vein, the sub-Nyquist, CR-based approach from [39] incorporates the handshaking of spectral information between the two systems. The CR configuration is key to spectrum sharing since the radar transceiver can adapt its transmission to available bands, thus achieving coexistence with communication signals. Suppose the set of all frequencies of the available common system spectrum is given by $\mathcal{F}$.The communication and radar systems occupy the subsets $\mathcal{F}_C$ and $\mathcal{F}_R$ of $\mathcal{F}$, respectively. The goal is to design the radar waveform and its support $\mathcal{F}_R$, conditional on the fact that the communication occupies frequencies $\mathcal{F}_c$, which are unknown to the radar transceiver [39]. To detect the bands left vacant by the communication signals, spectrum sensing needs to be performed over a large bandwidth. Such a task has recently received tremendous interest in the communication community, which faces a bottleneck in terms of spectrum availability. To increase the efficiency of spectrum managment, a dynamic opportunistic exploitation of temporarily vacant spectral bands by secondary users has been considered, called *cognitive radio* (*CRo*) [80], [81].

To increase the efficiency of spectrum management, a dynamic opportunistic exploitation of temporarily vacant spectral bands by secondary users has been considered, called *cognitive radio*.

A spectrum sharing paradigm using Xampling techniques, the spectral coexistence via Xampling (SpeCX) system [39] is composed of a sub-Nyquist CRo receiver [81] to detect the occupied communication bands so that the radar transmitter may subsequently exploit the spectral holes. In this setting, the received signal at the communication receiver is given by

$$x(t) = x_C(t) + x_R(t), \tag{42}$$

where $x_R(t) = r_{Tx}(t) + r_{Rx}(t)$ is the radar signal sensed by the communication receiver, composed of the transmitted and received radar signals. The goal is therefore to recover the support of $x_C(t)$, given the known support of $x_R(t)$, which is shared by the radar transmitter with the communication receiver. This can be formulated as a sparse recovery with partial-support knowledge, studied under the framework of a modified CS [82].

Once $\mathcal{F}_C$ is identified, the communication receiver provides a spectral map of occupied bands to the radar. Equipped with the detected spectral map and known radio environmental map, the objective of the radar is to identify an appropriate transmit-frequency set $\mathcal{F}_R \subset \mathcal{F} \setminus \mathcal{F}_C$ such that the radar's probability of detection P_d is maximized. For a fixed probability of false alarm P_{fa}, the P_d increases with a higher signal-to-interference plus-noise ratio (SINR) [83]. Hence, the frequency selection process can, alternatively, choose to maximize the SINR or minimize the spectral power in the undesired parts of the spectrum. To find available bands with the least amount of interference, a structured sparsity framework [84] is adopted in [39]. Additional requirements of transmit-power constraints, range-sidelobe levels, and minimum separation between the bands can also be imposed. At the receiver of this spectrum-sharing radar, the sub-Nyquist processing method of [30] recovers the delay-Doppler map from the subset of Fourier coefficients defined by $\mathcal{F}_R$.

This CR system leads to three main advantages. First, the CS reconstruction, as presented in [30] on the transmitted fragmented bands, achieves the same resolution as traditional Nyquist processing over a significantly smaller bandwidth. Second, by concentrating all of the available power in the transmitted narrow bands rather than over a wide bandwidth, the CR increases the SNR. Finally, this technique allows for a dynamic form of the transmitted signal spectrum, in which only a small portion of the whole bandwidth is used at each transmission, thereby enabling spectrum sharing with communication signals, as illustrated in Figure 10(d). There, the coexistence between radar-transmitted bands in red and existing communication bands in white is shown.

SpeCX prototype

The SpeCX prototype, presented in Figure 10, demonstrates radar and communication spectrum sharing. It is composed of a CRo receiver and a CR transceiver. At the heart of the CRo system lies the proprietary modulated wideband converter board [29] that implements a sub-Nyquist analog front-end receiver, which processes signals with Nyquist rates up to 6 GHz. The card first splits the wideband signal into $M = 4$ hardware channels with an expansion factor of $q = 5$, yielding $M_q = 20$ virtual channels after digital expansion (see [85]). In each channel, the signal is mixed with a periodic sequence $p_i(t)$, which are truncated versions of Gold codes [86], generated on a dedicated field-programmable gate array, with a periodic frequency $f_p = 20$ MHz.

Next, the modulated signal passes through an analog antialiasing LPF. Finally, the low-rate analog signal is sampled by an NI ADC operating at $f_s = (q + 1)f_p = 120$ MHz (with intended oversampling), leading to a total sampling rate of 480 MHz. The digital receiver is implemented on an NI PXIe-1065 computer with a dc-coupled ADC. Since the digital processing is performed at the low rate of 120 MHz, very low computational load is required to achieve real-time recovery. The prototype is fed with RF signals composed of up to $N_{\text{sig}} = 5$ real communication transmissions, i.e., ten spectral bands with a total bandwidth occupancy of up to 200 MHz and varying support, with a Nyquist rate of 6 GHz.

The input transmissions then go through an RF combiner, resulting in a dynamic multiband input signal that enables fast carrier switching for each of the bands. This input is specially designed to allow the testing of the system's ability to rapidly sense the input spectrum and adapt to changes, as required by modern CRo and shared spectrum standards, e.g., in the SSPARC program. The system's effective sampling rate, equal to 480 MHz, is only 8% of the Nyquist rate. Support recovery is digitally performed on the low-rate samples. The prototype successfully recovers the support of the communication transmitted bands, as demonstrated in Figure 10(b) and (c). Once the support is recovered, the signal itself can be reconstructed from the sub-Nyquist samples. This step is performed in real time, reconstructing the signal bands one sample at a time.

The CR receiver system is identical to the sub-Nyquist sampling prototype of [30], [31] and [40]. In the cognitive case, the transmitter only transmits over $N_b = 4$ bands, which constitute 3.2% of the original wideband signal bandwidth after the spectrum-sensing process has been completed by the communication receiver. Figure 10(d) illustrates the coexistence between the radar-transmitted bands in red and the existing communication bands in white. The gain in power is demonstrated in Figure 10(e); the wideband radar spectrum is shown in blue, the CR in red, and the noise in yellow on a logarithmic scale. The true and recovered range-velocity maps are presented in Figure 10(f). All of the $L = 10$ targets are perfectly recovered, and the clutter, depicted in blue, is discarded. Below the map, the range-recovery accuracy is shown for three scenarios: from

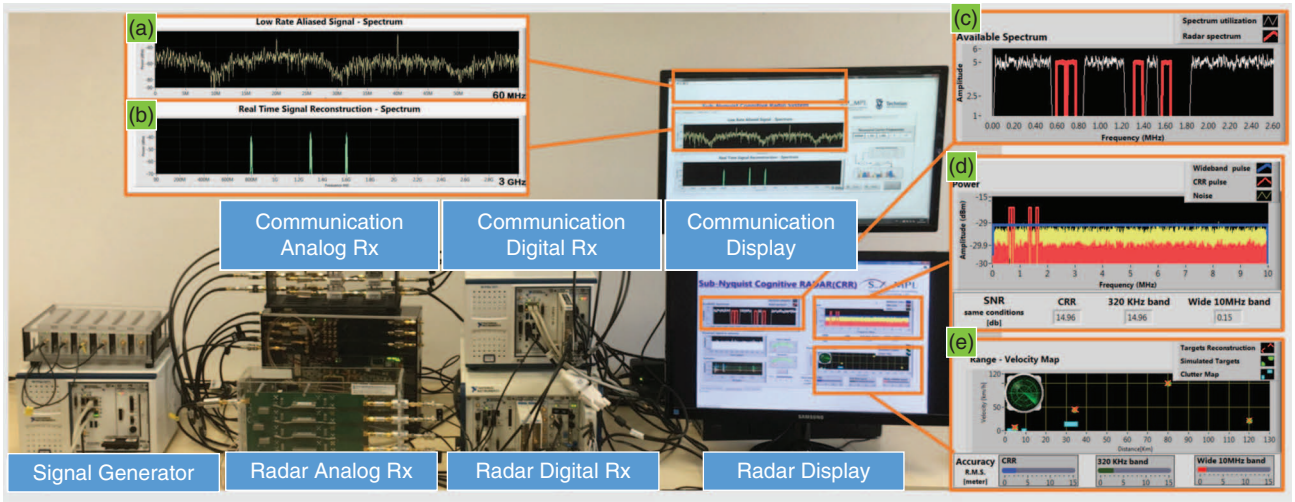

FIGURE 10. The SpeCX prototype. The system consists of a signal generator, a CRo communication analog receiver, including the modulated wideband converter (MWC) analog front-end board, a communication digital receiver, a CR analog, and a receiver. The SpeCX communication system display shows (a) low-rate samples acquired from one MWC channel at a rate of 120 MHz and (b) a digital reconstruction of the entire spectrum from sub-Nyquist samples. The SpeCX radar display shows (c) the coexisting communication and CR, (d) the CR spectrum compared with the full-band radar, and (e) the range-velocity display of both the detected and true locations of the targets [39].

left to right, the CR in blue (2.5 m), the four adjacent bands with the same bandwidth (12.5 m), and the wideband (4 m). The second configuration selects four adjacent frequency bands with the same bandwidth as the CR (with nonadjacent bands) for transmission. Its poor resolution stems from its small aperture. The CR system with nonadjacent bands yields better resolution than traditional wideband transmission, sampling, and processing at the Nyquist rate, due to the increased SNR.

Compressed MIMO radar

Compressed radar methods have recently been extended to MIMO settings, in which their impact may be even greater than for single-antenna configurations. MIMO radar systems belong to the family of array radars, which allow for the simultaneous recovery of the targets' ranges, Dopplers, and azimuths. This 3-D recovery results in high digital processing complexity. One of the main challenges of MIMO radar is therefore coping with complicated systems in terms of cost, high computational load, and hardware implementation. CS has been naturally applied to MIMO to reduce the processing complexity on the digital side as well as allow for spatial compression, in addition to the time compression achieved in single-antenna systems. In MIMO radars, the array aperture, which depends on the number of antennas, dictates the azimuth resolution. Since the aperture is determined by the number of antennas in traditional virtual ULAs, high-azimuth resolution requires a large number of antennas.

Increased resolution

As in single-antenna radar systems, CS has first been exploited to increase the parameter resolution. Here, the MIMO array is composed of T transmitters and R receivers so as to achieve the desired aperture $Z = TR/2$, as shown in Figure 2. The transmitted signal at the mth transmit antenna is given by (12), and each receiver samples the received signal at the Nyquist rate, as in a traditional MIMO. Assuming a sparse target scene, in which the ranges, Dopplers, and azimuths lie on a predefined grid, the work of [15] is extended to MIMO architectures in [18] and [19]. The transmit and receive array manifolds respectively, are given by

$$\mathbf{a}_T(\theta) = \left[e^{j2\pi\xi_1\theta}, e^{j2\pi\xi_2\theta}, \ldots, e^{j2\pi\xi_T\theta}\right]^T, \tag{43}$$

and

$$\mathbf{a}_R(\theta) = \left[e^{j2\pi\zeta_1\theta}, e^{j2\pi\zeta_2\theta}, \ldots, e^{j2\pi\zeta_R\theta}\right]^T, \tag{44}$$

where ξ_m and ζ_q are the relative mth transmit and qth receive antenna spacings. The $R \times N$ received signal matrix from a unit strength target at direction θ, with delay τ and Doppler ν is defined as

$$\mathbf{Z} = \mathbf{a}_R(\theta)\mathbf{a}_T^T(\theta)\mathbf{S}^T(\tau, \nu). \tag{45}$$

Here, $\mathbf{S}^T(\tau,\nu)_{i,m} = s_m(t_i - \tau)e^{j2\pi\nu t_i}$, where t_i are the sampling times and m indexes the transmitted waveforms. In this case, the columns of the dictionary $\mathbf{A}$ are given by $\mathrm{vec}(\mathbf{Z})$ for all possible combinations of θ, ν, and τ on a predefined grid.

The targets' parameters are recovered by matching the received signal with dictionary atoms. To achieve measurement diversity, random waveforms may be used, while the antenna locations are deterministic. ULAs are considered in [18] for both the transmit and receive arrays that do not benefit from the virtual array configuration. Alternatively, deterministic waveforms can be used, e.g., Kerdock codes [19], while the antenna locations are selected uniformly at random over the aperture $Z = TR/2$. Bounds on N, with respect to the number of antennas T and R and the number of samples that ensure targets' parameters recovery, are provided in [18] and [19].

A similar approach extends the framework of [15] to the MIMO setting by adding an azimuth matrix to the time-shift and frequency-modulation matrices $\mathbf{T}$ and $\mathbf{M}$, respectively, as defined in (14). In this case, each target lying on the grid is represented by a time shift, a frequency modulation, and an angle $\mathbf{A}_{q,m} = e^{j\theta(\xi_m + \zeta_q)}$ [87].

In both works, assuming N grid points in each dimension, the number of columns of $\mathbf{A}$ is N^3. The processing efficiency is thus penalized by a very large dictionary that contains every parameter combination. Note that the previously mentioned works focus on increased parameter resolution and do not deal with reduced time/spatial sampling and processing rates.

Reduced processing

Fast time compression is performed in [23]–[25], in which the Nyquist rate samples are compressed in each antenna before being forwarded to the central unit. A circular array is adopted in [23], with transmit and receive nodes uniformly distributed on a disk with a small radius. At each receive antenna, linear projections of the measurement vector are retained so that the resulting samples are compressed in both the slow and fast time domains. Both individual reconstruction at each receiver and joint processing at a fusion center are proposed, using CS recovery methods. The actual sampling is still performed at the Nyquist rate.

The MIMO matrix completion (MIMO-MC) radar [24], [25] employs MC techniques to avoid parameter discretization, which is typically used in CS methods. Two configurations are proposed for azimuth-Doppler recovery in a range bin of interest. In the first scenario, each receiver performs an MF and forwards the maximum of each MF output to the fusion center. The samples from the pth pulse transmitted to the fusion center can then be written in matrix form as

$$\mathbf{X}_p = \mathbf{A}_R \boldsymbol{\Sigma} \mathbf{D}_p \mathbf{A}_T^T, \tag{46}$$

where $\mathbf{X}_p$ is the $R \times T$ matrix that contains the maximum of the MF output for each transmitter and each receiver. For ULA configurations, the lth column of the $T \times L$ transmitter-steering matrix $\mathbf{A}_T$ is given by $(\mathbf{A}_T)_l = \left[1, e^{j\frac{2\pi}{\lambda}d_T\sin(\theta_l)}, \ldots, e^{j\frac{2\pi}{\lambda}(T-1)d_T\sin(\theta_l)}\right]^T$, where d_T is the interelement spacing. The steering matrix $\mathbf{A}_R$ at the receiver is similarly defined. The diagonal matrix $\boldsymbol{\Sigma}$ contains the targets' RCS α_l, and the diagonal matrix $\mathbf{D}$ contains the targets' Dopplers such that

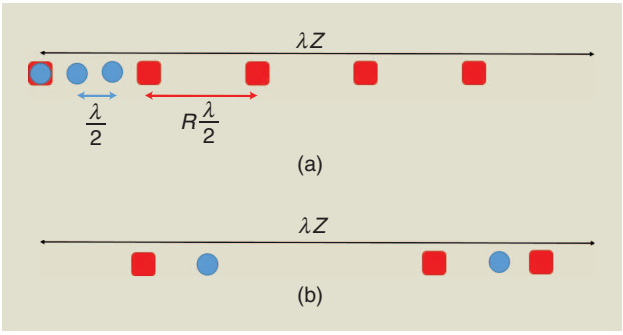

FIGURE 11. An illustration of MIMO arrays: (a) the standard array and (b) the random thinned array [32].

$\mathbf{D}_{p(l,l)} = e^{j\frac{2\pi}{\lambda}2\nu_l(p-1)\tau}$. In this scheme, each receiver transmits the output of a few randomly chosen MFs to the fusion center so that $\mathbf{X}_p$ is only partially known.

In the second scenario, the receivers forward Nyquist samples to the fusion center without performing the MF. In this case, the samples are written as

$$\mathbf{X}_p = \mathbf{A}_R \Sigma \mathbf{D}_p \mathbf{A}_R^T \mathbf{S}, \tag{47}$$

where $\mathbf{S}$ is the $T \times N$ matrix that contains the Nyquist rate samples of each transmitted waveform $s_m(t)$. In this scheme, each receive antenna randomly acquires a subset of the Nyquist samples and transmits these to the fusion center. In both cases, the fusion center performs MC before parametric estimation methods are applied to extract θ_l and ν_l, such as multiple signal classification, also known as *MUSIC* [88]. In these works, sampling and processing rate reduction are not addressed since compression is performed in the digital domain after sampling, and the missing samples are reconstructed before recovering the targets' parameters. Instead, these approaches are aimed at reducing the communication overhead between the receivers and the fusion center.

Spatial compression

Several recent works have considered applying CS to MIMO radar to reduce the number of antennas or the number of samples per receiver without degrading resolution. The problem of azimuth recovery of targets all in the same range-Doppler bin is investigated in [28]. Spatial compressive sampling is performed, in which the number of antennas is reduced while preserving the azimuth resolution. The classic MIMO virtual array configuration requires receivers with maximum spacing $\lambda/2$ and transmitters with spacing $R\lambda/2$ (or vice versa). The product RT thus scales linearly with aperture, which sets the azimuth resolution. Spatial compression is achieved by using a sparse random-array architecture [28], in which a low number of transmit and receive elements are placed at random over the same aperture Z, achieving similar resolution as a filled array, but with significantly fewer elements. The random-array configuration is illustrated in Figure 11. Beamforming is applied on the time-domain samples obtained from the thinned array at the Nyquist rate, and the azimuths are recovered using CS techniques. Recovery guarantees and guidelines concerning the choice of the product RT and the antenna locations are provided. The methods for choosing the antenna locations using deep networks are investigated in [89].

Spatial compression is achieved by using a sparse random-array architecture, in which a low number of transmit and receive elements are placed at random over the same aperture Z, achieving similar resolution as a filled array, but with significantly fewer elements.

Time and spatial compression

In all of the previously discussed works, recovery is performed in the time domain on acquired or reconstructed Nyquist rate samples for each antenna. The sub-Nyquist MIMO radar (SUMMeR) system, presented in [32], extends the Xampling concept to MIMO configurations and breaks the link between the aperture and the number of antennas, similar to [28]. The concept of Xampling is applied both in space (antenna deployment) and in time (sampling scheme) to simultaneously reduce the required number of antennas and samples per receiver, while preserving time and spatial resolution. In particular, targets' azimuths, ranges, and Dopplers are recovered from compressed samples in both space and time, while keeping the same resolution induced by Nyquist rate samples obtained from a full virtual array with low computational cost.

The SUMMeR system implements a collocated MIMO radar system with $M < T$ transmit and $Q < R$ receive antennas, whose locations are chosen uniformly at random within the aperture of the virtual array described previously in this section, i.e., $\{\xi_m\}_{m=0}^{M-1} \sim \mathcal{U}[0,Z]$ and $\{\zeta_q\}_{q=0}^{Q-1} \sim \mathcal{U}[0,Z]$, respectively, as shown in Figure 11. Note that, in principle, the antenna locations may be chosen on the ULAs' grid; however, this configuration is less robust than range-azimuth ambiguity and leads to coupling between these parameters in the presence of noise [32]. An FDMA framework is adopted so that spatial compression, which, in particular reduces the number of transmit antennas, removes the corresponding transmitting frequency bands as well. The transmitted signals are illustrated in Figure 12 in the frequency domain. Figure 12(a) and (b) shows a standard FDMA transmission for $T = 5$ and the resulting signal after spatial compression for $M = 3$.

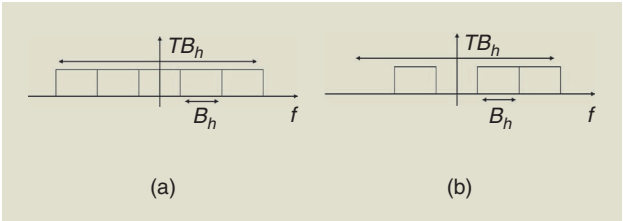

FIGURE 12. The frequency division multiple access transmissions: (a) standard and (b) spatial compression [32].

The transmitted pulses, defined in (12), are reflected by the targets and collected at the receive antennas. Under the assumptions described in "Targets' Assumptions," the received signal $\tilde{x}_q(t)$ at the qth antenna is the sum of time-delayed, scaled replicas of the transmitted signals:

$$\tilde{x}_q(t) = \sum_{m=0}^{T-1} \sum_{l=1}^{L} \alpha_l s_m\left(\frac{c+v_l}{c-v_l}\left(t - \frac{R_{l,mq}}{c+v_l}\right)\right), \quad (48)$$

where $R_{l,mq}$ is the sum of the distances from the mth transmitter and qth receiver to the lth target, which account for the array geometry. After demodulation to the baseband, the received signal can be further simplified to

$$x_q(t) = \sum_{p=0}^{P-1} \sum_{m=0}^{M-1} \sum_{l=1}^{L} \alpha_l h_m(t - p\tau - \tau_l) e^{j2\pi\beta_{mq}\vartheta_l} e^{j2\pi f_l^D p\tau}, \quad (49)$$

where $\beta_{mq} = (\zeta_q + \xi_m)(f_m \frac{\lambda}{c} + 1)$, with f_m the mth-transmission-carrier frequency and λ the signal wavelength. The goal is to estimate the targets' ranges, azimuths, and velocities, i.e., to estimate τ_l, ϑ_l, and f_l^D from low-rate samples of $x_q(t)$, and small numbers m and Q of the antennas.

Similar to the Xampling processing in [30], SUMMeR considers the Fourier coefficients of the received signal $x_q^p(t)$ at the qth antenna. To jointly recover the targets' ranges, azimuths, and Doppler frequencies, the concept of Doppler focusing from [30] (see "Doppler Focusing") is applied to the MIMO setting, and the CS algorithms are extended to simultaneous matrix recovery [32]. The minimal number of channels required for the recovery of L targets' parameters in noiseless settings is $MQ \geq 2L$, with a minimal number of $MK \geq 2L$ samples per receiver and $P \geq 2L$ pulses per transmitter [32]. The SUMMeR system has been implemented in hardware, as described in the following section.

Hardware prototype

The cognitive SUMMeR prototype [41], [42] extends the Doppler focusing, Xampling-based prototype [40] to the MIMO configuration. It simultaneously recovers the targets' delays, Dopplers, and azimuths from sub-Nyquist samples. More specifically, it implements a receiver with a maximum of eight transmit and ten receive antenna elements. The same hardware is used for each receive element and serially feeds the signals of all $R = 10$ receivers to the same prototype.

To avoid the use of an overwhelmingly large number of ADCs and bandpass filters for an 8×10 array, a cognitive transmission is adopted wherein each transmit signal lies in $N_b = 8$ disjoint, narrow slices over a 15-MHz band. Each subband is the width of 375 kHz, leading to a total signal bandwidth of 3 MHz. The transmit subbands, locations were chosen so that all can be subsampled using a single low-rate ADC without aliasing between them [41]. This allows the reduction of the number of samplers. The signal is subsampled at 7.5 MHz, whereas a noncognitive signal would have occupied the entire 15-MHz spectrum requiring a Nyquist sampling rate of 30 MHz. Therefore, the use of cognitive transmission enables spectral sampling reduction by a factor of four for each channel. The effective signal bandwidth is reduced by a factor of five (= 15/3 MHz), respectively, for each channel.

The system may be configured to operate in various array configurations simulating different numbers and locations of the antennas. The hardware switches off the inactive channels and does not sample any data over the corresponding ADCs. This governs the spatial compression by reducing the number of receivers and transmitters. In its baseline configuration, the system uses only half of the antennas with respect to the full virtual array, i.e., $M = 4$ transmitters and $Q = 5$ receivers. Figure 13 shows the sub-Nyquist MIMO prototype, user interface, and

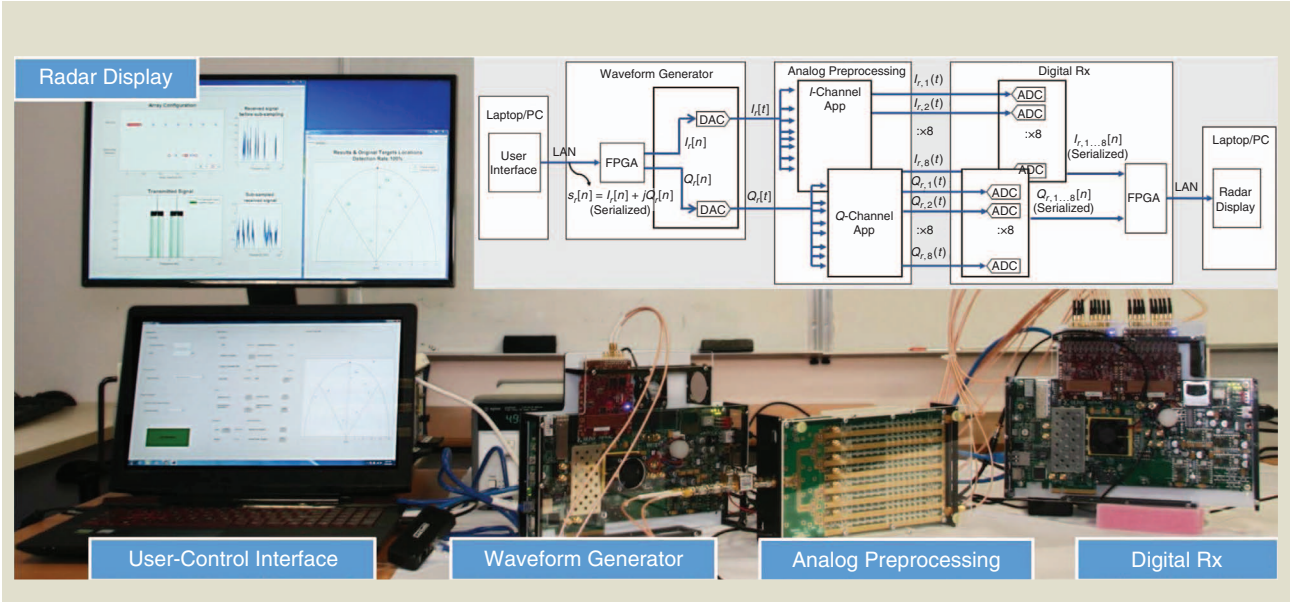

FIGURE 13. The sub-Nyquist MIMO prototype and user interface. The analog preprocessor module consists of two cards mounted on opposite sides of a common chassis. The inset shows the simplified block diagram of the system. The subscript r represents the received signal samples for the rth receiver. Wherever applicable, the second subscript corresponds to a particular transmitter. The square brackets (parentheses) are used for digital (analog) signals [41], [42].

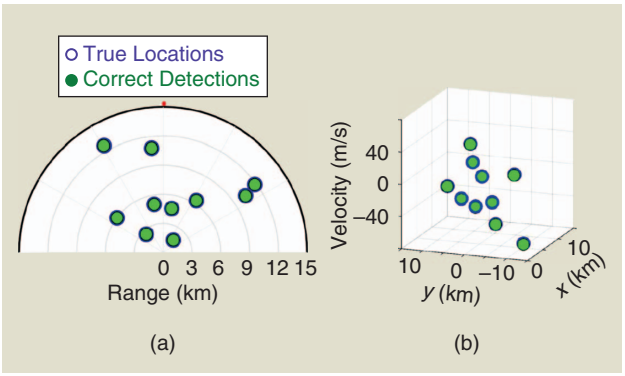

FIGURE 14. The SUMMeR prototype recovery performance: (a) the plan position indicator (PPI) display. The origin is the location of the radar. The red dot indicates the north direction relative to the radar. The positive/negative distances along the horizontal axis correspond to the east/west direction of the radar. Similarly, the positive/negative distances along the vertical axis correspond to the north/south direction of the radar. The estimated targets are plotted over the ground truth. (b) The range-azimuth-Doppler map for the same targets. The lower axes represent the Cartesian coordinates of the polar representation of the PPI plots from (a). The vertical axis represents the Doppler spectrum [32].

radar display. The inset graph depicts the signal flow through a simplified block diagram.

The experimental process consists of the following steps. The simulated radar scenario is stored in a custom-designed waveform generator. The scenario includes pulse-transmission modeling, accurate power loss due to wave propagation in a realistic medium, and interaction of a transmit signal with the target. A large variety of scenarios, consisting of different targets' parameters, i.e., delays, Doppler frequencies, and amplitudes, and array configurations, i.e., the number of transmitters and receivers and antenna locations, may be examined using the prototype. The waveform generator board then produces an analog signal corresponding to the synthesized radar environment, which is amplified and routed to the MIMO radar-receiver board. The prototype samples and processes the signal in real time. The physical array aperture and simulated target response correspond to an X-band ($f_c = 10$ GHz) radar.

Figure 14 presents some recovery results from the prototype. In the experiment, $P = 10$ pulses were transmitted at a uniform PRF of 100 μHz. The received signal corresponding to the echoes from $L = 10$ targets, placed at arbitrary ranges with azimuths and with arbitrary velocities, was injected into the transmit waveform generator. In the experiment, when the angular spacing (in terms of the sine of azimuth) between any two targets was greater than 0.025 and the signal SNR $= -8$ dB, the recovery performance of the compressed configuration in time and in space was equivalent to that of a full array, i.e., with eight transmitters and ten receivers. The figure shows the obtained plan position indicator plot and range-azimuth-Doppler maps for both true and recovered targets. Here, a successful detection (green circle) occurs when the estimated target is within one range cell, one azimuth bin, and one Doppler bin of the ground truth (blue circle). More experiments in [41] and [42] demonstrate that the prototype performance is robust, with SNRs dropping to as low as –10 dB, and the time and spatial resolution are preserved by simulating couples of close targets in range, Doppler, and azimuth.

Conclusions and future challenges

In this article, we reviewed several compressed radar systems that aim to reduce complexity while preserving parameter resolution. Throughout this article, we considered different popular radar systems, including pulse-Doppler and step-frequency radars as well as MIMO configurations. In particular, we showed that temporal, spectral, and spatial compression can be implemented without decreasing Doppler, range, and azimuth resolution. To recover these parameters for L targets, the minimum number of required samples per pulse, the minimum number of pulses, and the minimum number of channels are each equal to $2L$. These are determined by the actual number of DoF of the parameter estimation problem, governed by L, rather than a function of design parameters, such as signal bandwidth, CPI, or aperture. This is essential since the latter determine range, Doppler, and azimuth resolution and are increased for higher performance. By breaking the traditional links between the sampling rate, number of pulses, and antennas on the one hand and parameter estimation on the other hand, increased performance may be achieved without increasing sampling and processing rates.

An advantage of the Xampling system is that traditional radar-processing algorithms can be easily adapted and applied directly to the sub-Nyquist samples. For example, clutter-cancellation techniques have been implemented on the Xampling radar prototypes. These significantly enhance the performance of compressed radars without requiring the reconstruction of Nyquist rate samples. In addition, while CS-based methods traditionally do not perform well in the presence of large noise, since they inherently reduce SNR, Doppler focusing, applied to samples obtained using Xampling, enjoys an SNR improvement that scales linearly with the number of pulses, obtaining good detection at low SNRs.

The transmit subbands, locations were chosen so that all can be subsampled using a single low-rate ADC without aliasing between them.

An essential part of the approach adopted in this survey is the relation between the theoretical algorithms and practical hardware implementation, demonstrating real-time target detection from compressed samples in the fast and slow time domains, as well as in space. The prototypes presented here were built from off-the-shelf components, paving the way to enable commercial, compressed radar systems. To this end, such hardware prototypes should be further extended to implement radar transmit and receive systems and deploy them to be tested on real data. This would permit assessing their performance in real-world conditions, including different types of noise, clutter, and interference.

Authors

Deborah Cohen (deborah.co88@gmail.com) received her B.S. degree in electrical engineering (summa cum laude) in 2010 and her Ph.D. degree in electrical engineering in 2016, both from the Technion–Israel Institute of Technology, Haifa. She has been granted several awards, including the Meyer Foundation Excellence Award, the David and Tova Freud and Ruth Freud-Brendel Memorial Scholarship, the Sandor Szego Award, the Vivian Konigsberg Award, and the Muriel and David Jacknow Award for Excellence in Teaching. Since 2014, she has been an Azrieli fellow. She is currently a research scientist with Google Israel. Her research interests include theoretical aspects of signal processing, compressed sensing, reinforcement learning, and machine learning for dialogues.

Yonina C. Eldar (yonina@ee.technion.ac.il) is a professor in the Department of Electrical Engineering at the Technion–Israel Institute of Technology, Haifa, where she holds the Edwards chair in engineering. She is also an adjunct professor at Duke University, Durham, North Carolina, and a research affiliate with the Research Laboratory of Electronics at the Massachusetts Institute of Technology, Cambridge, and she was a visiting professor at Stanford University, California. She is a member of the Israel Academy of Sciences and Humanities and a Fellow of the IEEE and the European Association for Signal Processing. She has received many awards for excellence in research and teaching, including the IEEE Signal Processing Society Technical Achievement Award, the IEEE/Aerospace and Electronic Systems Society Fred Nathanson Memorial Radar Award, the IEEE Kiyo Tomiyasu Award, the Michael Bruno Memorial Award from the Rothschild Foundation, the Weizmann Prize for Exact Sciences, and the Wolf Foundation Krill Prize for Excellence in Scientific Research. She is the editor-in-chief of *Foundations and Trends in Signal Processing* and serves the IEEE on several technical and award committees.

References

[1] X. P. Masbernat, M. G. Amin, F. Ahmad, and C. Ioana, "An MIMO-MTI approach for through-the-wall radar imaging applications," in *Proc. IEEE Int. Waveform Diversity and Design Conf.*, 2010, pp. 188–192.

[2] D. J. Daniels, *Ground Penetrating Radar*. London, U.K.: Institution of Engineering and Technology, 2004.

[3] K. Schuler, M. Younis, R. Lenz, and W. Wiesbeck, "Array design for automotive digital beamforming radar system," in *Proc. IEEE Int. Radar Conf.*, 2005, pp. 435–440.

[4] V. Bringi and V. Chandrasekar, *Polarimetric Doppler Weather Radar: Principles and Applications*. Cambridge, U.K.: Cambridge Univ. Press, 2001.

[5] M. Skolnik, *Radar Handbook*. New York: McGraw-Hill, 1970.

[6] P. Z. Peebles, *Radar Principles*. Hoboken, NJ: Wiley, 2007.

[7] M. A. Richards, *Fundamentals of Radar Signal Processing*. New York: McGraw-Hill, 2014.

[8] E. Fishler, A. Haimovich, R. Blum, D. Chizhik, L. Cimini, and R. Valenzuela, "MIMO radar: An idea whose time has come," in *Proc. IEEE Radar Conf.*, 2004, pp. 71–78.

[9] J. Li and P. Stoica, *MIMO Radar Signal Processing*. Piscataway, NJ: IEEE Press, 2009.

[10] Y. C. Eldar, *Sampling Theory: Beyond Bandlimited Systems*. Cambridge, U.K.: Cambridge Univ. Press, 2015.

[11] J. C. Curlander and R. N. McDonough, *Synthetic Aperture Radar*. New York: Wiley, 1991.

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

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

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

[15] M. A. Herman and T. Strohmer, "High-resolution radar via compressed sensing," *IEEE Trans. Signal Process.*, vol. 57, no. 6, pp. 2275–2284, 2009.

[16] Y. Chi, R. Calderbank, and A. Pezeshki, "Golay complementary waveforms for sparse delay-Doppler radar imaging," in *Proc. IEEE Int. Workshop Computational Advances Multi-Sensor Adaptive Processing*, 2009, pp. 177–180.

[17] S. Shah, Y. Yu, and A. Petropulu, "Step-frequency radar with compressive sampling (SFR-CS)," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, 2010, pp. 1686–1689.

[18] T. Strohmer and H. Wang, "Sparse MIMO radar with random sensor arrays and Kerdock codes," in *Proc. IEEE Int. Conf. Sampling Theory and Applications*, 2013, pp. 517–520.

[19] T. Strohmer and B. Friedlander, "Analysis of sparse MIMO radar," *Appl. Comput. Harmon. Anal.*, vol. 37, no. 3, pp. 361–388, 2014.

[20] W. U. Bajwa, K. Gedalyahu, and Y. C. Eldar, "Identification of parametric underspread linear systems and super-resolution radar," *IEEE Trans. Signal Process.*, vol. 59, no. 6, pp. 2548–2561, 2011.

[21] O. Teke, A. C. Gurbuz, and O. Arikan, "A robust compressive sensing-based technique for reconstruction of sparse radar scenes," *Digit. Signal Process.*, vol. 27, pp. 23–32, Apr. 2014.

[22] J. H. Ender, "On compressive sensing applied to radar," *Signal Process.*, vol. 90, no. 5, pp. 1402–1414, 2010.

[23] Y. Yu, A. P. Petropulu, and H. V. Poor, "MIMO radar using compressive sampling," *IEEE J. Sel. Topics Signal Process.*, vol. 4, no. 1, pp. 146–163, 2010.

[24] D. S. Kalogerias and A. P. Petropulu, "Matrix completion in colocated MIMO radar: Recoverability, bounds & theoretical guarantees," *IEEE Trans. Signal Process.*, vol. 62, no. 2, pp. 309–321, 2014.

[25] S. Sun, W. U. Bajwa, and A. P. Petropulu, "MIMO-MC radar: A MIMO radar approach based on matrix completion," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 51, no. 3, pp. 1839–1852, 2015.

[26] E. Ertin, L. C. Potter, and R. L. Moses, "Sparse target recovery performance of multi-frequency chirp waveforms," in *Proc. IEEE European Signal Processing Conf.*, 2011, pp. 446–450.

[27] C. Liu, F. Xi, S. Chen, Y. D. Zhang, and Z. Liu, "Pulse-Doppler signal processing with quadrature compressive sampling," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 51, no. 2, pp. 1217–1230, 2015.

[28] M. Rossi, A. M. Haimovich, and Y. C. Eldar, "Spatial compressive sensing for MIMO radar," *IEEE Trans. Signal Process.*, vol. 62, no. 2, pp. 419–430, 2014.

[29] 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.

[30] O. Bar-Ilan and Y. C. Eldar, "Sub-Nyquist radar via Doppler focusing," *IEEE Trans. Signal Process.*, vol. 62, no. 7, pp. 1796–1811, 2014.

[31] D. Cohen, A. Dikopoltsev, R. Ifraimov, and Y. C. Eldar, "Towards sub-Nyquist cognitive radar," in *Proc. IEEE Radar Conf.*, 2016. doi: 10.1109/RADAR.2016.7485122.

[32] D. Cohen, D. Cohen, Y. C. Eldar, and A. M. Haimovich, "SUMMeR: Sub-Nyquist MIMO radar," *IEEE Trans. Signal Process.*, vol. 66, no. 16, pp. 4315–4330, 2018.

[33] L. C. Potter, E. Ertin, J. T. Parker, and M. Cetin, "Sparsity and compressed sensing in radar imaging," *Proc. IEEE*, vol. 98, no. 6, pp. 1006–1020, 2010.

[34] M. Cetin, I. Stojanovic, N. O. Onhon, K. R. Varshney, S. Samadi, W. C. Karl, and A. S. Willsky, "Sparsity-driven synthetic aperture radar imaging: Reconstruction, autofocusing, moving targets, and compressed sensing," *IEEE Signal Process. Mag.*, vol. 31, no. 4, pp. 27–40, 2014.

[35] L. Zhao, L. Wang, L. Yang, A. M. Zoubir, and G. Bi, "The race to improve radar imagery: An overview of recent progress in statistical sparsity-based techniques," *IEEE Signal Process. Mag.*, vol. 33, no. 6, pp. 85–102, 2016.

[36] H. Griffiths, L. Cohen, S. Watts, E. Mokole, C. Baker, M. Wicks, and S. Blunt, "Radar spectrum engineering and management: Technical and regulatory issues," *Proc. IEEE*, vol. 103, no. 1, pp. 85–102, 2015.

[37] M. P. Fitz, T. R. Halford, I. Hossain, and S. W. Enserink, "Towards simultaneous radar and spectral sensing," in *Proc. IEEE Int. Symp. Dynamic Spectrum Access Networks*, 2014, pp. 15–19.

[38] J. Bernhard, J. Reed, J. M. Park, A. Clegg, A. Weisshaar, and A. Abouzeid. (2010). Final report of the National Science Foundation workshop on enhancing

access to the radio spectrum (EARS). Nat. Sci. Foundation. Arlington. [Online]. Available: https://www.nsf.gov/mps/ast/nsf_ears_workshop_2010_final_report.pdf

[39] D. Cohen, K. V. Mishra, and Y. C. Eldar, "Spectrum sharing radar: Coexistence via Xampling," *IEEE Trans. Aerosp. Electron. Syst.,* vol. 54, pp. 1279–1296, 2018.

[40] E. Baransky, G. Itzhak, I. Shmuel, N. Wagner, E. Shoshan, and Y. C. Eldar, "Sub-Nyquist radar prototype: Hardware and algorithms," *IEEE Trans. Aerosp. Electron. Syst*., vol. 50, no. 1, pp. 809–822, 2014.

[41] K. V. Mishra, E. Shoshan, M. Namer, M. Meltsin, D. Cohen, R. Madmoni, S. Dror, R. Ifraimov, et al., "Cognitive sub-Nyquist hardware prototype of a collocated MIMO radar," in *Proc. Compressed Sensing Theory and its Applications to Radar, Sonar and Remote Sensing (CoSeRa)*, 2016, pp. 56–60.

[42] D. Cohen, K. V. Mishra, D. Cohen, E. Ronen, Y. Grimovich, M. Namer, M. Meltsin, and Y. C. Eldar, "Cognitive sub-Nyquist MIMO radar prototype with Doppler processing," in *IEEE Radar Conf.*, 2016, pp. 1179–1184.

[43] J. Yoo, C. Turnes, E. B. Nakamura, C. K. Le, S. Becker, E. A. Sovero, M. B. Wakin, M. C. Grant, et al., "A compressed sensing parameter extraction platform for radar pulse signal acquisition," *IEEE J. Emerg. Sel. Topics Circuits Syst.*, vol. 2, no. 3, pp. 626–638, 2012.

[44] E. Fishler, A. Haimovich, R. S. Blum, and L. J. Cimini, "Spatial diversity in radars-models and detection performance," *IEEE Trans. Signal Process*., vol. 54, pp. 823–838, Mar. 2006.

[45] C. Cook, *Radar Signals: An Introduction to Theory and Application*. Amsterdam, The Netherlands: Elsevier, 2012.

[46] J. Li and P. Stoica, "MIMO radar with collocated antennas," *IEEE Signal Process. Mag.*, vol. 24, no. 5, pp. 106–114, 2007.

[47] A. M. Haimovich, R. S. Blum, and L. J. Cimini, "MIMO radar with widely separated antennas," *IEEE Signal Process. Mag.*, vol. 25, no. 1, pp. 116–129, 2008.

[48] D. Bliss and K. Forsythe, "Multiple-input multiple-output (MIMO) radar and imaging: Degrees of freedom and resolution," in *Proc. IEEE Asilomar Conf. Signals, Systems and Computers*, 2003, pp. 54–59.

[49] C.-Y. Chen, "Signal processing algorithms for MIMO radar," Ph.D. dissertation, California Inst. of Technol., Pasadena, 2009.

[50] P. Vaidyanathan, P. Pal, and C.-Y. Chen, "MIMO radar with broadband waveforms: Smearing filter banks and 2D virtual arrays," in *Proc. IEEE Asilomar Conf. Signals, Systems and Computers*, , pp. 188–192.

[51] Y. C. Pati, R. Rezaiifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in *Proc. IEEE Asilomar Conf. Signals, Systems and Computers*, 1993, pp. 40–44.

[52] G. M. Davis, S. G. Mallat, and Z. Zhang, "Adaptive time-frequency decompositions with matching pursuit," *SPIE*, vol. 2242, pp. 402–414, Mar. 1994. doi: 10.1117/12.170041.

[53] T. Blumensath and M. Davies, "Gradient pursuits," *IEEE Trans. Signal Process.*, vol. 56, pp. 2370–2382, Jun. 2008.

[54] Y. C. Eldar, R. Levi, and A. Cohen, "Clutter removal in sub-Nyquist radar," *IEEE Signal Process. Lett.*, vol. 22, no. 2, pp. 177–181, 2015.

[55] R. Baraniuk and P. Steeghs, "Compressive radar imaging," in *Proc. IEEE Radar Conf*., 2007, pp. 128–133.

[56] W. O. Alltop, "Complex sequences with low periodic correlations," *IEEE Trans. Inf. Theory*, vol. 26, no. 3, pp. 350–354, 1980.

[57] Z. Yang, C. Zhang, and L. Xie, "Robustly stable signal recovery in compressed sensing with structured matrix perturbation," *IEEE Trans. Signal Process.*, vol. 60, no. 9, pp. 4658–4671, 2012.

[58] M. A. Hadi, S. Alshebeili, K. Jamil, and F. E. A. El-Samie, "Compressive sensing applied to radar systems: An overview," *Signal, Image Video Process*., vol. 9, no. 1, pp. 25–39, 2015.

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

[60] J. Yoo, S. Becker, M. Monge, M. Loh, E. Candes, and A. Emami-Neyestanak, "Design and implementation of a fully integrated compressed-sensing signal acquisition system," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, 2012, pp. 5325–5328.

[61] H. Liu, A. Ghafoor, and P. H. Stockmann, "A new quadrature sampling and processing approach," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 25, no. 5, pp. 733–748, 1989.

[62] X. Song, S. Zhou, and P. Willett, "The role of the ambiguity function in compressed sensing radar," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, 2010, pp. 2758–2761.

[63] K. Gedalyahu, R. Tur, and Y. C. Eldar, "Multichannel sampling of pulse streams at the rate of innovation," *IEEE Trans. Signal Process.*, vol. 59, no. 4, pp. 1491–1504, 2011.

[64] R. Tur, Y. C. Eldar, and Z. Friedman, "Innovation rate sampling of pulse streams with application to ultrasound imaging," *IEEE Trans. Signal Process.*, vol. 59, no. 4, pp. 1827–1842, 2011.

[65] D. Cohen and Y. C. Eldar, "Reduced time-on-target in pulse Doppler radar: Slow time domain compressed sensing," in *Proc. IEEE Radar Conf.*, 2016. doi: 10.1109/RADAR.2016.7485243.

[66] T. Wimalajeewa, Y. C. Eldar, and P. K. Varshney. (2013, 11 Nov.). Recovery of sparse matrices via matrix sketching. arXiv. [Online]. Available: http://arxiv.org/abs/1311.2448

[67] R. J. Doviak and D. S. Zrnic, *Doppler Radar & Weather Observations*. New York: Academic, 2014.

[68] A. Ferrari, G. Alengrin, and C. Theys, "Doppler ambiguity resolution using staggered PRF with a new chirp sweep-rate estimation algorithm," *IEEE Proc. Radar, Sonar Navigation*, vol. 142, no. 4, pp. 191–194, 1995.

[69] V. Venkatesh, L. Li, M. McLinden, G. Heymsfield, and M. Coon, "A frequency diversity pulse-pair algorithm for extending Doppler radar velocity Nyquist range," in *Proc. IEEE Radar Conf.*, 2016, pp. 1–6.

[70] A. Ludloff and M. Minker, "Reliability of velocity measurement by MTD radar," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-21, no. 4, pp. 522–528, 1985.

[71] G. Trunk and S. Brockett, "Range and velocity ambiguity resolution," in *Proc. IEEE Nat. Radar Conf.*, 1993, pp. 146–149.

[72] X. Liu, D. Cohen, T. Huang, Y. Liu, G. Winerich, L. Shani, and Y. C. Eldar, "Unambiguous delay-Doppler recovery from phased coded pulses," 2016, to be published.

[73] Q. Cao, G. Zhang, R. D. Palmer, and L. Lei, "Detection and mitigation of second-trip echo in polarimetric weather radar employing random phase coding," *IEEE Trans. Geosci. Remote Sens.*, vol. 50, no. 4, pp. 1240–1253, 2012.

[74] S. Haykin, "Cognitive radar: A way of the future," *IEEE Signal Process. Mag.*, vol. 23, no. 1, pp. 30–40, 2006.

[75] K. V. Mishra, D. Cohen, S. Tsiper, S. Stein, E. Shoshan, M. Namer, M. Meltsin, R. Madmoni, E. Ronen, Y. Grimovich, and Y. C. Eldar, "Xampling-enabled coexistence in spectrally crowded environments," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, 2017, pp. 6580–6581.

[76] G. M. Jacyna, B. Fell, and D. McLemore, "A high-level overview of fundamental limits studies for the DARPA SSPARC program," in *Proc. IEEE Radar Conf*., 2016. doi: 10.1109/RADAR.2016.7485100.

[77] P. Stinco, M. S. Greco, and F. Gini, "Spectrum sensing and sharing for cognitive radars," *IET Radar, Sonar Navigation*, vol. 10, no. 3, pp. 595–602, 2016.

[78] N. Nartasilpa, D. Tuninetti, N. Devroye, and D. Erricolo, "Let's share CommRad: Effect of radar interference on an uncoded data communication system," in *Proc. IEEE Radar Conf.*, 2016. doi: 10.1109/RADAR.2016.7485064.

[79] P. Kumari, N. Gonzalez-Prelcic, and R. W. Heath, "Investigating the IEEE 802.11ad standard for millimeter wave automotive radar," in *Proc. Vehicular Technol. Conf.*, 2015. doi: 10.1109/VTCFall.2015.7390996.

[80] J. Mitola and C. Q. Maguire, Jr., "Cognitive radio: Making software radios more personal," *IEEE Personal Commun.*, vol. 6, no. 4, pp. 13–18, 1999.

[81] D. Cohen, S. Tsiper, and Y. C. Eldar, "Analog-to-digital cognitive radio: Sampling, detection, and hardware," *IEEE Signal Process. Mag.*, vol. 35, no. 1, pp. 137–166, 2018.

[82] N. Vaswani and W. Lu, "Modified-CS: Modifying compressive sensing for problems with partially known support," *IEEE Trans. Signal Process.*, vol. 58, no. 9, pp. 4595–4607, 2010.

[83] S. M. Kay, *Fundamentals of Statistical Signal Processing, Volume II: Detection Theory*. Englewood Cliffs, NJ: Prentice Hall, 1998.

[84] J. Huang, T. Zhang, and D. Metaxas, "Learning with structured sparsity," *J. Mach. Learn. Res.*, vol. 12, pp. 3371–3412, Nov. 2011.

[85] 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.

[86] R. Gold, "Optimal binary sequences for spread spectrum multiplexing (Corresp.)," *IEEE Trans. Inf. Theory*, vol. 13, no. 4, pp. 619–621, 1967.

[87] C.-Y. Chen and P. Vaidyanathan, "Compressed sensing in MIMO radar," in *Proc. IEEE Asilomar Conf. Signals, Systems and Computers*, 2008, pp. 41–44.

[88] R. O. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Trans. Antennas Propag.*, vol. 34, no. 4, pp. 276–280, 1986.

[89] A. M. Elbir, K. V. Mishra, and Y. C. Eldar. (2018). Cognitive radar antenna selection via deep learning. arXiv. [Online]. Available: https://arxiv.org/abs/1802.09736