

1 Sub-Nyquist Radar: Principles and Prototypes

Kumar Vijay Mishra* and Yonina C. Eldar**

1.1 Introduction

Radar remote sensing has advanced tremendously since 1950 and is now applied to diverse areas such as military surveillance, meteorology, geology, collision avoidance, and imaging [1]. In monostatic pulse-Doppler radar systems, an antenna transmits a periodic train of known narrowband pulses within a defined coherent processing interval (CPI). When the radiated wave from the radar interacts with moving targets, the amplitude, frequency, and polarization states of the scattered wave change. By monitoring this change, it is possible to infer the targets' size, location, and radial Doppler velocity. The reflected signal received by the radar antenna is a linear combination of echoes from multiple targets; each of these is an attenuated, time-delayed, and frequency-modulated version of the transmit signal. The delay in the received signal is linearly proportional to the target's range or its distance from the radar. The frequency modulation encodes the Doppler velocity of the target. The complex amplitude or target's reflectivity is a function of the target's size, geometry, propagation, and scattering mechanism. Radar signal processing is aimed at detecting the targets and estimating their parameters from the output of this linear, time-varying system.

Traditional radar signal processing employs matched filtering (MF) or pulse compression [2] in the digital domain, wherein the sampled received signal is correlated with a replica of the transmit signal in the delay-Doppler plane. The MF maximizes the signal-to-noise ratio (SNR) in the presence of additive white Gaussian noise. In some specialized systems, this stage is replaced by a *mismatched filter* with a different optimization metric such as minimization of peak-to-sidelobe ratio of the output. Here, the received signal is correlated with a signal that is close but not identical to the transmit signal [3–5]. While all of these techniques reliably estimate target parameters, their resolution is inversely proportional to the support of the ambiguity function of the transmit pulse, thereby restricting ability to super-resolve targets that are closely spaced.

The digital MF operation requires the signal to be sampled at or above the Nyquist sampling rate, which guarantees perfect reconstruction of a bandlimited analog signal [6]. Many modern radar systems use wide bandwidths, typically ranging from hundreds

* K.V.M. acknowledges partial support via the Andrew and Erna Finci Viterbi Postdoctoral Fellowship and the Lady Davis Postdoctoral Fellowship.

** This work is supported by the European Union's Horizon 2020 research and innovation program under grant agreement no. 646804-ERC-COG-BNYQ.

of MHz to GHz, in order to achieve fine radar range resolution. Since the Nyquist sampling rate is twice the baseband bandwidth, the radar receiver requires expensive, high-rate analog-to-digital converters (ADCs). The sampled signal is then also processed at high rates, resulting in significant power, cost, storage, and computational overhead. Recently, in order to mitigate this rate bottleneck, new methods have been proposed that sample signals at sub-Nyquist rates and yet are able to estimate the targets' parameters [6,7].

Analogous trade-offs arise in other aspects of radar system design. For example, the number of transmit pulses governs the resolution in Doppler velocity. The estimation accuracy of target parameters is greatly affected by the radar's *dwell time* [1], i.e., the time duration a directional radar beam spends hitting a particular target. Long dwell times imply a large number of transmit pulses and, therefore, high Doppler precision. But, simultaneously, this negatively affects the ability of the radar to look at targets in other directions. Sub-Nyquist sampling approaches have, therefore, been suggested for the pulse dimension or "slow-time" domain in order to break the link between dwell time and Doppler resolution [8–10].

Finally, radars that deploy antenna arrays deal with similar sampling problems in the spatial domain. A *phased array* radar antenna consists of several radiating elements that form a highly directional radiating beam pattern. Without requiring any mechanical motion, a phased array accomplishes beam-steering electronically by adjusting the relative phase of excitation in the array elements. The operational advantage is the agile scanning of the target scene, ability to track a large number of targets, and efficient search-and-track in the regions of interest [11]. The beam pattern of individual array elements, array geometry, and its size define the overall antenna pattern [12,13], wherein high spatial resolution is achieved by large array apertures. As per the Nyquist Theorem, the array must not admit fewer than two signal samples per spatial period (i.e., radar's operating wavelength) [14]. Otherwise, it introduces *spatial aliasing* or multiple beams in the antenna pattern, thereby reducing its directivity. Often an exceedingly large number of radiating elements are required to synthesize a given array aperture in order to enhance the radar's ability to unambiguously distinguish closely spaced targets; the associated cost, weight, and area may be unacceptable. It is therefore desirable to apply sub-Nyquist techniques to thin a huge array without causing degradation in spatial resolution [15–17].

Sub-Nyquist sampling leads to the development of low-cost, power-efficient, and small-size radar systems that can scan faster and acquire larger volumes than traditional systems. Apart from design benefits, other applications of such systems have been envisioned recently, including imparting hardware-feasible cognitive abilities to the radar [18,19], a role in devising spectrally coexistent systems [20], and extension to imaging [21]. In this chapter, we provide an overview of sub-Nyquist radars, their applications, and hardware realizations.

The outline of the chapter is as follows. In the next section, we overview various reduced-rate techniques for radar system design and explain the benefits of our approach to sub-Nyquist radars. In Section 1.3, we describe the principles, algorithms, and hardware realization of temporal sub-Nyquist monostatic pulse-Doppler radar. Section 1.4

presents an extension of the sub-Nyquist principle to slow time. We then introduce the cognitive radar concept based on sub-Nyquist reception in Section 1.5 and show an application to coexistence in a spectrally crowded environment. Section 1.6 is devoted to spatial sub-Nyquist applications in multiple-input multiple-output (MIMO) array radars. Finally, we consider sub-Nyquist synthetic aperture radar (SAR) imaging in Section 1.7, followed by concluding remarks in Section 1.8.

1.2 Prior Art and Historical Notes

There is a large body of literature on reduced-rate sampling techniques for radars. Most of these works employ compressed sensing (CS) methods, which allow recovery of sparse, undersampled signals from random linear measurements [7]. A pre-2010 review of selected applications of CS-based radars can be found in [22]. A qualitative, system-level commentary from the point of view of operational radar engineers is available in [23], while CS-based radar imaging studies are summarized in [24]. An excellent overview on sparsity-based SAR imaging methods is provided in [25]. The review in [26] recaps major developments in this area from a nonmathematical perspective. In the following, we review the most significant works relevant to the sub-Nyquist formulations presented in this chapter.

On-Grid CS The earliest application of CS toward recovering time delays with sub-Nyquist samples in a noiseless case was formulated in [27]. CS-based parameter estimation for both delay and Doppler shifts was proposed in [28] with samples acquired at the Nyquist rate. These and similar later works [29–31] discretize the delay-Doppler domain, assuming that targets lie on a grid. Subsequently, these ideas were extended to colocated [32,33] and distributed [34] MIMO radars where targets are located on an angle-Doppler-range grid. In practice, target parameters are typically continuous values whose discretization may introduce gridding errors [35]. In particular, [28] constructs a dictionary that exhaustively considers all possible delay-Doppler pairs, thereby rendering the processing computationally expensive. Noise and clutter mitigation are not considered in this literature. Simulations show that such systems typically have poor performance in clutter-contaminated noisy environments.

Off-Grid CS A few recent works [36,37] formulate the radar parameter estimation for off-grid targets using atomic norm minimization [38,39]. However, these methods do not address direct analog sampling, the presence of noise, and clutter. Further details on this approach are available in Chapter 7 (Super-resolution radar imaging via convex optimization) of this book.

Parametric Recovery A different approach was suggested in [40], which treated radar parameter estimation as the identification of an underlying linear, time-varying system [41]. The proposed two-stage recovery algorithm, largely based on [42], first estimates target delays and then utilizes these recovered delays to estimate Doppler velocities and complex reflectivities. They also provide guarantees for system identification in terms

of the minimum time-bandwidth product of the input signal. However, this method does not handle noise well.

Matrix Completion In some radar applications, the received signal samples are processed as data matrices, which, under certain conditions, are low rank. In this context, general works have suggested retrieving the missing entries using matrix completion methods [10,43]. The target parameters are then recovered through classic radar signal processing. These techniques have not been exhaustively evaluated for different signal scenarios and their practical implementations have still not been thoroughly examined.

Finite-Rate-of-Innovation (FRI) Sampling The received radar signal from L targets can be modeled with $3L$ degrees of freedom because three parameters – time delay, Doppler shift, and attenuation coefficient – characterize each target. The classes of signals that have finite degrees of freedom per unit of time are called finite-rate-of-innovation (FRI) signals [44]. Low-rate sampling and signal recovery strategies for FRI signals have been studied in detail in the past [6, chapter 15]. In [45], a temporal sub-Nyquist radar was proposed to recover delays relying on the FRI model. The Xampling framework [6] was used to obtain Fourier coefficients from low-rate samples with a practical hardware prototype. Similar techniques were later studied for delay channel estimation problems in ultra-wideband communication systems [46,47] and for ultrasound imaging [48]. In [49], *Doppler focusing* was added to the FRI-Xampling framework to recover both delays and Dopplers. Doppler focusing is a narrowband technique that can be interpreted as low-rate beamforming in the frequency domain, and was applied earlier to ultrasound imaging [50,51]. It considers a chosen center frequency with a band of frequencies around it and coherently processes multiple echoes in this focused region to estimate the delays from low-rate samples.

Extensions of Sub-Nyquist Radars The system proposed in [49] reduces samples only in time and not in the Doppler domain. Since the set of frequencies for Doppler focusing is usually fixed a priori, the resultant Doppler resolution is limited by the focusing; it remains inversely proportional to the number of pulses P , as is also the case with conventional radar. In [8], sub-Nyquist processing in slow time was introduced to recover the target range and Doppler by simultaneously transmitting few pulses in the CPI and sampling the received signals at sub-Nyquist rates. Later, [52] proposed a whitening procedure to mitigate the presence of clutter in a sub-Nyquist radar. Spatial-domain compressed sensing (SCS) was examined for a MIMO array radar in [16] and later for phased arrays in [15]. Recently, [17] proposed Xampling in time and space to recover delay, Doppler, and azimuth of the targets by *thinning* a colocated MIMO array and collecting low-rate samples at each receive element. This sub-Nyquist MIMO radar (SUMMeR) system was also conceptually demonstrated in hardware [18,53]. The formulation in [54] proposes tensor-based 3D sub-Nyquist radar (TenDSuR) that performs thinning in all three domains and recovers the signal via tensor-based recovery. Finally, an extension to SAR was demonstrated in [21]. Table 1.1 summarizes these developments.

Table 1.1 Sub-Nyquist radars and their corresponding reduction domains.

Sub-Nyquist system	Temporal	Doppler	Spatial
Monostatic pulsed radar [45]	Yes	No	No
Monostatic pulse-Doppler radar [49]	Yes	No	No
Reduced time-on-target radar [8]	Yes	Yes	No
MIMO SCS [16]	No	No	Yes
Phased array SCS [15]	No	No	Yes
SUMMeR [15,17]	Yes	No	Yes
TenDSuR [54]	Yes	Yes	Yes
Sub-Nyquist SAR [21]	Yes	No	No

1.3 Temporal Sub-Nyquist Radar

Consider a standard pulse-Doppler radar that transmits a pulse train

$$r_{Tx}(t) = \sum_{p=0}^{P-1} h(t - p\tau), \quad 0 \leq t \leq P\tau, \quad (1.1)$$

consisting of P uniformly spaced known pulses $h(t)$. The interpulse transmit delay τ is the pulse repetition interval (PRI) or fast time; its reciprocal is the pulse repetition frequency (PRF). The entire duration of P pulses in (1.1) is known as the CPI or slow time. The radar operates at carrier frequency f_c so that its wavelength is $\lambda = c/f_c$, where $c = 3 \times 10^8$ m/s is the speed of light.

In a conventional pulse-Doppler radar, the pulse $h(t) = h_{Nyq}(t)$ is a time-limited baseband function whose continuous-time Fourier transform (CTFT) is $H_{Nyq}(f) = \int_{-\infty}^{\infty} h_{Nyq}(t)e^{-j2\pi ft} dt$. It is assumed that most of the signal's energy lies within the frequencies $\pm B_h/2$, where B_h denotes the effective signal bandwidth, such that the following approximation holds:

$$h_{Nyq}(t) \approx \int_{-B_h/2}^{B_h/2} H_{Nyq}(f)e^{j2\pi ft} df. \quad (1.2)$$

The total transmit power of the radar is defined as

$$\int_{-B_h/2}^{B_h/2} |H_{Nyq}(f)|^2 df = P_T. \quad (1.3)$$

1.3.1 Received Signal Model

Assume that the radar target scene consists of L non-fluctuating point-targets, according to the Swerling-0 target model [1]. The transmit signal is reflected back by the L targets and these echoes are received by the radar processor. The latter aims at recovering the following information about the L targets from the received signal: time delay τ_l , which is linearly proportional to the target's range $d_l = c\tau_l/2$; Doppler frequency ν_l ,

proportional to the target’s radial velocity $\nu_l = c\nu_l/4\pi f_c$; and complex amplitude α_l . The target locations are defined with respect to the polar coordinate system of the radar and their range and Doppler are assumed to lie in the unambiguous time-frequency region, i.e., the time delays are no longer than the PRI, and Doppler frequencies are up to the PRF.

Typically, the radar assumes the following operating conditions, which leads to a simplified expression for the received signal [49]:

- A1** “Far targets”: assuming $\nu_l \ll 2\pi f_c \tau_l / P\tau$, target radar distance is large compared to the distance change during the CPI over which attenuation α_l is allowed to be constant.
- A2** “Slow targets”: assuming $\nu_l \ll 2\pi f_c / P\tau B_h$, target velocity is small enough to allow for constant τ_l during the CPI. In this case, the following piecewise-constant approximation holds $\nu_l t \approx \nu_l p\tau$, for $t \in [p\tau, (p+1)\tau]$.
- A3** “Small acceleration”: assuming $d\nu_l/dt \ll c/2f_c(P\tau)^2$, target velocity remains approximately constant during the CPI allowing for constant ν_l .
- A4** “No time or Doppler ambiguities”: The delay-Doppler pairs (τ_l, ν_l) for all $l \in [1, L]$ lie in the radar’s unambiguous region of delay-Doppler plane defined by $[0, \tau] \times [-\pi/\tau, \pi/\tau]$.
- A5** The pairs in the set (τ_l, ν_l) for all $l \in [1, L]$ are unique.

Under these assumptions, the received signal can be written as

$$r_{R_X}(t) = \sum_{p=0}^{P-1} \sum_{l=0}^{L-1} \alpha_l h(t - \tau_l - p\tau) e^{-j\nu_l t} + w(t), \quad (1.4)$$

for $0 \leq t \leq P\tau$, where $w(t)$ is a zero mean wide-sense stationary random signal with autocorrelation $r_w(s) = \sigma^2 \delta(s)$. It is convenient to express $r_{R_X}(t)$ as a sum of single frames

$$r_{R_X}(t) = \sum_{p=0}^{P-1} r_{R_X}^p(t) + w(t), \quad (1.5)$$

where

$$r_{R_X}^p(t) = \sum_{l=0}^{L-1} \alpha_l h(t - \tau_l - p\tau) e^{-j\nu_l p\tau}, \quad (1.6)$$

for $p\tau \leq t \leq (p+1)\tau$ is the return signal from the p th pulse.

A classical radar signal processor samples each incoming frame $r_{R_X}^p(t)$ at the Nyquist rate B_h to yield the digitized samples $r_{R_X}^p[n], 0 \leq n \leq N-1$, where $N = \tau B_h$. The signal enhancement process employs an MF for the sampled frames $r_{R_X}^p[n]$. This is then followed by Doppler processing where a P -point discrete Fourier transform (DFT) is performed on slow-time samples. By stacking all the N DFT vectors together, a delay-Doppler map is obtained for the target scene. Finally, the time delays τ_l and Doppler shifts ν_l of the targets are located on this map using, e.g., a constant false-alarm rate detector [55].

1.3.2 Sub-Nyquist Delay-Doppler Recovery

Traditional radar systems sample the received signal at the Nyquist rate, determined by the baseband bandwidth of $h(t)$. Our goal is to recover $r_{R_X}^p(t)$ from its samples taken far below this rate. We note that over the interval τ , $r_{R_X}^p(t)$ is completely specified by $\{\alpha_l, \tau_l, \nu_l\}_{l=1}^L$, and is an FRI signal with rate of innovation $3L/\tau$. Hence, in the absence of noise, one would expect to be able to accurately recover $r_{R_X}^p(t)$ from only a few samples per time τ . Since radar signals tend to be sparse in the time domain, simply acquiring a few data samples at a low rate will not generally yield adequate recovery. Indeed, if the separation between samples is larger than the effective spread in time, then with high probability many of the samples will be close to zero. This implies that presampling analog processing must be performed on the frequency-domain support of the radar signal in order to smear the signal in time before low-rate sampling.

The approach we adopt follows the Xampling architecture designed for sampling and processing of analog inputs at rates far below Nyquist, whose underlying structure can be modeled as a union of subspaces (UoS). The input signal belongs to a single subspace, a priori unknown, out of multiple, possibly even infinitely many, candidate subspaces. Xampling consists of two main functions: low-rate analog-to-digital conversion (ADC), in which the input is compressed in the analog domain prior to sampling with commercial devices, and low-rate digital signal processing, in which the input subspace is detected prior to digital signal processing. The resulting sparse recovery is performed using CS techniques adapted to the analog setting. This concept has been applied to both communications [56–59] and radar [49,60], among other applications.

Time-varying linear systems, which introduce both time shifts (delays) and frequency shifts (Doppler shifts), such as those arising in surveillance point-target radar systems, fit nicely into the UoS model. Here, a sparse target scene is assumed, allowing the reduction of the sampling rate without sacrificing delay and Doppler resolution. The Xampling-based system is composed of an ADC, which filters the received signal to predetermined frequencies before taking point-wise samples. These compressed samples, or “Xamples,” contain the information needed to recover the desired signal parameters.

To demonstrate sub-Nyquist sampling, we begin by deriving an expression for the Fourier coefficients of the received signal and show that the target parameters are embodied in them. Let \mathcal{F}_R and f_{Nyq} be the set of all frequencies in the received signal spectrum and the corresponding Nyquist sampling rate, respectively. Consider the Fourier series representation of the aligned frames $r_{R_X}^p(t + p\tau)$, with $r_{R_X}^p(t)$ defined in (1.6):

$$c_p[k] = \int_0^\tau r_{R_X}^p(t + p\tau) e^{-j2\pi kt/\tau} dt = \frac{1}{\tau} H[k] \sum_{l=0}^{L-1} \alpha_l e^{-j2\pi k\tau_l/\tau} e^{-j\nu_l p\tau}, \quad (1.7)$$

for $k \in \kappa$, where $\kappa = \{k = \lfloor \frac{f}{f_{\text{Nyq}}} N \rfloor \mid f \in \mathcal{F}_R\}$. From (1.7), we see that the unknown parameters $\{\alpha_l, \tau_l, \nu_l\}_{l=0}^{L-1}$ are embodied in the Fourier coefficients $c_p[k]$. We can estimate these parameters using only a small number of Fourier coefficients, which translates to a low sampling rate.

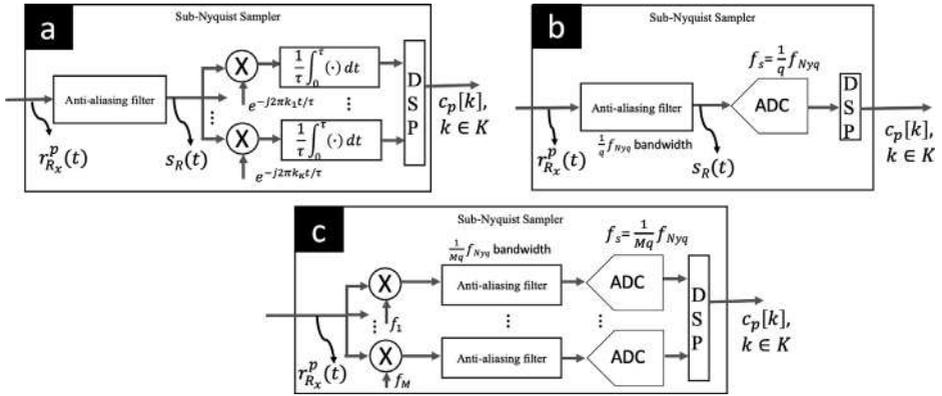


Figure 1.1 Sub-Nyquist sampling methods: (a) direct sampling; (b) low frequencies only; (c) multiband sampling.

There are several ways to implement a sub-Nyquist sampler [47,61] in order to obtain a set of Fourier coefficients from low-rate samples of the signal. For simplicity, consider $|\kappa| = K$ such that $q = N/K$ is an integer defining the sampling reduction factor. In *direct sampling* (Figure 1.1.a), the signal $r_{R_x}(t)$ obtained after the anti-aliasing filter is passed through as many analog chains as the number of sub-Nyquist coefficients K . Each branch is modulated by a complex exponential, followed by integration over τ and necessary digital signal processing (DSP). This technique provides the largest flexibility in choosing the Fourier coefficients, but is expensive in terms of hardware. Another approach is to limit the bandwidth of the anti-aliasing filter such that only the *lowest* K frequencies are free of aliasing (Figure 1.1.b). We then sample these lowest K frequencies. Here, the measurements are correlated and a modification in the analog hardware is also required so that the anti-aliasing filter has reduced passband. In the *multiband* sampling method shown in Figure 1.1.c, M disjoint randomly chosen groups of consecutive Fourier coefficients are obtained such that the total number of sampled coefficients is K . This translates to splitting the signal across M branches, passing the downconverted signal through reduced-passband anti-aliasing filters, and then sampling each band with a low-rate ADC. This method can be easily implemented but requires M low-rate ADCs. The sub-Nyquist hardware prototypes developed in [45,49] adopt multiband sampling using four groups of consecutive coefficients. In practice, the specific Fourier coefficients are chosen through extensive software simulations to provide low mutual coherence [6] for CS-based signal recovery.

Our goal now is to recover $\{\alpha_l, \tau_l, \nu_l\}_{l=0}^{L-1}$ from $c_p[k]$ for $k \in \kappa$ and $0 \leq p \leq P - 1$. To that end, [49] adopts the Doppler focusing approach. Consider the DFT of the coefficients $c_p[k]$ in the slow-time domain:

$$\tilde{\Psi}_\nu[k] = \sum_{p=0}^{P-1} c_p[k] e^{j\nu p \tau} = \frac{1}{\tau} H[k] \sum_{l=0}^{L-1} \alpha_l e^{-j2\pi k \tau_l / \tau} \sum_{p=0}^{P-1} e^{j(\nu - \nu_l) p \tau}. \quad (1.8)$$

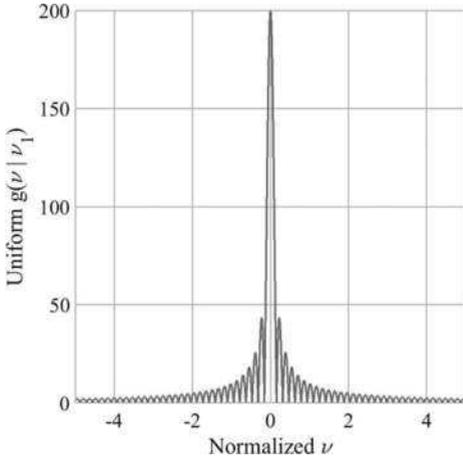


Figure 1.2 Sum of exponents $|g(\nu|\nu_l)|$ for $P = 200$, $\tau = 1$ s, and $\nu_l = 0$ [20,49]. ©2018 IEEE. Reprinted, with permission, from [20].

The key to Doppler focusing follows from the approximation:

$$g(\nu|\nu_l) = \sum_{p=0}^{P-1} e^{j(\nu-\nu_l)p\tau} \approx \begin{cases} P & |\nu - \nu_l| < \pi/P\tau \\ 0 & |\nu - \nu_l| \geq \pi/P\tau, \end{cases} \quad (1.9)$$

as illustrated in Figure 1.2. Denote the normalized focused measurements by

$$\Psi_\nu[k] = \frac{\tau}{PH[k]} \tilde{\Psi}_\nu[k]. \quad (1.10)$$

As in traditional pulse-Doppler radar, suppose we limit ourselves to the Nyquist grid so that $\tau_l/\tau = r_l/N$, where r_l is an integer satisfying $0 \leq r_l \leq N - 1$. Then, (1.10) can be approximately written in vector form as

$$\Psi_\nu = \mathbf{F}_\kappa \mathbf{x}_\nu, \quad (1.11)$$

where $\Psi_\nu = [\Psi_\nu[k_0] \dots \Psi_\nu[k_{K-1}]]$, $k_i \in \kappa$ for $0 \leq i \leq K - 1$, \mathbf{F}_κ is composed of the K rows of the $N \times N$ Fourier matrix indexed by κ , and \mathbf{x}_ν is an L -sparse vector that contains the values α_l at the indices r_l for the Doppler frequencies ν_l in the “focus zone,” that is, $|\nu - \nu_l| < \pi/P\tau$. It is convenient to write (1.11) in matrix form, by vertically concatenating the vectors Ψ_ν , for ν on the Nyquist grid, namely $\nu = -\frac{1}{2\tau} + \frac{1}{P\tau}$, into the $K \times P$ matrix Ψ , as

$$\Psi = \mathbf{F}_\kappa \mathbf{X}, \quad (1.12)$$

where \mathbf{X} is formed similarly by vertically concatenating the vectors \mathbf{x}_ν . Note that the matrix \mathbf{F}_κ is not square and, as a result, the system of linear equations (1.12) is underdetermined. The system in (1.12) can be solved using any CS algorithm, such as orthogonal matching pursuit (OMP) and ℓ_1 minimization [6,7].

A Nyquist receiver needs $B_h\tau$ samples to recover the targets. However, as stated shortly in Theorem 1.3.1, the number of samples required by the sub-Nyquist receiver

depends only on the number of targets present and not on B_h . This shows that a sub-Nyquist radar breaks the link between range resolution and transmit bandwidth. In general, only a few targets are present in the radar coverage region leading to a significant reduction in sampling rate.

THEOREM 1.3.1 [49] *The minimal number of samples required for perfect recovery of $\{\alpha_l, \tau_l, \nu_l\}_{l=0}^L$ in a noiseless environment is $4L^2$. In addition, the number of samples per period is at least $2L$, and the number of periods $P \geq 2L$.*

The Doppler focusing operation (1.8) is a continuous operation on the variable ν , and can be performed for any Doppler frequency up to the PRF. With Doppler focusing there are no inherent blind speeds, i.e., target velocities that are undetectable, as occurs with a classic moving target indicator [1]. Since strong amplitudes are indicative of true target existence as opposed to noise, Doppler focusing recovery searches for large magnitude entries and marks them as detections. After detecting each target, its influence is removed from the set of Fourier coefficients in order to reduce masking of weaker targets and to remove spurious targets created by processing sidelobes. It is important to note that our dictionary in (1.12) is indifferent to the Doppler estimation. CS methods, which estimate delay and Doppler simultaneously [28], require a dictionary that grows with the number of pulses. Here, by separating delay and Doppler estimation, the CS dictionary is not a function of P .

Moreover, the performance of the sub-Nyquist radar in the presence of noise improves with Doppler focusing. The following theorem states that Doppler focusing increases the per-target SNR by a factor of P . This linear scaling is similar to that obtained using an MF.

THEOREM 1.3.2 [49] *Let the prefocusing SNR of the l th target be $\Gamma_p^l[k] = \frac{|c_p^l[k]|^2}{\mathbb{E}[|w_p[k]|^2]}$ where $c_p^l[k]$ and $w_p[k]$ are the signal and white noise Fourier coefficients. Then, the focused SNR for the l th target at center frequency ν is $P\Gamma_p^l[k]$.*

A continuous-value parameter recovery using Doppler focusing is described in [49]. For practical considerations of computational efficiency, Doppler focusing can be performed on a uniform grid of frequencies so that focused coefficients are computed efficiently using the fast Fourier transform (FFT). Algorithm 1 in this section outlines this approach to solving the P equations (1.12) simultaneously, where, in each iteration, the maximal projection of the observation vectors onto the measurement matrix is retained. The algorithm termination criterion follows from the generalized likelihood ratio test (GLRT) framework presented in [62]. For each iteration, the alternative and null hypotheses in the GLRT problem define the presence or absence of a candidate target, respectively. In Algorithm 1, $Q\chi_2^2(\rho)$ denotes the right-tail probability of the chi-square distribution function with 2 degrees of freedom, Λ^C is the complementary set of Λ and

$$\rho = \frac{P_T}{\sigma^2 |\mathcal{F}_R|} \quad (1.13)$$

is the SNR with σ^2 the noise variance and P_T the total transmit power.