

 Open access • Journal Article • DOI:10.1109/TSP.2014.2304917

Sub-Nyquist Radar via Doppler Focusing — [Source link](#)

Omer Bar-Ilan, Yonina C. Eldar

Institutions: Technion – Israel Institute of Technology

Published on: 01 Apr 2014 - IEEE Transactions on Signal Processing (IEEE)

Topics: Continuous-wave radar, Bistatic radar, Doppler radar, Radar imaging and Radar

Related papers:

- [Sub-Nyquist radar prototype: Hardware and algorithm](#)
- [Compressed sensing : theory and applications](#)
- [Sampling Theory: Beyond Bandlimited Systems](#)
- [Xampling: Analog to digital at sub-Nyquist rates](#)
- [Innovation Rate Sampling of Pulse Streams With Application to Ultrasound Imaging](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/sub-nyquist-radar-via-doppler-focusing-4c9aj33kcd>

Sub-Nyquist Radar via Doppler Focusing

Omer Bar-Ilan and Yonina C. Eldar, *Fellow, IEEE*

Abstract—We investigate the problem of a monostatic pulse-Doppler radar transceiver trying to detect targets sparsely populated in the radar’s unambiguous time-frequency region. Several past works employ compressed sensing (CS) algorithms to this type of problem but either do not address sample rate reduction, impose constraints on the radar transmitter, propose CS recovery methods with prohibitive dictionary size, or perform poorly in noisy conditions. Here, we describe a sub-Nyquist sampling and recovery approach called Doppler focusing, which addresses all of these problems: it performs low rate sampling and digital processing, imposes no restrictions on the transmitter, and uses a CS dictionary with size, which does not increase with increasing number of pulses P . Furthermore, in the presence of noise, Doppler focusing enjoys a signal-to-noise ratio (SNR) improvement, which scales linearly with P , obtaining good detection performance even at SNR as low as -25 dB. The recovery is based on the Xampling framework, which allows reduction of the number of samples needed to accurately represent the signal, directly in the analog-to-digital conversion process. After sampling, the entire digital recovery process is performed on the low rate samples without having to return to the Nyquist rate. Finally, our approach can be implemented in hardware using a previously suggested Xampling radar prototype.

Index Terms—Compressed sensing, rate of innovation, radar, sparse recovery, sub-Nyquist sampling, delay-Doppler estimation.

I. INTRODUCTION

WE consider target detection and parameter estimation in a pulse-Doppler radar system, using sub-Nyquist sampling rates. The radar is a single transceiver, monostatic, narrow-band system. Targets are non-fluctuating point targets, sparsely populated in the radar’s unambiguous time-frequency region: delays up to the Pulse Repetition Interval (PRI), and Doppler frequencies up to its reciprocal, the Pulse Repetition Frequency (PRF). We propose a recovery method which can detect and estimate targets’ time delay and Doppler frequency, using a linear, non-adaptive sampling technique at a rate significantly lower than the radar signal’s Nyquist frequency, assuming the number of targets L is small.

Current state-of-the-art radar systems sample at the signal’s Nyquist rate, which can be hundreds of MHz and even up to several GHz. Systems exploiting sub-Nyquist sampling rates benefit from a lower rate analog-to-digital conversion (ADC), which requires less computational power. Moreover, sampling

at the Nyquist rate may not always be feasible due to high power consumption, heat dissipation, cost, or other practical considerations. Finally, offline radar systems which record samples for subsequent processing can gain substantial storage capacity reduction if they were to sample at sub-Nyquist rates.

The goal of this work is to present some steps towards breaking the link between radar signal bandwidth and sampling rate. We mainly focus on the setting of Swerling-0 point targets with additive white Gaussian noise (AWGN), in order to bring forth what we believe are useful ideas and algorithms for sample rate reduction. We also briefly treat target dynamic range and clutter. However, a full analysis of these important issues is left to future work. The sub-Nyquist Xampling (“compressed sampling”) [1] method we use is an ADC which performs analog prefiltering of the signal before taking point-wise samples. These compressed samples (“Xamples”) contain the information needed to recover the desired signal parameters using compressed sensing (CS) algorithms. This work expands on the results of [2], adding Doppler to the target model and proposing a new method to estimate it. The same sampling technique and hardware that were used in [2], [3] are also applicable here, while the digital processing we suggest is adapted to moving targets, and low signal-to-noise ratio (SNR).

Several past works employ CS algorithms to this type of problem, but do not address sample rate reduction and continue sampling at the Nyquist rate [4], [5]. Other ideas combine radar and CS in order to reduce the receiver’s sampling rate, but in doing so impose constraints on the radar transmitter and do not treat noise [6], or do not handle noise well [7]. The work in [7] first estimates target delays and then uses these recovered delays to estimate Doppler frequencies and amplitudes. We refer to this technique as “two-stage recovery” in subsequent sections. Another approach proposes single stage CS recovery methods with dictionary size proportional to the product of delay and Doppler grid sizes, making them infeasible for many realistic scenarios [4], [8].

Our approach is based on the observation that the received radar signal can be modeled with $3L$ degrees of freedom (DOF): a delay, Doppler frequency and amplitude for each of the L targets. Signals which can be described with a fixed number of DOF per unit of time are known as Finite Rate of Innovation (FRI) [9] signals. The proposed recovery process estimates these DOF from low rate samples. The concept of FRI together with the Xampling methodology enables sub-Nyquist rates using practical hardware [1], [10].

At the crux of our proposed recovery method is a coherent superposition of time shifted and modulated pulses, the Doppler focusing function $\Phi(t; \nu)$. For any Doppler frequency ν , this function combines the received signals from different pulses so that targets with appropriate Doppler frequencies come together in phase. For each sought after ν , $\Phi(t; \nu)$ is processed

Manuscript received September 22, 2013; revised January 04, 2014; accepted January 14, 2014. Date of publication February 06, 2014; date of current version March 10, 2014. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Antonio De Maio.

The authors are with the Department of Electrical Engineering, Technion-Israel Institute of Technology, Haifa 32000, Israel (e-mail: omerba@tx.technion.ac.il; yonina@ee.technion.ac.il).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSP.2014.2304917

as a simple one-dimensional CS problem and the appropriate time delays are recovered. The gain from this method is both in terms of SNR and Doppler resolution. For P pulses adding coherently, we obtain a factor P SNR improvement over white noise (which adds incoherently, i.e., in power), as will be proved in Section VI. Such an SNR improvement is optimal as it is the improvement obtained by a matched filter (MF). This result is established in Section III. We also analyze the minimal number of samples required for perfect recovery without noise, and show that Doppler focusing attains this lower bound in terms of number of samples per pulse. In addition, denoting the PRI as τ , we demonstrate in Section IV that the width of the Doppler focus for each $\Phi(t; \nu)$ is $2\pi/P\tau$, meaning that delays of targets separated in Doppler by more than $2\pi/P\tau$ create almost no interference with each other.

The idea of Doppler focusing comes from a similar function used in the context of ultrasound beamforming in [11], [12]. There, in a method named ‘‘Dynamic focusing’’, the signal returned to a set of linearly aligned transceivers is focused in a manner similar to how we focus pulses, where the Doppler frequency ν is replaced by a spatial direction θ . In both cases, the advantages of focusing are retained with sub-Nyquist processing since it can be performed on the low rate Xamples.

Simulations provided in Section VIII show that when sampling at one tenth the Nyquist rate, our Doppler focusing recovery method outperforms both classic MF recovery, described in Section III, and two-stage CS recovery, described in Section VII. When the SNR reaches -25 dB, our method achieves the performance of classic recovery operating at the full Nyquist rate.

The main merits of our proposed method are as follows:

- 1) **Low rate ADC and DSP**—using Xampling and the proposed recovery method, we are able to acquire the sub-Nyquist samples containing information needed for target recovery, and then digitally recover the unknown target parameters using low rate processing, without returning to the higher Nyquist rate.
- 2) **Transmitter compatibility**—our recovery method does not impose any restrictions on the transmitted signal, provided it meets the assumptions stated in Section II.
- 3) **Scaling with problem size**—many CS delay-Doppler estimation methods depend upon constructing a CS dictionary with a column for each delay-Doppler hypothesis. For even moderate size problems, this requires a huge dictionary, making them infeasible for many systems, especially real-time ones. Our Doppler focusing based method avoids this problem by separating the Doppler from delay recovery, making each CS delay recovery indifferent to the underlying Doppler.
- 4) **Robustness to noise and clutter**—The SNR achieved using Doppler focusing scales linearly with the number of received pulses P , as does an optimal MF, providing good performance in AWGN. Regarding clutter, Doppler focusing includes inherent isolation between signals with different Doppler frequencies. Therefore, unless target and clutter have very similar Doppler frequency, clutter interference can be significantly reduced.

The remainder of this paper is organized as follows. In Section II we describe the radar model and the assumptions

we use for its simplification. Section III reviews classic MF processing. We explain the Doppler focusing concept in Section IV and sub-Nyquist delay recovery in Section V. The delay-Doppler recovery method using Doppler focusing is described in Section VI, along with an analysis of noiseless recovery and some practical considerations. We provide a comparison with other CS recovery methods in Section VII. Numerical results are presented in Section VIII.

We denote vectors by boldface lower case letters, e.g., \mathbf{c} , and matrices by boldface capital letters, e.g., \mathbf{A} . We say a vector \mathbf{c} is L -sparse if $\|\mathbf{c}\|_0 \leq L$, i.e., at most L of its elements are nonzero. The n th component of a vector is written as c_n , and the ij th element of a matrix is denoted by \mathbf{A}_{ij} . Non-boldface variables represent scalars or functions, where continuous functions are denoted with round brackets, e.g., $x(t)$ and discrete functions with square brackets, e.g., $h[n]$. The cardinality of a set κ is denoted by $|\kappa|$.

II. RADAR MODEL

We consider a radar transceiver that transmits a pulse train

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

consisting of P equally spaced pulses $h(t)$. The pulse-to-pulse delay τ is referred to as the PRI, and its reciprocal $1/\tau$ is the PRF. The entire span of the signal in (1) is called the coherent processing interval (CPI). The pulse $h(t)$ is a known time-limited baseband function with continuous-time Fourier transform (CTFT) $H(\omega) = \int_{-\infty}^{\infty} h(t)e^{-j\omega t} dt$. We assume that $H(\omega)$ has negligible energy at frequencies beyond $B_h/2$ and we refer to B_h as the bandwidth of $h(t)$.

The target scene is composed of L non-fluctuating point targets (Swirling-0 model, see [13]), where we assume that L is known, although this assumption can easily be relaxed. The pulses reflect off the L targets and propagate back to the transceiver. Each target ℓ is defined by three parameters: a time delay τ_ℓ , proportional to the target’s distance from the radar; a Doppler radial frequency ν_ℓ , proportional to the target-radar closing velocity; and a complex amplitude α_ℓ , proportional to the target’s radar cross section (RCS), dispersion attenuation and all other propagation factors. We limit ourselves to defining targets in the radar’s radial coordinate system.

Throughout, we make the following assumptions on the targets’ location and motion, which leads to a simplified expression for the received signal. To this end we introduce a few auxiliary parameters. Denote the radar’s carrier frequency as f_c and the speed of light as c . We use the time-distance equivalence $r_\ell = c\tau_\ell/2$ and the non-relativistic Doppler radial frequency-velocity equivalence $\dot{r}_\ell = \nu_\ell c/4\pi f_c$.

A1. ‘‘Far targets’’—assuming $\nu_\ell \ll 2\pi f_c \tau_\ell / P\tau$, target-radar distance is large compared to the distance change during the CPI which allows for constant α_ℓ .

A2. ‘‘Slow targets’’—assuming $\nu_\ell \ll 2\pi f_c / P\tau B_h$, target velocity is small enough to allow for constant τ_ℓ during the CPI.

A3. “Small acceleration”—assuming $\ddot{r}_\ell \ll c/2f_c(P\tau)^2$, target velocity remains approximately constant during the CPI allowing for constant ν_ℓ .

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. For example, consider a $P = 100$ pulse radar system with PRI $\tau = 100 \mu\text{sec}$, bandwidth $B_h = 30 \text{ MHz}$ and carrier frequency $f_c = 3 \text{ GHz}$, tracking cars traveling up to 120 km/hour. The maximal distance change over the CPI is approximately 0.33 m, so if the targets’ minimal distance from the radar is more than a few meters, then **A1.** is satisfied. As for **A2.**, the maximal Doppler frequency is approximately 667 Hz, which is much smaller than $f_c/P\tau B_h = 10 \text{ KHz}$. An extreme acceleration of 10 m/sec² would cause a velocity change of 0.1 m/sec over the CPI, easily satisfying **A3.** As for airborne targets, care must be taken to ensure compliance.

Based on these three assumptions, we can write the received signal as

$$x(t) = \sum_{p=0}^{P-1} \sum_{\ell=0}^{L-1} \alpha_\ell h(t - \tau_\ell - p\tau) e^{-j\nu_\ell p\tau}. \quad (2)$$

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

$$x(t) = \sum_{p=0}^{P-1} x_p(t) \quad (3)$$

where

$$x_p(t) = \sum_{\ell=0}^{L-1} \alpha_\ell h(t - \tau_\ell - p\tau) e^{-j\nu_\ell p\tau}. \quad (4)$$

In reality $x(t)$ will be contaminated by AWGN. The main source of this noise is receiver thermal noise, but it also represents other wideband analog distortions and imperfections common to radio-frequency (RF) hardware. Another important source of distortion is clutter, originating from large objects in the targets’ vicinity reflecting the transmitted signal. We take both into account in our analysis in Section VI and simulations in Section VIII.

Since the transmitted signal (1) is a finite periodic pulse train, it is invariant to the transformation $\tau_\ell \rightarrow \tau_\ell + k_1\tau$, $\nu_\ell \rightarrow \nu_\ell + 2\pi k_2/\tau$ where $k_1, k_2 \in \mathbb{Z}$, except on its boundaries. Therefore the radar’s unambiguous time-frequency region, where it can resolve targets with no ambiguity, is $[0, \tau] \times [-\pi/\tau, \pi/\tau]$ respectively. Targets outside this region will be measured in delay and Doppler frequency modulo τ and $2\pi/\tau$ accordingly, as with any fixed-PRI pulse radar. To resolve delay or Doppler ambiguity, information between several CPIs must be shared, for example by employing the Chinese remainder theorem [14]. We make the following further assumptions on targets’ delay and Doppler:

A4. No time ambiguity: $\{\tau_\ell \in I \subset [0, \tau)\}_{\ell=0}^{L-1}$ where I is a continuous time interval in $[0, \tau)$, so that $x_p(t) = 0, \forall t \notin [p\tau, (p+1)\tau]$.

A5. No Doppler ambiguity: $\{\nu_\ell \in [-\pi/\tau, \pi/\tau)\}_{\ell=0}^{L-1}$.

A6. The pairs in the set $\{\tau_\ell, \nu_\ell\}_{\ell=0}^{L-1}$ are unique.

Our goal in this work is to accurately detect the L targets, i.e., to estimate the $3L$ DOF $\{\alpha_\ell, \tau_\ell, \nu_\ell\}_{\ell=0}^{L-1}$ in (2), using the least possible number of digital samples.

III. CLASSIC PULSE-DOPPLER PROCESSING

Classic radar processing samples and processes the received signal at its Nyquist rate B_h using a MF [13]. In modern systems the MF operation is performed digitally, and therefore requires an ADC capable of sampling at B_h , which can be hundreds of MHz and even up to several GHz. In order to evaluate our sampling and reconstruction method, we compare it to classic radar processing, which in general consists of the following stages:

- 1) **ADC**—sample each incoming frame $x_p(t)$ at its Nyquist rate B_h , equal to $h(t)$ ’s bandwidth, creating $x_p[n], 0 \leq n < N$, where $N = \tau B_h$.
- 2) **Matched filter**—for each $x_p[n]$, create $y_p[n] = x_p[n] * h^*[-n]$, where $h[n]$ is a sampled version of the transmitted pulse $h(t)$. The time resolution attained in this step is $1/B_h$, corresponding to the width of the autocorrelation of the pulse $h(t)$.
- 3) **Doppler processing**—for each discrete time index n , perform a P -point DFT along the pulse dimension: $z_n[k] = DFT_P\{y_p[n]\} = \sum_{p=0}^{P-1} y_p[n] e^{-j2\pi pk/P}$ for $0 \leq k < P$. The frequency resolution attained in this step is $2\pi/P\tau$, proportional to the inverse of the CPI.
- 4) **Delay-Doppler map**—stacking the vectors \mathbf{z}_n , and taking absolute value, we obtain a delay-Doppler map $\mathbf{Z} = \text{abs}[\mathbf{z}_0 \dots \mathbf{z}_{N-1}] \in \mathbb{R}^{P \times N}$.
- 5) **Peak detection**—a heuristic detection process, where knowledge of number of targets, target power, clutter location, etc. may help discover target positions. For example, if we know there are L targets, then we can choose the L strongest points in the map.

The Doppler processing stage can be viewed as MF in the pulse dimension to a constant radial velocity target. As such, it increases the SNR by P compared to the SNR of a single pulse [15], [16]. Since a MF is the linear time-invariant (LTI) system which maximizes SNR, it follows that a factor P increase is optimal for P pulses. In Section VI we show that the SNR achieved with Doppler focusing also scales linearly with P , while other CS methods either achieve sub-linear SNR increase, or require prohibitive computational cost, as shown in Section VII.

Classic processing requires sampling the received signal at its Nyquist rate B_h , which is inversely proportional to the system’s time resolution. The required computational power is P convolutions of a signal of length $N = \tau B_h$ and N FFTs of length P —both also proportional to B_h . The growing demands for improved estimation accuracy and target separation dictate an ever growing increase in signal’s bandwidth. The goal of this work is to *present some concrete steps towards breaking the link between radar signal bandwidth and sampling rate*, and to allow low rate sampling and processing of radar signals, regardless of their bandwidth, retaining the same SNR scaling.

We achieve this goal by utilizing the combination of ideas of FRI and Xampling. Previous papers have already used these complementary concepts together. The work in [3], [9] creates a mathematical framework for sub-Nyquist sampling of pulse streams and defines lower bounds on the sampling rate

needed for perfect reconstruction. Practical sampling methods achieving these bounds are explained in [2], [11], [17] in the context of ultrasound and radar, both without Doppler. Another work [7] investigates the delay-Doppler estimation problem, but recovers the delays and Doppler frequencies in a two-stage process which achieves an SNR increase which is sub-linear in P . In Section VI we prove that in noisy scenarios, Doppler focusing recovery is superior to the two-stage method, as it achieves an SNR increase equal to a pulse dimension MF. We verify these results via simulations in Section VIII. Combining Xampling and Doppler focusing, our $3L$ DOF input signal (2) enjoys the benefits of sub-Nyquist rate Xampling and accurate digital recovery of the target scene, effectively breaking the link between signal bandwidth and sampling rate.

IV. DOPPLER FOCUSING

We now introduce and explain the main idea in this paper, called Doppler Focusing. This processing technique uses target echoes from different pulses to create a single superimposed pulse. This improves the SNR for robustness against noise and implicitly estimates targets' Doppler frequency in the process. We point out that stages 2) and 3) in classic processing can be viewed as delay and Doppler processing accordingly. Since they are both LTI, they can be interchanged, performing the DFT before MF. In classic processing, applying a MF before DFT decreases computation latency, so most practical systems carry them out in the noted order.

When interchanging steps 2) and 3), the DFT is simply a discrete equivalent of the following time shift and modulation operation on the received signal:

$$\begin{aligned}
 \Phi(t; \nu) &= \sum_{p=0}^{P-1} x_p(t + p\tau) e^{j\nu p\tau} \\
 &= \sum_{p=0}^{P-1} \sum_{\ell=0}^{L-1} \alpha_\ell h(t - \tau_\ell) e^{j(\nu - \nu_\ell)p\tau} \\
 &= \sum_{\ell=0}^{L-1} \alpha_\ell h(t - \tau_\ell) \sum_{p=0}^{P-1} e^{j(\nu - \nu_\ell)p\tau} \\
 &= \sum_{\ell=0}^{L-1} \alpha_\ell h(t - \tau_\ell) g(\nu | \nu_\ell)
 \end{aligned} \quad (5)$$

where we used (4) and denoted $g(\nu | \nu_\ell) = \sum_{p=0}^{P-1} e^{j(\nu - \nu_\ell)p\tau}$.

We now analyze the sum of exponents in (5). For any given ν , targets with Doppler frequency ν_ℓ in a band of width $2\pi/P\tau$ around ν , i.e., in $\Phi(t; \nu)$'s "focus zone", will achieve coherent integration and an SNR boost of approximately

$$|g(\nu | \nu_\ell)| = \left| \sum_{p=0}^{P-1} e^{j(\nu - \nu_\ell)p\tau} \right|_{|\nu - \nu_\ell| \leq \pi/P\tau} \cong P \quad (6)$$

compared with a single pulse. 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 approximately cancel out. Thus $|g(\nu | \nu_\ell)| \cong 0$ for $|\nu - \nu_\ell| > \pi/P\tau$, where **A5**. implies that $|\nu - \nu_\ell| < \pi/\tau$. One can also use windowing in order to improve out of focus attenuation, as described in

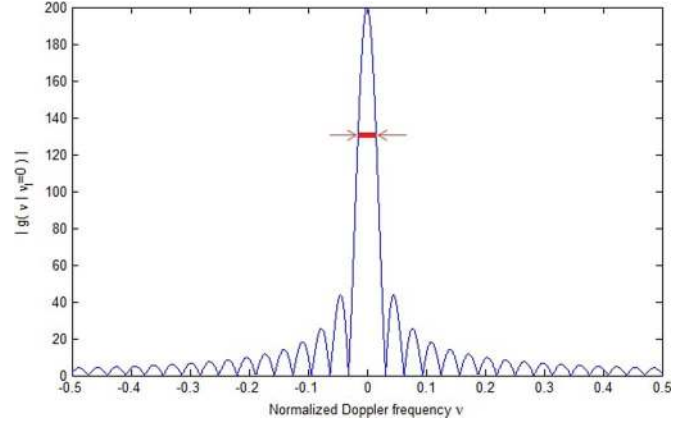


Fig. 1. Example of $|g(\nu | \nu_\ell)|$ for $P = 200$ pulses and $\nu_\ell = 0$. Arrows mark the focus zone, i.e., $|\nu| < \pi/P\tau$. Frequencies outside the focus zone are severely attenuated.




Fig. 2. Schematic delay-Doppler map. Dark area indicates out-of-focus region. Only one target is in focus for current choice of ν .

Section VI. An example of $|g(\nu | \nu_\ell)|$ is depicted in Fig. 1. Therefore, we can approximate (5) by

$$\Phi(t; \nu) \cong P \sum_{\ell \in \Lambda(\nu)} \alpha_\ell h(t - \tau_\ell) \quad (7)$$

where $\Lambda(\nu) = \{\ell : |\nu - \nu_\ell| < \pi/P\tau\}$.

Instead of trying to estimate delay and Doppler jointly, we reduced our problem to delay only estimation for a small range of Doppler frequencies, with increased amplitude for improved performance against noise. To emphasize this, consider the case of trying to detect and estimate parameters of two targets with very closely spaced delays but with different Doppler frequencies (see for example the two helicopters in Fig. 2). Algorithms whose time resolution is coarser than the targets' delay separation are likely to encounter various problems recovering this target scene, the most likely of which is identification of a single target instead of two. With Doppler focusing we achieve an extra dimension of potential separation, regardless of the underlying delay recovery algorithm, enabling improved recovery performance. Fig. 2 illustrates this concept by showing various targets spanning a delay-Doppler region. When focusing at some ν , only targets in ν 's focus zone (white region) come into view, while all other targets (dark region) disappear. In Section VIII we demonstrate this point via simulation.

To enable sub-Nyquist recovery, in the next sections we sample our signal in the time domain but extract frequency domain information. We now show how Doppler focusing can

also be performed in the frequency domain, paving the way towards sub-Nyquist Doppler focusing.

Using (4), and denoting $X_p(\omega)$ as the CTFT of $x_p(t + p\tau)$,

$$X_p(\omega) = H(\omega) \sum_{\ell=0}^{L-1} \alpha_\ell e^{-j\omega\tau_\ell} e^{-j\nu_\ell p\tau}. \quad (8)$$

Taking the CTFT of $\Phi(t; \nu)$ as a function of t we obtain:

$$\begin{aligned} \Psi(\omega; \nu) &= \text{CTFT}(\Phi(t; \nu)) = \sum_{p=0}^{P-1} X_p(\omega) e^{j\nu p\tau} \\ &= H(\omega) \sum_{\ell=0}^{L-1} \alpha_\ell e^{-j\omega\tau_\ell} \sum_{p=0}^{P-1} e^{j(\nu - \nu_\ell)p\tau}. \end{aligned} \quad (9)$$

All $3L$ of the problem's parameters appear in (9). Its structure is that of a delay estimation problem, as shown in Section V, combined with the familiar sum of exponents term from (5).

We have seen that Doppler focusing reduces a delay-Doppler estimation problem to a delay-only estimation problem for a specific Doppler frequency. In the next section we describe delay recovery from sub-Nyquist sampling rates using Xampling. In Section VI we revisit the Doppler focusing concept, combining it with these low rate Xamples to address the delay-Doppler radar problem at sub-Nyquist rates.

V. SUB-NYQUIST DELAY RECOVERY

The problem of recovering the $2L$ amplitudes and delays in

$$\phi(t) = \sum_{\ell=0}^{L-1} \alpha_\ell h(t - \tau_\ell), \quad 0 \leq t < \tau \quad (10)$$

from sub-Nyquist samples has been previously studied in [2], [3], [9], [17]. Since Doppler focusing yields such a problem, we now review how Xampling can be used to solve (10) at a sub-Nyquist sampling rate.

A. Xampling

The concept of Xampling, introduced in [1], [18], [19], describes analog-to-digital conversion which acquires samples at sub-Nyquist rates while preserving the ability to perfectly reconstruct the signal. Xampling can be interpreted as ‘‘compressed sampling’’, in the sense that we are performing data compression inherently in the sampling stage. To do this, we do not simply reduce sampling rate, since this is bound to cause loss of information. Instead, we perform an analog prefiltering operation on our signal and only then sample it, in order to extract the required information for recovery. We now show how the signal's Fourier series coefficients are related to the problem's unknown parameters [3], [7], [9], [17]. We then describe how to acquire these coefficients via Xampling.

Since $\phi(t)$ is confined to the interval $t \in [0, \tau]$, it can be expressed by its Fourier series

$$\phi(t) = \sum_{k \in \mathbb{Z}} c[k] e^{j2\pi kt/\tau}, \quad t \in [0, \tau], \quad (11)$$




Fig. 3. Multichannel direct sampling of the Fourier series coefficients, from [3].

where

$$\begin{aligned} c[k] &= \frac{1}{\tau} \int_0^\tau \phi(t) e^{-j2\pi kt/\tau} dt \\ &= \frac{1}{\tau} \sum_{\ell=0}^{L-1} \alpha_\ell \int_0^\tau h(t - \tau_\ell) e^{-j2\pi kt/\tau} dt \\ &= \frac{1}{\tau} H(2\pi k/\tau) \sum_{\ell=0}^{L-1} \alpha_\ell e^{-j2\pi k\tau_\ell/\tau}. \end{aligned} \quad (12)$$

From (12) we see that the unknown parameters $\{\alpha_\ell, \tau_\ell\}_{\ell=0}^{L-1}$ are embodied in the Fourier coefficients $c[k]$ in the form of a complex sinusoid problem. For these problems, without noise, $2L$ samples are enough to recover the unknown α 's and τ 's [9], i.e., $|\kappa| \geq 2L$. These parameters can be found using spectral analysis methods which require sampling a consecutive subset of coefficients, such as the annihilating filter [20], matrix pencil [21], or ESPRIT [22]. An alternative approach is to use MUSIC [23], which does not require consecutive coefficients. The lower bound on $|\kappa|$ can be achieved only when the noise is negligible and computational complexity is not of concern. When there is substantial noise in the problem, having more than $2L$ coefficients improves recovery.

Our signals exist in the time domain, and therefore we do not have direct access to $c[k]$. To this end, we can use the direct multichannel sampling scheme described in [3] in order to obtain any arbitrary set of Fourier series coefficients, as illustrated in Fig. 3. The analog input $x(t)$ is split into $|\kappa|$ channels, where in each channel k it is mixed with the harmonic signal $e^{-j2\pi kt/\tau}$, integrated over the PRI duration, and then sampled.

In [2] an actual radar system was built using a similar yet more practical technique, where a set of mixers, band-pass filters and low rate ADCs sampled different spectral bands of the signal. The matching Fourier coefficients were then created digitally. The radar model there included delay only without Doppler. An alternative Xampling method uses the Sum of Sincs filter described in [17]. All of these methods can be used to obtain arbitrary Fourier series coefficients. The number of Fourier coefficients extracted per pulse is a design parameter which controls the tradeoff between sampling rate and recovery performance. In our numerical experiments, presented in Section VIII, we demonstrate a Xampling rate of one tenth of the Nyquist rate.

B. Compressed Sensing Recovery

In the previous subsection we showed that $2L$ samples are enough to solve (10), when there is no noise. We now describe a CS-based delay recovery method, operating on the Xamples $c[k]$, which is often more robust to noise.

Assume the time delays are aligned to a grid

$$\tau_\ell = n_\ell \Delta_\tau, \quad 0 \leq n_\ell < N_\tau \quad (13)$$

where we choose Δ_τ so that $N_\tau = \tau/\Delta_\tau$ is an integer. Since in real scenarios target delays are not necessarily aligned to a grid, local interpolation around detected grid points can be used to reduce quantization errors. In our simulations in Section VIII we perform a parabolic fit around detected indices. The analysis in [24] can be used to quantify these off-the-grid errors. Choosing a set of indices $\kappa = \{k_0, \dots, k_{|\kappa|-1}\}$, we define the corresponding vector of Fourier coefficients

$$\mathbf{c} = [c[k_0] \dots c[k_{|\kappa|-1}]]^T \in \mathbb{C}^{|\kappa|}. \quad (14)$$

We can then write (12) in vector form as

$$\mathbf{c} = \frac{1}{\tau} \mathbf{H} \mathbf{V} \mathbf{x} \quad (15)$$

where \mathbf{H} is a $|\kappa| \times |\kappa|$ diagonal matrix with diagonal elements $H(2\pi k_i/\tau)$ and \mathbf{V} is a $|\kappa| \times N_\tau$ Vandermonde matrix with elements $\mathbf{V}_{mq} = e^{-j2\pi k_m n/N_\tau}$, i.e., it is composed of $|\kappa|$ rows of the $N_\tau \times N_\tau$ DFT matrix. The target delay vector $\mathbf{x} \in \mathbb{C}^{N_\tau}$ is L -sparse, where each index n contains the amplitude of a target with delay $n\Delta_\tau$ if it exists, or zero otherwise. Defining the CS dictionary $\mathbf{A} = \frac{1}{\tau} \mathbf{H} \mathbf{V} \in \mathbb{C}^{|\kappa| \times N_\tau}$ we obtain the CS equation

$$\mathbf{c} = \mathbf{A} \mathbf{x}. \quad (16)$$

Estimating targets' delays can be carried out by solving (16) and finding \mathbf{x} 's support - any nonzero index n denotes a target with delay $n\Delta_\tau$.

Given a set of sampled Fourier coefficients, a variety of CS techniques can be employed for recovery [11], for instance orthogonal matching pursuit (OMP) [25], iterative hard thresholding (IHT) [26], or ℓ_1 minimization (see [27] and references within). Choosing the coefficients at random produces favorable conditions for CS, aiding recovery in the presence of noise. When the indices in κ are selected uniformly at random, it can be shown that if $|\kappa| \geq cL(\log N_\tau)^4$, for some positive constant c , then \mathbf{A} obeys the desired Restricted Isometry Property (RIP) with large probability [28]. By satisfying the condition for RIP we are able to recover \mathbf{x} , using a CS recovery algorithm.

VI. DELAY-DOPPLER RECOVERY USING DOPPLER FOCUSING

In Section IV we introduced the concept of Doppler focusing, and in Section V we reviewed how to Xample and recover the $2L$ unknowns of the delay estimation problem (10). We now return to our original delay-Doppler problem (2).

We begin by showing how Xampling can be performed on the multi pulse signal (2). We then describe Doppler focusing, and analyze two aspects of the algorithm: the effect of multiple pulses on SNR when noise exists, and the minimal number of samples required for perfect recovery without noise. Finally we discuss some practical considerations.

A. Xampling

We can apply the Xampling technique of Section V to each of the pulses $x_p(t)$ of the multi-pulse signal (2) to obtain $c_p[k]$. Since $x_p(t)$ is confined to the interval $t \in [p\tau, (p+1)\tau]$, we replace $t \rightarrow t + p\tau$ and $\alpha_\ell \rightarrow \alpha_\ell e^{-j\nu_\ell p\tau}$ in (12) resulting in

$$c_p[k] = \frac{1}{\tau} H(2\pi k/\tau) \sum_{\ell=0}^{L-1} \alpha_\ell e^{-j\nu_\ell p\tau} e^{-j2\pi k\tau_\ell/\tau}, \quad (17)$$

where we used the fact that since both $k, p \in \mathbb{Z}$ we have $e^{-j2\pi kp} \equiv 1$. If we assume the time delays are aligned to a grid (13), similarly to (15), then for each pulse p we obtain

$$\mathbf{c}_p = \frac{1}{\tau} \mathbf{H} \mathbf{V} \mathbf{x}_p. \quad (18)$$

From (17) we see that all $3L$ unknown parameters $\{\alpha_\ell, \tau_\ell, \nu_\ell\}_{\ell=0}^{L-1}$ are embodied in the Fourier coefficients $c_p[k]$ in the form of a complex sinusoid problem. The number of Fourier coefficients sampled in each pulse, $|\kappa|$, controls the trade-off between sample rate and robustness to noise.

B. Applying Doppler Focusing and CS Recovery

Having acquired $c_p[k]$ using Xampling, we now perform the Doppler focusing operation for a specific frequency ν :

$$\begin{aligned} \Psi_\nu[k] &= \sum_{p=0}^{P-1} c_p[k] e^{j\nu p\tau} \\ &= \frac{1}{\tau} H(2\pi k/\tau) \sum_{\ell=0}^{L-1} \alpha_\ell e^{-j2\pi k\tau_\ell/\tau} \sum_{p=0}^{P-1} e^{j(\nu - \nu_\ell)p\tau}. \end{aligned} \quad (19)$$

From (9) we see that $\Psi_\nu[k] = \tau \Psi(\omega; \nu)|_{\omega=2\pi k/\tau}$.

Following the same arguments as in (6), for any target ℓ satisfying $|\nu - \nu_\ell| < \pi/P\tau$ we have

$$\sum_{p=0}^{P-1} e^{j(\nu - \nu_\ell)p\tau} \cong P. \quad (20)$$

Therefore, Doppler focusing can be performed on the low rate sub-Nyquist samples:

$$\Psi_\nu[k] \cong \frac{P}{\tau} H(2\pi k/\tau) \sum_{\ell \in \Lambda(\nu)} \alpha_\ell e^{-j2\pi k\tau_\ell/\tau}. \quad (21)$$

Equation (21) is identical in form to (12) except it is scaled by P , increasing SNR for improved performance with noise. Furthermore, we reduced the number of active delays.

For each ν we now have a delay estimation problem, which can be solved using MUSIC, annihilating filter, or any other spectral analysis method. Furthermore, (21) can be written in vector form using the same notations of Section V as

$$\Psi_\nu = \frac{P}{\tau} \mathbf{H} \mathbf{V} \mathbf{x}_\nu \quad (22)$$

where \mathbf{x}_ν is L -sparse and

$$\Psi_\nu = [\Psi_\nu[k_0] \dots \Psi_\nu[k_{|\kappa|-1}]]^T \in \mathbb{C}^{|\kappa|}. \quad (23)$$

This is exactly the delay estimation problem we have already shown how to solve in Section V, with one important difference.

In Section V all L targets appear in the same delay estimation problem. Here, since targets are spread out over the entire set of Doppler frequencies, we do not immediately treat each \mathbf{x}_ν in (22) as having L valid entries. Instead, we perform matching pursuit over all Doppler frequencies, i.e., we jointly process (22) for all values of ν , recovering a single delay in each iteration, and removing its influence from the Xamples. Algorithm 1 describes this process in detail.

Algorithm 1: Doppler Focusing

Input: Xamples $C = \{c_p[k]\}_{0 \leq p < P, k \in \kappa}$, number of targets L

Output: Estimated target parameters $\{\hat{\alpha}_\ell, \hat{\tau}_\ell, \hat{\nu}_\ell\}_{\ell=0}^{L-1}$

for $\ell = 0$ to $L - 1$ **do**

 % For any ν , C can be focused on ν using (19) to

 % create Ψ_ν , from which \mathbf{x}_ν can be recovered.

$(\hat{\ell}, \hat{\nu}_\ell) \leftarrow \arg \max_{\substack{0 \leq \ell < L \\ -\pi/\tau \leq \nu < \pi/\tau}} |\mathbf{x}_\nu[l]|$ using (22)

$\hat{\alpha}_\ell \leftarrow \mathbf{x}_{\hat{\nu}_\ell}[\hat{\ell}]$

$\hat{\tau}_\ell \leftarrow$ delay associated with $\mathbf{x}_{\hat{\nu}_\ell}[\hat{\ell}]$

for $k \in \kappa$ and $-\pi/\tau \leq \nu < \pi/\tau$ **do**

$\Psi_\nu[k] \leftarrow \Psi_\nu[k] - \frac{1}{\tau} H(2\pi k/\tau) \hat{\alpha}_\ell e^{-j2\pi k \hat{\tau}_\ell/\tau} \sum_{p=0}^{P-1} e^{j(\nu - \hat{\nu}_\ell)p\tau}$

end for

end for

The Doppler focusing operation (19) 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 which are undetectable, as occurs with classic Moving Target Indication (MTI) [15]. Define the set of Fourier coefficients $C = \{c_p[k]\}_{0 \leq p < P, k \in \kappa}$, and Ψ_ν as the vector of focused coefficients (23) obtained from C using (19). It follows that \mathbf{x}_ν can be recovered from Ψ_ν for any ν . Since strong amplitudes are indicative of true target existence as opposed to noise, Doppler focusing recovery searches for large (in magnitude) entries of \mathbf{x}_ν 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. A similar subtraction is performed in many iterative algorithms such as OMP or the Clean Process of [29]. Detection is performed iteratively until all targets have been detected, if L is known, or until an amplitude threshold is met, if the model order is unknown. The latter case has been studied extensively [30]–[32] as a problem of estimating the number of sinusoids in a noisy sequence. In our simulations, since we wish to eliminate model order errors which influence recovery performance, we assume L is known.

It is important to note that our dictionary is indifferent to the Doppler estimation. CS methods which estimate delay and Doppler simultaneously [4], [8], require a dictionary which

grows with the number of pulses. Here by separating delay and Doppler estimation, the CS dictionary is not a function of P .

Since \mathbf{V} is a partial Fourier matrix, the problem defined in (22), after normalizing by \mathbf{H}^{-1} , becomes a problem of recovering frequencies from a sum of complex exponentials. Many methods exist for solving such a spectral analysis problem (see [20] for a review), and they can be used instead of CS. These methods offer different combinations of robustness to noise, minimal sample rate, and sensitivity to grid errors. CS is our method of choice when we are interested in low-SNR scenarios. However, since Doppler focusing is independent of the underlying delay estimation, in different scenarios CS can be exchanged for alternative delay recovery methods. For example, in the upcoming noiseless recovery subsection, when noise is not a concern, we use the annihilating filter approach.

C. SNR Analysis

To analyze the effect of Doppler focusing on SNR, we add noise to (2):

$$\tilde{x}(t) = x(t) + w(t), \quad (24)$$

where $w(t)$ is a zero mean wide-sense stationary random signal with autocorrelation $r_w(s) = \sigma^2 \delta(s)$. The Fourier coefficients in (17) then become

$$\tilde{c}_p[k] = c_p[k] + w_p[k], \quad (25)$$

where

$$w_p[k] = \frac{1}{\tau} \int_{p\tau}^{(p+1)\tau} w(t) e^{-j2\pi k t/\tau} dt \quad (26)$$

is a zero mean complex random variable with variance

$$\begin{aligned} E[|w_p[k]|^2] &= \frac{1}{\tau^2} \int_{p\tau}^{(p+1)\tau} dt \int_{p\tau}^{(p+1)\tau} E[w(t)w^*(t')] e^{-j2\pi k(t-t')/\tau} dt' \\ &= \frac{1}{\tau^2} \int_{p\tau}^{(p+1)\tau} dt \int_{p\tau}^{(p+1)\tau} \sigma^2 \delta(t-t') e^{-j2\pi k(t-t')/\tau} dt' \\ &= \frac{1}{\tau^2} \sigma^2 \int_{p\tau}^{(p+1)\tau} dt = \sigma^2/\tau. \end{aligned} \quad (27)$$

We can write $\tilde{c}_p[k]$ as the disjoint effect of L targets and noise

$$\tilde{c}_p[k] = \sum_{\ell=0}^{L-1} c_p^\ell[k] + w_p[k] \quad (28)$$

where $c_p^\ell[k] = \frac{1}{\tau} H(2\pi k/\tau) \alpha_\ell e^{-j\nu_\ell p\tau} e^{-j2\pi k \tau_\ell/\tau}$. The SNR of target ℓ in $\tilde{c}_p[k]$ is then the power ratio

$$\begin{aligned} \Gamma_p^\ell[k] &= \frac{|c_p^\ell[k]|^2}{E[|w_p[k]|^2]} \\ &= \frac{|\frac{1}{\tau} H(2\pi k/\tau) \alpha_\ell e^{-j\nu_\ell p\tau} e^{-j2\pi k \tau_\ell/\tau}|^2}{\sigma^2/\tau} \\ &= |H(2\pi k/\tau)|^2 \frac{|\alpha_\ell|^2}{\tau \sigma^2}. \end{aligned} \quad (29)$$

Analyzing the SNR in the focused Fourier coefficients using (19) we obtain

$$\begin{aligned}\tilde{\Psi}_\nu[k] &= \sum_{p=0}^{P-1} \tilde{c}_p[k] e^{j\nu p\tau} \\ &= \sum_{p=0}^{P-1} (c_p[k] + w_p[k]) e^{j\nu p\tau} \\ &= \Psi_\nu[k] + w_\nu[k]\end{aligned}\quad (30)$$

where $w_\nu[k] = \sum_{p=0}^{P-1} w_p[k] e^{j\nu p\tau}$. Using (21) we have

$$\begin{aligned}\Psi_\nu[k] + w_\nu[k] &\cong \sum_{\ell \in \Lambda(\nu)} \frac{P}{\tau} H(2\pi k/\tau) \alpha_\ell e^{-j2\pi k\tau\ell/\tau} + w_\nu[k] \\ &= \sum_{\ell \in \Lambda(\nu)} \Psi_\nu^\ell[k] + w_\nu[k],\end{aligned}\quad (31)$$

where $\Psi_\nu^\ell[k] = \frac{P}{\tau} H(2\pi k/\tau) \alpha_\ell e^{-j2\pi k\tau\ell/\tau}$. We then define the focused SNR for target ℓ as

$$\Xi_\nu^\ell[k] = \frac{|\Psi_\nu^\ell[k]|^2}{E[|w_\nu[k]|^2]}.\quad (32)$$

Due to the independence of $w(t)$ in different time intervals, $w_p[k]$ is a discrete time white random sequence with respect to p , so that $w_\nu[k]$ is a sum of P independent random variables. Therefore, using (27), we have

$$E[|w_\nu[k]|^2] = \sum_{p=0}^{P-1} |e^{j\nu p\tau}|^2 E[|w_p[k]|^2] = P\sigma^2/\tau.\quad (33)$$

Substituting (33) into (32) we obtain

$$\begin{aligned}\Xi_\nu^\ell[k] &= \frac{\left| \frac{P}{\tau} H(2\pi k/\tau) \alpha_\ell e^{-j2\pi k\tau\ell/\tau} \right|^2}{P\sigma^2/\tau} \\ &= P |H(2\pi k/\tau)|^2 \frac{|\alpha_\ell|^2}{\tau\sigma^2} = P\Gamma_p^\ell[k].\end{aligned}\quad (34)$$

It is evident that the focused SNR (34) is P times greater than before Doppler focusing (29). We have thus obtained a linear SNR improvement with increasing number of pulses, as does an optimal MF.

D. Noiseless Recovery

The following theorems analyze the minimal number of samples required for perfect recovery when there is no noise.

Theorem 1: The minimal number of samples required for perfect recovery of L targets when there is no noise, is at least $4L^2$, with $|\kappa|$ and P at least $2L$ each.

Proof: Recall (17) and denote normalized delay $q_\ell = 2\pi\tau\ell/\tau$ and normalized Doppler frequency $g_\ell = \nu_\ell\tau$, to obtain a more symmetric form

$$c_p[k] = \frac{1}{\tau} H(2\pi k/\tau) \sum_{\ell=0}^{L-1} \alpha_\ell e^{-j(g_\ell p + q_\ell k)}.\quad (35)$$

Since there are no constraints on either delay or Doppler frequency, let us examine the case where all targets have identical Doppler $g_\ell = g$:

$$c_p[k] = \frac{1}{\tau} H(2\pi k/\tau) e^{-jgp} \sum_{\ell=0}^{L-1} \alpha_\ell e^{-jq_\ell k}.\quad (36)$$

If it were possible to solve (36) with less than $|\kappa| = 2L$ samples, for example by utilizing information from different values of p , then we could use this to bring (12) to the form of (36) by multiplying $c[k]$ by arbitrary values of e^{-jgp} . Thus we could solve (10) with less than $2L$ samples, in contradiction with [9]. Therefore $|\kappa| \geq 2L$.

Inspecting the case where all targets have the same delay $q_\ell = q$, we obtain:

$$c_p[k] = \frac{1}{\tau} H(2\pi k/\tau) e^{-jqk} \sum_{\ell=0}^{L-1} \alpha_\ell e^{-jg_\ell p}.\quad (37)$$

Applying the same logic which deduced $|\kappa| \geq 2L$ from (36), we infer $P \geq 2L$ from (37). \square

Theorem 2: Suppose target Doppler frequencies are aligned to a grid $\{\tilde{\nu}_m = 2\pi m/\tau M\}_{m=-M/2}^{M/2-1}$, with no restriction on target delays. Then the minimal number of samples required for perfect recovery of L targets when there is no noise, is $2L \min(M, 2L)$.

Proof: For any $k \in \kappa$ we use (17) to write $c_p[k] = \frac{1}{\tau} H(2\pi k/\tau) \sum_{\ell=0}^{L-1} \beta_\ell^k e^{-j\nu_\ell p\tau}$ where $\beta_\ell^k = \alpha_\ell e^{-j2\pi k\tau\ell/\tau}$. We obtain a standard CS problem by writing $c[k] = \frac{1}{\tau} \mathbf{H} \mathbf{F} \boldsymbol{\beta}[k]$, where \mathbf{H} is as defined in (15), $\mathbf{F} \in \mathbb{C}^{P \times M}$ has elements $\mathbf{F}_{pm} = e^{-j2\pi mp/M}$, and $\boldsymbol{\beta}[k] \in \mathbb{C}^M$ is an L -sparse vector.

If $M \geq 2L$ then there is no gain compared with the continuous setting of Theorem 1, and the minimal number of samples remains $4L^2$.

On the other hand, if $M < 2L$, then for $P \geq M$ the CS system is overdetermined and can be solved with the pseudoinverse $\boldsymbol{\beta}[k] = (\mathbf{A}^H \mathbf{A})^{-1} \mathbf{A}^H c[k]$, where $\mathbf{A} = \frac{1}{\tau} \mathbf{H} \mathbf{F}$. For different values of k , each element m of $\boldsymbol{\beta}[k]$ describes a spectral analysis problem for Doppler frequency $2\pi m/\tau M$, containing no more than L harmonics. Thus, we require $|\kappa| \geq 2L$ to complete the recovery. \square

Theorem 3: Under the conditions of Theorem 2, the minimal number of samples required for perfect recovery of L targets using Doppler focusing is $2LM$, with $|\kappa| \geq 2L$ and $P \geq M$.

Proof: In this setting target Doppler frequencies are aligned to a grid $\{\tilde{\nu}_m = 2\pi m/\tau M\}_{m=-M/2}^{M/2-1}$. If we perform Doppler focusing on the same grid, then the analysis of Section IV, and specifically (6), (7) are exact:

$$\begin{aligned}g(\tilde{\nu}_m | \nu_\ell) &= \sum_{p=0}^{M-1} e^{j(\tilde{\nu}_m - \nu_\ell)p\tau} \\ &= \sum_{p=0}^{M-1} e^{j(2\pi m/\tau M - 2\pi m_\ell/\tau M)p\tau} \\ &= \sum_{p=0}^{M-1} e^{j2\pi(m - m_\ell)pM} \\ &= \begin{cases} M, & \text{if } m = m_\ell \\ 0, & \text{otherwise} \end{cases}\end{aligned}\quad (38)$$

where $m_\ell \in [-M/2, \dots, M/2 - 1]$ satisfies $\nu_\ell = \tilde{\nu}_{m_\ell}$.

Each target ℓ appears in exactly one delay estimation problem, for m_ℓ satisfying $\nu_\ell = \tilde{\nu}_{m_\ell}$, and completely cancels out for the other $M - 1$ problems. The resulting set of focused Fourier coefficients can be concatenated and written as

$$\begin{bmatrix} \Psi_{-M/2} \\ \vdots \\ \Psi_{M/2+1} \end{bmatrix} = \frac{P}{\tau} \begin{bmatrix} \mathbf{H}\mathbf{V} & & \\ & \ddots & \\ & & \mathbf{H}\mathbf{V} \end{bmatrix} \begin{bmatrix} \mathbf{x}_{-M/2} \\ \vdots \\ \mathbf{x}_{M/2+1} \end{bmatrix} \quad (39)$$

where we used (22). Since $[\mathbf{x}_{-M/2} \dots \mathbf{x}_{M/2+1}]^T$ is L -sparse, the dictionary $\mathbf{Y} = \text{diag}\{\mathbf{H}\mathbf{V}\}$ must have *spark* greater than $2L$ for perfect recovery [33]. The spark of a block-diagonal matrix equals the maximal spark of any single block, so $\text{spark}(\mathbf{Y}) = \text{spark}(\mathbf{H}\mathbf{V})$. Since $\mathbf{H}\mathbf{V}$ has $|\kappa|$ rows, its spark is upper bounded by $|\kappa| + 1$, and thus $|\kappa|$, the minimal number of Fourier coefficients required per pulse, must be at least $2L$.

Having shown $|\kappa| \geq 2L$, we now show how to recover the L unknown target positions using Doppler focusing. We choose a consecutive set κ so that $H(2\pi k/\tau) \neq 0$ for all $k \in \kappa$. Next we perform Algorithm 1, and solve the resulting delay estimation problems in each iteration using an annihilating filter. Each delay estimation problem (22), after normalizing by \mathbf{H}^{-1} , becomes a spectral analysis problem with $|\kappa|$ samples. When there is no noise, the annihilating filter with $2L$ samples can perfectly recover up to L frequencies [9]. This guarantees successful recovery of all L targets, regardless of their distribution in different Doppler frequencies. Thus, the total number of samples required for perfect recovery using Doppler focusing comes to $|\kappa|P = 2LM$. \square

The minimal rate requirement exists separately on the number of sampled Fourier coefficients $|\kappa|$ and the number of sampled pulses P , and not for their product. This shows that in terms of minimal sampling rate, samples in the coefficient dimension k cannot be replaced by samples in the pulse dimension p , and vice versa.

These theorems show that the requirement of Doppler focusing for $|\kappa| \geq 2L$ matches the general lower bound on the number of samples required in each pulse. Furthermore, when $M = O(L)$, the number of pulses required for Doppler focusing is within order of magnitude of the lower bound. Finally, the result in Theorem 1 coincides with the minimal sampling rate for two dimensional spectral analysis [34].

E. Practical Considerations

We now describe a few practical issues, starting with computational efficiency.

If one wishes to perform Doppler focusing on a uniform grid of M Doppler frequencies, i.e., $\{\tilde{\nu}_m = 2\pi m/\tau M\}_{m=-M/2}^{M/2-1}$, then $\Psi_\nu[k]$ can be computed efficiently using a length M DFT or FFT of a series of length P :

$$\begin{aligned} \Psi_m[k] &\triangleq \Psi_{\tilde{\nu}_m}[k] = \sum_{p=0}^{P-1} c_p[k] e^{j\tilde{\nu}_m p \tau} \\ &= \sum_{p=0}^{P-1} c_p[k] e^{j2\pi m p / M} = \text{DFT}_M\{c_p[k]\}. \end{aligned} \quad (40)$$

Furthermore, if we assume there is only one target in each Doppler frequency, then the delay estimation problem in the $\arg \max$ step of Algorithm 1 can be simplified to a correlator, as described in Algorithm 2.

Algorithm 2: Doppler Focusing with Grid

Input: Samples $C = \{c_p[k]\}_{0 \leq p < P}^{k \in \kappa}$, number of targets L , Doppler grid size M

Output: Estimated target parameters $\{\hat{\alpha}_\ell, \hat{\tau}_\ell, \hat{\nu}_\ell\}_{\ell=0}^{L-1}$

Create $\{\Psi_m\}_{m=-M/2}^{M/2-1}$ from C using FFT (40).

for $\ell = 0$ to $L - 1$ **do**

$$(\hat{\tau}_\ell, \hat{m}_\ell) \leftarrow \arg \max_{\substack{0 \leq t < \tau \\ -M/2 \leq m < M/2}} \left| \sum_{k \in \kappa} \Psi_m[k] e^{j2\pi k t / \tau} \right|$$

$$\hat{\nu}_\ell \leftarrow 2\pi \hat{m}_\ell / \tau M$$

$$\hat{\alpha}_\ell \leftarrow \frac{\tau}{P|\kappa|} \sum_{k \in \kappa} \Psi_{\hat{m}_\ell}[k] e^{j2\pi k \hat{\tau}_\ell / \tau} / H(2\pi k / \tau)$$

for $k \in \kappa$ **do**

$$\begin{aligned} &[\Psi_{\frac{-M}{2}}[k] \dots \Psi_{\frac{M}{2}-1}[k]] \leftarrow [\Psi_{\frac{-M}{2}}[k] \dots \Psi_{\frac{M}{2}-1}[k]] - \\ &\quad \frac{1}{\tau} H(2\pi k / \tau) \hat{\alpha}_\ell e^{-j2\pi k \hat{\tau}_\ell / \tau} \text{FFT}_M\{e^{-j2\pi \hat{m}_\ell p / M}\} \end{aligned}$$

end for

end for

Another practical concern is target dynamic range. Since target amplitudes can differ by several orders of magnitude, care must be taken so that strong targets do not mask weaker ones. When focusing on some Doppler frequency ν , targets with Doppler frequencies ν_ℓ satisfying $|\nu_\ell - \nu| > \pi/P\tau$ are considered undesirable and we seek to minimize their effect. These targets can be viewed as “out-of-focus”, since they are not matched to ν and their responses from different pulses do not combine coherently; they will combine in phase for different ν 's satisfying $|\nu_\ell - \nu| < \pi/P\tau$. We can add to (19) a user defined window function $w[p]$, $p = 0, 1, \dots, P - 1$ (e.g., Hann, Blackman, etc.) which is designed to mitigate the impact of these out-of-focus targets:

$$\begin{aligned} \Psi_\nu[k] &= \sum_{p=0}^{P-1} c_p[k] e^{j\nu p \tau} w[p] \\ &= \frac{1}{\tau} H(2\pi k / \tau) \sum_{\ell=0}^{L-1} \alpha_\ell e^{-j2\pi k \tau \ell / \tau} \sum_{p=0}^{P-1} e^{j(\nu - \nu_\ell) p \tau} w[p]. \end{aligned} \quad (41)$$

The drawback of windowing is that it increases the frequency's focus zone, potentially including more targets in each delay estimation problem. In Fig. 4 we see an example of how windowing can reduce the effect of out-of-focus targets compared with no windowing (constant $w[p]$). For a comprehensive review of windowing function design considerations see [35]. When attempting to support a set of targets amplitudes comprising a large dynamic range, a situation fairly common in real scenarios, the focusing operation must be performed with aggressive windowing in order for the strong targets to be sufficiently attenuated. In Section VIII we show a simulation demonstrating the improved dynamic range attained using Doppler focusing.

So far we assumed only Swerling-0 point targets. We briefly comment on targets having nonzero Doppler spread—i.e., the




Fig. 4. DFT of windowing functions $w[p]$ compared with no windowing (constant $w[p]$) for $P = 100$ pulses. Attenuation of targets with Doppler frequencies far enough from the nominal frequency ($\nu = 0$ here) increases significantly with proper windowing. Also, focus zone changes for different $w[p]$'s.

micro-Doppler phenomenon [36]. This effect is the modulation of the target's main Doppler frequency caused by motion of the structure of the target around its main trajectory. For example, missile wingtips will exhibit different Doppler frequencies compared with the missile body when the missile is rotating around its central axis. Also, the missile body itself, when undergoing strong vibrations, can generate a spectrum of Doppler frequencies around its central velocity.

Our concern when considering this effect is whether the sparsity assumption remains valid. At the heart of our model is the assumption that the target scene is composed of a small number of targets with discrete Doppler frequencies. Recent works [36], [37] show that for cases of micro-Doppler, there are a small number of dominant frequencies in the continuous Doppler spectrum, caused by distinct vibration modes, rotation rates or resonant frequencies. Therefore, in most cases, targets exhibiting micro-Doppler can be treated as a superposition of several closely spaced targets, and the sparse target scene assumption remains valid.

F. Clutter

To complete this section, we add an analysis of the effects of clutter, considered the next major source of interference after thermal noise. Doppler focusing appears to have inherent clutter rejection capabilities, suggesting that special prefiltering operations such as MTI may not be required.

Clutter refers to unwanted echoes reaching the radar receiver from objects which are not the sought after targets: land, sea, buildings, etc. Due to the relative sizes of the objects, clutter echoes are usually several orders of magnitude stronger than target echoes, so if not treated properly, they can mask the target signals and prevent detection. Furthermore, as opposed to random noise which can be mitigated using coherent integration as in (34), clutter echoes are (deterministic) scaled, shifted and modulated replicas of the transmitted signal. They exhibit the benefits of coherent integration in exactly the same way

as the target signal does. Therefore, the signal-to-clutter ratio (SCR) cannot be increased by increasing the CPI.

The most common method to allow detection in clutter ridden scenarios, is to utilize the fact that clutter, as opposed to most targets, is mostly static. If we assume the radar transceiver itself is also stationary, then clutter echoes will be received with zero Doppler frequency. This is the reason that classic anti-clutter methods (e.g., MTI) are basically a notch filter blocking the Doppler frequency generated by the radar's own motion. We now show that Doppler focusing includes inherent target-clutter Doppler separation, so that it does not require any prefiltering or modifications in order to allow target detection when facing clutter.

The Doppler focusing operation (19) can be viewed as passing the Xamples $c_p[k]$ through a bandpass filter bank, where each filter has a pass-band of width $2\pi/P\tau$. The filters' attenuation can be controlled using windowing (41), at the cost of increasing the pass-band width. This creates adjustable isolation between delay estimation problems (22) for targets with Doppler frequencies separated by more than the pass-band width. Therefore, if clutter were to be primarily concentrated around some specific frequency, targets with Doppler frequencies shifted away by more than approximately $2\pi/P\tau$ could be detected without interference. This emphasizes that increasing the number of pulses P may be used to improve isolation between targets and clutter. In Section VIII we demonstrate via simulation successful target detection using Doppler focusing in a clutter ridden scenario.

VII. COMPARISON TO PREVIOUS APPROACHES

We now compare Doppler focusing to two other methods for delay-Doppler estimation.

A. Simultaneous Delay-Doppler Recovery

Previous methods for delay-Doppler estimation [4], [8] discretize the delay-Doppler plane, and construct a CS dictionary with a column for each two dimensional grid point. This approach can be seen as a discrete MF, where as opposed to classic processing, matching is performed to the entire pulse train rather than to a single pulse. Like a MF, these techniques perform well in noisy conditions, since they achieve an SNR which scales linearly with P .

The critical drawback of this approach is that for any realistic problem size, the dictionary size grows rapidly and becomes too large to store or process. For even moderate size problems, the number of delay or Doppler grid points can easily be on the order of 10^3 , with a number of measurements of similar order. This requires the dictionary to possess 10^9 elements, and occupy many gigabytes of memory, which requires high-end processing capabilities. Doppler focusing enjoys the same SNR improvement as simultaneous delay-Doppler recovery, with a dictionary whose size depends on delay estimation parameters only and remains fixed for any number of received pulses. We could not compare our method to this type of recovery since the Intel Core i5 PC with 12 GB of RAM used in Section VIII could not store the dictionary.

This drawback is not only a problem of memory, but also of computations. We now analyze the computational complexity

of Algorithm 1 on a uniform grid with M Doppler frequencies, and compare it to [4]. In terms of number of samples, Doppler focusing requires $|\kappa|P$ samples while [4] samples at the Nyquist rate and therefore requires $B_h\tau P$ samples. The Doppler focusing operation performs $L|\kappa|$ FFTs of length M , solves LM CS delay recovery problems with dictionary size $|\kappa| \times B_h\tau$, and performs $L|\kappa|P$ sample update operations. If we assume the complexity of a CS problem is proportional to its dictionary size, then the total complexity of Doppler focusing is $O(L|\kappa|M \log(M) + LM|\kappa|B_h\tau + L|\kappa|P)$. In [4] a single CS delay-Doppler recovery problem is solved, but with dictionary size $(PB_h\tau) \times (PB_h^2\tau^2)$, so that recovery complexity is proportional to B_h^3 . Thus, even if the dictionary can be stored, processing scales poorly with increasing bandwidth, which is a critical parameter for high-resolution radar.

B. Two-Stage Recovery

To overcome the problematic scaling of the simultaneous recovery dictionary, two-stage recovery techniques [7], [38] separate delay and Doppler estimation, performing them sequentially rather than in parallel. A common approach to performing two-stage recovery uses the multiple measurement vector (MMV) framework, as performed in [39] in the context of undersampling of sparse wideband signals. MMV recovery jointly processes (18) for $0 \leq p < P$ by stacking the sampled Fourier coefficient vectors and sparse target delay vectors as $\mathbf{C} = [\mathbf{c}_0 \cdots \mathbf{c}_{P-1}] \in \mathbb{C}^{|\kappa| \times P}$ and $\mathbf{X} = [\mathbf{x}_0 \cdots \mathbf{x}_{P-1}] \in \mathbb{C}^{N_\tau \times P}$ accordingly, obtaining

$$\mathbf{C} = \frac{1}{\tau} \mathbf{H} \mathbf{V} \mathbf{X}, \quad (42)$$

where the dictionary $\mathbf{A} = \frac{1}{\tau} \mathbf{H} \mathbf{V}$ is as in (16). MMV recovery algorithms (e.g., Simultaneous OMP [33]) exploit the joint sparsity of \mathbf{X} , i.e., the fact that the support of \mathbf{x}_p remains constant for all p , so that \mathbf{X} has at most L nonzero rows. This joint sparsity is used by taking the norm of the rows of $\mathbf{A} \mathbf{C}$, which can be seen as a form of non-coherent integration, since the phase information is destroyed by the norm operator. The norm is a non-linear operation which mixes together signal and noise components. Therefore, defining a simple SNR measure as we did in (29) and (32) is difficult. Non-coherent integration is a common practice in radar and has been analyzed extensively. Several sources develop approximations of the SNR increase for multiple pulses using non-coherent integration: [40] estimates it at \sqrt{P} for $P > 4$, while [15], [16] estimates P^β where $0.5 < \beta < 0.833$, with β decreasing towards 0.5 as P increases. Regardless of the exact value of β , Doppler focusing, which compensates for the exact phase differences, generates an SNR increase linear with P , better than two-stage recovery. We compare our method to this type of recovery in Section VIII, and show that the difference in SNR is substantial as predicted by the theory.

VIII. SIMULATION RESULTS

We now discuss how the user defined performance metric influences grid size and Fourier coefficient selection, and then show numerical examples comparing our method to other recovery techniques.

A. Performance Metric

Our problem lies in a continuous, analog world. When we choose to solve it using CS, which is an approach developed for discrete problems, we must discretize the delay grid, denoting grid step as Δ_τ . As real world targets delays do not lie on any predefined grid, but our CS recovery assumes they do, it seems we should take $\Delta_\tau \rightarrow 0$ in order to minimize quantization errors. Computational requirements aside, such a decrease in grid step renders the columns in the CS dictionary \mathbf{A} increasingly coherent, where the coherence is defined as

$$\mu(\mathbf{A}) = \max_{i \neq j} \frac{|\langle \mathbf{a}_i, \mathbf{a}_j \rangle|}{\|\mathbf{a}_i\|_2 \|\mathbf{a}_j\|_2}. \quad (43)$$

A basic premise of CS ties low coherence to successful recovery [18], which contradicts taking a small step size.

Here, we relinquish this basic assumption, and argue that depending on the chosen performance metric, high coherence can actually aid recovery. For example, assume we are interested in delay recovery but are tolerant of some small error τ_{\max} . In radar applications, a common performance metric is the hit-or-miss criterion on the estimated delays $\{\hat{\tau}_\ell\}_{\ell=0}^{L-1}$:

$$e_\ell(\hat{\tau}_\ell) = \begin{cases} 0, & \text{if } |\tau_\ell - \hat{\tau}_\ell| < \tau_{\max} \\ 1, & \text{otherwise} \end{cases}$$

$$e = \sum_{\ell=0}^{L-1} e_\ell. \quad (44)$$

Translating τ_{\max} to a condition on support recovery, (44) tolerates an error of no more than $K = \lceil \tau_{\max} / \Delta_\tau \rceil \in \mathbb{N}$ places in the recovered indices from (16). Instead of designing \mathbf{A} so that each column is as non-correlative with the other columns as possible, we should design the dictionary so that each column is correlative with its K nearest neighbors. This will improve recovery performance in noisy scenarios since in cases where the correct column is not recovered, any one of its similar neighboring columns still has a chance to overcome the noise and produce a ‘‘hit’’. Graphs in the next subsection show that such a coherent \mathbf{A} actually improves recovery performance compared with a less coherent dictionary in very noisy scenarios.

We can control the level of \mathbf{A} 's coherence by choosing different sets of Fourier coefficient κ in (14). Fig. 5 shows an example of the column correlation pattern for two sets of Fourier coefficients: a consecutive set and a random set, where all coefficients were chosen in $[-B_h/2, B_h/2]$. We define the column correlation function for some column \mathbf{a}_i as

$$\mu_i[j] = \frac{|\langle \mathbf{a}_i, \mathbf{a}_j \rangle|}{\|\mathbf{a}_i\|_2 \|\mathbf{a}_j\|_2}. \quad (45)$$

The consecutive set is better suited for criteria which allow some error in support recovery, while the random set achieves better performance when exact recovery is required.

B. Numerical Results

We now present some numerical experiments illustrating the recovery performance of a sparse target scene. We corrupt the received signal $x(t)$ with an additive white Gaussian noise $n(t)$ with power spectral density $S_n(f) = N_0/2$, bandlimited to




Fig. 5. Column correlation pattern $\mu_0[j]$ of the CS dictionary from (16) for two sets κ . The consecutive set achieves coherence of 0.9 due to the many correlative columns in the center, while the random set has coherence of 0.3.

$x(t)$'s bandwidth B_h . We define the signal to noise power ratio for target ℓ as

$$\text{SNR}_\ell = \frac{\frac{1}{T_p} \int_0^{T_p} |\alpha_\ell h(t)|^2 dt}{N_0 B_h}, \quad (46)$$

where T_p is the pulse time. The scenario parameters used were number of targets $L = 5$, number of pulses $P = 100$, PRI $\tau = 10 \mu\text{sec}$, and $B_h = 200 \text{ MHz}$. Target delays and Doppler frequencies are spread uniformly at random in the appropriate unambiguous regions, and target amplitudes were chosen with constant absolute value and random phase. The classic time and frequency resolutions (Nyquist bins), defined as $1/B_h$ and $1/P\tau$, are 5 nsec and 1 KHz accordingly. In order to demonstrate a 1:10 sampling rate reduction, our sub-Nyquist Xampling scheme generated 200 Fourier coefficients per pulse, as opposed to the 2000 Nyquist rate samples. We tested Doppler focusing with two types of Fourier coefficient sets κ , a consecutive set and a random set. We compared Doppler focusing with classic processing and the two-stage recovery method of [7] (where we use a CS algorithm instead of ESPRIT) using the following criteria:

- 1) **Hit-Rate**—we define a “hit” as a delay-Doppler estimate which is circumscribed by an ellipse around the true target position in the time-frequency plane. We used an ellipse with axes equivalent to ± 3 times the time and frequency Nyquist bins.
- 2) **Recovery RMS error**—for estimates classified as hits, we measure the root mean square error in both time and frequency.

As noted in the previous section, a single stage CS recovery method using Nyquist bins spacing consumes a prohibitive amount of memory and was not able to run on any computer at hand, since the CS dictionary required storing $4 \cdot 10^9$ elements (occupying 32 GB of memory using standard IEEE double precision): $\frac{2\pi/\tau}{2\pi/P\tau} \frac{\tau}{1/B_h} = P\tau B_h = 2 \cdot 10^5$ columns and $P\tau B_h/10 = 2 \cdot 10^4$ measurements per column.

For CS-based techniques, the delay grid step $\Delta\tau$ was chosen as half a Nyquist bin. For Doppler focusing, the Doppler frequency region was discretized with uniform steps of half a




Fig. 6. Hit Rate for classic processing, two-stage CS recovery and Doppler focusing. Sub-Nyquist sampling rate was one tenth the Nyquist rate.

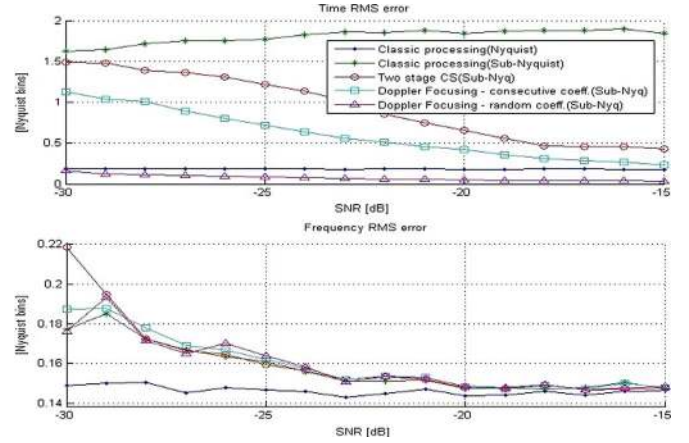


Fig. 7. RMS error of time and frequency estimates for classic processing, two-stage CS recovery and Doppler focusing. Sub-Nyquist sampling rate was one tenth the Nyquist rate.

Nyquist bin. To provide a fair comparison, classic processing was performed using identical bin sizes.

Figs. 6 and 7 demonstrate the hit-rate and RMS error performance of the different recovery methods for various SNR values. It is evident that Doppler focusing is superior to the other sub-Nyquist recovery techniques. Between the two Doppler focusing approaches, consecutive coefficients are better suited for lower SNR, while choosing coefficients randomly improves performance as SNR increases. Since both sets κ produce CS dictionaries with column correlation functions which are not matched to the “hit-or-miss” performance criteria used, we have no reason to assume one should be better than the other. Also, random coefficients, when producing a hit, have very small delay errors (even compared with Nyquist rate classic processing) due to low CS dictionary coherence. As opposed to Doppler focusing recovery performance which decreases gracefully with sample rate reduction, classic processing suffers significantly when sample rate is reduced below Nyquist.

Fig. 8 shows the same hit rate graph for classic processing, but this time the waveform used for Doppler focusing had its CTFT adjusted so that energy was transferred from frequencies which were not sampled, to those that were. This was performed by




Fig. 8. Hit Rate for classic processing and Doppler focusing at one tenth the Nyquist rate, where the waveform used for Doppler focusing had its entire energy contents concentrated in the sampled frequencies.




Fig. 9. Probability to distinguish two separate closely spaced identical targets, where classic processing uses a tenfold wideband signal compared with Doppler focusing.

passing the signal through a low pass filter and rescaling its amplitude so that target SNR (46) remains constant. Since Doppler focusing imposes no restrictions on the transmitter, we can use a signal with the same total energy, but have it spread out in a manner which is more favorable to the frequency domain sampling used in Xampling. As performance for Doppler focusing improves significantly, we are able to obtain excellent recovery results at much lower SNR values, surpassing classic processing which uses ten times as many samples. This shows that the performance degradation caused by sub-Nyquist sampling can be compensated for using a suitable transmitter. The drawback of using such a narrowband signal can be seen in Fig. 9, where we examine resolution in terms of the ability to separate two closely spaced identical targets with equal Doppler frequencies. We see that for very close targets, classic processing using a wideband signal is able to distinguish the two targets far better than Doppler focusing recovery using a narrowband signal.

Fig. 10 shows the sparse target scene on a time-frequency map for a -28 dB SNR scenario, where each target is displayed along with its hit rate ellipse, together with the various sub-Nyquist recovery methods' estimates and hit rates. As noted in Section IV, only Doppler focusing is able to distinguish between

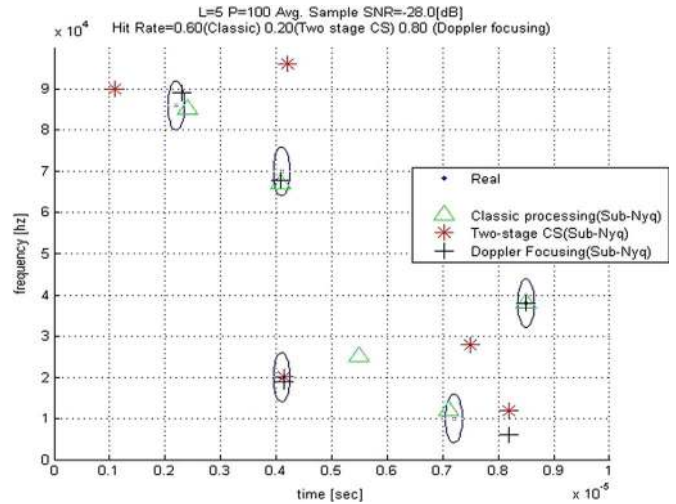


Fig. 10. Real target positions along with various estimates. Doppler focusing achieves highest hit rate among sub-Nyquist methods. Only Doppler focusing detects two targets around $4.2 \mu\text{sec}$. (Hit rate ellipses were enlarged for clarity).




Fig. 11. Target scene composed of two closely spaced targets, the target on the left 20 dB more powerful than the target on the right. MF processing at both Nyquist and one tenth the Nyquist rate recovers only one of the two targets, while Doppler focusing recovers both. No windowing was used ($w[p] = 1$).

the two targets having almost identical delays (around $4.2 \mu\text{sec}$) but different Doppler frequencies.

Regarding target dynamic range, Fig. 11 demonstrates the advantage of using Doppler focusing to resolve closely spaced targets with different powers. In this scenario, two targets satisfying $20 \log_{10}(|\alpha_1|/|\alpha_2|) = 20[\text{dB}]$ are placed adjacently, so that their hit rate ellipses intersect. Doppler focusing based recovery at one tenth the Nyquist rate generates two hits, while MF processing at Nyquist rate recovers only one of the targets.

C. Clutter

We demonstrate the robustness of Doppler focusing to clutter with the following scenario. We simulate nine targets with Doppler frequencies spread uniformly in $[-\pi/\tau, \pi/\tau]$. Clutter is modeled as 4000 Swerling-0 scatterers, with Doppler frequencies distributed uniformly in a single Nyquist bin around




Fig. 12. Histogram of target vs. clutter Xamples $c_p[k]$. Average SCR for a single Xample is -50 dB.

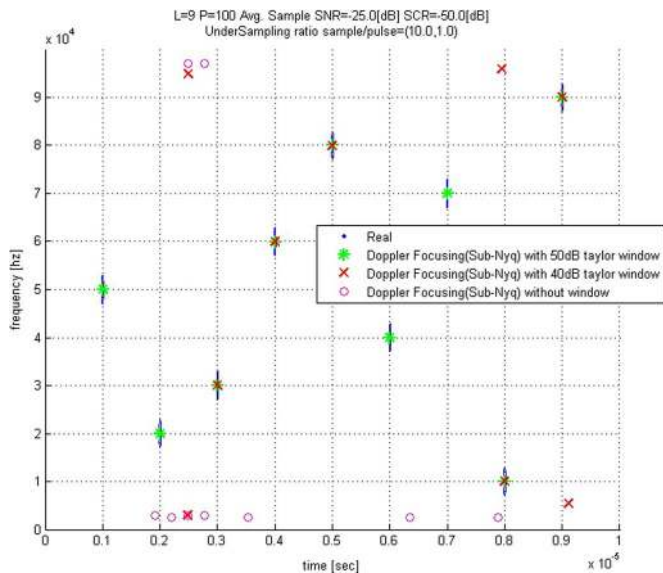


Fig. 13. Recovered target scene with nine targets and almost static clutter. Without windowing clutter sidelobes permeate the nonzero Doppler frequency area and cause misdetections. With 40 dB windowing five out of nine targets are “hits”. With 50 dB windowing the entire scene is detected correctly.

DC, to allow for small relative velocity, e.g., waves in the sea. Clutter scatterers are distributed uniformly at random at all delays in $[0, \tau]$. The SCR is -50 dB, as can be seen in Fig. 12, and the targets’ SNR (46) is -25 dB. All other system parameters are as described in the previous subsection.

We perform detection using Algorithm 1 using CS for delay recovery, ignoring the clutter ridden DC Doppler bin, and due to leakage effects, also its two nearest neighbours on each side. Since for this very low SCR, clutter sidelobes still cover the targets, we use a Taylor window with -50 dB attenuation to improve Doppler frequency isolation. The recovered target scene with and without windowing is shown in Fig. 13. This simulation shows that any interference which is localized in Doppler, can be manipulated using (41), so it has a negligible effect on targets in other Doppler frequencies.




Fig. 14. The 4-channel radar receiver board.

IX. RADAR EXPERIMENT

In this section we present a real experiment of our radar receiver hardware prototype. Our setup includes a custom made sub-Nyquist radar receiver board (see Fig. 14) which implements sub-Nyquist Xampling and digital recovery using Doppler focusing, while the analog input signal (2) was synthesized using National Instruments (NI) hardware.

To evaluate the board we make use of NI PXI equipment for both system synchronization and signal sources. The entire component ensemble is wrapped in an NI chassis. The RF front end and board we use are identical to the ones used in [2], but the digital recovery method accounting for target Doppler frequencies is different. Additional information regarding the system’s configuration and synchronization issues can be found in [2].

The experimental process consists of the following steps. We begin by using the AWR software, which enables us to examine a large variety of scenarios, comprised of different target parameters, i.e., delays, Doppler frequencies and amplitudes. With the AWR software we simulate the complete radar scenario, including the pulse transmission and accurate power loss due to wave propagation in a realistic medium. The AWR also contains a model of a realistic RF receiver, which simulates the demodulation of the RF signal to IF frequencies, and saves the output to a file. Next, the simulation result is loaded to the AWG module, which produces an analog signal. This signal is amplified using the NI 5690 low noise amplifier and then routed to our radar receiver board, which has 4 parallel input channels. Each channel samples a different frequency band, in the following manner: each channel is fed by a local oscillator (LO), which modulates the desired frequency band of the channel to the central frequency of a narrow 80 KHz bandwidth band pass filter (BPF). A fifth LO, common to all 4 channels, modulates the BPF output to a low frequency band, and it is sampled with a standard low rate ADC. The LOs are created using three NI 5781 baseband transceivers, acting as trigger based signal generators with a constant and known phase, controlled by NI Flex Rio FPGAs. The AWG also triggers the ADC to sample 250 samples in each sampling cycle, per channel. These samples are fed into the chassis controller and a MATLAB function is launched that runs the Doppler focusing reconstruction algorithm. Our system contains a fully detailed interface implemented in the LabView environment, which allows simple activation of the process.

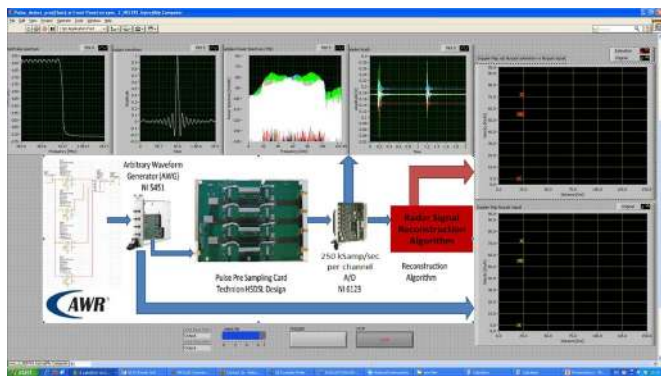


Fig. 15. The LabView experimental interface. From left to right: $H(\omega)$, $h(t)$, the frequency response of each channel, the 4 signals detected in each channel, at the top- the reconstructed target scene, at the bottom- the original target scene.

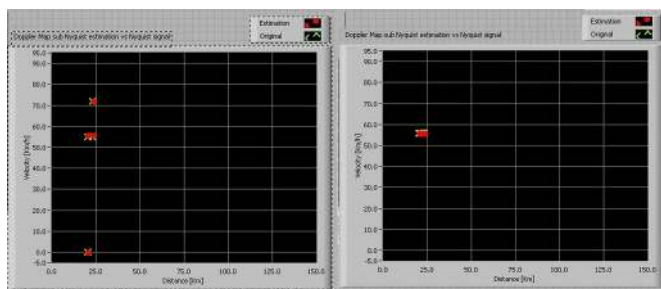


Fig. 16. Two additional target scenes. On the left, all four targets have closely spaced delays, and two of the targets also have close Doppler frequencies. On the right, all four targets have very similar delays and Doppler frequencies. Doppler focusing based recovery is successful in both cases.

Various target scenes, with different target delays, Doppler frequencies and amplitudes, are recovered successfully using this setup. Screenshots of the interface are depicted in Fig. 15 and in Fig. 16.

This experimental prototype proves that the sub-Nyquist methodology described in this paper is actually feasible in practice. The recovery method proposed here not only describes digital recovery, but also addresses the problem of sampling the analog signal at a low rate, in a way which is feasible with standard RF hardware.

X. CONCLUSION

We demonstrated a radar sampling and recovery method called Doppler focusing, which employs the techniques of Xampling and CS, and is independent of the radar signal's bandwidth. Doppler focusing allows for low rate sampling and digital processing, and imposes no constraints on the transmitted signal. It also leads to CS recovery with dictionary size scaling with delay grid size only, and provides SNR scaling which is linear in the number of received pulses, identical to an optimal MF. We compared our method to other sub-Nyquist recovery techniques and have seen its clear advantage in simulations. When sampling at one tenth the Nyquist rate, and for SNR above -25 dB, Doppler focusing achieves results almost equal to classic recovery working at the Nyquist rate.

We are currently working on enhancing Doppler focusing to handle the case of an unknown number of targets, and increasing the algorithm's dynamic range by improving treatment of strong targets and the sidelobes they introduce.

ACKNOWLEDGMENT

We would like to thank the anonymous reviewers for their constructive remarks, which helped improve the paper.

REFERENCES

- [1] 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, Jan. 2011.
- [2] E. Baransky, G. Itzhak, I. Shmuel, N. Wagner, E. Shoshan, and Y. C. Eldar, "A sub-Nyquist radar prototype: Hardware and algorithms," *IEEE Trans. Aerosp. Electron. Syst., Special Issue on Compressed Sensing for Radar*, Aug. 2012, submitted for publication.
- [3] 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, Apr. 2011.
- [4] M. A. Herman and T. Strohmer, "High-resolution radar via compressed sensing," *IEEE Trans. Signal Process.*, vol. 57, no. 6, pp. 2275–2284, June 2009.
- [5] B. Demissie, "High-resolution range-Doppler imaging by coherent block-sparse estimation," in *Proc. Int. Workshop Compress. Sens. Appl. Radar*, May 2012.
- [6] R. G. Baraniuk and P. Steeghs, "Compressive radar imaging," presented at the IEEE Radar Conf., Waltham, MA, USA, Apr. 2007.
- [7] 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, June 2011.
- [8] J. Zhang, D. Zhu, and G. Zhang, "Adaptive compressed sensing radar oriented toward cognitive detection in dynamic sparse target scene," *IEEE Trans. Signal Process.*, vol. 60, no. 4, pp. 1718–1729, Apr. 2012.
- [9] M. Vetterli, P. Marziliano, and T. Blu, "Sampling signals with finite rate of innovation," *IEEE Trans. Signal Process.*, vol. 50, no. 6, pp. 1417–1428, 2002.
- [10] M. Mishali and Y. C. Eldar, "Sub-Nyquist sampling: Bridging theory and practice," *IEEE Signal Process. Mag.*, vol. 28, no. 6, pp. 98–124, Nov. 2011.
- [11] N. Wagner, Y. C. Eldar, and Z. Friedman, "Compressed beamforming in ultrasound imaging," *IEEE Trans. Signal Process.*, vol. 60, no. 9, pp. 4643–4657, Sept. 2012.
- [12] T. Chernyakova and Y. C. Eldar, "Fourier domain beamforming: The path to compressed ultrasound imaging," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, to be published.
- [13] M. Skolnik, *Radar Handbook*. New York, NY, USA: McGraw-Hill, 1970.
- [14] *The Chinese Remainder Theorem and Multi-PRF Radars*, HMSO, Royal Signals and Radar Research Establishment, Tech. Memo., Feb. 1984.
- [15] M. A. Richards, *Fundamentals of Radar Signal Processing*. New York, NY, USA: McGraw-Hill, 2005.
- [16] M. A. Richards, *Noncoherent Integration Gain, and Its Approximation, Noncoherent Integration Gain, and Its Approximation*, Georgia Inst. of Technology, Tech. Memo., Jun. 2010.
- [17] 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, Apr. 2011.
- [18] M. Mishali and Y. C. Eldar, "Xampling: Compressed sensing for analog signals," in *Compressed Sensing: Theory and Applications*. Cambridge, U.K.: Cambridge Univ. Press, 2012.
- [19] M. Mishali, Y. C. Eldar, and A. Elron, "Xampling: Signal acquisition and processing in union of subspaces," *IEEE Trans. Signal Process.*, vol. 59, no. 10, pp. 4719–4734, Oct. 2011.
- [20] P. Stoica and R. Moses, *Introduction to Spectral Analysis*. Englewood Cliffs, NJ, USA: Prentice-Hall, 2000.
- [21] T. K. Sarkar and O. Pereira, "Using the matrix pencil method to estimate the parameters of a sum of complex exponentials," *IEEE Antennas Propag. Mag.*, vol. 37, no. 1, pp. 48–55, Feb. 1995.

- [22] R. Roy and T. Kailath, "ESPRIT-estimation of signal parameters via rotational invariance techniques," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 37, no. 7, pp. 984–995, Jul. 1989.
- [23] R. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Trans. Antennas Propag.*, vol. 34, no. 3, pp. 276–280, Mar. 1986.
- [24] Y. Chi, L. L. Scharf, A. Peshzki, and A. R. Calderbank, "Sensitivity to basis mismatch in compressed sensing," *IEEE Trans. Signal Process.*, vol. 59, no. 5, pp. 2182–2195, May 2011.
- [25] S. G. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Trans. Signal Process.*, vol. 41, no. 12, pp. 3397–3415, Dec. 1993.
- [26] T. Blumensath and M. Davies, "Iterative hard thresholding for compressive sensing," *Appl. Comput. Harmon. Anal.*, vol. 27, no. 3, pp. 265–274, 2009.
- [27] Y. C. Eldar and G. Kutyniok, Eds., *Compressed Sensing: Theory and Applications*. Cambridge, U.K.: Cambridge Univ. Press, May 2012.
- [28] M. Rudelson and R. Vershynin, "On sparse reconstruction from Fourier and Gaussian measurements," *Commun. Pure Appl. Math.*, vol. 61, no. 8, pp. 1025–1045, 2008.
- [29] J. Tsao and B. D. Steinberg, "Reduction of sidelobe and speckle artifacts in microwave imaging: The CLEAN technique," *IEEE Trans. Antennas Propag.*, vol. 36, no. 4, pp. 543–556, Feb. 1988.
- [30] G. Schwartz, "Estimating the dimension of a model," *Ann. Statist.*, vol. 6, no. 2, pp. 461–464, 1978.
- [31] M. Wax and T. Kailath, "Detection of signals by information theoretic criteria," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 33, no. 2, pp. 387–392, Apr. 1985.
- [32] J. J. Fuchs, "Estimating the number of sinusoids in additive white noise," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 36, no. 12, pp. 1846–1853, Dec. 1988.
- [33] M. F. Duarte and Y. C. Eldar, "Structured compressed sensing: From theory to applications," *IEEE Trans. Signal Process.*, vol. 59, pp. 4053–4085, Sep. 2011.
- [34] Y. Hua, "Estimating two-dimensional frequencies by matrix enhancement and matrix pencil," *IEEE Trans. Signal Process.*, vol. 40, no. 9, pp. 2267–2280, Sep. 1992.
- [35] F. J. Harris, "On the use of windows for harmonic analysis with the discrete Fourier transform," *Proc. IEEE*, vol. 66, no. 1, pp. 51–83, Jan. 1978.
- [36] V. C. Chen, F. Li, S. S. Ho, and H. Wechsler, "Micro-Doppler effect in radar: Phenomenon, model, and simulation study," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 42, no. 1, pp. 2–21, Jan. 2006.
- [37] V. C. Chen, F. Li, S. S. Ho, and H. Wechsler, "Analysis of Micro-Doppler signatures," *Proc. Inst. Electr. Eng.—Radar, Sonar, Navig.*, vol. 150, no. 4, pp. 271–276, Aug. 2003.
- [38] O. Bar-Ilan and Y. C. Eldar, "Sub-Nyquist radar," presented at the 9th Int. ITG Conf. Syst., Commun., Coding, Munich, Germany, Jan. 2013.
- [39] M. Mishali and Y. C. Eldar, "Blind multiband signal reconstruction: Compressed sensing for analog signals," *IEEE Trans. Signal Process.*, vol. 57, no. 3, pp. 993–1009, Mar. 2009.
- [40] J. Minkoff, *Signal Processing Fundamentals and Applications for Communications and Sensing Systems*. Norwood, MA, USA: Artech House, 2002.



Omer Bar-Ilan received the B.Sc. degree in electrical engineering and the B.Sc. degree in physics both from the Technion—Israel Institute of Technology, Haifa, in 2008. He is currently pursuing the M.Sc. degree in electrical engineering at the Technion—Israel Institute of Technology. His research interests include theoretical aspects of signal processing, sampling theory, and advanced signal processing methods for radar.



Yonina C. Eldar (S'98–M'02–SM'07–F'12) received the B.Sc. degree in physics and the B.Sc. degree in electrical engineering both from Tel-Aviv University (TAU), Tel-Aviv, Israel, in 1995 and 1996, respectively, and the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, in 2002.

From January 2002 to July 2002, she was a Postdoctoral Fellow at the Digital Signal Processing Group at MIT. She is currently a Professor in the

Department of Electrical Engineering at the Technion—Israel Institute of Technology, Haifa and holds the Edwards Chair in Engineering. She is also a Research Affiliate with the Research Laboratory of Electronics at MIT and a Visiting Professor at Stanford University, Stanford, CA. Her research interests are in the broad areas of statistical signal processing, sampling theory and compressed sensing, optimization methods, and their applications to biology and optics.

Dr. Eldar was in the program for outstanding students at TAU from 1992 to 1996. In 1998, she held the Rosenblith Fellowship for study in electrical engineering at MIT, and in 2000, she held an IBM Research Fellowship. From 2002 to 2005, she was a Horev Fellow of the Leaders in Science and Technology program at the Technion and an Alon Fellow. In 2004, she was awarded the Wolf Foundation Krill Prize for Excellence in Scientific Research, in 2005 the Andre and Bella Meyer Lectureship, in 2007 the Henry Taub Prize for Excellence in Research, in 2008 the Hershel Rich Innovation Award, the Award for Women with Distinguished Contributions, the Muriel & David Jacknow Award for Excellence in Teaching, and the Technion Outstanding Lecture Award, in 2009 the Technion's Award for Excellence in Teaching, in 2010 the Michael Bruno Memorial Award from the Rothschild Foundation, and in 2011 the Weizmann Prize for Exact Sciences. In 2012 she was elected to the Young Israel Academy of Science and to the Israel Committee for Higher Education, and elected an IEEE Fellow. In 2013 she received the Technion's Award for Excellence in Teaching, the Hershel Rich Innovation Award, and the IEEE Signal Processing Technical Achievement Award. She received several best paper awards together with her research students and colleagues. She is the Editor in Chief of *Foundations and Trends in Signal Processing*. In the past, she was a Signal Processing Society Distinguished Lecturer, member of the IEEE Signal Processing Theory and Methods and Bio Imaging Signal Processing technical committees, and served as an associate editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING, the *EURASIP Journal of Signal Processing*, the *SIAM Journal on Matrix Analysis and Applications*, and the *SIAM Journal on Imaging Sciences*.