A Novel Multiplicative Phase Dithering Scheme for 1-bit Compressive Radar

Thomas Feuillen[1], Luc Vandendorpe[2] and Laurent Jacques[2]
[1]*SPARC/SnT*, *UniLu*
Luxembourg
[2]*ICTEAM*, *UCLouvain*
Louvain-la-Neuve, Belgium

Abstract*—In this paper, we tackle the issue of implementing a dithering procedure for the 1-bit quantization of radar signals that is able to generate high-quality estimates while remaining a low-complexity and cost-efficient solution. Classic Dithering techniques that add a signal before the Analog to Digital Converter (ADC) have been used in many acquisition chain designs and have been studied as a way to *shape* the quantization noise more favourably. Here, we stray away from this additive dithering, which, as will be made clear later, induces a complex and high-cost implementation. Instead, we propose the use of a multiplicative phase dithering. This process can leverage already existing radar architectures of Frequency Modulated Continuous Wave (FMCW) radars and can thus be efficiently implemented. The efficiency of this multiplicative dithering is first studied theoretically, and its link to another coarse quantization scheme, namely the Phase-Only acquisition, is highlighted. The performances of this novel dithering scheme are then extensively tested using Monte Carlo simulations and are thoroughly compared to their additive counterparts. A hardware-relaxed version of the random phase dithering is also introduced and compared to the other 1-bit schemes. The observations made in simulations are then validated using actual radar measurements at $24\mathrm{GHz}$. Combined with the simulations, these measurements show that the multiplicative dithering is an appealing alternative to the additive random dithering in a low number of measurement setting. Specifically, we show that this procedure is a good trade-off between strong theoretical guarantees and reconstruction quality for low-complexity hardware.

I. Introduction

Recent years have seen the advent of smaller form factor radars [1], [2]. These small-footprint sensors have opened the way for numerous new applications, from healthcare [3], [4] to automotive safety [2], [5]. More generally, having small contactless sensors enables the ubiquitous sensing of the environment around us. In applications with little available space (*e.g.*, car bumper), this smaller footprint also enables the use of more complex transmitter topologies, such as MIMO radars, to increase the resolution of the system. This new breadth of applications, with higher sensor density, creates new challenges in acquiring and processing the data they generate. Data storage and transfer is becoming a major bottleneck in many applications. This paper seeks to remedy this for FMCW radar sensors by proposing a new extremely coarse 1-bit radar system architecture.

Reducing or compressing data streams has been studied in radar, as well as for other modalities, along two main axes: *a posteriori*, *i.e.*, compressing after the acquisition of the data stream or signal and *a priori* by modifying the sensor system itself to generate a compressed stream. The latter is the focus of this work.

High-quality Data and Compression. Data compression is ubiquitous when dealing with images on digital systems, *e.g.*, [6] and has also been investigated for Synthetic Aperture Radars (SAR) [7], [8]. In this context, the processing load to synthesise the SAR image exceed the one available on board the satellite; compressing the data to transmit it to the ground station is thus necessary. This approach, however, still generates a large quantity of data at the sensor stage. In this paper we focus on reducing the data rate directly at the sensor. Compressive sensing proposes, in its more classic form, to do so by lowering the number of measurements taken [9], [10]. In this framework, the sensor is modelled as $\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{x}$, with $\boldsymbol{y} \in \mathbb{C}^m$ the measurements resulting from the projection of a vector $\boldsymbol{x} \in \mathbb{C}^N$ by a so-called measurements matrix $\boldsymbol{\Phi} \in \mathbb{C}^{m\times N}$. Depending on the sensor, the properties of $\boldsymbol{\Phi}$ and the nature of the scenes $\boldsymbol{x}$ it observes (*e.g.*, sparsity), the number of measurements m can be dramatically lowered. The estimation from this *compressed* $\boldsymbol{y}$ is performed thanks to a wide range of algorithms tailored to different settings [9], [11], [12]. This has been demonstrated in imaging systems such as in the single pixel camera [13], in MRI in [14], [15], and in radar systems [16]–[18]. In the context of Frequency Modulated Continuous Wave (FMCW) radars, the sensor can be seen as acquiring a Fourier transform ($\boldsymbol{\Phi}$) of the range profile $\boldsymbol{x}$ its antenna illuminates.

Compression by Coarse Acquisition Compressive Sensing, in its classic linear form ($\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{x}$), implicitly considers High-Resolution ADCs with no quantization noise. The compression is thus left only to the sensor, *i.e.*, the dimension of $\boldsymbol{\Phi}$ and its design. Considering the ADC explicitly adds another dimension for the compression to take effect; the resolution of the ADCs can be lowered. The system is now modelled as $\boldsymbol{y} = \mathcal{Q}(\boldsymbol{\Phi}\boldsymbol{x})$; both the acquisition $\mathcal{Q}(\cdot)$ and the system $\boldsymbol{\Phi}$ can perform the compression. In this paper, we consider the most extreme acquisition setting, 1-bit ADCs .

1-bit Systems, Design & Limits. 1-bit acquisition schemes, whereby the signal is only recorded using binary value (*e.g.*, ± 1 or *low* and *high*) have been heavily investigated in multiple applications such as wireless communications [19], [20], time of flight sensors [21], and radar sensors [22], [23]. Indeed, recording signals in this fashion allows for low-cost or higher sampling frequency ADCs [24], [25], enabling the creation of more efficient systems in terms of cost, power, computation and or storage. The 1-bit quantizer used in this paper is defined as $\mathcal{Q}_\epsilon(\boldsymbol{y}) := \epsilon\,\mathrm{sign}(\Re\{\boldsymbol{y}\}) + \mathrm{i}\epsilon\,\mathrm{sign}(\Im\{\boldsymbol{y}\})$ with $\epsilon > 0$ the resolution in volts. 1-bit ADCs are assigned to both the real and imaginary part of the radar signal. The implementation is extremely simplified, but at a cost in terms of performance, as shown in [23], [26]–[28]. In [27] for Bernoulli matrices and in [23], [28] for Fourier matrices, authors showed that applying this extremely low-resolution quantizer on noiseless data can result in ambiguous scenarios, preventing high-quality reconstruction. One such scenario arises when two different sparse vectors $\boldsymbol{x}_0$ and $\boldsymbol{x}_1$, once measured by the measurement matrix $\boldsymbol{\Phi}$ of the sensor and then quantized, are sent exactly to the same bits. This removes any possibility to distinguish them after quantization. In the context of radar, the system cannot distinguish two different scenes, *i.e.*, it is partly blind.

Additive Dithering. Dither, whereby a random noise is added to the signal before the quantization, has been used since the early 1960s to decorrelate the quantisation noise from the signal [29]. Adding a random noise before the sign operator is tantamount to changing the threshold where the bit flips. The dithering can be modelled by a vector $\boldsymbol{\xi} \in \mathbb{C}^m$ and the dithered quantizer becomes

$$\mathcal{Q}^+_\epsilon(\boldsymbol{y}) := \mathcal{Q}_\epsilon(\boldsymbol{y} + \boldsymbol{\xi}).$$

This principle has also been leveraged in $\Sigma\Delta$ quantizer, where instead of adding a random noise, a feedback loop controls the thresholds in order to *shape the noise* generated by the quantizer by pushing it outside of the bandwidth of the signal of interest [30].

Authors in [31]–[36] proposed generating these varying thresholds according to the previous measurements to extract the most information possible from those coarse 1-bit measurements. While this process of generating *tailored* thresholds iteratively induces a steep performance gain in the reconstruction, it implies some memory and added processing for the feedback to take place. Authors in [26], [37] proposed the use of a memory-less dither that is generated randomly according to a uniform distribution $\boldsymbol{\xi} \sim \mathcal{U}^m(-\frac{\epsilon}{2}, \frac{\epsilon}{2})$, with $\epsilon \geq 2\|\boldsymbol{\Phi x}\|_\infty$. They were able to show that by leveraging the effect of this dither on the quantized data, one can upperbound the ℓ_2-reconstruction of the Projected Back-Projection (PBP) algorithm [37]. Authors in [38] used a Majorization–Minimization to recover sums of sinudoidals from dithered 1-bit measurements generated this way. Memory-less dither only needs to be generated with no control or influence by the reconstruction algorithm and signal; it is purely random.

Limits of Additive Dithering. While it is clear that adding a uniform random variable gives robust recovery guarantees [37], actually implementing this dither in practice, however, comes with numerous challenges. Indeed, and as stated before, the added dither must follow a uniform distribution that spans the dynamic of the measured signal, *i.e.*, $\|\boldsymbol{y}\|_\infty$. The problems associated with this constraint are threefold. First, one must find a way of generating this uniform random variable in hardware at a cost or power consumption that is still attractive compared to high-resolution ADCs. This means finding a hardware component that can generate this distribution or using high-resolution DACs. However, trading high-resolution ADCs for high-resolution DACs does not really simplify the acquisition process. Second, having this random variable span the signal's dynamic range implies that this dynamic is partially known. This assumption is rather strong given the high variability of received powers in radar systems. Indeed, for a constant RCS (Radar Cross Section), the received power varies as $\mathcal{O}(R^{-4})$ in a monostatic radar [39], with R the target's range. This means that there is more than a $10\mathrm{dB}$ loss of amplitude every time the range of a target doubles. Finally, from a computational standpoint, the last issue with this architecture is the fact that specific algorithms, such as the QIHT algorithm, the quantized version of *Iterative Hard Thresholding*[11], require the exact knowledge of the dither used for the reconstruction.

From Amplitude Dithering to Phase Dithering. Given the issues associated with this additive dithering combined with the need to go beyond the simple deterministic quantization, we set out to find another way of dithering signals coming from FMCW radars. We introduce the multiplicative dither that modifies, *i.e.*, dithers, the phases of the measured signals before the 1-bit quantization according to a unitary vector with a random phase in a cost and complexity-efficient manner. We do so by linking this new quantizer to the so-called Phase Only (PO), which only records the phase of a complex signal, discarding its amplitude.

Contributions. In this work, we significantly extend [40] in several key directions[1] through an end-to-end study of the proposed 1-bit radar architecture.

First, we show that using a multiplicative dither in the context of 1-bit quantization is a viable trade-off between reconstruction performances and hardware implementation complexity. Moreover, we demonstrate that, instead of dithering with a random phase [40], one can use a structured dither without sacrificing performance while simplifying further the hardware implementation of the dithering process. This structured dither enables the use of consistency-based algorithms (such as QIHT) without any added complexity commonly found in the additive quantization scheme.

Second, using this multiplicative dithering scheme, we theoretically bound the reconstruction error of the naive Projected Back-Projection (PBP) algorithm with the error guarantees achieved by the same algorithm when applied to phase-only acquisition. We also theoretically show the limit of the PO acquisition for general Fourier systems by demonstrating the existence of ambiguous scenarios that cannot be distinguished after being acquired. We then show that these ambiguous scenarios are not relevant in the context of FMCW radar systems.

Third, these results are then verified through extensive Monte-Carlo simulations. Specifically, we study the ℓ_2-error,

[1] Sec.IV-B, Sec.V, parts of Sec.VI and its proof in Sec., parts of Sec.VII-A, Sec.VII-B, and Sec.VIII, were not present in [40] and more details are given throughout the manuscript

and the support recovery of the different schemes and show the resilience of the proposed phase-dithered 1-bit quantization in the context of an unknown dynamic range. Finally, real radar measurements are presented to confirm in practice the feasibility and performance of the proposed low-resolution acquisition scheme.

Structure. This paper is structured as follows. Section I introduces the issue of implementing an efficient dithering procedure in the context of radar systems. Sec. II introduces the notations used throughout the paper. Sec. III introduces the sensor whose model is leveraged to obtain an efficient and cost-effective dither, *i.e.*, FMCW radars. Sec. IV introduces the multiplicative dithering studied in this paper and highlights its advantages compared to the additive dither in terms of practical implementation. Sec. V links the multiplicative dithering to the PO acquisition and studies its limits in terms of uniform recovery guarantees of Fourier-based measurements and its impact on radar systems. Sec. VI introduces a bound on the ℓ_2 reconstruction error of PBP using the proposed scheme and studies a non-uniform ℓ_2-bound on the discrepancy between the reconstruction obtained with this multiplicative dither and the PO-CS acquisition. The details of the different proofs can be found in the appendices. The results of Monte-Carlo simulations are presented in Sec. VII. These performances are then validated in Sec. VIII, where real radar measurements are used. Finally, conclusions and future works are presented in Sec. IX.

II. Notations

We denote matrices and vectors with bold symbols, *e.g.*, $\boldsymbol{\Phi} \in \mathbb{C}^{m\times N}$, $\boldsymbol{x} \in \mathbb{C}^N$, and scalar values with light symbols. We will often use the following quantities: $[d] := \{1, \cdots, d\}$ with $d \in \mathbb{N}$; the complex number i such that $\mathsf{i}^2 = -1$; $\mathbb{R}\{\lambda\}$ (or $\lambda^{\mathbb{R}}$) and $\mathbb{I}\{\lambda\}$ (or $\lambda^{\mathbb{I}}$) are the real and imaginary part of $\lambda \in \mathbb{C}$, respectively, and λ^* is its complex conjugate; $\boldsymbol{A}^H$ is the conjugate transpose of $\boldsymbol{A}$; $\operatorname{supp} \boldsymbol{x}$ is the support of $\boldsymbol{x} \in \mathbb{C}^d$; $|\mathcal{S}|$ is the cardinality of a finite set $\mathcal{S}$; $\langle \boldsymbol{x}, \boldsymbol{y}\rangle = \sum_{i=1}^d x_i y_i^*$ is the scalar product between two vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{C}^d$; the ℓ_p-norm of $\boldsymbol{x}$ ($p \geq 1$) is defined as $\|\boldsymbol{x}\|_p = (\sum_{i=1}^d |x_i|^p)^{1/p}$, with $\|\boldsymbol{x}\|_\infty = \max_i |x_i|$ and $\|\boldsymbol{u}\|_0 := |\operatorname{supp}(\boldsymbol{u})|$; $\mathbb{B}^n := \{\boldsymbol{u} \in \mathbb{C}^n : \|\boldsymbol{u}\| \leq 1\}$; the Hadamard product is $\odot$; and the angle operator (applied componentwise onto vectors) reads $\angle(ce^{\mathsf{i}\phi}) = \phi$ for $c > 0$ and $\phi \in [-\pi, \pi]$. An s-sparse $\boldsymbol{x}$ vector belongs to the set $\bar{\Sigma}_s^N := \{\boldsymbol{u} \in \mathbb{C}^N, \|\boldsymbol{u}\|_0 \leq s\}$, similarly $\tilde{\Sigma}_s^N = \bar{\Sigma}_s^N \cap \bar{\mathbb{B}}^N$. We define $\operatorname{sign}_{\mathbb{C}}$ as the Phase-Only operator, which removes the amplitude out of each component of the complex signal, $\operatorname{sign}_{\mathbb{C}}(\lambda) = \frac{\lambda}{|\lambda|}$ for $\lambda \in \mathbb{C}$. The window function is defined as $\Gamma_L(\cdot)$ with $\Gamma_L(l) = 1$ for $0 < l \leq L$, 0 elsewhere.

III. FMCW Radar Signal Model

In this section, before introducing the low-resolution acquisition model, we first focus on the sensor model from which the signals will be collected. This paper considers an FMCW radar with one transmitting and one receiving antenna. The radar's transmitting antenna emits a signal $s(t)$ modelled as

$$s^{RF}(t) = \sqrt{P_t}\exp\big(\mathsf{i}\,2\pi F_c(t) + \mathsf{i}\,\phi_0\big)\Gamma_{T_c}(t \bmod T_p), \quad (1)$$

where P_t is the transmitted power, $f_c(t)$ the transmitted frequency pattern, $F_c(t) := \int_0^t f_c(\xi)\mathrm{d}\xi$ the integrated frequency, and $\phi_0 \in [0, 2\pi]$ is the initial phase of the oscillator. The chirp has a duration T_c and is repeated at a *Pulse Repetition Frequency* (PRF) of $\frac{1}{T_p}$. The super-script $(\cdot)^{RF}$ is added to highlight the fact that this signal is centred around a carrier frequency (*e.g.*, f_0), which, in typical radar applications such as automotive, is in the order of tens of GHz.

The carrier frequency pattern $f_c(t)$ of an FMCW radar can be characterized as a saw-tooth function; see Fig. 1:

$$f_c(t) = f_0 + \frac{B}{T_c}(t \bmod T_p), \quad (2)$$

with f_0 the central frequency, and B the spanned bandwidth. Note that, in practical applications, B is not a design parameter but a constraint imposed by government regulations. Considering, for now, a simple ranging problem where the scene is made of only one target located at a range R from the radar with one receiving antenna. The received signal is a time-delayed version of the transmitted signal (1),

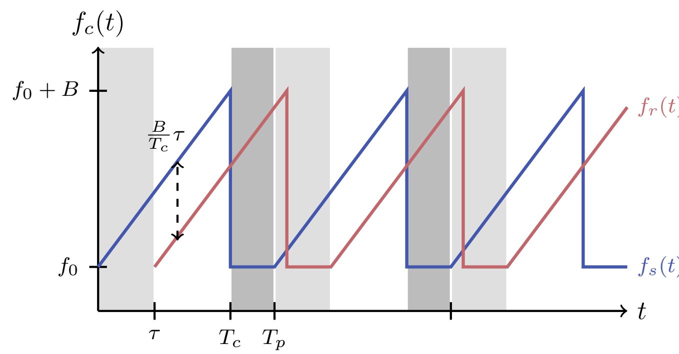

Fig. 1: Representation of the frequency content of the transmitted (**blue**) and received (**red**) chirps with a delay τ

$$r^{RF}(t) = \alpha s(t - \tau),$$

and is represented in Fig. 1 (see the **red** curve). The delay τ is simply the round-trip between the radar and the target and can be thus expressed as $\tau = \frac{2R}{\mathsf{c}}$, with c the speed of light. The variable $\alpha \in \mathbb{C}$ represents the amplitude of the received signal, this encompasses the Radar Cross Section (RCS) of the target but also all the other losses (*e.g.*, path losses, and gains in the acquisition chain, ...). For the sake of simplicity, in the rest of our presentation, the complex value α will refer to a global constant amplitude that may change from one line to the other in the description of the reception and demodulation processes. Indeed, each of the these stages is associated with its respective complex gain. In Fig. 1 interestingly, the delay τ in the time domain between $s(t)$ and $r(t)$ can also be observed as a frequency shift $\frac{B}{T_c}\tau$ in the frequency domain. From (1), the received analytical signal in radio frequency can be expressed as:

$$r^{RF}(t) = \alpha\exp\big(\mathsf{i}\,2\pi F_c(t-\tau) + \mathsf{i}\,\phi_0\big)\Gamma_{T_c}(t-\tau \bmod T_p). \quad (3)$$

The spectrum of (3), although narrow-band, is centred around f_0 which is, in the case of *Ka-band* radar, in the tens of

This article has been accepted for publication in IEEE Journal of Selected Topics in Signal Processing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/JSTSP.2025.3606387

GHz (*i.e.*, in Radio-Frequency). Directly sampling this signal is inadvisable as it would require a prohibitively high sampling rate that would then incur complex hardware and high costs. One possibility is to instead sample the signal around its central frequency $f_0 + \frac{B}{2}$ by multiplying the received signal in (3) by the carrier signal $\exp\big(-\mathrm{i}\,2\pi(f_0 + \frac{B}{2})t\big)$. The sampling requirements have lowered from the order of the carrier frequency (in GHz) to the order of the bandwidth B of the modulated signal, *i.e.*, in the order of the hundreds of MHz.

In essence, using this demodulation process, we are now only recording the changes that are imparted by the environment. However, sampling at multiple hundreds of MHz might still be prohibitively high compared to the observation made earlier in Fig. 1; for small ranges, the delay of interest manifests itself as a frequency shift between the received and transmitted signals. In order to leverage that fact, FMCW radars use coherent demodulation, where to recover this frequency shift, the received signal $r^{RF}(t)$ is demodulated using the transmitted signal $s^{RF}(t)$ itself. The architecture required to perform this coherent demodulation is presented in Fig. 2, where the I and Q channels represent the real and imaginary part of the demodulated received signal $r(t)$ respectively. The

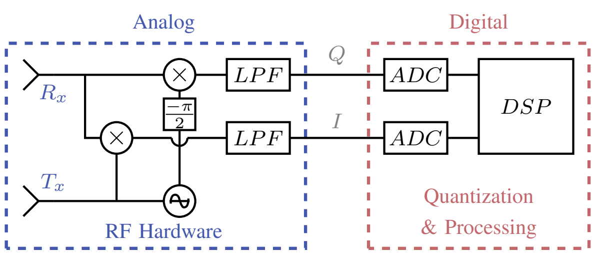

Fig. 2: Representation of a IQ coherent demodulation for an FMCW radar with one transmit and one receive antenna

process represented in Fig. 2, can be expressed mathematically simply as

$$r(t) = r^{RF}(t)(s^{RF}(t))^*. \tag{4}$$

Anticipating the rest of the section, notice that after the coherent demodulation, the signals are not in the radio-frequency domain any more; the $\cdot^{RF}$ superscript is dropped. Using the expression of the transmitted signal defined in (1), (4) reduces to

$$\begin{aligned} r(t) = \alpha \exp\Big(-\mathrm{i}\,2\pi \int_{t-\tau}^{t} f_c(\xi)\mathrm{d}\xi\Big) \\ \times \Gamma_{T_c}\big((t-\tau) \bmod T_p\big)\Gamma_{T_c}(t \bmod T_p). \end{aligned} \tag{5}$$

As mentioned earlier, assuming a delay $\tau \ll T_c$ small enough such that $\Gamma_{T_c}(t - \tau \bmod T_p) \approx \Gamma_{T_c}(t \bmod T_p)$ and using the saw-tooth model (2), the integral in (5) becomes

$$\int_{t-\tau}^{t} f_c(\xi)\mathrm{d}\xi = \int_{t-\tau}^{t} (f_0 + \frac{B}{T}\xi)\mathrm{d}\xi = \tau f_c(t) - \frac{B}{2T}\tau^2. \tag{6}$$

Combining (5) with (6) allows us to express the received signal $r(t)$ in base-band, *i.e.*, we have

$$r(t) = \alpha \exp\big(-\mathrm{i}\,2\pi\tau f_c(t)\big)\Gamma_{T_c}(t \bmod T_p), \tag{7}$$

where α also encompasses the static phase-shift $-\frac{B}{2T}\tau^2$ in (6). In words, (7) shows that the coherent demodulation expresses the time difference τ coming from the target as a frequency shift between the transmitted and received signals. This frequency shift linked to the range is represented in Fig. 1.

Interestingly, this frequency shift $\frac{B}{T_c}\tau$ is often several orders of magnitude below the bandwidth B. This fact motivates the use of coherent demodulation. The sampling frequency is now linked to $\frac{B}{T_c}\tau$ instead of B. For example, in the context of K-band radar with $f_0 = 24$GHz and $B = 250$MHz, a target at a range of 300m will generate a frequency shift of 200kHz for a PRF of 12kHz. Thanks to the property of the Fourier transform, the resolution δ_R can be simply expressed as $\delta_R = \frac{\mathsf{c}}{2B}$.

Considering a simple additive model, the received base-band signal from a scene with s targets can be expressed as

$$r(t) = \sum_{i}^{s} \alpha_i \exp\big(-\mathrm{i}\,2\pi\frac{2R_i}{\mathsf{c}} f_c(t)\big)\Gamma_{T_c}(t \bmod T_p), \tag{8}$$

where α_i and R_i are the received powers and ranges from each of the s targets measured in the scene. If the duration of the acquisition is restricted to $t \in [0, T_c]$, one can see that the problem of estimating the ranges of these s targets is tantamount to estimating the frequency content of $r(t)$ and by associating each peak in the frequency domain to a specific range.

Discretizing the ranges over an N sample grid (with sampling period δ_R) and sampling r as $r[j] := r(\frac{T_c}{N}j)$, $j \in [N]$, provides the following linear model $\boldsymbol{r} = \boldsymbol{\Psi x}$, where $\boldsymbol{x} \in \bar{\Sigma}_K^N$ (*i.e.*, x_j equals α_k if the j-th range bin contains the k-th target, and is 0 otherwise), $\boldsymbol{\Psi}$ is the $N \times N$ discrete Fourier matrix.

We can reach a compressive sensing scheme [41]

$$\boldsymbol{r} = \boldsymbol{\Phi x} \in \mathbb{C}^M,$$

by (randomly) subsampling $\boldsymbol{\Psi x}$ over $m \leq N$ frequencies, *i.e.*, $\boldsymbol{\Phi} := \boldsymbol{P\Psi}$ with a subsampling matrix $\boldsymbol{P} \in \{0,1\}^{m \times N}$. Since we consider quantized signals, we allow for oversampling, $m > N$, by accumulating $G = \lceil m/N \rceil$ FMCW ramps, and subsampling the last one over $m - (G-1)N < N$ frequencies, *i.e.*, $\boldsymbol{\Phi}^\top = [\boldsymbol{\Psi}^\top, \ldots, \boldsymbol{\Psi}^\top, (\boldsymbol{P\Psi})^\top]^\top$.

IV. Multiplicative Phase Dithering

A. Random Multiplicative Dither in Hardware

In this paper, instead of dithering the amplitudes of the received signal in the real and imaginary domain, we propose to dither their phases by multiplying the signal with a unitary signal with a random phase. The proposed quantization process can be represented as:

$$\mathcal{Q}^{\odot}(\boldsymbol{y}) = c_{\odot}\,\mathcal{Q}(\boldsymbol{y} \odot \boldsymbol{\xi}) \odot \boldsymbol{\xi}^*,$$

with $c_{\odot} = \frac{\pi}{4}$ and $\mathcal{Q}(\cdot) = \mathcal{Q}_1(\cdot)$. In hardware, the 1-bit ADCs acquire the phase dithered version of $\boldsymbol{y}$ represented by $\mathcal{Q}(\boldsymbol{y} \odot \boldsymbol{\xi})$. This complex binary signal now has a completely random phase that is then partially compensated in software by the application of the complex conjugate of the dither $\boldsymbol{\xi}$, *i.e.*, by applying $\cdot \odot \boldsymbol{\xi}^*$. The constant $c_{\odot}$ is added as a scaling factor in the processing (see Proposition 1). A complete study of the performance of this scheme is presented in Section VI.

This way of dithering solves the first hardware issue brought up by the previous section regarding the additive scheme. It is amplitude independent and can be thus applied without the knowledge of the signal's dynamic $\|\boldsymbol{y}\|_\infty$. Furthermore, dithering the phase can be achieved using multiple already existing radar system architectures. Using, for example, *Non-zero-if* demodulation [42], such as in Fig. 3, one needs to generate the dither $\xi^*(t)$, the radar transmits the original signal $s^{RF}(t)$ and then the received RF signal $r^{RF}(t)$ is demodulated by $s^{RF}(t)\xi^*(t)$. The coherent demodulation in (3) then becomes

$$r_{n0}(t) = r^{RF}(t)\big(s^{RF}(t)\xi^*(t)\big)^* = r(t)\xi(t);$$

sampling the time signal $r_{n0}(t)$ finally gives $\boldsymbol{r} \odot \boldsymbol{\xi}$.

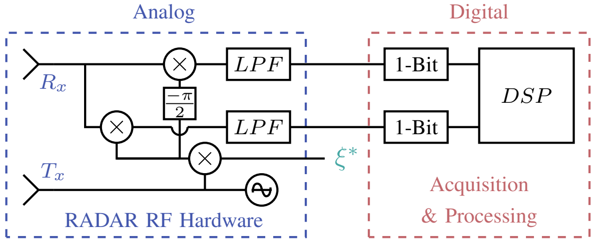

Fig. 3: Example of FMCW radar architecture with non zero IF demodulation and 1-bit quantization

Another way of randomly modifying the received signal's phase is to modify the signal after its coherent demodulation. In Fig. 4, the base-band signals are multiplied by the real and imaginary parts of the dither $\xi^*(t) = \xi^{\mathbb{R}}(t) + \mathrm{i}\,\xi^{\mathbb{I}}(t)$ and combined in the following way,

$$\begin{aligned} r(t)\xi(t) =& r^{\mathbb{R}}(t)\xi^{\mathbb{R}}(t) - r^{\mathbb{I}}(t)\xi^{\mathbb{I}}(t) \\ &+ \mathrm{i}\left(r^{\mathbb{R}}(t)\xi^{\mathbb{I}}(t) + r^{\mathbb{I}}(t)\xi^{\mathbb{R}}(t)\right). \end{aligned}$$

This process can be easily implemented with *off-the-shelf* components. Indeed, it only requires base-band components operating thus at low frequencies. These are both inexpensive and require low power compared to their RF counterparts. A future publication will study an actual 1-bit radar prototype based on this architecture. Compared to its additive counter-

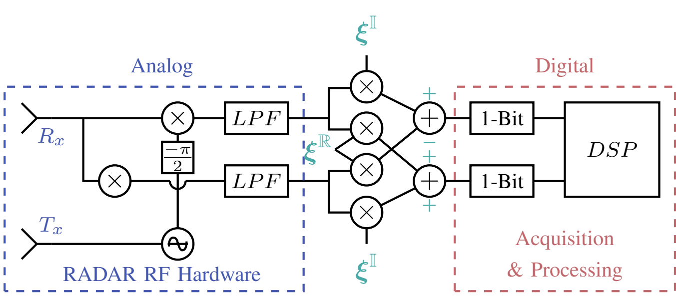

Fig. 4: Example of FMCW radar architecture with multiplicative dithering and 1-bit quantization

part, multiplicatively dithering the phase offers different ways of simplifying the hardware implementations of the dithering operation.

B. Relaxed Structured Multiplicative Dither

While the application of the dither has been simplified in the hardware architecture found in Fig. 3 and Fig. 4 compared to their additive counterparts, there is still some complexity involved in the generation of the dither $\boldsymbol{\xi}$ itself as well as how to leverage it in the reconstruction (*i.e.*, the software). Indeed, in both dithering methods, be it additive or multiplicative, the dither needs to follow a specific distribution that one must generate before applying it to the radar signals. Although the multiplicative dither does not require the knowledge of the signal's dynamic range, its random phase still requires the use of high-resolution DACs for its accurate generation.

We now propose to relax this constraint on the randomness of the phase of the dither by replacing it with a deterministic function that can be easily implemented in hardware. The unitary dither with a random phase is replaced by

$$\xi(t) := \exp\left(\mathrm{i}\,2\pi\Delta_f t + \chi\right). \tag{9}$$

This tremendously simplifies the hardware implementation. In the context of Fig. 4, one only needs to create two signals, namely $\cos(2\pi\Delta_f t)$ and $\sin(2\pi\Delta_f t)$ using a known frequency Δ_f. Sine and cosine functions are easily generated using based-band components without using high-resolution DACs [43].

Aside from the obvious hardware simplification this structured dither provides, it can also be leveraged in algorithms that enforce consistency between the measurements and the estimate, such as QIHT [23], [44], *i.e.*, removing contraints on the hardware also simplifies the software. Indeed, only the knowledge of Δ_f and the instants at which the signal is sampled (*e.g.*, t in (8)) are needed to recreate this dither and use it in the reconstruction. The reconstructed sparse signal will just be phase-shifted by the unknown phase χ, which is inconsequential in most radar applications. This also means there is no need for synchronisation or feedback between the dither, the radar and the processing.

The multiplicative dithering and its generation act as a completely separate system, simplifying its implementation even further compared to systems that require the control or knowledge of the dither to enforce consistency [23], [31]. It is also interesting to note that although this multiplication by a deterministic single tone dither might increase the frequency content of the acquired signal, its total bandwidth remains unchanged, which means the sampling frequency remains constant. This echoes the acquisition via undersampling of bandpass signals [45]. This deterministic dithering only works for models where increasing the number of measurements corresponds to acquiring (before the quantization) repetitions of the signals, like in ranging applications with FMCW radars, where this would correspond to sampling multiple chirps.

V. Link to and Limitations of Phase-Only Acquisition

While advantageous in its implementation, using a multiplicative dither, as we will see in the following sections, does not provide as strong theoretical guarantees as the one provided by the additive dithering. The performances of the additive dithering are upper-bounded by the classic high-resolution measurements, $\mathbb{E}\{\mathcal{Q}^{+}_{\epsilon}(\boldsymbol{y})\} = \boldsymbol{y}$ [37]. However, acquiring the measurements using $\mathcal{Q}^{\odot}(\cdot)$ does not capture all of the information about the signal of interest. This is best exhibited by

Proposition 1 (from [40])**.** *For $\boldsymbol{a} \in \mathbb{C}^m$ and $\xi_j \sim_{\text{i.i.d.}} e^{\mathrm{i}\mathcal{U}(0,2\pi)}$, $j \in [m]$, we have*

$$\mathbb{E}_{\xi}\mathcal{Q}^{\odot}(\boldsymbol{a}) = \operatorname{sign}_{\mathbb{C}}(\boldsymbol{a}).$$

Here lies the big difference with the additive scheme: when combined with one-bit quantization, in expectation, the multiplicative dither only capture the phase of the signal. One of its strengths, being independent of the amplitude, induces weaker reconstruction guarantees as part of the information about $\boldsymbol{\Phi x}$ is lost and cannot be retrieved regardless of the number of measurements m used.

Because the multiplicative dither is linked to the Phase-Only (PO) acquisition, we now focus in this section on the limitations that this acquisition process might encounter with Fourier-based measurements. For any algorithm to be able to reconstruct all sparse vectors from their measurements, be it quantized or other acquisition modality, these measurements need to be different from one another. This statement is obvious in the case of classic linear measurements where, if $\boldsymbol{\Phi} \in \mathbb{C}^{m\times N}$ follows the (ℓ_2,ℓ_2)-RIP$(\delta, 2s)$ [9], two different sparse vectors $\boldsymbol{x}_0 \neq \boldsymbol{x}_1 \in \Sigma^N_s$ once measured can be said to follow

$$(1-\delta)\|\boldsymbol{x}_0-\boldsymbol{x}_1\|_2^2 \le \tfrac{1}{m}\|\boldsymbol{\Phi x}_0-\boldsymbol{\Phi x}_1\|_2^2 \le (1+\delta)\|\boldsymbol{x}_0-\boldsymbol{x}_1\|_2^2,$$

which simply implies that $\boldsymbol{\Phi x}_0 \neq \boldsymbol{\Phi x}_1$.

This simple observation cannot be assumed to be true when using harsh deterministic quantizer (*e.g.*, $\mathcal{Q}_{\epsilon}(\cdot)$) [46].

Although the Phase-Only operator can be seen as having an infinite resolution in the phase domain where the other quantizers are limited to a fixed set of values (*e.g.*, number of quadrants), this acquisition can also suffer from ambiguous measurements similar to the one experienced by $\mathcal{Q}_{\epsilon}(\cdot)$ in the radar context in [23], [28].

For example, one can find 2 vectors $\boldsymbol{x}_0, \boldsymbol{x}_1 \in \mathbb{C}^N$, with $\boldsymbol{x}_0 \neq \boldsymbol{x}_1$, such that

$$\operatorname{sign}_{\mathbb{C}}(\boldsymbol{\Phi x}_0) = \operatorname{sign}_{\mathbb{C}}(\boldsymbol{\Phi x}_1), \tag{10}$$

which simply means that $\angle(\boldsymbol{\Phi x}_0) = \angle(\boldsymbol{\Phi x}_1)$. One obvious case is when the two signals are the same but with different amplitude, *e.g.*, $\boldsymbol{x}_0 = a\boldsymbol{x}_1$, with $a \in \mathbb{R}_+$. These ambiguous scenarios are, however, inconsequential in the setting of radar estimation as the absolute amplitude of reconstructed radar scene is rarely of interest compared to the relative power between elements of $\boldsymbol{x}$. Other pairs $(\boldsymbol{x}_0, \boldsymbol{x}_1)$ that satisfy (10) can be found using the properties of the Fourier transform. Indeed, any pairs of vectors $(\boldsymbol{x}_0, \boldsymbol{x}_1)$ will give ambiguous measurements if one vector can be expressed as the convolution of the other with a vector whose Fourier transform is real and positive. For two measured signals to have the same phase, there exists a vector $\boldsymbol{H} \in \mathbb{R}^m_+$, such that

$$\operatorname{sign}_{\mathbb{C}}(\boldsymbol{\Phi x}) = \operatorname{sign}_{\mathbb{C}}(\boldsymbol{H} \odot \boldsymbol{\Phi x}), \tag{11}$$

for a given vector $\boldsymbol{x} \in \mathbb{C}^N$. Given the properties of Fourier transforms, it means that this filter $\boldsymbol{H} \in \mathbb{R}^m_+$ can be expressed from the measurements domain to the signal domain as $\boldsymbol{H} \odot \boldsymbol{\Phi x} = \boldsymbol{\Phi}(\boldsymbol{h} * \boldsymbol{x})$, with $\boldsymbol{h} \in \mathbb{C}^N$ being the Fourier transform of $\boldsymbol{H}$.

Fig. 5 illustrates this ambiguous scenario for a 1-sparse vector $\boldsymbol{x}$ and a filter $H(t) = 1 + \alpha\cos(2\pi\frac{t}{3T_s})$, with $|\alpha| \le 1$ and $\frac{1}{T_s}$ the sampling frequency. Because all components of the resulting $\boldsymbol{H}$ are positive, the phase of the measured signals, $\boldsymbol{r}_0 = \boldsymbol{\Phi x}$ and $\boldsymbol{r}_1 = \boldsymbol{H} \odot \boldsymbol{\Phi x}$, are identical. The

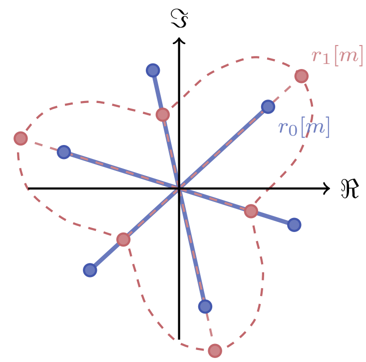

Fig. 5: Example of an ambiguous scenario where the PO measurements from the **blue** and **red** signals are identical

ambiguous scenario presented in the measurements domain in Fig. 5 is also represented in the signal domain in Fig. 6. The convolution between the vector $\boldsymbol{x}$ and $\boldsymbol{h}$ is represented in **red**, where the sparsity of the resulting vector has increased.

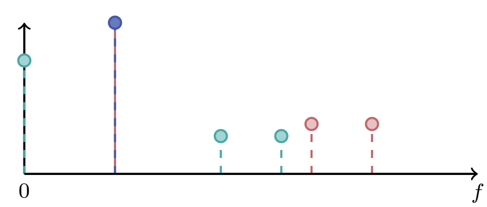

Fig. 6: Example in the frequency domain, of an ambiguous scenario where the PO measurements from the $\boldsymbol{x}$ and $\boldsymbol{h} * \boldsymbol{x}$ signals are identical, with the filter $\boldsymbol{h}$.

However, in realistic scenarios with radar systems, these ambiguous cases induced by Fourier-based systems rarely prevent the signal estimation in PO acquisition, and consequently, in the quantized model $\mathcal{Q}^{\odot}$ as well.

Indeed, while the theoretical guarantees of the type obtained for PO-CS in [47]–[49] for complex Gaussian matrices cannot be extended to Fourier measurements (because of (10)), as shown in the next proposition, we can leverage the fact that the power reflected by each target depends on different factors, to the point that their phases are often modelled as random [39].

Proposition 2. *If two vectors $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ in $\mathbb{C}^N$ have uniform, i.i.d. random phases, that is if $\boldsymbol{x}_0 = diag(\boldsymbol{\gamma}_0)\boldsymbol{u}_0$ and $\boldsymbol{x}_1 = diag(\boldsymbol{\gamma}_1)\boldsymbol{u}_1$ with amplitudes $\boldsymbol{u}_0, \boldsymbol{u}_1 \in \mathbb{R}^+_N$ and i.i.d. random phases $(\gamma_0)_j, (\gamma_1)_k \sim_{\text{i.i.d.}} \exp(\mathrm{i}\mathcal{U}(0,2\pi))$ for $j,k \in [N]$, then, provided that $\boldsymbol{u}_0, \boldsymbol{u}_1 \neq \boldsymbol{0}$,*

$$\mathbb{P}\big[\operatorname{sign}_{\mathbb{C}}(\boldsymbol{\Phi x}_0) = \operatorname{sign}_{\mathbb{C}}(\boldsymbol{\Phi x}_1)\big] = 0.$$

Proof. See App. C □

Besides this observation, notice that, from (11) and the structure of $\boldsymbol{h}$, it is clear the support size of $\boldsymbol{x}$ and $\boldsymbol{h} * \boldsymbol{x}$ will differ, Fig. 6. In the context of the estimation of sparse radars scenes, if the lower sparsity level signal is the one observed by the radar, then sparsity-promoting procedures, *e.g.*, hard-thresholding, will favour this vector over the larger sparsity one during the reconstruction. These two points make us confident that using 1-bit sensing with multiplicative phase dithering can provide a good trade-off between performance and the cost of implementation.

VI. Reconstruction Guarantee

We now study the reconstruction guarantees of the proposed quantization scheme using the Projected Back-Projection (PBP) algorithm. To that end, we compare the reconstruction of the dithered phase quantizer $\mathcal{Q}^{\odot}(\cdot)$ to the PO acquisition $\operatorname{sign}_{\mathbb{C}}(\cdot)$. We draw a parallel between these and the interplay between the 1-bit acquisition with additive dithering and the classic acquisition without quantization. The former tends to the latter for a high number of measurements [37]. In our case, $\mathcal{Q}^{\odot}(\cdot)$ is an unbiased estimator of the PO measurements, as seen in Lemma 1.

We start by defining the algorithm used for our study before introducing a global result about the recovery guarantee achieved by PBP when using Phase-Dithered 1-bit measurements. This result will then structure the rest of the section.

Definition (Projected-Back-Projection, PBP)**.** For a measurement vector $\boldsymbol{y} \in \mathbb{C}^m$ and given a linear model with a measurement matrix $\boldsymbol{\Phi} \in \mathbb{C}^{m\times N}$, one can estimate a s-sparse vector using

$$\hat{\boldsymbol{x}} = \tfrac{1}{m}\mathsf{H}_s\big(\boldsymbol{\Phi}^H\boldsymbol{y}\big),$$

where $\mathsf{H}_s(\cdot)$ is the hard-thresholding operator that keeps the s biggest elements in amplitude.

Given this algorithm, one can show the following reconstruction guarantee, which highlights the connection between the PO and $\mathcal{Q}^{\odot}$ measurements.

Theorem 1 (ℓ_2 recovery using PBP with 1-bit Phase Dither)**.** *For a given complex s-sparse vector $\boldsymbol{x} \in \bar{\Sigma}_s^N$, and a measurement matrix $\boldsymbol{\Phi} \in \mathbb{C}^{m\times N}$ that follows the (ℓ_2,ℓ_2)-RIP($2s$, δ) and with the PO ($\operatorname{sign}_{\mathbb{C}}(\cdot)$) and the 1-bit quantization with multiplicative phase dithering ($\mathcal{Q}^{\odot}(\cdot)$), the ℓ_2 reconstruction error using PBP can be bounded by*

$$\|\boldsymbol{x}-\hat{\boldsymbol{x}}\|_2 \le 2\|\boldsymbol{x} - \tfrac{1}{m}\big(\boldsymbol{\Phi}^H \operatorname{sign}_{\mathbb{C}}(\boldsymbol{\Phi}\boldsymbol{x})\big)_{\mathcal{T}}\|_2 + \epsilon,$$

with $\hat{\boldsymbol{x}}$ the PBP estimate, $\mathcal{T} := \operatorname{supp}(\boldsymbol{x}) \cup \operatorname{supp}(\hat{\boldsymbol{x}})$, and with $\epsilon = \mathcal{O}(m^{-\frac{1}{2}})$.

Proof. See App. A. □

Thm.1 formalises the intuition given in Proposition 1, which linked the two modalities in the measurement domain, into a reconstruction guarantee using the PBP algorithm. The inequality highlights two terms in Thm1, first the reconstruction error using PBP with Phase-Only measurements and then a ϵ degradation term between the PBP estimates from 1-bit measurements and from the PO measurements. Given that ϵ discrepancy follows $\mathcal{O}(m^{-\frac{1}{2}})$, the ℓ_2 degradation between the PO and multiplicative dithered measurements can be made arbitrarily small given a sufficiently high number of measurements.

The rest of the section focuses on each term of the bound presented in Theorem 1.

ℓ_2 reconstruction using PO measurements While there is no bound on the reconstruction of PBP using Phase-Only measurement, reconstructing signals from the phase of their Fourier measurements has been a subject of interest for decades. Oppenheim in various publications [50], [51] tackled this issue for images. They proposed a reconstruction algorithm that can reconstruct images from the phases of their Fourier transform. This algorithm, however, relies on the fact that the considered images are real. More recently, authors in [47], [48] studied the recovery guarantees that one can expect from Phase-Only measurement using the framework of Compressive Sensing. Their results applied to complex sparse vectors but assumed a linear model where the measurement matrix follows the (ℓ_1,ℓ_2)-RIP, which cannot be directly applied to Fourier-based measurements. The reconstruction of complex sparse vectors from PO Fourier measurements is still an open and interesting question. However, and as will be highlighted in the simulations in Section VII, the theoretical limitations of Phase-Only acquisition do not mean that the reconstructions from these measurements are of poor quality.

Study of the degradation between PO and $\mathcal{Q}^{\odot}$ We now provide more details on the ϵ discrepancy introduced in Thm.1. Let us now study a non-uniform bound on the ℓ_2-degradation between the PO and 1-bit with multiplicative dithering. As highlighted in the previous section, the Phase-Only acquisition, using Fourier-based measurements, cannot guarantee uniform reconstruction results. Consequently, we only focus on a non-uniform guarantee for the difference between the Phase-Only and 1-bit quantization with multiplicative dithering.

For the coming proof (whose details are provided in the appendices), we consider the canonic multiplicative dither whose phase follows a uniform distribution *i.e.*, $\angle\boldsymbol{\xi} \sim \mathcal{U}^m(0,2\pi)$. The difference between this unstructured random dither and the structured single tone dithering introduced in Sec. IV will be thoroughly studied in the simulations in Section VII.

Theorem 2. *For a given complex s-sparse vector $\boldsymbol{x} \in \bar{\Sigma}_s^N$, and a measurement matrix $\boldsymbol{\Phi} \in \mathbb{C}^{m\times N}$ that follows the (ℓ_2,ℓ_2)-RIP($2s$, δ) and with the PO ($\operatorname{sign}_{\mathbb{C}}(\cdot)$) and the 1-bit quantization with multiplicative dithering ($\mathcal{Q}^{\odot}(\cdot)$), and for all support set $\mathcal{T}$ of size $2s$, we get the following bound with probability exceeding $1-\exp\left(-\frac{m\gamma^2}{4c_{\odot}^2}\right)$,*

$$\|\big(\boldsymbol{\Phi}^H(\operatorname{sign}_{\mathbb{C}}(\boldsymbol{y}) - \mathcal{Q}^{\odot}(\boldsymbol{y}))\big)_{\mathcal{T}}\|_2 < 2m\gamma, \tag{12}$$

provided $m \ge 8\gamma^{-2}s(\log(\frac{eN}{2s}) + 2(1+\frac{\pi}{\gamma\sqrt{2}}))$.

Proof. See App. B. □

The bound in (12) combined with the condition imposed on the number of measurement m shows that the discrepancy between the reconstruction using quantized and PO measurements follows $\mathcal{O}(m^{-\frac{1}{2}})$. As mentioned above, this means that this discrepancy between the two (linked) quantizations can be controlled by modifying the number of measurements m.

VII. Simulations

We now assess the performances of the developed multiplicative scheme by carrying out Monte Carlo simulations. 100 runs are performed for each set of parameters. At each run, a s-sparse vector is generated. The support is set according to a uniform distribution; the amplitudes of each non-zero component are then generated as random uniform variables and given a random phase. The resulting s-sparse vector is then normalized. The linear measurements are generated by the multiplication of these sparse vectors by matrices made of elements of a Fourier transform. We define the dimensions of the sparse vector $\boldsymbol{x} \in \mathbb{C}^N$ with $N = 256$, and we vary the number of measurements as $m = \mu N$ with $\mu \in [2^{-6}, 2^4]$. When $\mu \leq 1$, we sub-sample the Fourier transform, which, in the context of the FMCW radar model, amounts to sub-sampling the elements of one demodulated chirp. When $\mu > 1$, whole repetitions of the Fourier matrix are taken, corresponding thus to multiple consecutive chirps. Five different acquisitions are compared: the classic linear acquisition (●), the 1-bit acquisition with additive dithering $\mathcal{Q}_\epsilon^+$ (⊕) where the additive dithering is generated according to $\boldsymbol{\xi} = \boldsymbol{\xi}_\mathbb{R} + \imath\boldsymbol{\xi}_\mathbb{I}$, with both $\boldsymbol{\xi}_\mathbb{R}$ and $\boldsymbol{\xi}_\mathbb{I}$ following $\mathcal{U}^m(-\frac{\epsilon}{2}, \frac{\epsilon}{2})$, the Phase-Only acquisition $\mathrm{sign}_\mathbb{C}$ (⊖), and 1-bit acquisition with multiplicative phase dithering $\mathcal{Q}^\odot$ (⊙). These 1-bit schemes with dithering are also compared with the deterministic 1-bit acquisition $\mathcal{Q}_\epsilon$ (○). Those quantized data are then processed with two different algorithms, namely PBP and QIHT, introduced hereafter.

Definition ((Q)-Iterative-Hard-Thresholding, (Q)IHT [11])**.** The k^{th} iteration of the (Q)IHT algorithm is expressed as

$$\hat{\boldsymbol{x}}^k = \mathsf{H}_s\big(\hat{\boldsymbol{x}}^{k-1} + \tfrac{1}{m}\boldsymbol{\Phi}^H(\boldsymbol{y} - \mathcal{A}(\boldsymbol{\Phi}\hat{\boldsymbol{x}}^{k-1}))\big),$$

where $\hat{\boldsymbol{x}}^0$ is the PBP estimate and $\mathcal{A}(\cdot)$ is the considered quantizer (*e.g.*, $\mathcal{Q}^+$,$\mathcal{Q}^\odot$, and others).

We use a consistency based stopping criterion of QIHT. The algorithm will run between 20 and $100s$ iterations with an early stop when the consistency defined as $m^{-1}\sum_k(\mathcal{A}(\boldsymbol{\Phi}_k\hat{\boldsymbol{x}}) = y_k)$ exceeds 95% or actually decreases from the previous iteration.

A. ℓ_2 reconstruction

Results for PBP Fig. 7 compares the different acquisition schemes using the PBP algorithm for a sparsity level of $s = 10$. One can first see that the additive dithering behaves as expected by the theory and has a ℓ_2-error that decreases as $\mathcal{O}(m^{-\frac{1}{2}})$ [37]. While this behaviour guarantees an arbitrarily low reconstruction error for a number of measurements m high enough for a low number of measurements (*e.g.*, $m \leq 2^4N$), it is clearly outperformed by the other low-resolution schemes. The multiplicative phase dithering achieves a better reconstruction quality than the additive dithering for any number of measurements presented in Fig. 7. For $m = N$, the phase dithered quantizer exhibit a gain of 6dB. For $\mu \leq 1$, it closely follows the non-dithered curve ($\mathcal{Q}_\epsilon$) before outperforming it and resolving to the performances of the Phase-Only acquisition. This is consistent with Proposition 1 and the reconstruction bound developed in Theorem 1. The only caveat to these performances is that for higher values of m, the additive scheme reaches the performances of the linear classic acquisition without quantization, thus possibly reaching a perfect reconstruction while the multiplicative scheme's performances will still be bound to the one of the PO. It is important to note that increasing the number of measurements also has a price in terms of data transfer. Thus, the quality of the reconstruction of the multiplicative scheme combined with the relatively low values of m required for its success makes it an appealing alternative.

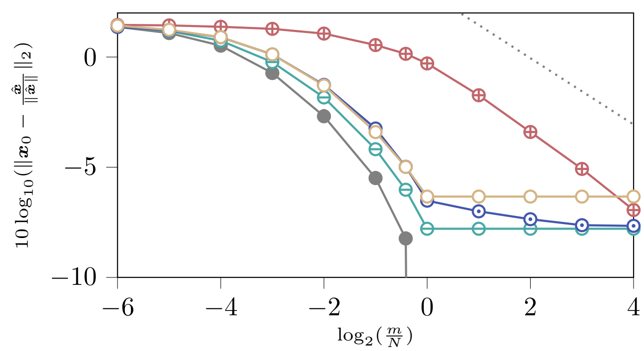

Fig. 7: Comparison of different reconstructions using PBP between different quantization schemes for $s = 10$, 1-bit with additive dither (—⊕—); 1-bit without dither in (—○—); 1-bit with multiplicative dither (—⊙—); Phase-Only acquisition (—⊖—); without quantization (—●—); (⋯⋯) represents $\mathcal{O}(m^{-\frac{1}{2}})$.

Results for Consistency Enforcing Algorithm In Fig. 8, we now study the performances of an iterative algorithm, namely QIHT. One can observe that the PO acquisition used in conjunction with the QIHT algorithm yields high-quality reconstructions. Consequently the multiplicative dithering that has its performances bounded by the PO reconstruction, outperforms the additive dither for any number of measurements below 2^4N. Indeed, it follows the deterministic quantization and then continues to exhibit a rate of $\mathcal{O}(m^{-1})$ while the deterministic quantization plateaus at $\mu \geq 1$. It is interesting to see that this rate of $\mathcal{O}(m^{-1})$ has also been observed in other 1-bit settings but with stronger theoretical guarantees [46]. The trade-off between strong theoretical guarantees and performances is clearly highlighted. Our proposed scheme outperforms the additive dithering approach, exhibiting a 4dB gain while boasting a low-complexity hardware dither that does not require the knowledge of the signal's dynamic range compared to the additive dithering procedure.

Impact of the Hardware Relaxed Dither The simulations in Fig. 7 and Fig. 8 showcase the performances of an ideal multiplicative dithering that has a random phase distributed uniformly in the complex domain. In Fig. 9, we compare

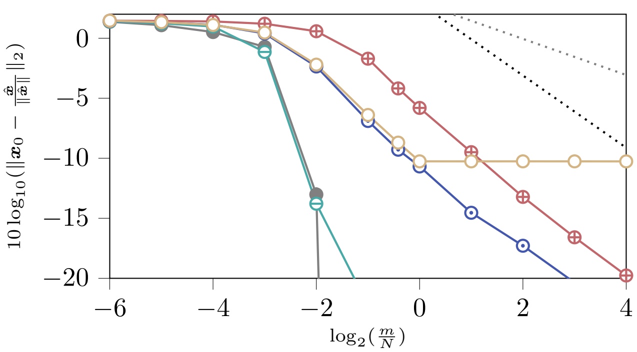

Fig. 8: Comparison of different reconstructions using QIHT between different quantization schemes for $s = 10$, 1-bit with additive dither (—⊕—) in **red**; 1-bit without dither (—○—) in **yellow**; 1-bit with multiplicative dither (—⊙—) in **blue**; Phase-Only acquisition (—⊖—) in **green**; without quantization (—●—) in **gray**; the dotted curve (⋯⋯) in **gray** represent $\mathcal{O}(m^{-\frac{1}{2}})$, the black dotted curve (⋯⋯) $\mathcal{O}(m^{-1})$.

this random dithering with the proposed structured single tone complex exponential defined in (9). At each run, the frequency of this complex exponential Δ_f is chosen at random as $\mathcal{U}([-\frac{1}{2T_s}, \frac{1}{2T_s}])$. This known Δ_f is then used to enforce the consistency between the estimated and received radar signal, as only this value is needed to reconstruct the dither used in the quantization. The curves in Fig. 9 are almost identical. This shows that the hardware implementation can be further simplified by removing the randomness of the phase and simply using *off-the-shelf* sine and cosine generators at a known frequency. One can now use consistency-promoting algorithms without actually measuring or controlling exactly the dither that is applied to the measurements before the quantization thanks to the proposed Hardware-Relaxed Phase-Dither approach.

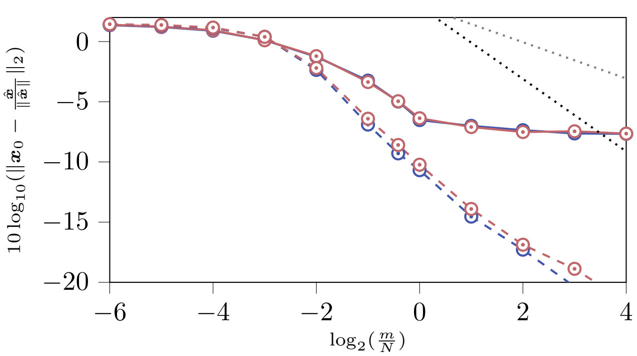

Fig. 9: Comparison between the random dithering using PBP (—⊙—) and QIHT (- ⊙ -) in **blue** and the deterministic and structured dither using PBP (—⊙—) and QIHT (- ⊙ -)in **red** for $s = 10$; for the PBP algorithm in solid; and QIHT in dashed; the dotted curve (⋯⋯) in **gray** represent $\mathcal{O}(m^{-\frac{1}{2}})$, the black dotted curve (⋯⋯) $\mathcal{O}(m^{-1})$.

Impact of Unknown Dynamic Range In Fig. 10, the resilience of the different 1-bit quantization schemes is tested against the imperfect knowledge of the dynamic of the signal Φx, *i.e.*, $\|\Phi x\|_\infty$. As mentioned in Section IV, the additive dithering procedure requires an estimate of the signal's dynamic range in order to perform according to the theory developed in [37]. The quality of this estimate thus directly impact its performance. In contrast, our proposed scheme, multiplicative phase dithering, does not require such priors. To mimick the effect of the imperfect knowledge of the signal's dynamic, the size of the additive dither ϵ is randomly varied following $\frac{\epsilon}{2} = \|\Phi x\|_\infty 10^\beta$, with $\beta \sim \mathcal{U}(-1, 1)$. One can

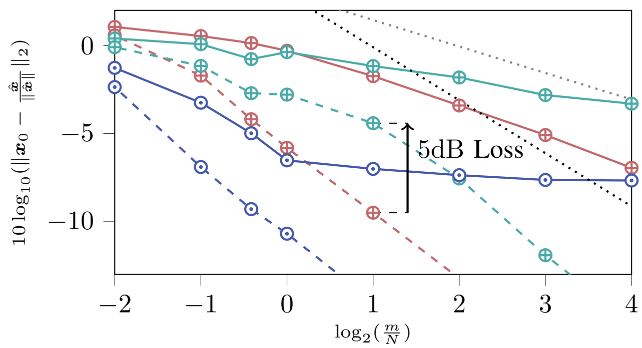

Fig. 10: Comparison for $s = 10$ of different 1-bit scheme; for a *a-priori* known dynamic range with $\mathcal{Q}^+$ using PBP (—⊕—) and QIHT (- ⊕ -), the results stemming from an unknown dynamic range using PBP (—⊕—) and QIHT (- ⊕ -), the 1-bit with multiplicative dither using PBP (—⊙—) and QIHT (- ⊙ -); the gray dotted curve (⋯⋯) represents $\mathcal{O}(m^{-\frac{1}{2}})$, the black dotted curve (⋯⋯) $\mathcal{O}(m^{-1})$.

see that this small unknown on the dynamic range induces a deterioration of almost 5dB on the additive scheme for both the PBP and QIHT algorithms. This highlights clearly the robustness provided by the multiplicative phase dither. As only the phase is impacted, the scheme does not require nor suppose anything on the measured signal's amplitude, simplifying the implementation of this scheme again in an actual acquisition board.

B. Support Recovery Performances

The previous simulations focused on the reconstruction error using the ℓ_2-norm. In some applications, however, recovering the position of the different targets is the only focus, not the actual value of the RCS of the different targets. To that end, we use the TPR (True Positive Rate) to assess the support recovery. The TPR is defined as

$$\text{TPR} = \frac{|\{i \in [N] \text{ s.t. } |\hat{x}_i| > 0 \text{ and } |x_i| > 0\}|}{s}.$$

In Fig. 11 and Fig. 12, we compare the TPR for the different algorithms and 1-bit acquisition processes for $s = 4$ and $s = 20$. As seen in [23], the additive scheme is not robust to an increase of the sparsity for a low number of measurements. For both algorithms, the multiplicative dither, again, follows the deterministic quantization before outperforming it for a higher number of measurements. Using this multiplicative dither, according to the simulations, ensures that the reconstruction cannot be worse than the deterministic quantization and can see up to a $10-15\%$ increases in TPR, especially using QIHT in Fig. 12.

VIII. Radar Measurements

To assess the quality of the proposed scheme, we used the dataset featured in [23]. In this paper, the authors used

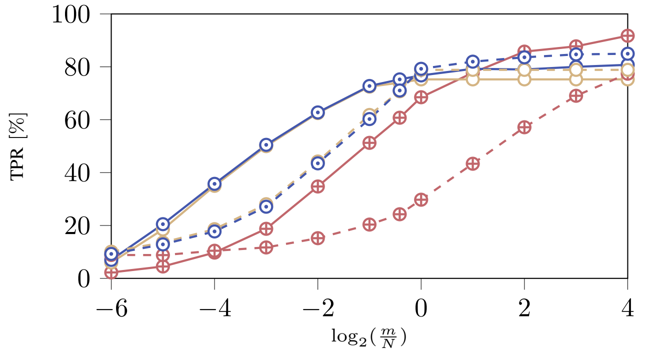

Fig. 11: Comparison of PBP in terms of TPR for $s = 4$ solid, $s = 20$ in dashed; 1-bit with additive dither (—⊕— & - ⊕ -) ; 1-bit without dither (—○— & - ○ -); 1-bit with multiplicative structured dither (—⊙— & - ⊙ -).

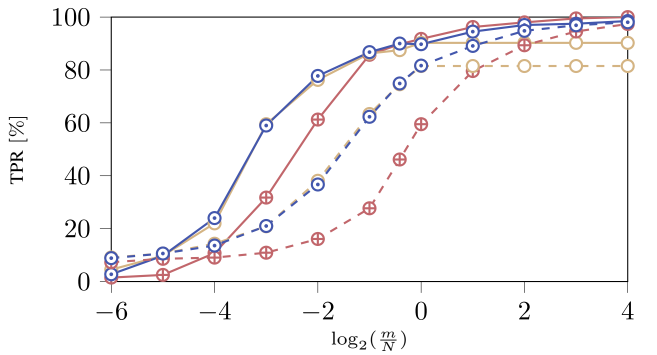

Fig. 12: Comparison of QIHT in terms of TPR for $s = 4$ solid, $s = 20$ in dashed, 1-bit with additive dither (—⊕— & - ⊕ -); 1-bit without dither (—○— & - ○ -); 1-bit with multiplicative structured dither (—⊙— & - ⊙ -).

a 24GHz radar made by RF-Beam[2] to measure the signals reflected by 2-target simulators[3]. These devices are able to reflect the signal transmitted by the radar back to itself with a specific delay and gain, thus simulating targets at specific ranges.

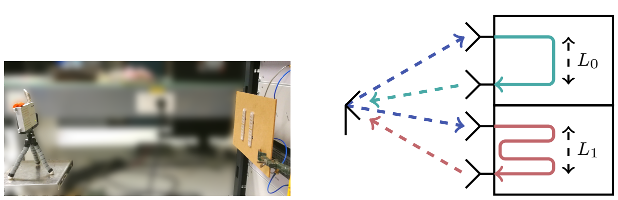

(a) (b)

Fig. 13: Radar measurements set-up: (a) the KMD2 radar in front of the target simulator; (b) its functional representation.

A complete explanation of the set-up can be found in [23]. To study the effect of a high number of targets, these 2-sparse measurements are combined to generate realistic high-sparsity measurements by assuming a simple additive model. For each sparsity level, a 100 Monte-Carlo runs are performed. Where, in order to simulate a s-sparse vector, $\frac{s}{2}$ experiments with disjoint supports are combined. The vector $\boldsymbol{x}$ is defined as the back-projection of the mean of the linear measurements to which an oracle hard-thresholding is applied. The support is simply defined as the ranges set on the target simulator. Indeed, only a control of the range and of the relative power is available; the phase and amplitude are observed *a posteriori*.

ℓ_2 Error In Fig. 14, we recreated the setting of Fig. 7 using the same parameters and reconstruction algorithm, *i.e.*, PBP. The observations made in the simulations, in Fig. 7 and Fig. 8, are confirmed using real measurements. Indeed, there is a stark degradation of the additive scheme for high sparsity level ($s = 10$) compared to the multiplicative dithering for a low number of measurements (*e.g.*, $m \leq 4N$). Although done in a lab setting using target simulators, these measurements are not exactly noiseless, as seen by the ℓ_2-reconstruction reached by the classical linear measurements, which saturate around -7dB, giving thus an approximation of the SNR during the acquisition. This value limits the performance of the other schemes and explains why the non-dithered scheme varies with the number of measurements. Indeed, the noise still acts as an imperfect dither and thus impacts the reconstruction. Future publication will study theoretically the impact of the noise on this multiplicative phase dithering acquisition scheme. Although this noise is not, for now, accounted for in our acquisition model, the reconstruction from phase dithered measurements follows the behaviour of Fig. 7 with the gain from the repetition (*i.e.*, $\mu \geq 1$) with no adverse effect of the noise.

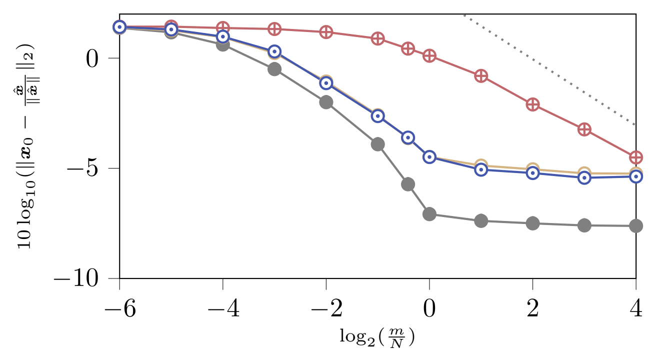

Fig. 14: ℓ_2 reconstruction for $s = 10$ with PBP using actual radar measurements, 1-bit with additive dither (—⊕—); 1-bit without dither (—○—); 1-bit with multiplicative structured dither (—⊙—) ; without quantization (—●—); the dotted curve (⋯⋯) represents $\mathcal{O}(m^{-\frac{1}{2}})$.

TPR Error Similarly to the previous section, in Fig. 15, the different acquisitions are compared using the TPR. Again the observations made in Fig. 11 are confirmed here. For low sparsity levels, the multiplicative dither outperforms the deterministic quantization. Increasing the sparsity level to $s = 20$, one can also observe a drop in performances of the additive scheme that is not experienced by the multiplicative one that only resolves in the worst case — as also observed in Fig. 14 — to the deterministic 1-bit quantizer.

The results showcased in this section clearly show that using a multiplicative dither provides the best trade-off between hardware requirements and reconstruction performances. We confirmed in a practical setting that the proposed scheme benefits from both the robustness of the deterministic scheme for high sparsity levels, and from high-quality reconstructions

[2]https://rfbeam.ch/product/k-md2-engineering-sample/
[3]AMG-043-007, https://www.amg-microwave.com/-Test-et-Mesure-

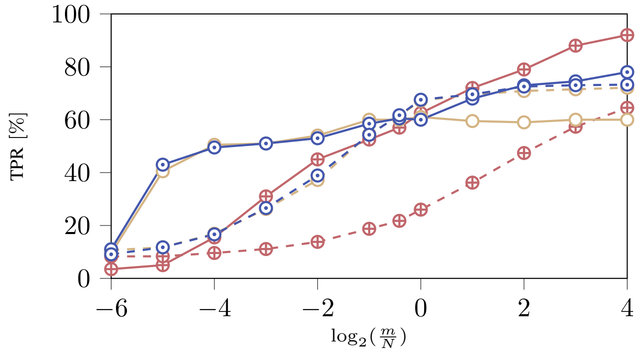

Fig. 15: Comparison of different TPR using actual radar measurements, for $s = 4$ solid, $s = 20$ in dashed, 1-bit with additive dither (—⊕— & - ⊕-); 1-bit without dither (—○— & - ○-); 1-bit with multiplicative structured dither (—⊙— & - ⊙-).

similar to the additive dithering but for a lower number of measurements.

IX. Conclusions & Discussion

This paper introduces the multiplicative phase dithering and its associated quantizer $\mathcal{Q}^{\odot}$. For radar systems, we further showed that it is a viable trade-off between performances, hardware complexity and theoretical guarantees. Indeed, this work stems from the fact that using an additive dither dramatically increases the complexity of the hardware, negating the gains of the low-resolution 1-bit acquisition. The multiplicative dithering of radar signals can be implemented efficiently, relies on already existing hardware architectures and does not require any side information on the signal's dynamic range. Although the theory uses a dither with random phase, we showed that this condition can be relaxed to a single-tone complex exponential, simplifying the architecture further and relaxing the conditions on reconstruction algorithms that use consistency. The developed theory showed that the reconstruction error of this scheme is linked to the one of the Phase-Only acquisition. Moreover, a non-uniform bound on the discrepancy between these two modalities was also studied and shown to behave as $\mathcal{O}(m^{-\frac{1}{2}})$. Next, the simulations showed that the proposed multiplicative dithering 1-bit scheme provides a good trade-off between the number of measurements needed and the sensitivity to the sparsity level of the signal. These results were then confirmed using real radar measurements.

Future works will focus on this paper's two main aspects: the proposed scheme's practicality and the theoretical guarantees and limitations of the PO acquisition. The architecture for the multiplicative dithering presented in Fig. 4 will be implemented in a radar prototype and its effectiveness demonstrated. The model's extension to the arrival angle will also be studied as well as its robustness to noise. This paper also showed that uniform reconstruction guarantees for Fourier-based measurements using the Phase-Only acquisition cannot be developed. This was shown by highlighting that PO measurements can be ambiguous for pairs of sparse vectors. It would be, however, interesting to study this problem further by developing guarantees that can properly deal with these ambiguities.

Appendix A
Proof of Theorem 1

Theorem 1. *For a given complex s-sparse vector $\boldsymbol{x} \in \bar{\Sigma}_s^N$, and a measurement matrix $\boldsymbol{\Phi} \in \mathbb{C}^{m\times N}$ that follows the (ℓ_2, ℓ_2)-RIP($2s$, δ) and with the PO ($\mathrm{sign}_{\mathbb{C}}(\cdot)$) and the 1-bit quantization with multiplicative phase dithering ($\mathcal{Q}^{\odot}(\cdot)$), the ℓ_2 reconstruction error using PBP can be bounded by*

$$\|\boldsymbol{x} - \hat{\boldsymbol{x}}\|_2 \le 2\|\boldsymbol{x} - \tfrac{1}{m}\big(\boldsymbol{\Phi}^H \mathrm{sign}_{\mathbb{C}}(\boldsymbol{\Phi x})\big)_{\mathcal{T}}\|_2 + \epsilon,$$

with $\hat{\boldsymbol{x}}$ the PBP estimate, $\mathcal{T} := \mathrm{supp}\,(\boldsymbol{x}) \cup \mathrm{supp}\,(\hat{\boldsymbol{x}})$, and with $\epsilon = \mathcal{O}(m^{-\frac{1}{2}})$.

Proof. For a support $\mathcal{S}$ defined such that $\mathrm{supp}(\hat{\boldsymbol{x}}) = (\boldsymbol{\Phi}^H \boldsymbol{y})_{\mathcal{S}}$, with $\boldsymbol{y} = \mathcal{Q}^{\odot}(\boldsymbol{\Phi x})$, the PBP reconstruction can be expressed as:

$$\begin{aligned}\|\boldsymbol{x} - \hat{\boldsymbol{x}}\|_2 =& \|\boldsymbol{x} - \tfrac{1}{m}\big(\boldsymbol{\Phi}^H \mathcal{Q}^{\odot}(\boldsymbol{\Phi x})\big)_{\mathcal{S}}\|_2 \\ \le& \|\boldsymbol{x} - \frac{1}{m}(\boldsymbol{\Phi}^H \mathcal{Q}^{\odot}(\boldsymbol{\Phi x}))_{\mathcal{T}}\|_2 \\ +&\frac{1}{m}\|(\boldsymbol{\Phi}^H \mathcal{Q}^{\odot}(\boldsymbol{\Phi x}))_{\mathcal{S}} - (\boldsymbol{\Phi}^H \mathcal{Q}^{\odot}(\boldsymbol{\Phi x}))_{\mathcal{T}}\|_2 \\ \le& 2\|\boldsymbol{x} - \tfrac{1}{m}\big(\boldsymbol{\Phi}^H \mathcal{Q}^{\odot}(\boldsymbol{\Phi x})\big)_{\mathcal{T}}\|_2 \end{aligned} \tag{13}$$

with $\mathcal{T} := \mathrm{supp}\,(\boldsymbol{x}) \cup \mathcal{S}$. The upper-bound in (13) uses the triangle inequality and the fact that, compared to $\boldsymbol{x} \in \bar{\Sigma}_s^N$, $(\boldsymbol{\Phi}^H \mathcal{Q}^{\odot}(\boldsymbol{\Phi x}))_{\mathcal{S}}$ is the best s-term approximation of $(\boldsymbol{\Phi}^H \mathcal{Q}^{\odot}(\boldsymbol{\Phi x}))_{\mathcal{T}}$. Furthermore, we can show that

$$\begin{aligned}\|\boldsymbol{x} - \hat{\boldsymbol{x}}\|_2 \le& 2\|\boldsymbol{x} - \tfrac{1}{m}\big(\boldsymbol{\Phi}^H \mathcal{Q}^{\odot}(\boldsymbol{\Phi x})\big)_{\mathcal{T}}\|_2 \\ \le& 2\|\boldsymbol{x} - \tfrac{1}{m}\big(\boldsymbol{\Phi}^H \mathrm{sign}_{\mathbb{C}}(\boldsymbol{\Phi x})\big)_{\mathcal{T}}\|_2 \\ +& \tfrac{2}{m}\|\big(\boldsymbol{\Phi}^H (\mathrm{sign}_{\mathbb{C}}(\boldsymbol{\Phi x}) - \mathcal{Q}^{\odot}(\boldsymbol{\Phi x}))\big)_{\mathcal{T}}\|_2. \end{aligned} \tag{14}$$

We then use Thm 2 on the second term in (14) to finish the proof. □

Appendix B
Proof of Theorem 2

We first develop a lemma that, along with Prop 1, will be used in the proof of Thm 2.

Lemma 1 (Bound related to phase error)**.** *For $\mathcal{Q}^{\odot}(\cdot)$ and $\mathrm{sign}_{\mathbb{C}}(\cdot)$ for any $a \in \mathbb{C}$, one can bound*

$$|\mathcal{Q}^{\odot}(a) - \mathrm{sign}_{\mathbb{C}}\,(a)| \le \frac{\pi}{4}.$$

Proof. Let us define, without loss of generality, $a = e^{j\phi}$ and $\xi = e^{j\psi}$, one can first bound

$$\begin{aligned}|\angle\mathcal{Q}^{\odot}(a) - \angle\,\mathrm{sign}_{\mathbb{C}}\,(a)| &= |\angle\mathcal{Q}(e^{j\phi+j\psi})e^{-j\psi} - \angle e^{j\phi}| \\ &= |\angle\mathcal{Q}(e^{j\phi'}) - \angle e^{j\phi'}| \le \frac{\pi}{4},\end{aligned}$$

with $\phi' := \phi + \psi$.

Considering now a phase difference of β such that $\angle\mathcal{Q}^{\odot}(a) - \angle\mathrm{sign}_{\mathbb{C}}\,(a) = \beta$, one can bound

$$|\mathcal{Q}^{\odot}(a) - \mathrm{sign}_{\mathbb{C}}\,a| = |\sqrt{2}\frac{\pi}{4}e^{\mathrm{i}\beta} - 1| = 1 + 2\left(\frac{\pi}{4}\right)^2 - 2\sqrt{2}\frac{\pi}{4}\cos\beta.$$

Setting $\beta = \frac{\pi}{4}$ finally gives the desired bound

$$|\mathcal{Q}^{\odot}(a) - \mathrm{sign}_{\mathbb{C}}(a)| = 1 - (1 - \frac{\pi}{4})\frac{\pi}{2} \le \frac{\pi}{4}. \quad (15)$$

□

Upper bounding (15) by $\frac{\pi}{4}$ is made out of convenience and to highlight the fact that $\mathcal{Q}^{\odot}(\cdot)$ effectively quantizes the phase of the signal, not its amplitude.

Lemma 1 and Prop 1 combined with *Hoeffding*'s inequality, are used in the following theorem to upper-bound the ℓ_2 discrepancy between the multiplicative and PO reconstruction using PBP.

Theorem 2. *For a given complex s-sparse vector $\boldsymbol{x} \in \tilde{\Sigma}_s^N$, and a measurement matrix $\boldsymbol{\Phi} \in \mathbb{C}^{m\times N}$ that follows the (ℓ_2, ℓ_2)-RIP($2s$, δ) and with the PO ($\mathrm{sign}_{\mathbb{C}}(\cdot)$) and the 1-bit quantization with phase dithering dithering ($\mathcal{Q}^{\odot}(\cdot)$), and for all support set $\mathcal{T}$ of size $2s$, we get the following bound with probability exceeding $1 - \exp\left(-\frac{m\gamma^2}{4c_{\odot}^2}\right)$,*

$$\|\big(\boldsymbol{\Phi}^H(\mathrm{sign}_{\mathbb{C}}(\boldsymbol{y}) - \mathcal{Q}^{\odot}(\boldsymbol{y}))\big)_{\mathcal{T}}\|_2 < 2m\gamma,$$

provided $m \ge 8\gamma^{-2}s(\log(\frac{eN}{2s}) + 2(1 + \frac{\pi}{\gamma\sqrt{2}}))$.

Proof. Starting with

$$\|\big(\boldsymbol{\Phi}^H(\mathrm{sign}_{\mathbb{C}}(\boldsymbol{y}) - \mathcal{Q}^{\odot}(\boldsymbol{y}))\big)_{\mathcal{T}}\|_2 = \sup_{\substack{\boldsymbol{u}\in\mathbb{B}^N \\ \mathrm{supp}(\boldsymbol{u})\in\mathcal{T}}} \langle \boldsymbol{\Phi u}, \mathrm{sign}_{\mathbb{C}}(\boldsymbol{y}) - \mathcal{Q}^{\odot}(\boldsymbol{y})\rangle. \quad (16)$$

Using Prop 1 one can leverage Hoeffding's inequality to upper-bound (16). Each quantized i^{th} element of the scalar product are bounded by $|\boldsymbol{\Phi}_i \boldsymbol{u}\mathcal{Q}^{\odot}(\boldsymbol{y}_i)| \le |\boldsymbol{\Phi}_i\boldsymbol{u}|c_{\odot}\sqrt{2}$. For a given $\boldsymbol{u}$, (16) becomes

$$\mathbb{P}\{\|(\boldsymbol{\Phi}^T(\mathrm{sign}_{\mathbb{C}}(\boldsymbol{y}) - \mathcal{Q}^{\odot}(\boldsymbol{y})))_{\mathcal{T}}\|_2 > m\gamma\} \le 2\exp\left(-\frac{m^2\gamma^2}{\|\boldsymbol{\Phi u}\|_2^2 c_{\odot}^2}\right).$$

Leveraging the $(\ell_2$,$\ell_2)$-RIP($2s$, δ) defined in [9] we have that $\|\boldsymbol{\Phi u}\|_2^2 \le m(1+\delta)$, the expression then becomes

$$\mathbb{P}\{\|\big(\boldsymbol{\Phi}^T(\mathrm{sign}_{\mathbb{C}}(\boldsymbol{y}) - \mathcal{Q}^{\odot}(\boldsymbol{y}))\big)_{\mathcal{T}}\|_2 > m\gamma\} \le 2\exp\left(-\frac{m\gamma^2}{2c_{\odot}^2}\right). \quad (17)$$

This relationship is only valid for one given $\boldsymbol{u}$, one needs thus to extend it to all possible $\boldsymbol{u}$ using a covering and a union bound argument. One can define a ρ-covering $\mathcal{J}_\rho$ of the $2s$-sparse space in $\tilde{\Sigma}_{2s}^N$ as defined in [9]. Using this covering, one can express any $\boldsymbol{u}$ as $\boldsymbol{u} = \boldsymbol{a} + \boldsymbol{r}$, with $\boldsymbol{a} \in \mathcal{J}_\rho \in \tilde{\Sigma}_{2s}^N$ and $\boldsymbol{r} \in \tilde{\Sigma}_{2s}^N$ with $\|\boldsymbol{r}\|_2 \le \rho$. The inequality is then extended as follows

$$\begin{aligned}\|(\boldsymbol{\Phi}^H(\mathrm{sign}_{\mathbb{C}}(\boldsymbol{y}) - \mathcal{Q}^{\odot}(\boldsymbol{y})))_{\mathcal{T}}\|_2 &= \langle \boldsymbol{\Phi}(\boldsymbol{a}+\boldsymbol{r}), \mathrm{sign}_{\mathbb{C}}(\boldsymbol{y}) - \mathcal{Q}^{\odot}(\boldsymbol{y})\rangle \\ &\le m\gamma + \langle \boldsymbol{\Phi r}, \mathrm{sign}_{\mathbb{C}}(\boldsymbol{y}) - \mathcal{Q}^{\odot}(\boldsymbol{y})\rangle, \quad (18)\end{aligned}$$

where (17) is extended to all the elements of $\mathcal{J}_\rho$, thanks to a union bound argument. This holds with a probability of failure upper bounded by $2(\frac{eN}{2s})^{2s}(1+\frac{2}{\rho})^{4s}\exp\left(-\frac{m\gamma^2}{2c_{\odot}}\right)$ using the bound of $|\mathcal{J}_\rho|$ in [9]. One just needs to bound the last term depending on $\boldsymbol{r}$ to conclude the proof. Given its norm and the fact that $\boldsymbol{\Phi}$ follows the (ℓ_2, ℓ_2)-RIP($2s, \delta$), one can bound the second term in (18) using Lemma 1 and the *Cauchy-Schwarz* inequality

$$\begin{aligned}\|(\boldsymbol{\Phi}^H(\mathrm{sign}_{\mathbb{C}}(\boldsymbol{y}) - \mathcal{Q}^{\odot}(\boldsymbol{y})))_{\mathcal{T}}\|_2 &\le m\gamma + \|\boldsymbol{\Phi r}\|_2\frac{\pi}{4}\sqrt{m} \\ &\le m\gamma + \sqrt{2}\rho\frac{\pi}{4}m.\end{aligned}$$

Setting $\rho = \frac{4}{\sqrt{2}\pi}\gamma$, finally gives

$$\|(\boldsymbol{\Phi}^H(\mathrm{sign}_{\mathbb{C}}(\boldsymbol{y}) - \mathcal{Q}^{\odot}(\boldsymbol{y})))_{\mathcal{T}}\|_2 \le 2m\gamma,$$

with a probability of failure upper-bounded by

$$p_\gamma \le 2(\frac{eN}{2s})^{2s}(1+\frac{\pi}{\gamma\sqrt{2}})^{4s}\exp\left(-\frac{m\gamma^2}{2c_{\odot}^2}\right). \quad (19)$$

Finally, provided that $m \ge 8\gamma^{-2}s(\log(\frac{eN}{2s}) + 2(1+\frac{\pi}{\gamma\sqrt{2}}))$, one can upperbound the probability in (19) by

$$(\frac{eN}{2s})^{2s}(1+\frac{\pi}{\gamma\sqrt{2}})^{4s}\exp\left(-\frac{m\gamma^2}{2c_{\odot}^2}\right) \le \exp\left(-\frac{m\gamma^2}{4c_{\odot}^2}\right).$$

□

Appendix C
Proof of Proposition 2

The probability characterized in Prop. 2 is equivalent to this other probability

$$p = \mathbb{P}\big[\forall l \in [m] : (\boldsymbol{\Phi D}_0\boldsymbol{x}_0)_l^*(\boldsymbol{\Phi D}_1\boldsymbol{x}_1)_l \in \mathbb{R}_+\big],$$

with $\boldsymbol{D}_0 = \mathrm{diag}(\boldsymbol{\gamma}_0)$ and $\boldsymbol{D}_1 = \mathrm{diag}(\boldsymbol{\gamma}_1)$. However, by hypothesis,

$$(\boldsymbol{\Phi D}_0\boldsymbol{x}_0)_l = \boldsymbol{e}_l^\top\boldsymbol{\Phi}\mathrm{diag}(\boldsymbol{u}_0)\boldsymbol{\gamma}_0 = \boldsymbol{v}_l^\top\boldsymbol{\gamma}_0,$$

and $\boldsymbol{\Phi D}_1\boldsymbol{x}_1 = \boldsymbol{w}_l^\top\boldsymbol{\gamma}_1$, with $\boldsymbol{e}_l$ the l-th element of the canonical basis, $\boldsymbol{v}_l = \mathrm{diag}(\boldsymbol{u}_0)\boldsymbol{\Phi}^\top\boldsymbol{e}_l$, and $\boldsymbol{w}_l = \mathrm{diag}(\boldsymbol{u}_1)\boldsymbol{\Phi}^\top\boldsymbol{e}_l$. Since $\boldsymbol{\Phi}$ is a subsampling of the rows of the DFT matrix and $\boldsymbol{u}_0, \boldsymbol{u}_1 \ne \boldsymbol{0}$, $\boldsymbol{v}_l$ and $\boldsymbol{w}_l$ are non-zero vectors. Therefore, passing to the complementary probability, we get

$$\begin{aligned}1 - p &= \mathbb{P}\big[\exists l \in [m] : (\boldsymbol{v}_l^\top\boldsymbol{\gamma}_0)^*(\boldsymbol{w}_l^\top\boldsymbol{\gamma}_1) \notin \mathbb{R}_+\big] \\ &\ge \max_l \mathbb{P}\big[(\boldsymbol{v}_l^\top\boldsymbol{\gamma}_0)^*(\boldsymbol{w}_l^\top\boldsymbol{\gamma}_1) \notin \mathbb{R}_+\big].\end{aligned}$$

Fixing l, both (independent) complex random variables $a = \boldsymbol{v}_l^\top\boldsymbol{\gamma}_0$ and $b = \boldsymbol{w}_l^\top\boldsymbol{\gamma}_1$ are circularly symmetric, *i.e.*, $\boldsymbol{v}_l^\top\boldsymbol{\gamma}_0$ and $e^{\mathrm{i}\,\alpha}\boldsymbol{v}_l^\top\boldsymbol{\gamma}_0$ have the same distribution for any phase α. Therefore the phases of a and b are uniformly distributed over $[0, 2\pi]$, which implies that $\mathbb{P}\big[a^*b \notin \mathbb{R}_+\big] = 1$ since elements of $\mathbb{R}_+$ have a zero phase; a set of measure zero. This shows that $1 - p \ge 1$, and $p = 0$ as announced.

Acknowledgment

The authors thank Thomas Fromenteze, Mike E. Davies, Maxime Drouguet, Danielle Vanhoenacker, Christophe Craeye, and the reviewers for their insightful comments.

REFERENCES

[1] J. Lien, N. Gillian, M. E. Karagozler, *et al.*, "Soli: Ubiquitous gesture sensing with millimeter wave radar," *ACM Trans. Graph.*, vol. 35, no. 4, 142:1–142:19, Jul. 2016, ISSN: 0730-0301. DOI: 10.1145/2897824.2925953.

[2] S. Saponara, M. S. Greco, and F. Gini, "Radar-on-chip/in-package in autonomous driving vehicles and intelligent transport systems: Opportunities and challenges," *IEEE Signal Processing Magazine*, vol. 36, no. 5, pp. 71–84, Sep. 2019, ISSN: 1558-0792. DOI: 10.1109/MSP.2019.2909074.

[3] G. Paterniani, D. Sgreccia, A. Davoli, *et al.*, "Radar-based monitoring of vital signs: A tutorial overview," *Proceedings of the IEEE*, vol. 111, no. 3, pp. 277–317, 2023. DOI: 10.1109/JPROC.2023.3244362.

[4] G. Beltrão, R. Stutz, F. Hornberger, *et al.*, "Contactless radar-based breathing monitoring of premature infants in the neonatal intensive care unit," en, *Scientific Reports*, vol. 12, no. 1, p. 5150, Mar. 2022, ISSN: 2045-2322. DOI: 10.1038/s41598-022-08836-3.

[5] C. Aydogdu, M. F. Keskin, G. K. Carvajal, *et al.*, "Radar interference mitigation for automated driving: Exploring proactive strategies," *IEEE Signal Processing Magazine*, vol. 37, no. 4, pp. 72–84, 2020. DOI: 10.1109/MSP.2020.2969319.

[6] D. Taubman and M. Marcellin, "Jpeg2000: Standard for interactive imaging," *Proceedings of the IEEE*, vol. 90, no. 8, pp. 1336–1357, Aug. 2002, ISSN: 1558-2256. DOI: 10.1109/JPROC.2002.800725.

[7] S. Rane, P. Boufounos, A. Vetro, and Y. Okada, "Low complexity efficient raw sar data compression," *Proc SPIE*, vol. 8051, May 2011. DOI: 10.1117/12.884034.

[8] U. Benz, K. Strodl, and A. Moreira, "A comparison of several algorithms for sar raw data compression," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 33, no. 5, pp. 1266–1276, Sep. 1995, ISSN: 1558-0644. DOI: 10.1109/36.469491.

[9] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*. Birkhäuser Basel, 2013, ISBN: 0817649476.

[10] E. J. Candes and M. B. Wakin, "An introduction to compressive sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 21–30, Mar. 2008, ISSN: 1558-0792. DOI: 10.1109/MSP.2007.914731.

[11] T. Blumensath and M. E. Davies, "Iterative hard thresholding for compressed sensing," *Applied and Computational Harmonic Analysis*, vol. 27, no. 3, pp. 265–274, 2009, ISSN: 1063-5203. DOI: https://doi.org/10.1016/j.acha.2009.04.002. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1063520309000384.

[12] D. Needell and J. Tropp, "Cosamp: Iterative signal recovery from incomplete and inaccurate samples," *Applied and Computational Harmonic Analysis*, vol. 26, no. 3, pp. 301–321, 2009, ISSN: 1063-5203. DOI: https://doi.org/10.1016/j.acha.2008.07.002. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1063520308000638.

[13] M. F. Duarte, M. A. Davenport, D. Takhar, *et al.*, "Single-pixel imaging via compressive sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 83–91, Mar. 2008, ISSN: 1558-0792. DOI: 10.1109/MSP.2007.914730.

[14] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse mri: The application of compressed sensing for rapid mr imaging," en, *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007, ISSN: 1522-2594. DOI: 10.1002/mrm.21391.

[15] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly, "Compressed sensing mri," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 72–82, Mar. 2008, ISSN: 1558-0792. DOI: 10.1109/MSP.2007.914728.

[16] M. A. Herman and T. Strohmer, "High-resolution radar via compressed sensing," *IEEE Transactions on Signal Processing*, vol. 57, no. 6, pp. 2275–2284, 2009. DOI: 10.1109/TSP.2009.2014277.

[17] J. Ender, "A brief review of compressive sensing applied to radar," in *2013 14th International Radar Symposium (IRS)*, vol. 1, 2013, pp. 3–16.

[18] E. Baransky, G. Itzhak, I. Shmuel, N. Wagner, E. Shoshan, and Y. C. Eldar, "A sub-nyquist radar prototype: Hardware and algorithms," no. arXiv:1208.2515, Aug. 2012, arXiv:1208.2515 [cs]. DOI: 10.48550/arXiv.1208.2515. [Online]. Available: http://arxiv.org/abs/1208.2515.

[19] S. Krone and G. Fettweis, "Capacity of communications channels with 1-bit quantization and oversampling at the receiver," in *2012 35th IEEE Sarnoff Symposium*, May 2012, pp. 1–7. DOI: 10.1109/SARNOF.2012.6222713. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/6222713.

[20] M. S. Stein, S. Bar, J. A. Nossek, and J. Tabrikian, "Performance analysis for channel estimation with 1-bit adc and unknown quantization threshold," *IEEE Transactions on Signal Processing*, vol. 66, no. 10, pp. 2557–2571, May 2018, ISSN: 1053-587X, 1941-0476. DOI: 10.1109/TSP.2018.2815022.

[21] A. Bhandari, M. H. Conde, and O. Loffeld, "One-bit time-resolved imaging," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, no. 7, pp. 1630–1641, Jul. 2020, ISSN: 1939-3539. DOI: 10.1109/TPAMI.2020.2986950.

[22] F. Xi, Y. Xiang, S. Chen, and A. Nehorai, *Gridless parameter estimation for one-bit mimo radar with time-varying thresholds*, 2019. arXiv: 1908.10487 [eess.SP].

[23] T. Feuillen, C. Xu, J. Louveaux, L. Vandendorpe, and L. Jacques, "Quantity over quality: Dithered quantization for compressive radar systems," *2019 IEEE Radar Conference (RadarConf)*, pp. 1–6, 2018.

[24] B. Le, T. Rondeau, J. Reed, and C. Bostian, "Analog-to-digital converters," *IEEE Signal Processing Magazine*, vol. 22, no. 6, pp. 69–77, 2005. DOI: 10.1109/MSP.2005.1550190.

[25] R. Walden, "Analog-to-digital converter survey and analysis," *IEEE Journal on Selected Areas in Communications*, vol. 17, no. 4, pp. 539–550, 1999. DOI: 10.1109/49.761034.

[26] L. Jacques and V. Cambareri, "Time for dithering: fast and quantized random embeddings via the restricted isometry property," *Information and Inference: A Journal of the IMA*, vol. 6, no. 4, pp. 441–476, Apr. 2017, ISSN: 2049-8764. DOI: 10.1093/imaiai/iax004. eprint: https://academic.oup.com/imaiai/article-pdf/6/4/441/22874344/iax004.pdf. [Online]. Available: https://doi.org/10.1093/imaiai/iax004.

[27] Y. Plan and R. Vershynin, *One-bit compressed sensing by linear programming*, 2012. arXiv: 1109.4299 [cs.IT].

[28] T. Feuillen, C. Xu, L. Vandendorpe, and L. Jacques, "1-bit localization scheme for radar using dithered quantized compressed sensing," *2018 5th International Workshop on Compressed Sensing applied to Radar, Multimodal Sensing, and Imaging (CoSeRa)*, 2018.

[29] L. Roberts, "Picture coding using pseudo-random noise," *IRE Transactions on Information Theory*, vol. 8, no. 2, pp. 145–154, Feb. 1962, ISSN: 2168-2712. DOI: 10.1109/TIT.1962.1057702.

[30] P. Aziz, H. Sorensen, and J. vn der Spiegel, "An overview of sigma-delta converters," *IEEE Signal Processing Magazine*, vol. 13, no. 1, pp. 61–84, Jan. 1996, ISSN: 1558-0792. DOI: 10.1109/79.482138.

[31] U. Kamilov, A. Bourquard, A. Amini, and M. Unser, "One-bit measurements with adaptive thresholds," *IEEE Signal Processing Letters*, vol. 19, pp. 607–610, 2012.

[32] S. Dirksen, "Quantized compressed sensing: A survey," in *Compressed Sensing and Its Applications: Third International MATHEON Conference 2017*, H. Boche, G. Caire, R. Calderbank, G. Kutyniok, R. Mathar, and P. Petersen, Eds. Cham: Springer International Publishing, 2019, pp. 67–95, ISBN: 978-3-319-73074-5. DOI: 10.1007/978-3-319-73074-5_2. [Online]. Available: https://doi.org/10.1007/978-3-319-73074-5_2.

[33] R. Baraniuk, S. Foucart, D. Needell, Y. Plan, and M. Wootters, "Exponential decay of reconstruction error from binary measurements of sparse signals," *IEEE Transactions on Information Theory*, vol. 63, pp. 3368–3385, 2017.

[34] J. Li, M. M. Naghsh, S. J. Zahabi, and M. Modarres-Hashemi, "Compressive radar sensing via one-bit sampling with time-varying thresholds," in *2016 50th Asilomar Conference on Signals, Systems and Computers*, 2016, pp. 1164–1168. DOI: 10.1109/ACSSC.2016.7869554.

[35] A. Eamaz, F. Yeganegi, D. Needell, and M. Soltanalian, "Harnessing the power of sample abundance: Theoretical guarantees and algorithms for accelerated one-bit sensing," *IEEE Transactions on Information Theory*, vol. 70, no. 9, pp. 6690–6713, 2024. DOI: 10.1109/TIT.2024.3422918.

[36] A. Eamaz, F. Yeganegi, and M. Soltanalian, "Covariance recovery for one-bit sampled non-stationary signals with time-varying sampling thresholds," *IEEE Transactions on Signal Processing*, vol. 70, pp. 5222–5236, 2022. DOI: 10.1109/TSP.2022.3217379.

[37] C. Xu and L. Jacques, "Quantized compressive sensing with RIP matrices: the benefit of dithering," *Information and Inference: A Journal of the IMA*, vol. 9, no. 3, pp. 543–586, Nov. 2019, ISSN: 2049-8772. DOI: 10.1093/imaiai/iaz021. eprint: https://academic.oup.com/imaiai/article-pdf/9/3/543/33720723/iaz021.pdf. [Online]. Available: https://doi.org/10.1093/imaiai/iaz021.

[38] J. Ren, T. Zhang, J. Li, and P. Stoica, "Sinusoidal parameter estimation from signed measurements via majorization–minimization based relax," *IEEE Transactions on Signal Processing*, vol. 67, no. 8, pp. 2173–2186, 2019. DOI: 10.1109/TSP.2019.2899804.

[39] M. Skolnik, *Introduction to Radar Systems /2nd Edition/*, 2nd ed. New York: McGraw Hill Book Co., 1980.

[40] T. Feuillen, L. Vandendorpe, and L. Jacques, "1-bit compressive radar with phase dithering," in *2024 32nd European Signal Processing Conference (EUSIPCO)*, 2024, pp. 2692–2696. DOI: 10.23919/EUSIPCO63174.2024.10715097.

[41] S. Foucart, "Flavors of compressive sensing," in *Approximation Theory XV: San Antonio 2016*, G. E. Fasshauer and L. L. Schumaker, Eds., Cham: Springer International Publishing, 2017, pp. 61–104, ISBN: 978-3-319-59912-0.

[42] B. Meyer, J. B. de Swardt, and P. van der Merwe, "Comparing baseband and intermediate frequency fmcw radar receivers," in *2017 IEEE AFRICON*, 2017, pp. 563–568. DOI: 10.1109/AFRCON.2017.8095543.

[43] J. Graeme, G. E. Tobey, and L. Huelsman, "Operational amplifiers; design and applications," 1971.

[44] L. Jacques, K. Degraux, and C. De Vleeschouwer, "Quantized iterative hard thresholding: Bridging 1bit and HighResolution quantized compressed sensing," in *10th international conference on Sampling Theory and Applications (SampTA 2013)*, Bremen, Germany, Jul. 2013, pp. 105–108.

[45] R. Vaughan, N. Scott, and D. White, "The theory of bandpass sampling," *IEEE Transactions on Signal Processing*, vol. 39, no. 9, pp. 1973–1984, 1991. DOI: 10.1109/78.134430.

[46] L. Jacques, J. N. Laska, P. T. Boufounos, and R. G. Baraniuk, "Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors," *IEEE Transactions on Information Theory*, vol. 59, no. 4, pp. 2082–2102, 2013. DOI: 10.1109/TIT.2012.2234823.

[47] L. Jacques and T. Feuillen, *The importance of phase in complex compressive sensing*, 2021. DOI: 10.1109/TIT.2021.3073566.

[48] T. Feuillen, M. E. Davies, L. Vandendorpe, and L. Jacques, *(l1,l2)-rip and projected back-projection reconstruction for phase-only measurements*, 2019. arXiv: 1912.02880 [eess.SP].

[49] J. Chen, Z. Liu, M. K. Ng, and J. Scarlett, "Robust instance optimal phase-only compressed sensing,"

no. arXiv:2408.06275, Aug. 2024, arXiv:2408.06275 [cs]. DOI: 10 . 48550 / arXiv . 2408 . 06275. [Online]. Available: http://arxiv.org/abs/2408.06275.

[50] A. V. Oppenheim and J. S. Lim, “The importance of phase in signals,” *Proceedings of the IEEE*, vol. 69, no. 5, pp. 529–541, May 1981, ISSN: 1558-2256. DOI: 10.1109/PROC.1981.12022.

[51] A. V. Oppenheim, M. H. Hayes, and J. S. Lim, “Iterative Procedures For Signal Reconstruction From Fourier Transform Phase,” *Optical Engineering*, vol. 21, no. 1, pp. 122–127, 1982. DOI: 10.1117/12.7972871. [Online]. Available: https://doi.org/10.1117/12.7972871.