1-bit Compressive Radar with Phase Dithering

Thomas Feuillen[1], Luc Vandendorpe[2] and Laurent Jacques[2]

[1]*SPARC, SnT, UniLu*, Luxembourg. [2]*ICTEAM, UCLouvain.*

***Abstract*—Dithering techniques, through the *addition* of a (random) signal before the Analog to Digital Converter (ADC), are widely used in acquisition chain designs for their ability to *shape* the quantization noise, *e.g.*, to improve the estimation and detection of the sampled signal. In this paper, we stray away from this additive dithering scheme—that often induces a complex and high-cost implementation—and adopt a phase dithering technique applicable to 1-bit quantization of radar signals. Our approach leverages already existing radar architecture of Frequency Modulated Continuous Wave (FMCW) radars and can be efficiently implemented. We first study theoretically the efficiency of our scheme by linking it to Phase-Only (PO) acquisition, a coarse quantization scheme that only records the signal phase. We then show, through extensive Monte Carlo simulations, that phase dithering is an interesting alternative to the additive random dithering for low-measurement regimes.**

I. Introduction

The recent years have seen the advent of smaller form factor radars [1]. These low-footprint sensors have opened the way for numerous new applications (from healthcare [2] to automotive safety [3]) and, more generally, for a more ubiquitous way of sensing the environment around us. This creates new challenges in acquiring and processing the data generated by this new breadth of applications, which often have higher sensor density. One way of reducing the amount of data to transmit and process is to lower its resolution using coarser Analog to Digital Converters (ADCs). Recently, 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 [4], [5], as well as radar sensors [6]–[8].

Given a (discretized) *scene* $\boldsymbol{x} \in \mathbb{C}^N$, the measurement vector $\boldsymbol{y}$ collected by a (linear) sensor—associated with a sensing matrix $\boldsymbol{\Phi} \in \mathbb{C}^{M\times N}$—can often be modelled as $\boldsymbol{y} = \mathcal{Q}(\boldsymbol{\Phi x}) \in \mathcal{C}$, where the quantizer $\mathcal{Q} : \mathbb{C}^M \to \mathcal{C}$, for some codebook $\mathcal{C}$ (with $|\mathcal{C}| = M2^B$ for a B-bit scalar quantizer), digitizes the *signal* $\boldsymbol{r} := \boldsymbol{\Phi x}$ generated by the sensor. Combining the model $\boldsymbol{\Phi}$ (often represented by a Fourier transform for FMCW radars) with the quantized measurements $\boldsymbol{y}$, enables the estimation of $\boldsymbol{x}$ through reconstruction algorithms. In applications where 1-bit quantization is considered, *i.e.*, with binary real and imaginary components, the quantizer reads $\mathcal{Q}(\cdot) = \mathcal{Q}_{\rm b}(\cdot) := \operatorname{sign}\Re(\cdot) + {\rm i}\operatorname{sign}\Im(\cdot)$. Using this type of binary acquisition aims to lower the requirement on the hardware by either lowering its cost or its power consumption [9], [10]. Although the implementation is extremely simplified, applying crudely this binary quantizer often comes at a cost in terms of performance as shown in [6], [11]–[13]. Indeed, in [13] for Bernoulli matrices and in [6], [11] for Fourier matrices, authors showed that applying this extremely low-resolution quantizer on noiseless data can result in ambiguous scenarios, preventing high-quality reconstructions. One such scenario arises when two different sparse vectors, once quantized, are sent exactly to the same bits, removing any possibility of distinguishing them after quantization. Fortunately, using an additive dithering before quantization often solves this issue, even in highly quantized settings. Mathematically, this can be modelled by a vector $\boldsymbol{\xi} \in \mathbb{C}^m$ that is added to the signal before the quantization; the quantizer becomes $\mathcal{Q}^+(\cdot) := \mathcal{Q}(\cdot + \boldsymbol{\xi})$ (see [11], [14], [15] and others). Several schemes have been considered to generate such a dithering. Authors in [16]–[18] proposed to generate it according to the previous measurements (following a $\Sigma\Delta$ quantizer principle [17, Sec. 6]) in order to extract the most information possible out of those coarse 1-bit measurements. While this *tailored* process induces a steep performance gain in the signal reconstruction, it also implies a feedback loop between the reconstruction and the 1-bit acquisition that can be quite expensive, be it in cost or power, to implement.

Authors in [15] proposed a randomly generated dither, generated uniformly at random inside the quantization bin, not depending of the reconstruction algorithm. They were able to show that by leveraging the effect of this dither on the quantized data, one can upperbound the ℓ_2-reconstruction error of the Projected-Back Projection (PBP) algorithm [15].

Implementing this dither in practice, however, comes with several challenges. Indeed, in the case of binary quantization, the random uniform dither has to span the dynamic of the measured signal, *i.e.*, $\|\boldsymbol{r}\|_\infty$. Therefore, *(i)* this uniform random variable must be generated in hardware at a reduced cost or power consumption compared to high-resolution ADCs, without using high-resolution DACs (Digital to Analog Converters); *(ii)* we need an estimate of the signal's dynamic, a strong assumption in radar signals given the high variability of their received powers [19]; and *(iii)* since certain signal reconstruction algorithms (such as the quantized iterative hard thresholding [20]) require knowing the sampled dithering to enforce the consistency of the estimated signal with the quantized measurement [21], [22], the dither must be controlled (for reproducibility) or measured before or during the acquisition, which imposes further constraints on the hardware.

Given these issues, and still considering the need of randomized schemes to mitigate the quantization error, we set out

The authors thank T. Fromenteze, M. E. Davies, M. Drouguet, D. Vanhoenacker, and C. Craeye for their insightful comments. TF's work was supported by FNR CORE SURF Project, ref C23/IS/18158802/SURF and European Research Council (ERC, SpeckleCARS, 101052911), LJ's work by FRS-FNRS (QuadSense, T.0160.24).

to find another way of dithering signals coming from FMCW radars. We introduce a phase dithering that modifies the phases of the measured signals before the 1-bit quantization according to unitary vector with a random phase.

The contributions of this paper are the following: *(i)* we show that using phase dithering in the context of 1-bit quantization is a viable trade-off between reconstruction performances and complexity of implementation in hardware; *(ii)* we demonstrate that this dithering scheme has reconstruction performances bounded by what is achievable by Phase-Only (PO) acquisition scheme [23]–[25]; *(iii)* finally, we confirm these results and show that phase dithering compares favorably to other dithering schemes at low-measurement regimes.

This paper is structured as follows. Sec. I explains the difficulty of implementing an efficient dithering procedure in the context of radar systems. Sec. II introduces the sensor whose model is leveraged to obtain an efficient and cost-effective dither, *i.e.*, FMCW radars. Sec. III introduces the phase dithering studied in this paper and highlights its advantages compared additive dithering in terms of practical implementation. Sec. IV links phase dithering of 1-bit observations to PO acquisition and studies its limits in terms of uniform recovery guarantees of Fourier based measurements. The results of Monte-Carlo simulations are presented in Sec. V. Finally, conclusions and future works are presented in Section VI.

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 $\mathrm{i}^2 = -1$; $\Re\{\lambda\}$ (or $\lambda^{\Re}$) and $\Im\{\lambda\}$ (or $\lambda^{\Im}$) are the real and imaginary part of $\lambda \in \mathbb{C}$, respectively, and λ^* is its complex conjugate; $\boldsymbol{A}^*$ is the conjugate transpose of $\boldsymbol{A}$; $\operatorname{supp}\boldsymbol{x}$ is the support of $\boldsymbol{x} \in \mathbb{C}^d$; $\langle \boldsymbol{x}, \boldsymbol{y}\rangle = \sum_{i=1}^d x_i y_i^*$ 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 $\|\boldsymbol{x}\|_\infty = \max_i |x_i|$. The Hadamard product is $\odot$, and the angle operator (applied componentwise onto vectors) reads $\angle(ce^{\mathrm{i}\phi}) = \phi$ for $c > 0$ and $\phi \in [-\pi, \pi]$. An K-sparse $\boldsymbol{x}$ vector belongs to the set $\bar{\Sigma}_K^N := \{\boldsymbol{u} \in \mathbb{C}^N, \|\boldsymbol{u}\|_0 \leq K\}$. 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 $\lambda \in \mathbb{C}$.

II. FMCW Radar Signal Model

In this section, before introducing the low-resolution acquisition model, we first focus on the model of the sensor from which the signals will be collected. In this paper, we consider a Frequency Modulated Continuous Wave radar with one transmitting and one receiving antenna.

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

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

with f_0 the central frequency, T_c the chirp duration, and B the spanned bandwidth. Considering, for now, a simple ranging problem where the scene is made of only one target located at a range R from the radar which has one receiving antenna.

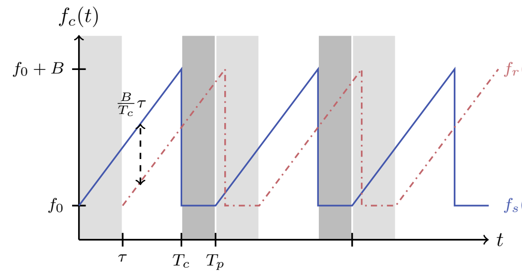

Fig. 1: Representation of the frequency content of the transmitted (plain, **blue**) and received (dash dotted, **red**) chirps with a delay τ

The received signal is a time-delayed version of the frequency modulated transmitted signal $s(t)$,

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

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 $\tau = \frac{2R}{\mathrm{c}}$. 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 simplicity's sake 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. It is interesting to note that in Fig. 1, 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.

FMCW radars use coherent demodulation, where, to recover this frequency shift, the received signal $r^{RF}(t)$ is demodulated using the transmitted signal $s(t)$ itself. In essence, using this demodulation process, we are now only recording the changes that are imparted by the environment. 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 process represented in Fig. 2, can be expressed

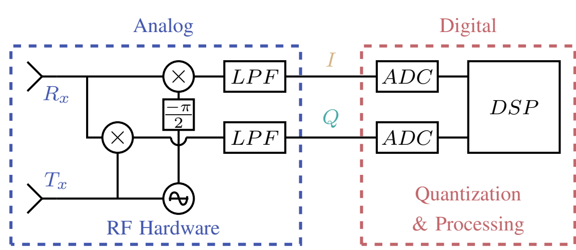

Fig. 2: Representation of a IQ coherent demodulation for an FMCW radar with one transmit and one receive antenna

mathematically simply as $r(t) = r^{RF}(t)s^*(t)$. Leveraging the linear change in frequency of the transmitted signal allows us to express the received signal $r(t)$ in base-band, *i.e.*, we have

$$r(t) = \alpha \exp\big(-\mathrm{i}\,2\pi\tau f_c(t)\big). \tag{1}$$

In words, (1) 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.

It is interesting to note that this frequency shift $\frac{B}{T_c}\tau$ is often several orders of magnitude below the bandwidth B. This fact motivates the use of the coherent demodulation as the sampling frequency is now linked to $\frac{B}{T_c}\tau$ instead of B. The range resolution δ_R can be simply expressed as $\delta_R = \frac{\mathsf{c}}{2B}$ thanks to the property of the Fourier transform.

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

$$r(t) = \sum_{k=1}^{K} \alpha_k \exp\big(-\,\mathrm{i}\,2\pi \tfrac{2R_k}{\mathsf{c}} f_c(t)\big),$$

where α_k and R_k are the received powers and ranges from each of the K targets measured in the scene. If the duration of the acquisition is restricted to $t \in [0, T_c]$, the problem of estimating the ranges of these K targets is tantamount to estimating the frequency content of $r(t)$ (*e.g.*, 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 $\boldsymbol{r} = \boldsymbol{\Psi}\boldsymbol{x}$, where $\boldsymbol{x} \in \bar{\Sigma}_K^N$ (*i.e.*, x_j equals α_k if the j-th pixel 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 [26]

$$\boldsymbol{r} = \boldsymbol{\Phi}\boldsymbol{x} \in \mathbb{C}^M, \tag{2}$$

by (randomly) subsampling $\boldsymbol{\Psi}\boldsymbol{x}$ over $M \leq N$ frequencies, *i.e.*, $\boldsymbol{\Phi} := \boldsymbol{P}\boldsymbol{\Psi}$ 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.*, $\boldsymbol{\Phi}^\top = [\boldsymbol{\Psi}^\top, \ldots, \boldsymbol{\Psi}^\top, (\boldsymbol{P}\boldsymbol{\Psi})^\top]^\top$.

III. Phase Dithering

In this paper, instead of dithering the amplitudes of the received signal in (2) in the real and imaginary domain, we propose to dither their phases by multiplying the signal with a unitary signal $\boldsymbol{\xi}$ with a random phase, *i.e.*, $\xi_j \sim_{\text{i.i.d.}} e^{\mathrm{i}\,\mathcal{U}(0,2\pi)}$, $j \in [M]$. Defining $\mathcal{Q}_{\mathrm{b}}(\cdot) := \operatorname{sign}\Re(\cdot) + \mathrm{i}\operatorname{sign}\Im(\cdot)$, the proposed quantization process can be represented as:

$$\boldsymbol{y} := \mathcal{Q}^{\odot}(\boldsymbol{r} = \boldsymbol{\Phi}\boldsymbol{x}) := c_{\odot}\mathcal{Q}_{\mathrm{b}}(\boldsymbol{r} \odot \boldsymbol{\xi}) \odot \boldsymbol{\xi}^*, \quad c_{\odot} := \tfrac{\pi}{4}.$$

This way of dithering solves the first issue brought up by the previous section regarding the additive scheme. It is indeed amplitude independent and can be thus applied without the knowledge of the dynamic of the signal $\|\boldsymbol{y}\|_\infty$. Furthermore, dithering the phase can be achieved using different and already existing architectures. Using a non *zero-if* demodulation architecture [27] such as in Fig. 3. One needs to generate the dither $\xi^*(t)$, the radar transmits the original signal $s(t)$ and then the received RF signal $r^{RF}(t)$ is demodulated by $s(t)\xi^*(t)$. The coherent demodulation process then becomes

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

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

Another way of randomly modifying the phase of the received signal is to modify the signal after its coherent demodulation. In Fig. 4, the base-band signals are multiplied by

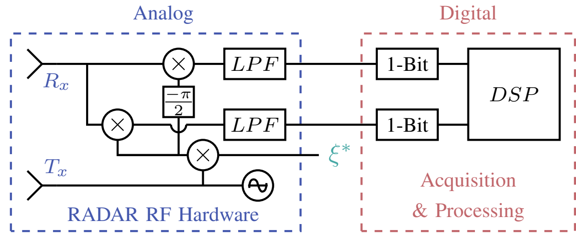

Fig. 3: Example of FMCW radar architecture with non zero IF demodulation and 1-bit quantization

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,

$$\begin{aligned} r(t) \times \xi(t) = {} & r^{\mathbb{R}}(t) \times \xi^{\mathbb{R}}(t) \;-\; r^{\mathbb{I}}(t) \times \xi^{\mathbb{I}}(t) \\ & +\; \mathrm{i}\,\big(r^{\mathbb{R}}(t) \times \xi^{\mathbb{I}}(t) \;+\; r^{\mathbb{I}}(t) \times \xi^{\mathbb{R}}(t)\big). \end{aligned}$$

This process can be easily implemented with *off-the-shelf* components; it only requires base-band components operating thus at low frequencies. These are both inexpensive and require low power compared to their RF counterparts.

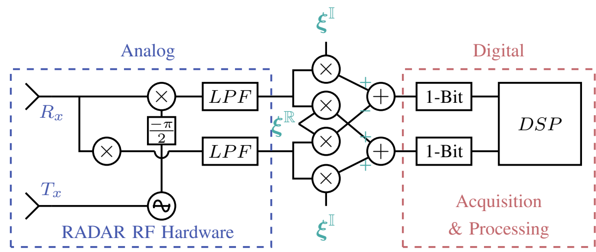

Fig. 4: Possible FMCW radar architecture with phase dithering and 1-bit quantization.

IV. Reconstruction Guarantees and Connections with Phase-Only acquisition

We now study the reconstruction guarantees of the proposed quantization scheme using a specific algorithm, the projected back projection (PBP). It produces the following estimate, $\hat{\boldsymbol{x}} = \mathcal{H}_K(\boldsymbol{\Phi}^*\boldsymbol{y})$, where $\mathcal{H}_K(\cdot)$ is a hard thresholding zeroing out all but the K largest components (in amplitude) of its input.

Let us assume that $\boldsymbol{x} \in \bar{\Sigma}_K^N$. Given phase dithered 1-bit data $\boldsymbol{y} = \mathcal{Q}^{\odot}(\boldsymbol{\Phi}\boldsymbol{x})$, the backprojected vector $\boldsymbol{a} = \frac{1}{m}\boldsymbol{\Phi}^*\boldsymbol{y}$ (with $\hat{\boldsymbol{x}} = \mathcal{H}_K(\boldsymbol{a})$), the support $\mathcal{S} := \operatorname{supp}(\hat{\boldsymbol{x}})$ of the back-projection estimate, and $\mathcal{T} := \operatorname{supp}(\boldsymbol{x}) \cup \mathcal{S}$, the reconstruction error can be bounded by $\|\boldsymbol{x} - \hat{\boldsymbol{x}}\|_2 = \|\boldsymbol{x} - \boldsymbol{a}_{\mathcal{S}}\|_2 \leq \|\boldsymbol{x} - \boldsymbol{a}_{\mathcal{T}}\|_2 + \|\boldsymbol{a}_{\mathcal{T}} - \boldsymbol{a}_{\mathcal{S}}\|_2$. Since $\boldsymbol{a}_{\mathcal{S}} = \mathcal{H}_K(\boldsymbol{a})$ is the best K-term approximation of both $\boldsymbol{a}$ and $\boldsymbol{a}_{\mathcal{T}}$, we have $\|\boldsymbol{x} - \hat{\boldsymbol{x}}\|_2 \leq 2\|\boldsymbol{x} - \boldsymbol{a}_{\mathcal{T}}\|_2$ and this last term can be further split into

$$\|\boldsymbol{x} - \hat{\boldsymbol{x}}\|_2 \leq \mathcal{E}_{\mathrm{PO}} + \mathcal{E}_{\mathrm{q}}, \tag{3}$$

with $\mathcal{E}_{\mathrm{PO}} := 2\|\boldsymbol{x} - \frac{1}{m}\big(\boldsymbol{\Phi}^* \operatorname{sign}_{\mathbb{C}}(\boldsymbol{\Phi}\boldsymbol{x})\big)_{\mathcal{T}}\|_2$ accounting for the error achieved by PBP on Phase-Only measurements $\operatorname{sign}_{\mathbb{C}}(\boldsymbol{\Phi}\boldsymbol{x})$, and $\mathcal{E}_{\mathrm{q}} := \frac{2}{m}\|\big(\boldsymbol{\Phi}^*(\operatorname{sign}_{\mathbb{C}}(\boldsymbol{\Phi}\boldsymbol{x}) - \mathcal{Q}^{\odot}(\boldsymbol{\Phi}\boldsymbol{x}))\big)_{\mathcal{T}}\|_2$ the quantization distortion induced by 1-bit measurements compared to the PO measurements.

The error $\mathcal{E}_{\mathrm{PO}}$ is known to decay as $O(\sqrt[4]{K/M})$ when M increases [25] if the matrix $\boldsymbol{\Phi}$ satisfies the (ℓ_p, ℓ_2)-RIP with $p = 1$ and level $2K$—*i.e.*, $\|\boldsymbol{\Phi}\boldsymbol{u}\|_p$ is close to $\|\boldsymbol{u}\|_2$ up to small distortion $\delta = O(\sqrt{K/M})$ for all $2K$-sparse vectors

$\boldsymbol{u}$, which is the case, with high probability (*w.h.p.*), if $\boldsymbol{\Phi}$ is an Gaussian random matrix. Extending this error bound for Fourier-based sensing matrix is an open question; but provably efficient reconstruction methods exist for real signals $\boldsymbol{x}$ [28].

Regarding the second error term in (3), $\mathcal{E}_{\rm q}$, one can show that if $\boldsymbol{\Phi}$ satisfies the (ℓ_2, ℓ_2)-RIP of level $2K$ with small distortion, which is verified *w.h.p.* for Gaussian and partial random Fourier matrices $\boldsymbol{\Phi}$ [26], and if the dithered phases are drawn uniformly at random on $[0, 2\pi)$, then, given $\gamma > 0$, *w.h.p.*, $\mathcal{E}_{\rm q} \leq \gamma$ provided $M \geq M_0$ with $M_0 = O(\gamma^{-2}K)$ up to log factors [29]. Without proving this last bound here, this fact is related to the following property: in expectation, a phase dithered 1-bit quantized vector matches its PO observations.

Proposition 1. *For $\boldsymbol{a} \in \mathbb{C}^m$ and $\xi_j \sim_{\rm i.i.d.} e^{\mathrm{i}\,\mathcal{U}(0,2\pi)}$, $j \in [M]$, we have $\mathbb{E}_\xi \mathcal{Q}^{\odot}(\boldsymbol{a}) = \mathrm{sign}_{\mathbb{C}}(\boldsymbol{a})$.*

Proof. Since $\mathcal{Q}^{\odot}$ is applied componentwise onto vectors, it suffices to show the result for scalars. Given $a = |a|e^{\mathrm{i}\,\phi}$, and the dither $\xi = e^{\mathrm{i}\,u}$ with $u \sim \mathcal{U}(0, 2\pi)$, $\mathbb{E}_\xi[\mathcal{Q}^{\odot}(a)]$ is given by $\frac{c_\odot}{2\pi}\int_0^{2\pi} \mathcal{Q}(|a|e^{\mathrm{i}(\phi+u)})e^{-\mathrm{i}\,u}\mathrm{d}u$, or $\frac{c_\odot}{2\pi}e^{\mathrm{i}\,\phi}\int_0^{2\pi} \mathcal{Q}(e^{\mathrm{i}\,u})e^{-\mathrm{i}\,u}\mathrm{d}u$, the integrated function being periodic in u. However, the function $g(u) := Q(e^{\mathrm{i}\,u})e^{-\mathrm{i}\,u}$ respects $g(u + \frac{\pi}{2}) = g(u)$ since $\mathcal{Q}(\mathrm{i}\,z) = \mathrm{i}\,\mathcal{Q}(z)$. Therefore, $\mathbb{E}_\xi[\mathcal{Q}^{\odot}(a)] = 4\frac{c_\odot}{2\pi}e^{\mathrm{i}\,\phi}\int_0^{\pi/2} \mathcal{Q}(e^{\mathrm{i}\,u})e^{-\mathrm{i}\,u}\mathrm{d}u$, and since $\mathcal{Q}(e^{\mathrm{i}\,u}) = \sqrt{2}e^{\mathrm{i}\,\frac{\pi}{4}}$ in this integration range, the result follows from the fact that $c_\odot = \frac{\pi}{4}$ and $\int_0^{\pi/2} e^{-\mathrm{i}\,u}\mathrm{d}u = \sqrt{2}e^{-\mathrm{i}\,\frac{\pi}{4}}$. $\square$

This last proposition shows a clear difference with the additive dithering scheme. While this one cancels out the quantizer in expectation (over a certain domain, depending on the quantizer range and resolution [15]), in expectation, phase dithered one-bit quantization can only recover signal's phase. One of its strengths—being independent of the amplitude—induces a weaker reconstruction guarantees. As shown by (3), the PBP error in the proposed 1-bit scheme is bounded by both the performances of the PO acquisition ($\mathcal{E}_{\rm PO}$) and the information loss between PO and quantized observation ($\mathcal{E}_{\rm q}$). However, as explained below, this fact should not discount this novel dithering scheme, and the hardware simplifications it brings, as a viable alternative to its additive counterpart.

V. Simulations

We now assess the developed phase dithered 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 a matrix 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 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 one bit-acquisition with additive dithering $\mathcal{Q}^+$ (⊕), the phase only acquisition $\mathrm{sign}_{\mathbb{C}}$ (⊖), and 1-bit acquisition with phase dithering $\mathcal{Q}^{\odot}$ (⊙). These 1-bit schemes with dithering are also compared with the deterministic 1-bit acquisition $\mathcal{Q}$ (○). Those quantized data are then processed with two different algorithms, namely PBP and QIHT.

Fig. 5 compares the different acquisition schemes using the PBP algorithm. One can first see that the additive dithering behaves as expected by the theory and has a ℓ_2 error that decreases as $O(M^{-\frac{1}{2}})$ (see [15]). While this behaviour guarantees an arbitrarily low reconstruction error for a number of measurements M large enough, for a low number of measurements (*e.g.*, $M \leq 4N$) it is clearly outperformed by the other low-resolution schemes. The random phase dithering achieves a reconstruction quality that is better than the additive dithering for any number of measurements presented in Fig. 5. Indeed, for $\mu \leq 1$ it follows closely the non-dithered curve ($\mathcal{Q}$) before outperforming it and resolving to the performances of the Phase-Only acquisition. This is consistent with Prop. 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 phase dithered scheme will still be bounded by the PO performances. It is important to note that, increasing the number of measurements has also a price in terms of data transfer, thus, the quality of the reconstruction of the phase dithered scheme combined with the relative low values of M required for its success makes it an appealing alternative.

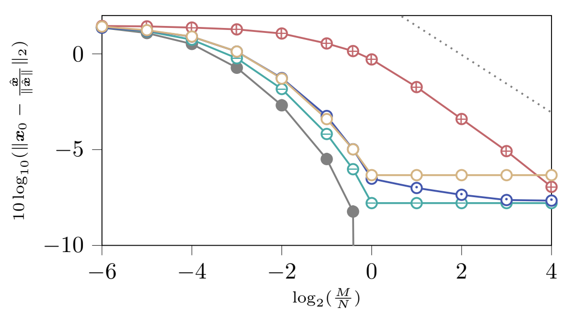

Fig. 5: Comparison of different reconstructions using PBP between different quantization schemes for $s = 10$: 1-bit with additive dither (⊕ in **red**); 1-bit without dither (○ in **yellow**); 1-bit with phase dither (⊙ in **blue**); Phase-Only acquisition (⊖ in **green**); without quantization (● in **gray**); the dotted curve (in **gray**) represent $O(M^{-\frac{1}{2}})$.

In Fig. 6, we now study the performances of an iterative algorithm, namely QIHT [21]. One can observe that the PO acquisition used in conjunction with QIHT algorithm yield high quality reconstructions. Consequently phase dithering, whose performances is bounded by PO reconstruction, outperforms the additive dither for any number of measurements below $2^4 N$. Indeed, it follows the deterministic quantization and then continues to exhibits a rate of $O(M^{-1})$ while the deterministic quantization plateaus at $\mu \geq 1$. It is interesting to see that his rate of $O(M^{-1})$ has also been observed in other 1-bit settings but with stronger theoretical guarantees [30]. We clearly see that, while the phase dithered case does not have as strong guarantees as the additive dither, it still has impressive performances and allows good reconstruction while keeping the number of measurements low, especially given the fact that it requires less complex hardware to apply the dither.

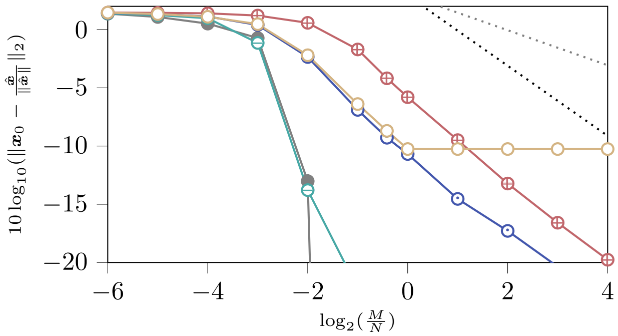

Fig. 6: Comparison of different reconstructions using QIHT between different quantization schemes for $s = 10$: 1-bit with additive dither (⊕ in **red**); 1-bit without dither (○ in **yellow**); 1-bit with phase dither (⊙ in **blue**); Phase-Only acquisition (⊖ in **green**); without quantization (● in **gray**); the dotted curve (in **gray**) represent $O(M^{-\frac{1}{2}})$, the black dotted curve $O(M^{-1})$.

VI. Conclusion

In this paper, we showed that using phase dithering with 1-bit quantized radar measurements is a viable trade-off between performances, hardware complexity and theoretical guarantees. Indeed, this work stems from the fact that using an additive dither increases dramatically the complexity of the hardware, negating the gains of the low resolution 1-bit acquisition. The phase dithering of radar signals can be implemented efficiently and relies on already existing hardware architectures. The developed theory showed that the reconstruction error of this scheme is linked to the one of the Phase-Only acquisition process. Next, the simulations showed that the proposed phase dithered 1-bit scheme provides a good trade-off between the number of measurements needed and the sensitivity to the sparsity level of the signal. A future publication will study an actual 1-bit radar prototype based on the proposed phase-dithered architecture.

References

[1] M. Arsalan, A. Santra, and V. Issakov, “Radarsnn: A resource efficient gesture sensing system based on mm-wave radar,” *IEEE Transactions on Microwave Theory and Techniques*, vol. 70, no. 4, pp. 2451–2461, 2022.

[2] 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.

[3] 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.

[4] A. Khalili, F. Shirani, E. Erkip, and Y. C. Eldar, *Mimo networks with one-bit adcs: Receiver design and communication strategies*, 2021.

[5] J. Liu, Z. Luo, and X. Xiong, “Low-resolution adcs for wireless communication: A comprehensive survey,” *IEEE Access*, vol. 7, pp. 91 291–91 324, 2019.

[6] 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.

[7] X. Dong and Y. Zhang, “A map approach for 1-bit compressive sensing in synthetic aperture radar imaging,” *IEEE Geoscience and Remote Sensing Letters*, vol. 12, no. 6, pp. 1237–1241, 2015.

[8] X. Wang, G. Li, Y. Liu, and M. G. Amin, “Enhanced 1-bit radar imaging by exploiting two-level block sparsity,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 2, pp. 1131–1141, 2019.

[9] 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.

[10] R. Walden, “Analog-to-digital converter survey and analysis,” *IEEE Journal on Selected Areas in Communications*, vol. 17, no. 4, pp. 539–550, 1999.

[11] 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.

[12] 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.

[13] Y. Plan and R. Vershynin, *One-bit compressed sensing by linear programming*, 2012.

[14] 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.

[15] 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.

[16] 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.

[17] 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. Cham: Springer International Publishing, 2019, pp. 67–95.

[18] 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.

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

[20] 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.

[21] 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.

[22] M. Friedlander, H. Jeong, Y. Plan, and Ö. Yilmaz, “Nbiht: An efficient algorithm for 1-bit compressed sensing with optimal error decay rate,” *ArXiv*, vol. abs/2012.12886, 2020.

[23] 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.

[24] L. Jacques and T. Feuillen, *The importance of phase in complex compressive sensing*, 2021.

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

[26] S. Foucart, “Flavors of compressive sensing,” in *Approximation Theory XV: San Antonio 2016* G. E. Fasshauer and L. L. Schumaker, Cham: Springer International Publishing, 2017, pp. 61–104.

[27] 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.

[28] 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.

[29] T. Feuillen, L. Vandendorpe, and L. Jacques, “A novel phase dithering scheme for 1-bit compressive radar,” *(in preparation)*, 2024.

[30] 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.