Blind ℓ_1-Norm-Based Phase Estimation for Cross-QAM Robust to Finite Precision Effects

Marcos Álvarez-Díaz*, *Student Member, IEEE,* Roberto López-Valcarce, *Member, IEEE,* and Iago Gómez Alonso

Abstract

We introduce a new blind criterion for carrier phase recovery of QAM modulated data, based on the maximization of the average of the q-th power of the phase-compensated received data vector ℓ_1-norm. In order to improve the performance of state-of-the-art higher-order estimators for cross-QAM constellations, which are sensitive to finite precision effects, our focus is on low-order methods ($q = 1$ and 2). Fixed-point iterations with good local convergence properties are given for these cases; different existing phase estimators can be used as initializers, depending on the operation environment and application constraints. A variance analysis of the novel estimators is derived, which shows that they are competitive against higher-order estimators. More importantly, Monte Carlo simulations show that the novel schemes remain robust to quantization effects and finite precision implementation, and thus they constitute an attractive choice in realistic scenarios.

Index Terms

Carrier offset, cross-QAM constellations, blind phase estimation, fixed point implementation.

* Corresponding author. M. Álvarez-Díaz and R. López-Valcarce are with the Department of Signal Theory and Communications, University of Vigo, 36310 Vigo, Spain (e-mail: {malvarez,valcarce}@gts.tsc.uvigo.es). I. Gómez Alonso is with the Galician Research and Development Center in Advanced Telecommunication (GRADIANT), ETSI Telecomunicación, 36310 Vigo, Spain (email: igomez@gradiant.org).

This work was supported by the European Commission through the IST Programme under Contract IST-027393 SatNEx-II and by the Spanish Government under projects SPROACTIVE (ref. TEC2007-68094-C02-01/TCM) and COMONSENS (CONSOLIDER-INGENIO 2010 CSD2008-00010). Parts of this paper were presented at the 13th IEEE Workshop on Statistical Signal Processing (SSP 2005), Bourdeaux, France, July 2005.

Blind ℓ_1-Norm-Based Phase Estimation for Cross-QAM Robust to Finite Precision Effects

I. INTRODUCTION

Carrier phase recovery is a key element in bandpass digital communication receivers. Many authors have addressed this topic over the last years, yielding a wide variety of recovery methods depending on the modulation scheme and application constraints [1]-[2]. The evolution of digital communications has fostered the use of denser constellations such as Quadrature Amplitude Modulation (QAM) at higher transmission rates, posing new challenges to the receiver operation in general; and in particular to the estimation of the unknown carrier phase offset, a task that can be performed resorting to either trained or non-data-aided (NDA, or *blind*) methods. Our focus is on NDA phase estimation schemes (since they provide better spectral efficiency) for QAM communication systems.

Existing NDA methods can be roughly classified into two categories. Histogram-based schemes [1], [3] estimate the probability density function (pdf) of the phase of a suitable nonlinear transformation of the observed data. The phase offset affects this pdf as a cyclic shift, from which it can be estimated. Although these methods present good performance, their computational load is relatively high. In addition, they require previous estimation of the channel gain [1] or the SNR operation point [3].

On the other hand, higher-order statistics (HOS)-based methods obtain the phase estimate as a function of the sample averages of nonlinear transformations of the data. This class includes the classical fourth-power (4P) estimator [4], the Viterbi and Viterbi (V&V) family of estimators [5], which were originally proposed for Phase-Shift Keying (PSK) modulation, and eighth-order estimates such as that of Cartwright (C8) [6] and the so-called concentration ellipse orientation (CEO) method [7]. All of these apply fixed nonlinear transformations to the observed data, and do not require previous gain control or knowledge of the SNR. Generalizations of the V&V family were proposed in [8], [9], where the nonlinearity is matched to the particular constellation and optimized in terms of the estimation variance. This approach provides good performance, but similarly to histogram-based methods, it requires knowledge of the overall gain and of the effective SNR.

Low computational complexity makes the standard 4P estimator very attractive for practical implementation. Its asymptotic variance was analyzed in [4], [2], showing that for a given sample size, it does not decrease with increasingly high SNR (except in the case of a QPSK constellation). This error floor

is induced by the so-called *self-noise* due to the multimodulus nature of QAM constellations. Despite this drawback, performance is considered acceptable for square QAM constellations (it is known that the NDA Maximum Likelihood estimator reduces to the 4P estimator as the SNR approaches zero [4]); however, its self-noise-induced error floor is much higher for cross-QAM constellations, in which several corner symbols in the constellation (those with larger modulus) are missing. For these cases, eighth-order methods such as C8 and CEO constitute a better choice, since they provide lower variance floors at medium to high SNR, at the expense of an increase in complexity.

Dense QAM constellations are very sensitive to phase errors, as seen in Fig. 1, which shows the symbol error rate (SER) versus the SNR per bit ($\mathrm{SNR_b}$) curves, obtained by simulation[1] for uncoded 32- and 128-QAM in additive white Gaussian noise (AWGN) under a number of phase offsets. Clearly, even relatively small phase errors may incur a substantial penalty in terms of $\mathrm{SNR_b}$ (e.g. 1.25 and 5.25 dB for 32- and 128-QAM respectively, for a phase offset of 3° at a target raw SER of 10^{-3}).

Practical implementation of digital receivers must usually face hardware limitations imposed by cost, size, speed, and power consumption constraints. In those situations, fixed point arithmetic devices may be the only available choice. Similarly, if high-speed analog-to-digital converters (ADCs) are required, the available resolution may not be sufficiently high so as to ignore quantization effects in the design. In particular, these effects will impact different phase estimation algorithms in a different way, so that the choice of an estimator will be determined by its behavior with quantized data and fixed point arithmetic processing, rather than by its theoretical asymptotic variance under infinite precision. Although estimation performance is difficult to analyze under finite wordlength constraints, intuitively one would expect that estimators requiring fewer, lower-order operations on the observed data be more robust in these situations. For example, the performance of the 4P phase estimate is expected to degrade more gracefully than that of eighth-order methods such as C8 and CEO.

These considerations motivate the search for low-complexity NDA phase estimators for cross-QAM systems. In this paper we develop two such schemes by maximizing certain cost functions arising from geometrical considerations based on the square-like shape of QAM constellations. Although it is not possible to derive closed-form expressions for these maximizers, we present simple iterations which locally converge to the desired values, and which can be initialized by any suitable scheme providing a coarse phase estimate. The choice of the initializer should be tailored to the operation environment (constellation size, number of samples, availability of training data, etc.).

[1]Closed-form expressions for the SER of cross-QAM in AWGN are known only for zero phase offset [10], [11].

The paper is organized as follows. Section II presents the signal model and reviews the classical one-shot estimators that will be considered in the paper. Maximization of the ℓ_1-norm as a criterion for phase estimation is justified in Section III, whereas the two proposed estimators are derived in Section IV. Section V provides an analysis in terms of variance, initialization and computational complexity. Simulation results, for both floating point and fixed point implementations, are shown in Section VI, and conclusions are drawn in Section VII.

Notation is as follows: for a complex number z, the real and imaginary parts are denoted as $\bar{z}$ and $\tilde{z}$ respectively, whereas $\arg\{z\}$ denotes the phase. Thus, $z = \bar{z} + j\tilde{z} = |z| \cdot \mathrm{e}^{j\arg\{z\}}$. Superscript * denotes complex conjugation. The *complex sign* of z is defined as $\mathrm{csgn}(z) \doteq \mathrm{sgn}(\bar{z}) + j\mathrm{sgn}(\tilde{z})$, where for real x, $\mathrm{sgn}(x) = x/|x|$ if $x \neq 0$ and zero otherwise. Throughout the paper we use the term *fixed point* with two different meanings. A point $x_0 \in \mathbb{R}$ is said to be a fixed point of a transformation $f : \mathbb{R} \to \mathbb{R}$ if $f(x_0) = x_0$. On the other hand, we refer to fixed point (as opposed to floating point) implementations of a given algorithm as the format used to store and manipulate numbers within a processing device. Which of the two meanings applies in each case should be clear from the context.

II. System model and classical estimators

Consider the receiver of QAM system, in which the received signal is sampled at the baud rate after front-end processing. Assuming that carrier frequency recovery has been previously established, the observed data can be written as

$$r_k = a_k \mathrm{e}^{j\theta} + n_k, \qquad k = 0, 1, \ldots, L-1, \tag{1}$$

where $\{a_k\}$ are the transmitted symbols, drawn equiprobably from a constellation $\mathcal{A}$ with size M and variance E_a, and $\{n_k\}$ are the complex-valued noise samples. The noise is assumed zero-mean, circular white Gaussian with variance σ^2, and independent of the symbols. The phase offset θ is assumed to vary slowly; thus, it can be taken as constant within the block of L samples. The goal is to identify θ without knowledge of the symbols a_k or the variances E_a, σ^2. Due to the quadrant symmetry of QAM constellations, this phase offset can only be blindly identified up to a four-fold ambiguity; hence, we assume that $|\theta| < \pi/4$.

The V&V family of estimators [5] embraces a wide range of choices. Of particular interest are power-law estimators

$$\hat{\theta}_{\text{V\&V-}p} = \frac{1}{4}\arg\left\{ -\sum_{k=0}^{L-1} |r_k|^p \mathrm{e}^{j4\arg\{r_k\}} \right\}, \tag{2}$$

where p is an integer. For $p = 0$, the corresponding estimate $\hat{\theta}_{\text{VV-0}}$ discards the envelope information of the received signal. On the other hand, for $p = 4$ one obtains the 4P estimator [4], [2]:

$$\hat{\theta}_{\text{4P}} = \frac{1}{4}\arg\left\{-\sum_{k=0}^{L-1} r_k^4\right\}, \tag{3}$$

In order to improve the performance of the 4P method with cross-QAM constellations, Cartwright [6] proposed the eighth-order estimator

$$\hat{\theta}_{\text{C8}} = \frac{1}{4}\tan^{-1}\left(\frac{\sum_k A_k B_k \sum_l A_l C_l - \sum_k A_k^2 \sum_l B_l C_l}{\sum_k A_k B_k \sum_l B_l C_l - \sum_k B_k^2 \sum_l A_l C_l}\right), \tag{4}$$

where $A_k \doteq \text{Re}\{r_k^4\}$, $B_k \doteq \text{Im}\{r_k^4\}$ and $C_k \doteq |r_k|^4$, and all summations in k and l in (4) run from 0 to $L-1$. We note that the four quadrant inverse tangent function is required in (4). The price to pay for the reduced variance of (4) with respect to that of (3) is a higher computational load. A different eighth-order estimator (CEO) was proposed in [7]; its computational complexity and performance with cross-QAM constellations are similar to those of (4).

III. Phase estimation based on ℓ_1-norm maximization

Consider the mapping $f : \mathbb{C} \to \mathbb{R}^2$ that assigns to each complex number $z = \bar{z} + j\tilde{z}$ the real vector $\boldsymbol{z} \doteq f(z) = [\bar{z} \quad \tilde{z}]^T$. Under f, multiplication by a phase term $\text{e}^{j\theta}$ in $\mathbb{C}$ becomes multiplication by an orthogonal rotation matrix in $\mathbb{R}^2$. Given a candidate estimate $\hat{\theta}$, let us define the de-rotated samples

$$y_k = \text{e}^{-j\hat{\theta}} r_k = \bar{y}_k + j\tilde{y}_k, \tag{5}$$

and consider the effect of such de-rotation on the vector $\boldsymbol{y}_k = f(y_k) = [\bar{y}_k \quad \tilde{y}_k]^T$. In particular, we focus on the resulting ℓ_1-norm of the phase-compensated vector

$$||\boldsymbol{y}_k||_1 = |\bar{y}_k| + |\tilde{y}_k|. \tag{6}$$

We claim that maximization of the ℓ_1-norm with respect to $\hat{\theta}$ is a meaningful criterion for obtaining a phase estimate with QAM constellations. To see this, it is instructive to consider the geometry of the simplest case, i.e. a QPSK constellation ($M = 4$) $\mathcal{A} = \{\sqrt{E_a} e^{j(\frac{\pi}{4} + n\frac{\pi}{2})},\ 0 \le n \le 3\}$ in the noiseless case. Recall that the ℓ_1-ball defined by $||\boldsymbol{y}||_1 = \gamma$ is a diamond with corners at $[0 \quad \pm\gamma]^T$ and $[\pm\gamma \quad 0]^T$. It is clear from Fig. 2 that any rotation of $\mathcal{A}$ by an angle that is not of the form $n\pi/2$ with n integer can only result in a smaller average ℓ_1-norm of the rotated constellation. That is, the ℓ_1-norm is maximized when the positioning of the constellation is restored to its original regular grid.

Observe that for QPSK the average ℓ_1-norm of the constellation coincides with the ℓ_1-norm of any of its elements, for all phase offsets. This is because this constellation can be obtained by picking any of its elements and all possible rotations by an integer multiple of $\pi/2$ rad, and the ℓ_1-norm is preserved by these operations, as it can be easily checked[2]. For larger constellations this property does not hold anymore, although it can be checked that the average ℓ_1-norm still attains its maximum for a zero phase offset. This is shown in Figs. 3 and 4, which plot the costs

$$J_q(\hat{\theta}) \doteq E\{||\boldsymbol{y}||_1^q\} = E\{(|\bar{y}| + |\tilde{y}|)^q\} \tag{7}$$

for $q = 1$ and 2, in terms of the residual phase error $\theta - \hat{\theta}$. The jagged appearance of J_1 and J_2 is due to the fact that the ℓ_1-norm is not everywhere differentiable. Besides the desired global maximum at zero phase offset, these costs present additional local maxima for 16-QAM as well as for cross-QAM constellations; for square QAM with $M > 16$, both are unimodal.

IV. Fixed point iterations for phase estimation

Since no closed-form expressions for the maxima of $J_q(\hat{\theta})$ are available, one must resort to numerical optimization strategies, e.g. gradient ascent or Newton's method. We propose alternative fixed point iterations which are computationally simpler than Newton's method and, in contrast to gradient ascent, do not require stepsize tuning. Our focus is on $J_q(\hat{\theta})$ for $q = 1$ and 2, as their lower order yields computationally simpler schemes robust to finite precision effects, as will be shown in Section VI.

A. Maximization of $J_1(\hat{\theta})$

The goal is to find a maximum of $J_1(\hat{\theta}) \doteq \mathrm{E}\{|\bar{y}_k| + |\tilde{y}_k|\}$. If J_1 is differentiable at such point, then $\partial J_1(\hat{\theta})/\partial\hat{\theta}$ must vanish. From (5), note that

$$\frac{\partial y_k}{\partial\hat{\theta}} = -jy_k \quad \Rightarrow \quad \frac{\partial\bar{y}_k}{\partial\hat{\theta}} = \tilde{y}_k,\ \frac{\partial\tilde{y}_k}{\partial\hat{\theta}} = -\bar{y}_k. \tag{8}$$

Therefore, at the points at which J_1 is differentiable,

$$\frac{\partial J_1(\hat{\theta})}{\partial\hat{\theta}} = \mathrm{E}\left\{\mathrm{sgn}\left(\bar{y}_k\right)\tilde{y}_k - \mathrm{sgn}\left(\tilde{y}_k\right)\bar{y}_k\right\}. \tag{9}$$

If $\hat{\theta}_*$ is a maximum of J_1, then (9) equals zero (if it exists), i.e.

$$\mathrm{E}\left\{\mathrm{sgn}\left(\bar{y}_k\right)\tilde{y}_k\right\} = \mathrm{E}\left\{\mathrm{sgn}\left(\tilde{y}_k\right)\bar{y}_k\right\}. \tag{10}$$

[2]Conjugation is another ℓ_1-norm-preserving operation.

Now, if we write $r_k = \bar{r}_k + j\tilde{r}_k$, then from (5)

$$\bar{y}_k = \bar{r}_k \cos\hat{\theta} + \tilde{r}_k \sin\hat{\theta}, \quad \tilde{y}_k = \tilde{r}_k \cos\hat{\theta} - \bar{r}_k \sin\hat{\theta}, \tag{11}$$

and we arrive at

$$\begin{aligned} \tan\hat{\theta}_* &= \frac{\mathrm{E}\left\{\mathrm{sgn}\left(\bar{y}_k\right)\tilde{r}_k - \mathrm{sgn}\left(\tilde{y}_k\right)\bar{r}_k\right\}}{\mathrm{E}\left\{\mathrm{sgn}\left(\bar{y}_k\right)\bar{r}_k + \mathrm{sgn}\left(\tilde{y}_k\right)\tilde{r}_k\right\}} && (12)\\ &= \frac{-\mathrm{E}\left\{\mathrm{Im}\left\{\mathrm{csgn}\left(y_k\right)\cdot r_k^*\right\}\right\}}{\mathrm{E}\left\{\mathrm{Re}\left\{\mathrm{csgn}\left(y_k\right)\cdot r_k^*\right\}\right\}} && (13)\\ &= \tan\left(-\mathrm{arg}\left\{\mathrm{E}\left\{\mathrm{csgn}\left(y_k\right)\cdot r_k^*\right\}\right\}\right). && (14)\end{aligned}$$

Note that (14) characterizes the extrema of J_1 only implicitly, since the right-hand side of (14) depends on $\hat{\theta}_*$. Nevertheless, this condition suggests a fixed point iteration to obtain $\hat{\theta}_*$. Given the observations $\{r_k\}$ and a suitable initialization $\hat{\theta}_0^{(1)}$, we substitute the expectation in (14) by a sample mean and then iteratively compute

$$\hat{\theta}_{n+1}^{(1)} = -\mathrm{arg}\left\{\sum_{k=0}^{L-1} \mathrm{csgn}\left(r_k \cdot \mathrm{e}^{-j\hat{\theta}_n^{(1)}}\right)\cdot r_k^*\right\}. \tag{15}$$

B. Maximization of $J_2(\hat{\theta})$

The cost J_2 can be written as

$$J_2(\hat{\theta}) = \mathrm{E}\left\{\left(|\bar{y}_k| + |\tilde{y}_k|\right)^2\right\} = \mathrm{E}\left\{|y_k|^2\right\} + 2\cdot\mathrm{E}\left\{|\bar{y}_k|\cdot|\tilde{y}_k|\right\}. \tag{16}$$

Since $\mathrm{E}\left\{|y_k|^2\right\}$ is invariant under phase rotations, maximizing $J_2(\hat{\theta})$ amounts to maximizing $\mathrm{E}\left\{|\bar{y}_k\tilde{y}_k|\right\}$. Using basic trigonometric relations, it is found that

$$\bar{y}_k\tilde{y}_k = \bar{r}_k\tilde{r}_k\cos(2\hat{\theta}) - \frac{1}{2}(\bar{r}_k^2 - \tilde{r}_k^2)\sin(2\hat{\theta}). \tag{17}$$

Therefore, at the points at which J_2 is differentiable,

$$\frac{\partial J_2(\hat{\theta})}{\partial\hat{\theta}} = -\mathrm{E}\left\{\mathrm{sgn}\left(\bar{y}_k\tilde{y}_k\right)\ \left(\bar{r}_k^2 - \tilde{r}_k^2\right)\right\}\cos(2\hat{\theta}) - \mathrm{E}\left\{\mathrm{sgn}\left(\bar{y}_k\tilde{y}_k\right)\ 2\bar{r}_k\tilde{r}_k\right\}\sin(2\hat{\theta}). \tag{18}$$

Observe that

$$\begin{aligned} \bar{r}_k^2 - \tilde{r}_k^2 &= \mathrm{Re}\left\{r_k^2\right\}, \quad 2\bar{r}_k\tilde{r}_k = \mathrm{Im}\left\{r_k^2\right\}, && (19)\\ s_k(\hat{\theta}) &\doteq \mathrm{sgn}\left(\bar{y}_k\tilde{y}_k\right) = \mathrm{sgn}\left(\mathrm{Im}\left\{r_k^2\mathrm{e}^{-j2\hat{\theta}}\right\}\right). && (20)\end{aligned}$$

If $\hat{\theta}_*$ is a maximum of J_2 at which J_2 happens to be differentiable, then (18) must be zero, i.e.,

$$\begin{aligned} \tan(2\hat{\theta}_*) &= -\frac{\mathrm{E}\left\{\mathrm{Re}\left\{r_k^2\right\}s_k(\hat{\theta}_*)\right\}}{\mathrm{E}\left\{\mathrm{Im}\left\{r_k^2\right\}s_k(\hat{\theta}_*)\right\}} && (21)\\ &= \tan\left(\mathrm{arg}\left\{-j\mathrm{E}\left\{r_k^2\ s_k(\hat{\theta}_*)\right\}\right\}\right). && (22)\end{aligned}$$

Since the right-hand side of (22) depends on $\hat{\theta}_*$, this condition does not provide the desired solution in closed form, but similarly to (15), it suggests a fixed point iteration. Given the observations $\{r_k\}$ and a suitable initialization $\hat{\theta}_0^{(2)}$, the expectation in (22) is replaced by a sample mean and then $\hat{\theta}_n^{(2)}$ is iteratively computed as follows:

$$s_k(\hat{\theta}_n^{(2)}) = \text{sgn}\left(\text{Im}\left\{r_k^2 \text{e}^{-j2\hat{\theta}_n^{(2)}}\right\}\right),\ k = 0, \ldots, L-1, \tag{23a}$$

$$\hat{\theta}_{n+1}^{(2)} = \frac{1}{2}\ \text{arg}\left\{\sum_{k=0}^{L-1}\left[r_k^2 s_k(\hat{\theta}_n^{(2)})\right]\right\} - \frac{\pi}{4}. \tag{23b}$$

V. Discussion and Analysis

A. Algorithm initialization

The costs $J_q(\hat{\theta})$ are multimodal in general, due to finite sample effects (with cross-QAM, multimodality is the rule even as $L \to \infty$), and hence the proposed iterative schemes may experience convergence to undesirable solutions. Initialization must be good enough to set the starting point within the domain of attraction of the global maximum. Potential initializers include those methods discussed in Section II. The use of the novel iterative methods as refinements of the standard estimators of Section II is justified only if they are able to provide better performance than the initializer, and at a reasonable computational cost. Both issues are discussed next.

B. Computational complexity

Consider the standard 4P estimator (3). Obtaining the averaged value $\sum_k r_k^4$ requires $6L$ real multiplications and $4L$ real additions. On the other hand, computation of the C8 estimate (4) requires $11L$ real multiplications and $8L$ real additions. Thus, the C8 estimate is about twice as expensive as the standard fourth-phase estimate.

However, quantifying the computational load associated to a given estimator is not always such a straightforward task. Consider for instance the estimate $\hat{\theta}_{\text{VV-0}}$ from (2), which must compute $\sum_k \text{e}^{j4\text{arg}\{r_k\}} = (\sum_k \cos(4\text{arg}\,\{r_k\}) + j\sum_k \sin(4\text{arg}\,\{r_k\})$. This in turn requires (i) extraction of the phase of the complex numbers r_k, i.e., Cartesian-to-Polar (C2P) conversion; (ii) computation of the terms $\cos(4\text{arg}\,\{r_k\})$ and $\sin(4\text{arg}\,\{r_k\})$; and (iii) averaging these values. The last step amounts to $2L$ real additions; however, the first two steps are likely to be implemented as look-up table (LUT) operations. Thus, complexity comparisons of the 4P, C8, and V&V-0 estimators will depend on the relative costs assigned to memory

area, hardware multipliers, etc.[3]

In a similar vein, iteration (15) requires L phase rotations (in order to obtain csgn $\left(r_k \cdot \mathrm{e}^{-j\hat{\theta}_n^{(1)}}\right)$) and $2L$ real additions per iteration. Phase rotations can be performed very efficiently (with shift-and-add operations only) with CORDIC hardware [12], [13]. Alternatively, with a previous C2P module providing $|r_k|$ and arg $\{r_k\}$, determination of the csgn terms only requires L additions per iteration. The scheme (23) is slightly more costly than (15) due to the need to precompute r_k^2, although it can also benefit from C2P/CORDIC modules for the computation of $s_k(\hat{\theta}_n^{(2)})$. Complexity evaluation of the proposed estimators should also take into account the need for proper initialization (and also the fact that some of the partial results of the initialization stage can be reused in subsequent iterations).

C. Asymptotic Variances

Using a small-error analysis, it is possible to obtain closed-form expressions for the asymptotic variances of the estimates $\hat{\theta}^{(q)} \doteq \lim_{n\to\infty} \hat{\theta}_n^{(q)}$, $q = 1, 2$. Details are given in the Appendix, and the results are summarized in the following theorem.

Theorem 5.1: Assuming that the iterations (15) and (23) are initialized sufficiently close to the true phase offset θ and that the Signal-to-Noise Ratio $\eta \doteq E_a/\sigma^2$ is sufficiently high, the estimates $\hat{\theta}^{(1)}$ and $\hat{\theta}^{(2)}$ are asymptotically unbiased, and their asymptotic variances are given by

$$\lim_{L\to\infty} L\mathrm{E}\left\{(\hat{\theta}^{(1)} - \theta)^2\right\} = \frac{1}{2d_1^2}\left((2 - d_2) + \frac{2}{\eta}\right), \tag{24}$$

$$\lim_{L\to\infty} L\mathrm{E}\left\{(\hat{\theta}^{(2)} - \theta)^2\right\} = \frac{1}{4\tilde{g}_2^2}\left(\frac{c_4 + \bar{b}_4}{2} + \frac{2}{\eta} + \frac{1}{\eta^2}\right), \tag{25}$$

where b_4, c_4, d_n, g_2 are constellation-dependent constants:

$$b_4 \doteq \frac{\mathrm{E}\{a^4\}}{E_a^2}, \qquad c_4 \doteq \frac{\mathrm{E}\{|a|^4\}}{E_a^2},$$

$$d_n \doteq \frac{\mathrm{E}\{(a^*\mathrm{csgn}(a))^n\}}{E_a^{n/2}}, \qquad g_2 \doteq \frac{\mathrm{E}\{a^2\mathrm{sgn}(\mathrm{Im}\{a^2\})\}}{E_a}.$$

The asymptotic variances (24)-(25) depend on the constellation and the SNR, but not on the value of θ; this is also the case for the 4P estimate, whose asymptotic variance was given in [2]. Using that expression and (24)-(25), one can obtain the ratio of limits (as the SNR tends to infinity) of the asymptotic variances of the new estimates to that of the 4P estimate

$$\alpha_q \doteq \frac{\lim_{\eta,L\to\infty} L\mathrm{E}\left\{(\hat{\theta}^{(q)} - \theta)^2\right\}}{\lim_{\eta,L\to\infty} L\mathrm{E}\left\{(\hat{\theta}_{4\mathrm{P}} - \theta)^2\right\}}. \tag{26}$$

[3]Note that the operation counts given above for the 4P and C8 schemes will change if a previous LUT-based C2P stage is applied to the observed data, which makes direct comparison with V&V-0 even more difficult.

These ratios are given in Table I, showing that the new schemes perform similarly to the 4P estimator with square QAM, but they consistently offer a sizable improvement (about one order of magnitude reduction of the asymptotic variance) with cross-QAM constellations.

VI. Simulation Results

We present the results obtained with the novel iterative estimators, using Monte Carlo simulations, under both floating point and fixed point implementations. The focus is on 32- and 128-QAM cross constellations.

A. Floating point precision

Convergence properties were investigated first, with a phase offset $\theta = 20^\circ$. Fig. 5 shows the Root Mean Square Error (RMSE) as a function of the number of iterations. Both estimators were initialized using the 4P method (the RMSE shown for iteration zero corresponds to $\hat{\theta}_{4\mathrm{P}}$), and different values of the SNR per bit [defined as $\mathrm{SNR_b} \doteq E_a/(\sigma^2 \log_2 M)$] are considered. Convergence slows down with decreasing $\mathrm{SNR_b}$ as expected, but nevertheless, for practical situations a few iterations suffice.

Note from Fig. 5 that for 128-QAM and $\mathrm{SNR_b} = 30$ dB, the RMSE actually *increases* after the second iteration. This effect is due to the fact that, occasionally, the initializer ($\hat{\theta}_{4\mathrm{P}}$ in this case) will fail to set the starting point within the domain of attraction of the desired solution, so that convergence to an undesired local maximum of the cost J_q takes place. The nonzero probability of these events pulls up the final RMSE of the estimates.

Fig. 6 shows the RMSE of the proposed estimators (fixing the number of iterations to five), together with those of $\hat{\theta}_{4\mathrm{P}}$ and $\hat{\theta}_{\mathrm{C8}}$, as a function of the SNR per bit. For 32-QAM, the iterative estimators are initialized at $\hat{\theta}_{4\mathrm{P}}$; they perform similarly to $\hat{\theta}_{\mathrm{C8}}$ in the high $\mathrm{SNR_b}$ region, although the latter degrades more gracefully as the noise increases. Good agreement with the theoretical variance predicted by (24)-(25) is observed for $\mathrm{SNR_b} \geq 14$ dB.

Whereas $\hat{\theta}_{4\mathrm{P}}$ seems to be a sufficiently good initializer with 32-QAM, the situation is different for 128-QAM: occasional convergence to local maxima prevents the RMSE of the iterative estimators from reaching their theoretical asymptotic values. Another option is to use $\hat{\theta}_{\mathrm{C8}}$ as initializer. Note from Fig. 6 that the theoretical RMSE of the iterative estimators with 128-QAM in high SNR lies below the observed RMSE of $\hat{\theta}_{\mathrm{C8}}$, so that it makes sense to apply these schemes as refinements to $\hat{\theta}_{\mathrm{C8}}$. As shown, this initializer avoids the problem of convergence to local maxima. Of course, all these considerations depend on the number of samples L used for the estimation; too low a value of L could result in $\hat{\theta}_{\mathrm{C8}}$ not being a good

initializer either. But in general, it can be said that the novel estimators have the potential to improve the performance of either fourth- or eighth-order schemes with low computational complexity. This potential is even greater under finite precision constraints, as discussed next.

B. Fixed point implementation

The performance of the different estimators is evaluated under quantization and fixed point arithmetic conditions. The goal is to test the robustness of each method under realistic implementation constraints. The receiver ADC quantizes each real-valued (in-phase and quadrature) data sample to B bits. The binary point is set to the right of the first bit, which is the sign bit. The full-scale value of the ADC is chosen in order to set the probability of clipping at 10^{-3} under AWGN, measured at a reference SNR for which a raw SER of 10^{-1} is obtained.

When working with fixed point systems, care must be taken in order to minimize the effects of underflow and overflow. Appropriate scalings must be introduced at certain stages in order to obtain a suitable tradeoff between the probabilities of these two undesirable events, denoted by P_{u} and P_{o} respectively. The locations and values (usually powers of two, for efficiency reasons) of the scaling factors must be determined by careful study of P_{u} and P_{o} at each of the intermediate computations. Of the estimators considered in this paper, it was found that scaling was beneficial for the C8 and 4P schemes when $B \leq 18$ and 9 bits respectively, but not for the proposed iterative methods. This confirms what one would expect by considering the number of products involved in each case.

In the simulations, wordlengths ranging from $B = 30$ down to 8 bits have been tested, and the scalings were optimized for each estimator. When B is sufficiently high, the convergence and the RMSE of all estimators follow the floating point results shown in the previous section. However, the performance of the methods tested degrades below a certain value of B, which depends on the estimator and the constellation. It was observed that in general $\hat{\theta}_{\mathrm{C8}}$ is more sensitive to finite wordlength effects than the rest of estimators, whereas $\hat{\theta}^{(1)}$ and $\hat{\theta}^{(2)}$ are much more robust in this sense. Next we review in detail the results of our fixed point simulations in terms of bias and RMSE for the 32-QAM and 128-QAM constellations.

1) Fixed point bias: All estimators are unbiased for sufficiently large wordlengths, but as B is decreased an offset-dependent bias appears, whose magnitude grows inversely to B. Figs. 7 and 8 show the bias for 32- and 128-QAM respectively in terms of θ, at $\mathrm{SNR_b} = 30$ dB and for different wordlengths. At such high SNR values, the bias can only be attributable to finite precision effects. Clearly, $\hat{\theta}_{\mathrm{C8}}$ is more severely affected than the rest of estimators. For 32-QAM, $\hat{\theta}_{\mathrm{C8}}$ already shows bias at $B = 14$ bits, and

for $B \leq 10$ bits it becomes useless. The remaining estimators degrade more gracefully, with $\hat{\theta}_{4\text{P}}$ and $\hat{\theta}^{(2)}$ exhibiting the smallest peak bias (about 0.3° and 1.0° for $B = 10$ and 8 bits respectively). For 128-QAM, $\hat{\theta}_{\text{C8}}$ is seen to have a peak bias of 2° already for $B = 14$ bits. Note that a 2° phase error would incur an SNR_b penalty of 3.0 dB for a SER of 10^{-3} (see Fig. 1b). The bias of the other schemes becomes noticeable only for $B \leq 12$ bits.

2) Fixed point RMSE: Results for 32-QAM are given in Fig. 9 for $L = 1024$ symbols, $\theta = 20^\circ$ and different values of B. With $B > 12$ bits, all methods perform similarly to the floating point case. For $B = 12$ bits, only $\hat{\theta}_{\text{C8}}$ shows a small degradation, which becomes very large for smaller wordlengths. In contrast, $\hat{\theta}_{4\text{P}}$ is not as sensitive (the small decrease in RMSE with respect to the finite precision curve is achieved in exchange for a slight bias), similarly to the novel estimators $\hat{\theta}^{(1)}$, $\hat{\theta}^{(2)}$, which achieve a much smaller asymptotic RMSE.

Results for 128-QAM, see Fig. 10, roughly follow those for 32-QAM, but with all estimators suffering larger performance losses. For $B = 14$ and 12 bits, $\hat{\theta}_{\text{C8}}$ exhibits a large RMSE of about 2° (due to its large bias, see Fig. 8). Nevertheless, it still provides a good initializer for $\hat{\theta}^{(1)}$ and $\hat{\theta}^{(2)}$, which reach their asymptotic variances at high SNR_b. For $B = 10$ the high RMSE of $\hat{\theta}_{\text{C8}}$ (3.3°) pulls up the RMSEs of $\hat{\theta}^{(1)}$ and $\hat{\theta}^{(2)}$ when they start from $\hat{\theta}_{\text{C8}}$. The performance of $\hat{\theta}_{4\text{P}}$ is not affected by finite wordlengths as small as $B = 10$ bits, and thus when this estimate is used to initialize $\hat{\theta}^{(1)}$ and $\hat{\theta}^{(2)}$, these achieve the same RMSE as in the floating point implementation (see Fig. 6b).

VII. Conclusions

The relatively poor behavior of classical NDA phase estimators with cross-QAM constellations motivated the development of higher-order schemes with improved performance. These, however, are more sensitive to finite precision issues. The estimators proposed in this paper attempt to sidestep this problem: they achieve a low variance and are less prone to these harmful effects due to their low order, as simulations have confirmed. Hence, these new schemes constitute appealing choices for receiver design under tight wordlength constraints.

A key issue in the proposed iterative estimators is their initialization. We have seen that standard one-shot methods (eighth- and fourth-order) can provide adequate coarse estimators for this purpose. Different strategies are also possible: for example, the low complexity and adequate performance of the 0-th order Viterbi&Viterbi estimator makes it attractive as initializer for 32-QAM. Yet another possibility arises whenever training symbols are inserted in the transmitted frame, as it is usually the case in practice. A data-aided phase estimator based just on these training symbols may well present a variance too large

for decoding purposes (due to a limited number of training data), but sufficiently low to initialize the proposed iterative schemes. Such *semiblind* approach would significantly improve the quality of the plain data-aided phase estimate.

APPENDIX

PROOF OF THEOREM 5.1

Let $\hat{\theta}^{(1)}$ be a fixed point of (15). In order to prove (24), let us define

$$\hat{\rho} \doteq \frac{1}{L}\sum_{k=0}^{L-1} r_k^* \mathrm{csgn}\left(r_k \mathrm{e}^{-j\hat{\theta}^{(1)}}\right). \tag{27}$$

Then $\hat{\theta}^{(1)}$ satisfies $\hat{\theta}^{(1)} = -\arg\{\hat{\rho}\}$. Let us also define

$$\rho \doteq \mathrm{E}\left\{r_k^* \mathrm{csgn}\left(r_k \mathrm{e}^{-j\theta}\right)\right\} = \bar{\rho} + j\tilde{\rho}, \tag{28}$$

where θ is the true phase offset. Then, similarly to [2], we can write

$$\hat{\theta}^{(1)} = -\tan^{-1}\left(\frac{\mathrm{Im}\{\hat{\rho}\}}{\mathrm{Re}\{\hat{\rho}\}}\right) = -\tan^{-1}\left(\frac{\mathrm{Im}\{\rho\} - \mathrm{Im}\{\hat{\rho}-\rho\}}{\mathrm{Re}\{\rho\} - \mathrm{Re}\{\hat{\rho}-\rho\}}\right), \tag{29}$$

so that a first-order approximation of the argument in (29) yields

$$\hat{\theta}^{(1)} \approx -\tan^{-1}\left(\frac{\tilde{\rho}}{\bar{\rho}} + \epsilon\right), \tag{30}$$

$$\epsilon \doteq \frac{1}{\bar{\rho}}\left(\mathrm{Im}\{\hat{\rho}-\rho\} - \frac{\tilde{\rho}}{\bar{\rho}}\mathrm{Re}\{\hat{\rho}-\rho\}\right). \tag{31}$$

Note that $r_k \mathrm{e}^{-j\theta} = a_k + m_k$, where $m_k = n_k \mathrm{e}^{-j\theta}$ is a noise process with the same statistical properties as n_k. Hence, if the SNR is sufficiently high, $\mathrm{csgn}\left(r_k \mathrm{e}^{-j\theta}\right) = \mathrm{csgn}\,(a_k + m_k) \approx \mathrm{csgn}\,(a_k)$, so that

$$\rho \approx \mathrm{E}\{r_k^* \mathrm{csgn}\,(a_k)\} \tag{32}$$

$$= \underbrace{\mathrm{E}\{a_k^* \mathrm{csgn}\,(a_k)\}}_{=\sqrt{E_a}d_1} \mathrm{e}^{-j\theta} + \underbrace{\mathrm{E}\{n_k^* \mathrm{csgn}\,(a_k)\}}_{=0}, \tag{33}$$

and the approximation error in (32) decays exponentially fast with the SNR. The statistic d_1 turns out to be real and positive for QAM constellations, and therefore $-\tan^{-1}(\tilde{\rho}/\bar{\rho}) \to \theta$ as the SNR goes to infinity. Then a first-order expansion of (30) yields

$$\hat{\theta}^{(1)} \approx \theta - \epsilon \cdot \cos^2\theta. \tag{34}$$

With the asymptotic approximation $\rho \approx \sqrt{E_a} d_1 \mathrm{e}^{-j\theta}$, (31) becomes

$$\epsilon \approx \mathrm{Im}\left\{(\hat{\rho}-\rho)\mathrm{e}^{j\theta}\right\} / (\bar{\rho}\cos\theta). \tag{35}$$

Note that if the SNR is sufficiently large and the phase error $\theta - \hat{\theta}^{(1)}$ is sufficiently small, then $\hat{\rho}$ is an unbiased estimate of ρ:

$$\begin{aligned}
\mathrm{E}\left\{\hat{\rho}\right\} &= \frac{1}{L}\sum_{k=0}^{L-1}\mathrm{E}\left\{r_k^*\mathrm{csgn}\left(r_k\mathrm{e}^{-j\hat{\theta}^{(1)}}\right)\right\} \\
&= \frac{1}{L}\sum_{k=0}^{L-1}\mathrm{E}\left\{r_k^*\mathrm{csgn}\left(a_k\mathrm{e}^{j(\theta-\hat{\theta}^{(1)})}+n_k\mathrm{e}^{-j\hat{\theta}^{(1)}}\right)\right\} \\
&\approx \frac{1}{L}\sum_{k=0}^{L-1}\mathrm{E}\left\{r_k^*\mathrm{csgn}\left(a_k\right)\right\} \;\approx\; \rho, \qquad (36)
\end{aligned}$$

Hence, from (34)-(36), $\hat{\theta}^{(1)}$ is asymptotically unbiased.

In order to compute the asymptotic variance, some straightforward manipulations show that

$$\mathrm{Im}^2\left\{(\hat{\rho}-\rho)\mathrm{e}^{j\theta}\right\} = \frac{1}{2}\left[|\hat{\rho}-\rho|^2 - \mathrm{Re}\left\{(\hat{\rho}-\rho)^2\mathrm{e}^{j2\theta}\right\}\right]. \qquad (37)$$

Therefore, from (34)-(35) and (37),

$$\begin{aligned}
\lim_{L\to\infty} L\mathrm{E}\left\{(\hat{\theta}^{(1)}-\theta)^2\right\} &= \cos^4\theta \lim_{L\to\infty} L\mathrm{E}\left\{\epsilon^2\right\} \qquad (38)\\
&= \frac{\cos^2\theta}{2\bar{\rho}^2}\lim_{L\to\infty} L\left[s_1 - \mathrm{Re}\left\{s_2\mathrm{e}^{j2\theta}\right\}\right], \qquad (39)
\end{aligned}$$

where s_1, s_2 are given by

$$\begin{aligned}
s_1 &\doteq \mathrm{E}\left\{|\hat{\rho}-\rho|^2\right\} = \mathrm{E}\left\{|\hat{\rho}|^2 - |\rho|^2\right\} \\
&\approx \frac{1}{L^2}\sum_{k=0}^{L-1}\sum_{l=0}^{L-1}\mathrm{E}\left\{r_k^* r_l \mathrm{csgn}\left(r_k\mathrm{e}^{-j\theta}\right)\mathrm{csgn}^*\left(r_l\mathrm{e}^{-j\theta}\right)\right\} - |\rho|^2 \qquad (40)\\
&= \frac{(L^2-L)}{L^2}\left|\mathrm{E}\left\{r_k^*\mathrm{csgn}\left(r_k\mathrm{e}^{-j\theta}\right)\right\}\right|^2 + \frac{1}{L^2}\sum_{k=0}^{L-1}\mathrm{E}\left\{|r_k|^2\left|\mathrm{csgn}\left(r_k\mathrm{e}^{-j\theta}\right)\right|^2\right\} - |\rho|^2 \qquad (41)\\
&= \left(1-\frac{1}{L}\right)|\rho|^2 + \frac{2}{L}\mathrm{E}\left\{|r_k|^2\right\} - \rho^2 \qquad (42)\\
&= \frac{1}{L}\left(2E_a + 2\sigma^2 - |\rho|^2\right), \qquad (43)
\end{aligned}$$

$$\begin{aligned}
s_2 &\doteq \mathrm{E}\left\{(\hat{\rho}-\rho)^2\right\} = \mathrm{E}\left\{\hat{\rho}^2 - \rho^2\right\} \\
&\approx \frac{1}{L^2}\sum_{k=0}^{L-1}\sum_{l=0}^{L-1}\mathrm{E}\left\{r_k^* r_l^* \mathrm{csgn}\left(r_k\mathrm{e}^{-j\theta}\right)\mathrm{csgn}\left(r_l\mathrm{e}^{-j\theta}\right)\right\} - \rho^2 \qquad (44)\\
&= \frac{(L^2-L)}{L^2}\mathrm{E}^2\left\{r_k^*\mathrm{csgn}\left(r_k\mathrm{e}^{-j\theta}\right)\right\} + \frac{1}{L^2}\sum_{k=0}^{L-1}\mathrm{E}\left\{\left(r_k^*\mathrm{csgn}\left(r_k\mathrm{e}^{-j\theta}\right)\right)^2\right\} - \rho^2 \qquad (45)\\
&\approx \left(1-\frac{1}{L}\right)\rho^2 + \frac{1}{L}\mathrm{E}\left\{\left(r_k^*\mathrm{csgn}\left(a_k\right)\right)^2\right\} - \rho^2 \qquad (46)\\
&= \frac{1}{L}\left(E_a d_2 \mathrm{e}^{-j2\theta} - \rho^2\right). \qquad (47)
\end{aligned}$$

Substituting (43) and (47) in (39), and noting that $\bar{\rho}^2 = E_a d_1^2 \cos^2\theta$, the final result (24) is obtained. The proof of (25) follows analogous steps.

REFERENCES

[1] C. N. Georghiades, "Blind carrier phase acquisition for QAM constellations", *IEEE Trans. Commun.*, vol. 45, no. 11, pp. 1477–1486, Nov. 1997.

[2] E. Serpedin, P. Ciblat, G. B. Giannakis, and P. Loubaton, "Performance analysis of blind carrier phase estimators for general QAM constellations", *IEEE Trans. Signal Processing*, vol. 49, no. 8, pp. 1816–1823, August 2001.

[3] G. Panci, S. Colonnese, S. Rinauro, and G. Scarano, "Gain-control-free near-efficient phase acquisition for QAM constellations", *IEEE Trans. Signal Processing*, vol. 56, no. 7, pp. 2849–2864, 2008.

[4] M. Moeneclaey and G. de Jonghe, "ML-oriented NDA carrier synchronization for general rotationally symmetric signal constellations", *IEEE Trans. Commun.*, vol. 42, no. 8, pp. 2531–2533, August 1984.

[5] A. J. Viterbi and A. M. Viterbi, "Nonlinear estimation of PSK-modulated carrier phase with application to burst digital transmission", *IEEE Trans. Inform. Theory*, vol. IT-29, pp. 543–551, July 1983.

[6] K. V. Cartwright, "Blind phase recovery in cross QAM communication systems with eighth-order statistics", *IEEE Signal Processing Lett.*, vol. 8, no. 12, pp. 304–306, December 2001.

[7] P. Campisi, G. Panci, S. Colonnese, and G. Scarano, "Blind phase recovery for QAM communication systems", *IEEE Trans. Signal Processing*, vol. 53, no. 4, pp. 1348–1358, 2005.

[8] Y. Wang and E. Serpedin, "A class of blind phase recovery techniques for higher order QAM modulations: Estimators and bounds", *IEEE Signal Processing Lett.*, vol. 9, no. 10, pp. 301–304, October 2002.

[9] Y. Wang, E. Serpedin, and P. Ciblat, "Optimal blind nonlinear least-squares carrier phase and frequency offset estimation for general QAM modulations", *IEEE Trans. Wireless Commun.*, vol. 2, no. 5, pp. 1040–1054, 2003.

[10] N.C. Beaulieu and Y. Chen, "Closed-form expressions for the exact symbol error probability of 32-cross-QAM in AWGN and in slow Nakagami fading", *IEEE Commun. Lett.*, vol. 11, no. 4, pp. 310–312, April 2007.

[11] Jian Li, Xian-Da Zhang, and N.C. Beaulieu, "Precise calculation of the SEP of 128- and 512-cross-QAM in AWGN", *IEEE Commun. Lett.*, vol. 12, no. 1, pp. 1–3, January 2008.

[12] Y. H. Hu, "CORDIC-based VLSI architectures for digital signal processing", *IEEE Signal Processing Mag.*, vol. 9, no. 3, pp. 16–35, July 1992.

[13] K. Maharatna, A. Troya, S. Banerjee, and E. Grass, "Virtually scaling-free adaptive CORDIC rotator", *IEE Proc. Comput. Digital Tech.*, vol. 151, no. 6, pp. 448–456, November 2004.

TABLE I
ASYMPTOTIC VARIANCE RATIOS (26) FOR SEVERAL QAM CONSTELLATION SIZES

M	16	32	64	128	256	512	1024	2048
α_1	2.01	0.06	0.93	0.06	0.81	0.05	0.79	0.05
α_2	2.01	0.09	0.98	0.08	0.86	0.08	0.84	0.08

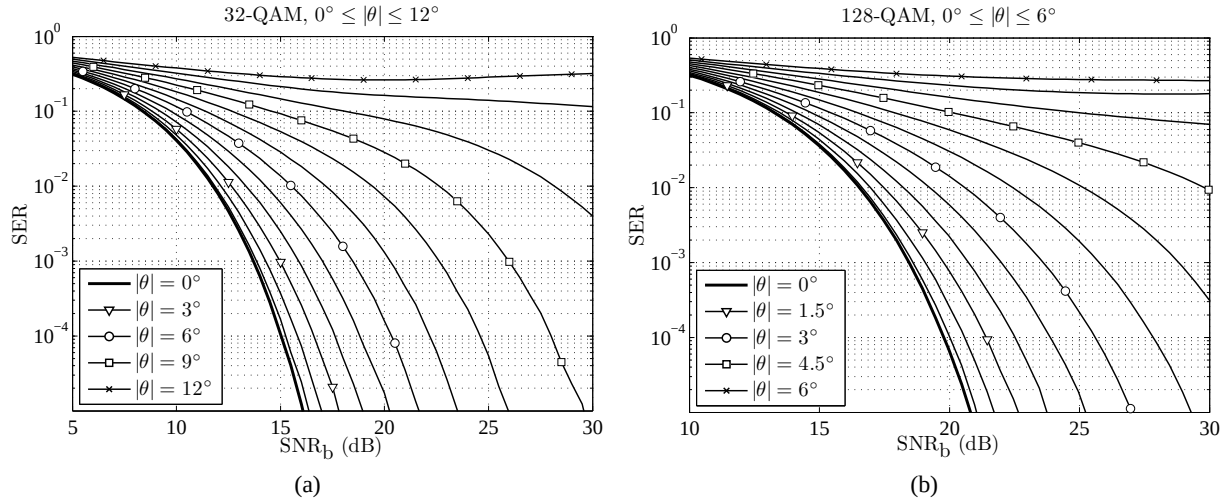

Fig. 1. SER versus $\mathrm{SNR_b}$ for (a) 32-QAM and (b) 128-QAM with different static phase offsets

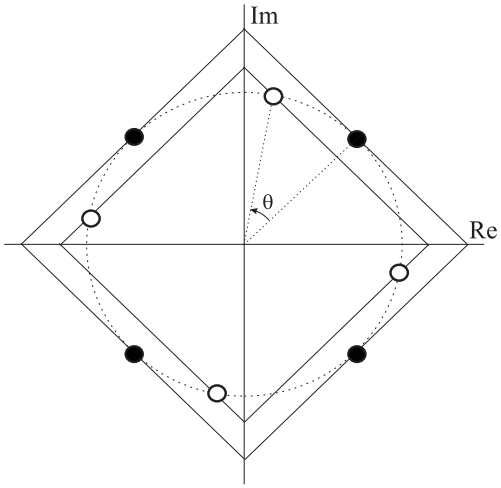

Fig. 2. Effect of a rotation on the average ℓ_1-norm of a QPSK constellation. Solid dots: original constellation; empty dots: rotated constellation

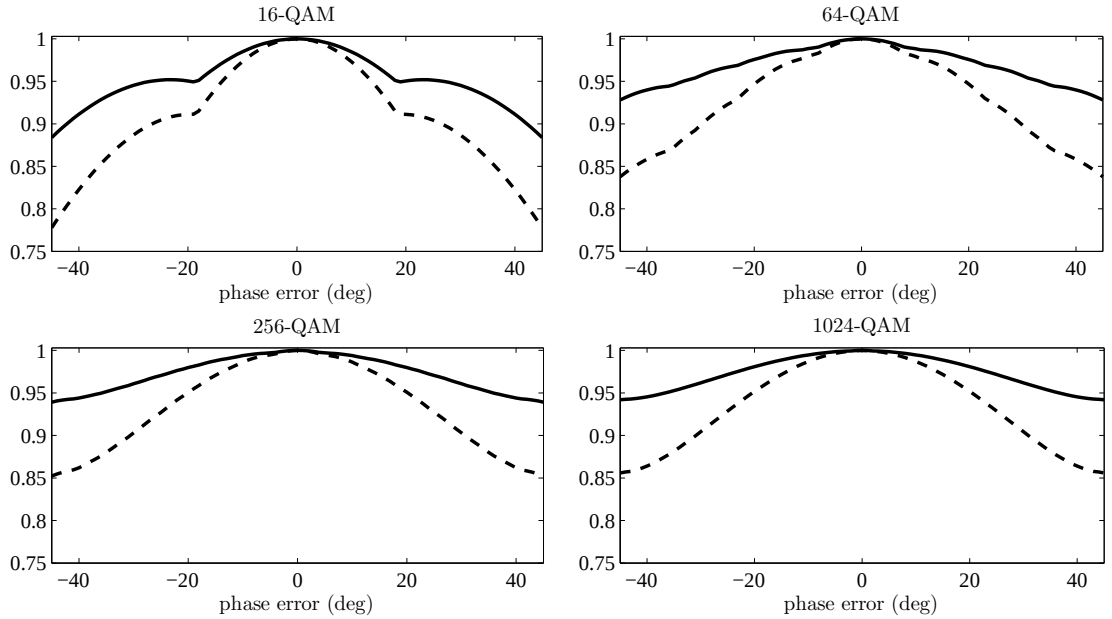

Fig. 3. Square QAM constellations: behavior of the ℓ_1-norm based costs $J_q(\hat{\theta})$ (normalized to their maximum value). $q = 1$ (solid) and 2 (dashed)

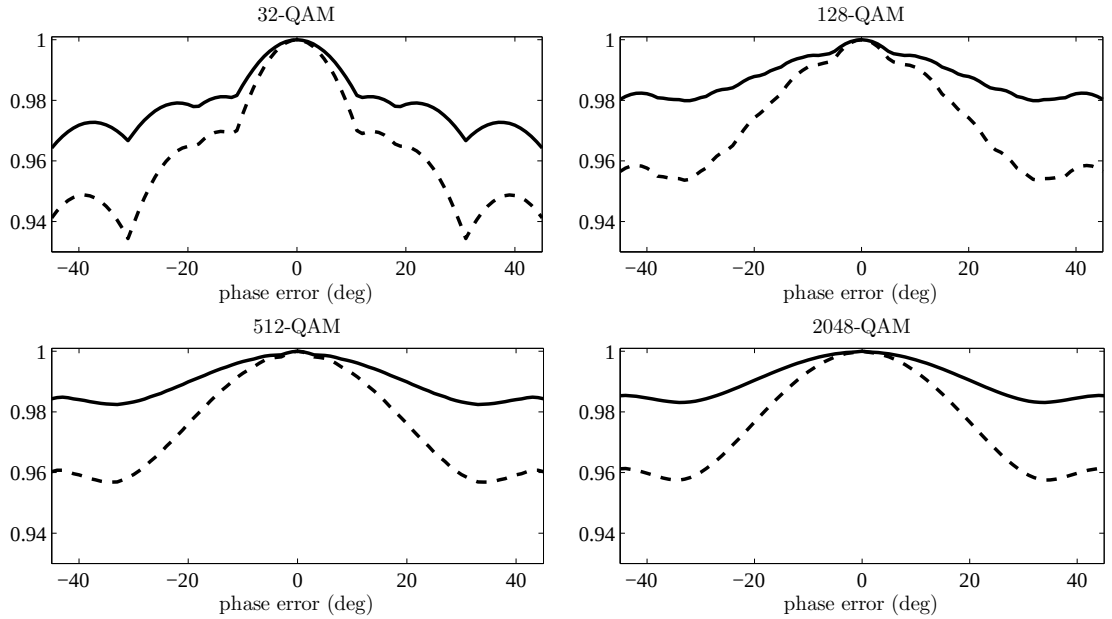

Fig. 4. Cross QAM constellations: behavior of the ℓ_1-norm based costs $J_q(\hat{\theta})$ (normalized to their maximum value). $q = 1$ (solid) and 2 (dashed)

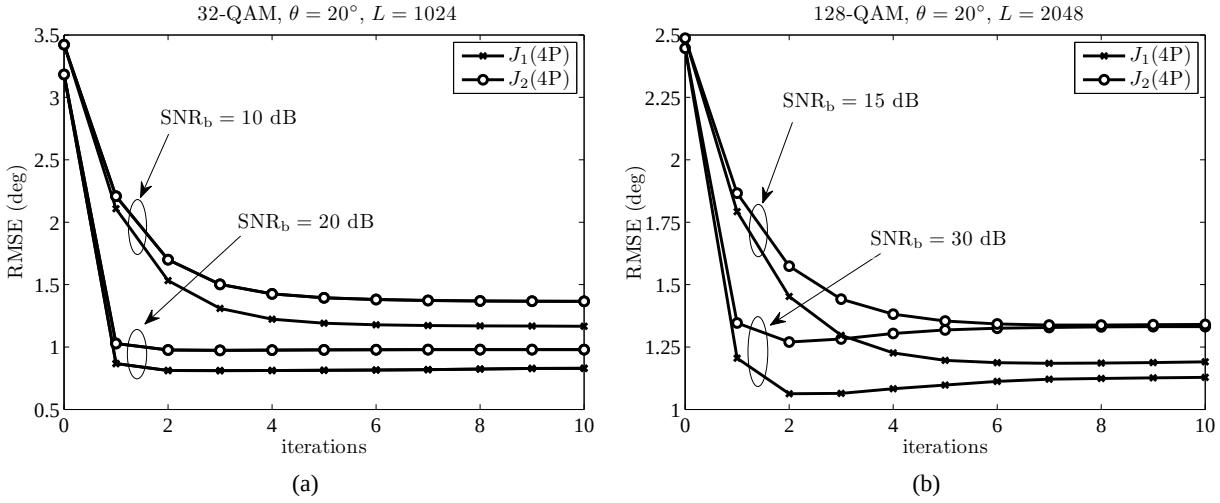

Fig. 5. RMSE vs. number of iterations for (a) 32-QAM and (b) 128-QAM. The method in parentheses after J_1 and J_2 indicates which estimator was used as initializer for the respective iteration

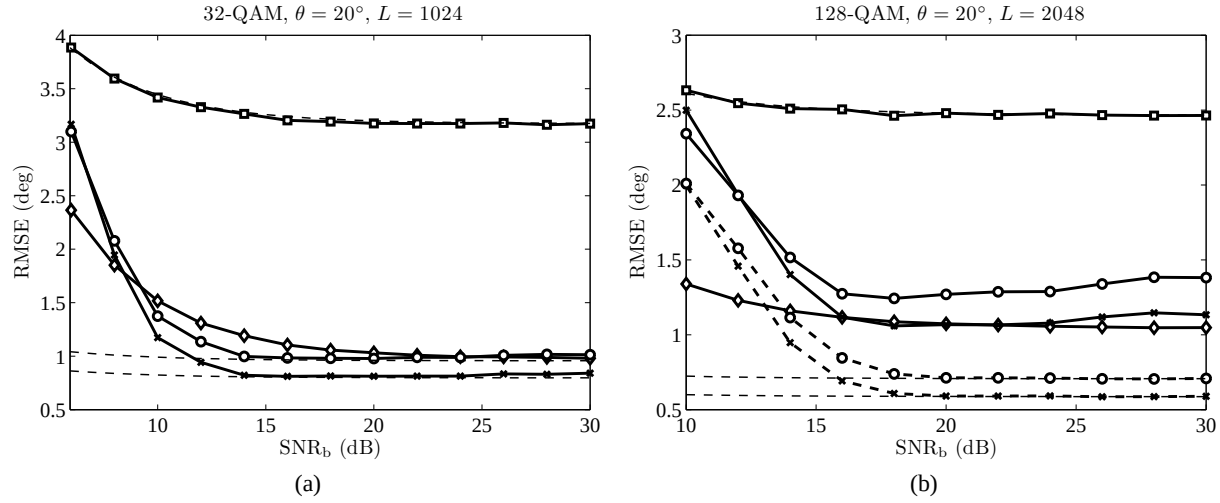

Fig. 6. RMSE vs. SNR per bit for (a) 32-QAM and (b) 128-QAM. Markers are as follows: 4P (□), C8 (◇), J_1 (×), J_2 (○). J_1 and J_2 are initialized with 4P (solid line) and C8 (dashed line). Dashed thin lines indicate the theoretical RMSE values for estimators $\hat{\theta}_{4\mathrm{P}}$, $\hat{\theta}^{(1)}$ and $\hat{\theta}^{(2)}$

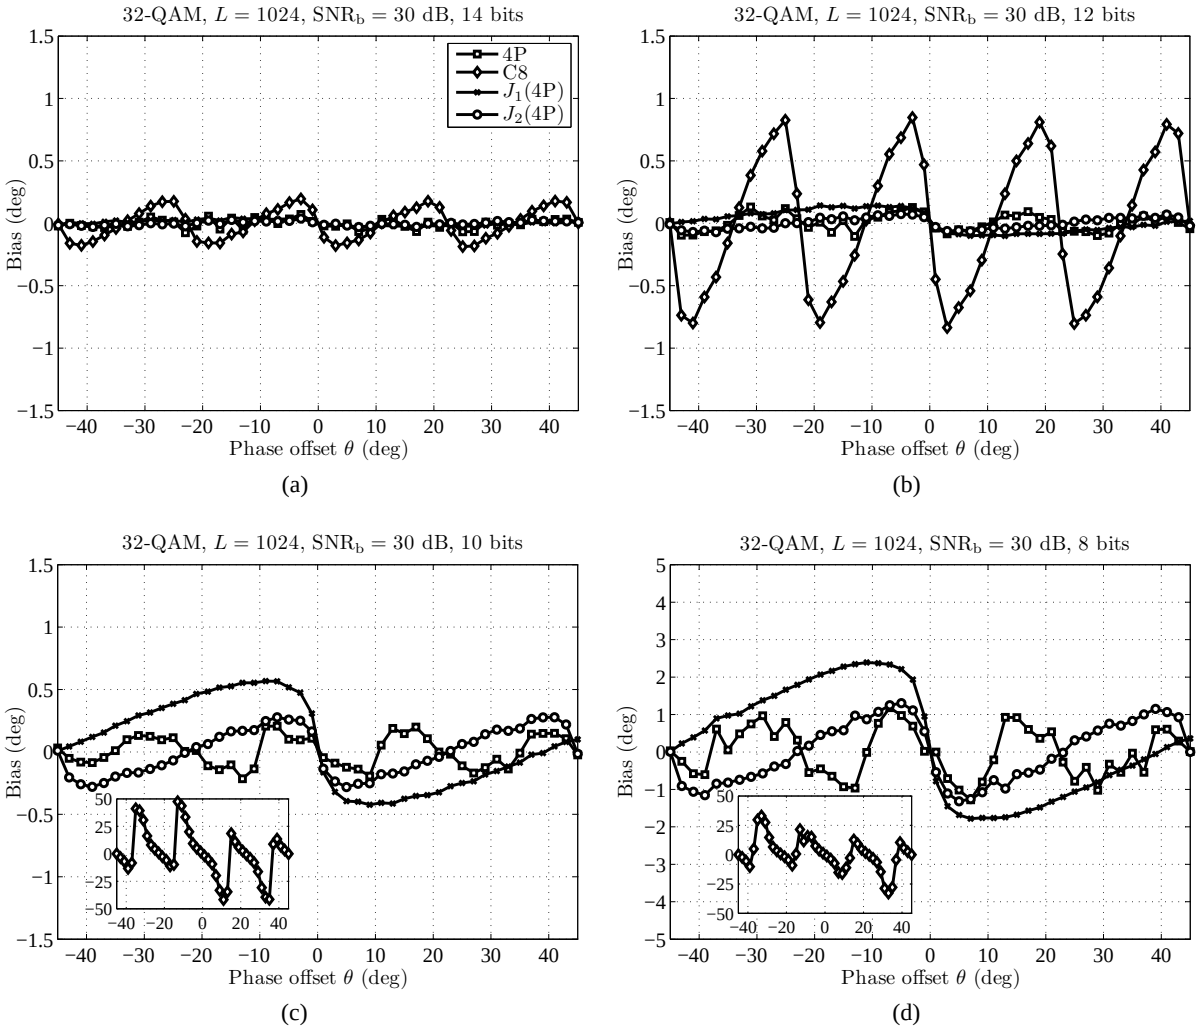

Fig. 7. Bias versus θ. 32-QAM, $L = 1024$, $\mathrm{SNR_b} = 30$ dB. (a) $B = 14$, (b) $B = 12$, (c) $B = 10$, and (d) $B = 8$ bits. Markers are common to all figures as indicated in (a)

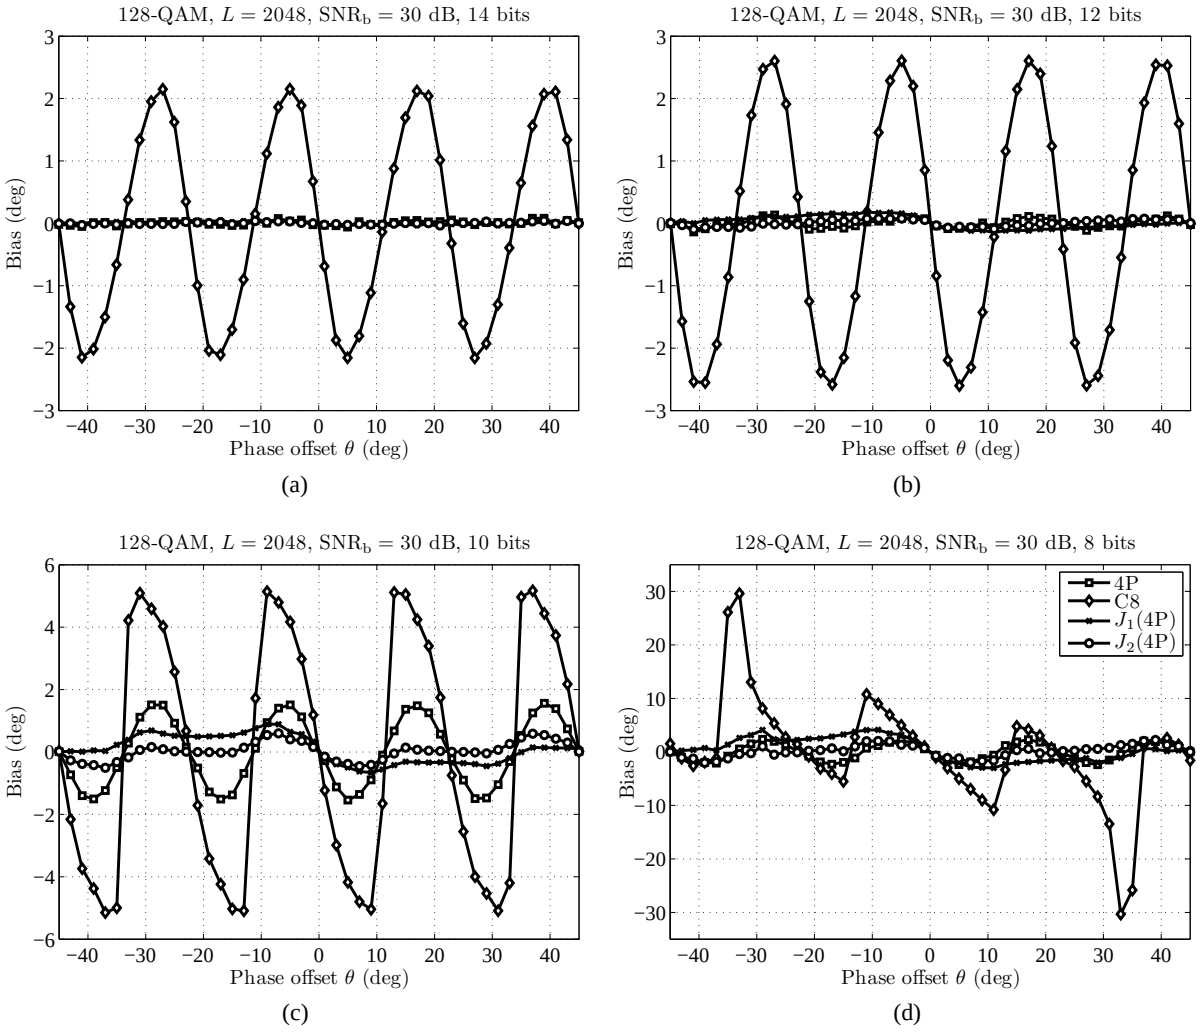

Fig. 8. Bias versus θ. 128-QAM, $L = 2048$, $\mathrm{SNR_b} = 30$ dB. (a) $B = 14$, (b) $B = 12$, (c) $B = 10$, and (d) $B = 8$ bits. Markers are common to all figures as indicated in (d)

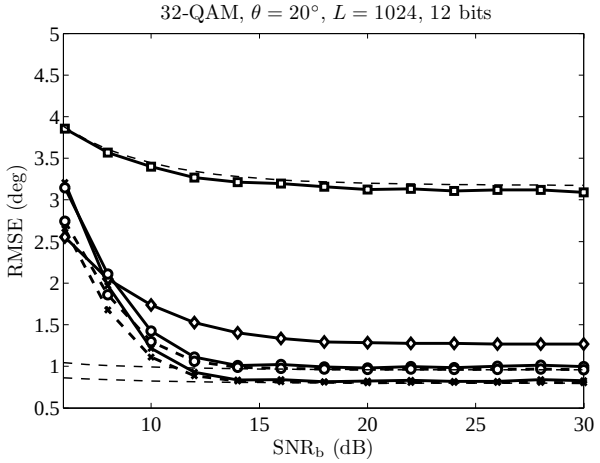

(a)

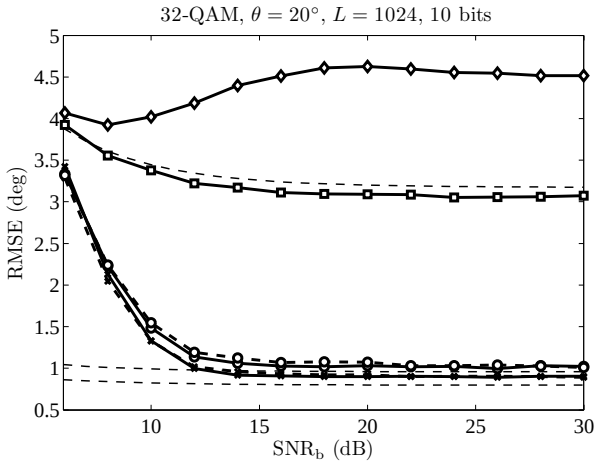

(b)

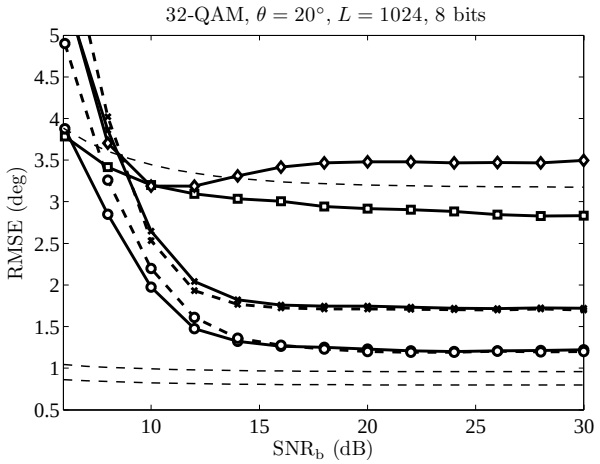

(c)

Fig. 9. RMSE versus $\mathrm{SNR_b}$. 32-QAM, $L = 1024$, $\mathrm{SNR_b} = 30$ dB. (a) $B = 12$, (b) $B = 10$, and (c) $B = 8$ bits. Markers are as follows: 4P (□), C8 (◇), J_1 (×), J_2 (○). J_1 and J_2 are initialized with 4P (solid line) and C8 (dashed line)

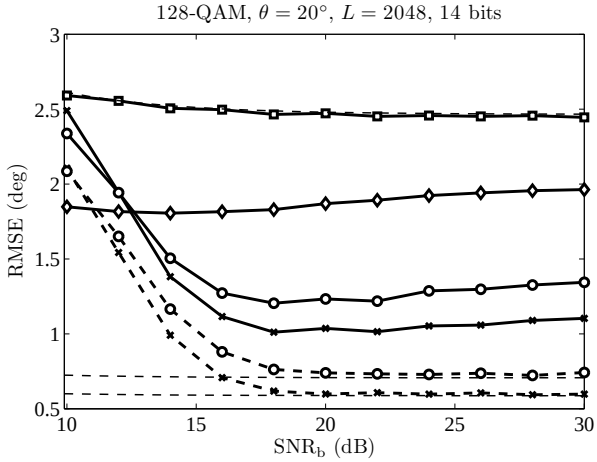

(a)

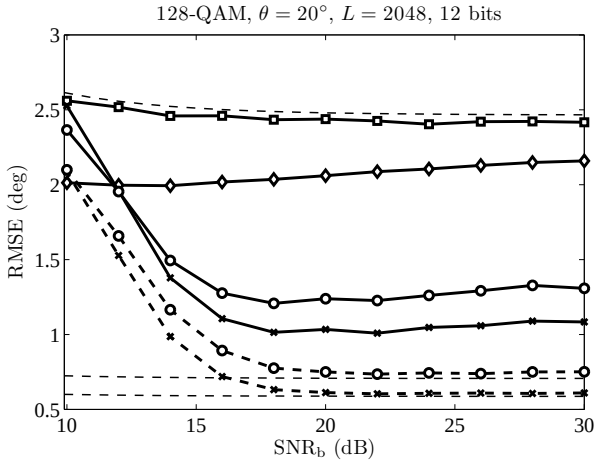

(b)

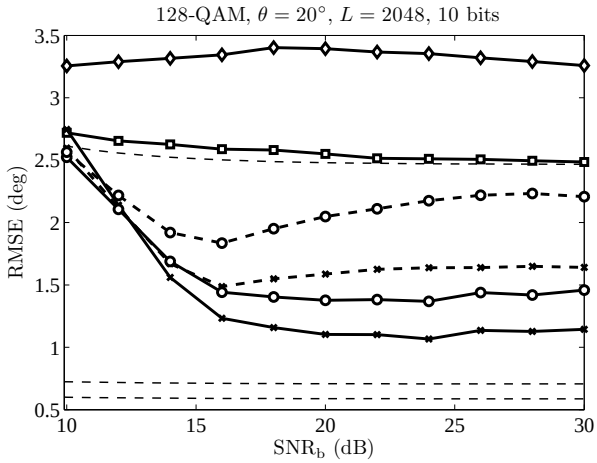

(c)

Fig. 10. RMSE versus $\mathrm{SNR_b}$. 128-QAM, $L = 2048$, $\mathrm{SNR_b} = 30$ dB. (a) $B = 14$, (b) $B = 12$, and (c) $B = 10$ bits. Markers are as follows: 4P (□), C8 (◇), J_1 (×), J_2 (○). J_1 and J_2 are initialized with 4P (solid line) and C8 (dashed line)