On Block Noncoherent Communication with Low-Precision Phase Quantization at the Receiver

Jaspreet Singh and Upamanyu Madhow *

Abstract-We consider communication over the block noncoherent AWGN channel with low-precision Analog-to-Digital Converters (ADCs) at the receiver. For standard uniform Phase Shift Keying (PSK) modulation, we investigate the performance of a receiver architecture that quantizes only the phase of the received signal; this has the advantage of being implementable without automatic gain control, using multiple 1-bit ADCs preceded by analog multipliers. We study the structure of the transition density of the resulting channel model. Several results, based on the symmetry inherent in the channel, are provided to characterize this transition density. A low complexity procedure for computing the channel capacity is obtained using these results. Numerical capacity computations for QPSK show that 8-bin phase quantization of the received signal recovers more than 80-85% of the capacity attained with unquantized observations, while 12-bin phase quantization recovers above 90-95% of the unquantized capacity. Dithering the constellation is shown to improve the performance in the face of drastic quantization.

I. INTRODUCTION

As communication systems scale up in speed and bandwidth, the cost and power consumption of high-precision Analog-to-Digital Conversion (ADC) becomes the limiting factor in modern receiver architectures based on Digital Signal Processing (DSP) [1]. One possible approach for the design of such DSP-centric architectures is to reduce the precision of the ADC. In our prior work [2], [3], we analyzed the impact of low-precision quantization on the capacity of the ideal real baseband discrete-time Additive White Gaussian Noise (AWGN) channel. In this paper, we consider a block noncoherent complex baseband AWGN channel that models the effect of carrier asynchronism. If the receiver's local oscillator is not synchronized with that of the transmitter, the phase after downconversion is a priori unknown, but, for practical values of carrier offset, well approximated as constant over a block of symbols.

The classical approach to noncoherent communication is to approximate the phase as constant over two symbols, and to apply differential modulation and demodulation. Divsalar and Simon [4] were the first to point out the gains that may be achieved by performing multiple symbol differential demodulation over a block of L > 2 symbols. More recent work [5], [6], [7] has shown that block demodulation, even for large values of L, can be implemented efficiently, and exhibits excellent performance for both coded and uncoded systems.

In this work, we study the effect of low-precision receiver quantization for the block noncoherent AWGN channel, under M-ary Phase Shift Keying (MPSK) modulation. Since PSK encodes the information in the phase of the transmitted symbol, we investigate an architecture in which the receiver simply quantizes the phase of the received signal, disregarding the amplitude information. Such phase quantization can be efficiently implemented using 1-bit ADCs preceded by analog multipliers: the use of 1-bit ADCs is attractive since it results in significant power savings and also eliminates the need for Automatic Gain Control (AGC). We study the structure of the inputoutput relationship of the resulting phase quantized-block noncoherent AWGN channel. Based on the symmetry inherent in the channel model, we derive several results characterizing the output probability distribution over a block of symbols, both conditioned on the input, and without conditioning. These results are used to provide a low-complexity procedure for computing the capacity of the channel (brute force computation has complexity exponential in block length L). As in prior work on the unquantized block noncoherent channel, our capacity computations assume that the channel phase is independent from block to block, (this yields a pessimistic estimate of performance, since the phase correlation across blocks can, in principle, be exploited to improve performance). Numerical results are provided for Quaternary Phase Shift Keying (QPSK) with 8-bin and 12-bin phase quantization at the receiver, and compared with the unquantized capacity obtained earlier in [8]. We also provide results that indicate that dithering the constellation improves performance in the face of drastic quantization.

Notation: Throughout the paper, we denote random variables by capital letters, and the specific value they take using small letters. Bold faced notation is used to denote vectors of random variables. \mathbb{E} is the expectation operator.

II. CHANNEL MODEL AND RECEIVER ARCHITECTURE

The received signal over a block of length L, after quantization is represented as

$$Z_l = \mathsf{Q}(S_l e^{j\Phi} + N_l) \ , \ l = 0, 1, \cdots, L - 1, \tag{1}$$

where,

- $\mathbf{S} := [S_0 \ S_1 \ \cdots \ S_{L-1}]$ is the transmitted vector,
- Φ is an unknown constant with uniform distribution on $[0, 2\pi)$,
- $\mathbf{N} := [N_0 \cdots N_{L-1}]$ is a vector of i.i.d. complex Gaussian noise with variance $\sigma^2 = N_0/2$ in each dimension,
- Q: C → K = {0, 1, · · · , K − 1} denotes a quantization function that maps each point in the complex plane to one of the K quantization indices, and

^{*}This work was supported in part by the National Science Foundation under grants CCF-0729222 and CNS-0832154, and by the Office for Naval Research under grant N00014-06-1-0066. The authors are with the ECE Department, UC Santa Barbara, CA 93106, USA. {jsingh, madhow}@ece.ucsb.edu.



Fig. 1. Receiver architecture for 8-sector quantization.

• $\mathbf{Z} := [Z_0 \ Z_1 \ \cdots \ Z_{L-1}]$ is the vector of quantized received symbols, so that each $Z_l \in \mathcal{K}$.

Each S_l is picked in an i.i.d. manner from a uniform M-PSK constellation denoted by the set of points \mathcal{A} = $\{e^{j\theta_0}, e^{j\theta_1}, \cdots, e^{j\theta_{M-1}}\}$, where $\theta_m = (\theta_{m-1} + \frac{2\pi}{M})^{-1}$, for $m = 1, 2, \cdots, M - 1.$

We now introduce the random vector \mathbf{X} = $[X_0 \ X_1 \ \cdots \ X_{L-1}]$, with each X_i picked in an i.i.d. manner from a uniform distribution on the set $\{0, 1, \dots, M-1\}$. Our channel model (1) can now equivalently be written as

$$Z_l = \mathsf{Q}(e^{j\theta_{X_l}}e^{j\Phi} + N_l) , \ l = 0, 1, \cdots, L - 1 , \qquad (2)$$

with every output symbol $Z_l \in \{0, 1, \dots, K-1\}$ as before, and every input symbol $X_l \in \{0, 1, \dots, M-1\}$. The set of all possible input vectors is denoted by \mathcal{X} , while \mathcal{Z} denotes the set of all possible output vectors.

We consider K-bin (or K-sector) phase quantization: our quantizer divides the interval $[0, 2\pi)$ into K equal parts, and the quantization indices go from 0 to K-1 in the counterclockwise direction. Fig. 1(b) depicts the scenario for K=8. Thus, our quantization function is $Q(c) = \lfloor \arg(c) | (\frac{2\pi}{K}) \rfloor$, where $c \in \mathbb{C}$, and $\lfloor p \rfloor$ denotes the greatest integer less than or equal to p. Such phase quantization can be implemented using 1-bit ADCs preceded by analog multipliers which provide linear combinations of the I and Q channel samples. For instance, employing 1-bit ADC on I and Q channels results in uniform 4-sector phase quantization, while uniform 8-sector quantization can be achieved simply by adding two new linear combinations, I+Q and I-Q, corresponding to a $\pi/4$ rotation of I/Q axes (no analog multipliers needed in this case), as shown in Fig. 1(a).

Note: Throughout the paper, we will assume that the PSK constellation size M, and the number of quantization bins K, are such that K = aM for some positive integer a.

III. INPUT-OUTPUT RELATIONSHIP

In this section, we study the relationship between the channel input and output, and present results that govern the structure of the output probability distribution, both conditioned on the input (i.e., $P(\mathbf{Z}|\mathbf{X})$), and without conditioning (i.e., $P(\mathbf{Z})$). These distributions are integral to computing the channel capacity (our focus in this paper), as well as for soft decision decoding (not considered here). While brute force



Fig. 2. QPSK with 8-sector quantization (i.e., M=4, K=8). a) depicts how the unknown channel phase ϕ results in a rotation of the transmitted symbol (square : original constellation , circle : rotated constellation). (b) and (c) depict the circular symmetry induced in the conditional probability $P(z|x, \phi)$ due to the circular symmetry of the complex Gaussian noise. (b) shows that increasing ϕ by $2\pi/K = (\pi/4)$ and z by 1 will keep the conditional probability unchanged, i.e., $P(z = 3|x, \phi) = P(z = 4|x, \phi + 2\pi/K)$. (c) shows that increasing x by 1 and z by 2 = (K/M) will keep the conditional probability unchanged, i.e., $P(z = 2|x, \phi) = P(z = 4|x + 1, \phi)$.

computation (computing $P(\mathbf{z}|\mathbf{x})$ for every $\mathbf{z} \in \mathcal{Z}$ and every $\mathbf{x} \in \mathcal{X}$) of these distributions has exponential complexity in the block length, we show that their inherent structure can be exploited to obtain significant complexity reduction. We illustrate our results throughout with the running example of QPSK with 8-sector quantization, depicted in Fig. 2(a).

Conditioned on the channel phase Φ , $P(\mathbf{Z}|\mathbf{X}, \Phi)$ is a product of individual symbol probabilities $P(Z_l|X_l, \Phi)$. We therefore begin by analyzing the symmetries in the latter.

A. Properties of $P(Z_l|X_l, \Phi)$

We have that $P(z_l|x_l, \phi)$ is the probability that $\arg(e^{j(\theta_{x_l}+\phi)}+N_l)$ belongs to the interval $[\frac{2\pi}{K}z_l \frac{2\pi}{K}(z_l+1))$. In other words, it is the probability that the complex Gaussian noise N_l takes the point $e^{j(\theta_{x_l}+\phi)}$ on the unit circle, to another point whose phase belongs to $\left[\frac{2\pi}{K}z_l \quad \frac{2\pi}{K}(z_l+1)\right)$. Due to the circular symmetry of the complex Gaussian noise, this is the same as the probability that N_l takes the point $e^{j(\theta_{x_l}+\phi+\frac{2\pi}{K}i)}$ on the unit circle, to another point whose phase belongs to $\left[\frac{2\pi}{K}(z_l+i) \quad \frac{2\pi}{K}(z_l+1+i)\right)$, where *i* is an integer. We thus get our first two results.

Property A-1: $P(z_l|x_l, \phi) = P(z_l + i|x_l, \phi + i\frac{2\pi}{K}).$ Property A-2: $P(z_l|x_l, \phi) = P(z_l + ia|x_l + i, \phi).$ Note that $\theta_{x_l+i} = \theta_{x_l} + \frac{2\pi}{M}i = \theta_{x_l} + \frac{2\pi}{K}(ia)$, which gives Property A-2.

Property A-2 simply states that if we jump from one point in the M-PSK constellation to the next, then we must jump $a = \frac{K}{M}$ quantization sectors in order to keep the conditional probability invariant. This is intuitive, since the separation between consecutive points in the input constellation is $2\pi/M$, while each quantization sector covers an angle of $2\pi/K$. For QPSK with K = 8, Figs. 2(b) and 2(c) depict example scenarios for the two properties.

If we put $i = -x_l$ in Property A-2, we get the following special case, which relates the conditioning on a general x_l to the conditioning on 0.

Property A-3: $P(z_l|x_l, \phi) = P(z_l - ax_l|0, \phi).$

To motivate our final property, we consider our example of QPSK with K = 8. While we have 8 distinct quantization sectors, if we look at Fig. 2(a), the orientation of these 8 sectors relative to the 4 constellation points (shown as squares) can be

¹Unless stated otherwise, any arithmetic operations for phase angles are assumed to be performed modulo 2π . For the output symbols Z_l , the arithmetic is modulo K, while for the input symbols X_{l} (introduced immediately after in the text), it is modulo M.

described by dividing the sectors into 2 groups : $\{0, 2, 4, 6\}$, and $\{1, 3, 5, 7\}$. For instance, the positioning of the first sector (z = 0) w.r.t. x = 0 is identical to the positioning of the third sector (z = 2) w.r.t. x = 1 (and similarly z = 4 w.r.t x = 2, and z = 6 w.r.t x = 3). On the other hand, the positioning of the second sector (z = 1) w.r.t. x = 0 is identical to the positioning of the fourth sector (z = 3) w.r.t. x = 1 (and similarly z = 5 w.r.t x = 2, and z = 7 w.r.t x = 3). In terms of the conditional probabilities, this implies, for example, that we will have $P(z_l = 7 | x_l = 3, \phi) = P(z_l = 1 | x_l = 0, \phi)$, and similarly, $P(z_l = 6 | x_l = 3, \phi) = P(z_l = 0 | x_l = 0, \phi)$. In general, we can relate the conditional probability of every odd z_l with that of $z_l = 1$, and similarly of every even z_l with that of $z_l =$ 0, with corresponding rotations of the symbol x_l . For general values of K and M, the number of groups equals $a = \frac{K}{M}$, and we can relate the probability of any z_l with that of $z_l \mod a$.

Property A-4: Let $z_l = q_l a + r_l$, where q_l is the quotient on dividing z_l by a, and r_l is the remainder, i.e, $r_l = z_l \mod a$. Then, $\mathsf{P}(z_l | x_l, \phi) = \mathsf{P}(z_l \mod a | x_l - q_l, \phi)$.

While this result follows directly from Property A-2 by putting $i = -q_l$, it is an important special case, as it enables us to restrict attention to only the first *a* sectors $(Z_l \in \{0, 1, \dots, a-1\})$, rather than having to work with all the *K* sectors. As detailed later, this leads to significant complexity reduction in capacity computation.

We now use these properties to present results for $P(\mathbf{Z}|\mathbf{X})$.

B. Properties of $P(\mathbf{Z}|\mathbf{X})$

Property B-1: Let 1 denote the row vector with all entries as 1. Then $P(\mathbf{z}|\mathbf{x}) = P(\mathbf{z} + i\mathbf{1}|\mathbf{x})$.

Proof: For a fixed x, increasing each z_i by the same number i leaves the conditional probability unchanged, because the phase Φ in the channel model (1) is uniformly distributed in $[0, 2\pi)$. A detailed proof follows. We have

$$P(\mathbf{z}|\mathbf{x}) = \mathbb{E}_{\Phi} \left(\mathsf{P}(\mathbf{z}|\mathbf{x}, \Phi) \right) = \mathbb{E}_{\Phi} \left(\prod_{l=0}^{L-1} \mathsf{P}(z_l|x_l, \Phi) \right)$$
$$= \mathbb{E}_{\Phi} \left(\prod_{l=0}^{L-1} \mathsf{P}(z_l+i|x_l, \Phi+i\frac{2\pi}{K}) \right)$$
$$= \mathbb{E}_{\hat{\Phi}} \left(\prod_{l=0}^{L-1} \mathsf{P}(z_l+i|x_l, \hat{\Phi}) \right)$$
$$= \mathbb{E}_{\hat{\Phi}} \left(\mathsf{P}(\mathbf{z}+i\mathbf{1}|\mathbf{x}, \hat{\Phi}) \right) = \mathsf{P}(\mathbf{z}+i\mathbf{1}|\mathbf{x}).$$

The second equality follows by the fact that the components of **Z** are independent conditioned on **X** and Φ . Property *A*-1 gives the third equality. A change of variables, $\hat{\Phi} = \Phi + i \frac{2\pi}{K}$ gives the fourth equality (since Φ is uniformly distributed on $[0, 2\pi)$, so is $\hat{\Phi}$), thereby completing the proof.

Remark 1: For the rest of the paper, we refer to the operation $\mathbf{z} \rightarrow \mathbf{z} + i\mathbf{1}$ as *constant addition*.

Our next result concerns the observation that the conditional probability remains invariant under an *identical* permutation of the components of the vectors z and x.

Property B-2: Let Π denote a permutation operation, and $\Pi \mathbf{x}$ ($\Pi \mathbf{z}$) the vector obtained on permuting \mathbf{x} (\mathbf{z}) under this operation. Then, $\mathsf{P}(\mathbf{z}|\mathbf{x}) = \mathsf{P}(\Pi \mathbf{z}|\Pi \mathbf{x})$.

Proof: As in the proof of Property 1, the idea is to condition on Φ and work with the symbol probabilities $\mathsf{P}(z_l|x_l, \Phi)$. Consider $\mathsf{P}(\mathbf{z}|\mathbf{x}, \Phi) = \prod_{l=0}^{L-1} \mathsf{P}(z_l|x_l, \Phi)$, and $\mathsf{P}(\Pi \mathbf{z}|\Pi \mathbf{x}, \Phi) = \prod_{l=0}^{L-1} \mathsf{P}((\Pi \mathbf{z})_l | (\Pi \mathbf{x})_l, \Phi)$. Since multiplication is a commutative operation, we have $\mathsf{P}(\mathbf{z}|\mathbf{x}, \Phi) = \mathsf{P}(\Pi \mathbf{z}|\Pi \mathbf{x}, \Phi)$. Taking expectation w.r.t. Φ completes the proof.

The next two results extend properties A-3 and A-4.

Property B-3: Define the input vector $\mathbf{x}_0 = [0 \cdots 0]$. Then, $\mathsf{P}(\mathbf{z}|\mathbf{x}) = \mathsf{P}(\mathbf{z} - a\mathbf{x}|\mathbf{x}_0)$, where $a = \frac{K}{M}$, and the subtraction is performed modulo K.

Property B-4: Let $z_l = q_l a + r_l$, where q_l is the quotient on dividing z_l by a, and r_l is the remainder, i.e, $r_l = z_l \mod a$. Define $\mathbf{q} = [q_0, \cdots, q_{L-1}]$, and, $\mathbf{z} \mod a = [z_0 \mod a \cdots z_{L-1} \mod a]$. Then $\mathsf{P}(\mathbf{z}|\mathbf{x}) = \mathsf{P}(\mathbf{z} \mod a \mid \mathbf{x} - \mathbf{q})$.

Proofs: The properties follow from A-3 and A-4 respectively, by first noting that the vector probability $P(\mathbf{z}|\mathbf{x}, \Phi)$ is the product of the scalar ones, and then integrating over Φ .

C. Properties of $P(\mathbf{Z})$

We now consider the unconditional distribution $P(\mathbf{z})$. The first result states that $P(\mathbf{z})$ is invariant under constant addition. *Property C-1:* $P(\mathbf{z}) = P(\mathbf{z} + i\mathbf{1})$.

Proof: Using Property B-1, this follows directly by taking expectation over \mathbf{X} on both sides.

On similar lines, we now extend Property B-2 to show that P(z) is invariant under any permutation of z.

Property C-2: $P(\mathbf{z}) = P(\Pi \mathbf{z})$.

Proof: We have $P(\mathbf{z}) = \frac{1}{M^L} \sum_{\mathbf{x} \in \mathcal{X}} P(\mathbf{z}|\mathbf{x})$. Using Property *B*-2, we get $P(\mathbf{z}) = \frac{1}{M^L} \sum_{\mathbf{x} \in \mathcal{X}} P(\Pi \mathbf{z}|\Pi \mathbf{x})$. Since Π is just a permutation operation, every unique choice of $\mathbf{x} \in \mathcal{X}$ results in a unique $\Pi \mathbf{x} \in \mathcal{X}$. Hence, we can rewrite the last equation as $P(\mathbf{z}) = \frac{1}{M^L} \sum_{\mathbf{x} \in \mathcal{X}} P(\Pi \mathbf{z}|\mathbf{x}) = P(\Pi \mathbf{z})$.

Our final result extends Property B-4.

Property C-3: Let $a = \frac{K}{M}$. Then $P(\mathbf{z}) = P(\mathbf{z} \mod a)$.

Proof: Using the same notation as in Property *B*-4, we have $P(\mathbf{z}|\mathbf{x}) = P(\mathbf{z} \mod a \mid \mathbf{x} - \mathbf{q})$. Noting that the transformation $\mathbf{x} \rightarrow \mathbf{x} - \mathbf{q}$ is a one-to-one mapping, the proof follows on the same lines as the proof of Property *C*-2.

Example: For QPSK with K = 8 and L = 4, $P(z = [5 \ 7 \ 2 \ 4]) = P(z = [1 \ 1 \ 0 \ 0]).$

We now apply these results for low complexity capacity computation.

IV. EFFICIENT CAPACITY COMPUTATION

We wish to compute the mutual information

$$I(\mathbf{X}; \mathbf{Z}) = H(\mathbf{Z}) - H(\mathbf{Z}|\mathbf{X}).$$

We first discuss computation of the conditional entropy.

A. Computation of the conditional entropy $H(\mathbf{Z}|\mathbf{X})$

We have $H(\mathbf{Z}|\mathbf{X}) = \sum_{\mathcal{X}} H(\mathbf{Z}|\mathbf{x}) \mathsf{P}(\mathbf{x})$, where $H(\mathbf{Z}|\mathbf{x}) = -\sum_{\mathcal{Z}} \mathsf{P}(\mathbf{z}|\mathbf{x}) \log \mathsf{P}(\mathbf{z}|\mathbf{x})$ is the entropy of the output when the input vector \mathbf{X} takes on the specific value \mathbf{x} . Our main result in this section is that $H(\mathbf{Z}|\mathbf{x})$ is constant $\forall \mathbf{x}$.

Property D-1: $H(\mathbf{Z}|\mathbf{x})$ is a constant.

Proof: We show that for any input vector \mathbf{x} , $H(\mathbf{Z}|\mathbf{x}) = H(\mathbf{Z}|\mathbf{x}_0)$, where $\mathbf{x}_0 = [0 \cdots 0]$ as defined before. We have

$$H(\mathbf{Z}|\mathbf{x}) = - \sum_{\mathcal{Z}} \mathsf{P}(\mathbf{z}|\mathbf{x}) \log \mathsf{P}(\mathbf{z}|\mathbf{x})$$
$$= - \sum_{\mathcal{Z}} \mathsf{P}(\mathbf{z} - a\mathbf{x}|\mathbf{x}_0) \log \mathsf{P}(\mathbf{z} - a\mathbf{x}|\mathbf{x}_0) , (3)$$

where the second equality follows from Property *B*-3. Now, since $\mathbf{z} \rightarrow \mathbf{z} - a\mathbf{x}$ is just a subtraction operation, it is easy to see that every unique choice of $\mathbf{z} \in \mathcal{Z}$ results in a unique choice of $\mathbf{z} - a\mathbf{x} \in \mathcal{Z}$. Hence, we can rewrite (3) as

$$H(\mathbf{Z}|\mathbf{x}) = -\sum_{\mathcal{Z}} \mathsf{P}(\mathbf{z}|\mathbf{x}_0) \log \mathsf{P}(\mathbf{z}|\mathbf{x}_0) = H(\mathbf{Z}|\mathbf{x}_0)$$
(4)

Thus, $H(\mathbf{Z}|\mathbf{X}) = H(\mathbf{Z}|\mathbf{x}_0)$, but brute force computation of $H(\mathbf{Z}|\mathbf{x}_0)$ still has exponential complexity, $\mathsf{P}(\mathbf{Z}|\mathbf{x}_0)$ must be computed for each of the K^L possible output vectors \mathbf{Z} . However, we show that it suffices to compute $\mathsf{P}(\mathbf{Z}|\mathbf{x}_0)$ for a much smaller set of \mathbf{Z} vectors.

Using Property B-2, we have $P(\mathbf{z}|\mathbf{x}_0) = P(\Pi \mathbf{z}|\Pi \mathbf{x}_0)$. Since $\mathbf{x}_0 = [0..0]$, any permutation of \mathbf{x}_0 gives back \mathbf{x}_0 . Hence, $P(\mathbf{z}|\mathbf{x}_0) = P(\Pi \mathbf{z}|\mathbf{x}_0)$. Combined with Property B-1, we thus get that it suffices to compute $P(\mathbf{z}|\mathbf{x}_0)$ for a set of vectors $S_{\mathbf{Z}}$ in which no vector can be obtained from another by performing the operations of constant addition and permutation. For K = 8 and $L = \{3, 4, 5, 6, 7\}$, the cardinality of the entire set of \mathbf{Z} vectors, K^L , evaluates to $\{512, 4096, 32768, 2.6 \times 10^5, 2.1 \times 10^6\}$, while the cardinality of $S_{\mathbf{Z}}$ is $\{15, 43, 99, 217, 429\}$, illustrating the large reduction in complexity. For simplicity of exposition, we do not delve into the exact details of how we can obtain the set $S_{\mathbf{Z}}$. Fast algorithms to do this, and their associated complexity are currently being investigated. More details are available from the authors upon request, and will be provided in future publications as well.

Once we have the set $S_{\mathbf{Z}}$, we can numerically compute the probability $\mathsf{P}(\mathbf{z}|\mathbf{x}_0)$ for every vector in $S_{\mathbf{Z}}$. The entropy $H(\mathbf{Z}|\mathbf{x}_0)$ can then be obtained as follows. For $\mathbf{z} \in S_{\mathbf{Z}}$, let $n(\mathbf{z})$ denote the number of distinct vectors that can be generated from it by performing the operations of constant addition and permutation. This is straightforward to compute. The conditional entropy then is $H(\mathbf{Z}|\mathbf{x}_0) = -\sum_{S_{\mathbf{z}}} n(\mathbf{z})\mathsf{P}(\mathbf{z}|\mathbf{x}_0) \log \mathsf{P}(\mathbf{z}|\mathbf{x}_0)$.

B. Computation of the output entropy $H(\mathbf{Z})$

The output entropy is $H(\mathbf{Z}) = -\sum_{\mathcal{Z}} P(\mathbf{z}) \log P(\mathbf{z})$. A brute force computation requires us to know $P(\mathbf{z}) \forall \mathbf{z} \in \mathcal{Z}$, which clearly has exponential complexity. However, using Properties C-1, C-2 and C-3, we get that it is sufficient to compute $P(\mathbf{z})$ for a set of vectors $\tilde{S}_{\mathbf{Z}}$ in which no vector can be obtained from another one by performing the operations of constant addition and permutation, and also, the vector components $\in \{0, 1, \dots, a-1\}$. This is similar to the situation we encountered earlier in the last subsection, except that the vector components there were allowed to be in $\{0, 1, \dots, K-1\}$. *Example:* For QPSK with 8 sectors (so a = 2), the relevant vectors for block length 2 are $\begin{bmatrix} 0 & 0 \end{bmatrix}$ and $\begin{bmatrix} 0 & 1 \end{bmatrix}$.

Computation of $P(\mathbf{Z})$: For each of the vectors in the set $\hat{S}_{\mathbf{Z}}$ defined above, we now need to obtain $P(\mathbf{z}) =$ $\sum_{\mathbf{x}\in\mathcal{X}}\mathsf{P}(\mathbf{z}|\mathbf{x})\mathsf{P}(\mathbf{x}).$ A brute force approach is to compute $P(\mathbf{z}|\mathbf{x})$ for each x. However, we can exploit the structure in z to reduce the number of vectors \mathbf{x} for which we need $P(\mathbf{z}|\mathbf{x})$. Specifically, we have that each $z_i \in \{0, 1, \dots, a-1\}$. Since there are only a different types of components in z, for block length L > a, some of the components in z will be repeated. For any x, we can then use Property B-2 to rearrange the components at those locations for which the components in z are identical, without changing the conditional probability. For instance, let $z_m = z_n$ for some m, n. Then, $\mathsf{P}(\mathbf{z}|\mathbf{x}) =$ $P(\mathbf{z}|\Pi \mathbf{x})$, where $\Pi \mathbf{x}$ is obtained from \mathbf{x} by rearranging the components at locations m and n. To sum up, we can restrict attention to a set of vectors $S_{\mathbf{X}}$ in which no vector can be obtained from another one by permutations between those locations for which the elements in z are identical. While for large a, the potential reduction in complexity may not be large, for small values of a (which is the paradigm of interest in this work), the savings will be significant. As before, the algorithmic details for obtaining the set $S_{\mathbf{X}}$ will be provided in upcoming publications.

V. NUMERICAL RESULTS

We now present capacity results (obtained using the low-complexity procedure outlined in the last section) for QPSK with 8-sector and 12-sector phase quantization, for different block lengths L. For all our results, we normalized the mutual information $I(\mathbf{X}; \mathbf{Z})$ by L-1 to obtain the per symbol capacity, since in practice the successive blocks can be overlapped by one symbol due to slow phase variation from one block to the next.

8-sector quantization: In Fig. 3, we plot the channel capacity with 8-sector quantization, at different SNR values .(To avoid clutter, we show the results for L = 6 only.) Also shown for reference are the capacity values for the coherent case, and for the block noncoherent case without any quantization. We see that, for SNR > 2-3 dB, our simple 8-sector quantization scheme recovers more than 80-85% of the spectral efficiency obtained with unquantized observations. This is encouraging, given that our work is targeted towards future high bandwidth systems (such as those operating in the 60 GHz mm-wave band), for which a small reduction in spectral efficiency is acceptable. On the other hand, if we measure the power loss for fixed spectral efficiency, we see that at rates of up to about 1.2 bits/channel use, there is a loss of about 1-1.5 dB compared to the unquantized case. However, the loss is more significant as SNR increases: the capacity approaches 2 bits/channel use rather slowly at high SNR. Since the input entropy $H(\mathbf{X})$ is constant, this in turn implies that $H(\mathbf{X}|\mathbf{Z})$ falls off very slowly as SNR $\rightarrow \infty$. A more detailed analysis of the likelihood ratio $P(\mathbf{Z}|\mathbf{X}, \Phi)$ (omitted here due to lack of space) provides insight into this behavior. We find that in addition to the symmetries in $P(\mathbf{Z}|\mathbf{X}, \Phi)$ that we exploited to reduce the complexity of



Fig. 3. Performance comparison for QPSK with block length L = 6: plots depict the capacity of the block noncoherent channel without quantization, and with 8-sector quantization (with and without dithering). Also shown is the capacity for coherent QPSK.

capacity computations, there are certain other symmetries with adverse consequences as well : they make it impossible to distinguish between the effect of the unknown phase offset and the phase modulation on the received signal. More specifically, if we consider the maximum likelihood (ML) estimator argmax $P(\mathbf{z}|\mathbf{x}, \phi)$, we find that for certain outputs \mathbf{z} , irrespective of the SNR (and also the block length), the estimator always returns two distinct equally likely solutions (\mathbf{x}_1, ϕ_1) and

 (\mathbf{x}_2, ϕ_2) . In an information-theoretic sense, this ambiguity indicates a significant conditional entropy $H(\mathbf{X}|\mathbf{Z})$. As SNR $\rightarrow \infty$, the probability of these ambiguous outputs does go to zero, but very slowly, leading to a slow decrease in $H(\mathbf{X}|\mathbf{Z})$ as well.

Possible ways to break the undesirable symmetries could be to use non-uniform phase quantization, or to employ dithering across symbols in a block. Here we investigate the role of the latter. We can dither either at the transmitter by rotating the QPSK constellation points, or at the receiver by using analog pre-multipliers to shift the phase quantization boundaries. We use a simple transmit dither scheme in which we rotate the QPSK constellation by an angle of $\frac{1}{L}(\frac{2\pi}{K})$ from one symbol to the next. Fig. 4(a) shows this scheme for block length L=2 and K=8. The constellation used for the second symbol (shown by the diamond shape) is dithered from the constellation used for the first symbol by an angle of $\pi/8$. With this choice of transmit constellations, we find that the ambiguity in the ML estimator is removed, and hence the performance is expected to improve. The plot in Fig. 3 shows the performance improvement for L=6. 2

While the preceding simple transmit dither scheme has improved the performance for 8-sector quantization, we hasten to add that there is no optimality associated with it. A more detailed investigation of different dithering schemes and their potential gains is therefore an important topic for future research.

12-sector quantization: In Fig. 5, we plot the performance curves for QPSK with 12-sector quantization, for block length L=2,4,6,8. Also shown for reference are the coherent and unquantized block noncoherent performance curves. For



Fig. 4. (a) Standard PSK : the same constellation (the one shown) is used for both symbols in the block. (b) Dithered-PSK : the constellations used for the two symbols are not identical, but the second one is a dithered version of the first one.



Fig. 5. Performance comparison for QPSK : plots depict the capacity of the coherent channel, unquantized block noncoherent channel (different block lengths), and the 12-sector quantized block noncoherent channel (different block lengths).

identical block lengths, the loss in capacity (at a fixed SNR > 2-3 dB) compared to the unquantized case is less than 5-10 %, while the loss in power efficiency (for fixed capacity) varies between 0.5-2 dB, and dithering is not required.

VI. CONCLUSIONS

We have investigated the capacity limits imposed by the use of low-precision phase quantization at a carrier-asynchronous receiver. The symmetries in input-output relationship of the resulting channel have been exploited to reduce the complexity of capacity computation. Important topics for future research include a more detailed investigation of different dithering schemes (motivated by the performance improvement obtained using the simple scheme considered here), as well as development of practical capacity-approaching coded modulation strategies. An important practical issue is determining whether timing synchronization (which is assumed in the model here) can also be attained using phase-quantized samples, or whether some form of additional information (perhaps using analog techniques prior to the ADC) is required.

REFERENCES

- R. Walden, Analog-to-Digital Converter survey and analysis, IEEE J. Select. Areas Comm., 17(4):539–550, Apr. 1999.
- [2] J. Singh, O. Dabeer and U. Madhow, Capacity of the discrete-time AWGN channel under output quantization, In Proc. ISIT 2008.
- [3] J. Singh, O. Dabeer and U. Madhow, On the limits of communication with lowprecision analog-to-digital conversion at the receiver, Submitted for publication.
- [4] D. Divsalar and M. K. Simon, Multiple-symbol differential detection of MPSK, IEEE Trans. on Comm., vol. 38, pp. 300-308, Mar. 1990.
 [5] D. Weise and U. Machen, Scatterful, "Every character provide the symposium terms of terms of the symposium terms of t
- [5] D. Warrier and U. Madhow, Spectrally efficient noncoherent communication, IEEE Trans. Info. Theory, vol. 48, pp. 651-668, Mar. 2002.
- [6] W. Sweldens, Fast block noncoherent decoding, IEEE Comm. Letters, Apr. 2001.
- [7] R-R Chen, R. Koetter, U. Madhow and D. Agrawal, Joint noncoherent demodulation and decoding for the block fading channel: a practical framework for approaching Shannon capacity, IEEE Trans. on Comm., Oct. 2003.
- [8] M. Peleg and S. Shamai, On the capacity of the blockwise incoherent MPSK channel, IEEE Trans. Comm., vol. 46, pp. 603-609, May 1998.

²Since the low-complexity procedure outlined in Section IV does not work once we dither, we used Monte Carlo simulations to compute the capacity with dithering.