Design of Discrete Constellations for Peak-Power-Limited Complex Gaussian Channels

Wasim Huleihel
MIT
wasimh@mit.edu

Ziv Goldfeld
MIT
zivg@mit.edu

Tobias Koch
Universidad Carlos III de Madrid
koch@tsc.uc3m.es

Mokshay Madiman
University of Delaware
madiman@udel.edu

Muriel Médard
MIT
medard@mit.edu

Abstract—The capacity-achieving input distribution of the complex Gaussian channel with both average- and peak-power constraint is known to have a discrete amplitude and a continuous, uniformly-distributed, phase. Practical considerations, however, render the continuous phase inapplicable. This work studies the backoff from capacity induced by discretizing the phase of the input signal. A sufficient condition on the total number of quantization points that guarantees an arbitrarily small backoff is derived, and constellations that attain this guaranteed performance are proposed.

I. Introduction

The most common channel model in the information theory literature is, arguably, the additive white Gaussian noise (AWGN) channel. Due to practical considerations, the input is typically constrained in some manner. For an average power constraint P_{avg}, it is well-known that the channel capacity of the discrete-time, complex-valued AWGN channel with noise variance σ^2 is $\log(1 + P_{\text{avg}}/\sigma^2)$, and the capacity-achieving input distribution is zero-mean Gaussian with variance P_{avg} [1]. Gaussian inputs, however, suffer from several drawbacks which limit their use in practical systems. One main drawback is that they have unbounded and continuous support, hence an infinite number of bits is needed to represent the signal points.

This impracticality is alleviated when considering the complex-valued AWGN channel with both average- and peak-power constraints. In this case, it was shown by Shamai and Bar-David [2] that the capacity-achieving input distribution is discrete in amplitude and continuous in phase. Efficient algorithms for calculating the capacity-achieving input distribution were proposed in [3]. Furthermore, Wu and Verdú [4] studied the information rates achievable over the Gaussian channel when the input takes value in a finite constellation with N signal points. For every fixed SNR, they showed that the difference between the capacity and the achievable rate tends to zero exponentially in N.

In practice, discrete constellations such as phase-shift-keying (PSK) or quadrature amplitude modulation (QAM) are often used instead, despite the fact that they may produce suboptimal performance with a significant backoff from capacity. Discrete signal constellations are also required for coded modulation methods such as multi-level coding (MLC). To this end, [5] proposed an ad-hoc approach to approximate the semi-continuous capacity-achieving input distribution by discretizing its phase. The authors in [5] took the following approach: given a budget $K > 0$ on the total number of desired constellation points, their resulting discrete signal set, referred to as K-ary amplitude and phase-shift keying (K-APSK), consists of $|\mathcal{A}|$ different PSK constellations with radii $a_i \in \mathcal{A}$ (each with probably $P_A(a_i)$) with k_i equiprobable signal points on each (such that $\sum_{i=1}^{|\mathcal{A}|} k_i = K$). Choosing k_i's to be proportional to $P_A(a_i)$, [5] empirically showed that K-APSK significantly outperforms K-QAM constellations for the peak-power-limited complex-valued AWGN channel. However, no theoretical guarantees of performance were provided. To the best of our knowledge, it is yet unknown how to convert theoretical capacity-achieving input distributions for the peak-power-limited AWGN channel into practically applicable modulation schemes.

In this paper, we propose distributions that provably approach capacity, while being discrete both in amplitude and phase. To this end, we analyze the capacity loss incurred by discretizing the phase of the semi-continuous capacity-achieving input distribution. We propose an optimized K-APSK constellation and derive a sufficient condition on the total number of quantization points K that guarantees a given backoff from capacity. We also briefly discuss the case of re-discretizing the amplitude A, or, equivalently, shrinking the set $\mathcal{A}$, which might be essential when the cardinality of $\mathcal{A}$ is large. The main technical tool in our analysis is based on a recent result by Polyanskiy and Wu [6], which states that when smoothed by Gaussian noise, mutual information is Lipschitz continuous with respect to the Wasserstein distance.

II. Preliminaries and Model Formulation

A. Notation

Throughout this paper we use the following notation. Differential entropy and mutual information are denoted by $h(\cdot)$ and $I(\cdot;\cdot)$, respectively. The Euclidean and ℓ_1 norms of $x^n \in \mathbb{R}^n$ are denoted by $\|x\| \triangleq (\sum_{i=1}^n x_i^2)^{1/2}$ and $\|x\|_1 \triangleq \sum_{i=1}^n |x_i|$, respectively. Given two probability measures P and Q on $\mathbb{R}^n$, their p-Wasserstein distance ($p \geq 1$) on the Euclidean space is defined as $W_p(P, Q) \triangleq \inf \left(\mathbb{E} \|X - Y\|^p\right)^{1/p}$, where the

infimum is taken over all couplings of P and Q, i.e., joint distributions P_{XY} of two random variables X and Y, whose marginals satisfy $P_X = P$ and $P_Y = Q$. The complex conjugate, real part, and imaginary part of a complex number $z \in \mathbb{C}$ are denoted by $z^\star$, $\mathrm{Re}(z)$, and $\mathrm{Im}(z)$, respectively. Finally, the unnormalized sinc function is defined as $\mathrm{sinc}(x) \triangleq \sin(x)/x$, for $x \neq 0$, and $\mathrm{sinc}(0) = 1$.

B. The Model

Consider a discrete-time, complex-valued AWGN channel, where the channel output Y_i at time $i \in \mathbb{N}$ is given by

$$Y_i = X_i + \sigma \cdot W_i, \tag{1}$$

where X_i is the time-i channel input, $\{W_i,\ i \in \mathbb{N}\}$ is a sequence of independent and identically distributed (i.i.d.), centered, unit-variance, circularly-symmetric, complex Gaussian random variables, and $\sigma \in \mathbb{R}_+$ denotes the standard deviation of the additive noise. The noise sequence $\{W_i, i \in \mathbb{N}\}$ is independent of the channel input sequence $\{X_i,\ i \in \mathbb{N}\}$.

Since the channel is memoryless, the channel capacity (in nats per channel-use) under both average and peak-power constraints is given by

$$C(P_{\mathrm{avg}}, P_{\mathrm{max}}) = \sup_{\substack{P_X:\ |X| \leq \sqrt{P_{\mathrm{max}}}\ \text{a.s.,} \\ \mathbb{E}|X|^2 \leq P_{\mathrm{avg}}}} I(X;Y) \tag{2}$$

where the supremum is over all input distributions with essential support in $\{x \in \mathbb{C}\ :\ |x| \leq \sqrt{P_{\mathrm{max}}}\}$ that satisfy $\mathbb{E}|X|^2 \leq P_{\mathrm{avg}}$, for some $P_{\mathrm{max}}, P_{\mathrm{avg}} > 0$. It was demonstrated by Shamai and Bar-David [2] that the capacity-achieving input distribution is compactly supported on the complex plane, satisfies spherical symmetry, and has amplitudes supported on a finite subset of $[0, \infty)$ that contains 0. In other words, for every $\sigma > 0$, the capacity-achieving input distribution is discrete in amplitude and uniform in phase, with the number of mass points growing as $\sigma \to 0$. In [3] an algorithm for computing the optimal amplitude distribution was devised.

We next introduce some additional notation. Throughout this papers, the complex-valued transmitted signal is represented by $X = A \cdot e^{j\varphi}$, where the amplitude (or, modulo) A and phase φ take values in $\mathbb{R}_+$ and $[0, 2\pi)$, respectively; $j \triangleq \sqrt{-1}$ denotes the imaginary unit. As described above, for the capacity-achieving input distribution, the random variable A is discrete taking values in a finite set $\mathcal{A} \subset \mathbb{R}$, and φ is uniformly distributed on $[0, 2\pi)$, i.e., $\varphi \sim \mathrm{Unif}[0, 2\pi)$. Furthermore, A and φ are statistically independent.

Remark 1. *While not explicit in our notation, the distribution of A, as well as the cardinality of the set $\mathcal{A}$, depend on the average and maximal signal-to-noise ratio, namely, on* $\mathrm{SNR}_{\mathrm{avg}} \triangleq P_{\mathrm{avg}}/\sigma^2$ *and* $\mathrm{SNR}_{\mathrm{max}} \triangleq P_{\mathrm{max}}/\sigma^2$.

C. Phase Quantization

Our goal is to design discrete input distributions that mimic the performance of the capacity-achieving input distribution. To this end, we discretize the phase φ. Specifically, for any $a \in \mathcal{A}$, we define a *phase quantizer* as a map $Q_a : [0, 2\pi) \to \mathcal{Q}_a$, where $\mathcal{Q}_a$ is a finite set with cardinality $k_a \triangleq |\mathcal{Q}_a|$. To

wit, for any given value of $A = a \in \mathcal{A}$, $Q_a(\varphi)$ maps the phase into $\mathcal{Q}_a$. Given $(k_a)_{a \in \mathcal{A}}$, a natural choice of $Q_a(\cdot)$ is

$$Q_a(\varphi) = \frac{2\pi}{k_a} \cdot (\ell - 1),\ \forall a \in \mathcal{A} \tag{3}$$

where $\varphi \in [0, 2\pi)$ and $\ell \in [1 : k_a]$ is such that

$$\frac{2\pi}{k_a} \cdot (\ell - 1) \leq \varphi < \frac{2\pi}{k_a} \cdot \ell. \tag{4}$$

For convenience, for each $a \in \mathcal{A}$ and $\ell \in [1 : k_a]$, we set

$$\Phi_{a,\ell} \triangleq \left\{ \varphi \in [0, 2\pi) :\ \frac{2\pi}{k_a} \cdot (\ell - 1) \leq \varphi < \frac{2\pi}{k_a} \cdot \ell \right\}. \tag{5}$$

The above choice corresponds to discretizing the "circle" of radius $A = a$ using k_a points (namely, the corresponding roots of unity). Geometrically, we divide the circle of radius $A = a$ in the complex plane into equal-lengthed k_a arcs, each subtending an angle of $2\pi/k_a$ at the origin and centered around the k_a-th root of unity. In the sequel, X_Q denotes the input with discretized phase, i.e., $X_Q = A \cdot e^{jQ_A(\varphi)}$. Accordingly, Y_Q denotes the output of the AWGN channel (1) when X_Q is transmitted, namely, $Y_Q = X_Q + \sigma \cdot W$.

Deviating from optimality, we study the loss incurred by the above pre-processing. More precisely, we focus on the loss in terms of mutual information when the quantized input distribution has at most K constellation points. Indeed, let $\mathbf{k} \triangleq (k_a)_{a \in \mathcal{A}} \in \mathbb{N}^{|\mathcal{A}|}$ be such that $\sum_{a \in \mathcal{A}} k_a \leq K$, and define

$$\mathcal{L}_K(\mathbf{k}) \triangleq |I(X;Y) - I(X_Q;Y_Q)| = |h(Y) - h(Y_Q)|. \tag{6}$$

The definition in (6) calls for an optimization over $\mathbf{k}$. Namely, given a budget $K > 0$ on $\mathbf{k}$, one would like to find the best phase-quantizer through the optimization problem

$$\mathcal{L}_K^\star \triangleq \min_{\mathbf{k}:\ \|\mathbf{k}\|_1 \leq K} \mathcal{L}_K(\mathbf{k}). \tag{7}$$

Unfortunately, obtaining a closed-form expression for (7) seems out of reach, since even $I(X;Y)$ cannot be evaluated in closed form. To circumvent this difficulty we derive and study an upper bound on (7) in Section III.

D. Amplitude and Phase Quantization

The previous subsection considers only phase quantization. However, one may also wish to re-discretize A, or, equivalently, shrink the set $\mathcal{A}$. For example, if $|\mathcal{A}|$ is large, it may be desirable to revert to a smaller set of amplitudes, while keeping the power and average constraints satisfied. To do so, let $\mathfrak{D}$ be the set of all maps $\Delta : \mathcal{A} \to \mathbb{R}$ with $|\mathrm{Image}(\Delta)| \leq |\mathcal{A}|$, and

$$\sum_{a \in \mathcal{A}} P_A(a)\Delta^2(a) \leq P_{\mathrm{avg}}, \tag{8a}$$

$$|\Delta(a)| \leq \sqrt{P_{\mathrm{max}}},\ \forall a \in \mathcal{A}. \tag{8b}$$

The set $\mathfrak{D}$ comprises all possible *amplitude quantizers* Δ. We assume that $0 \in \mathrm{Image}(\Delta)$ for any quantizer in $\mathfrak{D}$. This assumption comes without loss of generality since one can always add 0 as an amplitude without increasing the average or the maximal power of the signal while improving the result

of subsequently stated optimization problem (see (10)). For simplicity of notation, we let $\mathcal{D} \triangleq \mathrm{Image}(\Delta)$.

Given $\Delta \in \mathfrak{D}$, the phase quantizer is defined as in Subsection II-C with A replaced by $\Delta(A)$. In the sequel, $X_{Q,\Delta}$ denotes the resulting amplitude-phase quantized input, i.e., $X_{Q,\Delta} = \Delta(A) \cdot e^{jQ_{\Delta(A)}(\varphi)}$. Accordingly, we let $Y_{Q,\Delta}$ be the output of our AWGN channel when $X_{Q,\Delta}$ is transmitted.

We aim to find the optimal pair of amplitude-phase quantizers, given a constraint on the total number of constellation points. More precisely, for fixed $K \in \mathbb{N}_0$, $\Delta \in \mathfrak{D}$, and $\mathbf{k}_\Delta \triangleq (k_a)_{a \in \mathcal{D}} \in \mathbb{N}^{|\mathcal{D}|}$ with $\|\mathbf{k}_\Delta\|_1 \leq K$, let

$$\bar{\mathcal{L}}_K(\mathbf{k}_\Delta, \Delta) \triangleq |I(X;Y) - I(X_{Q,\Delta};Y_{Q,\Delta})| . \qquad (9)$$

The optimal information loss due to amplitude and phase quantization is then

$$\bar{\mathcal{L}}_K^\star \triangleq \min_{\substack{\Delta \in \mathfrak{D}, \\ \mathbf{k}_\Delta : \|\mathbf{k}_\Delta\|_1 \leq K}} \bar{\mathcal{L}}_K(\mathbf{k}_\Delta, \Delta) . \qquad (10)$$

In this paper, we focus on the phase quantization part, but several results on amplitude and phase quantization are briefly discussed at the end of Section III.

III. Main Results

In this section we present our main results. Proof sketches of some of the results appear in Section IV.

A. Theoretical Bounds

Our first result concerns the phase-quantization scenario. We establish a lower-bound on the total number of phase-quantization points K that ensures that the loss function $\mathcal{L}_K^\star$ in (7) does not exceed a given $\varepsilon > 0$.

Theorem 1 (Sufficient Conditions for Proximity to Capacity). *For any $\varepsilon > 0$, we have $\mathcal{L}_K^\star \leq \varepsilon$ if any of the three sufficient conditions holds:*

$$K > \frac{\pi}{\varepsilon\sqrt{3}}|\mathcal{A}| \left[1 + \sqrt{\mathrm{SNR}_{\mathrm{avg}}\mathrm{SNR}_{\mathrm{max}}} \right] , \qquad (11)$$

$$K > \frac{\pi}{\varepsilon\sqrt{3}} \sqrt{\mathbb{E}\left[\frac{A^2}{P_{\mathrm{avg}} \cdot P_A^2(A)} \right]} \sqrt{\mathrm{SNR}_{\mathrm{avg}}\mathrm{SNR}_{\mathrm{max}}}, \qquad (12)$$

or

$$K > \frac{\pi}{\varepsilon\sqrt{3}} \sqrt{\mathbb{E}\left[\frac{A^2}{P_{\mathrm{avg}} \cdot f^2(A)} \right]} \sqrt{\mathrm{SNR}_{\mathrm{avg}}\mathrm{SNR}_{\mathrm{max}}}, \qquad (13a)$$

where

$$f(a) \triangleq \frac{\sqrt[3]{a^2 \cdot P_A(a)}}{\sum_{\bar{a} \in \mathcal{A}} \sqrt[3]{\bar{a}^2 \cdot P_A(\bar{a})}}, \quad \forall a \in \mathcal{A}. \qquad (13b)$$

To the best of our knowledge, Theorem 1 gives the first theoretical performance guarantee of discrete constellations for the peak-power-limited complex Gaussian channel. The sufficient conditions in the theorem follow by upper bounding the information loss $\mathcal{L}_K^\star$ using three different (possibly suboptimal) phase quantizers. In particular, (11) is obtained by setting $k_a = \lfloor K/|\mathcal{A}| \rfloor$, for any $a \in \mathcal{A}$. We refer to this quantizer as the *uniform quantizer*. To get (12), we set $k_a = \lfloor P_A(a) \cdot K \rfloor$, for any $a \in \mathcal{A}$, as proposed in [5]. The

third condition (13a) is achieved by choosing k_a as in (16) below.

The proof of Theorem 1 relies on the following upper bound on the loss function $\mathcal{L}_K(\mathbf{k})$. As previous works mostly focused on numerical evaluations of the information loss, we find this theoretical bound to be of independent interest.

Lemma 1 (General Upper Bound on the Information Loss). *For any $\mathbf{k} \in \mathbb{N}^{|\mathcal{A}|}$ with $\sum_{a \in \mathcal{A}} k_a \leq K$,*

$$\mathcal{L}_K(\mathbf{k}) \leq \sqrt{2}\sqrt{\mathrm{SNR}_{\mathrm{avg}}\mathrm{SNR}_{\mathrm{max}}}$$
$$\times \sqrt{1 - \frac{1}{P_{\mathrm{avg}}}\mathbb{E}\left[A^2 \cdot \mathrm{sinc}\left(\frac{\pi}{k_A} \right) \right]}, \qquad (14a)$$

which can be loosened to

$$\mathcal{L}_K(\mathbf{k}) \leq \frac{\pi}{\sqrt{3}}\sqrt{\mathrm{SNR}_{\mathrm{avg}}\mathrm{SNR}_{\mathrm{max}}}\sqrt{\frac{1}{P_{\mathrm{avg}}}\mathbb{E}\left[\frac{A^2}{k_A^2} \right]} . \qquad (14b)$$

Furthermore, allowing $\mathbf{k}$ to take values in $\mathbb{R}_+^{|\mathcal{A}|}$, the vector $\tilde{\mathbf{k}} \triangleq (\tilde{k}_a)_{a \in \mathcal{A}}$ given by

$$\tilde{k}_a = \frac{\sqrt[3]{a^2 \cdot P_A(a)}}{\sum_{\bar{a} \in \mathcal{A}} \sqrt[3]{\bar{a}^2 \cdot P_A(\bar{a})}} \cdot K, \quad \forall a \in \mathcal{A}, \qquad (15)$$

minimizes the right-hand side (r.h.s.) of (14b).

The upper bounds in Lemma 1 are general in the sense that they apply to any choice of quantizer $\mathbf{k} \in \mathbb{N}^{|\mathcal{A}|}$ and any amplitude distribution (not necessarily the optimal one). Most coding techniques assume a uniform use of constellation points. As the optimal amplitude distribution is not necessarily uniform, one may mimic uniformity of the constellation by quantizing the phase while allowing multiple copies of some points. Doing so gives rise to a trade-off between the number of constellation points and the approximation accuracy of the optimal amplitude distribution. The generality of Lemma 1 enables a theoretic analysis of such scenarios.

Remark 2 (Relation Between Lemma 1 and (13a)). *The vector $\tilde{\mathbf{k}}$ in* (15) *is obtained by relaxing a discrete-valued optimization problem to a continuous domain. A natural choice for the phase-quantization vector $(k_a)_{a \in \mathcal{A}}$ would be to round each $\tilde{k}_a$ given in* (15) *to the next smallest integer, namely,*

$$k_a^\star = \left\lfloor \frac{\sqrt[3]{a^2 \cdot P_A(a)}}{\sum_{\bar{a} \in \mathcal{A}} \sqrt[3]{\bar{a}^2 \cdot P_A(\bar{a})}} \cdot K \right\rfloor, \quad \forall a \in \mathcal{A}. \qquad (16)$$

This choice yields the sufficient condition (13a) *in Lemma 1. Note that this is not necessarily the best choice. For example, if $\sum_{a \in \mathcal{A}} k_a^\star < K$ then one may allocate the unused phase-quantization points to any of the circles corresponding to the different amplitude values $a \in \mathcal{A}$, which would decrease the information loss.*

As mentioned in the introduction, [5] proposed an ad-hoc approach to approximate the semi-continuous capacity-achieving input distribution. Specifically, the authors of [5] employed the same constellation used herein but with $k_a = \lfloor P_A(a) \cdot K \rfloor$. Our results suggest, however, that the dependence of the phase quantizer on the amplitude is as appears in (16).
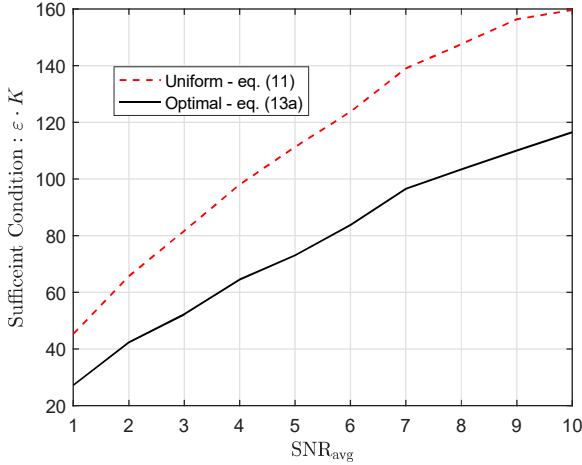
Fig. 1. Comparison between the number of quantization points times the gap from capacity, i.e., $\varepsilon \cdot K$, obtained for the uniform quantizer (see, eq. (11)) and the optimal quantizer (see, eq. (13a)), as a function of $\mathrm{SNR_{avg}}$, where $\sigma^2 = 1$ and $P_{\max} = 10$.

B. Numerical Comparison of the Sufficient Conditions

The sufficient conditions in Theorem 1 depend on the distribution of A. Unfortunately, little is known about the cardinality or the peak amplitude of the optimal input. To evaluate the conditions in Theorem 1, we use an efficient algorithm proposed in [3] to numerically approximate the optimal distribution of A.

Fig. 1 presents a comparison between the sufficient conditions from (11) and (13a). We plot $\varepsilon \cdot K$ (the number of quantization points times the gap from capacity) as a function of $\mathrm{SNR_{avg}}$, for $\sigma^2 = 1$ and $P_{\max} = 10$. The dashed and solid curves correspond to (11) and (13a), respectively. While the sufficient condition in (12) is easily evaluated using the same method, it yields large values of $\varepsilon \cdot K$ (spanning from 300 to around 1700 points). Indeed, since (12) depends on the inverse of P_A, small values of $P_A(a)$ result in a very large contribution to K. To keep a reasonable scale in Fig. 1, we therefore decided not to plot the curve corresponding to (12).

Fig. 1 shows that the quantizer proposed in (16) significantly outperforms the uniform quantizer. The latter sets $k_a = \lfloor K/|\mathcal{A}| \rfloor$, for all $a \in \mathcal{A}$, thus allocating the same number of points to all amplitude values. Roughly speaking, as each amplitude corresponds to a circle in the complex plane, the uniform strategy implies that constellation points may be too sparse on larger circles and too dense on smaller ones. The quantizer from (16), on the other hand, scales the number of constellation points on a circle of radius a according to $\sqrt[3]{a^2 \cdot P_A(a)}$. Consequently, larger and more probable circles are allocated with more points, while smaller and less probable ones contain less constellation points.

C. Extensions to Amplitude and Phase Quantization

Finally, we discuss briefly the scenario where both amplitude and phase quantizers are used. In this case, in the spirit of Lemma 1, the following can be shown.

For any $K \in \mathbb{N}_0$, and given quantizers $\Delta \in \mathfrak{D}$ and $\mathbf{k}_\Delta \in \mathbb{N}^{|\mathcal{D}|}$, with $\mathcal{D} = \mathrm{Image}(\Delta)$,

$$\bar{\mathcal{L}}_K(\mathbf{k}_\Delta, \Delta) \leq \sqrt{2}\sqrt{\mathrm{SNR_{avg}}\mathrm{SNR_{max}}}$$
$$\times \left[1 - \frac{1}{P_{\mathrm{avg}}} \mathbb{E}\left(A \cdot \Delta(A) \cdot \mathrm{sinc}\left(\frac{\pi}{k_{\Delta(A)}} \right) \right) \right]^{\frac{1}{2}}. \quad (17)$$

Using the same arguments as in the proof of Lemma 1, one can show that $\bar{\mathcal{L}}_K(\mathbf{k}_\Delta, \Delta)$ can be further upper bounded by

$$\bar{\mathcal{L}}_K(\mathbf{k}_\Delta, \Delta) \leq \sqrt{2}\sqrt{\mathrm{SNR_{avg}}\mathrm{SNR_{max}}}$$
$$\times \left[1 - \frac{1}{P_{\mathrm{avg}}} \mathbb{E}\left(A \cdot \Delta(A) \cdot \left(1 - \frac{\pi^2}{6k_{\Delta(A)}^2} \right) \right) \right]^{\frac{1}{2}}, \quad (18)$$

and that choosing $\mathbf{k}_\Delta$ according to

$$k_\Delta^{\mathrm{opt}}(a) = \frac{\sqrt[3]{a \cdot \Delta(a) \cdot P_A(a)}}{\sum_{\bar{a} \in \mathcal{D}} \sqrt[3]{\bar{a} \cdot \Delta(\bar{a}) \cdot P_A(\bar{a})}} \cdot K, \ \forall a \in \mathcal{D} \quad (19)$$

minimizes the r.h.s. of (18) when the optimization domain is relaxed to $\mathbb{R}^{|\mathcal{D}|}$.

To further upper bound (18), consider the following (possibly suboptimal) amplitude quantizer

$$\Delta(a) = \begin{cases} \frac{a}{\sqrt{\sum_{a \in \mathcal{M}_{|\mathcal{D}|}} P_A(a) \cdot a^2}} \sqrt{P}, & a \in \mathcal{M}_{|\mathcal{D}|} \\ 0, & \mathrm{otherwise} \end{cases}, \quad (20)$$

where $\mathcal{M}_{|\mathcal{D}|}$ consists of the $|\mathcal{D}|$ elements in $\mathcal{A}$ with the highest values of $a \cdot P_A(a)$. It can be shown that this choice minimizes the information loss when only the amplitude is quantized but not the phase. It therefore seems plausible that $\Delta(A)$ performs also well when both amplitude and phase are quantized. Substituting (19)-(20) into (18), it would remain to optimize over the allocation of amplitude and phase quantization points.

IV. PROOF OF MAIN RESULTS

A. Proof of Theorem 1

To prove Theorem 1, we apply (14b) in Lemma 1 with two possibly suboptimal choices of $\mathbf{k}$. We have

$$\mathcal{L}_K(\mathbf{k}) \leq \frac{\pi}{\sqrt{3}} \sqrt{\mathrm{SNR_{avg}}\mathrm{SNR_{max}}} \sqrt{\frac{1}{P_{\mathrm{avg}}} \mathbb{E}\left[\frac{A^2}{k_A^2} \right]}, \quad (21)$$

which holds for all $\mathbf{k} = (k_a)_{a \in \mathcal{A}}$. Setting $k_a = \tilde{K} \triangleq \left\lfloor \frac{K}{|\mathcal{A}|} \right\rfloor$ into the expected value on the r.h.s. of (21) gives

$$\mathbb{E}\left[\frac{A^2}{k_A^2} \right] = \frac{1}{\tilde{K}^2} \mathbb{E}\left[A^2 \right] \leq P_{\mathrm{avg}} \cdot \left(\frac{K}{|\mathcal{A}|} - 1 \right)^{-2}. \quad (22)$$

Inserting back to (21), we have

$$\mathcal{L}_K^\star \leq \frac{\pi}{\sqrt{3}} \sqrt{\mathrm{SNR_{avg}}\mathrm{SNR_{max}}} \cdot \frac{|\mathcal{A}|}{K - |\mathcal{A}|}. \quad (23)$$

Requiring that the r.h.s. is upper bounded by $\varepsilon > 0$ and isolating K, we obtain

$$K > \frac{\pi}{\varepsilon\sqrt{3}} |\mathcal{A}| \left[1 + \sqrt{\mathrm{SNR_{avg}}\mathrm{SNR_{max}}} \right], \quad (24)$$

as required. The proof of (12) and (13a) follows the same steps but using $k_a = \lfloor P_A(a) \cdot K \rfloor$, $a \in \mathcal{A}$, and $\mathbf{k}$ as given in (16), respectively.

B. Proof of Lemma 1

To prove (14a) we use the following result from [6] concerning the continuity of differential entropy with respect to the quadratic Wasserstein distance between sufficiently regular probability density functions.

Lemma 2. *[6, Proposition 5] Let B be an $\mathbb{R}^n$-valued random vector satisfying $\|B\| \leq \sqrt{nP}$ almost surely, and let $G \sim \mathcal{N}(0, \sigma^2 I_n)$. Assume that B and G are independent, and let $V = B + G$. For any $\mathbb{R}^n$-valued random vector U,*

$$h(U) - h(V) \leq \frac{1}{2\sigma^2} \left[\mathbb{E} \|U\|^2 - \mathbb{E} \|V\|^2 + 2\sqrt{nP} \cdot W_1(P_U, P_V) \right]. \tag{25}$$

We apply Lemma 2 to upper bound the information loss $\mathcal{L}_K(\mathbf{k})$. Treating complex random variables as real-valued, two-dimensional, random vectors, we set $U = Y$, $V = Y_Q$ and we further note that $\mathbb{E} \|U\|^2 = \mathbb{E} \|V\|^2$. This gives

$$\mathcal{L}_K(\mathbf{k}) \leq \frac{1}{\sigma} \sqrt{\mathrm{SNR_{max}}} \cdot W_1(P_Y, P_{Y_Q}). \tag{26}$$

Thus, upper bounding $\mathcal{L}_K(\mathbf{k})$ is tantamount to upper bounding the Wasserstein distance between P_Y and P_{Y_Q}.

By Hölder's inequality one readily gets $W_1(P_Y, P_{Y_Q}) \leq W_2(P_Y, P_{Y_Q})$. Furthermore, by definition, the Wasserstein distance is non-increasing under convolutions, so

$$W_2(P_Y, P_{Y_Q}) \leq W_2(P_X, P_{X_Q}). \tag{27}$$

Consequently, to upper bound the information loss $\mathcal{L}_K(\mathbf{k})$ it suffices to estimate $W_2(P_X, P_{X_Q})$. By the definition of Wasserstein distance, any coupling of P_X and P_{X_Q} yields an upper bound. We use the natural coupling where the conditional distribution of X_Q given X is the uniform distribution on the arc of appropriate width around X. Formally, let $A \sim P_A$ and $\varphi \sim \mathrm{Unif}[0, 2\pi)$ be independent random variables over the same probability space. We set $\bar{X} = A \cdot e^{j\varphi}$ and $\bar{X}_Q = A \cdot e^{jQ_A(\varphi)}$. We now have

$$\begin{aligned} W_2^2(P_X, P_{X_Q}) &\leq \mathbb{E} \left\| \bar{X} - \bar{X}_Q \right\|^2 \\ &= 2P_{\mathrm{avg}} - 2\mathrm{Re}\left(\mathbb{E} \bar{X} \bar{X}_Q^\star \right) \\ &= 2 \left(P_{\mathrm{avg}} - \mathbb{E} \left[A^2 \cos(\varphi - Q_A(\varphi)) \right] \right). \end{aligned} \tag{28}$$

The expected value can be written as

$$\begin{aligned} &\mathbb{E}[A^2 \cdot \cos(\varphi - Q_A(\varphi))] \\ &= \sum_{a \in \mathcal{A}} a^2 P_A(a) \mathbb{E} \cos(\varphi - Q_a(\varphi)) \\ &= \sum_{a \in \mathcal{A}} a^2 P_A(a) \frac{1}{2\pi} \int_{[0, 2\pi]} \mathrm{d}\theta \cos(\theta - Q_a(\theta)). \end{aligned} \tag{29}$$

Recalling the definition of $\Phi_{a,\ell}$ in (5), the integral on the r.h.s. of (29) evaluates to

$$\frac{1}{2\pi} \int_{[0, 2\pi]} \mathrm{d}\theta \cos(\theta - Q_a(\theta))$$

$$= \frac{1}{2\pi} \sum_{\ell=1}^{k_a} \int_{\Phi_{a,\ell}} \mathrm{d}\theta \cos(\theta - Q_a(\theta))$$

$$= \frac{k_a}{2\pi} \int_{\Phi_{a,1}} \mathrm{d}\theta \cos(\theta) = \mathrm{sinc}\left(\frac{\pi}{k_a} \right), \tag{30}$$

where the second equality uses the symmetry of the phase quantizer in (3). Combining (26), (29) and (30), we obtain

$$W_2^2(P_X, P_{X_Q}) \leq 2 \left[P_{\mathrm{avg}} - \mathbb{E}_A \left(A^2 \cdot \mathrm{sinc}\left(\frac{\pi}{k_a} \right) \right) \right]. \tag{31}$$

Together with (27) and (26), this proves (14a). The upper bound in (14b) follows by further upper bounding the r.h.s. of (31) using $\sin(x) \geq x - x^3/6$.

We conclude by showing that (15) minimizes the r.h.s. of (14b), if we relax the constraint that $(k_a)_{a \in \mathcal{A}}$ are integer-valued. To this end, we solve the optimization problem

$$\min_{\mathbf{k} \in \mathbb{R}_+^{|\mathcal{A}|} : \sum_{a \in \mathcal{A}} k_a \leq K} \sum_{a \in \mathcal{A}} \frac{a^2 \cdot P_A(a)}{k_a^2}. \tag{32}$$

The minimization problem in (32) can be solved by resorting to Lagrange multipliers. It follows that

$$\tilde{k}_a = \frac{\sqrt[3]{a^2 \cdot P_A(a)}}{\sum_{\bar{a} \in \mathcal{A}} \sqrt[3]{\bar{a}^2 \cdot P_A(\bar{a})}} \cdot K, \quad \forall a \in \mathcal{A}, \tag{33}$$

minimizes (32) and, hence, also the r.h.s. of (14b).

REFERENCES

[1] C. E. Shannon, "A mathematical theory of communication," *Bell Syst. Tech. J*, vol. 27, no. 3, pp. 379–423, Jul. 1948.

[2] S. Shamai and I. Bar-David, "The capacity of average and peak-power-limited quadrature Gaussian channels," *IEEE Transactions on Information Theory*, vol. 41, no. 4, pp. 1060–1071, Jul 1995.

[3] J. Huang and S. P. Meyn, "Characterization and computation of optimal distributions for channel coding," *IEEE Transactions on Information Theory*, vol. 51, no. 7, pp. 2336–2351, July 2005.

[4] Y. Wu and S. Verdú, "The impact of constellation cardinality on Gaussian channel capacity," in *in Proc. 48th Allerton Conf. Comm., Contr. and Comp., Allerton H., Monticello*, Sept. 2010, pp. 40–53.

[5] R. R. Muller, U. Wachsmann, and J. B. Huber, "Multilevel coding for peak power limited complex Gaussian channels," in *Proceedings of IEEE International Symposium on Information Theory*, Jun 1997, p. 103.

[6] Y. Polyanskiy and Y. Wu, "Wasserstein continuity of entropy and outer bounds for interference channels," *IEEE Transactions on Information Theory*, vol. 62, no. 7, pp. 3992–4002, July 2016.