On Quantization of Log-Likelihood Ratios for Maximum Mutual Information

Andreas Winkelbauer and Gerald Matz

Institute of Telecommunications, Vienna University of Technology
Gusshausstrasse 25/389, 1040 Vienna, Austria
email: {andreas.winkelbauer, gerald.matz}@nt.tuwien.ac.at

***Abstract*—We consider mutual-information-optimal quantization of log-likelihood ratios (LLRs). An efficient algorithm is presented for the design of LLR quantizers based either on the unconditional LLR distribution or on LLR samples. In the latter case, a small number of samples is sufficient and no training data are required. Therefore, our algorithm can be used to design LLR quantizers during data transmission. The proposed algorithm is reminiscent of the famous Lloyd-Max algorithm and is not restricted to any particular LLR distribution.**

I. Introduction

Quantization is well studied in lossy source coding, where quantizers are designed to minimize the average distortion between the input signal and the output signal. However, such quantizer designs may not be appropriate in the context of communications, where the aim is to maximize the data rate rather than to represent a signal with small distortion.

In this paper, we consider the quantization of log-likelihood ratios[1] (LLRs) that is optimal in the sense of maximum mutual information. This is motivated by the fact that many modern receiver designs in digital communications use LLRs to represent reliability information. Note that mutual information quantifies the achievable data rate. Application examples for LLR quantization include bit-interleaved coded modulation [1], iterative receivers [2], and relay networks [3].

We present an efficient quantizer design algorithm that maximizes mutual information and resembles the famous Lloyd-Max algorithm [4], [5]. While the basic version of our algorithm presupposes knowledge of the unconditional LLR distribution, we also present an extension that only requires a relatively small number of actual LLRs but *no* training data.

In contrast to [6], our approach is not restricted to the Gaussian case and involves simple, closed-form expressions. Apart from also being restricted to the Gaussian case, the work in [7] is based on a flawed reformulation of the quantizer design problem. The algorithm in [8] finds an encoding using LLR samples and training data but does not yield the optimal quantization regions.

The paper is organized as follows. In Section II we introduce the problem setup and in Section III we present an algorithm for mutual-information-optimal LLR quantizer design. The convergence of our algorithm is analyzed in Section IV. Generalizations and discussions are provided in Section V. Finally, conclusions are given in Section VI.

Funding by WWTF Grant ICT12-054 and by the EU FP7 project NEWCOM# (GA 318306).

[1]Although some authors prefer the term *L-value*, we use *LLR* throughout.

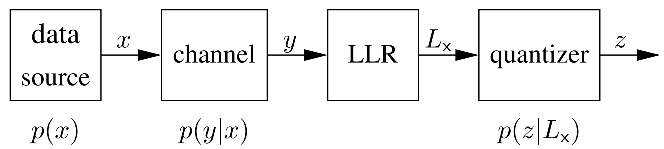

Figure 1: System model for LLR quantizer design. The quantizer is designed to maximize the mutual information $I(\mathsf{x};\mathsf{z})$.

Notation: We use boldface letters for vectors, upright sans-serif letters for random variables, and calligraphic letters for sets. Markov chains are denoted as $\mathsf{x} \leftrightarrow \mathsf{y} \leftrightarrow \mathsf{z}$ which implies that $p(x,z|y) = p(x|y)p(z|y)$. We write $\mathbb{P}\{\cdot\}$ and $\mathbb{E}\{\cdot\}$ for probability and expectation, respectively. The indicator function $\mathbb{1}\{\cdot\}$ equals 1 if its argument is a true statement and it equals 0 otherwise. Entropy, mutual information, and relative entropy are denoted by $H(\cdot)$, $I(\cdot\,;\cdot)$, and $D(\cdot\|\cdot)$, respectively. A Gaussian (normal) distribution with mean μ and variance σ^2 is denoted by $\mathcal{N}(\mu,\sigma^2)$ and log denotes the natural logarithm.

II. Problem Statement

We consider scalar quantization of LLRs in the setting depicted in Figure 1. The binary data $x \in \mathcal{X} = \{-1, 1\}$ is transmitted over the channel $p(y|x)$, yielding the channel output $y \in \mathcal{Y}$. The (posterior) LLR for x is then given by

$$L_{\mathsf{x}}(y) = \log \frac{\mathbb{P}\{\mathsf{x}=1|\mathsf{y}=y\}}{\mathbb{P}\{\mathsf{x}=-1|\mathsf{y}=y\}}. \tag{1}$$

Here we have assumed a memoryless channel to simplify notation. However, the only difference for channels with memory is that the probabilities in (1) have to be conditioned on the entire channel output sequence. The quantizer $q\colon \mathbb{R} \to \mathcal{Z}$ maps the LLR $L_{\mathsf{x}}(y)$ to the quantizer output $z = q(L_{\mathsf{x}})$, where $|\mathcal{Z}|$ is the number of quantization levels.

The optimal quantizer $q(\cdot)$ with n quantization levels is defined by

$$p^{\star}(z|L_{\mathsf{x}}) = \operatorname*{arg\,max}_{p(z|L_{\mathsf{x}})} I(\mathsf{x};\mathsf{z}) \qquad \text{subject to} \quad |\mathcal{Z}| = n, \tag{2}$$

where $I(\mathsf{x};\mathsf{z})$ is the mutual information, defined by

$$I(\mathsf{x};\mathsf{z}) = \sum_{x\in\mathcal{X}} p(x) \sum_{z\in\mathcal{Z}} p(z|x) \log \frac{p(z|x)}{p(z)} \tag{3}$$

with

$$p(z|x) = \int_{\mathbb{R}} p(z|L_{\mathsf{x}}) p(L_{\mathsf{x}}|x) \mathrm{d}L_{\mathsf{x}}. \tag{4}$$

In (4), we have used the fact that $\mathsf{x} \leftrightarrow \mathsf{L}_{\mathsf{x}} \leftrightarrow \mathsf{z}$ is a Markov chain. The quantizer in (2) is described by the probabilistic mapping $p^\star(z|L_{\mathsf{x}})$. Throughout we assume that $p(x)$ and $p(y|x)$ are fixed and known. Furthermore, we focus on continuous output channels in what follows. For binary-input discrete memoryless channels (2) can be solved using dynamic programming [9]. In the following, we restrict ourselves to finding a locally optimal solution of (2).

It is important to note that (2) is substantially different from distortion-based quantizer design. Specifically, we note that (a) the reproducer values are immaterial since mutual information depends only on probability distributions, (b) the problem in (2) is a convex *maximization* problem, and (c) a third random variable in addition to the quantizer input and the quantizer output is involved in (2). Due to (a) it suffices to choose $\mathcal{Z} = \{1,\dots,n\}$ and (b) entails that the solution of (2) is a *deterministic* quantizer, i.e., we have $p^\star(z|L_{\mathsf{x}}) \in \{0,1\}$. Deterministic quantizers can be characterized in terms of the associated *quantization regions*

$$\mathcal{L}_z = \{L_{\mathsf{x}} \in \mathbb{R} : p(z|L_{\mathsf{x}}) = 1\}, \quad z = 1,\dots,n. \tag{5}$$

Using (3) and (5), we can rewrite (2) as

$$\max_{\mathcal{L}_1,\dots,\mathcal{L}_n} \sum_{x\in\mathcal{X}} p(x) \sum_{z\in\mathcal{Z}} \int_{\mathcal{L}_z} p(L_{\mathsf{x}}|x)\mathrm{d}L_{\mathsf{x}} \log \frac{\int_{\mathcal{L}_z} p(L_{\mathsf{x}}|x)\mathrm{d}L_{\mathsf{x}}}{\int_{\mathcal{L}_z} p(L_{\mathsf{x}})\mathrm{d}L_{\mathsf{x}}}. \tag{6}$$

Finding an optimal quantizer thus amounts to finding the optimal quantization regions. In general, problems of the type (6) are difficult to solve since the regions $\mathcal{L}_z$, $z = 1,\dots,n$, need not be convex (see, e.g., [10]). However, in our case [11, Theorem 1] ensures that each set $\mathcal{L}_z$ is convex. We can thus rewrite (6) as

$$\max_{\boldsymbol{g}} I(\boldsymbol{g}), \tag{7}$$

where $\boldsymbol{g} = (g_1 \cdots g_{n-1})^{\mathrm{T}}$ and

$$I(\boldsymbol{g}) = \sum_{x\in\mathcal{X}} p(x) \sum_{i=1}^{n} \int_{g_{i-1}}^{g_i} p(L_{\mathsf{x}}|x)\mathrm{d}L_{\mathsf{x}} \log \frac{\int_{g_{i-1}}^{g_i} p(L_{\mathsf{x}}|x)\mathrm{d}L_{\mathsf{x}}}{\int_{g_{i-1}}^{g_i} p(L_{\mathsf{x}})\mathrm{d}L_{\mathsf{x}}} \tag{8}$$

(we use the convention $g_0 = -\infty$ and $g_n = \infty$). Hence, the optimal quantizer is determined solely by the $n-1$ *quantizer boundaries* $g_1,\dots,g_{n-1}$. In the following we distinguish the cases where the unconditional LLR distribution $p(L_{\mathsf{x}})$ is either known or unknown. In the latter case, we assume that LLR samples are available.

We note that (2) could be solved by a suitably modified version of the information bottleneck (IB) algorithm [12]. This approach is useful, e.g., for the design of channel-optimized vector quantizers maximizing mutual information [13]. However, the IB algorithm does not take advantage of the special structure of our problem and is computationally more expensive than the algorithm we propose next.

III. Quantizer Design

Before we derive our quantizer design algorithm, we recall a few properties of LLRs (see, e.g., [14, Section 3.3]). The posterior probability of x can be expressed in terms of $L_{\mathsf{x}}(y)$ as follows:

$$\mathbb{P}\{\mathsf{x}=x|\mathsf{y}=y\} = \frac{1}{1+e^{-xL_{\mathsf{x}}(y)}}, \quad x \in \mathcal{X}. \tag{9}$$

A basic property of LLRs is that $L_{\mathsf{x}}\big(L_{\mathsf{x}}(y)\big) = L_{\mathsf{x}}(y)$ [14, Lemma 3.2]. Hence, conditioning on the LLR instead of the observation in (9) does not change the result and we thus have

$$\mathbb{P}\{\mathsf{x}=x|\mathsf{L}_{\mathsf{x}}=L_{\mathsf{x}}\} = \frac{1}{1+e^{-xL_{\mathsf{x}}}}, \quad x \in \mathcal{X}. \tag{10}$$

Using Bayes' rule together with (10) allows us to rewrite the conditional distributions of L_{x} given $x \in \mathcal{X}$ as

$$p(L_{\mathsf{x}}|\mathsf{x}=x) = \frac{1}{1+e^{-xL_{\mathsf{x}}}} \frac{p(L_{\mathsf{x}})}{p(x)}. \tag{11}$$

We note that (11) couples the three distributions $p(L_{\mathsf{x}})$, $p(L_{\mathsf{x}}|\mathsf{x}=1)$, and $p(L_{\mathsf{x}}|\mathsf{x}=-1)$ such that any one of them is sufficient to express the other two. Rewriting the objective function (8) using (11) yields

$$I(\boldsymbol{g}) = \sum_{x\in\mathcal{X}} \sum_{i=1}^{n} \int_{g_{i-1}}^{g_i} \frac{p(L_{\mathsf{x}})}{1+e^{-xL_{\mathsf{x}}}}\mathrm{d}L_{\mathsf{x}} \log \frac{\int_{g_{i-1}}^{g_i} \frac{p(L_{\mathsf{x}})}{1+e^{-xL_{\mathsf{x}}}}\mathrm{d}L_{\mathsf{x}}}{p(x)\int_{g_{i-1}}^{g_i} p(L_{\mathsf{x}})\mathrm{d}L_{\mathsf{x}}}. \tag{12}$$

We emphasize that the expression in (12) is obtained without invoking a Gaussian assumption and involves only the *unconditional* LLR distribution.

A. Known LLR Distribution

Next, we derive our quantizer design algorithm for the case where $p(L_{\mathsf{x}})$ is known. We observe that the ratio of the integrals in the logarithm of (12) equals the posterior probabilities $p(x|z)$, i.e., we have

$$\mathbb{P}\{\mathsf{x}=x|\mathsf{z}=i\} = \frac{\int_{g_{i-1}}^{g_i} \frac{p(L_{\mathsf{x}})}{1+e^{-xL_{\mathsf{x}}}}\mathrm{d}L_{\mathsf{x}}}{\int_{g_{i-1}}^{g_i} p(L_{\mathsf{x}})\mathrm{d}L_{\mathsf{x}}}, \quad i = 1,\dots,n. \tag{13}$$

We next associate the ith quantizer output, $i = 1,\dots,n$, with its corresponding LLR for x which is given by

$$L_i^\star = \log \frac{\mathbb{P}\{\mathsf{x}=1|\mathsf{z}=i\}}{\mathbb{P}\{\mathsf{x}=-1|\mathsf{z}=i\}} = \log \frac{\int_{g_{i-1}}^{g_i} \frac{p(L_{\mathsf{x}})}{1+e^{-L_{\mathsf{x}}}}\mathrm{d}L_{\mathsf{x}}}{\int_{g_{i-1}}^{g_i} \frac{p(L_{\mathsf{x}})}{1+e^{L_{\mathsf{x}}}}\mathrm{d}L_{\mathsf{x}}}. \tag{14}$$

Using (13) and (14) we can rewrite (12) as follows:

$$I(\boldsymbol{g}) = -\sum_{x\in\mathcal{X}}\sum_{i=1}^{n}\int_{g_{i-1}}^{g_i}\frac{p(L_{\mathsf{x}})}{1+e^{-xL_{\mathsf{x}}}}\mathrm{d}L_{\mathsf{x}}\log p(x)\big(1+e^{-xL_i^\star}\big). \tag{15}$$

Note that $L_i^\star$, $i=1,\ldots,n$, in (15) just serves as a shorthand notation for the expression on the right-hand side of (14) that actually depends only on the quantizer boundaries.

Next, we let $L_i \in \mathbb{R}$, $i=1,\ldots,n$, be *arbitrary* real numbers and we define the following modified objective function:

$$I(\boldsymbol{g},\boldsymbol{L}) = -\sum_{x\in\mathcal{X}}\sum_{i=1}^{n}\int_{g_{i-1}}^{g_i}\frac{p(L_{\mathsf{x}})}{1+e^{-xL_{\mathsf{x}}}}\mathrm{d}L_{\mathsf{x}}\log p(x)\big(1+e^{-xL_i}\big), \tag{16}$$

where $\boldsymbol{L} = (L_1 \cdots L_n)^{\mathrm{T}}$. The following result establishes a relation between (15) and (16).

Theorem 1. *The functions $I(\boldsymbol{g})$ and $I(\boldsymbol{g},\boldsymbol{L})$ are related as*

$$I(\boldsymbol{g}) = \max_{\boldsymbol{L}} I(\boldsymbol{g},\boldsymbol{L}), \tag{17}$$

where the maximum in (17) is achieved by (14).

Proof: Writing the difference $I(\boldsymbol{g}) - I(\boldsymbol{g},\boldsymbol{L})$ in terms of relative entropy yields $D\big(p(x|z)p(z)\|f_z(x)p(z)\big)$, where $f_z(x) = \big(1+e^{-xL_z}\big)^{-1}$. Due to the information inequality we have $I(\boldsymbol{g}) \geq I(\boldsymbol{g},\boldsymbol{L})$. Furthermore, we have $I(\boldsymbol{g}) = I(\boldsymbol{g},\boldsymbol{L})$ iff $f_z(x) = p(x|z)$, i.e., iff $L_z = L_z^\star$, $z=1,\ldots,n$. ■

Due to Theorem 1, the quantizer design problem (7) can be rewritten as

$$\max_{\boldsymbol{g}} I(\boldsymbol{g}) = \max_{\boldsymbol{g}}\max_{\boldsymbol{L}} I(\boldsymbol{g},\boldsymbol{L}). \tag{18}$$

We note that the approach of rewriting the original problem as a double maximization is similar to that used to derive the Blahut-Arimoto algorithm [15], [16]. The result in (18) allows us to approach the quantizer design problem via alternating maximization. To this end, we compute the following partial derivatives ($i=1,\ldots,n-1$):

$$\frac{\partial I(\boldsymbol{g},\boldsymbol{L})}{\partial g_i} = -p_{\mathsf{L}_{\mathsf{x}}}(g_i)\sum_{x\in\mathcal{X}}\frac{1}{1+e^{-xg_i}}\log\frac{1+e^{-xL_i}}{1+e^{-xL_{i+1}}}. \tag{19}$$

Setting the right-hand side of (19) to zero, assuming $p_{\mathsf{L}_{\mathsf{x}}}(g_i) > 0$, and solving for g_i yields

$$g_i = \log\frac{\log\dfrac{1+e^{L_{i+1}}}{1+e^{L_i}}}{\log\dfrac{1+e^{-L_i}}{1+e^{-L_{i+1}}}}, \qquad i=1,\ldots,n-1. \tag{20}$$

Our algorithm starts with an initialization for $\boldsymbol{g}$ such that $p_{\mathsf{L}_{\mathsf{x}}}(g_i) > 0$, $i=1,\ldots,n-1$. Next, $\boldsymbol{L}$ is updated using (14) and subsequently $\boldsymbol{g}$ is updated using (20). In this manner, our algorithm alternatingly updates $\boldsymbol{L}$ and $\boldsymbol{g}$ until the relative increase in $I(\mathsf{x};\mathsf{z})$ between two iterations is below a prescribed threshold or a certain number of iterations has been performed. Algorithm 1 summarizes the proposed algorithm.

Algorithm 1 *Scalar LLR quantizer design for maximum mutual information (known LLR distribution).*

Input: $\mathcal{Z}$, $p(L_{\mathsf{x}})$, $\varepsilon > 0$, $M \in \mathbb{N}$

Initialization: $\mathcal{X} \leftarrow \{-1,1\}$, $\eta \leftarrow \infty$, $m \leftarrow 1$, $n \leftarrow |\mathcal{Z}|$, choose $\boldsymbol{g}^{(0)}$ such that $g_1^{(0)} < \cdots < g_{n-1}^{(0)}$ and $p(g_i^{(0)}) > 0$, $i=1,\ldots,n-1$, $g_0^{(0)} \leftarrow -\infty$, $g_n^{(0)} \leftarrow \infty$

$$f_i^{(0)}(x) \leftarrow 1 + \left(\frac{\int_{g_{i-1}^{(0)}}^{g_i^{(0)}}\frac{p(L_{\mathsf{x}})}{1+e^{-L_{\mathsf{x}}}}\mathrm{d}L_{\mathsf{x}}}{\int_{g_{i-1}^{(0)}}^{g_i^{(0)}}\frac{p(L_{\mathsf{x}})}{1+e^{L_{\mathsf{x}}}}\mathrm{d}L_{\mathsf{x}}}\right)^{-x}, \quad i=1,\ldots,n,\ x\in\mathcal{X}$$

$$I^{(0)} \leftarrow -\sum_{x\in\mathcal{X}}\sum_{i=1}^{n}\int_{g_{i-1}^{(0)}}^{g_i^{(0)}}\frac{p(L_{\mathsf{x}})}{1+e^{-xL_{\mathsf{x}}}}\mathrm{d}L_{\mathsf{x}}\log p(x) f_i^{(0)}(x)$$

while $\eta \geq \varepsilon$ **and** $m \leq M$ **do**

$g_0^{(m)} \leftarrow -\infty$, $g_n^{(m)} \leftarrow \infty$

$$g_i^{(m)} \leftarrow \log\frac{\log\dfrac{f_{i+1}^{(m-1)}(-1)}{f_i^{(m-1)}(-1)}}{\log\dfrac{f_i^{(m-1)}(1)}{f_{i+1}^{(m-1)}(1)}}, \quad i=1,\ldots,n-1$$

$$f_i^{(m)}(x) \leftarrow 1 + \left(\frac{\int_{g_{i-1}^{(m)}}^{g_i^{(m)}}\frac{p(L_{\mathsf{x}})}{1+e^{-L_{\mathsf{x}}}}\mathrm{d}L_{\mathsf{x}}}{\int_{g_{i-1}^{(m)}}^{g_i^{(m)}}\frac{p(L_{\mathsf{x}})}{1+e^{L_{\mathsf{x}}}}\mathrm{d}L_{\mathsf{x}}}\right)^{-x}, \quad i=1,\ldots,n,\ x\in\mathcal{X}$$

$$I^{(m)} \leftarrow -\sum_{x\in\mathcal{X}}\sum_{i=1}^{n}\int_{g_{i-1}^{(m)}}^{g_i^{(m)}}\frac{p(L_{\mathsf{x}})}{1+e^{-xL_{\mathsf{x}}}}\mathrm{d}L_{\mathsf{x}}\log p(x) f_i^{(m)}(x)$$

$\eta \leftarrow \big(I^{(m)} - I^{(m-1)}\big)/I^{(m)}$

$m \leftarrow m+1$

end while

Output: quantizer boundaries $\boldsymbol{g}^{(m-1)}$

B. Unknown LLR Distribution

Next, we assume that K independent LLR samples $L_{\mathsf{x},k}$, $k=1,\ldots,K$, are available but the underlying distribution $p(L_{\mathsf{x}})$ is unknown. In this case, one could first compute an estimate of $p(L_{\mathsf{x}})$ that is then used in Algorithm 1. However, we instead propose to compute estimates of (14) and (16) such that we do not require $p(L_{\mathsf{x}})$ in our algorithm. To this end, we rewrite the LLRs $L_i^\star$, $i=1,\ldots,n$, as follows (cf. (14)):

$$L_i^\star = \log\frac{\mathbb{E}\left\{\dfrac{\mathbb{1}\{\mathsf{L}_{\mathsf{x}} \in [g_{i-1},g_i)\}}{1+e^{-\mathsf{L}_{\mathsf{x}}}}\right\}}{\mathbb{E}\left\{\dfrac{\mathbb{1}\{\mathsf{L}_{\mathsf{x}} \in [g_{i-1},g_i)\}}{1+e^{\mathsf{L}_{\mathsf{x}}}}\right\}} \approx \log\frac{\sum_{k\in\mathcal{K}_i}\big(1+e^{-L_{\mathsf{x},k}}\big)^{-1}}{\sum_{k\in\mathcal{K}_i}\big(1+e^{L_{\mathsf{x},k}}\big)^{-1}}. \tag{21}$$

Here, $\mathcal{K}_i = \big\{k \in \{1,\ldots,K\} : L_{\mathsf{x},k} \in [g_{i-1},g_i)\big\}$ denotes the set of samples that belong to the ith quantization region. Due to the law of large numbers, the estimate on the right-hand side of (21) converges to $L_i^\star$ as $K \to \infty$. Similarly, we obtain

Algorithm 2 *Scalar LLR quantizer design for maximum mutual information (unknown LLR distribution).*

Input: $\mathcal{Z}$, $\{L_{\mathsf{x},k}\}_{k=1}^{K}$, $\varepsilon > 0$, $M \in \mathbb{N}$

Initialization: $\mathcal{X} \leftarrow \{-1, 1\}$, $\eta \leftarrow \infty$, $m \leftarrow 1$, $n \leftarrow |\mathcal{Z}|$, choose $\boldsymbol{g}^{(0)}$ such that $g_1^{(0)} < \cdots < g_{n-1}^{(0)}$ and $|\mathcal{K}_i^{(0)}| > 0$, $i = 1, \ldots, n-1$, $g_0^{(0)} \leftarrow -\infty$, $g_n^{(0)} \leftarrow \infty$

$$f_i^{(0)}(x) \leftarrow 1 + \left(\frac{\sum_{k \in \mathcal{K}_i^{(0)}} \left(1 + e^{-L_{\mathsf{x},k}}\right)^{-1}}{\sum_{k \in \mathcal{K}_i^{(0)}} \left(1 + e^{L_{\mathsf{x},k}}\right)^{-1}} \right)^{-x}, \quad i = 1, \ldots, n, \quad x \in \mathcal{X}$$

$$I^{(0)} \leftarrow -\sum_{x \in \mathcal{X}} \sum_{i=1}^{n} \frac{\log p(x) f_i^{(0)}(x)}{K} \sum_{k \in \mathcal{K}_i^{(0)}} \frac{1}{1 + e^{-x L_{\mathsf{x},k}}}$$

while $\eta \geq \varepsilon$ **and** $m \leq M$ **do**

$\quad g_0^{(m)} \leftarrow -\infty$, $g_n^{(m)} \leftarrow \infty$

$$g_i^{(m)} \leftarrow \log \frac{\log \frac{f_{i+1}^{(m-1)}(-1)}{f_i^{(m-1)}(-1)}}{\log \frac{f_i^{(m-1)}(1)}{f_{i+1}^{(m-1)}(1)}}, \quad i = 1, \ldots, n-1$$

$$\mathcal{K}_i^{(m)} \leftarrow \left\{ k \in \{1, \ldots, K\} : L_{\mathsf{x},k} \in \left[g_{i-1}^{(m)}, g_i^{(m)}\right) \right\}$$

$$f_i^{(m)}(x) \leftarrow 1 + \left(\frac{\sum_{k \in \mathcal{K}_i^{(m)}} \left(1 + e^{-L_{\mathsf{x},k}}\right)^{-1}}{\sum_{k \in \mathcal{K}_i^{(m)}} \left(1 + e^{L_{\mathsf{x},k}}\right)^{-1}} \right)^{-x}, \quad i = 1, \ldots, n, \quad x \in \mathcal{X}$$

$$I^{(m)} \leftarrow -\sum_{x \in \mathcal{X}} \sum_{i=1}^{n} \frac{\log p(x) f_i^{(m)}(x)}{K} \sum_{k \in \mathcal{K}_i^{(m)}} \frac{1}{1 + e^{-x L_{\mathsf{x},k}}}$$

$\quad \eta \leftarrow \left(I^{(m)} - I^{(m-1)}\right) / I^{(m)}$

$\quad m \leftarrow m + 1$

end while

Output: quantizer boundaries $\boldsymbol{g}^{(m-1)}$

the following unbiased and consistent estimate of $I(\boldsymbol{g}, \boldsymbol{L})$:

$$I(\boldsymbol{g}, \boldsymbol{L}) \approx -\sum_{x \in \mathcal{X}} \sum_{i=1}^{n} \frac{\log p(x)\left(1 + e^{-x L_i}\right)}{K} \sum_{k \in \mathcal{K}_i} \frac{1}{1 + e^{-x L_{\mathsf{x},k}}}. \tag{22}$$

The estimates in (21) and (22) allow us to modify Algorithm 1 to perform quantizer design based on LLR samples only (cf. Algorithm 2).

IV. Convergence

In this section, we prove the convergence of Algorithm 1 to a local optimum of (7). The output of Algorithm 2 depends on the samples $L_{\mathsf{x},k}$, $k = 1, \ldots, K$, which are taken at random. However, since the estimates (21) and (22) are consistent, Algorithm 2 is asymptotically equivalent to Algorithm 1.

To prove the convergence of Algorithm 1, we show that the updates (14) and (20) do not decrease the value of the objective function. Specifically, we establish the following inequalities:

$$I\big(\boldsymbol{g}^{(m)}, \boldsymbol{L}^{(m)}\big) \leq I\big(\boldsymbol{g}^{(m+1)}, \boldsymbol{L}^{(m)}\big) \leq I\big(\boldsymbol{g}^{(m+1)}, \boldsymbol{L}^{(m+1)}\big). \tag{23}$$

Since the objective function is upper bounded as $H(\mathsf{x}) \geq I\big(\boldsymbol{g}^{(m)}, \boldsymbol{L}^{(m)}\big)$, the inequalities in (23) imply convergence to a local optimum as $m \to \infty$.

The second inequality in (23) follows immediately from Theorem 1. To prove that the first inequality in (23) holds, we show that $\boldsymbol{g}^{(m+1)}$ corresponds to a local maximum. To this end, we note that (cf. (19))

$$\frac{\partial^2}{\partial g_i \partial g_j} I\big(\boldsymbol{g}^{(m+1)}, \boldsymbol{L}^{(m)}\big) = 0, \qquad i \neq j, \tag{24}$$

and therefore we have to show that

$$\frac{\partial^2}{\partial g_i^2} I\big(\boldsymbol{g}^{(m+1)}, \boldsymbol{L}^{(m)}\big) < 0, \qquad i = 1, \ldots, n-1. \tag{25}$$

The inequality in (25) is equivalent to the following inequality (we suppress the iteration index in what follows):

$$-\sum_{x \in \mathcal{X}} \frac{x e^{-x g_i}}{(1 + e^{-x g_i})^2} \log \frac{1 + e^{-x L_i}}{1 + e^{-x L_{i+1}}} < 0. \tag{26}$$

The above inequality can in turn be shown to be equivalent to $L_i < L_{i+1}$, which is a true statement since $\mathbb{P}\{\mathsf{x}{=}1 | \mathsf{L}_{\mathsf{x}}{=}L_{\mathsf{x}}\}$ is monotonically increasing in L_{x}. Hence, the unique stationary point $\boldsymbol{g}^{(m+1)}$ is a local maximum and we have therefore established the first inequality in (23). This concludes the proof of the convergence of Algorithm 1.

V. Discussion

Finally, we give some remarks and discuss generalizations.

Remark 1: The initialization for $\boldsymbol{g}$ in our algorithm may affect the resulting quantizer. We have found that initializing $\boldsymbol{g}$ using the maximum output entropy quantizer [17] yields good results.

Remark 2: Algorithm 2 yields excellent results with few samples and may therefore be used to design LLR quantizers online during data transmission. Figure 2 shows the deviation of the resulting mutual information $I(\mathsf{x}; \mathsf{z})$ from the optimal value (in percent) versus the number of available samples. Here, the LLRs are conditionally Gaussian, i.e., $\mathsf{L}_{\mathsf{x}} | x \sim \mathcal{N}(x\mu, 2\mu)$, $\mathsf{x} \in \mathcal{X}$ is equally likely, we have $|\mathcal{Z}| = 4$, and 10^6 realizations have been generated for each value of K. We observe that the average error is below $1\,\%$ even for as few as 10 samples. About 40 (60) samples are required to keep the error below $1\,\%$ for $95\,\%$ $(99\,\%)$ of all realizations.

Remark 3: Algorithm 1 and 2 require that the LLRs are computed exactly, which is not always feasible. Using our algorithm with approximate LLRs will generally entail a performance penalty. However, it is straightforward to reformulate Algorithms 1 and 2 such that they respectively operate on the conditional LLR distributions and on LLR samples with training data.

Remark 4: Once an optimal LLR quantizer has been designed, it is not necessary to perform the LLR computation (1) prior to quantization. Indeed, for any LLR quantizer

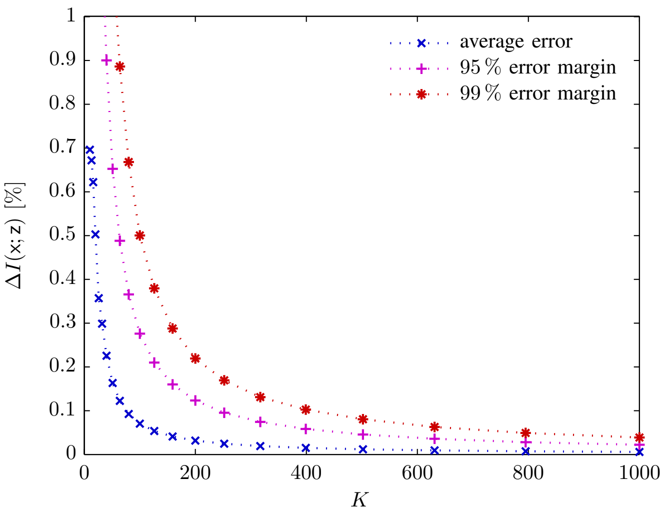

Figure 2: Relative deviation of $I(\mathsf{x};\mathsf{z})$ from the optimal value versus the number of available samples (10^6 realizations).

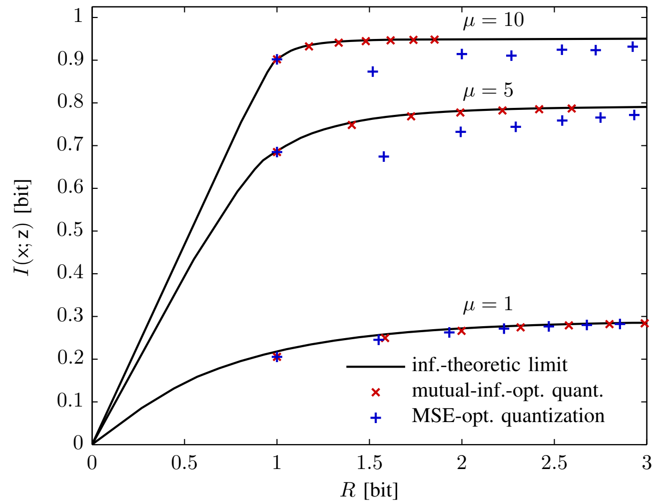

Figure 3: Comparison of scalar quantizers with 2 to 8 quantization levels to the information-theoretic limit.

$q\colon \mathbb{R} \to \mathcal{Z}$ there exists an equivalent quantizer $\tilde{q}\colon \mathcal{Y} \to \mathcal{Z}$ with $\tilde{q}(y) = q\big(L_{\mathsf{x}}(y)\big)$. Note that the quantization regions of $\tilde{q}$ are given by $\mathcal{Y}_i = \{y \in \mathcal{Y} : L_{\mathsf{x}}(y) \in [g_{i-1}, g_i)\}$, $i = 1, \ldots, n$.

Remark 5: Our algorithm can be generalized to channels with nonbinary input. A generalization of Algorithm 1 to the nonbinary case is given in [14, Section 5.3].

Remark 6: In Figure 3, we compare the mutual information $I(\mathsf{x};\mathsf{z})$ achieved by scalar LLR quantizers with $2, \ldots, 8$ quantization levels ('$\times$' markers) to the information-theoretic limit (solid lines; computed using the IB algorithm [12]) for conditionally Gaussian LLRs with $\mu \in \{1, 5, 10\}$ (cf. Remark 2). Note that $R = H(\mathsf{z})$. We observe that the quantizers closely approach the information-theoretic limit. Hence, vector quantization can only provide a negligible performance gain in this setting. Furthermore, Figure 3 shows that MSE-optimal quantizers ('$+$' markers) are inferior to mutual-information-optimal quantizers.

Remark 7: A MATLAB implementation of our algorithm and code to reproduce the above figures is available at [18].

VI. Conclusions

We have studied LLR quantizer design in a communications context with the aim to maximize the data rate. We have formulated the optimization problem for mutual-information-optimal quantizer design in terms of the quantizer boundaries. The corresponding objective function can be rewritten such that it only involves the unconditional LLR distribution. We have then introduced a modified objective function that enables the quantizer design using alternating maximization. The resulting algorithm resembles the famous Lloyd-Max algorithm and is shown to converge to local optimum. A modification of our algorithm allows us to perform quantizer design based on a small number of LLR samples. This is relevant because it enables on-the-fly quantizer design during data transmission. Furthermore, we briefly mention the extension to approximate LLRs and the generalization to the nonbinary case.

References

[1] C. Novak, P. Fertl, and G. Matz, "Quantization for soft-output demodulators in bit-interleaved coded modulation systems," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, July 2009, pp. 1070–1074.

[2] H. Wymeersch, *Iterative receiver design*. Cambridge Univ. Press, 2007.

[3] A. Winkelbauer and G. Matz, "Joint network-channel coding in the multiple-access relay channel: Beyond two sources," in *Proc. 5th Int. Symp. Commun., Control, Signal Process. (ISCCSP)*, May 2012.

[4] S. Lloyd, "Least squares quantization in PCM," *IEEE Trans. Inf. Theory*, vol. 28, no. 2, pp. 129–137, March 1982.

[5] J. Max, "Quantization for minimum distortion," *IEEE Trans. Inf. Theory*, vol. 6, no. 2, pp. 7–12, 1960.

[6] W. Rave, "Quantization of log-likelihood ratios to maximize mutual information," *IEEE Signal Process. Lett.*, vol. 16, no. 4, pp. 283–286, April 2009.

[7] Y. Samayoa and J. Osterman, "A Lloyd-Max based quantizer of L-values for AWGN and Rayleigh fading channel," in *Proc. 6th Int. Conf. Wireless Commun. Signal Process. (WCSP)*, Oct. 2014.

[8] M. Danieli, S. Forchhammer, J. Andersen, L. Christensen, and S. Christensen, "Maximum mutual information vector quantization of log-likelihood ratios for memory efficient HARQ implementations," in *Proc. Data Compression Conf. (DCC)*, March 2010, pp. 30–39.

[9] B. Kurkoski and H. Yagi, "Concatenation of a discrete memoryless channel and a quantizer," in *Proc. IEEE Inf. Theory Workshop (ITW)*, Jan. 2010.

[10] G. Zeitler, A. Singer, and G. Kramer, "Low-precision A/D conversion for maximum information rate in channels with memory," *IEEE Trans. Commun.*, vol. 60, no. 9, pp. 2511–2521, Sept. 2012.

[11] D. Burshtein, V. D. Pietra, D. Kanevsky, and A. Nadas, "Minimum impurity partitions," *Ann. Stat.*, vol. 20, no. 3, pp. 1637–1646, Sept. 1992.

[12] N. Tishby, F. Pereira, and W. Bialek, "The information bottleneck method," in *Proc. 37th Annu. Allerton Conf. Commun., Control, Comput.*, Sept. 1999, pp. 368–377.

[13] A. Winkelbauer, G. Matz, and A. Burg, "Channel-optimized vector quantization with mutual information as fidelity criterion," in *Proc. 47th Asilomar Conf. Signals, Systems, Computers*, Nov. 2013, pp. 851–855.

[14] A. Winkelbauer, "Blind performance estimation and quantizer design with applications to relay networks," Ph.D. dissertation, Vienna University of Technology, Dec. 2014.

[15] R. E. Blahut, "Computation of channel capacity and rate-distortion functions," *IEEE Trans. Inf. Theory*, vol. 18, pp. 460–473, 1972.

[16] S. Arimoto, "An algorithm for computing the capacity of arbitrary discrete memoryless channels," *IEEE Trans. Inf. Theory*, vol. 18, pp. 14–20, 1972.

[17] D. Messerschmitt, "Quantizing for maximum output entropy," *IEEE Trans. Inf. Theory*, vol. 17, no. 5, pp. 612–612, Sept. 1971.

[18] [Online]. Available: https://github.com/andreaswinkelbauer/spawc2015