

Quantization for Soft-Output Demodulators in Bit-Interleaved Coded Modulation Systems

Clemens Novak, Peter Fertl, and Gerald Matz

Institut für Nachrichtentechnik und Hochfrequenztechnik, Vienna University of Technology

Email: {cnovak, pfertl, gmatz}@nt.tuwien.ac.at

Abstract—We study quantization of log-likelihood ratios (LLR) in bit-interleaved coded modulation (BICM) systems in terms of an equivalent discrete channel. We propose to design the quantizer such that the quantizer outputs become equiprobable. We investigate semi-analytically and numerically the ergodic and outage capacity over single- and multiple-antenna channels for different quantizers. Finally, we show bit error rate simulations for BICM systems with LLR quantization using a rate 1/2 low-density parity-check code.

I. INTRODUCTION

Bit-interleaved coded modulation (BICM) is an attractive scheme for wireless communications where a block of information bits is mapped to transmit symbols via a channel encoder and a symbol mapper separated by a code bit interleaver [1]. At the receiver side, a demodulator (demapper, detector) calculates *log-likelihood ratios (LLR)* for the code bits, which are de-interleaved and passed to the channel decoder.

Theoretically, one real-valued LLR per code bit needs to be computed and stored by the receiver. Clearly, practical digital implementations can only use finite word-length approximations of real numbers, which motivates the study of LLR quantization. We note that LLR quantization is also relevant for wireless (relay) networks that perform distributed turbo and network coding by exchanging soft information between different nodes. Optimal LLR quantization maximizing information rate for the special case of BPSK modulation over an AWGN channel was considered in [2]. However, an extension of this approach to other channels and modulations appears infeasible. Thus, we consider a different quantizer design in this paper which allows for simple implementation while only slightly degrading information rate. More specifically, our contributions are as follows:

- We propose to quantize the LLRs such that the quantizer outputs become equiprobable and provide appropriate reliability information to the channel decoder.
- We investigate the impact of the proposed LLR quantization on ergodic and outage rate, using a semi-analytical approach for single-input single-output (SISO) systems with BPSK and Monte-Carlo simulations otherwise.
- We develop a method for designing and implementing the proposed quantizer during data transmission.
- We provide bit error rate (BER) simulations for BICM systems with LLR quantization using *low-density parity-check (LDPC)* codes.

The paper is organized as follows. Section II presents the system model and Section III discusses the proposed LLR quantization based on an equivalent discrete channel. In Sections IV and V, we study the system capacity of SISO- and MIMO-BICM systems, respectively. The estimation of the quantizer parameters is addressed in Section VI and BER results are provided in Section VII.

II. SYSTEM MODEL

We consider a MIMO-BICM system with M_T transmit antennas and M_R receive antennas (SISO-BICM can be viewed as special case with $M_T = M_R = 1$). A block diagram is shown in Fig. 1. A sequence of information bits $b[n']$ is encoded using an error-correcting code, passed through a *bitwise interleaver* Π and then scrambled by a pseudo-random sequence $p_l[n]$. The uniformly distributed interleaved and scrambled code bits are demultiplexed into M_T antenna streams (“layers”). In each layer, groups of m code bits are mapped to (complex) data symbols $x_k[n] \in \mathcal{A}$, $k = 1, \dots, M_T$; here, \mathcal{A} denotes the symbol alphabet of size $|\mathcal{A}| = 2^m$. The transmit vector at symbol time n is given by $\mathbf{x}[n] \triangleq (x_1[n] \dots x_{M_T}[n])^T$ and carries $R_0 = mM_T$ interleaved code bits $c_l[n]$, $l = 1, \dots, R_0$.

Assuming flat fading, the length- M_R receive vector equals

$$\mathbf{y}[n] = \mathbf{H}[n]\mathbf{x}[n] + \mathbf{w}[n]. \quad (1)$$

Here, $\mathbf{H}[n]$ is the $M_R \times M_T$ MIMO channel matrix and $\mathbf{w}[n] \sim \mathcal{CN}(\mathbf{0}, \sigma^2 \mathbf{I})$ denotes the complex Gaussian noise vector. In the following, we will omit the time index n to simplify notation.

At the receiver, the max-log demodulator calculates LLRs for the code bits c_l according to [3]

$$\Lambda_l = \frac{1}{\sigma^2} \left[\min_{\mathbf{x} \in \mathcal{X}_l^0} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|^2 - \min_{\mathbf{x} \in \mathcal{X}_l^1} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|^2 \right]. \quad (2)$$

Here, \mathcal{X}_l^b denotes the set of transmit vectors for which $c_l = b$. These LLRs (or approximate/quantized versions thereof) are de-scrambled by the sequence $\bar{p}_l[n] = 1 - 2p_l[n]$, de-interleaved and used by the channel decoder to obtain bit estimates $\hat{b}[n]$. The symmetric noise distribution and the use of the scrambler yield the symmetries $f_\Lambda(\xi) = f_\Lambda(-\xi)$ and $f_{\Lambda|c}(\xi|c=1) = f_{\Lambda|c}(-\xi|c=0)$ for the (un)conditional LLR distribution. Hence, knowledge of $f_{\Lambda|c}(\xi|c=1)$ is sufficient for characterizing Λ .

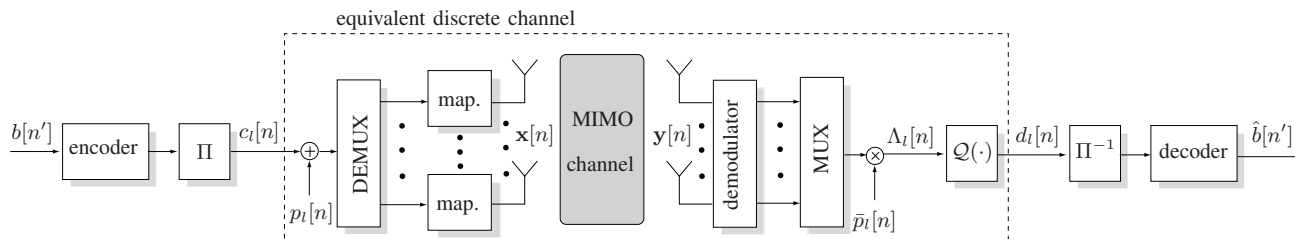


Fig. 1. Block diagram of a MIMO-BICM system.

III. LLR QUANTIZATION

The LLRs in (2) can attain any real value. We next study how to quantize these LLRs. While in practice the demodulator will directly deliver quantized LLRs, the efficient calculation of quantized LLRs is out of the scope of this paper.

We consider a q -bit quantizer characterized by $K = 2^q$ bins $\mathcal{I}_k = [i_{k-1}, i_k]$, $k = 1, \dots, K$. We use the convention $i_0 = -\infty$, $i_K = \infty$ and assume symmetric bins (this is motivated by the symmetry of the LLR distributions), with boundaries i_k sorted in ascending order. The quantizer $\mathcal{Q}(\cdot)$ maps the LLR Λ_l to a discrete LLR d_l according to

$$d_l = \mathcal{Q}(\Lambda_l) = \lambda_k \quad \text{if } \Lambda_l \in \mathcal{I}_k.$$

Here, $\lambda_k \in \mathcal{I}_k$ is the k th quantization level.

In the following, we consider the equivalent discrete channel with binary input $c \in \{0, 1\}$ and K -ary output $d \in \{\lambda_1, \dots, \lambda_K\}$. Here, c and d are obtained by randomly picking a bit position $l = 1, \dots, R_0$ according to a uniform distribution. This models a situation where the outer channel code is “blind” to the bit positions within the symbol labels. The crossover probabilities $p_{bk} = \Pr\{d = \lambda_k | c = b\} = \Pr\{\Lambda \in \mathcal{I}_k | c = b\}$ of this channel are given by

$$p_{bk} = \int_{\mathcal{I}_k} f_{\Lambda|c}(\xi|b) d\xi,$$

where $f_{\Lambda|c}(\xi|b)$ is the conditional probability density function (pdf) of the LLR Λ given that $c = b$ (averaged with respect to bit position l). Note that $\Pr\{d = \lambda_k\} = \Pr\{\Lambda \in \mathcal{I}_k\} = \frac{1}{2}(p_{0k} + p_{1k})$. The mutual information (capacity) $I = I(c; d)$ of this discrete channel is given by [4]

$$I = \frac{1}{2} \sum_{b=0}^1 \sum_{k=1}^K p_{bk} \log_2 \frac{2p_{bk}}{p_{0k} + p_{1k}}. \quad (3)$$

If the LLR distribution $f_{\Lambda|c}(\xi|b)$ and hence the transition probabilities p_{bk} are averaged with respect to the statistics of the physical channel \mathbf{H} (reflecting fast fading), the quantity I describes the ergodic rate achievable over the equivalent channel (cf. [5]). Otherwise (quasi-static fading), the transition probabilities p_{bk} , and thus the rate I , change with every realization of the channel matrix \mathbf{H} . Here, the probability

$$p_{\text{out}}(r) = \Pr\{I \leq R\}, \quad 0 \leq R \leq R_0 \quad (4)$$

characterizes the rate (denoted R) versus outage trade-off [5].

Designing the quantizer to maximize the mutual information $I(c; d)$ appears analytically infeasible in general (for BPSK

and no fading the solution is given in [2]). Hence, we propose a different approach: since $c - \Lambda - d$ is a Markov chain, the data processing inequality implies $I(c; d) \leq I(c; \Lambda)$. In order for $I(c; d)$ to be as close as possible to $I(c; \Lambda)$ (for fixed K), our proposed quantizer maximizes the mutual information $I(\Lambda; d)$. With $H(\cdot)$ denoting entropy, it follows that $I(\Lambda; d) = H(d) - H(d|\Lambda)$ and $H(d|\Lambda) = 0$ because d is a deterministic function of Λ . $H(d)$ is maximized by a uniform distribution of d and therefore, the quantizer boundaries i_k^* , $k = 1, \dots, K-1$, have to ensure that

$$\Pr\{d = \lambda_k\} = \frac{p_{0k} + p_{1k}}{2} = \frac{1}{K}, \quad k = 1, \dots, K. \quad (5)$$

Using the unconditional cumulative LLR distribution $F_\Lambda(\lambda) = \Pr\{\Lambda \leq \lambda\} = \frac{1}{2} \int_{-\infty}^{\lambda} [f_{\Lambda|c}(\xi|c=0) + f_{\Lambda|c}(\xi|c=1)] d\xi$, the optimal boundaries can be obtained by finding the arguments for which $F_\Lambda(\lambda) = k/K$, i.e.,

$$i_k^* = F_\Lambda^{-1}\left(\frac{k}{K}\right) \quad k = 1, \dots, K-1. \quad (6)$$

We note that for the capacity in (3) only the bins (boundaries) are relevant, i.e., the actual quantization levels λ_k do not influence the achievable rate. However, these values are important in order to provide the channel decoder (e.g., a belief propagation decoder) with correct reliability information [6]. In view of the equivalent discrete channel, we hence propose to choose the quantization levels as corresponding LLRs

$$\lambda_k^* = \log \frac{\Pr\{c = 1 | d = \lambda_k\}}{\Pr\{c = 0 | d = \lambda_k\}} = \log \frac{p_{1k}}{p_{0k}}. \quad (7)$$

We finally note that $\lambda_k^* \in \mathcal{I}_k$.

IV. SISO-BICM SYSTEMS WITH BPSK MODULATION

We next study in more detail the case of a SISO system ($M_T = M_R = 1$) with BPSK modulation ($R_0 = 1$ bpcu) in Rayleigh fading¹. Here, the system model (1) becomes real-valued and simplifies to $y = hx + w$, with $h \sim \mathcal{N}(0, 1)$, $w \sim \mathcal{N}(0, \sigma^2/2)$, and $x = 2c - 1 \in \{-1, 1\}$. Then, the LLR Λ in (2) equals

$$\Lambda = \frac{hy}{\sigma^2} = \frac{1}{\sigma^2} h(hx + w). \quad (8)$$

¹The results in this section also apply to the inphase and quadrature phase of SISO systems with Gray-labeled QPSK and to the two layers of BPSK-modulated 2×2 MIMO systems.

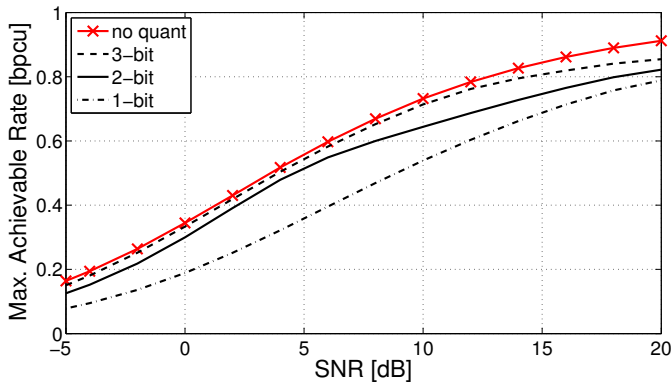


Fig. 2. Comparison of ergodic capacity for SISO-BICM with BPSK and different quantizer word-lengths.

A. Ergodic Capacity

Conditioned on $c = x = 1$, the LLR can be rewritten as $\Lambda = \frac{1}{\sigma^2} \mathbf{z}^T \mathbf{A} \mathbf{z}$, where $\mathbf{z} = (h \frac{\sqrt{2}w}{\sigma})^T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and

$$\mathbf{A} = \begin{pmatrix} 1 & \sigma/2 \\ \sigma/2 & 0 \end{pmatrix}.$$

Using the eigenvalue decomposition $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{U}^T$, with \mathbf{U} orthogonal and $\mathbf{\Sigma} = \text{diag}\{\sigma_1, \sigma_2\}$, where $\sigma_{1,2} = \frac{1 \pm \sqrt{1 + \sigma^2}}{2}$, we further obtain

$$\Lambda = \frac{1}{\sigma^2} \tilde{\mathbf{z}}^T \mathbf{\Sigma} \tilde{\mathbf{z}} = \frac{1}{\sigma^2} [\sigma_1 \tilde{z}_1^2 + \sigma_2 \tilde{z}_2^2].$$

Here, $\tilde{\mathbf{z}} = \mathbf{U}^T \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ due to the orthogonality of \mathbf{U} . Thus, Λ is a linear combination of two independent chi-square random variables with one degree of freedom. The distribution $f_{\Lambda|c}(\xi|c=1)$ can thus be shown to be given by (cf. [7])

$$f_{\Lambda|c}(\xi|c=1) = \frac{\sigma}{\pi} \exp(-\xi\sqrt{1+\sigma^2}) K_0(|\xi|), \quad (9)$$

where $K_0(\cdot)$ denotes the modified Bessel function of the second kind and order 0.

Using (9), one can determine the LLR distribution, the LLR quantization (cf. (6)), and the ergodic capacity of the equivalent channel. Numerical results for the rate in bits per channel use (bpsu) versus SNR achievable with our proposed LLR quantizers of different word-length q are shown in Fig. 2. As a reference, we also show the capacity of non-quantized max-log demodulation (labeled 'no quant'). Hard-output demodulation (i.e., 1-bit quantization) incurs a significant performance loss compared to non-quantized demodulation (more than 5 dB at rate 1/2 bpsu). With 2-bit and 3-bit LLR quantization, performance remains within 1 dB of the non-quantized case up to rates of approximately 1/2 bpsu and 3/4 bpsu, respectively.

B. Outage Capacity

Additionally conditioning on the channel coefficient h , it follows straightforwardly that $\Lambda|c \sim \mathcal{N}(x\gamma, 2\gamma)$ with $\gamma = h^2/\sigma^2$. This allows to calculate the transition probabilities of the equivalent channel as

$$p_{bk} = Q\left(\frac{i_{k-1} - (2b-1)\gamma}{\sqrt{2\gamma}}\right) - Q\left(\frac{i_k - (2b-1)\gamma}{\sqrt{2\gamma}}\right).$$

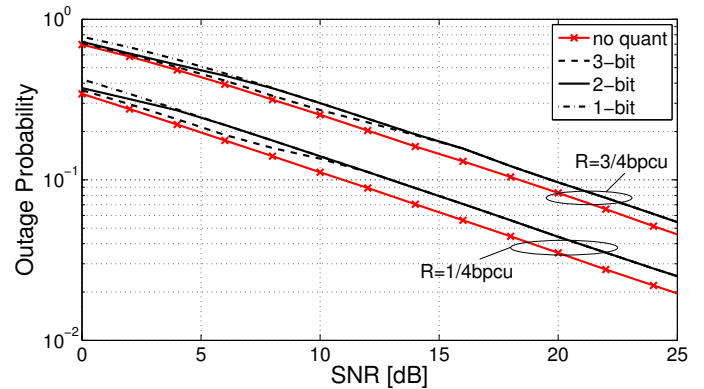


Fig. 3. Outage probability for quasi-stationary fading for SISO-BICM with BPSK for rate $R=1/4$ bpsu and $R=3/4$ bpsu.

The outage probability can thus be evaluated according to (4). Numerical results of $p_{\text{out}}(r)$ versus SNR for quasi-stationary fading with rate $R=1/4$ bpsu and with $R=3/4$ bpsu are shown in Fig. 3. LLR quantization with more than 2 bits is required to offer performance gains at medium-to-high outage probability. At high SNR, the gap between the non-quantized case and all quantized demodulators ($q=1, 2, 3$) is 2.5 dB and 1.5 dB for $R=1/4$ bpsu and $R=3/4$ bpsu, respectively. Here, $q > 3$ is required to close this gap and to reach outage probabilities close to the non-quantized case.

V. MIMO SYSTEMS AND HIGHER-ORDER MODULATION

In the following, we investigate LLR quantization for MIMO systems and higher-order constellations. Since in this case analytical expressions for the LLR distribution are hard to obtain in general, the remaining discussion is based exclusively on numerical results. For the capacity results in this section, we used empirical LLR distributions obtained from Monte-Carlo simulations to determine the bins \mathcal{I}_k such that $\Pr\{\Lambda \in \mathcal{I}_k\} = 1/K$ (cf. (5)). In the remainder of the paper, we will consider a 2×2 MIMO system with Gray-labeled 16-QAM modulation (here, $R_0 = 8$ bpsu).

A. Ergodic Capacity

We evaluated the capacity in (3) under the assumption of ergodic spatio-temporally i.i.d. fast Rayleigh fading for various quantizer word-lengths q . To this end, we estimated the transition probabilities p_{bk} by means of Monte-Carlo simulations after having determined the optimal bins based on 10^5 channel realizations. Fig. 4 shows the results obtained. It can be seen that the extreme case of 1-bit quantization yields a considerable performance loss in comparison to the non-quantized case (e.g., 3 dB SNR loss at 4 bpsu). Increasing the number of quantization levels reduces this gap significantly (to 0.5 dB and 0.1 dB for 2-bit and 3-bit quantization, respectively, at 4 bpsu). Even though at higher rates slightly increasing SNR gaps are observed, these results suggest that 3-bit LLR quantization is sufficient for practical purposes.

In order to illustrate the impact of the LLR quantizer design on ergodic capacity, we consider the same 2×2 MIMO-BICM

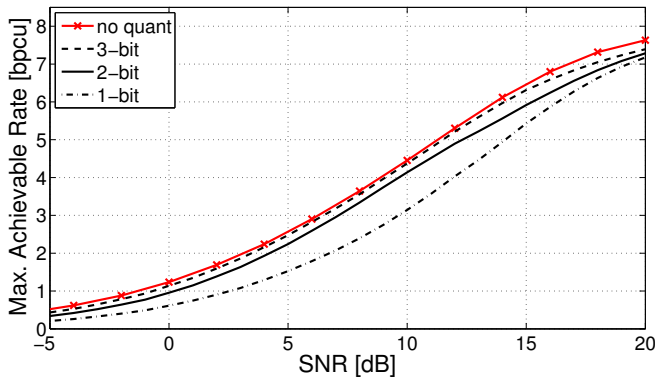


Fig. 4. Ergodic capacity of a 2×2 MIMO-BICM system with Gray-labeled 16-QAM for different LLR quantization word-lengths.

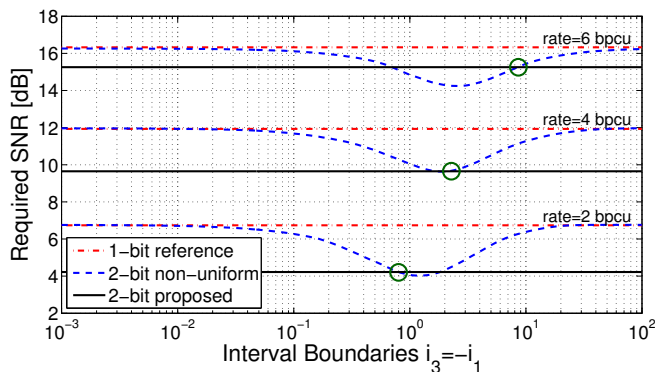


Fig. 5. SNR required for a target rate of 2, 4, 6 bpcu vs. quantizer boundary i_3 of a 2-bit LLR quantizer (2×2 MIMO-BICM system with Gray-labeled 16-QAM). Boundaries i_3 of proposed quantizer are marked by green dots.

system with Gray-labeled 16-QAM modulation for various 2-bit (i.e., 4-level) LLR quantizers. Since symmetric quantization here amounts to $i_2 = 0$ and $i_1 = -i_3$, the boundary i_3 is sufficient to index all quantizers in this case. Fig. 5 plots the SNR required to achieve target rates of 2 bpcu, 4 bpcu and 8 bpcu versus i_3 , respectively. As a reference we also show the required SNR using a 2-bit quantizer with uniform distribution and the required SNR for 1-bit quantization (hard demodulation). It can be seen that for rates of 2 bpcu and 4 bpcu the 2-bit quantizer with uniform distribution requires the same SNR as the 2-bit quantizer with optimal choice of i_3 (this is the quantizer proposed in [2]). For rates of 6 bpcu the SNR loss of the 2-bit quantizer with uniform distribution is about 1 dB compared to the optimal quantizer.

B. Outage Capacity

We next provide numerical results for the outage probability in (4) for quasi-static fading. Fig. 6 shows the outage probability $p_{\text{out}}(r)$ versus SNR for different quantizer word-lengths and $R = 2$ bpcu and $R = 6$ bpcu. From the asymptotic slopes of these curves it is seen that the diversity order equals 2 in all cases. For $R = 2$ bpcu and $R = 6$ bpcu, hard demodulation ($q = 1$) is respectively 4.8 dB and 1.8 dB away from the non-quantized case at high SNR. LLR quantization with 2 and 3 bits performs only slightly better at very low

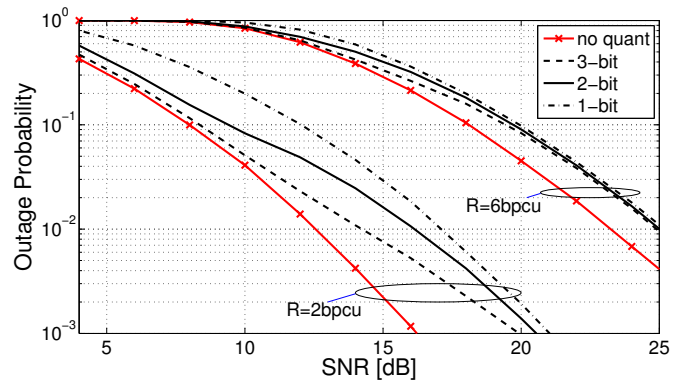


Fig. 6. Outage probability for quasi-stationary fading for 2×2 MIMO with 16-QAM using Gray labeling for rate $R = 2$ bpcu and $R = 6$ bpcu.

outage probability, but offer significant gains at medium-to-high outage probability. For $R = 2$ bpcu, the SNR loss of LLR quantization with 1, 2, and 3 bits at $p_{\text{out}} = 10^{-1}$ equals 4 dB, 1.4 dB, and 0.4 dB, respectively.

VI. ESTIMATION OF QUANTIZATION PARAMETERS

The computation of the quantization boundaries i_k^* and quantization levels λ_k^* according to (6) and (7), respectively, requires the LLR distributions $f_{\Lambda}(\xi)$ and $f_{\Lambda|c}(\xi|c)$, which in general are unknown. We thus address on-the-fly estimation of the quantization parameters. The boundaries can be estimated by using an empirical estimate of the unconditional LLR distribution $F_{\Lambda}(\xi)$, which can be obtained from a reasonable number of non-quantized LLRs.

In contrast, determining the quantization levels λ_k^* by estimating $f_{\Lambda|c}(\xi|c)$ is more difficult since the code bits are unknown at the receiver. Hence, we propose to use the following simple parametric model, which is motivated by numerical results for the 2×2 case with 16-QAM (other system parameters may require a different model):

$$f_{\Lambda|c}(\xi|c=1) = \begin{cases} \frac{\alpha\beta}{\alpha+\beta} \exp(\alpha\xi) & \xi < 0, \\ \frac{\alpha\beta}{\alpha+\beta} \exp(-\beta\xi) & \xi \geq 0. \end{cases} \quad (10)$$

To estimate the two parameters $\alpha > 0$ and $\beta > 0$, we choose two bins $\bar{\mathcal{I}}_1$ and $\bar{\mathcal{I}}_2$ and use the non-quantized LLRs Λ to obtain empirical estimates \hat{P}_i , $i = 1, 2$, of the probabilities

$$P_i(\alpha, \beta) = \Pr\{\Lambda \in \bar{\mathcal{I}}_i\} = \int_{\bar{\mathcal{I}}_i} f_{\Lambda}(\xi) d\xi,$$

with $f_{\Lambda}(\xi) = [f_{\Lambda|c}(\xi|c=0) + f_{\Lambda|c}(\xi|c=1)]/2$. The system of equations $P_i(\alpha, \beta) = \hat{P}_i$ can then be solved numerically to obtain estimates of α and β . The transition probabilities of the equivalent channel and the quantization levels are then computed based on (10) using the estimates of α and β .

VII. NUMERICAL BER RESULTS

To verify the capacity results, we performed BER simulations for SISO- and MIMO-BICM systems in ergodic Rayleigh fast fading. The channel code was a regular LDPC code² with rate 1/2 and block length 64000.

²The LDPC code was designed using the EPFL web-tool at <http://lthcwww.epfl.ch/research/ldpcopt>.

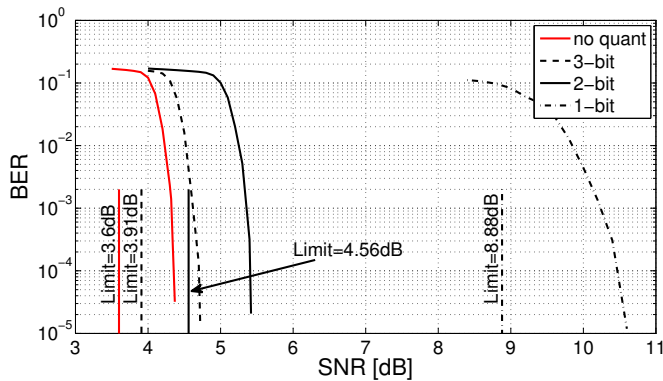


Fig. 7. BER performance for a rate-1/2 LDPC coded SISO-BICM system with BPSK and different LLR quantization word-lengths.

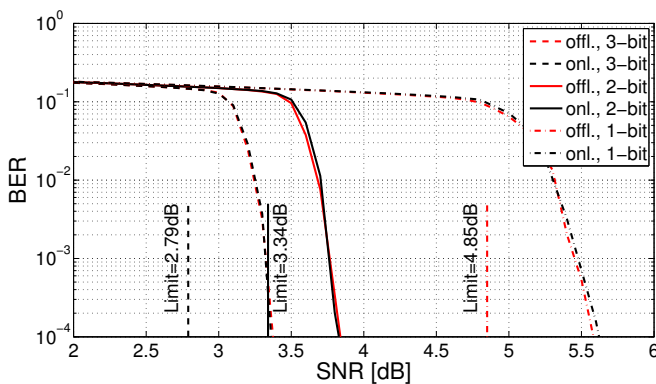


Fig. 8. BER performance for a rate-1/2 LDPC coded 2×2 MIMO system with Gray-labeled 16-QAM and different LLR quantization word-lengths.

A. SISO-BICM

We first consider a BPSK-modulated SISO-BICM system with LLR quantizers designed using the analytical results from Section IV. Fig. 7 shows the BER for our proposed LLR quantizers with different word-length together with the theoretical SNR thresholds (obtained from Fig. 2). All BER curves are reasonably close to the respective SNR thresholds (obtained from Fig. 2 and indicated by vertical lines). The gaps of 1-bit, 2-bit, and 3-bit LLR quantization to the non-quantized case respectively equal 6.2 dB, 1.1 dB, and 0.4 dB.

B. MIMO-BICM

Fig. 8 shows two (strongly overlapping) sets of BER curves for the 2×2 MIMO-BICM system with Gray-labeled 16-QAM and different LLR quantization word-lengths. One set of curves (labeled ‘offl.’) pertains to an offline design of the LLR quantizer, whereas the other set (labeled ‘onl.’) estimates the quantization parameters on-the-fly according to Section VI.

The gap to the theoretical SNR thresholds (obtained from Fig. 4 and indicated by vertical lines) equals 0.6 dB for 3-bit and 2-bit quantization and 1 dB for 1-bit quantization (hard demodulation). Furthermore, the proposed on-the-fly estimator for the LLR quantizer parameters performs extremely well in this setup (virtually indistinguishable from the offline design).

To illustrate the importance of the correct choice of the LLR quantization levels, Fig. 9 shows BER versus quantization

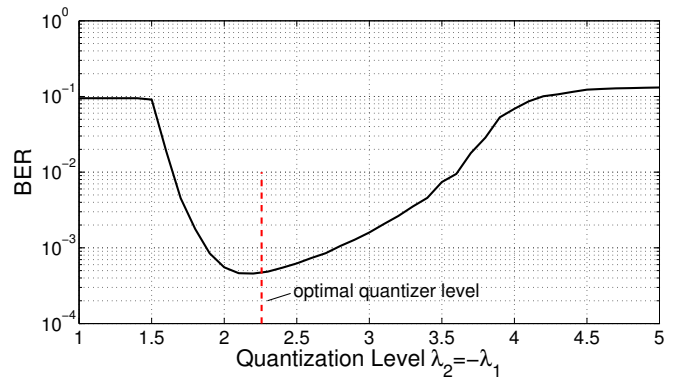


Fig. 9. BER versus quantization level for 1-bit quantization at an SNR of 12.8 dB using a rate-1/2 LDPC code (2×2 MIMO, 16-QAM, Gray labeling).

level $\lambda_2 = -\lambda_1$ for the same MIMO system as before with 1-bit LLR quantization at an SNR of 12.8 dB. Here, the optimal quantizer level $\lambda_2^* = 2.26$ (indicated by a dashed vertical line) achieves a BER of $4.5 \cdot 10^{-4}$. It is seen that the BER achieved by the belief propagation decoder is quite sensitive to the choice of λ_2 ; for $\lambda_2 \leq 1.5$ or $\lambda_2 \geq 4.3$, BER has deteriorated to about 10^{-1} (i.e., by more than 2 orders of magnitude).

VIII. CONCLUSION

We considered bit-interleaved coded modulation systems with demodulators providing quantized log-likelihood ratios (LLR). We provided design rules which lead to easily implementable LLR quantizers and studied the information rates of the equivalent discrete channel in the ergodic and outage regime. Numerical results for capacity and bit error rate showed that LLR quantization using a small number of bits is often sufficient. We also proposed simple procedures to estimate the quantizer parameters during data transmission.

ACKNOWLEDGEMENTS

The authors are grateful to Jossy Sayir for drawing their attention to the problem of quantized soft information and to Joakim Jaldén for helpful comments. This work was supported by the STREP project MASCOT (IST-026905), the Network of Excellence NEWCOM++ (IST-216715), and by the FWF Grant S10606 ‘‘Information Networks’’.

REFERENCES

- [1] A. Guillén i Fàbregas, A. Martínez, and G. Caire, ‘‘Bit-interleaved coded modulation,’’ *Foundations and Trends in Communications and Information Theory*, vol. 5, no. 1-2, pp. 1–153, 2008.
- [2] W. Rave, ‘‘Quantization of log-likelihood ratios to maximize mutual information,’’ *IEEE Signal Processing Letters*, vol. 16, pp. 283–286, Apr. 2009.
- [3] S. H. Müller-Weinfurtner, ‘‘Coding approaches for multiple antenna transmission in fast fading and OFDM,’’ *IEEE Trans. Signal Processing*, vol. 50, pp. 2442–2450, Oct. 2002.
- [4] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.
- [5] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Boston (MA): Cambridge University Press, 2005.
- [6] S. Schwandt, P. Fertl, Clemens Novak, and G. Matz, ‘‘Log-likelihood ratio clipping in MIMO-BICM systems: Information geometric analysis and impact on system capacity,’’ in *Proc. IEEE ICASSP-2009*, pp. 2433–2436, 2009.
- [7] M. K. Simon, *Probability Distributions Involving Gaussian Random Variables: A Handbook for Engineers and Scientists*. New York: Springer, 2002.