

Quantized Message Passing for LDPC Codes

Michael Meidlinger*, Alexios Balatsoukas-Stimming†, Andreas Burg†, and Gerald Matz*

*Vienna University of Technology, Austria

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

†EPFL, Switzerland

Email: {alexios.balatsoukas, andreas.burg}@epfl.ch

Abstract—We propose a quantized decoding algorithm for low-density parity-check codes where the variable node update rule of the standard min-sum algorithm is replaced with a look-up table (LUT) that is designed using an information-theoretic criterion. We show that even with message resolutions as low as 3 bits, the proposed algorithm can achieve better error rates than a floating-point min-sum decoder. Moreover, we study in detail the effect of different decoder design parameters, like the design SNR and the LUT tree structure on the performance of our decoder, and we propose some complexity reduction techniques, such as LUT re-use and message alphabet downsizing.

I. INTRODUCTION

Low-density parity-check (LDPC) codes have excellent error rate performance and can be efficiently decoded using message passing (MP) schemes like the sum-product (SP) and the min-sum (MS) algorithms, both of which involve real-valued (infinite precision) messages. By contrast, practical implementations require finite-precision message representations, i.e., the decoder messages have to be quantized and represented using typically 4 to 7 bits per message. Lower message resolutions tend to deteriorate the error rate performance of the code severely, especially in the error floor regime at high signal-to-noise ratio (SNR) [1]. Recent work on quantized MP decoders [2]–[4] has shown that a significant reduction of the message resolution is possible if the decoding algorithm is explicitly tailored to finite message alphabets.

In this paper, we present a novel “min-LUT” algorithm that replaces the variable node (VN) update of the MS algorithm with a look-up table (LUT) designed to maximize the local information flow through the code’s Tanner Graph [5]. In our previous work [3], we have shown that an actual implementation of the min-LUT decoder reduced the hardware complexity and increases the decoder throughput relative to a conventional adder-based MS implementation. This paper complements [3] by providing an in-depth discussion of the algorithm and decoder design. Specifically, we discuss in detail the symmetry requirements and the information-theoretic construction of the LUTs for the message updates. Furthermore, we examine the effects of design SNR and LUT tree structure on the error performance and develop additional complexity reduction techniques such as LUT reuse and alphabet downsizing. We demonstrate that LUT reuse is attractive for implementations and can even improve error rate performance. Finally, we show simulation results illustrating the design and performance trade-offs.

Funding by WWTF Grant ICT12-054.




Fig. 1: VN update (a) and CN update (b) for $\mathcal{N}(n) = \{m, m_1, \dots, m_{d_v-1}\}$ and $\mathcal{N}(m) = \{n, n_1, \dots, n_{d_c-1}\}$

II. LDPC CODES AND MP DECODING

A (d_v, d_c) -regular LDPC code with parity-check matrix $\mathbf{H} \in \{0, 1\}^{M \times N}$ can be represented by a Tanner graph consisting of N VNs and M check nodes (CNs). Every CN is connected to d_c VNs and every VN is connected to d_v CNs, where a connection between two nodes is indicated by a non-zero entry in the parity check matrix.

LDPC codes are traditionally decoded using MP algorithms, where messages are exchanged between VNs and CNs over the course of several decoding iterations. Let \mathcal{M}_i denote the message alphabet at iteration i . At each iteration the messages from VN n to CN m are computed using the mapping $\Phi_v^{(i)} : \mathcal{L} \times \mathcal{M}_{i-1}^{d_v-1} \rightarrow \mathcal{M}_i$, which is defined as¹

$$\mu_{n \rightarrow m} = \Phi_v^{(i)}(L_n, \bar{\mu}_{\mathcal{N}(n) \setminus m \rightarrow n}), \quad (1)$$

where $\mathcal{N}(n)$ is the set of neighbours of node n in the Tanner graph, $\bar{\mu}_{\mathcal{N}(n) \setminus m \rightarrow n} \in \mathcal{M}_{i-1}^{d_v-1}$ is a vector containing the incoming messages from all neighboring CNs except m , and $L_n \in \mathcal{L}$ denotes the channel log-likelihood ratio (LLR) at VN n . Similarly, the CN-to-VN messages are computed via the mapping $\Phi_c^{(i)} : \mathcal{M}_i^{d_c-1} \rightarrow \mathcal{M}_i$ defined as

$$\bar{\mu}_{m \rightarrow n} = \Phi_c^{(i)}(\mu_{\mathcal{N}(m) \setminus n \rightarrow m}). \quad (2)$$

Fig. 1 illustrates the message updates in the Tanner graph. The decision for a code bit c_n is computed with a mapping $\Phi_d : \mathcal{L} \times \mathcal{M}_I^{d_v} \rightarrow \{0, 1\}$ based on the incoming check node messages and the channel LLR according to

$$\hat{c}_n = \Phi_d(L_n, \bar{\mu}_{\mathcal{N}(n) \rightarrow n}). \quad (3)$$

For the MS algorithm, the mappings read

$$\Phi_v^{\text{MS}}(L, \bar{\mu}) = L + \sum_i \bar{\mu}_i, \quad (4)$$

$$\Phi_c^{\text{MS}}(\mu) = \text{sign}(\mu) \min |\mu|, \quad (5)$$

¹To simplify notation, we suppress the iteration index i in the messages.

with $\min|\boldsymbol{\mu}|$ denoting the minimum of the absolute values of the vector elements and $\text{sign}(\boldsymbol{\mu}) = \prod_j \text{sign}(\mu_j)$. The MS mappings remain unchanged for all iterations and all message alphabets are taken to be the reals, $\mathcal{M}_i = \mathcal{M} = \mathcal{L} = \mathbb{R}$. The bit decision Φ_d is based on the sign of the a-posteriori LLRs,

$$\Phi_d^{\text{MS}}(L, \bar{\boldsymbol{\mu}}) = \frac{1}{2} \left(1 - \text{sign} \left(L + \sum_i \bar{\mu}_i \right) \right). \quad (6)$$

III. THE MIN-LUT ALGORITHM

A. Basic Idea

Since floating-point arithmetic is not feasible for practical hardware implementations, the real-valued messages of the MS algorithm are usually discretized using a small number of uniformly spaced quantization levels. Together with the well-established two's complement and sign-magnitude binary encoding, the uniform quantization leads to efficient arithmetic circuits but leads to degraded error-rate performance.

Recently, efforts have been made to design decoders that explicitly account for finite message and LLR alphabets [2], [4]. Instead of arithmetic computations such as (4) and (5), the update rules for these decoders are implemented as look-up tables (LUTs). There are numerous approaches to the design of such LUTs. In the following, we present an algorithm that combines the conventional MS algorithm and the purely LUT-based approach of [4]. In this min-LUT algorithm, the VN updates are realized as LUTs, whereas the CN updates follow (5). This is motivated by the following observations:

- The CN degree is larger than the VN degree, especially for high code rates. Consequently, without further simplifications, the CN LUTs are far more complex than VN LUTs as the size of the LUTs grows exponentially in the number of inputs.
- For the MS algorithm, the VN update (4) typically increases the dynamic range of the messages whereas the CN update (5) preserves the dynamic range. Replacing the VN update (4) with a LUT eliminates the need for a message representation that can be interpreted as a numeric value. As will be explained in Section III-B, the outputs of the LUT-based VNs can be sorted in such a way that the CN update (4) can be performed based on the LUT output labels.

The LUT design for the VN updates is based on [4] and follows a density evolution (DE) approach. Given the CN-to-VN message distributions of the previous iterations, one can design the VN LUTs for each iteration in a way that maximizes the mutual information between the VN output messages and the codeword bit corresponding to the VN in question.

In order to initialize the DE procedure, we first characterize the LLR distribution at the decoder input in Section III-B. Furthermore, Section III-B discusses the relevance of symmetry conditions for the min-LUT algorithm. After these prerequisites, we present the actual evolution of message probability mass functions (PMFs) in Section III-C.

B. Channel Model and Symmetry Conditions

Throughout this paper, we focus on a binary input additive white Gaussian noise (BI-AWGN) channel $p_{y|x}(y|x)$ with noise variance σ^2 followed by a quantizer $Q_L : \mathbb{R} \rightarrow \mathcal{L}$. The quantizer uses an even number of levels $|\mathcal{L}|$ and the quantization regions are symmetric about the origin. The quantized LLRs are derived from the output of the BI-AWGN channel via $L = Q_L(-\frac{2y}{\sigma^2})$, inducing a symmetric pmf $p_{L|x}(L|x)$ that can in turn be used to define the reproducer values of the quantized LLRs as

$$L \triangleq \log \frac{p_{L|x}(L|0)}{p_{L|x}(L|1)}, \quad (7)$$

hence $p_{L|x}(-L|0) = p_{L|x}(L|1)$. Similarly, we can assign reproducer values

$$\mu \triangleq \log \frac{p_{m|x}^{(i)}(\mu|0)}{p_{m|x}^{(i)}(\mu|1)} \quad (8)$$

to the output message labels of the VN LUTs at iteration i . We again assume that the number $|\mathcal{M}|$ of messages is even. When the reproducer values are in an ascending order,

$$L_1 < L_2 < \dots < L_{|\mathcal{L}|}, \quad \mu_1 < \mu_2 < \dots < \mu_{|\mathcal{M}|}, \quad (9)$$

the identities

$$L_k \equiv -L_{|\mathcal{L}|-k+1}, \quad \mu_j \equiv -\mu_{|\mathcal{M}|-j+1}, \quad (10)$$

follow from the symmetry of $p_{L|x}(L|x)$ and the MP algorithm (cf. [6], Definition 1) and associate each label $k \in \{1, \dots, |\mathcal{L}|\}$ and $j \in \{1, \dots, |\mathcal{M}|\}$ with a sign. Based on this association and the ordering (9), the MS CN update (5) can be performed directly on the message labels; the reproducer values (7) and (8) are not needed for decoding. However, (8) bears an interesting interpretation: As the messages become more informative over the course of iterations —implying more concentrated densities $p_{m|x}^{(i)}(\mu|x)$ — the reproducer values grow in magnitude. Using different LUTs for different iterations is thus similar to using different message representations for different iterations, an approach which has already been used successfully in [1].

The symmetry of the MP algorithm discussed above is guaranteed, if the VN LUT at any iteration i satisfies

$$\Phi_v^{(i)}(-L, -\bar{\mu}_1, \dots, -\bar{\mu}_{d_v-1}) = -\Phi_v^{(i)}(L, \bar{\mu}_1, \dots, \bar{\mu}_{d_v-1}). \quad (11)$$

This identity can be reformulated based on (10) as a symmetry relation involving only labels.

Whereas our decoder design is exemplified for the BI-AWGN channel, it applies to any symmetric binary input channel followed by a symmetric quantizer. As an example, the channels characterized in [7] could be used to design decoders for bit-interleaved coded modulation (BICM) systems.

C. Density Evolution and LUT Design

In this section, we show how the message PMFs evolve over the course of iterations. We first describe how the distribution of the CN-to-VN messages can be computed based on the distribution of the incoming VN-to-CN messages. If the Tanner graph is cycle-free, then the individual input messages of a CN at iteration i are iid conditioned on the transmitted bit x , and their distribution is denoted by $p_{\mathbf{m}|x}^{(i)}(\boldsymbol{\mu}|x)$. The joint distribution of the $(d_c - 1)$ incident messages conditioned on the transmitted bit value corresponding to the recipient VN (cf. Fig. 1) reads

$$p_{\mathbf{m}|x}^{(i)}(\boldsymbol{\mu}|x) = \left(\frac{1}{2}\right)^{d_c-2} \sum_{\mathbf{x}: \bigoplus \mathbf{x} = x} \prod_{j=1}^{d_c-1} p_{\mathbf{m}|x}^{(i)}(\mu_j|x_j), \quad (12)$$

where $\bigoplus \mathbf{x}$ denotes the modulo-2 sum of the components of \mathbf{x} . Using the update rule (5), the distribution of the outgoing CN-to-VN message is then given by

$$p_{\bar{\mathbf{m}}|x}^{(i)}(\bar{\boldsymbol{\mu}}|x) = \sum_{\boldsymbol{\mu}: \text{sign}(\boldsymbol{\mu}) \min |\boldsymbol{\mu}| = \bar{\boldsymbol{\mu}}} p_{\mathbf{m}|x}^{(i)}(\boldsymbol{\mu}|x). \quad (13)$$

Let $\bar{\boldsymbol{\mu}} = (\bar{\mu}_1, \dots, \bar{\mu}_{d_v-1})$ denote the $(d_v - 1)$ incident CN-to-VN messages that are involved in the update of a certain VN. Then, the joint distribution of the VN input messages and the LLR is given by

$$p_{\mathbf{L}, \bar{\mathbf{m}}|x}^{(i)}(L, \bar{\boldsymbol{\mu}}|x) = \sum_{\mathbf{x}: x_0 = \dots = x_{d_v-1} = x} p_{\mathbf{L}|x}(L|x_0) \prod_{j=1}^{d_v-1} p_{\bar{\mathbf{m}}|x}^{(i)}(\bar{\mu}_j|x_j). \quad (14)$$

Given this distribution, we can construct an update rule Φ_v that maximizes the mutual information $I_i(\mathbf{m}; \mathbf{x})$ between \mathbf{m} and \mathbf{x} :

$$\Phi_v^{(i)} = \arg \max_{\Phi} I_i(\mathbf{m}; \mathbf{x}) = \arg \max_{\Phi} I_i(\Phi(\mathbf{L}, \bar{\mathbf{m}}); \mathbf{x}). \quad (15)$$

Here, the maximization is over all deterministic mappings Φ in the form of (1) that respect the symmetry condition (11). Hence, the resulting update rule $\Phi_v^{(i)}$ maximizes the local information flow between the CNs and the VNs. An algorithm that solves (15) with complexity $O(|\mathcal{L}|^3 |\mathcal{M}|^{3(d_v-1)})$ was provided in [5]. Using the update rule (15), we can compute the conditional distribution of the messages in the next iteration:

$$p_{\mathbf{m}|x}^{(i+1)}(\boldsymbol{\mu}|x) = \sum_{(L, \bar{\boldsymbol{\mu}}): \Phi_v^{(i)}(L, \bar{\boldsymbol{\mu}}) = \boldsymbol{\mu}} p_{\mathbf{L}, \bar{\mathbf{m}}|x}^{(i)}(L, \bar{\boldsymbol{\mu}}|x). \quad (16)$$

The noise threshold σ^* of a (d_v, d_c) -regular LDPC code ensemble with at most I decoding iterations is defined as

$$\sigma^* = \sup\{\sigma \geq 0: I_i(\mathbf{m}; \mathbf{x}) > 1 - \epsilon \text{ for some } i \leq I\}. \quad (17)$$

Algorithm 1 summarizes the individual steps of a bisection algorithm that uses the DE algorithm to calculate σ^* .

Algorithm 1 Density Evolution based LUT design

Input: Search interval $[\sigma_{\min}, \sigma_{\max}]$, precision $\Delta\sigma > 0$, ϵ , I

```

1: while  $\sigma_{\max} - \sigma_{\min} > \Delta\sigma$  do
2:    $\sigma \leftarrow (\sigma_{\max} + \sigma_{\min})/2$ 
3:   Get  $p_{\mathbf{L}|x}(L|x)$  corresponding to BI-AWGN( $\sigma^2$ )
4:   achievable  $\leftarrow$  false
5:   for  $i = 1, \dots, I$  do
6:     Update CN-to-VN distribution (12) and (13)
7:     Build the product distribution (14)
8:     Design LUT update  $\Phi_v^{(i)}$  (15)
9:     Update VN-to-CN distribution (16)
10:    if  $I(\mathbf{m}^{(i)}; \mathbf{x}) > 1 - \epsilon$  then
11:      achievable  $\leftarrow$  true
12:    break
13:  end if
14: end for
15: if achievable then
16:    $\sigma_{\min} \leftarrow \sigma$ 
17: else
18:    $\sigma_{\max} \leftarrow \sigma$ 
19: end if
20: end while
21:  $\sigma^* \leftarrow \sigma$ 

```

Output: Threshold σ^* , LUT sequence $\Phi_v^{(1)}, \dots, \Phi_v^{(I)}$

IV. DESIGN AND PERFORMANCE TRADE-OFFS FOR PRACTICAL DECODERS

Algorithm 1 is well suited to determine the asymptotic performance of the min-LUT algorithm for large block length N and many decoding iterations (large I). In order to design practical min-LUT decoders with N and I not too large, we propose the following approach:

- 1) Choose a practical number of maximum iterations I .
- 2) Define a reuse pattern $\mathcal{I} = \{i_1, \dots, i_r\} \subseteq \{1, \dots, I\}$.
- 3) Choose a LUT tree structure, cf. Section IV-B
- 4) Choose a design SNR γ such that the corresponding noise level σ is below the threshold σ^* .
- 5) For the chosen σ , run the inner loop of Algorithm 1, (lines 3 to 14). However, only design a new LUT if $i \in \mathcal{I}$. If $i \notin \mathcal{I}$, reuse the LUT from the previous iteration.
- 6) Check the performance of the results by error rate simulations; possibly repeat the procedure with adjusted parameters.

The resulting LUTs can be used to synthesize a decoder that outperforms a conventional MS decoder in error rate performance, throughput, and hardware complexity [3]. Since for the above procedure there are several design parameters to be chosen, we next give an overview of the performance impact of each of the individual parameters. We support our discussion with comprehensive simulation results that illustrate the design and performance trade-offs. All simulations have been conducted using the (6, 32)-regular LDPC code (block length $N = 2048$, rate $R = 13/16$) defined for the 10 Gbit/s Ethernet standard [8].

A. Design SNR

The information-theoretic LUT design depends strongly on the initial LLR distribution and thus in turn on the design SNR and the LLR quantizer. Our simulations indicate that even though the min-LUT decoder is designed for one particular SNR, excellent performance is maintained over a wide range of actual noise levels, cf. Fig. 3. Re-designing the LLR quantizer or the entire decoder for different SNRs would further improve the performance but simultaneously would substantially increase the implementation cost. For this reason, we kept both the LLR quantizer and the decoder fixed over the range of simulated SNRs.

We next discuss how the choice of the design SNR affects decoder performance. As can be seen in Fig. 3, by increasing the design SNR, we can trade off performance in the waterfall region against performance in the error floor region. The interpretation is straight-forward: decoders that are designed for bad channels work better for bad channels and vice versa. Another interpretation can be found in terms of decoding iterations: a lower design SNR implies that the decoder is operating closer to the DE threshold and thus DE convergence is much slower as compared to the case of a higher design SNR well beyond the threshold. If, however, decoders designed for low and high γ use the same number of iterations, the lack of convergence translates into a higher residual error for low design SNRs.

B. LUT Reuse and Alphabet-Downsizing

Algorithm 1 produces a distinct VN LUT for each iteration. While this does not affect silicon complexity for an unrolled decoder architecture, non-unrolled decoders would need to implement multiple LUTs for the VNs. Contrary to our expectations, we found in our simulations that reusing LUTs for multiple iterations does not necessarily degrade the performance and can even lead to an improvement. As an example, Fig. 4 shows that with a reuse pattern $\mathcal{I} = \{1, 5\}$ with only $r = 2$ different LUTs, we can improve the error rate compared to a decoder that uses distinct LUTs for every iteration. An explanation for this effect is still an open issue to be explored. At this point, we can only conjecture that the effect originates from the overly optimistic message distributions of DE, which tends to overestimate the speed of convergence for practical codes that are not cycle-free.

Another means of reducing LUT complexity is message downsizing, i.e., reducing the size of the message set,

$$|\mathcal{M}_{i'}| \leq |\mathcal{M}_i| \quad \text{for } i' > i.$$

The idea here is that the messages undergo a gradual hardening while being passed through the decoder before culminating into the binary-output decision mapping (3). As can be observed in Fig. 4, a decoder with down-sized LUTs using decaying message resolutions from 3 to 2 to 1 bits over the range of $I = 8$ iterations performs only slightly worse than a comparable min-LUT decoder with fixed resolution of 3 bits. LUT reuse and LUT downsizing cannot be combined

arbitrarily, i.e., reducing the message resolution in a certain iteration prevents reuse of the corresponding LUT.

C. LUT Tree Structure

Since the number of input configurations for the VN update $\Phi_v^{(i)}$ equals $|\mathcal{L}||\mathcal{M}_i|^{d_v-1}$, a full-fledged LUT would be prohibitively complex for codes with high VN degree d_v . A similar problem occurs with the decision LUTs (3). To overcome this limitation, we restrict ourselves to nested update rules, e.g., for $d_v = 6$ a possible nesting could take the form

$$\Phi(\bar{\mu}_1, \dots, \bar{\mu}_5) = \Phi_1(\Phi_2(\bar{\mu}_1, \bar{\mu}_2, \bar{\mu}_3), \Phi_3(\bar{\mu}_4, \bar{\mu}_5), L).$$

Obviously, any such nesting can be represented graphically by a directed tree, cf. Fig. 2, tree T_2 for this particular example. Since we assume iid messages, the ordering of the arguments in the nesting is immaterial and we consider nesting that differ only in the ordering as equivalent.

While the nested structure clearly reduces complexity, it is not clear a priori, which tree structure to prefer over another. In what follows, we provide guidelines on how to choose the tree structure based on information-theoretic arguments and a heuristic metric. For the moment, we do not distinguish between messages $\bar{\mu}$ and channel input L ; the discussion of the location of L within the tree is deferred to Section IV-C3.

1) *Partial ordering*: Let the tree T_1 represent a specific nesting and let T_2 be a refinement² of T_1 . Furthermore, let \mathcal{Q}_j denote the set of all LUTs that respect the nesting induced by some tree T_j . By construction, any LUT in \mathcal{Q}_2 also conforms with the nesting associated with T_1 . Thus, $\mathcal{Q}_1 \supseteq \mathcal{Q}_2$ and

$$\max_{\Phi \in \mathcal{Q}_1} I_i(\Phi(L, \bar{\mathbf{m}}); \mathbf{x}) \geq \max_{\Phi \in \mathcal{Q}_2} I_i(\Phi(L, \bar{\mathbf{m}}); \mathbf{x}).$$

Consequently, tree refinement defines a partial ordering $\geq_{\mathcal{T}}$, effectively inducing a hierarchy in terms of maximum information flow. However, since the totality axiom is not fulfilled, not all tree structures can be compared in terms of the relation $\geq_{\mathcal{T}}$, cf. Fig. 2.

2) *A heuristic metric*: The data processing inequality states that processing can only reduce mutual information. Therefore, for maximum information flow the paths from the input leaves to the root output should be as short as possible. We thus define the cumulative depth $\lambda(T)$ of a tree T as the sum of distances of all leaf nodes to the root node. DE simulations confirmed that cumulative depth is useful in ranking tree structures. Table I shows how a larger λ corresponds with a lower DE threshold. However, the threshold differences are small and our simulations have shown that all the trees presented here perform similar in terms of error rate. While there were small differences conforming with the ordering discussed above, they are not significant enough to serve as a basis for choosing the tree structure. Rather, we recommend choosing the tree based on its silicon complexity. Trees that are close to complete binary trees are preferable because they have short critical paths with low complexity LUTs and at the same time have small cumulative depth λ .

²Graphically, a refinement of nesting corresponds to the placement of new nodes between parent and child nodes.




Fig. 2: Six different LUT tree structures. Note that $T_1 \geq_{\mathcal{T}} T_4 \geq_{\mathcal{T}} T_6$, $T_2 \geq_{\mathcal{T}} T_5$, $T_3 \geq_{\mathcal{T}} T_5$, and $T_3 \geq_{\mathcal{T}} T_6$. However, we cannot compare T_2 with T_3 or T_5 with T_6 using the relation $\geq_{\mathcal{T}}$.

TABLE I: Comparison of cumulative depth and DE threshold for various tree structures (cf. Fig. 2). Here, all LUTs had a resolution of 3 bit.

T	T_1	T_2	T_3	T_4	T_5	T_6
λ	10	11	11	14	16	19
σ^*	0.5330	0.5328	0.5327	0.5313	0.5309	0.5305

3) *Position of the channel LLR:* The mutual information between the CN-to-VN messages and the coded bits is initially zero and increases over the course of iterations until at some iteration $I_{i'}(\bar{m}; x) \geq I(L; x)$. Using a similar argument as before, we can conclude that until iteration i' the channel LLR should be placed close to the root node to ensure a large information flow. After iteration i' , the CN-to-VN messages tend to carry more information than the channel LLR and thus should be placed closer to the leaf nodes. Our simulations show that this strategy indeed provides the best FER performance; however, the loss as compared to the case where the channel LLR stays at the root node is only relevant for a large number of iterations ($I > 20$).

D. Comparison with MS

As can be seen in Fig. 4, the min-LUT decoder with a message resolution of 3 bits outperforms a conventional MS decoder using a message resolution of 4 bits by a significant margin and even beats a floating point MS decoder. The gain is even larger for the case of LUT reuse. We conclude that our min-LUT decoder is an attractive alternative to the conventional MS decoder.

V. CONCLUSION

In this paper, we presented the min-LUT algorithm for decoding LDPC codes. Contrary to the min-sum algorithm, the min-LUT decoder is custom-designed to work with discrete messages of very low resolution. Hence, it constitutes an attractive choice for practical hardware implementations. Using the 10 Gbit/s Ethernet code, we furthermore exemplified that the min-LUT error rate performance can be superior to min-sum decoding in spite of small message resolutions.




Fig. 3: FER versus channel SNR for min-LUT decoder at different design SNRs ($I = 8$, LUT tree T_1 , message resolution Q_μ down-sized from 3 to 2 bit, $Q_L = 3$ bit per channel LLR).




Fig. 4: Performance comparisons for different decoders ($I = 8$, LUT tree T_6 , $Q_L = 4$, $\gamma = 4.2$ dB).

REFERENCES

- [1] X. Zhang and P. Siegel, "Quantized Iterative Message Passing Decoders with Low Error Floor for LDPC Codes," *IEEE Transactions on Communications*, vol. 62, no. 1, pp. 1–14, Jan. 2014.
- [2] S. Planjery, D. Declercq, L. Danjean, and B. Vasic, "Finit Alphabet Iterative Decoders – Part I: Decoding Beyond Belief Propagation on the Binary Symmetric Channel," *IEEE Transactions on Communications*, Oct. 2013.
- [3] A. Balatsoukas-Stimming, M. Meidlinger, R. Ghanaatian, G. Matz, and A. Burg, "A Fully Unrolled LDPC Decoder based on Quantized Message Passing," in *IEEE Workshop on Signal Processing Systems (SiPS)*, Hangzhou, China, Oct. 2015.
- [4] B. Kurkoski, K. Yamaguchi, and K. Kobayashi, "Noise thresholds for discrete LDPC decoding mappings," in *Proc. IEEE Global Telecommunications Conf. (GLOBECOM)*, Nov. 2008.
- [5] B. Kurkoski and H. Yagi, "Quantization of binary-input discrete memoryless channels," *IEEE Transactions on Information Theory*, vol. 60, no. 8, pp. 4544–4552, Aug. 2014.
- [6] T. Richardson and R. Urbanke, "The capacity of low-density parity-check codes under message-passing decoding," *IEEE Transactions on Information Theory*, vol. 47, no. 2, pp. 599–618, Feb. 2001.
- [7] A. Alvarado, L. Szczecinski, R. Feick, and L. Ahumada, "Distribution of L-values in gray-mapped M2-QAM: closed-form approximations and applications," *IEEE Transactions on Communications*, vol. 57, no. 7, pp. 2071–2079, Jul. 2009.
- [8] "IEEE Standard for Information Technology – Telecommunications and Information Exchange between Systems – Local and Metropolitan Area Networks – Specific Requirements Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications," IEEE Std. 802.3an, Sep. 2006.