An Efficient Fixed-Complexity Sphere Decoder with Quantized Soft Outputs

Sandra Roger, *Student Member, IEEE,* Alberto Gonzalez, *Senior Member, IEEE,* Vicenc Almenar, *Member, IEEE,* and Gerald Matz, *Senior Member, IEEE*

***Abstract*—In practical multiple-input multiple-output bit-interleaved-coded-modulation (MIMO-BICM) systems, demodulators have to deliver finite word-length (quantized) log-likelihood-ratios (LLRs). In this letter, we propose an efficient modification of the fixed-complexity sphere decoder for MIMO-BICM systems working with quantized LLRs. Our approach reduces the complexity of previously proposed schemes via pruning strategies that exploit the clipping and quantization of LLRs. Numerical results confirm that our scheme achieves a significant complexity reduction (by 37% for the case of 2 bits per LLR and by 31% for the case of 3 bits per LLR) with negligible degradation in bit error rate performance.**

***Index Terms*—MIMO-BICM, FSD, LLR quantization.**

I. Introduction

In multiple-input multiple-output (MIMO) systems with bit-interleaved-coded-modulation (BICM), the receiver consists of a demodulator (or soft detector) and a subsequent channel decoder. The demodulator provides reliability information (soft outputs) about the transmitted coded bits in the form of real-valued log-likelihood ratios (LLRs). Next, these values are used by the channel decoder to obtain final decisions on the transmitted code bits. Many demodulators can be found in the literature, some of them based on sphere-decoding (SD). Examples include the *list-based SD* (LSD) scheme [1], the *single tree-search* (STS) algorithm [2], the *soft-output fixed-complexity SD* (SFSD) [3] and the *smart ordering and candidate adding* (SOCA) algorithms [4][5]. These methods offer a trade-off between performance and complexity (related to the size of the candidate list). Both the SFSD and SOCA algorithms have lower complexity than the LSD and are attractive because, unlike the STS, they exhibit fixed complexity.

In actual hardware implementations, LLRs have to be represented with a finite word length. LLR quantization has been studied in [6], which provides guidelines for the quantizer design, i.e., number of bits, quantization intervals, and quantization levels.

In this paper, we propose an efficient SFSD scheme for MIMO-BICM systems working with quantized LLRs. Our approach uses the fact that the quantized LLRs belong to an a priori known finite set (quantization levels) in order to avoid some of the computations carried out by the SFSD. In addition, LLR clipping is included to save complexity when a certain LLR value exceeds a previously selected clipping threshold. We also propose to make use of the quantizer parameters to set the clipping threshold. We finally provide numerical evidence illustrating the performance-complexity trade-off of the proposed SFSD scheme.

Manuscript received July 30, 2012. The associate editor coordinating the review of this letter and approving it for publication was M. Lentmaier.

S. Roger, A. Gonzalez, and V. Almenar are with the Institute of Telecommunications and Multimedia Applications, Universitat Politècnica de València, Spain (e-mail: sanrova@iteam.upv.es).

G. Matz is with the Institute of Telecommunications, Vienna University of Technology, Vienna, Austria.

This work was supported in part by the PROMETEO/2009/013 and TEC2009-13741 projects, by the AP2007-01417 FPU grant, and by FWF grant S10606.

Digital Object Identifier 10.1109/LCOMM.2012.091712.121702

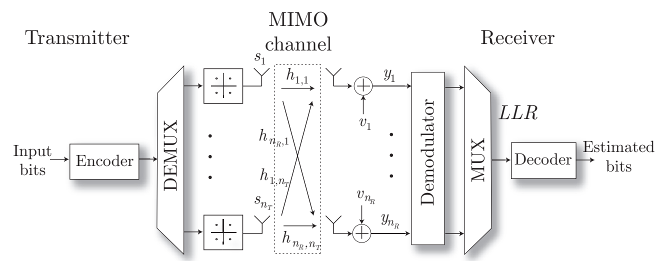

Fig. 1. Block diagram of a MIMO-BICM system.

II. Fixed-Complexity Soft-Output Detection

Let us consider a MIMO-BICM system with n_{T} transmit antennas and n_{R} receive antennas (see Fig. 1). We assume $n_{\mathrm{R}} \geq n_{\mathrm{T}}$.

In this system, the sequence of information bits is encoded using an error-correcting code and passed through a bitwise interleaver before being demultiplexed and mapped to complex-valued transmit symbol vector $\mathbf{s} = (s_1, \ldots, s_{n_{\mathrm{T}}})^T$. The symbols s_i are taken from a constellation Ω of size $|\Omega| = M$ and hence carry $\log_2 M$ code bits each. The baseband equivalent model for received vector is given by

$$\mathbf{y} = \mathbf{H}\mathbf{s} + \mathbf{v}. \tag{1}$$

Here, $\mathbf{H}$ is an $n_{\mathrm{R}} \times n_{\mathrm{T}}$ matrix modeling the Rayleigh fading MIMO channel and $\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \sigma^2\mathbf{I})$ is circularly symmetric complex Gaussian noise. We assume that the channel $\mathbf{H}$ is known at the receiver.

At the receive side, the demodulator uses the model (1) to compute soft information about the code bits in terms of LLRs. Using the max-log approximation [1], the LLR of the bth bit of the symbol in layer j equals

$$L_{j,b} = \frac{1}{\sigma^2}\left[\min_{\mathbf{s}\in\mathcal{X}_{j,b}^{(0)}} \|\mathbf{y}-\mathbf{H}\mathbf{s}\|^2 - \min_{\mathbf{s}\in\mathcal{X}_{j,b}^{(1)}} \|\mathbf{y}-\mathbf{H}\mathbf{s}\|^2\right], \tag{2}$$

where $x_{j,b}$ denotes the bth bit in the bit label of symbol s_j. Furthermore, $\mathcal{X}_{j,b}^{(c)}$ denotes the set of symbol vectors for which

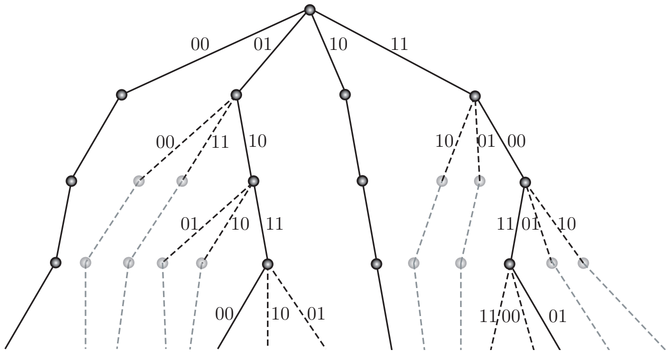

Fig. 2. Decoding tree of the SFSD algorithm for a 4×4 MIMO system with QPSK symbols and $N_{\mathrm{iter}} = 2$.

the bth bit in layer j equals c. The hard-output *maximum-likelihood* (ML) detection problem can be shown to provide one of the two minima in (2), i.e.,

$$\mathbf{s}^{\mathrm{ML}} = \arg\min_{\mathbf{s}\in\Omega^{n_{\mathrm{T}}}} \|\mathbf{y} - \mathbf{H}\mathbf{s}\|^2, \quad d^{\mathrm{ML}} = \|\mathbf{y} - \mathbf{H}\mathbf{s}^{\mathrm{ML}}\|^2. \quad (3)$$

Denote by $x_{j,b}^{\mathrm{ML}}$ the bth bit associated with s_j^{ML}. For each j and b, the second minimum in (2) can be computed as

$$\bar{d}_{j,b} = \min_{\mathbf{s}\in\mathcal{X}_{j,b}^{(\overline{x}_{j,b}^{\mathrm{ML}})}} \|\mathbf{y} - \mathbf{H}\mathbf{s}\|^2, \quad (4)$$

where $\overline{x}$ denotes the complement of bit x. Note that $\mathbf{s} \in \mathcal{X}_{j,b}^{(\overline{x}_{j,b}^{\mathrm{ML}})}$ represents the counter-hypothesis to the ML solution for bit b in layer j. Once (3) and (4) have been calculated, the LLRs are obtained as

$$L_{j,b} = \frac{1}{\sigma^2}\left(d^{\mathrm{ML}} - \bar{d}_{j,b}\right)\left(1 - 2x_{j,b}^{\mathrm{ML}}\right), \quad (5)$$

where the term at the end adjusts the sign depending on whether d^{ML} corresponds to the first or the second minimum in (2).

There are several methods to avoid an exhaustive search over the n_{T}-dimensional set $\Omega^{n_{\mathrm{T}}}$ in the ML problem (3). An approximate solution method that is interesting in our context is the *fixed-complexity sphere decoder* (FSD) [7]. The FSD performs a predetermined tree search composed of two different stages: a full search in the first (highest) T levels of the tree and a single-path search in the remaining tree levels using succesive-interference-cancellation (SIC). The symbols are detected following a specific ordering described in [7]. The SFSD [3] extends the FSD tree search in order to obtain the minimum distances in (4) after (3) has been solved. Fig. 2 shows an example search-tree of the SFSD for $n_{\mathrm{T}} = 4$, QPSK symbols, and $T = 1$. Note that the SFSD starts from the candidate list obtained by the hard-output FSD (4 SIC branches of 4 symbols each) and adds new candidates for the counter-hypotheses. All the necessary values to compute the LLRs of the symbol bits in the first full expanded level are already available. To begin the list extension, the best N_{iter} paths are selected from the initial hard-output FSD list (in this example, $N_{\mathrm{iter}} = 2$). This is motivated by the heuristics that the lowest-distance paths may be candidates differing from the best paths in only a few bits. The symbols belonging to these N_{iter} paths are picked up from the root up to a certain level l; at level $l-1$, additional $\log_2 M$ branches are explored, each of them having one of the bits of the initial path symbol negated. Afterwards, these new partial paths are completed as done in the hard-output FSD scheme. The same operation is repeated until the lowest level of the tree is reached. The values that achieved almost max-log performance for a 4×4 system in [3] were $N_{\mathrm{iter}} = 2, 4, 6$ for QPSK, 16-QAM, and 64-QAM, respectively.

The SOCA algorithm [4] combines a smart-ordered QR decomposition with smart candidate adding and parallel layer-by-layer search in the tree. It also includes optional sort-and-select pruning as the K-Best algorithm defined in a vector $\mathbf{m}$. This method exhibits a meaningful performance-complexity tradeoff based on the selection of vector $\mathbf{b} = (b_1, b_2, \ldots, b_{n_T})^T$, which sets the number of nodes expanded per survivor path. Note that, unlike in the SFSD, there are no full tree-paths available until the algorithm ends, due to the layer-by-layer tree-expansion.

III. Proposed SFSD with Quantized Outputs

Quantizer designs for the LLRs delivered by the demodulator of a MIMO-BICM system have been introduced in [6]. A q-bit quantizer uses $K = 2^q$ quantization intervals $\{I_1, I_2, ..., I_K\}$. These intervals are specified by $I_k = [i_{k-1}, i_k]$ where $\{i_0, i_1, i_2, ..., i_K\}$ are the quantizer thresholds ($i_0 = -\infty$ and $i_K = \infty$ by convention). The quantized LLR value $\Lambda_{j,b}$ is given by

$$\Lambda_{j,b} = \mathcal{Q}(L_{j,b}) = \lambda_k, \quad \text{if} \quad L_{j,b} \in I_k, \quad (6)$$

where λ_k is the kth quantization level. For fixed q, the ideal quantizer maximizes the mutual information between the input code bits and the quantized LLRs. Since the design of this ideal quantizer is difficult in practice, [6] proposed an alternative quantizer design that maximizes the mutual information between the LLRs before and after quantization [6], i.e.,

$$\{i_k^{\mathrm{opt}}\}_{k=1}^{K-1} = \arg\max_{\{i_k\}_{k=1}^{K-1}} I\left(L_{j,b}; \Lambda_{j,b}\right). \quad (7)$$

Following [6], we restrict to even K and symmetric quantizers in which case $\{i_{K/2}^{\mathrm{opt}}\} = 0$ and $i_k = -i_{K-k+1}$.

A. Quantization-based Pruning

If the MIMO-BICM system uses quantized LLRs, the final demodulator output is known in advance to belong to the finite set $\{\lambda_1, \ldots, \lambda_K\}$. This fact allows us to reduce the size of the candidate list used by the SFSD to refine the initial LLRs.

As described in Section II, once d^{ML} has been found at the hard-output FSD stage, the SFSD adds $\log_2 M$ new tree paths to the candidate list. Slightly abusing notation, we denote by $\bar{d}_{j,b}$ the current estimate of the minimum distance associated with the counter-hypothesis for bit b at level j. Every new path starting from level j is intended to update $\bar{d}_{j,b}$ by negating the bth bit. In fact, any refinement of the distances in the candidate list can only modify $\bar{d}_{j,b}$ by a smaller value than the current one. The proposed approach decides whether a certain $\bar{d}_{j,b}$ needs to be updated or not by testing whether the magnitude of the associated LLR satisfies the following condition:

$$|L_{j,b}| = \frac{1}{\sigma^2}\left|d^{\mathrm{ML}} - \bar{d}_{j,b}\right| \leq i_{K/2+1}^{\mathrm{opt}}. \quad (8)$$

This condition holds for those LLR values that lie within either $I_{K/2} = [i_{K/2-1}, i_K/2]$ or $I_{K/2+1} = [i_{K/2}, i_K/2+1]$. Since d^{ML} and $\bar{d}_{j,b}$ both are greater than zero and $d^{\mathrm{ML}} \leq \bar{d}_{j,b}$ for all $\{j,b\}$, (8) is equivalent to

$$\bar{d}_{j,b} \leq (d^{\mathrm{ML}} + \sigma^2 i^{\mathrm{opt}}_{K/2+1}). \tag{9}$$

This condition can be easily tested in the SFSD path extension stage to avoid unnecessary updates of $\bar{d}_{j,b}$. However, the same idea is not straightforwardly applicable to the LSD, STS and SOCA methods since their candidate lists are not built as an extension of an initial list. Thus, their final d^{ML} value is not known until the algorithm ends and using directly the same prunning as in (9) may lead to detection errors. Therefore other approaches should be investigated to exploit quantization-based pruning.

B. Clipping-based Pruning

In order to further reduce the number of path extensions in the SFSD, we propose to combine the previously described pruning with LLR clipping. The use of LLR clipping has been widely adopted as a mechanism to reduce the complexity of tree-search detectors [1], [2]. LLR clipping imposes a constraint on the dynamic range of the LLRs to enable fixed-point implementations:

$$L_{j,b} = \begin{cases} L_{j,b}, & |L_{j,b}| \leq L_{\mathrm{clip}}, \\ L_{\max}\,\mathrm{sign}(L_{j,b}), & \text{else}. \end{cases} \tag{10}$$

Here, L_{clip} is the LLR clipping threshold and $L_{\max}$ is the LLR value given to all the LLRs greater than L_{clip}. Similar to (9), the clipping can be exploited in the SFSD search to avoid unnecessary updates of $\bar{d}_{j,b}$ whenever the current LLR estimate $L_{j,b}$ satisfies $|L_{j,b}| > L_{\mathrm{clip}}$, which is equivalent to

$$\bar{d}_{j,b} > (d^{\mathrm{ML}} + \sigma^2 L_{\mathrm{clip}}). \tag{11}$$

Note that, since we are working with quantized LLRs, both L_{clip} and $L_{\max}$ are determined by the optimization given in (7), i.e. $L_{\mathrm{clip}} = i^{\mathrm{opt}}_{K-1} = -i^{\mathrm{opt}}_1$ and $L_{\max} = \lambda_K = -\lambda_1$. This approach gives lower values than the traditional uniform quantization solution [1] increasing in this way the number of pruned paths. This clipping threshold could be also applied to other demodulators such as the LSD, STS and SOCA methods considering the temporary ML candidate instead of the final one, as done in [2].

IV. Results

We consider a 4×4 MIMO-BICM system with a 16-QAM symbol alphabet. The channel code is a parallel concatenated turbo code with rate 1/2 and a block length of 6144 bits. The MIMO channel was i.i.d. block Rayleigh fading (staying constant for 16 MIMO symbols). Fig. 3 shows the bit error rate (BER) versus signal-to-noise ratio (SNR) for max-log demodulation, the plain SFSD, SFSD with clipping only, and SFSD with 2-bit and 3-bit LLR quantization. The performance of the SOCA method is included as a benchmark, where the method is configured to have the same complexity as the SFSD ($\mathbf{b} = (16,4,4,4)^T$). For the demodulator using LLR clipping only, we chose $L_{clip} = 8$ as in [1]. For the quantization-based demodulator, the LLR quantizer was designed as proposed in [6]. It can be observed that the plain SFSD achieves max-log performance and the SOCA method nearly reaches it. With 3-bit quantization, the BER degradation compared to the non-quantized case is negligible ($\simeq 0.1$dB). SFSD with 2-bit quantization performs slightly worse ($\simeq 0.5$dB).

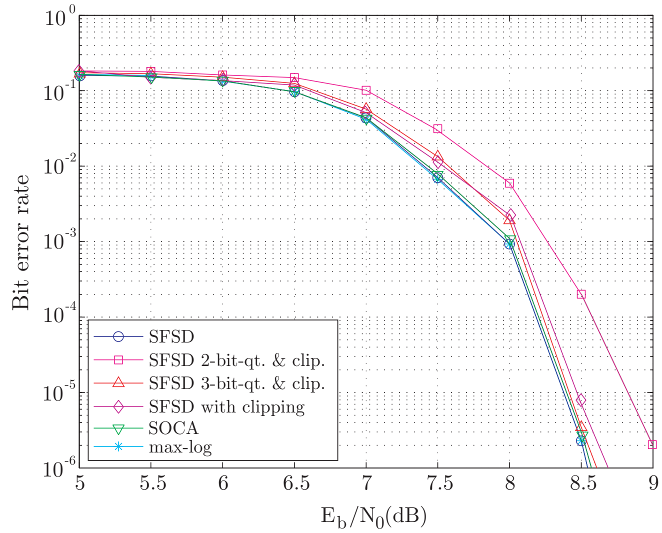

Fig. 3. BER curves for different SFSD schemes and for SOCA and max-log demodulation, all with a rate-1/2 turbo code in a 4×4 MIMO-BICM system using 16-QAM.

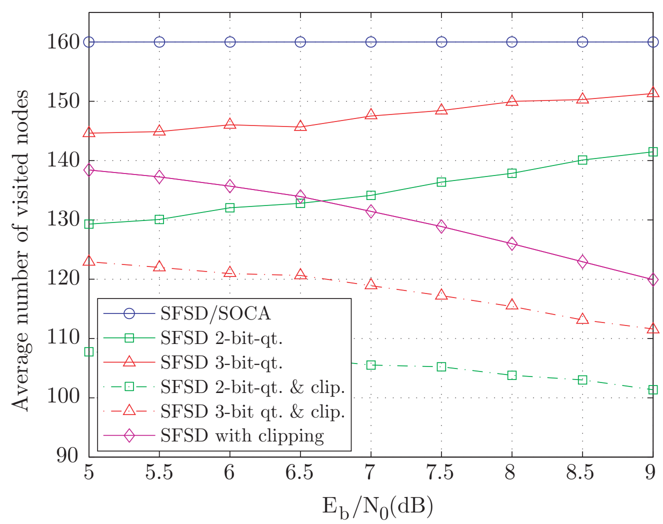

Fig. 4. Average number of nodes visited by the different SFSD schemes in a 4×4 MIMO-BICM system using 16-QAM.

Next, we study the complexity savings in terms of the average number of visited nodes, which is directly related to the number of partial distance calculations in the tree. Fig. 4 shows the average number of visited nodes for the SFSD scheme with quantization-based pruning only and with quantization- and clipping-based pruning. The complexity of plain SFSD detection without pruning and SOCA detection are included in the figure as a baseline reference. The quantization-based pruning is more significant at lower SNR whereas the complexity reduction due to clipping-based pruning increases as the SNR grows. Moreover, as expected, 2-bit quantization results in larger complexity savings than 3-bit quantization. Clearly, the complexity savings achieved by the combined pruning strategy are much larger than those achieved by either of the two pruning techniques alone, with complexity reductions of 37% and 31% for 2-bit- and 3-bit quantization, respectively, at an SNR of $9\,\mathrm{dB}$. The BER of the SOCA demodulators that have similar complexity than the proposed schemes was also evaluated. For instance, the SOCA demodulator that expands 112 nodes ($\mathbf{b} = (16, 2, 2, 2)^T$) was compared to the SFSD with 3-bit quantization and clipping at $8.5\,\mathrm{dB}$. It was observed that the SOCA performs $0.3\,\mathrm{dB}$ worse. Furthermore, since quantization-based pruning yields more pronounced savings at low SNR and clipping-based pruning yields larger savings at high SNR, their combination provides a more uniform complexity over the SNR range considered.

V. Conclusion

In this paper we considered MIMO-BICM systems working with LLR quantization. The well-known SFSD scheme was modified to take into account the subsequent quantization stage in order to prune in advance those tree paths that result in superfluous LLR refinements. LLR clipping is efficiently combined with quantization-based pruning to further reduce the average number of visited nodes in the SFSD tree search. Both pruning techniques are complementary, that is, savings from both methods add up to give a lower complexity, being the clipping-based pruning more important at high SNR. Further work includes finding quantization-based pruning techniques suitable for other soft detection methods.

References

[1] B. M. Hochwald and S. ten Brink, “Achieving near-capacity on a multiple-antenna channel,” *IEEE Trans. Commun.*, vol. 51, no. 3, pp. 389–399, Mar. 2003.

[2] C. Studer, A. Burg, and H. Bolcskei, “Soft-output sphere decoding: algorithms and VLSI implementation,” *IEEE J. Sel. Areas Commun.*, vol. 26, no. 2, pp. 290–300, Feb. 2008.

[3] L. G. Barbero, T. Ratnarajah, and C. Cowan, “A low-complexity soft-MIMO detector based on the fixed-complexity sphere decoder,” in *2008 IEEE International Conference on Acoustics, Speech and Signal Processing*.

[4] D. Milliner, E. Zimmermann, J. Barry, and G. Fettweis, “A fixed-complexity smart candidate adding algorithm for soft-output MIMO detection,” *IEEE J. Sel. Topics Signal Process.*, vol. 3, no. 6, pp. 1016–1025, Dec. 2009.

[5] X. Wu and J. S. Thompson, “A fixed-complexity soft-MIMO detector via parallel candidate adding scheme and its FPGA implementation,” *IEEE Commun. Lett.*, vol. 15, no. 2, pp. 241–243, Feb. 2011.

[6] C. Novak, P. Fertl, and G. Matz, “Quantization for soft-output demodulators in bit-interleaved-coded-modulation systems,” in *2009 IEEE International Symposium on Information Theory*.

[7] J. Jalden, L. G. Barbero, B. Ottersten, and J. S. Thompson, “The error probability of the fixed-complexity sphere decoder,” *IEEE Trans. Signal Process.*, vol. 57, no. 7, pp. 2711–2720, July 2009.