

EFFICIENT SOFT DEMODULATION IN MIMO-OFDM SYSTEMS WITH BICM AND CONSTANT MODULUS ALPHABETS

Dominik Seethaler, Gerald Matz, and Franz Hlawatsch

Institute of Communications and Radio-Frequency Engineering, Vienna University of Technology
Gusshausstrasse 25/389, A-1040 Vienna, Austria (Europe)
phone: +43 1 58801 38958, fax: +43 1 58801 38999, email: dominik.seethaler@tuwien.ac.at
web: http://www.nt.tuwien.ac.at

ABSTRACT

MIMO-OFDM using *bit-interleaved coded modulation* (BICM) is an attractive scheme for wireless communications in frequency-selective fading channels. The BICM decoder requires log-likelihood ratios (LLRs) whose exact computation is extremely costly. In this paper, we present a novel method termed *soft sphere-projection algorithm* (SSPA) that provides approximate LLRs for the case of constant modulus symbol alphabets. The SSPA has very low computational complexity and performs nearly as well as the list sphere decoder. These properties are demonstrated by numerical simulations using synthetic and measured MIMO channels.

1. INTRODUCTION

MIMO-OFDM with *bit-interleaved coded modulation* (BICM) is a promising technique for wireless communications over frequency-selective fading channels (e.g., [1–4]). Such systems usually employ a two-stage receiver consisting of a soft demodulator and a channel decoder. The soft demodulator computes log-likelihood ratios (LLRs)—i.e., soft values—for the coded bits; these LLRs are used by the decoder as bit metrics. Unfortunately, exact LLR calculation is computationally very costly. Various efficient MIMO-BICM soft demodulation algorithms providing approximate LLRs have therefore been proposed, such as the list extension of the Fincke-Phost sphere decoding (LFPSD) algorithm [5] and algorithms based on zero-forcing (ZF) or minimum mean-square error (MMSE) equalization [4, 6, 7].

In this paper, we present a novel soft demodulation method, termed *soft sphere-projection algorithm* (SSPA), that is applicable to constant modulus alphabets. The SSPA is based on the (hard-decision) sphere-projection algorithm (SPA) [8, 9] and achieves nearly the performance of the LFPSD at a small fraction of the LFPSD's computational complexity. We note that a similar approach using a different hard-detector was described in [10]. We also present methods that exploit subcarrier correlation for a further reduction of the computational complexity of the SSPA in MIMO-OFDM/BICM systems.

The paper is organized as follows. In Section 2, we provide some background on MIMO-OFDM using BICM. The SSPA is outlined in Section 3 and described in more detail in Section 5 (Section 4 reviews the SPA). Simulation results and conclusions are provided in Section 6 and 7, respectively.

2. MIMO-OFDM WITH BICM

System Model. We consider a BICM-based MIMO-OFDM system (see e.g. [2, 4]) with M_T transmit antennas, $M_R \geq M_T$ receive antennas, and K subcarriers. A block of information bits is passed

through a channel coder and an interleaver. The coded and interleaved bits are mapped to complex data symbols $d_{m,k} \in \mathcal{A}$ that are transmitted on the m th antenna and k th subcarrier ($m = 1, \dots, M_T$, $k = 0, \dots, K - 1$). Throughout this paper, we assume Gray labeling and a constant modulus symbol alphabet \mathcal{A} with $|d_{m,k}| = 1$.

The impulse response of the frequency-selective MIMO channel is assumed to have L taps $\mathbf{H}[l]$, $l = 0, \dots, L - 1$. The MIMO-OFDM modulator and demodulator convert the frequency-selective MIMO channel into K parallel flat fading MIMO channels acting as

$$\mathbf{r}_k = \mathbf{H}_k \mathbf{d}_k + \mathbf{w}_k, \quad k = 0, \dots, K - 1. \quad (1)$$

Here, \mathbf{r}_k denotes the received vector, $\mathbf{H}_k = \sum_{l=0}^{L-1} \mathbf{H}[l] e^{-j2\pi \frac{kl}{K}}$ is the channel matrix, $\mathbf{d}_k \triangleq (d_{1,k} \cdots d_{M_T,k})^T \in \mathcal{D} \triangleq \mathcal{A}^{M_T}$ is the transmit vector, and \mathbf{w}_k is additive white Gaussian noise.

Using \mathbf{r}_k and \mathbf{H}_k , the soft demodulator calculates (approximate) LLRs for each of the $M_T \log_2 |\mathcal{A}|$ code bits associated with \mathbf{d}_k . After deinterleaving, the LLRs from all layers and subcarriers serve as bit metrics for the soft-input channel decoder.

LLR Calculation. In what follows, we will omit the subcarrier index k to simplify notation. Let $b_{m,i}$, $i = 1, \dots, \log_2 |\mathcal{A}|$ denote the bit label of the symbol $d_m \in \mathcal{A}$ of the m th layer. We assume that the code bits $b_{m,i}$ are i.i.d. uniform. Using the log-sum approximation [2, 5] and ignoring scaling factors irrelevant to our setup, the LLR for $b_{m,i}$ equals

$$\Lambda_{m,i} \triangleq \lambda_{m,i}^0 - \lambda_{m,i}^1, \quad \text{with } \lambda_{m,i}^b \triangleq \min_{\mathbf{d} \in \mathcal{D}_{m,i}^b} \psi^2(\mathbf{d}). \quad (2)$$

Here, $\psi^2(\mathbf{d}) \triangleq \|\mathbf{r} - \mathbf{H}\mathbf{d}\|^2$ and $\mathcal{D}_{m,i}^b$ denotes the set of all transmit vectors whose label at layer m and bit position i equals $b \in \{0, 1\}$. Note that $\mathcal{D}_{m,i}^b$ consists of all data vectors \mathbf{d} with $d_m \in \mathcal{A}_i^b$ and $d_{m'} \in \mathcal{A}$, $m' \neq m$, where the reduced alphabet \mathcal{A}_i^b is obtained by retaining only those symbols of \mathcal{A} whose bit label at position i equals b . In practice, exact soft demodulation according to (2) may be unfeasible since its complexity scales exponentially with M_T .

Soft demodulation is closely related to the ML detector

$$\hat{\mathbf{d}}_{\text{ML}} = \underset{\mathbf{d} \in \mathcal{D}}{\operatorname{argmin}} \|\mathbf{r} - \mathbf{H}\mathbf{d}\|^2 = \underset{\mathbf{d} \in \mathcal{D}}{\operatorname{argmin}} \psi^2(\mathbf{d}). \quad (3)$$

Denoting the i th bit of the m th component of $\hat{\mathbf{d}}_{\text{ML}}$ as $(\hat{\mathbf{d}}_{\text{ML}})_{m,i} = b$, it follows that $\hat{\mathbf{d}}_{\text{ML}} \in \mathcal{D}_{m,i}^b$ and hence we immediately obtain

$$\lambda_{m,i}^b = \psi^2(\hat{\mathbf{d}}_{\text{ML}}). \quad (4)$$

However, (3) also has exponential complexity and one must still compute the term $\lambda_{m,i}^b$ in (2) ($\bar{b} = 1 - b$ denotes bit flipping).

3. BASIC IDEA OF THE SSPA

The idea behind the SSPA is as follows. The SPA [8,9] is an efficient approximation to the ML detector. It replaces \mathcal{D} with a *reduced search set* $\tilde{\mathcal{D}}$ to obtain a hard decision

$$\hat{\mathbf{d}}_{\text{SP}} \triangleq \arg \min_{\mathbf{d} \in \tilde{\mathcal{D}}} \psi^2(\mathbf{d}) \quad (5)$$

such that $\psi^2(\hat{\mathbf{d}}_{\text{SP}}) \approx \psi^2(\hat{\mathbf{d}}_{\text{ML}})$ (Section 4 will provide more details). This together with (4) yields the approximation

$$\lambda_{m,i}^b \approx \tilde{\lambda}_{m,i}^b \triangleq \psi^2(\hat{\mathbf{d}}_{\text{SP}}), \quad \text{for } b = (\hat{\mathbf{d}}_{\text{SP}})_{m,i}.$$

It remains to calculate a similar approximation for $\tilde{\lambda}_{m,i}^b$. A comparison of (2) and (4) reveals that such an approximation could in principle be obtained by applying the SPA with $\tilde{\mathcal{D}}_{m,i}^b$ instead of \mathcal{D} , i.e.,

$$\tilde{\lambda}_{m,i}^b \approx \tilde{\lambda}_{m,i}^b \triangleq \min_{\mathbf{d} \in \tilde{\mathcal{D}}_{m,i}^b} \psi^2(\mathbf{d}). \quad (6)$$

Here $\tilde{\mathcal{D}}_{m,i}^b$ is the reduced search set corresponding to $\tilde{\mathcal{D}}_{m,i}^b$. In general, computing (6) requires a second SPA pass. However, we will show in Section 5 that this second SPA pass can be circumvented and $\tilde{\lambda}_{m,i}^b$ can be calculated with little additional effort using intermediate results of the first SPA pass. In fact, all vectors $\mathbf{d} \in \tilde{\mathcal{D}}_{m,i}^b$ can be obtained by a *requantization* of the vectors $\mathbf{d} \in \tilde{\mathcal{D}}$ and the associated distances $\psi^2(\mathbf{d})$ can be computed via simple updates of $\psi^2(\mathbf{d})$ (see Section 5 and in particular (9)). The SSPA for one subcarrier can be summarized as follows:

Step 1: Use the SPA to calculate the reduced search set $\tilde{\mathcal{D}}$, the associated distance set $\Psi = \{\psi^2(\mathbf{d}), \mathbf{d} \in \tilde{\mathcal{D}}\}$, the detector output $\hat{\mathbf{d}}_{\text{SP}}$, and $\psi^2(\hat{\mathbf{d}}_{\text{SP}})$, the minimum element of Ψ .

Step 2: For each $m \in \{1, \dots, M_T\}$ and $i \in \{1, \dots, \log_2 |\mathcal{A}|\}$:

- determine $b = (\hat{\mathbf{d}}_{\text{SP}})_{m,i}$ and set $\tilde{\lambda}_{m,i}^b = \psi^2(\hat{\mathbf{d}}_{\text{SP}})$;
- calculate $\Psi_{m,i}^b = \{\psi^2(\mathbf{d}), \mathbf{d} \in \tilde{\mathcal{D}}_{m,i}^b\}$ by requantizing each $\mathbf{d} \in \tilde{\mathcal{D}}$ and performing the distance update (9);
- obtain $\tilde{\lambda}_{m,i}^b$ as the minimum element of $\Psi_{m,i}^b$;
- finally, calculate the approximate LLR

$$\Lambda_{m,i} \approx (1-2b)(\tilde{\lambda}_{m,i}^b - \tilde{\lambda}_{m,i}^b),$$


where the factor $1-2b \in \{-1, 1\}$ serves to adjust the sign.

We recall that the SSPA has to be performed for each subcarrier separately (cf. (1)). For $M_T = M_R$, the overall complexity can be shown to scale as $\mathcal{O}(KM_T^3 |\mathcal{A}| \log_2 |\mathcal{A}|)$.

4. REVIEW OF THE SPA

Before explaining the SSPA in more detail, we provide a short review of the hard-decision SPA (see [8,9] for further details). In essence, the SPA is an add-on to an arbitrary reference detector (in this paper we restrict to ZF and MMSE detectors) that uses a reduced search set $\tilde{\mathcal{D}}$ to approximate the ML detector (cf. (5)). The reduced search set is given by $\tilde{\mathcal{D}} \triangleq \{\hat{\mathbf{d}}\} \cup \mathcal{P}$. Here, $\hat{\mathbf{d}} = Q\{\mathbf{y}\}$ is the reference detector output (\mathbf{y} is the output of the ZF or MMSE equalizer and $Q\{\mathbf{y}\}$ denotes elementwise quantization with respect to \mathcal{A}) and \mathcal{P} is an additional search set (see next paragraph). Apart from $\tilde{\mathcal{D}}$, the SPA provides the corresponding distance set $\Psi = \{\psi^2(\mathbf{d}), \mathbf{d} \in \tilde{\mathcal{D}}\}$ and $\hat{\mathbf{d}}_{\text{SP}} = \arg \min_{\mathbf{d} \in \tilde{\mathcal{D}}} \psi^2(\mathbf{d})$. One can show that $|\tilde{\mathcal{D}}| \leq |\mathcal{A}|M_T + 2$ (e.g., for $M_T = 6$ and 4-QAM, the size of the reduced search set $\tilde{\mathcal{D}}$ is less than 26 while the full search set \mathcal{D} has

(a) Component circles:



(b) Angle intervals:




Figure 1: Illustration of SPA geometry for a 4×4 MIMO channel and 4-QAM modulation: (a) component circles \mathcal{C}_m and (b) associated angle intervals I^r (corresponding to \mathcal{P}) and \tilde{I}^r (corresponding to $\mathcal{P}_{3,2}^0$). Here, $\mathcal{A}_2^0 = \{\frac{1+j}{\sqrt{2}}, \frac{1-j}{\sqrt{2}}\}$.

size 4096). The computational complexity of the SPA (not counting the complexity of the reference detector) is $\mathcal{O}(|\mathcal{A}|M_T^2)$.

The construction of \mathcal{P} is based on approximating \mathbf{H} by an *idealized bad channel* (IBC). Furthermore, it exploits the constant modulus property which implies that all data vectors lie on an M_T -dimensional hypersphere \mathcal{H} , i.e., $\mathcal{D} \subset \mathcal{H}$. The center of this hypersphere is the origin and its radius equals $R = \sqrt{M_T}$. The IBC captures essential properties of a poorly conditioned channel \mathbf{H} by approximating the smallest singular value of \mathbf{H} with zero and the remaining singular values with the largest singular value.

In what follows, let \mathbf{v} denote the right singular vector of \mathbf{H} corresponding to the smallest singular value. With the IBC approximation, it is shown in [8,9] that $\psi^2(\mathbf{d})$ is small if and only if \mathbf{d} is close to the reference line

$$\mathcal{L}: \mathbf{y}_{\mathcal{L}}(\alpha) = \alpha \mathbf{v} + \mathbf{y}^{\perp}, \quad \alpha \in \mathbb{C}, \quad (7)$$

where $\mathbf{y}^{\perp} = (\mathbf{I} - \mathbf{v}\mathbf{v}^H)\mathbf{y}$ denotes the component of \mathbf{y} orthogonal to \mathbf{v} . For that reason, \mathcal{P} will be composed of data vectors that are close to \mathcal{L} . Here, two cases can be distinguished.

Case I. If \mathcal{L} intersects \mathcal{H} , \mathcal{P} is chosen to consist of data vectors close to the intersection circle $\mathcal{C} = \mathcal{L} \cap \mathcal{H}$ that is given by

$$\mathcal{C}: \mathbf{y}_{\mathcal{C}}(\phi) \triangleq \sqrt{R^2 - \|\mathbf{y}^{\perp}\|^2} e^{j\phi} \mathbf{v} + \mathbf{y}^{\perp}, \quad \phi \in [0, 2\pi).$$

Evidently, the data vector closest to $\mathbf{y}_{\mathcal{C}}(\phi)$ equals $Q\{\mathbf{y}_{\mathcal{C}}(\phi)\}$. Thus, \mathcal{P} consists of all data vectors obtained as $Q\{\mathbf{y}_{\mathcal{C}}(\phi)\}$ for all possible angles ϕ . For an efficient construction of \mathcal{P} , \mathcal{C} is viewed as the collection of M_T component circles \mathcal{C}_m in the symbol domain (see Fig. 1(a)). If we move along \mathcal{C} by varying ϕ , the m th component of $\mathbf{y}_{\mathcal{C}}(\phi)$ moves along \mathcal{C}_m . The first data vector in \mathcal{P} is obtained as $\mathbf{d}^1 = Q\{\mathbf{y}_{\mathcal{C}}(0)\}$. Then ϕ is increased until one of the components of $\mathbf{y}_{\mathcal{C}}(\phi)$ crosses the boundary of a symbol decision region; at that point, a second data vector \mathbf{d}^2 is added to \mathcal{P} . This procedure is continued until we have moved along the whole intersection circle \mathcal{C} . Thus, each data vector $\mathbf{d}^r \in \mathcal{P}$, $r = 1, \dots, |\mathcal{P}|$ is associated with a specific interval I^r of the angle parameter ϕ , i.e., $\mathbf{d}^r = Q\{\mathbf{y}_{\mathcal{C}}(\phi)\}$ for any $\phi \in I^r$ (cf. Fig. 1).

Case II. If \mathcal{L} and \mathcal{H} do not intersect, then \mathbf{y}^\perp is the point of \mathcal{L} closest to \mathcal{H} [8], and thus $\mathbf{d}^1 = Q\{\mathbf{y}^\perp\}$. The remaining data vectors in \mathcal{P} are chosen as the nearest neighbors of \mathbf{d}^1 , which can be easily found by substituting the two nearest-neighbor symbols for the individual components of \mathbf{d}^1 .

5. SSPA IMPLEMENTATION DETAILS

Preliminaries. Recall that $\mathcal{P}_{m,i}^b$ consists of all data vectors \mathbf{d} with $d_m \in \mathcal{A}_i^b$ and $d_{m'} \in \mathcal{A}$, $m' \neq m$. Thus, when determining the reduced search set $\tilde{\mathcal{P}}_{m,i}^b = \{\hat{\mathbf{d}}_{m,i}^b\} \cup \mathcal{P}_{m,i}^b$ for (6), the main difference to the initial SPA pass is that quantization and nearest neighbor search for layer m are limited to the reduced alphabet \mathcal{A}_i^b . The efficient implementation of the SSPA is based on the following fundamental observation. For Gray labeling and a constant modulus alphabet whose size is a power of two, it can be shown that quantization with respect to the reduced alphabet \mathcal{A}_i^b (denoted $Q_i^b\{\cdot\}$) remains unchanged when preceded by quantization with respect to \mathcal{A} , i.e.,

$$Q_i^b\{Q\{y\}\} = Q_i^b\{Q\{y\}\}. \quad (8)$$

This means that the decision boundaries for \mathcal{A}_i^b form a subset of the boundaries for \mathcal{A} . Thus, any quantization with respect to \mathcal{A}_i^b can be performed using \mathcal{A} -quantized values. This allows one to determine $\tilde{\mathcal{P}}_{m,i}^b$ solely by using $\tilde{\mathcal{P}}$ (see below).

The basic operation to obtain a data vector $\tilde{\mathbf{d}} \in \tilde{\mathcal{P}}_{m,i}^b$ is the *re-quantization* of the m th component of $\mathbf{d} \in \tilde{\mathcal{P}}$ (the other components remain unchanged):

$$\tilde{d}_{m'} = \begin{cases} Q_i^b\{d_m\}, & \text{if } m' = m, \\ d_{m'}, & \text{if } m' \neq m. \end{cases}$$

On the vector level, this re-quantization of the m th component will be denoted $\tilde{\mathbf{d}} = Q_{m,i}^b\{\mathbf{d}\}$. For the reference detector output, we thus have $\hat{\mathbf{d}}_{m,i}^b = Q_{m,i}^b\{\hat{\mathbf{d}}\}$. It remains to determine $\mathcal{P}_{m,i}^b$.

Construction of $\mathcal{P}_{m,i}^b$. If \mathcal{L} intersects \mathcal{H} (case I in Section 4), each data vector $\mathbf{d}^r \in \mathcal{P}$ is associated to a specific interval I^r , i.e., $\mathbf{d}^r = Q\{\mathbf{y}_{\mathcal{C}}(\phi)\}$ for any $\phi \in I^r$. For $\mathcal{P}_{m,i}^b$, we now use \mathcal{A}_i^b instead of \mathcal{A} for layer m . Thus, each data vector $\tilde{\mathbf{d}}^r \in \mathcal{P}_{m,i}^b$ is obtained as $\tilde{d}_m^r = Q_i^b\{y_{\mathcal{C},m}(\phi)\}$ and $\tilde{d}_{m'}^r = Q\{y_{\mathcal{C},m'}(\phi)\}$, $m' \neq m$, for ϕ in a specific interval \tilde{I}^r . Property (8) implies that an interval \tilde{I}^r either equals an interval I^r or is a union of adjacent intervals I^r (cf. Fig. 1(b)). This is due to the fact that some decision boundaries for \mathcal{A} and the corresponding intersections with the m th component circle \mathcal{C}_m are no longer present for the reduced alphabet \mathcal{A}_i^b . It follows that the *entire* set $\mathcal{P}_{m,i}^b$ can be obtained by re-quantizing every data vector $\mathbf{d}^r \in \mathcal{P}$, i.e., $\tilde{\mathbf{d}}^r = Q_{m,i}^b\{\mathbf{d}^r\}$. In that way, data vectors $\tilde{\mathbf{d}}^r$ associated to an interval \tilde{I}^r that is the union of two or more intervals I^r will be multiply obtained. We do not take this into account, since the resulting complexity reduction would be insignificant.

If \mathcal{L} and \mathcal{H} do not intersect (case II in Section 4), \mathcal{P} consists of $\mathbf{d}^1 = Q\{\mathbf{y}^\perp\}$ and all its nearest neighbors. Thus, $\tilde{\mathbf{d}}^1 \in \mathcal{P}_{m,i}^b$ is obtained as $\tilde{\mathbf{d}}^1 = Q_{m,i}^b\{\mathbf{d}^1\}$. Furthermore, for constant modulus alphabets it can be checked quite easily that $Q_{m,i}^b\{\mathbf{d}^r\}$, $\mathbf{d}^r \in \mathcal{P}$, $r = 2, \dots, |\mathcal{P}|$ yields the *entire* set of nearest neighbors of $\tilde{\mathbf{d}}^1$.

Efficient Calculation of $\Psi_{m,i}^b$. To calculate $\tilde{\lambda}_{m,i}^b$ according to (6), we need the distance set $\Psi_{m,i}^b$ associated to $\tilde{\mathcal{P}}_{m,i}^b$. The evaluation of

$\psi^2(\tilde{\mathbf{d}})$ has roughly complexity $\mathcal{O}(M_T M_R)$. We next show that the distances $\psi^2(\tilde{\mathbf{d}})$, $\tilde{\mathbf{d}} \in \tilde{\mathcal{P}}_{m,i}^b$ can be obtained via a low-complexity update of $\psi^2(\mathbf{d})$, $\mathbf{d} \in \tilde{\mathcal{P}}$, provided that the error vectors $\mathbf{e}(\mathbf{d}) \triangleq \mathbf{r} - \mathbf{H}\mathbf{d}$ were stored during the first SPA pass. Consider the re-quantized data vector $\tilde{\mathbf{d}} = Q_{m,i}^b\{\mathbf{d}\}$ and let $\Delta_{m,i}^b \triangleq \tilde{d}_m - d_m$ denote the difference between the m th components of $\tilde{\mathbf{d}}$ and \mathbf{d} . It can then be shown that

$$\psi^2(\tilde{\mathbf{d}}) = \psi^2(\mathbf{d}) + \|\mathbf{h}_m\|^2 |\Delta_{m,i}^b|^2 - 2 \operatorname{Re}\left\{ \mathbf{e}^H(\mathbf{d}) \mathbf{h}_m \Delta_{m,i}^b \right\}, \quad (9)$$

where \mathbf{h}_m denotes the m th column of \mathbf{H} . This update requires roughly $\mathcal{O}(M_R)$ operations.

Exploiting Subcarrier Correlation. In a MIMO-OFDM system with BICM, the SSPA has to be run for each subcarrier. Further computational savings are possible by exploiting the strong correlation of the MIMO channel matrices \mathbf{H}_k of nearby subcarriers [11] (now, we take the subcarrier index k again into account, see (1)). In particular, this allows us to use efficient interpolation schemes (e.g. [11]) to calculate the inverses of \mathbf{H}_k , $k = 1, \dots, K$ required by the ZF or MMSE reference detector.

For the SSPA, we further need the right singular vector \mathbf{v}_k corresponding to the smallest singular value of \mathbf{H}_k (see (7)). Since \mathbf{v}_k also is the eigenvector of the inverse Gram matrix¹ $\mathbf{G}_k = (\mathbf{H}_k^H \mathbf{H}_k)^{-1}$ associated to the largest singular value, it can be efficiently computed by means of a few iterations of the power method [12]:

$$\mathbf{v}_k^{(j)} = \frac{\mathbf{G}_k \mathbf{v}_k^{(j-1)}}{\|\mathbf{G}_k \mathbf{v}_k^{(j-1)}\|}, \quad j = 1, \dots, J.$$

Usually, these iterations are initialized with a random vector $\mathbf{v}_k^{(0)}$. However, since it can be expected that \mathbf{v}_k does not change dramatically from one subcarrier to another, we propose to initialize the power method with the singular vector computed for the preceding subcarrier, i.e., $\mathbf{v}_k^{(0)} = \mathbf{v}_{k-1}^{(J)}$. In general, this significantly reduces the number of iterations required. In our numerical simulations (see next section), we even used $J = 1$, i.e., just a *single* iteration (matrix-vector product plus normalization) for the singular vector update (apart from the first subcarrier where we used a random initialization and $J = 5$).

6. SIMULATION RESULTS

We next provide simulation results to illustrate the packet error rate (PER) performance and computational complexity of the SSPA for synthetic and measured frequency-selective MIMO channels. We simulated a BICM-based MIMO-OFDM system with $M_T = M_R = 4$ transmit and receive antennas, $K = 128$ subcarriers, a rate-1/2 16-state convolutional code with octal generators (23, 35), 4 bits trellis termination, a random block interleaver, 4-QAM symbol alphabet, and Gray labeling (thus, each MIMO-OFDM symbol carries 508 information bits). At the receiver, a Viterbi decoder with a traceback depth of 25 was employed for channel decoding. As soft demodulators, we used the SSPA in conjunction with ZF and MMSE detectors as reference detectors (denoted SSPA-Z and SSPA-M, respectively), the LFPSSD [5], and ZF-based and MMSE-based soft demodulation according to [6] and [7], respectively. SSPA-Z was considered besides SSPA-M since it has the advantage that knowledge of the noise variance is not required. The LFPSSD used LLR thresholds ± 8 and 32 candidate data vectors inside the hypersphere.

¹We note that the \mathbf{G}_k 's can also be computed efficiently using the interpolation techniques in [11].

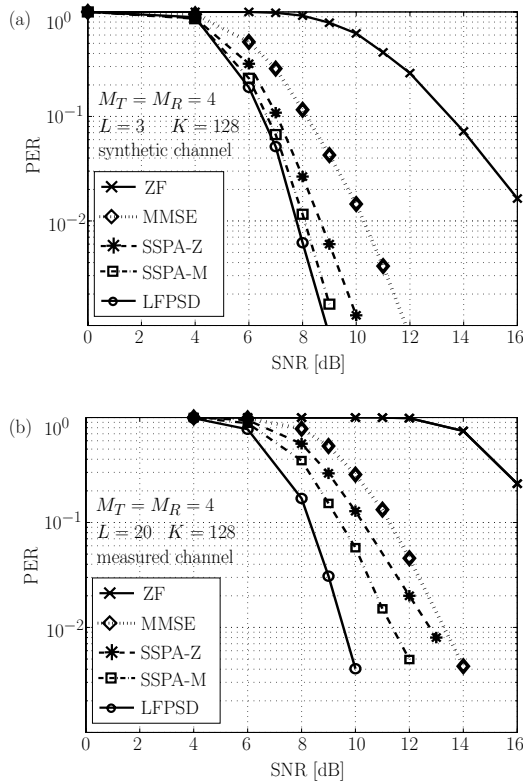


Figure 2: PER performance of various soft-demodulation schemes versus SNR for (a) a synthetic MIMO channel, (b) a measured MIMO channel.

Synthetic Channel. We first consider a synthetic MIMO channel with $L = 3$. The channel taps $\mathbf{H}[l]$ are spatially and temporally i.i.d. The PER obtained with the different soft demodulation algorithms versus SNR is shown in Fig. 2(a). For a PER of 10^{-2} , SSPA-M achieves virtually the same performance as LFPSPD and SSPA-Z performs less than 1 dB worse. ZF-based and MMSE-based demodulation perform more than, respectively, 8 dB and 2 dB worse than their SSPA counterparts. In contrast to ZF/MMSE-based demodulation, SSPA-M apparently does not incur any diversity loss.

Measured MIMO Channel. Next we evaluate all schemes using indoor MIMO channel measurements obtained at Vienna Airport.² The transmitter and receiver positions were fixed and there was no line of sight. The channel data comprised 4452 impulse response snapshots of length $L = 20$. Fig. 2(b) again shows PER versus SNR for the various soft demodulators. While in general PER performance for the measured channel (specifically the coding gain) is somewhat poorer, the relations of the individual algorithms to each other are quite similar to the synthetic case. To achieve a PER of 10^{-2} , an SNR of about 9.5 dB, 11.3 dB, 12.8 dB, 13.3 dB, and 20 dB is required for LFPSPD, SSPA-M, SSPA-Z, MMSE, and ZF, respectively. The gap between SSPA-M and LFPSPD now is larger, which can be attributed to spatial channel correlations that reduce the quality of the IBC approximation underlying the SSPA.

Computational Complexity. For a rough comparison of the computational complexity of the various demodulators, we provide MATLAB kflaps (per subcarrier and OFDM symbol) measured in

²The authors thank N. Czink for providing the measurement data. A detailed description of the measurement campaign can be found in [13].

the synthetic simulation setup within 100 OFDM symbols at an SNR of 7 dB. We did not count the calculations for the equalizers (required for ZF/MMSE and SSPA) and QR-decompositions (required for LFPSPD). With LFPSPD and SSPA, complexity depends on the individual channel realization. The LFPSPD required up to 118.2 kflaps with an average of 44 kflaps. With the SSPA, on average 4.2 and at most 6.7 kflaps were required. The complexity of ZF/MMSE based demodulation was 0.3 kflaps. We conclude that SSPA is more complex than ZF and MMSE, but an order of magnitude more efficient than LFPSPD.

7. CONCLUSIONS

The *soft sphere-projection algorithm* (SSPA) is a novel soft demodulation technique for MIMO-OFDM systems using bit-interleaved coded modulation (BICM). The SSPA is based on the (hard-decision) sphere-projection algorithm (SPA) and exploits intermediate SPA results to obtain approximate log-likelihood ratios with very low computational effort. Simulation results demonstrated that the performance of the SSPA is similar to that of Fincke-Phost list sphere decoding although the computational complexity is significantly smaller. Thus, the SSPA offers an excellent complexity-performance compromise for MIMO-OFDM/BICM systems.

REFERENCES

- [1] J. J. Boutros, F. Boixadera, and C. Lamy, "Bit-interleaved coded modulations for multiple-input multiple-output channels," in *Proc. IEEE Symp. Spread Spectrum Techn. Appl.*, (Parsippany, NJ), pp. 123–126, Sept. 2000.
- [2] 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.
- [3] R. Visoz, A. O. Berthet, and J. J. Boutros, "Reduced-complexity iterative decoding and channel estimation for space time BICM over frequency-selective wireless channels," in *Proc. IEEE PIMRC-02*, (Lisbon, Portugal), pp. 1017–1022, Sept. 2002.
- [4] K.-B. Song and S. A. Mujtaba, "A low complexity space-frequency BICM MIMO-OFDM system for next-generation WLANs," in *Proc. IEEE Globecom 2003*, (San Francisco, CA), pp. 1059–1063, Dec. 2003.
- [5] B. M. Hochwald and S. ten Brink, "Achieving near-capacity on a multiple-antenna channel," *IEEE Trans. Inf. Theory*, vol. 51, pp. 389–399, March 2003.
- [6] M. Butler and I. Collings, "A zero-forcing approximate log-likelihood receiver for MIMO bit-interleaved coded modulation," *IEEE Commun. Letters*, vol. 8, pp. 105–107, Feb. 2004.
- [7] D. Seethaler, G. Matz, and F. Hlawatsch, "An efficient MMSE-based demodulator for MIMO bit-interleaved coded modulation," in *Proc. IEEE Globecom 2004*, vol. IV, (Dallas, Texas), pp. 2455–2459, Dec. 2004.
- [8] H. Artés, D. Seethaler, and F. Hlawatsch, "Efficient detection algorithms for MIMO channels: A geometrical approach to approximate ML detection," *IEEE Trans. Signal Processing*, vol. 51, pp. 2808–2820, Nov. 2003.
- [9] D. Seethaler, H. Artés, and F. Hlawatsch, "Efficient near-ML detection for MIMO channels: The sphere-projection algorithm," in *Proc. IEEE Globecom 2003*, vol. IV, (San Francisco, CA), pp. 2089–2093, Dec. 2003.
- [10] D. Seethaler, G. Matz, and F. Hlawatsch, "Low-complexity soft demodulation of MIMO-BICM using the line search detector," in *Proc. ICC 2005*, vol. IV, (Seoul, Korea), pp. 2447–2451, May 2005.
- [11] M. Borgmann and H. Bölcskei, "Interpolation-based efficient matrix inversion for MIMO-OFDM receivers," in *Proc. Asilomar Conf. Signals, Systems and Computers*, vol. II, (Pacific Grove, CA), pp. 1941–1947, Nov. 2004.
- [12] G. H. Golub and C. F. Van Loan, *Matrix Computations*. Baltimore: Johns Hopkins University Press, 3rd ed., 1996.
- [13] M. Herdin, *Non-Stationary Indoor MIMO Radio Channels*. PhD thesis, Vienna University of Technology, Aug. 2004.