On the Complexity of Sphere Decoding

Ernesto Zimmermann, Wolfgang Rave, Gerhard Fettweis

Dresden University of Technology, Vodafone Chair Mobile Communications Systems, D-01062 Dresden, Germany

email: zimmere@ifn.et.tu-dresden.de

Abstract— Recently, sphere detection has emerged as a powerful means of finding the maximum likelihood solution to the detection problem for multiple-antenna (MIMO) systems. In this paper, we analyze the performance and computational complexity of different sphere detectors and compare it with that of already known MIMO receiver techniques, in uncoded as well as coded transmission[1]. We propose a number of techniques that allow for reducing the complexity of sphere detection – some at the expense of bit error performance. Our results confirm that for uncoded transmission and low target BERs, sphere detectors outperform all other known receiver techniques with only minor additional complexity. For coded transmission, the complexity of sphere detection essentially scales with the number of candidates, motivating for further research to find techniques that provide good soft outputs with very low numbers of candidates.

I. INTRODUCTION

The main challenge of receiver design for multiple antenna (MIMO) systems lies in the non-orthogonality of the transmission channel – the superposition of the signals from all transmit antennas at the receiver side. Optimum maximum likelihood detection (MLD) requires finding the signal point $\hat{x}$ of the transmitter vector signal set that minimizes the Euclidean distance with respect to the received signal vector y when transmitted over the channel, i.e., the closest lattice point in a transformed vector space:

$$\hat{x} \equiv \arg \min \|y - Hx\|^2 \qquad (1)$$

Unfortunately this problem is exponential in the number of possible constellation points, making MLD unsuitable for practical purposes when aiming at high spectral efficiencies.

A number of sub-optimum receivers of low to moderate complexity have been devised, yet all suffer from rather limited performance. The most commonly considered techniques are linear receivers and Successive Interference Cancellation (SIC) [1]. It has to be stressed that for both techniques, the overall performance is limited by the quality of the strongest detected signal, hence none of them achieves diversity in the number of receive antennas, in contrast to MLD. The concept of sphere detection (SD) was introduced in [2] and has been further discussed in various publications [3], [4], [5]. To avoid the exponential complexity of the MLD problem, the search for the closest lattice point is restricted to include only vector constellation points that fall within a certain search sphere. This approach allows for finding the ML solution with only polynomial complexity, for sufficiently high SNR [4].

In this paper, we analyze the performance and computational complexity of different MIMO receiver algorithms, for coded and uncoded transmission. We propose a number of complexity reduction techniques for sphere decoding and show that complexity becomes comparable to that of other MIMO receiver techniques in the high SNR regime and can be significantly reduced in the low SNR regime using layer ordering and MMSE filtering. Sphere detection outperforms all other approaches in terms of bit error performance.

The remainder of this document is structured as follows: Section II introduces different sphere detector implementations, as well as methods for reducing their complexity. Complexity evaluations

are presented in Section III and bit error performances studied in Section IV, before we draw conclusions in Section V.

II. SPHERE DETECTION

In its search for the ML solution, the sphere detector evaluates all transmit vector signals x fulfilling:

$$\|y - Hx\|^2 < R^2 \qquad (2)$$

where R is the search radius of the sphere. Obviously, the selection of R is a critical issue largely influencing the complexity of any SD algorithm. Choosing R too large leads to a sphere containing a very high number of hypotheses (also referred to as candidates) and hence to high detection complexity. Choosing R too small will result in an empty sphere and the search has to be restarted with an increased radius [3] – eventually leading to similar problems. The search for candidates fulfilling (2) is done by back-substitution algorithms. Towards this aim, Cholesky [3] or QR [5] decompositions of the channel matrix H may be equivalently used. Note, however, that using a Cholesky decomposition may be advantageous in systems with more transmit than receive antennas, since the size of the upper triangular matrix R is limited by $\min\{N_{Rx}, M_{Tx}\}$, which for $N_{Rx} < M_{Tx}$ leads to an underdetermined problem for the QR implementation of the SD.

In the following, we concentrate ourselves on the symmetric case, i.e., $M_{Tx} = N_{Rx}$ and use a QR decomposition of H, for practicality of implementation. With $H = QR$, where R is upper triangular and Q is unitary, (2) may be rewritten as follows:

$$
\begin{aligned}
\|y - Hx\|^2 &< R^2 \\
\|Q^H y - Rx\|^2 &< R^2 \\
\|y' - Rx\|^2 &< R^2.
\end{aligned}
\qquad (3)
$$

It is easily seen that (3) implies

$$\sum_{i=l}^{M_{Tx}} \left| y'_i - \sum_{j=i}^{M_{Tx}} r_{i,j} x_j \right|^2 < R^2, \qquad l = 1, \dots, M_{Tx}. \quad (4)$$

The detection process starts with the last layer $l = M_{Tx}$, for which (4) reduces to $|y'_{M_{Tx}} - r_{M_{Tx}, M_{Tx}} x_{M_{Tx}}|^2 < R^2$ and then works its way up until the first layer is detected. This process is quite similar to SIC techniques – the signals from previously detected layers are subtracted from the received signal before detection within the current layer is performed. However, in SD, detection at layer i essentially takes on the form

$$|c_{c,i} - x_i|^2 < \frac{R_{c,i}^2}{r_{i,i}^2}, \qquad (5)$$

where $c_{i,c}$ is the search center and $R_{c,i}$ the *remainder* search radius for the currently considered (incomplete) candidate c. Both depend on the estimates for previously detected layers and hence vary from candidate to candidate. The above inequality implies that not only a single, but several constellation points may be selected. The SD receiver hence performs its search in a tree like structure, which motivates the search for appropriate enumeration strategies in [5].

Since we eventually consider coded transmission where extracting good quality soft outputs requires finding a predefined minimum number of candidates $N_{c,min}$ [3], we take a different approach towards this problem. Our interest lies in ensuring a

certain minimum number of entries N_c at the bottom of the search tree. With the entries of $\mathbf{H}$ and the noise $\mathbf{n}$ at the receiver being random variables, N_c is obviously a random variable, too. If we allow for all hypotheses fulfilling (3) to be tested, we get an upper bound on the expected complexity. Remember that whenever an empty sphere is declared, the search must be restarted with an increased radius. Since the overall detection complexity is the sum of all taken attempts to find the ML point, it is unfavorable to choose a small initial search radius and then increasing it stepwise until $N_{c,min}$ candidates are found.

In the following we will describe some extensions and variants of SD that may be used to reduce its computational complexity, or at least trade bit error performance against implementation complexity.

A. Initial Search Radius

For our symmetric $M \times M$ MIMO system the proposal for R from [3] reduces to $R^2 = 2\sigma^2 K N_{Rx}$. Since in our system we normalize $2\sigma^2 = 1$ and then scale the transmitter signal constellation, this is further reduced to $R^2 = N_{Rx} K$ where the critical point lies in the selection of K. We found empirically that it is beneficial to use a small initial search radius at low SNR that is subsequently increased with rising SNR. In the course of our evaluations using

$$ K = \frac{1}{60} L R_c \sqrt{\frac{E_b}{N_0}}, \qquad (6) $$

where L is the number of bits per symbol and R_c is the code rate, proved to be a good choice that minimized detection complexity in a wide SNR range. The middle factor ensures that the scaling of the transmitter signal set is taken into account and the last term derives from the notion that in the high SNR regime, finding too many candidates becomes less probable and that the number of events with an empty sphere should be minimized. We increased the radius by a factor of 1.5 whenever an empty sphere was declared.
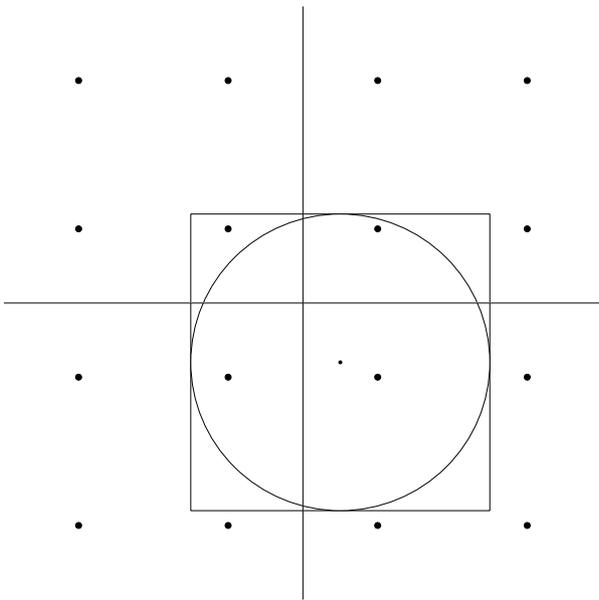
Fig. 1. Low complexity selection of candidates for sphere decoding: all points that fall within the search square are tentatively used as candidates (in this case, 4 points). This is easily done by checking only 4 bounds. Of the 4 tentative candidates, all that are inside the square but outside the circle (in this case, the upper left point) will be excluded from further search when calculating the remainder search radius.

B. Layer ordering

Looking at the tree like structure of the SD process, the overall detection complexity of a SD can be illustrated by the area "under"

this tree, i.e., the sum of the number of (incomplete) candidates for each layer, going from $l = M_{Tx}$ to 1. By using a sorted QR decomposition (SQRD) [6], we ensure that more reliably received layers are detected first, and hence the width of the search tree is reduced for layers $l = M_{Tx}, \ldots, 2$ (cf. [5]). Remember that with R fixed, the width N_c of the tree at its bottom is the same, irrespective of layer ordering. The width reduction becomes evident by looking again at (4) – the radius of the search circle in each detection step scales with $1/r_{i,i}$. Sorting the layers such that $i < j \Leftrightarrow |r_{i,i}| < |r_{j,j}|$ ensures that this radius is as small as possible for the first detected layers. As a result, the width of the search tree increases only slowly before reaching finally N_c. Thus, the area "under" the tree and hence detection complexity are reduced.

C. MMSE Extension

One major problem of the QR implementation of the SD is that for close-to-singular values of $r_{i,i}$, the search circle becomes very large in detection of a layer, leading to a high number of candidates. This effect is obviously more critical in the low SNR regime, since more constellation points fall in the increased circle. Following the approach in [7] we extend the channel matrix $\mathbf{H}$ such that the SQRD will take the noise variance at the receiver into account. This reduces the number of close-to-singular diagonal entries in $\mathbf{R}$ and thus detection complexity.

D. Detection

The search for constellation points as defined by (5) may be performed based on finding intersections of circles (cf. [3]). While being very efficient for M-PSK signal sets, this approach requires calculations of $tan^{-1}(\cdot)$ and $cos^{-1}(\cdot)$, several times per detection for M-QAM constellations that are usually employed for high spectral efficiency transmission.

Finding the constellation points fulfilling (5) may in fact be achieved in a very simple manner. We propose to approximate the search circle by a square (cf. Figure 1). This approach allows for finding candidates by using simple boundary controls: with $a = \Re\{c_{c,i}\}$ and $b = \Im\{c_{i,c}\}$ we only need to find all constellation points x_i for which

$$ \begin{aligned} |\Re\{x_i\} - a| &< R_{c,i}/r_{i,i}, \quad \text{and} \\ |\Im\{x_i\} - b| &< R_{c,i}/r_{i,i}. \end{aligned} \qquad (7) $$

Since $\Re\{x_i\}$ and $\Im\{x_i\}$ usually take on only very few discrete values, this is easily done by checking these against $a \pm R_{c,i}/r_{i,i}$ and $b \pm R_{c,i}/r_{i,i}$, respectively. The list of tentative candidates can then be constructed by combining all $\Re\{x_i\}$ and $\Im\{x_i\}$ lying within these bounds (cf. Figure 1 – two possible values per dimension result in 4 candidates).

Note that all points that lie inside the square but outside the circle (and hence are not valid candidates) can be easily discarded when calculating the remainder search radius $R_{c',i+1}(x_i)$ for the extended candidate. Namely, $R_{c',i+1}(x_i)$ is defined by

$$ \begin{aligned} R^2_{c',i+1}(x_i) &= R^2_{c,i} - |c_{c,i} - x_i|^2 r^2_{i,i} \\ &= R^2_{c,i} - \left(|\Re\{x_i\} - a|^2 + |\Im\{x_i\} - b|^2 \right) r^2_{i,i}. \end{aligned} $$

It is easily seen that for all points that lie inside the square but outside the circle, $R^2_{c',i+1}(x_i)$ will be negative. The additional complexity due to excess points is negligible, since the values for $|\Re\{s_k\} - x_c|^2$ and $|\Im\{s_k\} - y_c|^2$ have to be calculated anyway for the valid points and can be used to identify the invalid one(s). At the same time it is straightforward to see that the number of excess points diminishes with decreasing search radii (hence, increasing SNR). This notion is confirmed by the results in Figure 2 – the additional overhead is negligible in the high SNR regime.
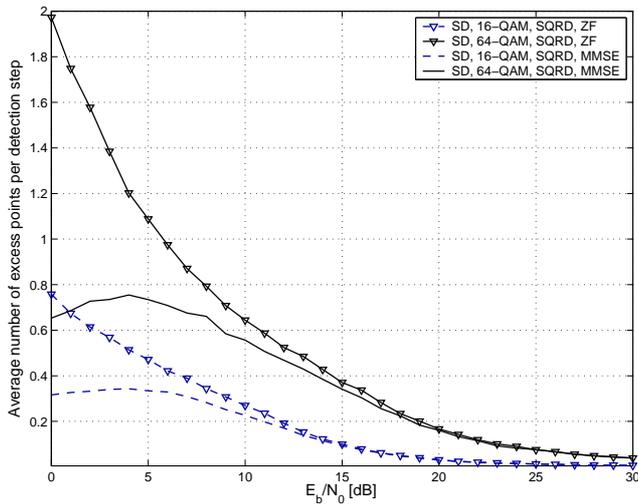
Fig. 2. Average number of excess constellation points found per detection step, for a 4×4 MIMO system, as a function of SNR and constellation size. In the high SNR regime, the number of additionally found points is negligible. It is also visible that by the MMSE extension of SQRD and the resulting reduction of the effective search radii in the different layers, this number can be further reduced.

E. QRD-M and SQRD-M

Based on the "M-algorithm" [8], the concept of "QRD-M" was discussed in [9]. The basic idea is similar to SQRD approaches for MIMO detection. However, instead of selecting only the closest constellation point in each dimension, a total of M metrics is considered in evaluation.

We propose a joint application of the two approaches from [7] and [9] to create a SD-like detector and denote the resulting algorithm "SQRD-M". The basic idea is to use a SD with MMSE filtering and a SQRD, but to fix the number of evaluated constellation points in each dimension to the closest M points. The main advantage over conventional SD is that $N_c = M^{M_{Tx}}$ is no longer a random variable, ensuring a fixed delay of the detector. However, restricting the search in such a fashion will lead to a performance deterioration with respect to MLD, since we incorporate no flexibility in the search algorithm and might thus exclude the part from the search tree in which the ML solution is to be found.

III. COMPLEXITY

We limit ourselves to the complexity required to find only the ML solution. We also assume the channel to be static over a sufficiently long period of time, such that the computational complexity of any preprocessing steps (decomposition of the channel matrix, layer ordering) is negligible. This approximation is motivated by a further fact: since we use the SQRD equivalently for SIC and SD, the preprocessing complexity of both receivers is obviously equal. Furthermore, since a QR decomposition may also be used to calculate an inverse matrix, the same holds for linear detection. Knowing that the detection complexity of SD can be expected to be higher than that of linear and SIC receivers, our comparison is in fact the worst case scenario for SD – the higher the complexity of the preprocessing steps in comparison to the actual detection, the lower the relative complexity of SD versus linear and SIC receivers. For purposes of exhibition, we assume in the following (complex) multiply-add-compare (MAC) operations and slicing to be of equal complexity.

A. Linear receiver

A linear receiver multiplies the received signal vector with a weighting matrix $\mathbf{G} = \mathbf{H}^+$ and feeds the output to a slicer.

The resulting detection complexity for a single transmitted vector symbol is hence

$$O_L = M_{Tx}^2 + M_{Tx}. \tag{8}$$

B. Successive Interference Cancellation

We concentrate on SIC approaches based on a QR-decompostion of $\mathbf{H}$ [6]. After multiplying the received signal with $\mathbf{Q}^{\mathbf{H}}$, the upper triangular matrix $\mathbf{R}$ is used for successive detection and interference cancellation. When detecting layer i, the signal estimates from all precedingly detected layers are weighted and subtracted from the received signal, the output weighted by $r_{i,i}$ and fed to a slicer. The overall detection complexity per vector symbol is hence

$$O_{SQRD} = M_{Tx}^2 + \sum_{i=1}^{M_{Tx}} \left((i-1) + 2 \right) = \frac{3}{2}(M_{Tx}^2 + M_{Tx}) \tag{9}$$

C. Sphere Detection

In each dimension, the detection problem comes down to finding all candidates $\hat{x}_i$ fulfilling (5) and updating R_c appropriately (cf. [5]). The detection complexity for sphere decoding (calculating the search center and radius, slicing operation, update of the remaining search radius for each found candidate) can then be shown to be:

$$O_{Sphere} = M^2 + \sum_{i=1}^{M} \left((i-1)N_{i-1} + 2 + N_i \right) \tag{10}$$

where N_i is the number of candidate points found in dimension i. Throughout our simulations, we generate statistics on N_i which enables us to derive O_{Sphere}. The additional complexity of running the SD multiple times before finding $N_{c,min}$ candidates is captured in these statistics.

D. Simulation results

Figure 3 shows the computational complexity of detection for different MIMO receivers relative to a linear detector, as a function of the SNR, for uncoded 16-QAM and 64-QAM transmission ($N_{c,min} = 1$ for SD).
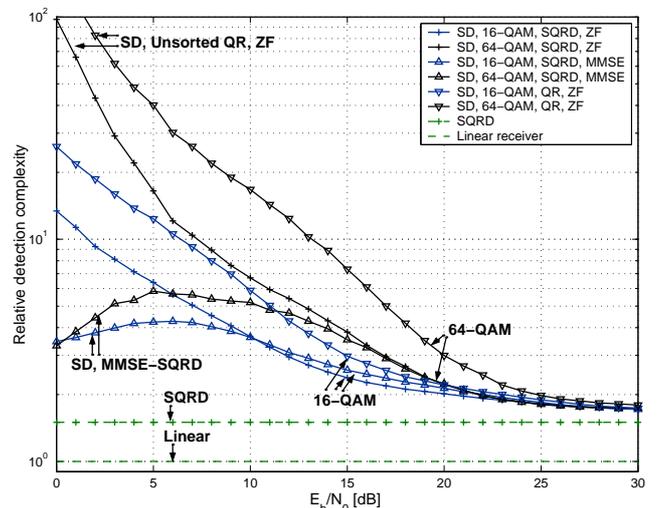
Fig. 3. Relative complexity of different MIMO receiver techniques, for a 4x4 antenna system. The complexity of linear and SQRD receivers is independent of constellation size and SNR. Complexity of sphere decoding very quickly diminishes as the SNR increases and is comparable to that of SQRD, in the high SNR regime.

We find that appropriate layer ordering (see results for unsorted QR vs. SQRD) is able to significantly lower complexity for SD, especially for higher order modulation. Even more pronounced is the effect of using the MMSE extension to remove close-to-singular values from $\mathbf{R}$ – the exponential increase in complexity

for lower SNR regions is thus avoided and detection complexity of SD limited to 6 times that of linear receivers, even under worst case conditions (64-QAM, SNR 3dB). This stability is especially important when considering an actual implementation where the complexity of detection must always be upper bounded in complexity.

When considering uncoded 16-QAM transmission, we see that in the regime of interest (e.g. at BER 10^{-3}, SNR = 16 dB) the complexity of SD is roughly double that of SQRD, while offering SNR gains of roughly 7 dB over BLAST (cf. Figure 5). For 64-QAM transmission, the gain is 9 dB (at 21dB, results for SQRD not shown) with only 30% increase in complexity. Manipulating the constant factor in (6) leads to better results in this SNR regime, at the expense of reduced stability, i.e., higher complexity in the low SNR regime. This motivates for further research on the selection of an appropriate initial search radius.
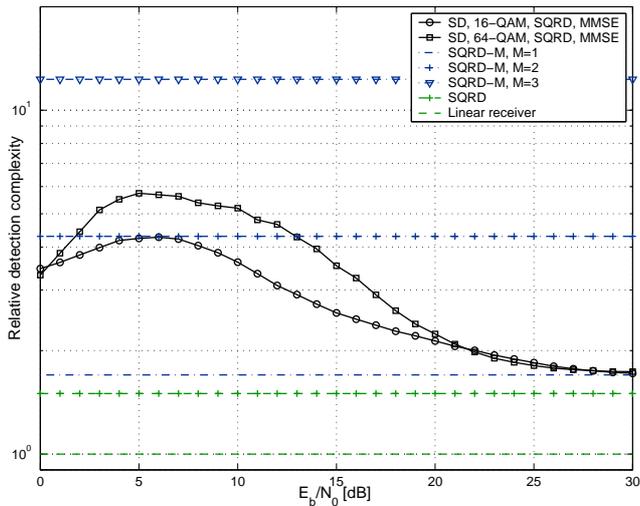
Fig. 4. Relative complexity of different SD and SD-like receiver techniques, for a 4x4 antenna system. The complexity of SQRD-M is independent of constellation size and SNR and slightly higher than for SQRD due to the required radius updates. However, complexity of SD is always inferior to SQRD-M, $M = 2$ for 16-QAM and for 64-QAM at least in the regime of interest (SNR > 15 dB).

Figure 4 compares the relative computational complexity of MMSE-SQRD based SD with that of SQRD-M, for 64-QAM. While having a complexity independent of SNR and constellation size, it is evident that SQRD-M approaches are not to be favored when considering only uncoded transmission. Especially values of $M \geqslant 3$ appear to be prohibitive in complexity. Note that the complexity of SD for 16-QAM is always inferior to that of SQRD-M with $M = 2$ while the former can be expected to have superior performance (cf. Figure 6 for the 64-QAM case). The complexity of SQRD-M, $M = 1$ is slightly higher than that of standard SQRD due to calculation of the remainder search radius, which eventually is used to calculate L-values at the output of the detector.

IV. PERFORMANCE

A. Uncoded Transmission

Figure 5 shows the bit error performance of different MIMO receiver architecturs for uncoded transmission, for 16-QAM and 64-QAM. It is clearly visible that linear as well as SIC based receivers are limited to first order diversity and are hence significantly outperformed by SD in the high SNR regime. Results for ZF as well as MMSE based SD are given to illustrate the fact that for SD the difference between the ZF and MMSE solution is to be found in the complexity domain, instead of the bit error performance. Both achieve a performance equaling that of MLD.

Gains from SD over BLAST at BER 10^{-3} are 7 for 16-QAM and 9 dB for 64-QAM (results not shown). These gains can be expected to increase by 7.5 dB every time we decrease the target BER by one order of magnitude, owed to the difference in the diversity order.
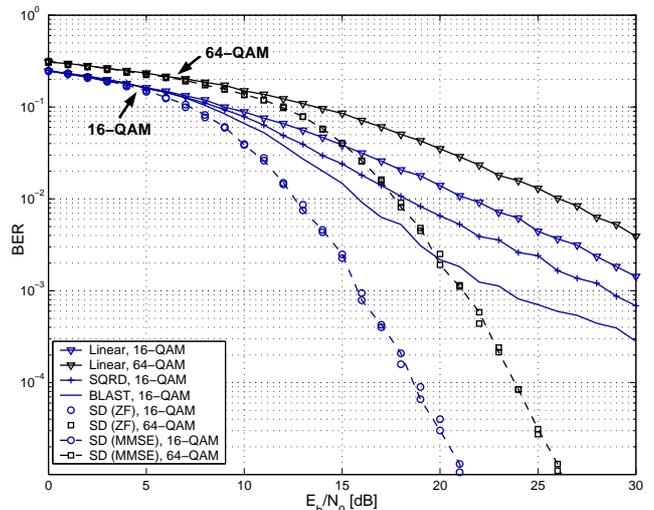
Fig. 5. Uncoded bit error performance of different receiver techniques for a 4x4 MIMO system using 16-QAM and 64-QAM. Spectral efficiency is 16 and 24 bit/channel use, respectively. SD clearly outperforms all other techniques in the moderate to high SNR regime and is the only technique achieving diversity in the number of receive antennas.

Figure 6 compares the bit error performance of SQRD-M with that of SD for uncoded 64-QAM transmission. As expected the SQRD-M approach is limited to first order diversity and SQRD-M, $M = 1$ is equivalent to standard SQRD SIC (cf. Figure 5). With increasing M, performance improves slightly but still falls far short of MLD performance. In the light of the detection complexity (cf. Figure 4), SQRD-M is no alternative to reduced complexity sphere detection.
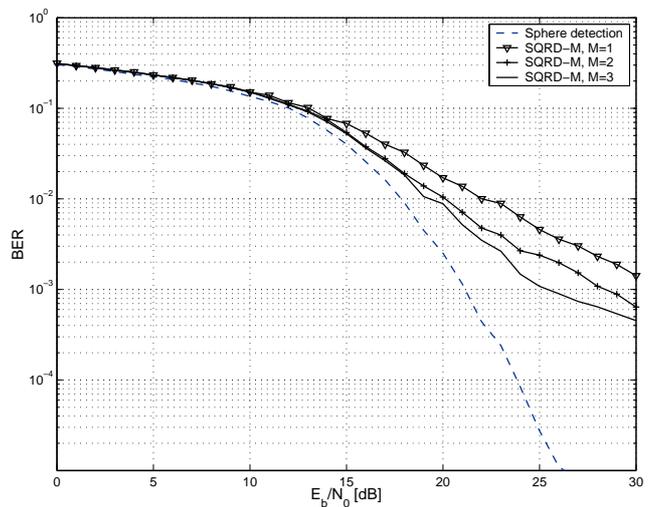
Fig. 6. Uncoded bit error performance of different sphere-like detectors for a 4x4 MIMO system using 64-QAM. M-SQRD is clearly limited to first order diversity and shows good performance only in the low-medium SNR regime, even for high numbers of M.

B. Coded Transmission

For coded transmission, we use the standard "off-the-shelf" Turbo Code (code polynoms $g = (7, 5)_o$ for constituent encoders)

without puncturing, yielding a code rate of $R_c = 1/3$. The block length is 9216 bits and we allow for 4 internal decoder iterations.

Figure 7 shows the bit error performance of SD as a function of the number of minimum candidates $N_{c,min}$ for coded 16-QAM and 64-QAM transmission. As expected, performance improves as we increase $N_{c,min}$. Note that the actual number of candidates may be significantly higher than the minimum number of candidates. This is owed to the fact that we increase the search radius whenever we find too few candidates – and then may find too many. We found that for higher $N_{c,min}$ and our selected initial search radius and radius increment, $E\{N_c\} \approx 3N_{c,min}$ in the regime of interest.

This also motivated for our selection of $N_{c,min}$ – with $E\{N_c\}$ being in the order of $3N_{c,min}$, requiring a minimum of 8 and 32 candidates yields a detection complexity comparable to that of SQRD-M, $M = 2$ and $M = 3$, respectively.
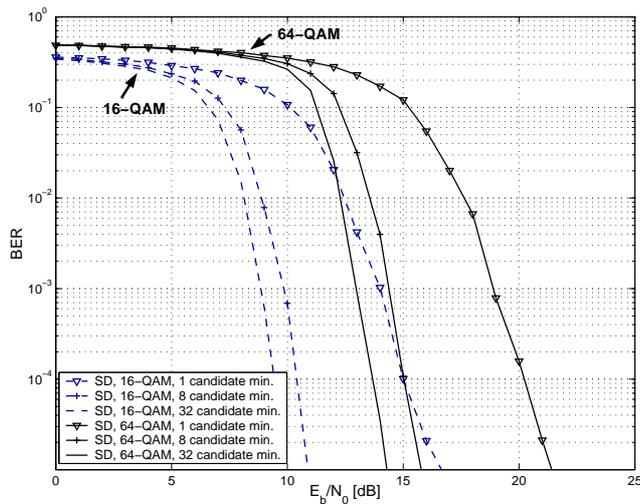
Fig. 7. Bit error performance of SD for $R_c = 1/3$ coded transmission in a 4x4 MIMO system using 16-QAM and 64-QAM. Spectral efficiency is 5.3 and 8 bit/channel use, respectively. Performance increases with rising number of required candidates $N_{c,min}$.

Figure 8 compares the error performance of SQRD-M with that of SD for coded 64-QAM transmission. The performance for SQRD-M, $M = 1$, and SD using $N_{c,min} = 1$ is comparable. Remember that this implies that coded SQRD-SIC and SD with only a single candidate have also comparable performance.

However, Figure 8 also reveals that the picture changes as soon as higher number of candidates are required to improve the quality of the output of the detector. Increasing the number of branches M in the SQRD-M technique over $M = 2$ does not yield any benefit, while detection complexity is substantially increased. The loss with respect to MMSE-SQRD based SD, $N_{c,min} = 8$ is roughly 2 dB. This result indicates that SQRD-M is not a viable alternative to SD, due to its lack of adaptivity.

V. RESULTS AND DISCUSSION

We evaluated the computational complexity required for detection as well as the performance of different MIMO receiver techniques for uncoded as well as Turbo coded transmission. Our results show that for uncoded transmission and low target BERs, sphere detectors have only minor additional complexity in comparison to other known receiver algorithms, while significantly outperforming them all in terms of bit error performance. We also proposed a number of techniques to reduce the complexity of sphere detection, most notably in the low SNR regime. This allows for using SD under a wide range of environments.

For coded transmission, the complexity of sphere detection grows linearly with the minimum number of candidates. So far,
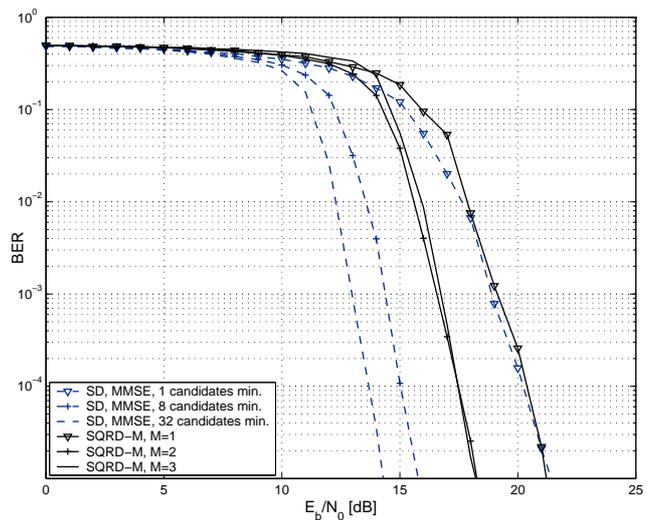
Fig. 8. Bit error performance of different SD for $R_c = 1/3$ coded transmission in a 4x4 MIMO system 64-QAM. M-SQRD loses roughly 2 dB with respect to conventional SD. Remarkably, increasing M does not improve performance.

rather high numbers of candidates are required to achieve good performance. This motivates for further research to find techniques that provide good soft outputs requiring only a low number of candidates.

REFERENCES

[1] P. Wolniansky, C. Foschini, G. Golden, and R. Valenzuela, "V-BLAST: an architecture for realizing very high data rates over the rich-scattering wireless channel," in *International Symposium on Signals, Systems and Electronics ISSSE98*, pp. 295–300.

[2] E. Viterbo and J. Boutros, "A Universal Lattice Code Decoder for Fading Channels," *IEEE Transactions on Information Theory*, vol. 45, no. 5, pp. 1639–1642, July 1999.

[3] B. M. Hochwald and S. ten Brink, "Achieving Near-Capacity on a Multiple-Antenna Channel," *IEEE Transactions on Communications*, vol. 51, no. 3, pp. 389–399, Mar. 2003.

[4] H. Vikalo and B. Hassibi, "The Expected Complexity of Sphere Decoding, Part I: Theory, Part II: Applications," *IEEE Transactions on Signal Processing, submitted for publication*, 2003.

[5] O. M. Damen, H. E. Gamal, and G. Caire, "On the Complexity of ML Detection and the Search for the Closest Lattice Point," *IEEE Transactions on Information Theory*, vol. 59, no. 10, pp. 2400–2414, Oct. 2003.

[6] D. Wuebben, J. Rinas, R. Boehnke, V. Kuehn, and K. Kammeyer, "Efficient Algorithm for Detecting Layered Space-Time Codes," in *4th International ITG Conference on Source and Channel Coding*, Berlin, Germany, Jan. 2002.

[7] D. Wuebben, R. Boehnke, V. Kuehn, and K. Kammeyer, "MMSE Extension of V-BLAST based on Sorted QR Decomposition," in *IEEE Semiannual Vehicular Technology Conference (VTC2003-Fall)*, Orlando, USA, Oct. 2003.

[8] L. Wei, L. Rasmussen, and R. Wyrwas, "Near optimum tree-search detection schemes for bit-synchronous multiuser CDMA systems over Gaussian and two-path Rayleigh-fading channels," *IEEE Transactions on Communications*, vol. 45, pp. 691–700, June 1997.

[9] K. J. Kim and R. A. Iltis, "Joint detection and channel estimation algorithms for QS-CDMA signals over time-varying channels," *IEEE Transactions on Communications*, vol. 50, pp. 845–855, May 2002.