Optimized Block-Diagonalization Precoding Technique Using Givens Rotations QR Decomposition

Victor-Florin Crâşmariu*, Marius-Octavian Arvinte, Andrei-Alexandru Enescu
Digital Networking
NXP Semiconductors
Bucharest, Romania
**e-mail: victor-florin.crasmariu@nxp.com*

Silviu Ciochină, *Member*, *IEEE*
Telecommunications Department
University Politehnica of Bucharest
Bucharest, Romania

Abstract—**The emerging 5G mobile communication standard aims to increase the throughput and, in the same time, to considerably increase the number of users serviced concurrently (Internet of Things). The key direction for achieving these requirements is to heavily extend the use of the spatial degrees of freedom, especially in the multi-user scenarios.**

Multi-User Massive MIMO (MU-MIMO) is one of the key technologies that responds well to the 5G needs. However, the use of MU-MIMO in the downlink direction raises the collaborative detection problem at the user side, thus the elimination of the inter-user interference becomes necessary.

The paper presents a reduced complexity linear transmitter precoding technique that cancels the inter-user interference in a downlink MU-MIMO system.

The reduced complexity is achieved through re-using as many low-level operations as possible. The method is suitable for implementation on any modern processor and proven to be scalable to a Massive MIMO scenario without any loss in performance.

Keywords—MU-MIMO, Massive MIMO, Precoding, Reduced complexity, Block-Diagonalization

I. INTRODUCTION

Multi-User Massive MIMO, where the same time-frequency resources are simultaneously used in communicating with a large number of users is one of the key enablers for 5G cellular communications [1]. While the asymptotic gain in capacity offered by MU-Massive MIMO is large, eliminating inter-user interference is a key aspect of exploiting this gain [2].

The linear Block-Diagonalization precoding is a well-known method for compensating the inter-user interference in MU-MIMO scenarios [3, 4, 5] leading to the effective decomposition in multiple parallel SU-MIMO scenarios, where another stage of precoding can be applied per user [6].

BD is essentially a generalization of the Zero Forcing (ZF) precoding scheme for a multi-user scenario [4] and has been extensively studied since its introduction in the literature. [7] introduces an improved performance modification of BD that takes the noise covariance into account as well, [8] introduces a low-complexity hybrid BD scheme suitable for Massive MIMO, while [9] investigates the scalability of BD precoding in a Massive MIMO scenario in terms of floating-point operation (FLOP) count and power usage for a typical modern digital signal processor and acknowledges that the required computational power is considerably large in terms of used number of processors.

While many variations of BD exist in the literature, the core operation in each case remains the QR decomposition [10]. The present paper describes a method of reducing the complexity of the BD precoding by reusing as many elementary operations as possible when successively computing the precoding matrix for every user through QR decompositions. This is achieved by noticing that the order in which the per-user matrices are computed affects the overall cost of the BD method. Closed-form expressions for the achieved reduction of complexity are derived and an optimal ordering of the users is presented, noting that the use of the proposed algorithm incurs no performance loss whatsoever.

The paper is organized as follows: the first section introduces the MU-MIMO inter-user interference problem and the BD solution to it. The second section describes the equations associated with a MU-MIMO system that uses BD precoding. The third section presents the proposed complexity reduction method in two different flavors and the fourth section performs a complexity analysis to evaluate the achieved reduction. The fifth and final section presents the conclusions and possible further directions of research.

II. BLOCK-DIAGONALIZATION MU-MIMO TECHNIQUE

Assuming a downlink MU-MIMO scenario involving one base station (BS) communicating with K mobile stations (users) the narrowband, flat-fading MIMO channel equation that describes the system is

$$\mathbf{y} = \mathbf{H}\mathbf{s} + \mathbf{w}, \tag{1}$$

where:

$\mathbf{s}$ is an $n_t \times 1$ vector formed by the symbols transmitted by the n_t antennas at the BS,

$\mathbf{y}$ is an $n_r \times 1$ vector formed by the received symbols, where n_r represents the sum of all the K user antennas and is given by $n_r = \sum n_{r,i}$, where $n_{r,i}$ stands for the number of antennas of the i-th user,

ISBN 978-0-9928626-7-1 © EURASIP 2017

w is an $n_r \times 1$ vector that represents the AWGN noise at each user antenna,

H is an $n_r \times n_t$ matrix that models the narrowband, flat-fading MIMO channel. The transmitter is assumed to have perfect channel knowledge.

Equation (1) can be expanded using a precoding operation at the BS, that generally aims to adapt the transmitted symbols to the channel conditions to obtain either a better performance and/or a reduced number of signal processing operations that must be performed in the receiver chain.

The precoding operation is modeled via the use of a precoding matrix that maps the data symbols on the n_t transmit antennas. The total number of symbols that can be mapped is n_l and is equal to the total number of spatial paths (otherwise termed *layers*) that can be employed in the transmission ($n_l \leq n_t$). Let $n_{l,i}$ denote the number of layers assigned to the i-th user.

When using a precoding operation, (1) becomes

$$\mathbf{y} = \mathbf{H}\mathbf{P}\mathbf{x} + \mathbf{w}, \qquad (2)$$

where:

x is the n_l x 1 vector formed by the transmitted symbols,

P is the $n_t \times n_l$ precoding matrix.

The **HP** (hereinafter denoted as $\mathbf{H}_e$) product is the equivalent channel that is created after the precoding operation and which is experienced by the elements of the **x** vector.

Block-Diagonalization precoding [4] is a method of suppressing the inter-user interference and providing similar performance to that of multiple parallel SU-MIMO scenarios [6], particularly when multi-antenna mobile stations are involved in the communication.

Denoting $\mathbf{P}_{BD}$ as the BD precoding matrix and $\mathbf{H}_{e,BD} = \mathbf{H}\mathbf{P}_{BD}$, then

$$\mathbf{H}_{e,BD} = \begin{bmatrix} \mathbf{H}_{e,1} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{H}_{e,2} & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{H}_{e,K} \end{bmatrix} \qquad (3)$$

where **0** represents the null matrix and $\mathbf{H}_{e,i}$ is the equivalent SU-MIMO link virtually created for the i-th user.

The $\mathbf{P}_{BD}$ matrix can be viewed as being formed from K groups of column vectors, where each group is formed of $n_{l,i}$ beamforming vectors that belong to the null space of the complementary channel matrix of the i-th user $\tilde{\mathbf{H}}_i$.

There are multiple ways of deriving the null space of a given matrix, the most common ones being the QR factorization and the singular value decomposition (SVD) [10]. The work in this paper focuses on reducing the computational complexity of the BD precoding method when used in conjunction with the QR factorization.

To find the null space of $\tilde{\mathbf{H}}_i$ the QR factorization is applied to its Hermitian

$$\tilde{\mathbf{H}}_i^{\mathbf{H}} = \mathbf{Q}_i \begin{bmatrix} \mathbf{R}_i \\ \mathbf{0} \end{bmatrix}, \qquad (4)$$

where:

$\mathbf{Q}_i$ is an n_t x n_t unitary matrix,

$\mathbf{R}_i$ is an $(n_r - n_{r,i})$ x $(n_r - n_{r,i})$ upper triangular matrix.

The null space of the $\tilde{\mathbf{H}}_i$ matrix is thus given by the rightmost $n_t - n_r + n_{r,i}$ column vectors of the $\mathbf{Q}_i^{\mathbf{H}}$ matrix [9] and the maximum number of layers that can be transmitted to the i-th user is

$$n_{l,i} = \min\left(n_{r,i}, n_t - n_r + n_{r,i}\right). \qquad (5)$$

Let $\mathbf{Q}_i^{(0)}$ be the partition of the $\mathbf{Q}_i$ matrix that forms a null space for the $\tilde{\mathbf{H}}_i$ matrix. Let $\tilde{\mathbf{Q}}_i^{(0)}$ denote the set of $n_{l,i}$ column vectors that are being selected out from the $\mathbf{Q}_i^{(0)}$ matrix to be used as precoding for the i-th user. As such, the $\mathbf{P}_{BD}$ precoding matrix is

$$\mathbf{P}_{BD} = \begin{bmatrix} \tilde{\mathbf{Q}}_1^{(0)} & \tilde{\mathbf{Q}}_2^{(0)} & \cdots & \tilde{\mathbf{Q}}_K^{(0)} \end{bmatrix}. \qquad (6)$$

The work in this paper uses the Givens rotations algorithm for performing the QR decomposition. The Givens algorithm performs successive multiplications of the input matrix with a set of elementary rotational matrices denoted $\mathbf{T}_{p,q}$. The effect of multiplying with a $\mathbf{T}_{p,q}$ matrix is nullifying the (p, q)-th element of the resulted matrix. Then

$$\mathbf{T}_{p,q} = \begin{bmatrix} 1 & \cdots & 0 & \cdots & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & \cdots & \vdots & \cdots & \vdots \\ 0 & \cdots & c_{p,q} & \cdots & s_{p,q} & \cdots & 0 \\ \vdots & \cdots & \vdots & \ddots & \vdots & \cdots & \vdots \\ 0 & \cdots & -s_{p,q}^* & \cdots & c_{p,q}^* & \cdots & 0 \\ \vdots & \cdots & \vdots & \cdots & \vdots & \ddots & \vdots \\ 0 & \cdots & 0 & \cdots & 0 & \cdots & 1 \end{bmatrix}, \qquad (7)$$

where:

$$c_{p,q} = \frac{|a_{q,q}|}{\sqrt{|a_{q,q}|^2 + |a_{p,q}|^2}}, \quad s_{p,q} = \frac{a_{q,q} a_{p,q}^*}{|a_{q,q}|\sqrt{|a_{q,q}|^2 + |a_{p,q}|^2}}, \qquad (8)$$

and $a_{i,j}$ is the (i, j)-th element of the input matrix upon which the Givens rotation is applied, denoted here as **A**.

The $\mathbf{T}_{p,q}$ matrices are applied successively in a column-wise manner, first nullifying the elements within the leftmost column of the input matrix and progressively advancing towards the rightmost column,

$$\mathbf{R}^{(n)} = \mathbf{T}^{(n)}\mathbf{R}^{(n-1)}, \qquad (9)$$

where:

ISBN 978-0-9928626-7-1

n is the current iteration index,

$\mathbf{R}^{(0)}$ is initialized with the input matrix,

$\mathbf{T}^{(n)}$ is the elementary Givens rotation matrix that is applied at the n^{th} iteration used to cancel the (p,q)-th element.

The algorithm is complete when all the sub-diagonal elements of the input matrix have been canceled. The $\mathbf{Q}$ matrix within the QR factorization method is obtained similarly as

$$\mathbf{Q}^{(n)} = \mathbf{T}^{(n)}\mathbf{Q}^{(n-1)}, \tag{10}$$

where $\mathbf{Q}^{(0)} = \mathbf{I}_N$, with N denoting the number of rows of the input matrix.

The following observations can be made about the $\mathbf{T}$ matrices:

i) The $\mathbf{T}_{p,1}$ matrices used to nullify the elements under the main diagonal of the first column of the input matrix are derived from the first column of the input matrix i.e. $\mathbf{T}_{p,1} = f(\mathbf{c}_1)$, where $\mathbf{c}_i$ denotes the i-th column of the input matrix.

ii) The $\mathbf{T}_{p,2}$ matrices used to nullify the elements under the main diagonal of the second column are derived from the second column of the input matrix. However, when the algorithm advances to the second column, it will find the input matrix already affected by the $\mathbf{T}_{p,1}$ matrices. Thus $\mathbf{T}_{p,2}$ also depends indirectly on the first column of the input matrix i.e. $\mathbf{T}_{p,2} = f(\mathbf{c}_1, \mathbf{c}_2)$.

iii) Following the same principle we have $\mathbf{T}_{p,3} = f(\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3)$, $\mathbf{T}_{p,4} = f(\mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_3, \mathbf{c}_4)$, etc.

The following proposition is thus true:

Proposition 1:

Given any two input matrices, **A** and **B**, having the leftmost k-1 columns equal, then:

$$\mathbf{T}^{A}_{p,q} = \mathbf{T}^{B}_{p,q}, \forall q < k,$$

where $\mathbf{T}^{A}_{p,q}$ and $\mathbf{T}^{B}_{p,q}$ denote the elementary rotation matrices applied to **A** and **B** respectively.

III. Proposed Algorithms

The present paper aims to exploit Proposition 1 in order to identify reduced complexity methods for computing the BD precoding matrix. This is achieved by reusing as many elementary rotation matrices as possible from one user to another when computing the null space.

For any given multiuser scenario with $K > 2$, the $\tilde{\mathbf{H}}_i$ and $\tilde{\mathbf{H}}_j$ matrices of any two users have in common multiple rows of the $\mathbf{H}$ channel matrix opens the possibility of reusing some of the previous operations when the BD precoding algorithm advances to the next user. In the rest of the section three optimized BD precoding matrix computation algorithms based on the Givens rotations method are proposed.

A. Algorithm 1 – Option A

This method uses of the Givens rotations method for finding the null-space of the user-specific $\tilde{\mathbf{H}}_i$ matrix by computing them in descending order relative to how the users are arranged in the channel matrix i.e. $i = \{K, K\text{-}1, \ldots, 1\}$.

TABLE I. Algorithm 1A pseudocode

```
for i=K to 1 do
    compute H̃_i from H
    if (i=0 or i=K) then // no reuse of T_{p,q} matrices expected
        apply the Givens rotation algorithm to the H̃_0^H matrix to
        obtain its null space
    else // reuse of previous T_{p,q} matrices to be assessed
        if (leftmost common columns exist) then
            save k as the position of the leftmost column that differs
            between H̃_i^H and H̃_{i-1}^H
        end if
        apply the Givens rotation algorithm to the H̃_i^H matrix to
        obtain its null-space, starting from column k; reuse all the T_{p,q}
        matrices with q < k from iteration i-1
    end if
    store the T_{p,q} matrices for potential later use
    update P_BD with n_{l,i} column vectors selected from the null-space of
    H̃_i^H
end for
```

As explained in the previous section, the null space of the $\tilde{\mathbf{H}}_i$ matrix can be found by applying the QR factorization method to its Hermitian $\tilde{\mathbf{H}}_i^{\mathbf{H}}$.

When the algorithm advances from one user to another, an assessment of the common columns between $\tilde{\mathbf{H}}_i^{\mathbf{H}}$ and $\tilde{\mathbf{H}}_{i-1}^{\mathbf{H}}$ is performed, searching for identical columns residing adjacently on the leftmost positions of the two involved matrices (corresponding to the current and previous users). If any such columns exist, then the algorithm reuses the $\mathbf{T}_{p,q}$ matrices that were applied to null out the sub-diagonal elements of these columns in the $\tilde{\mathbf{H}}_{i-1}^{\mathbf{H}}$ matrix.

Algorithm 1A is summarized below in the form of a pseudocode.

Depending on the system parameters (n_t, K, $n_{r,i}$), the number of $\mathbf{T}_{p,q}$ matrices that need to be computed at the i-th iteration (apart from the reused ones) is

$$f_c^{1A}(i) = \begin{cases} 0, & \forall i \in \{1, K\} \\ N_T(i) - \dfrac{\left(2n_t - \sum_{j=1}^{i-1} n_{r,j} - 1\right)\sum_{j=1}^{i-1} n_{r,j}}{2}, & \forall i \notin \{1, K\} \end{cases} \tag{11}$$

where $N_T(i)$ gives the total number of $\mathbf{T}_{p,q}$ matrices that are needed at the i-th iteration:

$$N_T(i) = \left(n_r - n_{r,i}\right)\left(n_t - \frac{n_r - n_{r,i} + 1}{2}\right). \tag{12}$$

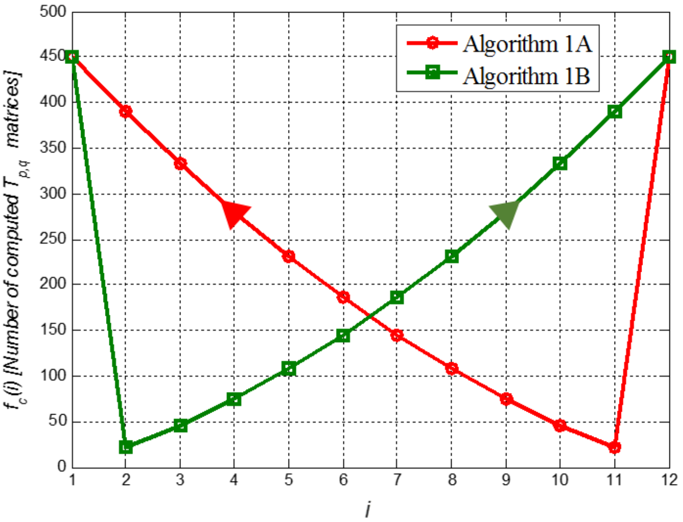

Figure 1. The cost functions for Algorithm 1A and Algorithm 1B. The MU-MIMO system parameters are: $K = 12$, $n_t = 32$, $n_{r,i} = 2$.

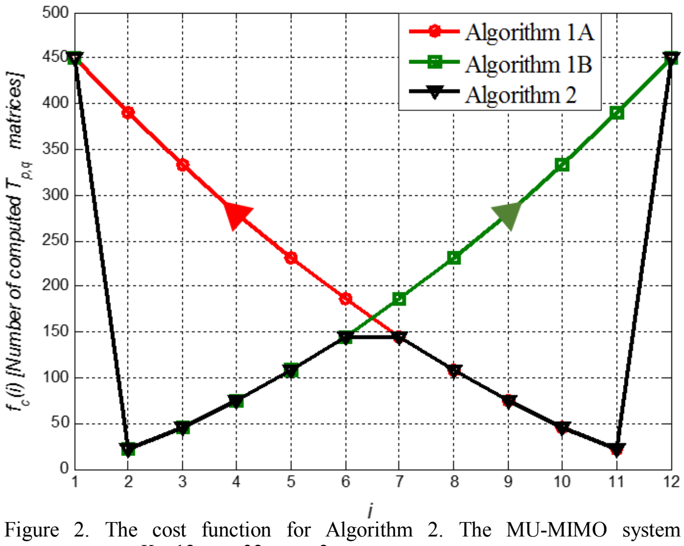

Figure 2. The cost function for Algorithm 2. The MU-MIMO system parameters are: $K = 12$, $n_t = 32$, $n_{r,i} = 2$.

$f_c^{1A}(i)$ is a monotonically decreasing sequence for $i \in \{2,3,...,K-1\}$. The maximum degree of reusing the $\mathbf{T}_{p,q}$ matrices is achieved when the cost function achieves its minimum i.e. for the (K-1)-th user, while the lowest reuse (excluding the first and the last users, where the QR decomposition needs to be computed in its entirety) is achieved for the second user.

Algorithm 2 is given in pseudocode form in Table II.

B. *Algorithm 1 – Option B*

This represents an alternative version of the first algorithm that provides the same degree of reuse and creates the premises for the final algorithm that will be described in section C. Algorithm 1B operates similarly to algorithm 1A, with two major differences:

i) The users are now iterated in ascending order.

ii) The channel matrix is up-down flipped and denoted hereinafter as $\mathbf{H}^{(f)}$. The user specific matrices that are derived from $\mathbf{H}^{(f)}$ are denoted as $\tilde{\mathbf{H}}_i^{(f)}$ with the mention that $\mathbf{P}_{BD}$ can be derived from the $\tilde{\mathbf{H}}_i^{(f)}$ matrices in absolutely the same way it would have been derived from the $\tilde{\mathbf{H}}_i$ matrices.

The cost function of algorithm 1B is

$$f_c^{1B}(i) = \begin{cases} 0, & \forall\, i \in \{1, K\} \\ N_T(i) - \dfrac{\left(2n_t - \sum\limits_{j=i+1}^{K} n_{r,j} - 1\right) \sum\limits_{j=i+1}^{K} n_{r,j}}{2}, & \forall i \notin \{1, K\} \end{cases} \quad (13)$$

In contrast with $f_c^{1A}(i)$, $f_c^{1B}(i)$ is a monotonically increasing sequence for $i \in \{2,3,...,K-1\}$. The maximal reuse of the $\mathbf{T}_{p,q}$ matrices is achieved for the second user.

Fig. 1 plots the cost functions of algorithms 1A and 1B for a given MU-MMIMO scenario. The arrows depict the order in which the per-user precoding matrices are computed.

C. *Algorithm 2*

The final proposed algorithm combines the two versions of the first algorithm in order to obtain a lower cost function for each user. The two functions introduced earlier have opposite monotonic tendencies for $i \in \{2,3,...,K-1\}$, intersecting in at most one point.

Algorithm 2 introduces the possibility to switch the processing from algorithm 1A to algorithm 1B at the user index where algorithm 1B starts providing a better $\mathbf{T}_{p,q}$ matrices reuse degree i.e. a lower cost function. Thus, the cost function for algorithm 2 combines discrete portions of $f_c^{1A}(i)$ and $f_c^{1B}(i)$.

TABLE II. ALGORITHM 2 PSEUDOCODE

i=K // start with algorithm 1A
while (*not reached the receive antenna midpoint*) **do**
run the *i*-th iteration of algorithm 1A
$i = i - 1$ *// the users are processed in descending order*
if (*midpoint has been reached*)
break
end if
end while
i=1 // switch to algorithm 1B
while (*users that haven't been processed exist*) **do**
run the *i*-th iteration of algorithm 1B *// use* $\mathbf{H}^{(f)}$ *and* $\tilde{\mathbf{H}}_i^{(f)}$ *matrices*
i = i+1 // the users are processed in ascending order
end while

The ideal switching point from algorithm 1A to algorithm 1B occurs at the user index where the inequality

ISBN 978-0-9928626-7-1 © EURASIP 2017

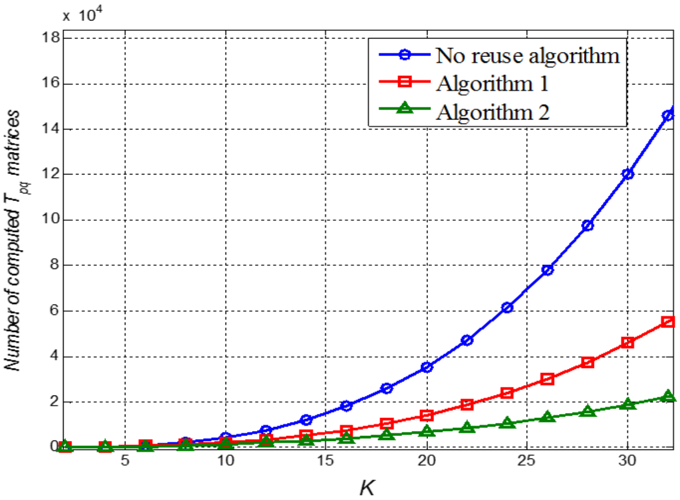

Figure 3. Complexity assessment of the number of Givens rotation matrices that need to be computed in each of the 3 cases: No reuse algorithm (classical approach), Algorithm 1 and Algorithm 2.

$$\sum_{j=i+1}^{K} n_{r,j} - \sum_{j=1}^{i-1} n_{r,j} < 0 \qquad (14)$$

is satisfied. When reaching the user corresponding to the midpoint of the total number of receive antennas, the switch from algorithm 1A to algorithm 1B takes place.

Considering the same MU-MIMO system dimensioning used in Fig. 1, then the cost function of algorithm 2 is plotted in Fig. 2. Algorithm 2 is given in pseudocode form in Table II.

IV. Complexity Analysis

To assess the benefits of the methods proposed in the present paper, the number of total $\mathbf{T}_{p,q}$ matrices that need to be computed to find the BD precoding matrix is analyzed. The results are compared to the case when no reuse algorithm would be in place.

The MU-MIMO system considered in the simulation is dimensioned as follows:

i) $n_t = 128$

ii) $n_{r,i} = 4$

iii) K = variable

Fig. 3 shows the simulation results that assess the total number of $\mathbf{T}_{p,q}$ matrices that had to be computed (i.e. that couldn't be reused from the previous iterations) to obtain the BD precoding matrix. The proposed algorithms are compared to the situation where the QR factorization is performed in a classical way and no $\mathbf{T}_{p,q}$ matrices reuse algorithm is used.

The simulations prove that the proposed methods bring an important complexity reduction through the reuse of the Givens rotations matrices from the previous iterations. Fig. 3 shows that for a 5G use case that involves a Massive MIMO system (large number of antennas at the BS side) and having 32 multi-antenna users participating in the MU-MIMO transmission, algorithm 1 achieves 62% reuse of the Givens rotations matrices, while algorithm 2 achieves 85%.

Simulations aimed at evaluating the link performance of the proposed methods, when compared to classical QR decomposition based approaches, have also been performed. The simulation results show no loss in performance when using the proposed algorithms. This is the expected behavior, as the proposed methods reproduce the same numerical outputs as when the QR decomposition method would be applied without any complexity optimizations.

Another cost of the proposed algorithms is that a real-world hardware or software implementation will require extra memory usage for storing the $\mathbf{T}_{p,q}$ matrices from one iteration to another. However, this drawback is ameliorated by the sparse shape of the $\mathbf{T}_{p,q}$ matrices. With suitably chosen hardware capabilities only two complex values per $\mathbf{T}_{p,q}$ matrix need to be stored, drastically lowering the required amount of memory.

V. Conclusions

Two low complexity methods of computing the BD precoding matrix have been presented. The proposed algorithms achieve a considerable reduction in computational complexity, while fully preserving the link performance.

The added algorithmic related complexity is low, consisting of simple logic and some memory management for storing the results from the previous iterations. Thus, the algorithms proposed in this paper raise minimal challenges in being implemented on modern-day processors.

The proposed BD precoding optimizations assume the use of the Givens rotations algorithm which is one of the most widely spread implementation for the QR factorization method [10], especially when used in conjunction with the CORDIC rotation algorithm in modern processors.

References

[1] T. L. Marzetta, "Noncooperative cellular wireless with unlimited numbers of base station antennas", IEEE Trans. on Wireless Communications, vol. 9, no. 11, pp. 3590-3600, 2010.

[2] T. L. Marzetta, E. G. Larsson, H. Yang and H. Q. Ngo, "Fundamentals of Massive MIMO – 1st edition", Cambridge University Press, December 2016

[3] M. Min, Y.-S. Jeon and G.-H. Im, "On achievable multiplexing gain of BD in MIMO broadcast channels with limited feedback", IEEE Trans. on Wireless Communications, vol. 15, no. 2, pp. 871-885, 2016.

[4] Q. H. Spencer, A. L. Swindlehurst and M. Haardt, "Zero-forcing methods for downlink spatial multiplexing in multiuser MIMO channels", IEEE Trans. on Signal Processing, vol. 52, no. 2, pp. 461-471, 2004.

[5] Z. Shen, R. Chen, J. G. Andrews, R. W. Heath, Jr., B. L. Evans, "Low complexity user selection for multiuser MIMO systems with block diagonalization", IEEE Trans. on Signal Processing, vol. 54, no. 9, pp. 3658-3663, 2006.

[6] V.-F. Crasmariu, M.-O. Arvinte, A.-A. Enescu and S. Ciochina, "Performance Analysis of the Singular Value Decomposition with Block-Diagonalization Precoding in Multi-User Massive MIMO Systems", 12th IEEE International Symposium on Electronics and Telecommunications (ISETC), 2016.

[7] K. Zu and R. C. de Lamare, "Low-Complexity Lattice Reduction-Aided Regularized Block Diagonalization for MU-MIMO Systems", IEEE Communication Letters, vol. 16, no. 6, pp. 925-928, 2016.

[8] W. Ni and X. Dong, "Hybrid block diagonalization for massive multiuser MIMO systems", IEEE Trans. on Communications, vol. 3, no. 1, pp. 201-211, 2015.

[9] V.-F. Crasmariu, M.-O. Arvinte, A.-A. Enescu and S. Ciochina, "Block-Diagonalization Applied in Multi-User Massive MIMO", 11th International Conference Communications (COMM), 2016.

[10] G. H. Golub, C. F. Van Loan, "Matrix Computations - Third Edition", Johns Hopkins University Press, 1996.