

A High Throughput Beamforming Architecture for MIMO Systems

Melissa Duarte and Ashutosh Sabharwal
Department of Electrical and Computer Engineering
Rice University
Houston, Texas 77005, USA
Email: {mduarte, ashu}@rice.edu

Chris Dick and Raghu Rao
Xilinx Inc.
San Jose, California 95124, USA
Email: {chris.dick, raghu.rao}@xilinx.com

Abstract—Quantized feedback in multiple antenna systems has potential to increase throughput or reduce probability of outage. In this paper, we present a method to generate quantization codebooks tailored for efficient implementation of beamforming-based MIMO transmissions. The proposed codebooks can reduce computational requirements significantly, making feasible an efficient architecture for a real-time transmit beamforming and receive combining MIMO orthogonal frequency division multiplexing (OFDM) system. Simulation results are used to validate the proposed codebook construction method and architecture.

I. INTRODUCTION

MIMO systems can be used to increase the data rate (multiplexing), improve reliability (diversity), or both in certain combinations. Full diversity can be achieved using either space-time codes or a system with beamforming at the transmitter and combining at the receiver. For narrow band channels it is known that compared to space time codes, beamforming and combining systems provide the same diversity order [1], more array gain [1], and lower probability of outage [2]. When transmission is done over a wide band channel, OFDM can be used to convert the wide band channel into parallel narrow-band channels, allowing the use of narrow band techniques per subcarrier.

In this paper we will investigate the implementation of a beamforming MIMO-OFDM system with low rate feedback channel. We observe that for quantization the number of multiplications and additions increases exponentially with the number of feedback bits per subcarrier, and linearly with the number of transmitter antennas and the number of receiver antennas. For example, a beamforming system using either Grassmannian [1], Unitary Space-Time Constellation (USTC) [3], or 802.16e [4] codebooks, with 6 transmitter antennas, 4 receiver antennas, and 5 feedback bits per subcarrier, would require a total 3,328 multiplications and 3,040 additions per subcarrier, leading to large area requirements to meet real-time constraints in a high throughput system. These observations motivate the design of codebooks tailored for hardware implementation. The main challenge is the design of codebooks that reduce the computational requirements and achieve good performance.

Our contributions are the following:

- 1) We propose a codebook mapping scheme that generates

codebooks with a structure that reduces the number of multiplications and additions required by allowing the implementation of several multipliers and adders to be done using multiplexers and negators (two's complements) only. Simulation results show that the codebooks generated using the proposed mapping scheme have very good performance, where the performance metric is the average bit error rate (BER).

- 2) We design a high throughput Quantizer architecture tailored for codebooks generated using the proposed mapping scheme. Our measure of throughput is the number of channels quantized per unit area.
- 3) We propose a Mixed Codebook (MxC) scheme that allows the use of the proposed high throughput Quantizer architecture in a standards-based system.

The rest of the paper is organized as follows. Section II is a review of beamforming in MIMO-OFDM systems. Section III shows the computational requirements for the Quantizer block and at the same time sets a framework for comparison with the computational requirements of a Quantizer that uses codebooks generated based on the proposed mapping scheme. Section IV explains the mapping scheme and shows the comparison of computational requirements. The MxC scheme is explained in section V. Section VI is a summary of results.

II. MIMO-OFDM AND BEAMFORMING

We consider a transmit beamforming and receive combining MIMO-OFDM system with M_t transmitter antennas, M_r receiver antennas, and K data subcarriers. The complex symbol transmitted on subcarrier k is represented by s_k . The $M_t \times 1$ vector used at the beamforming block for beamforming subcarrier k is represented by \mathbf{w}_{b_k} . The K beamforming vectors $\mathbf{w}_{b_1}, \dots, \mathbf{w}_{b_K}$ are chosen from a codebook \mathbf{W} of cardinality N , b_k specifies the index of the beamforming vector chosen to beamform subcarrier k , $1 \leq b_k \leq N$. Each subcarrier is transmitted with the same average power E_s , this constraint is met by setting $\mathbb{E}[|s_k|^2] = E_s$ and $\|\mathbf{w}_{b_k}\|_2 = 1$.

Assume that the time span of the channel impulse response is shorter than the time span of the cyclic prefix. In this case the data symbols go through parallel narrow-band channels and the baseband relationship between the symbol transmitted

on subcarrier k , and the corresponding received signal x_k , is given by

$$x_k = \mathbf{z}_k^H \mathbf{H}_k \mathbf{w}_{b_k} s_k + \mathbf{z}_k^H \mathbf{n}_k. \quad (1)$$

The $M_r \times 1$ noise vector \mathbf{n}_k has *i.i.d* entries, each entry is assumed to be circularly symmetric complex Gaussian with zero mean and variance N_0 per complex dimension. The $1 \times M_r$ vector used at the maximum ratio combining (MRC) block for combining subcarrier k is represented by \mathbf{z}_k . The channel matrix for subcarrier k is represented by the $M_r \times M_t$ matrix \mathbf{H}_k . The entries of this matrix are *i.i.d* and each entry is a circularly symmetric complex Gaussian random variable with zero mean and unit variance per complex dimension. The channel matrix \mathbf{H}_k is assumed to be known at the receiver.

The SNR for subcarrier k is

$$\text{SNR}_k = \frac{E_s |\mathbf{z}_k^H \mathbf{H}_k \mathbf{w}_{b_k}|^2}{\|\mathbf{z}_k\|_2^2 N_0}. \quad (2)$$

Because the model in (1) corresponds to a narrow band system, we know from [1] that the beamforming and combining vectors that maximize SNR_k are :

$$\mathbf{w}_{b_k} = \arg \max_{\mathbf{w} \in \mathbf{W}} \|\mathbf{H}_k \mathbf{w}\|_2^2 \quad (3)$$

and

$$\mathbf{z}_k = \frac{\mathbf{H}_k \mathbf{w}_{b_k}}{\|\mathbf{H}_k \mathbf{w}_{b_k}\|_2}. \quad (4)$$

With this choice of \mathbf{w}_{b_k} and \mathbf{z}_k the SNR for subcarrier k is

$$\text{SNR}_k = \frac{E_s}{N_0} \|\mathbf{H}_k \mathbf{w}_{b_k}\|_2^2. \quad (5)$$

The channel is assumed to remain constant during the transmission of a frame and to have independent realizations across different frames. The channel state information at the transmitter is provided by the receiver through a low rate, noiseless, zero delay feedback channel. The codebook \mathbf{W} is known to both transmitter and receiver, the receiver sends back to the transmitter the b_1, \dots, b_K indexes that specify the beamforming vectors, or codewords, to use for transmission. Since the codebook has cardinality $N = 2^B$, each index b_k is represented using B bits. This means that each subcarrier channel \mathbf{H}_k is quantized into B bits, this processing takes place at the Quantizer. In order to reuse hardware resources the Quantizer process one subcarrier at a time. The total number of feedback bits is equal to KB , which corresponds to B bits of feedback information per subcarrier. We assume that the feedback channel is able to convey this amount of information. If this is not the case, clustering strategies like the ones proposed in [5] are an option. Our proposed architecture is well suited for implementation of these strategies.

III. COMPUTATIONAL REQUIREMENTS OF QUANTIZER

At the receiver the K channels are quantized into KB bits. The Quantizer runs the exhaustive search in (3) for each subcarrier channel. The input output relationship for the Quantizer block is equal to

$$b_k = \arg \max_{1 \leq i \leq N} \|\mathbf{H}_k \mathbf{w}_i\|_2^2. \quad (6)$$

The computation of $\|\mathbf{H}_k \mathbf{w}_i\|_2^2$ can be performed in two steps. First compute \mathbf{Q}_i and then compute Qn_i , where:

$$\mathbf{Q}_i = \mathbf{H}_k \mathbf{w}_i \quad (7)$$

$$\text{Qn}_i = \|\mathbf{Q}_i\|_2^2 \quad (8)$$

The entries of \mathbf{H}_k and \mathbf{w}_i are complex numbers, so computing \mathbf{Q}_i requires a total of $4M_t M_r$ multiplications and $4M_t M_r - 2M_r$ additions, and computing Qn_i given \mathbf{Q}_i requires a total of $2M_r$ multiplications and $2M_r - 1$ additions. This means that to compute \mathbf{Q}_i and Qn_i for all N codewords requires a total of $4NM_t M_r + 2NM_r$ multiplications and $4NM_t M_r - N$ additions. After computing Qn_i for all N codewords the Quantizer searches for the codeword with the largest Qn_i , implementing a tree search requires $N - 1$ relational blocks, these are blocks that can compare two inputs and output the greatest of the two.

For the Quantizer, we propose the pipelined architecture shown in Fig. 1. In section IV we will introduce a modified version of this architecture that achieves a higher throughput by taking advantage of the structure of the codebooks generated using the mapping scheme.

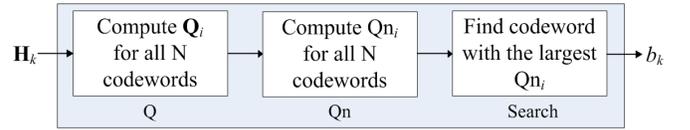


Fig. 1. Quantizer architecture.

IV. QUANTIZER ARCHITECTURE USING CODEBOOK MAPPING

The generation of codebooks using the mapping scheme is explained in section IV-A. Simulation results show that these codebooks have very good performance. Section IV-B shows the computational requirements for a Quantizer using these codebooks and the corresponding high throughput Quantizer architecture.

A. Codebook Mapping

A codebook \mathbf{W} can be represented as a matrix with N rows and M_t columns. We now introduce a simple mapping scheme that transforms any codebook matrix \mathbf{W} , into a codebook matrix \mathbf{W}_M , that can be decomposed as

$$\mathbf{W}_M = \mathbf{G} \mathbf{C}_M, \quad (9)$$

where \mathbf{G} is a $N \times N$ diagonal matrix whose entries are real numbers, and \mathbf{C}_M is a $N \times M_t$ matrix whose entries belong to $\{0, 1, -1, j, -j\}$. Any codebook that can be decomposed in this way will be called a Mapped Codebook (MC).

The mapping is determined by the mapping regions in Fig. 2 and the corresponding region values in Table I. Region zero, R_0 , has an area that depends on D , we have set

$$D = \frac{1}{2\sqrt{M_t}}, \quad (10)$$

this choice of D will be explained later.

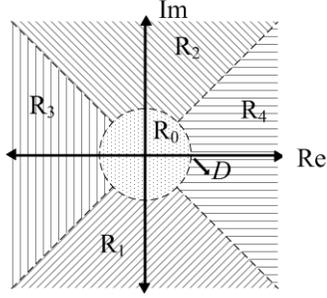


Fig. 2. Mapping Regions. R_1 , R_2 , R_3 , and R_4 , have the same area. The boundaries between these regions are defined by a 45° rotation of the complex plane axes. The area of R_0 is determined by D .

TABLE I

VALUE ASSIGNED TO EACH MAPPING REGION IN FIG.2

Region	R_0	R_1	R_2	R_3	R_4
Value	0	$-j$	j	-1	1

The mapping scheme starts with any codebook \mathbf{W} . We call this the Original Codebook (OC) which can be, for example, any of the codebooks proposed in [1] [3] or [4]. The mapping consists in projecting, via the mapping regions, each of the entries of the OC matrix \mathbf{W} . In other words, the mapping scheme consists in taking each of the NM_t entries of the OC matrix, one by one, and do the mapping as specified by the mapping regions and the region values in table I. For example, if entry $w_{i,t}$, the entry in the i -th row and t -th column of the OC matrix \mathbf{W} , falls in R_2 , the mapping assigns to this entry a value equal to j . This value is stored in the i -th row and t -th column of matrix \mathbf{C}_M , which has the same dimensions as the OC. Matrix \mathbf{C}_M can be a codebook matrix only if its rows (possible beamforming vectors) are normalized, this in order to meet the power constraint at the transmitter. If this is not the case, then \mathbf{C}_M can be transformed into a codebook by multiplying it by a diagonal matrix $\mathbf{G} = \text{diag}(g_1, g_2, \dots, g_N)$, where g_i denotes the reciprocal of the norm of the i -th row of \mathbf{C}_M . The result of this multiplication is a MC. Tables II and III show an example of \mathbf{W} and the corresponding \mathbf{W}_M matrix.

D in Fig.2 has been defined as shown in (10), we now explain this choice of D . The lower bound of D is zero, this makes the area of R_0 equal to zero. If the entries of the OC have all the same norm, then this norm must be equal to $1/\sqrt{M_t}$ in order to have rows with norm equal to 1. To avoid mapping all the entries of this codebook to R_0 , which maps to a value equal to zero, we must choose $D < 1/\sqrt{M_t}$, this is the upper bound for D . We choose D as shown in (10) because this value corresponds to the midpoint between the lower and upper bound for D . Simulation results shows that this is a good choice for D , we have not study other choices of D or their effect on performance.

Simulations using different MCs were run for a one sub-carrier model. These simulations represent the per subcarrier

TABLE II

OC MATRIX USED IN THE 802.16E STANDARD [4] FOR $M_t = 4$ AND $N = 8$

Antenna 1	Antenna 2	Antenna 3	Antenna 4
1	0	0	0
0.378	-0.2698-0.5668j	0.5957+0.1578j	0.1587-0.2411j
0.378	-0.7103+0.1326j	-0.235-0.1467j	0.1371+0.4893j
0.378	0.283-0.094j	0.0702-0.8261j	-0.2801+0.0491j
0.378	-0.0841+0.6478j	0.0184+0.049j	-0.3272-0.5662j
0.378	0.5247+0.3532j	0.4115+0.1825j	0.2639+0.4299j
0.378	0.2058-0.1369j	-0.5211+0.0833j	0.6136-0.3755j
0.378	0.0618-0.3332j	-0.3456+0.5029j	-0.5704+0.2113j

TABLE III

MC MATRIX OBTAINED BY MAPPING AND THEN NORMALIZING THE OC MATRIX IN TABLE II

Antenna 1	Antenna 2	Antenna 3	Antenna 4
1	0	0	0
0.5	-0.5j	0.5	-0.5j
0.5	-0.5	-0.5	0.5j
0.5	0.5	-0.5j	-0.5
0.5773	0.5773j	0	-0.5773
0.5	0.5	0.5	0.5
0.5773	0	-0.5773	0.5773
0.5	-0.5j	0.5j	-0.5

behavior in an OFDM system. The MCs were generated using the proposed mapping scheme. The OCs that were mapped were generated based on three different construction methods: Grassmannian packings [1] [6], USTC systematic generation [3], and the codebooks provided in the 802.16e standard [4]. Fig. 5 shows results for a system with four transmitter antennas and four receiver antennas transmitting 64QAM symbols. This figure shows that codebooks generated using the proposed mapping scheme have very good performance, we observe that for the scenario considered, the performance loss between using a MC and using an OC is never greater than 0.25dB.

B. Computational Requirements and Quantizer Architecture using a Mapped Codebook

Any row or beamforming vector \mathbf{w}_i of a MC is equal to $\mathbf{w}_i = g_i \mathbf{c}_i$, where \mathbf{c}_i is the i -th row of \mathbf{C}_M and g_i is the reciprocal of the norm of \mathbf{c}_i . In this case the input output relationship for the Quantizer block shown in (6) is equivalent to

$$b_k = \arg \max_{1 \leq i \leq N} \|\mathbf{H}_k \mathbf{c}_i\|_2^2 g_i^2. \quad (11)$$

Similar to the analysis done in section III, the computation of $\|\mathbf{H}_k \mathbf{w}_i\|_2^2 g_i^2$ can be performed in three steps. First compute

\mathbf{Q}_i^M , second compute \mathbf{Qn}_i^M , and third compute \mathbf{Qr}_i^M where:

$$\mathbf{Q}_i^M = \mathbf{H}_k \mathbf{c}_i \quad (12)$$

$$\mathbf{Qn}_i^M = \|\mathbf{Q}_i^M\|_2^2 \quad (13)$$

$$\mathbf{Qr}_i^M = \mathbf{Qn}_i^M g_i^2 \quad (14)$$

Multiplying an entry of the channel matrix times an entry of \mathbf{c}_i , can be done without actually using any multiplier, only two multiplexers and two negators are needed, as shown in Fig. 3. This implementation is possible because the entries of \mathbf{c}_i belong to $\{0, 1, -1, j, -j\}$ and can be identified by the region to which they belong (table I). Using the architecture in Fig. 3, the computation of \mathbf{Q}_i^M for all N codewords requires $2NM_tM_r$ multiplexers, $2NM_tM_r$ negators, and $2NM_tM_r - 2NM_r$ additions. Once \mathbf{Q}_i^M is known, computing \mathbf{Qn}_i^M for all codewords requires $2NM_r$ multiplications and $2NM_r - N$ additions. The computation of \mathbf{Qr}_i^M for all codewords given \mathbf{Qn}_i^M and g_i requires only N multiplications. The tree search to find the largest \mathbf{Qr}_i^M requires $N - 1$ relational blocks.

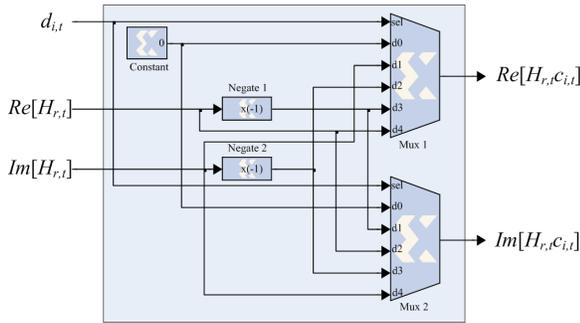


Fig. 3. Multiplication of an entry of \mathbf{c}_i times an entry of \mathbf{H}_k . $c_{i,t}$ denotes the t -th entry of vector \mathbf{c}_i , and $d_{i,t}$ denotes the corresponding region according to table I. $H_{r,t}$ denotes the entry in row r and column t of the channel matrix \mathbf{H}_k .

Table IV shows the total number of multiplications, multiplexers and additions that are needed for quantization when using an OC and when using a MC. We observe that using a MC reduces the number of multiplications and additions by trading off multipliers and adders by multiplexers. Using a MC, a system with 6 transmitter antennas, 4 receiver antennas, and 5 feedback bits per subcarrier, would require a total of 288 multiplications, 1,536 multiplexers, and 1,504 additions. This is a significant reduction in the number of multiplications and additions compared to the 3,328 multiplications and 3,040 additions that are needed if using an OC.

In order to quantify the efficiency in resource utilization when using a MC, we compute the increase in throughput, where our measure of throughput is the number of channels Quantized per unit area, and we define the increase in throughput as β . To compute β we assume that for the implementation of one multiplexer the area needed is the same as for the implementation of one multiplier, a very conservative assumption for an ASICs and a reasonable assumption for an FPGA that has embedded multipliers, like for example a Xilinx Virtex II Pro FPGA. Since multipliers are more expensive in

hardware than adders, we only include in the computation of β the trade off between multipliers and multiplexers; β is then equal to the ratio between the number of multipliers needed when using an OC, and the number of multipliers and multiplexers needed when using a MC,

$$\beta = \frac{\text{Multipliers (OC)}}{\text{Multipliers (MC) + Multiplexers (MC)}}. \quad (15)$$

For the example we have been considering, a system with 6 transmitter antennas, 4 receiver antennas and 5 feedback bits per subcarrier, we obtain $\beta = 1.82$. This means that for this example using the MC reduces the area utilized by a factor equal to 1.82.

TABLE IV
COMPARISON OF MULTIPLICATIONS, MULTIPLEXERS, AND ADDITIONS REQUIRED FOR QUANTIZATION USING AN OC AND QUANTIZATION USING A MC

	Quantizer using an OC	Quantizer using a MC
Multiplications	$4NM_tM_r + 2NM_r$	$2NM_r + N$
Multiplexers	0	$2NM_tM_r$
Additions	$4NM_tM_r - N$	$2NM_tM_r - N$

Another advantage of using a MC is the flexibility in resource utilization. Using a MC allows the implementation of the multiplication of an entry of \mathbf{H}_k times an entry of \mathbf{c}_i as shown in Fig. 3, but the designer can choose to implement some or all of these multiplications using multipliers. When implementing on an FPGA with embedded multipliers, this flexibility is an important tool for the designer, who would utilize as many embedded multipliers as possible and then implement the rest of the multiplications on the FPGA fabric as shown in Fig. 3.

The implementation of the Quantizer is based on the pipelined architecture in Fig. 1. For a Quantizer using a MC we propose the architecture shown in figure 4. Since using a MC increases the throughput in channels quantized per unit area we call this a high throughput Quantizer architecture.

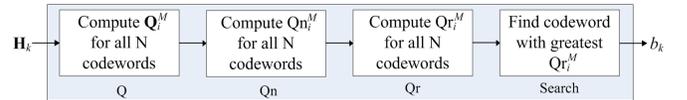


Fig. 4. Quantizer architecture tailored for a Quantizer using a MC.

V. USING THE QUANTIZER ARCHITECTURE WITH CODEBOOK MAPPING IN A STANDARDS-BASED SYSTEM

In a standards-based system like 802.16e the codebook is predefined. In this scenario it is not likely that a single equipment vendor will control both ends of the link, therefore, it may not be possible to impose the requirement of using a MC at the beamforming, MRC, and Quantizer blocks, instead of using the predefined standards codebook. We are motivated to consider the scenario in which the standards codebook is

used at the beamforming and MRC blocks, and the MC version of this codebook is used at the Quantizer. We call this a MxC scheme. The MC version of the standards codebook is obtained by following the mapping procedure explained in section IV-A. Simulation results are shown in Fig. 5. We observe that the performance loss by using the MxC scheme is about 1dB compared to the OC scheme, which uses the standards codebook at the beamforming, MRC, and Quantizer blocks. This can be interpreted as sacrificing 1dB in performance in order to be able to implement the proposed high throughput Quantizer architecture. This trade off between performance and computational complexity is a useful tool when designing a communications link.

An intuition behind the result in Fig. 5 is that the Quantizer chooses the index of the best possible codeword, the one that maximizes the SNR, based on the MC codebook, but the beamforming and MRC blocks are using the standards codebook, so for some channel realizations the best codeword chosen from the MC won't match the best possible codeword from the standards codebook. In these situations the codeword chosen by the Quantizer usually corresponds to the second or third best codeword to choose from the standards codebook.

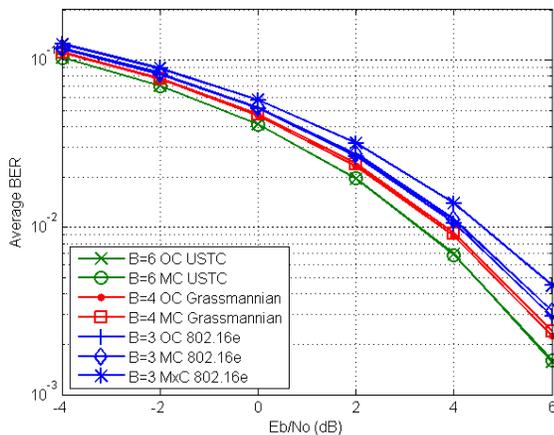


Fig. 5. The figure shows BER results for a four transmitter and four receiver antennas system with three feedback bits transmitting 64QAM symbols.

In the 802.16e standard a possible configuration is a system with 4 transmitter antennas, 4 receiver antennas and 3 feedback bits. For this system, using the mapped codebook yields an increase in throughput $\beta = 1.75$ and the number of multipliers reduces from 576 to 72.

Table V shows that even if we allow a resource reutilization we obtain an increase in throughput and a great reduction in the number of multipliers. Table V compares the number of multipliers, multiplexers, and adders required for a Quantizer implemented using an OC with the resources required for a Quantizer using a MC. We set the resource reutilization factor equal to $C_p = T_{OFDM}/(KT_{FPGA})$. This value of C_p is obtained by setting the constraint that all the K channels must be quantized in a time equal to the duration of an OFDM symbol, which is represented by T_{OFDM} . In other words, the

TABLE V

COMPARISON OF RESOURCE UTILIZATION BETWEEN A QUANTIZER USING AN OC AND A QUANTIZER USING A MC. THE RESOURCE ESTIMATE IS COMPUTED FOR AN 802.16E BASED SYSTEM WITH $M_t = 4$, $M_r = 4$, $N = 2^3$, AND $T_{OFDM} = 100.8\mu s$. WE ASSUME $T_{FPGA} = 10ns$

	$K = 512$		$K = 1024$		$K = 2048$	
	5MHz BW	10MHz BW	OC	MC	OC	MC
Multipliers	31	5	60	8	119	16
Multiplexers	0	14	0	27	0	53
Adders	26	13	52	26	104	52
β	1.63		1.71		1.72	

input of the Quantizer block is refreshed every T_{OFDM}/K seconds. This constraint is reasonable for an 802.16e based system operating in Frequency Division Duplexing (FDD) mode because in this case feedback is almost simultaneous with the downlink [4]. With an FPGA clock period equal to T_{FPGA} , each of the blocks in the Quantizer has C_p clock periods to perform the required computations.

From table V we observe that for 2048 subcarriers the increase in throughput β is equal to 1.72. Even by allowing some resource reutilization, using a MC reduces the area utilized. We also observe that for 2048 subcarriers, using the MC reduces the number of multipliers required from 119 to only 16. This is a reduction in the multiplier complexity by a factor of approximately 86%.

VI. SUMMARY OF RESULTS

We proposed a MC generation scheme that allows a high throughput Quantizer architecture. Simulation results show that the performance loss is less than 0.25dB when the MC is used at the beamforming, Quantizer and MRC blocks, and less than 1dB when the MC is used at the Quantizer only. The latter corresponds to the standards-based scheme, MxC scheme explained in section V.

REFERENCES

- [1] D. J. Love, R. W. Heath, and T. Strohmer, "Grassmannian beamforming for multiple-input multiple-output wireless systems," *IEEE Trans. Inform. Theory*, vol. 49, pp. 2735–2747, Oct. 2003.
- [2] K. K. Mukkavilli, A. Sabharwal, E. Erkip, and B. Aazhang, "On beamforming with finite rate feedback in multiple-antenna systems," *IEEE Trans. Inform. Theory*, vol. 49, pp. 2562–2579, Oct. 2003.
- [3] B. M. Hochwald, T. Marzetta, T. J. Richardson, W. Sweldens, and R. Urbanke, "Systematic design of unitary space-time constellations," *IEEE Trans. Inform. Theory*, vol. 46, pp. 1962–1963, 2000.
- [4] *Air Interface for Fixed and Mobile Broadband Wireless Access System - Amendment for Physical and Medium Access Control Layers for Combined Fixed and Mobile Operation in Licensed Bands - Corrigendum 1*, IEEE Std. 802.16e, 2006.
- [5] B. Mondal and R. W. Heath, "Algorithms for quantized precoding in MIMO OFDM beamforming systems," in *Proc. of Third SPIE Int. Symp. on Fluctuations and Noise*, Austin Tx, US, May 2005.
- [6] D. J. Love. Grassmannian subspace packing. Accessed 11/29/2006. [Online]. Available: <http://dynamo.ecn.purdue.edu/~djlove/grass.html>