ADMM-based Infinity Norm Detection for Large MU-MIMO: Algorithm and VLSI Architecture

Shahriar Shahabuddin and Markku Juntti
Centre for Wireless Communications, Finland
e-mail: firstname.lastname@oulu.fi

Christoph Studer
Cornell University, USA
e-mail: studer@cornell.edu

***Abstract*—We propose a novel data detection algorithm and a corresponding VLSI design for large multi-user (MU) multiple-input multiple-output (MIMO) wireless receiver. Our algorithm, referred to as ADMIN, performs alternating direction method of multipliers (ADMM)-based infinity norm constrained equalization. ADMIN is an iterative algorithm that outperforms linear detectors if the number of users is small compared to that of the antennas in base station (BS). ADMIN computes the linear minimum mean-square error (MMSE) solution in the first iteration. It is sufficient when the ratio between the numbers of BS antennas and users is rather large. We develop a time-shared and iterative VLSI architecture for LDL-decomposition based soft-output ADMIN. Our architecture achieves 685.71 Mb/s for linear MMSE and 212.38 Mb/s for ADMIN for a 16-user system that employs 64-QAM in a 28 nm CMOS technology.**

I. Introduction

Large (or massive) multiuser (MU) multiple-input multiple-output (MIMO) is a key technology for fifth-generation (5G) wireless communication systems that handles orders of magnitude more data traffic. The idea of massive MU-MIMO is to equip the base station (BS) with a large number of antenna elements that can serve a large number of user terminals in the same frequency band [1], [2]. The improvements in spectral efficiency of massive MU-MIMO systems come at the cost of higher computational complexity compared to that of conventional small-scale MIMO wireless systems. To enable high-speed communication in the uplink (users communicate to the BS), a variety of algorithms and VLSI architectures have been proposed recently [3]–[6]; all of these methods use an approximation of the linear minimum mean-square error (MMSE) equalizer. These algorithms provide high throughput, but entail a performance loss compared to exact inversion-based MMSE equalizers, especially in systems in which the number of users is comparable to the number of BS antennas. Castañeda *et al.* recently proposed an approximate semidefinite relaxation-based data detector, referred to as TASER, that achieves near maximum-likelihood performance in such "symmetric" systems. TASER, however, is limited to BPSK and QPSK modulation [7].

A. Contributions

In this paper, we propose a novel data detection algorithm and VLSI design based on the alternating direction method of multipliers (ADMM). Our algorithm is referred to as ADMM-based infinity-norm (ADMIN for short) and performs infinity-norm or box-constrained equalization, which outperforms linear detectors by a large margin if the ratio between the numbers of BS antennas and users is small (two or less). ADMIN is iterative by nature and performs linear MMSE equalization in the first iteration. Therefore, for systems in which the number of BS antennas is an order of magnitude more than the number of users, it is sufficient to perform one ADMIN iteration. We present a VLSI architecture for LDL-based soft-output ADMIN for 16 users transmitting data to a 16 antenna BS. The architecture is implemented in a 28 nm CMOS technology and compared to the only available 16-user ASIC implementations in [7] in terms of throughput, area, and hardware efficiency.

B. Notation

Boldface lowercase and boldface uppercase letters stand for column vectors and matrices, respectively. For a matrix, $\mathbf{A}$, we denote its Hermitian transpose by $\mathbf{A}^H$. We use $A_{k,l}$ for the entry in the kth row and lth column of the matrix $\mathbf{A}$. The real and imaginary part of a complex-valued matrix $\mathbf{A}$ are denoted by $\Re(\mathbf{A})$ and $\Im(\mathbf{A})$, respectively. The identity matrix is $\mathbf{I}$ and ℓ_2-norm of the vector $\mathbf{a}$ is $\|\mathbf{a}\|_2 = \sqrt{\sum_K |a_k|^2}$.

II. System model and Data Detection

We consider a massive MU-MIMO wireless uplink system that employs orthogonal frequency division multiplexing (OFDM). We assume that U single-antenna user terminals send data simultaneously to a BS with $B \geq U$ antennas over W subcarriers. The U users first encode their own bit stream with a channel encoding method (e.g., a convolutional code) and map the coded bit stream to constellation points in the finite alphabet set $\mathcal{O}$ (e.g., 16-QAM with Gray mapping) with an average transmit power E_s per symbol. We assume perfect channel state information (CSI) and synchronization at the receiver, as well as a sufficiently long cyclic prefix such that there is no intersymbol or interblock interference. By omitting the subcarrier index, the input-output relation per subcarrier can be written as $\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{n}$, where $\mathbf{y} \in \mathbb{C}^B$ is the received signal vector, $\mathbf{x} \in \mathbb{C}^U$ is the transmit symbol vector, $\mathbf{H} \in \mathbb{C}^{B\times U}$ is the channel matrix, and $\mathbf{n} \in \mathbb{C}^B$ is the circularly symmetric complex white Gaussian noise vector with zero mean and variance N_0 per complex entry.

Maximum likelihood (ML) detection is optimal in terms of minimizing the vector error rate assuming that all data vectors are equally likely. For the considered system model, the ML detector is given by

$$\hat{\mathbf{x}}_{\text{ML}} = \arg\min_{\mathbf{x}\in\mathcal{O}^U} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2^2. \tag{1}$$

Unfortunately, this problem is of combinatorial nature and requires prohibitive complexity for a large number of users [8]. To reduce complexity, one can use minimum mean-square error (MMSE) equalization, which can be viewed as a relaxed version of the ML problem in which the data set $\mathcal{D}$ underlying ML detection is relaxed to the set $\mathbb{C}^U$ with an additional regularization term as follows:

$$\hat{\mathbf{x}}_{\text{MMSE}} = \arg\min_{\mathbf{x}\in\mathbb{C}^U} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2^2 + \ N_0E_s^{-1}\|\mathbf{x}\|_2^2. \tag{2}$$

The regularization term $N_0E_s^{-1}\|\mathbf{x}\|_2^2$ prevents $\mathbf{x}$ from growing too large, because the transmit signals are taken from the discrete constellation $\mathcal{O}$ that is centered around zero. Since the objective function of (2) is quadratic in $\mathbf{x}$, MMSE equalization has a closed form solution and can be solved relatively efficiently in hardware [9].

III. ADMIN: ADMM-based Infinity Norm Detection

A. Infinity-Norm Constrained Equalization

Infinity norm or box-constrained equalization [10], [11] relaxes the finite-alphabet constraint $\mathbf{x} \in \mathcal{O}^U$ to the convex polytope $\mathcal{C}_\mathcal{O}$ around the constellation set $\mathcal{O}$ and solves the following convex optimization problem:

$$\hat{\mathbf{x}}_{\text{BOX}} = \arg\min_{\mathbf{x}\in\mathcal{C}_{\mathcal{O}}^{U}} \|\mathbf{y}-\mathbf{H}\mathbf{x}\|_2^2. \tag{3}$$

The convex polytope for QPSK and higher order QAM alphabets can be expressed as $\mathcal{C}_{\mathcal{O}} = \{x_R + jx_I : x_R, x_I \in [-\alpha, +\alpha]\}$ where $\alpha = \max_{u\in\mathcal{O}} \Re\{u\}$ is the tightest radius of the box around the square constellation. For example, the convex polytope for QPSK is given by a square box with radius $\alpha = 1$ around the square constellation of QPSK. Solving (3) exactly with conventional interior-point methods in hardware would result in prohibitive complexity. We therefore propose an ADMM-based solution for this box-constrained equalization problem.

B. ADMM-Based Infinity-Norm Detection

ADMM is a well known numerical method to solve a wide range of constrained convex and nonconvex optimization problems [12]. More specifically, ADMM solves the original convex optimization problem by breaking them into smaller sub-problems that can be solved efficiently. We rewrite (3) into the following equivalent form

$$\underset{\mathbf{x},\mathbf{z}\in\mathbb{C}^{U}}{\text{minimize}}\ \tfrac{1}{2}\|\mathbf{y}-\mathbf{H}\mathbf{x}\|_2^2 + g(\mathbf{z}) \quad \text{subject to } \mathbf{z}=\mathbf{x} \tag{4}$$

where $g(\mathbf{z})$ is the indicator function on the convex set $\mathcal{C}_{\mathcal{O}}$ such that

$$g(\mathbf{z}) = \begin{cases} 0, & \text{if } \mathbf{z}\in\mathcal{C}_{\mathcal{O}}^{U} \\ \infty, & \text{otherwise.} \end{cases}$$

The augmented Lagrangian for the problem in (4) is

$$L_{\beta}(\mathbf{x},\mathbf{z},\boldsymbol{\lambda}) = \tfrac{1}{2}\|\mathbf{y}-\mathbf{H}\mathbf{x}\|_2^2 + g(\mathbf{z}) + \tfrac{\beta}{2}|\mathbf{z}-\mathbf{x}-\boldsymbol{\lambda}\|_2^2 \tag{5}$$

where $\boldsymbol{\lambda}$ is the scaled dual variable associated with the constraint $\mathbf{z}=\mathbf{x}$ and $\beta > 0$ is a suitably chosen regularization parameter. We can now use ADMM to solve the augmented Lagrangian over $\mathbf{x}$ and $\mathbf{z}$. In the first step of an ADMM iteration, we minimize $\mathbf{x}$ while $\mathbf{z}$ is fixed. We take the derivative of (5) with respect to $\mathbf{x}$ and set it to zero to compute the first step of ADMIN as

$$\begin{aligned} &\mathbf{H}^{H}(\mathbf{y}-\mathbf{H}\mathbf{x}) - \beta(\mathbf{z}-\mathbf{x}-\boldsymbol{\lambda}) = 0 \\ \Rightarrow &\hat{\mathbf{x}} = (\mathbf{H}^{H}\mathbf{H}+\beta\mathbf{I})^{-1}(\mathbf{H}^{H}\mathbf{y}+\beta(\mathbf{z}-\boldsymbol{\lambda})). \end{aligned} \tag{6}$$

The $\mathbf{x}$-update of ADMIN solves a regularized least-squares problem. Thus, ADMIN can be viewed as a method for solving the box-constrained problem of (3) by iteratively carrying out regularized least-square computations. Note that initializing $\mathbf{z}$ and $\boldsymbol{\lambda}$ with zero and $\beta = N_0 E_s^{-1}$ at (6) yields the MMSE equalizer in the first iteration. The $\mathbf{z}$ update step can be expressed as

$$\hat{\mathbf{z}} = \arg\min_{\mathbf{z}\in\mathcal{C}_{\mathcal{O}}^{U}} \tfrac{1}{2}\|\mathbf{z}-(\hat{\mathbf{x}}+\boldsymbol{\lambda})\|_2^2. \tag{7}$$

Equation (7) is equivalent to an orthogonal projection of $\hat{\mathbf{x}}+\boldsymbol{\lambda}$ onto the convex polytope $\mathcal{C}_{\mathcal{O}}^{U}$. This projection is given by

$$\text{proj}_{\mathcal{C}_{\mathcal{O}}}(w) = \begin{cases} w, & \text{if } w\in\mathcal{C}_{\mathcal{O}} \\ \arg\min_{q\in\mathcal{C}_{\mathcal{O}}} |w-q|, & \text{otherwise.} \end{cases}$$

In words, if w is outside the set $\mathcal{C}_{\mathcal{O}}$, the projection outputs the value closest to w within the set $\mathcal{C}_{\mathcal{O}}$ in terms of the Euclidean distance. For example, if w is outside of a box with radius $\alpha = 1$ that encloses the square constellation of QPSK, then the projection outputs a value q that is closest to w within the box. The update for the Lagrange vector is given by

$$\boldsymbol{\lambda} \leftarrow \boldsymbol{\lambda} - \gamma(\hat{\mathbf{z}}-\hat{\mathbf{x}}), \tag{8}$$

where $0 < \gamma$ is a suitably chosen step size. Note that step sizes satisfying $0 < \gamma < 1$ ensure the convergence of the ADMM but larger values may improve the performance for small numbers of iterations.

Algorithm 1 ADMIN

```
inputs: $\mathbf{y}$, $\mathbf{H}$, $N_0$ and $E_s$
preprocessing
    $\beta = N_0 E_s^{-1}\epsilon$
    $\mathbf{G} = \mathbf{H}^H\mathbf{H} + \beta\mathbf{I}_U$
    $\mathbf{G} = \mathbf{L}\mathbf{d}\mathbf{L}^H$
    $\tilde{\mathbf{L}} = \mathbf{L}^{-1}$, $\tilde{\mathbf{D}} = \mathbf{D}^{-1}$
initialization
    $\mathbf{z} = \mathbf{0}$
    $\boldsymbol{\lambda} = \mathbf{0}$
detection
    $\mathbf{y}_{\text{MF}} = \mathbf{H}^H\mathbf{y}$
    for $i = 1 : K$
        $\hat{\mathbf{x}} \leftarrow \tilde{\mathbf{L}}^H\tilde{\mathbf{D}}\tilde{\mathbf{L}}(\mathbf{y}_{\text{MF}} + \beta(\mathbf{z}-\boldsymbol{\lambda}))$
        $\hat{\mathbf{z}} \leftarrow \text{proj}_{\mathcal{C}_{\mathcal{O}}}(\hat{\mathbf{x}}+\boldsymbol{\lambda}, \alpha)$
        $\boldsymbol{\lambda} \leftarrow \boldsymbol{\lambda} - \gamma(\hat{\mathbf{z}}-\hat{\mathbf{x}})$
        $\mathbf{z} \leftarrow \hat{\mathbf{z}}$
    end
output: $\hat{\mathbf{x}}$
```

C. LDL-Decomposition based Soft-output ADMIN

The inversion of the regularized Gramian matrix, $\mathbf{G} = \mathbf{H}^H\mathbf{H}+\beta\mathbf{I}_U$ is required to compute the $\mathbf{x}$-update of ADMIN. We can compute the exact inverse of regularized the Gramian matrix with LDL-decomposition in ADMIN. Fortunately, the computations required for $\mathbf{G}$, LDL decomposition, and inversion of $\mathbf{L}$ and $\mathbf{D}$ can be done during the preprocessing, and thereby equalization can be simplified. During the equalization stage, ADMIN computes the matched filter and then iteratively updates $\hat{\mathbf{x}}$, $\hat{\mathbf{z}}$, and $\boldsymbol{\lambda}$. The complete ADMIN algorithm is summarized in Algorithm 1.

The post-equalization SINR vector $\boldsymbol{\rho}$, that is required to compute the log-likelihood ratio (LLR) values, can be extracted from $\rho_i = 1/N_0 E_s^{-1} a_i$ where a_i is the i-th entry of the main diagonal of $\mathbf{G}^{-1}$. This can be done efficiently with the help of $\tilde{\mathbf{L}} = \mathbf{L}^{-1}$ and $\tilde{\mathbf{D}} = \mathbf{D}^{-1}$ as $\rho_i = (\tilde{\mathbf{l}}_i)^H \text{diag}(\tilde{\mathbf{D}})(\tilde{\mathbf{l}}_i)$ where $\tilde{\mathbf{l}}_i$ is the i-th column of $\tilde{\mathbf{L}}$. The number of complex-valued multiplication needed for each ADMIN iteration is $2(U^2+U)$ where U is the number of users. Therefore, the complexity of the detector scales with KU^2 where K is the number of ADMIN iterations.

D. Error-rate Performance

We simulate a typical 40 MHz IEEE 802.11n OFDM uplink scenario with a rate-3/4 convolutional code where the channel matrices are generated using WINNER-phase-2 model and the max-log BCJR algorithm is used for soft-input soft-output channel decoding. We show the (coded) packet error-rate (PER) for ADMIN as well as linear MMSE equalization, single-input multiple-output (SIMO) lower bound, TASER and box-constrained coordinate descent (CD) detector [6] in Fig. 1. ADMIN provides a significant performance improvement over MMSE equalization for a symmetric system with five iterations ($K = 5$) and thus, we choose ADMIN with $K = 5$ for error-rate performance. CD outperforms other MMSE approximations like the Neumann series approach [3] or the conjugate gradient (CG) [5] based detectors in terms of PER. For a fair comparison against ADMIN with $K = 5$, we run CD with five ($K = 5$) and ten ($K = 10$) iterations

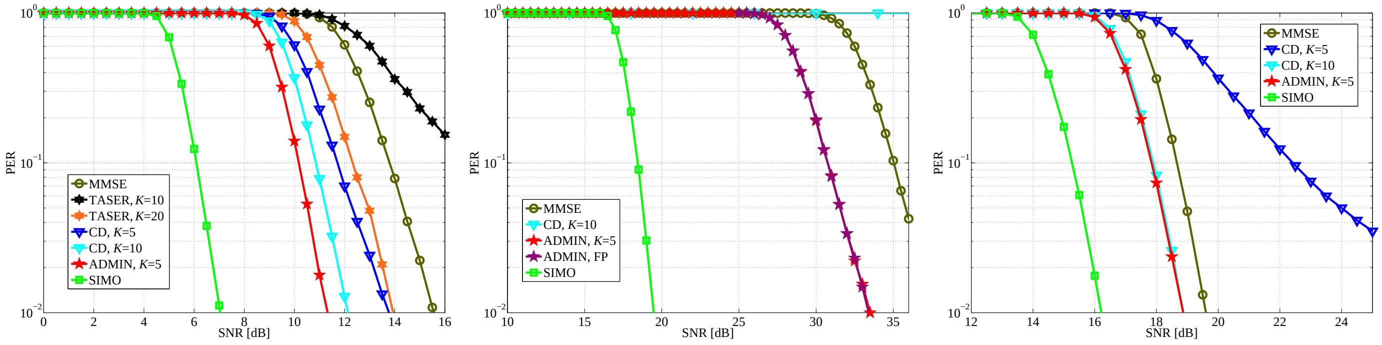

(a) 16 BS antennas, 16 users and QPSK.

(b) 16 BS antennas, 16 users and 64-QAM.

(c) 32 BS antennas, 16 users and 64-QAM.

Fig. 1: Packet error rate (PER) for a large MU-MIMO-OFDM system with rate-3/4 convolutional code and WINNER channel model.

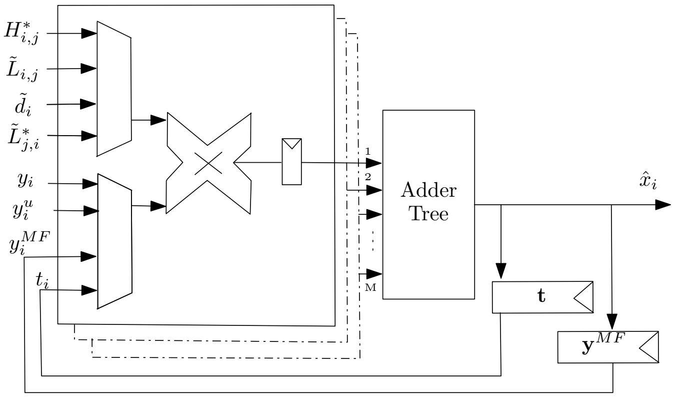

Fig. 2: VM unit: Computes vector-vector multiplication. It has $i = 1, 2, \ldots, M$ multiplier units in parallel. The adder tree sum the output of the multipliers.

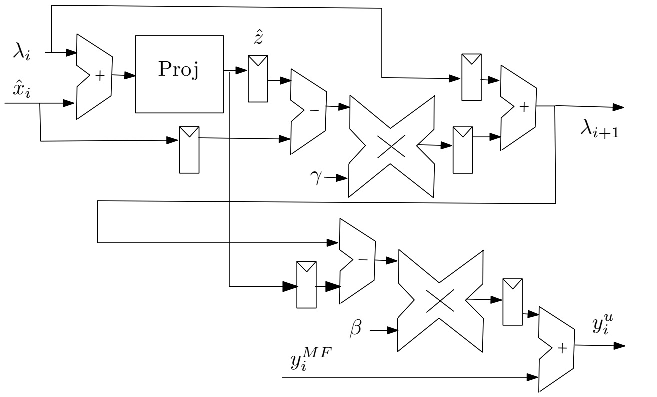

Fig. 3: MFU unit: Computes $\mathbf{z}$ minimization and $\boldsymbol{\lambda}$-update in pipelined fashion.

respectively. Figs. (1a) and (1b) illustrates the PER performance for a 16 users and 16 BS antenna system. For lower order modulation (e.g., QPSK), ADMIN achieves a substantial performance gain over MMSE. ADMIN also outperforms TASER and CD with a smaller number of iterations. For higher order modulation (e.g. 64-QAM), CD fails to detect the symbols for square systems. ADMIN provides approximately 5 dB gain compared to MMSE in this scenario. Fig. (1c) illustrates the PER performance for 16 users and 32 BS antennas. ADMIN with five and CD with ten iterations performs similar in this scenario. However, CD performs worse than MMSE with five iterations in this scenario.

IV. VLSI Architecture

A. Architecture Overview

We propose a VLSI architecture for ADMIN that takes $\mathbf{H}$, $\mathbf{y}$, $\tilde{\mathbf{L}}$, $\tilde{\mathbf{d}} = \text{diag}(\tilde{\mathbf{D}})$ as inputs. We use fixed-point arithmetic to optimize the efficiency of our hardware. We use 18 bits throughout the design where 12 bits are used for fractional part and 5 bits are used for the integer. The fixed-point performance is shown in Fig. 1(b). The architecture supports ADMIN detection (lines $6 - 16$) as summarized in Algorithm 1. The architecture is divided in two parts and they are explained in the following subsections.

B. Vector Multiplication Unit

The vector multiplication unit (VM) computes the $\mathbf{x}$ minimization step of ADMIN (line 12) of Algorithm 1. VM is designed with time-shared processing elements to compute vector-vector multiplication. VM consists of 16 complex multipliers followed by an adder tree. A register is placed between each complex multiplier and adder tree to reduce the critical path. The adder tree can sum 16 complex values in a single cycle. $\mathbf{H}$ is stored in a standard cell-based memory in such a way that each address can read a column of $\mathbf{H}$ in a single cycle. The VM unit first computes the matched filter $\mathbf{y}^{\text{MF}} = \mathbf{H}^H\mathbf{y}$ in 16 cycles. The same values of $\mathbf{y}^{\text{MF}}$ are needed for all the five ADMIN iterations and that is why they stored in a register array.

The lower triangular matrix $\tilde{\mathbf{L}}$ is also stored in another standard cell-based memory. The triangular memory is designed in such a way that it is possible to read an entire column or row of $\tilde{\mathbf{L}}$ in a single cycle. The $\tilde{\mathbf{L}}$ is read row-wise to compute the $\tilde{\mathbf{L}}\mathbf{y}^{MF}$ and the output is stored in a temporary register array, $\mathbf{t}$. The next step is to compute the elementwise multiplication between $\tilde{\mathbf{d}}$ and $\mathbf{t}$. The output from the multiplier array is written back to $\mathbf{t}$ unlike the previous computations. The triangular memory is read column-wise in next 16 cycles to compute $\tilde{\mathbf{L}}^H\mathbf{t}$ that results in $\hat{\mathbf{x}}$ which is the output of VM unit.

C. Matched Filter Update Unit

The matched filter update (MFU) unit computes the $\mathbf{z}$ minimization and $\boldsymbol{\lambda}$-update steps of ADMM. The output values from the VM unit, $\hat{x}_i$, where $i = 1, 2, \ldots, M$, are obtained sequentially. Therefore, we choose a pipelined architecture for MFU. The $\boldsymbol{\lambda}$ array is stored in a register array and initialized as zero. The projection unit compares the real and imaginary parts of the addition of $\hat{x}_i$ and λ_i and

TABLE I: Throughput and latency of ADMIN for K iterations and 64-QAM, and 16 BS antennas and 16 users

	K=1	K = 2	K = 3	K = 4	K = 5
Max. throughput (Mb/s)	685	440	324	256	212
Latency (µs)	0.14	0.21	0.29	0.37	0.45

outputs $\hat{z}$. The subtraction of $\hat{z}$ from $\hat{x}_i$ is multiplied with the scaling parameter γ. A shimming register is used after $\hat{x}_i$ to synchronize with $\hat{z}$. Similarly another shimming register is used to synchronize λ_i to add with $\gamma(\hat{x}_i - \hat{z})$ that results in an updated λ_{i+1}. The updated λ_{i+1} is stored back in the same register array designated for λ_i and they are used for the next iteration. The subtraction of $\hat{z}$ and λ_{i+1} is multiplied with the penalty parameter β. The output is added with the corresponding matched filter value y_i^{MF} for an updated y_i^u. The updated matched filter values $\mathbf{y}^u$ are stored in a register array and sent back to the VM unit to compute for the next ADMIN iteration.

V. Implementation Results

The ADMIN architecture is designed and optimized in VHDL on register-transfer level (RTL). The architecture is synthesized using Synopsys DC with a 28 nm CMOS standard cell library. ADMIN achieves a maximum clock frequency of 500 MHz and takes an area of 0.061 mm^2 which equals to 186.11 k gate equivalents. The critical path goes through the multiplier array to the temporary register, $\mathbf{t}$. The MFU unit takes a small area compared to the VM unit. The majority of the area in the VM unit is consumed by the complex multiplier array. The throughput and latency of different ADMIN iterations for 16-users and 64-QAM modulation scheme is provided in Table I.

The throughput decreases linearly with respect to the number of iterations K because the number of operations remains the same in each iteration of ADMIN. ADMIN needs five iterations to achieve the desired PER shown in Fig. 1. In the first 70 cycles, the architecture computes the first ADMIN iteration that provides the MMSE estimates. Here, 16 cycles are used for storing the inputs to $\mathbf{H}$ and $\tilde{\mathbf{L}}$ memory. The architecture needs 226 cycles to compute five ADMIN iterations that results in a throughput of 212.38 Mb/s.

A plethora of ASIC designs for small-scale MIMO detection is available in the literature. Their effectiveness is still unexplored for large scale MU-MIMO systems. A few FPGA designs for large scale MU-MIMO detection can be found in [4]–[6]. Yin *et al.* proposed a high-throughput ASIC for a Neumann series based large scale MU-MIMO detector in [3]. However, the FPGA and ASIC implementations support 8-user MU-MIMO and not comparable with our 16-user architecture. In Table II, we compare our ADMIN architecture with TASER [7], which is—to the best of our knowledge—the only ASIC that supports 16-user large MU-MIMO detection. Three separate TASER designs were presented in [7] to support different numbers of users for BPSK and QPSK MU-MIMO systems. We compare our results with the 16-user BPSK TASER and QPSK TASER. ADMIN provides higher throughput and superior hardware efficiency (measured throughput per cell area) than both TASER designs supporting 16-users. Furthermore, ADMIN achieves superior PER performance compared to the well-known MU-MIMO detectors. It should be noted that ADMIN and TASER designs do not include preprocessing circuitry, whereas the Neumann-series detector [3] includes the preprocessing circuitry.

VI. Conclusions

We have proposed ADMIN, a novel data-detection algorithm and a corresponding VLSI architecture. The algorithm outperforms linear MMSE equalization in terms of PER by a large margin when the ratio between BS antenna and number of user is rather small (two or less). ADMIN architecture also provides promising results for 16-user MIMO detectors in terms of throughput and energy efficiency. Thus, ADMIN enables a realistic large-scale MU-MIMO detector implementation for next generation communication systems.

TABLE II: Comparison of data detection ASICs for 16-user large scale MIMO systems

Detection algorithm	ADMIN	TASER	TASER
Modulation Scheme	64-QAM	BPSK	QPSK
Preprocessing	Not included	Not included	Not Included
Iteration	5	5	5
Technology [nm]	28	40	40
Supply Voltage [V]	1	1.1	1.1
Clock freq. [MHz]	500	598	560
Throughput^a [Mb/s]	212.38	74.8 (105.36^a)	72.6 (103.09^a)
Core area^a [mm^2]	0.061	0.482 (.236^a)	1.382 (0.676^a)
Cell area^b [kGE]	186.11	471.23	1427.96
Throughput/ area^a[b/(s×GE)]	1141.15	222.9	72.02

^{a} Technology scaling to 28 nm assuming $A \sim 1/s^2$ and $t \sim 1/s$.
^{b} Excluding the gate count of memories.

VII. Acknowledgements

This research is supported by Academy of Finland. We greatfully acknowledge Ilkka Hautala and Charles Jeon for their insightful comments on ADMIN and its architecture.

References

[1] T. L. Marzetta, “Noncooperative cellular wireless with unlimited numbers of base station antennas,” *IEEE Trans. Wireless Commun.*, vol. 9, no. 11, pp. 3590–3600, Nov. 2010.

[2] F. Rusek, D. Persson, B. K. Lau, E. Larsson, T. Marzetta, O. Edfors, and F. Tufvesson, “Scaling up MIMO: Opportunities and challenges with very large arrays,” *IEEE Signal Process. Mag.*, vol. 30, no. 1, pp. 40–60, Jan. 2013.

[3] B. Yin, M. Wu, G. Wang, C. Dick, J. R. Cavallaro, and C. Studer, “A 3.8 Gb/s large-scale MIMO detector for 3GPP LTE-Advanced,” in *Proc. IEEE Int'l Conf. Acoust., Speech and Signal Process. (ICASSP)*, May 2014, pp. 3907–3911.

[4] B. Yin, M. Wu, J. R. Cavallaro, and C. Studer, “VLSI design of large-scale soft-output MIMO detection using conjugate gradients,” in *Proc. IEEE Int'l Conf. Circuits Syst. (ISCAS)*, May 2015, pp. 1498–1501.

[5] M. Wu, C. Dick, J. R. Cavallaro, and C. Studer, “FPGA design of a coordinate descent data detector for large-scale MU-MIMO,” in *Proc. IEEE Int'l Conf. Circuits Syst. (ISCAS)*, May 2016, pp. 1894–1897.

[6] ——, “High-throughput data detection for massive MU-MIMO-OFDM using coordinate descent,” *IEEE Trans. Circuits Syst. I*, Dec. 2016.

[7] O. Castañeda, T. Goldstein, and C. Studer, “Data detection in large multi-antenna wireless systems via approximate semidefinite relaxation,” *IEEE Trans. Circuits Syst. I*, pp. 2659–2662, Dec. 2016.

[8] A. Paulraj, R. Nabar, and D. Gore, *Introduction to Space-Time Wireless Communications*. Cambridge Univ. Press, 2003.

[9] W.-K. Ma, T. N. Davidson, K. M. Wong, Z.-Q. Luo, and P.-C. Ching, “Quasi-maximum-likelihood multiuser detection using semi-definite relaxation with application to synchronous cdma,” *IEEE Trans. Signal Process.*, vol. 50, no. 4, pp. 912–922, Apr. 2002.

[10] P. H. Tan, L. K. Rasmussen, and T. J. Lim, “Constrained maximum-likelihood detection in CDMA,” *IEEE Trans. Commun.*, vol. 49, no. 1, pp. 142–153, Jan. 2001.

[11] C. Jeon, A. Maleki, and C. Studer, “On the performance of mismatched data detection in large MIMO systems,” in *Proc. IEEE Int'l Symp. Inf. Theory (ISIT)*, Jul. 2016, pp. 180–184.

[12] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, “Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers,” *Foundations and Trends in Machine Learning*, 2010.