

Hardware-Friendly Two-Stage Spatial Equalization for All-Digital mmWave Massive MU-MIMO

Oscar Castañeda¹, Sven Jacobsson^{2,3}, Giuseppe Durisi³, Tom Goldstein⁴, and Christoph Studer¹

¹ETH Zurich, Zurich, Switzerland; e-mail: caoscar@ethz.ch, studer@ethz.ch

²Ericsson Research, Gothenburg, Sweden; e-mail: sven.jacobsson@ericsson.com

³Chalmers University of Technology, Gothenburg, Sweden; e-mail: durisi@chalmers.se

⁴University of Maryland, College Park, MD; e-mail: tomg@cs.umd.edu

Abstract—Next generation wireless communication systems are expected to combine millimeter-wave communication with massive multi-user multiple-input multiple-output technology. All-digital base-station implementations for such systems need to process high-dimensional data at extremely high rates, which results in excessively high power consumption. In this paper, we propose two-stage spatial equalizers that first reduce the problem dimension by means of a hardware-friendly, low-resolution linear transform followed by spatial equalization on a lower-dimensional signal. We consider adaptive and non-adaptive dimensionality reduction strategies and demonstrate that the proposed two-stage spatial equalizers are able to approach the performance of conventional linear spatial equalizers that directly operate on high-dimensional data, while offering the potential to reduce the power consumption of spatial equalization.

I. INTRODUCTION

In order to meet the ever-growing demand for higher communication data-rates, next generation wireless systems are anticipated to combine millimeter-wave (mmWave) communication [1] with massive multi-user multiple-input multiple-output (MU-MIMO) technology [2]. The abundance of bandwidth at mmWave frequencies, together with the fine-grained beamforming capabilities of massive MU-MIMO systems, enables high-throughput communication with several user equipments (UEs) in the same time-frequency resource. Such emerging systems require base-stations (BSs) that are equipped with hundreds of antenna elements, where each antenna must process data at extremely high rates using computationally-complex baseband-processing circuitry. As a result, power consumption is a key concern when designing practical BSs for mmWave massive MU-MIMO systems.

To arrive at energy-efficient mmWave MU-MIMO BS designs, hybrid analog-digital solutions [3]–[5] have been proposed in the past. However, such hybrid approaches are limited in their spatial multiplexing capabilities [5]–[7], which

reduces their spectral efficiency. All-digital BS architectures [8]–[10] are able to overcome this issue, but are commonly ascertained as energy inefficient. Recent results [7], [9] have shown that, by reducing the resolution of the data converters, the power consumption of radio-frequency (RF) circuitry and data converters in all-digital BS architectures is comparable to that of hybrid solutions. Furthermore, hybrid architectures require a non-trivial design effort for the analog circuitry. In contrast, all-digital architectures minimize the amount of RF circuitry which simplifies migration to the newest CMOS technology nodes. Nonetheless, the power consumption of baseband processing in all-digital BS architectures is—at this time—largely unexplored.

A. All-Digital Spatial Equalization

In the uplink (UEs transmit to BS), all-digital spatial equalization processes the signals coming from the data converters at the B BS antennas to recover the signals transmitted by each of the U UEs. Linear spatial equalization amounts to performing a complex-valued matrix-vector product between a $U \times B$ equalization matrix with each vector sampled at the B BS antennas. For a system with $B = 256$ BS antennas and $U = 16$ UEs, a single matrix-vector product at a rate of 2 G samples/s already consumes 28 W in 28 nm CMOS [11]. For future systems that operate at higher bandwidths, or with more BS antennas or UEs, the power consumption will increase even further. Evidently, efficient spatial-equalization circuitry is necessary to reduce the power consumption of all-digital BS architectures without compromising spectral efficiency.

One way of reducing the power consumption of matrix-vector products is to decrease the bit resolution of their composing multiplications and additions. Existing work has mainly concentrated on using low-resolution (e.g., 1-bit to 8-bit) analog-to-digital converters at the BS antennas [5], [7]–[9], [12], [13], which reduces the precision of the received vectors. However, the spatial equalization matrix is typically represented with high-resolution numbers, e.g., 10-bit to 12-bit [14], [15]. In order to exploit this insight, we proposed in [11] to use matrices with low-resolution entries—a new paradigm we refer to as *finite-alphabet equalization*. When compared to traditional, high-resolution equalization, finite-alphabet equalization reduces power consumption by at least $3.9\times$ when

The work of OC was supported in part by ComSenTer, a Semiconductor Research Corporation (SRC) program, by SRC nCORE task 2758.004, and by a Qualcomm Innovation Fellowship. The work of SJ and GD was supported by the Swedish Foundation for Strategic Research under grant ID14-0022, and by the Swedish Governmental Agency for Innovation Systems (VINNOVA). The work of TG was supported in part by the US NSF under grant CCF-1535902 and by the US Office of Naval Research under grant N00014-17-1-2078. The work of CS was supported in part by Xilinx Inc. and by the US NSF under grants ECCS-1408006, CCF-1535897, CCF-1652065, CNS-1717559, and ECCS-1824379.

using 1 bit of resolution to represent the beamforming weights, and offers competitive performance when using as few as 3 bits. Unfortunately, for fewer than 3 bits of resolution, the performance loss of finite-alphabet equalization compared to high-resolution equalization is still considerable.

B. Contributions

To arrive at power-efficient all-digital spatial equalization, we propose a two-stage spatial equalization approach. The first stage uses a low-resolution matrix to reduce the dimensionality of the B -dimensional data received by the BS antennas. The second stage then performs spatial equalization on the low-dimensional output of the first stage with higher-resolution matrix-vector products. We study both adaptive and non-adaptive approaches to implement the first stage, and we compare the computational complexity and performance of two-stage equalizers with both high-resolution and finite-alphabet single-stage equalizers.

C. Notation

Uppercase and lowercase boldface letters represent matrices and vectors, respectively. For a matrix \mathbf{A} , the Hermitian transpose is \mathbf{A}^H and the Frobenius norm is $\|\mathbf{A}\|_F$. The $M \times M$ identity matrix is \mathbf{I}_M . For a vector \mathbf{a} , the k th entry is a_k , the ℓ_2 -norm is $\|\mathbf{a}\|_2$, and the entry-wise complex conjugate is \mathbf{a}^* . The floor function $\lfloor a \rfloor$ returns the largest integer such that $\lfloor a \rfloor \leq a$. The expectation with respect to the random vector \mathbf{x} is represented with $\mathbb{E}_{\mathbf{x}}[\cdot]$.

II. SYSTEM MODEL AND SPATIAL EQUALIZATION

A. Uplink System Model

We focus on the uplink of a massive MU-MIMO system in which U single-antenna UEs transmit data to a BS with $B \gg U$ antennas. We consider a narrowband scenario with input-output relation $\mathbf{y} = \mathbf{H}\mathbf{s} + \mathbf{n}$. Here, $\mathbf{y} \in \mathbb{C}^B$ is the vector received at the BS, $\mathbf{H} \in \mathbb{C}^{B \times U}$ is the uplink MU-MIMO channel matrix, $\mathbf{s} \in \mathcal{S}^U$ is a vector containing the symbols transmitted by the UEs, where \mathcal{S} is the constellation (e.g., 16-QAM), and $\mathbf{n} \in \mathbb{C}^B$ is an i.i.d. circularly-symmetric complex Gaussian noise vector with covariance matrix $\mathbf{C}_{\mathbf{n}} = \mathbb{E}_{\mathbf{n}}[\mathbf{n}\mathbf{n}^H] = N_0\mathbf{I}_B$. We assume that the symbols in \mathbf{s} are i.i.d. with zero mean and variance E_s , so that their covariance matrix $\mathbf{C}_{\mathbf{s}} = \mathbb{E}_{\mathbf{s}}[\mathbf{s}\mathbf{s}^H] = E_s\mathbf{I}_U$. Furthermore, we assume that the channel stays constant across several symbol transmissions, so that the BS is able to accurately estimate the channel matrix \mathbf{H} . For simplicity, we assume perfect channel knowledge at the BS.

B. L-MMSE Spatial Equalization

Spatial equalization is a critical task at the BS, which forms estimates $\hat{\mathbf{s}}$ of the transmitted vector \mathbf{s} given the received vector \mathbf{y} and knowledge of the channel matrix \mathbf{H} . Due to their simplicity, linear estimators are of interest for mmWave massive MU-MIMO systems. Hence, we focus on linear spatial equalizers of the form $\hat{\mathbf{s}} = \mathbf{W}^H\mathbf{y}$, where the equalization matrix $\mathbf{W}^H \in \mathbb{C}^{U \times B}$ is typically designed to minimize the mean squared-error (MSE) given by $\mathbb{E}_{\mathbf{s}, \mathbf{n}}[\|\hat{\mathbf{s}} - \mathbf{s}\|_2^2]$. The matrix

\mathbf{W}^H that minimizes the MSE is known as the linear minimum MSE (L-MMSE) equalizer, which corresponds to [16]

$$\mathbf{W}^H = (\mathbf{H}^H\mathbf{H} + \rho\mathbf{I}_U)^{-1}\mathbf{H}^H, \quad (1)$$

where $\rho = N_0/E_s$. In practical implementations, the entries of the L-MMSE equalizer usually have high resolution (e.g., 10 bits to 12 bits [14], [15]), which leads to power-hungry digital baseband-processing circuitry.

C. Finite-Alphabet Equalization

To reduce the power consumption of spatial equalization, we proposed finite-alphabet equalization in [11], an approach in which each spatial equalization matrix \mathbf{V}^H has the structure:

$$\mathbf{V}^H = \text{diag}(\boldsymbol{\beta}^*)\mathbf{X}^H. \quad (2)$$

Here, $\mathbf{X}^H \in \mathcal{X}^{U \times B}$ is a low-resolution matrix with entries coming from a low-cardinality finite alphabet \mathcal{X} (e.g., the “1-bit” alphabet is $\{\pm 1 \pm j\}$), and $\boldsymbol{\beta} \in \mathbb{C}^U$ contains per-UE high-resolution scaling factors. According to [11], the main problem when designing a finite-alphabet equalization matrix (2) is to find the low-resolution matrix \mathbf{X}^H . Once this matrix is determined, the scaling factors in $\boldsymbol{\beta}$ are computed using [11, Eq. (13)]. In [11], we outline two procedures to obtain the low-resolution matrix \mathbf{X}^H : (i) per-row uniform quantization of each entry of the L-MMSE matrix (called FL-MMSE), and (ii) an iterative algorithm to approximately solve the finite-alphabet minimum MSE equalization (FAME) problem in [11, Eq. (12)] using forward-backward splitting (FBS) for t_{\max} iterations (called FAME-FBS). As shown in [11], while both procedures have a computational complexity of $O(BU^2)$, FAME-FBS significantly outperforms FL-MMSE when using resolutions lower than 3 bits; otherwise, their performance is similar. Regardless of the procedure used to compute the finite-alphabet matrix, spatial equalizers with the structure in (2) reduce hardware complexity as the multiplication with \mathbf{X}^H can be implemented using low-resolution multipliers and adders.

III. TWO-STAGE SPATIAL EQUALIZATION

We will now present a novel spatial equalization approach that consists of two stages. In the first stage (S1), the spatial equalizer projects the B -dimensional received vector \mathbf{y} to a K -dimensional vector \mathbf{z} , where $U \leq K \ll B$. In the second stage (S2), the spatial equalizer performs high-resolution equalization over the S1 output \mathbf{z} to form the estimates $\hat{\mathbf{s}} \in \mathbb{C}^U$.

A. Spatial Equalization Stage One (S1)

To keep circuitry simple, we implement S1 as a linear operation, so that $\mathbf{z} = \mathbf{X}_1^H\mathbf{y}$, where $\mathbf{X}_1^H \in \mathcal{X}^{K \times B}$ and \mathcal{X} is a low-cardinality finite alphabet whose elements can be represented with few bits. As a result, S1 can be performed using energy-efficient, low-resolution hardware multipliers and adders. We consider two kinds of S1 matrices \mathbf{X}_1^H : non-adaptive and adaptive matrices.

Non-adaptive S1 matrices \mathbf{X}_1^H do not depend on channel state information and are not updated as the channel changes.

TABLE I
COMPLEXITY OF COMPUTING SINGLE-STAGE EQUALIZERS

Algorithm	Computational complexity	Scaling
L-MMSE	$6BU^2 + 2U^3 - 2BU - 2U + 1$	$O(BU^2)$
FL-MMSE	$10BU^2 + 2U^3 + 2U^2 + U + 1$	$O(BU^2)$
FAME-FBS	$(8t_{\max} + 4)BU^2 + (4t_{\max} + 2)BU + 2U^2 + (2t_{\max} + 3)U$	$O(BU^2)$

We consider two non-adaptive $K \times B$ matrices: (i) a sub-sampled Hadamard matrix, where we take rows $1, \lfloor B/K \rfloor + 1, 2\lfloor B/K \rfloor + 1, \dots, (K-1)\lfloor B/K \rfloor + 1$ of a $B \times B$ Hadamard matrix and (ii) a randomly-generated matrix, where the entries are chosen uniformly from the finite alphabet \mathcal{X} . Note that a matrix-vector product with these matrices can be performed with low-resolution constant multipliers, which consume lower power than low-resolution multipliers where both operands vary [17].

Adaptive S1 matrices \mathbf{X}_1^H depend on the channel state information and are updated as the channel fluctuates. We will consider three adaptive $K \times B$ matrices: (i) the low-resolution matrix \mathbf{X}^H of the FL-MMSE finite-alphabet equalizer, computed as detailed in Section II-C, (ii) the low-resolution matrix \mathbf{X}^H of the FAME-FBS equalizer, and (iii) the maximum ratio combining (MRC) equalizer \mathbf{H}^H , uniformly quantized on a per-row basis. We will refer to this last approach as finite-alphabet MRC (FMRC). Note that, for all three cases, $K = U$.

B. Spatial Equalization Stage Two (S2)

For S2, we focus on linear spatial equalizers so that $\hat{\mathbf{s}} = \mathbf{W}_2^H \mathbf{z}$, where \mathbf{W}_2^H is a $U \times K$ equalization matrix whose entries are not constrained to be low resolution. Since the second stage operates on the K -dimensional signal \mathbf{z} , where $K \ll B$, the complexity of S2 is significantly lower than that of a traditional spatial equalizer that operates directly on the B -dimensional received vector \mathbf{y} . Once the S1 matrix \mathbf{X}_1^H is fixed, we can compute the (infinite-resolution) matrix that minimizes the MSE by solving the following problem:

$$\mathbf{W}_2^H = \arg \min_{\tilde{\mathbf{W}}^H \in \mathbb{C}^{U \times K}} \mathbb{E}_{\mathbf{s}, \mathbf{n}} \left[\|\tilde{\mathbf{W}}^H \mathbf{z} - \mathbf{s}\|_2^2 \right] \quad (3)$$

$$= \arg \min_{\tilde{\mathbf{W}}^H \in \mathbb{C}^{U \times K}} \mathbb{E}_{\mathbf{s}, \mathbf{n}} \left[\|\tilde{\mathbf{W}}^H \mathbf{X}_1^H (\mathbf{H}\mathbf{s} + \mathbf{n}) - \mathbf{s}\|_2^2 \right] \quad (4)$$

$$= \arg \min_{\tilde{\mathbf{W}}^H \in \mathbb{C}^{U \times K}} \|\mathbf{H}^H \mathbf{X}_1 \tilde{\mathbf{W}} - \mathbf{I}_U\|_F^2 + \rho \|\mathbf{X}_1 \tilde{\mathbf{W}}\|_F^2. \quad (5)$$

The resulting L-MMSE matrix for S2 is given by

$$\mathbf{W}_2^H = \mathbf{H}^H \mathbf{X}_1 (\mathbf{X}_1^H \mathbf{H} \mathbf{H}^H \mathbf{X}_1 + \rho \mathbf{X}_1^H \mathbf{X}_1)^{-1}. \quad (6)$$

Note that, if $\mathbf{X}_1^H = \mathbf{I}_B$ and the Woodbury matrix identity [18] is applied, then the S2 L-MMSE equalizer in (6) is the same as the traditional L-MMSE equalizer in (1).

C. Computational Complexity

Tables I and II show the computational complexity for generating equalization matrices for single-stage and two-stage

TABLE II
COMPLEXITY OF COMPUTING TWO-STAGE EQUALIZERS

Algorithm	Computational complexity	Scaling
FL-MMSE	$6BU^2 + 2U^3 - 2BU - 2U + 1$	$O(BU^2)$
FAME-FBS	$8t_{\max}BU^2 + 4t_{\max}BU + 2t_{\max}U$	$O(BU^2)$
L-MMSE	$2BK^2 + 4BKU + 2K^3 + 6K^2U + K^2 - 2KU - 2K + 1$	$O(BK^2)$

TABLE III
COMPLEXITY OF APPLYING EQUALIZATION

Equalization	Real-valued multiplication count	
	High-resolution	Low-resolution
Traditional	$4BU$	0
Finite-alphabet	$4U$	$4BU$
Two-stage	$4KU$	$4BK$

equalizers, respectively. We characterize computational complexity as the number of required real-valued multiplications, where we count one complex-valued multiplication as four real-valued multiplications. The complexity counts for the single-stage equalizers in Table I are taken from [11, Table I].

For the two-stage equalizers, we consider the following S1 matrices \mathbf{X}_1^H : Random, sub-sampled Hadamard, FMRC, FL-MMSE and FAME-FBS. Since Random and Hadamard matrices are non-adaptive, they do not cause any computational complexity as they are generated once when designing the system. For FMRC, we only need to quantize the \mathbf{H}^H matrix, a complexity that we will ignore as it can be executed efficiently in hardware. For FL-MMSE and FAME-FBS, their complexity for S1 (Table II) is lower than for single-stage, finite-alphabet equalization (Table I), as S1 only needs the low-resolution matrix \mathbf{X}^H while the single-stage, finite-alphabet equalization also requires computation of the scaling factors in β .

Besides computing the low-resolution S1 matrix \mathbf{X}_1^H , two-stage spatial equalizers must also compute the high-resolution S2 matrix \mathbf{W}_2^H , which we obtain using S2 L-MMSE. From Table II, we observe that S2 L-MMSE has a complexity of $O(BK^2)$. However, we must consider that the $2BK^2 + 4BKU$ products required by S2 L-MMSE to compute $\mathbf{H}^H \mathbf{X}_1$ and $\mathbf{X}_1^H \mathbf{X}_1$ can be performed using simple low-resolution hardware. Hence, if we only consider high-resolution multiplications on Table II, then the complexity of S2 L-MMSE is $O(K^3)$. Thus, for the case where $K = U$ (e.g., when using adaptive S1 matrices), the computational complexity $O(U^3)$ of S2 L-MMSE is significantly lower than the complexity $O(BU^2)$ of traditional L-MMSE. In short, considering both stages, computing two-stage adaptive equalizers has the same asymptotic scaling of $O(BU^2)$ as computing single-stage equalizers. We note that the equalization matrices are recomputed each time that the channel changes.

Table III reports the computational complexity of applying spatial equalization on each new received vector \mathbf{y} . To summarize, single-stage finite-alphabet equalization uses

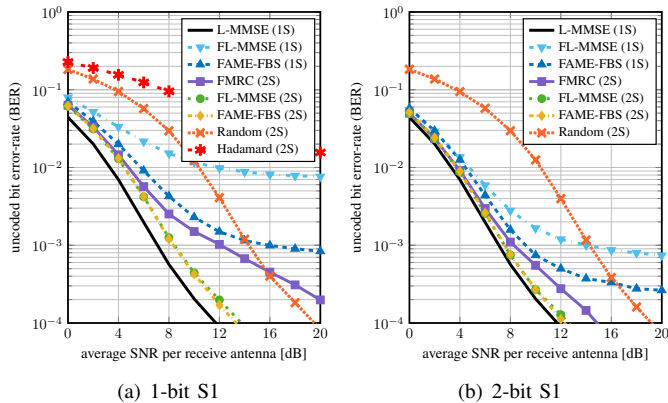


Fig. 1. Uncoded BER for a $B = 256$ BS antenna, $U = 16$ UE, 16-QAM system in a line-of-sight (LoS) mmWave channel. All two-stage (2S) equalizers use S2 L-MMSE \mathbf{W}_2^H for the second stage. FAME-FBS runs $t_{\max} \leq 20$ iterations and is initialized with \mathbf{H}^H .

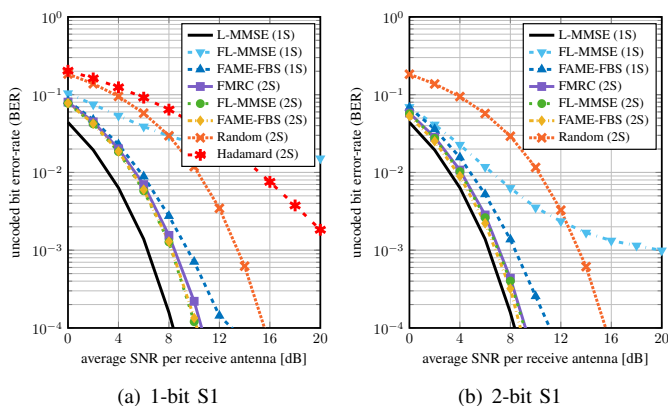


Fig. 2. Uncoded BER for a $B = 256$ BS antenna, $U = 16$ UE, 16-QAM system in a non-line-of-sight (non-LoS) mmWave channel. All two-stage (2S) equalizers use S2 L-MMSE \mathbf{W}_2^H . FAME-FBS runs $t_{\max} \leq 20$ iterations and is initialized with \mathbf{H}^H .

low-resolution operations to perform all the multiplications of a traditional equalizer, and then performs U additional high-resolution, complex-valued multiplications to bring the equalized signals onto the correct scale. Two-stage spatial equalization performs at least as many low-resolution multiplications as finite-alphabet equalization, but performs more high-resolution operations in order to achieve a better performance, as we will show in Section IV. Nonetheless, note that the number of high-resolution multipliers required by two-stage equalization is significantly lower than those required by traditional equalization. Hence, if K and the resolution of the S1 matrix \mathbf{X}_1^H are sufficiently small, two-stage spatial equalization reduces complexity when compared to traditional linear spatial equalization methods.

IV. SIMULATION RESULTS

Figures 1 and 2 show uncoded bit error rates (BERs) for various single-stage (1S) and two-stage (2S) spatial equalizers under mmWave line-of-sight (LoS) and non-LoS channels,

respectively, when using S1 matrices with 1 or 2 bits of resolution. The simulation results correspond to a $B = 256$ BS antenna, $U = 16$ UE, 16-QAM system. For the single-stage spatial equalizers, we consider the traditional, high-resolution L-MMSE, as well as the finite-alphabet equalizers FL-MMSE and FAME-FBS. For the two-stage spatial equalizers, S2 is always implemented using S2 L-MMSE \mathbf{W}_2^H , while S1 varies. For the non-adaptive and adaptive S1 matrices \mathbf{X}_1^H , we use $K = 64$ and $K = 16$, respectively. The mmWave channels are generated with the QuaDRiGa model [19], using a uniform linear array with half-wavelength spacing in the “mmMAGIC_UMi” scenario and a 60 GHz carrier frequency. We also simulate power control to ensure a maximum ± 3 dB power variation across UEs.

Figures 1 and 2 show that two-stage spatial equalizers with an adaptive S1 are able to outperform single-stage, finite-alphabet equalizers for all considered scenarios. In stark contrast, two-stage spatial equalizers with a non-adaptive S1 perform worse than finite-alphabet equalizers at low SNR. While FAME-FBS significantly outperforms FL-MMSE when considering single-stage, finite-alphabet equalization, they offer virtually the same performance in the two-stage approach. What is more, their performance is met by the simple FMRC-based two-stage spatial equalizer for non-LoS scenarios, although not for the LoS case. Finally, we note that two-stage spatial equalization with a 2-bit S1 matrix achieves a performance that is within 1 dB of the infinite-resolution L-MMSE equalizer.

V. CONCLUSIONS

We have analyzed the performance and complexity of two-stage spatial equalizers that consist of a linear transform to project the received data on a lower-dimensional space, followed by high-resolution spatial equalization. Since the first stage can be implemented using low-resolution hardware and the second-stage equalization is performed over a lower-dimensional signal, this two-stage approach promises lower power compared to traditional high-resolution equalization. Our simulation results have shown that 2-bit adaptive projections for the first stage and a high-resolution L-MMSE equalizer for the second stage achieves an error-rate performance within 1 dB of a traditional high-resolution, single-stage equalizer. Moreover, two-stage equalization outperforms the error-rate of finite-alphabet equalization, a single-stage approach that also targets power reduction in all-digital mmWave massive MU-MIMO spatial equalization. However, this improved performance comes at the cost of a higher computational complexity.

There are many avenues for future work. To quantify the energy-efficiency of all-digital, two-stage spatial equalizers, a hardware evaluation is in order. The application of other methods for the second stage, such as finite-alphabet equalization, is an interesting research direction. Finally, algorithms that jointly design the first-stage transform with the second-stage equalizer might further improve error-rate performance.

REFERENCES

- [1] A. L. Swindlehurst, E. Ayanoglu, P. Heydari, and F. Capolino, "Millimeter-wave massive MIMO: The next wireless revolution?" *IEEE Commun. Mag.*, vol. 52, no. 9, pp. 56–62, Sep. 2014.
- [2] E. G. Larsson, F. Tufvesson, O. Edfors, and T. L. Marzetta, "Massive MIMO for next generation wireless systems," *IEEE Commun. Mag.*, vol. 52, no. 2, pp. 186–195, Feb. 2014.
- [3] W. Roh, J.-Y. Seol, J. Park, B. Lee, J. Lee, Y. Kim, J. Cho, K. Cheun, and F. Aryanfar, "Millimeter-wave beamforming as an enabling technology for 5G cellular communications: Theoretical feasibility and prototype results," *IEEE Commun. Mag.*, vol. 52, no. 2, pp. 106–113, Feb. 2014.
- [4] B. Sadhu, Y. Tousi, J. Hallin, S. Sahl, S. Reynolds, Ö. Renström, K. Sjögren, O. Haapalahti, N. Mazor, B. Bokinge, G. Weibull, H. Bengtsson, A. Carlinger, E. Westesson, J. Thillberg, L. Rexberg, M. Yeck, X. Gu, D. Friedman, and A. Valdes-Garcia, "A 28GHz 32-element phased-array transceiver IC with concurrent dual polarized beams and 1.4 degree beam-steering resolution for 5G communication," in *IEEE Int. Solid-State Circuits Conf. (ISSCC)*, Feb. 2017, pp. 128–129.
- [5] A. Alkhateeb, J. Mo, N. González-Prelcic, and R. W. Heath Jr., "MIMO precoding and combining solutions for millimeter-wave systems," *IEEE Commun. Mag.*, vol. 52, no. 12, pp. 122–131, Dec. 2014.
- [6] E. Björnson, L. Van der Perre, S. Buzzi, and E. G. Larsson, "Massive MIMO in sub-6 GHz and mmWave: Physical, practical, and use-case differences," *IEEE Wireless Commun.*, vol. 26, no. 2, pp. 100–108, Apr. 2019.
- [7] S. Dutta, C. N. Barati, A. Dhananjay, D. A. Ramirez, J. F. Buckwalter, and S. Rangan, "A case for digital beamforming at mmWave," *IEEE Trans. Wireless Commun.*, vol. 19, no. 2, pp. 756–770, Feb. 2020.
- [8] J. Mo and R. W. Heath Jr., "Capacity analysis of one-bit quantized MIMO systems with transmitter channel state information," *IEEE Trans. Signal Process.*, vol. 63, no. 20, pp. 5498–5512, Oct. 2015.
- [9] K. Roth and J. A. Nossek, "Achievable rate and energy efficiency of hybrid and digital beamforming receivers with low resolution ADC," *IEEE J. Sel. Areas Commun.*, vol. 35, no. 9, pp. 2056–2068, Sep. 2017.
- [10] S. Jacobsson, G. Durisi, M. Coldrey, U. Gustavsson, and C. Studer, "Throughput analysis of massive MIMO uplink with low-resolution ADCs," *IEEE Trans. Wireless Commun.*, vol. 16, no. 6, pp. 4038–4051, Jun. 2017.
- [11] O. Castañeda, S. Jacobsson, G. Durisi, T. Goldstein, and C. Studer, "Finite-alphabet MMSE equalization for all-digital massive MU-MIMO mmWave communication," *IEEE J. Sel. Areas Commun.*, 2020, to appear.
- [12] C. Studer and G. Durisi, "Quantized massive MU-MIMO-OFDM uplink," *IEEE Trans. Commun.*, vol. 64, no. 6, pp. 2387–2399, Jun. 2016.
- [13] H. Yan, S. Ramesh, T. Gallagher, C. Ling, and D. Cabric, "Performance, power, and area design trade-offs in millimeter-wave transmitter beamforming architectures," *IEEE Circuits Syst. Mag.*, vol. 19, no. 2, pp. 33–58, May 2019.
- [14] C. Studer, S. Fateh, and D. Seethaler, "ASIC implementation of soft-input soft-output MIMO detection using MMSE parallel interference cancellation," *IEEE J. Solid-State Circuits*, vol. 46, no. 7, pp. 1754–1765, Jul. 2011.
- [15] M. Wu, B. Yin, G. Wang, C. Dick, J. R. Cavallaro, and C. Studer, "Large-scale MIMO detection for 3GPP LTE: Algorithms and FPGA implementations," *IEEE J. Sel. Topics Signal Process.*, vol. 8, no. 5, pp. 916–929, Oct. 2014.
- [16] A. Paulraj, R. Nabar, and D. Gore, *Introduction to space-time wireless communications*. Cambridge Univ. Press, 2003.
- [17] R. Zimmerman, "Computer arithmetic: Principles, architectures, and VLSI design," Integrated Systems Laboratory, ETH Zürich, Tech. Rep., 1999.
- [18] K. B. Petersen and M. S. Pedersen, "The matrix cookbook," Nov. 2012.
- [19] S. Jaeckel, L. Raschkowski, K. Börner, and L. Thiele, "QuaDRiGa: A 3-D multi-cell channel model with time evolution for enabling virtual field trials," *IEEE Trans. Antennas Propag.*, vol. 62, no. 6, pp. 3242–3256, Jun. 2014.