Quantized Massive MU-MIMO-OFDM Uplink

Journal Article

Author(s):

Studer, Christoph Durisi, Giuseppe

Publication date:

2016-06

Permanent link:

<https://doi.org/10.3929/ethz-b-000455414>

Rights / license:

In Copyright - Non-Commercial Use Permitted

Originally published in:

IEEE Transactions on Communications 64(6), <https://doi.org/10.1109/TCOMM.2016.2558151>

Quantized Massive MU-MIMO-OFDM Uplink

Christoph Studer, *Senior Member, IEEE*, and Giuseppe Durisi, *Senior Member, IEEE*

Abstract—Coarse quantization at the base station (BS) of a massive multi-user (MU) multiple-input multiple-output (MIMO) wireless system promises significant power and cost savings. Coarse quantization also enables significant reductions of the raw analog-to-digital converter (ADC) data that must be transferred from a spatially-separated antenna array to the baseband processing unit. The theoretical limits as well as practical transceiver algorithms for such quantized MU-MIMO systems operating over frequency-flat, narrowband channels have been studied extensively. However, the practically relevant scenario where such communication systems operate over frequency-selective, wideband channels is less well understood. This paper investigates the uplink performance of a quantized massive MU-MIMO system that deploys orthogonal frequency-division multiplexing (OFDM) for wideband communication. We propose new algorithms for quantized maximum a-posteriori (MAP) channel estimation and data detection, and we study the associated performance/quantization trade-offs. Our results demonstrate that coarse quantization (e.g., four to six bits, depending on the ratio between the number of BS antennas and the number of users) in massive MU-MIMO-OFDM systems entails virtually no performance loss compared to the infinite-precision case at no additional cost in terms of baseband processing complexity.

Index Terms—Analog-to-digital conversion, convex optimization, forward-backward splitting (FBS), frequency-selective channels, massive multi-user multiple-input multiple-output (MU-MIMO), maximum a-posteriori (MAP) channel estimation, minimum mean-square error (MMSE) data detection, orthogonal frequency-division multiplexing (OFDM), quantization.

I. INTRODUCTION

We investigate the effects of coarse quantization at the base-station (BS) on the performance of uplink data transmission in a massive multi-user (MU) multiple-input multiple-output (MIMO) system. To enable reliable communication over frequency-selective channels, we develop new methods for (optimal) channel estimation and data detection for the case when massive MU-MIMO is combined with orthogonal frequency-division multiplexing (OFDM). Fig. 1 illustrates the massive MU-MIMO-OFDM uplink transmission system studied in this paper. We consider a coded massive MU-MIMO system where U independent single-antenna user terminals communicate with a BS equipped with $B \gg U$ antennas. Our

model includes coarse quantization in the analog-to-digital converters (ADCs) at the in-phase and quadrature outputs of each radio-frequency (RF) chain at the BS side. The square-dashed box at the BS side encompasses the proposed channel-estimation and data-detection algorithms, which—as we shall see—enable reliable wideband data transmission, even when the ADCs have low precision (e.g., four to six quantization bits).

A. Benefits of Quantized Massive MIMO

In conventional multiple-antenna BSs, each RF port is connected to two high-precision ADCs (typically, the in-phase and the quadrature signal components are quantized with resolutions exceeding 10 bit). Scaling such architectures to massive MIMO with hundreds or thousands active antenna elements would result in prohibitively high power consumption and hardware costs, especially for systems operating in the millimeter-wave frequency bands. Since the power consumption of an ADC scales roughly exponentially in the number of quantization bits [1], reducing the precision of the ADCs provides an effective means to accommodate a massive array of active antenna elements, while—at the same time—keeping the power consumption and system costs within desirable limits. Furthermore, the use of low-precision ADCs allows for a relaxation of the quality requirements on the RF circuitry (e.g., low-noise amplifiers, oscillators, and mixers), which enables a further reduction in both the power consumption and system costs. The underlying reason is that RF circuitry needs to operate at precision levels “just above” the quantization noise floor. Apart from power and cost reductions, coarsely quantized massive MIMO also reduces the amount of bits per second that needs to be transferred from the antenna unit (usually located on top of the cell tower) to the baseband processing unit (usually located at the bottom of the cell tower). The deployment of low-precision ADCs at the BS mitigates this data-transfer bottleneck.

B. Relevant Prior Art

Massive MU-MIMO promises to increase the spectral efficiency at reduced signal-processing complexity compared to that of conventional, small-scale MIMO systems, and is therefore considered to be a key technology for next-generation wireless systems [2], [3]. Most theoretical results on the performance of the massive MU-MIMO uplink are for the scenario where the base-station is equipped with infinite-precision ADCs. The performance impact of low-precision ADCs has been analyzed only recently.

Björnsson *et al.* [4] modeled the aggregate effect of various residual hardware impairments (including quantization noise) as additive Gaussian noise that is independent of the transmit

C. Studer is with the School of Electrical and Computer Engineering, Cornell University, Ithaca, NY (e-mail: studer@cornell.edu).

G. Durisi is with the Department of Signals and Systems, Chalmers University, Göteborg, Sweden (e-mail: durisi@chalmers.se).

The work of C. Studer was supported in part by Xilinx Inc. and by the US National Science Foundation (NSF) under grants ECCS-1408006 and CCF-1535897.

The work of G. Durisi was supported by the Swedish Foundation for Strategic Research under grants SM13-0028 and ID14-0022, and by the Swedish Governmental Agency for Innovation Systems (VINNOVA) within the VINN Excellence center Chase.

The system simulator for the quantized massive MU-MIMO-OFDM uplink studied in this paper will be available on GitHub: https://github.com/quantizedmassivemimo/mu_mimo_ofdm/

Fig. 1. Overview of the considered quantized massive MU-MIMO-OFDM uplink system. Left: U single-antenna user terminals; right: massive MIMO base station with $B \gg U$ antennas and coarse quantization in the analog-to-digital converters (ADCs). The proposed system includes OFDM to simplify communication over frequency-selective, wideband channels.

signal. Based on this model, it was concluded that massive MU-MIMO exhibits a certain degree of robustness against hardware impairments. While the underlying Gaussian impairment model was demonstrated to be accurate for a diverse set of residual RF impairments [5], [6], this model is a poor match for *coarse* quantization, which causes deterministic distortions that are—by nature—dependent on the transmit signal.

Risi *et al.* [7] analyzed the performance of a massive MU-MIMO uplink system under the assumptions that the ADCs have 1-bit resolution, that the users transmit QPSK symbols, that the receiver performs maximum-ratio combining or zero forcing, and that pilot transmission followed by least-squares (LS) channel estimation at the BS is used to acquire channel-state information (CSI) at the receiver. Their result—albeit limited to QPSK—shows that massive MU-MIMO is resilient against coarse-quantization noise. The motivation for using LS channel estimation is that the maximum a-posteriori (MAP) estimator is claimed to have a computational complexity that scales exponentially in the number of active users.

The analysis in [7] has been extended recently to higher-order modulation schemes in [8]. There, it was shown that amplitude information about the transmitted signal can be recovered even in the presence of a single-bit quantizer, provided that the number of receive antennas is sufficiently large and the signal-to-noise ratio (SNR) is not too high. The observation that noise can help recovering magnitude information under 1-bit quantization has also been made in the compressive-sensing literature [9], [10].

Liang and Zhang [11] considered a mixed-ADC massive MU-MIMO architecture, where few high-resolution ADCs are used together with many 1-bit ADCs. The high-resolution ADCs allow the receiver to acquire accurate channel-state information (CSI). Generalized mutual information was used to characterize the performance of this architecture for the case of Gaussian codebooks and mismatched nearest-neighbor decoding. In [11], it is also noted that dithering prior to quantization may help improving the performance.

In summary, the results in [4], [7], [8], [11] suggest that massive MU-MIMO provides resilience against noise caused by coarse quantization. However, all these results pertain to massive MU-MIMO systems operating over frequency-flat

fading channels. Such channels are rarely encountered in modern broadband wireless communication systems, which instead experience a high degree of frequency selectivity. This paper investigates the impact of coarse ADC quantization on the error-rate performance in systems operating over frequency-selective channels, which is—as we believe—more relevant for next-generation wireless systems. Although a frequency-selective channel can be converted into a set of independent frequency-flat channels by means of OFDM, such an orthogonalization relies on the assumption that infinite-precision ADCs are used on the time-domain signals (see Section III-A for the details).

Outside the massive-MIMO literature, the performance of communication systems employing low-precision ADCs has been previously analyzed, mainly in the context of ultra-wideband and millimeter-wave systems. The structure of the capacity-achieving input distribution for point-to-point AWGN channels with low-precision ADCs has been characterized in [12]. An extension to fading channels in the limiting regime of 1-bit ADCs for the case of both no CSI and of CSI available only at the receiver was provided in [13]–[15]. The case of full CSI at both the transmitter and receiver has been analyzed recently in [16]. In the low-SNR regime, it was shown in [17] that a zero-threshold comparator is not optimal for 1-bit ADCs. Instead, the capacity-achieving strategy requires the combination of flash signaling [18, Def. 2] with a suitably-chosen quantization threshold.

The optimization of the ADC levels in single-antenna communication systems operating over deterministic (i.e., non-fading) frequency-selective channels was tackled in [19] using bit error rate as the performance metric, and in [20] using the mutual information. Sampling faster than the Nyquist rate may be beneficial when using low-precision ADCs (see, e.g., [21], [22] and [23, Sec. VI]). For example, it was shown in [24] that by oversampling at the receiver, one can reliably transmit 16-QAM symbols over a 1-bit-quantized single-antenna AWGN channel.

The strong nonlinearity introduced by low-precision ADCs renders the tasks of MIMO channel estimation and data detection challenging. Such problems can be cast in the more general framework of parameter estimation in the presence

of quantized output measurements [25]–[30]. The specific case of channel estimation has been discussed in [31], where the maximum-likelihood (ML) estimate is computed via expectation maximization. For the case of 1-bit ADCs combined with time-multiplexed pilot transmission, it was shown recently that the ML estimate has a closed-form expression [32]. Data detection algorithms have been discussed, e.g., in [33]–[36]. All these works, however, deal exclusively with systems operating over frequency-flat (or narrowband) fading channels.

C. Contributions

We investigate the effects of coarse quantization on a massive MU-MIMO-OFDM uplink system operating over frequency-selective, wideband channels. We develop new channel-estimation and data-detection algorithms, and investigate their performance/complexity trade-offs. Our main contributions are summarized as follows:

- We show that, under suitable assumptions on the statistics of the fading channel and of the noise, MAP channel estimation in the presence of (coarse) quantization is a convex problem that can be solved exactly using computationally efficient numerical methods.
- We formulate the problem of MAP data detection in the presence of quantization and develop a low-complexity minimum-mean square error (MMSE) data-detection algorithm.
- We develop channel-estimation and data-detection algorithms for two *mismatched* quantization models, which trade error-rate performance for computational complexity. This trade-off is investigated through a comprehensive set of numerical simulations.

Our results demonstrate that massive MIMO enables the use of coarse quantization at the BS, without a significant performance loss and without an increase in computational complexity compared to the case of systems with infinite-precision quantizers. Although we focus on the massive MU-MIMO-OFDM uplink, our algorithms also apply to traditional, small-scale MIMO-OFDM or even single-antenna OFDM systems.

D. Notation

Lowercase and uppercase boldface letters designate column vectors and matrices, respectively. For a matrix $\mathbf{A}$, we indicate its conjugate transpose by $\mathbf{A}^H$. The ℓ th column of the matrix $\mathbf{A}$ is denoted by $\mathbf{a}_\ell$, the entry on the k th row and on the ℓ th column is $A_{k,\ell}$, and the k th entry of a vector $\mathbf{a}$ is $[\mathbf{a}]_k = a_k$. The $M \times M$ identity matrix is denoted by $\mathbf{I}_M$ and the $M \times N$ all-zeros matrix by $\mathbf{0}_{M \times N}$. The real and imaginary parts of a complex scalar a are a^R and a^I , respectively. The multivariate complex-valued circularly-symmetric Gaussian probability density function with covariance matrix $\mathbf{K}$ is denoted by $\mathcal{CN}(\mathbf{0}, \mathbf{K})$. We use $p(\cdot)$ to denote generic probability mass functions, whereas $f(\cdot)$ is reserved for probability density functions.

E. Paper Outline

The rest of the paper is organized as follows. Section II introduces the quantization model. Section III describes the

proposed channel-estimation and data-detection algorithms for a simplified single-antenna OFDM system. Section IV deals with the extension of the proposed algorithms to the full-fledged quantized MU-MIMO-OFDM system illustrated in Fig. 1. Section V summarizes the numerical methods used to solve the convex optimization problems for channel estimation and data detection. Section VI provides numerical simulation results. We conclude in Section VII.

II. EXACT AND MISMATCHED QUANTIZATION MODELS

We now introduce the quantization model considered throughout the paper. We also propose a mismatched quantization model that enables suboptimal receiver algorithms requiring lower computational complexity.

A. Quantization of Complex-Valued Data

The in-phase and quadrature components of an RF-chain's outputs are typically converted to the digital domain using a pair of ADCs. After a sample-and-hold stage, the analog time-domain sample $s \in \mathbb{C}$ is mapped onto a finite-cardinality quantization alphabet according to $q = \mathcal{Q}(s)$, where $\mathcal{Q}(\cdot) : \mathbb{C} \rightarrow \mathcal{A}_c$ is a complex-valued scalar quantizer and $\mathcal{A}_c$ is the complex-valued quantization alphabet. In what follows, we focus on scalar quantizers that operate independently on the in-phase (real) and on the quadrature (imaginary) part of each sample. With a slight abuse of notation, we frequently apply the quantizer to a vector/matrix, with the understanding that quantization is applied entry-wise to the vector/matrix.

The real part s^R and the imaginary part s^I of each complex sample s are mapped onto a label from the quantization set $\mathcal{A} = \{1, \dots, Q\}$ (the same for both real and imaginary part) of cardinality $Q = |\mathcal{A}|$. Hence, we have that $\mathcal{A}_c = \mathcal{A} \times \mathcal{A}$. Each label describes the quantization bin in which s^R and s^I fall into. Consequently, the quantizer output $\mathcal{Q}(s)$ associated to a complex-valued sample $s = s^R + js^I$ consists of a pair $(q^R, q^I) \in \mathcal{A} \times \mathcal{A}$, or, equivalently, of a complex number $q = q^R + jq^I$. The labels are determined by comparing s^R and s^I with $Q + 1$ quantization bin boundaries $\{b_q\}$:

$$-\infty = b_1 < b_2 < \dots < b_Q < b_{Q+1} = +\infty.$$

Specifically, if $b_m \leq s^R < b_{m+1}$, $m = 1, \dots, Q$, then $q^R = m$. The imaginary part of the quantizer output q^I is obtained by applying the same rule to s^I . We use $Q_b = \log_2 |\mathcal{A}|$ to designate the number of quantization bits per real dimension. The total number of bits per complex-valued sample is $2Q_b$.

In what follows, we shall focus on the scenario where the thermal noise in the system (which appears before the quantizer) is a stationary memoryless circularly symmetric complex Gaussian process, whose samples have variance $\sigma^2 = N_0/2$ per real and imaginary part. With this assumption, we obtain the following nonlinear input-output relation for the quantizer:

$$q = \mathcal{Q}(z + n). \quad (1)$$

Here, $z \in \mathbb{C}$ is the noiseless, unquantized (time-domain) sample and $n \sim \mathcal{CN}(0, N_0)$ is the thermal (receive) noise that is added to the signal of interest prior to quantization. Throughout the paper, we make frequent use of the likelihood $p(q|z)$ of a

quantization label $q \in \mathcal{A}_c$ given a noiseless unquantized time-domain sample z . To compute this quantity, we use the facts that (i) the real and the imaginary parts of the thermal noise n are independent and (ii) the complex-valued quantizer $\mathcal{Q}(\cdot)$ operates independently on the real and the imaginary part. These two properties imply that

$$p(q | z = z^R + iz^I) = p(q^R | z^R)p(q^I | z^I). \quad (2)$$

We can now compute the probability of observing the real part q^R of the quantization label, given the real part z^R of the noiseless unquantized sample as follows [28], [34], [37]:

$$\begin{aligned} p(q^R | z^R) &= p(\ell(q^R) \leq z^R + n^R < u(q^R)) \\ &= \int_{\ell(q^R)}^{u(q^R)} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(z^R - \nu)^2}{2\sigma^2}\right) d\nu \\ &= \Phi\left(\frac{u(q^R) - z^R}{\sigma}\right) - \Phi\left(\frac{\ell(q^R) - z^R}{\sigma}\right). \end{aligned} \quad (3)$$

Here, $\Phi(a) = \int_{-\infty}^a (1/\sqrt{2\pi}) \exp(-\nu^2/2) d\nu$ is the cumulative distribution function of a (real-valued) standard normal random variable; $u(q^R) = b_{(q^R+1)}$ and $\ell(q^R) = b_{(q^R)}$ are the upper and the lower bin boundary positions associated with the quantized measurement $q^R \in \mathcal{A}$, respectively. The probability of observing the imaginary part q^I of the quantization label, given z^I is obtained analogously as

$$p(q^I | z^I) = \Phi\left(\frac{u(q^I) - z^I}{\sigma}\right) - \Phi\left(\frac{\ell(q^I) - z^I}{\sigma}\right). \quad (4)$$

B. First Mismatched Quantization Model

The above (exact) quantization model often leads to computationally expensive numerical algorithms that require high arithmetic precision. We therefore also consider the case where the receiver algorithms are designed on the basis of *mismatched* quantization models; such models enable the development of simpler and (often) faster algorithms at the cost of a small performance loss. By “mismatched quantization model,” we mean that algorithms for channel estimation and data detection are developed on the basis of a quantized input-output relation that does not match the exact one given in (1).¹ Obviously, such a mismatch yields an error-rate performance loss.

To describe our first mismatched quantization model (which we abbreviate as “Mismatch 1”), it is convenient to assign to each quantization label q a complex value $y(q)$ whose real and imaginary parts lie within the quantization bin boundaries $\mathcal{S}(q) = [\ell(q^R), u(q^R)) \times [\ell(q^I), u(q^I))$ associated to q . Let $f(v)$ denote the probability density function of the unquantized noisy signal $v = z + n$. Following [28], we take $y(q)$ as the centroid of $\mathcal{S}(q)$ under $f(v)$, i.e.,

$$y(q) = \frac{\int_{\mathcal{S}(q)} v f(v) dv}{\int_{\mathcal{S}(q)} f(v) dv}. \quad (5)$$

Now we can write

$$y(q) = z + n + e \quad (6)$$

¹The term “mismatched” is commonly used in the information-theoretic literature to denote the scenario where the decoder operates according to a specific, possibly suboptimal, rule; see e.g., [38] and references therein.

for some quantization error e that depends on the unquantized signal $z + n$. Note that, so far, we have just provided an alternative description of (1). We now obtain a mismatched quantization model by assuming that the quantization error e is *independent* of both z and n . Furthermore, we shall approximate its distribution by a circularly symmetric complex Gaussian distribution with zero mean and variance

$$\gamma^2(q) = \frac{\int_{\mathcal{S}(q)} |v - y(q)|^2 f(v) dv}{\int_{\mathcal{S}(q)} f(v) dv}, \quad (7)$$

which depends on the quantization label q . This approximation has been suggested in [28] in the context of sparse signal recovery from noisy, quantized measurements.

In this first mismatched quantization model, the likelihood $\tilde{p}(q | z)$ takes the following form (cf. (2) for the exact model):

$$\tilde{p}(q | z) = \frac{1}{\pi(N_0 + \gamma^2(q))} \exp\left(-\frac{|z - y(q)|^2}{N_0 + \gamma^2(q)}\right). \quad (8)$$

To compute (5) and (7), one can either approximate the probability density function $f(v)$ with a uniform distribution over the set $\mathcal{S}(q)$, or—if known—use the distribution of the unquantized signal $v = z + n$.

In summary, the first mismatched quantization model is obtained by assuming (i) that the quantization error does not depend on the signal z or the noise n , and (ii) that the quantization error is Gaussian. We shall see that despite the crudeness of these approximations, the first mismatched quantization model enables computationally efficient channel-estimation and data-detection algorithms, which closely approach the error-rate performance of (more complex) algorithms developed for the exact quantization model detailed in Section II-A.

III. QUANTIZED SINGLE-INPUT SINGLE-OUTPUT OFDM

We now explain the principles of the channel-estimation and data-detection algorithms proposed in this paper. In order to avoid that the key features of our algorithms be obfuscated by the unavoidably intricate notation needed to describe a quantized massive MU-MIMO-OFDM uplink system, we first focus on a simple, quantized single-input single-output (SISO) OFDM system. The full-fledged quantized massive MU-MIMO-OFDM uplink system will be discussed in Section IV.

A. Quantized SISO-OFDM System Model

Let the set $\mathcal{O}$ contain the points of the chosen digital modulation format (e.g., QAM or PSK). Assume that the transmitter converts a W -tone frequency-domain signal $\mathbf{s} \in \mathcal{O}^W$ into the time domain using an inverse discrete Fourier transform (DFT) according to $\mathbf{v} = \mathbf{F}^H \mathbf{s}$. Here, $\mathbf{F}$ denotes the $W \times W$ DFT matrix that satisfies $\mathbf{F}^H \mathbf{F} = \mathbf{I}$. After prepending the cyclic prefix, the signal is transmitted over a frequency-selective channel. The receiver quantizes the baseband time-domain signal and discards the cyclic prefix. The resulting time-domain input-output relation is given by

$$\mathbf{q} = \mathcal{Q}(\mathbf{H}\mathbf{F}^H \mathbf{s} + \mathbf{n}). \quad (9)$$

Here, the vector $\mathbf{q} \in \mathcal{A}_c^W$ contains the quantization labels, which are obtained by applying the quantizer $\mathcal{Q}(\cdot)$ element-wise to the vector $\mathbf{H}\mathbf{F}^H \mathbf{s} + \mathbf{n}$. The $W \times W$ channel matrix

$\mathbf{H}$ is circulant and contains the samples of the channel's impulse response. The vector $\mathbf{n} \in \mathbb{C}^W$ has independent and identically distributed (i.i.d.) circularly symmetric complex Gaussian entries with variance N_0 .

Since circulant matrices are diagonalized by the Fourier transform, we can rewrite (9) as

$$\mathbf{q} = \mathcal{Q}(\mathbf{F}^H \text{diag}(\mathbf{h})\mathbf{s} + \mathbf{n}) \quad (10)$$

where $\text{diag}(\mathbf{h}) = \mathbf{F}\mathbf{H}\mathbf{F}^H$ stands for a diagonal matrix containing the frequency-domain representation $\mathbf{h}$ of the channel impulse response on its main diagonal. Unfortunately, the nonlinear nature of the quantization operator $\mathcal{Q}(\cdot)$ implies that it is not possible to convert the time-domain input-output relation (10) into a *diagonalized* frequency-domain input-output relation, simply by computing a DFT at the receiver.

B. MAP Channel Estimation

We now show that pilot-based MAP channel estimation of the channel $\mathbf{h}$ in (10) can be formulated as a convex optimization problem. Specifically, we assume a training phase consisting of a single² OFDM symbol, which we model as follows:

$$\mathbf{q} = \mathcal{Q}(\mathbf{T}\mathbf{h} + \mathbf{n}). \quad (11)$$

Here, $\mathbf{T} = \mathbf{F}^H \text{diag}(\mathbf{t})$ where the vector $\mathbf{t} \in \mathbb{C}^W$ contains the pilot symbols (which are known to the receiver). To obtain (11) from (10), we set $\mathbf{s} = \mathbf{t}$ and used that $\text{diag}(\mathbf{h})\mathbf{t} = \text{diag}(\mathbf{t})\mathbf{h}$. By assuming prior knowledge of the probability density function $f(\mathbf{h})$ of the channel $\mathbf{h}$, we can formulate the following MAP channel-estimation problem

$$\hat{\mathbf{h}}^{\text{MAP}} = \arg \max_{\mathbf{h} \in \mathbb{C}^W} p(\mathbf{q} | \mathbf{T}\tilde{\mathbf{h}})f(\tilde{\mathbf{h}}). \quad (12)$$

We now show that, under mild assumptions on the probability density function $f(\mathbf{h})$, the optimization problem (12) can be solved *exactly* via convex optimization, despite the nonlinearity introduced by the quantizer $\mathcal{Q}(\cdot)$. To this end, we first take the negative logarithm of the objective function on the right-hand side (RHS) of (12), which allows us to rewrite (12) in the following equivalent form:

$$\hat{\mathbf{h}}^{\text{MAP}} = \arg \min_{\mathbf{h} \in \mathbb{C}^W} \left\{ -\log p(\mathbf{q} | \mathbf{T}\tilde{\mathbf{h}}) - \log f(\tilde{\mathbf{h}}) \right\}. \quad (13)$$

It is now key to realize that the first term $-\log p(\mathbf{q} | \mathbf{T}\tilde{\mathbf{h}})$ on the RHS of (13) is convex in $\tilde{\mathbf{h}}$. To show this, we require the following result.

Theorem 1: The negative log-likelihood $-\log p(q | z)$ in (2) is a smooth, convex function in $z \in \mathbb{C}$.

Proof: We start by noting that

$$-\log p(q | z) = -\log p(q^R | z^R) - \log p(q^I | z^I). \quad (14)$$

To prove the theorem, we shall show that both summands on the RHS of (14) are smooth convex functions in the variables

z^R and z^I . Following the line of reasoning in [28, Sec. III-A], we note that the function

$$\Phi(u - c) - \Phi(\ell - c) = \int_{\ell - c}^{u - c} \frac{1}{\sqrt{2\pi}} \exp(-\nu^2/2) d\nu$$

which appears in the expressions for $p(q^R | z^R)$ and $p(q^I | z^I)$ in (3) and (4), respectively, is log-concave in the real variable c , because it is a convolution between two log-concave functions: the Gaussian pdf and the indicator function of the interval $[\ell, u]$. This concludes the proof of the theorem. ■

Since $\mathbf{T}\tilde{\mathbf{h}}$ is linear in $\tilde{\mathbf{h}}$, Theorem 1 implies that the negative log-likelihood $-\log p(\mathbf{q} | \mathbf{T}\tilde{\mathbf{h}}) = \sum_{w=1}^W -\log p(q_w | [\mathbf{T}\tilde{\mathbf{h}}]_w)$ is convex in $\tilde{\mathbf{h}}$. Hence, the MAP channel-estimation problem (13) is convex as long as the negative logarithm of the prior term $-\log f(\tilde{\mathbf{h}})$ is convex. This property is satisfied by any *log-concave* prior distribution $f(\mathbf{h})$ on the channel vector $\mathbf{h}$. For example, if the entries of $\mathbf{h}$ are i.i.d. circularly symmetric complex Gaussian with variance C_0 (which corresponds to i.i.d. Rayleigh fading), we have $-\log f(\tilde{\mathbf{h}}) = \|\tilde{\mathbf{h}}\|_2^2/C_0 + \log(\pi C_0)$, which is a smooth, convex function in $\tilde{\mathbf{h}}$.

If we assume that no prior knowledge on the probability distribution of $\mathbf{h}$ is available, which is equivalent to letting $C_0 \rightarrow \infty$ in our example, we transform the MAP channel-estimation problem into an ML channel-estimation problem, which remains convex. The crucial observation that MAP and ML channel estimation in the presence of the quantizer $\mathcal{Q}(\cdot)$ can be formulated as convex problems implies that efficient numerical methods for channel estimation exist (see Section V for the details).

C. MAP and MMSE Data Detection

To perform data detection from quantized measurements, we first rewrite the input-output relation (10) as $\mathbf{q} = \mathcal{Q}(\tilde{\mathbf{H}}\mathbf{s} + \mathbf{n})$, where $\tilde{\mathbf{H}} = \mathbf{F}^H \text{diag}(\tilde{\mathbf{h}})$ and $\tilde{\mathbf{h}}$ stands for either the MAP or the ML channel estimate obtained solving the convex optimization problem described in Section III-B.

1) *Hard-Output MAP Detection:* By assuming that prior information on the probability mass function $p(\mathbf{s})$ of the transmit data vectors $\mathbf{s} \in \mathcal{O}^W$ is available, one could—in principle—minimize the probability that the vector $\mathbf{s}$ is detected erroneously by solving the following MAP data detection problem:³

$$\hat{\mathbf{s}}^{\text{MAP}} = \arg \max_{\mathbf{s} \in \mathcal{O}^W} p(\mathbf{q} | \tilde{\mathbf{H}}\hat{\mathbf{s}})p(\hat{\mathbf{s}}). \quad (15)$$

In absence of prior information, one can transform (15) into an ML data detection problem by assuming equally-likely transmit vectors, i.e., that $p(\mathbf{s}) = |\mathcal{O}|^{-W}$, $\forall \mathbf{s} \in \mathcal{O}^W$. Since the matrix $\tilde{\mathbf{H}}$ is, in general, not diagonal, this optimization problem is at least as hard as the MAP (or ML) data-detection problem in conventional, unquantized MIMO systems with W spatial streams. Specifically, solving (15) exactly entails an exhaustive search whose complexity grows exponentially in the number of subcarriers W (see [39], [40] for more details). Since W

²The number of degrees of freedom of frequency-selective channels is typically much smaller than the number of OFDM tones. This property ensures that a single training symbol is sufficient to acquire accurate channel estimates, even in the presence of coarse quantization. Nevertheless, our algorithms can easily be extended to systems that use more than one training symbol.

³To be precise, the data-detection problem in (15) is the MAP problem only under the assumption of perfect channel state information. Since we estimate the channel in our setup, this detector is actually mismatched. This observation holds for all detectors presented in the remainder of this section.

can be of the order of hundred or even thousand in modern OFDM systems, solving (15) directly is not feasible and one has to resort to approximate algorithms.

2) *Hard-Output and Soft-Output MMSE Detection*: To arrive at a low-complexity data-detection method, we relax the finite-alphabet constraint $\tilde{\mathbf{s}} \in \mathcal{O}^W$ so that the resulting data-detection problem becomes convex and can be solved efficiently. We make the common approximation that the entries of $\mathbf{s}$ are i.i.d. circularly symmetric complex Gaussian random variables with variance E_s (where E_s stands for the average energy per transmitted symbol); this allows us to formulate a quantized version of the well-known MMSE equalizer:

$$\hat{\mathbf{s}} = \arg \min_{\tilde{\mathbf{s}} \in \mathbb{C}^W} \left\{ -\log p(\mathbf{q} | \hat{\mathbf{H}}\tilde{\mathbf{s}}) + \frac{\|\tilde{\mathbf{s}}\|_2^2}{E_s} \right\}. \quad (16)$$

As a consequence of Theorem 1, this optimization problem is convex and can be solved efficiently. Its solution can either be mapped element-wise onto the closest constellation point as $\hat{s}_w^{\text{MMSE}} = \arg \min_{a \in \mathcal{O}} |a - \hat{s}_w|$, $w = 1, \dots, W$, to obtain a hard-output MMSE estimate $\hat{\mathbf{s}}^{\text{MMSE}}$, or used to compute soft-information (for a given mapping between bits and constellation symbols) in the form of max-log log-likelihood ratio (LLR) values [41], [42]:

$$L_{w,b} = \rho_w \min_{a \in \mathcal{O}_b^{(0)}} |\hat{s}_w - a|^2 - \rho_w \min_{a \in \mathcal{O}_b^{(1)}} |\hat{s}_w - a|^2. \quad (17)$$

Here, ρ_w denotes the post-equalization signal-to-interference-and-noise ratio (SINR) on the w th subcarrier, and $\mathcal{O}_b^{(0)}$ and $\mathcal{O}_b^{(1)}$ refer to the subsets of $\mathcal{O}$ consisting of all constellation symbols whose b th bit is 0 and 1, respectively.⁴ We note that solving the convex problem (16) yields, in general, no information on the SINR ρ_w , which makes an exact computation of (17) impossible. Throughout the paper, we approximate (17) by setting $\rho_w = 1$, unless stated otherwise. We note that this approximation was shown in [43], [44] to result in near-optimal performance for the case of massive MIMO data detection with max-log channel decoding.

D. Channel Estimation and Data Detection with the First Mismatched Quantization Model

We now turn our attention to the first mismatched quantization model introduced in Section II-B and derive the optimal channel estimator and data detector for the case in which the receiver is designed on the basis of this mismatched model. With this mismatched quantization model, the input-output relation (10) is replaced by

$$\mathbf{y}(\mathbf{q}) = \mathbf{F}^H \text{diag}(\mathbf{h})\mathbf{s} + \mathbf{n} + \mathbf{e}. \quad (18)$$

In (18), which is the OFDM generalization of the scalar input-output relation (6), the noise vector follows a $\mathcal{CN}(0, N_0 \mathbf{I}_W)$ distribution, and the entries $\{e_w\}_{w=1}^W$ of the quantization noise vector $\mathbf{e}$ are distributed according to $\mathcal{CN}(0, \gamma^2(q_w))$. The variance $\gamma^2(q_w)$, where q_w denotes the quantization label associated to the received signal on the w subcarrier, is defined

⁴The subsets $\mathcal{O}_b^{(0)}$ and $\mathcal{O}_b^{(1)}$ depend on the mapping from constellation symbols to bits (e.g., a Gray mapping).

similarly to (7). Proceeding as in the case of unquantized SISO-OFDM systems, we apply a DFT to the time-domain vector $\mathbf{y}(\mathbf{q})$ and obtain the following frequency-domain input-output relation:

$$\hat{\mathbf{y}} = \text{diag}(\mathbf{h})\mathbf{s} + \hat{\mathbf{n}} + \hat{\mathbf{e}}. \quad (19)$$

Here, $\hat{\mathbf{y}} = \mathbf{F} \mathbf{y}(\mathbf{q})$, $\hat{\mathbf{n}} \sim \mathcal{CN}(0, N_0 \mathbf{I}_W)$, and $\hat{\mathbf{e}} \sim \mathcal{CN}(0, \mathbf{K}_e)$ where $\mathbf{K}_e = \mathbf{F} \mathbf{T} \mathbf{F}^H$ is the covariance matrix of the quantization noise $\hat{\mathbf{e}}$, expressed in the frequency domain. Finally, $\mathbf{T} = \text{diag}(\gamma^2(q_1), \dots, \gamma^2(q_W))$. We note that the entries of the vector $\hat{\mathbf{e}}$ are, in general, correlated. As we shall discuss next, this has a negative impact on the complexity of MAP channel-estimation and data-detection algorithms.

1) *MAP Channel Estimation*: Analogously to (13), MAP channel estimation for the mismatched quantization model (19) can be formulated as follows:

$$\hat{\mathbf{h}}^{\text{MAP}} = \arg \min_{\tilde{\mathbf{h}} \in \mathbb{C}^W} \left\{ \|\mathbf{K}^{-1/2}(\hat{\mathbf{y}} - \text{diag}(\tilde{\mathbf{h}})\tilde{\mathbf{h}})\|_2^2 - \log p(\tilde{\mathbf{h}}) \right\}. \quad (20)$$

Here, $\mathbf{K} = N_0 \mathbf{I}_W + \mathbf{K}_e$ stands for the covariance matrix of the Gaussian additive noise plus the frequency-domain quantization noise, and $\tilde{\mathbf{h}} \in \mathbb{C}^W$ denotes the vector containing the pilot symbols used for channel estimation. If $p(\tilde{\mathbf{h}})$ is log-concave, the problem (20) is also convex and can be solved efficiently using numerical methods. In the case where $p(\tilde{\mathbf{h}})$ is Gaussian, (20) reduces to a W -dimensional LS problem.

2) *MAP and MMSE Data Detection*: Similarly to the MAP channel-estimation problem (20), we can formulate the MAP data-detection problem as follows:

$$\hat{\mathbf{s}}^{\text{MAP}} = \arg \min_{\tilde{\mathbf{s}} \in \mathcal{O}^W} \left\{ \|\mathbf{K}^{-1/2}(\hat{\mathbf{y}} - \text{diag}(\hat{\mathbf{h}})\tilde{\mathbf{s}})\|_2^2 - \log p(\tilde{\mathbf{s}}) \right\}. \quad (21)$$

Here, $\hat{\mathbf{h}}$ is the MAP channel estimate, obtained by solving (20). Since, as already mentioned, $\mathbf{K}$ is in general not diagonal, solving (21) entails prohibitive complexity (see the discussion in Section III-C).

To reduce the computational complexity of the MAP data-detection problem, we relax the alphabet constraint in (21) by proceeding as in Section III-C, and obtain the following mismatched version of the quantized MMSE equalizer in (16):

$$\hat{\mathbf{s}}^{\text{MMSE}} = \arg \min_{\tilde{\mathbf{s}} \in \mathbb{C}^W} \left\{ \|\mathbf{K}^{-1/2}(\hat{\mathbf{y}} - \text{diag}(\hat{\mathbf{h}})\tilde{\mathbf{s}})\|_2^2 + \frac{\|\tilde{\mathbf{s}}\|_2^2}{E_s} \right\}. \quad (22)$$

The optimization problem in (22) is a W -dimensional LS problem and its solution can be used to perform hard-output or soft-output detection, as described in Section III-C.

E. Second Mismatched Quantization Model

We now provide a second mismatched quantization model (abbreviated as ‘‘Mismatch 2’’), which yields a significant complexity reduction for channel estimation and data detection (cf. Section III-D). The key idea is to replace the noise vector $\mathbf{e}$ in (18), whose entries are independent but not identically distributed, with a vector $\tilde{\mathbf{e}}$ with i.i.d. entries.

Specifically, we let the entries of $\tilde{\mathbf{e}}$ be $\mathcal{CN}(0, \bar{\gamma})$ -distributed and set $\bar{\gamma} = \frac{1}{W} \sum_{w=1}^W \gamma^2(q_w)$, which is the average variance of the entries of the vector $\mathbf{e}$. This approach results in the following mismatched quantization model:

$$\tilde{\mathbf{y}}(\mathbf{q}) = \mathbf{F}^H \text{diag}(\mathbf{h})\mathbf{s} + \mathbf{n} + \tilde{\mathbf{e}}. \quad (23)$$

The advantage of this second mismatched quantization model is as follows: when we apply the DFT operator to $\tilde{\mathbf{y}}(\mathbf{q})$, the total noise remains uncorrelated because the correlation matrix of $\mathbf{n} + \tilde{\mathbf{e}}$ is a scaled identity (cf. (19)). This means that the MAP channel-estimation and data-detection problems decouple into W independent one-dimensional problems that can be solved at low computational complexity.⁵ In particular, one can deploy standard OFDM signal-processing techniques that operate tone-wise in the frequency domain. An immediate benefit of this mismatched quantization model is that we can easily extract the post-equalization SINR values ρ_w , which are required for LLR computation in (17).

IV. QUANTIZED MASSIVE MU-MIMO-OFDM UPLINK

We are now ready to present our quantized massive MU-MIMO-OFDM uplink scheme in full detail. Although the combination of MU-MIMO and OFDM renders the notation rather involved, the main ideas of our approach follow those discussed in Section III for the quantized SISO-OFDM case. We start by introducing the system model. We then derive the MAP channel-estimation and MMSE data-detection algorithms for the exact quantization model. Finally, we propose low-complexity algorithm variants based on the mismatched quantization models introduced in Sections II-B and III-E.

A. Quantized MIMO-OFDM System Model

We consider a coded MU-MIMO-OFDM uplink system that employs spatial multiplexing. The system model is depicted in Fig. 1 and consists of U independent single-antenna users⁶ communicating with a BS having $B \geq U$ receive antennas. We consider a frame-based OFDM transmission, where each OFDM symbol consists of W tones, similar to what is used in IEEE 802.11n [45]. We furthermore assume that communication is effected in two phases: (i) a training phase consisting of U OFDM symbols (which enables the training of every channel coefficient at least once) followed by (ii) a data transmission phase consisting of D OFDM symbols. The disjoint sets Ω_{data} , Ω_{pilot} , and Ω_{guard} contain the indices associated to data tones, pilot tones, and guard (or zero) tones, respectively. Clearly, $|\Omega_{\text{data}}| + |\Omega_{\text{pilot}}| + |\Omega_{\text{guard}}| = W$. For the sake of simplicity, we assume a sufficiently long cyclic prefix, perfect synchronization, and ignore hardware impairments such as phase noise, I/Q imbalance, and amplifier nonlinearities (see, e.g., [5], [46], [47] for more details on hardware impairments in OFDM systems).

⁵This holds provided that the priors are product distributions, i.e., that $f(\mathbf{h}) = \prod_{w=1}^W f(h_w)$ and that $p(\mathbf{s}) = \prod_{w=1}^W p(s_w)$.

⁶Our framework enables also the use of multi-antenna user equipments. In this case, forward error correction and channel decoding can be performed over the multiple antennas.

1) *User Terminal (Transmit Side)*: During the training phase, the data and the pilot tones contain QPSK training symbols that are known to the BS; the guard tones remain unused. During the data-transmission phase, each user encodes its information bits using a forward error correction (FEC) channel code and maps the coded bits onto constellation symbols $s_w \in \mathcal{O}$ for $w \in \Omega_{\text{data}}$ to be transmitted over the data-carrying tones. The pilot tones contain BPSK symbols; in a real system, they can be used to compensate for residual phase noise or frequency offset [47]. The guard tones are still left unused. For a given OFDM symbol, the frequency-domain transmit vector of each user terminal u , which is compactly described by the vector $\mathbf{s}_u \in \mathcal{O}^W$, is converted into the time domain using a W -point inverse DFT according to $\mathbf{v}_u = \mathbf{F}^H \mathbf{s}_u$. After prepending a cyclic prefix of length P , the time-domain signals $\mathbf{v}_u$ of all users $u = 1, \dots, U$ are transmitted simultaneously and in the same frequency band.

2) *Base Station (Receive Side)*: Each of the B antennas at the BS receive a noisy mixture of time-domain user data. Specifically, let $\mathbf{y}_b \in \mathbb{C}^W$ denote the unquantized time-domain received vector at antenna b , after removal of the cyclic prefix. The time-domain received vector at each antenna $b = 1, \dots, B$ is quantized according to $\mathbf{q}_b = \mathcal{Q}(\mathbf{y}_b)$, and then passed to the channel-estimation/data-detection unit. During the training phase, the channel-estimation unit generates estimates $\hat{\mathbf{H}}_w \in \mathbb{C}^{B \times U}$, $w = 1, \dots, W$, for all OFDM tones. During the data-transmission phase, the data-detection unit generates LLR values for each coded bit of each user, on the basis of the channel estimates and of the quantized received vectors. The resulting (approximate) soft-information is used to perform decoding on a per-user basis.

3) *Channel model*: To simplify the notation, we make frequent use of the mapping [48]

$$\mathcal{T}\{\mathbf{X}_w\}_{w=1}^W = \{\mathbf{X}'_b\}_{b=1}^B \quad (24)$$

between matrices $\mathbf{X}_w \in \mathbb{C}^{B \times T}$, $w = 1, \dots, W$, in *per-frequency* orientation and matrices $\mathbf{X}'_b \in \mathbb{C}^{W \times T}$, $b = 1, \dots, B$, in *per-BS-antenna* orientation. Here, $T = U + D$ is the total number of OFDM symbols (during the training and the data-transmission phases). Formally, the mapping (24) is defined as $[\mathbf{X}_w]_{b,\ell} = [\mathbf{X}'_b]_{w,\ell}$, where $w = 1, \dots, W$, $b = 1, \dots, B$, and $\ell = 1, \dots, T$ are the frequency (or time) index, the BS-antenna index, and OFDM-symbol index, respectively. To illustrate the effect of the mapping $\mathcal{T}$ in (24), consider the transmission of a *single* OFDM symbol ($T = 1$) consisting of W subcarriers to a BS equipped with B antennas. We can either represent the overall received signal as W vectors $\mathbf{x}_w \in \mathcal{O}^B$, each containing B symbols, or, equivalently, as B vectors $\mathbf{x}'_b \in \mathcal{O}^W$, each containing W symbols. The operator $\mathcal{T}$ allows us to conveniently switch from one representation to the other. A pictorial representation of this mapping is shown in Fig. 2.

Using the mapping $\mathcal{T}$ in (24) together with the single-antenna OFDM input-output relation (10), we can write the input-output

Fig. 2. Mapping $\mathcal{T}$ between per-frequency orientation (left) and per-BS-antenna orientation (right). The data cube (frequency tones $\times$ BS antennas $\times$ OFDM symbols) remains the same but data is associated differently to the matrices $\mathbf{X}$ and $\mathbf{X}'$.

relation of the quantized MU-MIMO-OFDM system as follows:

$$\begin{aligned} \{\mathbf{Z}_b\}_{b=1}^B &= \mathcal{T}\{\mathbf{H}_w \mathbf{S}_w\}_{w=1}^W \\ \mathbf{Q}_b &= \mathcal{Q}(\mathbf{F}^H \mathbf{Z}_b + \mathbf{N}_b), \quad b = 1, \dots, B. \end{aligned} \quad (25)$$

Here, the inputs of the channel are the W frequency-domain matrices $\mathbf{S}_w \in \mathbb{C}^{U \times T}$, each containing the data symbols, the pilots, and the guard tones for T OFDM symbols and U users. The outputs of the channel are the B quantized time-domain matrices $\mathbf{Q}_b \in \mathcal{A}_c^{W \times T}$ (one for each BS antenna), containing a W -dimensional time-domain signal for each of the T OFDM symbols. The frequency-domain channel matrices $\mathbf{H}_w \in \mathbb{C}^{B \times U}$ in (25), which are assumed to remain constant over the T OFDM symbols, can be obtained from the time-domain matrices $\mathbf{H}_t \in \mathbb{C}^{B \times U}$, $t = 1, \dots, W$, using the Fourier transform $\{\mathbf{H}_w\}_{w=1}^W = \mathcal{F}\{\mathbf{H}_t\}_{t=1}^W$, defined as follows [49]:

$$\mathbf{H}_w = \frac{1}{\sqrt{W}} \sum_{t=1}^W \mathbf{H}_t \exp\left(-i2\pi \frac{(t-1)(w-1-W/2)}{W}\right). \quad (26)$$

Here, the time-domain matrices $\{\mathbf{H}_t\}_{t=1}^W$, of dimension $B \times U$, contain the impulse responses from each user to each BS antenna. In what follows, we assume that the impulse responses are supported on at most $L < W$ taps, i.e., $\mathbf{H}_t = \mathbf{0}_{B \times U}$ for $t = L+1, \dots, W$; we furthermore assume that the cyclic-prefix length P satisfies $L \leq P$.

B. MAP Channel Estimation

We are now ready to formulate MAP channel estimation as a convex optimization problem. As outlined in Section III, we seek estimates $\hat{\mathbf{H}}_w$ of the channel matrices $\mathbf{H}_w$, $w = 1, \dots, W$, given the quantized measurements $\mathbf{Q}_b$, $b = 1, \dots, B$, and the known, orthogonal pilot matrices $\mathbf{T}_w \in \mathcal{O}^{U \times U}$, where $w \in \Omega_{\text{pilot}} \cup \Omega_{\text{data}}$. Specifically, our objective is to minimize the negative log-likelihood of the quantized data, given prior information on the channel matrices to be estimated. Since it is difficult in practice to acquire accurate prior channel statistics,⁷ we assume that the only information known to the receiver

about the channel is that its impulse response does not exceed the cyclic prefix length P . This results in the following convex optimization problem:

$$\text{(Q-CHE)} \quad \begin{cases} \text{minimize} & -\sum_{b=1}^B \log p(\mathbf{Q}_b | \mathbf{F}^H \mathbf{Z}_b) \\ \text{subject to} & \{\mathbf{Z}_b\}_{b=1}^B = \mathcal{T}\{\tilde{\mathbf{H}}_w \mathbf{T}_w\}_{w=1}^W \\ & \{\tilde{\mathbf{H}}_t\}_{t=1}^W = \mathcal{F}^{-1}\{\tilde{\mathbf{H}}_w\}_{w=1}^W \\ & \tilde{\mathbf{H}}_t = \mathbf{0}_{B \times U}, \quad t = P+1, \dots, W. \end{cases}$$

Here, Q-CHE stands for quantized channel estimation. The solution of this convex optimization problem is the set of estimates $\{\hat{\mathbf{H}}_w\}_{w=1}^W$ of the MU-MIMO-OFDM channel in the frequency domain. The ingredients of the optimization problem (Q-CHE) are as follows. The objective function, i.e., the negative log-likelihood of the quantization labels given the noiseless received signal is the MIMO generalization of the objective function given in (13) for the SISO-OFDM case.⁸ The first constraint simply transforms a per-frequency representation into a per-BS-antenna representation as illustrated in Fig. 2; this simplifies the form of the objective function. The second constraint transforms the channel matrices from the frequency domain to the time domain. The third constraint captures that the channel's impulse response has at most P nonzero taps. Similarly to the SISO-OFDM case, in (Q-CHE) the MAP channel estimation must be performed jointly over all W OFDM tones. Here, however, we additionally include the constraint that the channel's delay spread does not exceed P taps. This enables us to acquire accurate channel estimates using only few OFDM pilot symbols. Because of the convexity, (Q-CHE) can be solved exactly and efficiently using first-order methods (see Section V for the details).

C. MMSE Data Detection

Next, we shall extend the MMSE detector (16) to the MU-MIMO case. The key difference with respect to the SISO-OFDM case is that the resulting convex-optimization problem involves not only all subcarriers, but also all antennas. During the data-transmission phase, the following convex-optimization problem needs to be solved for each OFDM symbol:

$$\text{(Q-DET)} \quad \begin{cases} \text{minimize} & -\sum_{b=1}^B \log p(\mathbf{q}_b | \mathbf{F}^H \mathbf{z}_b) \\ & + \sum_{w \in \Omega_{\text{data}}} E_s^{-1} \|\mathbf{s}_w\|_2^2 \\ \text{subject to} & \{\mathbf{z}_b\}_{b=1}^B = \mathcal{T}\{\hat{\mathbf{H}}_w \mathbf{s}_w\}_{w=1}^W \\ & \mathbf{s}_w = \mathbf{t}_w, \quad w \in \Omega_{\text{pilot}}. \end{cases}$$

Here, Q-DET stands for quantized detection. The solution of the optimization problem (Q-DET) are the estimates $\hat{\mathbf{s}}_w \in \Omega_{\text{data}}$ of the transmitted data symbols; these estimates can then be used to compute hard or soft estimates as described in Section III-C. The ingredients of this optimization problem are as follows.

⁸Note that, differently from (13), the term corresponding to the prior distribution of the channel is missing from the objective function of (Q-CHE). This is because *a priori* information about the channel distribution is assumed unavailable at the receiver.

⁷Especially in MU cellular scenarios, where users may experience vastly different path losses.

The objective function is the MIMO generalization of (16) for the SISO-OFDM case. Specifically, we optimize jointly over the data vectors $\tilde{\mathbf{s}}_w$, $w \in \Omega_{\text{data}}$, transmitted by the U users, and we account for the fact that the base station is equipped with B receive antennas. The first constraint in (Q-DET) transforms a per-frequency representation into a per-BS-antenna representation as illustrated in Fig. 2. The second constraint captures that the symbols $\mathbf{t}_w$, $w \in \Omega_{\text{pilot}}$, transmitted on the pilot tones are known to the receiver (see Section IV-A). Similarly to the MAP channel estimation problem (Q-CHE), the problem (Q-DET) can be solved exactly and efficiently using first-order methods (see Section V for the details).

D. Mismatched Quantization Models

1) *First Mismatched Quantization Model:* We now briefly discuss the MAP channel-estimation and data-detection problems for the first mismatched quantization model introduced in Section II-B. In essence, we only replace the conditional probabilities $p(\cdot|\cdot)$ in (Q-CHE) and (Q-DET) with the mismatched conditional probability $\tilde{p}(\cdot|\cdot)$ in (8). While the dimensionality of the resulting problems remains unchanged, the algorithms that use the mismatched quantizer exhibit improved numerical stability and (often) faster convergence rates at the cost of a small degradation in performance (see Section V for the details).

2) *Second Mismatched Quantization Model:* For the second mismatched quantization model introduced in Section III-E, the MAP channel estimation and the MMSE data detection problems simplify drastically. Specifically, after a DFT operation at the receiver, the channel estimation problem decouples into $U \times B$ independent subproblems, each one involving the estimation of a W -dimensional channel vector. Specifically, each channel estimation subproblem is of the form

$$(\text{MQ-CHE}) \quad \begin{cases} \underset{\mathbf{h} \in \mathbb{C}^W}{\text{minimize}} & \sum_{w \in \{\Omega_{\text{pilot}} \cup \Omega_{\text{data}}\}} |\hat{y}_w - \tilde{h}_w|^2 \\ \text{subject to} & [\mathbf{F}^H \mathbf{h}]_t = 0, t = P + 1, \dots, W. \end{cases}$$

Here, MQ-CHE stands for mismatched-quantization channel estimation. In the above problem, $\hat{y}_w$, $w \in \{\Omega_{\text{pilot}} \cup \Omega_{\text{data}}\}$, are the frequency-domain entries of the matrices $\hat{\mathbf{H}}_w$, $w \in \{\Omega_{\text{pilot}} \cup \Omega_{\text{pilot}}\}$ that correspond to the u th user and the b th BS antenna, which are estimated using the orthonormal pilot matrices $\mathbf{T}_w$ (see Section IV-B). Although (MQ-CHE) has a closed-form solution,⁹ we will use a first-order method to reduce both computational complexity and memory requirements (see Section V). In practice, one can also obtain approximate solutions to (MQ-CHE) using more efficient methods (see e.g., [50]).

With the second mismatched quantization model, MMSE data detection can be performed *independently* per subcarrier, as it is the case for conventional unquantized MIMO-OFDM systems. Hence, MMSE data detection requires one to solve $|\Omega_{\text{data}}|$ independent $B \times U$ MMSE data detection problems,

⁹The problem (MQ-CHE) can be solved by computing a projection matrix of dimension $W \times W$, which involves inverting a matrix of the same dimension. Such a direct approach would not only result in excessive computational complexity but also consume large amounts of memory.

which can be done using efficient algorithms. Furthermore, the SINR values ρ_w that are necessary for soft-output data detection (17) can be obtained in this case at virtually no cost; see [41], [42], [51] for more details on unquantized (i.e., infinite-precision) soft-output MMSE data detection.

V. EFFICIENT NUMERICAL METHODS

The problems (Q-CHE), (Q-DET), and (MQ-CHE) require us to solve large-dimensional convex optimization problems. We now summarize a class of algorithms that can be used to solve these problems at low complexity. We refer the interested reader to [52] for more details on these algorithms.

A. Forward-Backward Splitting

In many fields, including (but not limited to) image processing, machine learning, and spectral clustering, one is interested in solving high-dimensional convex optimization problems of the following form [53]:

$$\underset{\mathbf{x} \in \mathbb{C}^N}{\text{minimize}} \quad h(\mathbf{A}\mathbf{x}) + g(\mathbf{x}). \quad (27)$$

Here, $\mathbf{A} \in \mathbb{C}^{M \times N}$, the function h is convex and differentiable, and g is convex, but not necessarily smooth or bounded. The generality of the function g prevents the use of simple gradient-descent algorithms. For many problems of interest, however, the so-called *proximal operator* [53]

$$\text{prox}_g(\mathbf{z}, \tau) = \arg \min_{\mathbf{x}} \left\{ \tau g(\mathbf{x}) + \frac{1}{2} \|\mathbf{x} - \mathbf{z}\|_2^2 \right\} \quad (28)$$

can be computed efficiently. In this case, one can deploy a computationally efficient first-order method to solve (27), generally referred to as forward-backward splitting (FBS). Specifically, FBS performs the following iteration for $k = 1, 2, \dots$, until convergence [53]:

$$\mathbf{x}^{(k+1)} = \text{prox}_g(\mathbf{x}^{(k)} - \tau^{(k)} \mathbf{A}^H \nabla h(\mathbf{A}\mathbf{x}^{(k)}), \tau^{(k)}). \quad (29)$$

FBS requires the computation of the proximal operator (28), of the gradient ∇h of the smooth function, the computation of matrix-vector products involving $\mathbf{A}$ and its adjoint $\mathbf{A}^H$, and multiplications by the step size $\tau^{(k)} > 0$. To reduce the complexity of FBS, one should (i) exploit fast transforms instead of matrix-vector products (if possible) and (ii) select an appropriate step-size $\tau^{(k)}$ in every iteration k . While a number of fast FBS implementations have been described in the literature, the numerical results reported in Section VI use the fast adaptive shrinkage/thresholding algorithm (FASTA) proposed in [52].

B. Implementation Details

We shall illustrate next that the optimization problems (Q-CHE), (Q-DET), and (MQ-CHE) described in Section IV are of the form given in (27). Indeed, in all these problems the objective function is smooth and convex. We can therefore identify these objective functions with h in (27). The (non-smooth) function g in (27) can be associated with the affine constraints listed in (Q-CHE), (Q-DET), and (MQ-CHE). To this end, let $\mathcal{C}$ be the feasible set defined by the affine constraints

of the optimization problems described in Section IV-A, where $\mathcal{X}_{\mathcal{C}}(x)$ is the characteristic function of $\mathcal{C}$, which we define to be zero if $x \in \mathcal{C}$ and infinity otherwise. We shall set $g(x) = \mathcal{X}_{\mathcal{C}}(x)$, which is a convex (but not smooth or bounded) function. Next we detail how to compute the proximal operator (28) and the gradient of h , which is needed in (29).

1) *Gradients*: For the exact quantization model introduced in Section II-A, the gradient of the negative log-likelihood function $-\log p(\mathbf{q} | \mathbf{a})$, for the case of real-valued $\mathbf{a}$, is given by [37]

$$[\nabla h(\mathbf{a})]_i = \frac{\exp\left(-\frac{(u(q_i)-a_i)^2}{2\sigma^2}\right) - \exp\left(-\frac{(\ell(q_i)-a_i)^2}{2\sigma^2}\right)}{\sqrt{2\pi\sigma^2}\left[\Phi\left(\frac{u(q_i)-a_i}{\sigma}\right) - \Phi\left(\frac{\ell(q_i)-a_i}{\sigma}\right)\right]}. \quad (30)$$

The complex-valued extension of (30), which is required in both (Q-CHE) and (Q-DET) can be readily obtained by computing the gradient separately for real and imaginary parts [37]. For the mismatched quantization model introduced in Section II-B, the (complex-valued) gradient of $-\log \tilde{p}(\mathbf{q} | \mathbf{a})$ is given by

$$[\nabla h(\mathbf{a})]_i = \frac{y(q_i) - a_i}{N_0/2 + \gamma^2(q_i)}. \quad (31)$$

A comparison of (30) and (31) reveals that gradient computation is easier and less computationally intensive for the mismatched quantization model.

2) *Proximal Operators*: As we shall show next, the proximal operators associated with the convex-optimization problems (Q-CHE), (Q-DET), and (MQ-CHE) can be evaluated efficiently. Specifically, the function g in (Q-CHE) corresponds to the affine constraint $\tilde{\mathbf{H}}_t = \mathbf{0}_{B \times U}$, $t = P+1, \dots, W$. This implies that the proximal operator (28) involves transforming the channel matrices from frequency domain to time domain, setting to zero all matrices with index $t = P+1, \dots, W$, and transforming the resulting matrices back to the frequency domain. By using the fast Fourier transform (FFT), the proximal operator can be evaluated efficiently. Similarly, (MQ-CHE) involves the affine constraint $[\mathbf{F}^H \tilde{\mathbf{h}}]_t = 0$, $t = P+1, \dots, W$. The computation of the proximal operator requires transforming the frequency-domain vector $\tilde{\mathbf{h}}$ into time domain, setting to zero all entries with index $t = P+1, \dots, W$, and transforming the resulting vector back to frequency domain. Finally, (Q-DET) involves the affine constraints $\mathbf{s}_w = \mathbf{t}_w$, $w \in \Omega_{\text{pilot}}$. The proximal operator simply requires one to set $\mathbf{s}_w$ equal to $\mathbf{t}_w$ for all $w \in \Omega_{\text{pilot}}$.

VI. SIMULATION RESULTS

We will now present numerical simulation results for the proposed quantized massive MU-MIMO-OFDM uplink system. Due to space constraints, we shall focus on a selected set of system parameters.¹⁰

A. Simulation Parameters

All simulation results are based on an IEEE 802.11n-like system [45] with 40 MHz bandwidth (similarly to the OFDM

system parameters considered in [48]). The number of OFDM subcarriers is $W = 128$; guard tones, pilot tones, and data tones are as defined in [45]. Each user encodes its information bits using a (rather weak) rate-5/6 convolutional code of constraint length 7 with random interleaving. Decoding at the base-station relies on a max-log soft-input Viterbi decoder. Channel training is performed as described in Section IV-A1. In the data-transmission phase, the transmitted symbols belong to a 16-QAM constellation. The channel matrices are generated according to (26), where the entries of the non-zero time-domain matrices $\mathbf{H}_t$, $t = 1, \dots, L$, are assumed i.i.d. circularly complex Gaussian with unit variance. We assume that the number of channel taps is $L = 4$; the cyclic prefix length is $P = 16$. The pilot matrices $\mathbf{T}_w$ are generated from a $U \times U$ orthonormal Hadamard transform matrix, where the signs of the columns and rows are randomly altered.¹¹ The quantization alphabet $\mathcal{A}_c$ and the quantization bin boundaries $\{b_q\}$ are set using the Lloyd-Max algorithm [54], [55]. For a given set of system parameters, the algorithm uses an estimate of the probability density function of the received signal, which is acquired offline.¹² We use the FASTA solver [56] with tolerance set to 10^{-6} and maximum number of iterations set to 500. The packet error rate (PER), which is averaged over all users, is obtained via Monte-Carlo simulations with 1000 packet transmissions for each simulated point.¹³

To characterize the system performance for a given number of quantization bits Q_b (see Section II-A), we compute the so-called “SNR operating point” proposed in [57], [58], which is the minimum average receive SNR required to achieve 1% PER. We consider three receiver architectures. The first one, referred to as “Quantizer,” performs MAP channel estimation and MMSE data detection as described in Section IV-B and Section IV-C, respectively. The second and third one, referred to as “Mismatch 1” and “Mismatch 2,” respectively, perform their computations as described in Section IV-D1 and Section IV-D2, respectively.

B. Results and Discussion

Fig. 3 illustrates the trade-off between the SNR operating point and the number of quantization bits Q_b for different BS-antenna-to-user ratios. As a reference, we show the performance of an infinite-precision single-input multiple-output (SIMO) system both for the case of perfect CSI at the receiver (CSIR), and for the case of channel estimation (CHEST). To obtain these curves, we replaced the negative log-likelihood of the quantized channel output with the one of the channel output before quantization in the objective functions of the iterative algorithms described Section IV.

¹¹We observed that the choice of the pilot matrices has a noticeable impact on the error-rate performance in systems with coarse quantization. In particular, orthogonalizing the users in time during pilot transmission results in higher channel-estimation errors. A detailed investigation of this phenomenon, which is similar to the one described in [8] in the context of modulation design for 1-bit massive MIMO systems, is part of ongoing work.

¹²In practice, one can precompute the quantization bin boundaries and store them in a look-up table.

¹³The trade-off curves are obtained by performing 1000 Monte-Carlo trials for each SNR point. The irregular behavior of some of the curves is a consequence of the chosen number of trials and the SNR resolution.

¹⁰Our simulation framework will be available for download from GitHub (https://github.com/quantizedmassivemimo/mu_mimo_ofdm/). The purpose is to enable interested readers to perform their own simulations with different system parameters and also to test alternative algorithms.

Fig. 3. Trade-off between the SNR operating point and the number of quantization bits Q_b for 8 users and $\{16, 32, 64, 128\}$ base station antennas. Only a few bits are sufficient to approach the performance of the SMO lower bound with infinite-precision ADCs. Furthermore, for systems with a large number of BS antennas (e.g., $B = 128$), no sophisticated signal processing is required in the presence of low-precision ADCs, i.e., the “Mismatch 2” receiver achieves near-optimal performance.

Fig. 3(a) compares the performance/quantization trade-off for the “Quantizer,” “Mismatch 1,” and “Mismatch 2” receiver architectures for a system with 16 BS antennas and 8 users (corresponding to a BS-antenna-to-user ratio of 2). We see that “Quantizer” outperforms both “Mismatch 1” and “Mismatch 2” when $3 \leq Q_b \leq 5$. When $Q_b < 3$, the system is not able to reach a PER of 1%. Interestingly, when $Q_b \geq 6$, “Mismatch 2” outperforms both “Quantizer” and “Mismatch 1”; this is because in “Mismatch 2” it is possible to compute the post-equalization SINR values ρ_w required for soft-output computation in (17), whereas in the other receiver architectures, we simply set $\rho_w = 1$.

Fig. 3(b) shows that increasing the number of BS antennas to 32 (corresponding to a BS-antenna-to-user ratio of 4), enables “Quantizer” and “Mismatch 2” to achieve a similar SNR operating point when $Q_b \geq 5$. This shows that the LLR approximation in Section III-C is accurate for such a BS-antenna-to-user ratio, confirming the observations made in [43], [44]. In this regime, one should use the “Mismatch 2” receiver architecture (which effectively ignores the presence of a quantizer), as it does not entail a complexity increase compared to conventional algorithms for infinite-precision MU-MIMO-OFDM systems. For lower values of Q_b , however, “Quantizer” significantly outperforms “Mismatch 2” (we do

not show “Mismatch 1” as it exhibits similar performance to “Quantizer”). Also note that with “Quantizer,” we are now able to achieve 1% PER for only two quantization bits (i.e., $Q_b = 2$).

Figs. 3(c) and 3(d) show that by further increasing the BS-antenna-to-user ratio to 8 and 16, respectively, “Mismatch 2” delivers the same performance as “Quantizer.” In addition, we see that for both receiver architectures, $Q_b = 4$ yields an SNR gap to the infinite-precision ADC case of only 0.25 dB. Furthermore, the SNR gap to the SMO system (with channel estimation, but infinite precision) is less than 1 dB. All these observations have far reaching consequences for quantized massive MU-MIMO-OFDM systems with BS-antenna-to-user ratio exceeding 8. In particular, we see that the simple and low-complexity “Mismatch 2” receiver architecture achieves near-optimal performance with only 4 bit precision at virtually no complexity increase in baseband processing compared to infinite-precision MU-MIMO-OFDM systems.

VII. CONCLUSIONS

We have studied the performance of a quantized massive MU-MIMO-OFDM uplink operating over frequency-selective, wideband channels. We have developed new channel-estimation and data-detection algorithms, and studied the associated performance/quantization trade-offs. Our simulation results show that

the use of four-bit ADCs is sufficient to achieve near-optimal performance in massive MU-MIMO-OFDM systems having a BS-antenna-to-user ratio of eight or higher. Remarkably, this comes at no additional costs in terms of baseband-processing complexity, compared to infinite-precision systems. Our results imply that coarse quantization enables low-cost and low-power massive MU-MIMO-OFDM system implementations that achieve near-optimal performance. In addition, reducing the precision from 10 bit to 4 bit yields $2.5\times$ lower ADC output rates, which is particularly relevant for massive MIMO systems where the antenna elements and baseband processing unit are spatially separated.

There are many avenues for future work. The presence of a quantizer renders synchronization and noise-variance estimation a difficult task; the development of corresponding methods is of paramount importance to enable the practical deployment of quantized massive MU-MIMO-OFDM systems. Extracting post-equalization SINR values for LLR computation directly via convex optimization may yield significant performance improvements; this is an open research problem. The application of the algorithms presented in this paper to single-carrier frequency-division multiple access (SC-FDMA) systems with coarse quantization, which is relevant to the 3GPP LTE uplink [59], is left for future work. One final open issue is whether the proposed quantized MU-MIMO-OFDM architecture is compatible with systems that use digital filters *after* the ADCs, as typically done in broadband wireless receivers to reduce costs and increase spectrum flexibility.

REFERENCES

- [1] B. Murmann, "ADC Performance Survey 1997-2014." [Online]. Available: <http://web.stanford.edu/~murmann/adcsurvey.html>
- [2] T. L. Marzetta, "Non-cooperative cellular wireless with unlimited numbers of base station antennas," *IEEE Trans. Wireless Comm.*, vol. 9, no. 11, pp. 3590–3600, Nov. 2010.
- [3] E. G. Larsson, F. Tufvesson, O. Edfors, and T. L. Marzetta, "Massive MIMO for next generation wireless systems," *IEEE Commun. Mag.*, pp. 186–195, Feb. 2014.
- [4] E. Björnson, J. Hoydis, M. Kountouris, and M. Debbah, "Massive MIMO systems with non-ideal hardware: Energy efficiency, estimation, and capacity limits," *IEEE Trans. Inf. Theory*, vol. 11, no. 60, pp. 7112–7139, Nov. 2014.
- [5] C. Studer, M. Wenk, and A. Burg, "MIMO transmission with residual transmit-RF impairments," in *Int. ITG Workshop on Smart Antennas (WSA)*, Bremen, Germany, Feb. 2010, pp. 189–196.
- [6] U. Gustavsson, C. Sánchez-Pérez, T. Eriksson, F. Athley, G. Durisi, P. Landin, K. Hausmair, C. Fager, and L. Svensson, "On the impact of hardware impairments on massive MIMO," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, Austin, TX, Dec. 2014, pp. 294–300.
- [7] C. Risi, D. Persson, and E. G. Larsson, "Massive MIMO with 1-bit ADC," Apr. 2014. [Online]. Available: <http://arxiv.org/abs/1404.7736>
- [8] S. Jacobsson, G. Durisi, M. Coldrey, U. Gustavsson, and C. Studer, "One-bit massive MIMO: channel estimation and high-order modulations," in *Proc. IEEE Int. Conf. Commun. (ICC)*, London, U.K., Jun. 2015.
- [9] M. A. Davenport, Y. Plan, E. van den Berg, and M. Wotter, "1-Bit matrix completion," *Information and Inference*, vol. 3, no. 189–223, July 2014.
- [10] K. Knudson, R. Saab, and R. Ward, "One-bit compressive sensing with norm estimation," Apr. 2014. [Online]. Available: <http://arxiv.org/abs/1404.6853>
- [11] N. Liang and W. Zhang, "Mixed-ADC massive MIMO," *IEEE J. Sel. Areas Commun.*, 2016, to appear.
- [12] J. Singh, O. Dabeer, and U. Madhow, "On the limits of communication with low-precision analog-to-digital conversion at the receiver," *IEEE Trans. Commun.*, vol. 57, no. 12, pp. 3629–3639, Dec. 2009.
- [13] A. Mezghani and J. Nossek, "On ultra-wideband MIMO systems with 1-bit quantized outputs: Performance analysis and input optimization," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Nice, France, Jun. 2007, pp. 1286–1289.
- [14] —, "Analysis of Rayleigh-fading channels with 1-bit quantized output," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Toronto, ON, Canada, Jul. 2008, pp. 260–264.
- [15] —, "Analysis of 1-bit output noncoherent fading channels in the low SNR regime," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Seoul, Korea, Jun. 2009, pp. 1080–1084.
- [16] 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.
- [17] T. Koch and A. Lapidoth, "At low SNR, asymmetric quantizers are better," *IEEE Trans. Inf. Theory*, vol. 59, no. 9, pp. 5421–5445, Sep. 2013.
- [18] S. Verdú, "Spectral efficiency in the wideband regime," *IEEE Trans. Inf. Theory*, vol. 48, no. 6, pp. 1319–1343, Jun. 2002.
- [19] R. Narasimha, M. Lu, N. Shanbhag, and A. Singer, "BER-optimal analog-to-digital converters for communication links," *IEEE Trans. Signal Process.*, vol. 60, no. 7, pp. 3683–3691, Jul. 2012.
- [20] G. Zeitler, A. Singer, and G. Kramer, "Low-precision A/D conversion for maximum information rate in channels with memory," *IEEE Trans. Commun.*, vol. 60, no. 9, pp. 2511–2521, Sep. 2012.
- [21] S. Shamai (Shitz), "Information rates by oversampling the sign of a bandlimited process," *IEEE Trans. Inf. Theory*, vol. 40, no. 4, pp. 1230–1236, Jun. 1994.
- [22] T. Koch and A. Lapidoth, "Increased capacity per unit-cost by oversampling," in *IEEE 26th Convention of Electrical and Electronics Engineers in Israel (IEEEI)*, Eliat, Israel, Nov. 2010, pp. 684 –688.
- [23] W. Zhang, "A general framework for transmission with transceiver distortion and some applications," *IEEE Trans. Commun.*, vol. 60, no. 2, pp. 384–399, Feb. 2012.
- [24] S. Krone and G. Fettweis, "Capacity of communications channels with 1-bit quantization and oversampling at the receiver," in *IEEE Sarnoff Symp. (SARNOFF)*, Newark, NJ, May 2012.
- [25] O. Dabeer and A. Karnik, "Signal parameter estimation using 1-bit dithered quantization," *IEEE Trans. Inf. Theory*, vol. 52, no. 12, pp. 5389–5405, Dec. 2006.
- [26] A. Mezghani, F. Antreich, and J. A. Nossek, "Multiple parameter estimation with quantized channel output," in *ITG Smart Antenna Workshop*, Bremen, Germany, Feb. 2010, pp. 143–150.
- [27] G. Zeitler, G. Kramer, and A. Singer, "Bayesian parameter estimation using single-bit dithered quantization," *IEEE Trans. Signal Process.*, vol. 60, no. 6, pp. 2713–2726, Jun. 2012.
- [28] A. Zymnis, S. Boyd, and E. Candès, "Compressed sensing with quantized measurements," *IEEE Signal Process. Lett.*, vol. 17, no. 2, pp. 149–152, Feb. 2010.
- [29] P. T. Boufounos and R. G. Baraniuk, "1-Bit compressive sensing," in *Proc. Annual Conf. Inf. Sciences Syst. (CISS)*, Mar. 2008, pp. 16–21.
- [30] Y. Plan and R. Vershynin, "One-bit compressed sensing by linear programming," *Comm. Pure Appl. Math.*, vol. 66, no. 8, pp. 1275–1297, Feb. 2013.
- [31] T. Lok and V.-W. Wei, "Channel estimation with quantized observation," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Cambridge, MA, Aug. 1998, p. 333.
- [32] M. T. Ivrlac and J. A. Nossek, "On MIMO channel estimation with single-bit signal quantization," in *ITG Smart Antenna Workshop*, Vienna, Austria, Feb. 2007.
- [33] S. Wang, Y. Li, and J. Wang, "Multiuser detection in massive spatial modulation MIMO with low-resolution ADCs," *IEEE Trans. Wireless Commun.*, vol. 14, no. 4, pp. 2156–2168, Apr. 2014.
- [34] A. Mezghani and J. A. Nossek, "Belief propagation based MIMO detection operating on quantized channel output," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Austin, TX, Jun. 2010, pp. 2113–2117.
- [35] A. Mezghani, M. Rouatbi, and J. Nossek, "An iterative receiver for quantized MIMO systems," in *IEEE Mediterranean Electrotechnical Conf. (MELECOM)*, Hammamet, Tunisia, Mar. 2012, pp. 1049–1052.
- [36] J. Choi, J. Mo, and R. W. Heath Jr, "Near maximum-likelihood detector and channel estimator for uplink multiuser massive MIMO systems with one-bit ADCs," *IEEE Trans. Commun.*, 2016, to appear.
- [37] D. Bellasi, L. Bettini, C. Benkeser, T. Burger, Q. Huang, and C. Studer, "VLSI design of a monolithic compressive-sensing wideband analog-to-information converter," *IEEE J. Emerg. Sel. Topics Circuits Syst.*, vol. 3, no. 4, pp. 552–565, Mar. 2013.
- [38] A. Ganti, A. Lapidoth, and I. Telatar, "Mismatched decoding revisited: general alphabets, channels with memory, and the wide-band limit," *IEEE Trans. Inf. Theory*, vol. 46, no. 7, pp. 2315–2328, Nov. 2000.

- [39] D. Seethaler, J. Jaldén, C. Studer, and H. Bölcskei, "On the complexity distribution of sphere decoding," *IEEE Trans. Inf. Theory*, vol. 57, no. 9, pp. 5754–5768, Sept. 2011.
- [40] J. Jaldén and B. Ottersten, "On the complexity of sphere decoding in digital communications," *IEEE Trans. Signal Process.*, vol. 53, no. 4, pp. 1474–1484, Apr. 2005.
- [41] 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.
- [42] D. Seethaler, G. Matz, and F. Hlawatsch, "An efficient MMSE-based demodulator for MIMO bit-interleaved coded modulation," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, vol. 4, Dallas, TX, Nov. 2004, pp. 2455–2459.
- [43] M. Wu, B. Yin, G. Wang, C. Dick, J. Cavallaro, and C. Studer, "Large-scale MIMO detection for 3GPP LTE: Algorithm and FPGA implementation," *IEEE J. Sel. Topics Signal Process.*, vol. 8, no. 5, pp. 916–929, Oct. 2014.
- [44] B. Yin, M. Wu, J. R. Cavallaro, and C. Studer, "Conjugate gradient-based soft-output detection and precoding in massive MIMO systems," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, Austin, TX, Dec. 2014, pp. 3696–3701.
- [45] *IEEE Draft Standard; Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications; Amendment 4: Enhancements for Higher Throughput*, P802.11n/D3.0, Sep. 2007.
- [46] T. C. Schenk, P. F. Smulders, and E. R. Fledderus, "Performance of MIMO OFDM systems in fading channels with additive TX and RX impairments," in *Proc. IEEE BENELUX/DSP Valley Signal Process. Symp.*, Antwerp, Belgium, Apr. 2005, pp. 41–44.
- [47] T. C. Schenk, *RF imperfections in high-rate wireless systems: impact and digital compensation*. Dordrecht, The Netherlands: Springer, 2008.
- [48] C. Studer and E. G. Larsson, "PAR-aware large-scale multi-user MIMO-OFDM downlink," *IEEE J. Sel. Areas Comm.*, vol. 31, no. 2, pp. 303–313, Feb. 2013.
- [49] H. Bölcskei, D. Gesbert, and A. Paulraj, "On the capacity of OFDM-based spatial multiplexing systems," *IEEE Trans. Commun.*, vol. 50, no. 2, pp. 225–234, Feb. 2002.
- [50] S. Haene, A. Burg, N. Felber, and W. Fichtner, "OFDM channel estimation algorithm and ASIC implementation," in *4th European Conference on Circuits and Systems for Communications (ECCSC 2008)*, 2008, pp. 270–275.
- [51] A. Paulraj, R. Nabar, and D. Gore, *Introduction to Space-Time Wireless Communications*. Cambridge Univ. Press, 2003.
- [52] T. Goldstein, C. Studer, and R. G. Baraniuk, "A field guide to forward-backward splitting with a FASTA implementation," Nov. 2014. [Online]. Available: <http://arxiv.org/abs/1411.3406>
- [53] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. Imag. Sci.*, vol. 2, no. 1, pp. 183–202, Jan. 2009.
- [54] J. Max, "Quantizing for optimal distortion," *IRE Trans. Info. Theory*, vol. 6, no. 1, pp. 7–12, Mar. 1960.
- [55] S. P. Lloyd, "Least squares quantization in pcm," *IEEE Trans. Inf. Theory*, vol. 28, no. 2, pp. 129–137, Mar. 1982.
- [56] T. Goldstein, C. Studer, and R. G. Baraniuk, "FASTA: A generalized implementation of forward-backward splitting," Aug. 2015. [Online]. Available: <http://arxiv.org/abs/1501.04979>
- [57] C. Studer and H. Bölcskei, "Soft-input soft-output single tree-search sphere decoding," *IEEE Trans. Inf. Theory*, vol. 56, no. 10, pp. 4827–4842, Oct. 2010.
- [58] C. Studer, A. Burg, and H. Bölcskei, "Soft-output sphere decoding: Algorithms and VLSI implementation," *IEEE J. Sel. Areas Commun.*, vol. 26, no. 2, pp. 290–300, Feb. 2008.
- [59] *3rd Generation Partnership Project; Technical Specification Group Radio Access Network; Evolved Universal Terrestrial Radio Access (E-UTRA); Multiplexing and channel coding (Release 9)*, 3GPP Organizational Partners TS 36.212, Rev. 8.3.0, May 2008.