

Neural-Network Optimized 1-bit Precoding for Massive MU-MIMO

Conference Paper**Author(s):**

Balatsoukas-Stimming, Alexios; [Castañeda Fernández, Oscar](#) ; Jacobsson, Sven; Durisi, Giuseppe; [Studer, Christoph](#) 

Publication date:

2019-07

Permanent link:

<https://doi.org/10.3929/ethz-b-000448922>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

<https://doi.org/10.1109/spawc.2019.8815519>

Neural-Network Optimized 1-bit Precoding for Massive MU-MIMO

Alexios Balatsoukas-Stimming^{*†}, Oscar Castañeda[†], Sven Jacobsson^{‡§}, Giuseppe Durisi[§], and Christoph Studer[†]

^{*}*Telecommunications Circuits Laboratory, École Polytechnique Fédérale de Lausanne, Switzerland; email: alexios.balatsoukas@epfl.ch*

[†]*School of Electrical and Computer Engineering, Cornell University, Ithaca, NY; email: oc66@cornell.edu, studer@cornell.edu*

[‡]*Ericsson Research, Gothenburg, Sweden; email: sven.jacobsson@ericsson.com*

[§]*Chalmers University of Technology, Gothenburg, Sweden; email: durisi@chalmers.se*

Abstract—Base station (BS) architectures for massive multi-user (MU) multiple-input multiple-output (MIMO) wireless systems are equipped with hundreds of antennas to serve tens of users on the same time-frequency channel. The immense number of BS antennas incurs high system costs, power, and interconnect bandwidth. To circumvent these obstacles, sophisticated MU precoding algorithms that enable the use of 1-bit DACs have been proposed. Many of these precoders feature parameters that are, traditionally, tuned manually to optimize their performance. We propose to use deep-learning tools to automatically tune such 1-bit precoders. Specifically, we optimize the biConvex 1-bit Precoding (C2PO) algorithm using neural networks. Compared to the original C2PO algorithm, our neural-network optimized (NNO-)C2PO achieves the same error-rate performance at $2\times$ lower complexity. Moreover, by training NNO-C2PO for different channel models, we show that 1-bit precoding can be made robust to vastly changing propagation conditions.

I. INTRODUCTION

Massive multiuser (MU) multiple-input multiple-output (MIMO) will be a core technology in fifth-generation (5G) wireless systems [1]–[3]. By equipping the base station (BS) with hundreds of antennas, massive MU-MIMO permits communication with tens of user equipments (UEs) in the same time-frequency resource using fine-grained beamforming. However, such systems with hundreds of BS antennas, in which each antenna is connected to a high-precision radio frequency (RF) chain, would result in prohibitively high system costs, power consumption, and interconnect bandwidth between the baseband processing unit and the remote radios. Thus, novel BS architectures, in combination with efficient baseband-processing algorithms, are necessary to reduce the cost, power, and interconnect bandwidth of massive MU-MIMO systems, while preserving high reliability and spectral efficiency.

A. Massive MU-MIMO with Low-Precision BS Architectures

The cost, power, and interconnect bandwidth of massive MU-MIMO BSs can be greatly reduced by using low-precision digital-to-analog converters (DACs). Such an architecture tolerates decreased linearity and noise requirements for the RF circuitry, enabling the use of low-cost and power-efficient analog circuits. However, advanced baseband algorithms are needed to attain high spectral efficiency.

The work of ABS was supported by the Swiss NSF under projects #184772 and #182621. The work of OC and 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.

In the downlink (i.e., when the BS transmits to the UEs), linear precoders (e.g., maximal-ratio transmission (MRT) or zero-forcing (ZF) precoding) followed by quantization result in low complexity algorithms that suppress MU interference using low-precision DACs, at the cost of a significant bit error-rate (BER) performance degradation [4]–[6]. In contrast, sophisticated nonlinear precoders [7]–[16] achieve superior BER performance (especially in the extreme case where a pair of 1-bit DACs per RF chain is used), at the cost of a higher complexity. Existing nonlinear precoding algorithms require manual parameter tuning to optimize the BER performance [13]. Parameter tuning is time-consuming and may yield suboptimal results in realistic systems with vastly changing propagation conditions. Furthermore, most existing precoders have only been evaluated in idealistic Rayleigh-fading conditions, and it remains unclear whether they also provide a satisfactory performance in realistic channels.

B. Contributions

In this paper, we propose to use neural-network optimization tools to tune the algorithm parameters of the nonlinear biConvex 1-bit Precoding (C2PO) algorithm proposed in [13]. Specifically, we unfold the C2PO iterations and use backpropagation for parameter tuning. We show that the resulting neural-network optimized C2PO (NNO-C2PO) algorithm delivers significantly improved performance when compared to C2PO for the same number of iterations, for a range of channel models. More interestingly, we show that NNO-C2PO can use the same set of learned per-iteration parameter values on vastly different channel models at a negligible performance loss compared to the case where the algorithm is separately trained for each channel model. In contrast, the original C2PO algorithm is unable to deal with different channel models, even when its single set of parameters is trained in the same way as NNO-C2PO. Overall, our results illustrate that, by using machine learning techniques, existing iterative nonlinear precoders can significantly improve their BER performance, while also becoming robust to changing channel conditions.

C. Relevant Prior Art

Several recent works have explored the application of machine-learning techniques to different problems that arise in the context of massive MU-MIMO. The majority of results have focused on data detection [17]–[23], while some have

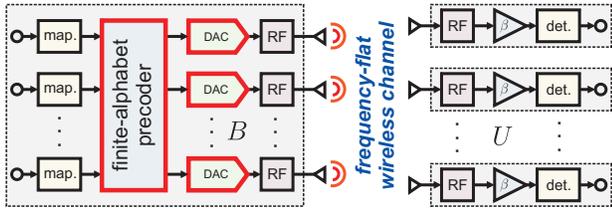


Fig. 1. Simplified view of a massive MU-MIMO downlink system with low-resolution DACs: A BS with B antennas uses a finite-alphabet precoder to transmit data to U single-antenna UEs through a wireless channel.

tackled channel estimation [18], [24], and only [25] has considered the case of (non-quantized) downlink precoding. However, most of these results [17]–[20], [25] consider deep learning models that are highly parametrizable (e.g., containing multiple dense layers), and thus, both difficult to train and too complex to implement in practice. The methods proposed in [21], [22] are the most closely related to our work, since they also unfold iterative algorithms and learn a small set of algorithm parameters from training data. We note that some of the parameters that are trained in [21], [22] are introduced arbitrarily into the algorithm, while the C2PO algorithm that we use already contains tunable parameters by design. To our knowledge, the unfolding technique has not been examined in the context of 1-bit precoding in massive MU-MIMO systems before. Moreover, the robustness of the learned communications systems (particularly, 1-bit massive MU-MIMO systems) to changing channel conditions is a practically-relevant issue that has not been addressed to date.

II. SYSTEM MODEL AND 1-BIT PRECODING

A. System Model

We focus on the downlink of a single-cell, narrowband massive MU-MIMO system as shown in Figure 1, where a BS equipped with B antennas serves $U \ll B$ UEs with one antenna each. This system is modeled as $\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{n}$, where $\mathbf{y} = [y_1, \dots, y_U]^T \in \mathbb{C}^U$ comprises the signals received by each UE, $\mathbf{H} \in \mathbb{C}^{U \times B}$ is the downlink MIMO channel matrix (which is assumed to be perfectly known at the BS), $\mathbf{n} \in \mathbb{C}^U$ models i.i.d. circularly-symmetric complex Gaussian noise with variance N_0 per complex entry, and $\mathbf{x} \in \mathcal{X}^B$ is the precoded vector with entries chosen from the discrete set \mathcal{X} determined by the DAC outputs. We assume that the precoded vector satisfies an instantaneous power constraint $\|\mathbf{x}\|_2^2 \leq P$.

B. 1-bit Precoding

For the case of 1-bit precoding, where both the in-phase and quadrature paths are fed by 1-bit DACs, we have $\mathcal{X} = \{+v + jv, +v - jv, -v + jv, -v - jv\}$, where $v \in \mathbb{R}_+$ is set to satisfy the power constraint. The objective of the precoder at the BS is to take a symbol vector $\mathbf{s} = [s_1, \dots, s_U]^T \in \mathcal{O}^U$ and to map it to a precoded vector $\mathbf{x} \in \mathcal{X}^B$ that is transmitted through the BS antennas. The vector \mathbf{x} is chosen so that, after transmission through the wireless channel, the received signal y_u at the u th UE is proportional to the transmitted symbol s_u . We assume that each UE is able to rescale the received signal

y_u by a common scalar $\beta \in \mathbb{C}$ to estimate the symbol s_u , i.e., the UEs compute $\hat{s}_u = \beta y_u$. For mean-squared error (MSE)-optimal precoders, the optimal precoding problem (OPP) can then be formulated as [13]:

$$(\text{OPP}) \quad \{\hat{\mathbf{x}}, \hat{\beta}\} = \arg \min_{\mathbf{x} \in \mathcal{X}^B, \beta \in \mathbb{C}} \|\mathbf{s} - \beta \mathbf{H}\mathbf{x}\|_2^2 + |\beta|^2 U N_0. \quad (1)$$

As shown in [8], the users can accurately estimate $\hat{\beta}$. For 1-bit massive MU-MIMO systems, several algorithms have been proposed to approximately solve (OPP) using convex [7], [8] and nonconvex [13] relaxations.

C. C2PO: Biconvex 1-bit Precoding

In this work, we will follow the C2PO algorithm put forward in [13], as it leads to a simple, yet efficient, algorithm that has been implemented in hardware. To arrive at low-complexity approximate solutions to (OPP), C2PO assumes that the system operates in the high-SNR regime, so that $N_0 \rightarrow 0$. Then, with an additional approximation [13, Eq. (2)], (OPP) is re-written:

$$(\text{OPP}^*) \quad \{\hat{\mathbf{x}}, \hat{\alpha}\} = \arg \min_{\mathbf{x} \in \mathcal{X}^B, \alpha \in \mathbb{C}} \|\alpha \mathbf{s} - \mathbf{H}\mathbf{x}\|_2^2. \quad (2)$$

The optimal value for α , can be computed in closed form as $\hat{\alpha} = \mathbf{s}^H \mathbf{H}\mathbf{x} / \|\mathbf{s}\|_2^2$. Substituting this result into (OPP*) leads to the following equivalent form:

$$(\text{OPP}^{**}) \quad \hat{\mathbf{x}} = \arg \min_{\mathbf{x} \in \mathcal{X}^B} \frac{1}{2} \|\mathbf{A}\mathbf{x}\|_2^2, \quad (3)$$

where $\mathbf{A} = (\mathbf{I}_U - \mathbf{s}\mathbf{s}^H / \|\mathbf{s}\|_2^2) \mathbf{H}$. The finite-alphabet constraint $\mathbf{x} \in \mathcal{X}^B$ is then relaxed to the convex hull \mathcal{B} of the points in \mathcal{X} . To avoid the trivial all-zeros solution, a concave regularizer $-\frac{\delta}{2} \|\mathbf{x}\|_2^2$ with $\delta > 0$ is added to the objective:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x} \in \mathcal{B}^B} \frac{1}{2} \|\mathbf{A}\mathbf{x}\|_2^2 - \frac{\delta}{2} \|\mathbf{x}\|_2^2. \quad (4)$$

C2PO uses forward-backward splitting (FBS) [26], [27] to solve (4), an efficient numerical method for convex optimization problems. As the objective in (4) is nonconvex, FBS is not guaranteed to converge to an optimal solution. However, C2PO has been shown to yield excellent performance in practice [13]. After applying FBS to (4), the C2PO algorithm is obtained (see [13] for a detailed explanation). To present C2PO, we need to introduce some auxiliary notation. Let $\rho = \frac{1}{1-\tau\delta}$ with $\delta \in \mathbb{R}$, $\xi = \sqrt{\frac{P}{2B}}$, and the proximal-operator:

$$\text{prox}_g(\mathbf{z}; \rho, \xi) = \text{clip}(\rho \Re\{\mathbf{z}\}, \xi) + j \text{clip}(\rho \Im\{\mathbf{z}\}, \xi), \quad (5)$$

where the clipping function $\text{clip}(\mathbf{z}, \gamma)$ applies the operation $\min(\max(z_i, -\gamma), \gamma)$ to each element of the vector \mathbf{z} .

Algorithm 1 (C2PO). Initialize $\mathbf{x}^{(0)} = \mathbf{H}^H \mathbf{s}$. Fix $\tau^{(t)}$ and $\rho^{(t)}$. For every iteration $t = 1, 2, \dots, t_{\max}$ compute:

$$\mathbf{z}^{(t)} = \mathbf{x}^{(t-1)} - \tau^{(t)} \mathbf{A}^H \mathbf{A} \mathbf{x}^{(t-1)} \quad (6)$$

$$\mathbf{x}^{(t)} = \text{prox}_g(\mathbf{z}^{(t)}; \rho^{(t)}, \xi), \quad (7)$$

Finally, quantize the output $\mathbf{x}^{(t_{\max})}$ to the set \mathcal{X}^B .

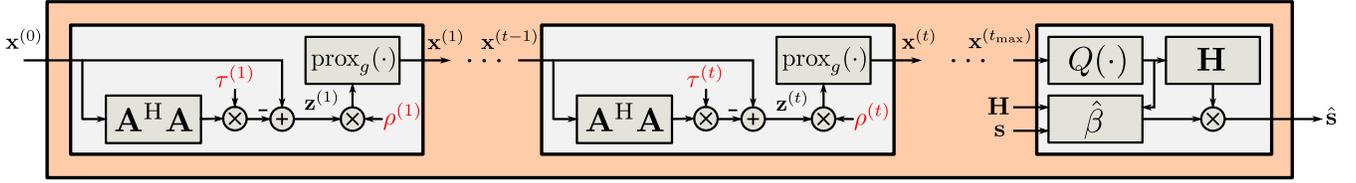


Fig. 2. Computation graph of the iteration-unfolded version of NNO-C2PO. All trainable parameters are highlighted in red color.

We note that in the original reference [13], constant values were used for $\tau = \tau^{(t)}$ and $\rho = \rho^{(t)}$ at every iteration because manually tuning $2t_{\max}$ distinct parameters seemed impractical.

III. NEURAL-NETWORK OPTIMIZED C2PO

Analogous to [21], [22], the idea behind our approach is to unfold the iterations of C2PO and to learn distinct parameters $\tau^{(t)}$ and $\rho^{(t)}$ for each iteration $t = 1, \dots, t_{\max}$. These parameters are learned offline using deep-learning methods, and then used in the C2PO algorithm as described in Algorithm 1. We refer to the overall procedure as neural-network optimized C2PO (NNO-C2PO).

A. Unfolding the Algorithm

As shown in Figure 2, the unfolded algorithm forms a computation graph, which can be described using standard deep learning frameworks (e.g., Keras or PyTorch). If all operations in the computation graph have well-defined gradients, then the gradients of the cost function with respect to $\tau^{(t)}$ and $\rho^{(t)}$ can be calculated efficiently (and easily, due to automatic differentiation) using backpropagation. These gradients can then be used to learn the values of $\tau^{(t)}$ and $\rho^{(t)}$.

The high-level architecture of the unfolded NNO-C2PO is shown in Figure 2. The algorithm takes \mathbf{H} , \mathbf{s} , $\mathbf{x}^{(0)}$, and $\mathbf{A}^H \mathbf{A}$ as inputs and produces $\hat{\mathbf{s}}$ as an output. We note that $\mathbf{x}^{(0)}$ and \mathbf{A} are functions of \mathbf{s} and \mathbf{H} , but they are pre-computed and given as additional inputs for convenience. The input to the first sub-block is $\mathbf{x}^{(0)} = \mathbf{H}^H \mathbf{s}$. Each successive sub-block takes $\mathbf{x}^{(t-1)}$ as input and produces $\mathbf{x}^{(t)}$ by means of (6) and (7). The final sub-block quantizes $\mathbf{x}^{(t_{\max})}$ to the finite alphabet \mathcal{X}^B as:

$$\mathbf{x}_Q^{(t_{\max})} = Q(\mathbf{x}^{(t_{\max})}; \xi), \quad (8)$$

where $Q(\cdot)$ is a 1-bit quantization function that is applied element-wise to $\mathbf{x}^{(t_{\max})}$ and is defined as:

$$Q(x; \xi) = \xi \text{sign}(\Re\{x\}) + j \cdot \xi \text{sign}(\Im\{x\}). \quad (9)$$

Moreover, the final block also computes the scalar value $\hat{\beta} = \|\mathbf{s}\|_2^2 / \mathbf{s}^H \mathbf{H} \mathbf{x}_Q^{(t_{\max})}$, and emulates the (noiseless, since by assumption $N_0 \rightarrow 0$) transmission over the channel as:

$$\hat{\mathbf{s}} = \hat{\beta} \cdot \mathbf{H} \mathbf{x}_Q^{(t_{\max})}. \quad (10)$$

B. Implementation Details

All operations involved in (6) are linear and thus have well-defined gradients. However, care needs to be taken for the operations in (7) as well as in the last stage in Figure 2. More

specifically, in (7), the gradient of the function $\text{clip}(z, \gamma)$ that is applied element-wise to $\mathbf{z}^{(t)}$ has a discontinuity for $z = \gamma$ and $z = -\gamma$, but since the gradients are evaluated numerically, this is not an issue in practice. Moreover, the gradient of $\text{clip}(z, \gamma)$ with respect to z is equal to 0 for $z < -\gamma$ and $z > \gamma$. However, as $\text{clip}(z, \gamma)$ is applied to B elements of $\mathbf{z}^{(t)}$ concurrently and over several training samples, the probability that the (averaged) gradient of $\text{clip}(z^{(t)}, \gamma)$ with respect to $\rho^{(t)}$ is equal to 0 is very low.

The quantization function $Q(\cdot)$ used in the last stage in Figure 2 is not differentiable. There are several ways to include quantization functions into deep-learning frameworks, such as using a soft sign function [17] or a $\tanh(\cdot)$ function [22]. Here, we follow the approach used by binarized neural networks (BNNs) [28], which applies $Q(\cdot)$ only during the forward propagation, and replaces $Q(x; \xi)$ with a clipping function $\text{clip}(x; \xi)$ (as a straight-through estimator) during backpropagation. All remaining operations in the last stage of Figure 2 have well-defined gradients and can be differentiated automatically.

We note that the updates in (6) and (7) are performed in the complex domain. Most deep learning tools are unable to deal with complex numbers. However, all complex operations in (6) and (7) can be easily recasted as real-valued operations. For example, the operation $\mathbf{x} = \mathbf{H} \mathbf{s}$ can be equivalently performed in the real domain by setting:

$$\mathbf{H}^{\mathbb{R}} = \begin{bmatrix} \Re\{\mathbf{H}\} & -\Im\{\mathbf{H}\} \\ \Im\{\mathbf{H}\} & \Re\{\mathbf{H}\} \end{bmatrix} \quad \text{and} \quad \mathbf{s}^{\mathbb{R}} = \begin{bmatrix} \Re\{\mathbf{s}\} \\ \Im\{\mathbf{s}\} \end{bmatrix}, \quad (11)$$

computing $\mathbf{x}^{\mathbb{R}} = \mathbf{H}^{\mathbb{R}} \mathbf{s}^{\mathbb{R}}$ and, finally, setting:

$$x_b = x_b^{\mathbb{R}} + j x_{b+B}^{\mathbb{R}}, \quad b = 1, \dots, B. \quad (12)$$

C. Training

When unfolding the iterations of an iterative algorithm, the corresponding computation graph can become very deep. As such, the well-known problem of vanishing gradients may make training difficult. This problem was identified in [22] in the context of massive MU-MIMO and an incremental training procedure was used to enable effective training. However, since we unfold only a small number of iterations in our experiments and we only learn a small set of parameters, we did not encounter this problem. Hence, performing a one-shot training of the entire network suffices. The training set consists of (the real-valued equivalents of) K training samples $\{\mathbf{s}_{(k)}, \mathbf{H}_{(k)}, \mathbf{x}_{(k)}^{(0)}, \mathbf{A}_{(k)}^H \mathbf{A}_{(k)}\}$, where $k = 1, \dots, K$. The desired output is $\mathbf{s}_{(k)}$ and, due to (1), the cost function

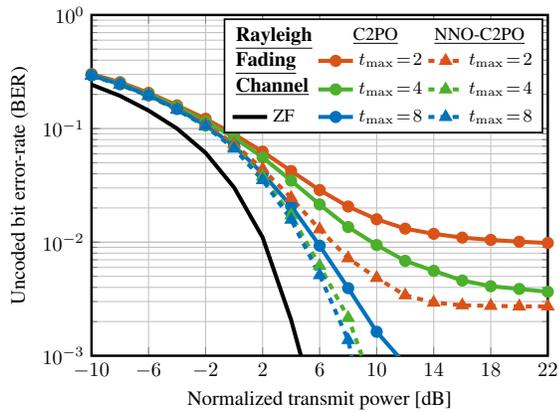


Fig. 3. Uncoded bit error-rate (BER) performance of C2PO and NNO-C2PO, using different number of iterations t_{\max} , for a massive MU-MIMO system with $B = 128$ BS antennas and $U = 8$ UEs, operating with 16-QAM in a Rayleigh-fading channel. Zero-forcing (ZF) precoding is included as reference.

to be minimized is the MSE between $\mathbf{s}_{(k)}$ and $\hat{\mathbf{s}}_{(k)}$:

$$C = \frac{1}{K} \sum_{k=1}^K \left(\|\mathbf{s}_{(k)} - \hat{\mathbf{s}}_{(k)}\|_2^2 \right). \quad (13)$$

Training the network essentially means choosing values for $\tau^{(t)}$ and $\rho^{(t)}$ that minimize (13). The parameters are initialized to $\tau^{(t)} = 2^{-8}$ and $\rho^{(t)} = 1.25$, which have been shown to work well in practice [13].

IV. RESULTS

We now compare the performance of the proposed NNO-C2PO algorithm with the C2PO algorithm. For C2PO, we learn its parameters using the same procedure as for NNO-C2PO, but we constrain the weights so that $\tau^{(t)} = \tau$ and $\rho^{(t)} = \rho$ for all iterations $t = 1, \dots, t_{\max}$. We provide BER results for Rayleigh-fading channels, as well as for more realistic line-of-sight (LoS) and non-line-of-sight (NLoS) channels generated with the QuaDRiGa channel model [29] using transmission at 2GHz with a 6.8kHz bandwidth in the “Berlin UMA” scenario. For all experiments, we consider a massive MU-MIMO system with $B = 128$ BS antennas, $U = 8$ UEs, and 16-QAM. As the number of trainable parameters is small, we generate a small dataset of $K = 500$ training samples for each channel model and we perform full batch training over $100t_{\max}$ epochs. The BER performance evaluation is carried out using $K_{\text{eval}} = 1000$ evaluation samples that are distinct from the training samples. The computation graph of Figure 2 is described using Keras [30] with a TensorFlow backend [31]. We use the Adam optimizer [32] with a learning rate of $\lambda = 10^{-4}$.

A. Rayleigh-Fading Channel

The BER performance of NNO-C2PO is compared to that of C2PO under Rayleigh-fading and for different values of t_{\max} in Figure 3, where we also show the infinite-precision ZF solution for reference. We observe that NNO-C2PO with $t_{\max} = 2$ and $t_{\max} = 4$ slightly outperforms C2PO with $t_{\max} = 4$ and $t_{\max} = 8$, respectively. This means that NNO-C2PO requires approximately 50% fewer iterations than C2PO

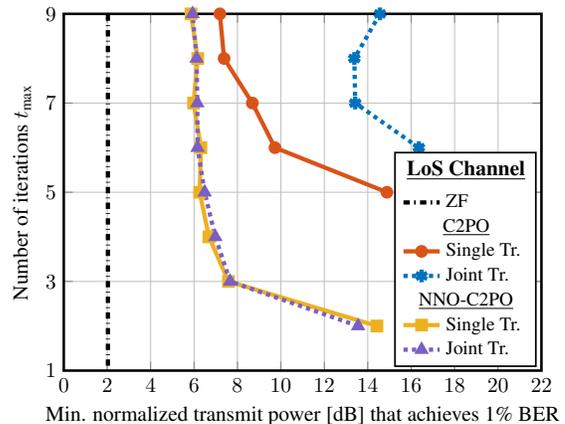


Fig. 4. Performance-iteration trade-offs for C2PO and NNO-C2PO for a massive MU-MIMO system with $B = 128$ BS antennas and $U = 8$ UEs, operating with 16-QAM in a realistic LoS channel. The vertical line shows the normalized transmit power at which ZF precoding achieves 1% BER.

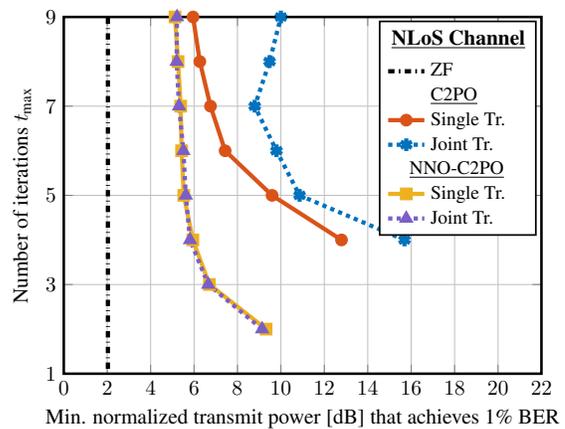


Fig. 5. Performance-iteration trade-offs for C2PO and NNO-C2PO for a massive MU-MIMO system with $B = 128$ BS antennas and $U = 8$ UEs, operating with 16-QAM in a realistic NLoS channel. The vertical line shows the normalized transmit power at which ZF precoding achieves 1% BER.

to achieve similar BER performance. Moreover, even when considering the case of $t_{\max} = 8$ where NNO-C2PO delivers diminishing returns, there is a difference of roughly 3 dB between the two algorithms at a BER of 0.1%. Note that this improvement comes at essentially no cost, since the C2PO hardware architecture of [13] can be modified to support NNO-C2PO with very little hardware overhead to operate with the per-iteration values of $\tau^{(t)}$ and $\rho^{(t)}$.

B. QuaDRiGa LoS and NLoS Channels

For the QuaDRiGa LoS and NLoS channel models, we present performance results using a different metric. Specifically, we fix the normalized transmit power, and then we determine the minimum number of C2PO and NNO-C2PO iterations that are required to achieve a target BER of 1%. Moreover, in order to examine the robustness of C2PO and NNO-C2PO to different channel conditions, we provide results for two scenarios. In the *single training* scenario, C2PO and NNO-C2PO are trained on a dataset that contains training data from only one channel model, and they are tested on data

from the same channel model they were trained for. In the *joint training* scenario, C2PO and NNO-C2PO are trained on a dataset that contains training data from both channel models, but they are tested on each channel model separately.

The results of the above experiments are shown in Figures 4 and 5, for the LoS and NLoS channels, respectively. We observe that, in the single training scenario, NNO-C2PO provides significant gains with respect to C2PO for both channel models. For example, on the LoS and NLoS channels, NNO-C2PO with $t_{\max} = 5$ achieves the target BER at 9 dB and 4 dB lower transmit power than C2PO, respectively. More interestingly, we observe that in the joint training scenario, the performance of C2PO on both channels becomes significantly worse, while the performance of NNO-C2PO remains nearly the same. The fact that NNO-C2PO is more parameterizable seems to make it sufficiently robust so that it can learn to operate on vastly different channel models with the same set of parameters $\{\tau^{(t)}, \rho^{(t)}\}$. This also indicates that NNO-C2PO may be robust to channel estimation errors.

V. CONCLUSIONS

We have shown that introducing per-iteration parameters to the C2PO precoding algorithm for 1-bit massive MU-MIMO systems and tuning them using deep-learning tools, enables significantly improved error-rate performance compared to a similarly-tuned C2PO variant that uses the same parameters for all iterations. Such improved performance is achieved without increasing the algorithmic complexity. Furthermore, by learning the algorithm parameters for different channel models, we have observed that our NNO-C2PO algorithm is robust to vastly changing propagation conditions, unlike its traditional C2PO counterpart. Our results indicate that tedious manual parameter tuning of nonlinear 1-bit precoding algorithms can be avoided in practice, while significantly improving their performance at the same time.

REFERENCES

- [1] F. Rusek, D. Persson, B. Kiong, E. G. Larsson, T. L. Marzetta, O. Edfors, and F. Tufvesson, "Scaling up MIMO: Opportunities and challenges with very large large arrays," *IEEE Signal Process. Mag.*, vol. 30, no. 1, pp. 40–60, Jan. 2013.
- [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] L. Lu, G. Ye Li, A. L. Swindlehurst, A. Ashikhmin, and R. Zhang, "An overview of massive MIMO: Benefits and challenges," *IEEE J. Sel. Topics Signal Process.*, vol. 8, no. 5, pp. 742–758, Oct. 2014.
- [4] A. Mezghani, R. Ghiat, and J. A. Nossek, "Transmit processing with low resolution D/A-converters," in *Proc. IEEE Int. Conf. Electron., Circuits, Syst. (ICECS)*, Dec. 2009, pp. 683–686.
- [5] A. K. Saxena, I. Fijalkow, and A. L. Swindlehurst, "Analysis of one-bit quantized precoding for the multiuser massive MIMO downlink," *IEEE Trans. Signal Process.*, vol. 65, no. 17, pp. 4624–4634, Sep. 2017.
- [6] Y. Li, C. Tao, A. L. Swindlehurst, A. Mezghani, and L. Liu, "Downlink achievable rate analysis in massive MIMO systems with one-bit DACs," *IEEE Commun. Lett.*, vol. 21, no. 7, pp. 1669–1672, Jul. 2017.
- [7] S. Jacobsson, G. Durisi, M. Coldrey, T. Goldstein, and C. Studer, "Quantized precoding for massive MU-MIMO," *IEEE Trans. Commun.*, vol. 65, no. 11, pp. 4670–4684, Nov. 2017.
- [8] —, "Nonlinear 1-bit precoding for massive MU-MIMO with higher-order modulation," in *Proc. Asilomar Conf. Signals, Syst., Comput.*, Nov. 2016, pp. 763–767.
- [9] H. Jedda, J. A. Nossek, and A. Mezghani, "Minimum BER precoding in 1-bit massive MIMO systems," in *IEEE Sensor Array and Multichannel Sig. Proc. Workshop*, Jul. 2016.
- [10] O. Tirkkonen and C. Studer, "Subset-codebook precoding for 1-bit massive multiuser MIMO," in *Conf. Inf. Sci. Syst.*, Mar. 2017.
- [11] A. L. Swindlehurst, A. K. Saxena, A. Mezghani, and I. Fijalkow, "Minimum probability-of-error perturbation precoding for the one-bit massive MIMO downlink," in *Proc. IEEE Intl. Conf. Acoust., Speech, Signal Process.*, Mar. 2017, pp. 6483–6487.
- [12] O. Castañeda, C. Studer, and T. Goldstein, "POKEMON: A non-linear beamforming algorithm for 1-bit massive MIMO," in *Proc. IEEE Intl. Conf. Acoust., Speech, Signal Process.*, Mar. 2017, pp. 3464–3468.
- [13] O. Castañeda, S. Jacobsson, G. Durisi, M. Coldrey, T. Goldstein, and C. Studer, "1-bit massive MU-MIMO precoding in VLSI," *IEEE J. Emerging Sel. Topics Circuits Syst.*, vol. 7, no. 4, pp. 508–522, Dec. 2017.
- [14] M. Shao, Q. Li, and W.-K. Ma, "One-bit massive MIMO precoding via minimum symbol-error probability design," in *Proc. IEEE Intl. Conf. Acoust., Speech, Signal Process.*, Mar. 2018, pp. 3579–3583.
- [15] A. Nedelcu, F. Steiner, M. Staudacher, G. Kramer, W. Zirwas, R. Sisava Ganesan, P. Baracca, and S. Wesemann, "Quantized precoding for multi-antenna downlink channels with MAGIQ," in *Int. ITG Workshop Smart Antennas*, Mar. 2017.
- [16] H. Jedda, A. Mezghani, A. L. Swindlehurst, and J. A. Nossek, "Quantized constant envelope precoding with PSK and QAM signaling," *IEEE Trans. Wireless Commun.*, vol. 17, no. 12, pp. 8022–8034, Oct. 2018.
- [17] N. Samuel, T. Diskin, and A. Wiesel, "Deep MIMO detection," in *Proc. IEEE Intl. Workshop Signal Process. Adv. Wireless Commun.*, Jul. 2017.
- [18] A. Klautau, N. González-Prelcic, A. Mezghani, and R. W. Heath Jr., "Detection and channel equalization with deep learning for low resolution MIMO systems," in *Proc. Asilomar Conf. Signals, Syst., Comput.*, Oct. 2018, pp. 1836–1840.
- [19] V. Corlay, J. J. Boutros, P. Ciblat, and L. Brunel, "Multilevel MIMO detection with deep learning," in *Proc. Asilomar Conf. Signals, Syst., Comput.*, Oct. 2018, pp. 1805–1809.
- [20] S. Khobahi, N. Naimipour, M. Soltanalian, and Y. C. Eldar, "Deep signal recovery with one-bit quantization," Nov. 2018. [Online]. Available: <https://arxiv.org/abs/1812.00797>
- [21] H. He, C.-K. Wen, S. Jin, and G. Y. Li, "A model-driven deep learning network for MIMO detection," in *Proc. IEEE Global Conf. Signal Inf. Process.*, Nov. 2018, pp. 584–588.
- [22] S. Takabe, M. Imanishi, T. Wadayama, and K. Hayashi, "Trainable projected gradient detector for massive overloaded MIMO channels: Data-driven tuning approach," Dec. 2018. [Online]. Available: <https://arxiv.org/abs/1812.10044>
- [23] S. Kim and S.-N. Hong, "Semi-supervised learning detector for MU-MIMO systems with one-bit ADCs," Feb. 2019. [Online]. Available: <https://arxiv.org/abs/1902.00866>
- [24] M. Borgerding, P. Schniter, and S. Rangan, "AMP-inspired deep networks for sparse linear inverse problems," *IEEE Trans. Signal Process.*, vol. 65, no. 16, pp. 4293–4308, Aug. 2017.
- [25] W. Xia, G. Zheng, Y. Zhu, J. Zhang, J. Wang, and A. Petropulu, "A deep learning framework for optimization of MISO downlink beamforming," Jan. 2019. [Online]. Available: <https://arxiv.org/abs/1901.00354>
- [26] 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>
- [27] 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.
- [28] I. Hubara, M. Courbariaux, D. Soudry, R. El-Yaniv, and Y. Bengio, "Binarized neural networks," in *Adv. Neural Inf. Process. Syst.*, Dec. 2016, pp. 4107–4115.
- [29] S. Jaeckel, L. Raschkowski, K. Borner, 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.
- [30] F. Chollet *et al.*, "Keras," <https://keras.io>, 2015.
- [31] M. Abadi *et al.*, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015, software available from tensorflow.org.
- [32] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. Intl. Conf. Learn. Representations*, May 2015.