

Robustness of Distributed Multi-User Beamforming: An Experimental Evaluation

Rahman Doost-Mohammady, Mehdi Zafari, and Ashutosh Sabharwal
Rice University

Abstract—In this paper, we study the robustness of distributed beamforming in the presence of hardware imperfections. In particular, we characterize the impact of carrier frequency offset (CFO) between distributed antennas on multi-user receive beamforming. We characterize the performance of distributed beamforming in the presence of CFO using a massive MIMO platform. Using datasets collected from the platform, we measure and compare the beamforming performance in the presence and absence of CFO. Further, we study the effect of variation in CFO among beamforming antennas on the performance through simulations. Interestingly, we observe that this drop in performance is higher as more distributed antennas are included in beamforming. Our results can be used to determine the optimal number of antennas in distributed beamforming systems.

Index Terms—Distributed MU-MIMO, Beamforming, Carrier Frequency offset

I. INTRODUCTION

Distributed beamforming has many applications, from distributed antenna systems (DAS) in infrastructure networks (e.g., cellular and WiFi) to wireless ad-hoc networks. In distributed beamforming, a collection of radios jointly transmit or receive messages through *coherent* beamforming which can increase communication range and improve rates. For coherent beamforming, it is important to have frequency and time synchronization among distributed antennas. However, this requirement tends to become a bottleneck as the number of distributed radios increases, and it is unclear how imperfections in synchronizations impact performance as the system scales.

In a distributed network of antennas, each antenna has an independent digital radio with its own local oscillator (LO). When operating at a carrier frequency F_c , naturally each radio will have a random offset from the true F_c , due to variations in manufacturing and other environmental factors such as temperature and age. Thus, the antennas must perform over-the-air (OTA) synchronization periodically to correct both their time and carrier frequency offset (CFO).

While time offsets can be accurately measured in signal-to-noise ratios (SNR) well below the communication range, the accuracy of CFO estimation can be severely impacted as SNR decreases. Fig. 1 shows the CFO estimation error in additive white Gaussian noise (AWGN) channel under different SNRs. At low SNRs, the CFO estimation error dramatically increases and can be as high as 10s of kHz even if the actual CFO is zero.

Recent advances in the design of LOs have led to significant improvement in their accuracy and cost-effectiveness. E.g., in OCXO and atomic clocks, the accuracy is down to 10s or




Fig. 1. The effect of SNR on CFO estimation error in an AWGN channel. For CFO estimation, we use an 802.11 long training (LTS) sequence [1] including two OFDM symbols of length 64 and add CFO and noise.

even single-digit parts per billion (ppb) [2], [3]. Thus, the newer generation of radio hardware employs LOs that exhibit a much lower average CFO at their operating radio frequency (RF) band compared to previous generations¹. The low CFO that is achieved by newer hardware leads us to the question that whether and to what extent small CFO quantities can be tolerated in distributed beamforming. To answer this question, the performance of distributed beamforming in the presence of CFO must be characterized in realistic settings. Past works, such as [4], [5], have theoretically modelled the effect of residual CFO on the performance of MU-MIMO with 2-4 antennas. Though to the best of our knowledge, no prior work has investigated the performance of multi-user beamforming in the presence of CFO through experiments, especially with a large number of beamforming antennas.

In this paper, we take the first step in the experimental study of distributed beamforming in the presence of CFO. In particular, we focus on the receive beamforming performance in time-division duplex (TDD) mode where the user terminals first transmit time-orthogonal pilots for channel estimation and subsequently transmit uplink data simultaneously toward up to 64 beamforming antennas. For our experiments, we use a massive MIMO platform that includes independent radios equipped with high precision LOs that exhibit sub-kHz CFO at the operating frequency of 3.6 GHz. Through analyzing tens of datasets collected from this platform, we compare the performance of receive beamforming in the presence and absence of frequency synchronization among beamforming antennas. Furthermore, we evaluate the performance of two different beamforming methods, i.e. conjugate and Zero Forcing. We

¹Here we are focusing mostly on sub-6 GHz bands since at higher bands such as mmWave, the CFO could be large even with today's high-accuracy oscillators.

use the average EVM-SNR per user as our performance metric. Our experiments show a significant drop (6-9 dB) in average EVM-SNR when there is CFO among beamforming antennas. Interestingly, we observe that when CFO is present, the performance of conjugate and ZF beamforming is quite similar. Lastly, given the fixed distribution of CFO in our platform, we perform simulations to evaluate the effect of other CFO distributions on the performance of distributed beamforming. Our simulation results show that in high CFO regimes, a higher number of beamforming antennas yield lower beamforming performance.

II. BACKGROUND

The most practical method for multi-user beamforming is *linear precoding*. In receive beamforming, the received signal vector from the antennas is multiplied with a beamforming matrix. To describe this process better, we assume the following signal model:

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{n}, \quad (1)$$

where $\mathbf{y}$ is a receive data vector of size $M \times 1$ where M is the number of beamforming antennas, $\mathbf{x}$ is a $K \times 1$ transmit vector from K users, $\mathbf{H}$ is an $M \times K$ channel matrix and $\mathbf{n}$ is the complex circularly Gaussian noise vector at the receiver antennas. To estimate $\mathbf{x}$ at the receiver, the beamforming matrix $\mathbf{W}$ is applied to the receive vector,

$$\tilde{\mathbf{x}} = \mathbf{W}^H \mathbf{y}. \quad (2)$$

where $(\cdot)^H$ denotes the hermitian operation. Some of the widely used beamforming methods are Conjugate and Zero Forcing (ZF) where their beamforming matrix is calculated as $\mathbf{W} = \mathbf{H}$ and $\mathbf{W} = \mathbf{H}(\mathbf{H}^H \mathbf{H})^{-1}$, respectively.

Linear methods are sub-optimal in performance but more practical to implement. Conjugate beamforming is the simplest method in terms of computational complexity but underperforms ZF in medium and high-SNR regimes. In the low-SNR regime, however, conjugate beamforming performs better than ZF as modelled in [6], [7].

III. EXPERIMENT DESIGN

A. Hardware Setup

We use the RENEW programmable massive MIMO platform [8] to perform experiments. The RENEW massive-MIMO base station uses the Iris SDR [9] as its building block with each SDR supporting two RF chains. The base station includes a central hub that by default distributes clock and time trigger signals to up to 48 Iris SDRs (96 antennas) for time and frequency synchronization. At each Iris, both the RF and sampling clocks are derived from the clock distributed by the hub. Therefore, all antennas are locked at carrier frequency F_c in the form of a *coherent* array. An Iris can also be used in standalone mode to emulate a user. The RENEW platform is designed to operate at the CBRS spectrum ranging from 3.55-3.7 GHz.

Apart from the capability to receive an external clock, each Iris is equipped with its own LO, a super-TCXO with a nominal ± 100 ppb accuracy. Through firmware settings, the Iris can be toggled to use its own LO to generate the RF clock

or use the external clock from the central hub. To emulate a distributed antenna array, we program each Iris in the base station to use its own LO. In this mode, there will be a CFO between the antennas. We call the array in this mode to be *non-coherent*. In the non-coherent mode, the antennas still receive the time trigger signal from the central hub as their time reference. This allows us to only focus on the effect of CFO.




Fig. 2. Our setup including a 64-antenna base station (red box) and 4 users (one shown in the yellow box).

B. Experiment setup

For our experiments, we use an indoor 64-antenna RENEW massive MIMO array, as shown in Fig. 2. We set up four Iris users in the vicinity of the massive MIMO array. Initially, we deploy the users in the same room as the base station, 5-7 meters away from the array. We change the location and orientation of the users during the experiments to capture as many channel conditions as possible. Additionally, we place the users in the adjacent room to capture non-line-of-sight (NLOS) channels.

We use the RENEWLab real-time software [10] to collect datasets in all user configurations. The RENEWLab software runs on a host server that is connected to the RENEW base station through fiber. Through a configurable frame schedule, RENEWLab software orchestrates the transmission and reception of pilot and data OFDM symbols at the users and from the base station antennas, respectively. Within each frame, one of the base station antennas transmits a synchronization beacon so that users can synchronize their frame time boundaries with the base station frame. We set up the frame schedule for the users so that each user sends time-orthogonal pilots and then all simultaneously transmit data OFDM symbols. We use the 802.11 LTS signal as the pilot symbol which is used for channel estimation and beamforming matrix calculation. The format of the data OFDM symbols also follows the 802.11 standard where the FFT size is 64. Out of the 64 subcarriers, 48 subcarriers carry data, 4 are pilot subcarriers, and the remaining are null subcarriers. The pilot subcarriers are used for residual phase offset correction. We try both QPSK and 16QAM for modulating data subcarriers. RENEWLab framework enables the recording of pilot and data OFDM symbols received at base station antennas in a dataset file with HDF5 format [11]. The HDF5 datasets can be processed offline with RENEWLab's Python-based post-processing library.




Fig. 3. Left: Distribution of the CFO measured at each beamforming antenna in our setup in non-coherent mode with respect to a reference beamforming antenna. Right: Distribution of the received signal SNR at the base station antennas from all our experiments.

During data collection, we keep the number of beamforming antennas at 64 and we use 1, 2, and 4 transmitting users. For each number of transmitting antennas, we collect data in 16 different user positions including LOS and NLOS positions. We collect 700 frames for each measurement with 10 data OFDM symbols in each frame. We set the carrier frequency to 3.6 GHz and the sample rate at 5 MHz. For each configuration, we collect data in both coherent and non-coherent modes. The collected datasets amount to 200 GB of data. We have made these datasets open-source [12].

C. Setup Characterization

Due to manufacturing variations, oscillators employed in Irises have random offsets from the true operating frequency. To characterize the CFO in the non-coherent array, we measure the CFO of each antenna with respect to a reference beamforming antenna. For each non-reference beamforming antenna, we collect 100 LTS pilots, each including two OFDM symbols with 64 samples each. For two consecutive pilots, we measure the CFO in Hz through the following operation:

$$\frac{\arg\left(\sum_{n=1}^N y^*[n]y[n+N]\right)/N}{2\pi rN} \quad (3)$$

where r is the sample rate, $N = 64$, and $\arg(\cdot)$ denotes the unwrapped angle. In Fig. 3 Left, a histogram of the CFO values for beamforming antennas is shown. We fit a normal distribution to the histogram with a mean of 83 Hz and a variance of 840 Hz. Based on this observation, we later consider a similar model for CFO in distributed antennas. More specifically, we model the CFO in distributed antennas with a normal distribution with zero mean and standard deviation σ_{CFO} . The standard deviation σ_{CFO} can be expressed either in Hz or ppb.

In addition to measuring the CFO in our setup, we also measure the SNR of the received pilots at each beamforming antenna from all users at different positions. We plot a histogram of the received SNR in Fig. 3 Right which shows a relatively high SNR regime with a small fraction of the channels showing low SNR. The low-SNR channels are all in NLOS measurements.




Fig. 4. Experimental results showing average EVM-SNR per user vs. number of ZF beamforming antennas in coherent and non-coherent array.




Fig. 5. Experimental results showing average EVM-SNR per user vs. number of conjugate beamforming antennas in coherent and non-coherent array.

IV. RESULTS

In this section, we provide both experimental and simulation results demonstrating the effect of CFO on the performance of distributed beamforming. First, we go through the results obtained from our dataset collection campaign explained in § III. Since the CFO distribution in our hardware platform is fixed, we resort to simulations to get a better understanding of how different CFO distributions will affect the beamforming performance. To evaluate performance, we calculate the EVM in the received constellation for each user and derive the EVM-SNR as the inverse of the EVM. We plot the average of EVM-SNR for all users. The definition of EVM for each user is as follows [13]:

$$\frac{\sum_{l=1}^L |x[l] - \tilde{x}[l]|^2}{\sum_{l=1}^L |x[l]|^2} \quad (4)$$

where $x[l]$ is the transmitted symbol, $\tilde{x}[l]$ is the equalized symbol, and L is the total number of received equalized data symbols.

A. Experimental Results

Our datasets include 64×1 , 64×2 , and 64×4 MIMO configurations for both coherent and non-coherent array. For each configuration, we collect datasets in up to 16 locations.




Fig. 6. Simulation results showing the average EVM-SNR per user in different MIMO configuration vs. CFO distributions.

We implement both Conjugate and ZF beamforming and apply them to each dataset. We calculate the EVM-SNR for each user's constellation after the beamforming operation and average that among the users. To evaluate the effect of CFO, we vary the number of beamforming antennas from 8 to 64 in steps of 8. At each point, we take a 64 random selection of existing antennas in each dataset and average the results.

For ZF beamforming, as shown in Fig. 4, the average EVM-SNR slowly improves with the number of beamforming antennas for the coherent array. Our results are in line with theoretical findings in [7] where it is shown that, when CSI is imperfect, the spectral efficiency in receive beamforming will converge to a non-zero value as the number of antennas goes up. For the non-coherent case, there is a significant drop in ZF performance, which widens as the number of users increases. A counter-intuitive observation is that, in single-user beamforming, the performance slightly drops as we increase the number of antennas. To better characterize this effect, we perform simulations in §IV-B to see the effect of different CFO scenarios on the average EVM-SNR.

Fig. 5 shows the EVM-SNR results for conjugate beamforming. In the coherent case, the results are again in line with analytical findings in earlier works such as [7]. Specifically in MU-MIMO (2 and 4 users), we observe a more significant performance drop compared to ZF, due to higher inter-user interference in conjugate beamforming. Note that, for 64×1 configuration, the performance is the same as ZF since conjugate and ZF are essentially the same beamformer for a single user. Interestingly, we observe similar performances for all MIMO configurations between conjugate and ZF beamforming in non-coherent mode. This observation demonstrates that ZF fails to effectively cancel the inter-user interference in the presence of CFO and therefore conjugate beamforming is more preferred due to less complexity.

B. Simulation Results

Our experimental results reveal interesting trends in the performance of distributed beamforming in the presence of CFO. We observed in §IV-A that the performance drops for single-user beamforming as we increase the number of beamforming antennas. Given the fixed distribution of CFO in our hardware, we simulate an OFDM-based distributed beamforming system with a similar setup and parameters as our experiments. The OFDM symbols follow the same format as our experimental setup in §III-B. We use a rayleigh fading channel model and an average SNR of 20 dB in the

simulations. We also enable the configuration of the number of beamforming antennas and users. We parametrize the distribution of CFO by varying σ_{CFO} as described in §III-C. We increase σ_{CFO} from zero (no CFO) to 300 ppb, equivalent to 1 kHz of CFO at 3.6 GHz carrier frequency. At each variance, we apply a randomly generated residual CFO to the pilot and data at each beamforming antenna. We apply both ZF and conjugate beamformers to the received data. Similar to our experiments, we use the average EVM-SNR per user as our performance metric. For each CFO distribution, we perform 250 random experiments and we average the results.

The plots in Fig. 6, show the average EVM-SNR obtained from ZF beamforming with respect to CFO error (σ_{CFO}) for 1, 2, and 4 users. For each case, we vary the number of beamforming antennas from 8 to 64. The results show the average EVM-SNR decreases for all beamforming antenna numbers and users as the CFO error increases. Interestingly, the results show that at higher CFOs, using more beamforming antennas reduces performance. For one and two users, we observe a cross-over point (roughly at 100 ppb for one user and 160 ppb for two users) after which more beamforming antennas result in lower performance. As the number of users increases, the cross-over point shifts to higher CFOs. We obtained similar results for conjugate beamforming that are omitted due to the lack of space. For conjugate, the cross-over point is the same for one user since, in this case, conjugate is the same beamformer as ZF. For two users, however, we observe the cross-over occurs at 200 ppb compared to 160 ppb for ZF with two users. Our conclusion from these results is that the knowledge of CFO distribution in the beamforming hardware can guide the setting of the number of antennas for a beamforming scenario. Moreover, we can determine what type of beamforming method can be more beneficial for each CFO distribution.

V. CONCLUSION & FUTURE WORK

The lack of frequency synchronization among antennas is a challenge in distributed beamforming. In this paper, we perform an experimental evaluation of the effect of CFO among beamforming antennas on the performance of distributed receive beamforming using a programmable massive MIMO platform. From the collected datasets, we quantify the beamforming performance degradation in the presence of CFO vs the zero CFO scenario. To evaluate the beamforming performance over a range of CFOs, we perform simulation with a similar setup to the experimental platform. Our simulation

results show that in large CFOs, performance degradation is higher as more distributed antennas are included in beamforming. In our future work, we will evaluate the effect of CFO on distributed transmit beamforming in TDD mode, including the effect of CFO on the accuracy of reciprocity calibration. We will also analytically study the performance of the distributed beamforming in the presence of CFO.

ACKNOWLEDGMENT

This work was sponsored in part by NSF PAWR Grant 1827940 for the POWDER-RENEW project, NSF CCRI Grant 2016727, and a grant from Army Research Labs.

REFERENCES

- [1] E. Sourour, H. El-Ghoroury, and D. McNeill. Frequency offset estimation and correction in the IEEE 802.11a wlan. In *IEEE 60th Vehicular Technology Conference, 2004. VTC2004-Fall. 2004*, volume 7, pages 4923–4927 Vol. 7, 2004.
- [2] Thomas Schmid, Zainul Charbiwala, Jonathan Friedman, Young H. Cho, and Mani B. Srivastava. Exploiting manufacturing variations for compensating environment-induced clock drift in time synchronization. *SIGMETRICS Perform. Eval. Rev.*, 36(1):97–108, jun 2008.
- [3] B.M. Sadler. Fundamentals of energy-constrained sensor network systems. *IEEE Aerospace and Electronic Systems Magazine*, 20(8):17–35, 2005.
- [4] Y.Y.L. Lebrun, K.K.Z. Zhao, S.S.P. Pollin, and et al. Performance analysis of distributed zf beamforming in the presence of cfo. *J Wireless Com Network*, 177, 2011.
- [5] R. Rogalin, O. Y. Bursalioglu, H. C. Papadopoulos, G. Caire, A. F. Molisch, A. Michaloliakos, H. V. Balan, and K. Psounis. Scalable synchronization and reciprocity calibration for distributed multiuser MIMO. *IEEE Transactions on Wireless Communications*, pages 1815–1831, 2014.
- [6] Hong Yang and Thomas L. Marzetta. Performance of conjugate and zero-forcing beamforming in large-scale antenna systems. *IEEE Journal on Selected Areas in Communications*, 31(2):172–179, 2013.
- [7] Hien Quoc Ngo, Erik G. Larsson, and Thomas L. Marzetta. Energy and spectral efficiency of very large multiuser mimo systems. *IEEE Transactions on Communications*, 61(4):1436–1449, 2013.
- [8] R. Doost-Mohammady, O. Bejarano, L. Zhong, J. R. Cavallaro, E. Knightly, Z. M. Mao, W. W. Li, X. Chen, and A. Sabharwal. Renew: Programmable and observable massive mimo networks. In *2018 52nd Asilomar Conference on Signals, Systems, and Computers*, pages 1654–1658, Oct 2018.
- [9] C. Shepard, J. Blum, R.E. Guerra, R. Doost-Mohammady, and L. Zhong. Design and implementation of scalable massive-mimo networks. *Open-Wireless'20*, page 7–13, 2020.
- [10] Rice University. Renew software git repository, 2022. <https://github.com/renew-wireless/RENEWLab>.
- [11] The HDF5 Group. The hdf5 library and file format, 2022. <https://www.hdfgroup.org/solutions/hdf5/>.
- [12] Rice University. Renew datasets, 2022. <https://renew-wireless.org/datasets.html>.
- [13] Rishad Ahmed Shafik, Md. Shahriar Rahman, and AHM Razibul Islam. On the extended relationships among evm, ber and snr as performance metrics. In *2006 International Conference on Electrical and Computer Engineering*, pages 408–411, 2006.