

Scalable Synchronization and Reciprocity Calibration for Distributed Multiuser MIMO

R. Rogalin¹, O. Y. Bursalioglu², H. Papadopoulos², G. Caire¹, A. Molisch¹,
A. Michaloliakos¹, V. Balan¹, and K. Psounis¹

Abstract

Large-scale distributed Multiuser MIMO (MU-MIMO) is a promising wireless network architecture that combines the advantages of “massive MIMO” and “small cells.” It consists of several Access Points (APs) connected to a central server via a wired backhaul network and acting as a large distributed antenna system. We focus on the downlink, which is both more demanding in terms of traffic and more challenging in terms of implementation than the uplink. In order to enable multiuser joint precoding of the downlink signals, channel state information at the transmitter side is required. We consider Time Division Duplex (TDD), where the *downlink* channels can be learned from the user *uplink* pilot signals, thanks to channel reciprocity. Furthermore, coherent multiuser joint precoding is possible only if the APs maintain a sufficiently accurate relative timing and phase synchronization.

AP synchronization and TDD reciprocity calibration are two key problems to be solved in order to enable distributed MU-MIMO downlink. In this paper, we propose novel over-the-air synchronization and calibration protocols that scale well with the network size. The proposed schemes can be applied to networks formed by a large number of APs, each of which is driven by an inexpensive 802.11-grade clock and has a standard RF front-end, not explicitly designed to be reciprocal. Our protocols can incorporate, as a building block, any suitable timing and frequency estimator. Here we revisit the problem of joint ML timing and frequency estimation and use the corresponding Cramer-Rao bound to evaluate the performance of the synchronization protocol. Overall, the proposed synchronization and calibration schemes are shown to achieve sufficient accuracy for satisfactory distributed MU-MIMO performance.

Index Terms

Distributed multiuser MIMO downlink, cooperative small cells, synchronization, TDD calibration.

¹ University of Southern California, Ming-Hsieh Department of Electrical Engineering, Los Angeles, CA 90089, email: rogalin, caire, molisch, michalol, hbalan, kpsounis@usc.edu

² DOCOMO Innovations, Inc., 3240 Hillview Avenue, Palo Alto, California 94304, email: obursalioglu, hpapadopoulos@docomoinnovations.com

I. INTRODUCTION

The explosive growth of wireless data traffic, spurred by powerful and data-hungry user devices such as smartphones and tablets, is putting existing wireless networks under stress and is calling for a technology paradigm shift. Two recent research trends in wireless networks have attracted significant attention as potential solutions: massive MIMO and very dense spatial reuse. The former refers to serving multiple users on the same time-frequency channel resource by eliminating multiuser interference via spatial precoding, using a very large number of antennas at each base station site [1], [2]. The latter pertains to the use of small cells [3], such that multiple short-range low-power links can co-exist on the same time-frequency channel resource thanks to sufficient spatial separation. Distributed Multiuser MIMO (MU-MIMO) unifies these two approaches, simultaneously obtaining both multiuser interference suppression through spatial precoding, and dense coverage by reducing the average distance between transmitters and receivers. This is achieved by coordinating a large number of Access Points (APs), distributed over a certain coverage region, through a wired backhaul network connected to a Central Server (CS), in order to form a distributed antenna system.

While distributed MU-MIMO may be applied in a variety of different scenarios, in this work we focus primarily on cost-effective consumer grade equipment and focus particularly on the downlink (DL), which is both more demanding in terms of traffic and more challenging in terms of implementation than the uplink (UL). We assume inexpensive APs connected to the CS via a conventional wired digital backbone (e.g., Ethernet), capable of transporting digital packetized data at sufficient speed from/to the CS, but not able to distribute a common clock (i.e., timing and phase information) to the APs, beyond the limitations of standard network synchronization protocols such as IEEE 1588 [4]. In order to achieve large spectral efficiencies through MU-MIMO downlink spatial precoding, channel state information at the transmitter (CSIT) is needed. Following the massive MIMO approach [1], [2], CSIT can be obtained from the users' uplink pilots by operating the system in Time-Division Duplexing (TDD) and exploiting the reciprocity of the radio channels. However, while the propagation channel from antenna to antenna is reciprocal,¹

¹This statement holds if the UL and DL transmissions take place within the coherence time and coherence bandwidth of the underlying doubly-selective channel. Using OFDM and TDD, the propagation channel is effectively reciprocal on each OFDM subcarrier over many consecutive OFDM symbols, since the coherence time for typical wireless local area networks with nomadic users may span ~ 0.1 s to 1s or more. In fact, 800 msec is commonly used as the simulation assumption for 801.11ac IEEE standard for WLAN [5]. In WARP software-radio based experiments, we observed the coherence time to be about 1 second. In contrast, Frequency Division Duplexing schemes have the UL and the DL separated by tens of MHz, well beyond the coherence bandwidth. Therefore, reciprocity for such systems typically does not hold.

the transmitter and receiver hardware are not, and introduce an unknown amplitude scaling and phase shift between the UL and DL channels [6]. Furthermore, in order to have all the remote APs behave as a single distributed multiple antenna system, the APs must maintain relative timing and phase synchronism throughout the duration of a DL precoded slot (see Section II).

Since each AP is driven by its own individual clock, and the wired backhaul network is not capable of providing a sufficiently accurate common timing and frequency reference, such synchronization must occur via over-the-air signaling. Also, since typical commercial grade radios are not equipped with a built-in self-calibration capability, the non-reciprocal effects of the receiver and transmitter hardware must be compensated explicitly via a TDD reciprocity calibration protocol. Note that, although in our treatment we focus on 802.11 grade hardware, the methods we propose are platform agnostic. Indeed, the same principles can be applied to any future cost-effective flexible network deployments, where practical hardware mismatches need to be eliminated in order to enable coherent reciprocity-based distributed MU-MIMO transmission.

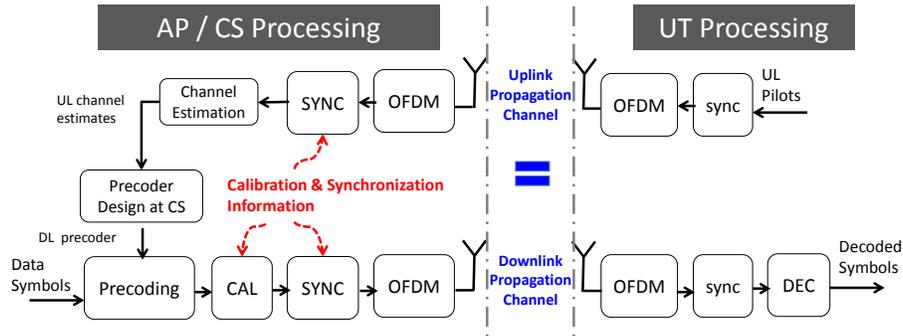


Fig. 1. Uplink training and downlink transmission for distributed channel-reciprocity based MU-MIMO.

Fig. 1 provides a high-level block diagram that illustrates the operation of channel-reciprocity based MU-MIMO on the hardware, including UL channel training, DL MU-MIMO transmission, and the compensation mechanisms (synchronization and calibration) needed to enable such coherent DL MU-

MIMO transmission. OFDM is assumed throughout this paper, since it is widely used in current WLAN (802.11a/g/n/ac) and cellular (LTE/802.16m) standards [7]. An UL-pilot/DL-precoded transmission cycle works as follows. At first, UL pilots are transmitted by the User Terminals (UTs). After receiving the UL pilots, each AP individually processes its observations through an OFDM front-end coupled with a standard synchronization block, as used in any WLAN current AP implementation. Estimates of the UL channels (between each AP and the pilot-transmitting UTs) are then formed at each AP. These estimates are then sent to the CS through the digital backhaul network and are jointly used at the CS to calculate the MU-MIMO precoding matrix [1], [2]. In the DL transmission phase, the packets of OFDM encoded symbols destined to the UTs are first precoded using the DL precoding matrix. Then, the precoded signal at each AP is calibrated through the “CAL” block, which compensates for the non-reciprocal amplitude scaling and phase rotations introduced by the transmitter and receiver AP hardware.² Finally, all the APs forming the distributed MU-MIMO network send simultaneously, on the same time-frequency slot, the precoded and calibrated data packets.

We discuss now the synchronization blocks indicated by “sync” and “SYNC” in Fig. 1 at the UT and AP side, respectively. Synchronization at the UT side (both transmitter and receiver) and at the AP receiver takes care of frame and carrier frequency synchronization in order to transmit and receive on the assigned time slots and demodulate the OFDM signals with negligible inter-carrier interference (ICI). This can be implemented by completely standard schemes used in current WLAN technology and forming a well-known and thoroughly studied subject [9]–[13]. Hence, such blocks need not be treated explicitly in this paper. In contrast, synchronization at the AP *transmitter* side plays a critical role in the distributed MU-MIMO architecture, since it needs to compensate for timing misalignment and the relative phase rotation of the DL data blocks, transmitted *simultaneously* by the jointly precoded APs. Such compensation is non-standard in current network implementations and it is *not* widely treated in the literature (see the existing-results overview provided below). In fact, in most theoretical works on distributed MU-MIMO (see for example [2], [14]–[17]) perfect transmitter synchronization is *assumed* without specifying how this is achieved. As shown in Section II-A, uncompensated synchronization errors at the AP transmitter side cannot be undone at the UT receivers, since the jointly precoded signals are mixed up and cannot be individually resampled (for timing) and phase-rotated (for carrier frequency

²There are several ways to apply calibration, some of which are not captured by the diagram in Fig. 1. For example, calibration can be applied to the UL channel estimates before the precoding matrix calculation [8]. While specific implementations may depend on the specific hardware at hand, all such solutions are conceptually equivalent to the scheme in Fig. 1.

offset). For a correct operation of the distributed MU-MIMO architecture it is therefore essential that: 1) the channel amplitude and phase rotations incurred by UL channel estimation can be compensated at the AP side (TDD calibration), and 2) the jointly transmitting APs send their precoded data slots with sufficiently accurate timing and carrier frequency coherence. This paper focuses on the design of signaling protocols and associated signal processing techniques for the CAL and SYNC blocks at the APs side of Fig. 1 that accomplish the above goals with sufficient accuracy so as to materialize a large fraction of the theoretical capacity gain promised by distributed MU-MIMO.

A. Literature Overview

Both synchronization of multiple transmitters and TDD reciprocity calibration have been treated to some extent in the existing literature. The synchronization of multiple transmitters for multiple-input single-output (MISO) cooperative beamforming is considered in [18], [19], and software-defined radio (SDR) implementations of such ideas are presented in [20]–[22]. SDR implementations of MU-MIMO have been recently presented by a few research teams (including our own) in [23]–[25], where the feasibility of over-the-air synchronization with enough accuracy to guarantee near-ideal multiuser interference suppression is demonstrated. These works are based on a master-slave protocol, requiring that all APs receive from a single master station broadcasting a beacon signal with sufficiently large power and thus limiting the size and topology of the network. Extensions to networks of arbitrary size and topology based on distributed consensus protocols are considered in [26], [27]. The slow rate of convergence, however, makes such schemes impractical for large networks.³ TDD reciprocity was considered in [6], presenting a calibration technique based on exchanging pilot signals between transmitter and receiver. In practice, it is much more desirable to design calibration protocols that do not assume the presence or collaboration of the UTs, which may be legacy devices not necessarily implementing the protocol. A novel TDD reciprocity calibration method, referred to as “Argos”, was presented in [28] as part of an SDR implementation of a TDD reciprocity-based massive MIMO base station. This method enables calibration by exploiting two-way signaling that involves only base station antennas, exchanging calibration pilot signals with a (possibly additional) *reference* antenna. It turns out that Argos is very sensitive to the placement of the

³For example, [26] reports a convergence time of ~ 100 seconds, widely insufficient to cope with the typical oscillator dynamics of the order of seconds. In addition, consensus schemes are targeted to infrastructure-less ad-hoc networks and involve over-the-air message passing between the nodes. In contrast, for the distributed MU-MIMO application considered in this work, we can effectively and efficiently exploit the presence of the wired backhaul network, which is needed anyway to share user data and CSIT across the APs.

reference antenna relative to the other antennas [28]. As a result, this scheme is not readily scalable and is not sufficiently accurate to enable large-scale MIMO in a distributed system.

B. Contributions of this Work

In this paper we present novel AP synchronization and TDD reciprocity calibration protocols that overcome the drawbacks of existing proposals. We assume that the APs form a connected network.⁴ For synchronization, a hierarchical scheme is defined where a subset of “anchor” APs is selected at the time of network deployment in order to form a connected cover of the network graph (see Section III). The anchor APs exchange pilot signals during special synchronization slots, periodically inserted in the frame structure. A pilot reuse scheme is used in order to avoid inter-pilot interference. The noisy timing and frequency estimates extracted from the pilots are sent to the CS via the wired backhaul, and centrally processed by solving a global *weighted* Least-Squares (LS) minimization. Then, the estimated timing and frequency correction factors are sent back to the anchor APs. All other APs get synchronized by their neighboring anchor APs, in a sort of “local master-slave” scheme. Unlike the master-slave topology, which requires having one AP (the master) with sufficiently high SNR to *all* other APs (clearly impossible in large networks), the proposed scheme has significantly more relaxed requirements. First, it requires a network of geographically dispersed anchor nodes, whereby each anchor AP is “connected” (has sufficiently high SNR) to at least another anchor AP, such that the anchor nodes form a connected graph. Second, it requires that the anchor density in the network is sufficiently high such that every other AP has at least one anchor in its vicinity, with respect to which it can be synchronized (e.g., via a local master-slave protocol).

For calibration we define a spanning connected subgraph of the network and have all nodes exchange calibration pilots with their neighbors on this subgraph. Noisy complex amplitude estimates of the received pilot signals are sent to the CS, which finds the calibration coefficients by solving a *constrained* LS problem. Then, the estimated calibration coefficients are sent back to the APs. The complexity of both the weighted LS and the constrained LS for synchronization and calibration is polynomial in the network size and it is easily affordable for networks of practical size. In order to extract timing and frequency information from the pilot signals, the APs may use any suitable estimator (e.g., [9]–[13]). In this paper, we consider joint ML timing and frequency estimation through an unknown multi-path channel and use

⁴Two APs are connected if they can receive each other’s pilot signals with sufficiently large SNR, where “sufficiently large” means above some design threshold determined by the system designer.

the corresponding Cramer-Rao Bound (CRB) as a good approximation of the estimation MSE achievable in practice [29], [30].

The rest of this paper is organized as follows. Section II presents a simplified but sufficiently accurate model for OFDM in the presence of timing, sampling frequency and carrier frequency offsets, and motivates the need for accurate compensation of these effects in a distributed MU-MIMO scenario. Section III discusses the high-level protocol architecture for scalable synchronization and TDD reciprocity calibration. In Sections IV and V we present the details of the synchronization and calibration protocols, respectively, and show their effectiveness through system simulations. Conclusions and discussion of future work are pointed out in Section VI, and Appendix A contains the derivation of the joint ML timing and frequency estimation and of the associated CRB.

II. A SIMPLE MODEL FOR OFDM WITH SYNCHRONIZATION ERRORS

Let f_0 and f_s be the nominal carrier and sampling frequencies of all nodes of a wireless network, respectively. The actual carrier and sampling frequencies of node i , indicated by $f_{0,i}$ and $f_{s,i}$, respectively, differ from their nominal values by some *deterministic unknown* carrier frequency offset (CFO) and sampling frequency offset (SFO), respectively. Also, each node i operates according to its local time axis with timing offset (TO) τ_i with respect to a common nominal time axis.

Motivated by the 802.11 standard [7], we assume that the sampling clock and the RF clock on each node hardware are derived from the same (local) oscillator, such that $f_{0,i}/f_{s,i} = \kappa \gg 1$ where κ is a constant factor that depends only on the hardware design but is independent of i . Then, we let $f_{s,i} = f_s + \epsilon_i$, where $\epsilon_i \ll f_s$ is the SFO at node i . Correspondingly, the CFO is given by $\kappa\epsilon_i$. Consistently with the typical accuracy of 802.11 commercial grade hardware, we assume that the CFO is much smaller than the signal bandwidth.

Consider the transmission of a sequence of OFDM symbols from node i to node j . For the sake of clarity of exposition, we assume that each node has a single antenna, although the synchronization and calibration schemes proposed in this work can be immediately applied to multi-antenna APs. We let $\mathbf{X}_i[n] = (X_i[n, 0], \dots, X_i[n, N - 1])$ denote the n -th block of frequency-domain symbols (i.e., the n -th OFDM symbol), where N is the number of OFDM subcarriers and $X_i[n, \nu]$ is the complex baseband symbol sent on subcarrier $\nu \in \{0, \dots, N - 1\}$ at OFDM symbol time n . The OFDM modulator performs an IFFT, producing the block of N time-domain “chips” $\mathbf{x}_i[n] = \text{IFFT}(\mathbf{X}_i[n]) = (x_i[n, 0], \dots, x_i[n, N - 1])$, where $x_i[n, k]$ denotes the k -th chip of the n -th time-domain block. Each block $\mathbf{x}_i[n]$ is cyclic-prefixed, forming $\tilde{\mathbf{x}}_i[n] = (\tilde{x}_i[n, -L], \dots, \tilde{x}_i[n, N - 1])$, where L is the length of the cyclic prefix (CP).

The sequence of cyclic-prefixed blocks is sent to the Digital-to-Analog Converter (DAC), forming the continuous-time complex baseband signal

$$\tilde{x}_i(t) = \sum_n \sum_{k=-L}^{N-1} \tilde{x}_i[n, k] \Pi((t - \tau_i)/T_{s,i} - n(N + L) - k), \quad (1)$$

where $\Pi(t)$ is the elementary interpolator DAC waveform, assumed for simplicity to be the rectangular pulse

$$\Pi(t) = \begin{cases} 1 & t \in [-1/2, 1/2] \\ 0 & \text{elsewhere} \end{cases} \quad (2)$$

and where $T_{s,i} = 1/f_{s,i}$ is the sampling interval of transmitter i . The complex baseband signal (1) is carrier-modulated in order to produce the transmitted signal $s_i(t) = \text{Re} \{ \tilde{x}_i(t) \exp(j2\pi f_{0,i}t) \}$.

Now consider the corresponding receiver operations at node j . Let $h_{i \rightarrow j}(t)$ denote the complex baseband equivalent impulse response of the channel between node i and node j , including the receiver hardware (LNA, Analog-to-Digital Converter (ADC), anti-aliasing filter, etc.). Since the CFO between node i and node j is small with respect to the signal bandwidth, the complex baseband signal after demodulation can be written as (see also [31])

$$\begin{aligned} \tilde{y}_{i \rightarrow j}(t) &= \left(\sum_n \sum_{k=-L}^{N-1} \tilde{x}_i[n, k] g_{i \rightarrow j}(t - (\tau_i - \tau_j) - (n(N + L) + k)T_{s,i}) \right) \\ &\times \exp(j2\pi(f_{0,i} - f_{0,j})t), \end{aligned} \quad (3)$$

where we define ⁵ $g_{i \rightarrow j}(t) = (\Pi_i \otimes h_{i \rightarrow j})(t)$, with $\Pi_i(t) = \Pi(t/T_{s,i})$.

The ADC at receiver j produces samples at rate $p/T_{s,j}$, where $p \geq 1$ is a suitable oversampling integer factor. The resulting discrete-time complex baseband signal, is given by $\tilde{y}_{i \rightarrow j}[\ell] = \tilde{y}_{i \rightarrow j}(\ell T_{s,j}/p)$. Since $T_{s,i}$ and $T_{s,j}$ are both assumed to be very close to the nominal chip interval T_s , over a single chip, the scale of the time axis at the two nodes is essentially identical. However, by accumulating this small time difference over many OFDM symbols, the SFO manifests itself as a timing misalignment of the time axes that grows linearly with the OFDM symbol index. Letting $\ell = m(N + L)p + q$ with $q \in \{-Lp, \dots, Np - 1\}$, extracting blocks of $p(N + L)$ output samples, and assuming $|\tau_i - \tau_j| \ll T_s L$, i.e., that the TO is significantly less than the duration of the CP, ⁶ we arrive at the following channel

⁵ $(a \otimes b)(t)$ denotes the continuous-time convolution of signals $a(t)$ and $b(t)$.

⁶As it will be clear from Section III, the synchronization protocol that we propose in this paper makes sure that this assumption is satisfied.

model without inter-block interference between consecutive OFDM symbols: for $q \in \{0, \dots, pN - 1\}$

$$\begin{aligned} \tilde{y}_{i \rightarrow j}[m, q] &= \left(\sum_{k=-L}^{N-1} \tilde{x}_i[m, k] g_{i \rightarrow j}((q - kp)T_s/p - (\tau_i - \tau_j) - m(N + L)(T_{s,i} - T_{s,j})) \right) \\ &\times \exp(j2\pi(f_{0,i} - f_{0,j})T_s(m(N + L) + q/p)). \end{aligned} \quad (4)$$

Since $|f_{0,i} - f_{0,j}|T_s \ll 1$, the phase rotation across the subcarriers of a single OFDM symbol is negligible. Hence, we can drop the dependence on q from the argument of the exponential in (4). This is equivalent to approximating the instantaneous phase of the carrier term in (4) as piecewise constant on each OFDM symbol of duration $(N + L)T_s$, and taking into account only the phase increment in discrete steps from one OFDM symbol to the next. This approximation is the time-domain equivalent of neglecting ICI, which is indeed a valid approximation when the CFO is much smaller than the OFDM subcarrier separation.⁷ Notice also that we make this approximation here for the sake of obtaining a simple model for deriving the proposed CFO/SFO/TO compensation scheme. However, in our simulations we work with the full signal model (4) such that the effect of residual ICI is taken into account.

The received discrete-time m -th signal block after CP removal can thus be expressed by

$$\begin{aligned} y_{i \rightarrow j}[m, q] &\approx \left(\sum_{k=-L}^{N-1} x_i[m, k \bmod N] g_{i \rightarrow j}((q - kp)T_s/p - (\tau_i - \tau_j) - m(N + L)(T_{s,i} - T_{s,j})) \right) \\ &\times \exp(j2\pi(f_{0,i} - f_{0,j})T_s(N + L)m), \quad q = 0, \dots, pN - 1. \end{aligned} \quad (5)$$

The first term in the above product is easily recognized to be the cyclic convolution of the up-sampled sequence $\{x_i[m, k] : k = 0, \dots, N - 1\}$ (with insertion of $p - 1$ zeros in between each sample), with the discrete-time impulse response

$$g_{i \rightarrow j}[m, q] = g_{i \rightarrow j}(qT_s/p - (\tau_i - \tau_j) - m(N + L)(T_{s,i} - T_{s,j})), \quad \text{for } q = 0, \dots, pN - 1. \quad (6)$$

We wish to obtain a frequency-domain representation of this cyclic convolution. To this purpose, notice that $g_{i \rightarrow j}[m, q]$ is obtained by sampling at rate p/T_s the impulse response $g_{i \rightarrow j}(t)$ delayed by $(\tau_i - \tau_j) + m(N + L)(T_{s,i} - T_{s,j})$. Using the well-known *spectral folding* relationship between the continuous-time Fourier transform and the discrete-time Fourier transform of the corresponding sampled signal [32], and

⁷This assumption is validated by the oscillator requirements of the 802.11-2012 standard, which allows 20 ppm frequency error. For a $f_o = 2.45$ GHz system, this corresponds to a maximum CFO value of 98 kHz, which is much less than 312.5 KHz, the subcarrier spacing in the 802.11 standard.

after some algebra we find the discrete-time Fourier transform of $g_{i \rightarrow j}[m, q]$ in the form

$$\begin{aligned} G_{i \rightarrow j}(m, \xi) &= \sum_{q=0}^{pN-1} g_{i \rightarrow j}[m, q] e^{-j2\pi\xi q} \\ &= \frac{p}{T_s} \sum_{\ell} \mathcal{G}_{i \rightarrow j} \left(\frac{\xi - \ell}{T_s/p} \right) \times \exp \left(-j2\pi(\xi - \ell) \frac{(\tau_i - \tau_j) + m(N+L)(T_{s,i} - T_{s,j})}{T_s/p} \right), \end{aligned} \quad (7)$$

where $\mathcal{G}_{i \rightarrow j}(f) = \int g_{i \rightarrow j}(t) e^{-j2\pi f t} dt$ is the continuous-time Fourier transform of $g_{i \rightarrow j}(t)$, and where $\xi \in [-1/2, 1/2]$ and $f \in \mathbb{R}$ denote the frequency variables for discrete-time and continuous-time signals, respectively.

Usually, systems are designed such that $N \gg L \gg 1$ and the receiver sampling frequency p/T_s is large enough in order to avoid significant spectral folding. This means that only the term for $\ell = 0$ is significant in the (discrete-time) frequency interval $\xi \in [-1/(2p), 1/(2p)]$, which is the interval of interest containing the N OFDM subcarriers of the transmitted signal. Hence, after applying the DFT and taking the samples corresponding to the set of discrete frequencies $\{\nu/(pN) : \nu = -N/2, \dots, N/2 - 1\}$, straightforward algebra (omitted for the sake of brevity) yields the OFDM frequency-domain demodulated signal in the form

$$Y_{i \rightarrow j}[m, \nu] = H_{i \rightarrow j}[\nu] X_i[m, \nu] E_{i,j}[m, \nu], \quad (8)$$

where $H_{i \rightarrow j}[\nu] = \frac{p}{T_s} \mathcal{G}_{i \rightarrow j} \left(\frac{\nu}{NT_s} \right)$, and where

$$\begin{aligned} E_{i,j}[m, \nu] &= \exp \left(-j2\pi \left[\frac{\nu}{N} \right] (\mu_i - \mu_j) \right) \\ &\quad \times \exp \left(j2\pi \left[\frac{\nu}{N} \right] (\delta_i - \delta_j) m \right) \\ &\quad \times \exp (j2\pi (\Delta_i - \Delta_j) m), \end{aligned}$$

with $\nu \in \{-N/2, \dots, N/2 - 1\}$ and where we define the normalized TO, SFO and CFO as

$$\mu_i \triangleq f_s \tau_i, \quad \delta_i \triangleq (N+L) T_s \epsilon_i, \quad \Delta_i \triangleq \kappa \delta_i, \quad (9)$$

respectively.

In (8) the term $H_{i \rightarrow j}[\nu] X_i[m, \nu]$ represents what we would expect from an ideal system without synchronization errors while the combined effects of the TO, CFO and SFO are captured by the multiplicative phase rotation term $E_{i,j}[m, \nu]$.⁸

⁸This model has been validated experimentally by the authors using the WARP software radio platform, as documented in [23], [24], and it is found to be accurate within the typical errors of 802.11 legacy APs.

A. Impact of Synchronization Errors on Distributed MU-MIMO

We consider a network formed by UTs $k = 1, \dots, N_u$ served in the DL by APs $i = 1, \dots, N_a$, using distributed MU-MIMO. From the analysis developed above, it is immediate to show that the received signal $1 \times N_u$ vector at the UT receivers, at OFDM symbol m and subcarrier ν , is given by

$$\mathbf{Y}[m, \nu] = \mathbf{X}[m, \nu] \mathbf{\Phi}[m, \nu] \mathbf{H}[m, \nu] \mathbf{\Theta}[m, \nu] + \mathbf{Z}[m, \nu], \quad (10)$$

where $\mathbf{H}[m, \nu]$ is the $N_a \times N_u$ channel matrix with (i, k) -th element $H_{i,k}[m, \nu]$, $\mathbf{X}[m, \nu]$ is the $1 \times N_a$ vector of frequency domain symbols transmitted by the N_a APs simultaneously, and $\mathbf{Z}[m, \nu]$ is the corresponding $1 \times N_u$ vector of Gaussian noise samples, i.i.d. $\sim \mathcal{CN}(0, N_0)$. In general, the MU-MIMO precoded signal vector is given by

$$\mathbf{X}[m, \nu] = \mathbf{U}[m, \nu] \mathbf{G}[m, \nu], \quad (11)$$

where $\mathbf{U}[m, \nu]$ is the $1 \times N_u$ vector of time-frequency encoded data symbols to be sent to the N_u UTs, and $\mathbf{G}[m, \nu]$ is the $N_u \times N_a$ MU-MIMO precoding matrix for OFDM symbol m and subcarrier ν . In this work we consider linear Zero-Forcing Beamforming (ZFBB) [2], [28], [33] and *conjugate beamforming* [1], [2], [28], the details of which are given later on in this section.

The two matrices $\mathbf{\Theta}[m, \nu]$ and $\mathbf{\Phi}[m, \nu]$ are diagonal of dimension $N_u \times N_u$ and $N_a \times N_a$, respectively, with diagonal elements given by⁹

$$\theta_{k,k}[m, \nu] = \exp\left(j2\pi \left[\frac{\nu}{N}\right] \tilde{\mu}_k\right) \exp\left(-j2\pi \left[\frac{\nu}{N}\right] \tilde{\delta}_k m\right) \exp\left(-j2\pi \tilde{\Delta}_k m\right), \quad (12)$$

and

$$\phi_{i,i}[m, \nu] = \exp\left(-j2\pi \left[\frac{\nu}{N}\right] \mu_i\right) \exp\left(j2\pi \left[\frac{\nu}{N}\right] \delta_i m\right) \exp\left(j2\pi \Delta_i m\right). \quad (13)$$

The effect of $\mathbf{\Theta}[m, \nu]$ can be compensated at each UT (see sync blocks at UTs in Fig. 1) by standard pilot-aided or data-aided timing and frequency synchronization techniques suited to OFDM (see [9]–[13]). On the other hand, the presence of $\mathbf{\Phi}[m, \nu]$ in (10) between the precoded transmit vector $\mathbf{X}[m, \nu]$ and the channel matrix $\mathbf{H}[m, \nu]$ yields a degradation of the MU-MIMO precoding performance that cannot be undone by UT processing. As anticipated in Section I, in this paper we focus on the compensation of the term $\mathbf{\Phi}$ at the APs transmitter side. As shown in Fig. 1, our approach includes the use of a central server (CS) module and SYNC blocks, responsible for carrying out the synchronization at the transmitter

⁹The equations derived above apply both when the nodes are APs and UTs. To distinguish APs from UTs we use tildes over the UT variables. In this case $\tilde{\mu}_k, \tilde{\delta}_k$ and $\tilde{\Delta}_k$ denote the normalized TO, SFO and CFO of UT k with respect to the nominal time and frequency references.

side. The MU-MIMO precoding matrix $\mathbf{G}[m, \nu]$ in (11) is constant over time throughout the transmission block (i.e., it is independent of m), and it is calculated from the noisy estimate of the “nominal” channel matrix $\mathbf{H}[0, \nu]$ obtained by the UL pilots at the beginning of the precoding block (here indicated as OFDM symbol $m = 0$). Hence, the precoder is ignorant of the matrix $\Phi[m, \nu]$ unless TO, SFO and CFO are explicitly taken into account. In the case of ZFBBF, we let $\mathbf{G}[m, \nu] = \mathbf{\Lambda}[0, \mu](\mathbf{H}^H[0, \nu]\mathbf{H}[0, \nu])^{-1}\mathbf{H}^H[0, \nu]$ for all m , and in the case of conjugate beamforming we let $\mathbf{G}[m, \nu] = \mathbf{\Gamma}[0, \nu]\mathbf{H}^H[0, \nu]$, where $\mathbf{\Lambda}[0, \mu]$ and $\mathbf{\Gamma}[0, \mu]$ are diagonal matrices that ensure that the rows of $\mathbf{G}[m, \nu]$ have unit norm. Notice also that the ZFBBF precoder requires $N_u \leq N_a$.

In order to motivate the need for the synchronization protocol proposed in this paper, we examine the performance degradation due to typical uncompensated SFO and CFO between the APs. We use lower bounds on the achievable sum-rates as a performance metric, obtained assuming an i.i.d. equal-power Gaussian coding ensemble, where the UTs treat the residual interference due to imperfect zero-forcing as noise.¹⁰ We assume DL blocks of $M = 60$ OFDM symbols, typical of 802.11 [7]. In order to focus on the impact of synchronization errors only, $\mathbf{H}[m, \nu]$ is assumed constant with respect to the time index $m = 0, \dots, M - 1$ over each block and, optimistically, we assume ideal TDD reciprocity and noiseless channel estimation. Provided that the relative TO between APs is within the length of the CP, the terms $\exp(-j2\pi [\frac{\nu}{N}] \mu_i)$ are automatically included as part of the channel estimated from the UL pilot symbols.¹¹ Under these assumptions, the CS computes its MU-MIMO precoding matrix based on the channel matrix $\Phi[0, \nu]\mathbf{H}[0, \nu]$ and uses it throughout the whole DL block comprising M symbols. Therefore, because of the time-varying matrix $\Phi[m, \nu]$, the precoder is more and more mismatched as m increases. Fig. 2 shows the achievable rates obtained by Monte Carlo simulation assuming $\mathbf{H}[0, \nu]$ with i.i.d. elements $\sim \mathcal{CN}(0, 1)$ (normalized independent Rayleigh fading) of a network with $N_a = N_u = 4$, SFO ϵ_i i.i.d. across the users and the DL blocks, uniformly distributed over $[-\epsilon_{\max}, \epsilon_{\max}]$, with $\epsilon_{\max} = 800$ Hz (20 ppm frequency error), and $f_s = 40$ MHz from the 802.11 specifications [7]. The OFDM modulation has parameters $N = 64$ and $L = 16$. The achievable rate

¹⁰Quantifying the system performance in terms of achievable rates is now common practice in modern applied communication works (e.g., see [23]–[25], [28]) and it is much more significant than evaluating traditional performance metrics such as self-interference variance, estimation MSE, or uncoded BER, for which we are typically unable to quantify their impact on the ultimate system performance, which makes use of powerful modulation and coding schemes.

¹¹It can be readily shown that it takes about 1 second for two completely uncompensated clocks to fall out of sync by more than the CP, if they are off by worst case frequency offsets dictated by 802.11. Hence, synchronization performed at intervals 1 second apart or faster suffices.

shown in Fig. 2 is the average rate across a frame of 60 OFDM symbols. Since the system loses synchronism progressively across each block, the achievable rate rapidly degrades as the OFDM symbol m increases, such that the average performance is severely degraded. As a comparison, we also show the performance of an ideally synchronized system (zero CFO/SFO) and the performance of a MISO cooperative beamforming scheme employing conjugate beamforming (denoted by Conj BF in the figure) to a single user [19], [28], with and without frequency offsets. Notice that MISO cooperative beamforming suffers much less from the lack of synchronization, and significantly outperforms the mismatched MU-MIMO ZFBF at high SNR. Thus the motivation of this work is to provide the APs with sufficiently accurate estimation of their SFO/CFO such that the large spectral efficiency promised by distributed MU-MIMO is effectively realized in practice.

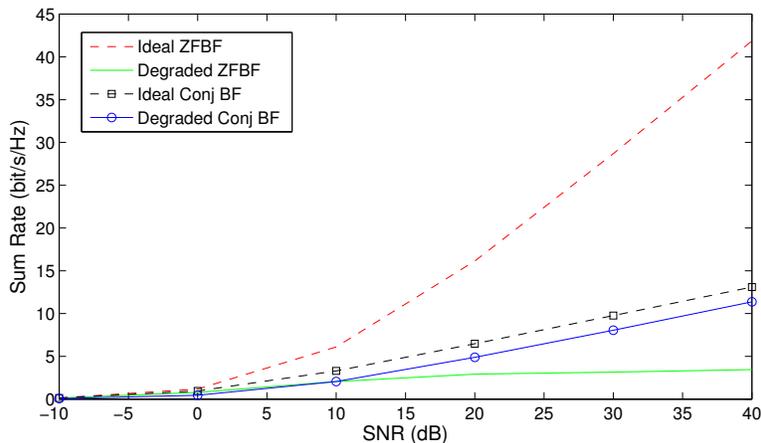


Fig. 2. Achievable average rates for a 4×4 distributed MIMO systems in ideal synchronization conditions and degraded by synchronization impairments (uncompensated free running oscillators).

III. SYSTEM ARCHITECTURE

At the network setup phase, each AP discovers its neighbor APs and measures the associated propagation delays. Since the APs are placed in fixed positions, the network structure (neighboring APs) and path delays change very slowly in time and it is sufficient to repeat such procedure at a very slow time scale (e.g., of the order of tens of minutes or slower) to safely assume that these parameters are known.

Calibration for TDD reciprocity (Section V) depends on the complex scaling factors introduced by the modulation/demodulation hardware (amplifiers, filters, ADC/DACs), that change slowly in time. Experimental SDR implementation shows that calibration must be repeated on a time-scale on the order

of 10 min [28] for a system where all of the APs share a common clock. In contrast, synchronization must track the variations of the clock frequencies. While in this paper we model the frequency errors $\{\epsilon_i\}$ as unknown constants, in practice they depend on temperature, and fluctuate around their nominal 0 value with a coherence time on the order of 1s. Hence synchronization must be performed frequently enough such that the errors do not significantly degrade the system throughput.

In a distributed MU-MIMO system where APs do not share a common clock, the very small residual frequency errors due to non-ideal synchronization cause slow relative drift among the different AP phasors. As these uncompensated phase rotations accumulate over time, and since they affect the extent of end-to-end reciprocity in the channel, calibration needs to be performed at the same rate as synchronization. The proposed protocol makes use of a frame structure comprised of synchronization, calibration and DL data slots, as shown in Fig. 3 where synchronization and calibration take place back to back.

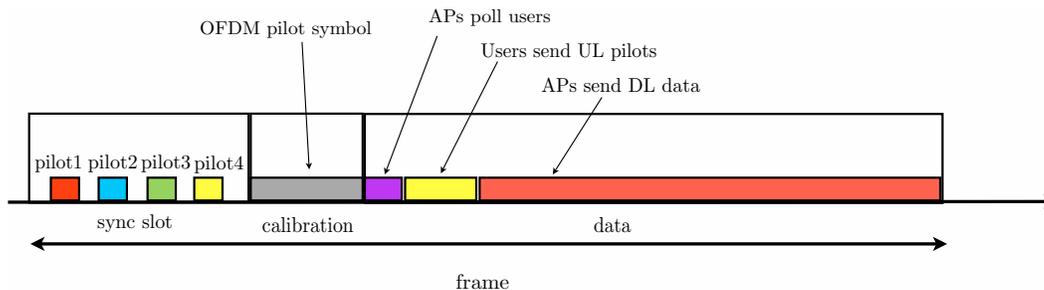


Fig. 3. Concept of the frame structure for the proposed distributed MU-MIMO architecture. The synchronization slot contains several pilot signals, orthogonal in time, and separated by guard intervals. The pilot signals are assigned to the anchor APs according to a graph coloring scheme, similarly to what is done with frequency reuse in cellular systems. In this example, we have 4 orthogonal pilot signals. The calibration slot comprises standard OFDM pilot symbols. During this slot, each AP transmits pilots over one or more time-frequency slots, allowing it to calibrate itself with other APs in its vicinity. Again, pilot patterns orthogonal in time-frequency are allocated to the APs and reused across the network according to a graph coloring scheme. The relative size of the sync and calibration slots with respect to the data slots in this figure do not reflect the actual duration in a real implementation. As a matter of fact, the protocol overhead is very small for typical channel coherence times in a WLAN nomadic users scenario.

Let $\mathcal{T} = \{i : i = 1, \dots, N_a\}$ denote the set of APs. We define a directed network graph with vertices \mathcal{T} and edges $\mathcal{E} = \{(i, j)\}$ for all ordered pairs of APs i, j such that the *average* SNR of the transmission $i \rightarrow j$ is larger than some suitably defined threshold. Since the distance-dependent pathloss is symmetric with respect to the AP indices, if $(i, j) \in \mathcal{E}$ then also $(j, i) \in \mathcal{E}$. Without loss of generality, we assume

that the graph $(\mathcal{T}, \mathcal{E})$ is connected.¹²

For the synchronization protocol, we define $\mathcal{A} \subseteq \mathcal{T}$ to form a *connected cover* [34], i.e., the subgraph formed by the APs $i \in \mathcal{A}$ and their associated edges $\mathcal{E}(\mathcal{A}) = \{(i, j) : i, j \in \mathcal{A}\}$ is connected and any other AP $i' \notin \mathcal{A}$ has at least one neighbor in \mathcal{A} . The APs in \mathcal{A} are referred to as “anchor” nodes. Fig. 4 shows a network with 6 anchor nodes. In a practical system deployment, we may imagine that each cluster corresponds to a geographically isolated area (e.g., a building) and the anchor nodes are placed in special positions (e.g., roof-top), such that they can exploit stronger propagation conditions (e.g., line of sight) to each other. During the network deployment phase, the network designer or some self-organizing scheme should find a good set of anchor APs. Such optimization, however, is beyond the scope of this paper and is left for future investigation.

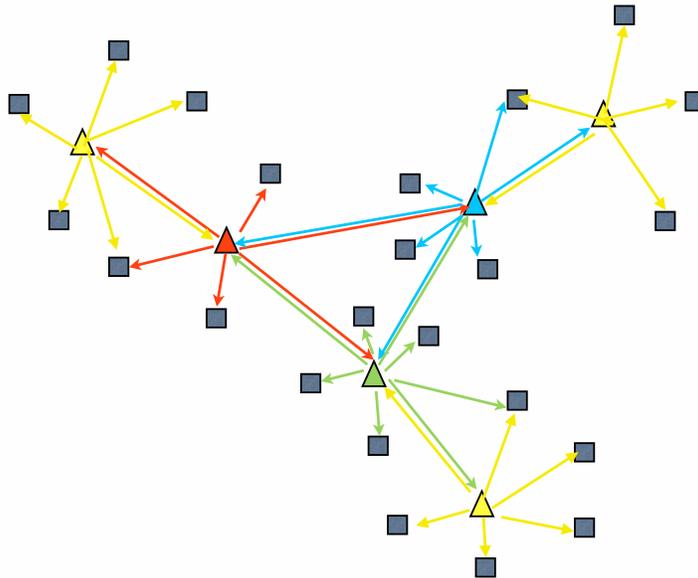


Fig. 4. Network graph with 6 anchor nodes (triangles) reusing 4 orthogonal pilot bursts (identified by different colors, as in Fig. 3). The non-anchor APs do not transmit sync pilots, and listen to their neighbor anchor node(s) pilot burst in order to correct their timing and frequency reference.

The synchronization slots are formed by a number of *pilot bursts* (sequences of time-domain chips) separated by guard intervals. We assume that a rough timing synchronization is provided through the wired backhaul network using a coarse frame synchronization protocol such as IEEE 1588 [4]. As a

¹²Otherwise, the same algorithms can be applied independently to each connected component.

result, the APs know where to look for the pilot bursts up to an accuracy of a few OFDM symbols, corresponding to the guard intervals inserted between the pilot bursts. Pilot bursts are assigned to the anchor nodes. We require that each anchor node is able to receive the pilot bursts from its neighboring anchor nodes without significant co-pilot interference. This is possible by finding an $L(1, 1)$ -labeling (see [35] and references therein) of the subgraph $(\mathcal{A}, \mathcal{E}(\mathcal{A}))$, and associating pilot bursts to the labeling colors. An $L(1, 1)$ -labeling of a graph is a coloring scheme for which all neighbors of any node $i \in \mathcal{A}$ have distinct colors. Using a number of orthogonal pilot bursts larger or equal to the number of colors of the $L(1, 1)$ -labeling, guarantees that each anchor AP is able to receive pilots from its neighbors without pilot collisions. Notice that the orthogonal pilots are re-used across the network. For example, the 6 anchor nodes of Fig. 4 require only 4 orthogonal pilots. More generally, when $(\mathcal{A}, \mathcal{E}(\mathcal{A}))$ is a cycle, if $|\mathcal{A}|$ is a multiple of 3, then 3 colors are sufficient, while if $|\mathcal{A}|$ is not multiple of 3 then 4 colors are needed. If $(\mathcal{A}, \mathcal{E}(\mathcal{A}))$ is a clique (fully connected graph), then $|\mathcal{A}|$ is the necessary and sufficient number of colors for an $L(1, 1)$ labeling, and if $(\mathcal{A}, \mathcal{E}(\mathcal{A}))$ is a tree, then it is known that $d_{\max} + 1$ colors are necessary and sufficient, where d_{\max} is the maximum node degree. The non-anchor APs in $\mathcal{T} - \mathcal{A}$ do not transmit pilot bursts, and just listen to the anchor nodes. Eventually, at the end of a synchronization slot, each anchor node has received a distinct and non-overlapping pilot burst from each of its neighbors in $(\mathcal{A}, \mathcal{E}(\mathcal{A}))$, and all non-anchor nodes have received one or more pilots from their neighboring anchor nodes. The estimates collected by the nodes are sent through the wired backhaul network to the CS, which computes correction factors as explained in Section IV.

The calibration slot is organized in a similar way, with the only difference being that in this case *all* nodes have to send and receive a pilot signal, as explained in Section V. If the network is formed by sufficiently isolated clusters, calibration can be organized in a hierarchical way such that all clusters calibrate in parallel, and then a second round of calibration is run between the anchor nodes (one per cluster) in order to calibrate the whole network (see [8]). For example, such hierarchical schemes prove to be very useful in the presence of multiple antenna APs, where the groups of antennas of each AP, driven by a common clock, can be calibrated “locally” and in parallel using the scheme of [28], and a second level of calibration is used to set the relative calibration coefficients between the APs. This again makes the proposed scheme to work in even very large networks without the requirement of a “master” node or calibration antenna connected with high SNR to all other nodes or antennas.

Thanks to the periodically repeated computation of the timing, frequency and calibration correction factors, the APs are time/frequency synchronous and calibrated and can operate (up to residual errors) as a large distributed and coherent multi-antenna system. Both synchronization and calibration must remain

sufficiently accurate over the duration of the data transmission slot in the frame, and until the next synchronization and calibration slot (see Fig. 3).

Notice that for network graphs of bounded degree (the typical case in wireless networks with propagation pathloss) the number of orthogonal pilots does not need to grow with the network size. In fact, the orthogonal pilot reuse of the proposed protocol is analogous to frequency reuse in cellular networks, where a finite (typically small) number of subchannels are reused throughout the network irrespectively of the number of cells, which can be arbitrarily large. Hence, the proposed synchronization and calibration protocols are *scalable*, in the sense that the pilot overhead is *independent of the network size*.

The data slots can be arranged in the classical way widely proposed in TDD-based MU-MIMO literature [1], [2], [23], [24], [28], [33]: the APs send a request for UL pilots. The polled users respond with their mutually orthogonal UL pilots. Then, the APs send the received UL pilot burst to the CS which estimates the UL channel, computes the corresponding DL MU-MIMO precoding matrix and sends the precoding coefficients (columns of the precoding matrix) along with the encoded data packets back to the APs for transmission. Finally, the APs pre-code the user data packets according to (11), and compensate for the TDD calibration factors (see Section V) and phase rotation factors due to timing and frequency offsets (see Section IV) before joint DL transmission.

IV. SYNCHRONIZATION

In order to overcome the impairments discussed in Section II, we observe that the effects of the error term (13) can be undone by each AP i transmitter by multiplying the transmitted frequency domain symbols $X_i[m, \nu]$ by the time and frequency dependent phase rotation factor

$$P_i[m, \nu] = \exp\left(-j2\pi\left(1 + \frac{1}{\kappa}\left[\frac{\nu}{N}\right]\right)\Delta_i^{corr}m\right), \quad (14)$$

and by adjusting the timing reference (origin of the transmitter time axis) by μ_i^{corr} . The timing and frequency corrections $\Delta_i^{corr}, \mu_i^{corr}$ are provided by the CS through the low-complexity centralized computation illustrated in the rest of this section.

We should also notice that a symbol phase rotation by $P_i^*[m, \nu]$, i.e., the complex conjugate of the term in (14), must be applied to signals *received* at the APs during the UL training from the users, in the data slots. In this way, all APs use the same time and frequency reference for the calculation of the DL channel matrix for MU-MIMO precoding. Performing such corrections during UL training and DL transmission is consistent with the diagram in Fig. 1, showing SYNC block modules in both UL and DL AP/CS processing.

In addition to the baseband symbol rotation as modeled in Section II, the SFO produces a contraction or dilation of the time axis of each AP. While this is a negligible effect over a few tens of OFDM symbols (typical duration of a data slot), it accumulates over the slots such that at some point the OFDM symbol misalignment between different APs becomes larger than the OFDM CP, thus producing inter-block interference and not only relative phase rotations. However, for typical values of SFO consistent with the 802.11 standard specifications, we have checked that the duration for which this symbol misalignment becomes significant with respect to the OFDM CP length is much larger than the frame length at which the synchronization protocol must be repeated. Therefore, thanks to the TO correction, the timing axis relative shift is reset at each frame and the SFO can be effectively compensated by the proposed baseband symbol rotation scheme.

Each anchor node collects pilot bursts from all its neighboring anchor nodes and produces estimates of the TO and CFO with respect to its neighbors. While any suitable estimation scheme can be used for this purpose, in Appendix A we derive the joint ML timing and frequency estimator and of the corresponding CRB in the case, relevant to our scenario, where the pilot signal is observed through a multipath channel with known path delays and unknown path coefficients.

In this section we focus on the centralized estimation of the correction factors Δ_i^{corr} and μ_i^{corr} from the noisy CFO and TO estimates

$$\widehat{\delta\Delta}_{j \rightarrow i} = \Delta_j - \Delta_i + z_{j \rightarrow i}, \quad (15)$$

and

$$\widehat{\delta\mu}_{j \rightarrow i} = \mu_j - \mu_i + w_{j \rightarrow i}, \quad (16)$$

for all $i \in \mathcal{A}$ and for all $j \in \mathcal{A}(i)$,¹³ where $z_{j \rightarrow i}$ and $w_{j \rightarrow i}$ are estimation errors with variance equal to the corresponding estimation MSE. As already noted, if $j \in \mathcal{A}(i)$ then $i \in \mathcal{A}(j)$. Therefore, these measurements always come in pairs.

The CS collects all these estimates and computes the correction terms $\Delta_i^{corr}, \mu_i^{corr}$. In the absence of estimation errors we would have $\widehat{\delta\Delta}_{j \rightarrow i} = \Delta_j - \Delta_i$. Hence, a reasonable approach consists of minimizing the *weighted LS* cost function

$$J_{\Delta}(\Delta_1, \dots, \Delta_{|\mathcal{A}|}) = \sum_{(i,j) \in \mathcal{E}(\mathcal{A})} \beta_{j,i} \left(\widehat{\delta\Delta}_{j \rightarrow i} - (\Delta_j - \Delta_i) \right)^2, \quad (17)$$

¹³In the following, $\mathcal{A}(i)$ denotes the neighborhood of AP i in the subgraph $(\mathcal{A}, \mathcal{E}(\mathcal{A}))$.

where $\beta_{j,i} = 1/\mathbb{E}[|z_{j \rightarrow i}|^2]$ is the inverse estimation MSE for the link $j \rightarrow i$. Notice that if the estimation errors $z_{j \rightarrow i}$ are Gaussian i.i.d., then the minimizer of (17) coincides with the joint ML estimator based on the observations (15).

Taking the partial derivative with respect to Δ_i , and setting it equal to zero, we obtain the set of linear equations

$$\sum_{j \in \mathcal{A}(i)} (\beta_{i,j} + \beta_{j,i}) (\Delta_i - \Delta_j) = \sum_{j \in \mathcal{A}(i)} \left(\beta_{i,j} \widehat{\delta \Delta}_{i \rightarrow j} - \beta_{j,i} \widehat{\delta \Delta}_{j \rightarrow i} \right), \quad (18)$$

that can be written in matrix form as $\mathbf{L}\mathbf{\Delta} = \mathbf{u}$, where $\mathbf{\Delta} = (\Delta_1, \dots, \Delta_{|\mathcal{A}|})^\top$, \mathbf{u} is the $|\mathcal{A}| \times 1$ vector with i -th element $\sum_{j \in \mathcal{A}(i)} \left(\beta_{i,j} \widehat{\delta \Delta}_{i \rightarrow j} - \beta_{j,i} \widehat{\delta \Delta}_{j \rightarrow i} \right)$, and \mathbf{L} is the $|\mathcal{A}| \times |\mathcal{A}|$ matrix with (i, j) -th elements

$$[\mathbf{L}]_{i,j} = \begin{cases} \sum_{j \in \mathcal{A}(i)} (\beta_{i,j} + \beta_{j,i}) & \text{for } j = i \\ -(\beta_{i,j} + \beta_{j,i}) & \text{for } j \neq i \end{cases} \quad (19)$$

In general, the actual estimation MSE may not be known. Hence, the CRB developed in Appendix A can be used to approximate the coefficients $\beta_{i,j}$, motivated by the fact that in the high SNR range of interest for WLAN applications, good estimators yield MSE essentially identical to the CRB (see results in Appendix A). A simple alternative consists of using $\beta_{i,j} = \text{const.}$ for all i, j , that is, using LS instead of weighted LS. In particular, by letting $\beta_{i,j} = 1/2$ for all $(i, j) \in \mathcal{E}(\mathcal{A})$ we have that \mathbf{L} is the Laplacian matrix of the connectivity graph of $(\mathcal{A}, \mathcal{E}(\mathcal{A}))$.¹⁴

It is immediate to see that \mathbf{L} has rank $|\mathcal{A}| - 1$ and the system of equations (18) is under-determined. In fact, $\mathbf{L}\mathbf{1} = \mathbf{0}$ ($\mathbf{1}$ being the all-one vector). This is to be expected, since only frequency differences can be observed. In order to find a solution, we choose a reference anchor AP, without loss of generality indexed by 1, and define the frequency difference vector $\boldsymbol{\xi}$ with elements $\xi_i = \Delta_i - \Delta_1$ for all $i \in \mathcal{A}$ with $i \neq 1$. With this substitution, we obtain the system of equations $\mathbf{L}_1 \boldsymbol{\xi} = \mathbf{u}$, where \mathbf{L}_1 is obtained from \mathbf{L} by eliminating the first column, yielding the canonical LS solution $\boldsymbol{\xi} = (\mathbf{L}_1^\top \mathbf{L}_1)^{-1} \mathbf{L}_1^\top \mathbf{u}$. The resulting correction factor Δ_i^{corr} is equal to 0 for $i = 1$ (reference anchor AP) and equal to ξ_i for $i \neq 1$.

The TO estimation problem is completely analogous, by defining the corresponding weighted LS cost function $\sum_{(j,i) \in \mathcal{E}(\mathcal{A})} \gamma_{j,i} \left(\widehat{\delta \mu}_{j \rightarrow i} - (\mu_j - \mu_i) \right)^2$, with $\gamma_{j,i} = 1/\mathbb{E}[|w_{j \rightarrow i}|^2]$. The details are omitted for brevity.

¹⁴The connectivity graph of the directed graph $(\mathcal{A}, \mathcal{E}(\mathcal{A}))$ is the undirected graph with nodes \mathcal{A} and edges connecting i, j whenever $(i, j), (j, i) \in \mathcal{E}(\mathcal{A})$. The Laplacian matrix is the $|\mathcal{A}| \times |\mathcal{A}|$ matrix with elements -1 in positions corresponding to i, j whenever these nodes are connected, element $|\mathcal{A}(i)|$ (the degree of i) in the diagonal position corresponding to node i and zero elsewhere.

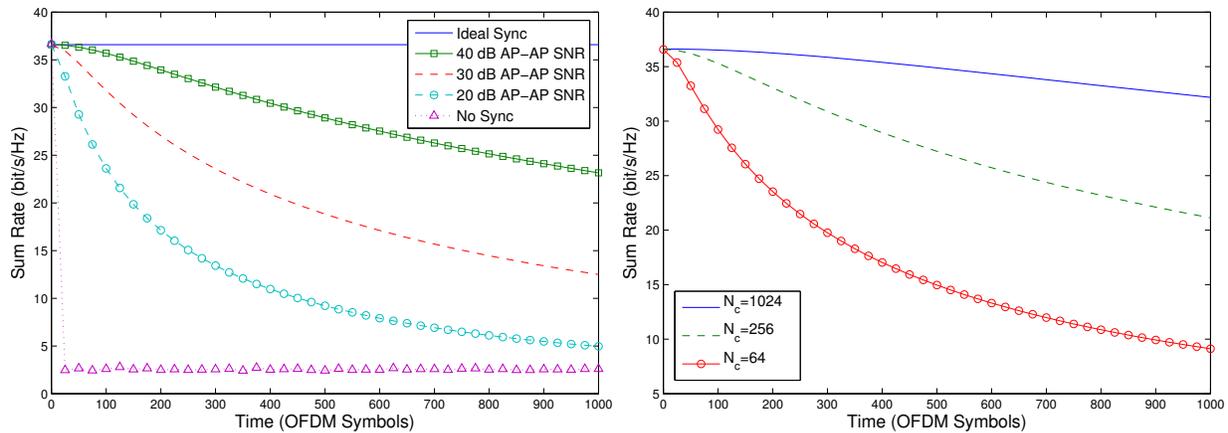


Fig. 5. The achievable rates versus time of a 4×4 system after synchronization and compensation, as a function of the OFDM index in the data slot. The left figure assumes $N_c = 256$, and the right figure assumes the AP-AP SNR is 30 dB.

A. Numerical Results

In order to evaluate the impact of TO, CFO and SFO compensation on the performance of distributed MU-MIMO, we examine the example of Section II-A under the proposed synchronization/compensation scheme. In this example there are 4 single-antenna APs serving 4 UTs. For this small network, we assume that all 4 APs are anchor nodes, in line of sight of each other (the channel is formed by a single path with an unknown but constant channel coefficient), with the same AP-to-AP SNR. We perform timing and frequency joint ML estimation, as derived in Appendix A to estimate the offsets for each pair of APs, and then use LS-based synchronization on the fully connected anchor graph. For simplicity, in all the numerical results of this paper we have used LS with uniform weights, assuming (conservatively) that the CS does not know the actual estimation MSEs.

The performance of the TO/CFO estimators between any pair of APs in this example are given in Fig. 9-b) in Appendix A, showing the MSE vs. SNR relationship of the proposed joint timing and frequency ML estimator along with the associated CRB. We notice that even for fairly low SNR this estimator performs very close to the CRB, which decreases inversely proportional to the receiver SNR.

Fig. 5 shows the achievable rate versus the index of the OFDM symbol for the case where the channel matrix is perfectly known at time $m = 0$ and, as explained before, the ZFBF precoder is calculated at time 0 and kept constant throughout the block.¹⁵ While the results of Fig. 2 are obtained for free-running

¹⁵Achievable rates are calculated assuming perfect calibration in order to isolate the effects of synchronization.

oscillators, here we assume that at time $m = 0$ the CFO is estimated according to the proposed LS scheme and the phase rotation factor (14) is applied throughout the sequence of OFDM symbols. We also assume that the SNR for the AP-to-UT channels is constant at 30 dB. For the pilot bursts, we used a structure formed by two repetitions of the same sequence of $N_c/4$ random QPSK time-domain chips, separated by $N_c/2$ zeros, for a total sequence length of N_c pilot chips (see also the example in Appendix A).

In Fig. 5 (left), we held the sequence length constant at $N_c = 256$ and vary the AP-to-AP SNR. Consistently with Fig. 9-b), the higher the SNR, the lower the estimation MSE and, as a consequence, the residual CFO is smaller and produces a less evident mismatch of the MU-MIMO precoder throughout the data block in Fig. 5 (left). In Fig. 5 (right), we held the AP-to-AP SNR constant at 30 dB while varying N_c . We observe that with $N_c = 1024$ (roughly corresponding to 13 OFDM symbols including the cyclic prefix) we can afford to send about 1000 OFDM symbols with small degradation. Considering data slots of ~ 100 OFDM symbols, we can afford 10 DL MU-MIMO precoded data slots in between each synchronization slot. Performance can be improved further by implementing a smoothing filter over time (e.g., Kalman filtering) in order to track the timing and frequency corrections factors through the sequence of frames instead of re-estimating anew at each frame. However, this requires modeling the variations of the clock frequency errors ϵ_i over time (e.g., a Gauss-Markov model) and goes beyond the scope of this paper.

V. CALIBRATION FOR TDD RECIPROCITY

In this section we present a novel TDD *relative* calibration technique that generalizes the approach of [28] to the case of an arbitrary distributed network topology. We re-consider the channel model (10) assuming that, for clarity of exposition, synchronization is perfect, i.e., the residual TO, CFO and SFO are equal to zero. Then, we focus on the effects of transmit/receive hardware mismatch on end-to-end channel reciprocity, and devise a scheme able to compensate for such mismatches. We remark here that while a centrally clocked MU-MIMO architecture (such as [28]) needs to perform calibration at a very low rate (e.g., one calibration round every 10 min), due to the inherent stability of the complex scaling introduced by the transmit/receive hardware at each AP, in the case of a distributed MU-MIMO architecture as considered in this paper the residual CFO after non-ideal compensation yields a phase error that accumulates over the OFDM symbols across the data blocks. Hence, calibration must be repeated together with synchronization, as evidenced by the frame structure of Fig. 3. The calibration slot is formed by pilot symbols arranged on the time-frequency plane induced by OFDM, as commonly

implemented for channel estimation purposes [7], [36], [37]. Since the non-reciprocal elements of the channel (due to the modulation/demodulation hardware) are smooth over the signal bandwidth both in amplitude and in phase, the calibration pilots can be arranged efficiently on the time-frequency plane such that a large number of mutually orthogonal pilots can be exchanged in a short slot. Furthermore, exploiting the correlation between the calibration coefficients at different subcarriers can be exploited to improve estimation accuracy. More details on efficient calibration pilot design are provided in [38], [39] and go beyond the scope of this work. Here, for the sake of clarity, we focus on a single subcarrier and drop the subcarrier index ν from the notation.

In order to motivate the need for TDD reciprocity calibration, let's consider a scheme where the DL channel is learned at the APs side through UL pilot symbols sent by the users (with reference to the frame structure of Fig. 3). In general, the signal sent from the APs to the UTs over the DL data slot at a given subcarrier and OFDM symbol takes on the form

$$\mathbf{Y} = \mathbf{X}\mathbf{H} + \mathbf{Z}, \quad (20)$$

where the DL channel matrix is given by $\mathbf{H} = \mathbf{T}\mathbf{B}\tilde{\mathbf{R}}$, where $\tilde{\mathbf{R}} = \text{diag}(\tilde{R}_1, \dots, \tilde{R}_{N_u})$ contains the complex coefficients introduced by the UT receiver hardware, $\mathbf{T} = \text{diag}(T_1, \dots, T_{N_a})$ contains analogous coefficients introduced by the APs transmitter hardware, and \mathbf{B} is the physical channel matrix, i.e., the channel coefficients due solely to the antenna-to-antenna propagation.

At each given frame, the UTs that have been polled, send their UL pilots for channel estimation on the UL pilot slot (see Fig. 3). The signal received by the APs is given by

$$\mathbf{Y}^{up} = \mathbf{H}^{up}\tilde{\mathbf{X}} + \mathbf{Z}^{up}, \quad (21)$$

where \mathbf{Z}^{up} is the UL Gaussian noise vector, $\tilde{\mathbf{X}}$ is a $N_u \times N_u$ unitary matrix of frequency domain UL pilot symbols, and where $\mathbf{H}^{up} = \mathbf{R}\mathbf{B}\tilde{\mathbf{T}}$ with $\tilde{\mathbf{T}} = \text{diag}(\tilde{T}_1, \dots, \tilde{T}_{N_u})$ containing the UTs transmitter coefficients, and $\mathbf{R} = \text{diag}(R_1, \dots, R_{N_a})$ containing the APs receiver coefficients.

Due to TDD (uplink and downlink are in the same frequency band) and to the reciprocity of the physical propagation channel, the same \mathbf{B} appears in both the UL pilot slot and the DL data slot [28].

Uplink channel estimation using the UL pilot slot \mathbf{Y}^{up} can be written as $\hat{\mathbf{H}}^{up} = \mathbf{Y}^{up}\tilde{\mathbf{X}}^H = \mathbf{H}^{up} + \tilde{\mathbf{Z}}^{up}$, with $\tilde{\mathbf{Z}}^{up} = \mathbf{Z}^{up}\tilde{\mathbf{X}}^H$. Neglecting the estimation error $\tilde{\mathbf{Z}}^{up}$ for the time being, we notice that if the CS computes the downlink MU-MIMO precoder assuming $\hat{\mathbf{H}}^{up}$ as the actual *downlink* channel, then this precoder will be mismatched with respect to the true downlink channel \mathbf{H} . For example, with ZFBF precoding, the MU-MIMO precoding matrix obtained from the UL pilots is given by the two-normalized

pseudo-inverse $\mathbf{G} = \mathbf{\Lambda} [(\mathbf{H}^{up})^H \mathbf{H}^{up}]^{-1} (\mathbf{H}^{up})^H$ (see Section II-A). Clearly, $\mathbf{G}\mathbf{H}$ is not a diagonal matrix unless $\mathbf{H} = \mathbf{H}^{up} \mathbf{A}$ for some invertible diagonal matrix \mathbf{A} .

In order to turn \mathbf{H} into a right diagonally scaled version of \mathbf{H}^{up} , each AP i multiplies its transmit signal by the coefficient $\alpha R_i/T_i$, for some $\alpha \neq 0$ (this is the operation indicated by the CAL block in Fig. 1). Using matrix notation, it can be easily checked that this calibration operation corresponds to

$$\alpha \mathbf{R}\mathbf{T}^{-1}\mathbf{H} = \alpha \mathbf{R}\mathbf{B}\tilde{\mathbf{R}} = \alpha \mathbf{H}^{up}\tilde{\mathbf{T}}^{-1}\tilde{\mathbf{R}}, \quad (22)$$

where the diagonal matrix $\mathbf{A} = \alpha \tilde{\mathbf{T}}^{-1}\tilde{\mathbf{R}}$ contains the scalar amplitude and phase shifts that are estimated at each UT receiver as part of their DL channels. It is therefore apparent that the calibration protocol consists of efficiently estimating the TDD calibration matrix $\alpha \mathbf{R}\mathbf{T}^{-1}$, defined up to an arbitrary multiplicative non-zero factor α . As for the synchronization protocol of Section IV, also in this case we devise a scheme that involves only the exchange of calibration pilots between APs, i.e., our scheme does not assume the participation of the UTs, which may be legacy devices not designed to handle distributed MU-MIMO.

During the calibration slot, pilot symbols are exchanged by the APs over a connected spanning subgraph $(\mathcal{T}, \mathcal{F})$ of the network graph $(\mathcal{T}, \mathcal{E})$ where the subset of links \mathcal{F} is such that if $(i, j) \in \mathcal{F}$ then also $(j, i) \in \mathcal{F}$. The optimization of the subgraph $(\mathcal{T}, \mathcal{F})$ is an interesting topic for future work.

Let $(i, j) \in \mathcal{F}$. Then, after a calibration slot, AP i gathers the observation

$$Y_{j \rightarrow i} = T_j B_{j \rightarrow i} R_i + Z_{j \rightarrow i}, \quad (23)$$

and AP j gathers the observation

$$Y_{i \rightarrow j} = T_i B_{i \rightarrow j} R_j + Z_{i \rightarrow j} \quad (24)$$

where we have assumed without loss of generality that the calibration pilot transmitted by the APs is equal to one. We can write

$$\begin{aligned} \begin{bmatrix} Y_{j \rightarrow i} \\ Y_{i \rightarrow j} \end{bmatrix} &= \begin{bmatrix} T_j R_i \\ T_i R_j \end{bmatrix} B_{i \rightarrow j} + \begin{bmatrix} Z_{j \rightarrow i} \\ Z_{i \rightarrow j} \end{bmatrix} \\ &= \begin{bmatrix} c_i \\ c_j \end{bmatrix} \rho_{i,j} + \begin{bmatrix} Z_{j \rightarrow i} \\ Z_{i \rightarrow j} \end{bmatrix}, \end{aligned} \quad (25)$$

owing to the fact that, by the physical channel reciprocity, $B_{i \rightarrow j} = B_{j \rightarrow i}$, and defining $\rho_{i,j} = \rho_{j,i} = T_i T_j B_{i \rightarrow j}$ and $c_i = R_i/T_i$. Our goal is to estimate the relative calibration coefficients c_i for $i =$

$1, \dots, N_a$, up to a common multiplicative non-zero constant α . Without loss of generality we assume that $\{c_i\}$ is a set of non-zero bounded complex numbers.¹⁶

If the observations $Y_{j \rightarrow i}, Y_{i \rightarrow j}$ were noiseless, we would have the equality $c_j Y_{j \rightarrow i} = c_i Y_{i \rightarrow j}$ for all $(i, j) \in \mathcal{F}_u$, where \mathcal{F}_u is the set of undirected edges corresponding to \mathcal{F} . Hence, in the presence of observation noise it makes sense to define the LS cost function

$$J_{cal}(c_1, c_2, \dots, c_{N_a}) = \sum_{(i,j) \in \mathcal{F}_u} |c_j Y_{j \rightarrow i} - c_i Y_{i \rightarrow j}|^2. \quad (26)$$

Letting $\mathbf{c} = (c_1, c_2, \dots, c_{N_a})$, notice that the LS objective function satisfies the scaling property $J_{cal}(\alpha \mathbf{c}) = |\alpha|^2 J_{cal}(\mathbf{c})$. In fact, given a vector of calibration coefficients \mathbf{c} , any scaled vector $\alpha \mathbf{c}$ for $\alpha \neq 0$ provides equally good calibration coefficients. In order to eliminate this ambiguity and at the same time exclude the all-zero solution in the minimization of (26), we impose the constraint $\|\mathbf{c}\|^2 = 1$. The Lagrangian function of the constrained LS problem is given by

$$F_{cal}(\mathbf{c}, \lambda) = J_{cal}(\mathbf{c}) - \lambda(\|\mathbf{c}\|^2 - 1). \quad (27)$$

where λ is the Lagrangian multiplier. Differentiating F_{cal} with respect to c_i^* , treating c_i and c_i^* as if they were independent variables [40], and then setting the partial derivatives to zero, we obtain

$$\frac{\partial}{\partial c_i^*} F_{cal}(c_1, c_2, \dots, c_{N_a}, \lambda) = \sum_{j: (i,j) \in \mathcal{F}_u} (c_i |Y_{i \rightarrow j}|^2 - c_j Y_{i \rightarrow j}^* Y_{j \rightarrow i}) - \lambda c_i. \quad (28)$$

In matrix form, we obtain $\mathbf{A} \mathbf{c} = \lambda \mathbf{c}$, where \mathbf{A} is the $N_a \times N_a$ matrix with (i, j) -th elements

$$[\mathbf{A}]_{i,j} = \begin{cases} \sum_{j: (i,j) \in \mathcal{F}_u} |Y_{i \rightarrow j}|^2 & \text{for } j = i \\ -Y_{i \rightarrow j}^* Y_{j \rightarrow i} & \text{for } j \neq i, (i, j) \in \mathcal{F}_u \\ 0 & \text{for } j \neq i, (i, j) \notin \mathcal{F}_u \end{cases} \quad (29)$$

Finally, we notice that the sought constrained LS solution \mathbf{c} is a unit-norm eigenvector associated to the eigenvalue of \mathbf{A} with the smallest magnitude. In fact, this yields the direction in the domain \mathbb{C}^{N_a} of $J_{cal}(\mathbf{c})$ with slowest growth, such that the value of $J_{cal}(\mathbf{c})$ is minimized at the intersection of the unit circle $\|\mathbf{c}\|^2 = 1$ and such eigenspace of \mathbf{A} .

It is interesting to notice the difference between the proposed method and the scheme in [28]. When the graph $(\mathcal{T}, \mathcal{F})$ is a star with (say) AP 1 at the center, [28] proposes to estimate the calibration coefficients as $c_j = \frac{Y_{1 \rightarrow j}}{Y_{j \rightarrow 1}} c_1$ for some arbitrary unit-magnitude c_1 . This is equivalent to minimize the objective function (26) subject to the constraint $|c_1| = 1$. The problem with this method is that the c_j 's are given by the

¹⁶If $c_i = 0$ or $1/c_i = 0$, the i -th node can be omitted as it is a ‘‘non-communicating’’ node.

ratio of two noisy pilot symbols. In a distributed topology when SNR to the center star may not be large, this yields an ill-behaved estimation where some coefficients may be near zero and others be very large. When using these coefficients in the MU-MIMO beamforming scheme, subject to a per-AP power constraint, the system performance incurs significant degradation unless all APs have very good SNR to the star center.

A final remark about scalability is in order. For general network graphs, the computational complexity of LS for both synchronization and TDD calibration is polynomial in the number of APs. For certain network graphs this complexity can be linear (e.g., in the star topology of the Argos scheme [28]). In any case, the computation involved in solving the LS estimation problems is easily affordable for network of practical size (up to a few tens of APs). It is also interesting to remark that the data exchange over the wired backhaul network required by the proposed protocol is also easily affordable. For example, suppose that the timing and frequency measurements (real numbers) and the calibration pilot symbol (complex numbers) are represented with 16 bits per real coefficient. This requires 64 bits per frame per AP (roughly speaking). For frames of 10 ms, this is equivalent to 6.4 kbit/s per AP of protocol data overhead. For a 1 Gb Ethernet backbone, as it would be meaningful for a distributed MU-MIMO system, this overhead is less than 5 orders of magnitude less than the backhaul capacity. Therefore, a system with 100 APs would consume less than 0.1% of the backhaul capacity.

A. Numerical Results

We provide a simulation-based comparison between the calibration scheme of [28] and our scheme. For convenience, we refer to the former as “Argos-Calibration” and to the latter as “LS-Calibration”. We consider a scenario comprising 64 single antenna APs distributed over a regularly spaced 8×8 squared grid, as shown in Fig. 6. The system serves 16 UTs simultaneously, using distributed MU-MIMO. The UTs are independently and randomly located with uniform probability over the square. The users’ achievable rates are calculated by Monte Carlo simulation, randomizing over several realizations of the users’ locations. One such realization is shown in Fig. 6. The distance between the two most distant nodes is 100m. Hence, the minimum distance in the regular grid arrangement between APs is equal to $100/7/\sqrt{2} \approx 10.1\text{m}$.

In order to isolate the impact of the TDD reciprocity calibration on the performance of MU-MIMO

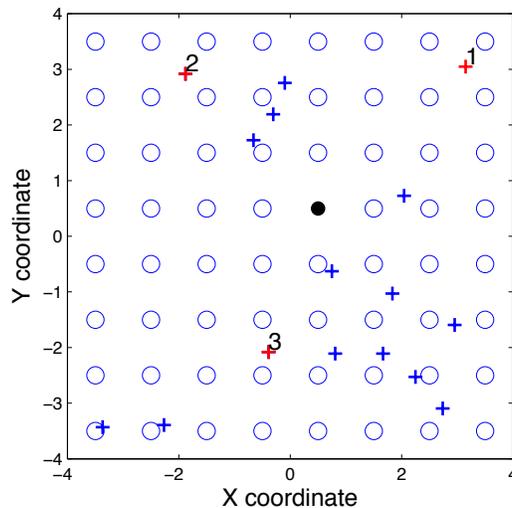


Fig. 6. Sample topology involving APs (depicted by “o”) arranged on an 8×8 grid, and 16 UTs (depicted by “+”). The reference antenna used for Argos calibration is denoted by a \bullet .

precoding, we assume perfect synchronization. The UL and DL channels are given by

$$Y_i^{up}[m, \nu] = R_i \sum_{j=1}^{N_U} \pi_{i,j} B_{i,j}[m, \nu] \tilde{T}_j X_j^{up}[m, \nu] + Z_i^{up}[m, \nu] \quad (30)$$

$$Y_j[m, \nu] = \tilde{R}_j \sum_{i=1}^{N_A} \pi_{i,j} B_{i,j}[m, \nu] T_i X_i[m, \nu] + Z_j[m, \nu], \quad (31)$$

where $Z_j[m, \nu]$ and $Z_j^{up}[m, \nu]$ are i.i.d. $\sim \mathcal{CN}(0, N_0)$ additive Gaussian noise samples, the real-nonnegative scalar $\pi_{i,j}$ denotes the large-scale path gain between AP i and UT j and we assume i.i.d. small-scale fading $B_{i,j}[m, \nu] \sim \mathcal{CN}(0, 1)$.

The large-scale path gains between AP-AP or AP-UT are both based on the WINNER model [41], where the pathloss (PL) is given in dB as a function of distance (d , in meters), carrier frequency (f_0 , in GHz), and log-normal shadowing (ρ_{dB} with variance σ_{dB}^2) as:

$$\text{PL}(d) = A \log_{10}(d) + B + C \log_{10}(f_0/5) + \rho_{\text{dB}}, \quad 3 \leq d \leq 100, \quad (32)$$

where the parameters A , B , C and σ_{dB}^2 are scenario-dependent¹⁷. We consider an indoor office scenario¹⁷ where $A = 18.7$, $B = 46.8$, $C = 20$, $\sigma_{\text{dB}}^2 = 9$ when in line of sight, otherwise $A = 36.8$, $B = 43.8$, $C =$

¹⁷These values correspond to the so-called A1 Indoor Office scenario with single, light walls in every path and where all of the UTs and APs are located on the same floor.

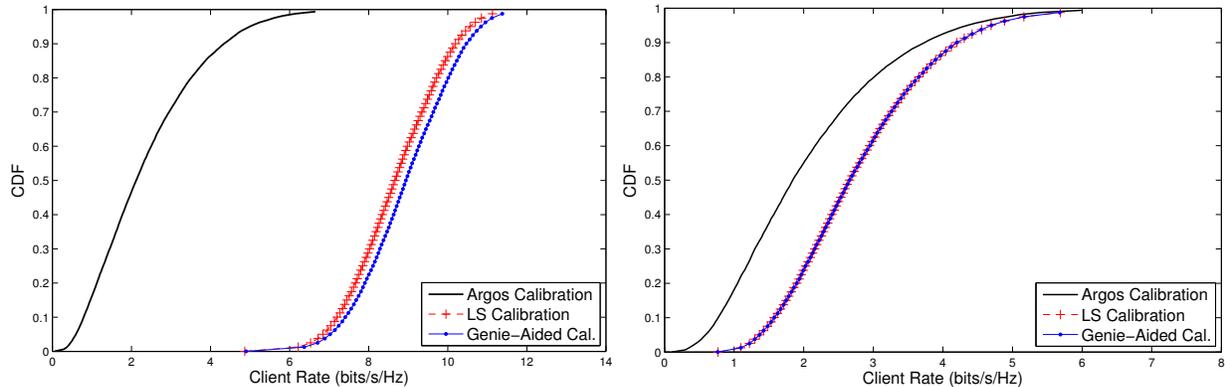


Fig. 7. The CDFs of the user rates, across the user locations and the calibration estimates. The left figure displays the results for ZFBF, while the right displays the results for conjugate beamforming.

$20, \sigma_{\text{dB}}^2 = 16$ when not in line of sight. For distances $d < 3\text{m}$, we conservatively extend the model by setting $\text{PL}(d) = \text{PL}(3)$. This is justified by the fact that the extremely high receive powers associated with shorter distances do not lead to higher link capacities, due to practical constraints on the modulation order as well as the receiver hardware (gain control and ADC range). The line of sight probability is given by a Bernoulli distribution with parameter p_{los} , which depends on distance as follows:

$$p_{\text{los}} = \begin{cases} 1 & d \leq 2.5m \\ 1 - 0.9(1 - (1.24 - 0.6 \log_{10}(d))^3)^{1/3} & \text{else.} \end{cases}$$

The large-scale path gain between AP i and UT j at distance $d_{i,j}$ is given by $\pi_{i,j} = 10^{-(\text{PL}(d_{i,j})/20)}$.

In the TDD calibration protocols, pilot signals are transmitted between APs. In particular, when AP j transmits a pilot symbol $X_j[m, \nu]$, AP i receives

$$Y_{j \rightarrow i}[m, \nu] = R_i \pi_{j \rightarrow i} B_{j \rightarrow i}[m, \nu] T_j X_j[m, \nu] + Z_{j \rightarrow i}[m, \nu]. \quad (33)$$

$\pi_{j \rightarrow i} = \pi_{i \rightarrow j}$ denotes the large-scale path gain between APs i and j and follows the same model used for $\pi_{i,j}$. The small scale fading coefficients $B_{i \rightarrow j}[m, \nu]$ have the same statistics as $B_{i,j}[m, \nu]$ and we have $B_{i \rightarrow j}[m, \nu] = B_{j \rightarrow i}[m, \nu]$ thanks to the TDD reciprocity.

The hardware-induced non-reciprocal coefficients $\{R_i, T_i\}$ and $\{\tilde{R}_j, \tilde{T}_j\}$ are modeled as i.i.d. complex random variables with uniformly distributed phase over $[-\pi, \pi]$ and uniformly distributed magnitude in $[1 - \varrho, 1 + \varrho]$ with ϱ chosen such that the standard deviation of the squared-magnitudes is equal to 0.1.

Fig. 7 shows the performance in terms of achievable rate cumulative distribution function (CDF), obtained by generating independent realizations of the channels, of the calibration estimation, and of the

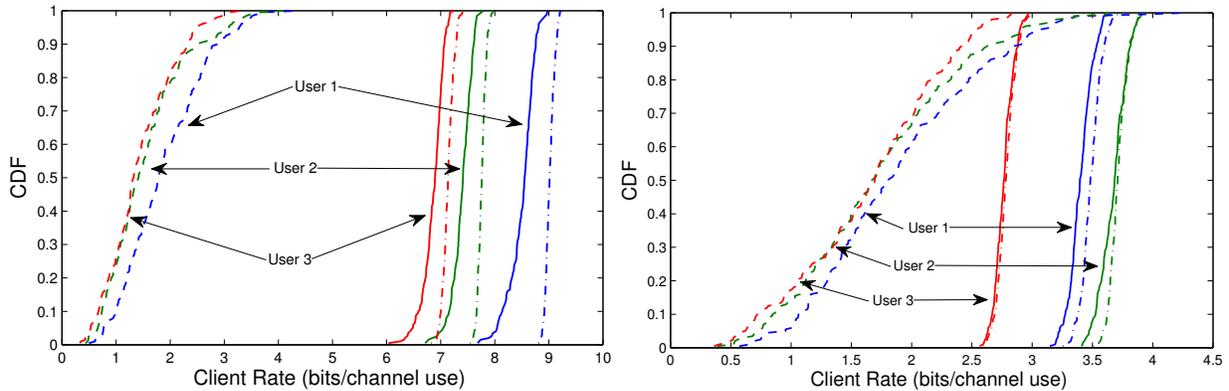


Fig. 8. Rate CDFs using ZFBF (left) and conjugate beamforming (right), with Argos-Calibration (dashed), LS-Calibration (solid) and genie-aided calibration (dash-dot), for the UTs indicated by “1”, “2” and “3” in Fig. 6.

UT positions. In these results we have assumed (for simplicity) that all UTs are served with equal power per user, and we have scaled the transmit signals by a common scaling factor such that the transmit power at each antenna does not exceed the per antenna power constraint of 90 dB.¹⁸ The reference antenna location used in Argos-Calibration is also indicated in Fig. 6. LS-Calibration is run using a fully connected graph. Genie-aided calibration uses the true values of T_i 's and R_i 's to calculate the calibration coefficients. It should be noted that the pilot overhead for the LS and Argos schemes are identical.

The performance of different calibration schemes is affected by the SNR value between antennas. As reported in [28], Argos-Calibration requires a careful placement of the reference antenna such that it has high enough SNR to all the other antennas. Argos, which was designed for co-located antennas, does not perform well in the distributed case due to the large pathlosses between the reference antenna and the more distant antennas. Since LS-Calibration does not depend on a single reference antenna, its performance is much less sensitive to the quality of a single AP-to-AP channel, and it achieves essentially genie-aided performance.

In Fig. 8, we present the rates of three particular users (labeled 1, 2 and 3 in Fig. 6) for the particular realization of the user locations show in Fig. 6.

¹⁸Backing off the transmit power in order to meet the per-AP power constraint is suboptimal. The optimal MU-MIMO zero-forcing precoding with per-antenna power constraint is discussed in detail in [42], and its optimization requires non-uniform power allocation per DL data stream and can be obtained using the theory of generalized inverses and convex relaxation. Here, we chose to use a more practical suboptimal approach for the sake of simplicity.

VI. CONCLUSIONS

In this work we presented scalable solutions for two of the main implementation hurdles of distributed MU-MIMO downlink. First, we motivated the need for timing and frequency synchronization, and detailed the performance degradation in the case of uncompensated frequency offsets between the jointly precoded APs. Next, we outlined a network-wide synchronization procedure based on a pilot burst exchange protocol, local timing and frequency offset estimation, and a centralized constrained LS solution for the overall timing and frequency correction factors. Finally, we introduced a novel method for TDD reciprocity calibration so as to enable DL MU-MIMO based on UL training, even though the APs are equipped with a standard, non-reciprocal, RF front-end. The proposed calibration protocol method, unlike its predecessors, yields near-ideal performance with distributed node deployment. Notice that both synchronization and calibration involve only the APs, and do not assume any UT collaboration. Therefore, the proposed schemes are suited to work with “legacy” user equipment, not explicitly designed for this purpose. Also, all the proposed schemes are immediately applicable to APs equipped with multiple antennas, where the groups of antennas belonging to the same AP are driven by a common clock and therefore need only to be calibrated but not mutually synchronized. Hierarchical calibration in multiple steps is naturally suited to this purpose, as evidenced in [8], [38]. It is also interesting to notice that the proposed synchronization scheme calculates the timing and frequency correction factors relative to some particular anchor AP (denoted by AP 1 in Section IV). Hence, a side benefit of our scheme is that network-wide timing and frequency stability can be achieved by having a single node (denoted as AP 1) equipped with a very stable oscillator, without the need to have a high-SNR connection from each node to this particular AP. Finally, extensive simulation results and experimental evidence provided by a software radio implementation [23], [24], we conclude that the proposed schemes effectively enable the implementation of distributed MU-MIMO network architectures, and can turn a cluster of small cell APs into a large distributed cooperative antenna system.

APPENDIX A

TIMING AND FREQUENCY JOINT ML ESTIMATION

We consider joint TO and CFO pilot-aided estimation for a transmitter-receiver pair, where the channel is multipath time-invariant with additive white Gaussian noise, with known path delays and unknown path coefficients. We denote by $\delta\tau$ and by δf the timing and carrier frequency differences between the transmitter and the receiver.

The received time-domain baseband signal corresponding to a pilot burst transmission is given by

$$y(t) = \left(\sum_{p=0}^{P-1} h_p s(t - \delta\tau - \rho_p) \right) e^{j2\pi\delta f t} + z(t) \quad (34)$$

where the multipath channel impulse response is $h(t) = \sum_{p=0}^{P-1} h_p \delta(t - \rho_p)$, and where the pilot signal is given by

$$s(t) = \frac{1}{\sqrt{T_s}} \sum_{n=0}^{N_c-1} s_n \Pi(t/T_s - n), \quad (35)$$

for the sequence of time-domain chip symbols $\mathbf{s} = (s_1, \dots, s_{N_c})$. In (34), $z(t)$ denotes the complex circularly-symmetric Gaussian white noise, with autocorrelation function $\mathbb{E}[z(t)z^*(t - \tau)] = N_0\delta(\tau)$. Letting $\nabla(t) = (\Pi \otimes \Pi)(t)$ and assuming that the receiver performs chip matched filtering and sampling with one sample per chip, we obtain the received discrete-time observation

$$y[m] = \left(\sum_{p=0}^{P-1} h_p \sum_{n=0}^{N_c-1} s_n \nabla \left(m - n - \frac{\delta\tau + \rho_p}{T_s} \right) \right) e^{j2\pi\delta f T_s m} + z[m] \quad (36)$$

where $z[m] \sim \mathcal{CN}(0, N_0)$ is a discrete-time complex circularly symmetric i.i.d. Gaussian process. As said in Section III, the receiver has a coarse knowledge of the frame timing such that it expects to find the pilot burst in a given interval including the guard times. The receiver collects M samples such that the sampled interval of duration MT_s contains the pilot burst. We collect the M received samples into an $M \times 1$ vector $\mathbf{y} = (y[0], \dots, y[M-1])^\top$. We define $\delta\xi = \delta f T_s$,¹⁹ and $\delta\mu = \delta\tau/T_s$, consistently with the normalized TO defined in (9). We define the diagonal matrix

$$\mathbf{\Xi}(\delta\xi) = \text{diag} \left(1, e^{j2\pi\delta\xi}, e^{j4\pi\delta\xi}, \dots, e^{j2\pi\delta\xi(M-1)} \right),$$

the $M \times (M - N_c + 1)$ convolution (tall Toeplitz) matrix

$$\mathbf{S} = \begin{bmatrix} s_0 & 0 & \cdots & 0 \\ s_1 & s_0 & & \vdots \\ \vdots & & \ddots & \\ s_{N_c-1} & & & \vdots \\ 0 & \ddots & & s_0 \\ \vdots & & & \vdots \\ 0 & \cdots & & s_{N_c-1} \end{bmatrix},$$

¹⁹This is related to Δ_i, Δ_j as defined in (9) by $\delta\xi = \frac{\Delta_i - \Delta_j}{N+L}$, when i is the transmitter and j the receiver.

and the matrix $\nabla(\delta\mu)$ of dimension $(M - N_c + 1) \times P$ whose p -th column (for $p = 0, \dots, P - 1$) has entries

$$\nabla(\ell - \delta\mu - \rho_p/T_s), \quad \text{for } \ell = 0, \dots, M - N_c.$$

In this way, the discrete-time sampled version of the convolution $(\nabla \otimes h)(t)$ is given by the product $\nabla(\delta\mu)\mathbf{h}$, where $\mathbf{h} = (h_0, \dots, h_{P-1})^\top$ is the vector of path coefficients. Eventually we arrive at the vector observation model

$$\mathbf{y} = \mathbf{\Gamma}(\delta\xi, \delta\mu)\mathbf{h} + \mathbf{z}, \quad (37)$$

where $\mathbf{\Gamma}(\delta\xi, \delta\mu) = \mathbf{\Xi}(\delta\xi)\mathbf{S}\nabla(\delta\mu)$ and $\mathbf{z} = (z[0], \dots, z[M - 1])^\top$.

The joint ML estimator for $\delta\xi, \delta\mu$ and \mathbf{h} , assuming N_0 and $\{\rho_p\}$ known, is obtained by maximizing the Likelihood Function

$$\Lambda(\delta\xi, \delta\mu, \mathbf{h}) = \frac{1}{(\pi N_0)^M} \exp\left(-\frac{1}{N_0} \|\mathbf{y} - \mathbf{\Gamma}(\delta\xi, \delta\mu)\mathbf{h}\|^2\right), \quad (38)$$

or equivalently by minimizing the square distance $\|\mathbf{y} - \mathbf{\Gamma}(\delta\xi, \delta\mu)\mathbf{h}\|^2$. For given $\delta\xi$ and $\delta\mu$, the minimization with respect to \mathbf{h} yields the classical LS solution

$$\hat{\mathbf{h}} = \left(\mathbf{\Gamma}^H(\delta\xi, \delta\mu)\mathbf{\Gamma}(\delta\xi, \delta\mu)\right)^{-1} \mathbf{\Gamma}^H(\delta\xi, \delta\mu)\mathbf{y}.$$

Replacing this into the objective function, after some simple algebra, we find that the joint ML estimator is obtained by maximizing with respect to $\delta\xi$ and $\delta\mu$ the quadratic form

$$F(\delta\xi, \delta\mu) = \mathbf{y}^H \mathbf{\Gamma}(\delta\xi, \delta\mu) \left(\mathbf{\Gamma}^H(\delta\xi, \delta\mu)\mathbf{\Gamma}(\delta\xi, \delta\mu)\right)^{-1} \mathbf{\Gamma}^H(\delta\xi, \delta\mu)\mathbf{y}. \quad (39)$$

This can be interpreted as maximizing the energy of the projection \mathbf{y} onto the column space of the signal matrix $\mathbf{\Gamma}(\delta\xi, \delta\mu)$. Eventually, the joint ML estimator is obtained by searching over a two-dimensional sufficiently fine grid of points with respect to the variables $\delta\xi$ and $\delta\mu$.

Next, we derive the CRB for this estimation problem. It is well-known that when the likelihood function is a multivariate Gaussian pdf $\mathcal{CN}(\mathbf{m}(\boldsymbol{\theta}), \boldsymbol{\Sigma})$ with the dependence on the parameters in the mean, then the Fisher information matrix $\mathbf{J}(\boldsymbol{\theta})$ has (a, b) elements [43]

$$J_{a,b}(\boldsymbol{\theta}) = 2\text{Re} \left\{ \left(\frac{\partial \mathbf{m}(\boldsymbol{\theta})}{\partial \theta_a} \right)^H \boldsymbol{\Sigma}^{-1} \frac{\partial \mathbf{m}(\boldsymbol{\theta})}{\partial \theta_b} \right\}.$$

In our case, we have $\boldsymbol{\Sigma} = N_0\mathbf{I}$, $\mathbf{m}(\boldsymbol{\theta}) = \mathbf{\Xi}(\delta\xi)\mathbf{S}\nabla(\delta\mu)\mathbf{h}$, and the vector of parameters $\boldsymbol{\theta}$ has elements

$$\theta_1 = \delta\xi, \theta_2 = \delta\mu, \theta_{2p+3} = \text{Re}\{h_p\}, \theta_{2p+4} = \text{Im}\{h_p\}, p = 0, \dots, P - 1.$$

Define $\mathbf{D} = \text{diag}(0, 1, \dots, M - 1)$ and let \mathbf{e}_p denote $P \times 1$ vector with all zeros and a single 1 in position p , for $p = 1, \dots, P$. Also let $\nabla'(\delta\mu)$ denote the derivative of $\nabla(\delta\mu)$ with respect to $\delta\mu$, whose

p -th column has only two non-zero elements equal to²⁰ -1 at component $\ell = \lfloor \delta\mu + \rho_p/T_s \rfloor$ and equal to $+1$ at component $\ell = \lfloor \delta\mu + \rho_p/T_s \rfloor + 1$. Then, we can write

$$\frac{\partial \mathbf{m}(\boldsymbol{\theta})}{\partial \delta\xi} = j2\pi \mathbf{D} \boldsymbol{\Xi}(\delta\xi) \mathbf{S} \nabla(\delta\mu) \mathbf{h}, \quad (40)$$

$$\frac{\partial \mathbf{m}(\boldsymbol{\theta})}{\partial \delta\mu} = \boldsymbol{\Xi}(\delta\xi) \mathbf{S} \nabla'(\delta\mu) \mathbf{h}, \quad (41)$$

$$\frac{\partial \mathbf{m}(\boldsymbol{\theta})}{\partial \text{Re}\{h_p\}} = \boldsymbol{\Xi}(\delta\xi) \mathbf{S} \nabla(\delta\mu) \mathbf{e}_p, \quad (42)$$

$$\frac{\partial \mathbf{m}(\boldsymbol{\theta})}{\partial \text{Im}\{h_p\}} = j \boldsymbol{\Xi}(\delta\xi) \mathbf{S} \nabla(\delta\mu) \mathbf{e}_p. \quad (43)$$

Defining the $M \times 2(P+1)$ matrix $\mathbf{K} = \left[\frac{\partial \mathbf{m}(\boldsymbol{\theta})}{\partial \theta_1}, \frac{\partial \mathbf{m}(\boldsymbol{\theta})}{\partial \theta_2}, \dots, \frac{\partial \mathbf{m}(\boldsymbol{\theta})}{\partial \theta_{2(P+1)}} \right]$, the Fisher information matrix is obtained as

$$\mathbf{J}(\boldsymbol{\theta}) = \frac{2}{N_0} \text{Re} \left\{ \mathbf{K}^H \mathbf{K} \right\}. \quad (44)$$

Eventually, the CRB for $\delta\xi$ and $\delta\mu$ are obtained by the first two diagonal elements of $\mathbf{J}^{-1}(\boldsymbol{\theta})$. It is easy to check that the CRB decreases as $O(1/\text{SNR})$ where SNR denotes the Signal-to-Noise Ratio at the receiver [43].

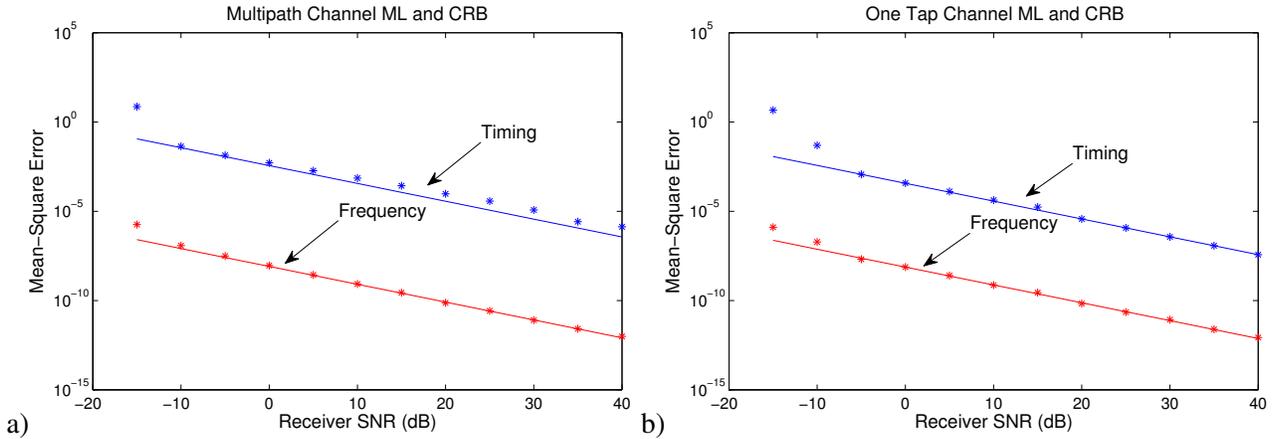


Fig. 9. CRB (solid line) and simulated estimation MSE (asterisks) of the joint ML estimation a) for a multi-path channel example b) single-path channel example.

Fig. 9 a) shows the CRB for $\delta\xi = 0$ and $\delta\mu = 0$, $\mathbf{h} = (1, 0.5, -0.2, 0.1, -0.05, -0.005)^T$, with path delays (normalized with respect to T_s) equal to $(0, 1.75, 3.56, 7.90, 10.72, 15.30)$, respectively, and for a pilot sequence \mathbf{s} formed by a block of 64 pseudo-random chips with values in the QPSK constellation,

²⁰This is the only step where we make explicit use of the fact that $\nabla(t)$ is a triangular waveform.

repeated twice and separated by 128 chips equal to 0, for a total pilot sequence length of 256 chips, corresponding to less than 4 64-carrier OFDM symbols (in fact, the OFDM symbols also include the cyclic prefix which here, for this time-domain pilot burst, are not necessary). The simulated MSE of the joint ML estimator derived before is also shown for comparison in the figure. As the figure reveals, the simulated estimation MSE is very close to the corresponding CRB even for very low SNR. Fig. 9 b) shows the associated MSEs of the joint ML estimator assuming a single-path channel between the two APs. The behavior of the MSE is very similar to the case of Fig. 9 a), showing that the joint ML estimator is remarkably insensitive to the actual channel delay-intensity profile.

REFERENCES

- [1] T. Marzetta, "Noncooperative cellular wireless with unlimited numbers of base station antennas," *IEEE Trans. on Wireless Commun.*, vol. 9, no. 11, pp. 3590–3600, Nov. 2010.
- [2] H. Huh, G. Caire, H. Papadopoulos, and S. Ramprasad, "Achieving "Massive MIMO" Spectral Efficiency with a Not-so-Large Number of Antennas," *IEEE Trans. on Wireless Commun.*, vol. 11, no. 9, pp. 3226–3239, Sep. 2012.
- [3] V. Chandrasekhar, J. Andrews, and A. Gatherer, "Femtocell networks: a survey," *IEEE Commun. Mag.*, vol. 46, no. 9, pp. 59–67, Sep. 2008.
- [4] "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems," *IEEE Std 1588-2008 (Revision of IEEE Std 1588-2002)*, pp. c1–269, 24 2008.
- [5] "IEEE 802.11-09/0308r12, TGac Channel Model Addendum," March 2010.
- [6] F. Kaltenberger, H. Jiang, M. Guillaud, and R. Knopp, "Relative channel reciprocity calibration in MIMO/TDD systems," in *Proc. Future Network and Mobile Summit, 2010*, Jun. 2010, pp. 1–10.
- [7] "IEEE Standard for Information technology–Telecommunications and information exchange between systems Local and metropolitan area networks–Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," *IEEE Std 802.11-2012 (Revision of IEEE Std 802.11-2007)*, pp. 1–2793, 29 2012.
- [8] R. Rogalin, O. Y. Bursalioglu, H. C. Papadopoulos, G. Caire, and A. F. Molisch, "Hardware-impairment compensation for enabling distributed large-scale MIMO," in *Information Theory and Applications Workshop (ITA)*. IEEE, 2013, pp. 1–10.
- [9] T. M. Schmidl and D. C. Cox, "Robust frequency and timing synchronization for OFDM," *IEEE Trans. on Commun.*, vol. 45, no. 12, pp. 1613–1621, Dec. 1997. [Online]. Available: <http://dx.doi.org/10.1109/26.650240>
- [10] J.-J. van de Beek, M. Sandell, M. Isaksson, and P. Ola Borjesson, "Low-complex frame synchronization in OFDM systems," in *Proc. IEEE International Conf. on Universal Personal Commun.*, Nov. 1995, pp. 982–986.
- [11] F. Tufvesson, O. Edfors, and M. Faulkner, "Time and frequency synchronization for OFDM using PN-sequence preambles," in *Proc. IEEE Vehic. Tech. Conf. (VTC)*, vol. 4, 1999, pp. 2203–2207.
- [12] P. Moose, "A technique for orthogonal frequency division multiplexing frequency offset correction," *IEEE Trans. on Commun.*, vol. 42, no. 10, pp. 2908–2914, Oct. 1994.
- [13] M.-O. Pun, M. Morelli, and C.-C. Kuo, "Maximum-likelihood synchronization and channel estimation for OFDMA uplink transmissions," *IEEE Trans. on Commun.*, vol. 54, no. 4, pp. 726–736, Apr. 2006.

- [14] H. Huh, A. M. Tulino, and G. Caire, "Network MIMO With Linear Zero-Forcing Beamforming: Large System Analysis, Impact of Channel Estimation, and Reduced-Complexity Scheduling," *IEEE Trans. on Inform. Theory*, vol. 58, no. 5, pp. 2911 – 2934, 2012.
- [15] S. A. Ramprasad and G. Caire, "Cellular vs. network MIMO: A comparison including the channel state information overhead," in *IEEE 20th International Symposium on Personal, Indoor and Mobile Radio Communications*. IEEE, 2009, pp. 878–884.
- [16] F. Boccardi, H. Huang, and A. Alexiou, "Network MIMO with reduced backhaul requirements by MAC coordination," in *42nd Asilomar Conference on Signals, Systems and Computers*. IEEE, 2008, pp. 1125–1129.
- [17] P. Marsch and G. Fettweis, "On base station cooperation schemes for downlink network MIMO under a constrained backhaul," in *IEEE Globecom*. IEEE, 2008, pp. 1–6.
- [18] Y.-S. Tu and G. Pottie, "Coherent cooperative transmission from multiple adjacent antennas to a distant stationary antenna through AWGN channels," in *Proc. IEEE Vehic. Tech. Conf. (VTC)*, vol. 1, 2002, pp. 130 – 134.
- [19] R. Mudumbai, D. Brown, U. Madhow, and H. Poor, "Distributed transmit beamforming: challenges and recent progress," *IEEE Commun. Mag.*, vol. 47, no. 2, pp. 102 –110, Feb. 2009.
- [20] M. M. Rahman, H. E. Baidoo-Williams, R. Mudumbai, and S. Dasgupta, "Fully wireless implementation of distributed beamforming on a software-defined radio platform," in *Proc. 11th International Conf. on Information Proc. in Sensor Networks*, ser. IPSN '12. New York, NY, USA: ACM, 2012, pp. 305–316. [Online]. Available: <http://doi.acm.org/10.1145/2185677.2185745>
- [21] F. Quitin, U. Madhow, M. Rahman, and R. Mudumbai, "Demonstrating distributed transmit beamforming with software-defined radios," in *Proc. IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, Jun. 2012, pp. 1 –3.
- [22] F. Quitin, M. Mahboob, R. Mudumbai, and U. Madhow, "Distributed beamforming with software-defined radios: frequency synchronization and digital feedback," in *Proc. IEEE GLOBECOM*, Hsinchu, Taiwan, Dec. 2012.
- [23] H. V. Balan, R. Rogalin, A. Michaloliakos, K. Psounis, and G. Caire, "Achieving high data rates in a distributed MIMO system," in *Proc. 18th Annual International Conf. on Mobile computing and networking*, ser. Mobicom '12. New York, NY, USA: ACM, 2012, pp. 41–52. [Online]. Available: <http://doi.acm.org/10.1145/2348543.2348552>
- [24] —, "AirSync: Enabling Distributed Multiuser MIMO With Full Spatial Multiplexing," *IEEE/ACM Trans. on Networking*, vol. PP, no. 99, 2012.
- [25] H. S. Rahul, S. Kumar, and D. Katabi, "JMB: scaling wireless capacity with user demands," in *Proc. ACM SIGCOMM*. New York, NY, USA: ACM, 2012, pp. 235–246. [Online]. Available: <http://doi.acm.org/10.1145/2342356.2342401>
- [26] M. Rahman, S. Dasgupta, and R. Mudumbai, "A distributed consensus approach to synchronization of RF signals," in *Proc. IEEE Statistical Signal Processing Workshop (SSP)*, Aug. 2012, pp. 281 –284.
- [27] S. Niranjayan and A. Molisch, "Ultra-wide bandwidth timing networks," in *Ultra-Wideband (ICUWB), 2012 IEEE International Conference on*, 2012, pp. 51–56.
- [28] C. Shepard, H. Yu, N. Anand, E. Li, T. Marzetta, R. Yang, and L. Zhong, "Argos: practical many-antenna base stations," in *Proc. 18th Annual International Conf. on Mobile computing and networking*, ser. Mobicom '12. New York, NY, USA: ACM, 2012, pp. 53–64. [Online]. Available: <http://doi.acm.org/10.1145/2348543.2348553>
- [29] S. Kay, "A fast and accurate single frequency estimator," *IEEE Trans. on Acoustics, Speech and Signal Process.*, vol. 37, no. 12, pp. 1987 –1990, Dec. 1989.

- [30] B. Quinn and P. Kootsookos, "Threshold behavior of the maximum likelihood estimator of frequency," *IEEE Trans. on Signal Process.*, vol. 42, no. 11, pp. 3291–3294, Nov. 1994.
- [31] M. García and C. Oberli, "Intercarrier Interference in OFDM: A General Model for Transmissions in Mobile Environments with Imperfect Synchronization," *EURASIP J. Wireless Comm. and Networking*, 2009.
- [32] A. V. Oppenheim, A. S. Willsky, and S. H. Nawab, *Signals & systems (2nd ed.)*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1996.
- [33] G. Caire, N. Jindal, M. Kobayashi, and N. Ravindran, "Multiuser MIMO achievable rates with downlink training and channel state feedback," *IEEE Trans. Inf. Theory*, vol. 56, no. 6, pp. 2845–2866, 2010.
- [34] Z. Zhang, X. Gao, and W. Wu, "Algorithms for connected set cover problem and fault-tolerant connected set cover problem," *Theoretical Computer Science*, vol. 410, no. 8-10, pp. 812 – 817, 2009. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0304397508008554>
- [35] T. Calamoneri, "The $L(h, k)$ -Labelling Problem: A Survey and Annotated Bibliography," *The Computer Journal*, vol. 49, no. 5, pp. 585–608, 2006.
- [36] S. A. Ramprashad and A. Benjebbour, "MU-MIMO in LTE: Performance and the challenges for future enhancement," *ICC'2012, Ottawa, Canada*, Jun. 2012.
- [37] A. F. Molisch, *Wireless communications*. Wiley. com, 2010, vol. 15.
- [38] H. Papadopoulos, "Towards Realizing Scalable Large-Scale MIMO with Non-Colocated Arrays," in *IEEE Communication Theory Workshop, Phuket, Thailand, June 23-26*, 2013.
- [39] H. Papadopoulos and O. Bursalioglu, "Efficient Hierarchical Reciprocity Calibration of Large TDD Networks," 2013.
- [40] A. Hjørungnes and D. Gesbert, "Complex-valued matrix differentiation: Techniques and key results," *IEEE Trans. on Signal Processing*, vol. 55, no. 6, pp. 2740–2746, Jun. 2007.
- [41] P. Kyösti, J. Meinilä, L. Hentilä, X. Zhao, T. Jämsä, C. Schneider, M. Narandzić, M. Milojević, A. Hong, J. Ylitalo, V.-M. Holappa, M. Alatossava, R. Bultitude, Y. de Jong, and T. Rautiainen, "WINNER II channel models," EC FP6, Tech. Rep., Sep. 2007. [Online]. Available: <http://www.ist-winner.org/deliverables.html>
- [42] H. Huh, H. Papadopoulos, and G. Caire, "Multiuser MISO transmitter optimization for intercell interference mitigation," *IEEE Trans. on Signal Processing*, vol. 58, no. 8, pp. 4272–4285, Aug. 2012.
- [43] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Englewood Cliffs, NJ: Prentice Hall, 1993.