Hardware-Impairment Compensation for Enabling Distributed Large-Scale MIMO

Ryan Rogalin*, Ozgun Y. Bursalioglu[†], Haralabos C. Papadopoulos[†], Giuseppe Caire*, and Andreas F. Molisch*

*Department of Electrical Engineering, University of Southern California, Los Angeles, CA 90089, USA

Email: {rogalin, caire, molisch}@usc.edu

[†]Wireless Systems Project, Docomo Innovations Inc, Palo Alto, CA 94304, USA

Email: {obursalioglu, hpapadopoulos}@docomoinnovations.com

Abstract—Distributed large-scale MIMO is a promising option for coping with the projected explosion in mobile traffic. It involves multiple Access Points (APs) that are connected to a central server via wired backhaul and act as a distributed MIMO transmitter, serving multiple users via spatial precoding. As is well known, large downlink (DL) spectral efficiencies can be achieved with TDD operation, pilots sent in the uplink (UL), and DL-UL channel reciprocity.

With APs made of inexpensive hardware and connected via, e.g., Ethernet, synchronization and reciprocity calibration are the main hurdle for implementing a truly distributed MU-MIMO system. This work studies mechanisms for RF calibration that can enable distributed high-performing large-scale MIMO operation. We propose methods for relative calibration of the APs in order to ensure TDD reciprocity while not relying on an explicitly self-calibrating RF design. As our analysis and simulations suggest, the proposed methods significantly outperform existing self calibration methods without requiring additional signaling overhead and can enable TDD reciprocity for calibration of noncolocated networks.

I. INTRODUCTION

The explosive growth of wireless device usage has pushed the demand for wireless data to unforeseen levels. While the wireless industry has made steady improvements in spectral efficiency (e.g. through adaptive coding and modulation), these improvements pale in comparison to the expected growth of demand. Two recent developments in communication theory have revealed a potential solution: Massive MIMO and very dense spectral reuse. The former refers to the use of a very large number of antennas at base stations in order to exploit spatial multiplexing on a large scale [1]. The latter pertains to the use of very small cells (i.e. femtocells [2]) such that users do not need to compete with a large number of other users for the use of the spectrum. However, the small cell solution increases the level of interference a user experiences due to the large number of cells that are transmitting to other users. Distributed MIMO unifies these two techniques, creating very large distributed antenna systems. By coordinating the access points (APs), all signal energy is useful, eliminating the interference problem of small cells. In addition, a network may contain massive levels of AP antennas (including multiantenna APs) accompanied by the enhanced spatial diversity that comes from being distributed. Distributed MIMO may be seen as a key contribution to the ultimate solution to the "spectrum crunch" problem.

While distributed MIMO may be applied in many different scenarios, in this work we focus primarily on cost-effective consumer grade equipment. We assume a network of APs, such as would be found on a corporate or academic campus, conference center, or airport, connected via a wired backbone (i.e., Ethernet) to a central server (CS) that processes data from (and to) all APs and acts as a gateway to other networks.

To achieve large spectral efficiencies over distributed largescale antenna deployments, there is a need for enabling highperforming multiuser MIMO (MU-MIMO) transmission with inexpensive hardware. To enable MU-MIMO, timely channel state information is needed at the transmitter (CSIT). The most promising MU-MIMO approach (in terms of CSIToverheads/performance trade-offs) relies on channel reciprocity and Time-Division Duplex (TDD) transmission. Uplink pilot transmissions from all user terminals (to be served with MU-MIMO) are first used to estimate the uplink channels at each AP. These estimates are then sent to the CS and used to calculate the precoder for the downlink MU-MIMO scheme within the coherence time and coherence bandwidth of the channel.

Accurate synchronization is a key requirement in these large-scale deployments for enabling sufficiently coherent MU-MIMO transmission. Software-defined radio implementations that can achieve such synchronization with over-the-air signaling are presented in [3]–[5]. They rely on the use of a master-slave protocol in order to achieve the required levels of synchronization in a distributed network of inexpensive oscillators.

Another key factor for enabling reciprocity-based distributed large-scale MU-MIMO is to guarantee that the RF chains of the radios used in the hardware implementation at the AP side do not violate the channel reciprocity requirements. In fact, while the uplink and downlink channels from antenna to antenna have identical impulse response in the same coherence interval, (recall that uplink and downlink take place at the same carrier frequency in TDD), the baseband-to-RF and RFto-baseband conversion chains need not be reciprocal and in general they are not, unless some specific self-calibrating design is used. As a result, the effective downlink baseband channel is not equal to the effective uplink baseband channel. Therefore, unless this mismatch is explicitly compensated for, learning the uplink channels is not enough to guarantee sufficient multiuser signal separation by joint precoding in the downlink. Since typical low-cost hardware designs, relying on off-the-shelf radios, have not been designed with reciprocity in mind, devising efficient and scalable signal processing schemes to achieve TDD reciprocity in a large-scale distributed MU-MIMO network is highly desirable.

In [6] such a calibration technique is presented, relying on pilot signaling among the APs and the user terminals involved in the transmission, and requiring feedback from each user terminal (UT) to the transmitting entities to enable the desired calibration. The requirement that UTs be involved in the calibration signaling and the inherent feedback overhead of this method renders it unsuitable for large-scale distributed MU-MIMO.

In [8] a proof-of-concept reciprocity-based Massive MIMO implementation, referred to as Argos, was presented along with a new TDD calibration method. One attractive feature of the Argos calibration scheme is that it only requires the APs to be involved in the calibration, i.e., it does not involve the UTs in the process. One important drawback of Argos calibration, however, is that it is very sensitive to the relative placement of the reference antenna used for calibration [8]. As a result, this scheme is not readily scalable, and is not suited for enabling large-scale MIMO in distributed antenna deployments.

In this paper we consider a new class of techniques for TDD reciprocity calibrating that can enable robust and efficient large-scale MU-MIMO operation. The techniques presented can be regarded as an extension of the Argos calibration methods [8], recovering it as a special case. As we demonstrate in this paper, the proposed scheme significantly outperforms Argos even in a co-located deployment. More important, unlike Argos, it enables effective spatial mutiplexing gains and high-performance in large-scale reciprocity-based distributed MU-MIMO system.

The remainder of the paper is organized as follows: Sec. II presents the system model. Sec. III describes the requirements for calibration, while Sec. IV establishes our new algorithm for achieving that calibration. Section V discusses the importance of synchronization, and briefly mentions algorithms that can achieve this. Section VI demonstrates the efficacy of the proposed calibration procedure by means of simulations, and Sec. VII provides some concluding remarks.

II. DISTRIBUTED RECIPROCITY-BASED MU-MIMO

We consider a network formed by user nodes $k = 1, ..., N_U$ served in the downlink by APs $i = 1, ..., N_A$, using distributed MU-MIMO and OFDM. We consider linear precoding methods, such as, e.g., linear zero-forcing beamforming (ZFBF) [9] and conjugate beamforming [8].

We focus our attention to the calibration problem, and, for the time being, assume that synchronization is perfectly achieved (see Section V for further discussion on synchronization). The downlink signal at the UT receivers, at OFDM symbol m and subcarrier ν , is given by the $1 \times N_{\rm U}$ vector

$$\mathbf{Y}[m,\nu] = \mathbf{X}[m,\nu]\mathbf{H}[m,\nu] + \mathbf{Z}[m,\nu].$$
(1)

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, and $\mathbf{Z}[m, \nu]$ is the corresponding $1 \times N_U$ vector of i.i.d. $\mathcal{CN}(0, N_0)$ noise samples. The downlink channel matrix $\mathbf{H}[m, \nu]$ is given by

$$\mathbf{H}[m,\nu] = \mathbf{T}[m,\nu]\mathbf{B}[m,\nu]\mathbf{R}[m,\nu], \qquad (2)$$

where $\widetilde{\mathbf{R}}[m,\nu] = \operatorname{diag}(\widetilde{R}_1[m,\nu],\ldots,\widetilde{R}_{N_{\mathrm{U}}}[m,\nu])$ and $\mathbf{T}[m,\nu] = \operatorname{diag}(T_1[m,\nu],\ldots,T_{N_{\mathrm{A}}}[m,\nu])$ are diagonal matrices of complex coefficients, introduced by the users' receiver chains, and by the APs' transmission chains, respectively. The matrix $\mathbf{B}[m,\nu]$ represents the discrete-time frequency domain physical channel at subcarrier ν and OFDM symbol m, containing the channel coefficients due solely to the antennato-antenna propagation.

In reciprocity-based MU-MIMO the downlink channel matrix is estimated at the APs based on uplink pilot signals transmitted by the user terminals. The relevant uplink channel at OFDM symbol m and subcarrier ν is given by

$$\mathbf{Y}^{\mathrm{up}}[m,\nu] = \mathbf{H}^{\mathrm{up}}[m,\nu]\widetilde{\mathbf{X}}[m,\nu] + \mathbf{Z}^{\mathrm{up}}[m,\nu], \qquad (3)$$

where $\mathbf{Z}^{up}[m,\nu]$ is the uplink Gaussian noise vector, and $\widetilde{\mathbf{X}}[m,\nu]$ is a $N_{\mathrm{U}} \times N_{\mathrm{U}}$ unitary matrix of frequency domain uplink pilot symbols. The uplink channel matrix, $\mathbf{H}^{up}[m,\nu]$ satisfies

$$\mathbf{H}^{\mathrm{up}}[m,\nu] = \mathbf{R}[m,\nu]\mathbf{B}[m,\nu]\mathbf{T}[m,\nu]$$
(4)

with $\widetilde{\mathbf{T}}[m,\nu] = \operatorname{diag}(\widetilde{T}_1[m,\nu],\ldots,\widetilde{T}_{N_{\mathrm{U}}}[m,\nu])$ denoting the matrix of user transmitter coefficients, and $\mathbf{R}[m,\nu] = \operatorname{diag}(R_1[m,\nu],\ldots,R_{N_{\mathrm{A}}}[m,\nu])$ denoting the matrix of AP receiver coefficients.

The key property exploited by reciprocity-based MU-MIMO is that the physical channel matrix $\mathbf{B}[m, \nu]$ is the *same* in both uplink and downlink, due to TDD (uplink and downlink are at the same carrier frequency) and the reciprocity of the physical propagation channel ¹. In the absence of RF impairments, i.e., in the case that

 $\mathbf{R}[m,\nu] = \mathbf{T}[m,\nu] = \mathbf{I}_{N_{\mathrm{A}}}, \quad \text{and} \quad \widetilde{\mathbf{R}}[m,\nu] = \widetilde{\mathbf{T}}[m,\nu] = \mathbf{I}_{N_{\mathrm{U}}}$

we have

$$\mathbf{H}^{\mathrm{up}}[m,\nu] = \mathbf{H}[m,\nu]$$

and hence, estimates of the uplink channel $\mathbf{H}^{up}[m,\nu]$ directly provide estimates of the downlink channel $\mathbf{H}[m,\nu]$. These channel estimates can then be directly used at the APs to calculate the MU-MIMO precoder for the downlink. In practice, however, $\mathbf{R}[m,\nu]$, $\mathbf{T}[m,\nu]$, $\mathbf{\tilde{R}}[m,\nu]$, $\mathbf{\tilde{T}}[m,\nu]$ are non-identity unknown diagonal matrices that vary slowly in time (m) and frequency (ν). These impairments have to be compensated in order to enable reciprocity-based MU-MIMO transmission.

¹As long as the interval between UL and DL is much smaller than a channel coherence time, which is typically between 1ms and 100 ms

III. CALIBRATION REQUIREMENTS FOR UPLINK-DOWNLINK CHANNEL RECIPROCITY

In this section we revisit the relative calibration requirements presented in [8]. First note that, since $\mathbf{R}[m,\nu]$, $\mathbf{T}[m,\nu]$, $\widetilde{\mathbf{R}}[m,\nu]$, and $\widetilde{\mathbf{T}}[m,\nu]$ vary very slowly in m (order of several minutes), essentially because of the temperature drift of the front-end electronic components, the calibration protocol can operate at a much slower time scale than the MU-MIMO uplink channel estimation (whose time-scale is dictated by the coherence-time of the physical propagation channel). Hence, for the sake of estimation, these matrices can be treated as unknown constants.

Without loss of generality we focus on a particular tone ν and drop the dependence of all variables on the OFDM symbol index m and subcarrier ν for notation simplicity. The uplink pilot burst \mathbf{Y}^{up} is sent to the CS, which estimates the uplink channel as²

$$\widehat{\mathbf{H}}^{\mathrm{up}} = \mathbf{Y}^{\mathrm{up}} \widetilde{\mathbf{X}}^{\mathsf{H}} = \mathbf{H}^{\mathrm{up}} + \widetilde{\mathbf{Z}}^{\mathrm{up}},$$

with $\widetilde{\mathbf{Z}}^{up} = \mathbf{Z}^{up} \widetilde{\mathbf{X}}^{H}$. Neglecting for the time being the estimation error $\widetilde{\mathbf{Z}}^{up}$, we have that if the CS computes the downlink multiuser MIMO precoder from $\widehat{\mathbf{H}}^{up}$ this is will be mismatched with respect to the downlink channel H because of the presence of the diagonal matrices $\mathbf{T}, \widetilde{\mathbf{R}}$ in lieu of $\mathbf{R}, \widetilde{\mathbf{T}}$.

The key observation made in [8] is that the downlink channel matrix $\mathbf{TB}\widetilde{\mathbf{R}}$ is not entirely needed to perform beamforming. In fact, only the column-space of this matrix is needed, that is, any matrix formed by

$$\mathbf{H}_{\text{alt}} = \mathbf{T}\mathbf{B}\mathbf{D},\tag{5}$$

with **D** some arbitrary invertible constant diagonal matrix, can be used as an alternative for any kind of beamforming. For example, consider Zero Forced Beamforming (ZFBF). We can calculate the ZFBF precoding matrix as

$$\mathbf{V}(\mathbf{H}_{\text{alt}}) = \mathbf{\Lambda}^{1/2} \left[\mathbf{H}_{\text{alt}}^{\mathsf{H}} \mathbf{H}_{\text{alt}} \right]^{-1} \mathbf{H}_{\text{alt}}^{\mathsf{H}}$$
(6)

where Λ is a diagonal matrix that imposes on each row of the matrix \mathbf{V} , the row $\|\mathbf{v}_i\|^2 = 1$, for all *i*. Hence, the received version of ZFBF precoded signal \mathbf{u} in the downlink becomes

$$\mathbf{Y} = \mathbf{u}\mathbf{V}\mathbf{T}\mathbf{B}\widetilde{\mathbf{R}} + \mathbf{Z}$$
(7)

$$= \mathbf{u} \mathbf{\Lambda}^{1/2} [\mathbf{D}^{\mathsf{H}} \mathbf{B}^{\mathsf{H}} \mathbf{T}^{\mathsf{H}} \mathbf{T} \mathbf{B} \mathbf{D}]^{-1} \mathbf{D}^{\mathsf{H}} \mathbf{B}^{\mathsf{H}} \mathbf{T}^{\mathsf{H}} \mathbf{T} \mathbf{B} \mathbf{\widetilde{R}} + \mathbf{Z} (8)$$

$$= \mathbf{u} \mathbf{\Lambda}^{1/2} \mathbf{D}^{-1} \mathbf{R} + \mathbf{Z}$$
(9)

We notice that the resulting channel matrix is diagonal, provided that $N_{\rm U} \leq N_{\rm A}$. As a result, the problem reduces to estimating **TB** up to the left multiplication by some matrix **D**, from the uplink training observation **RB** $\widetilde{\mathbf{T}}$.

In particular, as shown in [8], an estimate of the diagonal *relative calibration matrix* $\alpha \mathbf{RT}^{-1}$, for some non-zero scalar α , suffices for enabling spatial multiplexing with reciprocitybased MU-MIMO. Assuming that the uplink channel $\mathbf{H}^{up} = \mathbf{RBT}$ is provided by uplink channel estimation (ignoring estimation noise), and $\alpha \mathbf{RT}^{-1}$ is available for some arbitrary (and unknown) $\alpha \neq 0$, multiplying \mathbf{H}^{up} from the left by the inverse of $\alpha \mathbf{RT}^{-1}$ provides the CS with an matrix \mathbf{H}_{alt} of the form (5) with $\mathbf{D} = \alpha^{-1} \mathbf{\tilde{T}}$. As a result,

$$\mathbf{Y} = \mathbf{u} \left[\alpha \mathbf{\Lambda}^{1/2} \mathbf{T}^{-1} \widetilde{\mathbf{R}} \right] + \mathbf{Z}, \tag{10}$$

i.e., the effective downlink precoded channel matrix is diagonal. An alternative way of arriving to this operation involves pre-compensating for the effects of the RF impairments at each AP. In particular, if each AP *i* premultiplies its own transmit signal by the corresponding element $\alpha R_i/T_i$, the downlink channel can be turned into a "calibrated" downlink channel with matrix $\alpha RB\tilde{R}$. Suppose for example that the multiuser downlink precoding matrix is the uplink channel pseudo-inverse $\mathbf{V} = \mathbf{\Lambda}^{1/2} \left[(\mathbf{H}^{up})^{\mathsf{H}} \mathbf{H}^{up} \right]^{-1} (\mathbf{H}^{up})^{\mathsf{H}}$. Then, this matrix applied to the calibrated downlink channel yields \mathbf{Y} in the form (10), resulting again in a diagonal matrix of the precoded downlink channel.

In summary, spatial multiplexing with perfect user signal separation is possible, provided the calibration protocol allows (sufficiently accurate) estimation of the matrix $\alpha \mathbf{RT}^{-1}$, defined up to some arbitrary non-zero factor α .

IV. ROBUST CALIBRATION FOR UPLINK-DOWNLINK RECIPROCITY

In this section we present novel TDD relative calibration techniques that generalize the idea exposed in [8] to the case of an arbitrary distributed network topology. As outlined in Section III, the calibration protocol operates at much slower time scales than synchronization and channel estimation, i.e., the system inserts one to a few special calibration slots every many frames. Calibration slots are formed by pilot bursts designed to have a flat transmit power spectral density. For example, this can be obtained by sending some OFDM symbols formed by known frequency domain symbols. Calibration might be done independently at each subcarrier or, by exploiting the fact that the non-reciprocal elements of the channel (due to the transmit and receive chains) are typically smooth over the signal bandwidth both in amplitude and phase, it can be performed after some smoothing in the frequency domain, in order to gain noise margin. Here, we focus on a single subcarrier and neglect the possible improvement by frequency smoothing.

A. Prior Art: Argos Calibration [8]

We first review the relative calibration method of [8]. The goal consists of estimating $\alpha R_i/T_i$ for each AP i^3 , i.e., estimating R_i/T_i up to a (common for all *i*) multiplicative constant. Letting $c_i = R_i/T_i$ and setting this constant α equal to one of the c_i 's, e.g., $\alpha = c_1$, the task of calibration reduces to estimating all c_i 's relatively to c_1 .

In Argos [8], each AP is calibrated independently of all other APs with respect to a reference AP. In particular, the Argos calibration procedure is as follows

²Note that, as explained in Section IV-C, in our simulations we used improved MMSE-type estimates in place of $\widehat{\mathbf{H}}^{up}$.

³In the centralized case, each AP is essentially an array element

- (A1) Sequentially transmit calibration pilots, one pilot from each AP.
- (A2) Calibrate AP j with respect to (reference) AP 1, for each $j \neq 1$.

The observation collected by AP j when AP i transmits its pilot can be expressed as

$$Y_{i \to j} = T_i \, B_{i \to j} \, R_j + Z_{i \to j} \tag{11}$$

where $B_{i \rightarrow j}$ is the channel response from antenna *i* to antenna *j* that is solely due to the propagation environment, and $Z_{i \rightarrow j}$ represents thermal noise. The Argos calibration mechanism (A2) for AP *j* relies only on the observations collected by the pair of APs 1 and *j* during the associated pair of pilot transmissions. In particular, since $B_{j\rightarrow 1} = B_{1\rightarrow j}$, the ratio $Y_{1\rightarrow j}/Y_{j\rightarrow 1}$ provides an estimate of the ratio $(R_j/T_j)/(R_1/T_1)$, i.e., it provides an estimate of the desired relative calibration parameter.

B. Robust relative calibration for distributed MU-MIMO

In this section we present the proposed generalization of [8], enable scalable TDD calibration in a distributed MU-MIMO deployment.

Assume that the APs form a connected directed network graph $(\mathcal{T}, \mathcal{E})$, where $\mathcal{T} = \{1, \ldots, N_A\}$ and $(i, j) \in \mathcal{E}$ if the channel between APs *i* and *j* has sufficiently large SNR. During the calibration slots, pilot bursts are transmitted and received by the APs over a connected spanning subgraph $(\mathcal{T}, \mathcal{F})$ including all the APs and a subset of links $\mathcal{F} \subseteq \mathcal{E}$. Specifically, we have $(i, j) \in \mathcal{F}$ if there is a pair of observations $\{Y_{i \to j}, Y_{j \to i}\}$ of the form (11), due to calibration pilots transmitted by APs *i* and *j* on distinct OFDM symbols but within the same coherence-time of the channel. Hence, the subset of links is such that if $(i, j) \in \mathcal{F}$ then also $(j, i) \in \mathcal{F}$. For example, $(\mathcal{T}, \mathcal{F})$ could be obtained as a spanning tree of the underlying undirected network graph, where each edge of the spanning tree corresponds to two directed edges in \mathcal{F} .

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

$$Y_{j \to i} = T_j B_{j \to i} R_i + Z_{j \to i}, \tag{12}$$

and AP j gathers the observation given by (11). Grouping such measurements in pairs, we have

$$\begin{bmatrix} Y_{j \to i} \\ Y_{i \to j} \end{bmatrix} = \begin{bmatrix} T_j R_i \\ T_i R_j \end{bmatrix} B_{i \to j} + \begin{bmatrix} Z_{j \to i} \\ Z_{i \to j} \end{bmatrix}$$
$$= \begin{bmatrix} c_i \\ c_j \end{bmatrix} \beta_{ij} + \begin{bmatrix} Z_{j \to i} \\ Z_{i \to j} \end{bmatrix}, \quad (13)$$

owing to the fact that, by the physical channel reciprocity, $B_{i\rightarrow j} = B_{j\rightarrow i}$, and defining $\beta_{ij} = \beta_{ji} = T_i T_j B_{i\rightarrow j}$. Our goal is to estimate the relative calibration coefficients c_i for $i = 1, \ldots, N_A$, up to a common multiplicative non-zero constant. Without loss of generality we assume that $\{c_i\}$ is a set of nonzero bounded complex scalars (if $c_i = 0$ or $1/c_i = 0$, the *i*-th node can be omitted as it is a "non-communicating" node.).

Inspection of (13) reveals that if the observations $Y_{j\to i}, Y_{i\to j}$ were noiseless, we would have $c_j Y_{j\to i} = c_i Y_{i\to j}$ for all $(i, j) \in \mathcal{F}_u$, where \mathcal{F}_u is the set of undirected edges

corresponding to \mathcal{F} , i.e., $\mathcal{F}_u = \{(i, j); (i, j) \in \mathcal{F} \text{ and } (j, i) \in \mathcal{F}\}$. Hence, a natural approach in the presence of observation noise is to define the following LS cost function

$$J_{\rm cal}(c_1, c_2, \dots, c_{N_{\rm A}}) = \sum_{(i,j)\in\mathcal{F}_u} |c_j Y_{j\to i} - c_i Y_{i\to j}|^2, \quad (14)$$

and find the solution $\mathbf{c} = (c_1, c_2, \dots, c_{N_A})$ that minimizes (14). At this point, some observations are in order. First, observe that in order to exclude the trivial all-zero solution we need to impose a fixed value for $c_1 \neq 0$, e.g., $c_1 = 1$ (we chose without loss of generality AP 1 as the reference AP). Second, notice that the *constrained* non-trivial solution is defined up to an arbitrary multiplicative constant α of magnitude 1. Hence, there is no loss of generality in solving for c_2, \dots, c_{N_A} as a function of c_1 and finally replace some suitable value of c_1 with non-zero (e.g., unit) magnitude.

We wish to solve the minimization of (14) subject to, e.g., $c_1 = 1$ for an arbitrary topology. To this purpose, we differentiate J_{cal} with respect to c_i^* , treating c_i and c_i^* as if they were independent variables [10], and then set the partial derivatives equal to zero. We obtain

$$\frac{\partial}{\partial c_i^*} J_{\text{cal}}(c_1, c_2, \dots, c_{N_{\text{A}}}) = \sum_{j:(i,j)\in\mathcal{F}_u} \left(c_i |Y_{i\to j}|^2 - c_j Y_{i\to j}^* Y_{j\to i} \right).$$
(15)

In matrix form, we obtain Ac = 0, where A is the $N_A \times N_A$ matrix with element in its *i*-th row and *j*-th column given by

$$A_{i,j} = \begin{cases} \sum_{j:(i,j)\in\mathcal{F}_u} |Y_{i\to j}|^2 & \text{for } j=i\\ -Y_{i\to j}^*Y_{j\to i} & \text{for } j\neq i, \ (i,j)\in\mathcal{F}_u\\ 0 & \text{for } j\neq i, \ (i,j)\notin\mathcal{F}_u \end{cases}$$

Finally, we can solve for the variables $\tilde{\mathbf{c}} = (c_2, \dots, c_{N_A})^{\mathsf{T}}$ by letting $\mathbf{A} = [\mathbf{a}_1 | \mathbf{A}_1]$ where \mathbf{a}_1 is the first column of \mathbf{A} , such that

$$\widetilde{\mathbf{c}} = -(\mathbf{A}_1^{\mathsf{H}} \mathbf{A}_1)^{-1} \mathbf{A}_1^{\mathsf{H}} \mathbf{a}_1 c_1.$$
(16)

Hence, any value of c_1 with unit magnitude yields a solution of the relative calibration problem.

We next note that the Argos relative calibration [8] (reviewed in Section IV-A) coincides with the solution to (14) subject to the graph $(\mathcal{T}, \mathcal{F})$ being a star with AP 1 at the center. In fact, in this case the objective function is given by

$$J_{\rm cal}(c_1, c_2, \dots, c_{N_{\rm A}}) = \sum_{j \neq 1} |c_j Y_{j \to 1} - c_1 Y_{1 \to j}|^2, \quad (17)$$

where the constrained minimum is obviously achieved by letting $c_j = \frac{Y_1 \rightarrow j}{Y_j \rightarrow 1} c_1$ for some c_1 , as proposed in [8]. In general, however, we can obtain significantly better performance than [8] by considering topologies different from the star topology. As we demonstrate in Section VI via simulations, this is especially true in the case of APs distributed over a relatively large area, resulting in AP-to-AP channel SNRs that can vary significantly between different AP pairs.

C. MU-MIMO Operation

We next consider the MU-MIMO training and signaling operation based on a given set of estimates, $\{\hat{c}_k\}$ of the relative calibration parameters. First, observations of the form (3) are collected based on uplink pilots. These observations are then used at the CS to obtain an MMSE estimate of the \mathbf{H}^{up} , namely⁴, $\hat{\mathbf{H}}^{up}$. Assuming the set of $\{\hat{c}_k\}_{k=2}^{N_A}$ is computed via an equation of the form (16) with $c_1 = 1$, the matrix \mathbf{H}_{alt} is then constructed as follows:

$$\mathbf{H}_{\text{alt}} = \text{diag}\left(1, \hat{c}_2^{-1}, \dots, \hat{c}_{N_{\text{A}}}^{-1}\right) \,\widehat{\mathbf{H}}^{\text{up}} \,. \tag{18}$$

We remark that if we replace \hat{c}_i with c_i/c_1 , and $\hat{\mathbf{H}}^{up}$ with \mathbf{H}^{up} , the matrix \mathbf{H}_{alt} takes the desired form (5) with \mathbf{D} diagonal.

Consequently, given \mathbf{H}_{alt} from (18), and any given precoder function $\mathbf{V} = \mathbf{V}(\mathbf{H}_{\text{alt}})$, such as e.g., ZFBF in (6), the effective downlink channel is given by (7), with effective $N_{\text{U}} \times N_{\text{U}}$ channel matrix given by $\boldsymbol{\Phi} = \mathbf{VTBR}$. We then use as our performance metric for the *i*th user the instantaneous rate the quantity $\log_2(1 + \text{SINR}_i)$, with SINR_i computed in the usual manner.

The calibration performance of different calibration methods is evaluated via comparison of the associated user instantaneous rates (subject to a common MU-MIMO precoder method). As an upper bound we consider the performance with genie-aided calibration. Genie-aided calibration uses the same MU-MIMO training and signaling operation, with $\mathbf{H}_{\rm alt}$ in (18) replaced with

$$\mathbf{H}_{\text{alt}}^{\text{genie}} = \text{diag}\left(c_1^{-1}, c_2^{-1}, \dots, c_{N_{\text{A}}}^{-1}\right) \, \widehat{\mathbf{H}}^{\text{up}} \, . \tag{19}$$

D. Hierarchical calibration

In a large distributed network of nodes it may be necessary to provide relative calibration between large sets of nodes that are dispersed over wide areas. As inspection of (16) reveals, the calibration methods of the previous section require inversion of a matrix with dimensions equal to the cardinality of the size of the network. Clearly, with increasing network sizes, computationally, these methods may become prohibitively expensive.

In this section we consider an alternative approach that becomes attractive for calibration in large-scale networks, which relies on hierarchical calibration. In its simplest two-layer case, this involves first splitting the network in sufficiently small-size clusters, and using the techniques of the previous section to calibrate all nodes within each cluster. Subsequently a second, inter-cluster, calibration step is performed, which accomplishes relative calibration across clusters.

For convenience we re-index cluster nodes within each cluster, and denote by (i,m) the *m*-th node in cluster *i*. We also let $c_{i,m} = R_{i,m}/T_{i,m}$ denote the unknown parameter of interest. We will assume that for each *i*, sufficiently accurate



Fig. 1. Inter-cluster calibration among nine 3×3 clusters of APs, based on two-way measurements (double arrows) between APs in different clusters.

intra-cluster calibration has been performed using an algorithm of the form (16) so that each node has been calibrated relative to a reference node in cluster c_i . In particular, given that the algorithm (16) applied to cluster *i* has returned intra-cluster calibration estimates $\{\hat{c}_{i,m}\}_m$, we have

$$c_{i,m} \approx \hat{c}_{i,m} \, c_i,\tag{20}$$

for some unknown parameter c_i . We also let $Y_{(i,m)\to(j,n)}$ denote the observation at node (j,n) based on a pilot transmitted by node (i,m), i.e., an observation of the form (11), with i and j replaced by (i,m) and (j,n), respectively.

In order to study the inter-cluster calibration problem. we next consider clusters as nodes on a graph. Assume a connected cluster-network graph ($\mathcal{T}^{cl}, \mathcal{E}^{cl}$), pilot bursts are transmitted and received by APs across clusters over a connected spanning subgraph $(\mathcal{T}^{\mathrm{cl}},\mathcal{F}^{\mathrm{cl}})$ including all the clusters and a subset of links $\mathcal{F}^{cl} \subset \mathcal{E}^{cl}$. We also let \mathcal{F}^{cl}_u denote the set of undirected edges corresponding to \mathcal{F} . A pair $(i, j) \in \mathcal{F}_u^{cl}$, if there is at least one pair of observations $\{Y_{(i,m)\to(j,n)}, Y_{(j,m)\to(i,n)}\}$ of the form (11) that are to be used for calibration), and where the pair is due to a pair of calibration pilots transmitted by APs (i, m) and (j, n) on distinct OFDM symbols but within the coherence-time of the channel. We also let \mathcal{G}_{ij} denote the set of all (m, n) index pairs for which such bi-directional pairs of observations are available between APs (i, m) and (j, n), in clusters i and j respectively. Thus, $(i, j) \in \mathcal{F}_u^{cl}$ if and only if the set \mathcal{G}_{ij} is non-empty.

A visual interpretation of the hierarchical calibration problem is shown in Figures 1 and 2. Figure 1 depicts a 9×9 network of nodes. As shown in the figure, there are nine 3×3 clusters of APs. Assuming intra-cluster calibration has already

⁴In principle, determining the gain of the MMSE filter requires knowledge of the magnitudes of the RF-impairments at the AP receivers and the user terminal transmitters, i.e., quantities that are unknown. In our simulations we simply used the large-scale gains in the associated point-to-point channels as indicators of the large scale SNR in order to determine the MMSE filter gains.



Fig. 2. Subgraph based on which inter-cluster calibration is performed on the network of Figure 1.

been performed within each 3×3 cluster, the bidirectional arrows (where each arrow represents two-way measurements between APs in different clusters) represent a set of measurements that can be used for inter-cluster calibration.

Figure 2 shows the corresponding connected spanning subgraph ($\mathcal{T}^{cl}, \mathcal{F}^{cl}$) associated with the two-way measurements shown in Figure 1, on which the inter-cluster calibration is to be performed. The inter-cluster calibration problem can be tackled with a straightforward extension of the baseline methods of Section IV-B. The associated objective function for inter-cluster calibration can then be readily expressed as follows

$$J_{\rm h} = \sum_{(i,j)\in\mathcal{F}_u^{\rm cl}} \sum_{(m,n)\in\mathcal{G}_{ij}} \left| c_{i,m} Y_{(i,m)\to(j,n)} - c_{j,n} Y_{(j,n)\to(i,m)} \right|^2.$$
(21)

The solution can be readily derived by following the same steps as for the baseline LS-calibration. Letting

$$Y_{(i,m)\to(j,n)} = \hat{c}_{i,m} Y_{(i,m)\to(j,n)}$$
 (22)

and using (20) we can re-express $J_{\rm h}$ in (21), as a function of the c_i 's as follows

$$J_{\rm h} = \sum_{(i,j)\in\mathcal{F}_{u}^{\rm cl}} \sum_{(m,n)\in\mathcal{G}_{ij}} \left| c_i \tilde{Y}_{(i,m)\to(j,n)} - c_j \tilde{Y}_{(j,n)\to(i,m)} \right|^2$$
(23)

Letting $\tilde{\mathbf{c}} = (c_2, \ldots, c_{N_A})^{\mathsf{T}}$, with $N_{\mathrm{C}} = |\mathcal{T}^{\mathrm{cl}}|$, the vector $\tilde{\mathbf{c}}$ that minimizes (23) as a function of c_1 is given by (16), and where \mathbf{A} is the $N_{\mathrm{C}} \times N_{\mathrm{C}}$ matrix with element in its *i*-th row and *j*-th column given by

$$A_{i,j} \!=\! \begin{cases} \sum_{\substack{j:(ij)\in\mathcal{F}_u^{\mathrm{cl}} \\ (m,n)\in\mathcal{G}_{ij}}} \sum_{\substack{(m,n)\in\mathcal{G}_{ij} \\ -\sum_{(m,n)\in\mathcal{G}_{ij}}} \tilde{Y}^*_{(i,m)\to(j,n)} \tilde{Y}_{(j,n)\to(i,m)} & \text{for } j \!\neq\! i, \end{cases}$$

Notice that for some $i \neq j$, the coefficient $A_{i,j}$ may be zero, if \mathcal{G}_{ij} is empty, i.e., if $(i, j) \notin \mathcal{F}_u^{cl}$.

V. SYNCHRONIZATION ASPECTS

Since the APs are driven by individual commercial-grade clocks, they operate at carrier frequencies $f_{0,i}$ and sampling rate $f_{s,i}$ close but not identical to their nominal values f_0 and

 f_s , respectively. After some algebra (see [7]) and assuming that the carrier and sampling frequency offsets (CFO and SFO, respectively) relative to the nominal frequencies, are sufficiently small (e.g., ≈ 20 ppm as given by the specifications of the 802.11n protocol [11]), we obtain a discrete-time baseband model of the MU-MIMO OFDM channel including synchronization errors in the form

$$\mathbf{Y}[m,\nu] = \mathbf{X}[m,\nu]\mathbf{\Phi}[m,\nu]\mathbf{H}[m,\nu]\mathbf{\Theta}[m,\nu] + \mathbf{Z}[m,\nu],$$
(24)

where $\mathbf{H}[m, \nu]$ is the $N_{\rm A} \times N_{\rm U}$ channel matrix defined before, and the two matrices $\boldsymbol{\Theta}[m, \nu]$ and $\boldsymbol{\Phi}[m, \nu]$ are diagonal of dimension $N_{\rm U} \times N_{\rm U}$ and $N_{\rm A} \times N_{\rm A}$, respectively, with diagonal elements given by ⁵

$$\phi_{i,i}[m,\nu] = \exp\left(-j2\pi \left[\frac{\nu}{N}\right]\mu_i\right) \\ \times \exp\left(j2\pi \left[\frac{\nu}{N}\right]\delta_i m\right) \\ \times \exp\left(-j2\pi\Delta_i m\right),$$
(25)

for subcarrier index $\nu \in \{-N/2, ..., N/2 - 1\}$, where we define the following timing offset (TO), SFO and CFO normalized quantities:

$$\mu_i \stackrel{\Delta}{=} f_s \tau_i,\tag{26}$$

 τ_i being the TO between the AP *i* time reference and the nominal time 0 reference of the frame structure;

$$\delta_i \stackrel{\Delta}{=} \frac{(N+L)\epsilon_i}{f_s},\tag{27}$$

 $\epsilon_i = f_{s,i} - f_s$ being the SFO of AP *i* relative to the nominal sampling frequency, and where N is the number of subcarriers and L is the cyclic prefix length (in time-domain samples) of the OFDM modulation; and

$$\Delta_i \stackrel{\Delta}{=} \frac{\gamma(N+L)\epsilon_i}{f_s} = \gamma \delta_i, \tag{28}$$

where we assume, as justified by the 802.11 standard [11], that sampling clock and the RF clock on each AP hardware are derived from the same (local) oscillator, with the ratio $f_0 = \gamma f_s$ for some $\gamma \gg 1$.

While the effect of $\Theta[m, \nu]$ can be undone at each UT receiver by standard timing and frequency synchronization techniques suited to OFDM modulation [12]–[16], the presence of $\Phi[m, \nu]$ in between the precoded transmit vector $\mathbf{X}[m, \nu]$ and the channel matrix $\mathbf{H}[m, \nu]$ yields a degradation of the MU-MIMO precoder performance. In fact, the ZFBF precoding matrix is calculated based on an estimate of the "nominal" channel matrix $\mathbf{H}[0, \nu]$, at the beginning of the preceding block (here indicated as OFDM symbol m = 0), but is ignorant of the matrix $\Phi[m, \nu]$ unless TO, SFO and CFO are explicitly taken into account and compensated.

⁵The diagonal elements $\theta_{k,k}[m,\mu]$ take on an identical form with $\widetilde{\mu}_k, \widetilde{\delta}_k$ and $\widetilde{\Delta}_k$ replacing with μ_i, δ_i and Δ_i , respectively, where the "tilde" quantities denote normalized TO, SO and CFO at the user receivers side.

In order to motivate the need for accurate synchronization of the APs, we examine the performance degradation due to typical uncompensated SFO and CFO between the APs. Again, we use $\log_2(1 + \text{SINR}_i)$, to evaluate the achievable rate. We assume downlink blocks of M = 60 OFDM symbols, typical of 802.11 [11], and assume $\mathbf{H}[m,\nu]$ constant with respect to the time index $m = 0, \ldots, M - 1$ over each block. Optimistically, we assume ideal TDD reciprocity and noiseless uplink channel estimation through uplink pilot symbols. Provided that the relative TO between APs is within the length of the OFDM cyclic prefix, the terms $\exp\left(-j2\pi\left[\frac{\nu}{N}\right]\mu_i\right)$ in (25) can be included as part of the channel matrix provided by the uplink estimation. Under these assumptions, the CS computes the ZFBF precoder as seen before, using the nominal channel matrix $\Phi[0,\nu]\mathbf{H}[0,\nu]$ at the beginning of each slot and uses it throughout the slot of M symbols. Because of the phase rotations introduced by the (time-varying) matrix $\Phi[m,\nu]$, the precoder is mismatched for m > 0 and the performance may severely degrade because of the residual multiuser interference.

Figure 3 shows the achievable rates obtained by Monte Carlo simulation assuming $\mathbf{H}[0, \nu]$ with i.i.d. elements $\mathcal{CN}(0,1)$ (normalized independent Rayleigh fading) of a 4×4 network ($N_{\rm A} = N_{\rm U} = 4$), ϵ_i i.i.d. across the users and the downlink blocks, uniform over $[-\epsilon_{\max}, \epsilon_{\max}]$, with $\epsilon_{\rm max} = 20 \times 10^{-6}$ (20 ppm). The OFDM modulation has parameters N = 64 and L = 16. The achievable rate shown here is the average rate across the block. Since the system loses synchronism progressively across each block, the achievable rate rapidly degrades as the OFDM symbol index m increases, such that the average performance is quite poor. As a comparison, we also show the performance for a corresponding ideal system without CFO and SFO (i.e., average rate at m = 0 only), and the performance of a singleuser distributed beamforming scheme (conjugate beamforming [17]) with and without CFO and SFO. Notice that single-user beamforming suffers much less from the lack of synchronization, and significantly outperforms ZFBF at high SNR. Thus, in order to achieve the promised performance of MU-MIMO it is necessary to provide the APs with sufficiently accurate estimates of the CFO and CFO such that each AP i can apply a phase de-rotation factor $\phi_{ii}^*[m,\nu]$ to its OFDM base band frequency domain symbols before transmission and effectively eliminate the effect of $\Phi[m, \nu]$. The challenge here is represented by the fact that such compensation must be done at the transmitters, while producing the transmit signal. This is a quite different and significantly more challenging problem than standard synchronization at the receiver, where one can first sample the whole signal block, and then apply some standard timing and frequency estimation/compensation. In [7] we present a family of scalable and efficient joint timing and frequency estimation/synchronization schemes that operate over the air at the network level (no requirement for a common clock distribution by wire), and are able to transfer a single very precise AP clock reference to the whole network, achieving sufficient synchronization for near-ideal

performance of the ZFBF MU-MIMO precoder, even for large networks. Furthermore, in [3], [4] we presented a Softwaredefined radio implementations that can achieve such synchronization with over-the-air signaling in a small scale prototype scenario. Therefore, if the network is properly designed, nearideal synchronization can be assumed. In the next section we make this assumption and focus on the performance of the distributed calibration protocol proposed in this paper.



Fig. 3. Achievable rates for a 4×4 distributed MIMO systems with synchronization impairments (free running local clocks at the APs).

VI. SIMULATED PERFORMANCE ANALYSIS

In this section we provide a simulation-based comparison between the calibration scheme of Argos [8], and the scheme presented in Sec. IV-B, in two different AP-location scenarios. For convenience, we refer to the calibration schemes in [8] and Sec. IV-B as "Argos calibration" and "LS-Calibration," respectively. In the first scenario we consider, the APs are co-located, thereby forming a co-located MIMO array. The second scenario involves single-antenna APs, distributed over a square grid. In particular, the distributed case involves 64 APs on an 8×8 grid, covering the square region shown in Fig. 4. In the co-located case, we assume all of the antennas in the center of the square. For both scenarios, the users' positions are drawn randomly from a uniform distribution in a square region. The longest distance between APs is chosen as 100 m.

In order to isolate the impact of the proposed relative calibration algorithms on the distributed multiuser MIMO performance, in the calibration experiments we assume perfect synchronization. The uplink and downlink channels are given by

$$Y_{i}^{\rm up}[m,\nu] = \sqrt{P}R_{i}\sum_{j=1}^{N_{\rm U}} \pi_{ij} B_{ij}[m,\nu]\tilde{T}_{j} X_{j}^{\rm up}[m,\nu] + Z_{i}^{\rm up}[m,\nu]$$
(29)

$$Y_{j}[m,\nu] = \sqrt{P}\tilde{R}_{j}\sum_{i=1}^{N_{A}} \pi_{ij} B_{ij}[m,\nu]T_{i} X_{i}[m,\nu] + Z_{j}[m,\nu],$$
(30)

where $\{Z_j[m,\nu]\}\$ are i.i.d. unit-norm circularly-symmetric complex Gaussian random variables representing the additive



Fig. 4. Sample Topology involving APs (depicted by "o") arranged on an 8×8 grid, and 16 users (depicted by "+"). The reference antenna used for Argos calibration is denoted by a •.

noise. The real-nonnegative scalar π_{ij} denotes the large-scale path gain between AP *i* and user *j* and *P* denotes the transmit power. Both for the distributed and co-located case, we assume uncorrelated, symmetric, small-scale fading by setting $\{B_{ij}[m,\nu]\}$ to be i.i.d. unit-norm circularly-symmetric complex Gaussian random variables. Although this assumption can be partially true in the distributed case, it is an oversimplification for the co-located antenna case.

Large-scale path gains between AP-AP or AP-user are both based on the WINNER Model [18] 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 (ϕ_{dB} with variance σ_{dB}^2) (Eq 31). The parameters A, B C and ϕ_{dB} are scenario dependent constants.

The pathloss of the WINNER Model is given as follows for 3 < d < 100:

$$PL(d) = A \log_{10}(d) + B + C \log_{10}(f_0/5) + \phi_{dB}.$$
 (31)

We consider here the particular indoor office scenario⁶ where $A = 18.7, B = 46.8, C = 20, \sigma_{dB}^2 = 9$ when in line of sight, otherwise $A = 36.8, B = 43.8, C = 20, \sigma_{dB}^2 = 16$. The model is specified for $3m \le d \le 100m$; for distances d < 3m, we conservatively extend the model by setting PL(d) = PL(3). This is justified partly 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 Bernoulli distribution with parameter p_ℓ which again depends on distance as follows:

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

⁶These values correspond to the A1 Indoor Office scenario with single, light walls in every path and where all of the users and APs are located on the same floor.



Fig. 5. CDF of rate of the user depicted "2" in Fig. 4 when calibration is done by Argos, LS and genie-aided calibration for various κ values. The results are obtained for $P_C = 10^3$ and $P = 5 \times 10^{10}$. From this perspective, the two full-LS curves essentially lie on top of the genie-aided scheme.

For a given transmit power, P, the large-scale path gain between AP i and user j is given by $\pi_{ij}^2 = P10^{-(\text{PL}(d_{ij})/10)}$. For the distributed case, the large-scale path gain between any set of APs and a given user are assumed to be uncorrelated.

In the case of co-located antennas, we can model the large scale fading π_{ij} between a user j and any transmit antenna i to be the same for all i, implying a small antenna array size with omni-directional antennas.

In the context of relative calibration, signals are also transmitted between APs in order to implement the relative calibration methods presented in this paper. In particular, when AP *j* transmits a pilot symbol $X_j[m, \nu]$ on the ν -th carrier and the *m*-th OFDM symbol, AP *i* receives

$$Y_{j \to i}[m, \nu] = \sqrt{P_C R_i \pi_{j \to i} B_{j \to i}[m, \nu] T_j X_j[m, \nu] + Z_{j \to i}[m, \nu]}$$
(32)

where $Z_{j\to i}[m,\nu]$ are i.i.d. unit-norm circularly-symmetric complex Gaussian random variables and P_C is the transmit power for calibration procedure. Thanks to uplink-downlink reciprocity, $B_{i\to j}[m,\nu] = B_{j\to i}[m,\nu]$.

In the distributed-AP scenario, the real-nonnegative scalar $\pi_{j \to i} = \pi_{i \to j}$ denotes the large-scale path gain between APs *i* and *j* and follows the same model used for π_{ij} . In this case $B_{j \to i}[m, \nu]$ are i.i.d. unit-norm circularly-symmetric complex Gaussian random variables.

For the co-located AP scenario, we assume $\pi_{j \to i}$ is constant and $B_{j \to i}$ are assumed to be i.i.d. Rician random variables with parameter κ . This model allows us to study the effect of SNR variations on the performance of various calibration schemes.

The hardware-induced non-reciprocal coefficients $\{R_i, T_i\}$ and $\{\tilde{R}_j, \tilde{T}_j\}$ are modeled as i.i.d. random variables, with independent magnitude uniformly distributed on $[1 - \epsilon, 1 + \epsilon]$, with ϵ chosen such that the standard deviation of the squaredmagnitudes is 0.1, and phase uniformly distributed on $[-\pi, \pi)$.



Fig. 6. CDF of users rates with relative calibration based on Argos calibration (dashed), LS calibration based on all edges (solid), and genie-aided calibration (dash-dot), for the users depicted "1" through "3" in Figure 4.

To conduct performance comparisons for different calibration schemes, we use the CDF of user rates as our metric. In particular, given a set of RF impairment coefficients drawn according to the model, 200 relative calibration realizations are performed. For each calibration realization we assume calibration pilots are exchanged, and relative calibration coefficients c_i are calculated as in Section IV. These are used over many downlink channel realizations. Notice that different calibration realizations result in different calibration coefficient estimates, and, in general, result in different user-achievable rates.

Fig. 5 illustrates the net effect of calibration on the user achievable rate for a sample user. Specifically shown in the figure are the rate CDFs of user-3 (as shown in Fig. 4) for Argos calibration, LS calibration, and genie-aided calibration for $\kappa = 0$ (Rayleigh) and $\kappa = 1000$ when ZFBF is used as the MU-MIMO precoding method. The reference antenna location used in Argos calibration is also indicated in Fig. 4. 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 and select the MU-MIMO precoder as described in Sec. IV-C.

The performance of different calibration schemes is affected by the SNR value between antennas. As reported in [8], the Argos calibration scheme requires a careful placement of the reference antenna such that it can have high enough SNR with every other antenna for effective calibration. In the Rician model, as we increase κ from 0, the channel between any two APs becomes less random. As can be seen from the figure, Argos calibration quality and the rate performance is the lowest for $\kappa = 0$ (Rayleigh). On the other hand as κ increases, performance of the Argos calibration significantly improves but still remains at a constant gap from the genie-aided scheme. Since LS-calibration does not depend on a single reference antenna, its performance is much less dependent on the κ and, as the figure reveals, it gets very close to the upper bound provided by the genie-aided calibration performance.

In Fig. 6, we compare different calibration schemes for the distributed-AP scenario with $P = P_C = 5 \times 10^8$. The three user-rate CDFs depicted in the figure, correspond to the ones associated with users 1, 2 and 3 in Figure 4. The Argos calibration scheme, originally designed for co-located antennas, performs poorly in the distributed setting. This is intuitively expected, since, in a distributed-AP setting, it is much more likely that some of the AP antennas have bad links to the reference antenna. This is exacerbated by the fact that NLOS link s between the reference AP and other APs are much more likely, increasing the likelihood of AP-to-AP links without proper calibration. On the other hand, as the figure reveals, the LS calibration scheme based on a full-connectivity graph which is a yields user-rate performance close to the genie-aided upper bound. As such, the LS calibration schemes introduced in this work provide an important "diversity benefit" in the RF calibration impairment mechanisms used to enable reciprocity-based MU-MIMO.

VII. CONCLUSION

In this paper we have developed hardware-impairment compensation techniques that can enable spatial multiplexing gains with reciprocity-based multiuser MIMO in the downlink of large-scale distributed MIMO deployments with inexpensive radios. Reciprocity-based MU-MIMO in such networks promises large aggregate spectral efficiencies with manageable CSI acquisition overheads, by exploiting uplink-downlink propagation-channel reciprocity. Since, however, radios typically have slowly (and randomly) varying, non-reciprocal, impairments in their baseband-to-RF and RF-to-baseband chains, there is no end-to-end uplink-downlink channel reciprocity. In the paper we present calibration algorithms for these impairments which enable multiplexing gains in large-scale distributed MIMO. The RF caibration methods we present are robust extensions of the Argos calibration methods [8] (originaly developed for co-located Massive MIMO deployments), and can enable spatial multiplexing gains without requiring knowledge of the user-terminal impairments or user-terminal involvement in the calibration process.

REFERENCES

- T. Marzetta, "Noncooperative cellular wireless with unlimited numbers of base station antennas," *IEEE Trans. Wireless Comm.*, vol. 9, no. 11, pp. 3590 –3600, Nov. 2010.
- [2] V. Chandrasekhar, J. Andrews, and A. Gatherer, "Femtocell networks: a survey," *IEEE Comm. Magazine*, vol. 46, no. 9, pp. 59–67, Sep. 2008.
- [3] H. V. Balan, R. Rogalin, A. Michaloliakos, K. Psounis, and G. Caire, "Achieving high data rates in a distributed MIMO system," in *Proceedings of the 18th Annual International Conference* on Mobile Computing and Networking, Mobicom '12. New York, NY, USA: ACM, 2012, pp. 41–52. [Online]. Available: http://doi.acm.org/10.1145/2348543.2348552
- [4] —, "AirSync: Enabling Distributed Multiuser MIMO with Full Spatial Multiplexing," accepted to IEEE Trans. on Networking, July 2012.
- [5] H. S. Rahul, S. Kumar, and D. Katabi, "Jmb: scaling wireless capacity with user demands," in *Proceedings of the ACM SIGCOMM 2012 Conference on Applications, Technologies, Architectures, and Protocols* for Computer Communication, New York, NY, USA: ACM, 2012, pp. 235–246. [Online]. Available: http://doi.acm.org/10.1145/2342356. 2342401
- [6] F. Kaltenberger, H. Jiang, M. Guillaud, and R. Knopp, "Relative channel reciprocity calibration in mimo/tdd systems," in *Future Network and Mobile Summit*, June 2010, pp. 1 –10.

- [7] "Scalable Synchronization and Reciprocity Calibration for Distributed Multiuser MIMO," *in preparation.*
- [8] C. Shepard, H. Yu, N. Anand, E. Li, T. Marzetta, R. Yang, and L. Zhong, "Argos: practical many-antenna base stations," in *Proceedings* of the 18th Annual International Conference on Mobile Computing and Networking, Mobicom '12. New York, NY, USA: ACM, 2012, pp. 53– 64. [Online]. Available: http://doi.acm.org/10.1145/2348543.2348553
- [9] G. Caire, N. Jindal, M. Kobayashi, and N. Ravindran, "Multiuser MIMO achievable rates with downlink training and channel state feedback," *IEEE Trans. Inform. Theory*, vol. 56, no. 6, pp. 2845–2866, 2010.
- [10] A. Hjorungnes and D. Gesbert, "Complex-valued matrix differentiation: Techniques and key results," *IEEE Trans. Signal Process.*, vol. 55, no. 6, pp. 2740 –2746, June 2007.
- [11] "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.
- [12] T. M. Schmidl and D. C. Cox, "Robust frequency and timing synchronization for OFDM," *IEEE Trans. Comm.*, vol. 45, no. 12, pp. 1613–1621, Dec. 1997. [Online]. Available: http://dx.doi.org/10.1109/ 26.650240

- [13] J.-J. van de Beek, M. Sandell, M. Isaksson, and P. Ola Borjesson, "Lowcomplex frame synchronization in ofdm systems," in *Fourth IEEE International Conference on Universal Personal Communications*, Nov. 1995, pp. 982 –986.
- [14] F. Tufvesson, O. Edfors, and M. Faulkner, "Time and frequency synchronization for OFDM using pn-sequence preambles," in *IEEE Vehicular Technology Conference*, Fall 1999, pp. 2203 –2207 vol.4.
- [15] P. Moose, "A technique for orthogonal frequency division multiplexing frequency offset correction," *IEEE Trans. Comm.*, vol. 42, no. 10, pp. 2908 –2914, Oct. 1994.
- [16] M.-O. Pun, M. Morelli, and C.-C. Kuo, "Maximum-likelihood synchronization and channel estimation for OFDMA uplink transmissions," *IEEE Trans. Comm.*, vol. 54, no. 4, pp. 726 – 736, Apr. 2006.
- [17] R. Mudumbai, D. Brown, U. Madhow, and H. Poor, "Distributed transmit beamforming: challenges and recent progress," *IEEE Comm. Magazine*, vol. 47, no. 2, pp. 102 –110, Feb. 2009.
- [18] 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