

Utility-based Downlink Pilot Assignment in Cell-Free Massive MIMO

Giovanni Interdonato^{*†}, Pål Frenger^{*}, Erik G. Larsson[†]

^{*}Ericsson Research, 581 12 Linköping, Sweden

[†]Department of Electrical Engineering (ISY), Linköping University, 581 83 Linköping, Sweden
{giovanni.interdonato, pal.frenger}@ericsson.com, erik.g.larsson@liu.se

Abstract—We propose a strategy for orthogonal downlink pilot assignment in cell-free massive MIMO (multiple-input multiple-output) that exploits knowledge of the channel state information, the channel hardening degree at each user, and the mobility conditions for the users. These elements, properly combined together, are used to define a *user pilot utility* metric, which measures the user’s real need of a downlink pilot for efficient data decoding. The proposed strategy consists in assigning orthogonal downlink pilots only to the users having a pilot utility metric exceeding a predetermined threshold. Instead, users that are not assigned with an orthogonal downlink pilot decode the data by using the statistical channel state information. The utility-based approach guarantees higher downlink net sum throughput, better support both for high-speed users and shorter coherent intervals than prior art approaches.

I. INTRODUCTION

Cell-free massive MIMO [1] comes out from the combination of massive MIMO [2], distributed architecture, and cell-less network topology. In such a system, a very large number of base station antennas, named as access points (APs) herein, are geographically distributed and connected to one (or more) central process unit (CPU) through a fronthaul network. All the APs jointly and coherently serve a smaller number of users in the same time-frequency resources.

The use of coherent cooperation as key enabling to mitigate the inter-cell interference has been fully exploited in literature, and it results in concepts known as network MIMO [3], coordinated multipoint with joint processing (CoMP-JP) [4], and multi-cell MIMO cooperative networks [5]. Cell-free massive MIMO can be considered the scalable way to implement network MIMO, since channel state information (CSI) acquisition and channel estimation can be performed locally at each AP by exploiting the *channel reciprocity* property of time-division duplex (TDD) systems [6].

Compared to conventional (cellular centralized) massive MIMO, cell-free massive MIMO guarantees an increased macro-diversity gain, since each UE receives the same signal from very different paths. Moreover, the co-processing over the APs alleviates the interference and yields higher spectral efficiency to cell-edge users [1]. Hence, each user experiences that there are no cell boundaries since it is surrounded by many serving APs (hence the term cell-free). However, these benefits come at the price of increased fronthaul network requirements.

This paper was supported by the European Union’s Horizon 2020 research and innovation programme under grant agreement No 641985 (5Gwireless).

Furthermore, cell-free massive MIMO is characterized by a low degree of *channel hardening* [7], [8].

The channel hardening property is crucial in conventional massive MIMO to eliminate the effects of small-scale fading. The channel variations averaged out when the same signal is transmitted by a massive number of co-located antennas over multiple stochastic channels. This is a direct consequence of the *law of the large numbers*. Thanks to the channel hardening, the user sees a channel that behaves almost as constant, and it can reliably decode data by using only long-term statistical CSI. Hence, downlink training that facilitates the users to acquire instantaneous CSI is not needed in massive MIMO [9].

In contrast, in cell-free massive MIMO the channel hardening is, in general, less pronounced due to the distributed architecture. Basically, only a subset of APs, i.e., the closest APs, effectively serve a given user because the other contributes are attenuated by larger path-loss. Hence, the number of APs actually involved in the transmission is smaller and the law of the large numbers is not always applicable. Consequently, users with low degree of channel hardening need to estimate the CSI in a short-term basis from downlink pilots [7]. High-speed users also need downlink pilots to acquire instantaneous CSI, regardless of the channel hardening, because of the fast variations of their channel. However, downlink pilot transmission requires further radio resources which are subtracted from the data. Hence, a resource-saving approach while guaranteeing better system performance is needed when assigning the downlink pilots to the users.

Contributions: We propose an utility-based strategy for orthogonal downlink pilot assignment in cell-free massive MIMO, and provide different variants of pilot utility metrics as function of user speed, channel hardening degree, user priority, and rate estimates. We give an analytic definition of channel hardening degree for cell-free massive MIMO systems with i.i.d. Rayleigh fading channel. A performance comparison between the proposed scheme and the prior art is provided. This study focus to the case of mutually orthogonal pilots. The case with pilot reuse will be investigated in future work.

II. PILOT ASSIGNMENT IN MASSIVE MIMO

Acquiring high quality CSI to facilitate phase-coherent processing at multiple antennas is a crucial activity in massive MIMO. By operating in TDD mode, massive MIMO benefits from the channel reciprocity property, according to which

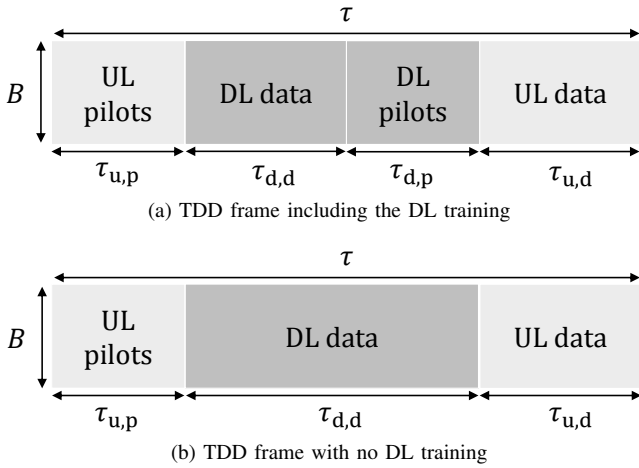


Fig. 1. The TDD frame consists of $\tau = T_c B$ symbols, where B and T_c are the coherence bandwidth and the coherence time, respectively.

the channel responses are the same in both uplink (UL) and downlink (DL). The BS can estimate the channels from a-priori known signals, called *pilots*, transmitted by the UEs in the UL (UL training). Thanks to the channel reciprocity, these estimates are then valid for both the UL and the DL. Similarly, in the DL, the UEs may need to estimate the channel from DL pilots sent by the BS (DL training).

A. TDD protocol and Resource Allocation

The TDD frame is generally designed to be smaller or equal to the channel coherence interval of all UEs. Here, we assume that a TDD frame is equal to the channel coherence interval, and that they are interchangeable terms. A TDD frame may consist of either three or four phases whether the DL training is performed or not, respectively. The TDD frame including the DL training phase consists of: (i) UL training (UL pilots); (ii) DL payload data transmission (DL data); (iii) DL training (DL pilots); and (iv) UL payload data transmission (UL data), as shown in Figure 1a. The TDD frame structure for the case with no DL training is shown in Figure 1b. The TDD frame is $\tau = T_c B$ symbols long, where B indicates the coherence bandwidth and T_c is the coherence time. Let τ_p denote the total number of symbols per TDD frame spent on transmission of pilots, $\tau_{d,d}$ and $\tau_{u,d}$ denote the number of symbols per TDD frame spent on transmission of DL and UL payload data, respectively. Hence, the length of the TDD frame is given by $\tau = \tau_p + \tau_{d,d} + \tau_{u,d}$.

In many practical systems, the pilot and data transmissions are interleaved in the time-frequency domain, as depicted schematically in Figure 2. Assuming a time-frequency grid of resource elements as in the 3GPP standards LTE and the upcoming New Radio (NR), a UL pilot (DL pilot) may consist of several resource elements (REs). Data as well as pilots in UL (DL) are transmitted during the UL data (DL data) phase. In the example in Figure 2, there are 10 UL pilots and 5 DL pilots, each consisting of 4 REs. If the UL and DL pilots have to be orthogonal in order not to generate interference, the pilot REs cannot be used for data transmission by any




Fig. 2. The data and pilot transmissions are typically interleaved in the time-frequency grid. In this example, 10 orthogonal UL pilots consumes 40 REs in the UL data (e.g., 1 pilot consists of 4 REs). The DL data contains 5 pilots consuming 20 REs.

other user. In addition, the assignment of UL and DL pilot resources is semi-static. Consequently, even when there is only one single UE active in a UL or DL data the REs reserved for the UL or DL pilots cannot be allocated for different purposes. Adapting the number of REs reserved for the pilots to the dynamic changes of the network requires a significant amount of control signaling which both consumes radio resources and introduces delay. Therefore, minimizing the pilot overhead without affecting the performance represents a relevant aspect when designing the transmission protocol.

B. Pilot Contamination and Orthogonal Pilot Transmission

Massive MIMO performance is profoundly affected by a basic phenomenon named as *pilot contamination* [10]. Pilot contamination arises when two or more UEs send the same, or in general non-orthogonal, pilot sequences. More specifically, sending non-orthogonal pilots causes mutual interference which deteriorates the quality of the corresponding channel estimates. In addition, pilot contamination does not vanish as the number of BS antennas grows large, thus it is an impairment that remains asymptotically. Ideally, every UE in a massive MIMO system should be assigned an orthogonal UL pilot sequence in order not to create interference and contaminate the UL channel estimates. However, the maximum number of orthogonal pilot sequences that can be assigned to the UEs is upper-bounded by the number of symbols in the channel coherence interval. Since this number is finite, and depending on propagation environments, the assignment of mutually orthogonal UL pilots might be physically impossible.

The radio resources become even more limited if DL pilots need to be allocated. As mentioned earlier, in massive MIMO the DL pilots are not needed because they do not introduce relevant gain in term of system capacity, by virtue of the combined effect of channel hardening and channel reciprocity. However, this holds under the assumption of low/moderate-speed user. High-speed UEs suffer from inaccurate and outdated CSI (CSI aging) as their channel conditions change very quickly. Therefore, they need to be trained in the DL to know the instantaneous CSI, regardless of the channel hardening.

Conversely, DL training is also in general more beneficial in cell-free massive MIMO due to the lower degree of channel hardening [7].

The DL training scheme introduced in [11] for conventional massive MIMO, and repurposed in [7] for cell-free massive MIMO, consists in assigning one orthogonal DL pilot per UE, and using conjugate beamforming to send the pilot, instead of simply broadcasting. This allows to sensibly reduce both the mutual pilot interference and the pilot overhead, which now scales as the number of the UEs instead of the number of antennas. However, if the number of active UEs exceeds the number of symbols that can be afforded for DL pilots, mutually orthogonal DL pilots cannot be assigned, and pilot reuse is necessary.

In this paper, we assume that only mutually orthogonal pilots can be assigned. If the number of the active UEs is larger than the number of symbols per frame reserved for pilots, then the exceeding users will not be served in the current frame but scheduled in the next frame. Adopting orthogonal pilots with no pilot reuse eliminates the pilot contamination and:

- makes the interference from DL signals towards different users possible to identify and report. If a UE experiences a large amount of DL interference, it can try to correlate the received signal with the different candidate pilot sequences and identify the interference source. The interfered UE may then report the interfering pilot index to the serving BS (or APs) which can eliminate the interference, e.g. by using more robust link adaptation when both interfered and interfering UEs are active, or by separating the UEs in time/frequency;
- enables a simpler overall implementation of the system, since no special measures need to be taken to handle the pilot contamination. The handling of the potential error cases associated with pilot reuse or non-orthogonal pilots increases the complexity at the BS/AP;
- reduces the overhead due to the pilot reconfiguration process. When pilots are reused, the BS (or the APs) may repeatedly re-assign the UL/DL pilots to reduce the pilot contamination between UEs spatially correlated (e.g., assigning orthogonal pilots to neighbor UEs and reusing the same pilot for distant UEs). With orthogonal pilots, each UE can keep the same pilot during its lifetime in the network. Hence, orthogonal pilots are *UE specific*.

III. SYSTEM MODEL

We consider a TDD cell-free massive MIMO system with M single-antenna APs and K single-antenna UEs, $M > K$. We assume both the APs and the UEs are randomly spread out in a large area without boundaries. The APs are connected to one or more CPUs through a fronthaul network with different network topologies, as shown in Figure 3. Each AP locally acquires CSI and estimates the channel from mutually orthogonal UL pilots sent by the UEs. These estimates are used in the DL for precoding, assuming perfect channel reciprocity calibration. We focus on *conjugate beamforming*, also known as *maximum-ratio transmission*, as DL precoding




Fig. 3. In cell-free massive MIMO, a large number of APs serve simultaneously a smaller number of UEs. The CPUs handle the APs cooperation.

technique. Although it does not represent the optimal precoder, performing such a linear processing offers low operational complexity with inexpensive hardware components. In addition, conjugate beamforming does not require CSI sharing among APs and CPU, with benefits on the fronthaul network load. Therefore, the CPU is responsible to collect and elaborate only payload data, and to implement advanced techniques for DL power control and DL/UL pilot assignment.

The TDD frame length is equal to the coherence interval. Hence, the channel is assumed to be static within a frame and variable independently for each frame. Let g_{mk} denote the channel gain between the m th AP and the k th user (and vice versa) defined as

$$g_{mk} = \sqrt{\beta_{mk}} h_{mk} \sim \mathcal{CN}(0, \beta_{mk}), \quad (1)$$

where h_{mk} represents the small-scale fading, and β_{mk} is the large-scale fading. We assume that $\{h_{mk}\}$ coefficients follow a Rayleigh distribution, i.e. $h_{mk} \sim \mathcal{CN}(0, 1)$ and i.i.d. RVs, for $m = 1, \dots, M$, $k = 1, \dots, K$. The large-scale fading addresses path-loss and shadow-fading. We assume the large-scale fading coefficients are estimated a-priori and known whenever required.

The AP estimates the UL channel by correlating the received UL pilot signal with the corresponding known pilot sequence and performing MMSE. The resulting UL channel estimated from the orthonormal UL pilot sequence $\varphi_k \in \mathbb{C}^{\tau_{u,p} \times 1}$, $\|\varphi_k\|^2 = 1$, is Gaussian RV distributed as $\hat{g}_{mk} \sim \mathcal{CN}(0, \gamma_{mk})$ where

$$\gamma_{mk} = \frac{\tau_{u,p} \rho_{u,p} \beta_{mk}^2}{\tau_{u,p} \rho_{u,p} \beta_{mk} + 1}, \quad (2)$$

and $\rho_{u,p}$ represents the normalized transmit signal-to-noise ratio (SNR) associated to the UL pilot symbol.

Remark 1: The transmission of orthogonal UL pilots requires $\tau_{u,p} = K$. Each AP is able to resolve K spatial dimensions since the channel estimates are orthogonal each other. The estimates are not affected by the pilot contamination.

Remark 2: γ_{mk} gives a measure of the estimation quality. In fact, $\gamma_{mk} \leq \beta_{mk}$, with equality if the estimation is perfect.

In the DL, each AP serves all the UEs implementing power control, and using the UL estimates to perform conjugate beamforming. The data signal sent by the m th AP is

$$x_m = \sqrt{\rho_d} \sum_{k=1}^K \sqrt{\eta_{mk}} \hat{g}_{mk}^* q_k, \quad (3)$$

where q_k is the unit-power data symbol, i.e., $\mathbb{E}\{|q_k|^2\} = 1$, for the k th UE, and ρ_d is the normalized transmit SNR related to the data symbol. The term \hat{g}_{mk}^* represents the *precoding factor*, and $\{\eta_{mk}\}$, $m = 1, \dots, M$, $k = 1, \dots, K$, are the power control coefficients satisfying the per-AP power constraint

$$\mathbb{E}\{|x_m|^2\} \leq \rho_d \implies \sum_{k=1}^K \eta_{mk} \gamma_{mk} \leq 1, \text{ for all } m. \quad (4)$$

The effective DL channel gain seen by the k th UE is given by

$$a_{kk'} \triangleq \sum_{m=1}^M \sqrt{\eta_{mk'}} g_{mk} \hat{g}_{mk'}^*, \quad k' = 1, \dots, K. \quad (5)$$

The k th UE can detect the data symbol q_k only if it has a sufficient knowledge of a_{kk} .

IV. ACHIEVABLE DOWNLINK RATE

In the DL, the UEs can either rely on statistical CSI or estimate the instantaneous CSI to perform data decoding. The DL channel can be estimated, as in the UL, from DL pilot beamformed by each AP, as described in [7].

A. Fading Channel with Statistical CSI at the Receiver

If each UE has knowledge of the channel statistics but not of the channel realizations, a lower bound on the DL ergodic capacity can be calculated as described in [2, Sec. 2.3.4], assuming that Gaussian codebooks are used.

An achievable DL rate of the transmission from the APs to the k th UE in the cell-free massive MIMO system with conjugate beamforming, mutually orthogonal UL pilots, for any finite M and K , is given by [1] and equal to

$$R_k^{\text{CSI}} = \log_2 \left(1 + \frac{\rho_d \left(\sum_{m=1}^M \sqrt{\eta_{mk}} \gamma_{mk} \right)^2}{\rho_d \varsigma_{kk} + \rho_d \sum_{k' \neq k} \varsigma_{kk'} + 1} \right), \quad (6)$$

where $\varsigma_{kk'} \triangleq \sum_{m=1}^M \eta_{mk'} \beta_{mk'} \gamma_{mk'}$, $k' = 1, \dots, K$. The term $\rho_d \varsigma_{kk}$ represents the so-called *beamforming uncertainty gain*. It comes from the UEs' lack of the instantaneous CSI. Hence, this gain gives an indirect measure of the channel hardening: the more the channel hardens for UE k , the smaller ς_{kk} is.

B. Fading Channel with Side Information

If each UE has imperfect knowledge of the instantaneous CSI, a lower bound on the DL ergodic capacity can be calculated by invoking the capacity-bounding technique for channel with *side information* (i.e., the channel estimates) as in [12] and [2, Sec. 2.3.5]. A valid approximation for such a DL ergodic capacity lower bound is given in [7] by observing

that the effective DL channel gains $\{a_{kk'}\}$, $k = 1, \dots, K$, can be approximated as Gaussian RVs. Hence, an achievable DL rate of the transmission from the APs to the k th UE in the cell-free massive MIMO system with conjugate beamforming, mutually orthogonal UL and DL pilots, for any finite M and K , is given by [7] as follows

$$R_k^{\text{CSI}} \approx \mathbb{E} \left\{ \log_2 \left(1 + \frac{\rho_d |\hat{a}_{kk}|^2}{\rho_d \frac{\varsigma_{kk}}{\tau_{d,p} \rho_{d,p} \varsigma_{kk} + 1} + \rho_d \sum_{k' \neq k} \varsigma_{kk'} + 1} \right) \right\}, \quad (7)$$

where $\rho_{d,p}$ represents the SNR associated to the DL pilot symbol, \hat{a}_{kk} is the estimate of the effective DL channel gain a_{kk} , given by

$$\hat{a}_{kk} = \frac{\tau_{d,p} \rho_{d,p} \varsigma_{kk} a_{kk} + \sqrt{\tau_{d,p} \rho_{d,p} \varsigma_{kk}} w_{dp,k} + \sum_{m=1}^M \sqrt{\eta_{mk}} \gamma_{mk}}{\tau_{d,p} \rho_{d,p} \varsigma_{kk} + 1}. \quad (8)$$

The expectation in (7) is taken with respect to the channel estimate \hat{a}_{kk} . In (8), the second term in the numerator is Gaussian RV as $w_{dp,k} \sim \mathcal{CN}(0, 1)$ is AWGN at the k th UE during the DL pilot signaling.

V. UTILITY-BASED DL PILOT ASSIGNMENT

A. Proposed Scheme

The detailed description of the proposed scheme is illustrated from the signaling diagram in Figure 4. All the UEs send their own UL pilot at the beginning of the coherence interval, as shown in Figure 1a. The AP estimates the UL channel and sends the corresponding CSI to the CPU. The CPU collects this information from each AP, and computes the pilot utility metric for each user. Based on this information, the CPU groups the UEs to two categories: (i) UEs requiring DL pilots, having a pilot utility metric above a predetermined threshold (UE₁ in Figure 4); (ii) UEs not requiring DL pilots, having a pilot utility metric below a predetermined threshold (UE₂ in Figure 4). Once these two groups are defined, the CPU sends to all the APs the *DL pilot configuration* message, which conveys the DL-pilot-to-UE mapping. Hence, the AP forwards part of this message (it might be just one bit) to each UE. After the data transmission, the APs send the orthogonal DL pilots only to the UEs of the first category (UE₁). UE₁ knows a priori the switching time from the DL data to the DL training phase since the TDD frame structure is determined on large time scale. If a DL pilot is assigned, the UE estimates the DL channel based on the DL pilot and demodulates the data on the remaining resources assigned to the DL transmission. If a DL pilot is not assigned, the UE assumes that the DL channel can be approximated by a constant, and it demodulates the data on all resources assigned to the DL transmission by exploiting the statistical CSI knowledge. No resources are assigned to UE₂ during the DL training phase, in order to avoid data-to-pilot interference. The TDD frame for the proposed scheme is depicted in Figure 5.

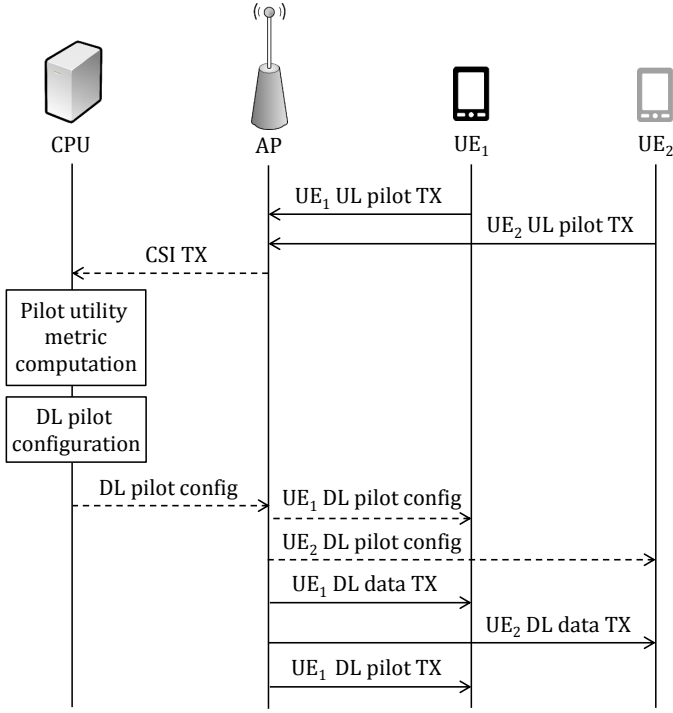


Fig. 4. In the utility-based DL pilot assignment strategy, the CPU collects from the APs all the information needed to compute the pilot utility metric. Afterwards, it sends the DL pilot config message to the APs to inform which UE is assigned an orthogonal DL pilot.




Fig. 5. In the utility-based DL pilot assignment strategy, the DL training phase is performed only for those users that need to estimate the instantaneous CSI the most (UE₁, in this example). The rest of the UEs (UE₂, in this example) rely on statistical CSI. In any case, $\tau_{d,p}$ symbols cannot be used for data, in order not to create data-to-pilot interference.

The achievable DL rate, subject to the assumption that the UEs transmit Gaussian message-bearing symbols, for the utility-based DL pilot assignment scheme is given by

$$R_k^{\text{ubPA}} = \begin{cases} R_k^{\text{iCSI}}, & \text{if UE } k \in \text{UE}_1, \\ R_k^{\text{sCSI}}, & \text{otherwise.} \end{cases} \quad (9)$$

B. Pilot utility metric

The pilot utility metric can be defined in different manner according to the purposes. It is a function of the UE mobility,

the channel hardening degree, and the CSI. In general, it must guarantee an orthogonal DL pilot to: (i) high-speed UEs, which necessarily need the instantaneous CSI to face the CSI aging; (ii) low/moderate-speed UEs experiencing a low degree of channel hardening, which require instantaneous CSI to face the unreliability of the statistical CSI. On the other hand, all low/moderate-speed UEs for which the channel hardens sufficiently will not require DL training, and they may rely only on statistical CSI. Hence, UEs may be assigned with an orthogonal DL pilot if the Doppler spread value is above a pre-determined threshold and/or the channel hardening degree is below a pre-determined threshold. Such a pilot utility metric can be defined as

$$\text{pu}_k = wD_k + (1-w)(1 - \text{ChD}_k), \quad (10)$$

where $D_k \in [0, 1]$ is the Doppler spread value, ChD_k indicates the *channel hardening degree* at the k th UE, and $w \in [0, 1]$ is the weight to prioritize differently the UE mobility and the channel hardening degree.

The channel hardening degree is defined as the ratio between the instantaneous channel gain and its average value. By following a similar methodology as in [9], we obtain the channel hardening degree expression at the k th UE for a cell-free massive MIMO system,

$$\text{ChD}_k = 1 - \frac{\text{Var} \left\{ \sum_{m=1}^M |g_{mk}|^2 \right\}}{\left(\mathbb{E} \left\{ \sum_{m=1}^M |g_{mk}|^2 \right\} \right)^2}. \quad (11)$$

The larger ChD_k is, the more the channel hardens for UE k .

The pilot utility metric may also account for the UE priority in terms of *Quality of Service* (QoS) requirements. This needs the APs collect and send to the CPU the QoS requirements for each UE, and the CPU determines the UE priorities accordingly. In this case, the pilot utility metric is

$$\text{pu}_k = \begin{cases} wD_k + (1-w)(1 - \text{ChD}_k) + \alpha_k, & \text{or} \\ \alpha_k [wD_k + (1-w)(1 - \text{ChD}_k)], \end{cases} \quad (12)$$

where $\alpha_k \in [0, 1]$ is the UE priority. The larger α_k is, the more QoS the UE k requires.

Alternatively to the channel hardening degree estimation, a DL pilot may be assigned to those users that would increase their rates the most by taking advantage of the DL pilot. In this case, the pilot utility metric may be defined as a function of the achievable rates (or throughputs) estimated by using the CSI knowledge at each AP. The following pilot utility metrics

$$\text{pu}_k = \begin{cases} \alpha_k [wD_k + (1-w) (R_k^{\text{iCSI}} - R_k^{\text{sCSI}})], \\ \alpha_k [wD_k + (1-w) (T_k^{\text{iCSI}} - T_k^{\text{sCSI}})], \\ \alpha_k [wD_k + (1-w) \left(\frac{R_k^{\text{iCSI}} - R_k^{\text{sCSI}}}{R_k^{\text{iCSI}}} \right)], \\ \alpha_k [wD_k + (1-w) \left(\frac{T_k^{\text{iCSI}} - T_k^{\text{sCSI}}}{T_k^{\text{iCSI}}} \right)], \\ \alpha_k [wD_k + (1-w) (R_k^{\text{sCSI}})^{-1}], \end{cases} \quad (13)$$

are valid alternatives, according to different aims: absolute rate increase, absolute throughput increase, relative rate increase, relative throughput increase, and low-rate prioritization, respectively. The *per-user net throughput* (bit/s) is proportional to the per-user rate, and takes into account the performance loss due to the DL and UL pilots transmission (as they subtract resources to the data). For a symmetric TDD frame, i.e., $\tau_{u,d} = \tau_{d,d}$, it is given by

$$T_k = \frac{B}{2} \left(\frac{1 - \tau_p}{\tau} \right) R_k, \quad (14)$$

where B is the bandwidth, and τ_p is the pilot overhead.

VI. NUMERICAL RESULTS

Our simulations aim to compare the performance, in terms of DL net sum throughput and average DL net per-user throughput, provided by our solution and the existing solutions of the state-of-the-art. The DL net sum throughput, measured in bits/sec, is equal to the sum of the DL net throughputs per UE defined in (14), and it is given by $T = \sum_{k=1}^K T_k$. Hence, the average DL net per-user throughput is $T_{\text{avg}} = T/K$.

A. Simulation Scenario

Assuming the channel model in (1), we describe the path-loss and shadow-fading considered in our performance evaluation. The coefficients $\{\beta_{mk}\}$ are defined as follows

$$\beta_{mk} = \text{PL}_{mk} \cdot 10^{\frac{\sigma_{sh} z_{mk}}{10}}, \quad (15)$$

where PL_{mk} represents the path loss, $10^{\frac{\sigma_{sh} z_{mk}}{10}}$ is the shadow-fading with standard deviation σ_{sh} , and $z_{mk} \sim \mathcal{N}(0, 1)$, i.e., we assume uncorrelated shadow-fading. The path-loss is modeled as three-slope [1], where the loss exponent can assume three different values according to the UE-to-AP distance. More specifically, the loss exponent is 3.5 if the distance between UE k and AP m (d_{mk}) is greater than a first reference distance d_1 ; it is equal to 2 if d_{mk} is greater than a second reference distance d_0 and less or equal to d_1 . In the case d_{mk} is less than d_0 the loss exponent is 0. The chosen propagation model is the Hata-COST231. Hence, the path-loss in dB is given by

$$\text{PL}_{mk} = \begin{cases} -L - 35\log_{10}(d_{mk}) & \text{if } d_{mk} > d_1, \\ -L - 15\log_{10}(d_1) - 20\log_{10}(d_{mk}) & \text{if } d_0 < d_{mk} \leq d_1, \\ -L - 15\log_{10}(d_1) - 20\log_{10}(d_0) & \text{if } d_{mk} \leq d_0, \end{cases} \quad (16)$$

where

$$L = 46.3 + 33.9 \log_{10}(f) - (1.1 \log_{10}(f) - 0.7) h_{\text{UE}} - 13.82 \log_{10}(h_{\text{AP}}) + (1.56 \log_{10}(f) - 0.8). \quad (17)$$

In our simulations, $d_1 = 50$ meters, $d_0 = 10$ meters, the carrier frequency f (in MHz) is 2 GHz, the height of the UE and AP antennas are 1.65 meters and 5 meters, respectively. In addition, the shadow-fading coefficients have standard deviation 8 dB.

The other simulation parameters are set as follows: $M = 200$ APs and $K = 50$ UEs are randomly and uniformly

TABLE I
SIMULATION SETTINGS

Description	Value	Description	Value
APs/UEs	unif. rand. distr.	M	200
simulation area	1 km ²	K	50
τ	200 symbols	time-slot	1 ms
path-loss	three-slope	d_1, d_0	50, 10 m
carrier frequency	2 GHz	bandwidth B	20 MHz
shadow-fading	uncorrelated	antenna gain	0 dBi
shadow-fading std	8 dB	noise figure	9 dB
AP antenna height	5 m	$\bar{\rho}_d, \bar{\rho}_{d,p}$	200 mW
UE antenna height	1.65 m	$\bar{\rho}_u, \bar{\rho}_{u,p}$	100 mW

distributed within a square of size 1 km²; the bandwidth B is 20 MHz; the antenna gains are 0 dBi; the noise figure (both for UL and DL) is 9 dB; the radiated power for data and pilot is 200 mW for APs, and 100 mW for UEs ($\bar{\rho}_d, \bar{\rho}_{d,p}, \bar{\rho}_u, \bar{\rho}_{u,p}$, respectively). The corresponding normalized transmit SNRs ($\rho_{d,p}, \rho_d, \rho_u, \rho_{u,p}$), defined in Section III, are obtained by dividing the radiated powers by the noise power, which is given by

$$\text{noise power} = B \times k_B \times T_0 \times \text{noise figure (W)},$$

where k_B is the Boltzmann constant, and $T_0 = 290$ (Kelvin) is the noise temperature. We take the length of the TDD frame $\tau = 200$ symbols, which corresponds to a coherence bandwidth of 200 kHz and a time-slot of 1 ms. To simulate a cell-free network topology (i.e., no cell-edge effects), we wrap the simulation area around with eight twin neighbor areas. The UL and DL orthogonal pilot sequences are randomly assigned to the users. Finally, the power control coefficients $\{\eta_{mk}\}$ are set by performing *max-min fairness power control*, as described in [1, Section IV-B]. The simulation settings are summarized in Table I.

B. Performance Evaluation

Firstly, we measure the channel hardening degree in such a cell-free massive MIMO system. The cumulative distribution function (cdf) of the channel hardening degree is shown in Figure 6, for cell-free massive MIMO and the cellular centralized massive MIMO. The cdf gives more insights about the distribution of the channel hardening degree given the random variables of the AP/UE positions in the deployment area. In cell-free massive MIMO the channel hardening degree is given by (11). Instead, in conventional massive MIMO, for i.i.d. Rayleigh fading channels, the channel hardening degree simply scales as $1/M$ [9], and in this case it is given by $1 - 1/M = 0.995$, regardless of the AP/UE position. Figure 6 confirms that in cell-free massive MIMO the channel hardening degree is significantly lower than in conventional massive MIMO.

Next, we introduce an example to show the benefits provided by our solution with respect to the state-of-the-art. Let us consider the following case studies:

- 1) *statistical CSI* (sCSI), prior art: the coherence interval is structured as in Figure 1b. No DL training is performed




Fig. 6. The ChD in cell-free massive MIMO is much lower compared to that in conventional massive MIMO. Here, we assume 200 APs randomly distributed in an area of 1 km², and i.i.d. Rayleigh fading channel.

and all the active UEs rely only on statistical CSI to decode data. Each UE is assigned with an orthogonal UL pilot. Hence, $\tau_{u,p} = 50$, $\tau_{d,p} = 0 \implies \tau_p = 50$.

- 2) *instantaneous CSI* (iCSI), prior art: the coherence interval is structured as in Figure 1a. The DL training is performed for all UEs. Each UE is assigned with an orthogonal UL pilot and receives an orthogonal DL pilot from which it estimates the instantaneous CSI. Hence, $\tau_{u,p} = \tau_{d,p} = 50 \implies \tau_p = 100$.

- 3) *utility-based pilot assignment* (ubPA), proposed scheme: the coherence interval is structured as in Figure 5. Each UE is assigned with an orthogonal UL pilot. All the UEs having a pilot utility metric exceeding a predetermined threshold are assigned with an orthogonal DL pilot. The remaining UEs rely only on statistical CSI.

Remark 3: In a realistic scenario, the assignment of UL/DL resources occurs in a quasi-static fashion, since a dynamic re-allocation of pilots, frame by frame, requires a huge amount of control signaling to let the UEs know the exact UL/DL switching time. This also introduces delay. Therefore, the predetermined threshold is actually a design setting varying in large time scale.

In view of the above, let us assume the following UL/DL resources assignment: $\tau_{u,p} = 50$, $\tau_{d,p} = 25 \implies \tau_p = 75$. Hence, according to the proposed scheme, only the 25 UEs out of 50 with the highest pilot utility metric receive an orthogonal DL pilot. The remaining 25 UEs have no allocated resource during the DL training phase, in order to avoid further interference, as shown in Figure 5.

With the purpose to maximize the DL net sum throughput, we set the pilot utility metric as in the first expression of (13) (absolute rate increase). For simplicity, we finally assume that the length of the TDD frame supports the speed of all the




Fig. 7. Performance comparison between the sCSI, iCSI, and ubPA scheme. The utility-based pilot assignment strategy guarantees higher average DL net per-user throughput than the state-of-the-art schemes.

active users in the network, and all the UEs have equal QoS priority, that is $w = 0$, $\alpha_k = 1$, respectively.

Remark 4: Since the number of UEs that can be served is proportional to the time spent sending pilots, while the sum throughput is proportional to the number of UEs served, it follows that a good rule to allocate the UL/DL pilot resources is simply setting the number of symbols spent on transmission of UL and DL pilots not exceeding half the length of the coherence interval.

According to this, and the simulation scenario, i.e., $\tau = 200$, the maximum number of symbols dedicated for the pilots is $\tau_p = 100$. Under this assumption, in the iCSI case all the UEs are served in the current frame. Conversely, if the number of UEs would be $K > 50$, then $\tau_p > 100$, and in the iCSI case the exceeding UEs would be scheduled in the next frame.

The simulation results in Figure 7 show that ubPA scheme performs better than the prior art schemes in terms of average net per-user throughput, providing about 12% improvement over case sCSI, and about 16% improvement over case iCSI.

The gain introduced by the proposed scheme in terms of DL net sum throughput is shown in Figure 8. As we can see, the ubPA scheme provides a gain over scheme sCSI by about 11%, 12% and 13% at the 5th percentile, median and 90th percentile, respectively. These gains are larger over scheme iCSI: 17% at the 5th percentile, 16% at the median point and 15% at the 90th percentile. The results also show that in such a scenario with many APs and UEs, and short coherence interval, transmitting one orthogonal DL pilot per UE (scheme iCSI) is less efficient than letting the UEs rely on statistical CSI (scheme sCSI). In fact, the performance loss due to the pilots transmission overhead overcomes the benefits, in terms of higher effective SINRs, introduced by the DL training.

Lastly, we compare the DL net sum throughputs provided by three different pilot utility metrics: (i) absolute rate increase,




Fig. 8. Cdf of the DL net sum throughput provided by the sCSI, iCSI, and ubPA scheme. The utility-based pilot assignment strategy guarantees higher DL net sum throughput than the state-of-the-art schemes.




Fig. 9. Cdf of the DL net sum throughput provided by the absolute rate and throughput increase, and ChD-aware pilot utility metric. The absolute rate increase approach guarantees higher performance than other approaches.

first expression in (13); (ii) absolute throughput increase, second expression in (13); (iii) ChD-aware pilot utility, second expression in (12). As shown in Figure 9, the absolute rate increase pilot utility metric guarantees the highest DL net sum throughput even though the gain over the absolute throughput increase approach is negligible.

On the other hand, the gain with respect to the ChD-aware approach may reach up to 5%. Estimating the rates/throughputs can give more information about the UE's real need of a DL pilot rather than estimating the ChD. In fact, the

pilot utility metric based on the rate/throughput estimates takes into account both the ChD by estimating the beamforming uncertainty gain, and the amount of inter-user interference. Both factors affect user's performance and they need to be jointly evaluated when assigning the DL pilots.

VII. CONCLUSION

In this paper, we proposed a strategy for orthogonal DL pilot assignment based on the so-called pilot utility metric, which is a function of the channel hardening degree or the rate estimate, user mobility, user QoS priority, and CSI.

The proposed strategy consists in assigning the orthogonal DL pilots only to the users that benefits the most when decoding data by exploiting the instantaneous CSI. In general, this approach guarantees an orthogonal DL pilot to users with high mobility, to cope with the CSI aging, and users with low/moderate mobility experiencing low channel hardening, for which the reliance on statistical CSI in the decoding yields poor performance. The utility-based pilot assignment scheme guarantees higher DL per-user and sum throughput, better support for high-speed users and shorter coherent interval compared to the state-of-the-art.

REFERENCES

- [1] H. Q. Ngo, A. Ashikhmin, H. Yang, E. G. Larsson, and T. L. Marzetta, "Cell-free massive MIMO versus small cells," *IEEE Transactions on Wireless Communications*, vol. 16, no. 3, pp. 1834–1850, March 2017.
- [2] T. L. Marzetta, E. G. Larsson, H. Yang, and H. Q. Ngo, *Fundamentals of Massive MIMO*. Cambridge University Press, 2016.
- [3] 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 Transactions on Information Theory*, vol. 58, no. 5, pp. 2911–2934, May 2012.
- [4] R. Irmer, H. Droste, P. Marsch, M. Grieger, G. Fettweis, S. Brueck, H. P. Mayer, L. Thiele, and V. Jungnickel, "Coordinated multipoint: Concepts, performance, and field trial results," *IEEE Communications Magazine*, vol. 49, no. 2, pp. 102–111, February 2011.
- [5] D. Gesbert, S. Hanly, H. Huang, S. S. Shitz, O. Simeone, and W. Yu, "Multi-cell MIMO cooperative networks: A new look at interference," *IEEE Journal on Selected Areas in Communications*, vol. 28, no. 9, pp. 1380–1408, December 2010.
- [6] E. Björnson, R. Zakhour, D. Gesbert, and B. Ottersten, "Cooperative multicell precoding: Rate region characterization and distributed strategies with instantaneous and statistical CSI," *IEEE Transactions on Signal Processing*, vol. 58, no. 8, pp. 4298–4310, Aug 2010.
- [7] G. Interdonato, H. Q. Ngo, E. G. Larsson, and P. Frenger, "How much do downlink pilots improve cell-free massive MIMO?" in *2016 IEEE Global Communications Conference (GLOBECOM)*, Dec 2016, pp. 1–7.
- [8] Z. Chen and E. Björnson, "Channel hardening and favorable propagation in cell-free massive MIMO with stochastic geometry," *CoRR*, 2017. [Online]. Available: <http://arxiv.org/abs/1710.00395>
- [9] H. Q. Ngo and E. G. Larsson, "No downlink pilots are needed in TDD massive MIMO," *IEEE Transactions on Wireless Communications*, vol. 16, no. 5, pp. 2921–2935, May 2017.
- [10] J. Jose, A. Ashikhmin, T. L. Marzetta, and S. Vishwanath, "Pilot contamination and precoding in multi-cell TDD systems," *IEEE Transactions on Wireless Communications*, vol. 10, no. 8, pp. 2640–2651, Aug. 2011.
- [11] H. Q. Ngo, E. G. Larsson, and T. L. Marzetta, "Massive MU-MIMO downlink TDD systems with linear precoding and downlink pilots," in *Proc. Allerton Conference on Communication, Control, and Computing*, Oct. 2013, pp. 293–298.
- [12] M. Medard, "The effect upon channel capacity in wireless communications of perfect and imperfect knowledge of the channel," *IEEE Transactions on Information Theory*, May 2000.