

Subset MMSE Receivers for Cell-Free Networks

Masoud Attarifar, *Student Member, IEEE*, Aliazam Abbasfar, *Senior Member, IEEE*,
Angel Lozano, *Fellow, IEEE*

Abstract—This paper formulates linear MMSE receivers that are both network- and user-centric for the uplink of cell-free wireless networks with centralized processing. Precisely, every user's reception involves a distinct subset of access points (APs) while every AP participates in the reception of a distinct subset of users, hence the moniker *subset MMSE receivers*. These subsets, defined on the basis of the large-scale channel gains between users and APs, capture the most relevant signal and interference contributions while disregarding those whose processing is cost-ineffective and whose associated channel estimations would incur unnecessary overheads. With that, subset reception approaches the performance of network-wide MMSE reception, offering a multiple-fold improvement over cellular and matched-filtering counterparts, while being scalable in terms of cost and channel estimation. Moreover, because the subsets overlap considerably, they can sometimes be advantageously combined and the computation of the corresponding receivers can share a hefty amount of processing.

Index Terms—Cell-free networks, C-RAN, MMSE, ultradense networks, user-centric networks

I. INTRODUCTION

The reigning paradigm in the organization of wireless networks is the cellular structure, in which every user is served by a single access point (AP) while being regarded as interference by the rest of APs. This paradigm has served us well. By virtue of the pathloss exponent η being substantially above 2 (typically $\eta \approx 3.5$ –4) almost everywhere, universal or quasi-universal frequency reuse is possible without other-cell interference becoming overwhelming. Networks have densified drastically over the years, always within the confines of the cellular paradigm, with only limited add-ons to allow for some cooperation among adjacent cells [1]–[3].

Densification is bound to reach a transition point beyond which the pathloss behaves differently, with $\eta \approx 2$ over the span of multiple cells and with an ensuing surge in other-cell interference [4]–[7]. This motivates the interest in alternative forms of network organization, and specifically in cell-free structures where every user is potentially served by every AP. These can be interpreted as deconstructions of a cellular structure whereby the associations between users and cells are released, or else as cellular networks whose cells are truly and seamlessly cooperative [8]–[11].

Capitalizing on extensive fronthaul, cell-free networks promise a host of advantages over their cellular counterparts:

- A reinterpretation of the notion of interference that turns the aforescribed pathloss behavior into a virtue. Rather than suffer stronger interference, a cell-free network can potentially recover additional signal power. Densification can proceed beyond the point where cellular networks might collapse under their own interference.
- Large-scale diversity by virtue of the fact that each user is not linked to a single AP, but rather to many, facilitating a more uniform service.
- Amenability to cloud-based implementations of the base-band processing, so-called C-RAN, with all the benefits of software-defined implementations, the convenience of maintenance and upgrades, and the elasticity in the use of resources [12]–[19].

Altogether, the research interest in cell-free networks is expanding [20]–[31] and it is reasonable to anticipate that, at least for ultra-dense deployments, cell-free structures might become the norm [32].

Focusing on the uplink, the AP observations must be combined to extract the user signals, hence the processing inevitably requires some degree of centralization [33]. A basic combining method is matched filtering, which amounts to desired-signal beamforming and performs satisfactorily when the number of users per time-frequency signaling resource is much smaller than the number of APs; as this is reminiscent of massive MIMO, cell-free networks with matched filtering can be viewed as a deconstruction of cellular massive MIMO.

By embracing complete centralization, possibly edge-cloud-based, more sophisticated receivers can be envisioned and, with those, a drastically higher number of users. Specifically, the optimum linear receiver—in the sense of maximizing the signal-to-interference-plus-noise ratio (SINR), and subsequently the spectral efficiency—is the MMSE receiver. A network-wide MMSE receiver, however, is not scalable and, in large deployments with thousands of APs and users, it becomes outright unfeasible. Moreover, a network-wide MMSE receiver is an unnecessary overkill because (i) a user's transmission elicits negligible signal power at faraway APs, and it causes negligible interference to faraway users, and (ii) reliable channel estimates at faraway APs are so burdensome overhead-wise that it is preferable to forgo those estimations altogether [34]. This suggests a degree of confinement in the channel estimation and in the data collection, with the aim of approaching the network-wide MMSE performance in a scalable manner. Important advances in this direction have been made in [35]–[42], which discuss curtailing the users and APs whose observations should be jointly processed. The present paper builds on these works as follows:

- Linear MMSE receivers that are both AP-centric and

M. Attarifar and A. Abbasfar are with the University of Tehran, North Kargar St., Tehran, Iran (e-mail: m.attarifar@ut.ac.ir; abbasfar@ut.ac.ir). Their work is supported by the Iranian National Science Foundation (INSF) Grant #98012057.

A. Lozano is with Univ. Pompeu Fabra, 08018 Barcelona (e-mail: angel.lozano@upf.edu). His work is supported by the European Research Council under the H2020 Framework Programme/ERC grant agreement 694974, and by MINECO's Projects RTI2018-102112 and RTI2018-101040.

user-centric, complete with AP and user subset selection policies, are formulated and evaluated on a very large network. Since they involve subsets of APs and users, we compactly refer to them as *subset MMSE receivers*.

- The tradeoff is established between performance, computational cost, and number of channel coefficients to estimate (which, in turn, determines the pilot overhead). Subset receivers are shown to be scalable, in the sense that their cost and channel estimation requirements do not grow faster than linearly with the size of the network. In large networks, subset receivers approach the network-wide MMSE performance with a tiny fraction of its cost and its channel estimation requirements.
- Two schemes are proposed that exploit the overlap among subsets to further improve the tradeoff between performance and cost while drawing connections with existing notions of cell clustering [43]–[52].

The document is organized as follows. Section II describes the network and channel models that the formulation adheres to. Then, Section III establishes the scalability criteria adopted in the sequel, Section IV succinctly derives the channel-estimation-based receiver observations, and Section V lays down various benchmarks, chiefly the network-wide MMSE receiver. The heart of the paper is Section VI, where subset receivers are described and interpreted, subset selection policies are proposed and validated, and scalability is discussed; the section finalizes with a large-network evaluation. Subsequently, Section VII elaborates on further improvements that capitalize on the overlaps among subsets, and Section VIII concludes the paper.

II. NETWORK AND CHANNEL MODELS

The networks under consideration feature N APs and K users, all equipped with a single omnidirectional antenna.

Some time-frequency resource units are reserved for pilot transmissions from the users, based on which the channels are estimated by the APs. The remaining resource units are available for data transmission.

A. Large-scale Modeling

The conjunction of distance-dependent pathloss with exponent η and shadowing gives rise to a local-average channel gain $G_{n,k}$ between the k th user and the n th AP.

Under the premise of AP positions that are agnostic to the radio propagation, shadow fading renders the network approximately Poisson-like from the vantage of any user [53]–[55]. This approximation sharpens as the shadowing strengthens, being precise for values of interest [54]–[57]. Relying on this result, we draw the AP positions uniformly; these positions thus conform to a binomial point process, which, as the network grows, converges to a Poisson point process, avoiding the explicit modeling of shadowing as it is then already captured by the network geometry. The user positions are also drawn uniformly.

Letting P and σ^2 denote a user's radiated power and the noise power, respectively, the local-average SNR of user k at the n th AP is $\text{SNR}_{n,k} = G_{n,k}P/\sigma^2$. We set P/σ^2 such that,

for $\eta = 4$, $\text{SNR}_{n,k} = 25$ dB at a distance d where d would be the inter-AP distance if the network were arranged as a hexagonal grid with the same spatial density that our network. Under reasonable values for the transmit power, bandwidth, and pathloss intercept [58], this is compatible with ultradense deployments ($d \approx 5$ –20 m) and it would ensure interference-limited operation should the network operate in cellular mode.

The local-average SNRs, and all the large-scale parameters, are stable over each network snapshot.

B. Small-scale Modeling

Besides $G_{n,k}$, the channel between the k th user and the n th AP features a small-scale fading coefficient $h_{n,k} \sim \mathcal{N}_C(0, 1)$, independent across users and APs. The fading coherence can be visualized as tiles of size $T_c \times B_c$ in the time-frequency plane, a block-wise structure that has been shown to be equivalent to smooth autocorrelation functions as far as the linear MMSE estimation of the fading is concerned [59], [60]. For any snapshot of the large-scale parameters, under the premise that channel estimation errors and interference are treated by the decoder as additional Gaussian noise, user k can achieve a spectral efficiency of [61]

$$C_k = \mathbb{E}[\log_2(1 + \text{sinr}_k)] \quad (1)$$

where sinr_k denotes the SINR of user k and the expectation is over the small-scale fading. Adding (1) over all K users, we obtain the sum spectral efficiency. Unless otherwise stated, the spectral efficiencies in this paper are *gross*, meaning that pilot overheads are yet to be subtracted out.

III. SCALABILITY

We consider scalability in terms of those aspects that are inherent to cell-free networks, namely (i) receiver obtainment cost, (ii) linear combining cost, and (iii) channel estimation. The decoding and remaining post-processing tasks are as in a cellular network, one chain per user, hence inherently scalable.

We measure the cost in number of complex multiply-and-accumulate (MA) operations per time-frequency coherence block, respectively M_{rx} and M_c for the obtainment of the receivers and for the linear combining, while we denote by L the number of channel coefficients to be estimated per coherence block. For $K, N \rightarrow \infty$ with fixed K/N , we want M_{rx}/N , M_c/N , and L/N , to be $\mathcal{O}(1)$ thereby ensuring that, as the network grows, the per-AP cost and overhead remain roughly constant, as in a cellular system.

Measured in MA operations, the cost of $N \times N$ matrix inversions or multiplications is αN^ν for $2.37 \leq \nu \leq 3$ [62]. Likewise, the product of $N \times K$ and $K \times N$ matrices costs $\alpha K^{\nu-2} N^2$. In both cases, a naïve procedure leads to $\nu = 3$; pushing the exponent below $\nu = 3$ comes at some expense in numerical stability, in memory, and in an increased factor α .

We hasten to emphasize that our goal is not to present a detailed complexity analysis, which would require positing specific implementations, but rather to establish scalability and to gauge the magnitude of the cost of each solution. With this in mind, these measures suffice and simpler operations such as additions can be neglected.

IV. CHANNEL ESTIMATION AND DATA TRANSMISSION

Disregarding pilot contamination for now, the linear MMSE channel estimate $\hat{h}_{n,k}$ gathered by the network upon observation at the n th AP of a pilot transmission from user k satisfies $h_{n,k} = \hat{h}_{n,k} + \tilde{h}_{n,k}$ where

$$\mathbb{E}[|\hat{h}_{n,k}|^2] = \frac{\text{SNR}_{n,k}}{1 + \text{SNR}_{n,k}} \quad (2)$$

while

$$\tilde{h}_{n,k} \sim \mathcal{N}_{\mathbb{C}}\left(0, \frac{1}{1 + \text{SNR}_{n,k}}\right) \quad (3)$$

is independent error. Generalizing (2)–(3) to multiple—as opposed to one—pilot transmissions per user and per fading coherence block, or to power-booster pilots, entails a straightforward scaling of $\text{SNR}_{n,k}$ therein [63]. Ultimately, the overhead corresponding to the resource units consumed by pilots must be discounted from the gross spectral efficiency.

Subsequently, upon data transmission, on a given time-frequency resource unit the n th AP observes

$$y_n = \sum_{k=0}^{K-1} \sqrt{G_{n,k}} h_{n,k} x_k + v_n \quad (4)$$

$$= \sum_{k=0}^{K-1} \sqrt{G_{n,k}} \hat{h}_{n,k} x_k + \underbrace{\sum_{k=0}^{K-1} \sqrt{G_{n,k}} \tilde{h}_{n,k} x_k}_{z_n} + v_n \quad (5)$$

where x_k is the signal from user k , satisfying $\mathbb{E}[|x_k|^2] = P$, while $v_n \sim \mathcal{N}_{\mathbb{C}}(0, \sigma^2)$.

The transmit-receive relationship between users and APs can be vectorized into

$$\mathbf{y} = \begin{bmatrix} y_0 \\ \vdots \\ y_{N-1} \end{bmatrix} = \sum_{k=0}^{K-1} \hat{\mathbf{c}}_k x_k + \mathbf{z} \quad (6)$$

where

$$\hat{\mathbf{c}}_k = \begin{bmatrix} \sqrt{G_{0,k}} \hat{h}_{0,k} \\ \vdots \\ \sqrt{G_{N-1,k}} \hat{h}_{N-1,k} \end{bmatrix} \quad \mathbf{z} = \begin{bmatrix} z_0 \\ \vdots \\ z_{N-1} \end{bmatrix} \quad (7)$$

and $\mathbb{E}[\mathbf{z}\mathbf{z}^*] = \sigma^2 \mathbf{D}$ with \mathbf{D} a diagonal matrix whose (n, n) th entry is

$$[\mathbf{D}]_{n,n} = 1 + \sum_{k=0}^{K-1} \frac{\text{SNR}_{n,k}}{1 + \text{SNR}_{n,k}}. \quad (8)$$

V. BENCHMARKS: CELLULAR, MATCHED-FILTERING, AND NETWORK-WIDE MMSE

The most basic benchmark is a cellular network where each user connects with the AP to which it has the strongest large-scale gain. If multiple users select the same AP, then multiple access in time and/or frequency takes place.¹ The user or users connected to a given AP are interfered by all users connecting to other APs.

¹This is somewhat different from the cellular operation in [20], where, once an AP has been selected by a user, it becomes unavailable, forcing other users to connect to more distant APs.

A second baseline is provided by matched-filtering cell-free operation. The matched filter for user k is $\mathbf{w}_k^{\text{MF}} = \hat{\mathbf{c}}_k$. For given fading realizations, $\mathbf{w}_0^{\text{MF}}, \dots, \mathbf{w}_{K-1}^{\text{MF}}$ yield

$$\begin{aligned} \text{sinr}_k^{\text{MF}} &= \frac{\mathbb{E}[|\mathbf{w}_k^{\text{MF}*} \hat{\mathbf{c}}_k x_k|^2 | \hat{\mathbf{c}}_k]}{\sum_{\ell \neq k} \mathbb{E}[|\mathbf{w}_k^{\text{MF}*} \hat{\mathbf{c}}_\ell x_\ell|^2 | \hat{\mathbf{c}}_\ell] + \mathbb{E}[|\mathbf{w}_k^{\text{MF}*} \mathbf{z}|^2]} \\ &= \frac{|\hat{\mathbf{c}}_k^* \hat{\mathbf{c}}_k|^2}{\sum_{\ell \neq k} |\hat{\mathbf{c}}_k^* \hat{\mathbf{c}}_\ell|^2 + \frac{\sigma^2}{P} \hat{\mathbf{c}}_k^* \mathbf{D} \hat{\mathbf{c}}_k} \quad k = 0, \dots, K-1. \end{aligned} \quad (9)$$

Since $\hat{\mathbf{c}}_0, \dots, \hat{\mathbf{c}}_{K-1}$ are readily available, the receiver obtainment cost is nil. However, applying the length- N matched filters $\mathbf{w}_0^{\text{MF}}, \dots, \mathbf{w}_{K-1}^{\text{MF}}$ to \mathbf{y} entails $M_c = KN$ complex MA operations, hence the linear combining cost per AP is $M_c/N = K$, which does not scale. Likewise, the number of channel estimates per AP is $L/N = K$, also not scalable. The matched-filter solution can be rendered scalable by curbing the number of APs serving each user [35], [38]–[41].

Now, let $\hat{\mathbf{C}} = [\hat{\mathbf{c}}_0 \dots \hat{\mathbf{c}}_{K-1}]$. The network-wide MMSE receiver for all K users is [64, sec. 6.4]

$$\mathbf{W}^{\text{MMSE}} = (\hat{\mathbf{C}} \hat{\mathbf{C}}^* + \frac{\sigma^2}{P} \mathbf{D})^{-1} \hat{\mathbf{C}} \quad (10)$$

$$= \mathbf{D}^{-1} \hat{\mathbf{C}} (\hat{\mathbf{C}}^* \mathbf{D}^{-1} \hat{\mathbf{C}} + \frac{\sigma^2}{P} \mathbf{I})^{-1}, \quad (11)$$

which gives

$$\text{sinr}_k^{\text{MMSE}} = \frac{|\mathbf{w}_k^* \hat{\mathbf{c}}_k|^2}{\sum_{\ell \neq k} |\mathbf{w}_k^* \hat{\mathbf{c}}_\ell|^2 + \frac{\sigma^2}{P} \mathbf{w}_k^* \mathbf{D} \mathbf{w}_k} \quad (12)$$

where $\mathbf{w}_k = [\mathbf{W}^{\text{MMSE}}]_{:,k}$ is the k th column of \mathbf{W}^{MMSE} . With a bit of algebra, the above can be seen to equal [64, sec. 6.4]

$$\text{sinr}_k^{\text{MMSE}} = \frac{\hat{\mathbf{c}}_k^* (\hat{\mathbf{C}} \hat{\mathbf{C}}^* + \frac{\sigma^2}{P} \mathbf{D})^{-1} \hat{\mathbf{c}}_k}{1 - \hat{\mathbf{c}}_k^* (\hat{\mathbf{C}} \hat{\mathbf{C}}^* + \frac{\sigma^2}{P} \mathbf{D})^{-1} \hat{\mathbf{c}}_k} \quad (13)$$

$$= \frac{P/\sigma^2}{\left[(\hat{\mathbf{C}}^* \mathbf{D}^{-1} \hat{\mathbf{C}} + \frac{\sigma^2}{P} \mathbf{I})^{-1} \right]_{k,k}} - 1. \quad (14)$$

In its more economic form for $K \leq N$, the one in (11), the obtainment cost of a network-wide MMSE receiver satisfies $M_{\text{rx}}/N = \mathcal{O}(K^2/N)$. It is decidedly not scalable. The linear combining cost and the number of channel estimates are as with matched filtering, again not scalable.

VI. SUBSET MMSE RECEPTION

As mentioned, a network-wide MMSE receiver is generally inordinate because, due to pathloss and shadowing, only a small share of APs capture substantial power from user k and only a small share of other users cause substantial interference to user k . Mathematically, this is manifested by \mathbf{C} having most of its mass concentrated on a small share of its entries, a condition that is best appreciated by examining either $\mathbf{C}\mathbf{C}^*$ or $\mathbf{C}^*\mathbf{C}$: with proper AP and user indexing, both these matrices are band-dominated, meaning that most of their mass lies on a diagonal band. The band-dominated nature of $\mathbf{C}\mathbf{C}^*$, for instance, is visualized in Fig. 1 for the large network considered later in the paper. A network-wide MMSE receiver, however, ignores this nature and treats all entries of \mathbf{C} as equally relevant.




Fig. 1. Visualization of the entries of $\mathbf{C}\mathbf{C}^*$ that represent most of the mass in a network with $K = 4000$ and $N = 5000$. The reverse Cuthill-McKee ordering [65] has been applied to permute rows and columns (which amounts to re-indexing APs and users) so as to concentrate high-mass entries around the diagonal. The average thickness of the diagonal band depends on the represented percentage of mass while its irregularities reflect those of the AP and user locations.

Focusing on $\mathbf{C}\mathbf{C}^*$, the diagonal entries quantify the power captured by the APs while the off-diagonals are a measure of the interference. If $\mathbf{C}\mathbf{C}^*$ were purely diagonal, a cellular structure would be optimum. If it were block-diagonal, then a partition of the APs into disjoint clusters would be fitting. But the actual nature of $\mathbf{C}\mathbf{C}^*$ is band-dominated, which invites a tailored receiver structure involving only the relevant entries of \mathbf{C} while excluding other entries that are exceedingly weak. In essence, this amounts to zeroing out those weak entries, converting \mathbf{C} into a sparse matrix. However, the entries to zero out must be chosen with care to ensure balance across rows and columns, such that (i) no users are overly favored while others are poorly serviced or even disconnected, and (ii) no APs are rendered inoperative.

Some channel sparsification approaches have been proposed in the past, say in [35] by applying—in that case to the downlink—a family of diagonal matrices to the columns of \mathbf{C} , or by alternative means in [38]. Our approach, presented next, takes a different route towards the goal of conveniently determining which channel entries to retain and which ones to zero out.

A. Formulation

Let us restrict to a subset, denoted by \mathcal{K}_n , the users whose channels are estimated by the n th AP; users not in \mathcal{K}_n are treated as noise by such AP. Then, (5) can be rewritten as

$$y_n = \sum_{k \in \mathcal{K}_n} \sqrt{G_{n,k}} \hat{h}_{n,k} x_k + \underbrace{\sum_{k \notin \mathcal{K}_n} \sqrt{G_{n,k}} \hat{h}_{n,k} x_k}_{\mathbf{z}_n} + z_n \quad (15)$$

from which

$$\mathbf{y} = \begin{bmatrix} y_0 \\ \vdots \\ y_{N-1} \end{bmatrix} = \hat{\mathbf{C}} \mathbf{x} + \mathbf{z} \quad (16)$$

with $\mathbf{z} = [z_0 \cdots z_{N-1}]^T$,

$$[\hat{\mathbf{C}}]_{n,k} = \begin{cases} [\hat{\mathbf{C}}]_{n,k} & k \in \mathcal{K}_n \\ 0 & \text{otherwise,} \end{cases} \quad (17)$$

and $\mathbb{E}[\mathbf{z}\mathbf{z}^*] = \sigma^2 \mathbf{D}$ where

$$[\mathbf{D}]_{n,n} = \sum_{k \notin \mathcal{K}_n} \frac{\text{SNR}_{n,k}}{1 + \text{SNR}_{n,k}} + [\mathbf{D}]_{n,n} \quad (18)$$

$$= 1 + \sum_{k \in \mathcal{K}_n} \frac{\text{SNR}_{n,k}}{1 + \text{SNR}_{n,k}} + \sum_{k \notin \mathcal{K}_n} \text{SNR}_{n,k}. \quad (19)$$

Next, let us curb to a subset \mathcal{N}_k the APs that participate in the reception of user k . Denoting by $|\cdot|$ the cardinality of a set, user k is then received based on the decimated $|\mathcal{N}_k| \times 1$ vector

$$\mathbf{y}_k = \hat{\mathbf{C}}_k \mathbf{x} + \mathbf{z}_k, \quad (20)$$

where $\hat{\mathbf{C}}_k$ and \mathbf{z}_k contain the rows of $\hat{\mathbf{C}}$ and \mathbf{z} corresponding to APs within \mathcal{N}_k . The subset MMSE receiver for user k , based on the observations of the APs within \mathcal{N}_k and conditioned on the channel estimates available at each, is the $|\mathcal{N}_k| \times 1$ vector

$$\mathbf{w}_k^{\text{MMSE}} = \left(\hat{\mathbf{C}}_k \hat{\mathbf{C}}_k^* + \frac{\sigma^2}{P} \mathbf{D}_k \right)^{-1} \hat{\mathbf{c}}_k \quad (21)$$

where $\hat{\mathbf{c}}_k = [\hat{\mathbf{C}}_k]_{:,k}$. This receiver delivers, for user k ,

$$\text{snr}_k^{\text{MMSE}} = \frac{\hat{\mathbf{c}}_k^* (\hat{\mathbf{C}}_k \hat{\mathbf{C}}_k^* + \frac{\sigma^2}{P} \mathbf{D}_k)^{-1} \hat{\mathbf{c}}_k}{1 - \hat{\mathbf{c}}_k^* (\hat{\mathbf{C}}_k \hat{\mathbf{C}}_k^* + \frac{\sigma^2}{P} \mathbf{D}_k)^{-1} \hat{\mathbf{c}}_k}. \quad (22)$$

We note that, for $|\mathcal{K}_n| = K$ and $|\mathcal{N}_k| = N$, $\mathbf{w}_0^{\text{MMSE}}, \dots, \mathbf{w}_{K-1}^{\text{MMSE}}$ stretch to coincide with the columns of the network-wide MMSE receiver in (10). What remains now is to see how to determine the composition of the subsets \mathcal{N}_k and \mathcal{K}_n to approach the corresponding performance in a scalable fashion.

B. AP Selection Policy

Our desiderata for the selection of the APs that constitute $\mathcal{N}_0, \dots, \mathcal{N}_{K-1}$ are:

- Based on large-scale quantities, hence relatively stable in time and frequency, to avoid having to reselect at the small scales of milliseconds and hundreds of kilohertz.
- Close-to-optimum performance.
- Scalable. For user k , the number of potential AP subsets of size $|\mathcal{N}_k|$ is $\frac{N!}{|\mathcal{N}_k|! (N - |\mathcal{N}_k|)!}$; the total number of subsets of all sizes and for all K users is thus

$$\sum_{k=0}^{K-1} \sum_{|\mathcal{N}_k|=1}^N \frac{N!}{|\mathcal{N}_k|! (N - |\mathcal{N}_k|)!} = K (2^N - 1) \quad (23)$$

and an exhaustive search is not only unscalable, but out of the question even for modest values of K and N .

The policy we propose to meet the above requisites, inspired by [37], [39], [40], is to have \mathcal{N}_k contain the $|\mathcal{N}_k|$ APs




Fig. 2. Average sum spectral efficiency (bits/s/Hz per AP) as a function of $|\mathcal{N}|$ with subset MMSE receivers: proposed AP subset selection policy (solid line) vs exhaustive search (circles) for $\eta = 4$, $N = 24$, $K = 19$. Each AP has channel estimates for all users.

whose $G_{n,k}$ are largest, i.e., to which user k has the strongest links. Although intuitive, this policy need not be optimal under MMSE reception, hence we proceed to validate it for a network small enough that an exhaustive search is feasible. Specifically, for $N = 24$ APs and $K = 19$ users, we examine the sum spectral efficiency averaged over the user and AP positions and normalized by N . Shown in Fig. 2 is such average sum spectral efficiency with subset MMSE receivers as a function of the subset size $|\mathcal{N}|$, identical for all users, both with the proposed AP subset selection policy as well as with an exhaustive search over all possible subsets (about 319 million). The proposed policy proves to be highly effective, tracking very closely the upper bound represented by the exhaustive search and leaving very little room for improvement.

The figure also shows the performance improving promisingly with $|\mathcal{N}|$, but we hasten to emphasize that the absolute performance in small networks is not fully representative because of the scant amount of interference. For a meaningful performance evaluation, larger networks must be contemplated as done in Section VI-F.

C. User Subset Selection

To complete the specification of $\mathbf{w}_0^{\text{MMSE}}, \dots, \mathbf{w}_{K-1}^{\text{MMSE}}$ in (21), we need to determine $\mathcal{K}_0, \dots, \mathcal{K}_{N-1}$, i.e., the channel estimates to be acquired by each AP. The earlier desiderata applies, and hence a dual policy of the one adopted for \mathcal{N}_k seems like a sensible solution for \mathcal{K}_n . However, an added requirement exists, namely that the n th AP should always estimate the channels of users in whose reception it participates. To account for this, we propound that \mathcal{K}_n contain the union of:

- A fixed number of the users whose $G_{n,k}$ are largest, i.e., with strongest links to AP n .
- All users for which $n \in \mathcal{N}_k$, i.e., all users being received by AP n .

We evaluate the efficacy of this policy directly on a large network, in Section VI-F.

D. Interpretation

Before proceeding, let us pause to convey some intuition on how the receiver structure spawned by the AP and user subsets naturally emanates from that of \mathbf{C} itself. The basis is $\hat{\mathbf{C}}$ in (17), which is a sparse version of \mathbf{C} . While $\mathbf{C}\mathbf{C}^*$ and $\mathbf{C}^*\mathbf{C}$ are merely band-dominated, $\hat{\mathbf{C}}\hat{\mathbf{C}}^*$ and $\hat{\mathbf{C}}^*\hat{\mathbf{C}}$ are true band matrices. In $\hat{\mathbf{C}}$, only the dominant entries of \mathbf{C} have been preserved; the proposed AP and user selection policies enable identifying those dominant entries in a way that is both efficient and balanced, ensuring that every row and column retain a minimum of entries. All non-dominant entries—the vast majority in a large network—have been zeroed out and moved into the noise \mathbf{z} in (16), and the corresponding fading coefficients need not be estimated. Then, $\mathbf{w}_k^{\text{MMSE}}$ is obtained from $\hat{\mathbf{C}}_k$, which is a version of $\hat{\mathbf{C}}$ decimated row-wise for user k , and:

- If any other user has its channels estimated by all the APs in $\mathbf{w}_k^{\text{MMSE}}$, the interference from such user is fully rejected.
- Users whose channels are estimated by some (but not all) of the APs in $\mathbf{w}_k^{\text{MMSE}}$ are partially deflected.
- Users whose channels are not estimated by any AP within $\mathbf{w}_k^{\text{MMSE}}$ are treated as noise.

Noting that the rejection is in the MMSE sense (i.e., down to the noise floor), the interference management with our proposed subset selection policies exhibits this very desirable feature: the stronger an interference term, the more resources devoted to its mitigation. The $|\mathcal{N}_k|$ degrees of freedom that $\mathbf{w}_k^{\text{MMSE}}$ possesses in terms of interference avoidance are allocated to reject very strong terms, deflect weaker ones, and ignore the rest.

This is exemplified in Fig. 3 for a toy network. Receivers 0 and 1 fully reject their users' mutual interference while only partially deflecting the interference from user 2, and ignoring the interference from user 3. Receiver 2, in turn, partially deflects the interference from users 0 and 1 while ignoring the interference from user 3. Finally, receiver 3 partially deflects the interference from users 1 and 2.

E. Scalability

Provided the size of the subsets $\mathcal{K}_0, \dots, \mathcal{K}_{N-1}$ and $\mathcal{N}_0, \dots, \mathcal{N}_{K-1}$ do not grow with K and N , scalability is guaranteed. Precisely, the cost of obtaining $\mathbf{w}_0^{\text{MMSE}}, \dots, \mathbf{w}_{K-1}^{\text{MMSE}}$ satisfies

$$\frac{M_{\text{rx}}}{N} = \mathcal{O}\left(\frac{\sum_{k=0}^{K-1} |\mathcal{N}_k|^\nu}{N}\right), \quad (24)$$

which is $\mathcal{O}(1)$ for fixed K/N . Turning to the cost of the linear combinings, we have that

$$\frac{M_{\text{c}}}{N} = \frac{\sum_{k=0}^{K-1} |\mathcal{N}_k|}{N} \quad (25)$$




Fig. 3. Toy network with $\mathcal{N} = 4$ (common to all APs); every AP estimates the channels of the union of (i) its two strongest users, and (ii) users connecting to it. The APs that comprise $\mathbf{w}_0^{\text{MMSE}}$, $\mathbf{w}_1^{\text{MMSE}}$, $\mathbf{w}_2^{\text{MMSE}}$, and $\mathbf{w}_3^{\text{MMSE}}$, are delineated in (a) while the channels to estimate are indicated in (b).

which is also $\mathcal{O}(1)$ for fixed K/N . Finally, the number of channel estimates satisfies

$$\frac{L}{N} = \frac{\sum_{n=0}^{N-1} |\mathcal{K}_n|}{N}, \quad (26)$$

which once more is $\mathcal{O}(1)$. And $\text{sinr}_0^{\text{MMSE}}, \dots, \text{sinr}_{K-1}^{\text{MMSE}}$, necessary for link adaptation purposes, do not demand any matrix operation in addition to those required by $\mathbf{w}_0^{\text{MMSE}}, \dots, \mathbf{w}_{K-1}^{\text{MMSE}}$.

F. Evaluation in a Large Network

Let us now contemplate a network having $N = 5000$ APs, which, in ultradense deployments, could cover 1–2 km² (say a mall, large sports arena, or concert hall). For that purpose, we resort to a wrapped-around (i.e., without boundaries) universe with the APs elevated 2 m above the users to avoid distance singularities. Thanks to the wrap-around, all APs and users contribute equally to the statistics. The number of network snapshots is adjusted to push the 95% confidence interval below 0.01 b/s/Hz.

The AP subset size \mathcal{N} is identical for all users, while the user subsets $\mathcal{K}_0, \dots, \mathcal{K}_{N-1}$ contain the $\frac{K}{N}|\mathcal{N}|$ users with strongest channels to the respective APs in union with the




Fig. 4. Average sum spectral efficiency (b/s/Hz per AP) as a function of K/N for $\eta = 4$ and $N = 5000$: network-wide MMSE vs matched-filtering and cellular reception. Also shown, for $K/N = 0.8$, is the performance with subset MMSE receivers for $|\mathcal{N}| = 30$ and $|\mathcal{N}| = 58$.

users being received by each such AP. The cardinality of the user subsets therefore varies across APs.

To begin with, Fig. 4 depicts the average sum spectral efficiency as a function of K/N . The load $K/N = 0.8$ is where the maximum takes place for network-wide MMSE reception and, since this is the ultimate benchmark, we adopt this load henceforth. The performance of subset MMSE reception is represented by the dotted line. As we move up that line, we progress from cellular operation ($|\mathcal{N}| = 1$) to network-wide MMSE reception ($|\mathcal{N}| = 5000$), trading cost and channel estimation requirements for performance. For $|\mathcal{N}| = 30$, subset MMSE reception almost quadruples the cellular spectral efficiency, reaching 80% of what a network-wide MMSE receiver would attain. For $|\mathcal{N}| = 58$, that share goes up to 90%. The network-wide MMSE performance is closely approached, and at a radically lower cost as seen next.

Taking the cost parameters to be $\alpha = 1$ and $\nu = 3$, and ignoring the cost of inverting and applying the diagonal matrix \mathbf{D} , the cost M_{rx} of obtaining the network-wide MMSE receiver in (11) satisfies

$$\frac{M_{\text{rx}}}{N} = \frac{K^3 + 2NK^2}{N} \quad (27)$$

$$= 4.5 \cdot 10^7 \quad (28)$$

while the combining cost satisfies $M_c/N = K = 4000$ and the number of channel estimates per AP is $L/N = K = 4000$. In contrast, obtaining $\mathbf{w}_0^{\text{MMSE}}, \dots, \mathbf{w}_{K-1}^{\text{MMSE}}$ with $|\mathcal{N}| = 30$ entails (see Appendix A)

$$\frac{M_{\text{rx}}}{N} = \frac{2.3 \cdot 10^6 + 4000(30^3 + 30^2)}{5000} \quad (29)$$

$$= 22780 \quad (30)$$

while the combinings require $M_c/N = K|\mathcal{N}|/N = 24$. The average number of channel estimates per AP, generated by the




Fig. 5. Average sum spectral efficiency (b/s/Hz per AP) as a function of $|\mathcal{N}|$ for $\eta = 4$, $N = 5000$, and $K = 4000$; in solid, with the proposed subsetting policies of both users and APs; in dashed, with the proposed subsetting of APs, but with all user channels estimated at every AP.

simulator, is

$$\frac{L}{N} = \frac{\sum_{n=0}^{N-1} |\mathcal{K}_n|}{N} \quad (31)$$

$$= 26.4. \quad (32)$$

Altogether, 80% of the network-wide MMSE performance is achieved with 0.05% of the cost and 0.6% of the channel estimates. Similarly, for $\mathcal{N} = 58$, subset MMSE reception attains 90% of the performance of network-wide MMSE with 0.3% of the cost and 1.2% of the channel estimates.

A different outlook is provided in Fig. 5, which shows the average sum spectral efficiency as a function of $|\mathcal{N}|$. The performance improves rapidly with $|\mathcal{N}|$ and, as mentioned, for $|\mathcal{N}| = 30$ and $|\mathcal{N}| = 58$ it respectively reaches 80% and 90% of what a network-wide MMSE receiver would deliver. This figure also validates our user subset selection policy, whose performance very closely tracks—with drastically lower cost and exceedingly fewer estimates—what would be achieved if every AP estimated the channel of all K users.

Next, let us complement the foregoing perspective by going beyond performance averages. In particular, Fig. 6 presents the cumulative distribution function (CDF) over the user and AP locations of $\mathbb{E}[\text{sinr}_k]$, expected over the small-scale fading. With cellular operation, about 30% of users fall below -5 dB (which is reasonable threshold for coverage) whereas, with subset receivers and $|\mathcal{N}| = 30$, that share drops below 1%. Also, this plot evidences that comparing cellular and MMSE receivers at the same K/N is inappropriate. The proper comparison should be at the respective values of K/N that guarantee a certain coverage, and ensuring $\mathbb{E}[\text{sinr}_k] < -5$ dB with small probability in a cellular network requires K/N far below 0.8, e.g., through a frequency reuse factor. Referring to Fig. 4, the MMSE advantage is then substantially larger than indicated by the dotted line therein.

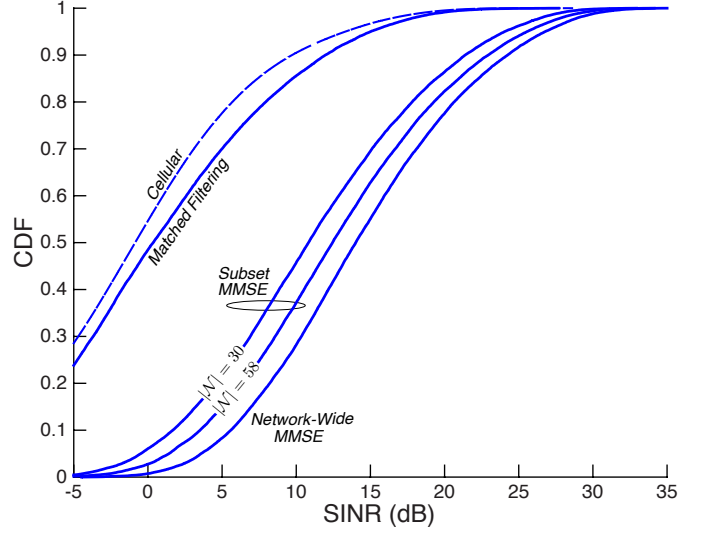


Fig. 6. CDF of $\mathbb{E}[\text{sinr}_k]$, expected over the small-scale fading, for $\eta = 4$, $N = 5000$, and $K = 4000$: network-wide MMSE, matched filtering, and subset MMSE receivers for $|\mathcal{N}| = 30$ and $|\mathcal{N}| = 58$. Also shown, in dashed, is the CDF for the corresponding cellular network.

G. Dual-Slope Pathloss

The performance of MMSE subset reception, as presented in the previous section, is conservative because of the single-slope pathloss, which was applied precisely as a precaution against performance overstatements. Now, let us redo the comparisons with a dual-slope model: $\eta = 2$ up to d , and $\eta = 4$ beyond, with d being as usual the inter-AP distance we would have in a hexagonal grid arrangement, and with P/σ^2 adjusted such that the SNR at distance d is preserved. The results, presented in Fig. 7, confirm that the cellular performance degrades and the advantage of cell-free solutions increases, even with simple matched filtering. This trend becomes more pronounced with longer pathloss breaking distances, as considered in [20], or equivalently with denser networks.

VII. FURTHER ELABORATIONS

An inspection of the AP subsets formed by our proposed policy reveals that, as one would expect, they overlap substantially. Therefore, the matrices $\hat{\mathbf{C}}_0, \dots, \hat{\mathbf{C}}_{K-1}$ from which $\mathbf{w}_0^{\text{MMSE}}, \dots, \mathbf{w}_{K-1}^{\text{MMSE}}$ are obtained via (21) exhibit substantial commonality: if users k and ℓ are in the vicinity of each other, $\hat{\mathbf{C}}_k$ and $\hat{\mathbf{C}}_\ell$ are likely to share most of their rows, as is the case with users 0, 1, and 2 in the toy network of Fig. 3. This opens the door to further cost reductions, especially in light of the fact that the inversion of $\hat{\mathbf{C}}_k \hat{\mathbf{C}}_k^*$ is what dominates the cost of obtaining $\mathbf{w}_k^{\text{MMSE}}$. We propose two distinct ideas to capitalize on these overlaps.

A. Subset Mergers

If the AP subsets for two users exhibit sufficient commonality, cost-wise it pays off to merge them into a single subset containing the union of the respective APs. Precisely, if $|\mathcal{N}_k \cup \mathcal{N}_\ell|^\nu < |\mathcal{N}_k|^\nu + |\mathcal{N}_\ell|^\nu$, the merger of \mathcal{N}_k and \mathcal{N}_ℓ

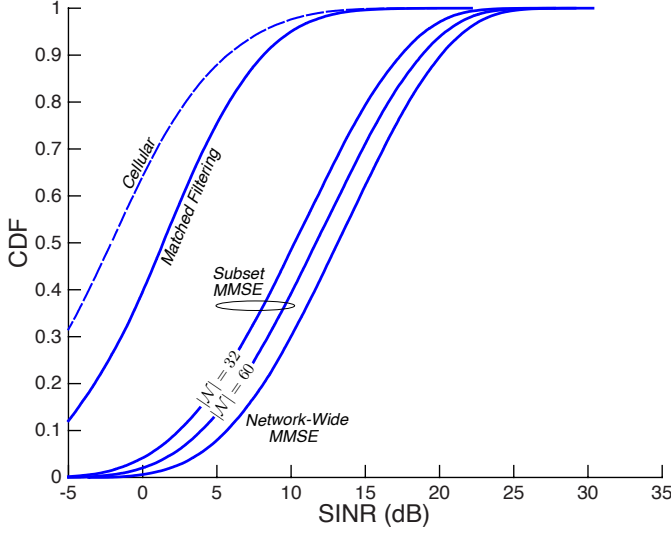


Fig. 7. CDF of $\mathbb{E}[\text{sinr}_k]$, expected over the small-scale fading, for a dual-slope pathloss model, $N = 5000$, and $K = 4000$: network-wide MMSE, matched filtering, and subset MMSE receivers for $|\mathcal{N}| = 30$ and $|\mathcal{N}| = 58$. Also shown, in dashed, is the CDF for the corresponding cellular network.

gives a new subset from which the receivers for users k and ℓ can be obtained—recall (21)—with a single matrix inversion. The aggregate cost is lowered while the performance can only improve, with the only downside of a possibly increased number of channel estimates.

The idea of merging subsets connects with the notion of arranging the network into clusters—an approach on which there is a rich literature (e.g., [43]–[52] and references therein). However, we merge subsets only when appropriate, e.g., when various users concentrate in close proximity relative to the rest; otherwise, the original subsets stand, unique to each user. And we note that whatever larger subsets may be formed through mergers have different sizes in general, may overlap with other subsets, and may deflect interference from outside users. This can be appreciated in the toy example of Fig. 3, where users 0 and 1 have identical AP subsets whose merger can be interpreted as a cluster of four APs serving two users, with two of the APs also participating in the subset that receives user 2, and with partial deflection of the interference from such user 2.

Subset mergers provide a mechanism for cell-free networks to take advantage of user concentrations.

B. Shared Processing

Once the appropriate subsets have been advantageously merged, a second idea can be applied: when various matrices have portions in common, their inversions share many of the operations. In particular, as shown in Appendix B, if K subsets have N_{cm} APs in common, then, rather than $\sum_{k=0}^{K-1} \alpha |\mathcal{N}_k|^\nu$, the total cost of inverting the corresponding $\hat{\mathbf{C}}_0, \dots, \hat{\mathbf{C}}_{K-1}$ is, at most,

$$\alpha N_{\text{cm}}^\nu + \sum_{k=0}^{K-1} \alpha \left(|\mathcal{N}_k|^\nu - \xi_k N_{\text{cm}}^\nu \right). \quad (33)$$

TABLE I
SHARED PROCESSING TOY EXAMPLE

N_{cm}	Paired Subsets	Common APs	Cost Savings
5	$\mathcal{K}_0, \mathcal{K}_1$	$\{1, 2, 3, 4, 5\}$	$\alpha (\xi_0 + \xi_1 - 1) 5^\nu$
4	$\mathcal{K}_0, \mathcal{K}_2$	$\{1, 2, 3, 4\}$	$\alpha (\xi_0 + \xi_2 - 1) 4^\nu$
4	$\mathcal{K}_0, \mathcal{K}_3$	$\{1, 2, 3, 4\}$	$\alpha (\xi_0 + \xi_3 - 1) 4^\nu$
3	$\mathcal{K}_4, \mathcal{K}_5$	$\{2, 3, 8\}$	$\alpha (\xi_4 + \xi_5 - 1) 3^\nu$
2	$\mathcal{K}_0, \mathcal{K}_4$	$\{2, 3\}$	$\alpha (\xi_0 + \xi_4 - 1) 2^\nu$

where

$$\xi_k = 1 - (2 + \rho^{\nu-3} - \nu) \rho - (3\rho^{\nu-3} - \binom{\nu}{2}) \rho^2 - (\rho^{\nu-3} - \binom{\nu}{3}) \rho^3 + \binom{\nu}{4} \rho^4 \quad (34)$$

with $\rho = \frac{|\mathcal{N}_k|}{N_{\text{cm}}} - 1$ and $\binom{\nu}{k} = \frac{\nu(\nu-1)\dots(\nu-k+1)}{k!}$. In particular, for $\alpha = 1$ and $\nu = 3$, the cost of inverting K matrices of size $|\mathcal{N}| \times |\mathcal{N}|$ goes from $K |\mathcal{N}|^3$ to, at most,

$$N_{\text{cm}}^3 + K(|\mathcal{N}|^3 - N_{\text{cm}}^3) = K |\mathcal{N}|^3 - (K - 1) N_{\text{cm}}^3. \quad (35)$$

For given α and ν , the extent of the possible cost savings depends on the degree of commonality among the post-merger AP subsets throughout the network.

In what follows, we present a specific scheme that exploits subset commonalities. It is a sequential procedure that starts by identifying all pairs of user receivers whose subsets share N_{cm} APs. From (33), the cost reduction associated with this commonality when obtaining $\mathbf{w}_k^{\text{MMSE}}$ and $\mathbf{w}_\ell^{\text{MMSE}}$ is, at least, $\alpha (\xi_k + \xi_\ell - 1) N_{\text{cm}}^\nu$. If more than two receivers share the same N_{cm} APs, the savings add up as per (33). The procedure continues with N_{cm} decreased by 1, by identifying pending receivers whose subsets have N_{cm} APs in common with any already obtained one, and so on until N_{cm} is smaller than the minimum commonality, at which point all the receivers have been computed.

In Table I, we illustrate the procedure with a toy example featuring $K = 6$ users and these AP subsets:

$$\begin{aligned} \mathcal{K}_0 &= \{1, 2, 3, 4, 5\} \\ \mathcal{K}_1 &= \{1, 2, 3, 4, 5\} \\ \mathcal{K}_2 &= \{1, 2, 3, 4, 6\} \\ \mathcal{K}_3 &= \{1, 2, 3, 4, 7\} \\ \mathcal{K}_4 &= \{2, 3, 8, 9, 10\} \\ \mathcal{K}_5 &= \{2, 3, 8, 11, 12\}. \end{aligned}$$

At a minimum, a total cost reduction equal to the sum of the entries in the rightmost column of Table I is possible. In an actual implementation, the procedure may be stopped once N_{cm} is small enough that the savings become negligible.

It behooves us to indicate that savings higher than those achieved with the above procedure may be possible with alternative schemes, which makes the processing sharing problem a fertile ground for further research.

C. Evaluation

To test the subset merging and the shared processing schemes, we apply them to the large network ($N = 5000$, $K = 4000$) studied in Section VI-F, with $\alpha = 1$ and $\nu = 3$ as

considered therein. First, all subset mergers resulting in cost reductions are executed, and then the receivers are computed via the shared processing procedure. The improvements are as follows:

- Thanks to subset merging, the value $\mathcal{N} = 30$ required to achieve 80% of the network-wide MMSE sum spectral efficiency shrinks to $\mathcal{N} = 26$. (This is the pre-merger subset size, with the subsets then swelling to various sizes with an average of 33 APs. (Although some merging does take place, a majority of the initial subsets remain unmerged.) The cost declines by 35% while the number of channel estimates per AP increases slightly, from $L/N = 26.5$ to $L/N = 32$.
- With shared processing, the cost is reduced by an additional 45% without incurring extra overheads.

Altogether, the cost of achieving 80% of the network-wide MMSE performance becomes about 1/3 of the value without further elaborations. And the savings expand as higher performance levels are targeted: for 90% of network-wide MMSE, the cost drops to about 1/4 what it would be without subset mergers and shared processing.

VIII. CONCLUSION

Subset MMSE reception is a flexible and natural solution for the wireless uplink. Rather than conform to an a-priori cellular structure, it configures itself depending on the AP and user locations, as well as the propagation. Subset reception can perform arbitrarily close to network-side MMSE reception while being scalable in terms of cost and channel estimation. In large networks, where network-wide MMSE becomes unfeasible, subset MMSE reception is fully viable.

Subset reception would surely perform even better if the subset cardinalities were separately optimized, rather than being equal as in our results; this constitutes a subject of follow-up research. As a step in this direction, we have presented a procedure to merge subsets.

We have conducted, with pilot overheads and pilot contamination disregarded, a clean assessment of the performance, computational cost, and channel estimation requirements of subset reception relative to network-wide MMSE, matched filtering, and cellular. Further accounting for the overheads would amount to postulating a fading coherence, $B_c T_c$, and then scaling the spectral efficiencies by $(1 - \frac{N_p}{B_c T_c})$ where N_p is the number of resource units reserved for pilots. While it is certainly of interest to further refine our results accounting for overheads, and for contamination as well, it is important to emphasize that ultradense deployments exhibit (because of the short range) very high frequency coherences and, most likely, also very high time coherences. Altogether, $B_c T_c$ is in the thousands of resource units, a very benevolent situation as far as overheads and contamination are concerned. For instance, for very conservative coherences of $B_c \approx 2$ MHz and $T_c \approx 5$ ms, we have $B_c T_c \approx 10^4$ and, allocating only 10% of those to pilots, we would have $N_p \approx 1000$. This is already an order of magnitude more than the $N_p \approx 100$ required for contamination to become unimportant with matched filtering and random pilot assignment [66], and let alone with active contamination

mitigation procedures such as the ones propounded in [20, sec. IV] or in [66]–[68].

The ratio P/σ^2 and the pathloss model affect the performance surge over matched filtering and cellular, as well as the subset sizes needed for a certain performance, but they do not alter the qualitative findings. As one would expect, and [20] confirms, shadow fading correlations within the large-scale gains are bound to worsen the performance of every type of receiver, but the qualitative conclusions are again upheld (with the cell-free advantage over cellular slightly diminished).

Power control is a key interference management mechanism in cellular and matched-filtering cell-free networks [20], [69]–[71], and it would be appropriate to replicate the comparisons of subset reception with such baselines in the presence of power control. This requires devising power control policies that are suitable for subset reception, hence this is another follow-up research topic. Importantly though, power control is far less critical with MMSE receivers than it is with matched filters because of the inherent interference mitigation abilities of the former.

Although not quantified in the paper, subset reception relaxes the fronthaul quantization specifications relative to network-wide MMSE, as weak signal and interference terms no longer need to be captured. This is bound to increase the compatibility with some of the functional splits proposed for C-RAN implementations of 5G and beyond [72], [73]. Also, subset reception is highly amenable to parallelization, as it entails the inversion of many small matrices as opposed to a giant one. Moreover, these small matrices exhibit major overlaps, which can be exploited to further reduce the cost; some ideas in this direction have been put forth.

Finally, the generalization to multiantenna APs and users is yet another avenue for subsequent research.

APPENDIX A

SUBSET RECEIVER OBTAINMENT COST

Let \mathbf{A} be a binary matrix that is a mask of $\hat{\mathbf{C}}$, i.e., such that

$$a_{n,k} = \begin{cases} 1 & \text{if } [\hat{\mathbf{C}}]_{n,k} \neq 0 \\ 0 & \text{if } [\hat{\mathbf{C}}]_{n,k} = 0 \end{cases} \quad (36)$$

whereby $a_{n,k} = 1$ indicates that the n th AP does estimate the channel from user k . Then, the number of complex MA operations required to compute $\hat{\mathbf{C}}\hat{\mathbf{C}}^*$ equals the sum of the entries of $\mathbf{A}\mathbf{A}^T$, i.e., $\sum_{n=0}^{N-1} \sum_{m=0}^{N-1} \sum_{k=0}^{K-1} a_{n,k} a_{m,k}$. However, since $\mathbf{A}\mathbf{A}^T$ is Hermitian, that number goes down to

$$\frac{1}{2} \sum_{n=0}^{N-1} \sum_{m=0}^{N-1} \sum_{k=0}^{K-1} a_{n,k} a_{m,k} + \frac{1}{2} \text{tr}(\mathbf{A}\mathbf{A}^T). \quad (37)$$

From $\hat{\mathbf{C}}\hat{\mathbf{C}}^*$, we can readily produce $\hat{\mathbf{C}}_0\hat{\mathbf{C}}_0^*, \dots, \hat{\mathbf{C}}_{K-1}\hat{\mathbf{C}}_{K-1}^*$ through row and column elimination. From there on, obtaining $\mathbf{w}_0^{\text{MMSE}}, \dots, \mathbf{w}_{K-1}^{\text{MMSE}}$ entails K matrix inversions and K multi-

plications with vectors of dimensions $\mathcal{N}_0, \dots, \mathcal{N}_{K-1}$, with the total cost adding up to

$$M_{\text{rx}} = \frac{1}{2} \sum_{n=0}^{N-1} \sum_{m=0}^{N-1} \sum_{k=0}^{K-1} a_{n,k} a_{m,k} + \frac{1}{2} \text{tr}(\mathbf{A} \mathbf{A}^T) + \sum_{k=0}^{K-1} \alpha \left(|\mathcal{N}_k|^\nu + |\mathcal{N}_k|^{\nu-1} \right). \quad (38)$$

The two leading terms, which are computed by the simulator, become progressively less significant as $|\mathcal{N}_0|, \dots, |\mathcal{N}_{K-1}|$ grow large, and (38) becomes dominated by $\alpha \sum_{k=0}^{K-1} |\mathcal{N}_k|^\nu$.

APPENDIX B

INVERSION OF MATRICES WITH COMMON ENTRIES

Let \mathbf{B} be a $Q \times Q$ matrix where $Q = Q_0 + Q_1$. This matrix consists of four submatrices, namely

$$\mathbf{B} = \left[\begin{array}{cc} \mathbf{B}_{0,0} & \mathbf{B}_{0,1} \\ \mathbf{B}_{1,0} & \mathbf{B}_{1,1} \end{array} \right] \left\{ \begin{array}{l} Q_0 \\ Q_1 \end{array} \right\}. \quad (39)$$

The j th column of \mathbf{B}^{-1} , denoted by $\mathbf{r}_j = [\mathbf{r}_{j0} \ \mathbf{r}_{j1}]^T$, satisfies

$$\left[\begin{array}{cc} \mathbf{B}_{0,0} & \mathbf{B}_{0,1} \\ \mathbf{B}_{1,0} & \mathbf{B}_{1,1} \end{array} \right] \left[\begin{array}{c} \mathbf{r}_{j0} \\ \mathbf{r}_{j1} \end{array} \right] = \left[\begin{array}{c} \mathbf{e}_{j0} \\ \mathbf{e}_{j1} \end{array} \right], \quad (40)$$

where \mathbf{r}_{j0} and \mathbf{r}_{j1} have, respectively, lengths Q_0 and Q_1 . In turn, $\mathbf{e}_j = [\mathbf{e}_{j0}^T \ \mathbf{e}_{j1}^T]^T$ is a standard basis vector having 1 in the j th position and zeroes elsewhere. Hence,

$$\begin{aligned} \mathbf{r}_{j1} &= (\mathbf{B}_{1,1} - \mathbf{B}_{1,0} \mathbf{B}_{0,0}^{-1} \mathbf{B}_{0,1})^{-1} (\mathbf{e}_{j1} - \mathbf{B}_{1,0} \mathbf{B}_{0,0}^{-1} \mathbf{e}_{j0}) \\ \mathbf{r}_{j0} &= \mathbf{B}_{0,0}^{-1} (\mathbf{e}_{j0} - \mathbf{B}_{0,1} \mathbf{r}_{j1}) \\ &= \mathbf{B}_{0,0}^{-1} \mathbf{e}_{j0} - \mathbf{B}_{0,0}^{-1} \mathbf{B}_{0,1} (\mathbf{B}_{1,1} - \mathbf{B}_{1,0} \mathbf{B}_{0,0}^{-1} \mathbf{B}_{0,1})^{-1} \\ &\quad \cdot (\mathbf{e}_{j1} - \mathbf{B}_{1,0} \mathbf{B}_{0,0}^{-1} \mathbf{e}_{j0}) \end{aligned} \quad (41)$$

The cost of obtaining \mathbf{r}_{j0} and \mathbf{r}_{j1} can be tallied from the breakdown in Table II. Altogether, the cost of obtaining \mathbf{B}^{-1} exploiting its submatrix structure is seen to be

$$\text{cost}\{\mathbf{B}^{-1}\} = \alpha \left(Q_0^\nu + Q_1^\nu + 2Q_0^{\nu-1} Q_1 + Q_0^2 Q_1^{\nu-2} + 3Q_0 Q_1^{\nu-1} \right) \quad (42)$$

while the cost of a direct inversion can be expanded into

$$\alpha (Q_0 + Q_1)^\nu = \alpha Q_0^\nu \left(1 + \sum_{k=1}^{\infty} \binom{\nu}{k} \left(\frac{Q_1}{Q_0} \right)^k \right) \quad (43)$$

where $\binom{\nu}{k} = \frac{\nu(\nu-1)\dots(\nu-k+1)}{k!}$ and where the expansion holds for both integer and non-integer ν . Letting $\epsilon = -\binom{\nu}{4} Q_0^{\nu-4} Q_1^4$,

$$\begin{aligned} (Q_0 + Q_1)^\nu &\geq Q_0^\nu + \nu Q_0^{\nu-1} Q_1 + \binom{\nu}{2} Q_0^{\nu-2} Q_1^2 \\ &\quad + \binom{\nu}{3} Q_0^{\nu-3} Q_1^3 - \epsilon, \end{aligned} \quad (44)$$

which holds because, for $2.37 \leq \nu \leq 3$, higher order terms have alternating signs and are decreasing in absolute value.

TABLE II
COST OF THE OPERATIONS REQUIRED TO OBTAIN \mathbf{B}^{-1}

Operation	Cost
$\mathbf{B}_{0,0}^{-1}$	αQ_0^ν
$\mathbf{B}_{0,0}^{-1} \mathbf{B}_{0,1}$	$\alpha Q_0^{\nu-1} Q_1$
$\mathbf{B}_{1,0} \mathbf{B}_{0,0}^{-1}$	$\alpha Q_0^{\nu-1} Q_1$
$\Xi = (\mathbf{B}_{1,1} - \mathbf{B}_{1,0} \mathbf{B}_{0,0}^{-1} \mathbf{B}_{0,1})^{-1}$	$\alpha (Q_0 Q_1^{\nu-1} + Q_1^\nu)$
$\Xi \mathbf{B}_{1,0} \mathbf{B}_{0,0}^{-1}$	$\alpha Q_0 Q_1^{\nu-1}$
$\mathbf{B}_{0,0}^{-1} \mathbf{B}_{0,1} \Xi$	$\alpha Q_0 Q_1^{\nu-1}$
$\mathbf{B}_{0,0}^{-1} \mathbf{B}_{0,1} \Xi \mathbf{B}_{1,0} \mathbf{B}_{0,0}^{-1}$	$\alpha Q_0^2 Q_1^{\nu-2}$

Combining (42) and (44),

$$\begin{aligned} (Q_0 + Q_1)^\nu &\geq \frac{1}{\alpha} \text{cost}\{\mathbf{B}^{-1}\} - (2 + \beta - \nu) Q_0^{\nu-1} Q_1 \\ &\quad - \left(3\beta - \binom{\nu}{2} \right) Q_0^{\nu-2} Q_1^2 \\ &\quad - \left(\beta - \binom{\nu}{3} \right) Q_0^{\nu-3} Q_1^3 - \epsilon \end{aligned} \quad (45)$$

with $\beta = \left(\frac{Q_0}{Q_1} \right)^{3-\nu}$. Using $\text{cost}\{\mathbf{B}_{0,0}^{-1}\} = \alpha Q_0^\nu$, we have that

$$\text{cost}\{\mathbf{B}^{-1}\} - \text{cost}\{\mathbf{B}_{0,0}^{-1}\} \leq \alpha \left((Q_0 + Q_1)^\nu - \xi Q_0^\nu \right) \quad (46)$$

where

$$\begin{aligned} \xi &= 1 - (2 + \beta - \nu) \frac{Q_1}{Q_0} - \left(3\beta - \binom{\nu}{2} \right) \left(\frac{Q_1}{Q_0} \right)^2 \\ &\quad - \left(\beta - \binom{\nu}{3} \right) \left(\frac{Q_1}{Q_0} \right)^3 - \epsilon Q_0^{-\nu}. \end{aligned} \quad (47)$$

Now, from (46) and the decomposition

$$\text{cost}\{\mathbf{B}^{-1}\} = \left(\text{cost}\{\mathbf{B}^{-1}\} - \text{cost}\{\mathbf{B}_{0,0}^{-1}\} \right) + \text{cost}\{\mathbf{B}_{0,0}^{-1}\} \quad (48)$$

we can assert that, if K matrices have $\mathbf{B}_{0,0}$ in common, then the aggregate cost of their inversions is not $K\alpha(Q_0 + Q_1)^\nu$, but rather it is, at most,

$$\alpha Q_0^\nu + K\alpha \left((Q_0 + Q_1)^\nu - \xi Q_0^\nu \right). \quad (49)$$

REFERENCES

- [1] R. Irmer, J. 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 Commun. Magazine*, vol. 49, no. 2, pp. 102–111, 2011.
- [2] M. Sawahashi, Y. Kishiyama, A. Morimoto, D. Nishikawa, and M. Tanno, "Coordinated multipoint transmission/reception techniques for LTE-advanced," *IEEE Wireless Commun.*, vol. 17, no. 3, p. 26, 2010.
- [3] D. Lee, H. Seo, B. Clerckx, E. Hardouin, D. Mazzaresse, S. Nagata, and K. Sayana, "Coordinated multipoint transmission and reception in LTE-advanced: deployment scenarios and operational challenges," *IEEE Commun. Magazine*, vol. 50, no. 2, pp. 148–155, 2012.
- [4] X. Zhang and J. G. Andrews, "Downlink cellular network analysis with multi-slope path loss models," *IEEE Trans. Commun.*, vol. 63, no. 5, pp. 1881–1894, May 2015.
- [5] M. Ding, P. Wang, D. López-Pérez, G. Mao, and Z. Lin, "Performance impact of LoS and NLoS transmissions in dense cellular networks," *IEEE Trans. Wireless Commun.*, vol. 15, no. 3, pp. 2365–2380, 2015.
- [6] J. G. Andrews, X. Zhang, G. D. Durgin, and A. K. Gupta, "Are we approaching the fundamental limits of wireless network densification?" *IEEE Commun. Magazine*, vol. 54, no. 10, pp. 184–190, Oct. 2016.

- [7] M. Ding and D. López-Pérez, "Performance impact of base station antenna heights in dense cellular networks," *IEEE Trans. Wireless Commun.*, vol. 16, no. 12, pp. 8147–8161, 2017.
- [8] M. K. Karakayali, G. J. Foschini, and R. A. Valenzuela, "Network coordination for spectrally efficient communications in cellular systems," *IEEE Wireless Commun. Mag.*, vol. 3, no. 14, pp. 56–61, Aug. 2006.
- [9] S. Venkatesan, A. Lozano, and R. Valenzuela, "Network MIMO: Overcoming intercell interference in indoor wireless systems," in *Asilomar Conf. on Signals, Systems and Computers*, 2007, pp. 83–87.
- [10] S. Venkatesan, H. Huang, A. Lozano, and R. Valenzuela, "A WiMAX-based implementation of network MIMO for indoor wireless systems," *EURASIP J. on Adv. in Signal Processing*, vol. 2009, p. 9, 2009.
- [11] D. Gesbert, S. Hanly, H. Huang, S. Shamai, O. Simeone, and W. Yu, "Multi-cell MIMO cooperative networks: A new look at interference," *IEEE J. Sel. Areas Commun.*, vol. 28, no. 9, pp. 1380–1408, 2010.
- [12] A. Gudipati, D. Perry, L. E. Li, and S. Katti, "SoftRAN: Software defined radio access network," in *ACM SIGCOMM workshop on Hot topics in software defined networking*, 2013, pp. 25–30.
- [13] A. Checko, H. L. Christiansen, Y. Yan, L. Scolari, G. Kardaras, M. S. Berger, and L. Dittmann, "Cloud RAN for mobile networks? a technology overview," *IEEE Commun. Surveys & Tutorials*, vol. 17, no. 1, pp. 405–426, 2014.
- [14] P. Rost, C. J. Bernardos, A. De Domenico, M. Di Girolamo, M. Lalam, A. Maeder, D. Sabella, and D. Wübben, "Cloud technologies for flexible 5G radio access networks," *IEEE Commun. Magazine*, vol. 52, no. 5, pp. 68–76, 2014.
- [15] S. Bhaumik, S. P. Chandrabose, M. K. Jataprolu, G. Kumar, A. Muralidhar, P. Polakos, V. Srinivasan, and T. Woo, "CloudIQ: A framework for processing base stations in a data center," in *Int'l Conf. Mobile Computing and Netw.*, 2012, pp. 125–136.
- [16] H. ElSawy, H. Dahrouj, T. Y. Al-Naffouri, and M.-S. Alouini, "Virtualized cognitive network architecture for 5G cellular networks," *IEEE Commun. Magazine*, vol. 53, no. 7, pp. 78–85, 2015.
- [17] C. I. J. Huang, R. Duan, C. Cui, J. . Jiang, and L. Li, "Recent progress on C-RAN centralization and cloudification," *IEEE Access*, vol. 2, pp. 1030–1039, 2014.
- [18] J. Wu, Z. Zhang, Y. Hong, and Y. Wen, "Cloud radio access network (C-RAN): a primer," *IEEE Network*, vol. 29, no. 1, pp. 35–41, Jan. 2015.
- [19] P. Rost, S. Talarico, and M. C. Valenti, "The complexity?rate tradeoff of centralized radio access networks," *IEEE Trans. Wireless Commun.*, vol. 14, no. 11, pp. 6164–6176, Nov. 2015.
- [20] H. Q. Ngo, A. Ashikhmin, H. Yang, E. G. Larsson, and T. L. Marzetta, "Cell-free massive MIMO versus small cells," *IEEE Trans. on Wireless Communications*, vol. 16, no. 3, pp. 1834–1850, 2017.
- [21] G. Interdonato, H. Q. Ngo, E. G. Larsson, and P. Frenger, "How much do downlink pilots improve cell-free massive MIMO?" in *IEEE Global Communications Conf. (GLOBECOM'16)*, 2016, pp. 1–7.
- [22] H. Q. Ngo, L.-N. Tran, T. Q. Duong, M. Matthaiou, and E. G. Larsson, "On the total energy efficiency of cell-free massive MIMO," *IEEE Trans. Green Commun. and Networking*, vol. 2, no. 1, pp. 25–39, 2017.
- [23] E. Nayeibi, A. Ashikhmin, T. L. Marzetta, H. Yang, and B. D. Rao, "Precoding and power optimization in cell-free massive MIMO systems," *IEEE Trans. Wireless Commun.*, vol. 16, no. 7, pp. 4445–4459, 2017.
- [24] Z. Chen and E. Björnson, "Channel hardening and favorable propagation in cell-free massive MIMO with stochastic geometry," *IEEE Trans. Commun.*, vol. 66, no. 11, pp. 5205–5219, 2018.
- [25] T. C. Mai, H. Q. Ngo, M. Egan, and T. Q. Duong, "Pilot power control for cell-free massive MIMO," *IEEE Trans. Vehic. Techn.*, vol. 67, no. 11, pp. 11 264–11 268, 2018.
- [26] M. Bashar, K. Cumanan, A. G. Burr, M. Debbah, and H. Q. Ngo, "On the uplink max-min SINR of cell-free massive MIMO systems," *IEEE Trans. Wireless Commun.*, vol. 18, no. 4, pp. 2021–2036, Apr. 2019.
- [27] H. Q. Ngo, H. Tataria, M. Matthaiou, S. Jin, and E. G. Larsson, "On the performance of cell-free massive MIMO in Ricean fading," in *Asilomar Conference on Signals, Systems, and Computers*, 2018, pp. 980–984.
- [28] M. Attarifar, A. Abbasfar, and A. Lozano, "Modified conjugate beamforming for cell-free massive MIMO," *IEEE Wireless Commun. Letters*, vol. 8, no. 2, pp. 616–619, Apr. 2019.
- [29] J. Zhang, S. Chen, Y. Lin, J. Zheng, B. Ai, and L. Hanzo, "Cell-free massive MIMO: A new next-generation paradigm," *IEEE Access*, vol. 7, pp. 99 878–99 888, 2019.
- [30] H. Masoumi and M. J. Emadi, "Performance analysis of cell-free massive MIMO system with limited fronthaul capacity and hardware impairments," *IEEE Trans. Wireless Commun.*, vol. 19, no. 2, pp. 1038–1053, 2020.
- [31] P. Liu, K. Luo, D. Chen, and T. Jiang, "Spectral efficiency analysis of cell-free massive MIMO systems with zero-forcing detector," *IEEE Trans. Wireless Commun.*, vol. 19, no. 2, pp. 795–807, 2020.
- [32] S. Perlman and A. Forenza, "An introduction to pCell," Artemis Networks LLC, White paper, Tech. Rep., Feb. 2015. [Online]. Available: <http://www.readen.com/artemis/An-Introduction-to-pCell-White-Paper-150224.pdf>
- [33] E. Björnson and L. Sanguinetti, "Making cell-free massive MIMO competitive with MMSE processing and centralized implementation," *IEEE Trans. Wireless Commun.*, vol. 19, no. 1, pp. 77–90, 2020.
- [34] A. Lozano, R. W. Heath, and J. G. Andrews, "Fundamental limits of cooperation," *IEEE Trans. on Inform. Theory*, vol. 59, no. 9, pp. 5213–5226, 2013.
- [35] E. Björnson, N. Jaldén, M. Bengtsson, and B. Ottersten, "Optimality properties, distributed strategies, and measurement-based evaluation of coordinated multicell OFDMA transmission," *IEEE Trans. Signal Process.*, vol. 59, no. 12, pp. 6086–6101, 2011.
- [36] J. Wang and L. Dai, "Downlink rate analysis for virtual-cell based large-scale distributed antenna systems," *IEEE Trans. Wireless Commun.*, vol. 15, no. 3, pp. 1998–2011, 2015.
- [37] E. Nayeibi, A. Ashikhmin, T. L. Marzetta, and B. D. Rao, "Performance of cell-free massive MIMO systems with MMSE and LSFD receivers," in *Asilomar Conf. on Signals, Systems and Computers*, Nov. 2016, pp. 203–207.
- [38] S. Buzzi and C. D'Andrea, "Cell-free massive MIMO: User-centric approach," *IEEE Wireless Commun. Letters*, vol. 6, no. 6, pp. 706–709, Dec. 2017.
- [39] S. Govindasamy and I. Bergel, "Uplink performance of multi-antenna cellular networks with co-operative base stations and user-centric clustering," *IEEE Trans. Wireless Commun.*, vol. 17, no. 4, pp. 2703–2717, Apr. 2018.
- [40] S. Zaidi, O. Ben Smida, S. Affes, U. Vilaipornsawai, L. Zhang, and P. Zhu, "User-centric base-station wireless access virtualization for future 5G networks," *IEEE Trans. Commun.*, vol. 67, no. 7, pp. 5190–5202, 2019.
- [41] E. Björnson and L. Sanguinetti, "Scalable cell-free massive MIMO systems," *arXiv:1908.03119*, 2019.
- [42] S. Buzzi, C. D'Andrea, A. Zappone, and C. D'Elia, "User-centric 5G cellular networks: Resource allocation and comparison with the cell-free massive MIMO approach," *IEEE Trans. Wireless Commun.*, vol. 19, no. 2, pp. 1250–1264, 2020.
- [43] A. Papadogiannis, D. Gesbert, and E. Hardouin, "A dynamic clustering approach in wireless networks with multi-cell cooperative processing," in *IEEE Int'l Conf. Commun. (ICC'08)*, 2008, pp. 4033–4037.
- [44] J. Zhang, R. Chen, J. G. Andrews, A. Ghosh, and R. W. Heath, "Networked MIMO with clustered linear precoding," *IEEE Trans. Wireless Commun.*, vol. 8, no. 4, pp. 1910–1921, 2009.
- [45] C. Ng and H. Huang, "Linear precoding in cooperative MIMO cellular networks with limited coordination clusters," *IEEE J. Sel. Areas Commun.*, vol. 28, no. 9, pp. 1446–1454, 2010.
- [46] J.-M. Moon and D.-H. Cho, "Inter-cluster interference management based on cell-clustering in network MIMO systems," in *IEEE Veh. Techn. Conf. (VTC'11 Spring)*, 2011.
- [47] P. Baracca, F. Boccardi, and V. Braun, "A dynamic joint clustering scheduling algorithm for downlink CoMP systems with limited CSI," in *Int'l Symp. Wireless Commun. Syst. (ISWCS'12)*, 2012.
- [48] M. Hong, R. Sun, H. Baligh, and Z.-Q. Luo, "Joint base station clustering and beamformer design for partial coordinated transmission in heterogeneous networks," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 2, pp. 226–240, 2013.
- [49] N. Lee, R. W. Heath, D. Morales-Jimenez, and A. Lozano, "Base station cooperation with dynamic clustering in super-dense cloud-RAN," in *IEEE Globecom Workshops (GC'13 Wkshps)*, Dec. 2013, pp. 784–788.
- [50] B. Dai and W. Yu, "Sparse beamforming and user-centric clustering for downlink cloud radio access network," *IEEE Access*, vol. 2, pp. 1326–1339, 2014.
- [51] C. Zhu and W. Yu, "Stochastic modeling and analysis of user-centric network MIMO systems," *IEEE Trans. Commun.*, vol. 66, no. 12, pp. 6176–6189, 2018.
- [52] G. Interdonato, P. Frenger, and E. G. Larsson, "Scalability aspects of cell-free massive MIMO," in *IEEE Int'l Conf. Commun. (ICC'19)*, May 2019.
- [53] N. Ross and D. Schuhmacher, "Wireless network signals with moderately correlated shadowing still appear Poisson," *IEEE Trans. Inform. Theory*, vol. 63, no. 2, pp. 1177–1198, 2016.

- [54] B. Błaszczyszyn, M. K. Karray, and H. P. Keeler, "Wireless networks appear Poissonian due to strong shadowing," *IEEE Trans. on Wireless Communications*, vol. 14, no. 8, pp. 4379–4390, Aug. 2015.
- [55] H. P. Keeler, N. Ross, and A. Xia, "When do wireless network signals appear poisson?" *Bernoulli*, vol. 24, no. 3, pp. 1973–1994, 2018.
- [56] G. George, R. K. Mungara, A. Lozano, and M. Haenggi, "Ergodic spectral efficiency in MIMO cellular networks," *IEEE Trans. on Wireless Communications*, vol. 16, no. 5, pp. 2835–2849, 2017.
- [57] G. George, A. Lozano, and M. Haenggi, "Massive MIMO forward link analysis for cellular networks," *IEEE Trans. Wireless Commun.*, vol. 18, no. 6, pp. 2964–2976, 2019.
- [58] 3GPP TS 36.814, "Further advancements for E-UTRA physical layer aspects (Release 9)," 3GPP, Tech. Rep., Mar. 2017.
- [59] N. Jindal and A. Lozano, "A unified treatment of optimum pilot overhead in multipath fading channels," *IEEE Trans. on Communications*, vol. 58, no. 10, pp. 2939–2948, Oct. 2010.
- [60] A. Lozano and D. Porrat, "Non-peaky signals in wideband fading channels: achievable bit rates and optimal bandwidth," *IEEE Trans. on Wireless Communications*, vol. 11, no. 1, pp. 246–257, Jan. 2012.
- [61] A. Lapidoth and S. Shamai, "Fading channels: How perfect need 'perfect side information' be?" *IEEE Trans. on Inform. Theory*, vol. 48, no. 5, pp. 1118–1134, May 2002.
- [62] J. Huang, T. M. Smith, G. M. Henry, and R. A. van de Geijn, "Strassen's algorithm reloaded," in *Int'l Conf. High Performance Computing, Networking, Storage and Analysis*, 2016, p. 59.
- [63] Q. Sun, D. C. Cox, A. Lozano, and H. C. Huang, "Training-based channel estimation for continuous flat fading BLAST," in *IEEE Int'l Conf. Commun. (ICC'02)*, 2002.
- [64] R. W. Heath Jr and A. Lozano, *Foundations of MIMO Communication*. Cambridge University Press, 2019.
- [65] E. Cuthill and J. McKee, "Reducing the bandwidth of sparse symmetric matrices," in *24th National Conference, ACM*, 1969, pp. 157–172.
- [66] M. Attarifar, A. Abbasfar, and A. Lozano, "Random vs structured pilot assignment in cell-free massive MIMO wireless networks," in *IEEE Int'l Conf. Commun. Workshops (ICCW'18)*, 2018, pp. 1–6.
- [67] O. Y. Bursalioglu, C. Wang, H. Papadopoulos, and G. Caire, "RRH based massive MIMO with 'on the fly' pilot contamination control," in *IEEE Int'l Conf. on Communications (ICC'16)*, 2016, pp. 1–7.
- [68] Y. Zhang, H. Cao, P. Zhong, C. Qi, and L. Yang, "Location-based greedy pilot assignment for cell-free massive MIMO systems," in *IEEE Int'l Conf. on Computer and Commun. (ICCC'18)*, Dec. 2018, pp. 392–396.
- [69] R. Nikbakht and A. Lozano, "Uplink fractional power control for cell-free wireless networks," *IEEE Int'l Conf. Commun. (ICC'19)*, June 2019.
- [70] —, "Unsupervised-learning power control for cell-free wireless systems," in *IEEE Int'l Symp. Personal, Indoor and Mobile Radio Commun. (PIMRC'19)*, Sep. 2019.
- [71] R. Nikbakht, R. Mosayebi, and A. Lozano, "Uplink fractional power control and downlink power allocation for cell-free networks," *IEEE Wireless Commun. Letters*, 2020.
- [72] "Study on new radio access technology: Radio access architecture and interfaces, V14.0.0 (2017–03)," 3GPP TR 38.801, Tech. Rep., 2017.
- [73] L. M. P. Larsen, A. Checko, and H. L. Christiansen, "A survey of the functional splits proposed for 5G mobile crosshaul networks," *IEEE Commun. Surveys & Tutorials*, vol. 21, no. 1, pp. 146–172, 2018.