

Spatial Blanking and Inter-Tier Coordination in Massive-MIMO Heterogeneous Cellular Networks

Ansuman Adhikary, Harpreet S. Dhillon, and Giuseppe Caire

Abstract—This paper studies inter-tier interference coordination in a heterogeneous cellular network (HetNet) consisting of “massive-MIMO” macrocells and multi-antenna small cells. We assume that the users are concentrated at certain areas of the cell forming *hotspots*. Assuming the hotspot size to be much smaller than the macrocell size, the users of a given hotspot are seen under a relatively narrow angular spread from the macrocell, thus appearing almost co-located to the macrocell. This gives rise to directional channel vectors, which can be exploited to simplify precoder design and obtain *spatial blanking*, a means of concentrating energy in the direction of scheduled hotspots, while mitigating interference caused to the *active* small cells located in the other directions. We further show that significant throughput gains can be achieved by complementing spatial blanking with active interference coordination strategies, such as turning OFF a small cell when it suffers from or causes excessive interference, or offloading macrocell hotspots to small cells.

Index Terms—Massive-MIMO, heterogeneous cellular networks, hotspots, interference coordination, spatial blanking.

I. INTRODUCTION

The mobile data traffic is almost getting doubled every year, leading to predictions of nearly 1000x increase in cellular capacity needed over the next decade [1]. A sensible way to handle this demand is by network densification, which can be efficiently achieved by deploying low-power low-cost base stations, often called *small cells*. This however comes at a price of increased interference power, which threatens to negate the gains achieved by network densification in certain regimes. As a result, both industry and academia have focused extensively on interference coordination strategies to enable the coexistence of macro and small cells in HetNets [2]. This is also the objective of 3GPP’s enhanced Inter-Cell Interference Coordination (eICIC) under which orthogonalization over both frequency and time can be achieved by fractional frequency reuse and almost-blank subframes, respectively [2]. In this paper, we present an alternate viewpoint and argue that orthogonalization over space leads to significantly more efficient interference coordination with much lower protocol overhead.

A macrocell base station equipped with a large number of antennas (massive-MIMO) and located at an elevated position (e.g., tower-mounted, or deployed on a building roof) “sees” both its own hotspots as well as the ones served by small cells under a relatively narrow angular spread. This gives rise to directional channel vectors that can be exploited to jointly provide spatial multiplexing to hotspots served by a macrocell and

perform implicit cross-tier interference mitigation to hotspots served by small cells. This is achieved by using Joint Spatial Division and Multiplexing (JSDM) technique of [3], where the basic principle is to partition the user space into groups of users with approximately similar covariances, and split the downlink beamforming into two stages. The first stage consists of a pre-beamformer that is responsible for spatial separation of the user groups and depends only on the second order statistics, i.e., the covariances of the user channels. The second stage is comprised of a standard multiuser MIMO precoder for spatial multiplexing on the effective channel including the pre-beamforming. As a consequence of pre-beamforming, interference mitigation can be achieved explicitly by nulling certain selected small cells’ spatial directions or implicitly by serving user groups that are not in the direction of the *active* small cells. This approach is termed *spatial blanking*, in analogy with the almost-blank subframe approach of eICIC.

In this paper, we consider a HetNet scenario, where (i) users are concentrated at certain areas in the cell, thereby forming hotspots; (ii) small cells are deployed in the vicinity of some of the hotspots (capturing current capacity-driven deployment); and (iii) hotspots devoid of a small cell are directly served by the macrocell. Assuming that the hotspot sizes are much smaller than the macrocell radius, a massive-MIMO macrocell can focus its energy in the direction of hotspots that it serves, while allowing simultaneous transmission of small cells located in the other directions, thus resulting in *spatial blanking*. Even though this provides an implicit interference mitigation, the *random* locations of the small cells may still give rise to a high cross-tier interference, which can be mitigated using two low-complexity *active* interference coordination strategies proposed in this work. The first strategy relies on the small cells making intelligent decisions on whether to transmit or not depending on the amount of cross-tier interference they receive or cause to the scheduled macrocell hotspots. The second strategy “offloads” some of the macrocell hotspots to small cells, in order to reduce the burden on the macrocell and simultaneously increase fairness in the rates of all the hotspots.

II. SYSTEM MODEL

Consider a HetNet formed by the small cells coexisting in the same coverage area and frequency band with the macrocells. Each macro and small cell base station is equipped with M and L antennas, respectively. We assume the regime of *massive MIMO* [4], [5] for which $M \gg L \gg 1$, and the number of users served by each base station in a given time-frequency slot (denoted hereafter as a *transmission resource*

A. Adhikary is with Ericsson, Santa Clara, CA (email: ansuman.adhikary@ericsson.com). H. S. Dhillon is with Wireless@VT, Virginia Tech, Blacksburg, VA (email: hdhillon@vt.edu). G. Caire is with the Technical University of Berlin, Germany (email: caire@tu-berlin.de).

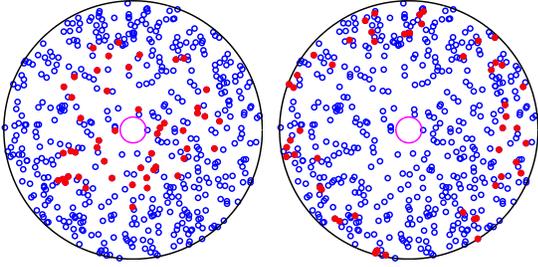


Fig. 1. The hollow “blue” circles denote the user groups that do not contain a small cell, and the filled “red” circles denote the user groups containing a small cell. (first) Cell-interior deployment. (second) Cell-edge deployment. The macrocell is located at the center of the “magenta” circle, which denotes an exclusion ball of radius R_{excl} for the macrocell.

block) is significantly less than the corresponding number of transmit antennas. For simplicity, we consider a single-cell scenario with a macrocell located at the center of a disk of radius R_{mc} , where R_{mc} denotes the macrocell coverage radius.

We focus on a non-uniform user distribution in which users are clustered into several high user-density *hotspots*, referred hereafter as *user groups*. We assume N_{u} user groups uniformly and independently distributed in the macrocell area. Assuming each group is concentrated over an area much smaller than the macrocell disk, it is reasonable to assume that the users in a given user group appear *almost* co-located to the macrocell. Therefore, we assume that while the users of a given user group are separated by several wavelengths, they all have the same scattering “landscape”. This ensures that the user channel vectors from users of the same group to the macrocell antenna array are mutually independent but identically distributed with the same covariance matrix. Refer to [3], [6] for more details.

We let $N_{\text{f}} \leq N_{\text{u}}$ denote the number of small cells in the system, each of which covers a user group. The set of such user groups is denoted by \mathcal{S} . For convenience, the small cell is assumed to be located at the center of its user group, reflecting a deployment at known and persistent hotspots such as airport lounges or coffee shops. The remaining $N_{\text{u}} - N_{\text{f}}$ user groups, denoted by the set \mathcal{M} , can be served either by the macrocell or by some neighboring small cell through *offloading*, the exact details of which will appear in Section IV.

Due to spatial blanking, it is not straightforward *a priori* whether to deploy small cells uniformly across the whole cell or more concentrated on the cell edge rather than towards the cell center. Therefore, we consider the following three deployment scenarios for the small cells: (i) *uniform deployment*: small cells are randomly and uniformly assigned to N_{f} user groups in the macrocell, i.e., user groups lying in $B(0, R_{\text{mc}})$, where $B(c, r)$ is a 2-dimensional ball of radius r centered at c ; (ii) *cell-interior deployment*: small cells are randomly and uniformly assigned to N_{f} user groups in $B\left(0, \frac{R_{\text{mc}}}{\sqrt{2}}\right)$; and (iii) *cell-edge deployment*: small cells are randomly and uniformly assigned to N_{f} user groups in $B(0, R_{\text{mc}}) \setminus B\left(0, \frac{R_{\text{mc}}}{\sqrt{2}}\right)$. While the uniform deployment is straightforward to visualize, a snapshot of the other two is shown in Fig. 1.

We assume each user group is surrounded by a circular ring of scatterers of radius $R_{\text{u}} \ll R_{\text{mc}}$, giving rise to the one-ring

channel model between the macrocell and each user group [6]. At every transmission resource block, the macrocell schedules $G \leq N_{\text{u}}$ user groups. For each scheduled user group g , a subset of the users S_g is served by spatial multiplexing, where the exact number of users selected depends on the rank of the user group channel covariance matrix.

The instantaneous channel between a user k in group g (denoted by g_k) and the macrocell, over any given transmission resource block, is an $M \times 1$ Gaussian random vector denoted by $\mathbf{h}_{g_k,0}$, written using Karhunen-Loeve representation as

$$\mathbf{h}_{g_k,0} = \mathbf{U}_g \mathbf{\Lambda}_g^{1/2} \mathbf{w}_{g_k}, \quad (1)$$

where $\mathbf{R}_g = \mathbf{U}_g \mathbf{\Lambda}_g \mathbf{U}_g^{\text{H}}$ is the channel covariance matrix of rank r_g common to all users in group g , \mathbf{U}_g is the tall unitary matrix of eigenvectors of dimension $M \times r_g$, $\mathbf{\Lambda}_g$ is the $r_g \times r_g$ diagonal positive definite matrix of covariance eigenvalues. The $r_g \times 1$ random vector $\mathbf{w}_{g_k} \sim \mathcal{CN}(0, \mathbf{I}_{r_g})$ is independent for different users and corresponds to the randomness due to the small-scale multipath fading components.

Following the one-ring scattering model [3], we have $\mathbf{R}_g = \mathbf{R}(\theta_g, \Delta_g)$ where, for a user group located at an angle of arrival θ_g and having angular spread Δ_g , the element (m, n) of $\mathbf{R}(\theta_g, \Delta_g)$ is given by

$$[\mathbf{R}(\theta_g, \Delta_g)]_{m,n} = \frac{a_{g,0}}{2\Delta_g} \int_{\theta_g - \Delta_g}^{\theta_g + \Delta_g} e^{-j\pi(m-n)\sin(\alpha)} d\alpha \quad (2)$$

assuming a uniform linear array at the macrocell. Note that $a_{g,0}$ represents the path-loss due to the propagation environment, which is given in (3). The total macrocell transmit power is denoted by P_0 . For analytical simplicity, we consider equal power allocation such that all the macrocell downlink data streams are transmitted with the same power $\frac{P_0}{S}$, where $S \leq \sum_{g=1}^G S_g$ is the total number of downlink streams, i.e., the total number of users served across all groups.

In each transmission resource block, a small cell serves \bar{S}_f users in the group by spatial multiplexing.¹ *Inter-tier interference coordination* strategies between the small cells and the macrocell are discussed in Section IV. All the active small cells transmit at their peak power P_1 . The presence of multiple users gives rise to intra-cell interference, which is handled by zero forcing beamforming. Since *user selection* [7] yields negligible gains in the regime of massive MIMO, we assume that each small cell simply schedules a random set of \bar{S}_f users in its group with uniform probability, achieving proportional fairness (which in this case coincides with equal air-time) across all its users.

Small cells are typically deployed at low elevation, therefore, the scattering geometry of the channels between users and the small cell array can be considered isotropic. Hence, the channel between a user k in group g and a small cell in group f is modeled as a $L \times 1$ vector $\mathbf{h}_{g_k,f}$ with i.i.d. entries $\sim \mathcal{CN}(0, a(g, f))$, where $a(g, f)$ is the distance-dependent path-loss coefficient between the users in group g and the

¹We assume that the user groups are fully loaded, i.e., they contain a sufficient number of users (much larger than \bar{S}_f).

small cell co-located with group f , and is given by

$$a(g, f) = \frac{w^{n_w(g, f)}}{1 + \left(\frac{d(g, f)}{d_0}\right)^\alpha}, \quad (3)$$

where $d(g, f)$ denotes the distance between user groups g and f , d_0 is the cutoff distance, α the path-loss exponent, w the wall penetration loss, and $n_w(g, f)$ denotes the number of walls between user groups g and f , which we model as

$$n_w(g, f) = \begin{cases} 0 & g = f \\ 1 & g \in \mathcal{M}, f \in \mathcal{S} \text{ OR } g \in \mathcal{S}, f \in \mathcal{M} \\ 2 & g, f \in \mathcal{S}, g \neq f \end{cases} \quad (4)$$

III. DOWNLINK TRANSMISSION

In this work, we focus on the downlink of both the macro and the small cells tiers. The macrocell serves G user groups in each transmission resource block using JSDM. The received signal vector \mathbf{y}_g of the users located in group g can be expressed as

$$\begin{aligned} \mathbf{y}_g &= \mathbf{H}_{g,0}^H \mathbf{B}_g \mathbf{P}_g \mathbf{d}_g + \sum_{g' \neq g} \mathbf{H}_{g,0}^H \mathbf{B}_{g'} \mathbf{P}_{g'} \mathbf{d}_{g'} \\ &+ \sum_{f \in \mathcal{S}_A} \mathbf{H}_{g,f}^H \mathbf{Q}_{f,f} \mathbf{s}_f + \mathbf{z}_g, \end{aligned} \quad (5)$$

where \mathbf{d}_g is the $S_g \times 1$ vector of transmitted data symbols to the group g users, $\mathbf{P}_g \in \mathbb{C}^{b_g \times S_g}$ and $\mathbf{B}_g \in \mathbb{C}^{M \times b_g}$ are the precoding and pre-beamforming matrices² for group g of the JSDM scheme, and $\mathbf{H}_{g,0} = [\mathbf{h}_{g1,0} \dots \mathbf{h}_{gS_g,0}]$ is the channel matrix between the macrocell antenna array and the served users in group g . Notice the structure of the two stage JSDM precoder: it is split into the product $\mathbf{B}_g \mathbf{P}_g$ where \mathbf{B}_g depends only on the channel second-order statistics information $\{\mathbf{R}_g : g = 1, \dots, G\}$, and \mathbf{P}_g on the instantaneous realization of the projected channels $\{\mathbf{H}_{g,0}^H \mathbf{B}_g : g = 1, \dots, G\}$. The matrix $\mathbf{H}_{g,f} = [\mathbf{h}_{g1,f} \dots \mathbf{h}_{gS_g,f}]$ contains the channels between users in group g and small cell f , $\mathbf{Q}_{f,f}$ is the precoding vector used by small cell f , \mathbf{s}_f is the vector of data symbols transmitted by small cell f , \mathcal{S}_A is the set of active small cells resulting from the various inter-tier interference coordination strategies discussed in Section IV, and \mathbf{z}_g denotes the additive white Gaussian noise, with i.i.d. components $\sim \mathcal{CN}(0, 1)$.

For JSDM, we assume that the macrocell performs *per group processing* (PGP) (see [3]), which has the advantage of significantly reducing the channel state information feedback requirement while ensuring that the second stage precoding matrix \mathbf{P}_g can be independently designed across all groups g . PGP results in an additional inter-group interference term, given by the sum over $g' \neq g$ in (5), which can be eliminated by block diagonalization or by serving groups of users with disjoint angular support using DFT pre-beamforming, in the limit of very large M . For more details, please refer to [3]. We consider zero forcing beamforming for \mathbf{P}_g , given by

the Moore-Penrose pseudoinverse of $\mathbf{B}_g^H \mathbf{H}_{g,0}$, up to a power normalization scalar factor ζ_g . Specifically, we have

$$\mathbf{P}_g = \zeta_g \mathbf{B}_g^H \mathbf{H}_{g,0} \left(\mathbf{H}_{g,0}^H \mathbf{B}_g \mathbf{B}_g^H \mathbf{H}_{g,0} \right)^{-1}. \quad (6)$$

Notice that (5) also includes a term that captures the interference caused by all the active small cells (in the set \mathcal{S}_A) to the users of group g .

Small cells serve \bar{S}_f users using zero forcing beamforming. For tractability, we let \bar{S}_f to be the same in all the small cells, and equal to a certain fraction of the number of antennas L , i.e., $\bar{S}_f = \bar{S} = \beta L \forall f \in \mathcal{S}$, where β is a design parameter that depends on the precoding scheme. The received signal of users in group f served by its corresponding small cell is

$$\begin{aligned} \bar{\mathbf{y}}_f &= \mathbf{H}_{f,f}^H \mathbf{Q}_{f,f} \mathbf{s}_f + \sum_{f' \neq f, f' \in \mathcal{S}_A} \mathbf{H}_{f,f'}^H \mathbf{Q}_{f',f'} \mathbf{s}_{f'} \\ &+ \sum_{g=1}^G \mathbf{H}_{f,0}^H \mathbf{B}_g \mathbf{P}_g \mathbf{d}_g + \bar{\mathbf{z}}_f \end{aligned} \quad (7)$$

where $\mathbf{H}_{f,f'} = [\mathbf{h}_{f1,f'} \dots \mathbf{h}_{f\bar{S}_f,f'}]$ is the channel matrix for the users in group f and the small cell array of group f' , $\mathbf{Q}_{f,f}$ is the zero forcing precoding matrix of small cell f (given as the column-normalized Moore Penrose pseudo inverse of $\mathbf{H}_{f,f}$) and $\bar{\mathbf{z}}_f$ is the additive white Gaussian noise. The users in group f suffer interference from all the other active small cells (sum over $f' \neq f$ in (7)), along with interference from the macrocell (sum over g in (7)). We also assume that the small cells transmit at peak power P_1 , and all users' data stream are allocated equal power P_1/\bar{S} .

IV. INTER-TIER COORDINATION STRATEGIES

In this section, we outline three different inter-tier coordination strategies in order to reduce interference experienced by both the macro and small cell user groups.

A. Spatial Blanking

In this scheme, implicit inter-tier coordination is achieved by spatial blanking. In each transmission resource block, the macrocell selects and serves a subset of size G of user groups from the set \mathcal{M} , according to some fair group selection algorithm ensuring that every user group in the set \mathcal{M} is given equal air time. Our user group selection algorithm does not take into account the small scale fading statistics and makes decisions based only on the second order statistics of the user groups, thus eliminating the need for explicit channel state feedback from all the users at every scheduling slot. The feedback is only required to design the PGP precoder \mathbf{P}_g for every group g , which depends on the instantaneous effective channel $\mathbf{B}_g^H \mathbf{H}_{g,0}$, as stated earlier.

The algorithm works by maintaining a priority vector for every user group in set \mathcal{M} , i.e., the user groups that do not contain a small cell, and updating the priority vector so as to maintain a high priority for user groups that have not been served in the recent time slots. We first initialize the priority of all user groups to 1 and, at the end of the scheduling slot, we increment the priorities of the non selected user groups

²The pre-beamforming dimension b_g determines the amount of channel state information to be fed back to the transmitter and should be carefully optimized [3].

by 1, while keeping the priorities of the selected user groups unchanged. This guarantees an equal opportunity to all the user groups to be served by the macrocell, so that no user group is left starving. At every scheduling slot, the macrocell first forms a set of user groups that have disjoint angular support with the already selected user groups. Then, from this set, it greedily selects the one that has the highest priority and causes the least interference to the already selected user groups. More formal details of the algorithm are omitted due to space constraints but can be found in the extended version of this paper [9].

After selecting user groups, the macrocell selects (uniformly at random) S_g users from each selected group g , and serves them using JSDM with PGP and zero-forcing precoding in each group to eliminate the intra-group interference as explained before. From the results in [3], for large number of antennas M and covariance group rank r_g , it is known that the optimal value of S_g is given by βr_g for some design parameter $\beta < 1$ that can be optimized depending on the scattering geometry. In order to keep the problem tractable and obtain meaningful results we use the same value of β for all groups.

Remark 1. In implementing the user selection algorithm, we set the pre-beamforming matrix \mathbf{B}_g of every user group being served by the macrocell as $\mathbf{B}_g = \mathbf{U}_g^*$, where \mathbf{U}_g^* contains the eigenvectors corresponding to the dominant eigenvalues of the corresponding channel covariance matrix \mathbf{R}_g . This greatly simplifies the algorithm as opposed to performing block diagonalization (or approximate block diagonalization), which requires recomputing the pre-beamformers every time a new user group is selected in the scheduled pool of the macrocell. Such a simplification does not come at an overly pessimistic performance price when the number of antennas M is very large, since in this case the channel covariances of two user groups with disjoint angular support are approximately orthogonal to each other and therefore, (approximate) block diagonalization is implicitly achieved.

In this scheme, all the small cells are active on all transmission resource blocks and transmit to their own users using zero forcing beamforming at their own peak total power. Recall that small cells serve $\bar{S} = \beta L$ users in every transmission resource block, selected at random in order to give equal air time to all their users. When the number of user groups G served by the macrocell is not too large, thanks to the inherent directionality in pre-beamforming achieved by JSDM, the macrocell implicitly mitigates the interference at the small cells that are not aligned in the direction of the pre-beamforming vectors. Thus, compared to a naive uncoordinated scheme that serves users isotropically instead of co-located user groups, our strategy achieves some non-trivial interference suppression benefits as a result of spatial blanking. This is illustrated next.

Example 1 (Spatial Blanking). Consider a toy example with two active user groups: one served by the macrocell and the other by a small cell. The macro user group is located at a distance of 0.2 km from the macrocell and the small cell user group's location is parameterized by (r, θ) as shown in the inset of the first figure of Fig. 2. When $\theta = 0$, both the user groups are aligned when seen from the macrocell,

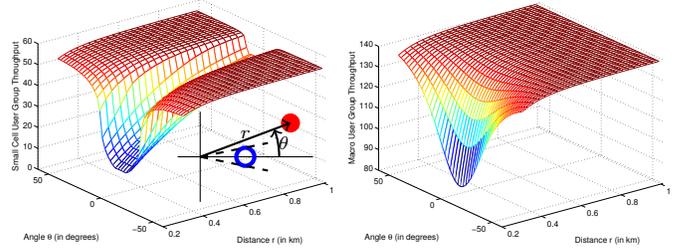


Fig. 2. The toy example to demonstrate spatial blanking in Example 1. (first) The throughput of the small cell user group. (second) The throughput of the macrocell user group. The inset in the first figure shows the layout used in this example. The hollow and filled circles denote macro and small cell user groups, respectively.

which results in a high cross-tier interference at the small cell user group. As $|\theta|$ is increased, this interference almost vanishes resulting in high small cell throughput (first figure of Fig. 2) because of “spatial blanking”. Also note that the distance between the two user groups dictates the cross-tier interference seen at the macrocell user group (second figure of Fig. 2). When this distance is reduced, the macro user rate drops significantly, which necessitates active interference coordination strategies to complement spatial blanking. The simulation parameters used for Fig. 2 are tabulated in Table I.

B. Spatial Blanking with ON/OFF Small Cells

In this coordination scheme, the macrocell first chooses a set of user groups to be served using the same user group selection algorithm of Section IV-A. The small cells then use this information along with the knowledge of the cross-tier interference and useful signal strengths to implement a simple ON/OFF strategy for the given scheduling slot. Explicitly, a small cell decides to shut down its transmission based on the amount of cross-tier interference it receives or causes to the scheduled macrocell user groups. In order to make this decision, the small cell compares the interference that its users receive from the macrocell with the strength of their useful signal, as well as the interference that it causes to the macrocell user groups with the strength of their useful signal.

More formally, a small cell serving a user group f decides to transmit or shut down according to the following criterion

$$\text{Small Cell } f = \begin{cases} \text{ON} & \frac{P_0}{S} J_{f,0}^{\text{mc}} \leq \epsilon_1 D_{f,f}^{\text{sc}} \frac{P_1}{S} \text{ and} \\ & J_{g,f}^{\text{sc}} P_1 \leq \epsilon_2 D_{g,0}^{\text{mc}} \frac{P_0}{S} \quad \forall g \in \mathcal{G} \\ \text{OFF} & \text{otherwise} \end{cases} \quad (8)$$

where ϵ_1 and ϵ_2 are design parameters that can be set to achieve a desired tradeoff between the macrocell and small cell throughputs. $D_{f,f}^{\text{sc}}$ denotes the useful signal strength to a user in group f that is being served by a small cell and $J_{f,0}^{\text{mc}}$ denotes the amount of interference caused by the macrocell to the same. $D_{g,0}^{\text{mc}}$ is the useful signal strength of a user in group g served by the macrocell and $J_{g,f}^{\text{sc}}$ is the interference caused by small cell f to this particular user. Expressions for $J_{f,0}^{\text{mc}}$, $D_{f,f}^{\text{sc}}$, $J_{g,f}^{\text{sc}}$ and $D_{g,0}^{\text{mc}}$ can be computed using the techniques of deterministic equivalents [10], [3] from (5) and (7). They are omitted here due to space limitations but can be found in the extended version of this paper [9].

TABLE I
LIST OF SIMULATION PARAMETERS

Parameter	Value
No. of user groups, N_u	500
Cell radius, R_{mc}	1 km
Cut off distance, d_0	50 m
Macrocell exclusion radius, R_{excl}	100 m
Path loss exponent, α	3.5
Cell edge SNR	10 dB
Wall loss, w	5 dB
Loading factor, β	0.8
ON/OFF algorithm thresholds, ϵ_1, ϵ_2	0.1
OFFLOAD parameter, γ	1
No. of user groups served by macrocell, G	5

C. Spatial Blanking with Offloading

In the ON/OFF coordination strategy, the cross-tier interference is mitigated by the small cells deciding on their transmissions. This may be disadvantageous to a user group being served by a small cell located in the vicinity of another user group not covered by a small cell, because the small cell would always shut down whenever the macrocell schedules to serve such an uncovered user group. We alleviate this problem by associating some of these “bottleneck” user groups devoid of small cell to the nearest small cell, so that the small cell can now serve both the user groups using TDMA. We call this the “OFFLOAD” strategy, because some of the macrocell user groups are being offloaded to neighboring small cells in order to be able to be served more efficiently and cause less harm in terms of their imposed interference constraints. This strategy is intended to increase the throughput of macrocell user groups which are located at the cell edge and have a small cell close to them, since such user groups at the cell edge already receive a low useful signal strength from the macrocell and, in addition, are likely to block the transmission of the neighboring small cells. However, the offloading approach results in a decrease of the rates observed in the small cells, because of the fact that the small cells may now serve two or more user groups in TDMA, thereby reducing the average throughput of its own users by a factor equal to the number of associated user groups. Formally, a small cell f decides to serve a user group g which normally would have been served by the macrocell according to the following condition:

$$\begin{aligned} D_{g,f}^{sc} \frac{P_1}{S} &> \gamma D_{g,0}^{mc} \frac{P_0}{S_g} && \text{OFFLOAD} \\ D_{g,f}^{sc} \frac{P_1}{S} &\leq \gamma D_{g,0}^{mc} \frac{P_0}{S_g} && \text{NO OFFLOAD,} \end{aligned} \quad (9)$$

where $D_{g,f}^{sc}$ is the direct link gain between a small cell f and a user in group g and $D_{g,0}^{mc}$ is the direct link gain of the same user from the macrocell, assuming that the macrocell serves the user group g in isolation. $D_{g,f}^{sc}$ and $D_{g,0}^{mc}$ can be computed from (5) and (7) using the techniques of deterministic equivalents. The parameter γ moderates the fraction of macrocell user groups being offloaded, e.g., a small value of γ indicates that more macrocell user groups will be offloaded.

It is worth mentioning that the offloading strategy also reduces the burden on the small cells compared to the ON/OFF strategy, where the small cells were required to make decisions on every scheduling slot³ by knowing the cross-tier interfer-

³A scheduling slot typically corresponds to the time over which the channel remains fairly constant, which for normal pedestrian speeds of 1 m/s and 3 GHz carrier frequency, corresponds to a coherence time of ~ 100 ms.

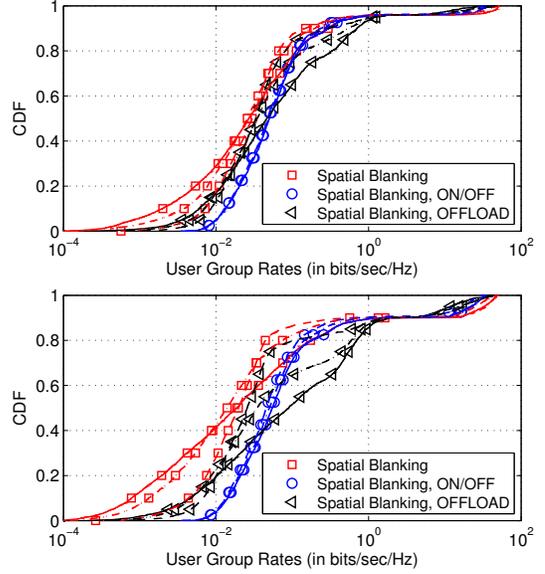


Fig. 3. CDF of user group rates for different inter-tier coordination schemes and different deployments. (first) $N_f = 20$, and (second) $N_f = 50$. Cell-edge: solid line, cell-interior: dashed line, and uniform: dash-dotted line.

ence. In fact, in this case, the small cells can decide in advance the user groups to serve and relay this information to the macrocell, so that the latter does not include these user groups in its scheduling selection process. Such information must be conveyed on a time scale of a few seconds (for low mobility scenarios, we can safely assume that the distribution of the user groups remains static over long periods of time when compared to the time slot for making scheduling decisions) without any significant protocol overhead.

V. NUMERICAL RESULTS

The default values of the simulation parameters are listed in Table I. We present results for $N_f = 20$ and $N_f = 50$, corresponding to a low and high density of small cells, respectively. For each case, we further consider three ways the small cells can be deployed: (i) uniform, (ii) cell-interior, and (iii) cell edge, which were formally defined in Section II. In all the figures, the results corresponding to the uniform, cell-interior and cell-edge deployments are presented using dash-dotted, dashed and solid lines respectively. The peak power of the small cells is set to $P_1 = \frac{P_0}{100}$, corresponding to a value that is 20 dB less than the peak power used by the macrocell. The macrocell power P_0 is calculated from the cell edge SNR. Due to space limitation, we only present *representative* results here, and delegate the rest to the extended version [9].

We first present the CDFs of the rates achieved by all the user groups in Fig. 3. The low and high rates respectively correspond to the rates observed by the macro and small cell user groups. As expected, the rates of macrocell user groups can be significantly improved by using active interference coordination (ON/OFF or OFFLOAD) in addition to spatial blanking. While the benefits of ON/OFF strategy are more prominent for the low-rate user groups, the OFFLOAD strategy is more beneficial for the medium-rate user groups. The

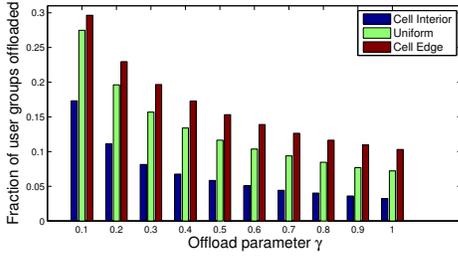


Fig. 4. Fraction of users offloaded from the macrocell to the small cells for different deployments and varying γ ($N_f = 20$).

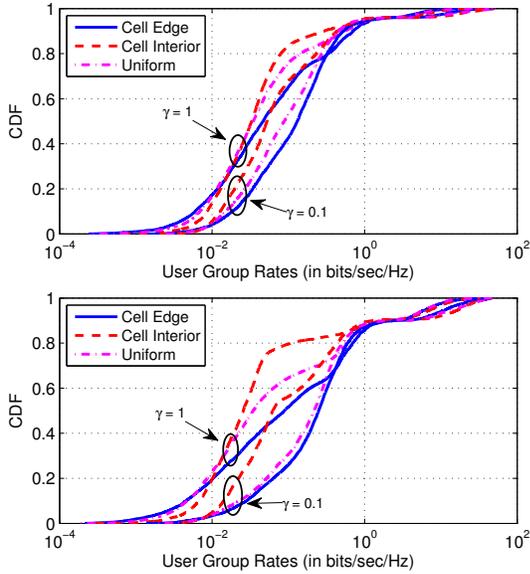


Fig. 5. CDF of user group rates for coordination scheme ‘OFFLOAD’ and different deployments. (first) $N_f = 20$, and (second) $N_f = 50$.

advantages of offloading are more prominent when more small cells are deployed at the edge, which is because of the fact that a user group located at the cell edge has a greater chance of being offloaded from the macrocell to a small cell because it receives a stronger signal from the latter. Note that both these strategies reduce small cell user group rates slightly because the small cells are either turned OFF for some fraction of time (ON/OFF) or use resources to serve macrocell user groups (OFFLOAD). In general, the macrocell user group rates show a decreasing trend on going from a cell interior deployment to a cell edge deployment of the small cells, because of the reduced signal strength to the macrocell user groups located far from the macrocell. The sharp transition in the plots at the higher end of the user group rates is because of the drastic difference in the rates observed at user groups served by the macrocell and those served by the small cell.

In Figures 4 and 5, we study the effect of varying the offload parameter γ . In Fig. 4, we note that the fraction of offloaded user groups decreases with an increase in the parameter γ . Therefore, γ can be tuned to achieve a desired objective. Note that for the same γ , the fraction of offloaded user groups increases as we go from a cell-interior to a cell-edge deployment of the small cells, because of the decrease in

the signal strength of the macrocell users as we go towards the cell edge. In Fig. 5, we present the CDFs of user group rates for $\gamma = 1$ and 0.1. As expected, a lower value of γ leads to more aggressive offloading, thereby increasing the rates of the offloaded macro user groups, although at the cost of reduced small cell rates. Note that due to limited backhaul capacity of the small cells, these reduced small cells rates are not a problem as long as they are higher than the bottleneck.

VI. CONCLUSION

For a HetNet with massive-MIMO macrocells and multi-antenna small cells, we presented a new interference coordination scheme, termed *spatial blanking*, under which the macrocell concentrates its transmission energy in the direction of its scheduled hotspots using JSMD, thereby creating transmission opportunities for the small cells located in the other directions. This implicit interference suppression is further complemented by two active low-overhead interference coordination strategies: (i) turn OFF a small cell intelligently if it causes or suffers from excessive interference, and (ii) offload macrocell hotspots to small cells to improve fairness in rates across hotspots. This presents a more efficient alternative to popular eICIC techniques such as fractional frequency reuse and almost-blank subframes, which orthogonalize interfering transmissions over time or frequency, thereby leading to an inefficient utilization of precious resources. Two promising extensions of this work are: (i) to include the effect of common scatterers, which would likely require more complex scheduling strategies, and (ii) to include isolated users (i.e., not a part of a hotspot) in the model, which would require resource sharing with the hotspots, leading to throughput degradation.

REFERENCES

- [1] Cisco, “Cisco visual networking index: Global mobile data traffic forecast update, 2011 - 2016,” white paper, Feb. 2012.
- [2] D. Lopez-Perez, I. Guvenc, G. D. L. Roche, M. Kountouris, T. Quek, and J. Zhang, “Enhanced intercell interference coordination challenges in heterogeneous networks,” *IEEE Wireless Commun.*, vol. 18, no. 3, pp. 22 – 30, Jun. 2011.
- [3] A. Adhikary, J. Nam, J.-Y. Ahn, and G. Caire, “Joint spatial division and multiplexing: The large-scale array regime,” *IEEE Trans. on Info. Theory*, vol. 59, no. 10, pp. 6441 – 6463, Oct. 2013.
- [4] T. Marzetta, “Noncooperative cellular wireless with unlimited numbers of base station antennas,” *IEEE Trans. on Wireless Commun.*, vol. 9, no. 11, pp. 3590 – 3600, Nov. 2010.
- [5] H. Huh, G. Caire, H. C. Papadopoulos, and S. A. Ramprasad, “Achieving ‘massive MIMO’ spectral efficiency with a not-so-large number of antennas,” *IEEE Trans. on Wireless Commun.*, vol. 11, no. 9, pp. 3226 – 3239, Sep. 2012.
- [6] W.-Y. Lee, “Effects on correlation between two mobile radio base-station antennas,” *IEEE Trans. on Veh. Technology*, vol. 22, no. 4, pp. 130 – 140, Nov. 1973.
- [7] T. Yoo, N. Jindal, and A. Goldsmith, “Multi-antenna downlink channels with limited feedback and user selection,” *IEEE Journal on Sel. Areas in Commun.*, vol. 25, no. 7, pp. 1478–1491, Sep. 2007.
- [8] H. Shirani-Mehr, G. Caire, and M. J. Neely, “MIMO downlink scheduling with non-perfect channel state knowledge,” *IEEE Trans. on Commun.*, vol. 58, no. 7, pp. 2055–2066, Jul. 2010.
- [9] A. Adhikary, H. S. Dhillon, and G. Caire, “Massive-MIMO meets HetNet: Interference coordination through spatial blanking,” submitted to *IEEE Journal on Sel. Areas in Commun.*, July 2014. Available online: goo.gl/W7DAVX.
- [10] S. Wagner, R. Couillet, M. Debbah, and D. T. Slock, “Large system analysis of linear precoding in correlated MISO broadcast channels under limited feedback,” *IEEE Trans. on Info. Theory*, vol. 58, no. 7, pp. 4509 – 4537, Jul. 2012.