Distributed Interference Management Policies for Heterogeneous Small Cell Networks

Kartik Ahuja, *Student member, IEEE,* Yuanzhang Xiao, *Student member, IEEE,*
and Mihaela van der Schaar, *Fellow, IEEE*

Abstract—We study the problem of *distributed* interference management in a network of heterogeneous small cells with different cell sizes, different numbers of user equipments (UEs) served, and different throughput requirements by UEs. We consider the uplink transmission, where each UE determines when and at what power level it should transmit to its serving small cell base station (SBS). We propose a general framework for designing distributed interference management policies, which exploits weak interference among non-neighboring UEs by letting them transmit simultaneously (i.e., spatial reuse), while eliminating strong interference among neighboring UEs by letting them transmit in different time slots. The design of optimal interference management policies has two key steps. Ideally, we need to find all the subsets of non-interfering UEs, i.e., the maximal independent sets (MISs) of the interference graph, but this is NP-hard (non-deterministic polynomial time) even when solved in a centralized manner. Then, in order to maximize some given network performance criterion subject to UEs' minimum throughput requirements, we need to determine the optimal fraction of time occupied by each MIS, which requires global information (e.g., all the UEs' throughput requirements and channel gains). In our framework, we first propose a *distributed* algorithm for the UE-SBS pairs to find *a subset of* MISs in logarithmic time (with respect to the number of UEs). Then we propose a novel problem reformulation which enables UE-SBS pairs to determine the optimal fraction of time occupied by each MIS with only local message exchange among the neighbors in the interference graph. Despite the fact that our interference management policies are distributed and utilize only local information, we can *analytically* bound their performance under a wide range of heterogeneous deployment scenarios in terms of the competitive ratio with respect to the optimal network performance, which can only be obtained in a centralized manner with NP complexity. Remarkably, we prove that the competitive ratio is independent of the network size. Through extensive simulations, we show that our proposed policies achieve significant performance improvements (ranging from 150% to 700%) over state-of-the-art policies.

Index Terms—Heterogeneous Network, Small Cell, Interference Management, Power Control, Interference Graph.

I. INTRODUCTION

Dense deployment of low-cost heterogeneous small cells (e.g. picocells, femtocells) has become one of the most effective solutions to accommodate the exploding demand for wireless spectrum [1] [2] [3]. On one hand, dense deployment of small cells significantly shortens the distances between small cell base stations (SBSs) and their corresponding user equipments (UEs), thereby boosting the network capacity. On the other hand, dense deployment also shortens the distances between neighboring SBSs, thereby potentially increasing the inter-cell interference. Hence, while the solution provided by the dense deployment of small cells is promising, its success depends crucially on interference management by the small cells. Efficient interference management is even more challenging in heterogeneous small cell networks, due to the lack of central coordinators, compared to that in traditional cellular networks.

In this paper, we propose a novel framework for designing interference management policies in the uplink of small cell networks, which specify when and at what power level each UE should transmit[1]. Our proposed design framework and the resulting interference management policies fulfill *all* the following important requirements:

- *Deployment of heterogeneous small cell networks*:
  Existing deployments of small cell networks exhibit significant heterogeneity such as different types of small cells (picocells and femtocells), different cell sizes, different number of UEs served, different UEs' throughput requirements etc.
- *Interference avoidance and spatial reuse*: Effective interference management policies should take into account the strong interference among neighboring UEs, as well as the weak interference among non-neighboring UEs. Hence, the policies should effectively avoid interference among neighboring UEs and use spatial reuse to take advantage of the weak interference among non-neighboring UEs.
- *Distributed implementation with local information and message exchange*: Since there is no central coordinator in small cell networks, interference management policies need to be computed and implemented by the UEs in a distributed manner, by exchanging only local information through local message exchanges among neighboring UE-SBS pairs.
- *Scalability to large networks*: Small cells are often deployed over a large scale (e.g., in a city). Effective interference management policies should scale in large networks, namely achieve efficient network performance

[1]Although we focus on uplink transmissions in this paper, our framework can be easily applied to downlink transmissions.

while maintaining low computational complexity.

- *Ability to optimize different network performance criteria*: Under different deployment scenarios the small cell networks may have different performance criteria, e.g., weighted sum throughput or max-min fairness. The design framework should be general and should prescribe different policies to optimize different network performance criteria.
- *Performance guarantees for individual UEs*: Effective interference management should provide performance guarantees (e.g., minimum throughput guarantees) for individual UEs.

As we will discuss in detail in Section II, existing state-of-the-art policies for interference management cannot simultaneously fulfill all of the above requirements.

Next, we describe our key results and major contributions:

1. We propose a general framework for designing distributed interference management policies that maximizes the given network performance criterion subject to each UE's minimum throughput requirements. The proposed policies schedule maximal independent sets (MISs)[2] of the interference graph to transmit in each time slot. In this way, they avoid strong interference among neighboring UEs (since neighboring UEs cannot be in the same MIS), and efficiently exploit the weak interference among UEs in a MIS by letting them to transmit at the same time.

2. We propose a distributed algorithm for the UEs to determine a subset of MISs. The subset of MISs generated ensures that each UE belongs to at least one MIS in this subset. Moreover, the subset of MISs can be generated in a distributed manner in logarithmic time (logarithmic in the number of UEs in the network) for bounded-degree interference graphs[3]. The logarithmic convergence time is significantly faster than the time (linear or quadratic in the number of UEs) required by the distributed algorithms for generating subsets of MISs in [4]–[6].

3. Given the computed subsets of MISs, we propose a distributed algorithm in which each UE determines the optimal fractions of time occupied by the MISs with only local message exchange. The message is exchanged only among the UE-SBS pairs that strongly interfere with each other, i.e. among neighbors in the interference graph. The distributed algorithm will output the optimal fractions of time for each MIS such that the given network performance criterion is maximized subject to the minimum throughput requirements.

4. Under a wide range of conditions, we analytically characterize the competitive ratio of the proposed distributed policy with respect to the optimal network performance. Importantly, we prove that the competitive ratio is independent of the net-

work size, which demonstrates the scalability of our proposed policy in large networks. Remarkably, the constant competitive ratio is achieved even though our proposed policy requires only local information, is distributed, and can be computed fast, while the optimal network performance can only be obtained in a centralized manner with global information (e.g., all the UEs' channel gains, maximum transmit power levels, minimum throughput requirements) and NP (non-deterministic polynomial time) complexity.

5. Through simulations, we demonstrate significant (from 160% to 700 %) performance gains over state-of-the-art policies. Moreover, we show that our proposed policies can be easily adapted to a variety of heterogeneous deployment scenarios, with dynamic entry and exit of UEs.

The rest of the paper is organized as follows. In Section II we discuss the related works and their limitations. We describe the system model in Section III. Then we formulate the interference management problem and give a motivating example in Section IV. We propose the design framework in Section V, and demonstrate the performance gain of our proposed policies in Section VI. Finally, we conclude the paper in Section VII.

II. RELATED WORKS

State-of-the-art interference management policies can be divided into three main categories: policies based on power control, policies based on spatial reuse, and policies based on joint power control and spatial reuse. In the following, we discuss their limitations for the considered distributed interference management problem in heterogeneous small cell networks. We will list some representative references in this section; a detailed list can be found in the online report [7].

A. Distributed Interference Management Based on Power Control

Policies based on distributed power control (representative works [8]–[10]) have been used for interference management in both cellular and ad-hoc networks. In these policies, all the UEs in the network transmit at *constant* power levels all the time (provided that the system parameters remain the same)[4]. For this reason, we refer to them as constant power control policies in the rest of this paper. The major limitation of constant power control policies is the difficulty in providing minimum throughput guarantees for each UE, especially in the presence of strong interference. Some works [8]–[10] use pricing to mitigate the strong interference. However, they cannot strictly guarantee the UEs' minimum throughput requirements [8]–[10]. Indeed, the low throughput experienced by some users, caused by strong interference, is the fundamental limitation of constant power control policies, even for the optimal constant power control policy obtained by a central controller[5] [12]. Since strong interference is very

---

[2]Consider the interference graph of the network, where each vertex is a UE-SBS pair and each edge indicates strong interference between the two vertices. An independent set (IS) is a set of vertices in which no pair is connected by an edge. An IS is a MIS if it is not a proper subset of another IS.

[3]Bounded-degree graphs are the graphs whose maximum degree can be bounded by a constant independent of the size of the graph, i.e., $\Delta = \mathcal{O}(1)$. As we will show in Theorem 5, for the interference graphs that are not bounded-degree graphs, even the centralized solution, given all the MISs, cannot satisfy the minimum throughput requirements.

[4]Although some power control policies [8]–[10] go through a transient period of adjusting the power levels before the convergence to the optimal power levels, the users maintain constant power levels after the convergence.

[5]In the case of average sum throughput maximization given the minimum average throughput constraints of the UEs, the power control policies are inefficient if the feasible rate region is non-convex [11] .

common in dense small cell deployments (e.g., in offices and apartments where SBSs are installed close to each other [13]), constant power control policies do not perform well in these scenarios. Note that there exist a different strand of works based on [14], which proposes distributed algorithms to achieve the desired minimum throughput requirement for each UE with the objective of minimizing transmit power levels. These works cannot optimize network performance criteria such as weighted sum throughput, max-min fairness etc., and hence are suboptimal under these performance criteria.

B. Distributed Spatial Reuse Based on Maximal Independent Sets

An efficient solution to mitigate strong interference is spatial reuse, in which only a subset of UEs (those who do not significantly interfere with each other) transmit at the same time. Spatial Time reuse based Time Division Multiple Access (STDMA) has been widely used in existing works on broadcast scheduling in multi-hop networks [4] [6][6]. Specifically, these policies construct a cyclic schedule such that in each time slot an MIS of the interference graph is scheduled. The constructed schedule ensures that each UE is scheduled at least once in the cycle.

In terms of performance, STDMA policies [4]–[6] cannot guarantee the minimum throughput requirement of each UE, and usually adopt a fixed scheduling (i.e. follow a fixed order in which the MISs are scheduled), which may be very inefficient depending on the given network performance criteria. For example, the policies in [6] are inefficient in terms of fairness. In terms of complexity, for the distributed generation of the subsets of MISs, the STDMA policies in [4]–[6] require an ordering of all the UEs, and have a computational complexity (in terms of the number of steps executed by the algorithm) that scales as $\mathcal{O}(|V|)$ (in [5] [6]) or $\mathcal{O}(|V||E|))$ (in [4]), where $|V|$ and $|E|$ are the number of vertices/UEs and the number of edges in the interference graph, respectively. Hence, in large-scale dense deployments, the complexity grows superlinearly with the number of UEs, making the policies difficult to compute. By contrast, our proposed distributed algorithm for generating subsets of MISs does not require the ordering of all the UEs, and has a complexity that scales as $\mathcal{O}(\log |V|)$, namely logarithmically with the number of the UEs, for bounded-degree graphs.[7]

Finally, the STDMA policies in [4]–[6] are designed for the MAC layer and assume that all the UEs are homogeneous at the physical layer. In practice, different UEs are heterogeneous due to their different distances from their SBSs, their different maximum transmit power levels, etc. This heterogeneity is important, and will be considered in our design framework.

[6]These works [4]–[6] do not have exactly the same model as in our setting. However, these works can be adapted to our model. Hence, we also compare with these works to have a comprehensive literature review.

[7]As we will show in Theorem 5, for graphs which do not have bounded degrees, even a centralized solution based on all the MISs cannot satisfy the minimum throughput requirements.

C. Distributed Power Control and Spatial Reuse For Multi-Cell Networks

The works discussed in the above two subsections either focus on distributed power control in the physical layer [8]–[10] or focus on distributed spatial reuse in the MAC layer [4]–[6]. Similar to our paper, some works (see [15] [16] for representative works) adopted a cross-layer approach and proposed joint distributed power control and spatial reuse for multi-cell networks. Although these works schedule a subset of UEs to transmit at each time slot, the subset is not the MIS of the interference graph [15] [16] . For example, the policies in [15] [16], called power matched scheduling (PMS) policies, schedule one UE from each small cell at the same time, even if some UEs from different cells are very close to each other. In this case, these UEs will experience strong inter-cell interference. Hence, the works in [15] [16] cannot perfectly eliminate strong interference from neighboring cells and exploit weak interference from non-neighboring cells. Moreover, the works in [15] [16] cannot provide minimum throughput guarantees for the UEs.

III. SYSTEM MODEL

A. Heterogeneous Network of Small Cells

We consider a heterogeneous network of K small cells operating in the same frequency band (see Fig. 1), which represents a common deployment scenario considered in practice [2] [10] [17]. Note that the small cells can be of different types (e.g. picocells, femtocells, etc.) and thereby belong to different tiers in the heterogeneous network. Each small cell j has one SBS, (SBS j), which serves a set of UEs under a closed access scenario [10]. Denote the set of UEs by $\mathcal{U} = \{1, ..., N\}$. We write the association of UEs to SBSs as a mapping $T : \{1, ..., N\} \to \{1, .., K\}$, where each UE i is served by SBS $T(i)$. The interference graph G of the network has N vertices, each of which is a UE-SBS pair. There is an edge between two pairs/vertices if their cross interference is high (rules for deciding if interference is high will be discussed in Section V).

We focus on the uplink transmissions; the extension to downlink transmissions is straightforward when each SBS serves one UE at a time (e.g., TDMA among the UEs connected to the same SBS). Each UE-i chooses its transmit power p_i from a compact set $\mathcal{P}_i \subseteq \mathbb{R}_+$. We assume that $0 \in \mathcal{P}_i, \ \forall i \in \{1, ..., N\}$, namely any UE can choose not to transmit. The joint power profile of all the UEs is denoted by $\mathbf{p} = (p_1, ..., p_N) \in \mathcal{P} \triangleq \Pi_{i=1}^{N} \mathcal{P}_i$. Under the joint power profile $\mathbf{p}$, the signal to interference and noise ratio (SINR) of UE i's signal, experienced at its serving SBS $j = T(i)$, can be calculated as $\gamma_i(\mathbf{p}) = \frac{g_{ij}p_i}{\sum_{k=1, k \neq i}^{N} g_{kj}p_k + \sigma_j^2}$, where g_{ij} is the channel gain from UE i to SBS j, and σ_j^2 is the noise power at SBS j. Since the UEs do not cooperate to encode their signals to avoid interference, each UE-SBS pair treats the interference from other UEs as white noise. Hence, each UE i gets the following throughput [15], $r_i(\mathbf{p}) = \log_2(1 + \gamma_i(\mathbf{p}))$[8].

[8]We use the Shannon capacity here. However, our analysis is general and applies to the throughput models that consider the modulation scheme used.
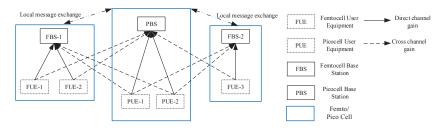
Figure 1. Illustration of a heterogeneous small cell network.

B. Interference Management Policies

The system is time slotted at $t = 0,1,2...$, and the UEs are assumed to be synchronized[9]. At the beginning of each time slot t, each UE i decides its transmit power p_i^t and obtains a throughput of $r_i(\mathbf{p}^t)$. Each UE i's strategy, denoted by $\pi_i : \mathbb{Z}_+ \triangleq \{0, 1, ..\} \to \mathcal{P}_i$, is a mapping from time t to a transmission power level $p_i \in \mathcal{P}_i$. The interference management policy is then the collection of all the UEs' strategies, denoted by $\boldsymbol{\pi} = (\pi_1, ..., \pi_N)$. The average throughput for UE-i is given by $R_i(\boldsymbol{\pi}) = \lim_{T\to\infty} \frac{1}{T+1} \sum_{t=0}^{T} r_i(\mathbf{p}^t)$, where $\mathbf{p}^t = (\pi_1(t), ..., \pi_N(t))$ is the power profile at time t. We assume that the channel gains are fixed over the considered time horizon as in [15] [18]–[21]. However, we will illustrate in Section VI that our framework performs well under time-varying channel conditions (e.g., due to fading) as well.

An interference management policy $\boldsymbol{\pi}^{const}$ is a constant power control policy [8]–[10] if $\boldsymbol{\pi}^{const}(t) = \mathbf{p}$ for all t. As we have discussed before, our proposed policy is based on MISs of the interference graph. Given an interference graph, we write $\mathbf{I} = \{I_1, ..., I_{N_{MIS}}\}$ as the set of all the MISs of the interference graph. Let $\mathbf{p}^{I_j}$ be a power profile in which the UEs in the MIS I_j transmit at their maximum power levels and the other UEs do not transmit, namely $p_k^{I_j} = p_k^{max} \triangleq \max \mathcal{P}_k$ if $k \in I_j$ and $p_k = 0$ otherwise. Let $\mathcal{P}^{MIS} = \{\mathbf{p}^{I_1}, ..., \mathbf{p}^{I_{N_{MIS}}}\}$ be the set of all such power profiles. Then $\boldsymbol{\pi}$ is a policy based on MIS if $\boldsymbol{\pi}(t) \in \mathcal{P}^{MIS}$ for all t. We denote the set of policies based on MISs by $\Pi^{MIS} = \{\boldsymbol{\pi} : \mathbb{Z}_+ \to \mathcal{P}^{MIS}\}$.

IV. PROBLEM FORMULATION

In this section, we formulate the interference management policy design problem.

A. The Interference Management Policy Design Problem

We aim to optimize a chosen network performance criterion $W(R_1(\boldsymbol{\pi}), ..., R_N(\boldsymbol{\pi}))$, defined as a function of the UEs' average throughput. We can choose any performance criterion that is concave in $R_1(\boldsymbol{\pi}), ..., R_N(\boldsymbol{\pi})$. For instance, W can be the weighted sum of all the UEs' throughput $\sum_{i=1}^{N} w_i R_i(\boldsymbol{\pi})$ with $\sum_{i=1}^{N} w_i = 1$ and $w_i \geq 0$. Alternatively, the network performance can be max-min fairness (i.e., the worst UE's throughput) $\min_i R_i(\boldsymbol{\pi})$. The policy design problem (PDP) can be then formalized as follows

---

Policy Design Problem (PDP)

$$\max_{\boldsymbol{\pi}} \quad W(R_1(\boldsymbol{\pi}), ..., R_N(\boldsymbol{\pi}))$$
$$\text{subject to} \quad R_i(\boldsymbol{\pi}) \geq R_i^{min}, \, \forall i \in \{1, ..., N\}$$

The above design problem is very challenging to solve even in a centralized manner (it is NP-hard [22] when W is the sum throughput). Denote the optimal value of the PDP as $\mathbf{W}_{opt}$. Our goal is to develop distributed, fast algorithms to construct policies that achieve a constant competitive ratio with respect to $\mathbf{W}_{opt}$, with the competitive ratio independent of the network size. We achieve our goal by focusing on policies based on MISs Π^{MIS}, among other innovations that will be described in Section V.

V. DESIGN FRAMEWORK FOR DISTRIBUTED INTERFERENCE MANAGEMENT

A. Proposed Design Framework

Our proposed design framework (see Fig. 2) consists of the following four steps.

Step 1. Identification of the interfering neighbors: In Step 1, each UE-SBS pair identifies the UE-SBS pairs that strongly interfere with it. Essentially, each pair obtains a *local* view (i.e., its neighbors) of the interference graph. Note that an edge exists between two pairs if at least one of them identifies the other as a strong interferer.

Specifically, each UE-SBS pair is first informed of other pairs in the geographical proximity by managing servers (e.g., femtocell controllers/gateways) [23] [24] [19] [20]. Then each pair can decide whether another pair is strongly interfering based on various rules, such as rules based on Received Signal Strength (RSS) in the *Physical Interference Model* [23] [19] [20], and rules based on the locations in the *Protocol Model* [18]. If one pair identifies another pair as strongly interfering, its decision can be relayed by the managing servers to the latter, such that any two pairs can reach consensus of whether there exists an edge between them.

Step 2. Distributed generation of MISs that span all the UEs: In Step 2, the UE-SBS pairs generate a subset of MISs in a distributed fashion. It is important that the generated subset spans all the UEs, namely every UE is contained in at least one MIS in the subset. Otherwise, some UEs will never be scheduled.

The key idea is that from a given list of colors, each UE has to choose a set of colors such that the choice does not conflict with its neighbors. We should ensure that each UE has at least

---

[9]Strict synchronization is required for inter-cell interference coordination (ICIC) in Release 10 of 3GPP [17] and is widely assumed in the literature as well [2] [4]–[6] [15] [16].

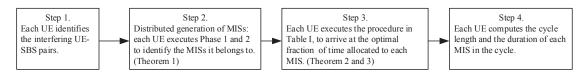| Step 1.
Each UE identifies the interfering UE-SBS pairs. | Step 2.
Distributed generation of MISs: each UE executes Phase 1 and 2 to identify the MISs it belongs to. (Theorem 1) | Step 3.
Each UE executes the procedure in Table I, to arrive at the optimal fraction of time allocated to each MIS. (Theorem 2 and 3) | Step 4.
Each UE computes the cycle length and the duration of each MIS in the cycle. |

Figure 2. Steps in the Design Framework.

one color. We call the set of UEs with the same color "a color class". In addition, we should also ensure that every color class is a MIS. This step is composed of two phases: first, distributed coloring of the interference graph based on [25], and second, extension of color classes to MISs. All the UEs are synchronized and carry out their computation simultaneously. We now explain the algorithm in detail. The pseudo-codes can be found in Table II and III in the Appendix.

Phase 1. Distributed coloring of the interference graph: Let H[10] be the maximum number of colors given to all SBSs at the installation and d_i be the degree (number of neighbors in the interference graph) of the i^{th} pair. The goal of this phase is to let each UE-SBS pair i choose *one* color from $C_i^0 \triangleq \{1, ...H\} \cap \{1, .., d_i+1\}$, such that no neighbors choose the same color. The distributed coloring works as follows.

i) At the beginning of each time slot t, each UE i chooses a color from the set of remaining colors C_i^t uniformly randomly, and informs its neighbors of its tentative choice. This information can be transmitted through the back-haul network/X2 interface that is used for ICIC [24].

ii) If the tentative choice of a UE does not conflict with any of its neighbor, then it fixes its color choice and informs the neighbors of its choice. This UE does not contend for colors any further in Phase 1. The neighbors delete the color chosen by i from their lists $C_j^{t+1}, \forall j \in \mathcal{N}(i)$, where $\mathcal{N}(i)$ is the set of i's neighbors.

iii) Otherwise, if there is a conflict, then the UE does not choose that color and repeats i) and ii) in the next time slot.

There are $\lceil c_1 \log_{\frac{4}{3}} N \rceil + 1$ time slots in Phase 1, where c_1 is the parameter given by the protocol. The number of time slots is known to the SBSs at installation. Phase 1 is successful if all the UEs acquire a color, which implies that the set of color classes (i.e., the set of UE-SBS pairs with the same color) spans all the UEs.

Phase 2. Extending color classes to the MISs: Each color class obtained at the end of Phase 1 is an independent set (IS) of the graph. In Phase 2, we extend each of these ISs to MISs and possibly generate additional MISs. After Phase 1, each UE has chosen one color and deleted some colors from its list. But there may still be remaining colors in its list that are not acquired by any of its neighbors. If the UEs can acquire these remaining colors without conflicting with its neighbors, then each color class will be a MIS. Phase 2 works as follows.

i) At each time slot in Phase 2, UE i chooses each color from

[10]The maximum number of colors H should be set to be larger than the maximum number of UE-SBS pairs interfering with any UE-SBS pair. The SBSs can determine H according to the deployment scenario. H in general will also include the number of UEs that use the same SBS who interfere with each other along with the other neighboring UEs. For example, H can be 10-15 in an office building with dense deployment of SBSs, and can be 3-5 in a residential area.

the remaining colors in its list independently with probability c. Each UE i then sends the set of its tentative choices to its neighboring UEs, and receives their neighbors' choices.

ii) For any tentative choice of color, if there is a conflict with at least one neighbor, then that color is not fixed; otherwise, it is fixed.

iii) At the end of each time slot, each UE deletes its set of fixed colors from its list, and transmits this set of fixed colors to its neighbors, who will delete these fixed colors from their lists as well. Note that a UE deletes a particular color if and only if the UE itself or some of its neighbors have chosen this color. Based on this key observation, we can see that if a color is not in any UE's list, the set of UEs with this color is a MIS. If all the UEs have an empty list, then for any color in the set $\{1, ..., H\}$, the set of UEs with this color is a MIS.

There are $\lceil c_2 \log_x N \rceil + 1$ time slots in Phase 2, where $x = \frac{1}{1-(c)^H(1-c)^{H^2}}$, and c_2 is the parameter given by the protocol. The number of time slots is known to the SBSs at installation. We say that Phase 2 is successful, if it finds H MISs, or equivalently if all the UEs have an empty list.

Example: We illustrate Step 2 in a network of 4 UE-SBS pairs, whose interference graph is shown in Fig. 3. At the start, each UE-SBS pair has a list of 3 colors {Red, Yellow, Green}. Phase 1 is run for $P1 = \lceil c_1 \log_{\frac{4}{3}} 5 \rceil$ time slots. At the end of Phase 1, UE 1 and UE 2 acquire Green and Yellow respectively, while UEs 3-4 acquire Red. Hence, UE 1 (UE 2) has an empty list, as Green (Yellow) is acquired by itself and Red, Yellow (Green) by its neighbors. UE 3 (UE 4) has Green (Yellow) color in its list of remaining colors. At the end of Phase 1, the Red color class is a MIS, while the Yellow and Green color classes are not. Phase 2 is run for $P2 = \lceil c_2 \log_x 5 \rceil + 1$ time slots. UE 3 (UE 4) acquires the remaining color Green (Yellow). At the end of Phase 2, the Green and Yellow color classes become MISs too.

The next theorem establishes the high success probability of Step 2.

Theorem 1. For any interference graph with the maximum degree $\Delta \leq H - 1$, the proposed algorithm in Table II and III outputs a set of H MISs that span all the UEs in $(\lceil c_1 \log_{\frac{4}{3}} N \rceil + \lceil c_2 \log_x N \rceil + 2)$ time slots with a probability no smaller than $(1 - \frac{1}{N^{c_1-1}})(1 - \frac{1}{N^{c_2-1}})$, where c_1 and c_2 are design parameters that trade-off the run time and the success probability.

See the Appendix for proof sketches of all our results, and see the Appendix of the online report [7] for detailed proofs.

Theorem 1 characterizes the performance of our proposed algorithm, in terms of the run time of the algorithm and the lower bound of the success probability. When the parameters c_1 and c_2 are larger, the lower bound of the success probability increases at the expense of a longer run time. When the
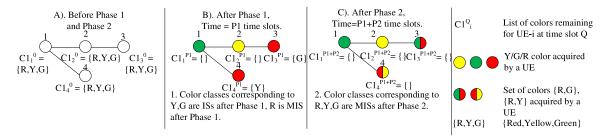
Figure 3. Illustration of the distributed generation of MISs in Step 2.

maximum degree of the interference graph is larger, we need to set a higher H, which results in a longer run time. This is reasonable, because it is harder to find coloring and MISs when the number of interfering neighbors is higher. Finally, we can see that the lower bound of the successful probability is very high even under smaller c_1 and c_2, especially if the number of UEs is large. Note that the exact successful probability should depend on the probability c in Phase 2, while the lower bound in Theorem 1 does not. Hence, our lower bound is robust to different system parameters. Note also that the interference graph here is a bounded-degree graph since the maximum degree is bounded by a given constant, $H - 1$. The algorithms in [4] [6] (require ordering of the vertices, work sequentially and have a higher complexity) can be used to output the MISs spanning all the UEs for arbitrary graphs. However, we will show in Theorem 5, that the restriction to bounded-degree graphs is a must to ensure that the minimum throughput requirement of each UE is satisfied for any MIS based policy.

Step 3. Distributed computation of the optimal fractions of time for each MIS: Let the set of MISs generated in Step 2 be $\{I'_1, ..., I'_H\}$. In Step 3, the UE-SBS pairs compute the fractions of time allocated to each MIS in a distributed manner.

When an MIS is scheduled, the UEs in this MIS transmit at their maximum power levels, and the other UEs do not transmit. Define R_i^k as the *instantaneous* throughput obtained by UE i in the MIS I'_k, which can be calculated as $\log_2(1 + \frac{g_{iT(i)}p_i^{I'_k}}{\sum_{r=1,r\neq i}^N g_{rT(i)}p_r^{I'_k}+\sigma^2_{T(i)}})$, where $p_i^{I'_k} = p_i^{max}$ if $i \in I'_k$ and $p_i^{I'_k} = 0$ otherwise. To determine R_i^k, the UE needs to know the total interference it experiences when transmitting in I'_k. This can be measured by having an initial cycle of transmissions of UEs in each MIS in the order of the indices of MISs/colors.

From now on, we assume that the network performance criterion $W(\mathbf{y})$ is concave in $\mathbf{y}$ and is separable, namely $W(y_1, ...y_N) = \sum_{i=1}^N W_i(y_i)$. Examples of separable criteria include weighted sum throughput and proportional fairness. Our framework can also deal with max-min fairness, although it is not separable (see the discussion in the Appendix of the online report [7]). The problem of computing the optimal fractions of time for the MISs is given as follows:

Coupled Problem (CP)

$$\max_\alpha \sum_{i=1}^N W_i\left(\sum_{k=1}^H \alpha^k R_i^k\right)$$

$$\text{subject to } \sum_{k=1}^H \alpha^k R_i^k \geq R_i^{min}, \forall i \in \{1,..,N\}$$

$$\sum_{k=1}^H \alpha^k = 1, \alpha^k \geq 0, \forall k \in \{1,..,H\}$$

Each UE i knows only its own utility function W_i and minimum throughput requirement R_i^{min}. Hence, it cannot solve the above problem by itself. We will first reformulate the above problem into a decoupled problem and then show that the reformulated problem can be solved in a distributed manner. Let each UE i have a local estimate β_i^k of the fractions of time allocated to each MIS I'_k (including those MISs that UE i does not belong to). We impose an additional constraint that all the UEs' local estimates are the same. Note that this constraint will be satisfied by our solution, and is not an assumption. Such a constraint is still global, because any two UEs, even if they are not neighbors, need to have the same local estimate. Hence, global message exchange among any pair of UEs is still needed to solve this problem with local estimates and global constraints[11]. To avoid global message exchange, we reformulate the CP into a decoupled problem (DP) that involves only local coupling among the neighbors and can be solved with local message exchange using Alternating Direction Method of Multipliers (ADMM) [27].

Write $\boldsymbol{\beta}_i = (\beta_i^1, ..., \beta_i^H)$ as UE i's local estimates of the fractions of time allocated to each MIS, and $\boldsymbol{R}_i = (R_i^1, ..., R_i^H)$ as UE i's throughput when each MIS is scheduled. Each UE i's local estimates should be in the polyhedron $\mathcal{B}_i \triangleq \{\boldsymbol{\beta}_i : \mathbf{1}^T\boldsymbol{\beta}_i = 1, \boldsymbol{\beta}_i \geq \mathbf{0}, \boldsymbol{\beta}_i^T \boldsymbol{R}_i \geq R_i^{min}\}$, where $()^T$ is the transpose. Let E be the set of edges, where each edge $e = \{i, j\}$ is an ordered set of the vertices, $i < j$ that are directly connected. As we will prove in Theorem 2, in a connected interference graph[12], the requirement that all UEs' local estimates are the same can be reduced to the requirement that every UE has the same local estimate as its neighbors, namely $\boldsymbol{\beta}_i = \boldsymbol{\beta}_j$ for i, j s.t. $\{i, j\} = e$, where $e \in E$. To make

[11]If the UEs could exchange messages globally, i.e. broadcast messages to all the UEs in the network, and if the network performance criterion is strictly concave, we could use standard dual decomposition with augmented Lagrangian in [26] to derive a distributed algorithm. However, in large networks, the UEs cannot exchange messages globally with other UEs, and the network performance criterion may not be strictly concave (e.g., the weighted sum throughput is linear).

[12]A graph is connected, if any two nodes are connected by a path of edges.

the problem solvable by ADMM, we rewrite the constraints by introducing auxiliary variables θ_{ei}^k, where $i \in e$ is one endpoint of the edge. Then the constraint for each edge $e = \{i, j\}$ can be rewritten as $\beta_i^k = \theta_{ei}^k$, $-\beta_j^k = \theta_{ej}^k$, $\theta_{ei}^k + \theta_{ej}^k = 0$. Hence, the auxiliary variable θ_{ei}^k can be interpreted as i's estimate of its neighbor j's estimate β_j^k. For $e = \{i, j\}$ define set of the auxiliary variables $\Theta_e^k = \{(\theta_{ei}, \theta_{ej}) \in \mathbb{R}^2 : \theta_{ei}^k + \theta_{ej}^k = 0, -1 \leq \theta_{ei} \leq 1, -1 \leq \theta_{ej} \leq 1\}$ and let $\Theta^k = \Pi_{e \in E} \Theta_e^k$. Also for each edge $e = \{i, j\}$ and for each $k \in \{1, .., H\}$ define $D_{ei}^k = 1$ and $D_{ej}^k = -1$. Then the decoupled problem is given as follows:

Decoupled Problem (DP)

$$\min_{\{\boldsymbol{\beta}_i \in \mathcal{B}_i\}_{i=1}^N, \{\boldsymbol{\theta}^k \in \Theta^k\}_{k=1}^H} - \sum_{i=1}^N W_i \left(\boldsymbol{\beta}_i^T \boldsymbol{R}_i \right)$$
$$\text{subject to } D_{eq}^k \beta_q^k = \theta_{eq}^k, \, \forall q \in e, \forall e \in E, \, \forall k \in \{1, .., H\}$$

Theorem 2: For any connected interference graph, the coupled problem (CP) is equivalent to the decoupled problem (DP).

The above theorem shows that the original problem (CP), which requires global information and global message exchange to solve, is transformed into an equivalent problem (DP), which as we will show, can be solved in a distributed manner with local message exchange

We denote the optimal solution to the DP by $\mathbf{W}_{\text{distributed}}^G$. We associate with each constraint $D_{eq}^k \beta_q^k = \theta_{eq}^k$ a dual variable λ_{eq}^k. The augmented Lagrangian for DP is $L_y \left(\{\boldsymbol{\beta}_i\}_i, \{\theta_{eq}^k\}_{k,e,q}, \{\lambda_{eq}^k\}_{k,e,q} \right) = -\sum_{i=1}^N W_i(\boldsymbol{\beta}_i^T \boldsymbol{R}_i) + \sum_{k=1}^H \sum_{e \in E} \sum_{q \in e} \left[\lambda_{eq}^k \left(D_{eq}^k \beta_q^k - \theta_{eq}^k \right) + \frac{y}{2} \left(D_{eq}^k \beta_q^k - \theta_{eq}^k \right)^2 \right]$. In the ADMM procedure (see Table IV in the Appendix), each UE i solves for its optimal local estimates $\boldsymbol{\beta}_i(t)$ that maximizes the augmented Lagrangian given the previous dual variables $\lambda_{ei}^k(t-1)$ and auxiliary variables $\theta_{ei}^k(t-1)$. Then it updates its dual variable $\lambda_{ei}^k(t)$ and auxiliary variable $\theta_{ei}^k(t)$ based on its local estimate $\beta_i^k(t)$ and its neighbor j's local estimate $\beta_j^k(t)$. This iteration of updating local estimates, dual variables, and auxiliary variables is repeated P times. Next, it is shown that this procedure will indeed converge.

Theorem 3: If DP is feasible[13], then the ADMM algorithm in Table IV converges to the optimal value $\mathbf{W}_{\text{distributed}}^G$ with a rate of convergence $\mathcal{O}(\frac{1}{P})$.

Step 4. Determining the cycle length and transmission times: At the end of Step 3, all the UEs have a consensus about the optimal fractions of time allocated to each MIS, namely $\boldsymbol{\beta}_i^* = \boldsymbol{\gamma}^* = (\gamma_1^*, ..., \gamma_H^*)$, $\forall i \in \{1, .., N\}$. The MISs transmit in the order of their indices (i.e., $\{1, .., H\}$) in cycles. In each cycle of transmission, MIS $I_k^{'}$ transmits for $\left\lceil \frac{\gamma_k^*}{\min_{i \in 1, ..., N} \gamma_i^*} \times 10^d \right\rceil$ slots, where we multiply by 10^d such that the rounding error is reduced or eliminated in case that $\frac{\gamma_k^*}{\min_{i \in 1, ..., N} \gamma_i^*}$ is not an integer.

B. A Motivation Example

Consider a network of 2 picocell base stations (PBS) and 2 femtocell base stations (FBS), each serving one UE. The network topology is shown in Fig. 4. We assume a path loss

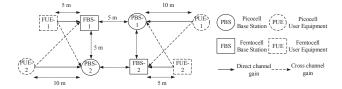[13]DP is feasible, if the feasible region resulting from the constraints in DP is non-empty.

Figure 4. A heterogeneous network of 2 PBS and 2 FBS and their corresponding UEs.

model for channel gains, with path loss exponent 4. The maximum transmit power of each UE is 80 mW, and the noise power at each SBS is 1.6×10^{-3} mW. UEs in different tiers have different minimum throughput requirements: FUE (femtocell UE) 1 and FUE 2 in the femtocells require a minimum throughput 0.4 bits/s/Hz, and PUE (picocell UE) 1 and PUE 2 in the picocells require 0.2 bits/s/Hz. The interference graph is constructed according to a distance based threshold rule similar to [18]. Specifically, an edge exists between two UE-BS pairs if the distance between any pair of SBSs is less than a threshold, which is set to be 1.2m here. There are two MISs. MIS 1 consists of FUE 1 and FUE 2, and MIS 2 consists of PUE 1 and PUE 2. We consider two performance criteria: the max-min fairness and the sum throughput. We will compare with the following state-or-the-art policies:

1. Distributed Constant Power Control Policies [8]–[10]**:** In these policies, all the UEs choose *constant* power levels determined by distributed algorithms utilizing information (e.g., power levels used by neighbors) made available through local/global message exchange.

2. Optimal Centralized Constant Power Policies: In these policies, all the UEs choose *constant* power levels determined by a central controller utilizing global information.

3. Distributed MIS STDMA-1 [6] and **STDMA**-2 [4]**:** These policies construct a subset of the MISs of the interference graph in a distributed manner and propose fixed schedules of the MISs. Different works adopt different schedules, and we differentiate them by referring to them as MIS STDMA-1 [6] and STDMA-2 [4].

4. Distributed Joint Power Control and Spatial Reuse [15] [16]**:** These policies choose one UE from each cell to form a subset, and schedule these subsets of UEs based on their channel gains to maximize the sum throughput. The policies are named power matched scheduling (PMS).

In Table 1, we compare the performance of our proposed policy with state-of-the-art policies for the same setup as in Fig. 4. We compute the optimal centralized constant power control policy by exhaustive search, which serves as the performance upper bound of the distributed constant power control policies [8]–[10] centralized constant power control policies [12]. In PMS policies [15] [16], UEs within the same cell are scheduled in a time-division multiple access (TDMA) fashion, and the active UEs in different cells transmit simultaneously. In this motivating example, there is one UE in each cell, which will be scheduled to transmit all the time. Therefore, the PMS policy reduces to a constant power control policy, and is worse than the optimal centralized

constant power control policy. We can see that our proposed policy outperforms all constant power control policies and distributed PMS policies by at least 375% and 32.8%, in terms of max-min fairness and sum throughput, respectively. The significant performance improvement over the constant power control policies results from the elimination of the high interference among the users through scheduling MISs. Our proposed policy also outperforms distributed STDMA policies by 30%-40%. As we will see in Section VI, the performance gain is even higher (160%-700%) in realistic deployment scenarios. Finally, in this motivating example, the proposed policy achieves the optimal performance of the benchmark problem defined in Section VI, which is a close approximation of the original problem (CP).

C. Performance Guarantees for Large Networks and Properties of Interference Graphs

In this subsection, we provide performance guarantees for our proposed framework described in Section V-A. Specifically, we prove that the network performance $\mathbf{W}^G_{\text{distributed}}$ achieved by the proposed distributed algorithms has a constant competitive ratio with respect to the optimal value $\mathbf{W}_{\text{opt}}$ of the PDP. Moreover, we prove that the competitive ratio does not depend on the network size. Our result is strong, because the solution to PDP needs to be computed by a centralized controller with global information and with NP complexity, while our proposed framework allows the UEs to compute the policy fast in a distributed manner with local information and local message exchange.

Before characterizing the competitive ratio analytically, we define some auxiliary variables. Define the upper and lower bounds on the UEs' maximum transmit power levels and throughput requirements as, $0 < p^{max}_{lb} \leq p^{max}_i \leq p^{max}_{ub}, \forall i \in \{1,...,N\}$ and, $0 < R^{min}_{lb} \leq R^{min}_i \leq R^{min}_{ub}, \forall i \in \{1,...,N\}$ respectively. Let D_{ij} is the distance between UE i and SBS j. Define upper and lower bounds on the distance between any UE and its serving SBS and the noise power at the SBSs as, $0 < D^{lb} \leq D_{iT(i)} \leq D^{ub}, \forall i \in \{1,...,N\}$ and, $\sigma^2_{lb} \leq \sigma^2_j \leq \sigma^2_{ub}, \forall j \in \{1,...,K\}$ respectively. We assume that the channel gain is $g_{ij} = \frac{1}{(D_{ij})^{np}}$, where np is the path loss exponent.

Definition 1 (Weak Non-neighboring Interference): The interference graph G exhibits ζ Weak Non-neighboring Interference (ζ-WNI) if for each UE i the maximum interference from its non-neighbors is bounded, namely $\sum_{j\notin\mathcal{N}(i),j\neq i} g_{jT(i)}p^{max}_j \leq (2^\zeta-1)\sigma^2_{ub}, \forall i \in \{1,...,N\}$.

Define $\Delta^{max} = \frac{\log_2(1+\frac{p^{max}_{lb}}{(D^{ub})^{np}2^\zeta\sigma^2_{ub}})}{R^{min}_{ub}} - 1$. Then we have the following theorem for the network performance criterion, sum throughput[14].

Theorem 4: For any connected interference graph, if the maximum degree $\Delta \leq \Delta^{max}$ and it exhibits ζ-WNI then, our proposed framework of interference management described in Section V-A achieves a performance $\mathbf{W}^G_{\text{distributed}} \geq \Gamma \cdot \mathbf{W}_{\text{opt}}$ with a probability no smaller than $(1 - \frac{1}{N^{c_1-1}})(1 - \frac{1}{N^{c_2-1}})$.

---

[14]We can extend this result for weighted sum throughput, with weights $w_i = \Theta(\frac{1}{N})$, it is not done to avoid complex notations.

Moreover, the competitive ratio $\Gamma = \frac{R^{min}_{ub}}{\log_2(1+\frac{p^{max}_{ub}}{(D^{lb})^{np}\sigma^2_{lb}})}$ is independent of the network size.

Note that the analytical expression of competitive ratio, $\Gamma = \frac{R^{min}_{ub}}{\log_2(1+\frac{p^{max}_{ub}}{(D^{lb})^{np}\sigma^2_{lb}})}$, does not depend on the size of the network. Our results are derived under the conditions that the interference graph has a maximum degree bounded by Δ^{max}, and that the interference from non-neighbors is bounded (i.e. ζ−WNI). These conditions do not restrict the size of the network (for more detail, see the Section V-C of the online report [7]). In addition, our results hold for any interference graph that satisfy the conditions in Theorem 4, regardless of how the graph is constructed.

Both Theorem 1 and 4 restricted the interference graph to be bounded-degree. We justify our restriction by showing that the bounded-degree property is necessary to fulfill the minimum throughput requirements of each UE. Specifically, we prove that if the maximum degree exceeds some threshold, then no MIS based policy in Π^{MIS} (which is a large policy space) is feasible. Suppose that the interference graph is constructed based on a distance based threshold rule similar to [18]: an edge exists between two UE-SBS pairs if and only if the distance between two SBSs is no greater than D^{th}. We define the threshold of the maximum degree as Δ^* (See the Appendix for the expression).

Theorem 5: If the maximum degree of the interference graph $\Delta \geq \Delta^*$, then any MIS based policy in Π^{MIS} fails to satisfy the minimum throughput requirements of the UEs.

The intuition behind Theorem 5 is that, if the degree of the interference graph is large then there must be a large number of UE-SBS pairs which interfere with each other strongly, which makes it impossible to allocate each UE enough transmission time to satisfy their minimum throughput requirements simultaneously.

D. Self-Adjusting Mechanism for Dynamic Entry/Exit of UEs

We now describe how the proposed framework can adjust to dynamic entry/exit by the UEs in the network without recomputing all the four steps. We allow the UEs to enter and exit, but number of SBSs is fixed. We only let one UE enter or leave the network in any time slot.

1. *UE leaves the network:* Suppose UE i, which was transmitting to SBS $T(i)$, leaves the network. If the UE i was transmitting in a set of colors C_i, then as soon as it leaves, these colors can be potentially used by some neighbors, $\mathcal{N}(i)$. SBS $T(i)$ can still be serving other UEs which are still in the network and transmitting. Then for each color $c' \in C_i$ it first searches among these UEs that are not already transmitting in c' and who also do not have a neighboring UE-SBS pair which is already transmitting in c'. Let the set of such UEs be $UE^{c'}_{i,left}$. SBS $T(i)$ allocates color c' to the UE whose index is $\arg\max_{j\in UE^{c'}_{i,left}} R^{c'}_j$. In case $UE^{c'}_{i,left}$ is empty then that color, c' is left unused.

2. *UE enters the network:* Suppose a UE i registered with SBS $T(i)$ enters the network. SBS $T(i)$ creates the list of colors $C^{valid}_{i,enter}$, which are either unused or the UEs transmitting

Table I
COMPARISONS IN TERMS OF MAX-MIN FAIRNESS & SUM THROUGHPUT CRITERION

Policies	Max-min throughput (bits/s/Hz)	Performance Gain %	Sum throughput (bits/s/Hz)	Performance Gain %
Distributed constant power control [8]–[10]	<0.28	>375 %	6.1	32.8 %
Distributed PMS [15], [16]	<0.28	>375%	6.1	32.8 %
Optimal centralized constant power control	0.28	375%	6.1	32.8 %
Distributed MIS STDMA-2/1 [4], [6]	0.96	38.5%	6.25	30.0 %
Proposed (Section-V)	1.33	-	8.12	-
Benchmark Problem (BP) (Section- VI)	1.33	-	8.12	-

in the colors are transmitting at more than their minimum throughput requirement. SBS $T(i)$ allocates some portions from the fractions of time allocated to the colors in $C_{i,enter}^{valid}$, to satisfy UE-i's throughput requirement to the best possible extent, making sure that the minimum throughput requirements of UEs transmitting to SBS $T(i)$ in $C_{i,enter}^{valid}$ are not violated. If the requirement of UE-i is not satisfied then, SBS $T(i)$ requests the neighboring UE-SBSs to announce the set of colors, which are either not being used or in which the UEs being served are operating at more than their throughput requirement. From the list of colors received, $T(i)$ chooses those in which UE i can transmit without conflicting with neighbors. For each of these colors it sends the request (portion of time needed) to the neighbors. SBS $T(i)$ and the neighbors go through a phase of communication (more detail in the Section V-D of online report [7]), based on which SBS $T(i)$ can decide how much time UE-i can transmit in these colors.

E. Extensions

In our model, UEs operate in the same frequency band. However, our methodology can be extended to scenarios where UEs operate in different frequency channels (frequency reuse) and transmit at the same time. In this case, the problem is to find the optimal frequency allocation with the same objective function and constraints as in PDP. To solve this problem, the first two steps of the framework remain the same. In Step 3, the UEs compute distributedly the optimal fractions of *bandwidth* to be allocated to each MIS. This step is equivalent to computing the optimal fraction of time allocated to each MIS as in our current formulation. In Step 4, the UEs compute the number of frequency channels allocated to each MIS based on the bandwidth allocation.

Note that we do not implement beamforming, although beamforming can be used in conjunction with our policy. If the UEs transmitting to the same SBS cooperate to do beamforming, we can delete the edge between them in the interference graph, and use the new interference graph in the scenario with beamforming.

VI. ILLUSTRATIVE RESULTS

In this section, we evaluate our proposed policy under a variety of scenarios with different levels of interference, large numbers of UEs, different performance criteria, time-varying channel conditions, and dynamic entry and exit of UEs.

We compare our policy with the optimal centralized constant power control policy, the distributed MIS STDMA-1 [6] and STDMA-2 [4], distributed PMS [15] [16], in
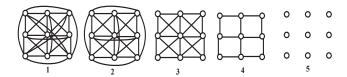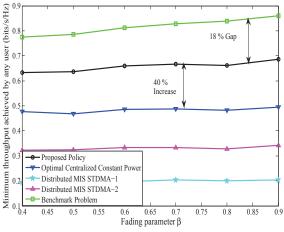

Figure 5. Different interference graphs for the 3 x 3 BS grid

terms of sum throughput and max-min fairness. We do not separately compare with distributed/centralized constant power control policies in [8]–[10] [12], because their performance is upper bounded by the optimal centralized power control. Since it is difficult to compute the solution to the NP-hard PDP, we define a benchmark problem, where we restrict our search to policies in which a UE either transmits at its maximum power level or does not transmit. The space of such policies can be written as $\Pi_{BC} = \{\boldsymbol{\pi} = (\pi_1, ..., \pi_N) : \pi_i : \mathbb{Z}_+ \to \{0, p_i^{max}\} \forall i \in \{1, .., N\}\}$. The policy space Π_{BC} is a subset of all policies Π and is a superset of MIS based policies Π^{MIS}. In other words, the benchmark problem has the same objective and constraints as PDP; the only difference is the policy space to search . Hence, the benchmark problem is a close approximation of the PDP. Note that the benchmark problem is also NP-hard (see the Appendix of the online report [7]).

A. Performance under time-varying channel conditions

Consider a 3x3 square grid of 9 SBSs with the minimum distance between any two SBSs being $d = 4.7$m. Each SBS serves one UE, who has a maximum power of 1000 mW and a minimum throughput requirement of 0.45 bits/s/Hz. The UEs and the SBSs are in two parallel horizontal hyperplanes, and each SBS is vertically above its UE with a distance of $\sqrt{10}m$. Then the distance from UE i to another SBS j is $D_{ij} = \sqrt{10 + (D_{ij}^{BS})^2}$, where D_{ij}^{BS} is the distance between SBSs i and j. The channel gain from UE i to SBS j is a product of path loss and Rayleigh fading $f_{ij} \sim Rayleigh(\beta)$, namely $g_{ij} = \frac{1}{(D_{ij})^2} f_{ij}$. The density function of $Rayleigh(\beta)$ is $v(z) = \frac{z}{\beta^2} e^{-\frac{z^2}{2\beta^2}}$ for $z \geq 0$, and $v(z) = 0$ for $z < 0$. The SBSs identify neighbors using a distance based rule with the threshold distance as in Section V-C with $D^{th} = 7$m. Note that different thresholds lead to different interference graphs, and hence different performance, which will be discussed next. Although, we use a distance based threshold rule, our framework is general and does not rely on a particular rule.
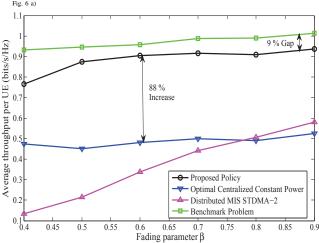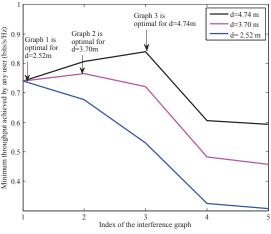
Fig. 6 a)

Fig. 6 b)

Figure 6. Comparison of the proposed policy with state of the art under different interference strength and time-varying channel conditions

The resulting interference graph for this setting is graph 3 shown in Fig. 5.

At the beginning, the UE-SBS pairs generate the set of MISs (Step 2 of the design framework in Section V), and compute the optimal fractions of time allocated to each MIS (Step 3). In our simulation, we assume a block fading model [28] and the fading changes every 100 time slots independently. To reduce complexity, the UEs do not recompute the interference graph and the MISs, but will recompute the optimal fractions of time under the new channel gains every 100 time slots. In Fig. 6, we compare the performance of the proposed policy with state of the art policies under different variances β of Rayleigh fading. We do not plot the performance of distributed PMS for this scenario since it is upper bounded by optimal centralized constant power control (because there is one UE per cell). We do not plot the distributed MIS STDMA -1 either, when the performance criterion is average throughput per UE (i.e., $\frac{\text{sum throughput}}{N}$), because it cannot satisfy the minimum throughput constraints. From Fig. 6, we can see that in terms of both average throughput and max-min fairness, our proposed policy achieves large performance gain (up to 88%) over existing policies, and achieves performance close to the benchmark (as close as 9%).
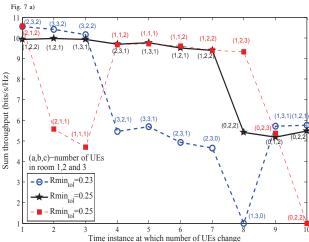

Fig. 7 a)

Fig. 7 b)

Figure 7. a) Comparison of max-min fairness under different grid sizes, b) Sample paths of sum throughput under dynamic entry/exit of UEs in the network
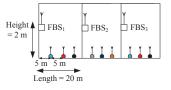

Figure 8. Illustration of setup with 3 rooms.

Selecting the Optimal Interference Graph : For different values of d, there can be five possible interference graphs, which are shown in Fig. 5. In Fig. 7 a) we show that as the grid size d decreases ($d = 4.7m, d = 3.7m$ and $d = 2.5m$), the levels of interference from the adjacent UEs increases, and as a result, the interference graph with higher degrees perform better (as d decreases, the optimal graph changes from graph 3 to graph 1) .

B. Performance scaling in large networks

Consider the uplink of a femtocell network in a building with 12 rooms adjacent to each other. Fig. 8 illustrates 3 of the 12 rooms with 5 UEs in each room. For simplicity, we consider a 2-dimensional geometry. Each room has a length
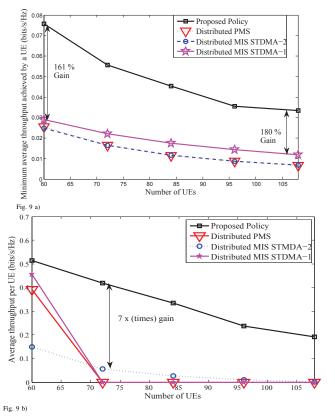
Fig. 9 a)

Fig. 9 b)

Figure 9. Comparison of max-min fairness and average throughput per UE against state of the art for large networks

of 20 meters. In each room, there are P uniformly spaced UEs, and one SBS installed on the left wall of the room at a height of 2m. The distance from the left wall to the first UE, as well as the distance between two adjacent UEs in a room, is $\frac{20}{(1+P)}$ meters. Based on the path loss model in [29], the channel gain from each SBS i to a UE j is $\frac{1}{(D_{ij})^2 \Delta^{n_{ij}}}$, where $\Delta = 10^{0.25}$ is the coefficient representing the loss from the wall, and n_{ij} is the number of walls between UE i and SBS j. Each UE has a maximum transmit power level of 50 mW, a minimum throughput requirement of $R_i^{min} = 0.025$ bits/s/Hz, and a noise power level of 10^{-11}mW at its receiver. Here, we consider that the UEs use a distance based threshold rule as in Section V-B with $D^{th} = 30$ m. This results in interference graphs which connects all the UE-SBS pairs within the room and in the adjacent rooms. We vary the number P of UEs in each room from 5 to 9 and compare the performance in Fig. 9. Note that the optimal centralized constant power policy cannot satisfy the feasibility conditions for any number of UEs in each room. Hence, only the performance of distributed MIS STDMA-1,2 and distributed PMS is shown in Fig. 9. We can see that under both criteria, the performance gain of our proposed policy is significant (from 160% to 700%). Note that since the number of UEs is large, it is impossible to solve the benchmark problem (which is NP-hard) is not possible.

C. Self-adjusting mechanism for dynamic entry/exit of the UEs

The self-adjusting mechanism proposed in Section V-D is aimed to provide incoming UEs with the maximum possible throughput without affecting the incumbent UEs, and to reuse the time slots left vacant by exiting UEs efficiently. Consider the same setup as in Section VI-B with 3 rooms and a maximum of $P = 3$ UEs in each room. Each UE has a maximum transmit power of 1000 mW and a minimum throughput requirement of 0.25 bits/s/Hz.

We assume that at a given time only one UE either enters or leaves the network. In Fig. 7 b) we show different sample paths of the sum throughput under different entry and exit processes. In the legends (i.e., $Rmin_{tol}$), we show the minimum throughput achieved at any point in the sample path. We repeated the same procedure 100 times. We can see that the self-adjusting mechanism works well by guaranteeing a worst-case minimum throughput requirement of 0.23 bits/s/Hz, which is just 0.02 bits/s/Hz below the original requirement more than 80% of the time.

VII. CONCLUSION

We proposed a design framework for distributed interference management in large-scale, heterogeneous networks, which are composed of different types of cells (e.g. femtocell, picocell), have different number of UEs in each cell, and have UEs with different minimum throughput requirements and channel conditions. Our framework allows each UE to have only local knowledge about the network and communicate only with its interfering neighbors. There are two key steps in our framework. First, we propose a novel distributed algorithm for the UEs to generate a set of MISs that span all the UEs. The distributed algorithm for generating MISs requires $\mathcal{O}(\log N)$ steps (which is much faster than state-of-the-art) before it converges to the set of MISs with a high probability. Second, we reformulate the problem of determining the optimal fractions of time allocated to the MISs in a novel manner such that the optimal solution can be determined by a distributed algorithm based on ADMM. Importantly, we prove that under wide range of conditions, the proposed policy can achieve a constant competitive ratio with respect to the policy design problem which is NP-hard. Moreover, we show that our framework can adjust to UEs entering or leaving the network. Our simulation results show that the proposed policy can achieve large performance gains (160%-700%).

APPENDIX A
PROOF OF THEOREM 1

The analysis of the success probability of Phase 1 is similar to [25] and is omitted here. For Phase 2, we first show that if every UE i's list $C1_i^n$ of remaining colors is empty, then Phase 2 has produced H MISs that span all the UEs. Let us assume otherwise, namely all the UEs' lists are empty, but there exists some color $h \in \{1, ..., H\}$ such that $I_h^{'}$ is not a MIS. $I_h^{'}$ has to be an IS (see [7]). Since $I_h^{'}$ is not maximal, there exists at least one UE $j \notin I_h^{'}$ that is not a neighbor of any UE in $I_h^{'}$. Since UE j's list is empty and $j \notin I_h^{'}$, UE j must have deleted color h because at least one of its neighbors $k \in \mathcal{N}(j)$ acquired it. However, this means that $k \in I_h^{'}$, which is contradictory to the fact that j is not a neighbor of any UE in $I_h^{'}$. Next, we derive the success probability of Phase 2. Let U^n be the number of UEs that have a non-empty list, and let

Table II
GENERATING MISs IN A DISTRIBUTED MANNER, ALGORITHM FOR UE i

Phase 1- Initialization: $\text{Tx}_{\text{tent}}^i = \phi$, $\text{Tx}_{\text{final}}^i = \phi$, tentative and final choice of UE i, $\text{Rx}_{\text{tent}}^{\mathcal{N}(i)} = \phi$, $\text{Rx}_{\text{final}}^{\mathcal{N}(i)} = \phi$ tentative and final choice made by the neighbors, $C_i^0 = \{1, ..., H\} \cap \{1, ..., d_i + 1\}$ the current list of subset of available colors, $C_i = \phi$, list of colors used by i, $\text{F}_{\text{colored}}^i = \phi$, $C1_i^0 = \{1, ..., H\}$, the current list of all available colors

for $n = 0$ to $\lceil c_1 \log_{\frac{4}{3}} N \rceil$
$\text{Tx}_{\text{tent}}^i = \phi$, $\text{Tx}_{\text{final}}^i = \phi$
if($\text{F}_{\text{colored}}^i = \phi$)
  $\text{Tx}_{\text{tent}}^i = \text{rand}\{C_i^n\}$, rand represents randomly selecting a color and informing the neighbors about it.
  $\text{Rx}_{\text{tent}}^{\mathcal{N}(i)} = \{\text{Tx}_{\text{tent}}^k, \forall k \in \mathcal{N}(i)\}$
   If($\text{Tx}_{\text{tent}}^i \neq \text{Rx}_{\text{tent}}^{\mathcal{N}(i)}(j), \forall j \in \mathcal{N}(i)$), here UE-$i$ checks if there is a conflict with any of the neighbor's choice
    $\text{Tx}_{\text{final}}^i = \text{Tx}_{\text{tent}}^i$, $C_i = \{\text{Tx}_{\text{final}}^i\}$, if there is no conflict then UE-i transmits its final color choice to the neighbors,
   else
    $\text{Tx}_{\text{final}}^i = \phi$
   end
end
$\text{Rx}_{\text{final}}^{\mathcal{N}(i)} = \{\text{Tx}_{\text{final}}^k, \forall k \in \mathcal{N}(i)\}$
$C_i^{n+1} = C_i^n \cap \{\text{Rx}_{\text{final}}^{\mathcal{N}(i)} \cup \text{Tx}_{\text{final}}^i\}^c$, $C1_i^{n+1} = C1_i^n \cap \{\text{Rx}_{\text{final}}^{\mathcal{N}(i)} \cup \text{Tx}_{\text{final}}^i\}^c$
if($\text{Tx}_{\text{final}}^i \neq \phi$)
$\text{F}_{\text{colored}}^i = 1$
end
end

Table III
PHASE 2 OF THE DISTRIBUTED MIS GENERATION

Phase 2-Initialization: $\text{Tx}_{\text{tent},i}^{\text{set}} = \phi$, $\text{Tx}_{\text{final},i}^{\text{set}} = \phi$, the set of tentative and final colors chosen by i, $\text{Rx}_{\text{tent},i}^{\text{set}} = \phi$, $\text{Rx}_{\text{final},i}^{\text{set}} = \phi$, the set of tentative and final colors chosen that are received from the neighbors, $x = \frac{1}{1-(c)^H(1-c)^{H^2}}$

for $n = \lceil c_1 \log_{\frac{4}{3}} N \rceil + 1$ to $\lceil c_1 \log_{\frac{4}{3}} N \rceil + \lceil c_2 \log_x N \rceil + 1$
$\text{Tx}_{\text{tent},i}^{\text{set}} = \phi$, $\text{Tx}_{\text{final},i}^{\text{set}} = \phi$,
  for $m = 1$ to $|C1_i^n|$
  with probability c, $\text{Tx}_{\text{tent},i}^{\text{set}}(m) = C1_i^n(m)$, randomly selecting and informing the neighbors about tentative choice
  with probability $1 - c$, $\text{Tx}_{\text{tent},i}^{\text{set}}(m) = \phi$
  end
$\text{Rx}_{\text{tent},i}^{\text{set}} = \cup_{k \in \mathcal{N}(i)} \text{Tx}_{\text{tent},k}^{\text{set}}$, set of tentative color choices of the neighbors of i
  for $r = 1$ to $|\text{Tx}_{\text{tent},i}^{\text{set}}|$
  If($\text{Tx}_{\text{tent},i}^{\text{set}}(r) \neq \text{Rx}_{\text{tent},i}^{\text{set}}(j) \,\forall j \in \mathcal{N}(i)$)
  $\text{Tx}_{\text{final},i}^{\text{set}}(r) = \text{Tx}_{\text{tent},i}^{\text{set}}(r)$
  else
  $\text{Tx}_{\text{final},i}^{\text{set}}(r) = \phi$
  end
$C_i = C_i \cup \text{Tx}_{\text{final},i}^{\text{set}}$
$\text{Rx}_{\text{final},i}^{\text{set}} = \cup_{k \in \mathcal{N}(i)} \text{Tx}_{\text{final},k}^{\text{set}}$, set of final color choices of the neighbors of i
$C1_i^{n+1} = C1_i^n \cap \{\text{Rx}_{\text{final},i}^{\text{set}} \cup \text{Tx}_{\text{final},i}^{\text{set}}\}^c$
end

$Tn(U^n)$ be the number of time slots needed before all the UEs have an empty list. Consider an arbitrary UE i with a non-empty list at time n. The probability that it will have an empty list at time $n + 1$ is no smaller than the probability that it picks all the remaining colors and none of its neighbors picks any of them, which is $c^{|C1_i^n|} \prod_{j \in \mathcal{N}(i)}(1-c)^{|C1_j^n|}$. Since $|C1_i^n| \leq H$ for all i and $\mathcal{N}(i) \leq \Delta < H$, the above probability is lower bounded by $c^H \left[(1-c)^H\right]^H = c^H(1-c)^{H^2}$. From this, we get $E(U^{n+1}) \leq \left(1 - c^H(1-c)^{H^2}\right) U^n \triangleq \frac{1}{x} U^n$ and $Tn(U^n) = 1 + Tn(U^{n+1})$. Assume the worst case where Phase 2 starts with all N UEs having non-empty lists. From [30], we get $P(Tn(N) \geq \lceil c_2 \log_x N \rceil) \leq \frac{1}{N^{c_2 - 1}}$. This gives the lower bound on the success probability of Phase 2. ■

APPENDIX B
PROOF OF THEOREM 2

We first prove that the coupled problem (CP) is equivalent to the following problem:

Global Primal Problem (GPP)
$$\max_{\{\beta_i^k\}_{i,,k}} \sum_{i=1}^N W_i\left(\sum_{k=1}^H \beta_i^k R_i^k\right)$$
$$\text{subject to} \sum_{k=1}^H \beta_i^k R_i^k \geq R_i^{min}, \sum_{k=1}^H \beta_i^k = 1, \forall i$$
$$\beta_i^k = \beta_l^k, \; \forall i \neq l, \forall k \in \{1, ..., H\}, \; \beta_i^k \geq 0, \; \forall i, \forall k$$

In the GPP, we let each UE have a local estimate of the fractions α^k in the CP, and require all the UEs' estimates to be the same. We show that there is a one-to-one mapping between

Table IV
ADMM UPDATE ALGORITHM FOR UE i

Initialization: arbitrary $\boldsymbol{\beta}_i(0) \in \mathcal{B}_i$, $\theta_{ei}^k(0)$ such that $\boldsymbol{\theta}^k \in \boldsymbol{\Theta}^k$, and $\lambda_{ei}^k(0) = 0$, $\forall k \in \{1, ..., H\}$, $\forall e$ such that $i \in e$
For $t = 0$ to $P - 1$
$\boldsymbol{\beta}_i(t+1) = \arg\min_{\boldsymbol{\beta}_i \in \mathcal{B}_i} -\sum_{i=1}^{N} W_i(\boldsymbol{\beta}_i^T \boldsymbol{R}_i) + \sum_{k=1}^{H} \sum_{e \in E} \sum_{q \in e} \left[\lambda_{eq}^k \left(D_{eq}^k \beta_q^k - \theta_{eq}^k \right) + \frac{y}{2} \left(D_{eq}^k \beta_q^k - \theta_{eq}^k \right)^2 \right]$
$\boldsymbol{\beta}_i(t+1)$ is transmitted to all of its neighbors in $\mathcal{N}(i)$.
$\lambda_{ei}^k(t)$ is transmitted to its neighbor connected with edge e, $\forall k \in \{1, ..., H\}$ and $\forall e$ such that $i \in e$
Update $\forall k \in \{1, ..., H\}$ and $\forall e$ such that $i \in e$
$\lambda_{ei}^k(t+1) = \frac{1}{2}(\lambda_{ei}^k(t) + \lambda_{ej}^k(t)) - \frac{y}{2}(D_{ei}^k \beta_i^k(t+1) + D_{ej}^k \beta_j^k(t+1))$, where j is the other endpoint of e.
$\theta_{ei}^k(t+1) = \frac{1}{y}(\lambda_{ei}^k(t+1) - \lambda_{e,i}^k(t)) + D_{ei}^k \beta_i^k(t+1)$
end

the optimal solution $\boldsymbol{\beta}^* = (\boldsymbol{\beta}_1^*, ..., \boldsymbol{\beta}_N^*)$ to the GPP and the optimal solution $\boldsymbol{\alpha}^*$ to the CP. Suppose $\boldsymbol{\beta}^* = (\boldsymbol{\beta}_1^*, ..., \boldsymbol{\beta}_N^*)$ is the optimal solution to the GPP. From the constraints, we have $\boldsymbol{\beta}_i^* = \boldsymbol{\beta}_j^*)$ component-wise $\forall i, j \in \{1, ..., N\}$. Let $\boldsymbol{\alpha}^* = \boldsymbol{\beta}^*$. Then $\boldsymbol{\alpha}^*$ satisfies the constraints in the CP and hence is a feasible solution to the CP. Assume that the optimal solution to the CP is $\boldsymbol{\alpha}' \neq \boldsymbol{\alpha}^*$. Then we can define $\boldsymbol{\beta}_i' = \boldsymbol{\alpha}' \, \forall i \in \{1, ..., N\}$, such that the objective of the GPP at $\boldsymbol{\beta}'$ is higher than that at $\boldsymbol{\beta}^*$, which contradicts the fact that $\boldsymbol{\beta}^*$ is the optimal solution to the GPP. Similarly, we can uniquely define the optimal solution to the GPP from the optimal solution to the CP. Finally, note that if either of CP or GPP is infeasible, then the other problem is infeasible as well. Hence, the CP is equivalent to the GPP.

Under connected graphs, the requirements that all the UEs' local estimates are the same are equivalent to the requirements that all the neighboring UEs' local estimates are the same. With the addition of the auxiliary variables θ_{ei}^k, we get the equivalent problem (DP). ∎

APPENDIX C
PROOF OF THEOREM 3

According to [27], the ADMM algorithm converges with rate $O(1/P)$ if the DP is feasible and if the feasible set is compact. Since $\mathcal{B}_i$ and Θ^k are all closed and bounded polyhedrons, the feasible set is compact. ∎

APPENDIX D
PROOF OF THEOREM 4

To obtain the competitive ratio, we will derive a lower bound of the performance of our proposed distributed policy and an upper bound of the optimal network performance. To obtain the lower bound of the performance of our policy, we first show that our policy is feasible by constructing a feasible solution. In Step 2, the maximum number of colors used is $\Delta + 1$, since each UE selects colors from subset of $\{1, ..., \Delta + 1\} \cap \{1, ..., d_i + 1\}$. After Step 2, we get $\Delta + 1$ MISs that span all the UEs. If the fraction of time assigned to each of these $\Delta + 1$ MISs is $\frac{R_{ub}^{min}}{\log_2(1 + \frac{p_{lb}^{max}}{(D^{ub})^{np} 2\zeta \sigma_{ub}^2})}$, then such an assignment is feasible since $\Delta \leq \Delta^{max} \implies (\Delta + 1) \frac{R_{ub}^{min}}{\log_2(1 + \frac{p_{lb}^{max}}{(D^{ub})^{np} 2\zeta \sigma^2})} \leq 1$. Since the network exhibits ζ-WNI, the minimum instantaneous throughput that can be obtained by UE i is $\log_2(1 + \frac{p_i^{max}}{(D_{iT(i)})^{np} 2\zeta \sigma_{ub}^2})$, and the minimum instantaneous

throughput of any UE is $\log_2(1 + \frac{p_{ub}^{max}}{(D^{ub})^{np} 2\zeta \sigma_{ub}^2})$. Thus, each UE i's throughput requirement is satisfied, because $\frac{R_{ub}^{min}}{\log_2(1 + \frac{p_{lb}^{max}}{(D^{ub})^{np} 2\zeta \sigma_{ub}^2})} \log_2(1 + \frac{p_i^{max}}{(D_{iT(i)})^{np} 2\zeta \sigma_{ub}^2}) \geq R_{ub}^{min}$. In other words, there exists a feasible solution. Since our policy is the optimal feasible solution, it is feasible, too. Based on the minimum throughput requirements R_{ub}^{min}, we get the lower bound $N R_{ub}^{min}$ for the sum throughput obtained by our proposed policy. The upper bound on the optimal network performance is obtained by summing up the maximum instantaneous throughput $\log_2(1 + \frac{p_{ub}^{max}}{(D^{lb})^{np} \sigma_{lb}^2})$ of all the UEs. Based on the lower bound of our policy and the upper bound of the optimal network performance, we get the competitive ratio $\Gamma = \frac{R_{ub}^{min}}{\log_2(1 + \frac{p_{ub}^{max}}{(D^{lb})^{np} \sigma^2})}$. ∎

APPENDIX E
PROOF OF THEOREM 5

Let $\Delta^* = 6\eta$ with $\eta = \lceil \frac{\log_2(1 + \frac{1}{(D^{lb})^{np} \sigma_{lb}^2} p_{ub}^{max})}{R_{lb}^{min}} \rceil$. We assume that the interference graph is constructed using a distance threshold rule (Section V-B). Since each UE's minimum throughput requirement is at least R_{lb}^{min}, and since the maximum instantaneous throughput of any UE is $\log_2(1 + \frac{p_{ub}^{max}}{(D^{lb})^{np} \sigma_{lb}^2})$, each UE needs at least $\frac{R_{lb}^{min}}{\log_2(1 + \frac{p_{ub}^{max}}{(D^{lb})^{np} \sigma_{lb}^2})}$ fraction of time. First, we need to show that if there exists a clique (subset of UEs which are mutually connected) in the interference graph of size $X > \eta$, then the minimum throughput constraints cannot be satisfied. For such a clique, any MIS based scheduling policy will allocate non-overlapping fractions of time to each UE in the clique, which leads to a total fraction $X \frac{R_{lb}^{min}}{\log_2(1 + \frac{p_{ub}^{max}}{(D^{lb})^{np} \sigma_{lb}^2})}$ of time allocated to this clique. However, since $X > \eta$, we have $X \frac{R_{lb}^{min}}{\log_2(1 + \frac{p_{ub}^{max}}{(D^{lb})^{np} \sigma_{lb}^2})} > 1$, which leads to infeasibility. Next, if $\Delta \geq \Delta^*$, we will have at least one clique satisfying this condition. If $\Delta \geq \Delta^*$ there exists a UE i with a degree $d_i \geq 6\eta$, which implies that there exists at least 6η SBSs within a radius of D^{th} around SBS $T(i)$. This circle around SBS $T(i)$ can be partitioned into 6 sectors subtending $\frac{\pi}{3}$ at the center. The distance between any two points located in the sector is no larger than D^{th} (See [7]). Hence, any two points in the sector are connected. If there are at least 6η SBSs in 6 sectors, then there are more than η SBSs in at least one sector (due to Pigeonhole principle), which implies the existence of a clique of size $X > \eta$. ∎

References

[1] E. Hossain, L. B. Le, and D. Niyato, "Self-organizing small cell networks," in *Radio Resource Management in Multi-Tier Cellular Wireless Networks*. Hoboken, NJ, USA: John Wiley Sons, Inc., 2013.

[2] A. Ghosh, N. Mangalvedhe, R. Ratasuk, B. Mondal, M. Cudak, E. Visotsky, T. A. Thomas, J. G. Andrews, P. Xia, H. S. Jo *et al.*, "Heterogeneous cellular networks: From theory to practice," *Communications Magazine, IEEE*, vol. 50, no. 6, pp. 54–64, 2012.

[3] J. G. Andrews, H. Claussen, M. Dohler, S. Rangan, and M. C. Reed, "Femtocells: Past, present, and future," *Selected Areas in Communications, IEEE Journal on*, vol. 30, no. 3, pp. 497–508, 2012.

[4] R. Ramaswami and K. K. Parhi, "Distributed scheduling of broadcasts in a radio network," in *INFOCOM'89. Proceedings of the Eighth Annual Joint Conference of the IEEE Computer and Communications Societies. Technology: Emerging or Converging, IEEE*. IEEE, 1989, pp. 497–504.

[5] I. Cidon and M. Sidi, "Distributed assignment algorithms for multihop packet radio networks," *Computers, IEEE Transactions on*, vol. 38, no. 10, pp. 1353–1361, 1989.

[6] A. Ephremides and T. V. Truong, "Scheduling broadcasts in multihop radio networks," *Communications, IEEE Transactions on*, vol. 38, no. 4, pp. 456–460, 1990.

[7] K. Ahuja, Y. Xiao, and M. van der Schaar, "Online report for distributed interference management in smallcell networks," 2014. [Online]. Available: "http://arxiv.org/pdf/1411.5102v2.pdf"

[8] J. Huang, R. A. Berry, and M. L. Honig, "Distributed interference compensation for wireless networks," *Selected Areas in Communications, IEEE Journal on*, vol. 24, no. 5, pp. 1074–1084, 2006.

[9] C. U. Saraydar, N. B. Mandayam, and D. Goodman, "Pricing and power control in a multicell wireless data network," *Selected Areas in Communications, IEEE Journal on*, vol. 19, no. 10, pp. 1883–1892, 2001.

[10] E. J. Hong, S. Y. Yun, and D.-H. Cho, "Decentralized power control scheme in femtocell networks: A game theoretic approach," in *Personal, Indoor and Mobile Radio Communications, 2009 IEEE 20th International Symposium on*, Sept 2009, pp. 415–419.

[11] S. Stanczak, M. Wiczanowski, and H. Boche, *Fundamentals of resource allocation in wireless networks: theory and algorithms*. Springer, 2009, vol. 3.

[12] M. Chiang, C. W. Tan, D. P. Palomar, D. O'Neill, and D. Julian, "Power control by geometric programming," *Wireless Communications, IEEE Transactions on*, vol. 6, no. 7, pp. 2640–2651, 2007.

[13] J. Ling, D. Chizhik, and R. Valenzuela, "On resource allocation in dense femto-deployments," in *Microwaves, Communications, Antennas and Electronics Systems, 2009. COMCAS 2009. IEEE International Conference on*. IEEE, 2009, pp. 1–6.

[14] G. J. Foschini and Z. Miljanic, "A simple distributed autonomous power control algorithm and its convergence," *Vehicular Technology, IEEE Transactions on*, vol. 42, no. 4, pp. 641–646, 1993.

[15] S. G. Kiani and D. Gesbert, "Optimal and distributed scheduling for multicell capacity maximization," *Wireless Communications, IEEE Transactions on*, vol. 7, no. 1, pp. 288–297, 2008.

[16] D. Gesbert, S. G. Kiani, A. Gjendemsjo *et al.*, "Adaptation, coordination, and distributed resource allocation in interference-limited wireless networks," *Proceedings of the IEEE*, vol. 95, no. 12, pp. 2393–2409, 2007.

[17] "Access, evolved universal terrestrial radio, further advancements for e-utra physical layer aspects (release 10)," *3GPP TS V9*. [Online]. Available: "http://www.3gpp.org/specifications/releases/70-release-10"

[18] K. Jain, J. Padhye, V. N. Padmanabhan, and L. Qiu, "Impact of interference on multi-hop wireless network performance," *Wireless networks*, vol. 11, no. 4, pp. 471–487, 2005.

[19] H. Li, X. Xu, D. Hu, X. Qu, X. Tao, and P. Zhang, "Graph method based clustering strategy for femtocell interference management and spectrum efficiency improvement," in *Wireless Communications Networking and Mobile Computing (WiCOM), 2010 6th International Conference on*. IEEE, 2010, pp. 1–5.

[20] S. Uygungelen, G. Auer, and Z. Bharucha, "Graph-based dynamic frequency reuse in femtocell networks," in *Vehicular Technology Conference (VTC Spring), 2011 IEEE 73rd*. IEEE, 2011, pp. 1–6.

[21] P. Lee, T. Lee, J. Jeong, and J. Shin, "Interference management in lte femtocell systems using fractional frequency reuse," in *Advanced Communication Technology (ICACT), 2010 The 12th International Conference on*, vol. 2. IEEE, 2010, pp. 1047–1051.

[22] C. W. Tan, S. Friedland, and S. H. Low, "Spectrum management in multiuser cognitive wireless networks: Optimality and algorithm," *Selected Areas in Communications, IEEE Journal on*, vol. 29, no. 2, pp. 421–430, 2011.

[23] K. Han, S. Woo, D. Kang, and S. Choi, "Automatic neighboring bs list generation scheme for femtocell network," in *Ubiquitous and Future Networks (ICUFN), 2010 Second International Conference on*. IEEE, 2010, pp. 251–255.

[24] D. Lopez-Perez, I. Guvenc, G. De La Roche, M. Kountouris, T. Q. Quek, and J. Zhang, "Enhanced intercell interference coordination challenges in heterogeneous networks," *Wireless Communications, IEEE*, vol. 18, no. 3, pp. 22–30, 2011.

[25] Ö. Johansson, "Simple distributed $\Delta+1$ coloring of graphs," *Information Processing Letters*, vol. 70, no. 5, pp. 229–232, 1999.

[26] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and distributed computation: numerical methods*. Prentice-Hall, Inc., 1989.

[27] E. Wei and A. Ozdaglar, "On the o (1/k) convergence of asynchronous distributed alternating direction method of multipliers," *arXiv preprint arXiv:1307.8254*, 2013.

[28] A. Goldsmith, *Wireless communications*. Cambridge university press, 2005.

[29] S. Y. Seidel and T. S. Rappaport, "914 mhz path loss prediction models for indoor wireless communications in multifloored buildings," *Antennas and Propagation, IEEE Transactions on*, vol. 40, no. 2, pp. 207–217, 1992.

[30] R. M. Karp, "Probabilistic recurrence relations," *Journal of the ACM (JACM)*, vol. 41, no. 6, pp. 1136–1150, 1994.

Kartik Ahuja received the B.tech-M.Tech dual degree in Electrical Engineering from Indian Institute of Technology Kanpur in 2013, and is currently pursuing the Ph.D. degree in the Electrical Engineering Department at University of California, Los Angeles. His research interests include game theory, optimization, mathematical modelling and network science.

Yuanzhang Xiao received the B.E. and M.E. degrees in electrical engineering from Tsinghua University in 2006 and 2009, and the Ph.D. degree in electrical engineering from UCLA in 2014. His research interests include multi-agent sequential decision making, game theory, and optimization, with applications to cyber-physical systems and wireless communications.

Mihaela van der Schaar is Chancellor's Professor of Electrical Engineering at UCLA. She was Distinguished Lecturer of the Communications Society and Editor in Chief of IEEE Transactions on Multimedia. She is also the founding director of the UCLA Center for Engineering Economics, Learning, and Networks. Her research interests include engineering economics and game theory, network science, expert and social networks, online reputation and social media, wireless networks, online learning, and real-time stream mining.