

A semi-distributed mechanism for Inter-cell Interference Coordination exploiting the ABSF paradigm

Vincenzo Sciancalepore^{*†‡}, Ilario Filippini[‡], Vincenzo Mancuso^{*}, Antonio Capone[‡], Albert Banchs^{*†}
* IMDEA Networks Institute † Universidad Carlos III de Madrid ‡ DEIB - Politecnico di Milano

Abstract—Inter-Cell Interference Coordination (ICIC) has been identified for LTE as the main instrument for interference control. With ICIC, quality requirements can be guaranteed while avoiding the complexity of coordinated baseband processing approaches. However, most ICIC schemes proposed so far rely on centralized multi-cell scheduling algorithms that involve very heavy signaling overhead and, as a result, cannot be used for dense cellular layouts. In this paper, we propose $H_2(IC)_2$, a novel ICIC scheme that, in contrast to previous approaches, incurs very low overhead and is practical for dense deployments. $H_2(IC)_2$ is based on the Almost Blank SubFrame (ABSF) approach specified by 3GPP, which controls interference by avoiding data transmission in some subframes. Our scheme follows a two-tier approach, consisting of (i) the *local schedulers*, which perform the scheduling decisions locally and compute ABSF patterns, and (ii) a *central coordinator*, which supervises ABSF decisions. As a result of such a two-tier design, the scheme requires very light signaling to drive the local schedulers to globally efficient operating points. We analyze the convergence of distributed ABSF/scheduling decisions by using game theoretical tools and show that $H_2(IC)_2$ performs fairly close to the benchmark provided by a centralized omniscient scheduler.

I. INTRODUCTION

The steady increase of traffic demand in LTE networks requires new network architectures able to handle such demands. These architectures involve the deployment of more devices than traditional cellular networks, including relay nodes, small-scale base stations and remote radio heads, as recommended by recent 3GPP releases [1]. The resulting increase of network device density clearly exacerbates interference issues. In this framework, addressing interference issues is key to the success of future LTE networks.

The Almost Blank SubFrame (ABSF) has been introduced by 3GPP as the main Inter-Cell Interference Coordination (ICIC) technique in LTE-Advanced [2]. With ABSF, the activity of a base station can be prevented in some parts of the LTE frame (*subframe blanking*), thus limiting interference.

A number of valid solutions using ABSF have been proposed in the literature [3], [4]. However, these solutions mainly approach the problem from a centralized point of view (e.g., a macro base station optimizes the activities of small base stations) and require a huge exchange of Channel State Information (CSI) messages to evaluate the potential

The research leading to these results has received funding from the European Union's Seventh Framework Programme (FP7-ICT-2013-EU-Japan) under grant agreement 608637 (MiWEBBA) and grant agreement 318115 (CROWD). The work has also been partially supported by the Madrid Regional Government through the TIGRE5-CM program (S2013/ICE-2919).

978-1-4673-7331-9/15/\$31.00 © 2015 European Union

interference caused by base station transmissions on scheduled users. As a consequence, they do not scale with the number of users or base stations and result in unpractical approaches [5], [6]. Another drawback of these solutions is that, as they select blanking patterns without specifying which users have to be scheduled, they typically assume worst case interference conditions, resulting in suboptimal performance [7].

To overcome the above limitations, in this paper we present a novel scheme called *hybrid two-tier inter-cell interference coordination* ($H_2(IC)_2$). $H_2(IC)_2$ provides a lightweight coordination by following a two-tier approach that consists of: (i) the *base stations* acting as *local schedulers*, each of which self organizes its transmission activities by jointly scheduling users and selecting its ABSF pattern based on local CSI; and (ii) a *central coordinator*, which supervises ABSF decisions of the base stations and drives the system to the best possible performance without imposing centralized decisions on ABSF patterns. Hence, our proposed solution relies on a semi-distributed approach that offloads and reduces the computational burden from a centralized controller while drastically abating the signaling overhead. This makes our approach a first step towards a practical and effective solution to ABSF that can be implemented in real networks.

We validate the proposed scheme via simulation by comparing its performance against an omniscient scheme based on brute force optimization. We also analyze its convergence through game theory. Our results show that $H_2(IC)_2$ achieves near-optimal performance and exhibits significant advantages over existing schemes in terms of efficiency, complexity, fairness, and throughput.

In addition, valuable comparisons with existing power control schemes reveal that complex approaches like [8] bring little additional gain with respect to $H_2(IC)_2$ and behave less fairly, whereas low-complexity solutions like [9] exhibit lower efficiency with respect to our proposal. Moreover, our results also show that $H_2(IC)_2$ is able to dynamically adapt to changing conditions by quickly converging to the stable point of operation (in general, it reaches a point of operation within 5% of the stable point of operation in only a few iterations).

II. RELATED WORK

A wide range of solutions have been proposed for interference coordination in LTE networks. For instance, a very fine interference control can be obtained by cooperative transmission schemes leveraging beamforming techniques; these approaches are usually centralized and require a huge exchange

of control signaling [3], [4]. Less signaling is required by coordinated user scheduling among base stations [10], [11], [12] and joint user scheduling and power control or fractional load [13], [14], [15]. However, the exchange of CSI messages may generate a considerable amount of traffic.

In the last years, the ABSF technique has gained momentum in the scientific community because of its ability to trade-off between performance improvement and low implementation complexity [16], [17]. Quantitative approaches, like in [18], aim at determining the best muted (blank) subframe density according to a given traffic distribution. Further solutions face the HetNet scenarios where a macro base station and several small base stations under its coverage coordinate on ABSF patterns to follow [19], [20], [21]. Some more sophisticated approaches additionally consider the user association by properly throttling the Cell Selection Bias [10], which has been shown to further improve network spectral efficiency [5], [6]. The main drawbacks of these existing ABSF solutions are twofold: they require to gather either per-user CSI or topological information. In both cases, scalability problems arise as the number of base stations and users increases. Moreover, collecting topological information works on long timescales, since it requires to tune propagation and drive tests.

Recently, an interesting solution for resource management in ODFMA femtocells has been presented in [22]. Although not explicitly referring to the ABSF mechanism, the work proposes a mechanism where each base station individually detects the best region of the time-frequency space where transmitting by using a probe-and-adapt algorithm. The lack of coordination among base stations, however, may not permit a full exploitation of the available transmission opportunities, and thus, might limit the practicality of the proposal. Similarly to [22], also in [23] the ABSF mechanism is not mentioned, but the very similar concept of reuse patterns for base station ON/OFF activities is used. With the goal of maximizing the total user throughput, the authors of [23] determine the best temporal duration of each pattern, given a set of chosen patterns. However, the design of the set of chosen patterns is not addressed, although it strongly influences the performance of the proposed algorithm. In addition, differently to our proposed solution, the blind maximization of the total throughput can lead to highly unfair user allocations. The authors of [24] present a game theoretical approach to ICIC. Their approach addresses the coordination among base stations over a set of finite resources as a non-cooperative game. However, they only target the minimization of the perceived interference, and do not take into account user scheduling.

None of the above proposals embodies the set of features that characterize our approach and can be summarized as follows: (i) the distributed ABSF interference coordination problem has been formalized and its convergence investigated, (ii) the proposed mechanism is distributed and the complexity of the centralized coordination is abated and split among base stations requiring little signaling exchange with the central coordinator, (iii) the proposed mechanism is adaptive and can adjust its parameters according to traffic dynamics, and (iv) despite the simplicity of the proposed mechanism, our results

show remarkable near-optimal performance figures.

III. ICIC PROBLEM FORMULATION

The goal of ICIC is to improve system spectral efficiency. To this end, ICIC optimally orchestrates base station activities and performs user scheduling on a time-slot basis, i.e., per Transmission Time Interval (TTI). Here, we cast the ICIC problem into an LTE-Advanced network that implements the ABSF mechanism. With this mechanism, each base station uses an ABSF pattern, which is a bitmap that specifies which TTIs must be blanked by the base station.

In what follows, we first formulate the ICIC problem from a *centralized* scheduling perspective, which is practically unfeasible due to computational and signaling overhead. Then we show how to abate and distribute the computational load of the ICIC problem over the base stations. However, as we show in the final part of this section, introducing such a *fully distributed* approach requires some game theory tools, and does not always guarantee that base stations' decisions *converge*. To solve these issues, in Section IV we propose our two-tier semi-distributed mechanism.

For the sake of simplicity, problem formulations presented in this section consider downlink traffic only; however, very similar techniques could be used for uplink traffic. Additionally, we focus on elastic traffic (i.e., traffic for which there are no stringent requirements in terms of latency and bandwidth) since it represents the most common traffic type in mobile data networks. Following current cellular deployments, we consider that base stations transmit at fixed power; therefore, it is sufficient to know which base stations are active to determine the level of interference suffered by a transmission.

A. Centralized problem

The main assumption behind the centralized problem is that users' CSI is perfectly known. Such information is gathered and updated by a centralized controller, which uses it to compute the optimal scheduling. Specifically, the centralized controller maps each user u onto any available TTI t and issues the resulting user scheduling information to every base station i . In this way, transmissions during each frame are entirely controlled by the centralized controller. The time horizon of the optimization consists in a set of TTIs $\mathcal{T} = 1..T$, in which base stations' activities are coordinated. While this scheme is clearly unpractical, the optimal solution to this problem provides us with the benchmark corresponding to the best possible performance of any implementable algorithm.

The objective function to be maximized by the centralized problem, $\hat{\eta}$, is the sum of the utilities of the individual base stations. Following the widely accepted *max-min* fairness criterion, we define the utility of base station i as the minimum rate of all the users in the base station.¹ With this, we can formulate the centralized optimization problem with the following Integer Linear Programming (ILP) model.

¹Note that the selected objective function provides a trade-off between maximizing the spectral efficiency and guaranteeing a minimum level of service quality, as pointed out, e.g., in [7]. Nevertheless, different objective functions can be considered as well, without substantially changing the proposed approach and the following analysis.

Problem CENTRAL:

$$\begin{aligned}
& \text{maximize } \hat{\eta} = \sum_{i \in \mathcal{N}} \left(\min_{(u,t) \in \mathcal{U}_i \times \mathcal{T}} R^r \cdot x_u^{r,t} \right), \\
& \text{s.t. } \sum_{u \in \mathcal{U}_i, r \in \mathcal{R}} x_u^{r,t} \leq y_{i,t}, \quad \forall i \in \mathcal{N}, t \in \mathcal{T}, \\
& \frac{P G_{u,i}}{N_0 + \sum_{k \in \mathcal{N}: k \neq i} P G_{u,k} \cdot y_{k,t}} \geq \gamma^r \cdot x_u^{r,t}, \\
& \quad \forall i \in \mathcal{N}, u \in \mathcal{U}_i, r \in \mathcal{R}, t \in \mathcal{T}, \\
& y_{i,t}, x_u^{r,t} \in \{0; 1\}, \quad \forall i \in \mathcal{N}, u \in \mathcal{U}_i, r \in \mathcal{R}, t \in \mathcal{T};
\end{aligned} \tag{1}$$

where \mathcal{N} is the set of base stations, \mathcal{U}_i is the set of users associated to base station i , $G_{u,i}$ is the channel gain between user u and the base station i , P is the transmitting power of any base station and N_0 is the background noise. Binary decision variables $x_u^{r,t}$ indicate whether user u is scheduled or not into TTI t , and binary variables $y_{i,t}$ take value 1 if base station i is active at TTI t , and 0 if the TTI is blanked. When scheduled, u 's transmission rate is r (the available rates each user can adopt are listed in the vector R^r corresponding to the Modulation and Coding Schemes (MCSs) of the 3GPP standard [25], where set \mathcal{R} is the index set of vector R^r).

The first set of constraints of Problem CENTRAL impose that at most one user per base station is accommodated in a single TTI t (note that this assumption can be easily relaxed in order to address different user schedulers). The second set of constraints ensure that, when setting a rate r for a user in a TTI, the experienced signal-to-noise-and-interference ratio (SINR) is not lower than the activation threshold γ^r for this rate. Note that, while these constraints are not linear, they can be straightforwardly linearized in order to solve the problem with state-of-the-art solvers (see [26]).

Problem CENTRAL can be reduced to a bin-packing problem in which the sum of interferences cannot exceed a threshold. Therefore, this problem is NP-hard [27]. Moreover, it involves a very high overhead to deliver CSI information to the centralized controller, which needs this information to select the ABSF patterns and compute the user scheduling. Thus, while the centralized approach can be an attractive option for small networks, a less complex and more distributed approach is required to deal with the case of very dense wireless networks consisting of hundreds of base stations and thousands of wireless nodes.

B. Distributed problem

We next present a distributed formulation of Problem CENTRAL, whose implementation distributes the computational burden of the original problem over the base stations present in the network. Specifically, to reduce complexity, in the distributed problem each base station only optimizes the scheduling of its own users and considers that other base stations use fixed ABSF patterns. However, this approach needs an iterative mechanism to find the optimal ABSF pattern of all base stations. Note that, with the distributed approach, *the complexity of the problem to solve is dramatically reduced, while the number of iterations required to converge will be shown to grow at most quadratically with the network size.*

To formulate the distributed approach, the original problem is split into several smaller instances, which are solved locally

by each base station. To solve a problem instance, the base station is provided with the activity pattern declared by other base stations. This is given by *ABSF patterns*, $ABSF_{i,t}$, which are exchanged among base stations ($ABSF_{i,t} = 0$ if base station i blanks TTI t). Such information is needed by each base station to estimate the interference in each TTI suffered by any possible candidate scheduled user. With the above information, and without explicitly forcing any additional constraint, each base station i would schedule users selfishly in the entire set of T TTIs, in order to optimize the local utility. Therefore, to avoid that base stations use all available TTIs, in the distributed problem formulation, we grant a single base station i access to up to M_i TTIs over T available TTIs; such M_i value plays a key role in the distributed mechanism, as it will be clarified in Section IV.

The above description corresponds to the following instance of the local problem for base station i , which can be formulated as an ILP model as follows:

Problem LOCAL:

$$\begin{aligned}
& \text{maximize } \hat{\eta}_i = \min_{(u,t) \in \mathcal{U}_i \times \mathcal{T}} R^r \cdot x_u^{r,t} \cdot a_{u,i}, \\
& \text{s.t. } \sum_{u,r} x_u^{r,t} \cdot a_{u,i} \leq 1, \quad \forall t \in \mathcal{T}, \\
& \frac{P G_{u,i}}{N_0 + \sum_{k \in \mathcal{N}: k \neq i} P G_{u,k} \cdot ABSF_{k,t}} \geq \gamma^r \cdot x_u^{r,t}, \\
& \quad \forall u \in \mathcal{U}_i, r \in \mathcal{R}, t \in \mathcal{T}, \\
& \sum_{u,t,r} x_u^{r,t} \leq M_i, \\
& x_u^{r,t} \in \{0; 1\}, \quad \forall u \in \mathcal{U}_i, r \in \mathcal{R}, t \in \mathcal{T};
\end{aligned} \tag{2}$$

where all parameters and constraints have the same meaning as in Problem CENTRAL, except for the third constraint, which limits the number of usable TTIs to M_i . Note that a feasible solution of Problem LOCAL can be computed by using any available max-min scheduling heuristic (see, e.g., [28]).

As stated above, each base station i is in charge of solving Problem LOCAL, by computing the optimal user scheduling into available TTIs. Note that the solution of this problem depends on the solutions computed by the other base stations, since the SINR of each user is given by the interference generated by the other base stations in the system when they are active. Therefore, in the distributed approach formulation, each base station simply schedules local users in order to maximize the objective function defined in Problem LOCAL. However, the schedule defines the activity of the base station, and the interference generated towards other base stations, which, in turn, can react readjusting their scheduling in order to adapt to changed interference conditions. A new scheduling may cause new interference levels, therefore each base station must iteratively solve Problem LOCAL, until the system converges to a stable solution.

C. Convergence analysis of the distributed approach

In the following, we analyze the (fully) *distributed* approach formulated above from a game theoretic standpoint and show that its *convergence is not guaranteed*. Building on this result, later in Section IV we propose a *semi-distributed* approach that guarantees the convergence of the game.

Based on game theory, the distributed approach can be modeled as a game where base stations iteratively play in order to maximize their utility. Let us define this game as an *Interference Coordination Game* Γ , where each base station i acts as a player (the terms “player” and “base station” are indistinctly used in the rest of the paper). The set of strategies of each player \mathbb{S}_i consists in the set of pairs (user, TTI), $(u, t) : u \in \mathcal{U}_i, t \in \mathcal{T}$, available for each base station according to constraints in Problem LOCAL.

In order to analyze the convergence of the above game, we rely on the concept of *Bottleneck Matroid Congestion Game* (for a detailed discussion, we refer the reader to [29]). A *Bottleneck Congestion Game* is a class of games where resources are shared among players. The utility of each player depends on the utility of the resources she chooses and the number of players choosing the same resources: the higher the congestion, the lower the utility. In particular, the individual player utility is the minimum of the utilities of the resources chosen in her strategy. In sequential improvement dynamics, players act selfishly and play a Best Response strategy (BR) $S_i^* \in \mathbb{S}_i$, i.e., the strategy that maximizes their individual utility function, given the strategies played by other players.

In addition to the above, regular congestion games can be generalized in *player-specific congestion games* and *weighted congestion games*. In the former, every player has her own utility function for every resource. In a weighted congestion game, every player affects the other players strategies with a different weight, namely, she causes a different level of congestion.

The following theorem shows that our game falls in the intersection between the above categories, and hence existing results on these classes of games can be applied to our problem.

Theorem 1. *The Interference Coordination Game Γ is a Weighted Player-specific Bottleneck Matroid Congestion Game.*

Proof: Here we provide the reader with a sketch of the proof. The Interference Coordination Game Γ is player-specific since utility is player-specific as it depends on received interference, and it is a congestion game in which congestion weights are given by the interference caused by the scheduled users in each TTI. Moreover, strategies’ constraints induced by constraints in Problem LOCAL make the strategy space a matroid, thus Γ is a Matroid Congestion Game. ■

Regular bottleneck congestion games have been proven to satisfy the *finite improvement property*, which states that an arbitrary BR sequence played by each player during the game always converges to an equilibrium in a finite number of steps [29]. However, the generalizations of player-specificity and different congestion weights introduce many degrees of freedom, which weakens the game structure and its convergence guarantees. Indeed, the following theorem shows that Weighted Player-specific Bottleneck Matroid Congestion Games do not satisfy the finite improvement property.

Theorem 2. *Weighted player-specific matroid bottleneck congestion games do not exhibit the finite improvement property*

TABLE I: Example of weighted player-specific matroid bottleneck congestion game that does not converge

Rate	Alone	With BS 1	With BS 2	With BS 3
$c_{u_1,t}$	2.0	—	1.5	1.1
$c_{u_2,t}$	2.0	1.1	—	1.5
$c_{u_3,t}$	2.0	1.5	1.1	—

in best-response improvement dynamics.

Proof: Let us consider a scenario with $T = 2$ TTIs and 3 base stations, each of them associated with $|\mathcal{U}_i| = 1$ distinct user. For each player i , the strategy space \mathbb{S}_i is defined as $\mathbb{S}_i = \{\{(u_i, t_1)\}; \{(u_i, t_2)\}; \{(u_i, t_1), (u_i, t_2)\}\}$. Let us assume an upper bound on available TTIs per base station $M_i = 1, \forall i \in N$ and a user rate $c_{u,t}$, expressed as bits/symb/TTI, according to Table I. Now we consider the sequence of strategies taken by each player, described by Table II.

Whenever a player i chooses a new strategy at the k^{th} step in order to maximize the utility function (bold-marked), the value of utility function calculated by the other players may decrease and they may want to change their strategy. This leads to a loop where players sequentially return on the same strategies indefinitely, such as strategies at step k and strategies at step $k + 6$. Hence, players playing arbitrary best responses do not necessarily converge to a Nash equilibrium in Weighted Player-specific Bottleneck Matroid Congestion Games, and thus, a finite improvement property does not always exist. ■

The above analysis has shown that the distributed approach may not converge.² Moreover, it does not ensure that M_i values are selected according to a global fairness. In order to address these shortcomings, in the next section we propose a semi-distributed two-level mechanism where a central coordinator controls the behavior of the distributed game.

IV. CONTROLLER-AIDED DISTRIBUTED MECHANISM

Building on the results of the previous section, in the following we design a *semi-distributed approach* that relies on a central coordinator. The global scheme of the system mechanism is depicted in Fig. 1. As shown in the figure, the scheme operates at two different timescales:

- On a long-term timescale (in the order of seconds), a central coordinator, the *Controller*, is in charge of adjusting the M_i value of each base station, where M_i gives the maximum number of TTIs that base station i can use to schedule its users within the time horizon T by solving Problem LOCAL. In addition, adapting M_i is used to react to traffic changes in the system.
- At a shorter timescale, base stations play the Interference Coordination Game Γ by sequentially exchanging their scheduling decisions in terms of ABSF patterns.³ As

²It is worthwhile noting that the somehow pathological scheduling behavior considered in theorem’s proof does not commonly exhibit in networks; indeed, according to the simulations conducted for typical realistic scenarios, the interference coordination game Γ reaches an equilibrium with very high probability. Nevertheless, we still need to design an algorithm whose convergence is guaranteed.

³Note that there is no need to announce which specific user will be scheduled in a specific TTI, since base stations transmit at a fixed power and thus their activity causes the same level of interference independently of the scheduled user. Therefore, it is sufficient to propagate a binary string of T bits containing the ABSF pattern.

TABLE II: State evolution for a weighted player-specific matroid bottleneck congestion game that does not converge (example used in the proof of Theorem 2)

$S_i^{*(s)}$, $f(S_i^*, S_{-i})$	$s = k - 1$	$s = k$	$s = k + 1$	$s = k + 2$	$s = k + 3$	$s = k + 4$	$s = k + 5$	$s = k + 6$
BS 1	$\{u_1, t_1\}, 2$	$\{u_1, t_1\}, (1,1)$	$\{u_1, t_2\}, 1.5$	$\{u_1, t_2\}, 2$	$\{u_1, t_2\}, (1,1)$	$\{u_1, t_1\}, 1.5$	$\{u_1, t_1\}, 2$	$\{u_1, t_1\}, (1,1)$
BS 2	$\{u_2, t_2\}, 2$	$\{u_2, t_2\}, 2$	$\{u_2, t_2\}, (1,1)$	$\{u_2, t_1\}, 1.5$	$\{u_2, t_1\}, 2$	$\{u_2, t_1\}, (1,1)$	$\{u_2, t_2\}, 1.5$	$\{u_2, t_2\}, 2$
BS 3	-	$\{u_3, t_1\}, 1.5$	$\{u_3, t_1\}, 2$	$\{u_3, t_1\}, (1,1)$	$\{u_3, t_2\}, 1.5$	$\{u_3, t_2\}, 2$	$\{u_3, t_2\}, (1,1)$	$\{u_3, t_1\}, 1.5$




Fig. 1: Hybrid two-level mechanism for intercell interference coordination. In the short-term level (bottom side of the figure), the game is played amongst the base stations, while in the long-term level (top side) the controller decides the number of available TTIs per base station.

described in the following, the central coordinator does not directly participate in the game, but it controls its convergence by limiting the number of iterations.

The remaining challenge for coordinator-aided approach is the design of the algorithms executed by the central coordinator to (i) ensure convergence, and (ii) adjust the values M_i . In the following we address the design of those algorithms, which aim at driving the system behavior to an optimal state in the long run.

A. Convergence control of game Γ

In order to guarantee the convergence of the game, the central coordinator imposes a deadline of Z TTIs, with $Z < T$: if the game has not finished by this deadline, it is terminated by the central coordinator.

When the game finishes before the deadline, the resulting scheduling corresponds to an equilibrium of the game, which ensures that resources are fairly shared among base stations. In contrast, when the game is terminated by the central coordinator, base stations use the scheduling that they computed in the latest iteration of the game, which does not correspond to an equilibrium. Thus, in the latter case some base stations could potentially have a better scheduling (i.e., more resources) than the others. However, as shown by our results of Section V, we have observed that in practice the game can be interrupted after only a very few iterations without negatively impacting fairness or performance in a significant manner.

The deadline Z has been chosen in order to have a valid scheduling before the current period T finishes: the resulting

scheduling (and the corresponding ABSF pattern) will then be used for the next period. During the game, transmissions and users are scheduled according to the result of the previous period. Note that the iterations of game Γ do not need to be synchronized with the TTIs; they can be much faster, allowing for more than Z iterations within Z TTIs. Indeed, the execution of one iteration only requires passing the “current” ABSF patterns from one base station to another. As shown in Section V-C, Z can be chosen in the range $[|\mathcal{N}|, |\mathcal{N}|^2]$.

B. Dynamic adjustment of TTI bounds M_i

One critical aspect for the performance of the proposed mechanism is the setting of the M_i parameters, which give the maximum number of non-blank TTIs available to each base station. Indeed, if the M_i values are too small, performance is degraded because, even if base stations can be scheduled one at a time with low interference, the number of TTIs available for transmitting can be too small to accommodate all users. Conversely, if the M_i values are too large, performance is degraded as a result of too many base stations scheduled together and creating high interference. Thus, performance is maximized when the M_i parameters are optimally set to values that are neither too large nor too small. In the rest of this section, we design an adaptive algorithm that follows an *additive-increase multiplicative-decrease* (AIMD) strategy [30] to find the optimal M_i setting.

In addition to optimally setting M_i to improve the performance of the network, the adaptive algorithm also aims at dynamically adjusting the M_i configuration to follow the changes in traffic and interference. From this perspective, the adaptive algorithm is a long-term process. In contrast, the distributed game is a short-term process played once per each period of T TTIs. This implies that the duration of the period T cannot exceed a few hundreds frames, which corresponds to a few seconds during which traffic and channel conditions remain practically unchanged.

From a high level perspective, the algorithm works as follows. At the end of each period of T TTIs, the controller gathers from the base stations the performance resulting from the M_i values (and the corresponding ABSF pattern) used during the period. The metric chosen to represent the performance of a base station is given by the average user rate experienced by users of base station i in the period⁴, i.e.:

$$\eta_i = \frac{1}{|\mathcal{U}_i|} \sum_{(u,t) \in \mathcal{U}_i \times \mathcal{T}} c_{u,t}. \quad (3)$$

The controller then uses the sum of the individual performance metrics, $\eta = \sum_{i \in \mathcal{N}} \eta_i$, to keep track of the global

⁴Note that, since user allocation is carried out according to Problem LOCAL, the max-min objective tends to assign rates with limited variance; as a consequence, the average user rate and the rate of the worst-off user are likely to be similar.

system performance and drive M_i to the setting that maximizes η . The algorithm to find such M_i setting follows an AIMD strategy: the M_i values are increased as long as performance is improved, and, when performance stops improving, then the M_i values are decreased. After each update of the M_i values, these are distributed to the base stations and used in the following period (i.e., the following iteration of game Γ).

The specific algorithm executed to calculate the new set of TTI bounds M_i is described in Algorithm 1. Each iteration of the algorithm is identified by an index k . At the initial step ($k = 0$), the controller initializes the system performance metrics η to 0 and assigns the initial TTI bounds $M_i^* = \lceil T/|\mathcal{N}| \rceil$ for every base station. This initial M_i^* setting has been chosen to allow base stations to schedule their users in disjoint portions of the period, which helps the convergence of the algorithm in case of very high mutual interference between all base stations. The M_i^* also provide a lower bound for M_i .

Algorithm 1 Resource Sharing Algorithm: Adaptive algorithm to dynamically design M_i . Called at the end of $(k-1)^{th}$ ABSF pattern

Input: $\mathcal{N}, T, M_i^*, \eta^{(k-1)}$
Initialization: $\eta^{(k)} \leftarrow 0; M_i \leftarrow M_i^*, \forall i \in \mathcal{N}$
Procedure
1: $\mathcal{V} \leftarrow \{\eta_i, \forall i \in \mathcal{N}\}$
2: Order \mathcal{V} non-increasing
3: $\eta^{(k)} = \sum_{i \in \mathcal{N}} \eta_i$
4: **if** $\eta^{(k)} > \eta^{(k-1)}$ **then**
5: **while** $\mathcal{V} \neq \emptyset$ **do**
6: $e = \text{pop}(\mathcal{V})$
7: **Consider** index i of element e
8: **if** $M_i^{(k-1)} < T$ **then**
9: $M_i^{(k)} = M_i^{(k-1)} + 1$
10: **break**
11: **end if**
12: **end while**
13: **else**
14: **while** $\mathcal{V} \neq \emptyset$ **do**
15: $e = \text{pop}(\mathcal{V})$
16: **Consider** index i of element e
17: **if** $M_i^{(k-1)} > M_i^*$ **then**
18: $M_i^{(k)} = \max \left\{ M_i^*; \left\lceil M_i^{(k-1)} / 2 \right\rceil \right\}$
19: $\eta^{(k)} = 0$
20: **break**
21: **end if**
22: **end while**
23: **end if**

At each step, the controller collects the performance metrics η_i from base stations and checks whether the performance of this period, $\eta^{(k)}$, has improved with respect to the previous period, $\eta^{(k-1)}$ (line 3). If this is the case, this means that system performance is raising and the controller increases TTI bounds M_i as follows. The controller increases by 1 unit the M_i of the base station with the smallest η_i whose M_i is below T (lines 8-9). Once one M_i value is increased, step k of the algorithm terminates (line 10).

If no M_i can be increased, which means that all base stations are active in all TTIs, then no adjustment of the M_i values is made as long as the system performance does not degrade. In case performance degrades, i.e., $\eta^{(k)}$ decreases, (line 13), the controller drastically reduces the M_i . Specifically, the controller looks at the base station i with the largest η_i whose

TABLE III: Overhead of centralized and $H_2(IC)_2$ semi-distributed approaches

Interface	centralized approach	$H_2(IC)_2$ approach
I_C	$64 \cdot \mathcal{U} \cdot \mathcal{N} + T \cdot \mathcal{N} $	$64 \cdot \mathcal{N} $
I_B	0	$T \cdot k \cdot \mathcal{N} $

M_i is above M_i^* . It sets the new M_i value of this station equal to the minimum between the half of the current M_i value and the lower bound M_i^* (lines 17-18). If $M_i = M_i^*$ for all i , no change is carried out.

The rationale behind using AIMD to adjust the M_i values is that, similar to what happens with TCP, increasing the utilization of the system (i.e., increasing M_i values) may lead to congestion (in our case, this corresponds to excessive interference), which causes user rates to drop. In this case, a quick reaction is required by the controller to drive the system to a safe point of operation, by properly adjusting TTI bounds M_i . Also similar to TCP, the additive increase of TTI bounds M_i allows to gracefully approach the optimal utilization of the system. Furthermore, since the problem may admit more than one local maximum, using multiplicative decrease for the TTI bounds M_i helps our heuristic to escape from a local maximum where the optimization function may be trapped in.

As a side comment, we point out that the proposed algorithm could accommodate different goals, such as, e.g., maximum throughput or proportional fairness, by simply replacing the function that gives the global system performance, η , by another function that reflects performance according to the objective pursued.

C. Control overhead

We conclude the analysis with the evaluation of the control overhead introduced by $H_2(IC)_2$. To this aim, we identify two different *interfaces*: one between central coordinator and base stations, namely I_C , and one between distinct base stations, namely I_B . They may be both implemented using, e.g., the LTE X2 interface [2].

In the centralized solution, the central coordinator requires message exchanges over I_C only. In particular, per each pair (*user, base station*), it requires the transmission of an average channel quality indicator (e.g., the *RSRP* value in the LTE-Advanced networks [25]) which can be encoded in double precision floating point format, e.g., 64 bits. Then, the controller issues a scheduling pattern (a string of T bits) per each base station.

In the $H_2(IC)_2$ mechanism, the controller requires to receive the average user rate η_i per base station over I_C at the end of each game Γ , consisting in a binary string of fixed length (e.g., 64 bits for a double precision floating point number). Regarding the interface I_B between different base stations, $H_2(IC)_2$ needs a sequential exchange of ABSF scheduling patterns (strings of T bits) during the interference coordination game Γ , until the game reaches a convergence state or the convergence deadline expires.

We can therefore summarize the total load in terms of bits for each interface as reported in Table III. In the table, k is the number of rounds the interference coordination game plays before reaching the convergence, and $|\mathcal{U}| = \sum_i |\mathcal{U}_i|$ is the total number of users in the system. We can easily observe that the

overhead of $H_2(IC)_2$ is lower than that of the centralized mechanisms when the following inequality holds:

$$|U| > 1 + \frac{T}{64}(k-1) \cong \frac{T|\mathcal{N}|^2}{64}, \quad (4)$$

where we have considered that the number of rounds k in the worst case is a function of $|\mathcal{N}|$ (i.e., at most $k = |\mathcal{N}|^2$ iterations are enough to converge, when convergence exists, as proven mathematically in [29] and empirically shown in Section V) and both T and $|\mathcal{N}|$ are (much) greater than 1. Therefore, our semi-distributed approach is convenient as soon as the number of users exceeds a threshold that depends on T and $|\mathcal{N}|$ (i.e., the threshold is $O(T|\mathcal{N}|^2)$). For example, in an (sub-)urban environment with $T = 70$ and $|\mathcal{N}| = 7$, as in our simulations described later, $H_2(IC)_2$ results convenient with as few as 54 users or more, while in a dense-urban environment with $|\mathcal{N}| = 30$, our approach exhibits a practical implementation starting with ~ 1000 users in the entire network. Those values are pretty low, revealing how *our semi-distributed approach drastically reduces the signaling overhead for existing cellular network size.*

V. PERFORMANCE EVALUATION

In this section, we use numerical simulations to show that our proposal performs near optimally and boosts achievable rates in the whole network, not just for topologically disadvantaged users. All simulations are carried out by means of MATLAB® with all parameters summarized in Table IV. The average quality of the user channel is computed as function of the distance from the base station (according to the propagation model suggested by 3GPP specifications, Table A.2.1.1-3 of TR.25.814 v7.1.0), and Rayleigh fading is considered. Based on user channel qualities, each simulated base station solves the local optimization problem by means of a remote call to a commercial solver, i.e., IBM CPLEX OPL®. Additionally, we show that game Γ quickly approaches its Nash equilibrium, which enables $H_2(IC)_2$ to easily follow changes occurring dynamically in the network.

A. Benchmarks

We benchmark $H_2(IC)_2$ against the optimal solution, obtained by solving Problem CENTRAL by means of an ILP solver. Additionally, we compare $H_2(IC)_2$ to the case of uncontrolled base stations using the same frequencies (No ICIC) and to a traditional frequency reuse 3 scheme, in which the available band is split into three orthogonal sub-bands. For the sake of completeness, we also compare $H_2(IC)_2$ with two existing approaches fully based on a power control schemes, showing how $H_2(IC)_2$ can achieve high network performance at a bargain price of complexity. In the first scheme, namely Utility-Based Power Control (UBPC) [8], base stations are allocated in all available TTIs by tuning properly the transmitted power to reduce interference. The algorithm suggested in [8] maximizes the user net utility by ensuring that the signal-to-noise-ratio of each transmission is greater than a minimum threshold γ_i (in our simulations we assume γ_i as the minimum MCS with nonzero rate). However, UBPC allows for multiple transmissions to different users in

TABLE IV: List of Parameters for the LTE-A wireless scenarios used in the experiments

$ \mathcal{N} $	Number of Base Stations	7
$ U_i $	Number of UEs per Base Station	10
T	ABSF Pattern Length	70 TTIs
BW	Spectrum Bandwidth	20 MHz
P	Transmitting Power	1 Watt
ISD	Inter-Site Distance	200 m
N_0	Background Noise	1.085×10^{-14}

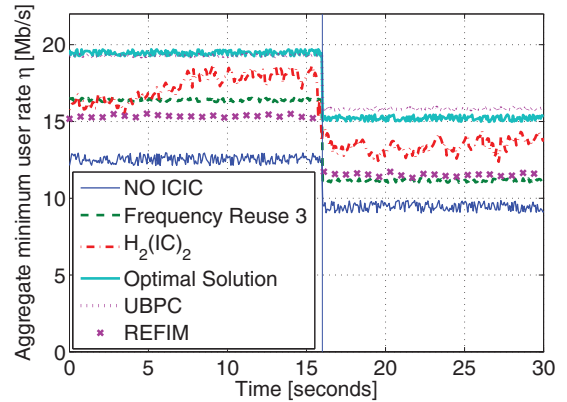


Fig. 2: Dynamic behaviour of $H_2(IC)_2$ applied to a changing scenario. On the left side, the scenario has $|\mathcal{N}| = 7$ base stations, $|U_i| = 10$ users and $T = 70$ TTIs. On the right side, the number of users is increased up to $|U_i| = 20$ users. $H_2(IC)_2$ quickly adapts to a network change while keeping high the accuracy of the solution (w.r.t. Optimal solution and no ICIC).

the same cell, which is not doable in schedule-based cellular networks. Therefore, to force the scheduling of a single user per cell on a per-TTI basis, we simply modify the original UBPC algorithm by setting the interference to infinite when two or more users from the same cell are scheduled. We refer the reader to [8] for more details on UBPC. While UBPC provides a rigorous centralized solution for the power allocation problem at the expense of a huge amount of information exchanged, a second power control scheme recently developed, namely REFERENCE based Interference Management (REFIM) [9], proposes a low-complex distributed scheme by exploiting the notion of reference user. The authors of [9] aim at simplifying the analysis of the impact of neighbouring cells by replacing all of them with a single virtual user, selected as the user with the worst channel condition belonging to the surrounding cells. This abstraction leads to a drastic reduction of the control signal overhead resulting in a practical implementation of the power control solution, which exhibits a conservative behaviour.

B. Utility and fairness performance

We start by evaluating the system utility η , which, according to the formulation of Problem CENTRAL, is the sum of minimum user rates experienced in the network. Fig. 2 shows η averaged over the time horizon of $T = 70$ TTIs⁵ for the case of $|\mathcal{N}| = 7$ base stations and 10 users per base station.

Due to the adaptive nature of the algorithm, $H_2(IC)_2$

⁵Typical values for the ABSF pattern length are between 60 and 80. We use 70, which yields a round number for M_i^* in our simulated scenario consisting of 7 cells.

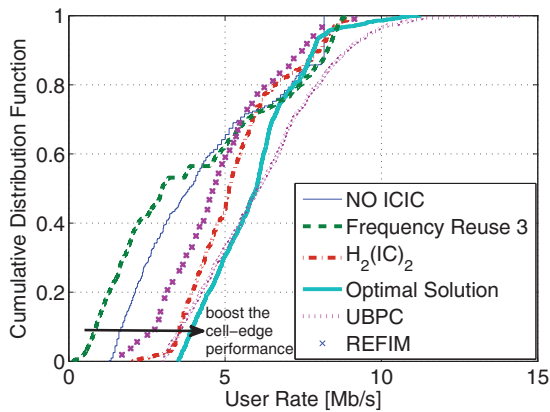


Fig. 3: CDF of average user rates with 7 base stations and 10 users per base station. The time horizon is set to $T = 70$ TTIs.

shows a dynamic behavior. Specifically, it takes a few seconds for $H_2(IC)_2$ to reach its stable operating point, after which it follows quite fast the evolution of channel and traffic conditions. In particular, at time $t = 16$ s, the number of users in the network doubles abruptly, but it takes only a fraction of a second for $H_2(IC)_2$ to adapt. In general, $H_2(IC)_2$ largely outperforms the No ICIC scheme and achieves significant gain over frequency reuse 3. Indeed, $H_2(IC)_2$ halves the distance between the optimal performance and the one of frequency reuse 3. Notably, after the initial adaptation period, the utility achieved by $H_2(IC)_2$ lies within 85% and 90% of the one achieved with the optimal solution for Problem CENTRAL. REFIM and UBPC results show the real potentials of power control schemes. UBPC can even go slightly beyond the performance of the optimal solution without power control, although it requires higher complexity in terms both of execution and device hardware. REFIM, notwithstanding a low-complexity scheme, shows lower performance with respect to frequency reuse 3 for a particular set of user populations due to the conservative assumption taken on interfering cells. Therefore, our approach $H_2(IC)_2$ perfectly lies in between an impermissible efficient power control scheme and a practical doable distributed power control solution.

Besides utility η , we want to evaluate the fairness achieved by the different schemes. To this aim, Fig. 4 presents the CDF of achieved user rates (averaged over the time horizon T). The figure clearly shows that the optimal solution, $H_2(IC)_2$ and UBPC behave similarly and exhibit two main advantages: (i) they achieve user rates in a compact interval of possible values (which is symptom of fairness according to Jain fairness definition), and (ii) with high probability, they guarantee a minimum rate which is several times higher than the one guaranteed by No ICIC or frequency reuse 3 (which is symptom of max-min fairness). In addition, REFIM shows a similar behaviour to the optimal solution in terms of fairness, even though its curve stays on the left side of the graph due to the critical user rates experienced by the users. For instance, with 95% probability, $H_2(IC)_2$ guarantees 3.7 Mbps per user, REFIM guarantees 2.4 Mbps per user, while No ICIC only guarantees 0.9 Mbps. Jain fairness achieved under the different schemes under evaluation is depicted in Fig. 4 as a function of the number of users per base station. As expected,

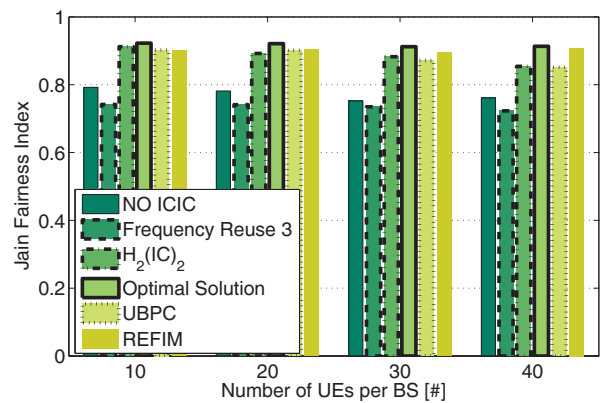


Fig. 4: Jain fairness indexes achieved with 7 base stations and a variable number of users per base station.

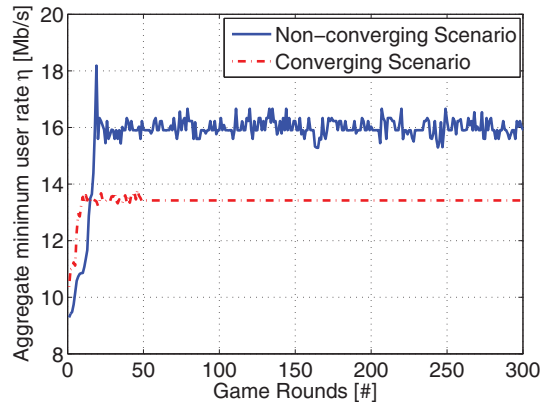


Fig. 5: Game convergence behavior considering two different cases with $|\mathcal{N}| = 7$ base stations.

REFIM presents a stable behaviour over different values of user population, due to its strong correlation with worst users. Also in this case, $H_2(IC)_2$ achieves near-optimal results, and it outperforms UBPC.

C. Game convergence speed

A key feature of $H_2(IC)_2$ is its ability to adapt quickly to network changes. Such a feature relies on quick ABSF pattern computation, which follows the rule of the Interference Coordination Game Γ . The game evolves over time as illustrated in Fig. 5. In this example two different cases are considered: the dashed line represents a case of convergence, while the solid line is for a rare case in which the game does not converge to a Nash equilibrium point. In both cases, 7 base stations are considered, and the TTI bounds M_i are fixed.

In case of convergence, which occurs in about $|\mathcal{N}|^2$ rounds, it is clear that a few game rounds suffice to approximate the performance achieved at the Nash equilibrium with an error smaller than 3%. Notably, also in case the game fails to converge, after a few rounds the utility starts fluctuating around a stable value, with small oscillations (about $\pm 5\%$).

Fig. 6 illustrates the CDF of the number of rounds needed to converge for a few different user populations (note that we use the value 1 to indicate that the game did not converge). The figure shows how the majority of the games converge much before $|\mathcal{N}|^2$ rounds (vertical line in the figure), and very few cases do not converge at all. We have observed very

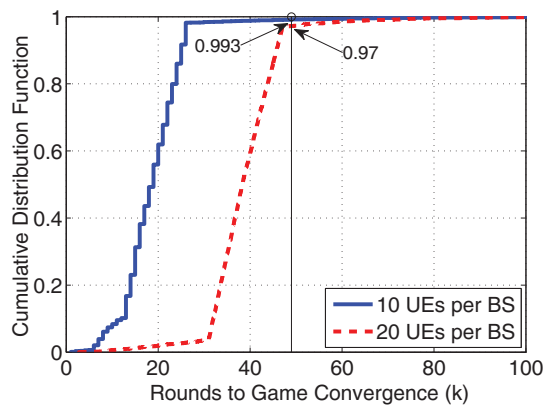


Fig. 6: CDF of number of rounds needed for game convergence with 7 base stations and different user populations.

similar behavior for the majority of the cases analyzed in our experiments, so we conclude that reasonably high utilities can be achieved by stopping the game after a number of rounds comprised between $|\mathcal{N}|$ and $|\mathcal{N}|^2$.

Overall, our results show that $H_2(IC)_2$ not only achieves near-optimal results according to the definition of utility given in the formulation of Problem CENTRAL, but also achieves high levels of max-min and Jain fairness, and significantly boosts average rates in the entire cellular network.

VI. CONCLUSIONS

We proposed $H_2(IC)_2$, a practical ICIC scheme leveraging the ABSF paradigm in LTE-A networks. $H_2(IC)_2$ uses a two-tier semi-distributed architecture, which allows base stations to jointly compute ABSF patterns and operate the local scheduling. The behavior of the base stations can be mapped onto the one of the players of a distributed game, for which we proved convergence. The architecture is semi-distributed because it requires the presence of a central coordinator to drive the system to the best achievable network utility by controlling the number of TTIs to blank at each base station. Due to the simplicity of our approach and its limited control overhead, at best of our knowledge, this is the first attempt to move towards a practical, efficient, scalable and adaptive implementation of ABSF in real networks. Indeed, our numerical results show that $H_2(IC)_2$ achieves near-optimal results with respect to a centralized omniscient network scheduler, and achieves performance levels similar to advanced schemes using complex power control approaches.

REFERENCES

- [1] Third Generation Partnership Project (3GPP), "Requirements for further advancements for Evolved Universal Terrestrial Radio Access (E-UTRA)," 3GPP TR 36.913 v 10.0.0, Mar. 2011.
- [2] —, "Evolved Universal Terrestrial Radio Access Network (E-UTRAN); X2 application protocol (X2AP)," 3GPP, Tech. Spec. Group Radio Access Network.
- [3] X. Hou, E. Bjornson, C. Yang, and M. Bengtsson, "Cell-grouping based distributed beamforming and scheduling for multi-cell cooperative transmission," in *IEEE PIMRC 2011*, Sept 2011, pp. 1929–1933.
- [4] R. Irmer, H. Droste, P. Marsch, M. Grieger, G. Fettweis, S. Brueck, H. Mayer, L. Thiele, and V. Jungnickel, "Coordinated multipoint: Concepts, performance, and field trial results," *IEEE Communications Magazine*, vol. 49, no. 2, pp. 102–111, 2011.
- [5] S. Singh and J. Andrews, "Joint resource partitioning and offloading in heterogeneous cellular networks," *IEEE Trans. on Wireless Communications*, vol. 13, no. 2, pp. 888–901, February 2014.

- [6] Y. Jin and L. Qiu, "Joint user association and interference coordination in heterogeneous cellular networks," *IEEE Communications Letters*, vol. 17, no. 12, pp. 2296–2299, December 2013.
- [7] V. Sciancalepore, V. Mancuso, and A. Banchs, "BASICS: Scheduling base stations to mitigate interferences in cellular networks," in *IEEE WoWMoM 2013*, June 2013.
- [8] M. Xiao, N. Shroff, and E. K. P. Chong, "A utility-based power-control scheme in wireless cellular systems," *Networking, IEEE/ACM Transactions on*, vol. 11, no. 2, pp. 210–221, Apr 2003.
- [9] K. Son, S. Lee, Y. Yi, and S. Chong, "REFIM: A practical interference management in heterogeneous wireless access networks," *Journal on Selected Areas in Communications*, vol. 29, no. 6, pp. 1260–1272, 2011.
- [10] S. Deb, P. Monogioudis, J. Miernik, and J. Seymour, "Algorithms for enhanced inter-cell interference coordination (eICIC) in LTE hetnets," *IEEE/ACM Trans. on Networking*, vol. 22, no. 1, pp. 137–150, Feb 2014.
- [11] T. Bonald, S. Borst, and A. Proutiere, "Inter-cell coordination in wireless data networks," *Wiley European Transactions on Telecommunications*, vol. 17, no. 3, pp. 303–312, 2006.
- [12] B. Rengarajan and G. De Veciana, "Architecture and abstractions for environment and traffic-aware system-level coordination of wireless networks," *IEEE/ACM Trans. on Networking*, pp. 721–734, 2011.
- [13] R. Combes, Z. Altman, M. Haddad, and E. Altman, "Self-optimizing strategies for interference coordination in OFDMA networks," in *IEEE ICC 2011*, June 2011, pp. 1–5.
- [14] A. Gjendemsj, D. Gesbert, G. Oien, and S. Kiani, "Binary power control for sum rate maximization over multiple interfering links," *IEEE Trans. on Wireless Communications*, vol. 7, no. 8, pp. 3164–3173, 2008.
- [15] M. Dirani and Z. Altman, "A cooperative reinforcement learning approach for inter-cell interference coordination in OFDMA cellular networks," in *IEEE WiOpt 2010*, May 2010, pp. 170–176.
- [16] Y. Wang and K. Pedersen, "Time and power domain interference management for lte networks with macro-cells and HeNBs," in *IEEE VTC 2011*, 2011, pp. 1–6.
- [17] A. Ghosh, N. Mangalvedhe, R. Ratasuk, B. Mondal, M. Cudak, E. Vitsosky, T. Thomas, J. Andrews, P. Xia, and H. e. a. Jo, "Heterogeneous cellular networks: From theory to practice," *IEEE Communications Magazine*, vol. 50, no. 6, pp. 54–64, 2012.
- [18] M. Cierny, H. Wang, R. Wichman, Z. Ding, and C. Wijting, "On number of almost blank subframes in heterogeneous cellular networks," *IEEE Trans. on Wireless Comm.*, vol. 12, no. 10, pp. 5061–5073, Oct. 2013.
- [19] L. Jiang and M. Lei, "Resource allocation for eICIC scheme in heterogeneous networks," in *IEEE PIMRC 2012*, Sept 2012, pp. 448–453.
- [20] M. Kamel and K. M. Elsayed, "Performance evaluation of a coordinated time-domain eICIC framework based on ABSF in heterogeneous LTE-Advanced networks," in *IEEE GLOBECOM 2012*, 2012, pp. 5326–5331.
- [21] M. Kamel and K. Elsayed, "ABSF offsetting and optimal resource partitioning for eICIC in LTE-Advanced: Proposal and analysis using a nash bargaining approach," in *IEEE ICC 2013*, pp. 6240–6244.
- [22] J. Yoon, M. Arslan, K. Sundaresan, S. Krishnamurthy, and S. Banerjee, "Self-organizing resource management framework in OFDMA femto-cells," *To appear on IEEE Trans. on Mobile Computing*, 2014.
- [23] K. Son, Y. Yi, and S. Chong, "Utility-optimal multi-pattern reuse in multi-cell networks," *IEEE Trans. on Wireless Communications*, vol. 10, no. 1, pp. 142–153, 2011.
- [24] J. Ellenbeck, C. Hartmann, and L. Berlemann, "Decentralized inter-cell interference coordination by autonomous spectral reuse decisions," in *IEEE EW 2008*, 2008, pp. 1–7.
- [25] Third Generation Partnership Project (3GPP), "Physical Layer Measurements," 3GPP TS 36.214 v 11.1.0, Dec. 2012.
- [26] A. Capone, G. Carello, I. Filippini, S. Gualandi, and F. Malucelli, "Routing, scheduling and channel assignment in wireless mesh networks: Optimization models and algorithms," *Ad Hoc Networks*, vol. 8, no. 6, pp. 545–563, Aug. 2010.
- [27] O. Goussevskaia, R. Wattenhofer, M. M. Haldrsson, and E. Welzl, "Capacity of arbitrary wireless networks," in *INFOCOM'09*, 2009.
- [28] D. Bertsekas and R. Gallager, *Data Networks (II Ed.)*. Prentice-Hall.
- [29] T. Harks, M. Hoefler, M. Klimm, and A. Skopalik, "Computing pure nash and strong equilibria in bottleneck congestion games," *Springer Mathematical Programming*, vol. 141, no. 1-2, pp. 193–215, 2013.
- [30] C. Liu and E. Modiano, "On the performance of additive increase multiplicative decrease (AIMD) protocols in hybrid space-terrestrial networks," *Comput. Netw.*, vol. 47, no. 5, pp. 661–678, Apr. 2005.