

BASICS: Scheduling Base Stations to Mitigate Interferences in Cellular Networks

Vincenzo Sciancalepore
Institute IMDEA Networks, and
University Carlos III of Madrid
Madrid, Spain

Vincenzo Mancuso
Institute IMDEA Networks
Madrid, Spain

Albert Banchs
Institute IMDEA Networks, and
University Carlos III of Madrid
Madrid, Spain

Abstract—The increasing demand for higher data rates in cellular network results in increasing network density. As a consequence, inter-cell interference is becoming the most serious obstacle towards spectral efficiency. Therefore, considering that radio resources are limited and expensive, new techniques are required for efficient radio resource allocation in next generation cellular networks. In this paper, we propose a pure frequency reuse 1 scheme based on base station scheduling rather than the commonly adopted user scheduling. In particular, we formulate a base station scheduling problem to determine which base stations can be scheduled to simultaneously transmit, without causing excessive interference to any user of any of the scheduled base stations. We show that finding the optimal base station scheduling is NP-hard, and formulate the BASICS (BAse Station Inter-Cell Scheduling) algorithm, a novel heuristic to approximate the optimal solution at low complexity cost. The proposed algorithm is in line with the ABSF (almost blank sub-frame) technique recently standardized at the 3GPP. By means of numerical and packet-level simulations, we prove the effectiveness and superiority of BASICS as compared to the state of the art of inter-cell interference mitigation schemes.

I. INTRODUCTION

Fourth-generation (4G) broadband wireless technologies such as 3GPP Long Term Evolution Advanced (LTE-A) [1] and IEEE 802.16m [2] are the main solutions addressing the foreseen future capacity needs. Since the number of mobile users is growing swiftly and applications require every day higher data rates, the new generation of mobile systems has to design new solutions to satisfy the demand of mobile users. Therefore, one of the most important design goals is achieving high spectral efficiency. In this context, it has been shown that frequency reuse 1 can provide substantial improvements in terms of efficient utilization of the scarce and expensive wireless resources. This implies that neighboring base stations (BSs) should be allowed to transmit on all available time-frequency resource blocks simultaneously, thus causing strong interference to each other's users. This contrasts with interference mitigation and/or cancellation techniques that have been used for many years in the past, which basically exploited orthogonality of frequency and/or spatial resources [3]. More recently, advanced solutions have been designed which actively reduce or cancel interference when orthogonality cannot be guaranteed [4], [5].

In this paper, we focus on the coordination among adjacent base stations to achieve high spectral efficiency.

The research leading to these results received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement N° 257263 (FLAVIA).

Differently from most of the work available in the literature, we tackle the problem of inter-cell interference mitigation from the perspective of scheduling *base stations* rather than *users*. In particular, we propose to coordinate base station downlink activities in order to mitigate the interference caused to neighboring cells. To do this, we propose a method to map base stations' activities onto subframes with regular patterns. We aim to mitigate interference by limiting the activity of a given base station to some subframes while forcing it to remain silent in the other subframes. Noticeably, our work is suitable for driving sub-frame blanking decisions in self organising networks, e.g., as per the ABSF (almost blank sub-frame) technique recently defined by 3GPP [6]. Our simulation results show that scheduling the base station activity achieves significant gain in terms of spectral efficiency.

A key feature of our proposal is that it incurs a very reduced signaling load between base stations. Indeed, we propose to coordinate base station downlink activities in order to limit the interference caused to *any* possible user in the system under *any* possible user scheduling decision taken by the base stations. Therefore, base stations do not need to exchange information with a per-user granularity but rather on a much coarser basis. Another key feature of our design is that we decouple the problem of mitigating downlink inter-cell interference from the problem of optimizing the user scheduling. While the focus of this paper is on the first problem, our proposal can easily be combined with existing user scheduling schemes to further improve spectral efficiency (as we show in the performance evaluation section).

The key contributions of this work are as follows: (*i*) we formulate a novel base station scheduling problem and show that it is NP-hard in strong sense; (*ii*) we design an algorithm, called BASICS (BAse Station Inter-Cell Scheduling), that runs in polynomial time and scales with the number of users; (*iii*) we show that BASICS not only achieves better throughput performance with respect to state of the art schedulers, but also significantly improves fairness among users.

The rest of the paper is structured as follows: Section II compares different solutions addressed in other works. Section III provides the mathematical formulation of the base station scheduling problem. Afterwards, Section IV presents our heuristic algorithm, and comments on the differences with other well-known algorithms proposed to solve similar problems. Section V introduces the tools used in this paper to evaluate the effectiveness of our solution. A complete performance evaluation study is provided in the Section VI. Finally, Section VII concludes the paper.

II. RELATED WORK

The authors of [3] provide an overview of techniques that can be exploited to mitigate inter-cell interference in OFDM-based networks. Interestingly, they do not identify base station scheduling as a possible tool to reduce interference, and limit their discussion to beamforming, coding and decoding techniques, opportunistic spectrum access, interference cancellation, power control and (fractional) frequency reuse.

Most of the work available in the literature focuses on *user scheduling*, in terms of beamforming, cooperative transmission (CoMP), and power allocation. For instance, random beamforming has been proposed to reduce the need for BS-to-BS CQI information exchange [4]. It enhances the per-cell throughput at limited cost, at least in small networks. Fast distributed beamforming in multi-cell environments was also proposed in [5], in which scheduling is performed in two steps: first each base station chooses the proper beamforming that minimizes intercell interference and then a particular user is scheduled in each cell. However, none of these two works address fairness. CoMP shows a very good gain by mitigating interference exploiting cooperation between sector transmitters or different base stations [7]. CoMP proposals range from complex distributed MIMO solutions to advanced beamforming mechanisms. E.g., Multi-Cell Joint Transmission [8] proposes to share the same data to transmit across multiple base stations, while Coordinated Beamforming Scheduling [9] proposes a method to choose transmission beam patterns in coordination between base stations. In both cases, a large backhaul capacity is required for inter-base station communication, over both data and control planes, which realistically prevents implementation in real systems. In contrast, the implementation of BASICS only requires control plane operations.

As for the work focusing on *base stations* instead of users, proposed solutions are available for frequency reuse and fractional frequency reuse schemes where the entire available set of frequencies is divided a priori and assigned to adjacent cells, or portions of cells, to avoid strong interference between neighboring cells. The disadvantage of such schemes lies in their scarce flexibility to adapt to changing cell load conditions. A scheme proposed in [10] suggests to divide the cell into two zones, namely edge and center, and assigns users to each zone dynamically based on their channel state information; however, the gain is limited due to the low number of available frequency bandwidths. Although originally proposed for WiMAX systems, RADION represents another valid approach to intercell interference mitigation [11]. RADION is a distributed resource management framework that manages interference across femtocells and enables femtocells to opportunistically find the available resources in a distributed manner, by performing three different actions: client categorization, resource decoupling and two-phase adaptation and allocation. A similar and interesting approach is presented in [12] where a management system, called FERMI, is introduced for OFDMA-based femtocell networks. It performs a clients classification in order to identify which clients need resource isolation and those that require just link adaptation,

then it incorporates a frame zoning structure that supports the coexistence of clients from both categories. Afterwards, the system allocates orthogonal sub-channels of the OFDMA spectrum in a fair manner. The authors of [13] use dynamic programming to optimize the transmission probability of base stations in a TDMA network. Although their approach might recall ours, in terms of base station scheduling, their analysis only yields a probabilistic scheduling model for base stations and cannot be realistically applied to existing cellular technologies. Additionally, their solution can only be used in case of networks consisting of two base stations, or with homogeneous topologies. In contrast, our algorithm is suitable for all topologies.

Recently, 3GPP has standardized the ABSF (almost blank sub-frame) technique [6], which allows base stations to blank data transmissions over one or several subframes. However, the standard does not specify the algorithm to decide which base stations remain silent and which ones do not in a given subframe. The algorithm that we propose in this paper addresses precisely this question, and can be used to drive the ABSF mechanism.

III. PROBLEM DESCRIPTION

In this section, we describe the system that we focus on in this paper, and formulate our base station scheduling problem for this system. We then prove that this problem is NP-hard by mapping it onto a well-known NP-hard optimization problem, namely the Multidimensional Vector Bin-Packing Problem. Finally, we derive some bounds for the problem's solution.

A. System Model

We consider a multicellular LTE-like environment with N base stations and U mobile users. We address only downlink transmissions, for which no power control is adopted, as in the majority of state of the art proposals. Each base station schedules its users across subframes, as specified in LTE systems [14], each subframe lasting $1ms$. The duplexing scheme adopted is FDD, with 20MHz bandwidth for each transmission direction. Unless otherwise specified, all base stations use the same frequencies. Users associate to the base station from which they receive the strongest signal, and transmission rates are selected, in each subframe, according to the Signal-plus-Noise Interference Ratio (SINR), see Table I. The SINR for a certain user $u \in \{1, \dots, U\}$ is defined as follows:

$$\text{SINR} := \frac{S_b^u}{N_0 + \sum_{j \neq b} I_j^u},$$

where N_0 is the background noise, S_b^u is the useful signal received by the current user u from the serving base station b (hereafter defined as S^u) and I_j^u is the interference sensed by the user u from any other base station j in the system, when that base station is scheduled. Signal and interferences received by each user are affected by Rayleigh fading, and user channels are assumed to be independent.

In our system, we focus on mitigating interference by deciding whether a base station can transmit during a given subframe. We refer to this decision as *base station scheduling*, to be distinguished from legacy *user scheduling* which occurs at each base station when it is allowed to transmit.

B. Base Station Scheduling Optimization Problem

We aim at guaranteeing a minimum SINR to every user in the system by allocating in each subframe a subset of the available base stations while minimizing the number of subframes needed in order to schedule the complete set of base stations. We will discuss in Section VI-A how to select that minimum SINR. Meanwhile, here we focus on the first part of the problem: decide whether a base station can transmit in a given subframe.

The goal of our proposal is to schedule base station transmissions in each subframe so that the SINR is greater than a threshold Th for every user u in the network (i.e., for any user that can receive a transmission from a scheduled base station):

$$\frac{S_b^u}{N_0 + \sum_{j \neq b} I_j^u} \geq \text{Th}. \quad (1)$$

For convenience of notation, we now call I_b^u the signal received by user u from its BS, i.e., $I_b^u = S_b^u = S^u$, therefore Eq. (1) can be rewritten as follows:

$$\sum_{j=1}^n I_j^u \leq \frac{S^u}{\text{Th}} - N_0 + S^u \triangleq \text{Th}^u. \quad (2)$$

With the above, the problem of minimizing the total number Z of subframes used to schedule once all N base stations in the system, for a given minimum SINR (or threshold Th), is formulated as follows:

$$\begin{cases} \text{minimize} & Z = \text{number of subframes needed to} \\ & \text{allocate all base stations once,} \\ \text{subject to} & \sum_j I_j^u x_{ij} \leq \text{Th}^u, u \in \{1, \dots, U\}, i \in \{1, \dots, Z\}, \\ & \sum_i x_{ij} = 1, \quad j \in \{1, \dots, N\}, \\ & x_{ij} \in \{0, 1\}, \quad i, j \in \{1, \dots, N\}, \end{cases} \quad (3)$$

where

$$x_{ij} = \begin{cases} 1, & \text{if base station } j \text{ is scheduled into subframe } i, \\ 0, & \text{otherwise.} \end{cases}$$

Note that Z is the analogous in the time domain of the frequency reuse factor. However, our approach has two main advantages over frequency reuse schemes. First, we need only one frequency allocated to the system, which is less expensive than using frequency reuse. Second, differently from the frequency reuse factor, which is static, Z can vary from time to time when network conditions change (i.e., with users' arrival or departure).

Theorem 1: Minimizing Z as defined in Problem (3) is NP-hard in strong sense.

Proof: The number of subframes needed to allocate all base stations at least once is upper bounded by the number of base stations, i.e., $Z \leq N$. Consider now the problem of scheduling each base station exactly once in Z consecutive subframes. If $Z < N$, then $N - Z$ subframes are left empty. Let us define a set of N binary variables y_i , which indicate whether a subframe $i = 1, \dots, N$ is used or empty:

$$y_i = \begin{cases} 1, & \text{if subframe } i = 1, \dots, N \text{ is used,} \\ 0, & \text{otherwise.} \end{cases}$$

Note that the concept of *empty subframe* is only an abstraction to simplify the description of the problem. In fact, $Z < N$ means that the scheduling of base stations over Z subframes is repeated cyclically, with period Z subframes. With the above notation, we can re-write Problem (3) as follows:

$$\begin{cases} \text{minimize} & Z = \sum_i y_i, \\ \text{subject to} & \sum_j I_j^u x_{ij} \leq \text{Th}^u y_i, u \in \{1, \dots, U\}, i \in \{1, \dots, N\}, \\ & \sum_i x_{ij} = 1, \quad j \in \{1, \dots, N\}, \\ & y_i \in \{0, 1\}, \quad i \in \{1, \dots, N\}, \\ & x_{ij} \in \{0, 1\}, \quad i, j \in \{1, \dots, N\}. \end{cases} \quad (4)$$

The above is the formulation of a k -dimensional vector bin-packing problem (kD-VBP) with $k = U$ being the number of mobile users in the system [15].

Our problem is therefore equivalent to a k -dimensional vector bin-packing problem, in which knapsacks represent subframes, items to be allocated are base stations, and the number of dimensions is given by the number of users, each imposing a constraint on its SINR as expressed in Eq. (2). Since our problem has been mapped onto kD-VBP, it can thus be classified as an NP-hard problem in strong sense. ■

C. Lower Bound

The value Z in Problem (3) determines the portion of time during which a base station is prevented from transmitting (i.e., each base station is scheduled with frequency $1/Z$). Therefore, the lower bound for Z in our problem represents the highest scheduling frequency that can be associated to base stations in the system with SINR not lower than Th for any of the users. Thus, in order to have high efficiency in the utilization of resources, our goal is to design an algorithm that finds the smallest possible value for Z .

In the following, we obtain a lower bound for Z , which bounds the best possible performance that we can achieve. This bound provides a benchmark against which we can evaluate the performance of our solution, as we do in the performance evaluation section.

Theorem 2: The lower bound $L \leq Z$ for Problem (3) with k users distributed over N base stations, and a guaranteed $\text{SINR} \geq \text{Th}$ for all users, is given by the following equation:

$$L = \max_u \left(\left\lceil \sum_{j=1}^N \frac{I_j^u}{\text{Th}^u} \right\rceil \right). \quad (5)$$

Proof: Our proof follows the same approach used in [16]. First, we recall that in vector bin-packing problems items cannot be *rotated*, i.e., the constraints are defined on a per-dimension basis, and dimensions cannot be rearranged. In our case, a dimension represents the interference caused to a given user, which justifies why dimensions cannot be rearranged. In particular, the minimum number of subframes (bins) needed to accommodate all the base stations (items) in such a way that the max interference (constraint) on the u th dimension is not violated is given by the ratio between the sum of all interferences caused by N base stations in the u th dimension, i.e., for the u th user, divided by the threshold Th^u , which represents

the capacity of the bin in the uth dimension. Of course, only integer numbers are allowed, hence we need $\left\lceil \sum_{j=1}^N I_j^u / Th^u \right\rceil$ bins to satisfy the constraint in the uth dimension. Since all dimensions are independent, the result follows. ■

IV. ALGORITHM DESIGN

To solve the problem formulated in Section III-B, we next propose a heuristic consisting in a greedy algorithm for the mapping of base stations to subframes. The algorithm is designed to dynamically mitigate inter-cell interference caused to any possible user in the system under any possible user scheduling decision taken by the base stations. As a result, our algorithm is *user-scheduling-agnostic* and does not require coordinated scheduling among base stations.

We propose a new heuristic rather than using existing heuristics for two main reasons. First, existing heuristics for multidimensional vector bin-packing problems are simple extensions of solutions designed for the one-dimension problem. Second, existing heuristics do not take into account the nature of the dimensions that describe the items to be allocated. In particular, they assume that the size of an object is the same in any of the possible combinations of items in a bin. In contrast, in our case, the size of an object is the interference caused to mobile users *belonging to the scheduled base stations only*. Therefore, the weight associated to a base station (i.e., its *size*) changes any time a base station is removed from the list of candidate transmitters (e.g., since it is allotted to a subframe).

In the following, we first briefly discuss existing algorithms for solving multidimensional vector bin-packing problems, then we present our novel solution and highlight the difference with existing proposals. Later, in Section VI, we prove, by means of empirical results, that our approach outperforms existing algorithms.

A. State of the art algorithms for bin-packing problems

The most commonly adopted algorithms for solving the bin-packing problem belongs to the family of FFD-based algorithms. The First-Fit Decreasing Algorithm (FFD) was proposed to solve the one-dimensional bin-packing problem [15]. With FFD, items are sorted by size, in decreasing order, and a number of empty bins—equal to the total number of items—is set. Then, items are inserted sequentially from the largest to the smallest in the first available bin with enough capacity left.

To cope with the case of multidimensional vector bin-packing problems, various greedy FFD extensions have been proposed in the literature [17]. Available heuristics collapse all dimensions into one, and then apply the FFD algorithm proposed for the resulting one-dimensional version of the problem. The name of each algorithm depends on how the dimensions are collapsed. When all dimensions are multiplied in order to get one unique *monodimensional size* for each item, the algorithm is called *FFDProd*, whereas the algorithm *FFDSum* uses a weighted sum of the original dimensions. Other algorithms such as *FFDAvgSum* or *FFDExpSum* use similar approaches to *FFDSum* [17].

As described in [18], the above algorithms can be classified as FFD item-centric since all items are allocated until there

are no items left to be placed. Another group of algorithms are classified as FFD bin-centric. The latter are algorithms which start with a single bin, and a new bin is initialized when there are no more items which can fit the previously used bins. As proved by empirical evaluations in [18], bin-centric algorithms (such as Dot-Product and Norm-based Greedy) outperform item-centric algorithms and they can sometimes reduce the number of required bins by up to 10%.

A common assumption of FFD-based algorithms is that the dimensions of an item do not change. In contrast, in the problem described in (3), dimensions (i.e., interferences caused by a base station transmission) *do* change with the set of items (i.e., base stations) that are included (i.e., scheduled) in the same bin (i.e., subframe). Indeed, the set of base stations scheduled in a subframe affects the set of mobile users that can receive interference, and therefore affects the number of dimensions of the problem in a given algorithm iteration. Therefore, legacy FFD-based approaches are not suitable for solving our problem, so we propose a novel approach, as described in the reminder of this Section.

B. BASICS

Interferences sensed by users play the role of dimensions in the optimization problem (3). Thus, we propose BASICS (BAsic Station Scheduling Inter-Cell Scheduling), a sum-based algorithm which solves kD-VBP problems by collapsing all problem dimensions (i.e., the interferences to different users) into one unique value. This value is computed for each base station, and consists in the total interference caused by the base station to users belonging to other scheduled base stations.

However, in our problem, the size of the items (base stations) to be allocated into bins (subframes) *changes* at any iteration of the algorithm. In particular, BASICS represents a modification of the *FFDSum* algorithm in which (i) the size of each item to be accommodated changes at each iteration, and (ii) items are accommodated into bins in order, beginning with the smallest one. As in bin-centric approaches, BASICS allocates a new bin only when there is no more room left in the old bins to accommodate the remaining items. Note that existing algorithms for kD-VBP would rather sort items from the largest to the smallest.

The rationale behind our approach is as follows. First, when we start allocating base stations from the least interfering one, we have a chance to schedule together the highest number of not-previously-allocated base stations in the same subframe. This eliminates the highest number of base station candidates for the next subframe allocation. In turn, considering a uniform distribution of users, this procedure eliminates the highest number of users from the set of interfered users in the next iteration of the algorithm. As a result, the cumulative interference over the remaining users, due to the remaining candidate base stations in the next iteration, is likely to be much lower than in the previous iteration. In contrast, if we removed a base station generating less interference, we would have a high probability that that base station interfered fewer users. Thus, removing the least interfering base station would not only bring less benefit to the current subframe, but also we

would not reduce much the impact of that base station in the next subframe allocation (since the set of potential interfered users did not change much). Interestingly, our interference sorting approach is similar to the one presented in [19], which focuses on groups of interfering users.

The details of the BASICS algorithm are presented in Algorithm 1, and described in the following. Initially, the algorithm computes the interference generated by any base station to any user in the system (lines 1 to 6). This computation is performed by means of a simple free space propagation model accounting for the transmission power of the base stations as well as the position of base stations and mobile users. Then, the algorithm checks whether the entire set of base stations can be active in the same subframe, i.e., the entire set of base stations forms the initial base station candidate set. This check is performed by comparing users' SINR thresholds against the SINR experienced when all base stations are active (line 10). At this point, if all SINR constraints are met, then all base stations are allocated into the current subframe, and the algorithm ends. Otherwise, the algorithm computes the overall interference figure due to each base station, and sorts the base stations in decreasing order. The overall interference figure of a base station is computed by summing up all interferences caused by that base station (line 11). Once base stations are sorted, the algorithm removes the most interfering base station from the set of candidate base stations (lines 12-13). The algorithm then re-checks SINR constraints for the subnetwork obtained by removing the most interfering base station and all its users from the original network. The procedure is repeated by removing the most interfering base station (and its users) at each iteration, until all SINR constraints are met. The resulting set of candidate base stations is allocated to the first subframe. Next, the algorithm has to run again for the subnet consisting of the base stations not previously allocated, i.e., the set of base stations that were removed during the first algorithm loop (line 15). The output of the i th algorithm loop is the list of base stations to be scheduled in the i th subframe. When all base stations are allocated, the algorithm ends returning the complete base station scheduling plan.

C. Computational complexity of BASICS

Next, we evaluate the computational complexity of the proposed algorithm. Let N be the number of base stations in the system and U the number of users (i.e., the number of dimensions for a multi-dimensional bin packing problem). Before the first round of checking for SINR constraints, the algorithm computes all signal strengths from N base stations to U users. This operation has computational cost $O(N \cdot U)$ (see lines 1-6 of the Algorithm 1). Then, each sub-frame is inspected in order to check if the candidate base stations meet the SINR constraints for each of their users. For the first sub-frame allocation, the algorithm will perform, for each user, N multiplications, $N - 1$ sums, and 1 division. In the worst case, each round of checking fails, which leads to the elimination of the most interfering base station and its users, which we assume to be uniformly spread over the set of base stations, i.e., we eliminate U/N users at each round. The

Algorithm 1 BASICS: heuristic to allocate base stations into subframes guaranteeing a minimum SINR for each user.

Input and variables

W : set of all base stations in the system
 A : set of base stations not yet allocated
 T_i : candidate set for subframe i
 U : set of all users
 N_0 : background noise
 Th : minimum SINR
 i : subframe index

Initialization

$A \leftarrow W$

Procedure

```

1: for each  $u \in U$  do
2:   Compute  $Th^u$  from  $N_0$ ,  $Th$ , and  $S_u$ 
3:   for each  $j \in W$  do
4:     Compute the signal strength  $I_j^u$  at user  $u$  from base station  $j$ 
5:   end for
6: end for
7:  $i \leftarrow 0$ 
8: while  $|A| > 0$  do
9:    $T_i \leftarrow A$ 
10:  while  $\exists u | \sum_{j \in T_i} I_j^u > Th^u$  do
11:     $\forall j \in T_i, I_j \leftarrow \sum_u I_j^u$ 
12:     $k \leftarrow \arg \max\{I_j\}$ 
13:     $T_i \leftarrow T_i \setminus \{k\}$ 
14:  end while
15:   $A \leftarrow A \setminus T_i$ 
16:   $i \leftarrow i + 1$ 
17: end while

```

largest number of rounds is $N - 1$. Henceforth, the number of complex operations to be performed (i.e., multiplications and divisions) for the first subframe allocation is at most z_1 , as computed in the following equation:

$$z_1 = U \cdot N + U \left(1 - \frac{1}{N}\right)(N-1) + \dots + U \left(1 - \frac{N-1}{N}\right) \cdot 1. \quad (6)$$

The following algorithmic step consists in allocating base stations for the second subframes. In the worst case, there are now $N - 1$ candidate base stations, with $U(1 - 1/N)$ users. Thereby the algorithm performs at most z_2 operations for this subframe:

$$z_2 = U \left(1 - \frac{1}{N}\right)(N-1) + \dots + U \left(1 - \frac{N-1}{N}\right) \cdot 1. \quad (7)$$

Similarly, for the allocation in the k -th subframe, the algorithm performs at most z_k operations:

$$z_k = U \left(1 - \frac{k-1}{N}\right)(N-k+1) + \dots + U \left(1 - \frac{N-1}{N}\right) \cdot 1. \quad (8)$$

Then, the following result is derived by taking into account the worst case, in which exactly N subframes are used:

$$\begin{aligned} \sum_{k=1}^N z_k &= 1(N \cdot U) + 2 \left[(N-1) \cdot U \left(1 - \frac{1}{N}\right) \right] + \dots \\ &+ N \left[1 \cdot U \left(1 - \frac{N-1}{N}\right) \right] = U \cdot \sum_{k=1}^N \frac{k(N-k+1)^2}{N}. \end{aligned} \quad (9)$$

Recalling the results for well-known sums $\sum_{k=1}^n k = \frac{n(n+1)}{2}$, $\sum_{k=1}^n k^2 = \frac{n(n+1)(2n+1)}{6}$, and $\sum_{k=1}^n k^3 = \frac{n^2(n+1)^2}{4}$, the computational complexity for our algorithm is $O(U \cdot N^3)$.

Therefore, our solution scales with the number of users, and, although it grows with N^3 , it can be used to optimize the scheduling of realistically small groups of interfering base stations (e.g., up to ~ 10 neighboring base stations).

V. EVALUATION TOOLS

In this Section we present the two software tools used to evaluate BASICS: (i) a mathematical tool, namely MATLAB, that gives a first evaluation of the impact of BASICS in multi-cellular environments without going into the intricacies of LTE detailed implementation,¹ and (ii) the OPNET Modeler simulator, which allows to evaluate the specific impact of LTE protocols onto our results, although such simulations can only be carried out in small network scenarios only due to the computational cost and time required by the simulator.²

A. MATLAB implementation

MATLAB provides a suitable set of mathematical tools to implement and evaluate BASICS and scheduling mechanisms without going into packet level simulations. Our MATLAB implementation operates as follows. In the first step, the system is initialized, i.e., the positions of base stations are chosen at random in a square area, and mobile stations are dropped in the same area according to a uniform spatial distribution. Users are associated to base stations based on the strongest average received power, i.e., based on distance, and do not change base station during the simulation. The average received power depends on the transmission power set at each base station and on the pathloss, according to the classical *Free Space formulation*. Fading is considered in the numerical simulations in addition to pathloss, through a random variable, expressed in dB, distributed as a zero-mean Gaussian with standard deviation equal to 2 dB.

The second phase is to run the BASICS algorithm to decide which base station has to transmit in which subframe. In the algorithm, we use a unique SINR threshold for all the users.

Eventually, in the last phase, we calculate the throughput received in 1000 consecutive frames by each user. We repeat the simulation with different random seeds, averaging the results.

Note that the throughput depends on the channel state simulated, which affects the transmission rate achievable in the current frame, and on the user scheduling mechanism adopted by the base station. As for transmission rates, LTE specifications define 16 different CQI indexes, which correspond to different Modulation and Coding Schemes (MCS's), as described in Table I. The mapping of SINR values onto CQI and transmission rates is done as follows. First, we compute the spectral efficiency η corresponding to the SINR using the Shannon's formula:

$$\eta = \log_2 \left(1 + \frac{\text{SINR}}{\Gamma} \right), \quad (10)$$

where Γ is a coefficient that depends on the target BER (which is equal to 0.00005 in our case): $\Gamma = -\ln(5 \cdot \text{BER})/1.5$ [20].

¹Available online: http://fourier.networks.imdea.org/~vincenzo_sciancapere/download/MATLAB_LTE_23072012.zip

²Available online: http://fourier.networks.imdea.org/~vincenzo_sciancapere/download/OPNET_patch_23072012.zip

TABLE I
LTE CQI INDEX AND EFFICIENCY

Modulation Scheme	Approximate code rate	CQI Index	Interval for η	Efficiency (bits/symbol)
No transm.	–	0	0	–
QPSK	0.076	1	$0 \div 0.15$	0.1523
	0.12	2	$0.15 \div 0.23$	0.2344
	0.19	3	$0.23 \div 0.38$	0.3770
	0.3	4	$0.38 \div 0.60$	0.6016
	0.44	5	$0.60 \div 0.88$	0.8770
16QAM	0.59	6	$0.88 \div 1.18$	1.1758
	0.37	7	$1.18 \div 1.48$	1.4766
	0.48	8	$1.48 \div 1.91$	1.9141
64QAM	0.6	9	$1.91 \div 2.40$	2.4063
	0.45	10	$2.40 \div 2.73$	2.7305
	0.55	11	$2.73 \div 3.32$	3.3223
	0.65	12	$3.32 \div 3.90$	3.9023
	0.75	13	$3.90 \div 4.52$	4.5234
	0.85	14	$4.52 \div 5.12$	5.1152
	0.93	15	≥ 5.12	5.5547

Second, using Table I, we find the interval for η that corresponds to the SINR, and use the efficiency reported in the rightmost column as the net rate per allocated symbol used in the subframe. Since each subframe is divided in 2 time slots, each time slot contains 100 Physical Resource Blocks (PRBs) and each PRB is structured in $7 \cdot 12 = 84$ OFDMA symbols, the maximum number of OFDMA symbols assigned to one user in a single subframe is $84 \cdot 100 \cdot 2 = 16800$. The maximum throughput we can get in this case will be exactly $16800 \cdot 5.5547 = 93318.96$ bit/subframe, i.e., 93.318 Mb/s as the subframe lasts 1 ms. Eventually, taking into account the adopted user scheduling scheme, MATLAB computes the number of PRBs to be allotted to each user, and computes the corresponding throughput.

As for the scheduling of users, we implement a basic round robin scheme, allotting equal airtime to each user in round robin order, and a state of the art proportional fairness scheduler [21]. The latter allots resources according to user priorities computed at the beginning of each subframe as the ratio between the achievable rate in that frame, and the average throughput received in the past.

Although the MATLAB implementation misses the impact of detailed LTE and network-layer protocols, e.g., the impact of mechanisms used for generating realistic traffic, or for computing link adaptation and physical resource allocation in OFDMA, it does provide a platform that allows to efficiently simulate large network scenarios while providing a reasonable level of accuracy.

B. OPNET Modeler

To evaluate the impact of real protocols on our proposal, we also modified the well-established OPNET simulator³.

OPNET already implements several LTE scenarios, for which nodes and functionalities are designed in a modular way. We modified the modules specifying the behavior of the base station, to simulate the control traffic needed to run BASICS, and the behavior of the physical channel, to account for dynamic fading effects not yet implemented on the simulator. Most importantly, we have implemented our base station scheduling into a central entity called Evolved Packet Core (EPC), which collects global interferences reported by users

³OPNET University Program, http://www.opnet.com/university_program/

to their base stations (CQI messages). For each simulation module, new layers and process model states have been added in order to perform the newly required operations.

Furthermore, to simulate the dynamic capabilities of our proposal, we have programmed an internal interrupt for each base station, in order to collect all interferences reported by the users, prepare a control message containing such information, and send it to the EPC component. The EPC component runs BASICS periodically, and enforces a new base station scheduling with a refresh interval of $2s$, i.e., 2000 subframes. The refresh interval has been selected to track channel quality variations, while keeping the signaling overhead low.

For our scope, we need an LTE scenario with several users associated to a group of base stations served by an EPC interface. To this end, the EPC is connected to a general gateway by a serial connection (providing a speed up to 2488 Mb/s), and a server is added to serve the users' demands.

We use an OPNET-predefined video conference application to generate traffic. Specifically, all users are adjusted to request a video streaming through a UDP connection characterized by 30 frames/s, where every frame has a resolution of 352×240 pixels (i.e., 253440 bytes). The server is able to respond to each demand thereby reaching the saturation in the transmission. According to LTE specification, a single user served by a base station can reach a very high throughput, about 90 Mb/s.

As OPNET is a very complex packet simulator, each simulation takes several minutes to run over our server, which is a Dell Optiplex 990 with a Intel(R) Core(TM) i7-2600 CPU at 3.40 GHz with 8 cores, 8 GB of RAM and Windows 7 Professional SP1 64 bit. That is the reason why it is not possible to use OPNET for a very large number of base stations. Therefore, we use OPNET for small network topologies only, while we use MATLAB for larger scenarios.

VI. PERFORMANCE EVALUATION & DISCUSSION

In this section, we evaluate the performance of BASICS by using the evaluation tools presented in the previous section and show that it achieves near-optimal results and outperforms existing solutions.

A. Selection of the BASICS Threshold Th

One of the key parameters upon which BASICS relies is the SINR threshold Th , i.e., the minimum SINR guaranteed to each scheduled user in the system. Since the most suitable threshold can fall in a continuous range of values, while the achievable rates are a discrete set corresponding to the available MCS values, we search the threshold Th among the MCS thresholds corresponding to the efficiency ranges listed in Table I. After evaluating all available thresholds, we select Th as the threshold that yields the best sum of logarithmic throughput; we use this well-known metric as it accounts for both throughput and fairness.

Fig. 1 compares the results obtained for different thresholds for a scenario with 9 base stations and 225 users. The lowest threshold reported in the figure represents the case of no minimum SINR requirements, i.e., it represents the normal network operation without BASICS. On the other hand, the

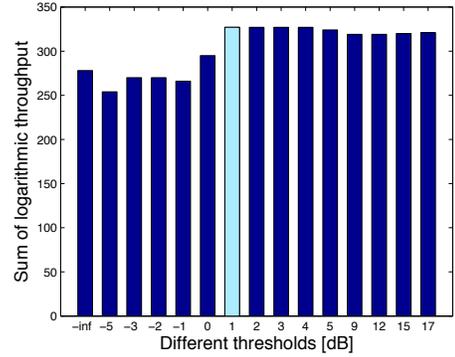


Fig. 1. Sum of logarithmic throughputs achieved with BASICS with different thresholds (dB) in a scenario with 9 base stations and 225 users.

highest threshold is higher than the minimum SINR required to use the highest MCS, which implies that the constraint expressed by Eq. (2) can be fulfilled only by scheduling no interfering base stations at all. As a result, the highest threshold corresponds to a pure TDM system in which each base station transmits in isolation. We observe from the figure that the optimal threshold lays in-between the two extremes, and provides a gain of several logarithmic units in comparison to other suboptimal thresholds.

B. Optimality of solution

In Section III-B, we have formulated our problem as to minimize the number of subframes needed to allocate all base stations once. In the following, we evaluate the performance of BASICS and compare it against the legacy FFD bin-centric algorithms, taking the number of subframes used as the evaluation metric. To this aim, we implemented in MATLAB the FFDSum algorithm as well as a simple tool that identifies the optimal base station mapping (in terms of achieved throughput) by means of a brute-force approach.

To compare performance achieved by BASICS, FFDSum and the optimal (brute force) algorithm, we simulate a network with 3 to 7 base stations, each having 25 users. Fig. 2 shows that BASICS finds the same number of subframes as the optimal solution, except for the case of 5 base stations in which it uses one extra subframe. Furthermore, BASICS uses at most one subframe more than the theoretical lower bound, obtained from Eq. (5). In contrast, FFDSum achieves significantly worse results (as expected from the discussion in Section IV). Fig. 3 further illustrates the throughput performance obtained from these algorithms, and reveals that FFDSum not only uses more subframes, but also provides worse throughput. In contrast, BASICS achieves near-optimal results.

C. Performance gain

We next evaluate the performance of BASICS in terms of throughput and fairness by using MATLAB, which allows to explore the impact of a large number of base stations and users under various network configurations.

In order to assess its performance, we compare BASICS against the following two approaches: (i) normal network operation, in which all base stations are allowed to transmit in any subframe (referred to as “Legacy” in the figures), and (ii) a frequency reuse 3 scheme that partitions the network into three parts. For all cases, two intra base station schedulers are

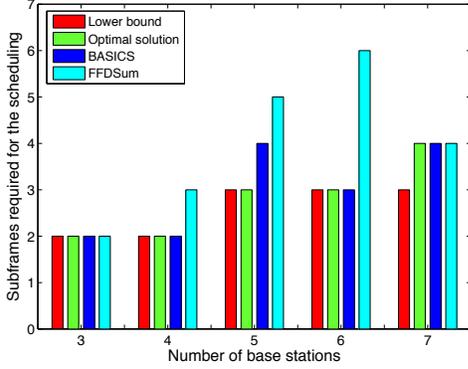


Fig. 2. Number of subframes used with BASICS and with the optimal base station scheduling (obtained via brute force search).

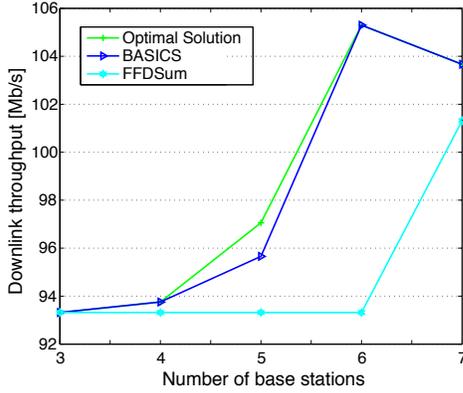


Fig. 3. Throughputs achieved with BASICS and with the optimal base station scheduling (obtained via brute force search).

considered: round robin and proportional fair scheduling (for clarity of presentation, for BASICS and frequency reuse 3, we only show results achieved with the proportional fair scheduler, which are slightly better than with the round robin scheduler). We measure network performance in terms of the sum of the logarithms of the throughputs, as this is a well accepted metric to compare different scheduling mechanisms in terms of efficiency as well as fairness. Note that for the case of frequency reuse 3, we normalize the throughput to the number of carriers utilized, i.e., 3.

Fig. 4 shows the performance of the above approaches when we fix the total number of users in the system to 150, and vary the number of base stations from 3 to 9. We observe from the figure that BASICS significantly improves network performance with respect to both legacy approaches and frequency reuse 3. Results achieved with frequency reuse 3 are similar to the ones achieved with BASICS only for scenarios with very few base stations.

In order to gain additional insights on the actual distribution of throughputs, Fig. 5 shows the CDF of user throughputs for the specific case of a network with 5 base stations and 150 users. Notably, with BASICS, the majority of users receive a throughput in the range 600 to 700 Kb/s, while other schemes yield a throughput distribution spread over large intervals (from few Kb/s to about 2 Mb/s). This translates into improved fairness levels when BASICS is adopted. Therefore, this shows that BASICS not only ameliorates the throughput, but also

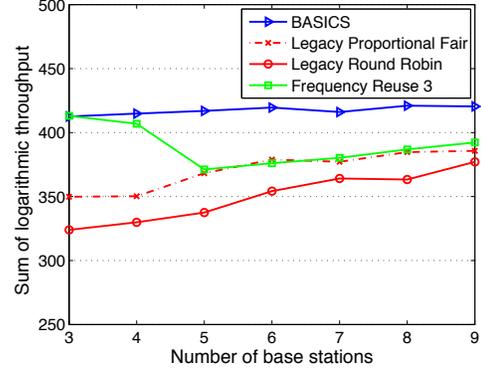


Fig. 4. Performance comparison with a fixed number of users (150 users).

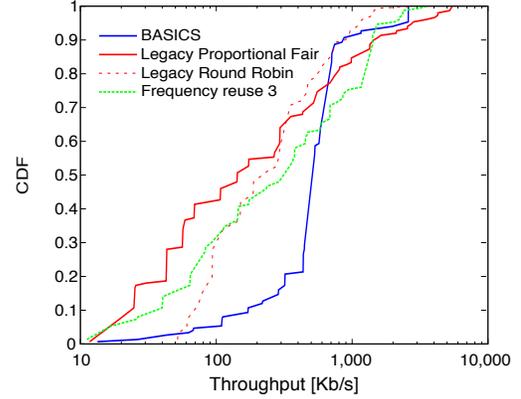


Fig. 5. CDF of per-user throughput with 5 base stations and 150 users.

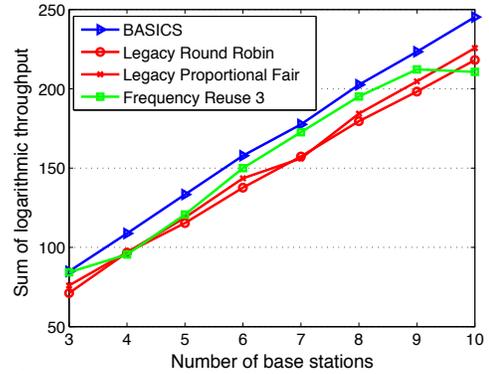


Fig. 6. Performance comparison with fixed number of users per base station (8 users per base station).

enhances fairness.

Next, we evaluate the impact of the number of users. To this end, we consider a network in which the number of users is proportional to the number of base stations. Specifically, we simulate 3 to 10 base stations with 8 users each. Fig. 6 depicts the sum of logarithmic throughputs as a function of the number of base stations. Also for this case, BASICS exhibits the best performance over all the other approaches. On the one hand, the gain of BASICS over the legacy schemes is of several logarithmic units (and hence substantial in a linear scale). On the other hand, the gain over frequency reuse 3 is lower until the number of base stations reaches 10 (which is explained by the fact that frequency reuse becomes less effective as the

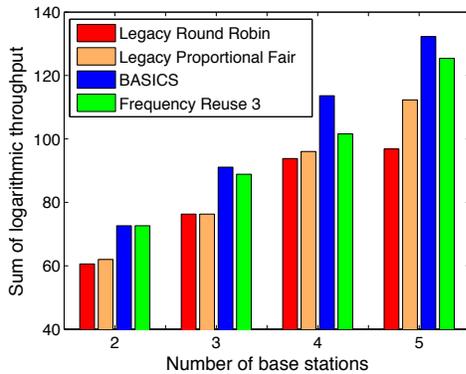


Fig. 7. Performance evaluation of BASICS (OPNET simulator).

network density grows). Taking into account that frequency reuse requires multiple carriers to achieve worse results, we conclude from these results that BASICS provides substantial improvements in performance also for this case.

D. Impact of LTE implementation details

In the previous subsections, we have evaluated the performance of BASICS based on our MATLAB tool, which does not take into account the impact of network protocols. In the following, we evaluate the performance of our proposal, of legacy schedulers, and of frequency reuse scheme in a more realistic scenario by using packet level simulations with the OPNET simulator, on top of which we implemented BASICS.

To assess the impact of the network protocols, we repeat the experiment of Fig. 6 (fixed number of users per base station) with OPNET. We limit the number of base stations to 5 due to the computational constraints of the packet level simulator. Fig. 7 depicts the sum of logarithmic throughputs achieved when each base station has 8 users. Results exhibit the same trend as Fig. 6, where BASICS outperforms all other approaches. A close look at the figures confirms that MATLAB results are very close to the OPNET ones.

Based on the above results, we draw the following two conclusions: (i) BASICS substantially boosts network performance in terms of throughput and fairness; and (ii) the gain provided by BASICS is not affected by transport-, network- and MAC-layer implementation details, which corroborates the MATLAB results presented throughout the paper.

VII. SUMMARY & CONCLUSIONS

In this paper, we have proposed BASICS, a scheme to coordinate neighboring base stations to minimize inter-cell interference while achieving high spectral efficiency. In contrast to previous works available in the literature, our approach leverages *base station* scheduling rather than *user* scheduling, which makes BASICS suitable to deploy the recently standardized ABSF mechanism. We have formulated the base station scheduling problem and proved that it is equivalent to a multi-dimensional vector bin-packing problem, which is NP-hard. We have then proposed a heuristic and have shown that it achieves near-optimal performance while scaling with the number of users. Our work has revealed that mapping base station activities over subframes yields significant gains in terms of spectral efficiency and also results in a good level

of fairness in the distribution of throughput among users. Additional advantages of BASICS are that it incurs a very low signaling overhead and that it does not require changes in the per-user scheduling policies implemented by the base stations.

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, March 2011.
- [2] IEEE 802.16 W.G., “IEEE Std. for Local and Metropolitan Area Networks. Part 16: Air Interface for Broadband Wireless Access Systems. Amendment 3: Advanced Air Interface”, IEEE 802.16m, May 2011.
- [3] G. Boudreau, J. Panicker, Ning Guo, Rui Chang, Neng Wang, and S. Vrzic, “Interference coordination and cancellation for 4G networks”, *IEEE Communications Magazine*, vol. 47, no. 4, pp. 74–81, April 2009.
- [4] X. Tang, S.A. Ramprasad, and H. Papadopoulos, “Multi-Cell User-Scheduling and Random Beamforming Strategies for Downlink Wireless Communications”, in *Proceedings of IEEE VTC 2009-Fall*, Sept. 2009.
- [5] R. Bendlin, Huang Y.-F., M.T. Ivrlac, and J.A. Nossek, “Fast Distributed Multi-Cell Scheduling with Delayed Limited-Capacity Backhaul Links”, in *Proceedings of IEEE ICC*, June 2009.
- [6] National Telecommunication Regulatory Authority Funded Project, “4G++: Advanced Performance Boosting Techniques in 4th Generation Wireless Systems”, WP4 Deliverable D4.1, 2012.
- [7] R. Irmer, H. Droste, P. Marsch, M. Grieger, G. Fettweis, S. Brueck, H.-P. Mayer, L. Thiele, and V. Jungnickel, “Coordinated multipoint: Concepts, performance, and field trial results”, *IEEE Communications Magazine*, vol. 49, no. 2, pp. 102–111, February 2011.
- [8] D. Lu and D. Li, “A Novel Precoding Method for MIMO Systems with Multi-Cell Joint Transmission”, in *Proceedings of IEEE VTC 2010-Spring*, May 2010.
- [9] E. Y. Kim and J. Chun, “Coordinated beamforming with reduced overhead for the downlink of multiuser MIMO systems”, *IEEE Communications Letters*, pp. 810–812, 2008.
- [10] S.M. Shukry, K. Elsayed, A. Elmoghazy, and A. Nassar, “Adaptive Fractional Frequency Reuse (AFFR) scheme for multi-cell IEEE 802.16 e systems”, in *Proceedings of IEEE HONET*, December 2009.
- [11] J. Yoon, M. Y. Arslan, K. Sundaresan, S. V. Krishnamurthy, and S. Banerjee, “A distributed resource management framework for interference mitigation in ofdma femtocell networks”, in *Proceedings of the thirteenth ACM international symposium on Mobile Ad Hoc Networking and Computing*, New York, NY, USA, 2012, MobiHoc ’12, pp. 233–242, ACM.
- [12] M. Y. Arslan, J. Yoon, K. Sundaresan, S. V. Krishnamurthy, and S. Banerjee, “FERMI: a femtocell resource management system for interference mitigation in OFDMA networks”, in *Proceedings of the 17th annual international conference on Mobile computing and networking*, New York, NY, USA, 2011, MobiCom ’11, pp. 25–36, ACM.
- [13] T. Bonald, S. Borst, and A. Proutière, “Inter-cell coordination in wireless data networks”, *European Transactions on Telecommunications*, vol. 17, no. 3, pp. 303–312, 2006.
- [14] Third Generation Partnership Project (3GPP), “Physical layer procedures (Release 10) for Evolved Universal Terrestrial Radio Access (E-UTRA)”, 3GPP TR 36.213 v 10.5.0, March 2012.
- [15] S. Martello and P. Toth, *Knapsack problems: algorithms and computer implementations*. John Wiley & Sons, Inc., 1990.
- [16] L.V. Kantorovich, “Mathematical methods of organizing and planning production”, *Management Science*, vol. 6, no. 4, pp. 366–422, 1960.
- [17] E.G. Coffman Jr., M.R. Garey, and D.S. Johnson, “Approximation algorithms for bin packing: A survey”, *Approximation Algorithms for NP-Hard Problems*, pp. 46–93, 1996.
- [18] R. Panigrahy, K. Talwar, L. Uyeda, and U. Wieder, “Heuristics for Vector Bin Packing”, <http://research.microsoft.com/apps/pubs/default.aspx?id=147927>.
- [19] M. C. Necker, “Interference Coordination in Cellular OFDMA Networks”, *IEEE Network*, vol. 22, no. 6, pp. 12, December 2008.
- [20] N. Baldo, M. Miozzo, M. Requena-Esteso, and J. Nin-Guerrero, “An open source product-oriented LTE network simulator based on ns-3”, in *Proceedings of ACM MSWiM*, October-November 2011.
- [21] P. Bender, P. Black, M. Grob, R. Padovani, N. Sindhushayana, and A. Viterbi, “CDMA/HDR: a bandwidth-efficient high-speed wireless data service for nomadic users”, *IEEE Communications Magazine*, vol. 38, no. 7, pp. 70–77, 2000.