

RESEARCH

Open Access

Autonomous algorithms for centralized and distributed interference coordination: a virtual layer-based approach

Martin Kasparick^{1*} and Gerhard Wunder²

Abstract

Interference mitigation techniques are essential for improving the performance of interference limited wireless networks. In this paper, we introduce novel interference mitigation schemes for wireless cellular networks with space division multiple access (SDMA). The schemes are based on a virtual layer that captures and simplifies the complicated interference situation in the network and that is used for power control. We show how optimization in this virtual layer generates gradually adapting power control settings that lead to autonomous interference minimization. Thereby, the granularity of control ranges from controlling frequency sub-band power via controlling the power on a per-beam basis, to a granularity of only enforcing average power constraints per beam. In conjunction with suitable short-term scheduling, our algorithms gradually steer the network towards a higher utility. We use extensive system-level simulations to compare three distributed algorithms and evaluate their applicability for different user mobility assumptions. In particular, it turns out that larger gains can be achieved by imposing average power constraints and allowing opportunistic scheduling instantaneously, rather than controlling the power in a strict way. Furthermore, we introduce a centralized algorithm, which directly solves the underlying optimization and shows fast convergence, as a performance benchmark for the distributed solutions. Moreover, we investigate the deviation from global optimality by comparing to a branch-and-bound-based solution.

Keywords: Cellular interference management; SDMA; Distributed algorithms; Autonomous inter-cell coordination; Power control; Network utility maximization

1 Introduction

Increasing bandwidth requirements, not least due to the fast growing popularity of handheld devices with high data rate consumption, bring cellular networks to the brink of their capacity. Until recently, an end to the growth of this demand is not yet in sight. In order to optimally exploit the available bandwidth, current cellular networks experienced a paradigm shift towards frequency reuse-1. Consequently, this leads to an increased susceptibility to interference such that current and future cellular networks are usually interference limited. This situation is aggravated by a trend towards ever smaller cell sizes. Especially users at the cell-edge are affected by high inter-cell interference (ICI). In theory, fully coordinated networks,

where neighboring base stations act as a large distributed antenna array, promise a vast boost in performance [1,2]. However, this makes great demands on synchronization and backhaul bandwidth. In fact, the promised gains from such schemes turn out to be hard to implement in practice [3]. As a consequence, distributed schemes for interference mitigation, incorporating joint scheduling and adaptive power allocation, are of utmost interest. However, due to mobile users and varying channel conditions, such algorithms have to be dynamic and able to operate autonomously^a. Moreover, future cellular systems, including pico and femto cells, must be self-organizing to maintain flexibility and scalability. Therefore, self-optimizing interference coordination schemes are needed.

In this paper, we introduce such schemes with special focus on cellular space division multiple access (SDMA) networks. All proposed algorithms can be seen as applications of a general radio resource management

*Correspondence: martin.kasparick@hhi.fraunhofer.de

¹ Fraunhofer Heinrich Hertz Institute, Einsteinufer 37, Berlin 10587, Germany
Full list of author information is available at the end of the article

framework that configures and optimizes network operation autonomously and that allows us to incorporate service requirements and performance targets. The framework combines the following three essential ingredients:

- (i) *Network utility maximization (NUM)*. By formulating a network utility maximization problem, we steer the system to a desired operating point. This includes fairness goals (like proportional fair and max-min).
- (ii) *Virtual model*. On top of regular scheduling and resource allocation, we maintain a virtual model of the network that captures and simplifies the complicated interference situation in the real world. It comprises a static ‘cooled-down’ version of the network based on long-term gains, thus suppressing the influence of fast-fading. We use this model to obtain granular power control decisions. Thus, it can be seen as a new layer for long-term resource allocation decisions in a cooperative way (including corresponding message exchange).
- (iii) *Suitable short-term scheduling*. Instantaneously, we employ a popular gradient scheduler which is known to asymptotically converge to the solution of the underlying utility maximization problem. The scheduler has to take the power constraints into account (which can be strict or average constraints) that are obtained in the virtual layer to manage interference.

To implement this approach, we equip each base station with an additional sector controller that requires additional (infrequent) long-term feedback from the mobiles. Note, that this is a practical assumption since currently, standardized schemes (such as long-term evolution (LTE)-advanced) also consider advanced feedback concepts that are not limited to the serving base station [4]. Moreover, we permit a limited message exchange between the sector controllers. Figure 1 depicts the general approach. Based upon the long-term feedback, the sector controllers create and update a ‘virtual model’ of the network. The sector controller is thereby just the entity in each base station that takes care of all (virtual) resource allocation and scheduling issues. The main idea of this virtual model is to optimize over ‘virtual resources’, based on long-term averaged versions of true variables, in order to distributively control a rapidly changing complex network. Thus, the virtual model captures the complicated interference interdependencies in the ‘real’ network and is used for the control and adaption of power allocations. The virtual model allows each sector to efficiently compute estimates of the gradient of the system utility function with respect to transmit powers of particular resources in the sector, thus allowing local maximization of the overall utility. More precisely these estimates

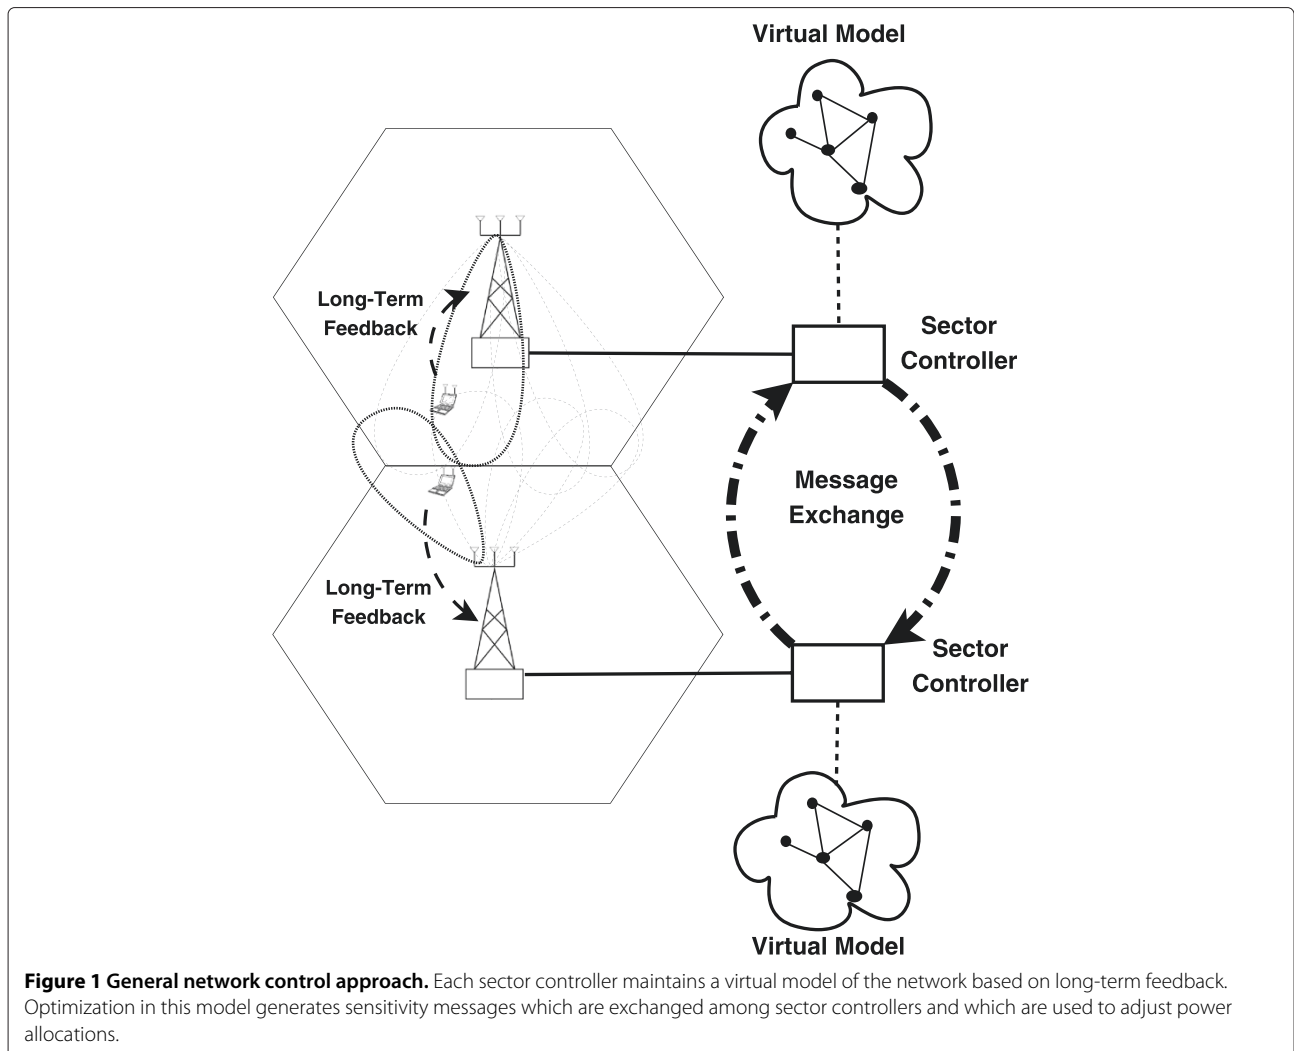
are eventually generated by a ‘virtual scheduling’ process. The virtual model and this virtual scheduling process can be considered as an additional virtual layer for resource allocation. The idea is that if user rates improve in the virtual layer, they also improve in the real network. Instead of exchanging channel state information (CSI) with all relevant base stations, only the (gradient) information (called *sensitivities*) that is obtained in the virtual layer needs to be exchanged. The virtual model, being based on average quantities, can be further justified since the goal of the algorithm is to adapt the transmit power levels to average interference levels and not to track fast fading.

As mentioned before, we focus on fixed codebook-based schemes in SDMA networks. In particular, we assume that each base station maintains a fixed codebook of a certain size comprising precoding vectors, called ‘beams’. These beams can be used to support multiple users on the same time-frequency resource. Using fixed codebooks is a practical assumption and allows us to compare our algorithms with practical schemes used in current cellular systems.

The optimization within the virtual layer can be organized either in a centralized or in a distributed manner. Clearly, a distributed implementation is favorable but in order to accurately quantify the tradeoffs involved, we also investigate a centralized solution. This does not only provide a valuable benchmark for the distributed algorithms, but may also be a feasible option for small networks, where a central controller is indeed possible. Our centralized baseline algorithm is based on an alternating optimization approach, solving scheduling and power optimization in the virtual control plane separately. Since user rates are strongly coupled via transmit powers, we employ a successive convex approximation technique to tackle the inherent non-convexity. Due to the non-convex nature of the underlying optimization problem, global optimality cannot be guaranteed. Therefore, we additionally assess the deviation from global optimality by comparing our approach to an optimal solution based on branch-and-bound (BNB) in a simplified setting.

1.1 Related work

There is a significant amount of research on interference mitigation in cellular networks [5], which is often treated as a specific aspect of self-organization in cellular networks. A comprehensive survey on this is provided in [6]. A straightforward approach to avoid interference is to use a frequency reuse factor greater than one or some fractional frequency reuse scheme [7]. Another line of research targets reuse-1 networks and interference mitigation by power control and resource allocation. There have been a variety of suggestions for joint multicell power



control and scheduling in cellular networks such as [8,9] (see also [10] for an overview). Multicell coordination via joint scheduling, beamforming, and power adaptation is considered in [11]. Thereby, fairness requirements (leading to concave utility functions) are fundamental for current and future cellular standards. The work [12] considers joint power allocation and user assignment to cells in the NUM context, taking into account a mixture of concave and non-concave utilities. In [13], a gradient algorithm-based scheme for self-organizing resource allocation in LTE systems is proposed. However, a multitude of information has to be exchanged between coordinating sectors.

Although many of the aforementioned references consider distributed schemes, none treats resource allocation and interference management in multiuser MIMO systems. By contrast, our framework explicitly aims to exploit the freedom in terms of resource and power allocation offered by SDMA. Thereby, our framework builds upon

and extends the framework introduced in [14,15] for single-antenna (SISO) networks.

It is commonly accepted that the underlying optimization problems, which are non-convex in general, can be solved optimally only for a limited set of problems and utilities in reasonable time. Existing solutions in cellular networks often rely on uplink downlink duality [16]. Since they attempt to solve such non-convex problems directly, successive approximation techniques become an increasingly popular tool to treat this non-convexity, used for example for power control in DSL [17] and multihop networks [18]. Global optimality is often achieved using branch-and-bound-based approaches [19,20] or monotonic optimization [21], however at a high computational complexity. Other directions include interference pricing [22,23] or game theoretic approaches [24]. Only recently, distributed coordination schemes in cellular networks have gained increasing popularity [9,14,25] due to the high complexity of centralized approaches. For example,

[26] considers the derivation of transmit beamformers, also based on interference prices, for dense small cell networks.

1.2 Organization

The paper is organized as follows. In Section 2, we introduce the considered system model, introduce the notation, and describe the optimization problem that we address. In Section 3, we present a virtual control layer for solving this problem and introduce three distributed algorithms that are based on different realizations of this virtual control plane. In Section 4, we propose an alternative centralized scheme, while in Section 5, we present an optimal solution based on branch-and-bound. In Section 6, we present system-level simulation results that evaluate the performance of the distributed algorithms and moreover investigate simpler scenarios to compare these to the centralized and the optimal baselines. Eventually, in Section 7, we state the most important conclusions.

2 System model and notation

We consider the downlink of a cellular OFDMA network, where each cell is sub-divided into three sectors. In total, we have M sectors $m \in \{1, \dots, M\}$. Each sector is served by a base station having n_T transmit antennas with a corresponding *sector controller* responsible for user selection and resource allocation. There are I users randomly distributed in the system, each equipped with n_R receive antennas. In the following, we assume $n_R = 1$. We assume that each user has pending data at all times. Let I_m be the number of users associated to sector m .

We assume slotted time with time slots $t = 1, 2, 3, \dots$, called transmission time interval (TTI). In reference to

LTE specifications, orthogonal frequency-division multiplexing (OFDM) sub-carriers are grouped into J sub-bands, called physical resource blocks (PRBs)^b. For the duration of a time slot, the system is in a fixed fading state from a finite set \mathcal{F} . Note, that the assumption of finite fading states is made to foster the analysis. In our simulations, however, we evaluate the performance using established 3GPP channel models (cf. Section 6). Nevertheless, the assumption of a finite set of fading states is often used in the literature [27] and can be further justified by the observation that, first, the measurement accuracy of the user devices is limited and, second, there is only a finite number of modulation and coding schemes that the base station can choose from. Fading state l , in turn, induces a finite set of possible scheduling decisions $k \in \mathcal{K}(l)$. We denote by $\pi(l)$ the probability of fading state l (with $\sum_l \pi(l) = 1$). The sector controllers perform linear precoding, where precoding vectors for beamforming are taken (as for example in LTE) from a *fixed* N -element codebook $\mathcal{C}_N := \{\mathbf{u}_1, \dots, \mathbf{u}_N\}$ which is publicly known. In the following, we identify a beamforming vector ($\mathbf{u}_b \in \mathbb{C}^{n_T}$) by its index $b \in \{1, \dots, N\}$. From the sector controllers' perspective, this makes a beam b on a specific PRB j a possible resource for user selection and power allocation. This relationship is depicted schematically in Figure 2. Let P_{jb}^m be the power assigned to beam b on PRB j by base station m . This value is determined differently in each of the presented approaches under comparison. In summary, the (non-trivial) task of each sector controller is to find a scheduling decision (being an assignment of available resources on PRBs— to users) and a suitable power allocation, such that a global network utility function is maximized.




Figure 2 Resource grid from the scheduler's perspective. Resource allocation and power control can be performed on a per-PRB-per-beam granularity.

Eventually, we have the following additional notational conventions. Let $\mathbf{h}_{ij}^m(t) \in \mathbb{C}^{n_T}$ be the vector of instantaneous complex channel gains from base station m to user i on PRB j (we assume frequency-flat channels within a PRB). Accordingly, $\langle \mathbf{h}_{ij}^m(t), \mathbf{u} \rangle$ denotes the scalar product of (beamforming) vector \mathbf{u} and channel $\mathbf{h}_{ij}^m(t)$. The noise power at the mobile terminal is denoted as σ^2 . Throughout the paper, we label vectors and matrices with bold-faced letters. For the ease of reference, the most important notation used throughout this paper is summarized in Table 1.

2.1 Problem statement

The overall goal is to devise autonomous network control schemes which maximize an overall increasing concave utility function U . The utility function is defined as the sum of sector utility functions U^m , which are in turn defined over *average* user rates $\bar{\mathbf{X}}^m$. Consequently, the problem to solve for each sector $m \in \{1, \dots, M\}$ is given by

$$\max_{\bar{\mathbf{X}}^m} \sum_m U^m(\bar{\mathbf{X}}^m) \quad (1)$$

$$\text{s.t. } \bar{\mathbf{X}}^m \leq \sum_l \pi(l) \sum_j \sum_{k \in \mathcal{K}(l)} \phi_{jk}^{lm} \sum_b \boldsymbol{\mu}_{jb}^l(k) \quad (2)$$

$$\sum_k \phi_{jk}^{lm} = 1, \text{ all } l, j, m \quad (3)$$

$$0 \leq \phi_{jk}^{lm} \leq 1 \quad (4)$$

Thereby, $\boldsymbol{\mu}_{jb}^l(k) = \left(\mu_{ijb}^l(k) \right)_{i \in \{1, \dots, I_m\}}$ is a vector comprising elements $\mu_{ijb}^l(k)$, which represent the rate that user i is assigned on PRB j and beam b when the system is in fading state l and scheduling decision k is chosen. They can be zero if the particular resource is not assigned to user i by decision k . ϕ_{jk}^{lm} denotes the fraction of time that scheduling decision k is chosen on PRB j , provided the system is in fading state l .

In a nutshell, we are not interested in a specific ‘snapshot’ of the system, but only in ergodic rates. Therefore, a possible control algorithm should not adapt to a specific system state but should be able to optimize the system performance *over time*.

Problem (1 to 4) is solved by applying a gradient scheduler [28] at each time instance. The gradient scheduler chooses the best scheduling decision k^* according to

$$k^*(t) \in \arg \max_{k \in \mathcal{K}(l)} \nabla U^m(\bar{\mathbf{X}}^m(t)) \sum_b \boldsymbol{\mu}_{jb}^{l(t)}(k). \quad (5)$$

Table 1 Important notation

Notation	Definition
P_{jb}^m	Power assigned to beam b on PRB j in sector m
\bar{P}_{jb}^m	Average power constraint (‘target’ power) of beam b on PRB j in sector m
$\bar{P}^m(t)$	Current total allocated power in sector m
P_{\max}	Maximum total base station power
$C_{jb}(k, \bar{P}_j^m)$	Power cost/consumption of beam b on PRB j in sector m
$C_{jb}^m(k)$	Virtual power cost/consumption of sector m on beam b and PRB j
$\pi(l)$	Probability of fading state l
$\mathcal{K}(l)$	Set of possible scheduling decisions given fading state l
$\mathbf{h}_{ij}^m(t)$	Channel vector from base station m to user i on PRB j at time t
$\bar{\mathbf{h}}_{ij}^m$	Average channel from base station m to user i on PRB j
$\boldsymbol{\mu}_{jb}^l(k)$	Vector of user rates on PRB j and beam b given decision k and fading state l
$\bar{\mathbf{X}}^m$	Vector of average total rates of users in sector m
ϕ_{jk}^{lm}	Fraction of time that scheduling decision k is chosen on PRB j given fading state l
ϕ_{ijb}^m	Fraction of time that user i in sector m is scheduled using beam b on PRB j
G_{ij}^m	Long-term gain of user i to base station m for his best beam on PRB j
G_{ijb}^m	Long-term gain of user i to base station m for beam b on PRB j
$R_{ij}^m(k)$	Virtual rate of user i on PRB j given decision k
\mathbf{R}_j	Vector of virtual user rates in PRB j
R_{ijb}^m	Virtual rate of user i on PRB j and beam b
X_i^m	Virtual average rate of user i in sector m
$\mathbf{X}(t/n_v)$	Vector of virtual average rates at t th virtual scheduling run
F_{ijb}^m	SINR of user i on beam b and PRB j
$D_{jb}^{(\hat{m}, m)}$	Sensitivity of sector m to a power change of beam b on PRB j in sector \hat{m}
D_{jb}^m	Sum sensitivity to a power change of beam b on PRB j in sector m
$n_j(k)$	Number of beams that are activated on PRB j if decision k is chosen
$\lambda_{jb}(t)$	Dual parameter: deviation of power on PRB j and beam b from the target value
$\alpha_{jb}^m(k)$	Scales target beam powers to instantaneous powers ‘costs’
$\alpha_{ijb}^m, \beta_{ijb}^m$	Approximation constants in concave lower bound

The gradient scheduler tracks average user rates $\bar{\mathbf{X}}^m(t)$ and updates them after each time slot according to

$$\bar{\mathbf{X}}^m(t+1) = (1 - \beta) \bar{\mathbf{X}}^m(t) + \beta J \sum_b \boldsymbol{\mu}_{jb}^{l(t)}(k^*(t)), \quad (6)$$

where the fixed parameter $\beta > 0$ determines the size of the averaging window. The gradient scheduler is known to asymptotically solve the problem (for $\beta \rightarrow 0$) without knowing the fading distribution. In case of logarithmic utilities, the gradient scheduler becomes the well-known proportional fair scheduler^c.

3 Autonomous distributed power control algorithms for interference mitigation

We now turn to the distributed power control schemes for autonomous interference management in cellular networks, which are designed to enable network entities to locally pursue optimization of the global network utility. We introduce three basic approaches. It is important to note that all algorithms are special cases of the well-known gradient algorithm [28,29], whose convergence behavior has been thoroughly analyzed. Therefore, we refrain from reproducing this theoretical analysis. However, in Section 6, for illustration purposes, we present numerical results indicating a fast convergence behavior. The main difference between the algorithms that we propose is the granularity of power control.

The first algorithm uses an opportunistic scheduler which only adapts the power per frequency sub-band, which is then distributed equally among activated beams. We call this *opportunistic algorithm* (OA). It leaves full choice to the actual scheduler as to which beam to activate at what time. The scheduler can therefore decide opportunistically (\Rightarrow opportunistic algorithm). However, the power budget per PRB which is distributed (equally) among activated beams is determined by an associated control scheme.

The second algorithm is the *virtual sub-band algorithm* (VSA), which enforces strict power constraints on each beam by requiring all beams to be switched on all the time (with power values given by the associated control). This has the advantage of making the interference predictable (assuming known power values). However, it leaves only limited freedom for the actual scheduler, whose task is reduced to user selection for each beam. Since a beam is always turned on, it can be treated as an independent resource for scheduling, just like a 'virtual' sub-band (\Rightarrow virtual sub-band algorithm).

The third is a hybrid approach, which permits opportunistic scheduling at each time instance but, in addition, enforces average power constraints per beam. We call it *cost-based algorithm* (CBA). It leaves more freedom for opportunistic scheduling than the virtual sub-band algorithm. In contrast to requiring all beams to be used at all times with strict power values, we only require the target beam power values to be kept *on average*. Thus, instantaneously, the scheduler is free to make opportunistic decisions based on the current system state. In order to assure that the average power constraints are kept, we introduce

an additional cost term into the utility maximization and the gradient scheduler (\Rightarrow cost-based algorithm).

We focus on the applicability in different fading environments, comparing the overall performance with respect to a network-wide utility function as well as the performance of cell-edge users. Thereby, we show that although the algorithms behave differently in different user mobility scenarios, in general, it is more beneficial to impose average rather than strict power constraints.

As we demonstrate later, the three algorithms perform differently with different mobility assumptions on the users. A problem that arises with increased mobility is that the (virtual) model lacks behind the actual network state. Especially when the controllers are restricted to a gradual power adaption process on a per-beam granularity, they might not always be able to fully exploit multiuser diversity. Since the proposed algorithms put different emphasis on opportunistic scheduling in power adaption and resource allocation decisions, they perform differently when facing user mobility and fast fading.

Let us now turn to the control plane, the virtual layer, of the considered algorithms. They all have the following general procedure in common. The goal of the control plane is to obtain estimates of the partial derivatives of the network utility with respect to the power allocation of particular resources and to adapt the power control policy accordingly. These estimates can be seen as estimates of the *sensitivity* of the network utility to changes of the allocation strategies. Thereby, the allocation can be, as in the OA of Section 3.1, the power allocation of a PRB which is then divided equally among activated beams. Or it can even be, as in the VSA of Section 3.2, the power allocation of an individual beam. Or it can also be, as in the CBA of Section 3.3, simply an average power constraint of an individual beam, which does not have to be kept at every single time instance.

A further similarity between all algorithms is that in addition to ordinary *short-term CSI*, the mobiles (not necessarily often) report long-term feedback to their base stations, which is then used to calculate (virtual) user rates and, accordingly, (virtual) average rates. These average rates are a good representation of the interference coupling throughout the network and are used to calculate the 'sensitivities' to power changes on particular resources in other sectors (and in the own sector). This sensitivity information is compiled in messages which are exchanged between the sectors^d. Upon reception of the message vectors, the sector controllers are now able to calculate the desired estimate of the *system* utility's sensitivity to power changes on particular resources and adjust powers (or power constraints) accordingly. It is important to always distinguish between actual rates and virtual rates. The actual average rates are the ones tracked by the ordinary (proportional fair) scheduler and determine

the sector's utility. The virtual average rates are based on long-term feedback of averaged channel gains and do not have an immediate physical meaning in the 'real world'. They are created by a 'virtual' scheduler based on 'virtual' scheduling decisions. The set of all these rates forms a 'virtual' model of the system which is used to derive power adaption decisions.

The question remains how the virtual average rates and accordingly the sensitivities to power changes are calculated. Besides the granularity of power control, this virtual model, which can be also seen as a virtual layer for interference mitigation above the actual short-term scheduling, is the main difference between the investigated algorithms. In the following, we will discuss this in detail.

3.1 Opportunistic algorithm

OA can be seen as a straightforward extension of the multi-sector gradient (MGR) algorithm in [14] to multi-antenna networks. Although it is designed for multi-antenna networks, it does not perform power control on a per-beam basis (as opposed to the other two algorithms) but gives complete freedom to the sector controllers with respect to the number and choice of beams that are active at every given time instance. The only value that is controlled is the power budget *per PRB*. Nevertheless, SDMA is applied where multiple users can be scheduled on the same PRB, however on different beams. Consequently, the long-term feedback of users comprises a codebook index (being the maximizing index b^* in (7)) as well as a corresponding gain

$$G_{ij}^m = \max_b \left(\bar{\mathbf{h}}_{ij}^m, \mathbf{u}_b \right), \quad (7)$$

where $\bar{\mathbf{h}}_{ij}^m$ is the channel from user i to its sector controller on PRB j , averaged to eliminate the influence of fast-fading.

The task of finding scheduling decision k now amounts to determining the best subset of users to be scheduled on a PRB, subject to the constraint that each user can only be scheduled exclusively on its reported beam. Let us define (virtual) user rates (time index omitted), assuming a user is scheduled on PRB j (otherwise $R_{ij}^m(k) = 0$), given by

$$R_{ij}^m(k) = \rho \left(F_{ij}^m \right), \text{ with } F_{ij}^m := \frac{G_{ij}^m \bar{P}_j^m(k)}{\sigma^2 + \sum_{m' \neq m} \bar{P}_j^{m'} G_{ij}^{m'}}. \quad (8)$$

In this paper, we use $\rho(x) = \log(1+x)$. Thereby, $\bar{P}_j^m(k)$ is the current power value of PRB j divided by the number of users scheduled, thus depending on decision k . Moreover, $\bar{P}_j^{m'} G_{ij}^{m'}$ represents a long-term estimate of the interference of sector m' on PRB j (which can be measured by the mobiles).

Having defined the virtual user rates \mathbf{R} , virtual *average* user rates \mathbf{X} and sensitivities are derived by *virtual scheduling* (based on a similar procedure in [14]) as follows. We run the following steps n_v times *per TTI* in each sector m and for any PRB j . Thereby, the parameter n_v determines how long the virtual scheduler runs before accepting the sensitivities. Consequently, a larger value means more overhead by the virtual layer but better results.

- We determine the virtual scheduling decision k^* using a gradient scheduler according to

$$k^* \in \arg \max_k \nabla U^m \left(\mathbf{X}^m \left(\frac{t}{n_v} \right) \right) \mathbf{R}_j \left(k, \frac{t}{n_v} \right).$$

- We update virtual average user rates according to

$$\mathbf{X} \left(\frac{t+1}{n_v} \right) = (1 - \beta_1) \mathbf{X} \left(\frac{t}{n_v} \right) + \beta_1 \mathbf{J} \mathbf{R}_j \left(k^*, \frac{t}{n_v} \right).$$

- We update sensitivities according to

$$D_j^{(\hat{m}, m)} \left(\frac{t+1}{n_v} \right) = (1 - \beta_2) D_j^{(\hat{m}, m)} \left(\frac{t}{n_v} \right) + \beta_2 \sum_{i=1}^{I_m} \frac{\partial U^m(X^m)}{\partial X_i^m} \frac{\partial R_{ij}^m \left(k^*, \frac{t}{n_v} \right)}{\partial \bar{P}_j^{\hat{m}}(t)}. \quad (9)$$

Thereby, β_1 and β_2 are small averaging parameters. Using (8), the derivatives in (9) are given by

$$\frac{\partial R_{ij}^m}{\partial P_j^{\hat{m}}} = \begin{cases} \rho' \left(F_{ij}^m \right) \frac{F_{ij}^m}{\bar{P}_j^m} & \hat{m} = m \\ -\rho' \left(F_{ij}^m \right) \frac{\left(F_{ij}^m \right)^2}{\bar{P}_j^m} \frac{G_{ij}^{\hat{m}}}{G_{ij}^m} & \hat{m} \neq m \end{cases}$$

Starting with equal power, the adaption of the PRB powers can be summarized as follows. From time to time, the sensitivities are exchanged and summed up by each sector controller for each beam and PRB. Since each $D_j^{(\hat{m}, m)}$ is an estimation of the sensitivity of sector m 's utility to a power change in sector \hat{m} , the summation gives an estimate of the *network* utility's sensitivity. Then, the power is increased on the PRB with the largest positive sum and decreased on the PRB with the largest negative sum.

3.2 Virtual sub-band algorithm

VSA requires long-term feedback that comprises average link gains *per beam* and sector from each mobile. We define

$$G_{ijb}^m = \left| \overline{\left(\mathbf{h}_{ij}^m, \mathbf{u}_b \right)} \right|^2 \quad (10)$$

to be the average link gain (the bar denotes empirical averaging over time) of mobile terminal i on beam b and PRB j to sector controller m . Based on the long-term

CSI, the control plane of VSA responsible for determining the power allocation per beam works as follows. Given average gains in (10), each sector controller calculates corresponding virtual rates according to

$$R_{ijb}^m =: \rho(F_{ijb}^m), \quad \text{with}$$

$$F_{ijb}^m = \frac{G_{ijb}^m P_{jb}^m}{\sigma^2 + \sum_{b' \neq b} G_{ijb'}^m P_{jb'}^m + \sum_{\hat{m} \neq m} \sum_{b'} P_{jb'}^{\hat{m}} G_{ijb'}^{\hat{m}}}. \quad (11)$$

Note that since all beams are activated at all times, we have an additional intra-sector interference term (as opposed to OA), since this interference can no longer be eliminated, e.g., by switching off beams. Let X_i^m be the virtual average rate of user i in sector m (not to be confused with actual average rates \bar{X} in (6)), defined as

$$X_i^m = \sum_{j=1}^J \sum_{b=1}^N \tilde{\phi}_{ijb}^m R_{ijb}^m. \quad (12)$$

Here, $\tilde{\phi}_{ijb}^m$ represent *optimal* time fractions of resource usage for sector m . They are determined as a solution to the following optimization problem (for fixed virtual user rates):

$$f(\mathbf{R}) := \max_{\phi_{ijb}^m} \sum_i U^m \left(\sum_j \sum_b \phi_{ijb}^m R_{ijb}^m \right) \quad (13)$$

$$\text{s.t.} \quad \sum_i \phi_{ijb}^m = 1, \quad 0 \leq \phi_{ijb}^m \leq 1.$$

We rely on an explicit solution to (13) since we cannot apply the virtual scheduling from [14]. This is because the resources for power control (which are now individual beams) are no longer orthogonal but cause interference to each other [30]. Having the virtual user rates, the sector controllers calculate sensitivities to power changes on beams for all sectors (including self) and beams, given by

$$D_{jb}^{(\hat{m}, m)} = \sum_i \frac{\partial U^m}{\partial X_i^m} \frac{\partial X_i^m}{\partial P_{jb}^{\hat{m}}} \quad (14)$$

$$= \sum_i \frac{\partial U^m}{\partial X_i^m} \sum_{b'} \left(\tilde{\phi}_{ijb'}^m \right)^{1-\varepsilon} \frac{\partial R_{ijb'}^m}{\partial P_{jb}^{\hat{m}}},$$

where

$$\frac{\partial R_{ijb}^m}{\partial P_{jb}^{\hat{m}}} = \begin{cases} \rho' \left(F_{ijb}^m \right) \frac{F_{ijb}^m}{P_{jb}^{\hat{m}}} & \hat{m} = m, b' = b \\ -\rho' \left(F_{ijb'}^m \right) \frac{\left(F_{ijb'}^m \right)^2}{P_{jb'}^{\hat{m}}} \frac{G_{ijb}^m}{G_{ijb'}^m} & \hat{m} = m, b' \neq b \\ -\rho' \left(F_{ijb'}^m \right) \frac{\left(F_{ijb'}^m \right)^2}{P_{jb'}^{\hat{m}}} \frac{G_{ijb}^{\hat{m}}}{G_{ijb'}^{\hat{m}}} & \hat{m} \neq m \end{cases}$$

Note that the small coefficient $\varepsilon > 0$ stems from an application of Theorem 1 (which can be found at the

end of Section 3.3) in order to ensure the differentiability of problem (13). The such generated sensitivities are exchanged from time to time between all sector controllers. Thereby, every sector k receives $J \cdot N$ sensitivity values from all other $(M - 1)$ sectors, in addition to the $J \cdot N$ values from its own sector. Thus, we have

$$D_{jb}^m = \sum_{\hat{m}=1}^M D_{jb}^{m, \hat{m}}, \quad (15)$$

summing up the sensitivities of all sectors (including itself) to a power change of beam b on PRB j in sector m and which can be either positive or negative. Note that sector indices m and \hat{m} in the RHS of (15) are interchanged compared with (14), since in (14), we are interested in how the beam in sector \hat{m} interferes with sector m , while in (15), it is of interest how the beam in sector m interferes with (all) sector(s) \hat{m} .

Since $D_{jb}^{m, \hat{m}}$ represent estimates of the sector utilities to a power change on jb in sector m , D_{jb}^m clearly is an estimate of the sensitivity of the *system's* utility to a power change on the respective beam. Depending on the D_{jb}^m , we can now make a power adjustment which steers the system operating point towards a greater utility in the virtual model.

The power adjustment is carried out in steps of $\Delta > 0$, which is small and fixed. Let $\bar{P}^m(t)$ denote the current total allocated power in sector m and P_{\max} the upper bound on the total sector powers. Then, the following procedure is applied.

- (1) Pick a virtual resource $(jb)_*$ (if there is one) such that $D_{(jb)_*}^{(m)}(t)$ is the smallest among all virtual resources jb with $D_{jb}^{(m)}(t) < 0$ and $P_{jb}^m(t) > 0$. Now, set

$$P_{(jb)_*}^m(t+1) = \max \left\{ P_{(jb)_*}^m(t) - \Delta, 0 \right\}.$$

- (2) If $\bar{P}^m(t) < P_{\max}$, pick $(jb)^*$ (if there is one) such that $D_{(jb)^*}^{(m)}(t)$ is the largest among those jb with $D_{jb}^{(m)}(t) > 0$. Set

$$P_{(jb)^*}^m(t+1) = P_{(jb)^*}^m(t) + \min \left\{ \Delta, P_{\max} - \bar{P}^m(t) \right\}.$$

- (3) If $\bar{P}^m(t) = P_{\max}$ and $\max_{jb} D_{jb}^{(m)}(t) > 0$, pick a pair $((jb)_*, (jb)^*)$ (if there is one) such that $D_{(jb)_*}^{(m)}(t)$ is the largest, and $D_{(jb)^*}^{(m)}(t)$ is the smallest among those virtual resources jb with $P_{jb}^m(t) > 0$ and $D_{(jb)_*}^{(m)}(t) < D_{(jb)^*}^{(m)}(t)$. Set

$$P_{(jb)_*}^m(t+1) = \max \left\{ P_{(jb)_*}^m(t) - \Delta, 0 \right\}, \quad \text{and}$$

$$P_{(jb)^*}^m(t+1) = P_{(jb)^*}^m(t) + \min \left\{ \Delta, P_{(jb)_*}^m(t) \right\}.$$

By intuition, the algorithm reallocates power to the beams with large positive utility-sensitivity. Note that for numerical reasons, it may be necessary to specify a certain minimum power per beam P_b^{\min} instead of allowing beam powers to be reduced to zero. In this case, the changes to the algorithmic notation above are straight forward, so we do not explicitly state them here.

3.3 Cost-based algorithm

Since CBA enforces average power constraints per beam, the following additional constraint to the NUM problem (1 to 4) is introduced.

$$\forall b : \bar{P}_{jb}^m \geq \sum_l \pi(l) \sum_k \phi_{jk}^{lm} c_{jb}(k, \bar{P}_j^m) \quad (16)$$

It includes cost term $c_{jb}(k, \bar{P}_j^m)$ which represents the power cost or power consumption of beam b (on PRB j) given scheduling decision k and \bar{P}_j^m , which is the total power budget of PRB j . We assume that each beam that is activated gets an equal share of the available total PRB power \bar{P}_j^m . Thus, if $n_j(k)$ is the number of beams that are activated on PRB j if decision k is chosen, the ‘cost’ of activating beam b on PRB j becomes

$$c_{jb}(k, \bar{P}_j^m) = \frac{1}{n_j(k)} \bar{P}_j^m,$$

with $\bar{P}_j^m = \sum_b \bar{P}_{jb}^m$.

To solve the modified problem (1 to 4 and 16), we also have to modify the gradient scheduler (5 to 6). The modified gradient scheduler now chooses the scheduling decision k^* according to

$$k^* \in \arg \max_k \left[\nabla_X U^m(\bar{X}^m(t)) \sum_b \mu_{jb}^{l(t)}(k) - \sum_b \lambda_{jb}(t) c_{jb}(k, \bar{P}_j^m) \right]. \quad (17)$$

Dual parameters $\lambda_{jb}(t)$, measuring the deviation of powers over time from the target power values on a particular beam, are updated according to the following rule:

$$\lambda_{jb}(t+1) = \left[\lambda_{jb}(t) + \beta_3 \left(c_{jb}(k^*, \bar{P}_j^m) - \bar{P}_{jb}^m(t) \right) \right]^+. \quad (18)$$

Average user rates $\bar{X}^m(t)$ are maintained and updated as in (6). The above algorithm can be seen as an application of the greedy primal dual (GPD) algorithm presented in [29].

The virtual control plane differs from VSA in the following. In contrast to the VSA, where every beam is switched on all the time, the situation is different here. To enable the calculation of derivatives of the rates with respect to

beam powers (needed in (22)), we introduce scaling factors $\alpha_{jb}^m(k)$, which scale target beam powers \bar{P}_{jb}^m to powers ‘costs’ $C_{jb}^m(k)$ that are instantaneously used by the virtual scheduler. Thus, $\alpha_{jb}^m(k) \bar{P}_{jb}^m = C_{jb}^m(k)$.

Given average gains (10) as in VSA, the sector controllers calculate (virtual) user rates given by

$$R_{ijb}^m(k) = \rho \left(F_{ijb}^m(k) \right), \quad (19)$$

with

$$F_{ijb}^m(k) = \frac{\alpha_{jb}^m(k) G_{ijb}^m \bar{P}_{jb}^m}{\sigma^2 + I_{\text{intra}} + I_{\text{inter}}}, \quad (20)$$

$I_{\text{intra}} = \sum_{b' \neq b} \alpha_{jb'}^m(k) G_{ijb'}^m \bar{P}_{jb'}^m$, and $I_{\text{inter}} = \sum_{\hat{m} \neq m} \sum_{\hat{j}} G_{ij\hat{j}}^{\hat{m}} \bar{P}_{j\hat{j}}^{\hat{m}}$.

Virtual average user rates are calculated by CBA as follows:

$$X_i^m = \sum_j \sum_k \tilde{\phi}_{jk}^m \sum_b R_{ijb}^m(k). \quad (21)$$

Again, $\tilde{\phi}_{jk}^m$ are optimal time fractions of resource usage for sector m ; however, in contrast to VSA where those time fractions were calculated explicitly, CBA uses the approach of OA to determine the virtual average rates (and implicitly the time fractions) through *virtual scheduling*. As before, to distinguish real and virtual scheduler, we use capital letters for virtual scheduler quantities whenever possible. In each TTI, the virtual scheduler performs n_v scheduling runs. In each run, the following steps are carried out on each PRB j :

- We determine the *virtual* scheduling decision k^* similar to (17).
- We update *virtual* average rates similar to (6).
- We update the *virtual* average power costs for each beam b similar to (18).
- We update sensitivities for each beam b and sector \hat{m} ($\beta_2 > 0$ small) according to

$$D_{jb}^{(\hat{m}, m)} \left(\frac{t+1}{n_v} \right) = (1 - \beta_2) D_{jb}^{(\hat{m}, m)} \left(\frac{t}{n_v} \right) + \beta_2 \sum_{i=1}^{I_m} \frac{\partial U^m(X^m)}{\partial X_i^m} \sum_{b'} \frac{\partial R_{ijb'}^m(k^*, \frac{t}{n_v})}{\partial \bar{P}_{jb'}^m(t)}. \quad (22)$$

Using (19) to (20), the derivatives in (22) are given by

$$\frac{\partial R_{ijb'}^m}{\partial \bar{P}_{jb'}^m} = \begin{cases} \rho' \left(F_{ijb'}^m \right) \frac{F_{ijb'}^m}{\bar{P}_{jb'}^m} & \hat{m} = m, b' = b \\ -\rho' \left(F_{ijb'}^m \right) \frac{\left(F_{ijb'}^m \right)^2}{\bar{P}_{jb'}^m} \frac{G_{ijb}^m \alpha_{jb}^m}{G_{ijb'}^m \alpha_{jb'}^m} & \hat{m} = m, b' \neq b \\ -\rho' \left(F_{ijb'}^m \right) \frac{\left(F_{ijb'}^m \right)^2}{\bar{P}_{jb'}^m} \frac{G_{ijb}^m \alpha_{jb}^{\hat{m}}}{G_{ijb'}^m \alpha_{jb'}^{\hat{m}}} & \hat{m} \neq m \end{cases}$$

The power adaption is then carried out similar to the other algorithms. For each resource, each sector sums up values $D_{jb}^{(m,m)}$ from all sectors and increases the power level on the beam with largest positive sum while decreasing the power level on the beam with largest negative sum. However, CBA adapts only power constraints per beam, not actually used beam powers.

Figure 3 gives a sketch of power trajectories using the simulation environment described below, in Section 6.1. To find out whether CBA really holds the average power constraints, we exponentially average instantaneously used powers per beam with the same time constant used for scheduling and compare the result to the target power values determined by the virtual model. The left side of Figure 3 shows the averaged powers, as actually used by the ‘real’ scheduler, while the right side shows the target power values determined by the virtual scheduling procedure. Note that since scheduling is opportunistic, instantaneously, the power levels fluctuate highly and beam powers can differ from the target values (or a beam can be completely turned off). This is illustrated by the ‘zoomed-in image’ in Figure 3 (left), where actual powers without averaging are shown. It turns out that *on average*, the power constraints are kept remarkably well.

Apart from the intuitive benefits of instantaneously allowing opportunistic scheduling, we observe that from a practical point of view, average power constraints are further justified since hybrid automatic repeat request (HARQ) coding is essentially performed over multiple successive transmissions.

In all three presented algorithms, naturally, questions arise regarding differentiability. Note that the problem to be solved by each sector controller is given in (13). Due to the maximum operator, it might not be differentiable everywhere even with the utility function being differentiable. Therefore, let us define a slightly modified version of problem (13), given by

$$f_\varepsilon(R) := \max_{\phi_{ijb}^m} \sum_i U^m \left(\sum_j \sum_b (\phi_{ijb}^m)^{1-\varepsilon} R_{ijb}^m \right). \quad (23)$$

One can show the following:

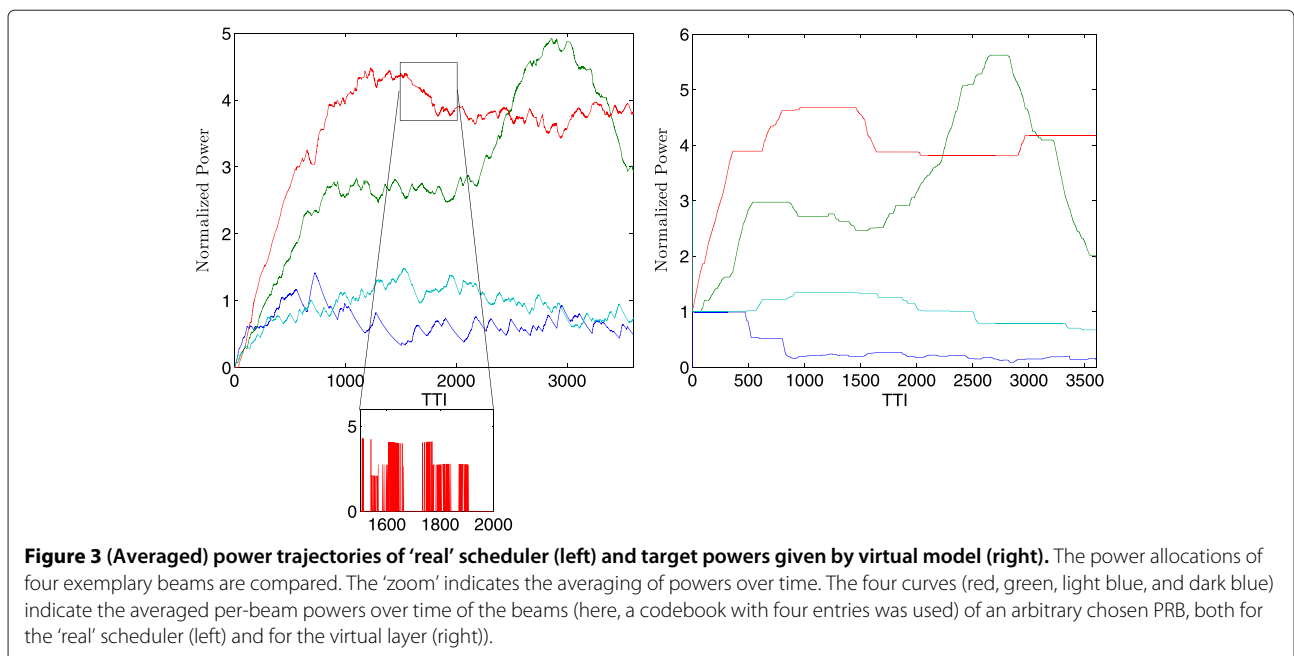
Theorem 1. *Let $0 < \varepsilon < 1$ be finite and U^m be an increasing concave utility function, defined in $(0, \infty)$. Then, the family of functions $f_\varepsilon(R)$ (defined by (23)) with $R_{ijb} \geq c > 0$ ($\forall i, j, b$) is differentiable everywhere and converges for any sequence $\varepsilon_n \rightarrow 0$ to f in (13) (which is continuous) in the uniform sense.*

Proof. The proof can be found in Appendix 1. \square

By Theorem 1, we can replace our utility function with a smooth, uniformly convergent approximation, which can be locally maximized in the power control loop. Note that this replacement is already incorporated in the calculation of sensitivities (Equation 14).

4 An alternating optimization-based approach

Given the distributed nature of the above presented algorithms, the question arises: Can a centralized controller do



better? Therefore, in the following, we want to compare the algorithms of Section 3 with a centralized solution.

The optimization problem which the virtual controller has to solve is given by

$$\max_{\mathbf{P}, \Phi} \sum_{m=1}^M \sum_{i=1}^{I_m} \log \left(\sum_{j=1}^J \sum_b \phi_{ijb}^m R_{ijb}^m(\mathbf{P}) \right) \quad (24)$$

$$\text{s.t. } \forall b, j, m : \sum_i \phi_{ijb}^m \leq 1 \quad (25)$$

$$\forall m : \sum_b \sum_j P_{jb}^m \leq P_{\max}, \quad (26)$$

with $R_{ijb}^m(\mathbf{P}) = \log(1 + F_{ijb}^m(\mathbf{P}))$ and F_{ijb}^m defined in (11). As described in Section 3, the power allocation problem is so far solved using a distributed gradient ascent procedure. However, when we allow a centralized controller for the network, we can instead solve (24 to 26) directly each time a power update is desired and use the resulting power allocation directly for actual resource allocation.

Obviously, problem (24 to 26) is highly non-convex. In the following, we try to solve the problem by alternating the optimization in \mathbf{P} (holding Φ constant) and Φ (holding \mathbf{P} constant). We therefore have a scheduling sub-problem and a power allocation (PA) sub-problem. The overall procedure is summarized in Algorithm 1.

Algorithm 1 Alternating optimization-based virtual layer

- 1: initialize counter $\tau = 0$, select feasible $\Phi(0)$ and $\mathbf{P}(0)$ arbitrarily
 - 2: **repeat**
 - 3: $\tau \leftarrow \tau + 1$
 - 4: solve scheduling sub-problem using fixed $\mathbf{P}(\tau - 1)$ and obtain $\Phi(\tau)$
 - 5: solve SCA-based power allocation sub-problem using fixed $\Phi(\tau)$ and obtain $\mathbf{P}(\tau)$ (Algorithm 2)
 - 6: **until** converged
 - 7: use $\mathbf{P}(\tau)$ to update power allocation in network
-

Since the scheduling sub-problem is convex (cf. Lemma 2), we only have to care about the power allocation sub-problem. We try to tackle this problem by a *successive convex approximation* (SCA) approach similar to [17]. The sub-problem in \mathbf{P} is still highly non-convex. However, using Lemma 3, we obtain a convexified version of the power allocation sub-problem.

Lemma 2. *With constant \mathbf{P} , optimization problem (24-26) is a convex optimization problem in Φ .*

Proof. The proof follows since non-negative weighted addition and scalar composition preserve concavity [31]. \square

Lemma 3. *Using a concave lower bound (assuming appropriately chosen constants) to the user rates, given by*

$$\tilde{R}_{ijb}^m := \alpha_{ijb}^m \log(F_{ijb}^m) + \beta_{ijb}^m \leq R_{ijb}^m, \quad (27)$$

(with equality holding when approximation constants are chosen as $\alpha_{ijb}^m = \frac{F_{ijb}^m}{1+F_{ijb}^m}$ and $\beta_{ijb}^m = \log(1 + F_{ijb}^m) - \alpha_{ijb}^m \log(F_{ijb}^m)$) and a logarithmic change of variables given by $\tilde{P}_{jb}^m = \log(P_{jb}^m)$, the modified PA sub-problem

$$\max_{\tilde{\mathbf{P}}} \sum_{m=1}^M \sum_{i=1}^{I_m} \log \left(\sum_{j=1}^J \sum_b \phi_{ijb}^m \tilde{R}_{ijb}^m(\tilde{\mathbf{P}}) \right) \quad (28)$$

$$\forall m : \sum_b \sum_j e^{\tilde{P}_{jb}^m} \leq P_{\max} \quad (29)$$

is a convex optimization problem.

Proof. The proof can be found in Appendix 2. \square

The SCA procedure is summarized in Algorithm 2.

Algorithm 2 SCA-based power allocation sub-problem

- 1: initialize counter $\tau = 0$, $\alpha(0) = \mathbf{1}$; $\beta(0) = \mathbf{0}$
- 2: **repeat**
- 3: solve (28 to 29), obtain $\tilde{\mathbf{P}}^*(\tau)$
- 4: tighten approximation ($\forall i, j, b, m$):

$$\alpha_{ijb}^m(\tau + 1) = \frac{F_{ijb}^m(\mathbf{P}^*(\tau))}{1 + F_{ijb}^m(\mathbf{P}^*(\tau))} \quad (30)$$

$$\beta_{ijb}^m(\tau + 1) = \log \left(1 + F_{ijb}^m(\mathbf{P}^*(\tau)) \right) - \alpha_{ijb}^m \log \left(F_{ijb}^m(\mathbf{P}^*(\tau)) \right)$$

- 5: $\tau \leftarrow \tau + 1$
 - 6: **until** converged
-

The algorithm is initialized as described in the first step of Algorithm 2, which is equivalent to a high signal-to-interference-plus-noise ratio (SINR) approximation (which can be seen when applying this initialization in (27)). The high-SINR approximation assumes that for large SINR, $\log(1 + \text{SINR}) \approx \log(\text{SINR})$. This is a common initialization for this kind of algorithm [17,32]. Each iteration of the algorithm comprises two steps, a maximize step and a tighten-step. In the maximize step, we find a solution to the current convexified version (28 to 29) of the power control problem. This solution is then used in the tighten step to update the convex approximation parameters α and β for each link according to (30).

The algorithm converges when the tighten step (30) does not produce any (significant) changes. Being based on an inner approximation framework by Marks and Wright [33], it can be shown that Algorithm 2 converges at least to a KKT point of the PA sub-problem.

5 Approaching global optimality: comparison with branch-and-bound

It is known that the underlying optimization problem is non-convex; thus, the gradient ascent-based algorithms presented in Section 3 as well as the alternating optimization-based algorithm of Section 4 will most likely converge to a local maximum. Thus, although simulation results (cf. Section 6) show already high gains in utility, the question remains how good the solution found actually is, that is how much of the achievable performance gains is actually realized?

To simplify the analysis in this section, we restrict ourselves to the single-antenna case (which reduces all algorithms to MGR in [14]). The main difficulty of the analysis is that even in this simple setting, the underlying problem

$$f(\mathbf{P}, \Phi) = \max_{\mathbf{P}, \Phi} \sum_{m=1}^M \sum_{i=1}^{I_m} \log \left(\sum_{j=1}^J \phi_{ij}^m R_{ij}^m \right) \quad (31)$$

$$\text{s.t. } \forall j, m: \sum_i \phi_{ij}^m \leq 1 \quad (32)$$

$$\forall m: \sum_j P_j^m \leq P_{\max} \quad (33)$$

with $R_{ij}^m = \log \left(1 + \frac{G_{ij}^m P_j^m}{\sigma^2 + \sum_{m' \neq m} G_{ij}^{m'} P_j^{m'}} \right)$, is highly non-convex (even in the two-sector, two-user case). Moreover, since this is essentially a joint optimization in time fractions Φ and power values \mathbf{P} , the complexity is significantly higher than with power control only. To gain insight into the deviation from the optimal solution of (31 to 33), we compare our algorithm's performance in the simplified setting with a (near) optimal solution based on BNB. Although computationally very expensive, it gives us a good impression on how much of the achievable performance is actually reached. We assume, without loss of generality, that the maximum sum power available in each cell in (33) is normalized to one.

Branch-and-bound is a standard algorithm for global optimization, which creates a search tree where at each node, an upper and a lower bound to the problem are evaluated. Details can be found for example in [34]. It is based on constantly sub-dividing the feasible parameter region and for each node of the resulting tree calculating the upper and lower bounds to the objective function. For ease of notation, we will combine all parameters to our objective function in one vector $\hat{\mathbf{P}} = (p_1, \dots, p_n)$. Let $\mathcal{P}^{(0)}$ denote our initial parameter region. This region forms a

convex n -dimensional polytope with V extreme points (or vertices) $\hat{\mathbf{P}}_v$. We collect the corner points of polytope $\mathcal{P}^{(0)}$ in set $\mathcal{V}^{(0)}$. Consequently, the parameter region associated with node l is denoted as $\mathcal{P}^{(l)}$ with associated vertices $\mathcal{V}^{(l)}$.

Due to the structure of the constraints, we have $\frac{n}{2}$ parameter pairs with independent constraints. Each of the parameter/constraint pairs form a (two-dimensional) simplex. When branching, we split the parameter region of node l , $\mathcal{P}^{(l)}$, in sub-regions $\mathcal{P}^{(l+1)}$ and $\mathcal{P}^{(l+2)}$ by splitting the longest edge among the edges of all ($\frac{n}{2}$) simplexes together. Doing so, we divide the parameter region into two sub-regions of equal size. Figure 4 visualizes the above-stated procedure by showing the simplex created by two interdependent power values together with an independent third dimension. The dashed line indicates the splitting into two sub-polytopes.

Below, we explain how the upper and lower bounds to the sub-problems are found. It is well known that the Shannon rate can be written as a difference of concave functions $f_{ij}^m(\mathbf{P}) - g_{ij}^m(\mathbf{P})$, with $f_{ij}^m(\mathbf{P}) = \log \left(\sigma^2 + \sum_{m'=1}^M G_{ij}^{m'} P_j^{m'} \right)$ and $g_{ij}^m(\mathbf{P}) = \log \left(\sigma^2 + \sum_{m' \neq m} G_{ij}^{m'} P_j^{m'} \right)$. Using this, our non-concave objective function (31) can be rewritten as

$$\hat{F}(\hat{\mathbf{P}}) = \sum_{m=1}^M \sum_{i=1}^{I_m} \log \left(\sum_{j=1}^J \phi_{ij}^m f_{ij}^m(\hat{\mathbf{P}}) - \phi_{ij}^m g_{ij}^m(\hat{\mathbf{P}}) \right). \quad (34)$$

Unfortunately, $\phi_{ij}^m f_{ij}^m(\hat{\mathbf{P}})$ (and thus also $\phi_{ij}^m g_{ij}^m(\hat{\mathbf{P}})$) is neither convex nor concave (note that $f(a, b) = a \cdot \log(1 + b)$

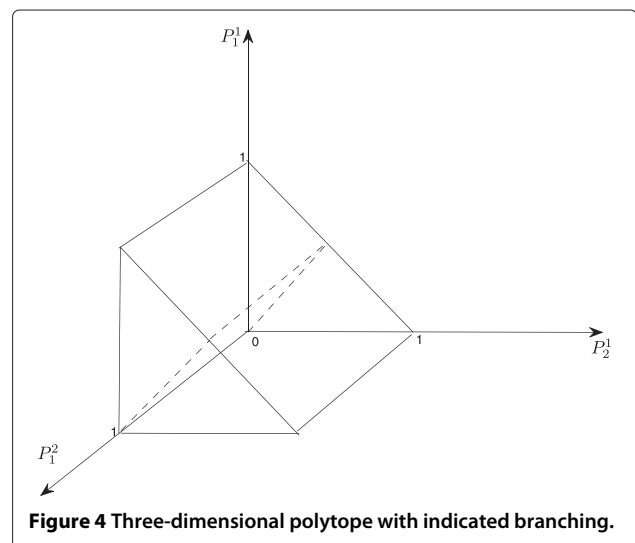


Figure 4 Three-dimensional polytope with indicated branching.

is not concave). However, we can define the concave function(s): $\tilde{f}_{ij}^m(\hat{\mathbf{P}}) = \phi_{ij}^m \log\left(\sigma^2 + \frac{1}{\phi_{ij}^m} \sum_{m'=1}^M G_{ij}^{m'} P_j^{m'}\right)$, and in the same way $\tilde{g}_{ij}^m(\hat{\mathbf{P}})$. Using this, (34) can be upper bounded by

$$\sum_{m=1}^M \sum_{i=1}^{I_m} \log\left(\sum_{j=1}^J \tilde{f}_{ij}^m(\hat{\mathbf{P}}) - \tilde{g}_{ij}^m(\hat{\mathbf{P}})\right). \quad (35)$$

Equation 35 is still not concave since it includes a sum of a concave function ($\tilde{f}_{ij}^m(\hat{\mathbf{P}})$) and a convex function ($-\tilde{g}_{ij}^m(\hat{\mathbf{P}})$). To get a concave objective function, we replace $\tilde{g}_{ij}^m(\hat{\mathbf{P}})$ by its *convex envelope* (cf. [34]), defined as follows:

Let $\mathbf{P}_1, \dots, \mathbf{P}_V$ be the vertices of a polytope \mathcal{P} . The convex envelope $\gamma(\mathbf{P})$ of a concave function $g(\mathbf{P})$ over \mathcal{P} can be expressed as

$$\gamma(\mathbf{P}) = \min_{\lambda} \sum_{v=1}^V \lambda_v g(\mathbf{P}_v),$$

subject to $\sum_{v=1}^V \lambda_v \mathbf{P}_v = \mathbf{P}$, $\sum_{v=1}^V \lambda_v = 1$, and $\lambda_v \geq 0$. Using this definition, the (convex) optimization problem we have to solve in order to obtain our upper bound (at node l) is given by

$$F^U = \max_{\lambda} \sum_{m=1}^M \sum_{i=1}^{I_m} \log\left(\sum_{j=1}^J \tilde{f}_{ij}^m(\hat{\mathbf{P}}) - \sum_{v=1}^V \lambda_v g(\hat{\mathbf{P}}_v)\right)$$

$$\text{s.t. } \hat{\mathbf{P}} = \sum_{v=1}^V \lambda_v \hat{\mathbf{P}}_v, \sum_{v=1}^V \lambda_v = 1, \lambda_v \geq 0 \ (\forall v),$$

with $\lambda = (\lambda_1, \dots, \lambda_V)$, $\hat{\mathbf{P}} \in \mathcal{P}^{(l)}$, $\hat{\mathbf{P}}_v \in \mathcal{V}^{(l)}$, and V being the number of vertices collected in $\mathcal{V}^{(l)}$. We use three different approaches to obtain a lower bound at node l . The first lower bound is obtained by taking the maximum of the objective function values at each of the corner points of the respective parameter region. Second, we evaluate the optimal parameter vector from calculating the upper bound. Third, we use a standard solver to compute a (local) optimum of the original non-convex problem. If one of the three methods leads to a higher value, the current global lower bound is replaced.

6 Numerical evaluation

In this section, we present numerical results to evaluate the distributed algorithms of Section 3 as well as the centralized and BNB-based solutions of Section 4 and Section 5, respectively.

6.1 System-level simulations

In order to compare the performance of the three algorithms in a setting close to practice, we conduct system-level simulations based on LTE.

6.1.1 Simulation setup

The general simulation setup can be found in Table 2. We employ a grid of seven hexagonal cells, each comprising three sectors (to ensure equal interference conditions for each sector, a wrap-around model is used at the system borders). Users are distributed randomly over the whole area; thus although we have, say, 210 user in total, leading to an average of ten users per sector, some sectors have a higher load than others. A detailed description of the precoding codebook can be found in [35]. As channel model, we use the WINNER model [36], Scenario C2. All simulations are performed with realistic link adaptation based on mutual information effective SINR mapping (MIESM); moreover, we use explicit modeling of HARQ (hybrid automatic repeat request) using chase combining. For efficiency reasons, we simulate a limited number of PRBs ($J = 8$) compared to a real-world system; however, the results are expected to scale to a higher number of PRBs.

6.1.2 Baseline algorithm

As baseline, we use a non-coordinative scheduling algorithm called greedy beam distance (GBD) algorithm, with a codebook size of 3 bits (eight beams). GBD requires feedback from each user, comprising a CDI and a channel quality information (CQI) value. Note that the CQI is only an estimate of the users' SINR values in case the user alone is scheduled on the PRB (with full power), since the scheduling decisions cannot be known in advance. Thus, it does not contain intra-sector interference. Moreover, it contains only an estimate of the inter-sector interference. On each PRB, the users are greedily scheduled on their best beams (using their proportionally fair weighted CQI feedback as utility). Thereby, a minimum beam distance has to be kept, in order to minimize the interference between users scheduled on the same PRB. This distance

Table 2 General simulation setup

Parameter	Value
Number of sectors (M)	21
Total number of terminals (l)	105 to 315
Mobile terminal velocity	0 km/h, 3 km/h
Number of PRBs (J)	8
Number of beams (B) (coordination)	4
Number of beams (B) (baseline)	8
Base station antennas (n_T)	4
Number of terminal antennas (n_R)	1
Simulation duration	10,000 TTI
Power adaption step size (Δ)	0.5% of the initial power
Traffic model	Full buffer

This table summarizes the main parameters that we use in the system-level simulations.

is based on a geometrical interpretation of beamforming. It means that given a user is scheduled on a certain beam, adjacent beams (up to a certain 'distance') are blocked and users that reported one of those beams are excluded from the list of scheduling candidates for the respective PRB. We use a minimum distance of 3, that is, with the used eight-beam codebook at most three users can be scheduled on the same PRB. Of course, no adaptive power allocation is performed; the power is distributed equally among the PRBs (and further among the thereon scheduled users).

6.1.3 Simulation results

We use two essential performance metrics for our comparison. First, since we have a system-wide proportional fair utility, we compare the geometric mean of average user rates (GAT) as a measure of the increase of the *overall* utility. The equivalence of maximizing GAT and sum utility can be seen when observing that the geometric mean of some set of values is equal to the exponential average of the logarithm of the values. Therefore, the sum of logarithms is maximized precisely if the geometric mean is maximized. Second, the performance of cell-edge users, which we measure by the 5% quantile of average user throughputs, is a natural benchmark for each distributed base station coordination algorithm.

Figure 5 shows cell-edge user throughput over GAT for the evaluated algorithms with user mobility (and therefore with fast fading), while Figure 6 depicts the simulation results in a setting without user mobility. We compare the three approaches of Section 3 to the uncoordinated baseline of Section 6.1.2. Thereby, for each algorithm, the performance is evaluated for different average number of

users per sector. More precisely, from left to right, the markers in the figure represent an average number of 5, 8, 10, 12, and 15 users per sector (corresponding to total numbers of 105, 168, 210, 252, and 315 users in the network, respectively). Obviously, increasing values on the ordinate corresponds to increased fairness (with respect to our 5% quantile metric), while higher values on the abscissa correspond to a higher sum utility (with respect to our GAT metric). In both figures, we observe that a higher number of users lead to increased sum utility (an effect called multiuser diversity) while the cell-edge users (represented by the 5% of users with lowest throughput) suffer from a reduced performance (due to an increased competition for resources).

In case of user mobility (Figure 5), we observe that CBA clearly outperforms the other algorithms. In fact, for every simulated user density, either the 5% fairness (low user densities) or both metrics are improved (higher user densities). For example, in the case of 210 users in the network (corresponding to the third marker on the curves in Figure 5), the CBA algorithm improves the performance with respect to GAT by about 10% and in cell-edge user throughput by more than 35%. It can also be observed that the VSA algorithm works best at high user densities, while in case of only a few users in the cell, the OA algorithm (which requires the least overhead with respect to additional feedback and signaling) and even the no-power control baseline show a better performance, at least with respect to the cell-edge performance. The decreased GAT shown by VSA in the fading case is basically the cost of leaving no freedom to the schedulers for opportunistic scheduling, but to strictly specify the powers to be used for all beams and time instances.

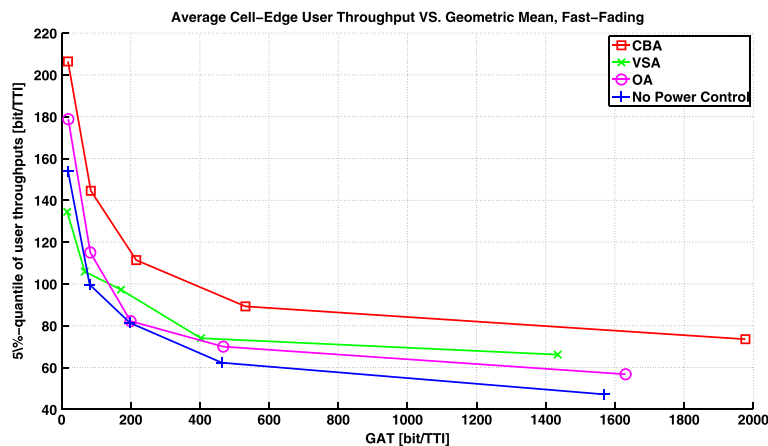


Figure 5 Cell-edge user throughput vs. GAT in a fading environment. The averaged cell-edge users' throughputs vs. the geometric mean of users' throughputs are plotted for a fast-fading scenario with users moving at 3 km/h. Thereby, the proposed three distributed algorithms are compared with a baseline algorithm without coordination. For each algorithm, the total number of users in the network is varied. In particular, the markers in the figure, from left to right, represent an average number of 5, 8, 10, 12, and 15 users per sector. This corresponds to total number of 105, 168, 210, 252, and 315 users in the network, respectively.

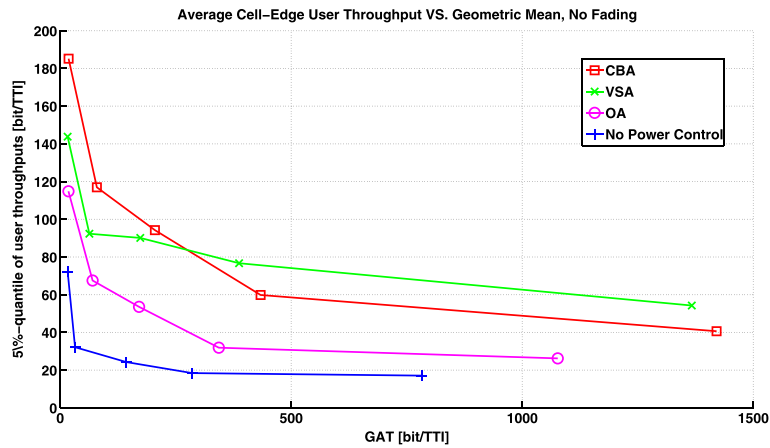


Figure 6 Cell-edge user throughput vs. GAT in a static environment. The averaged cell-edge users' throughputs vs. the geometric mean of users' throughputs are plotted for a scenario without user mobility. Thereby, the proposed three distributed algorithms are compared with a baseline algorithm without coordination. For each algorithm, the total number of users in the network is varied. In particular, the markers in the figure, from left to right, represent an average number of 5, 8, 10, 12, and 15 users per sector. This corresponds to total number of 105, 168, 210, 252, and 315 users in the network, respectively.

Figure 6 depicts the simulation results in a setting without user mobility. We can see that OA offers already quite high performance gains, both in GAT and in cell-edge user throughput, although being the least complex algorithm. However, the performance is even further improved by the other two algorithms. VSA can improve especially the gains of cell-edge users, for example, in the case of 210 users in the network to more than 200% compared with no-power control. However, CBA again shows the best performance, significantly improving the global utility compared with VSA (and even more compared with no-power control). Again, this gain increases with increasing number of users in the network.

Besides the performance, the speed of convergence is of large practical interest. To illustrate this, Figure 7 shows the GAT metric for all algorithms over time for

the case of stationary users (with 210 users in total). It can be seen that the fastest convergence is achieved by the OA and VSA algorithms, while the CBA algorithm, which in the end achieves the highest GAT, needs about 3,000 TTIs to converge. This is expectable, since the scheduler has to follow the power values obtained from the virtual layer only on average; thus, it always 'lags behind' the adaptation of the virtual model. Regardless of which algorithm performs best, it is of general importance that distributed power control algorithms do not only work in a specific environment but can adapt to changing network conditions (which, of course, is also related to the notion of convergence). Figure 8 demonstrates this autonomous adaptation ability by showing the cell-edge user throughput over time (exemplary for CBA in the static case) for multiple consecutive drops,


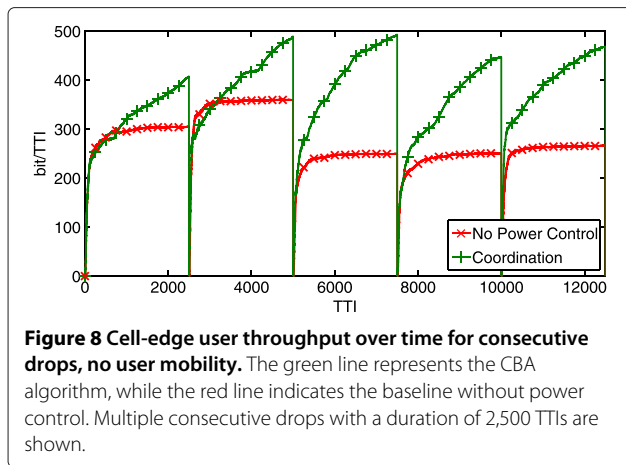


Figure 7 GAT over time in a static environment. The geometric mean of users' throughputs is plotted over the first 4,000 TTIs of the simulation for a scenario without user mobility. Thereby, the proposed three distributed algorithms are compared with a baseline algorithm without coordination.



where, in each drop, a new user location and assignment pattern are generated. It can be observed that the the performance is improving very fast compared to the baseline.

6.2 Performance of centralized solution based on alternating optimization

Subsequently, we investigate the performance of the centralized scheme based on alternating optimization (AO), as introduced in Section 4. For this, we compare the performance with that of the distributed solution. More precisely, we use the VSA algorithm of Section 3.2. Note that the only difference between the subsequently investigated algorithms is the power adaptation procedure. While in the distributed solution (in subsequent figures denoted as VSA), the power is adapted gradually based on the exchange of sensitivities, in the centralized solution (in subsequent figures denoted as AO), the solution to the optimization problem in the virtual layer

is directly used to globally update the power allocation. We are predominantly interested in the convergence speed of the geometric mean of average user rates in the different approaches and the relative performances in different signal-to-noise ratio (SNR) regimes. Moreover, we investigate both a static setting without mobility and a setting where a user moves at a velocity of 3 km/h.

In Figure 9, we compare the behavior of the AO and the distributed approach for constant channels. In general, it can be observed that the convergence speed of the average rates using AO is slightly faster. Often, both algorithms converge to the same solution; however, this is not necessarily always the case (as shown in Figure 9).

Figure 10 compares the performance of AO-based and gradient ascent-based solution in a setting with moderate user mobility. It can be observed that the convergence speed of the centralized scheme is still slightly higher than in the distributed case, and in most cases, the centralized solution outperforms the gradient-based approach (however, at a much higher complexity). An interesting effect that can be observed is that the influence of the initial state on the outcome of the distributed solution is significant. This is also depicted in Figure 10. Note that the only difference between the solid and the dashed red line in Figure 10 lies in the initial power values. In the first case, we start at a minimal power assigned to all sub-carriers (note that for technical reasons, it is not possible to assign exactly zero power to the resources) and in the latter case, we start at an allocation where power is distributed evenly among the resources.

In Figure 11, it can be observed that the gains from the centralized solution are higher in high-SNR regimes. While at low SNRs, when the system is essentially noise-limited, both approaches converge to the same solution; in high-SNR regimes, when the system is essentially




Figure 9 Comparison of the centralized and distributed virtual control without mobility. The geometric mean of the average user rates is shown in a scenario without mobility for the VSA algorithm (red) and the centralized baseline (blue).




Figure 10 Comparison of centralized and distributed virtual control at moderate mobility. The geometric mean of the average user rates is shown in a fast-fading scenario. The centralized baseline (blue) is compared to the VSA algorithm (red) with different initial power allocations.

interference-limited, the direct AO-based optimization outperforms the gradual power adaptation.

In summary, the centralized solution has some clear advantages. First, it is invariant to the initialization of the system and is thus able to adapt very fast to changing environmental conditions. Second, especially when the channels are varying fast and the SNR is high, the centralized scheme offers a higher performance in comparison to the distributed scheme. However, one should note that at realistic SNRs, the distributed solution still achieves almost the same performance and at the same time scales smoothly with the network size, while a centralized solution can only be implemented in very restricted scenarios.

6.3 Global optimality

In the following, we compare the performance of the distributed algorithm to an (nearly) optimal solution

based on BNB. To simplify the analysis in this section and due to the high complexity of the BNB algorithm, we restrict ourselves to the simple case of two sectors with two users (no mobility) and two PRBs and a single antenna.

Given a certain tolerance ε , BNB converges to an optimal solution of (31). However, it has exponential complexity and is therefore only feasible in very small settings. We compare our distributed coordination algorithm with an equal power non-cooperative scheduler and the BNB algorithm described in Section 5.

In Figures 12 and 13, the blue line represents the GAT obtained with the coordination algorithm, while the red line represents the outcome of a non-cooperative equal-power scheduler. The solid black line indicates the outcome of BNB. Finally, the dotted black line depicts the configured ε tolerance. Thus, the true optimum is guaranteed to lie in between the solid and dotted black lines. To get an overview of the performance in different scenarios, we investigate three different settings, namely, a setting with weak interference, a setting with approximately equal interference, and a setting with high interference. In Figure 12, the left marker corresponds to the weak interference case, the middle marker to the case of equal interference, while the right marker corresponds to the case of strong interference. The weak interference is characterized by $G_{ij}^m > G_{ij}^{m'}$ ($m \neq m'$) for all i, j, m, m' ; thus, for all users, the link gains to the serving base station are higher than the gains to the interfering base station. Note that this is, from a practical point of view, the more interesting case since due to handover algorithms, most of the time, the gains to the own base station are stronger than to the interferers. It can be observed that the distributed power control performs quite well since we see a significant improvement over equal power scheduling. Moreover, the distributed power control

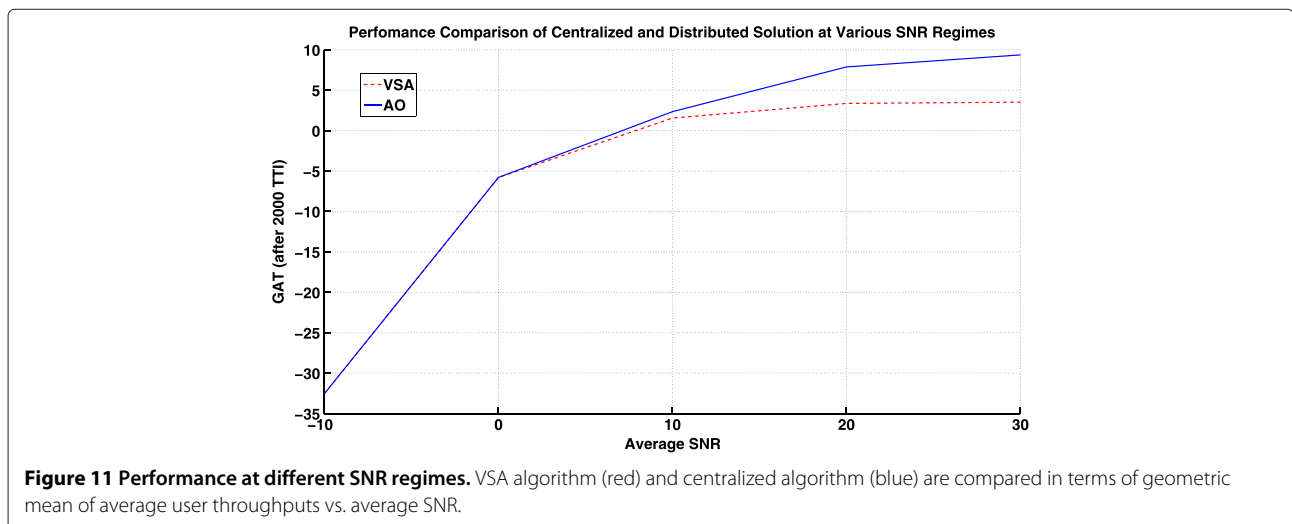
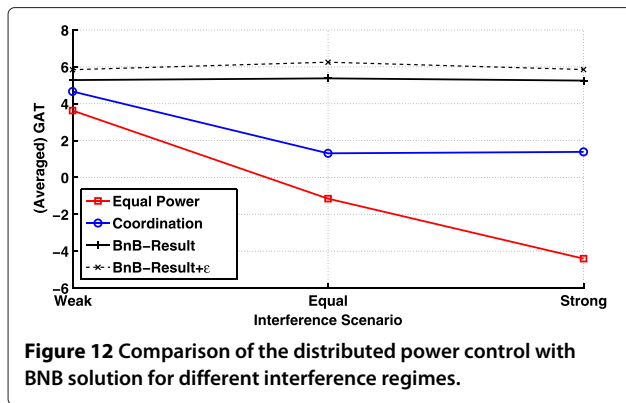


Figure 11 Performance at different SNR regimes. VSA algorithm (red) and centralized algorithm (blue) are compared in terms of geometric mean of average user throughputs vs. average SNR.

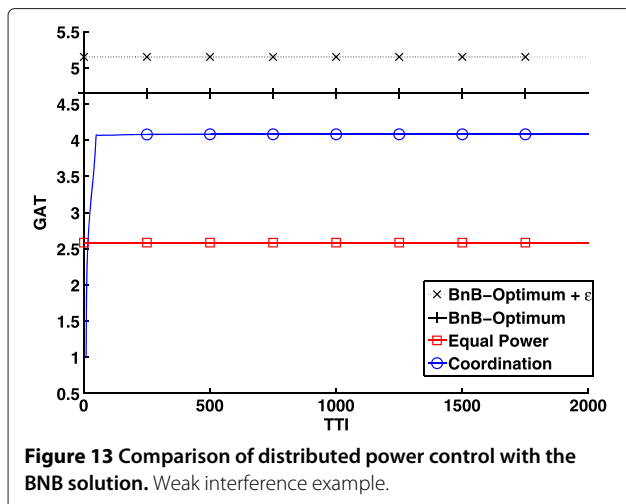


appears to realize already a large portion of the available gains.

In the ‘equal’ interference case, the gains to the interfering base stations are roughly the same as the gains to the own base stations; thus, $G_{ij}^m \approx G_{ij}^{m'}$ ($m \neq m'$) for all i, j, m, m' . It can be observed that the power control algorithm leaves a larger portion of the available performance gains unused, compared to the weak interference case, although it shows significant improvements over non-cooperative scheduling.

The strong interference case is characterized by $G_{ij}^m < G_{ij}^{m'}$ ($m \neq m'$), for all i, j, m, m' ; thus, for all users, the link gains to the interfering base stations are higher than the gains to the ‘own’ base stations. Here, the gap in the non-cooperative algorithm’s performance, compared to the BNB result, is significantly worse than in the other cases. By contrast, the power control algorithm shows roughly the same performance gap than in the case of equal gains.

Again, to illustrate the convergence, Figure 13 gives an example of the performance over time (here, for the weak



interference case). Obviously, the power control algorithm converges to a local maximum soon.

7 Conclusions

We proposed and compared three distributed algorithms for autonomous interference coordination in cellular SDMA networks. The algorithms are based on a virtual layer that models the interference interdependencies in the network and gradually adapts power control levels. The proposed algorithms differ in granularity of power control, required feedback, signaling overhead, and the virtual model itself. System-level simulations indicate high gains both in overall utility and in cell-edge user throughput for all three algorithms in static environments without user mobility. While VSA offers a very fine granularity of power control on a per-beam level, it suffers from the lack of freedom to instantaneously perform power allocation in an opportunistic manner. In fact, in an environment with significant user mobility, only CBA achieves significant gains in both metrics. This demonstrates the superiority of CBA’s approach to enforce average power constraints but instantaneously allowing opportunistic scheduling. Comparisons to a centralized benchmark scheme reveal that although the convergence of the centralized scheme is much faster than in the distributed case, the performance in overall utility is comparable.

Endnotes

^aNote, that in this paper, we use the term autonomous not in the sense of an autonomous operation of the different network entities, such as base stations, but to the ability of the network to find a suitable operating point without the needed of prior planning or human interaction.

^bEach PRB consists of a fixed number of OFDM symbols in time and has a total length of 1 TTI.

^cWe use $\sum_i \log(\bar{X}_i^m)$ as sector utility in this paper, leading to a proportional fair operating point.

^dTo reduce messaging overhead, sector controllers could limit the message exchange to strongest interferers, e.g., consider only neighboring base stations.

^eTo avoid confusion, we denote variables belonging to the real scheduler with lower case symbols, and variables from the virtual scheduler with corresponding upper case symbols.

Appendices

Appendix 1: proof of Theorem 1

First, we show that $f_\varepsilon(R)$ is everywhere differentiable. For this purpose, we use the following Lemma 4 and set $H(x, y) := U(R, \Phi)$.

Lemma 4. Consider a function $H(x, y)$, $x \in [a_1, a_2]$, $y \in [a_3, a_4]$, with some finite $a_1 < a_2$ and $a_3 < a_4$. Assume

that $H(x, y)$ and its partial derivative on x are continuous, and $H(x, y)$ is concave in x (for each y) and strictly concave in y (for each x). Then, the function

$$\varphi(x) = \max_y H(x, y) \quad (36)$$

is continuously differentiable in $[a_1, a_2]$. (This can be generalized to the multi-dimensional case. Namely, the domains $[a_1, a_2]$ and $[a_3, a_4]$ can be replaced by arbitrary convex compact sets in finite-dimensional vector spaces, and derivatives replaced by gradients).

Note, that Lemma 4 cannot be applied directly to (13) since strict concavity in y in this case is not given. Clearly, R and Φ are compact since $\phi_{ijb} \in [0, 1]$ and $R_{ijb} \in [c, B], B < \infty, c > 0$. Moreover, the concavity of U follows by definition. Now, we have to show that $f_\varepsilon(R)$ converges for any zero sequence $\varepsilon_n \rightarrow 0$ to f . First, fix a zero sequence ε_n^* . Moreover, we define a sequence of functions

$$f_n(R) := \max_{\phi_{ijb}^m} \sum_i U^m \left(\sum_j \sum_b (\phi_{ijb}^m)^{1-\varepsilon_n^*} R_{ijb}^m \right) \\ \hat{=} f_{\varepsilon_n^*}(R),$$

with $(n \in \mathbb{N})$. We can use Dini's Theorem to show that $\{f_n(R)\}_{n \in \mathbb{N}}$ converges uniformly to f . This theorem states that if $\{f_n\}_{n \in \mathbb{N}}$ ($f_n : K \rightarrow \mathbb{R}, n \in \mathbb{N}$, being a sequence of continuous functions and K being a compact metric space) converges pointwise to f ($f : K \rightarrow \mathbb{R}$ being a continuous function) and if $f_n(x) \geq f_{n+1}(x)$ ($\forall x \in K$ and $\forall n \in \mathbb{N}$), then $\{f_n\}_{n \in \mathbb{N}}$ converges uniformly to f .

The rate space $K := [c, B]^{IJN}$ is clearly compact, since $R_{ijb} \in [c, B], B < \infty, c > 0$. Moreover, for each $x^* \in K$, $f_n(x^*)$ converges to $f(x^*)$ when $n \rightarrow \infty$ (since $\varepsilon_n \rightarrow 0$), and since $\phi_{ijb} \in [0, 1]$, we have $\phi^{1-\varepsilon_n} \geq \phi^{1-\varepsilon_m}$ ($\forall i, j, b$) if $m \geq n$. Thus, $f_n(x) \geq f_{n+1}(x)$.

Appendix 2: proof of Lemma 3

Following the approach in [17], the convexity of optimization problem (28 to 29) can be easily shown by noting that \tilde{R}_{ijb}^m in (28) can be written as

$$\alpha_{ijb}^m \log(\tilde{F}_{ijb}^m) + \beta_{ijb}^m$$

where $\log(\tilde{F}_{ijb}^m)$ can be further decomposed into

$$\log(G_{ijb}^m) + \tilde{P}_{jb}^m - \log \left(\sigma^2 + \sum_{b' \neq b} G_{ijb'}^m e^{\tilde{P}_{jb'}^m} + \sum_{m' \neq m} \sum_b G_{ijb}^{m'} e^{\tilde{P}_{jb}^{m'}} \right).$$

Due to the convexity of log-sum-exp [31], the term

$$\log \left(\sigma^2 + \sum_{b' \neq b} G_{ijb'}^m e^{\tilde{P}_{jb'}^m} + \sum_{m' \neq m} \sum_b G_{ijb}^{m'} e^{\tilde{P}_{jb}^{m'}} \right)$$

is convex, and thus, \tilde{R}_{ijb}^m is concave. Noting that non-negative weighted addition and scalar composition preserve concavity concludes the proof.

Competing interests

The authors declare that they have no competing interests.

Author details

¹Fraunhofer Heinrich Hertz Institute, Einsteinufer 37, Berlin 10587, Germany.

²Technische Universität Berlin, Einsteinufer 25, Berlin 10587, Germany.

Received: 27 January 2014 Accepted: 28 June 2014

Published: 16 July 2014

References

1. D Lee, H Seo, B Clerckx, E Hardouin, D Mazzaresse, S Nagata, K Sayana, Coordinated multipoint transmission and reception in LTE-advanced: deployment scenarios and operational challenges. *IEEE Commun. Mag.* **50**(2), 148–155 (2012)
2. G Wunder, P Jung, M Kasparick, T Wild, F Schaich, Y Chen, S ten Brink, I Gaspar, N Michailow, A Festag, L Mendes, N Cassiau, D Ktenas, M Dryjanski, S Pietrzyk, B Eged, P Vago, F Wiedmann, 5G NOW: Non-orthogonal, asynchronous waveforms for future mobile applications. *IEEE Commun. Mag.* **52**(2), 97–105 (2014)
3. R Irmer, H Droste, P Marsch, M Grieger, G Fettweis, S Brueck, HP Mayer, L Thiele, V Jungnickel, Coordinated multipoint: concepts, performance, and field trial results. *IEEE Commun. Mag.* **49**(2), 102–111 (2011)
4. 3GPP, *Technical Report 36.819 V11.2.0: Coordinated Multi-point Operation for LTE Physical Layer Aspects*. (3GPP, Porto, 2013)
5. G Boudreau, J Panicker, N Guo, R Chang, N Wang, S Vrzic, Interference coordination and cancellation for 4G networks. *IEEE Commun. Mag.* **47**(4), 74–81 (2009)
6. O Aliu, A Imran, M Imran, B Evans, A survey of self organisation in future cellular networks. *IEEE Commun. Surv. & Tutor.* **PP**(99), 1–26 (2012)
7. RY Chang, Z Tao, J Zhang, CCJ Kuo, Dynamic fractional frequency reuse (D-FFR) for multicell OFDMA networks using a graph framework. *Wireless Commun. Mobile Comp.* **13**, 12–27 (2013). doi:10.1002/wcm.1088.
8. J Huang, VG Subramanian, R Agrawal, R Berry, Joint scheduling and resource allocation in uplink OFDM systems for broadband wireless access networks. *IEEE J. Select. Areas Commun.* **27**(2), 226–234 (2009)
9. L Venturino, N Prasad, X Wang, Coordinated scheduling and power allocation in downlink multicell OFDMA networks. *IEEE Trans. Vehic. Technol.* **58**(6), 2835–2848 (2009)
10. D Gesbert, SG Kiani, A Gjendemsj, GE Ien, Adaptation, coordination, and distributed resource allocation in interference-limited wireless networks. *Proc. IEEE.* **95**(12), 2393–2409 (2007)
11. W Yu, T Kwon, C Shin, Multicell coordination via joint scheduling, beamforming and power spectrum adaptation, in *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)* (Shanghai, 2011), pp. 2570–2578
12. S Borst, M Markakis, I Saniee, Distributed power allocation and user assignment in OFDMA cellular networks, in *Proceedings of the 49th Annual Allerton Conference on Communication, Control, and Computing* (Monticello, 2011), pp. 1055–1063
13. IH Hou, CS Chen, Self-organized resource allocation in LTE systems with weighted proportional fairness, in *Proceedings of the IEEE International Conference on Communications (ICC)* (Ottawa, 2012), pp. 5348–5353
14. AL Stolyar, H Viswanathan, Self-organizing dynamic fractional frequency reuse for best-effort traffic through distributed inter-cell coordination, in *Proceedings of the IEEE International Conference on Computer Communications (INFOCOM)* (Rio de Janeiro, 2009)
15. B Rengarajan, AL Stolyar, H Viswanathan, Self-organizing dynamic fractional frequency reuse on the uplink of OFDMA systems, in *Proceedings of the 44th Annual Conference on Information Sciences and Systems (CISS)* (Princeton, 2010), pp. 1–6
16. NUL Hassan, M Assaad, Downlink beamforming and resource allocation in multicell MISO-OFDMA systems. *Trans. Emer. Telecommun. Technol.* **25**(2), 173–182 (2012)

17. J Papandriopoulos, JS Evans, SCALE: a low-complexity distributed protocol for spectrum balancing in multiuser DSL networks. *IEEE Trans. Inform. Theory.* **55**(8), 3711–3724 (2009)
18. J Papandriopoulos, S Dey, J Evans, Optimal and distributed protocols for cross-layer design of physical and transport layers in MANETs. *IEEE/ACM Trans. Netw.* **16**(6), 1392–1405 (2008)
19. B Song, YH Lin, RL Cruz, Weighted max-min fair beamforming, power control, and scheduling for a MISO downlink. *IEEE Trans. Wireless Commun.* **7**(2), 464–469 (2008)
20. L Yu, E Karipidis, EG Larsson, Coordinated scheduling and beamforming for multicell spectrum sharing networks using branch and bound, in *Proceedings of the 20th European Signal Processing Conference (EUSIPCO)* (Bucharest, 2012), pp. 819–823
21. W Utschick, J Brehmer, Monotonic optimization framework for coordinated beamforming in multicell networks. *IEEE Trans. Signal Process.* **60**(4), 1899–1909 (2012)
22. C Shi, RA Berry, ML Honig, Distributed interference pricing for OFDM wireless networks with non-separable utilities, in *Proceedings of the 42nd Annual Conference on Information Sciences and Systems (CISS)* (Princeton, 2008), pp. 755–760
23. F Ahmed, AA Dowhuszko, O Tirkkonen, R Berry, A distributed algorithm for network power minimization in multicarrier systems, in *Proceedings of the IEEE 24th International Symposium on Personal Indoor and Mobile Radio Communications (PIMRC)* (London, 2013), pp. 1914–1918
24. S Sengupta, M Chatterjee, KA Kwiat, A game theoretic framework for power control in wireless sensor networks. *IEEE Trans. Comput.* **59**(2), 231–242 (2010)
25. K Son, S Lee, Y Yi, S Chong, Practical dynamic interference management in multi-carrier multi-cell wireless networks: a reference user based approach, in *Proceedings of the 8th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)* (Avignon, 2010)
26. AA Dowhuszko, F Ahmed, O Tirkkonen, Decentralized transmit beamforming scheme for interference coordination in small cell networks, in *Proceedings of the First International Black Sea Conference on Communications and Networking (BlackSeaCom)* (Batumi, 2013), pp. 121–126
27. Q Zhang, SA Kassam, Finite-state Markov model for Rayleigh fading channels. *Commun. IEEE Trans.* **47**(11), 1688–1692 (1999)
28. AL Stolyar, On the asymptotic optimality of the gradient scheduling algorithm for multiuser throughput allocation. *Oper. Res.* **1**, 12–25 (2005)
29. A Stolyar, Maximizing queueing network utility subject to stability: greedy primal-dual algorithm. *Queueing Syst.* **50**(4), 401–457 (2005)
30. G Wunder, M Kasparick, A Stolyar, H Viswanathan, Self-organizing distributed inter-cell beam coordination in cellular networks with best effort traffic, in *Proceedings of the 8th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)* (Avignon, 2010)
31. S Boyd, L Vandenberghe, *Convex Optimization*. (Cambridge University Press, Cambridge, 2004)
32. E Matskani, ND Sidiropoulos, L Tassiulas, Convex approximation algorithms for back-pressure power control. *IEEE Tran. Signal Process.* **60**(4), 1957–1970 (2012)
33. BR Marks, GP Wright, Technical note—a general inner approximation algorithm for nonconvex mathematical programs. *Oper. Res.* **26**(4), 681–683 (1978)
34. R Horst, PM Pardalos, NV Thoai, *Introduction to Global Optimization*, 2nd edn. (Kluwer, Dordrecht, 2000)
35. 3GPP R1-073937, Alcatel-Lucent, *3GPP TSG RAN WG1 #50bis: Comparison Aspects of Fixed and Adaptive Beamforming for LTE Downlink*. (3GPP, Shanghai, 2007)
36. Wireless World Initiative New Radio, *D5.4 Final Report on Link and System Level Channel Models, Technical report, IST-2003-507581*, (WINNER, 2005)

doi:10.1186/1687-1499-2014-120

Cite this article as: Kasparick and Wunder: **Autonomous algorithms for centralized and distributed interference coordination: a virtual layer-based approach.** *EURASIP Journal on Wireless Communications and Networking* 2014 **2014**:120.

Submit your manuscript to a SpringerOpen® journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Immediate publication on acceptance
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► springeropen.com
