Joint Power and Admission Control for Ad-hoc and Cognitive Underlay Networks: Convex Approximation and Distributed Implementation

M.Sc. Thesis
by
Ioannis Mitliagkas

Submitted to the Department of Electronic & Computer Engineering
Technical University of Crete
in partial fulfillment of the requirements for the M.Sc. Degree.

Advisor: Professor Sidiropoulos Nikolaos
Supervisory committee members:
Professor Liavas Athanasios
Assistant Professor Karystinos Georgios

July 13, 2010

Contents

1 Introduction **1**

2 Problem Formulation **4**

3 Convex Approximation **7**
3.1 Step 1: Single-stage Reformulation 7
3.2 Step 2: Isolating Non-convexity 8
3.3 Step 3: Semidefinite Programming Relaxation 8
3.4 Step 4: Approximation Algorithm 9

4 Imperfect Channel State Information **11**

5 Distributed Implementation **13**
5.1 Distributed Deflation and Feedback Requirements 16

6 Simulations **18**
6.1 Centralized Algorithm under Perfect CSI 18
6.2 Centralized Algorithm with Imperfect CSI 19
6.3 Distributed Algorithm 20

7 Discussion and Conclusions **26**

Acknowledgments

I want to thank N.D. Sidiropoulos and A. Swami for putting their ideas, time and expertise in [18], [19] and in the submitted journal paper containing this work.

I want to thank all the professors in my committee for the help and advice before and during my academic migration.

I want to thank my family, friends and Vanessa.

Abstract

Power control is important in interference-limited cellular, ad-hoc, and cognitive underlay networks, when the objective is to ensure a certain quality of service to each connection. Power control has been extensively studied in this context, including distributed algorithms that are particularly appealing in ad-hoc and cognitive settings. A long-standing issue is that the power control problem may be infeasible, thus requiring appropriate admission control. The power and admission control parts of the problem are tightly coupled, but the joint optimization problem is NP-hard. We begin with a convenient reformulation which enables a disciplined convex approximation approach. This leads to a centralized approximate solution that is numerically shown to outperform the prior art, and even yield close to optimal results in certain cases - at affordable complexity. The issue of imperfect channel state information is also considered. A distributed implementation is then developed, which alternates between distributed approximation and distributed deflation - reaching consensus on a user to drop, when needed. Both phases require only local communication and computation, yielding a relatively lightweight distributed algorithm with the same performance as its centralized counterpart.

Chapter 1

Introduction

Power control has been extensively studied in the context of cellular networks, as a way of mitigating intra-cell and inter-cell interference [29, 10]. Power control is also important in infrastructure-less ad-hoc wireless networks, where multiple co-channel links operate simultaneously, causing interference to one another. Originally motivated by the need to support circuit-switched-quality voice services (now voice-over-IP and other applications requiring guaranteed rate), the prevailing formulation of power control aims to ensure a certain quality of service, measured in terms of a link's signal to interference plus noise ratio (SINR), to every user in the network. A key difficulty that has long been recognized is that the problem is often infeasible: it is not possible to simultaneously satisfy all user demands in the same time or frequency slot. This brings up the issue of admission control, and a natural objective is to maximize the number (or weighted sum) of admitted users. The joint admission and power control problem is NP-hard, but important in practice [1, 7, 3].

The work to date on joint admission and power control has focused on *gradual removals* (e.g., [1, 7, 3]) until the problem becomes feasible, or *gradual admissions* (e.g., [28, 25, 2, 22, 23, 24]) when possible. In both cases, the issue is whether or not to remove or admit a single user, and adjust transmission powers if necessary. Distributed admission control algorithms that accept or reject an incoming call in a power-controlled cellular network can be found in [28] and [25]. Joint admission and power control strategies offering active user protection have been investigated in a series of papers [2, 22, 23, 24]. Active user protection makes sense from a customer experience point of view (e.g., few dropped calls). On the other hand, it can be far from optimal in terms of accommodating the maximum possible number of users, or other 'social' metrics; and it limits agility, which can be crucial in certain scenarios. Admission control for maximal throughput in power-controlled networks has been considered in [15].

Efficient utilization of the wireless spectrum has been a growing concern lately, owing to the inherent scarcity of the resource and the plethora of emerging mobile devices and services competing for bandwidth. It has become apparent that static regulatory allocations of parts of the spectrum to services/users is very inefficient. Cognitive radio has thus emerged as an adaptive cohabitation paradigm for wireless communication. Cognitive radio nodes sense their environment and adapt their transmission mode to enable efficient spectrum sharing. The idea is to enable secondary spectrum usage while avoiding or limiting interference to licensed primary users, in a way that is fair to other peers. Building upon the functionality offered by (then nascent) software radio, cognitive radio was conceived in the late '90's [20]. The concept started gaining momentum a few years later, after a U.S. Federal Communications Commission (FCC) Spectrum Policy Task Force report [8] highlighted that the typical utilization of licensed bands is under 20%. There is plenty of idle spectrum in most places, most of the time; the issue is how to discover it in a timely fashion and use it in an efficient manner. This realization sparked considerable research, regulatory, and standardization activity, starting in 2003 and growing fast nowadays.

Two basic modes of operation of cognitive radio have emerged so far [32, 31, 16]: *spectrum overlay*, in which secondary users seek idle time-frequency slots (transmission opportunities) and try to avoid colliding with the primary users (e.g., see [33]); and *spectrum underlay*, in which secondary users try to limit the amount of interference they cause to the primary users, but otherwise forego activity detection and may transmit 'at will' - even in the same time-frequency slot(s) as the primary users. Both modes require some level of situational awareness - spectrum sensing and activity detection for spectrum overlay, interference channel gain estimation for spectrum underlay - but at different accuracy and time scales. Overlay systems are collision-limited, but may transmit at relatively high power when transmission opportunities arise. Underlay systems require proper power control, but afford relatively seamless coexistence without stringent sensing requirements.

Taking advantage of spatial reuse, secondary spectrum underlay is closer in spirit to the traditional point of view of interference-limited wireless networks. This has facilitated migration of research results on power control, transmit beamforming, and scheduling from the cellular to the cognitive regime [12, 11, 30]. An uplink beamforming and power control scenario where the objective is to maximize the sum rate of the secondary users under interference constraints on the primary users has been considered in [30]. Explicit user admission is not needed in a sum-rate context. A downlink beamforming scenario for the secondary users is considered in [12], under SINR constraints on the primary and secondary users. Infeasibility and user selection issues were not dealt with in [12]. In the same context, a suboptimal user selection strategy was recently proposed in [11], based on pairwise orthogonality of the channel vectors.

The joint power and admission control problem is considered in this paper, for a cognitive underlay scenario where:

- Primary users must be guaranteed a premium service rate, measured by their signal to interference plus noise ratio (SINR);
- Secondary users, if admitted, should be provided with at least a basic service rate;
- The number of admitted secondary users should be maximized, and the total power required to serve them should be minimized.

The ad-hoc setting can be viewed as a special case wherein all users are peers, and there are no primary interference constraints. A disciplined *convex approximation approach* is adopted in this paper. Instead of aiming for the hard-to-get optimal solution or directly trying to approximate it, the idea is to approximate the problem *per se* by a suitable convex problem that is "close" to the original one. The solution of the convex problem is then used to guide the search for a good feasible solution of the original problem. In our particular context, linear programming relaxation is used for convex approximation, and the final approximate solution is obtained through a sequence of linear programs. The issue of imperfect channel state information (CSI) is also considered. Assuming bounded CSI errors, and insisting that the SINR constraints be met in the worst case, a robust reformulation of the joint power and admission control problem is obtained. This admits a second order cone programming (SOCP) relaxation, and approximate solution through a sequence of SOCP programs. Simulation results are included to illustrate the merits of the approach. Two scenarios are considered: with or without a primary user. In the latter, several good heuristic algorithms are available in the literature, and the prevailing one is used as a baseline. A brute-force enumeration algorithm is used in both cases to assess the gap from the optimal solution.

An appealing feature of classical power control solutions is that they lend themselves to distributed implementation. When the power control problem is feasible, the global optimum can be reached using only local updates. Each link uses local interference plus noise measurements at the receiver to update the corresponding power at the transmitter. Distributed implementation is important for a number of reasons, including scalability, agility (the ability to track changes in the operational environment), and reduced vulnerability to node failures. Depending on the kind of feedback required, distributed implementation can also be more lightweight in terms of signaling overhead. These considerations motivate distributed implementation of the proposed algorithm. This is the subject of the last part of the paper. The resulting implementation alternates between distributed approximation and distributed deflation - reaching consensus on a user to drop, when needed. The approximation phase uses dual decomposition - each node updates its local primal variables, while subgradient iterations are used to update the dual variables. The deflation phase employs a consensus-on-the-max algorithm to reach agreement on which user to drop, if needed. Both phases require only local communication and computation, yielding a relatively lightweight distributed algorithm that converges to the same approximate solution as its centralized counterpart.

This body of work is the expansion of my undergraduate diploma thesis under the supervision of Prof. Sidiropoulos. In that thesis, the work presented in [18], we focused on the centralized algorithm. Additions in this thesis include: a proof of NP-hardness for the original problem, a distributed implementation, a robust version of the problem and algorithm and comprehensive simulations.

Chapter 2

Problem Formulation

Consider a channel that is used by a single primary user U_0 and K secondary users $\mathcal{U} := \{1, \ldots, K\}$. By 'user' here we mean a transmitter - receiver pair (directed link). A single primary user is considered for brevity of exposition. It is straightforward to include additional constraints to account for more primary users; this does not change the structure of the problem in any way. User k transmits with power $p_k \leq P_k^{MAX}$. The primary user's transmission power p_0 is fixed, because cooperation cannot be assumed. For each link k we define c_k as the SINR threshold that must be attained for the link to meet its QoS requirement. Let σ_k^2 denote the thermal noise power at the reviver of link k and G_{ij} the link gain from the transmitter of link i to the receiver of link j.

Our purpose is to allow secondary users to use the channel without disrupting the primary user's communication. One way to achieve this is by controlling the secondary user transmission powers. When there are many secondary links competing for service and/or the SINR constraints are tighter than can be satisfied, power control alone cannot solve the problem. In this case we need to employ some form of admission control. Admission control should be optimized together with power allocation, because the two are intertwined.

The problem of interest can be described in two stages: maximize the number of secondary users that can be admitted, and then minimize the total power required to serve them. Mathematically, the first stage can be expressed as follows.

$$S_o = \text{argmax}_{S \subseteq \{1,\ldots,K\}, \{p_k \in \mathbb{R}_+\}_{k=1}^{K}} |S| \tag{2.1}$$

$$\text{s.t. } p_k \leq P_k^{MAX}, \quad \forall k \in \{1, \ldots, K\} \tag{2.2}$$

$$\frac{G_{kk} p_k}{\sum_{l=1,\, l \neq k}^{K} G_{lk} p_l \ + G_{0k} p_0 + \sigma_k^2} \geq c_k, \quad \forall k \in S \tag{2.3}$$

$$\frac{G_{00} p_0}{\sum_{l=1}^{K} G_{l0} p_l \ + \sigma_0^2} \geq c_0 \ . \tag{2.4}$$

Here (2.3) is the SINR constraint for the secondary users, and (2.4) is the SINR constraint for the primary user. Notice that the term $G_{0k} p_0$ in the denominator of (2.3) accounts for the interference caused by the primary user to user k.

Once a maximal admissible subset of secondary users is found, what remains is to adjust their powers to minimize the total transmitted power. This can be written as

$$\min_{\{p_k \in \mathbb{R}_+\}_{k=1}^K} \sum_{k \in S_o} p_k \tag{2.5}$$

$$\text{s.t. } p_k \le P_k^{MAX}, \quad \forall k \in S_o \tag{2.6}$$

$$\frac{G_{kk} p_k}{\sum_{l \ne k, l \in S_o} G_{lk} p_l \ + G_{0k} p_0 + \sigma_k^2} \ge c_k, \quad \forall k \in S_o \tag{2.7}$$

$$\frac{G_{00} p_0}{\sum_{l \in S_o} G_{l0} p_l \ + \sigma_0^2} \ge c_0 \tag{2.8}$$

Remark 1 *There may be multiple equivalent (in terms of cardinality) solutions of (2.1)-(2.4), which may lead to different sum-power in (2.5)-(2.8). If multiple solutions do exist, one may wish to solve (2.5)-(2.8) for each candidate solution of (2.1)-(2.4), and pick the one that yields the overall smallest sum power in the end. In the sequel, we will reformulate the overall problem in a way that will take us directly to the global minimum power solution through a single optimization problem.*

The power control problem in the second stage (2.5)-(2.8) is a Linear Program (LP) and thus easily solved - there even exist specialized solutions that are far more efficient than generic LP solvers for the particular problem in (2.5)-(2.8). The challenge lies in the first (subset selection) stage:

Claim 1 *The subset selection problem in (2.1)-(2.4) is NP-hard.*

Proof 1 *Consider the following special case of (2.1)-(2.4):*

$$S_o = \textit{argmax}_{S \subseteq \{1,\ldots,K\}, \{p_k \in [0,1]\}_{k=1}^K} |S| \tag{2.9}$$

$$\textit{s.t. } \frac{p_k}{\sum_{l=1,\ l \ne k}^K G_{lk} p_l \ + 1} \ge 1, \quad \forall k \in S \tag{2.10}$$

We will show that it contains the maximal independent set problem, which is known to be NP-hard [9]. Let $\Gamma = (V, E)$ be an undirected graph, with $|V| = K$ vertices, one for each user, and edges $e_{l,k} \in E$. A subset of vertices $S \subseteq V$ of Γ is independent when no two vertices in S are connected by an edge in E. Given any $\Gamma = (V, E)$, define a corresponding instance of (2.9)-(2.10) by setting

$$G_{lk} = \begin{cases} 1, & e_{l,k} \in E \\ 0, & \textit{otherwise} \end{cases} \tag{2.11}$$

Let S_i be a maximal independent set in Γ. Setting

$$p_k = \begin{cases} 1, & k \in S_i \\ 0, & \textit{otherwise} \end{cases} \tag{2.12}$$

will satisfy

$$\frac{p_k}{\sum_{l=1,\ l \ne k}^K G_{lk} p_l \ + 1} = 1, \quad \forall k \in S_i \tag{2.13}$$

because, by definition of independent set and G_{lk}, the nodes in S_i do not interfere with one another, and the power of any remaining nodes has been switched off. It follows that $|S_o| \geq |S_i|$. Conversely, let $\{p_k \in [0,1]\}_{k=1}^K$ be such that

$$\frac{p_k}{\sum_{l=1,\ l\neq k}^K G_{lk} p_l\ + 1} \geq 1, \quad \forall k \in S \tag{2.14}$$

for some $S \subseteq \{1, \ldots, K\}$. The only way for this to hold is to have $p_k = 1, \forall k \in S$, hence it must be that $G_{lk} = 0$ for all pairs $l \in S$, $k \in S$. By definition of G_{lk}, this implies that S is an independent set in Γ. This is true in particular for S_o, hence $|S_i| \geq |S_o|$. ■

Note that NP-hardness of joint admission and power control in a cellular context has been considered in [1], but the proof there is incomplete[1].

[1] [1] does not show that an *arbitrary* instance of the chosen NP-hard problem can be posed as an instance of (2.1)-(2.4).

Chapter 3

Convex Approximation

3.1 Step 1: Single-stage Reformulation

We next reformulate the two-stage problem in (2.1)-(2.4) and (2.5)-(2.8) into an *equivalent single-stage optimization* problem. This is in the spirit of the approach in [17], albeit it does not follow as a special case. Let us consider the following problem:

$$\min_{\{p_k \in \mathbb{R}_+, s_k \in \{-1,+1\}\}_{k=1}^K} \epsilon \sum_{k=1}^{K} p_k + (1-\epsilon) \sum_{k=1}^{K} \lambda_k (s_k+1)^2 \tag{3.1}$$

$$\text{s.t. } p_k \leq P_k^{MAX}, \quad \forall k \in \{1, \ldots, K\} \tag{3.2}$$

$$\frac{G_{kk} p_k + \delta_k^{-1}(s_k+1)^2}{\sum_{l=1,\ l \neq k}^{K} G_{lk} p_l \ + G_{0k} p_0 + \sigma_k^2} \geq c_k, \quad \forall k \in \{1, \ldots, K\} \tag{3.3}$$

$$\frac{G_{00} p_0}{\sum_{l=1}^{K} G_{l0} p_l \ + \sigma_0^2} \geq c_0 \tag{3.4}$$

We have introduced binary scheduling variables s_k which take the value -1 for an admitted user and 1 for a dropped one. Notice that variable s_k also appears in the SINR constraint of user k. For sufficiently small δ_k and $s_k = 1$, the SINR constraint of user k becomes inactive; whereas for $s_k = -1$ the constraint remains active. The cost function (3.1) accounts for both admission and power control. The admission control component of the cost is discrete-valued, whereas the power component is bounded. By choosing ϵ small enough, we can ensure that admission control has absolute priority over power control: dropping any user costs more than can possibly be saved in terms of transmission power for the rest. A ruler analogy in which the decimal ticks correspond to the discrete admission cost and the intervals between ticks are (partially) spanned by the power cost can be helpful to intuitively appreciate the following result:

Claim 2 *For $\lambda_k = 1, \forall k \in \{1, \cdots, K\}$, and*

$$0 < \epsilon < \frac{4}{\sum_{k=1}^{K} P_k^{MAX} + 4} \tag{3.5}$$

$$\delta_k \leq \frac{4}{c_k \left(\sum_{l=1,\ l \neq k}^{K} G_{lk} P_l^{MAX} + G_{0k} p_0 + \sigma_k^2 \right)} \tag{3.6}$$

the single-stage reformulation in (3.1)-(3.4) is equivalent *to solving the two-stage problem in (2.1)-(2.4) and (2.5)-(2.8). In fact, if there are multiple solutions to (2.1)-(2.4), solving (3.1)-(3.4) will yield the one of minimum sum power.*

The proof is by contradiction, similar to the line of argument in [17]. We skip it here for space considerations.

The reason for introducing the weights λ_k is that these can be used to promote 'social welfare' or 'fairness'. For example, setting λ_k proportional to the kth user's queue length will optimize system throughput; setting it inversely proportional to a running average estimate of the user's service rate will encourage fairness. Do note, however, that the equivalence to (2.1)-(2.4) and (2.5)-(2.8) is lost when the weights are not equal, as this differentiates the users.

3.2 Step 2: Isolating Non-convexity

The problem in (3.1)-(3.4) is of course also NP-hard[1] *and* not directly amenable to convex approximation. The following *equivalent* reformulation explicitly reveals the non-convex part of the problem, thus getting us closer to a convex one:

$$\min_{\{p_k \in \mathbb{R}_+, S_k \in \mathbb{R}^{2\times 2}\}_{k=1}^K} \epsilon \sum_{k=1}^K p_k + (1-\epsilon) \sum_{k=1}^K \lambda_k \mathrm{Tr}(1_{2\times 2} S_k) \tag{3.7}$$

$$\text{s.t. } p_k \leq P_k^{MAX}, \quad \forall k \in \{1, \ldots, K\} \tag{3.8}$$

$$\frac{G_{kk} p_k + \delta_k^{-1} \mathrm{Tr}(1_{2\times 2} S_k)}{\sum_{l=1,\ l\neq k}^K G_{lk} p_l \ + G_{0k} p_0 + \sigma_k^2} \geq c_k, \quad \forall k \in \{1, \ldots, K\} \tag{3.9}$$

$$\frac{G_{00} p_0}{\sum_{l=1}^K G_{l0} p_l \ + \sigma_0^2} \geq c_0 \tag{3.10}$$

$$S_k \geq 0, \mathrm{rank}(S_k) = 1, S_k(1,1) = S_k(2,2) = 1 \ \forall k \in \{1, \ldots, K\} \tag{3.11}$$

where $S_k \geq 0$ means that matrix S_k is positive semidefinite. Its diagonal elements are 1's and its off-diagonal elements hold the original scheduling variable s_k. Matrix $1_{2\times 2}$ is the 2×2 matrix of all 1's.

The rank-one constraint restricts the scheduling variables in the set $\{-1, +1\}$. This is the only source of non-convexity in (3.7)-(3.11).

3.3 Step 3: Semidefinite Programming Relaxation

Dropping the rank-one constraints (which is equivalent to allowing the s_k's to take any value in $[-1\ +1]$) leaves us with a Semidefinite Programing (SDP) [5] problem:

$$\min_{\{p_k \in \mathbb{R}_+, S_k \in \mathbb{R}^{2\times 2}\}_{k=1}^K} \epsilon \sum_{k=1}^K p_k + (1-\epsilon) \sum_{k=1}^K \lambda_k \mathrm{Tr}(1_{2\times 2} S_k) \tag{3.12}$$

[1] To see this, set $\lambda_k = 1, \forall k$, and send $\epsilon \to 0$ to recover (2.1)-(2.4). A formal proof can be constructed to show that it contains the maximal independent set problem, as per the proof of Claim 1

$$\text{s.t. } p_k \leq P_k^{MAX}, \quad \forall k \in \{1, \ldots, K\} \tag{3.13}$$

$$\frac{G_{kk}p_k + \delta_k^{-1}\mathrm{Tr}(1_{2\times 2}S_k)}{\sum_{l=1,\ l\neq k}^{K} G_{lk}p_l \ + G_{0k}p_0 + \sigma_k^2} \geq c_k, \quad \forall k \in \{1, \ldots, K\} \tag{3.14}$$

$$\frac{G_{00}p_0}{\sum_{l=1}^{K} G_{l0}p_l \ + \sigma_0^2} \geq c_0 \tag{3.15}$$

$$S_k \geq 0, S_k(1,1) = S_k(2,2) = 1 \ \forall k \in \{1, \ldots, K\} \tag{3.16}$$

In [26], it is shown that this rank relaxation yields the Lagrange bi-dual problem, which is the closest convex problem to (3.7)-(3.11) in a certain sense, thus motivating rank relaxation; see also [13] and [14] for further insights and motivation.

In our case, the relaxed problem (3.12)-(3.16) can be easily shown to be equivalent to the following linear program.

$$\min_{\{p_k \in \mathbb{R}_+, t_k \in \mathbb{R}_+\}_{k=1}^K} \epsilon \sum_{k=1}^{K} p_k + (1-\epsilon) \sum_{k=1}^{K} \lambda_k t_k \tag{3.17}$$

$$\text{s.t. } p_k \leq P_k^{MAX}, \quad \forall k \in \{1, \ldots, K\} \tag{3.18}$$

$$\frac{G_{kk}p_k + \delta_k^{-1} t_k}{\sum_{l=1,\ l\neq k}^{K} G_{lk}p_l \ + G_{0k}p_0 + \sigma_k^2} \geq c_k, \quad \forall k \in \{1, \ldots, K\} \tag{3.19}$$

$$\frac{G_{00}p_0}{\sum_{l=1}^{K} G_{l0}p_l \ + \sigma_0^2} \geq c_0 \tag{3.20}$$

$$0 \leq t_k \leq 4, \quad \forall k \in \{1, \ldots, K\} \tag{3.21}$$

which further simplifies computation. The solution of (3.17)-(3.21) yields a lower bound on the objective of (3.7)-(3.11), and thus a way to assess the quality of suboptimal solutions to (3.7)-(3.11). Still, solving the relaxed problem in (3.17)-(3.21) is certainly **not** equivalent to solving the original problem in (3.7)-(3.11). How to obtain a good approximate solution of (3.7)-(3.11) using (3.17)-(3.21) is addressed in the next section.

3.4 Step 4: Approximation Algorithm

The main idea is to employ *deflation* over (3.17)-(3.21). That is, solve (3.17)-(3.21), and check if all the original constraints are satisfied. If not, choose a user to drop and repeat until the problem becomes feasible.

Algorithm 1 L*inear* P*rogramming* D*eflation (LPD):*

1. $\mathcal{U} \leftarrow \{1, ..., K\}$
2. *Solve (3.17)-(3.21) for the users in* $\mathcal{U}$ *only.*
3. *If all links in* $\mathcal{U}$ *attain target SINR go to Step 4. Else use a heuristic (see text below) to choose a link, remove it from* $\mathcal{U}$ *and go to Step 2.*

A quite important factor for the performance of this algorithm is the heuristic employed to drop links at each iteration. We tried many, and the most promising one is as follows. At each step, after solving (3.17)-(3.21), we calculate a metric for each link. Let p_k^e be the excess transmission power needed for link k to attain its target SINR, assuming all other link powers are as calculated from (3.17)-(3.21). This excess transmission power for link k causes excess interference to all other links. Let $x_k^e = p_k^e \sum_{l \neq k} G_{kl}$ be the sum of excess interference powers caused to all other links due to p_k^e. Let $y_k^e = \sum_{l \neq k} G_{lk} p_l^e$ be the excess interference caused to link k due to the excess transmission powers of all other links. The link metric used for choosing the link to drop is $m_k := x_k^e + y_k^e$. The link that has the largest m_k is dropped, and the process continues by solving again (3.17)-(3.21) for the remaining links, until a feasible solution (requiring no excess power for any link) is found.

Chapter 4

Imperfect Channel State Information

An important issue in practice is what happens when the channel gains are not known exactly, but only estimates are available. Assuming that the estimation errors are bounded, it is possible to extend the basic approach to incorporate uncertainty, as explained next. The key is the LP relaxation in (3.17)-(3.21), for robust LP with bounded uncertainty in the constraint parameters is SOCP (see, e.g., section [4.4.2] in [5]).

The SINR constraints in (3.19) can be compactly written as

$$\bar{\boldsymbol{g}}_k^T \boldsymbol{p} - \frac{\delta_k^{-1}}{c_k} t_k \leq -\sigma_k^2, \ \forall k \in \{1, \ldots, K\} \tag{4.1}$$

where

$$\bar{\boldsymbol{g}}_k = [G_{0k} \ G_{1k} \ \ldots \ G_{(k-1)k} \ - \frac{G_{kk}}{c_k} \ G_{(k+1)k} \ \ldots \ G_{Kk}]^T \tag{4.2}$$

and the augmented power vector (note that p_0 is not an optimization variable)

$$\boldsymbol{p} = [p_0 \ p_1 \ldots \ p_K]^T. \tag{4.3}$$

Likewise, the primary user's SINR constraint in (3.20) can be expressed as

$$\bar{\boldsymbol{g}}_0^T \boldsymbol{p} \leq -\sigma_k^2, \ \forall k \in \{1, \ldots, K\} \tag{4.4}$$

where

$$\bar{\boldsymbol{g}}_0 = [-\frac{G_{00}}{c_0} \ G_{10} \ G_{20} \ldots \ G_{K0}]^T \tag{4.5}$$

Now, assume that the true vectors $\boldsymbol{g}_k$ and vector $\boldsymbol{g}_0$ lie inside ellipsoids $\mathcal{E}_k$ and $\mathcal{E}_0$ with centers the respective estimated values $\bar{\boldsymbol{g}}_k$ and $\bar{\boldsymbol{g}}_0$:

$$\boldsymbol{g}_k \in \mathcal{E}_k = \{\bar{\boldsymbol{g}}_k + E_k u \mid ||u||_2 \leq 1\}, \ \forall k \in \{0, \ldots, K\} \tag{4.6}$$

where matrix $E_k \in \mathbb{R}^{K+1 \times K+1}$ determines the size, shape and orientation of ellipsoid $\mathcal{E}_k$. The robust counterpart of (4.1) is

$$\boldsymbol{g}_k^T \boldsymbol{p} - \frac{\delta_k^{-1}}{c_k} t_k \leq -\sigma_k^2, \ \forall \boldsymbol{g}_k \in \mathcal{E}_k, \ \forall k \in \{1, \ldots, K\} \tag{4.7}$$

or equivalently, for each k,

$$
\begin{aligned}
\sup\left\{\boldsymbol{g}_k^T\boldsymbol{p} - \frac{\delta_k^{-1}}{c_k}t_k \mid \boldsymbol{g}_k \in \mathcal{E}_k\right\} &\leq -\sigma_k^2 \\
\sup\left\{\boldsymbol{g}_k^T\boldsymbol{p} \mid \boldsymbol{g}_k \in \mathcal{E}_k\right\} - \frac{\delta_k^{-1}}{c_k}t_k &\leq -\sigma_k^2 \\
\bar{\boldsymbol{g}}_k^T\boldsymbol{p} + \sup\left\{u^T E_k^T\boldsymbol{p} \mid ||u||_2 \leq 1\right\} - \frac{\delta_k^{-1}}{c_k}t_k &\leq -\sigma_k^2 \\
\bar{\boldsymbol{g}}_k^T\boldsymbol{p} + ||E_k^T\boldsymbol{p}||_2 - \frac{\delta_k^{-1}}{c_k}t_k &\leq -\sigma_k^2 \qquad (4.8)
\end{aligned}
$$

To ensure that the inequality holds when link k is not admitted, we have to pick a δ_k that satisfies it for $p_l = P_l^{MAX}, \forall l \neq k$, $p_k = 0$ and $t_k = 4$. For diagonal E_k, δ_k should satisfy

$$
\delta_k \leq \frac{4}{c_k\left(\sum_{l=0,\ l\neq k}^{K} G_{lk}P_l^{MAX} + ||E_k^T\boldsymbol{P}_{-k}^{MAX}||_2 + \sigma_k^2\right)}
$$

where $\boldsymbol{P}_{-k}^{MAX}$ is the vector of maximum link powers, including the primary user, with a zero in element k. Note that the primary user transmits with a fixed power $p_0 = P_0^{MAX}$. In the same manner, the robust counterpart of the primary user's SINR constraint (4.4) is

$$
\boldsymbol{g}_0^T\boldsymbol{p} \leq -\sigma_0^2,\ \forall \boldsymbol{g}_0 \in \mathcal{E}_0 \qquad (4.9)
$$

which can be reduced to

$$
\bar{\boldsymbol{g}}_0^T\boldsymbol{p} + ||E_0^T\boldsymbol{p}||_2 \leq -\sigma_0^2 \qquad (4.10)
$$

Replacing inequalities (4.1), (4.4) with their robust versions (4.8), (4.10) yields a SOCP problem. The overall approximation algorithm remains similar to LPD for the case of perfect CSI, except that the SOCP formulation is now employed in lieu of LP as the basic deflation step, and the robust constraints (4.8), (4.10) are used to check whether links attain their target SINR in the worst case.

In scenarios with severe uncertainty, we found that introducing an additional step (see below) helps prevent overestimating interference during the course of deflation, thus yielding significantly better results. The complete robust algorithm is as follows.

Algorithm 2 S*econd* O*rder* C*one* D*eflation (SOCD):*

1. $\mathcal{U} \leftarrow \{1, ..., K\}$
2. *Solve (3.17),(3.18),(4.8),(4.10),(3.21) for the users in* $\mathcal{U}$ *only.*
3. *If all links in* $\mathcal{U}$ *attain target SINR terminate.*
4. *Solve again only for the links that attained their SINR target and update their powers in the previous solution.*
5. *Use the heuristic on the full solution (resulting power vector) to choose a link, remove it from* $\mathcal{U}$ *and go to Step 2.*

Chapter 5

Distributed Implementation

The first obstacle in designing a distributed algorithm for (3.17)-(3.21) is that the constraints in (3.19)-(3.20) are coupled across users. Ideally, we would like each user to optimize its own variables (p_k and t_k), relying on low-rate feedback from other users to ensure that the solution converges to the global optimum. Towards this end, we will employ a dual decomposition approach [21, 34]. Let $\boldsymbol{p} = [p_1, p_2, \ldots p_K]^T$, $\boldsymbol{t} = [t_1, t_2, \ldots t_K]^T$ denote the primal variables, and $\boldsymbol{\mu} = [\mu_0, \mu_1, \ldots \mu_K]^T$ the vector of dual variables (bear in mind that the λ_k's are link weights defined in the original problem formulation; for this reason, the dual variables are denoted by μ_k.). Let us form the partial Lagrangian

$$L(\boldsymbol{p}, \boldsymbol{t}, \boldsymbol{\mu}) = \epsilon \sum_{k=1}^{K} p_k + (1-\epsilon) \sum_{k=1}^{K} \lambda_k t_k$$

$$+ \sum_{k=1}^{K} \mu_k \left(c_k \sum_{l=0, l \neq k}^{K} G_{lk} p_l + c_k \sigma_k^2 - G_{kk} p_k - \delta_k^{-1} t_k \right)$$

$$+ \mu_0 \left(c_0 \sum_{l=1}^{K} G_{l0} p_l + c_0 \sigma_0^2 - G_{00} p_0 \right)$$

$$= \epsilon \sum_{k=1}^{K} p_k + (1-\epsilon) \sum_{k=1}^{K} \lambda_k t_k + \sum_{k=0}^{K} \mu_k c_k \sum_{l=0, l \neq k}^{K} G_{lk} p_l$$

$$+ \sum_{k=0}^{K} \mu_k c_k \sigma_k^2 - \sum_{k=0}^{K} \mu_k G_{kk} p_k - \sum_{k=1}^{K} \mu_k \delta_k^{-1} t_k$$

All terms in this expression are separated (sums of individual user contributions), except for the third one. Notice, however, that this term may be rewritten as

$$\sum_{k=0}^{K} \mu_k c_k \sum_{l=0, l \neq k}^{K} G_{lk} p_l = \sum_{k=0}^{K} \sum_{l=0, l \neq k}^{K} \mu_k c_k G_{lk} p_l$$

$$= \sum_{l=0}^{K} \sum_{k=0, k \neq l}^{K} \mu_k c_k G_{lk} p_l = \sum_{l=0}^{K} p_l \sum_{k=0, k \neq l}^{K} \mu_k c_k G_{lk}$$

This is a key step towards distributing the computation. Swapping variables k and l and substituting back in the Lagrangian, we obtain

$$
\begin{aligned}
L(\boldsymbol{p}, \boldsymbol{t}, \boldsymbol{\mu}) = \epsilon \sum_{k=1}^{K} p_k + (1-\epsilon) \sum_{k=1}^{K} \lambda_k t_k + \sum_{k=0}^{K} p_k \sum_{l=0, l\neq k}^{K} \mu_l c_l G_{kl} \\
+ \sum_{k=0}^{K} \mu_k c_k \sigma_k^2 - \sum_{k=0}^{K} \mu_k G_{kk} p_k - \sum_{k=1}^{K} \mu_k \delta_k^{-1} t_k \\
= \sum_{k=0}^{K} p_k \left(\epsilon + \sum_{l=0, l\neq k}^{K} \mu_l c_l G_{kl} - \mu_k G_{kk} \right) \\
+ \sum_{k=1}^{K} t_k \left((1-\epsilon)\lambda_k - \mu_k \delta_k^{-1} \right) + \sum_{k=1}^{K} \mu_k c_k \sigma_k^2 \\
= \sum_{k=0}^{K} L_k(p_k, t_k, \boldsymbol{\mu})
\end{aligned}
$$

where for $k \in \{1, \dots, K\}$

$$
\begin{aligned}
L_k(p_k, t_k, \boldsymbol{\mu}) = p_k \left(\epsilon + \sum_{l=0, l\neq k}^{K} \mu_l c_l G_{kl} - \mu_k G_{kk} \right) \\
+ t_k \left((1-\epsilon)\lambda_k - \mu_k \delta_k^{-1} \right) + \mu_k c_k \sigma_k^2
\end{aligned}
\tag{5.1}
$$

and

$$
L_0(\boldsymbol{\mu}) = p_0 \left(\epsilon + \sum_{l=1}^{K} \mu_l c_l G_{0l} - \mu_0 G_{00} \right) + \mu_0 c_0 \sigma_0^2 \tag{5.2}
$$

Notice that L_0 is a function of just $\boldsymbol{\mu}$ since p_0 is constant and not included in $\boldsymbol{p}$ and there is no t_0 - the primary user is always admitted. Dual variable μ_k is the cost users have to pay to interfere with user k. We have rewritten the Lagrangian as the sum of $K + 1$ individual Lagrangians involving only local variables and the dual variables. The dual function can be split as well,

$$
d(\boldsymbol{\mu}) = \inf_{\boldsymbol{p}, \boldsymbol{t}} \sum_{k=0}^{K} L_k(p_k, t_k, \boldsymbol{\mu}) = \sum_{k=0}^{K} d_k(\boldsymbol{\mu})
$$

where we have suppressed the box constraints on $\boldsymbol{p}, \boldsymbol{t}$ for brevity, and for $k \in \{1, \dots, K\}$

$$
\begin{aligned}
d_k(\boldsymbol{\mu}) = \inf_{p_k, t_k} p_k \left(\epsilon + \sum_{l=0, l\neq k}^{K} \mu_l c_l G_{kl} - \mu_k G_{kk} \right) \\
+ t_k \left((1-\epsilon)\lambda_k - \mu_k \delta_k^{-1} \right) + \mu_k c_k \sigma_k^2
\end{aligned}
\tag{5.3}
$$

whereas

$$
d_0(\boldsymbol{\mu}) = p_0 \left(\epsilon + \sum_{l=1}^{K} \mu_l c_l G_{0l} - \mu_0 G_{00} \right) + \mu_0 c_0 \sigma_0^2 \tag{5.4}
$$

As expected d_0 is constant over $\boldsymbol{p}$ and $\boldsymbol{t}$. This is a consequence of the fact that the primary user has no local (i.e. primary) variables to optimize. The resulting dual problem is

$$\max_{\boldsymbol{\mu}\in\mathbb{R}_+^{K+1}} d(\boldsymbol{\mu}) \tag{5.5}$$

which can be solved in a distributed fashion using the projected subgradient method. The overall approach iterates between computing minimizers of (5.3) in *closed form*, using them to calculate subgradients of d, and updating costs $\boldsymbol{\mu}$.

In order to recover the solution of (3.17)-(3.21) (i.e., the optimal primal variables) from the dual problem, the objective of the primal problem should be strictly convex. The linear objective in (3.17) is convex, but not strictly convex. We may bypass this difficulty by approximating the objective in (3.17) with

$$\epsilon \sum_{k=1}^{K} p_k^{1+\theta} + (1-\epsilon) \sum_{k=1}^{K} \lambda_k t_k^{1+\theta} \tag{5.6}$$

where θ is a small positive constant which can be chosen to ensure that the solution of the modified problem is within specified tolerance from that of the original problem. With this modification, (5.1) becomes

$$\begin{aligned} L_k(p_k, t_k, \boldsymbol{\mu}) = p_k &\left(\epsilon p_k^{\theta} + \sum_{l=0, l\neq k}^{K} \mu_l c_l G_{kl} - \mu_k G_{kk} \right) \\ &+ t_k \left((1-\epsilon)\lambda_k t_k^{\theta} - \mu_k \delta_k^{-1} \right) + \mu_k c_k \sigma_k^2 \end{aligned} \tag{5.7}$$

whereas (5.3) becomes

$$\begin{aligned} d_k(\boldsymbol{\mu}) = \inf_{p_k, t_k} p_k &\left(\epsilon p_k^{\theta} + \sum_{l=0, l\neq k}^{K} \mu_l c_l G_{kl} - \mu_k G_{kk} \right) \\ &+ t_k \left((1-\epsilon)\lambda_k t_k^{\theta} - \mu_k \delta_k^{-1} \right) + \mu_k c_k \sigma_k^2 \end{aligned} \tag{5.8}$$

and both are strictly convex. Note that $L_k(p_k, t_k, \boldsymbol{\mu})$ contains a term depending only on p_k, another depending only on t_k, and separate interval constraints on p_k, t_k. It follows that minimization of $L_k(p_k, t_k, \boldsymbol{\mu})$ with respect to p_k, t_k amounts to two separate 1-D strictly convex subproblems. Taking partial derivatives with respect to p_k, t_k, and equating to zero, we obtain

$$p_k^* = \left(\frac{\mu_k G_{kk} - \sum_{l=0, l\neq k}^{K} \mu_l c_l G_{kl}}{\epsilon\,(1+\theta)} \right)^{1/\theta} \tag{5.9}$$

and

$$t_k^* = \left(\frac{\mu_k \delta_k^{-1}}{(1-\epsilon)\lambda_k\,(1+\theta)} \right)^{1/\theta} \tag{5.10}$$

followed by projection of p_k^* onto $[0\ P_k^{MAX}]$, and t_k^* onto $[0\ 4]$. In each iteration, user k updates p_k and t_k as above, then updates μ_k using a projected subgradient step

$$\mu_k = [\mu_k - \alpha \rho_k]_+ \tag{5.11}$$

where $[\cdot]_+$ denotes projection onto the positive half-space, α is a suitable step size, ρ_k is the positive slack from the SINR constraints, which for $k \in \{1, \ldots, K\}$ is given by

$$\rho_k(p_k, t_k) = G_{kk}p_k + \delta_k^{-1}t_k - c_k \sum_{l \neq k} G_{lk}p_l \ - c_k\sigma_k^2 \tag{5.12}$$

and

$$\rho_0 = G_{00}p_0 - c_0 \sum_{l \neq 0} G_{l0}p_l - c_0\sigma_0^2 \tag{5.13}$$

It has been shown (e.g., section 6.3 in [4], and [34]) that, if for every k and given $\boldsymbol{\mu}$, p_k^* and t_k^* are minimizers of L_k, the vector of slacks $\rho_k(p_k^*, t_k^*)$ makes up a subgradient of the negative dual function $-d$ at $\boldsymbol{\mu}$. Using the update rule in (5.11) results in minimizing $-d$ or, equivalently, solving our dual problem.

The convergence properties of the algorithm are dependent on the choice of step size α. There are various strategies for the step size choice in the literature. We chose $\alpha_i = \alpha_0/i$ where i is the iteration number and α_0 is the initial step size (this sequence is square summable but not summable). This ensures convergence to the optimal solution (e.g., Proposition 6.3.4 in [4]), however the speed of convergence depends heavily on the choice of α_0. Fig. 6.8 illustrates convergence of the primal t_k variables in an infeasible scenario with $K = 3$ nodes.

5.1 Distributed Deflation and Feedback Requirements

The algorithm used in the distributed setting is essentially the LPD algorithm described in Section 3.4, where the primal-dual method described in this section is used instead of a centralized LP solver for step 2. In each iteration of this primal-dual method, user $k \in \{1, \ldots, K\}$ updates its local variables using (5.9), (5.10), (5.12), (5.11) [or (5.13), (5.11) for $k = 0$]. The update in (5.9) requires that node k is aware of c_l, G_{kl} and the current price μ_l for each neighboring node l affected by interference from node k (i.e., for which $G_{kl} \neq 0$). A separate low-rate control channel can be used to pass around this information to neighboring nodes. The update in (5.12), (5.11), [or (5.13), (5.11) for $k = 0$] is lighter in terms of feedback, as it only requires measuring the received interference plus noise (i.e., the quantity $\sum_{l \neq k} G_{lk}p_l + \sigma_k^2$).

After convergence of the primal-dual method (end of step 2 in the algorithm), each link checks if its SINR constraint is satisfied. If not, a distributed consensus process to select a link to drop is initiated by any link, via the control channel. In order for the link dropping heuristic described in 3.4 to be used, again certain quantities need to be communicated over the control channel.

Let p_k^e be the excess power needed for link k to attain its target SINR, assuming all other link powers are those obtained upon convergence. Link k computes the sum of *excess* interference caused to and received from neighboring links, i.e., $m_k := p_k^e \sum_{l \neq k} G_{kl} + \sum_{l \neq k} G_{lk}p_l^e$. This requires that link k also knows G_{lk}, p_l^e for the links it receives interference from. This information can be locally shared using the control channel. A distributed consensus-on-the-max algorithm can then be employed over the control channel to reach agreement on the index of the link with maximum m_k and drop that link.

Distributed consensus algorithms have attracted considerable interest in signal processing lately, sparked by the work of Xiao and Boyd [27], among others. Distributed consensus has a longer history though, including the case of consensus-on-the-max and general functions; see [6] and references therein. A distributed flow that achieves consensus-on-the-max in finite time for strongly connected graphs is given in [6]. A conceptually simpler discrete-time approach is to let each node compute a local maximum at each time-step. If the graph is strongly connected, this will yield consensus on the global maximum in at most r steps, where r is the radius of the graph. This assumes that interim estimates are exchanged between neighbors at each time step, however it is easy to relax this requirement and still guarantee convergence, under mild assumptions.

Chapter 6

Simulations

We carried out three sets of experiments: centralized with perfect CSI, distributed with perfect CSI and robust centralized with CSI uncertainty. In each case we examined scenarios with and without a primary user, to cover cognitive radio and ad-hoc settings, respectively. In all our simulations we tested the ability of each algorithm to admit a close to optimal (as given by enumeration) number of users for varying K (user population), target SINR, or channel gain uncertainty in the robust case.

In the sequel, each figure reports Monte-Carlo (MC) average results for at least 300 MC runs. For each MC run, transmitter locations are uniformly drawn on a 2 Km $\times$ 2 Km square. For each transmitter location, a receiver location is drawn uniformly in a disc of radius 400 meters, excluding a radius of 10 meters. The power budget for any link k is given by $P_k^{MAX} = bP_k^{MIN}$, where P_k^{MIN} is the minimum power required for the link to satisfy its SINR constraint in the absence of any interference. The primary user's power is fixed to P_k^{MAX}. Link gains are calculated by $G_{ij} = 1/d_{ij}^4$ where d_{ij} is the Euclidean distance between transmitter i and receiver j, and receiver noise is set to -60 dBm. For our relaxation-based algorithms, the δ_k are kept close to the respective bounds (specifically at 0.999 times the value given by (3.6)) and ϵ is set to one order of magnitude smaller than the upper bound given by (3.5).

6.1 Centralized Algorithm under Perfect CSI

Results for this set of experiments are summarized in Figures 6.1 to 6.5. As a baseline for our LPD algorithm, we implemented the gradual removals GRN-DCPC algorithm of [1]. This algorithm was not developed for a cognitive radio scenario (it does not account for interference to the primary user). Despite its age, [1] still represents the state-of-art in the case when no primary users are considered. The heuristic used was 'SMART' as described in [1]. In the course of implementing this algorithm, we came up with an improved variant, which we also included in our simulations under the name 'GRN-DCPC SMART Modified'[1].

In order to include the ultimate upper bound in our comparisons, we also developed a carefully optimized stack-based enumeration algorithm that always finds the optimum solution for modest problem sizes (up to 20 secondary users). This works by either

[1]The modification consists of normalizing cross gains by the transmitter's self link gain, instead of the receiver's self link gain. Using the notation in [1] (beware of the reversed role of indices) this translates to $\alpha_{ij} = g_{ij}/g_{jj}$ instead of the original $\alpha_{ij} = g_{ij}/g_{ii}$ (for $j \neq i$).

growing or pruning the candidate set of users. In growing mode, once an infeasible set has been detected, its supersets are not tested; in pruning mode, once a feasible set has been found, its subsets are not tested. The code was verified against brute-force enumeration in extensive Monte-Carlo experiments for up to 12 users.

In all experiments in this section, except for figure 6.3 we set the power budget coefficient to $b = 5$. A comparison of LPD, the two flavors of GRN-DCPC, and the optimal solution (via enumeration) in terms of the average number of admitted users versus the user population, K, is provided in Figure 6.1 for $c_k = 0$ dB and $c_k = 8$ dB. Our modification of 'GRN-DCPC SMART' performs better than the original and LPD performs very close to optimal for the range considered. Figure 6.2 shows the average number of admitted users versus a larger number of candidate users, illustrating the increasing gap of LPD relative to both flavors of GRN-DCPC.

The transition from the power limited to the interference limited regime is illustrated in Figure 6.3, as the average number of admitted users over the power budget coefficient b. There we can see a 'law of diminishing returns'-type behavior, where gains from power are only reaped in the early stages of increasing the power budget.

Figures 6.4 and 6.5 depict results in a cognitive radio setting. The primary transmitter and receiver, when present, are located on an edge of the 2 Km $\times$ 2 Km square, 1 Km apart and symmetrically with respect to the edge midpoint. For Figure 6.4, a single primary user is present with $c_0 = 2dB$, and for the secondary users $c_k = 2dB$, or $c_k = 5dB$. Figure 6.5 shows the average number of admitted users versus the secondary user's SINR target, with or without a primary user (curves marked $P = 1$ or $P = 0$, respectively) with $c_0 = 2dB$. In this case, the number of admitted users decreases roughly linearly with respect to the SINR target in dB. In both figures we notice that our LPD algorithm performs close to optimal in the scenarios considered.

6.2 Centralized Algorithm with Imperfect CSI

In order to assess the performance of our robust SOCD algorithm, we use the same simulation setup as in our previous experiments. The new element lies in our modeling of channel gain uncertainty. As already described in Section 4, for any given receiver, the receiving gains are assumed to be lying in an ellipsoid centered on the nominal gain values. Furthermore, for the purpose of these simulations we assume diagonal ellipsoid matrices and perfect self link gain knowledge. Specifically, the entries of the ellipsoid matrix E_k are given by:

$$E_k(i,j) = \begin{cases} \eta_k G_{ik}, & i = j \text{ and } i \neq k \\ 0, & \text{otherwise} \end{cases},$$

where $\eta_k \in [0, 1)$ represents the level of uncertainty for the receiving gains estimated by receiver k. The amount of this uncertainty is a fraction of the actual gains, modeling an additive uncertainty for an estimate in dB.

The deflation algorithm employed here is the robust SOCD described in section 4. Only enumeration is available for comparison in the robust case. This is similar to the enumeration algorithm used in our earlier simulations, only this time using the SOCP formulation of Section 4 to test user subsets for admissibility. Figure 6.6 shows the average number of admitted users versus the total number of users for $c_k = 2$, no primary user present, and uncertainty coefficients $\eta_k = 0.1$ or $\eta_k = 0.9$. Figure 6.7 shows the average number of admitted users for 10 candidate users, versus the uncertainty coefficient η_k. For this figure $c_k = 0$, one set of curves is without a primary user and the other set includes a primary user with $c_0 = 2$ and higher estimation uncertainty than the more versatile secondary users ($\eta_0 = 2\eta_k$). Again our SOCD algorithm performs very close to optimal.

6.3 Distributed Algorithm

To assess the performance of our distributed LPD algorithm we will again compare to the two flavors of GRN-DCPC and enumeration, described in Section 6.1. We present indicative simulation results for both an ad-hoc scenario without a primary user (to enable comparison with [1]), and a cognitive radio scenario with a primary user present. In all experiments in this section we set the power budget coefficient to $b = 2$.

For the distributed-algorithm-specific parameters discussed in Section 5 we set $\theta = 0.2$, and the initial step-size was empirically set to $\alpha_0 = 1$. The dual variables were initialized as $\mu_k \sim 1/G_{kk}$, and the slacks ρ_k were normalized by $G_{kk}\delta_k^{-1}$ to bring the different links to scale and ensure approximately equal rates of convergence. A maximum of 5K iterations were allowed for the primal-dual distributed solver of the relaxed problem, followed by a final phase that linearly brings α_i to 0 in 500 iterations, thus damping any residual oscillation.

Figure 6.9 reports the average number of users admitted versus the total number of users for $c_k = 2$, for enumeration, the two flavors of GRN-DCPC, the centralized LPD algorithm, and its distributed counterpart, with or without a primary user with $c_0 = 2$. Since the GRN-DCPC algorithms are not applicable in scenarios with primary users, they are omitted in the second set of curves. Finally, figure 6.10 shows the average number of users admitted versus the secondary users' SINR target for 12 users and $c_0 = 2$.

We notice that our distributed LPD performs the same as the centralized LPD, which is a significant improvement over 'GRN-DCPC SMART'. Our modification of 'GRN-DCPC SMART' performs close to LPD in this simulation, however we would like to point again to the results in figure 6.2, which demonstrate the clear superiority of LPD for a large number of users.

For the purpose of discussing the communication requirements and solution speed of our distributed algorithm, let us give an example. Assume a 10 Mbps control channel. At every iteration, every user has to broadcast its dual variable μ_k. A conservative estimate of the message size including coding and user ID gives us a packet of 50 bits. Assuming a total of 10 users this translates to 20K iterations or approximately 4 link removals per second. Compared to this, simpler algorithms like GRN-DCPC [1] (or its improved variant proposed herein) take only a small fraction of the time, making the use of distributed deflation worth when we do admission control infrequently (for relatively longer transmission rounds) and/or in difficult scenarios where we need to squeeze-in the maximum possible number of users.

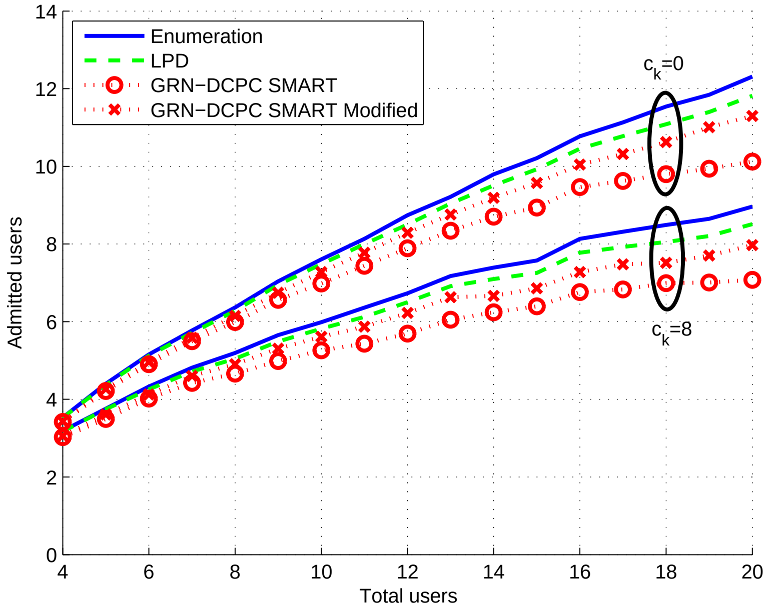

Figure 6.1: Mean number of admitted users vs. total number of candidate users, for $c_k = 0$ and $c_k = 8$ dB.

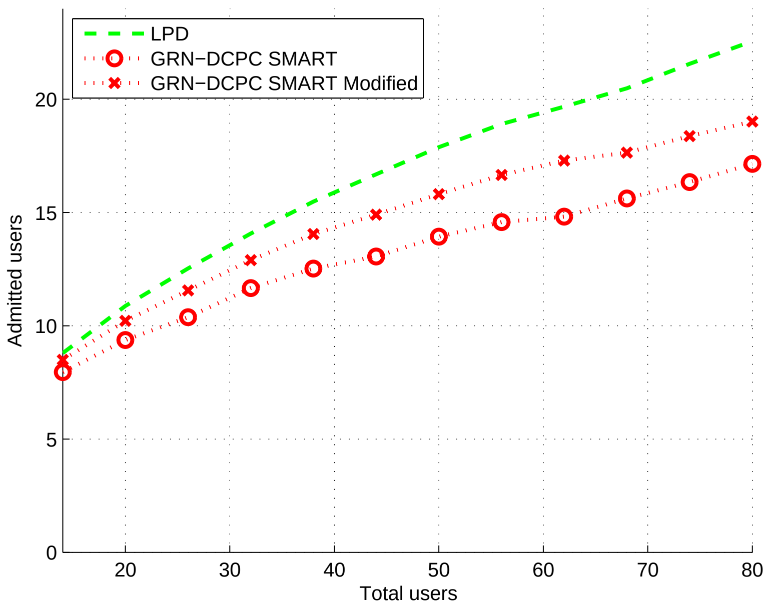

Figure 6.2: Mean number of admitted vs. total number of users, for a large candidate population and $c_k = 2$.

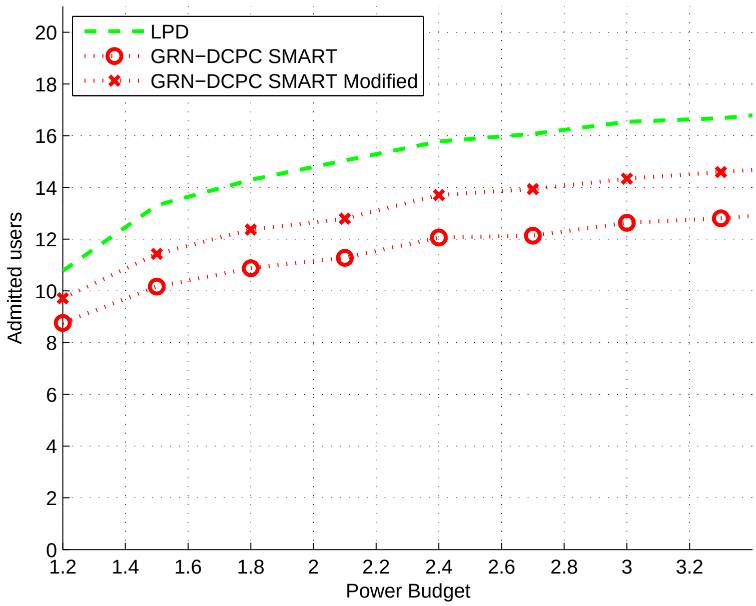

Figure 6.3: Mean number of admitted users vs. power budget coefficient b, for $c_k = 2$ and 50 users.

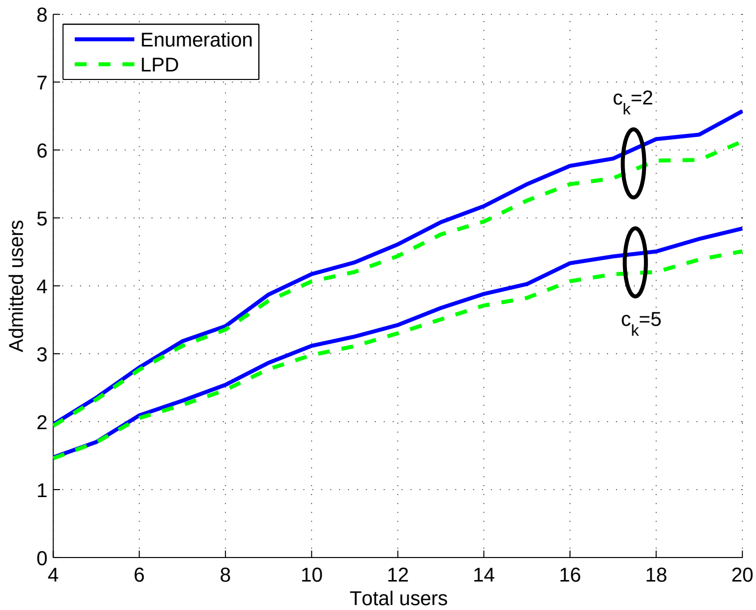

Figure 6.4: Mean number of admitted secondary users vs. total number of secondary users, for a single primary link with $c_0 = 2$ dB and secondary SINR $c_k = 2$ or $c_k = 5$.

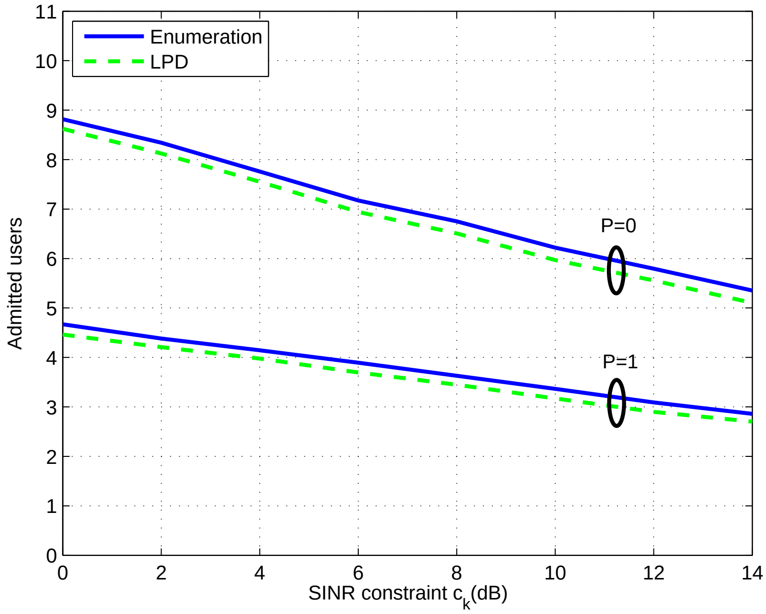

Figure 6.5: Mean number of admitted users vs. secondary user SINR constraint for 12 secondary users, with $(P = 1)$ / without $(P = 0)$ a primary user with $c_0 = 2$.

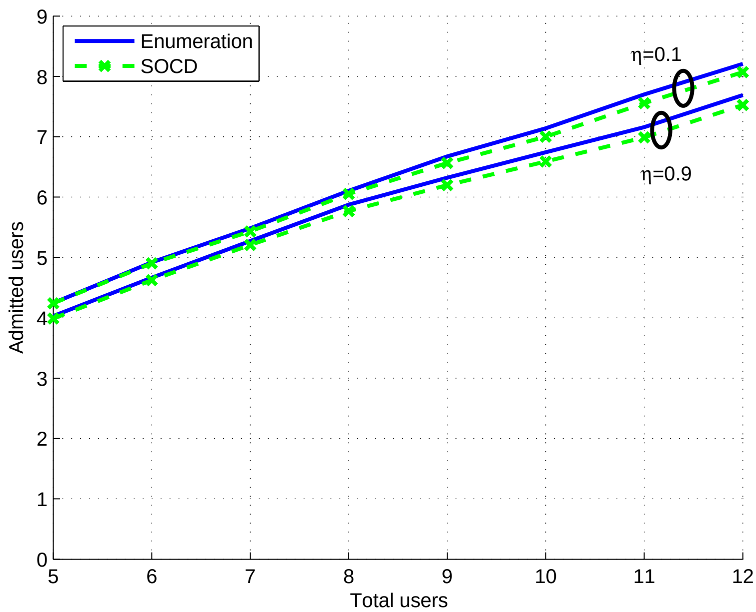

Figure 6.6: Robust case: Mean number of admitted users for $c_k = 2$, and $\eta_k = 0.1$ or $\eta_k = 0.9$.

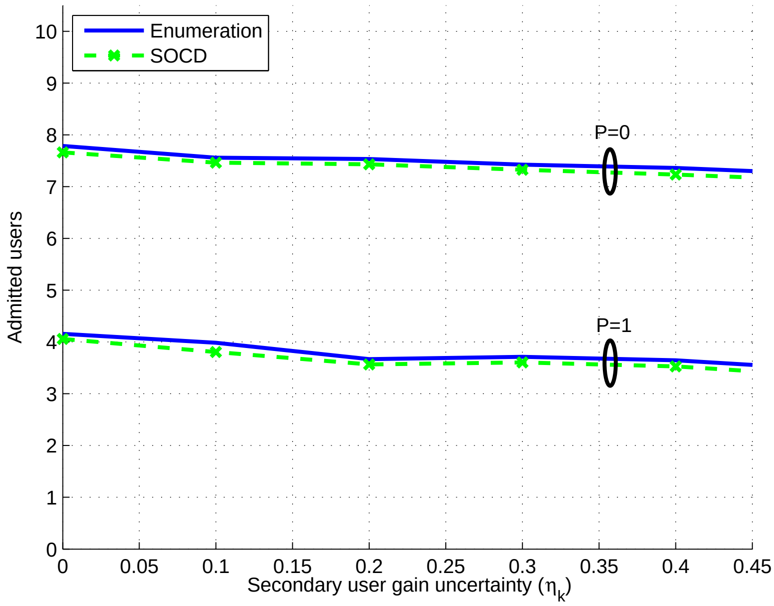

Figure 6.7: Robust case: Mean number of admitted users vs. uncertainty η_k, for 10 secondary users with $c_k = 0$, with $(P = 1)$ / without $(P = 0)$ a primary user. For the primary user $c_0 = 2$, gain uncertainty set to $\eta_0 = 2\eta_k$.

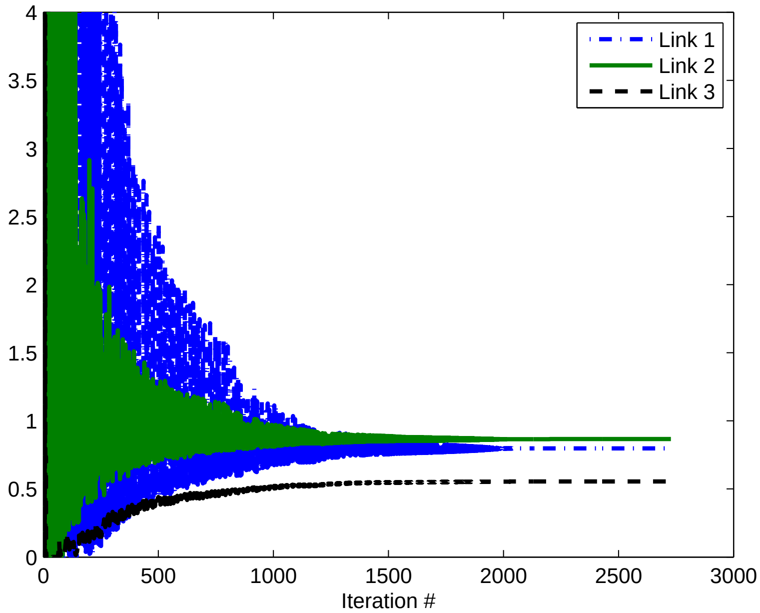

Figure 6.8: Distributed implementation: Convergence of primal t_k (infeasible scenario of 3 users).

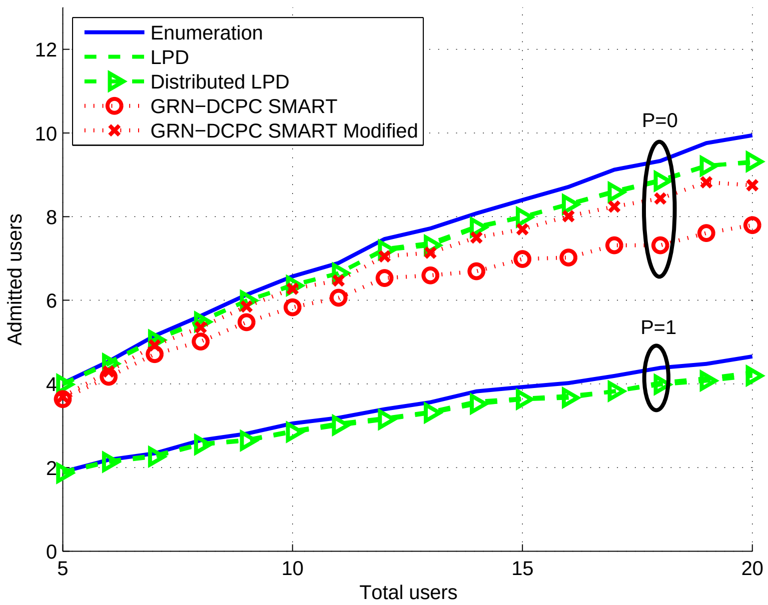

Figure 6.9: Distributed implementation: Mean number of admitted users vs. total number of secondary users with $c_k = 2$, with $(P = 1)$ / without $(P = 0)$ a primary user with $c_0 = 2$.

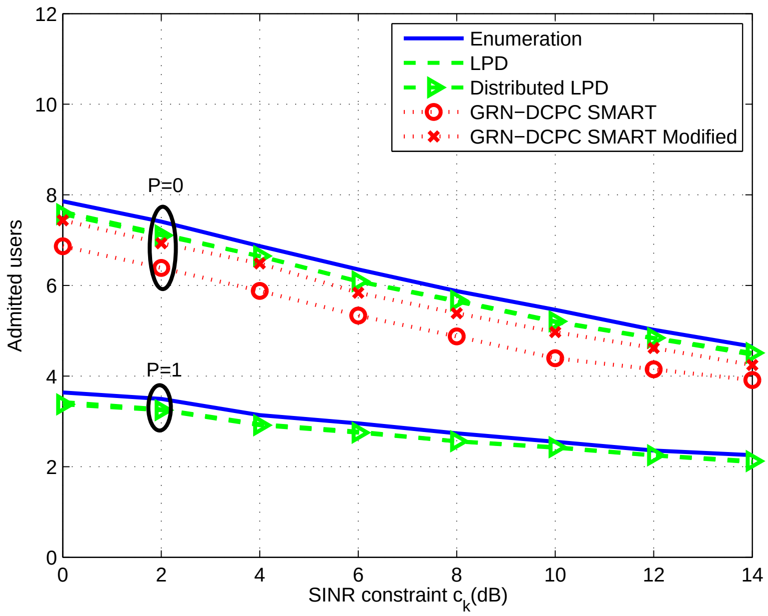

Figure 6.10: Distributed case: Mean number of admitted users vs. SINR target for secondary users, with $(P = 1)$ / without $(P = 0)$ a primary user with $c_0 = 2$.

Chapter 7

Discussion and Conclusions

Our results suggest that the proposed LPD algorithm is very promising. It indeed comes close to attaining the performance of the optimal solution at the cost of solving $O(K)$ LP problems in the worst case. This requires a fraction of a second on a current personal computer, as opposed to several minutes needed for enumeration for $K = 20$, which is modest. LPD clearly outperforms the state-of-art when no primary users are considered. This is already important, because the joint admission and power control problem has been under scrutiny for many years. Interestingly, our robust solution (the SOCD algorithm) appears to have an even smaller gap relative to the optimal robust solution.

We have also developed a distributed implementation of the joint admission and power control algorithm. The new implementation alternates between a distributed approximation phase and a distributed deflation phase. The latter employs consensus-on-the-max to select a link to drop, if needed. Both phases require local communication and computation. Still, communication and computation requirements are considerably higher than those of simpler heuristic solutions, making distributed deflation worth its cost in relatively challenging scenarios, or when we schedule for (and costs are amortized over) longer horizons.

Bibliography

[1] M. Andersin, Z. Rosberg, and J. Zander, “Gradual Removals in Cellular PCS with Constrained Power Control and Noise,” *Wireless Networks*, vol. 2, pp. 27–43, 1996.

[2] N. Bambos, S.C. Chen, and G.J. Pottie, “Channel access algorithms with active link protection for wireless communication networks with power control,” *IEEE/ACM Trans. on Networking*, vol. 8, no. 5, pp. 583–597, October 2000.

[3] A. Behzad, I. Rubin, and P. Chakravarty, “Optimum Integrated Link Scheduling and Power Control for Ad Hoc Wireless Networks,” in *Proc. 2005 Wireless And Mobile Computing, Networking And Communications, Conference* (WiMob2005), vol. 3, pp. 275–283, 2005.

[4] D. P. Bertsekas, *Convex Optimization Theory*, Athena Scientific, 2009.

[5] S. Boyd, and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004; see also http://www.stanford.edu/∼boyd/cvxbook.html.

[6] Jorge Cortes, “Distributed algorithms for reaching consensus on general functions,” *Automatica*, vol. 44, no. 3, pp. 726-737, March 2008.

[7] T. Elbatt and A. Ephremides, “Joint Scheduling and Power Control for Wireless *Ad Hoc* Networks,” *IEEE Trans. on Wireless Communications*, vol. 3, no. 1, pp. 74–85, 2004.

[8] Federal Communications Commission, “Spectrum Policy Task Force,” Rep. ET Docket no. 02-135, Nov. 2002.

[9] M.R. Garey and D.S. Johnson, *Computers and Intractability. A Guide to the Theory of NP-Completeness*. W.H. Freeman and Company, 1979.

[10] S.A. Grandhi, J. Zander, and R. Yates, “Constrained Power Control,” *Wireless Personal Commun.*, vol. 1, no. 4, April 1995.

[11] K. Hamdi, W. Zhang, and K.-B. Letaief, “Joint Beamforming and Scheduling in Cognitive Radio Networks,” in *Proc. GLOBECOM 2007*, pp. 2977-2981, Nov. 26-30, 2007.

[12] M. H. Islam, Y. C. Liang, and A. T. Hoang, “Joint beamforming and power control in the downlink of cognitive radio networks,” in *Proc. IEEE WCNC 2007*, Hong Kong, Mar. 2007.

[13] Z.-Q. Luo, N.D. Sidiropoulos, P. Tseng, S. Zhang, “Approximation Bounds for Quadratic Optimization with Homogeneous Quadratic Constraints,” *SIAM J. Optim.*, vol. 18, no. 1, pp. 1-28, Feb. 2007.

[14] Z.-Q. Luo, W.-K. Ma, A. M.-C. So, Y. Ye, and S. Zhang, “Nonconvex quadratic optimization, semidefinite relaxation, and applications,” *IEEE Signal Processing Magazine*, Special Issue on Convex Optimization for Signal Processing, May 2010 (to appear).

[15] Z. Marantz, P. Orenstein, and D.J. Goodman, “Admission control for maximal throughput in power limited CDMA systems,” in *Proc. IEEE WCNC 2005*, vol. 2, pp. 695–700, Mar. 2005.

[16] M. J. Marcus, “The underlay/overlay controversy,” *IEEE Wireless Commun.*, vol. 14 no. 5, pp. 4-5, Oct. 2007.

[17] E. Matskani, N.D. Sidiropoulos, Z.-Q. Luo, and L. Tassiulas,“Convex Approximation Techniques for Joint Multiuser Downlink Beamforming and Admission Control,” *IEEE Trans. on Wireless Communications*, vol. 7, no. 7, pp. 2682–2693, July 2008.

[18] I. Mitliagkas, N.D. Sidiropoulos, and A. Swami, “Convex Approximation-Based Joint Power and Admission Control for Cognitive Underlay Networks,” in *Proc. IEEE IWCMC’08*, pp. 28–32, 2008.

[19] I. Mitliagkas, N.D. Sidiropoulos, and A. Swami, “Distributed Joint Power and Admission Control for Ad-hoc and Cognitive Underlay Networks,” in *Proc. IEEE ICASSP 2010*, Mar. 14-19, 2010, Dallas, TX.

[20] J. Mitola III, and J.Q. Maguire, “Cognitive Raio: Making Software Radios More Personal,” *IEEE Personal Communications*, pp. 13-18, Aug. 1999.

[21] D. P. Palomar and M. Chiang, “A tutorial on decomposition methods for network utility maximization,” *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 8, pp. 1439–1451, 2006.

[22] S. Stanczak, M. Kaliszan, and N. Bambos, “Admission control for power-controlled wireless networks under general interference functions,” in *Proc. 42nd Asilomar Conference on Signals, Systems, and Computers*, Monterey, CA, USA, Oct. 2008.

[23] S. Stanczak, M. Kaliszan, and N. Bambos, “Admission Control for Autonomous Wireless Links with Power Constraints,” in *Proc. IEEE ICASSP 2010*, Mar. 14-19, 2010, Dallas, TX.

[24] S. Stanczak, M. Kaliszan, and N. Bambos, “Decentralized Admission Control for Power-Controlled Wireless Links,” submitted to *IEEE Trans. on Signal Processing*, 2009.

[25] J. T. Wang, “Admission control with distributed joint diversity and power control for wireless networks,” *IEEE Trans. Vehicular Technology*, vol. 58, no. 1, pp. 409–419, Jan. 2009.

[26] H. Wolkowicz, “Relaxations of Q2P,” Chapter 13.4 in *Handbook of Semidefinite Programming: Theory, Algorithms, and Applications*, H. Wolkowicz, R. Saigal, L. Vandenberghe (Eds.), Kluwer Academic Publishers, 2000.

[27] L. Xiao, and S. Boyd, “Fast Linear Iterations for Distributed Averaging,” *Systems and Control Letters*, vol. 53, pp. 65-78, 2004.

[28] M. Xiao, N. B. Shroff, and E. K. P. Chong, “Distributed admission control for power-controlled cellular wireless systems,” *IEEE/ACM Trans. on Networking*, vol. 9, no. 6, pp. 790–800, 2001.

[29] J. Zander, “Performance of Optimum Transmitter Power Control in Cellular Radio Systems,” *IEEE Trans. on Vehicular Technology*, vol. 41, pp. 57–62, 1992.

[30] L. Zhang, Y.-C. Liang, and Y. Xin, “Joint Beamforming and Power Allocation for Multiple Access Channels in Cognitive Radio Networks,” in *Journal on Selected Areas in Communications*, vol. 26, no. 1, pp. 38-51, Jan. 2008.

[31] Q. Zhao, and B. M. Sadler, “A survey of dynamic spectrum access,” *IEEE Signal. Proc. Magazine*, pp. 79-89, May 2007.

[32] Q. Zhao and A. Swami, “A survey of dynamic spectrum access: signal processing and networking perspectives,” in *Proc. IEEE ICASSP 2007*, vol. 4, 2007.

[33] Q. Zhao, L. Tong, A. Swami and Y. Chen, “Decentralized cognitive MAC for opportunistic spectrum access in ad hoc networks: a POMDP framework,” *IEEE J. Selected Areas in Communications*, vol. 25 no. 3, pp. 589-600, Apr. 2007.

[34] A. Zymnis, N. Trichakis, S. Boyd, and D. O’Neill, “An Interior-Point Method for Large Scale Network Utility Maximization,” in *Proc. Allerton Conference on Communication, Control, and Computing*, 2007.