Distributed Interference Compensation for Wireless Networks

Jianwei Huang, Randall Berry, Michael L. Honig

Abstract

We consider a distributed power control scheme for wireless ad hoc networks, in which each user announces a price that reflects compensation paid by other users for their interference. We present an asynchronous distributed algorithm for updating power levels and prices. By relating this algorithm to myopic best response updates in a fictitious game, we are able to characterize convergence using supermodular game theory. Extensions of this algorithm to a multi-channel network are also presented, in which users can allocate their power across multiple frequency bands.

I. Introduction

Mitigating interference is a fundamental problem in wireless networks. A basic technique for this is to control the nodes' transmit powers. In an ad hoc wireless network power control is complicated by the lack of centralized infrastructure, which necessitates the use of distributed approaches. This paper addresses distributed power control for rate adaptive users in a wireless network. We consider two models: a single channel spread spectrum (SS) network, where all users spread their power over a single frequency band, and a multi-channel model, where each user can allocate its power over multiple frequency bands. The latter model is motivated by multi-carrier transmission (e.g., Orthogonal Frequency Division Multiplexing (OFDM)), where each channel might represent a single carrier, or a group of adjacent carriers. In both cases, the transmission rate for each user depends on the received signal-to-interference plus noise ratio (SINR). Our objective is to coordinate user power levels to optimize overall performance, measured in terms of total network utility.

The authors are with the Department of ECE, Northwestern University, 2145 Sheridan Road, Evanston, IL 60208 USA, email: {jianweih,rberry,mh}@ece.northwestern.edu. This work was supported by the Northwestern-Motorola Center for Communications, ARO under grant DAAD19-99-1-0288, and NSF CAREER award CCR-0238382.

We study protocols in which the users exchange price signals that indicate the “cost” of received interference. Pricing mechanisms for allocating resources in networks have received considerable attention for both wire-line (e.g. [1], [15]) and wireless networks (e.g. [2]–[4]). The problem here differs from much of the previous work because, due to interference, the users’ objective functions are coupled, and the overall network objective may not be concave in the allocated resource (transmit power). Also, in most previous work, prices are Lagrange multipliers for some constrained resource such as power or bandwidth; here the prices reflect the interference or externalities among the users instead of a resource constraint. Our single channel model is similar to that considered in [5], which also discusses combined power and rate control. The power adaptation in [5] solves a similar problem to that considered here using gradient updates. Instead, we consider an approach based on supermodular game theory [6], which allows for a larger class of utility functions and appears to have faster convergence.

A variety of game-theoretic approaches have been applied to network resource allocation, as surveyed in [7]. Supermodular game theory, in particular, has been used to study power control in [8]–[10]. Our approach differs in that (i) we focus on an ad hoc instead of a cellular network; (ii) we consider a different functional form for the utilities than some authors, and (iii) we do not directly model the problem as a non-cooperative game. Instead, the users voluntarily cooperate with each other by exchanging interference information. We introduce a *fictitious game* and apply a strategy space transformation to view this algorithm as a supermodular game. Other work on power control in CDMA cellular and ad hoc networks includes [8], [9], [11]–[13]. In most prior work on ad hoc networks, a transmission is assumed to be successful if a fixed minimum SINR requirement is met. This is true for fixed-rate communications. However, this is not the case for “elastic” data applications, which can adapt transmission rates. In this paper, we focus on rate-adaptive users, where the goal of power control is to maximize total network performance instead of guarantee interference margins for each user.

For multi-channel networks, an additional consideration is how the users allocate their power across the available channels. We decompose this power allocation by introducing a “power price” for each user, which represents a dual variable corresponding to the user’s total power constraint. Each user must now take into account both the interference prices and their own power price.

We present a distributed gradient projection algorithm to solve for the optimal power prices. This is similar in spirit to the optimization flow control algorithm for wire-line networks in [15]. However, here the dual variables are not determined by each link in the network, but rather by each user. Also, the corresponding primal problem is not separable due to the interference.

In the next section, we describe and analyze our distributed price/power adjustment algorithm for a single channel network. We then turn to the multi-channel model in Sect. III. Simulation results are given in Sect. IV, and conclusions are presented in Sect. V.

II. Single Channel Networks

We consider a snap-shot of an ad hoc network with a set $\mathcal{M} = \{1, ..., M\}$ of distinct node pairs. As shown in Fig. 1, each pair consists of one dedicated transmitter and one dedicated receiver.[1] We use the terms "pair" and "user" interchangeably in the following. In this section, we assume that each user i transmits an SS signal spread over the total bandwidth of B Hz. Over the time-period of interest, the channel gains of each pair are fixed. The channel gain between user i's transmitter and user j's receiver is denoted by h_{ij}. Note that in general $h_{ij} \neq h_{ji}$, since the latter represents the gain between user j's transmitter and user i's receiver.

Each user i's quality of service is characterized by a utility function $u_i(\gamma_i)$, which is an increasing and strictly concave function of the received SINR,

$$\gamma_i(\boldsymbol{p}) = \frac{p_i h_{ii}}{n_0 + \frac{1}{B}\sum_{j \neq i} p_j h_{ji}}, \tag{1}$$

where n_0 is the background noise power and $\boldsymbol{p} = (p_1, \cdots, p_M)$ is a vector of the users' transmission powers. The users' utility functions are coupled due to mutual interference. An example utility function is a *logarithmic utility function* $u_i(\gamma_i) = \theta_i \log(\gamma_i)$, where θ_i is a user dependent priority parameter.[2]

[1] For example, this could represent a particular schedule of transmissions determined by an underlying routing and MAC protocol.

[2] In the high SINR regime, logarithmic utility approximates the Shannon capacity $\log(1 + \gamma_i)$ weighted by θ_i. For low SINR, a user's rate is approximately linear in SINR, and so this utility is proportional to the logarithm of the rate.

The problem we consider is to specify $\boldsymbol{p}$ to maximize the utility summed over all users, where each user i must also satisfy a transmission power constraint, $p_i \in \mathcal{P}_i = [P_i^{min}, P_i^{max}]$, i.e.,

$$\max_{\{\boldsymbol{p}: p_i \in \mathcal{P}_i \, \forall i\}} \sum_{i=1}^{M} u_i\left(\gamma_i(\boldsymbol{p})\right). \tag{P1}$$

Note that a special case is $P_i^{min} = 0$; i.e., the user may choose not to transmit.[3]

As a baseline distributed approach, consider the case where the users do not exchange any information and simply choose transmission powers to maximize their individual utilities. As in [8], this can be modeled as a *non-cooperative power (NCP) control game* $G_{NCP} = [\mathcal{M}, \{\mathcal{P}_i\}, \{u_i\}]$, where the players in the game correspond to the users in $\mathcal{M}$; each player picks a transmission power from the strategy set $\mathcal{P}_i$ and receives a payoff $u_i(\gamma_i)$. In this game $\boldsymbol{p}$ is the *power profile*, and the *power profile of user i's opponents* is defined to be $p_{-i} = (p_1, ..., p_{i-1}, p_{i+1}, ..., p_M)$, so that $\boldsymbol{p} = (p_i; p_{-i})$. Similar notation will be used for other quantities. User i's *best response* is $\mathcal{B}_i(p_{-i}) = \arg\max_{p_i \in P_i} u_i\left(\gamma_i(p_i, p_{-i})\right)$, i.e., the p_i that maximizes $u_i\left(\gamma_i\left(p_i, p_{-i}\right)\right)$ given a fixed p_{-i}. A power profile $\boldsymbol{p}^*$ is a *Nash Equilibrium (NE)* of G_{NCP} if it is a fixed point of the best responses, i.e. $u_i(\gamma_i(p_i^*; p_{-i}^*)) \geq u_i(\gamma_i(p_i'; p_{-i}^*))$ for any $p_i' \in \mathcal{P}_i$ and any user i.

Since each user's payoff $u_i\left(\gamma_i\left(p_i, p_{-i}\right)\right)$ is strictly increasing with p_i for fixed p_{-i}, and there is no penalty for high transmission power as long as $p_i \in \mathcal{P}_i$, it is easy to verify that the unique NE of G_{NCP} is $\boldsymbol{p}_{NCP}^* = (P_i^{\max})_{i=1}^{M}$, i.e., each transmitter uses its maximum power. This solution can be far from the socially optimal solution given by Problem P1.

Although $u_i(\cdot)$ is concave, the objective in Problem P1 may not be concave in $\boldsymbol{p}$. However, it is easy to verify that any local optimum, $\boldsymbol{p}^* = (p_1^*, ..., p_M^*)$, of this problem will be regular (see p. 309 of [16]), and so must satisfy the Karush-Kuhn-Tucker (KKT) necessary conditions:

Lemma 1 (KKT conditions:): For any local maximum $\boldsymbol{p}^*$ of Problem P1, there exist unique Lagrange multipliers $\lambda_{1,u}^*, ..., \lambda_{M,u}^*$ and $\lambda_{1,l}^*, ..., \lambda_{M,l}^*$ such that for all $i \in \mathcal{M}$,

$$\frac{\partial u_i\left(\gamma_i\left(\mathbf{p}^*\right)\right)}{\partial p_i} + \sum_{j \neq i} \frac{\partial u_j\left(\gamma_j\left(\mathbf{p}^*\right)\right)}{\partial p_i} = \lambda_{i,u}^* - \lambda_{i,l}^*, \tag{2}$$

[3]Occasionally, for technical reasons, we require $P_i^{min} > 0$; in these cases, P_i^{min} can be chosen arbitrarily small so that this restriction has little effect. Note that for certain utilities, e.g., $\theta_i \log(\gamma_i)$, all assigned powers must be strictly positive, since as $P_i \to 0$, the utility approaches $-\infty$.

$$\lambda_{i,u}^*(p_i^* - P_i^{max}) = 0,\ \lambda_{i,l}^*(P_i^{min} - p_i^*) = 0,\ \lambda_{i,u}^*, \lambda_{i,l}^* \geq 0. \tag{3}$$

Let

$$\pi_j\left(p_j, p_{-j}\right) = -\frac{\partial u_j\left(\gamma_j\left(p_j, p_{-j}\right)\right)}{\partial I_j\left(p_{-j}\right)}, \tag{4}$$

where $I_j\left(p_{-j}\right) = \sum_{k\neq j} p_k h_{kj}$ is the total interference received by user j (before bandwidth scaling). Here, $\pi_j\left(p_j, p_{-j}\right)$ is always nonnegative and represents user j's marginal increase in utility per unit decrease in total interference. Using (4), condition (2) can be written as

$$\frac{\partial u_i\left(\gamma_i\left(\mathbf{p}^*\right)\right)}{\partial p_i} - \sum_{j\neq i} \pi_j\left(p_j^*, p_{-j}^*\right) h_{ij} = \lambda_{i,u}^* - \lambda_{i,l}^*. \tag{5}$$

Viewing $\pi_j\ (= \pi_j\left(p_j, p_{-j}\right))$ as a *price* charged to other users for generating interference to user i, condition (5) is a necessary and sufficient optimality condition for the problem in which each user i specifies a power level $p_i \in \mathcal{P}_i$ to maximize the following surplus function

$$s_i\left(p_i; p_{-i}, \pi_{-i}\right) = u_i\left(\gamma_i\left(p_i, p_{-i}\right)\right) - p_i \sum_{j\neq i} \pi_j h_{ij}, \tag{6}$$

assuming fixed p_{-i} and π_{-i} (i.e., each user is a price taker and ignores any influence he may have on these prices). User i therefore maximizes the difference between its utility minus its payment to the other users in the network due to the interference it generates. The payment is its transmit power times a weighted sum of other users' prices, with weights equal to the channel gains between user i's transmitter and the other users' receivers. This pricing interpretation of the KKT conditions motivates the following asynchronous distributed pricing (ADP) algorithm.

A. Asynchronous Distributed Pricing (ADP) Algorithm

In the ADP algorithm, each user announces a single price and all users set their transmission powers based on the received prices. Prices and powers are asynchronously updated. For $i \in \mathcal{M}$, let $T_{i,p}$ and $T_{i,\pi}$, be two unbounded sets of positive time instances at which user i updates its power and price, respectively. User i updates its power according to $\mathcal{W}_i(p_{-i}, \pi_{-i}) = \arg\max_{\hat{p}_i \in \mathcal{P}_i}\ s_i\left(\hat{p}_i; p_{-i}, \pi_{-i}\right)$, which corresponds to maximizing the surplus in (6). Each user updates its price according to $\mathcal{C}_i(\boldsymbol{p}) = -\frac{\partial u_i(\gamma_i(\boldsymbol{p}))}{\partial I_i(p_{-i})}$, which corresponds to (4). Using these update rules, the ADP algorithm is given in Fig. 2. Note that in addition to being asynchronous across users, each user also need not update its power and price at the same time.[4]

[4]Of course, simultaneous updates of powers and prices per user or synchronous updating across all users is a special case.

In the ADP algorithm not only are the powers and prices generated in a distributed fashion, but also each user only needs to acquire limited information. To see this note that the power update function can be written as

$$\mathcal{W}_i(p_{-i},\pi_{-i}) = \max\left[\min\left[\frac{p_i}{\gamma_i\left(\boldsymbol{p}\right)}g_i\left(\frac{p_i}{\gamma_i(\boldsymbol{p})}\left(\sum_{j\neq i}\pi_j h_{ij}\right)\right), P_i^{\max}\right], P_i^{\min}\right],$$

where $\frac{p_i}{\gamma_i(\boldsymbol{p})}$ is independent of p_i, and

$$g_i\left(x\right) = \begin{cases} \infty, & 0 \leq x \leq u_i'\left(\infty\right), \\ \left(u_i'\right)^{-1}\left(x\right), & u_i'\left(\infty\right) < x < u_i'\left(0\right), \\ 0, & u_i'\left(0\right) \leq x. \end{cases}$$

Likewise, the price update can be written as $\mathcal{C}_i\left(\boldsymbol{p}\right) = \frac{\partial u_i(\gamma_i(\boldsymbol{p}))}{\partial \gamma_i(\boldsymbol{p})}\frac{(\gamma_i(\boldsymbol{p}))^2}{Bp_i h_{ii}}$. From these expressions, it can be seen that to implement the updates, each user i only needs to know: (i) its own utility u_i, the current SINR γ_i and channel gain h_{ii}, (ii) the "adjacent" channel gains h_{ij} for $j \in \mathcal{M}$ and $j \neq i$, and (iii) the price profile $\boldsymbol{\pi}$. By assumption each user knows its own utility. The SINR γ_i and channel gain h_{ii} can be measured at the receiver and fed back to the transmitter. Measuring the adjacent channel gains h_{ij} can be accomplished by having each receiver periodically broadcast a beacon; assuming reciprocity, the transmitters can then measure these channel gains. The adjacent channel gains account for only $1/M$ of the total channel gains in the network; each user does not need to know the other gains. The price information could also be periodically broadcast through this beacon. Since each user announces only a single price, the number of prices scales linearly with the size of the network. Also, numerical results show that there is little effect on performance if users only convey their prices to "nearby" transmitters, i.e., those generating the strongest interference.

Denote the set of fixed points of the ADP algorithm by

$$\mathcal{F}^{ADP} \equiv \left\{(\boldsymbol{p},\boldsymbol{\pi}) \,|\, (\boldsymbol{p},\boldsymbol{\pi}) = \left(\boldsymbol{\mathcal{W}}\left(\boldsymbol{p},\boldsymbol{\pi}\right), \boldsymbol{\mathcal{C}}(\boldsymbol{p})\right)\right\}, \tag{7}$$

where $\boldsymbol{\mathcal{W}}(\boldsymbol{p},\boldsymbol{\pi}) = (\mathcal{W}_k(p_{-k},\pi_{-k}))_{k=1}^M$ and $\boldsymbol{\mathcal{C}}(\boldsymbol{p}) = (\mathcal{C}_k(\boldsymbol{p}))_{k=1}^M$. Using the strict concavity of $u_i(\gamma_i)$ in γ_i, the following result can be easily shown.

Lemma 2: A power profile $\boldsymbol{p}^*$ satisfies the KKT conditions of Problem P1 (for some choice of Lagrange multipliers) if and only if $(\boldsymbol{p}^*, \mathcal{C}(\boldsymbol{p}^*)) \in \mathcal{F}^{ADP}$.

If there is only one solution to the KKT conditions, then it must be the global maximum and the ADP algorithm would reach that point if it converges.[5] In general, $\mathcal{F}^{ADP}$ may contain multiple points including local optima or saddle points.

B. Convergence Analysis of ADP Algorithm

We next characterize the convergence of the ADP algorithm by viewing it in a game theoretic context. A natural generalization of the NCP game is to consider a game where each player i's strategy includes specifying both a power p_i and a price π_i to maximize a payoff equal to the surplus in (6). However, since there is no penalty for user i announcing a high price, it can be shown that each user's best response is to choose a large enough price to force all other users transmit at P_i^{min}. This is certainly not a desirable outcome and suggests that the prices should be determined externally by another procedure.[6] Instead, we consider the following *Fictitious Power-Price (FPP) control game*, $G_{FPP} = [\mathcal{FW} \cup \mathcal{FC}, \left\{\mathcal{P}_i^{\mathcal{FW}}, \mathcal{P}_i^{\mathcal{FC}}\right\}, \left\{s_i^{\mathcal{FW}}, s_i^{\mathcal{FC}}\right\}]$, where the players are from the union of the sets $\mathcal{FW}$ and $\mathcal{FC}$, which are both copies of $\mathcal{M}$. $\mathcal{FW}$ is a *fictitious power player set*; each player $i \in \mathcal{FW}$ chooses a power p_i from the strategy set $\mathcal{P}_i^{\mathcal{FW}} = \mathcal{P}_i$ and receives payoff

$$s_i^{\mathcal{FW}}\left(p_i; p_{-i}, \pi_{-i}\right) = u_i\left(\gamma_i\left(\boldsymbol{p}\right)\right) - \sum_{j \neq i} \pi_j h_{ij} p_i. \tag{8}$$

$\mathcal{FC}$ is a *fictitious price player set*; each player $i \in \mathcal{FC}$ chooses a price π_i from the strategy set $\mathcal{P}_i^{\mathcal{FC}} = [0, \bar{\pi}_i]$ and receives payoff

$$s_i^{\mathcal{FC}}\left(\pi_i; \boldsymbol{p}\right) = -\left(\pi_i - \mathcal{C}_i\left(\boldsymbol{p}\right)\right)^2. \tag{9}$$

Here $\bar{\pi}_i = \sup_{\boldsymbol{p}} \mathcal{C}_i\left(\boldsymbol{p}\right)$, which could be infinite for some utility functions.

In G_{FPP}, each user in the ad hoc network is split into two fictitious players, one in $\mathcal{FW}$ who controls power p_i and the other one in $\mathcal{FC}$ who controls price π_i. Although users in the real network cooperate with each other by exchanging interference information (instead of choosing prices to maximize their surplus), each fictitious player in G_{FPP} is selfish and maximizes its own

[5] In the following section, we will give conditions under which this occurs.

[6] A similar situation arises in [2], where users in a multi-hop network announce prices charging other users for packets they forward. In that case, the prices also cannot be determined by individual surplus optimizations.

payoff function. In the rest of this section, a "user" refers to one of the M transmitter-receiver pairs in set $\mathcal{M}$, and a "player" refers to one of the $2M$ fictitious players in the set $\mathcal{FW} \cup \mathcal{FC}$.

In G_{FPP} the players' best responses are given by $\mathcal{B}_i^{\mathcal{FW}}(p_{-i}, \pi_{-i}) = \mathcal{W}_i(p_{-i}, \pi_{-i})$ for $i \in \mathcal{FW}$ and $\mathcal{B}_i^{\mathcal{FC}}(\boldsymbol{p}) = \mathcal{C}_i(\boldsymbol{p})$ for $i \in \mathcal{FC}$, where $\mathcal{W}_i$ and $\mathcal{C}_i$ are the update rules for the ADP algorithm. In other words, the ADP algorithm can be interpreted as if the players in G_{FPP} employ asynchronous *myopic best response (MBS)* updates, i.e. the players update their strategies according their best responses assuming the other player's strategies are fixed. It is known that the set of fixed points of MBS updates are the same as the set of NEs of a game [6, Lemma 4.2.1]. Therefore, we have:

Lemma 3: $(\boldsymbol{p}^*, \boldsymbol{\pi}^*) \in \mathcal{F}^{ADP}$ if and only if $(\boldsymbol{p}^*, \boldsymbol{\pi}^*)$ is a NE of G_{FPP}.

Together with Lemma 2, it follows that proving the convergence of asynchronous MBS updates of G_{FPP} is sufficient to prove the convergence of the ADP algorithm to a solution of KKT conditions. We next analyze this convergence using supermodular game theory [6].

We first introduce some definitions[7]. A real m-dimensional set $\mathcal{V}$ is a *sublattice* of $\mathbb{R}^m$ if for any two elements $a, b \in \mathcal{V}$, the component-wise minimum, $a \wedge b$, and the component-wise maximum, $a \vee b$, are also in $\mathcal{V}$. In particular, a compact sublattice has a (component-wise) smallest and largest element. A twice differentiable function f has *increasing differences* in variables (x, t) if $\partial^2 f / \partial x \partial t \geq 0$ for any feasible x and t.[8] A function f is *supermodular* in $\boldsymbol{x} = (x_1, .., x_m)$ if it has increasing differences in (x_i, x_j) for all $i \neq j$.[9] Finally, a game $G = [\mathcal{M}, \{\mathcal{P}_i\}, \{s_i\}]$ is *supermodular* if for each player $i \in \mathcal{M}$, (a) the strategy space $\mathcal{P}_i$ is a nonempty and compact sublattice, and (b) the payoff function s_i is continuous in all players' strategies, is supermodular in player i's own strategy, and has increasing differences between any component of player i's strategy and any component of any other player's strategy. The following theorem summarizes several important properties of these games.

Theorem 1: In a supermodular game $G = [\mathcal{M}, \{\mathcal{P}_i\}, \{s_i\}]$,

[7]More general definitions related to supermodular games are given in [6].

[8]If we choose x to maximize a twice differentiable function $f(x,t)$, then the first order condition gives $\partial f(x,t)/\partial x|_{x=x^*} = 0$, and the optimal value x^* increases with t if $\partial^2 f / \partial x \partial t > 0$.

[9]A function f is always supermodular in a single variable x.

(a) The set of NEs is a nonempty and compact sublattice and so there is a component-wise smallest and largest NE.

(b) If the users' best responses are single-valued, and each user uses MBS updates starting from the smallest (largest) element of its strategy space, then the strategies monotonically converge to the smallest (largest) NE.

(c) If each user starts from any feasible strategy and uses MBS updates, the strategies will eventually lie in the set bounded component-wise by the smallest and largest NE. If the NE is unique, the MBS updates globally converge to that NE from any initial strategies.

Properties (a) follows from Lemma 4.2.1 and 4.2.2 in [6]; (b) follows from Theorem 1 of [9] and (c) can be shown by Theorem 8 in [17].

Next we show that by an appropriate strategy space transformation certain instances of G_{FPP} are equivalent to supermodular games, and so Theorem 1 applies. We first study a simple two-user network, then extend the results to a M-user network.

1) Two-user networks: Let G^2_{FPP} be the FPP game corresponding to a two user network; this will be a game with four players, two in $\mathcal{FW}$ and two in $\mathcal{FC}$. First, we check whether G^2_{FPP} is supermodular. Each user $i \in \mathcal{FW}$ clearly has a nonempty and compact sublattice (interval) strategy set, and so does each user $i \in \mathcal{FC}$ if $\bar{\pi}_i < \infty$.[10] Each player's payoff function is (trivially) supermodular in its own one-dimensional strategy space. The remaining increasing difference condition for the payoff functions does *not* hold with the original definition of strategies $(\boldsymbol{p}, \boldsymbol{\pi})$ in G^2_{FPP}. For example, from (8), $\partial s_i^{\mathcal{FW}} / \partial p_i \partial \pi_j = -h_{ij} < 0$ for any $j \neq i$, e.g. a higher price leads the other users to decrease their powers. However, if we define $\pi'_j = -\pi_j$ and consider an equivalent game where each user $j \in \mathcal{FC}$ chooses π'_j from the strategy set $[-\bar{\pi}_j, 0]$, then $\partial s_i^{\mathcal{FW}} / \partial p_i \partial \pi'_j = h_{ij} > 0$, i.e. $s_i^{\mathcal{FW}}$ has increasing differences in the strategy pair $\left(p_i, \pi'_j\right)$ (or equivalently $(p_j, -\pi_j)$). If all the users' strategies can be redefined so that each player's payoff satisfies the increasing differences property in the transformed strategies, then the *transformed FPP game* is supermodular.

[10]When $P_i^{min} = 0$, this bounded price restriction is not satisfied for utilities such as $u_i(\gamma_i) = \theta_i \gamma_i^{\alpha} / \alpha$ with $\alpha \in [-1, 0)$, since $\pi_i = \theta_i \gamma_i^{\alpha+1} / (p_i h_{ii} B)$ is not bounded as $p_i \to 0$. However, as noted above, we can set P_i^{min} to some arbitrarily small value without effecting the performance.

Let $\gamma_i^{min} = \min\{\gamma_i(\boldsymbol{p}) : p_i \in \mathcal{P}_i \, \forall i\}$ and $\gamma_i^{max} = \max\{\gamma_i(\boldsymbol{p}) : p_i \in \mathcal{P}_i \, \forall i\}$. An increasing, twice continuously differentiable, and strictly concave utility function $u_i(\gamma_i)$ is defined to be

- Type I if $-\frac{\gamma_i u_i''(\gamma_i)}{u_i'(\gamma_i)} \in [1,2]$ for all $\gamma_i \in [\gamma_i^{min}, \gamma_i^{max}]$;
- Type II if $-\frac{\gamma_i u_i''(\gamma_i)}{u_i'(\gamma_i)} \in (0,1]$ for all $\gamma_i \in (\gamma_i^{min}, \gamma_i^{max}]$.

The term $-\gamma_i u_i''(\gamma_i)/u_i'(\gamma_i)$ is called the *coefficient of relative risk aversion* in economics [18] and measures the relative concaveness of $u_i(\gamma_i)$. Many common utility functions are either Type I or Type II, as shown in Table I. The logarithmic utility function is both Type I and II. A Type I utility function is "more concave" than a Type II one. Namely, an increase in one user's transmission power would induce the other users to increase their powers (i.e., $\partial^2 u_i(\gamma_i(\boldsymbol{p}))/\partial p_i \partial p_j \geq 0$ for $j \neq i$); a Type II utility would have the opposite effect (i.e., $\partial^2 u_i(\gamma_i(\boldsymbol{p}))/\partial p_i \partial p_j \leq 0$ for $j \neq i$). The strategy spaces must be redefined in different ways for these two types of utility functions to satisfy the requirements of a supermodular game.

Proposition 1: G_{FPP}^2 is supermodular in the transformed strategies $(p_1, p_2, -\pi_1, -\pi_2)$ if both users have Type I utility functions.

Proposition 2: G_{FPP}^2 is supermodular in the transformed strategies $(p_1, -p_2, \pi_1, -\pi_2)$ if both users have Type II utility functions.

The proofs of both propositions consist of checking the increasing differences conditions for each player's payoff function. These results along with Theorem 1 enable us to characterize the convergence of the ADP algorithm. For example, if the two users have Type I utility functions (and $\bar{\pi}_1, \bar{\pi}_2 < \infty$), then $\mathcal{F}^{ADP}$ is nonempty. In case of multiple fixed points, there exist two extreme ones $\left(\boldsymbol{p}^L, \boldsymbol{\pi}^L\right)$ and $\left(\boldsymbol{p}^R, \boldsymbol{\pi}^R\right)$, which are the smallest and largest fixed points in terms of strategies $(p_1, p_2, -\pi_1, -\pi_2)$. If users initialize with $(\boldsymbol{p}(0), \boldsymbol{\pi}(0)) = \left(P_1^{\min}, P_2^{\min}, \bar{\pi}_1, \bar{\pi}_2\right)$ or $(P_1^{\max}, P_2^{\max}, 0, 0)$, the power and prices converge monotonically to $\left(\boldsymbol{p}^L, \boldsymbol{\pi}^L\right)$ or $\left(\boldsymbol{p}^R, \boldsymbol{\pi}^R\right)$, respectively. If users start from arbitrary initial power and prices, then the strategies will eventually lie in the space bounded by $\left(\boldsymbol{p}^L, \boldsymbol{\pi}^L\right)$ and $\left(\boldsymbol{p}^R, \boldsymbol{\pi}^R\right)$. Similar arguments can be made with Type II utility functions with a different strategy transformation. Convergence of the powers for both types of utilities is illustrated in Fig. 3.

2) M-user Networks: Proposition 1 can be easily generalized to a network with $M > 2$:

Corollary 1: For an M-user network if all users have Type I utilities, G_{FPP} is a supermodular in the transformed strategies $(\boldsymbol{p}, -\boldsymbol{\pi})$.

In this case, Theorem 1 can again be used to characterize the structure of $\mathcal{F}^{ADP}$ as well as the convergence of the ADP algorithm. On the other hand, it can be seen that the strategy redefinition used in Proposition 2, can not be applied with $M > 2$ users so that the increasing differences property holds for every pair of users.

With logarithmic utility functions, it is shown in [5] that Problem P1 is a strictly concave maximization problem over the transformed variables $y_i = \log p_i$. In this case Problem P1 has a unique optimal solution, which is the only point satisfying the KKT conditions. It follows from Lemma 2 and Lemma 3 that G_{FPP} will have a unique NE corresponding to this optimal solution and the ADP algorithm will converge to this point from any initial choice of powers and prices.[11] With some minor additional conditions, the next proposition states that these properties generalize to other Type I utility functions. The proof is given in Appendix A.

Proposition 3: In an M-user network, if for all $i \in \mathcal{M}$:

a) $P_i^{\min} > 0$, and

b) $-\frac{\gamma_i u_i''(\gamma_i)}{u_i'(\gamma_i)} \in [a, b]$ for all $\gamma_i \in [\gamma_i^{\min}, \gamma_i^{\max}]$, where $[a, b]$ is a strict subset of $[1, 2]$;

then Problem P1 has a unique optimal solution, to which the ADP algorithm globally converges.

III. Multi-channel Networks

We now turn to a power control problem in a multi-channel network, where each user $i \in \mathcal{M}$ is able to transmit over a set of $\mathcal{K} = \{1, ..., K\}$ orthogonal channels. A superscript k denotes that a quantity refers to the kth channel, e.g. p_i^k is the ith user's power on channel k. We denote the vector of powers across users for a particular channel k by $\boldsymbol{p}^k = \left(p_i^k\right)_{i=1}^M$ and the vector of power across channels for a particular user i by $\boldsymbol{p}_i = \left(p_i^k\right)_{k=1}^K$. Finally, $\boldsymbol{p} = (\boldsymbol{p}_i)_{i=1}^M$ will denote the power profile of all users in all channels. The same notation is used for other quantities such as SINR and prices. Each user i's power allocation must lie in the set,

$$\mathcal{P}_i^{MC} = \left\{ \boldsymbol{p}_i : \sum_{k \in \mathcal{K}} p_i^k \leq P_i^{\max}, \text{ and } p_i^k \geq P_i^{\min}, \forall k \in \mathcal{K} \right\},$$

[11] Moreover, if each user $i \in \mathcal{M}$ starts from profile $(p_i(0), \pi_i(0)) = \left(P_i^{\min}, \theta_i/(n_0 B)\right)$ or $(P_i^{\max}, 0)$, then their strategies will monotonically converge to this fixed point

Where $P_i^{\max}$ is a total power constraint. User i's SINR on channel k is[12]

$$\gamma_i^k\left(p_i^k, p_{-i}^k\right) = \frac{p_i^k h_{ii}^k}{n_0^k + \sum_{j\neq i} h_{ji}^k p_j^k}.$$

In this section, we assume that each user has a "channel separable" utility, $u_i\left(\boldsymbol{\gamma}_i\left(\boldsymbol{p}\right)\right) = \sum_{k\in\mathcal{K}} u_i^k\left(\gamma_i^k\left(p_i^k, p_{-i}^k\right)\right)$, where u_i^k is an increasing and strictly concave function that represents the benefit user i receives from channel k. In other words, a user's utility is the sum of utilities from each channel. For example, this is appropriate when the utility is linear in the rate a user receives, and the total rate is the sum of the rate on each channel. Problem P1 then becomes

$$\max_{\{\mathbf{p}:\boldsymbol{p}_i\in\mathcal{P}_i^{MC},\forall i\}} \sum_{i\in\mathcal{M}}\sum_{k\in\mathcal{K}} u_i^k\left(\gamma_i^k\left(\boldsymbol{p}^k\right)\right). \tag{P2}$$

Next we discuss two generalizations of the ADP algorithm to this setting.

A. Multi-channel ADP (MADP)

The MADP algorithm is a direct generalization of the ADP algorithm in which each user i announces a vector of prices $\boldsymbol{\pi}_i$, one for each channel, and chooses a power vector $\boldsymbol{p}_i \in \mathcal{P}_i^{MC}$ to maximize the surplus function

$$s_i^{MC}\left(\boldsymbol{p}_i;\boldsymbol{p}_{-i},\boldsymbol{\pi}_{-i}\right) = u_i\left(\boldsymbol{\gamma}_i\left(\boldsymbol{p}_i;\boldsymbol{p}_{-i}\right)\right) - \sum_{k\in\mathcal{K}} p_i^k \sum_{j\neq i}\pi_j^k h_{ij}^k.$$

Specifically, for each user i, the MADP algorithm is exactly the same as the ADP algorithm except the scalars p_i and π_i are replaced by the corresponding vectors $\boldsymbol{p}_i$ and $\boldsymbol{\pi}_i$. The update functions $\mathcal{W}_i$ and $\mathcal{C}_i$ are also replaced by vector update rules $\boldsymbol{\mathcal{W}}_i(\boldsymbol{p}_{-i},\boldsymbol{\pi}_{-i})$ and $\boldsymbol{\mathcal{C}}_i(\boldsymbol{p}^k) = (\mathcal{C}_i^k(\boldsymbol{p}^k))_{k=1}^K$, where $\boldsymbol{\mathcal{W}}_i(\boldsymbol{p}_{-i},\boldsymbol{\pi}_{-i}) = \arg\max_{\hat{\boldsymbol{p}}_i\in\mathcal{P}_i^{MC}} s_i^{MC}\left(\hat{\boldsymbol{p}}_i;\boldsymbol{p}_{-i},\boldsymbol{\pi}_{-i}\right)$, and $\mathcal{C}_i^k\left(\boldsymbol{p}^k\right) = -\frac{\partial u_i^k\left(\boldsymbol{\gamma}_i^k\left(p_i^k;p_{-i}^k\right)\right)}{\partial I_i^k\left(p_{-i}^k\right)}$, with $I_i^k\left(p_{-i}^k\right) = \sum_{j\neq i} p_j^k h_{ji}^k$. Once again these updates may be asynchronous across users and among the price and power updates.

The single channel fictitious game G_{FPP} can also be generalized to the multi-channel setting so that each player's best response corresponds to the update steps in the MADP algorithm. We denote this game by $G_{MFPP} = \left[\mathcal{MFW}\cup\mathcal{MFC}, \left\{\mathcal{P}_i^{\mathcal{MFW}}, \mathcal{P}_i^{\mathcal{MFC}}\right\}, \left\{s_i^{\mathcal{MFW}}, s_i^{\mathcal{MFC}}\right\}\right]$. Again this game has two sets of players $\mathcal{MFW}$ and $\mathcal{MFC}$ both copies of $\mathcal{M}$. Each player in $\mathcal{MFW}$

[12]If there is any spreading on each channel as in multi-carrier CDMA, the factor $\frac{1}{B}$ can be absorbed into the channel gains.

chooses a power vector $\boldsymbol{p}_i$ from the strategy set $\mathcal{P}_i^{\mathcal{MFW}} = \mathcal{P}_i^{MC}$ and receives a payoff of $s_i^{\mathcal{MFW}}\left(\boldsymbol{p}_i;\boldsymbol{p}_{-i},\boldsymbol{\pi}_{-i}\right) = s_i^{MC}\left(\boldsymbol{p}_i;\boldsymbol{p}_{-i},\boldsymbol{\pi}_{-i}\right)$. Each player in $\mathcal{MFC}$ is chooses a price vector $\boldsymbol{\pi}_i$ from the strategy set $\mathcal{P}_i^{\mathcal{MFC}} = [0,\bar{\boldsymbol{\pi}}_i]$, where $\bar{\boldsymbol{\pi}}_i = \sup_{\boldsymbol{p}} \boldsymbol{\mathcal{C}}_i(\boldsymbol{p})$, and receives a payoff

$$s_i^{\mathcal{MFC}}(\boldsymbol{\pi}_i;\boldsymbol{p}) = -\sum_{k\in\mathcal{K}}\left(\pi_i^k - \mathcal{C}_i^k\left(\boldsymbol{p}^k\right)\right)^2.$$

Let $\mathcal{F}^{MADP}$ denote the set of fixed points of the MADP algorithm; i.e., the values of $(\boldsymbol{p},\boldsymbol{\pi})$ such that for all i, $\boldsymbol{\mathcal{W}}_i(\boldsymbol{p}_{-i},\boldsymbol{\pi}_{-i}) = \boldsymbol{p}_i$ and $\boldsymbol{\mathcal{C}}_i(\boldsymbol{p}^k) = \boldsymbol{\pi}_i$. By the same arguments as in the single channel case, we have:

Lemma 4: The following are equivalent: (1) A power profile $\boldsymbol{p}^*$ satisfies the KKT conditions of Problem P2; (2) $\left(\boldsymbol{p}^*,\boldsymbol{\mathcal{C}}^{MC}(\boldsymbol{p}^*)\right) \in \mathcal{F}^{MADP}$, and (3) $\left(\boldsymbol{p}^*,\boldsymbol{\mathcal{C}}^{MC}(\boldsymbol{p}^*)\right)$ is a NE of G_{MFPP}.

In a network with $K=2$ channels, certain instances of G_{MFPP} can again be transformed into equivalent supermodular games. Notice that due to the total power constraint, the strategy set $\mathcal{P}_i^{\mathcal{MFW}}$ is not a sublattice.[13] However, $\mathcal{P}_i^{\mathcal{MFW}}$ is a sublattice in transformed strategy $\left(p_i^1,-p_i^2\right)$,. Using this transformation, we can extend the results from Sect. II-B.

Corollary 2: In a network with $K=2$ channels, G_{MFPP} is supermodular in the transformed strategies $(\boldsymbol{p}^1,-\boldsymbol{p}^2,-\boldsymbol{\pi}^1,\boldsymbol{\pi}^2)$ if for all i and k, $u_i^k\left(\gamma_i^k\right)$ is Type I.

Corollary 3: In a network with $K=2$ channels and $M=2$ users, G_{MFPP} is supermodular in the strategies: $(p_1^1,-p_2^1,-p_1^2,p_2^2,\pi_1^1,-\pi_2^1,-\pi_1^2,\pi_2^2)$, if for all i and k, $u_i^k\left(\gamma_i^k\right)$ is Type II.

When G_{MFPP} is supermodular, the convergence of the MADP algorithm is again characterized by Theorem 1. Notice that Corollary 2 applies to a network with any number of users, while the strategy transformation in Corollary 3 does not generalize to $M>2$. In both cases, these transformations do not extend to $K>2$ channels.

B. Dual ADP (DADP) Algorithm

The DADP algorithm is another generalization of the ADP algorithm to multiple channels. This algorithm is based on relaxing each user i's total power constraint in Problem P2 by introducing a *power price* μ_i so that the objective function becomes $\sum_{k\in\mathcal{K}}\sum_{i\in\mathcal{M}}\left(u_i^k\left(\gamma_i^k\right) - \mu_i p_i^k\right)$. For a

[13]For example, $a = \left(P_i^{\min}, P_i^{\max} - P_i^{\min}\right) \in \mathcal{P}_i^{\mathcal{MFW}}$ and $b = \left(P_i^{\max} - P_i^{\min}, P_i^{\min}\right) \in \mathcal{P}_i^{\mathcal{MFW}}$ but $a \vee b = \left(P_i^{\max} - P_i^{\min}, P_i^{\max} - P_i^{\min}\right) \notin \mathcal{P}_i^{\mathcal{MFW}}$, assuming $P_i^{\max} > 2P_i^{\min}$, which is necessary for $\mathcal{P}_i^{\mathcal{MFW}}$ to contain for than one point.

given $\boldsymbol{\mu} = (\mu_i)_{i=1}^M$, the resulting problem is separable across channels, and so can be decomposed into K subproblems, one for each channel k, given by

$$\max_{\{\mathbf{p}^k : p_i^k \in \mathcal{P}_i, \forall i\}} \sum_{i \in \mathcal{M}} u_i^k \left(\gamma_i^k \left(\boldsymbol{p}^k\right)\right) - \mu_i p_i^k, \tag{P3}$$

where $\mathcal{P}_i = \left[P_i^{\min}, P_i^{\max}\right]$. A modified version of the (single channel) ADP algorithm can be applied to the subproblem P3 for each channel k, where the price update, $\mathcal{C}_i^{k,MC}\left(\boldsymbol{p}^k\right)$ is the same as in the MADP algorithm, and the power update is modified to be

$$\mathcal{W}_i^{k,MC}\left(p_{-i}^k, \pi_{-i}^k, \mu_i\right) = \arg\max_{p_i^k \in \mathcal{P}_i} \left(u_i^k \left(\gamma_i^k \left(p_i^k, p_{-i}^k\right)\right) - p_i^k \left(\sum_{j \neq i} \pi_j^k h_{ij}^k + \mu_i\right)\right),$$

which includes both the cost due to interference and user i's power price. For a given $\boldsymbol{\mu}$, any fixed point of this algorithm will satisfy the KKT conditions of subproblem P3.

In the DADP algorithm each user asynchronously updates its price and power for each channel using the above update rules. Additionally each user i periodically updates its own power price according to

$$\mu_i(t) = \left[\mu_i(t^-) + \kappa \left(\sum_{k \in \mathcal{K}} p_i^k(t^-)) - P_i^{max}\right)\right]^+, \tag{10}$$

where $\kappa > 0$ is a given constant and $[x]^+ = \max\{x, 0\}$. In other words, if the current power allocation is less (greater) than P_i^{max}, the user decreases (increases) its power price. The complete algorithm is given in Fig. 4, where $T_{i,p}^k$, $T_{i,\pi}^k$, and $T_{i,\mu}$ are unbounded sets of positive time instances at which each user i updates p_i^k, π_i^k, and μ_i, respectively. In this case, it can be seen that any fixed point of this algorithm will satisfy the KKT conditions of Problem P2.

We analyze the convergence of this algorithm under the following simplifying assumptions:

A1) Synchronous updates: the power prices are updated synchronously across all users.

A2) Separation of time-scales: between any two updates of the power prices, the updates in steps 3 and 4 of the algorithm converge to a fixed point.

Assumption A1 is for analytical convenience and can likely be relaxed using techniques as in [21]. Steps 3 and 4 of the algorithm are implementing the modified version of the ADP algorithm on each channel. If every utility satisfies the conditions as in Proposition 3, these updates will converge to a fixed point for any fixed $\boldsymbol{\mu}$. However, a large number of updates may be required for convergence; hence, A2 implies that there are many of these updates between any two power price updates. Numerical results in Sect. IV show that convergence can still be obtained when this assumption is dropped.

Theorem 2: In a network with M users and K channels, if for all $i \in \mathcal{M}$ and $k \in \mathcal{K}$, P_i^{min} and $u_i^k(\gamma_i^k)$ satisfy the conditions (a) and (b) in Proposition 3; then under assumptions A1 and A2, for small enough step size κ the DADP algorithm globally converges to the unique optimal solution to Problem P2.

Under these assumptions, it follows from Proposition 3 that for any $\boldsymbol{\mu}$ there is only one fixed-point, $\boldsymbol{p}^k(\boldsymbol{\mu}) = \left(p_i^k(\mu_i)\right)_{i=1}^M$, for each channel k which corresponds to the optimal solution of subproblem P3 for that channel. This fixed-point specifies the value of the following dual function for Problem P2,

$$D(\boldsymbol{\mu}) = \sum_{k\in\mathcal{K}} G^k(\boldsymbol{\mu}) + \sum_{i\in\mathcal{M}} \mu_i P_i^{\max}, \tag{11}$$

where $G^k(\boldsymbol{\mu}) = \sum_{i\in\mathcal{M}} \left(u_i^k\left(\gamma_i^k\left(\boldsymbol{p}^k(\boldsymbol{\mu})\right)\right) - \mu_i p_i^k(\mu_i)\right)$. In this setting the power price update can be viewed as a distributed gradient projection algorithm [16] for solving the dual problem:

$$\min_{\boldsymbol{\mu}\geq 0} \ D(\boldsymbol{\mu}). \tag{D}$$

The proof of this theorem, given in Appendix B, shows that (a) this algorithm converges to some $\boldsymbol{\mu}^*$ for small enough step-size κ, and (b) there is no duality gap and so $\boldsymbol{p}(\boldsymbol{\mu}^*)$ is the optimal solution to Problem P2. The proof of (b) uses a similar argument as in the proof of Proposition 3; the proof of (a) follows a similar argument as in [15], which requires showing that the gradient of the dual function is Lipschitz continuous. This is complicated here since the dual is not separable across users in each channel due to interference.

IV. Simulation Results

We provide some simulation results to illustrate the performance of the ADP and DADP algorithms. We simulate a network contained in a 10m×10m square area. Transmitters are randomly placed in this area according to a uniform distribution, and the corresponding receiver is randomly placed within $6\text{m}\times 6\text{m}$ square centered around the transmitter.

First we consider a single channel network with $M = 20$ users each with utility $u_i = \log(\gamma_i)$. The channel gains $h_{ij} = d_{ij}^{-4}$, P_i^{max}/n_0=$40dB$, and B=128. Figure 5 shows the convergence of the powers and prices for each user under the ADP algorithm for a typical realization, starting from random initializations. Also, for comparison we show the convergence of these quantities using a gradient-based algorithm as in [5] with a step-size of 0.001.[14]. Both algorithms converge to the optimal power allocation, but the ADP algorithm converges much faster; in all the cases we have simulated, the ADP algorithm converges about 10 times faster than the gradient-based algorithm (if the latter converges). The ADP algorithm, by adapting power according to the best response updates, is essentially using an "adaptive step-size" algorithm: users adapt the power in "larger" step-sizes when they are far away from the optimal solution, and use finer steps when close to the optimal.

Next we examine the convergence of DADP algorithm in a multi-channel network with $M = 50$ users and $K = 16$ channels. The other parameters are the same as in the single channel case, except here $h_{ij}^k = d_{ij}^{-4}\alpha_{ij}^k$, where α_{ij}^k is an unit mean exponential random variable that models frequency selective fading across channels. Here we simulate a version of the algorithm with step-size $\kappa = 0.05$ starting from a random initialization. All users synchronously update their power prices; the time between each update is referred to as a *dual iteration*. During each dual iteration, the users also synchronously perform both steps (3) and (4), which we refer to as a *primal update*. In Theorem 2, we assumed that there were arbitrarily many primal updates during each dual iteration. Here we investigate the case where only a small number of primal updates are used. Figure 6 shows the relative error between the current utility and the optimal value as a function of the number of dual iterations, with a maximum of 1,3,5, and 7 primal updates per iteration. Each point is averaged over 100 random topology realization. Even with only 1 primal update per iteration, the relative error quickly decreases. Figure 7 shows relative error as a function of the total number of primal updates; in this case, the number of updates per iteration appears to have little effect on the average performance.

V. CONCLUSIONS

We have presented distributed power control algorithms for both single channel and multi-channel wireless networks. In these algorithms users announce prices to reflect their sensitivities

[14] In our experiments, a larger step-size than 0.001 would often not converge

to the current interference levels and then adjust their power to maximize their surplus. In certain cases, we are able to characterize the convergence of there algorithms and show that they achieve an optimal power allocation. Some other desirable features of these algorithms are that they can be asynchronously implemented, they require only limited knowledge of channel gains by each user, and each users only announces a single price per channel. Also our numerical results show that the algorithms converge quickly, which also limits the required overhead. Here we focus on a static setting, where the communicating pairs and the channel conditions are fixed. An interesting future direction is to consider dynamic environments.

APPENDIX

A. Proof of Proposition 3:

As in [5], we use a logarithmic change of variables. Specifically, we show that in the variables $y_i = \log p_i$, Problem P1 becomes the optimization of a strictly concave objective over a compact, convex set. It follows that Problem P1 has a unique global optimum, which is the only solution to the KKT conditions. Furthermore, the solutions to the KKT conditions in the variables $\boldsymbol{y}$ have a one-to-one correspondence to solutions in the original variables $\boldsymbol{p}$. It follows that there is only one solution to the KKT conditions in the original variables, and hence by Lemma 2, $\mathcal{F}^{ADP}$ is a singleton set containing only the global optimum. Therefore, the ADP algorithm globally converges to this point.

All that remains is to show that Problem P1 has the desired properties in the variables $\boldsymbol{y}$. In the transformed variables, the constraint set becomes $\mathcal{Y} = \prod_{i \in \mathcal{M}} [\log P_i^{\min}, \log P_i^{\max}]$, which is clearly compact and convex. To show that the objective is strictly concave, we show that its Hessian is negative definite for all $\boldsymbol{y} \in \mathcal{Y}$.

Let $u_{tot}(\boldsymbol{y})$ denote the objective to Problem P1 in terms of the transformed variables. The Hessian matrix, $\boldsymbol{H}(\boldsymbol{y}) = \nabla_{\boldsymbol{yy}} u_{tot}(\boldsymbol{y})$ consists of diagonal elements:

$$H_{ii}(\boldsymbol{y}) = \gamma_i \left(u_i'' \gamma_i + u_i'\right) + \sum_{j \neq i} \gamma_j^2 \left(A_{ij}\right)^2 \left[u_j'' \gamma_j^2 + 2u_j' \gamma_j - u_j' \left(A_{ij}\right)^{-1}\right],$$

for all $i \in \mathcal{M}$, and off-diagonal elements,

$$H_{il}(\boldsymbol{y}) = -\gamma_i^2 A_{li} \left(u_i'' \gamma_i + u_i'\right) - \gamma_l^2 A_{il} \left(u_l'' \gamma_l + u_l'\right) + \sum_{j \neq i,l} \gamma_j^3 \left(A_{lj} A_{ij}\right) \left(u_j'' \gamma_j + 2u_j'\right),$$

for all $l \neq i$. Here $u_i' = \frac{\partial u_i(\gamma_i)}{\partial \gamma_i}$, $u_i'' = \frac{\partial^2 u_i(\gamma_i)}{\partial^2 \gamma_i}$, and $A_{jk} = \frac{h_{jk}\exp(y_j)}{h_{kk}\exp(y_k)}$. Since all users have Type I utilities, $u_i''\gamma_i + u_i' \leq 0$, and $u_i''\gamma_i + 2u_i' \geq 0$, for all i. It follows that $H_{il}(\boldsymbol{y}) \geq 0$, and

$$H_{ii}(\boldsymbol{y}) < \gamma_i\left(u_i''\gamma_i + u_i'\right) + \sum_{j\neq i}\gamma_j^3\left(A_{ij}\right)^2\left(u_j''\gamma_j + u_j'\right) \leq 0. \tag{12}$$

Using these relations, it can be shown that for all $i \in \mathcal{M}$ and all $\boldsymbol{y} \in \mathcal{Y}$,

$$|H_{ii}(\boldsymbol{y})| - \sum_{l\neq i}|H_{il}(\boldsymbol{y})| \geq \varepsilon_i, \tag{13}$$

where,

$$\varepsilon_i = u_i'\left(\gamma_i^{\max}\right)\frac{n_0\gamma_i^{\min}}{h_{ii}P_i^{\max}}(a-1) + \sum_{j\neq i}u_j'\left(\gamma_j^{\max}\right)\frac{h_{ij}P_i^{\min}\left(\gamma_j^{\min}\right)^3}{\left(h_{jj}P_j^{\max}\right)^2}(2-b).$$

Here a and b are the constants in the proposition. By assumption $(a-1) \geq 0$ and $(2-b) \geq 0$, and at least one of these inequalities is strict. It follows that $\varepsilon_i > 0$, i.e. $\boldsymbol{H}(\boldsymbol{y})$ is diagonal dominant. From Gersgorin's Theorem [20, page 344], the eigenvalues $\{\lambda_i\}_{i=1}^M$ of $\boldsymbol{H}(\boldsymbol{y})$ satisfy $|\lambda_j - H_{ii}| \leq \sum_{l\neq i}|H_{li}|$ for all i. Combining this with the diagonal dominance we have $\lambda_j \leq -\min_i \varepsilon_i < 0$ for all j. Since $\boldsymbol{H}(\boldsymbol{y})$ is real and symmetric and has all negative eigenvalues, it must be negative definite as desired. ■

B. Proof of Theorem 2

Consider the variable transformation $y_i^k = \log(p_i^k)$ for all i and k. By a similar argument as in the proof of Prop. 3, it follows that, under the conditions of Prop. 3, Problem P2 in the transformed variables is the optimization of a strictly concave objective over a bounded, convex set. Also, between each power price update, the DADP algorithm will converge to the unique fixed point with power allocation $\boldsymbol{p}(\boldsymbol{\mu})$ which maximizes the Lagrangian,

$$L\left(\boldsymbol{y},\boldsymbol{\mu}\right) = \sum_{k\in\mathcal{K}}\sum_{i\in\mathcal{M}}\left(u_i^k\left(\gamma_i^k\left(\boldsymbol{y}^k\right)\right) - \mu_i\exp\left(y_i^k\right)\right) + \sum_{i\in\mathcal{M}}\mu_i P_i^{\max},$$

over all $\mathbf{y}$ for which $\exp(y_i^k) \in \mathcal{P}_i$ for all i and k. This specifies the dual function $D(\boldsymbol{\mu})$ in (11). Since the primal is strictly concave in the transformed variables, there will be no duality gap between Problem P2 and the dual problem D [16, Prop. 5.3.1]. Therefore, given an optimal dual solution $\boldsymbol{\mu}^*$ to Problem D, $\boldsymbol{p}(\boldsymbol{\mu}^*)$ will be the optimal solution to Problem P2. Also, since the primal is strictly concave, $D\left(\boldsymbol{\mu}\right)$ is continuously differentiable everywhere [16, Prop. 6.1.1], and $\frac{\partial D(\boldsymbol{\mu})}{\partial \mu_i} = P_i^{\max} - \sum_{k\in\mathcal{K}} p_i^k(\mu_i)$, i.e., (10) is indeed a gradient projection update. All that remains to be shown is that (10) converges to an optimal dual value $\boldsymbol{\mu}^*$.

Let $\boldsymbol{H} = \nabla^2_{\boldsymbol{yy}} L(\boldsymbol{y},\boldsymbol{\mu})$ be the Hessian matrix of $L(\boldsymbol{y},\boldsymbol{\mu})$. Since $L(\boldsymbol{y},\boldsymbol{\mu})$ is separable across carriers, $\boldsymbol{H}$ will be a block diagonal matrix $diag\left(\boldsymbol{H}^1,\cdots,\boldsymbol{H}^K\right)$, where for each k, $\boldsymbol{H}^k = \left[\frac{\partial^2 L(\boldsymbol{y},\boldsymbol{\mu})}{\partial y_i^k \partial y_j^k}\right]$. From the same argument as in the Proof of Prop. 3, each matrix $\boldsymbol{H^k}$ will be negative definite and its eigenvalues $\{\lambda_j^k\}_{j=1}^M$ will satisfy $\max_{j\in\mathcal{M}} \lambda_j^k < -\varepsilon^k \equiv -\min_{i\in\mathcal{M}} \varepsilon_i^k$, where

$$\varepsilon_i^k = u_i^{k\prime}\left(\gamma_i^{k,\max}\right)\frac{n_0\gamma_i^{k,\min}}{h_{ii}^k P_i^{\max}}(a-1) + \sum_{j\neq i} u_j^{k\prime}\left(\gamma_j^{k,\max}\right)\frac{h_{ij}^k P_i^{\min}\left(\gamma_j^{k,\min}\right)^3}{\left(h_{jj}^k P_j^{\max}\right)^2}(2-b) > 0.$$

Therefore, $\boldsymbol{H}$ will be negative definite, and $\nabla^2 D(\boldsymbol{\mu}) = -\nabla g(\boldsymbol{y}(\boldsymbol{\mu}))' \boldsymbol{H}^{-1}\nabla g(\boldsymbol{y}(\boldsymbol{\mu}))$, where $\nabla g(\boldsymbol{y})$ is the gradient matrix of $\boldsymbol{g}(\boldsymbol{y}) = (g_i(\boldsymbol{y}_i))_{i=1}^M$, with $g_i(\boldsymbol{y}_i) = \sum_{k\in\mathcal{K}} e^{y_i^k} - P_i^{\max}$ [16, Sect. 6.1]. Note that $\nabla g(\boldsymbol{y}(\boldsymbol{\mu})) = \left[\boldsymbol{A}_{\boldsymbol{\mu}}^1 \cdots \boldsymbol{A}_{\boldsymbol{\mu}}^K\right]'$, where $\boldsymbol{A}_{\boldsymbol{\mu}}^k = diag\left(e^{y_1^k(\mu_1)},\cdots,e^{y_M^k(\mu_M)}\right)$. And so, $\nabla^2 D(\boldsymbol{\mu}) = -\sum_{k\in\mathcal{K}} \boldsymbol{A}_{\boldsymbol{\mu}}^k \left(\boldsymbol{H}^k\right)^{-1} \boldsymbol{A}_{\boldsymbol{\mu}}^k$. We use this to prove that $\nabla D(\boldsymbol{\mu})$ is Lipschitz continuous. Let $\|\boldsymbol{X}\|_2$ denote the Euclidean norm of matrix $\boldsymbol{X}$. Given any $\boldsymbol{\mu}$ and $\boldsymbol{\mu}'$, using Taylor's Theorem there exists some $\alpha \in [0,1]$ such that $\boldsymbol{\mu}'' = \alpha\boldsymbol{\mu} + (1-\alpha)\boldsymbol{\mu}'$ satisfies:

$$\|\nabla D(\boldsymbol{\mu}) - \nabla D(\boldsymbol{\mu}')\|_2 = \left\|\nabla^2 D(\boldsymbol{\mu}'')\right\|_2 \|\boldsymbol{\mu} - \boldsymbol{\mu}'\|_2, \tag{14}$$

where

$$\begin{aligned}\left\|\nabla^2 D(\boldsymbol{\mu})\right\|_2 &= \left\|-\sum_{k\in\mathcal{K}} \boldsymbol{A}_{\boldsymbol{\mu}}^k \left(\boldsymbol{H}^k\right)^{-1} \boldsymbol{A}_{\boldsymbol{\mu}}^k\right\|_2 \le \sum_{k\in\mathcal{K}} \|\boldsymbol{A}_{\boldsymbol{\mu}}^k\|_2 \left\|\left(\boldsymbol{H}^k\right)^{-1}\right\|_2 \|\boldsymbol{A}_{\boldsymbol{\mu}}^k\|_2 \\ &= \sum_{k\in\mathcal{K}} \left(\max_i p_i^k(\mu_i)\right)^2 \rho\left(\left(\boldsymbol{H}^k\right)^{-1}\right) \le \left(\max_i P_i^{\max}\right)^2 \sum_{k\in\mathcal{K}} \left(\frac{1}{\varepsilon^k}\right) \equiv J.\end{aligned} \tag{15}$$

These relations follow from because the Euclidean norm of a real, symmetric matrix is equal to its spectral radius [21, Prop. A.24], and the Euclidean norm of the inverse of a symmetric, nonsingular matrix is equal to the reciprocal of the smallest magnitude of an eigenvalue of the matrix [21, Prop. A.25]. Together (15) with (14) imply that $\nabla D(\boldsymbol{\mu})$ is Lipschitz continuous.

By a similar argument to the above, it can be shown that for small enough α, $\nabla^2 D(\boldsymbol{\mu}) - \alpha I$ is nonnegative definite for all $\boldsymbol{\mu}$. It follows that $\nabla D(\boldsymbol{\mu})$ is strongly convex [21, Prop. A.41]. Also, since Problem P2 has a finite maximum, the objective of Problem D is lower bounded. Combining these observations with the Lipschitz condition implies that there is a unique dual optimum $\boldsymbol{\mu}^*$,

and if $0 < \kappa < 2/J$ the gradient projection algorithm converges to $\boldsymbol{\mu}^*$ geometrically [21, p. 215].
$\blacksquare$

References

[1] F. P. Kelly, A. Maulloo, and D. Tan, "Rate control for communication networks: Shadow prices, proportional fairness and stability," *Journal of Operations Research Society*, vol. 49, no. 3, pp. 237–252, March 1998.

[2] Y. Qiu and P. Marbach, "Bandwidth allocation in ad hoc networks: A price-based approach," in *IEEE INFOCOM*, 2003.

[3] P. Marbach and R. Berry, "Downlink resource allocation and pricing for wireless networks," in *IEEE INFOCOM*, 2002.

[4] T. M. Heikkinen, "On congestion pricing in a wireless network," *Wireless Network*, vol. 8, pp. 347–357, 2002.

[5] M. Chiang, "Balancing transport and physical layers in wireless multihop networks: Jointly optimal congestion control and power control," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 1, pp. 104– 116, Jan 2005.

[6] D. M. Topkis, *Supermodularity and Complementarity*, Princeton University Press, 1998.

[7] E. Altman, T. Boulogne, R. El-Azouzi, T. Jimenez, and L. Wynter, "A survey on networking games in telecommunications," 2003, working paper.

[8] C. Saraydar, N. B. Mandayam, and D. J. Goodman, "Efficient power control via pricing in wireless data networks," *IEEE Trans. on Communications*, vol. 50, no. 2, pp. 291–303, Feb. 2002.

[9] E. Altman and Z. Altman, "S-modular games and power control in wireless networks," *IEEE Trans. on Automatic Control*, vol. 48, no. 5, pp. 839–842, May 2003.

[10] C. W. Sung and W. S. Wong, "Mathematical aspects of the power control problem in mobile communication systems," in *Lectures on Systems, Control and Information, Lectures at the Morningside Center of Mathematics*, AMS, 2000.

[11] T. Alpcan, T. Basar, R. Srikant, and E. Altman, "CDMA uplink power control as a noncooperative game," *Wireless Networks*, vol. 8, pp. 659–670, 2002.

[12] N. Shroff, M. Xiao, and E. Chong, "Utility based power control in cellular radio systems," in *IEEE INFOCOM*, 2001.

[13] T. Holliday, A. J. Goldsmith, , and P. Glynn, "Distributed power control for time varying wireless networks: Optimality and convergence," in *Proceedings: Allerton Conference on Communications, Control, and Computing*, 2003.

[14] F. P. Kelly, "Charging and rate control for elastic traffic," *European Transaction on Telecommun*, vol. 8, pp. 33–37, 1997.

[15] S. H. Low and D. E. Lapsley, "Optimization flow control, i: Basic algorithm and convergence," *IEEE/ACM Transactions on Networking*, vol. 7, no. 6, pp. 861–875, 1999.

[16] D. Bertsekas, *Nonlinear Programming*, 2nd ed. Belmont, Massachusetts: Athena Scientific, 1999.

[17] P. Milgrom and J. Roberts, "Rationalizability, learning and equilibrium in games with strategic complementarities," *Econometrica*, vol. 58, no. 6, pp. 1255–1277, 1990.

[18] A. Mas-Colell, M. D. Whinston, and J. R. Green, *Microeconomic Theory*, Oxford University Press, 1995.

[19] D. Xie, "Power risk aversion utility functions," *Annals of Economics and Finance*, vol. 1, pp. 265–282, 2000.

[20] R. A. Horn and C. R. Johnson, *Matrix Analysis*, Cambrige University Press, 1985.

[21] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: numerical methods*, Prentice Hall, 1989.

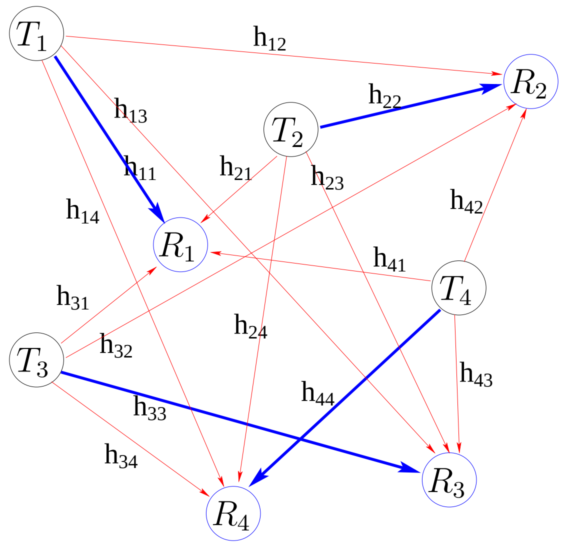

Fig. 1. An example wireless network with four users (pairs of nodes) (T_i and R_i denote the transmitter and receiver of "user" i, respectively).

(1.) INITIALIZATION: For each user $i \in \mathcal{M}$ choose some power $p_i(0) \in \mathcal{P}_i$ and price $\pi_i(0) \geq 0$.

(2.) POWER UPDATE: At each $t \in T_{i,p}$, user i updates its power according to

$$p_i(t) = \mathcal{W}_i \left(p_{-i}(t^-), \pi_{-i}(t^-) \right).$$

(3.) PRICE UPDATE: At each $t \in T_{i,\pi}$, user i updates its price according to

$$\pi_i(t) = \mathcal{C}_i \left(\boldsymbol{p}(t^-) \right).$$

Fig. 2. The ADP Algorithm.

TABLE I

EXAMPLES OF TYPE I AND II UTILITY FUNCTIONS.

Type I	Type II
$\theta_i \log(\gamma_i)$	$\theta_i \log(\gamma_i)$
$\theta_i \gamma_i^{\alpha}/\alpha$ (with $\alpha \in [-1, 0)$)	$\theta_i \gamma_i^{\alpha}/\alpha$ (with $\alpha \in (0, 1)$)
$1 - e^{-\theta_i \gamma_i}$ (with $\frac{1}{\gamma_i^{\min}} \leq \theta_i \leq \frac{2}{\gamma_i^{\max}}$)	$1 - e^{-\theta_i \gamma_i}$ (with $\theta_i \leq \frac{1}{\gamma_i^{\max}}$)
$a\left(\gamma_i\right)^2 + b\gamma_i$ (with $0 \leq -3a\gamma_i^{\max} \leq b \leq -4a\gamma_i^{\min}$)	$a\left(\gamma_i\right)^2 + b\gamma_i$ (with $b \geq -4a\gamma_i^{\max} > 0$)
$\frac{1}{\alpha}\left[1 - \exp\left(-\alpha\left(\frac{(\gamma_i)^{1-\sigma}-1}{1-\sigma}\right)\right)\right]$ (with $\sigma \in (0, 1]$ and $0 \leq (1-\sigma)\left(\gamma_i^{\min}\right)^{\sigma-1} \leq \alpha \leq (2-\sigma)\left(\gamma_i^{\max}\right)^{\sigma-1}$), or $\sigma \in (1, 2]$ and $0 \leq \alpha \leq (2-\sigma)\left(\gamma_i^{\min}\right)^{\sigma-1}$).[15]	$\frac{1}{\alpha}\left[1 - \exp\left(-\alpha\left(\frac{(\gamma_i)^{1-\sigma}-1}{1-\sigma}\right)\right)\right]$ (with $\sigma \in (0, 1]$ and $0 \leq \alpha \leq (1-\sigma)\left(\gamma_i^{\max}\right)^{\sigma-1}$)
	$\theta_i \log\left(1 + \gamma_i\right)$

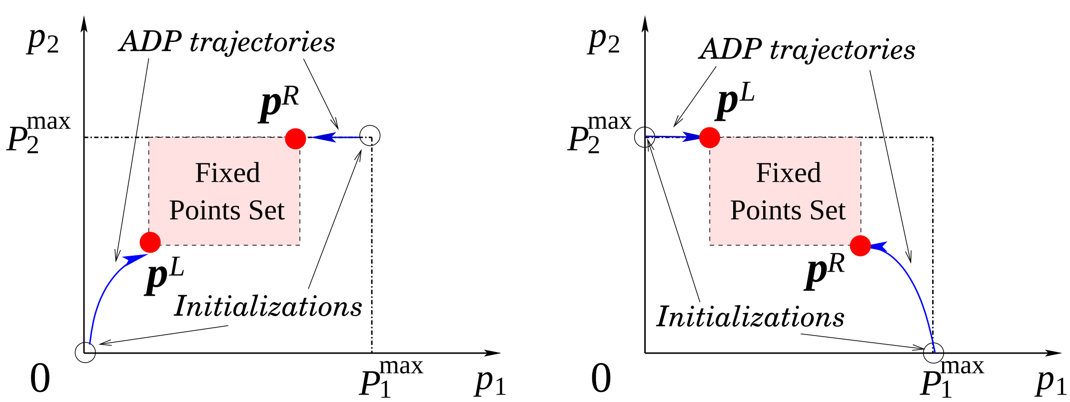

Fig. 3. Examples of the trajectories of the power profiles under the ADP algorithm for a two-user network with Type I (left) or Type II (right) utility functions. In both cases, from the indicated initializations the power profiles will monotonically converge to the indicated "corner" fixed points.

(1.) INITIALIZATION: For each user $i \in \mathcal{M}$ choose some power $\boldsymbol{p}_i(0) \in \mathcal{P}_i^{MC}$, interference price $\boldsymbol{\pi}_i(0) \geq 0$ and power price $\mu_i(0)$.

(2.) POWER PRICE UPDATE: At each $t \in T_{i,\mu}$, user i updates its power price according to

$$\mu_i(t) = \left[\mu_i\left(t^-\right) + \kappa\left(\sum_{k\in\mathcal{K}} p_i^k(t^-) - P_i^{\max}\right)\right]^+.$$

(3.) POWER UPDATE: At each $t \in T_{i,p}^k$, user i updates its power on carrier k according to

$$p_i^k(t) = \mathcal{W}_i^{k,MC}\left(p_{-i}^k(t^-), \pi_{-i}^k(t^-), \mu_i(t)\right).$$

(4.) INTERFERENCE PRICE UPDATE: At each $t \in T_{i,\pi}^k$, user i updates its interference price on carrier k according to

$$\pi_i^k(t) = \mathcal{C}_i^{k,MC}\left(\boldsymbol{p}^k(t^-)\right).$$

Fig. 4. The DADP Algorithm.

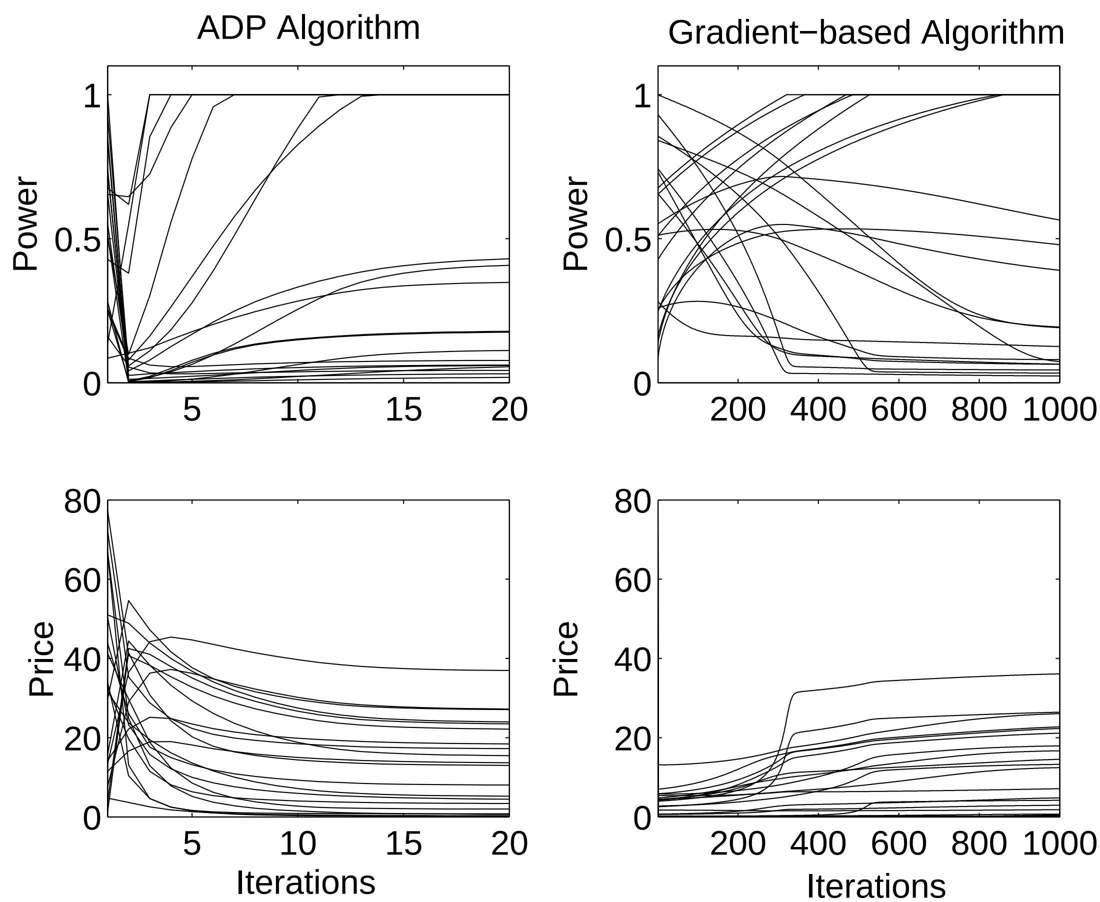

Fig. 5. Convergence of the prices and power for the ADP algorithm (left) and a gradient-based algorithm (right) in a network with 40 users and logarithmic utility functions. Each curve corresponds to the power or price for one user.

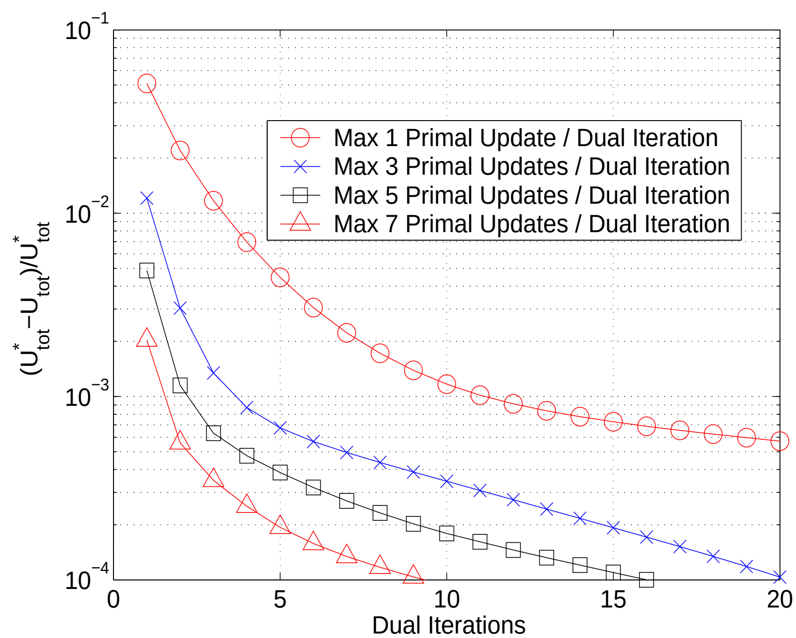

Fig. 6. Average relative difference between the total utility (U_{tot}) and the optimal total utility (U^*_{tot}) as a function of the number of dual iterations in the DADP algorithm.

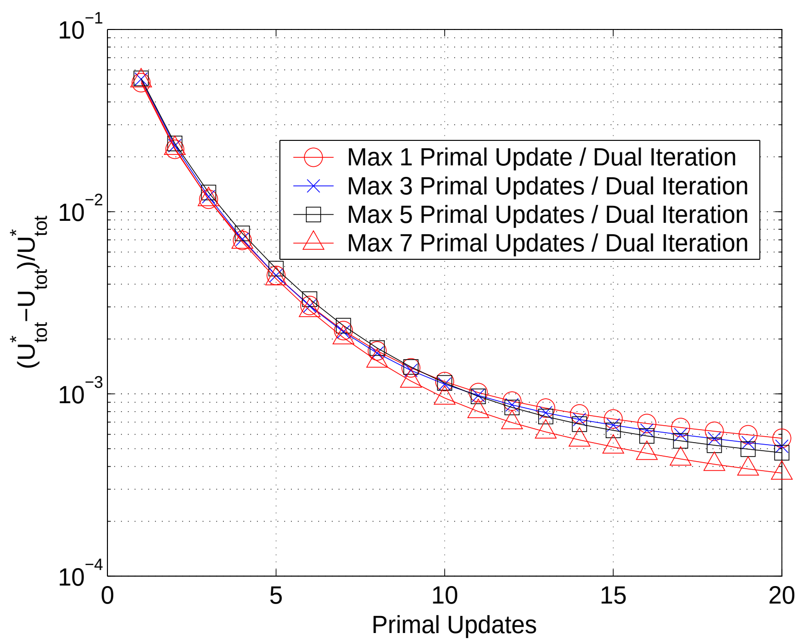

Fig. 7. Average relative difference between the total utility (U_{tot}) and the optimal total utility (U^*_{tot}) as a function of the total number of primal updates in the DADP algorithm.