Distributed Interference Compensation for Multi-channel Wireless Networks*

Jianwei Huang, Randall Berry, Michael L. Honig

Dept. of Electrical and Computer Engineering
Northwestern University
2145 Sheridan Road, Evanston, IL, 60208
{jianweih,rberry,mh}@ece.northwestern.edu

Abstract

We present distributed power control algorithms for a wireless peer-to-peer network with multiple channels per user. Users exchange "price" signals that indicate the negative effect of interference at the receivers in each channel. Given this set of prices, each transmitter chooses a power allocation across channels to maximize its net benefit (utility minus cost), subject to a total power constraint. We consider two specific algorithms for power and price updates, and establish global convergence for both algorithms to the unique globally optimal power allocation for a class of concave user utility functions. When the utility functions represent achievable rates, global convergence is not guaranteed; however, we show numerically that the proposed power control algorithms achieve much better performance than iterative water-filling, in which users maximize their own rates without exchanging price information.

1 Introduction

Interference in a wideband wireless network can be mitigated through the use of power control strategies that distribute power unevenly across available channels according to measured activity. Example scenarios include interference avoidance in wireless ad hoc networks with frequency-selective channels [1], power allocation across multiple cells in an Orthogonal Frequency Division Multiplexing (OFDM) network [2], and spectrum management in digital subscriber lines with crosstalk [3]. In each of those scenarios, power can be allocated across frequencies to minimize interference and optimize overall network performance.

Here we consider power control in a wideband ad hoc (peer-to-peer) network. Ideally, what is desired is a distributed algorithm, which does not rely on a centralized infrastructure, and has complexity that scales linearly with the network size. In this paper, we present distributed power control algorithms, motivated by the desire to maximize total utility summed over all users. Namely, each user is assigned a utility function, which is increasing and strictly concave in the received Signal-to-Interference plus Noise Ratio (SINR). Allocating power across users and channels to maximize total network utility is

*This work was supported by the Northwestern-Motorola Center for Communications, ARO under grant DAAD19-99-1-0288, and NSF CAREER award CCR-0238382.

typically difficult due to the effect of interference across channels, which can make the optimization objective non-concave. In addition, the total power constraints introduce dependencies across the power that can be allocated to different channels for each user.

To mitigate interference, we present power control algorithms in which the users exchange "price" signals that represent the marginal loss in utility due to a marginal increase in interference in a particular channel. Each transmitter then determines the power allocation across all channels autonomously by maximizing the user's surplus (utility minus cost), where the interference prices are used to compute the cost. The interference prices internalize the associated negative externalities among users, and therefore can be interpreted as a type of *Pigovian Tax* [4]. This is different from many previously proposed pricing mechanisms for resource allocation, both in wire-line networks (e.g. [5]) and wireless networks (e.g. [6,7]), where prices are Lagrange multipliers for a constrained resource, or heuristic signals to coordinate the behavior of different users.

In related work, we have characterized the convergence and performance of a distributed pricing mechanism for power control in a single-channel ad hoc network [8], and in a multi-channel ad hoc network where each user can choose only *one* channel on which to transmit [9]. In this paper, we consider the case where users can allocate power across *all* channels. We present two algorithms, a *primal* algorithm and a *dual* algorithm. In the *primal* algorithm, each transmitter determines the power allocation across all channels to maximize its surplus, subject to a total power constraint, taking into account the interference prices for each channel. The *dual* algorithm is based on the technique of Lagrangian relaxation, and allows us to decompose the network optimization problem into several subproblems, one for each channel. The dual variables are then updated to enforce the total power constraints.

We show that both algorithms converge to the globally optimal power allocation for a class of utility functions that are "sufficiently" concave. This condition is not satisfied when the users' utility functions correspond to achievable rates, so that convergence to the global optimum is not guaranteed in general. In that case, we show that the algorithms converge for a two-user two-channel network, and for a network with an arbitrary number of users, but with a constraint (upper bound) on cross-channel gains. We also present numerical results, which show that the pricing algorithms perform better than iterative water-filling [10] (in both low and medium SINR regimes), where users maximize their individual rates autonomously without exchanging information.

2 System Model

We consider a stationary wireless network with a set of $\mathcal{M} = \{1, ..., M\}$ distinct transceiver pairs. Each pair consists of one transmitter node and one receiver node. Although we focus on a peer-to-peer scenario, the model considered may also apply to certain multi-hop scenarios in which a particular schedule of transmissions has been determined by an underlying routing and MAC layer protocol. We will use the terms "pair" and "user" interchangeably. Each user $i \in \mathcal{M}$ is able to transmit over a set of $\mathcal{K} = \{1, ..., K\}$ nonoverlapping channels. Over the time-period of interest, we assume that the channel gains are fixed and that the users want to transmit continually. For channel $k \in \mathcal{K}$, the gain between user i's transmitter and user j's receiver is denoted by h_{ij}^k.[1] An example of

[1]Note that in general $h_{ij}^k \neq h_{ji}^k$, since the latter represents the gain in channel k between user j's transmitter and user i's receiver.

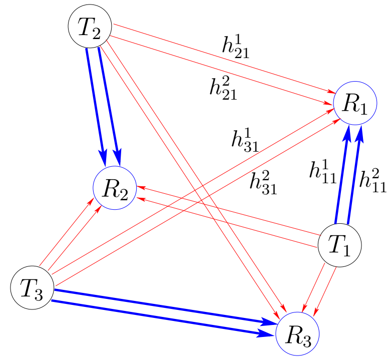

Figure 1: A multichannel network with three users (pairs of nodes) and $K = 2$ channels. T_i and R_i denote the transmitter and receiver for user i, respectively.

a network with three pairs of nodes and two channels is shown in Fig. 1.

User i is assigned the utility function

$$u_i\left(\boldsymbol{\gamma}_i\left(\boldsymbol{p}\right)\right) = \sum_{k\in\mathcal{K}} u_i^k\left(\gamma_i^k\left(p_i^k, p_{-i}^k\right)\right),$$

where u_i^k is an increasing and strictly concave function of user i's SINR on channel k,

$$\gamma_i^k\left(p_i^k, p_{-i}^k\right) = \frac{p_i^k h_{ii}^k}{n_0^k + \sum_{j\neq i} h_{ji}^k p_j^k}. \tag{1}$$

Here, n_0^k is the background noise power for channel k, p_i^k is user i's transmission power on channel k, and $p_{-i}^k = \left(p_j^k, j \in \mathcal{M}, j \neq i\right)$ is the vector of powers across all users except user i on channel k. We denote the vector of powers across users for channel k by $\boldsymbol{p}^k = \left(p_i^k, i \in \mathcal{M}\right) = \left(p_i^k, p_{-i}^k\right)$, and the vector of powers across channels for a particular user i by $\boldsymbol{p}_i = \left(p_i^k, k \in \mathcal{K}\right)$. Finally, $\boldsymbol{p} = \left(\boldsymbol{p}_i, i \in \mathcal{M}\right) = \left(\boldsymbol{p}^k, k \in \mathcal{K}\right)$ denotes the power profile across all users and channels. Our objective is to find a $\boldsymbol{p}$ that solves:

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

where user i's power is constrained to lie in the set

$$\mathcal{P}_i = \left\{\boldsymbol{p}_i \middle| \sum_{k\in\mathcal{K}} p_i^k \leq P_i^{\max},\ p_i^k \geq P_i^{\min} \geq 0, \forall k \in \mathcal{K}\right\}.$$

Let $\boldsymbol{\mathcal{P}} = \Pi_{i\in\mathcal{M}}\mathcal{P}_i$ denote the set of feasible power profiles, $\boldsymbol{p}$.

In what follows, we will take as particular examples the *rate* utility, $u_i^k\left(\gamma_i^k\left(\boldsymbol{p}^k\right)\right) = \theta_i \log\left(1 + \gamma_i^k\left(\boldsymbol{p}^k\right)\right)$, and the *log* utility $u_i^k\left(\gamma_i^k\left(\boldsymbol{p}^k\right)\right) = \theta_i \log\left(\gamma_i^k\left(\boldsymbol{p}^k\right)\right)$, where the θ_i's can represent user-dependent priorities. Of course, the log-utility approximates the rate utility in the high SINR regime. Although the rate utility is strictly concave in γ_i^k, the objective in Problem P1 may not be concave in $\boldsymbol{p}$. However, it is easy to verify that any locally optimal $\boldsymbol{p}^*$ must satisfy the following Kuhn-Tucker conditions (Prop. 3.3.1 of [11]):

Lemma 1 *For any locally optimal solution $\boldsymbol{p}^*$ of Problem P1, there exist unique Lagrange multiplier vectors $\boldsymbol{\mu}^* = (\mu_i^*, i \in \mathcal{M})$ and $\boldsymbol{\lambda}^* = \left(\lambda_i^{k*}, i \in \mathcal{M}, k \in \mathcal{K}\right)$ such that for all $i \in \mathcal{M}$ and $k \in \mathcal{K}$,*

$$\left.\left(\frac{\partial u_i^k\left(\boldsymbol{\gamma}_i^k\left(p_i^k, \boldsymbol{p}_{-i}^k\right)\right)}{\partial p_i^k} + \sum_{j \neq i} \frac{\partial u_j^k\left(\boldsymbol{\gamma}_j^k\left(p_j^k, \boldsymbol{p}_{-j}^k\right)\right)}{\partial p_i^k}\right)\right|_{\boldsymbol{p}^k = \boldsymbol{p}^{k*}} = \mu_i^* - \lambda_i^{k*}, \tag{2}$$

where $\lambda_i^{k*}, \mu_i^* \geq 0$, $\mu_i^*(\sum_{k \in \mathcal{K}} p_i^{k*} - P_i^{max}) = 0$, and $\lambda_i^{k*}\left(P_i^{\min} - p_i^{k*}\right) = 0$.

Let $\pi_j^k\left(\boldsymbol{p}^k\right) = -\frac{\partial u_j\left(\boldsymbol{\gamma}_j\left(p_j^k, \boldsymbol{p}_{-j}^k\right)\right)}{\partial I_j^k\left(\boldsymbol{p}_{-j}^k\right)}$, where $I_j^k\left(\boldsymbol{p}_{-j}^k\right) = \sum_{l \neq j} h_{lj}^k p_l^k$ is the total interference received by user j in channel k. Here, $\pi_j^k\left(\boldsymbol{p}^k\right)$ is always nonnegative and represents user j's sensitivity to its current interference level, i.e., how much its utility would increase if the interference is decreased by one unit. Equation (2) can then be written as

$$\left.\left(\frac{\partial u_i^k\left(\boldsymbol{\gamma}_i^k\left(p_i^k, \boldsymbol{p}_{-i}^k\right)\right)}{\partial p_i^k} - \sum_{j \neq i} \pi_j^k\left(\boldsymbol{p}^k\right) h_{ij}^k\right)\right|_{\boldsymbol{p}^k = \boldsymbol{p}^{k*}} = \mu_i^* - \lambda_i^{k*}, \forall i \in \mathcal{M}, \forall k \in \mathcal{K}. \tag{3}$$

Each $\pi_j^k\left(= \pi_j^k\left(\boldsymbol{p}^k\right)\right)$ can be viewed as a "price "that each user $i \neq j$ must pay for each unit of interference generated at user j on channel k. Condition (3) can then be interpreted as a necessary and sufficient optimality condition for determining the power vector $\boldsymbol{p}_i \in \mathcal{P}_i$, which maximizes user i's surplus function

$$s_i\left(\boldsymbol{p}_i, \boldsymbol{p}_{-i}; \boldsymbol{\pi}_{-i}\right) = \sum_{k \in \mathcal{K}} \left(u_i^k\left(\gamma_i^k\left(p_i^k, \boldsymbol{p}_{-i}^k\right)\right) - p_i^k \sum_{j \neq i} \pi_j^k h_{ij}^k\right), \tag{4}$$

given $\boldsymbol{p}_{-i} = \left(\boldsymbol{p}_{-i}^k, k \in \mathcal{K}\right)$ and prices $\boldsymbol{\pi}_{-i} = \left(\boldsymbol{\pi}_{-i}^k, k \in \mathcal{K}\right)$. The surplus is user i's utility minus its payment for the interference it generates. For each channel, the payment is user i's transmit power times a weighted sum of other users' prices, where the weights are the channels gains between the ith transmitter and the other users' receivers. The payment can viewed as compensation to other users for the interference associated with the transmission.

3 Asynchronous Distributed Pricing Algorithms

The pricing interpretation of the KKT conditions motivates the following multi-channel asynchronous distributed pricing (ADP) algorithms, in which users iteratively announce prices and update their transmit power allocations to achieve a solution that satisfies Lemma 1. We present two algorithms, a *primal* algorithm and a *dual* algorithm, and show the convergence of both under various utility and network assumptions. Most proofs can be found in [12] and are omitted here.

3.1 Primal ADP (PADP) Algorithm

In the PADP algorithm, each user i updates its power allocation to maximize its surplus as in (4); formally, we represent this via the update function,

$$\mathcal{W}_i\left(\boldsymbol{p}_{-i}, \boldsymbol{\pi}_{-i}\right) = \arg\max_{\boldsymbol{p}_i \in \mathcal{P}_i} s_i\left(\boldsymbol{p}_i, \boldsymbol{p}_{-i}, \boldsymbol{\pi}_{-i}\right).$$

Each user also updates each price π_i^k according to

$$\mathcal{C}_i^k\left(\boldsymbol{p}^k\right) = -\frac{\partial u_i^k\left(\boldsymbol{\gamma}_i^k\left(p_i^k;p_{-i}^k\right)\right)}{\partial I_i^k\left(p_{-i}^k\right)} = \frac{\partial u_i^k\left(\boldsymbol{\gamma}_i^k\left(p_i^k;p_{-i}^k\right)\right)}{\partial \boldsymbol{\gamma}_i^k\left(p_i^k;p_{-i}^k\right)} \frac{\left(\boldsymbol{\gamma}_i^k\left(p_i^k;p_{-i}^k\right)\right)^2}{p_i^k h_{ii}^k}.$$

The users iteratively update their prices and power allocations according to these function. We allow these updates to be done asynchronously, both across users as well as between price and power updates for a given user. To be precise, for each user i, let $T_{i,p}$ and $T_{i,\pi}^k$, $k \in \mathcal{K}$ be $K+1$ sets of infinite positive time instances at which user i updates its power allocation and price for channel k, respectively.[2] The complete PADP algorithm is then specified in Algorithm 1 (t^- denotes the time immediately before t).

Algorithm 1 PADP Algorithm

---

1. *Initialization*: at $t = 0$, each user $i \in \mathcal{M}$ chooses some initial power $\boldsymbol{p}_i(0) \in \mathcal{P}_i$ and price $\boldsymbol{\pi}_i(0) \geq 0$.
2. *Power Update:* At each $t \in T_{i,p}$, user i updates its power allocation according to $\boldsymbol{p}_i(t) = \mathcal{W}_i\left(\boldsymbol{p}_{-i}(t^-), \boldsymbol{\pi}_{-i}(t^-)\right)$.
3. *Price Update:* At each $t \in T_{i,\pi}^k$, user i updates its price on channel k according to $\pi_i^k(t) = \mathcal{C}_i^k\left(\boldsymbol{p}^k(t^-)\right)$.

It can be seen that in this algorithm each user's updates only require that user have the following limited information for each channel k: its own SINR γ_i^k and channel gain h_{ii}^k, channel gains h_{ij}^k for all $j \neq i$, and other users' prices π_{-i}^k. The γ_i^k and h_{ii}^k can be measured at its receiver and feed back to the transmitter. The adjacent channel gains h_{ij}^k account for only $1/M$ of the total channel gains, and can be measured by having each receiver periodically broadcast a beacon; assuming reciprocity, the transmitters can then measure these channel gains. The price information could also be periodically broadcast through these beacons.

When this algorithm converges, it can be shown that:

Lemma 2 *A power profile $\boldsymbol{p}^*$ satisfies the KKT conditions of Problem P1 if and only if $(\boldsymbol{p}^*, \boldsymbol{\mathcal{C}}(\boldsymbol{p}^*))$ is a fixed point of the PADP algorithm.*

Next, we study the convergence of this algorithm for different classes of utility functions. The first class we consider is defined in terms of their *coefficient of relative risk aversion*, $G_i^k\left(\gamma_i^k\right) = -\gamma_i^k u_i^{k\prime\prime}(\gamma_i^k)/u_i^{k\prime}(\gamma_i^k))$. This quantity is used in economics [4] and measures the relative concaveness of $u_i^k\left(\gamma_i^k\right)$ (a larger value indicates a "more concave" function). Let $\gamma_i^{\min} = \min\{\gamma_i^k(\boldsymbol{p}^k) : \boldsymbol{p} \in \boldsymbol{\mathcal{P}}, \forall k \in \mathcal{K}\}$ and $\gamma_k^{\max} = \max\{\gamma_i^k(\boldsymbol{p}^k) : \boldsymbol{p} \in \boldsymbol{\mathcal{P}}, \forall k \in \mathcal{K}\}$ for all user $i \in \mathcal{M}$. Our first class of utility functions is defined to be those for which $G_i^k\left(\gamma_i^k\right) \in [1,2]$ (e.g., $u_i^k\left(\gamma_i^k\right) = \theta_i \log\left(\gamma_i^k\right)$), in which case the objective function of Problem P1 can be shown to be strictly concave in the logarithmic transformed variables $\boldsymbol{y} = \mathbf{log}(\boldsymbol{p})$ under mild conditions. Thus, the KKT conditions become necessary and sufficient for the unique optimal solution to Problem P1; in this case the the PADP algorithm has a unique fixed-point that is optimal. With the restriction to only $K = 2$ channels, the next theorem states that the PADP algorithm will globally converge to this fixed point; we will discuss more $K > 2$ channels in Section 3.2.

[2]We do not require that updates be asynchronous; i.e. synchronous updates can simply be viewed as a special case.

Theorem 3 *Consider a network with only two channels. If for any $i \in \mathcal{M}$, $P_i^{\min} > 0$, and $G_i^k\left(\gamma_i^k\right) \in [a,b]$ for all $\gamma_i^k \in \left[\gamma_i^{\min}, \gamma_i^{\max}\right]$, where $[a,b]$ is a strict subset of $[1,2]$, then Problem P1 has a unique optimal solution to which the PADP algorithm globally converges.*

On the other hand, if $G_i^k\left(\gamma_i^k\right) \in (0,1)$ (e.g., rate utility $u_i^k\left(\gamma_i^k\right) = \theta_i \log\left(1+\gamma_i^k\right)$), then the objective of Problem P1 might not be concave in $\boldsymbol{p}$ (or $\boldsymbol{y} = \mathbf{log}\left(\boldsymbol{p}\right)$), and there may exist multiple local optimal solutions. As an example, we will next consider the performance of the PADP algorithm for the class of rate utilities. We first show that the PADP can still converge in this case with proper initialization, even if Problem P1 has multiple local optimal solutions. Let the maximum possible price of user i in channel k be $\bar{\pi}_i^k = \arg\max_{\{\boldsymbol{p}:\boldsymbol{p}_i \in \mathcal{P}_i,\, \forall i \in \mathcal{M}\}} \pi_i^k\left(\boldsymbol{p}^k\right)$, which is finite for rate utilities.

Theorem 4 *Consider a network with only two users and two channels with rate utility functions. If users initialize with $\left(p_1^1(0), p_1^2(0), p_2^1(0), p_2^2(0), \pi_1^1(0), \pi_1^2(0), \pi_2^1(0), \pi_2^2(0)\right)$ equal to*

$$\begin{aligned}&\left(P_1^{\max} - P_1^{\min}, P_1^{\min}, P_2^{\min}, P_2^{\max} - P_2^{\min}, \bar{\pi}_1^1, 0, 0, \bar{\pi}_2^2\right), \text{ or}\\ &\left(P_1^{\min}, P_1^{\max} - P_1^{\min}, P_2^{\max} - P_2^{\min}, P_2^{\min}, 0, \bar{\pi}_1^2, \bar{\pi}_2^1, 0\right),\end{aligned}$$

then the PADP algorithm converges.

Theorem 4 can be generalized to any utility function with $G_i^k\left(\gamma_i^k\right) \in (0,1]$ (for details, see [12]). To prove Theorems 3 and 4, we can map the PADP algorithm to the myopic best response updates of a "fictitious power-price control game," and show the convergence using supermodular game theory. Both results are restricted to the case of two channels in order for the resulting game to satisfy the requirements of a supermodular game. However, the simulation results in Section 4 show that the PADP algorithm converges in more general cases.

On the other hand, when the interferences are small enough, we can show the the convergence of the PADP algorithm for arbitrary number of channels using a contraction mapping argument. To simplify the discussion, here we only consider a particular synchronous update scheme, where $T_{i,p} = T_{j,p}$ and $T_{i,\pi}^k = T_{j,\pi}^{k'}$ for any $i \neq j$ and $k \neq k'$, i.e., the power updates and price updates are each done synchronously. For $j \neq i$, let $\alpha_{ji}^k = h_{ji}^k / h_{ii}^k$ be the normalized interference coefficient for user i from user j.

Theorem 5 *In a two-user K-channel network with symmetric ($\theta_1 = \theta_2$) rate utilities, there exists some constant $\xi > 0$ such that the PADP algorithm with synchronous updates globally and geometrically converges to the unique optimal solution of Problem P1, whenever*

$$\max_{i \in \{1,2\}, j \neq i, k \in \mathcal{K}} \alpha_{ji}^k \leq \xi.$$

The value ξ can be explicitly calculated and depends on the number of channels, K, the normalized noise, $n_i^k = n_0^k / h_{ii}^k$, and the power constraints of both users. This small interference condition can be satisfied when the receiving nodes are far enough away from any interfering transmission. We believe that the proof technique can be generalized to the case of more than two users as well as asymmetric utility functions.

3.2 Dual ADP (DADP) Algorithm

The DADP algorithm is based on the idea of relaxing each user i's total power constraint in Problem P1 by introducing a *dual price* μ_i, so that the objective function of Problem P1 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 given $\boldsymbol{\mu}=(\mu_i, i\in\mathcal{M})$, the resulting problem is separable across channels, and so can be decomposed into K subproblems, one for each channel k, given by

$$\max_{\{\boldsymbol{p}^k : p_i^k\in\mathcal{P}_i',\forall i\}} \sum_{i\in\mathcal{M}} u_i^k\left(\gamma_i^k\left(\boldsymbol{p}^k\right)\right)-\mu_i p_i^k, \tag{P2}$$

where $\mathcal{P}_i'=\left[P_i^{\min},P_i^{\max}\right]$. A single-channel version of the PADP algorithm can then be applied to the subproblem P2 for each channel k, where the price update, $\mathcal{C}_i^k\left(\boldsymbol{p}^k\right)$ is the same as in the PADP algorithm, and the power update is modified to be

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

which includes both the cost due to interference and user i's dual price[3]. For a given $\boldsymbol{\mu}$, any fixed point of this algorithm will satisfy the KKT conditions of subproblem P2. The dual prices $\boldsymbol{\mu}$ are then periodically updated to enforce the total power constraints. The complete DADP algorithm is given in Algorithm 2, 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, and $\kappa>0$ is a small constant.

Algorithm 2 DADP algorithm

---

1. *Initialization*: at $t=0$, each user $i\in\mathcal{M}$ chooses some initial power $\boldsymbol{p}_i(0)\in\mathcal{P}_i$, interference price $\boldsymbol{\pi}_i(0)\geq 0$ and dual price $\mu_i(0)\geq 0$.

2. *Dual Price Update*: at each $t\in T_{i,\mu}$, user i updates its dual price according to

$$\mu_i(t)=\max\left\{\mu_i\left(t^-\right)+\kappa\left(\sum_{k\in\mathcal{K}}p_i^k(t^-)-P_i^{\max}\right),0\right\}.$$

3. *Power Update*: at each $t\in T_{i,p}^k$, user i updates its power on carrier k according to

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

4. *Interference Price Update*: at each $t\in T_{i,\pi}^k$, user i updates its interference price on carrier k according to

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

It can be seen that any fixed point of the DADP algorithm will satisfy the KKT conditions of Problem P1. We analyze the convergence of this algorithm under two simplifying assumptions:

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

[3] To differentiate, we call π_i^k (for all i and k) the interference prices.

A2) Separation of time-scales: between any two updates of the dual 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 [13]. Steps 3 and 4 of the algorithm are implementing the single channel version of the PADP algorithm on each channel. If each utility function satisfies the conditions as in Theorem 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 dual price updates. Numerical results in Sect. 4 show that convergence can still be obtained when this assumption is dropped.

Theorem 6 *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 in Theorem 3; then under assumptions A1 and A2, for small enough step size κ the DADP algorithm globally and geometrically converges to the unique optimal solution to Problem P1.*

The proof of this is based on showing that under these assumptions the dual price update can be viewed as a distributed gradient projection algorithm [11] for solving a dual problem of Problem P1 and then proving the convergence of this algorithm.

4 Numerical Results

We illustrate the performance of the PADP and DADP algorithms through some numerical results. In all experiments, we let $P_i^{\max} = 1$, $P_i^{\min} = 0$, and $n_0^k = 10^{-2}$. The channel gains are modeled as $h_{ij}^k = d_{ij}^{-4}\alpha_{ij}^k$, where d_{ij} is the distance between transmitter i and receiver j, and the α_{ij}^k's are independent, unit-mean exponential random variables that model frequency-selective fading across channels.

Figure 2 illustrates the convergence of the PADP and DAPD algorithms. The transmitters are uniformly placed in a 10m×10m area, and the receivers are randomly placed within a 6m×6m square centered around the corresponding transmitters. Users are initialized with random powers and prices. The *left* subfigure shows utility versus iterations for the PADP algorithm in a network with 10 users and 20 channels. Each user has a rate utility function $u_i(\boldsymbol{\gamma}_i) = \sum_{k\in\mathcal{K}} \log\left(1+\gamma_i^k\right)$, and the curves correspond to different users. Although convergence is not guaranteed, according to Theorems 4 and 5, these results show that for the scenario simulated the algorithm converges within a few iterations. (An iteration represents one round of synchronous update of powers and prices.)

The *right* subfigure of Figure 2 shows normalized utility versus number primal updates for the DADP algorithm in a network with 50 users and 16 channels. For this set of plots the users have logarithmic utility functions $u_i(\boldsymbol{\gamma}_i) = \sum_{k\in\mathcal{K}} \log\left(\gamma_i^k\right)$, the dual update stepsize is $\kappa = 0.05$, and each point is averaged over 100 random topology realizations. All users synchronously update their dual prices, which marks a *dual iteration.* During each dual iteration, the users synchronously perform both steps (3) and (4) in Algorithm 2, which we refer to as a *primal update.* The figure shows plots corresponding to different numbers of primal updates within each dual iteration. The results show that the performance is insensitive to this parameter. Hence the separation of time scales, assumed in Theorem 6, which implies a large number of primal updates per dual iteration, is not necessary for convergence in practice.

Figure 3 compares the performance of the PADP algorithm with the iterative water-filling (IWF) algorithm [10], in which each transmitter allocates power to maximize

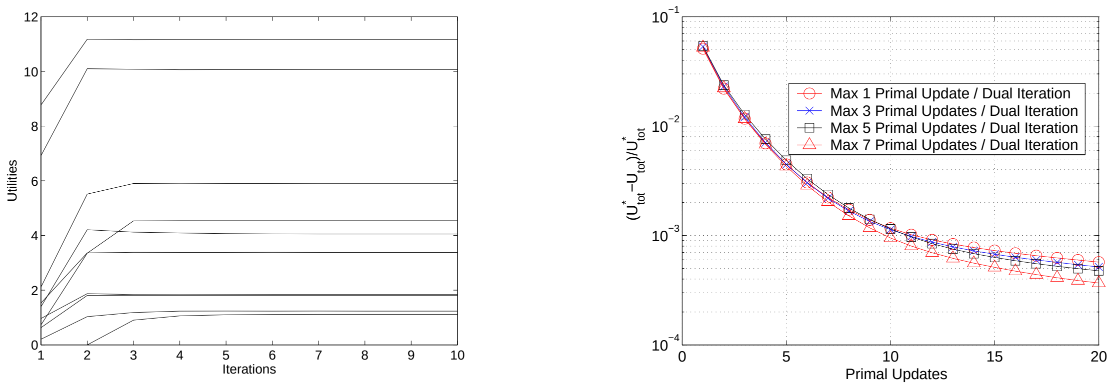

Figure 2: Algorithm convergence: *Left*: PADP, *right*: DAPP

the rate without exchanging information. In this example, each user has a rate utility function, and the placement of transmitters and receivers is the same as in Figure 2. For the *left* subfigure, there are two users and 20 channels. The light (yellow) bar denotes the allocated power, and the dark (blue) bar denotes the normalized noise plus interference $\left(n_0^k + h_{ji}^k p_j^k\right)/h_{ii}^k$, truncated to 0.2. For this example, the mean value of the interference channel gains $\left(h_{ji}^k\right)$ is half the mean value of the direct channel gains $\left(h_{ii}^k\right)$. In this case, the PADP algorithm assigns non-overlapping channels, whereas IWF assigns overlapping channels, which leads to significant interference.

The *right* subfigure shows averaged utility per user (averaged over 100 channel realizations) versus the number of users. Different curves are shown for different numbers of channels. When the network has only two channels, there is a large amount of interference on each channel (low SINR), and the PADP performs substantially better than IWF (e.g., a factor of three in averaged utility). As the number of channels increases, this performance gain diminishes, although it is still significant with 10 channels (around 50%).

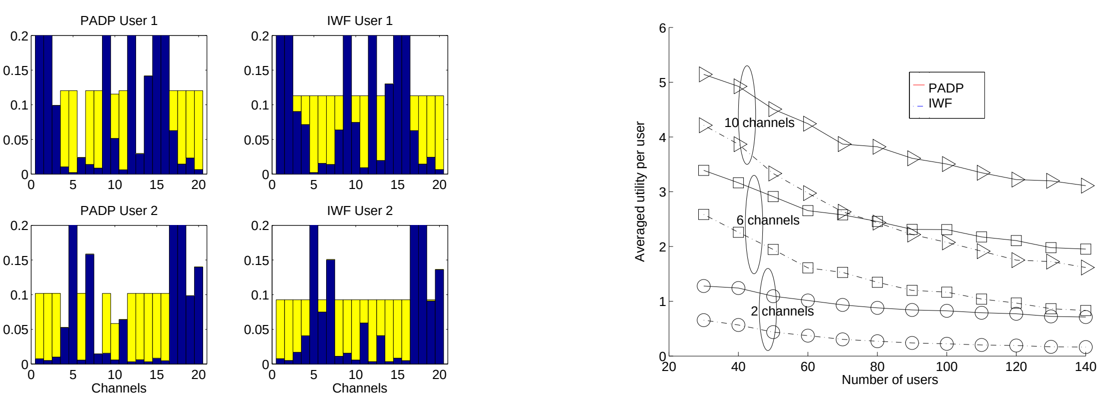

Figure 3: Performance of the PADP and IWF algorithms. *left*: power allocations for two users and 20 channels; *right*: averaged utility versus number of users.

5 Conclusions

We have presented two distributed power control algorithms for multi-channel ad hoc networks. The algorithms are based on exchanging price signals within each channel that measure the associated interference externalities. The DADP algorithm converges to the unique global optimal power allocation provided that the utility functions are "sufficiently" concave. Convergence of the PADP algorithm is more difficult to establish in general, although numerical results have shown that the algorithm converges with rate utility functions most of the time. Our numerical results also show that convergence of the DADP algorithm depends primarily on the total number of primal updates, and is insensitive to the number of primal updates per dual iteration. Here we have focused on a static setting, where the communicating pairs and channel conditions are fixed. An interesting future direction is to consider distributed power control in dynamic environments.

References

[1] O. Popescu, C. Rose, and D. C. Popescu, "Signal space partitioning versus simultaneous water filling for mutually interfering systems," in *Proceedings of IEEE Globecom*, 2004.

[2] W. Yu, R. Lui, and R. Cendrillon, "Dual optimization methods for multiuser orthogonal frequency division multiplex systems," in *Proceedings of IEEE Globecom*, vol. 1, 2004, pp. 225–229.

[3] R. Cendrillon, W. Yu, M. Moonen, J. Verlinden, and T. Bostoen, "Optimal multi-user spectrum management for digital subscriber lines," accepted by IEEE Transactions on Communications, 2005.

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

[5] 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.

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

[7] P. Marbach and R. Berry, "Downlink resource allocation and pricing for wireless networks," in *Proceedings of IEEE INFOCOM*, New York, NY, 2002.

[8] J. Huang, R. Berry, and M. L. Honig, "A game theoretic analysis of distributed power control for spread spectrum ad hoc networks," in *Proceedings of IEEE International Symposium on Information Theory*, 2005.

[9] ——, "Spectrum sharing with distributed interference compensation," in *Proceedings of IEEE Dynamic Spectrum Access Network (DySPAN)*, 2005.

[10] W. Yu, G. Ginis, and J. Cioffi, "Distributed multiuser power control for digital subscriber lines," *IEEE Journal on Selected Areas in Communication*, vol. 20, no. 5, pp. 1105–1115, June 2002.

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

[12] J. Huang, R. Berry, and M. L. Honig, "Distributed interference compensation in wireless networks," submitted to *IEEE Transactions on Selected Areas of Communications*, 2005.

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