Pricing Algorithms for Power Control and Beamformer Design in Interference Networks

David A. Schmidt, Changxin Shi, Randall A. Berry, Michael L. Honig, and Wolfgang Utschick

Full Paper Submission for the IEEE Signal Processing Magazine,
Special Issue on Game Theory

I. Introduction

Achieving high spectral efficiencies in wireless networks requires the ability to mitigate and manage the associated interference. This becomes especially important in networks where many transmitters and receivers are randomly placed, so that in the absence of coordination a particular receiver is likely to encounter significant interference from a neighboring transmitter. A challenge is then to provide a means for coordination, which allocates available resources, or Degrees of Freedom (DoFs), at each transmitter to avoid interference wherever possible.

In a wireless network DoFs generally refer to noninterfering modes of transmission, and can be defined in frequency, space, and time. DoFs in frequency and time typically correspond to multiple non-overlapping channels and time slots, respectively, and DoFs in space correspond to orthogonal spatial beams. *Resources* then include available power and beams at each transmitter, and the objective of the network operator is to allocate those resources among available DoFs to optimize an overall network objective.

In principle, optimal resource allocation in a wireless network can be achieved if all active nodes measure all channel gains to all other nodes, and pass that information to a resource manager. After collecting this information the resource manager would then determine allocations over all users, or transmitter-receiver pairs, and announce those to the various transmitters. Of course, such a centralized

D. Schmidt and W. Utschick are with the Associate Institute for Signal Processing at Technische Universität München, 80290 Munich, Germany. Email: {dschmidt|utschick}@tum.de

C. Shi, R. Berry, and M. Honig are with the Department of Electrical Engineering and Computer Science at Northwestern University, Evanston, Illinois 60208, USA. Email: cshi@u.northwestern.edu, {rberry|mh}@ece.northwestern.edu

scheme for resource allocation requires excessive information exchange and overhead for most practical networks.[1]

In this paper we discuss *distributed* resource allocation schemes in which each transmitter determines its allocation autonomously, based on the exchange of *interference prices*. These prices arise naturally when maximizing a sum utility objective. More specifically, to model user demands for service and associated priorities, we assume that each user is assigned a utility function, which depends on a Quality of Service metric such as Signal-to-Interference Plus Noise Ratio (SINR) or rate. The network objective is to maximize the sum utility over all users. This objective is flexible enough to accommodate a wide range of performance metrics through an appropriate assignment of utility functions. (For example, this includes the weighted sum of all user rates.)

Each interference price is associated with a particular receiver, and indicates the marginal decrease in utility due to a marginal increase in interference associated with a particular DoF. The cost of interfering with the DoF at that receiver is then the interference price times the received interference power. Given a set of interference prices from nearby receivers, a transmitter selects resources according to a *best response*, which maximizes its utility minus the total interference cost (summed over all neighboring receivers). Users then iterate between price and resource allocation updates. These schemes have been primarily motivated by the *commons model* for spectrum sharing in which a user or service provider may transmit in a designated band provided that they abide by certain rules (e.g., a standard such as 802.11) [1]. An attractive property of these schemes is that they are scalable, i.e., the information exchange and overhead can be adapted according to the size of the network.

A basic assumption for this class of distributed resource allocation schemes is that the users *cooperate* by following the rules for announcing prices and determining allocations. (For example, the pricing protocol may be built into approved devices.) This is in contrast to a network of *non-cooperative* users who may choose to deviate from the algorithm to increase their own utility. It is well known that such deviations from an optimal allocation (i.e., that maximizes sum utility) may lead to a different allocation (Nash equilibrium) for which most users are worse-off. This type of noncooperative behavior also arises in the absence of information exchange, since a transmitter would then presumably optimize its own resources ignoring the effects on neighboring receivers.

Applications of game theory to networking typically assume non-cooperative users, since those scenar-

[1]In addition, depending on the objective and specific resource constraints, the centralized optimization problem can be nonconvex with associated worst-case complexity that increases exponentially with the number of users and DoFs.

ios fit naturally within game theoretic frameworks. For example, a game theoretic approach to autonomous power control in cellular systems is discussed in [2], [3], [4]. Noncooperative adjustments of both power and bandwidth (spreading) for an interference channel are studied in [5]. The motivation there is to use game theory to model the underlying preferences of selfish users. In contrast, here we use game theory as an *engineering tool* to *design* the preferences of each user (or *agent*) so that "selfish actions" (best response updates in the resulting game) lead to desired cooperative behavior, i.e., convergence to an optimal (utility-maximizing) allocation. An analytical difficulty, which also motivates the application of game theory, is that the powers and beams can change substantially after each best response update. Hence we require different analytical tools from those typically used to analyze the convergence of gradient-based algorithms, which make incremental adjustments at each iteration [6], [7].

We start with a network consisting of narrowband, Single-Input Single-Output (SISO) links, and state some general conditions that guarantee convergence of the distributed pricing algorithm to a global optimum. We then extend this discussion to Multi-Input Single-Output (MISO) and Multi-Input Multi-Output (MIMO) channels. In general, a set of prices corresponding to all DoFs must be exchanged to achieve the centralized optimal allocation. The corresponding solution is contrasted with that achieved without information exchange (namely, iterative water-filling [8], [9]).

Finally, although the interference pricing schemes are discussed in the context of peer-to-peer networks, the approach can also be applied to cellular networks. Specifically, interference prices could be exchanged between neighboring cells to adjust power levels and DoFs to mitigate other-cell interference (e.g., see [10]). Other extensions and limitations of distributed pricing are discussed in Section V.

II. Peer-to-Peer System Model

We consider a wireless system consisting of a number of transmitter-receiver pairs. These pairs could, for example, be a base station and an associated mobile station in the downlink of a cellular system, or nodes in an ad-hoc network randomly placed within a geographic region, as shown in Figure 1. We will refer to each transmitter-receiver pair as a *user*. Each receiver is only interested in the signal from its associated transmitter; the signals from all other transmitters constitute the interference, which is assumed to be treated as noise. In addition to the interference, the receivers w.l.o.g. all experience the same level of background noise.

The conditions of the wireless channel are reflected in random channel gains between each transmitter and each receiver. We assume each node (receiver or transmitter) perfectly estimates all relevant channel gains. Also, each receiver must know its direct and cross-channel gains from neighboring transmitters,

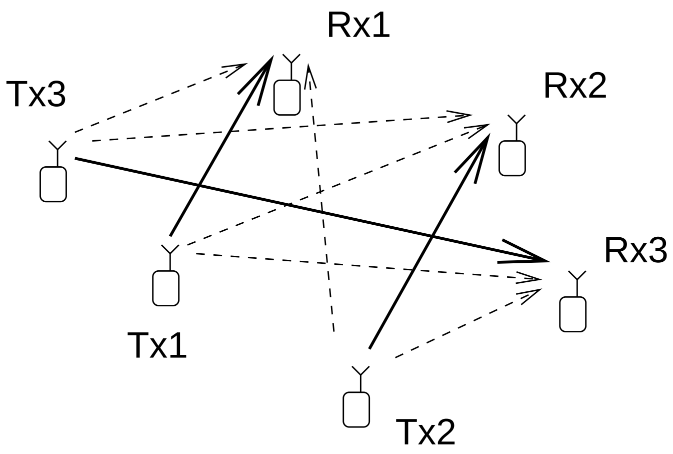

Fig. 1. A wireless network with multiple peer-to-peer transmissions.

and each transmitter must know the cross-channel gains to neighboring receivers. In practice, these gains can be estimated if each node periodically sends a known pilot.[2] Furthermore, we assume that channel conditions remain constant for the duration of the resource allocation procedures discussed corresponding to stationary users.

Given K users (transmitter/receiver pairs), the signal at receiver k corresponding to a particular symbol interval is

$$\boldsymbol{y}_k = \underbrace{\boldsymbol{H}_{kk}\boldsymbol{x}_k}_{\text{desired signal}} + \underbrace{\sum_{j \neq k} \boldsymbol{H}_{kj}\boldsymbol{x}_j}_{\text{interference}} + \underbrace{\boldsymbol{n}_k}_{\text{noise}}, \tag{1}$$

where in general $\boldsymbol{H}_{ki}$ is a complex matrix representing the channel from transmitter i to receiver k, $\boldsymbol{x}_k$ is the transmitted signal vector from transmitter k, and $\boldsymbol{n}_k$ is additive noise. Different assumptions concerning the number of available sub-channels and antennas at each node can be represented by changing the structure of the channel matrices. The simplest scenario corresponds to narrowband transmissions with single-antenna terminals (SISO model) in which case the channels and transmitted signals are scalars, i.e., $\boldsymbol{H}_{kj} = h_{kj}$ and $\boldsymbol{x}_k = x_k$,

[2]This may require some degree of synchronization among the nodes. Alternatively, a cross-channel gain at a transmitter might be measured when detecting an interference price from a particular receiver. Also, it is only necessary to measure cross-channel gains for strong interferers.

A. Optimization Problem

The performance perceived by user k is assumed to be represented by a utility function $u_k(\cdot)$. The argument can, in principle, be any quality of service metric, such as received Signal-to-Interference-Plus-Noise-Ratio (SINR), error rate, throughput, or packet delay. Here we will assume that the utility depends on a set SINRs over available DoFs. For example, for the SISO model the SINR at receiver k is

$$\mathrm{SINR}_k = \frac{|h_{kk}|^2 p_k}{\sum_{j \neq k} |h_{kj}|^2 p_j + \sigma^2}, \tag{2}$$

where p_k is the power of the transmit signal x_k and σ^2 is the power of the background noise. A user can therefore increase its SINR by increasing the transmit power, but then decreases the SINR for all other users. Typically the transmit power p_k is constrained to be no more than a maximum value P_k.

A prominent example of a utility function is $\log(1+\mathrm{SINR}_k)$, which corresponds to the Shannon capacity of the channel, since the interference is treated as additive Gaussian noise. Other utility functions, which have desirable properties, include $\log \mathrm{SINR}_k$, corresponding to the Shannon capacity at high SINRs, and the "α-fair" utility function SINR_k^α **[?]**, which "flattens out" at high SINRs when $\alpha \leq 1$, and therefore reflects an application that becomes insensitive to rate. In general, any sensible utility function should be non-decreasing in the SINR; more stringent criteria must be satisfied by the utility functions to prove convergence of the distributed pricing algorithms to be presented. Also, when a user transmits over multiple ($L > 1$) DoFs, its utility function depends on the set of SINRs over the DoFs. For example, the rate utility is given by $u_k(\mathrm{SINR}_{k,1}, \cdots, \mathrm{SINR}_{k,L}) = \sum_{i=1}^{L} \log(1 + \mathrm{SINR}_{k,i})$ where $\mathrm{SINR}_{k,i}$ is user k's SINR for the i-th DoF.

Given an assignment of utility functions to users, the overall system objective is to maximize the sum-utility across users. For the SISO model this can be posed as the following optimization problem:

$$\max_{p_1,\ldots,p_K} \sum_{k=1}^{K} u_k(\mathrm{SINR}_k) \qquad \text{s.t.:} \quad 0 \leq p_k \leq P_k \quad \forall k \in \{1, \ldots, K\}. \tag{3}$$

The properties of this optimization problem depend on the utility functions. In general this is a non-concave problem, due to the interference, so it may have many multiple locally optimal solutions, making the search for the global optimum a difficult task. However, as shown in [11] for the SISO model, for many utilities of interest this can be transformed into a convex problem by applying a logarithmic change of variables. Specifically, for a wide class of utility functions the objective in (3) is concave in terms of the variables $e_k = \ln(p_k)$ and the constraint set is convex. In those cases, (3) has a unique local optimum, which is also the global optimum, and can be solved using standard optimization techniques. A sufficient condition for this to hold is given in terms of the *coefficient of relative risk aversion* of the individual

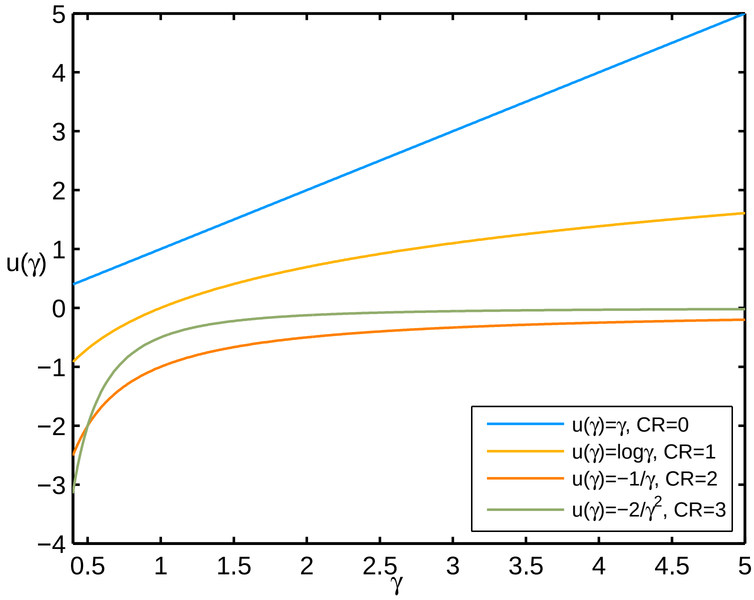

Fig. 2. Plots of utility functions having a constant coefficient of relative risk aversion (CR). As the CR increases, the functions exhibit stronger concavity.

utility functions. This is a measure of the relative concavity of a utility functions used in economics and is defined by

$$CR_k(x) = -\frac{u_k''(x)x}{u_k'(x)},$$

where $u_k'(x)$ and $u_k''(x)$ denote the first and second derivatives of the utility with respect to its argument. For any concave utility $CR_k(x) \geq 0$, with equality when $u_k(\cdot)$ is linear. As $CR_k(x)$ increases, the utility function becomes "more concave". This is illustrated in Figure 2, which shows examples of α-fair utility functions. For these utilities, $CR_k(x)$ is a constant (independent of x). The sufficient condition needed to transform (3) into a convex problem is that $CR_k(x) \geq 1$ for all feasible SINRs, i.e., the utilities must be sufficiently concave **??**. This holds for logarithmic utilities, but not for the rate-utility $\log(1 + \mathrm{SINR}_k)$.

B. Utility Region and Nash Equilibrium

First consider a centralized solution to (3). A resource manager would have to collect the necessary information about channel gains and utilities, solve the optimization problem, and then tell each user how to allocate their powers over DoFs. In this manner, a locally optimal solution will be found (which may be globally optimal if the problem can be transformed into a convex problem). The information exchange overhead required for this resource allocation scheme, however, is likely to be prohibitive for moderate-sized to large networks.

An alternative to the previous centralized solution is to allow each transmitter to adapt its own power autonomously. An extreme example of such a distributed allocation scheme, which does not require any information exchange, is for each transmitter to choose its power allocation to maximize its own utility. For the SISO model the SINR and utility increase with the transmit power. Hence the *best response* for each user is to transmit with as much power as possible, i. e., $p_k = P_k$, regardless of the other users' behavior. If we interpret the users as players in a game where the strategy is the choice of transmit power and the payoff is the resulting utility, then transmitting at full power will be a dominant strategy for each player and so this outcome is the unique *Nash equilibrium*.

It is well known that a Nash equilibrium need not maximize the sum-utility. Indeed it may not even be *Pareto optimal* in which case it would be possible to improve *all* of the players' payoffs by choosing a different set of allocations. In the context of the SISO problem, this means that a configuration in which one or more users transmit with reduced power can lead to an increase in all users' utilities. The set of Pareto optimal outcomes can be visualized as boundary points of the *utility region*, i.e., the set of all possible individual utilities that can be achieved for some feasible power allocation. For example, the union of all combinations of $\log(1 + \mathrm{SINR}_k)$-utilities is the achievable *rate region*. When a Nash equilibrium is not Pareto optimal, it lies strictly in the interior of this region.

To illustrate the preceding discussion, Fig. 3 shows a typical rate region for a two-user MISO interference network with two-antenna transmitters and single-antenna receivers. In this case, the channel matrices $\boldsymbol{H}_{kj}$ in (1) are row vectors, and the signal vector $\boldsymbol{x}_k = s\boldsymbol{v}_k$ where s is the scalar symbol and $\boldsymbol{v}_k$ is the beamforming vector of antenna weights (or beam) for user k. Fig. 3 shows the boundary of achievable rates over all possible choices of beams, subject to the same power constraint for each transmitter.

The boundary point in Fig. 3 corresponding to the maximum sum utility has a tangent line with slope -1. Also shown is the Nash equilibrium point, corresponding to non-cooperative users (or equivalently, zero information exchange). (The characterization of beams corresponding to the Nash equilibrium and the optimal sum rate point will be discussed in Section IV-A.) The Nash equilibrium point is far away from the Pareto boundary of the achievable rate region. Even a simple time-sharing scheme between the two users can yield better performance for each individual user than the non-cooperative outcome. Note that the rate region has distinct non-convexities; the convex hull of the rate region is also achievable if time sharing between strategies on the boundary is allowed. The shape of the rate region also depends on the SNR. In particular, the non-convexities shown in the figure correspond to high SNRs (e.g., see [**?**]).

For both the SISO and MISO channel models discussed so far, a transmitter's best response in the absence of information exchange does not depend on the powers or beams of other users. Computing the Nash equilibrium is therefore quite simple. In contrast, suppose that the channel matrices $\boldsymbol{H}_{kj}$ in the channel model (1) are diagonal with complex diagonal elements, which corresponds to multi-carrier transmission through frequency-selective channels (with single antennas). The diagonal elements of a channel matrix then correspond to the complex channel gains across frequency.

Each multi-carrier transmitter can allocate powers over the sub-channels to maximize utility subject to a total power constraint. For rate utilities the best response power allocation in the absence of information exchange is water-filling. This best response depends on the distribution of interference over sub-carriers, which in turn is determined by power allocations at neighboring transmitters. Hence, a Nash equilibrium can be determined by an *iterative water-filling* method in which users update their power allocations until the power allocations converge [**?**], [**?**]. (The updates could occur sequentially in any order or synchronously.) In fact, for the peer-to-peer networks considered, a Nash equilibrium need not exist and even when one does exist, iterative water-filling does not always converge. For the two-user channel this depends on the relative magnitude of the cross-channel gains [**?**]. In contrast, the distributed pricing algorithms to be discussed often converge even when iterative water-filling does not.

III. Distributed Pricing Algorithms

As shown in Fig. 3, having each user optimize its own utility without any information exchange can lead to poor performance. On the other hand, obtaining full information as in a centralized solution may require excessive overhead. We next introduce the *asynchronous distributed pricing* (ADP) algorithm [12] that enables each node to adapt their resources locally with the aid of *limited* information exchange. In this section, we describe this algorithm for the SISO model; extensions to multiple antennas and multi-carrier transmission will be described in subsequent sections.

The reason the non-cooperative approach in the previous section can result in poor performance is that when the users optimize only their own utility functions they do not account for the *dis-utility* they cause other users due to interference. In economic terms, a dis-utility of one agent due to the action of another is referred to as a *negative externality*. These negative externalities are the root of the inefficiencies of the non-cooperative approach. The main idea behind the ADP algorithm is to design a new payoff, which internalizes these externalities for each agent, i.e., so that when an agent optimizes the payoff, it is taking

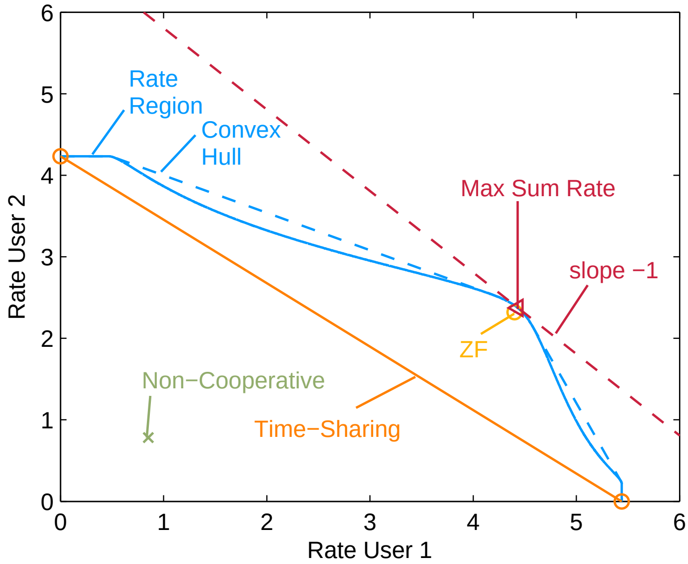

Fig. 3. Illustration of the achievable rate region for a two-user MISO network.

into account the interference to other users. For an agent to predict the exact effect of the interference it causes would again require excessive information exchange (essentially, each agent would need to know the entire global objective). Instead, in the ADP algorithm, each agent (receiver) announces a single *interference price*, which is the marginal cost of their own utility per unit interference. Specifically, the interference price announced by receiver j is given by

$$\pi_j = -\frac{\partial u_j(\mathrm{SINR}_j)}{\partial I_j} \tag{4}$$

where $I_j = \sum_{i \neq j} p_i |h_{ji}|^2$ is the total interference power at receiver j. If transmitter k has power p_k, then the "cost" of interfering with receiver j is then $\pi_j p_k |h_{jk}|^2$, i.e., the interference price times the received interference power. We can again view each user as a player in a game, only now instead of maximizing its own utility, the user's payoff is its utility minus the *total* cost (marginal decrease in utility) summed over all receivers,

$$\Pi_k(p_k; \boldsymbol{p}_{-k}) = u_k(\mathrm{SINR}_k) - p_k \sum_{j \neq k} \pi_j |h_{jk}|^2, \tag{5}$$

where $\boldsymbol{p}_{-k}$ indicates the vector of strategy choices (powers) over all user except for k.

In the ADP algorithm, the users iteratively adapt their power allocations and announce new interference prices. When a user adapts its power, it maximizes the payoff in (5) assuming that the power allocations and interference prices of the other users are fixed. The power update is therefore the best response of

1) INITIALIZATION: Each user k chooses an arbitrary initial price $\pi_k(0) \geq 0$ and power $0 \leq p_k(0) \leq P_k$.
2) POWER UPDATE: Each user in an arbitrary group updates its power to maximize (5).
3) PRICE UPDATE: Each user in an arbitrary group updates its interference price according to (4).
4) GOTO 2 and repeat.

Fig. 4. The ADP algorithm for the SISO model (narrowband transmissions with single-antenna terminals).

an agent in the resulting game. The complete algorithm is shown in Fig. 4. ("Arbitrary group" refers to a subset of users that update simultaneously.) In general this can be completely asynchronous, meaning that there are no restrictions on when a specific user updates its price or power except that each quantity must be updated infinitely often. This therefore includes simultaneous updates and round-robin updates as special cases.

To implement this algorithm, note that when updating the transmitted power according to (5), the transmitter needs to know the interference prices from other receivers, its own SINR, and the cross-channel gains to neighboring receivers. It does not need to know the other channel gains in the network and other users' utility functions. Also, to compute the interference price (4), the receiver must know the direct channel gain and the interference-plus-noise power (or equivalently, the SINR and the transmitted power). We next discuss the convergence of this algorithm.

A. Convergence Analysis

If the ADP algorithm has converged, then each user's best response using the payoff in (5) does not change the transmitted power. Additionally the interference prices should not change, i.e., they should represent the marginal cost of interference at the equilibrium. It can be shown that such a fixed point must satisfy the Karush-Kuhn-Tucker (KKT) optimality conditions for the centralized problem (3) [12]. (In words, the KKT conditions state that the marginal increase in utility to any user k that increases its power is equal to the total interference cost.) Thus any limit point of the algorithm satisfies the necessary conditions for a local or global optimum. Likewise, any local or global optimum of (3) must be a fixed point of the algorithm. If in addition we have $CR_k(x) \geq 1$ for each user k, then as discussed in the previous section, the overall problem can be transformed into a convex problem in which case the only allocation satisfying the KKT conditions (corresponding to the only limit point of the ADP algorithm) is the global optimum.

We next give conditions that guarantee convergence of the ADP algorithm. Showing convergence in this setting is different from the analysis of standard distributed optimization algorithms, such as distributed gradient or Newton-based methods (see e.g. [13], [**?**]), in that we do not place any step-size restriction on the power update of a user. This has the advantage of not requiring such a step-size, but also potentially complicates the analysis since the overall utility could change dramatically in a single time-step. Instead, the convergence proof of the ADP algorithm in [12] uses properties of best response updates in a game in which each user k is represented by two players: a power player, which chooses a feasible power allocation p_k to maximize the payoff in (5), and a price player, which chooses an interference price by maximizing a payoff that is optimized by the price in (4). A best response update of a power or price player in this game corresponds to a power or price update, respectively, in the ADP algorithm. At a Nash equilibrium of this game, no power or price player wishes to deviate, i.e. all Nash equilibria must be limit points of the ADP algorithm.

From the preceding discussion it can be seen that the convergence of the ADP algorithm is equivalent to showing that best response updates converge in the underlying game. For an arbitrary game, a Nash equilibrium need not exist and even if it does exist, best response updates need not converge to it. However, for the class of *supermodular games* (see side-bar in Appendix), much more is known about the convergence of such updates, even when done in an arbitrary asynchronous manner. For the ADP algorithm, if the coefficients of relative risk aversion of each user's utility satisfy $1 \leq CR_k(\mathrm{SINR}_k) \leq 2$ for all feasible SINR's, then the resulting game is supermodular. Moreover, in this case, (3) has a unique global optimum and so it follows that the ADP algorithm will globally converge to the optimum power allocation.

The main restriction for the preceding convergence result of the ADP algorithm is that the coefficient of relative risk aversion for each user's utility lie between 1 and 2. This can be interpreted as requiring that the utilities are sufficiently concave, but not too concave. If they are not sufficiently concave ($CR_k < 1$), then the overall problem may have multiple local optima. If the utilities are too concave ($CR_k > 2$), then the game is not supermodular and the best response updates may be too aggressive to guarantee convergence. One utility for which this condition applies is $u_k(x) = \log(x)$, the high-SINR approximation of the Shannon rate. This result does not apply for the rate utility $u_k(x) = \log(1 + x)$, which has $CR_k(x) < 1$. However, in [14] it shown that for rate utilities a version of the ADP algorithm with *fixed* interference prices converges to a unique Nash equilibrium provided that the channel is diagonally dominant, i.e., for each user k, $|h_{kk}|^2 \geq \sum_{j \neq k} |h_{jk}|^2$. Additionally, in [15], it is shown that the ADP monotonically converges for all utilities with $0 \leq CR_k \leq 2$ with sequential power updates provided that after each update every user announces a new interference price.

B. Multi-Carrier ADP

Consider now the channel model with diagonal channel matrices, corresponding to multi-carrier transmission. As for the SISO model, it is possible to improve upon the performance with non-cooperative users (iterative water-filling, or more generally, iterative utility maximization), by exchanging interference prices. An ADP algorithm for multi-carrier transmission is considered in [12] assuming that the utility for each user k is *separable*, meaning $u_k(\mathrm{SINR}_{k,1}, \cdots, \mathrm{SINR}_{k,M}) = \sum_m u_k(\mathrm{SINR}_{k,m})$, where $\mathrm{SINR}_{k,m}$ is the SINR for user k on sub-channel m. (For example, this applies to the rate-utility function.) In that case the sum utility can again be maximized by exchanging interference prices over all sub-channels, and using an iterative primal-dual algorithm to update powers and prices (taking into account power constraints at each transmitter).

For many applications, the utility function could be a nonlinear (e.g., increasing strictly concave) function of the total rate, in which case it cannot be expressed as the sum of utilities across sub-channels. Extensions of the ADP to this setting are considered in [**?**]. Namely, each user announces a set of interference prices $\pi_k^m = -\partial u_k[R_k(\mathbf{P})]/\partial I_k^m(\mathbf{P}_{-k})$, $m = 1, \cdots, M$, where I_k^m is the interference power for user k on sub-channel m, and $\mathbf{P}_{-k}$ denotes the set of all user powers except those for user k. The users then optimize surplus (a modified version of the payoff in (5)) as before. Although numerical examples have indicated that this algorithm typically converges rapidly to a locally optimal solution, obtaining general conditions on the utility functions that guarantee convergence of this algorithm is an open problem.

IV. Multiple Antennas

A. MISO System Model

The preceding discussion applies to interference networks consisting of single-antenna terminals. We now discuss the extension of those results to terminals with multiple antennas, which add spatial DoFs. We start by considering multiple antennas at the transmitter and continue to assume receivers with a single antenna; this is a common situation in the downlink of cellular systems.

As already mentioned in Section II-B, for the MISO channel model (1) the channel matrices become row vectors, so that we define $\boldsymbol{h}_{kj}^{\dagger} = \boldsymbol{H}_{kj}$, where '†' denotes Hermitian transpose. The multi-antenna transmitter k can now vary both its power and the beam $\boldsymbol{v}_k$, subject to the power constraint $\|\boldsymbol{v}_k\|_2^2 \leq P_k$. In analogy to the single-antenna case, we can define the SINR as

$$\mathrm{SINR}_k = \frac{|\boldsymbol{h}_{kk}^\dagger \boldsymbol{v}_k|^2}{\sum_{j\neq k} |\boldsymbol{h}_{kj}^\dagger \boldsymbol{v}_j|^2 + \sigma^2}. \tag{6}$$

Once again, the system objective is to maximize the sum utility, where each user is assigned a utility function $u_k(\mathrm{SINR}_k)$, which depends on the received SINR.

If the users are non-cooperative, and are only interested in maximizing their own utility, then each user k chooses $\boldsymbol{v}_k$ to maximize $|\boldsymbol{h}_{kk}^\dagger \boldsymbol{v}_k|^2$. Consequently, $\boldsymbol{v}_k$ should be a scaled version of $\boldsymbol{h}_{kk}$, using the maximum power. This is also called the *matched filter* solution, and corresponds to the Nash equilibrium (non-cooperative) point shown in Fig. 3. Even with very low background noise, the SINR and thus the utility for each user is strictly limited by the interference caused by the other users.

In contrast to the single-antenna system, the transmitters in the MISO system can make use of the spatial DoFs (corresponding to the choice of beams) to *avoid* interfering with other users. Consider, for example, the following *altruistic* scheme: the interference caused to every unintended receiver is forced to be zero, i. e., $\boldsymbol{h}_{jk}^\dagger \boldsymbol{v}_k = 0$ for all $j \neq k$. [3] When all users adhere to this scheme, there is no interference at any receiver and the SINR grows without bound as the background noise decreases. In fact, as the noise level diminishes to zero, this *zero-forcing* scheme optimizes the sum rate, so is clearly better for each user than the *egoistic* matched filter solution.

Referring to Fig. 3, the noise level is low, so that the zero-forcing solution is quite close to the point corresponding to maximum sum rate. However, as the noise power increases, the gain obtained by completely removing the interference decreases, and the non-cooperative matched filter scheme eventually performs better. Also, at each edge of the region, there is a section in which the boundary runs perpendicular to the coordinate axis; these sections correspond to strategies where one user employs the matched filter while the other user performs zero-forcing.

It is shown in [16] that for the case of two users and the rate utility function $\log(1 + \mathrm{SINR}_k)$, any Pareto optimal strategy requires that both users employ a linear combination of the altruistic and egoistic beams. The altruistic and egoistic solutions can thus be viewed as two extremes, in between which any desirable transmit strategy should lie. As in the single antenna case, our goal is to find a *distributed* transmit strategy that maximizes the sum utility (cf. (3)).

[3]Note that such a solution exists only if the transmitter has at least as many antennas as the number of users with which it interferes. Also note that in contrast to the matched filter solution, transmitter k must know all channel vectors $\boldsymbol{h}_{jk}$ to other receivers, requiring additional communication among the users.

B. MISO Pricing Algorithm

The ADP algorithm for a SISO network can be naturally extended to the MISO case. Following the same principle, each user k will again announce an interference price π_k representing the marginal decrease in that user's utility following a marginal increase in interference. Given these interference prices, we can again view each user as a player in a game with a payoff given by their utility minus the cost due to the interference they generate. However, in this case user's best response update is given by optimizing their payoff over beams instead of power, namely

$$\max_{\boldsymbol{v}_k} \; u_k(\mathrm{SINR}_k) - \sum_{j \neq k} \pi_j \cdot |\boldsymbol{h}_{jk}^{\dagger} \boldsymbol{v}_k|^2 \qquad \text{s. t.:} \quad \|\boldsymbol{v}_k\|_2^2 \leq P_k. \tag{7}$$

The users then iteratively update their beams and interference prices as before.

In the SISO case a user's best response update involves solving an optimization over a single variable, which can often be done in closed form. In (7), the optimization is over a vector and in general the solution cannot be obtained in closed-form. (Also, for reasonable utility functions the objective is likely to be neither convex nor concave.) **right?** Each user must therefore use an optimization algorithm to calculate its best response. To reduce this complexity, we can linearly approximate $u_k(\mathrm{SINR}_k)$ in (7) yielding the optimization problem

$$\max_{\boldsymbol{v}_k} \; \rho_k \cdot |\boldsymbol{h}_{kk}^{\dagger} \boldsymbol{v}_k|^2 - \sum_{j \neq k} \pi_j \cdot |\boldsymbol{h}_{jk}^{\dagger} \boldsymbol{v}_k|^2 \qquad \text{s. t.:} \quad \|\boldsymbol{v}_k\|_2^2 \leq P_k. \tag{8}$$

where $\rho_k = u_k'(\mathrm{SINR}_k)$ **right?**. The solution to this problem can be found by simply determining the dominant eigenvector of a matrix. It can be seen from the KKT conditions that stationarity in problem (8) implies stationarity in (7), and therefore both converge to a local optimum of the sum-utility problem. For the following convergence analysis, however, it is assumed that problem (7) is solved exactly.

To implement this MISO ADP algorithm, each user must still announce only a single (scalar) interference price. However, to calculate their best response, they now need to measure the vector of channel gains to each neighboring receiver ($\boldsymbol{h}_{jk}$) as well as their own, which increases the required overhead for channel estimation. Of course the information needed for alternative approaches such as a centralized solution also increases accordingly.

Fig. 5 shows how the schemes discussed perform in terms of sum rate versus the level of background noise. **reviewer asked for different line types to distinguish in B&W** A two user network is assumed with two antennas at each transmitter. The results are averaged over a large number of random channel realizations. While the non-cooperative strategy is optimal for strong background noise (i. e., low SNR), its performance deteriorates quickly compared to the other schemes when the background noise becomes

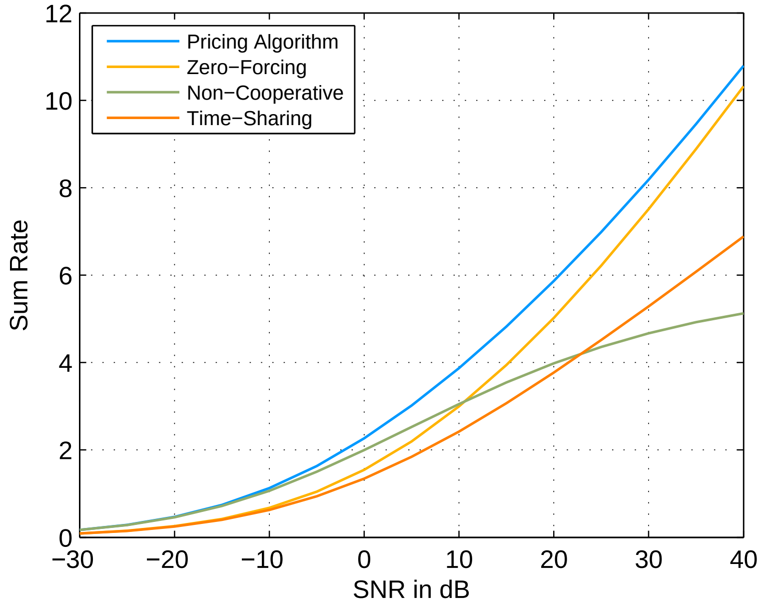

Fig. 5. Illustration of the sum rate versus SNR for different algorithms

insignificant. It is surpassed first by the zero-forcing scheme, which is optimal in the complete absence of background noise, but eventually also by a simple scheme in which the two users alternately transmit (time-sharing). The pricing algorithm performs best regardless of the noise conditions.

1) Convergence Analysis: As in the SISO case, it can be shown that any limit point of the MISO ADP algorithm satisfies the KKT conditions of the overall optimization problem. Once again such a limit point can be viewed as a Nash equilibrium of an underlying game. In the SISO case, when the utility functions satisfy certain conditions this game is supermodular. However, the argument cannot be directly extended to the MISO case. In particular, the notion of supermodular games can be extended to games in which the players have vector-valued strategies, but in this case it is required that the set of feasible strategies are a *lattice*, meaning that if any two strategies are allowable, then their componentwise maximum and minimum are also allowable. Due to the agents' power constraints, their strategies in terms of beam choices do not have this property. For the special case of two users, in [17] two alternative proofs of convergence are given, which are based on looking at an alternative formulation of the underlying game. We discuss each of these next.

The first re-formulation is based on the observation that in a two-user network, an optimal beamformer $\boldsymbol{v}_k$ can always be found that lies in the convex cone spanned by the corresponding channel-matched beamformer $\boldsymbol{h}_{kk}$ and zero-forcing beamformer $\boldsymbol{h}_k^{\perp} = P_{\boldsymbol{h}_{jk}}^{\perp}\boldsymbol{h}_{kk}$, and always consumes the maximum

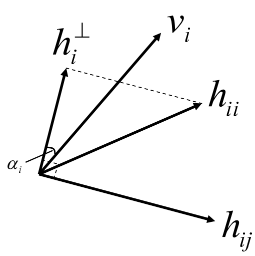

Fig. 6. Illustration of the optimal beamformer $\boldsymbol{v}_k$ in a two user network. The shaded area denotes the convex cone spanned by the channel matched filter and the zero-forcing beamformer. The boundary of this region corresponds to those beams that meet the power constraint.

power allowed (see Fig. 6). Using this observation, instead of viewing a user's strategy as the choice of a beam $\boldsymbol{v}$, we can think of a user's strategy as being the *angle* α between its beam and $\boldsymbol{h}_k^{\perp}$. This reduces the strategy set to a single scalar and then using similar ideas as in the SISO case, we can prove that the resulting game is supermodular if each user's utility satisfies $1 \leq CR(\mathrm{SINR}_k) \leq 2$ for all feasible $\mathrm{SINR_k}$. In fact, for 2 users, it can also be shown that the underlying game is supermodular if $0 < CR(\mathrm{SINR}) \leq 1$ for all feasible SINR. Thus, if we start from suitable initial conditions, the MISO-ADP algorithm can again be guaranteed to converge. (In this case, in general we do not know if the global optimization problem has a unique optimum and so global convergence can not be guaranteed).

The second reformulation in [17] is based on the observation that the only aspect of a user's strategy choice when solving (7) that effects the utility of any other user in the network is the total interference power received at that other user. Based on this we can instead view each user's action as determining the (scalar) interference powers received by other users. Given this power, a user's payoff is then given by optimizing his own beam subject to the chosen interference powers and his own power constraint. For a two user network, each user's strategy again becomes a scalar and we can again show that the resulting game is supermodular under the same restrictions on the two users' utility functions.

Both of the previous methods cannot be generalized to a MISO network with more than two users. For the method based on the angle parameterization, there is no clear generalization of this structure of the optimal solution with more than two users. The second formulation involving the interference power can be generalized to allow each agent to specify a vector of interference powers, one for each other receiver in the network. However, in this case we can construct examples which show that the resulting game is not always supermodular. This is due to the following key distinction between SISO and MISO networks. In a SISO network, when one user increases its strategy (i.e. power), it results in increased

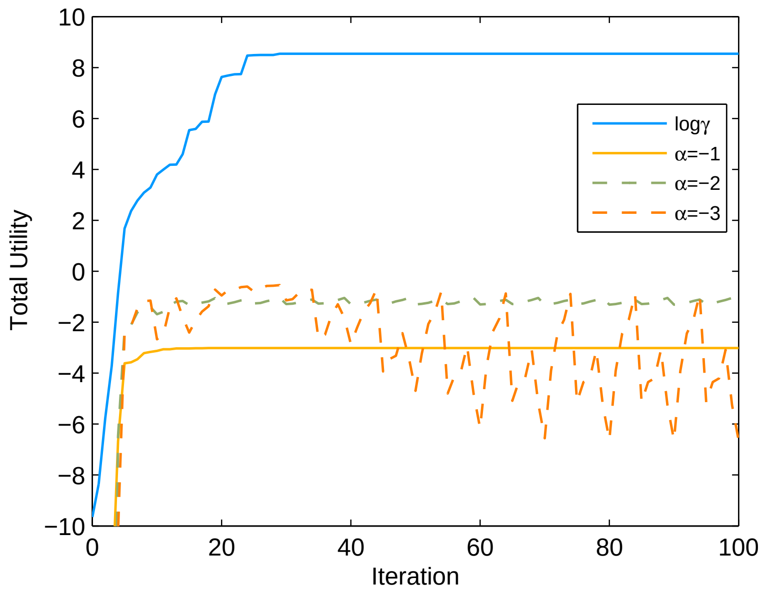

Fig. 7. Example of the convergence of the MISO-ADP algorithm for α-fair utilities.

interference at every other receiver. However, in a MISO network a user can increase the interference at one receiver while decreasing the interference at another. Hence there is no natural ordering of the users' strategies with regard to their effect on the other users, and makes it impossible to show the properties required by a supermodular game. Though convergence of the MISO ADP for more than two users cannot be proven using supermodularity, in [15] it is shown that, as for the SISO model, if users update their beams sequentially and exchange interference prices after each update, then the MISO ADP algorithm monotonically converges for all utilities with $0 \leq CR_k \leq 2$.

Figure 7 shows plots of the total utility versus the number of iterations (each corresponding to one beam update) for the MISO ADP algorithm in a network with 5 users and 3 transmit antennas. In these simulations each beam update is always followed by all users announcing their interference prices, so that the convergence results in [15] apply. Plots for four α-fair utilities are shown. As expected, for the $\log$ utility ($CR = 1$) and the $\alpha = -1$ utility ($CR = 2$) the algorithm converges. For the $\alpha = -2$ ($CR = 3$) and $\alpha = -3$ ($CR = 4$) cases it does not converge.

C. *MIMO Systems*

With MIMO channels the transmitted signal vector for user k becomes $\boldsymbol{x}_k = \boldsymbol{V}_k \boldsymbol{A}_k \boldsymbol{s}_k$, where $\boldsymbol{s}_k$ is a vector of transmitted symbols, $\boldsymbol{A}_k$ is a diagonal matrix of corresponding amplitudes, and $\boldsymbol{V}_k$ is a precoding matrix with unit-norm columns. The columns of $\boldsymbol{V}_k$ correspond to different spatial beams (DoFs), which multiplex different symbol streams. The centralized problem is then to select the set of precoding and amplitude matrices across users to maximize the sum utility subject to power constraints at each transmitter. For example, the utility for user k might be the achievable rate

$$u_k(\boldsymbol{Q}_k) = \log\det\left(\boldsymbol{I} + \boldsymbol{H}_{kk}^{\dagger}(\boldsymbol{R}_n + \sum_{i \neq k} \boldsymbol{H}_{ki}\boldsymbol{Q}_i\boldsymbol{H}_{ki}^{\dagger})^{-1}\boldsymbol{H}_{kk}\boldsymbol{Q}_k\right) \tag{9}$$

where $\boldsymbol{Q}_k = \boldsymbol{V}_k\boldsymbol{A}_k^2\boldsymbol{V}_k^{\dagger}$ is the transmit covariance matrix, $\boldsymbol{R}_n$ is the noise covariance matrix, and interference is again treated as additive Gaussian noise.

For the rate utility, it is straightforward to determine the best response strategies of non-cooperative users: each user k views all interference from other users as noise and selects the covariance matrix to maximize its own rate R_k. The solution to this single-user MIMO optimization is well-known: after a noise whitening filter, water-filling of powers over the singular values of the effective channel matrix is performed. However, as for multi-carrier channels discussed earlier, when one user adapts its transmit strategy in this way, the interference properties change for all other users, forcing them to readjust their strategies. This leads again to an iterative water-filling procedure for MIMO channels, which does not always converge [9], [8].

1) MIMO Pricing Algorithm: As for MISO channels, the non-cooperative strategy does not achieve full system potential, especially with low background noise. When averaged over many random channel realizations, the slope of the sum rate achieved by iterative waterfilling versus SNR shows a clear disadvantage compared to the optimal slope. Here we discuss distributed pricing schemes for MIMO channels that can improve performance.

There are several important differences from the MISO model, which makes finding a distributed solution to the utility maximization problem with MIMO channels considerably more complicated. First, the preceding Shannon rate with an optimal receiver is not a transparent function of a set of received interference powers (e.g., along different spatial directions), which makes it more difficult to define a set of interference prices that are used to update all precoders. Even with a centralized resource manager the globally optimal sum-rate strategy is difficult to find, since the problem is in general non-convex. A large number of locally optimal solutions has been observed for moderate system dimensions and low noise power [18].

A second complication is that the performance depends on the rank or multiplexing gain of each precoder matrix. For the MIMO interference channel a larger rank (multiplexing gain) generates more interference to neighboring receivers, since fewer DoFs are available for interference avoidance. Hence the ranks of the $\boldsymbol{V}_k$'s must be jointly optimized along with the columns and associated powers (set of $\boldsymbol{A}_k$'s). Finally, a third complication is that an optimal precoder matrix depends on the choice of neighboring receivers (e.g., optimal, corresponding to (9), linear, or decision-feedback). Furthermore, in the case of linear receivers the transmitter must know the cross-channel gains *combined with* the receive filters in order to determine the received interference (i.e., $\boldsymbol{G}_j^{\dagger}\boldsymbol{H}_{kj}$, where $\boldsymbol{G}_j$ is the receive filter at node j).

An approach presented in [19] is to view each column of $\boldsymbol{V}_k$ as a spatial "beam", which is used to carry a separate symbol stream, and assume that all receivers are linear. The utility for a user can then be defined as a function of the received SINRs across beams (e.g., sum rate over the beams), and interference prices can be defined and announced for each beam as before. Each user can then select each beam to maximize the associated utility minus cost. Of course, the achievable rate with linear receivers only approximates the optimal rate in (9). Also, the problem is still complicated by the joint optimization of precoder ranks and powers. Heuristic methods for beam updates are presented in [19], which trade off performance with the amount of information exchange. (Variations depend on the order in which beams are updated (e.g., sequentially versus all at once), how powers are allocated, and how often prices are announced.) The performance of the distributed algorithms is typically close to that obtained with centralized optimization, and achieves the optimal high-SNR slope. The performance can therefore be substantially better than iterative water-filling algorithm, especially at high SNRs.

V. Conclusions and Remaining Issues

Exchanging interference prices in a wireless network enables the transmitters to adjust their resources to optimize a network objective, as opposed to their individual single-link objectives. Although algorithms that iterate between price announcements and combined power/beam best response updates typically show rapid convergence to a limiting allocation, establishing general conditions that guarantee such convergence can be challenging. The concept of supermodularity, which arises in game theory, has been shown to be especially useful for this purpose. Specifically, convergence to the unique globally optimal allocation is guaranteed for SISO channels provided that the utility functions satisfy weak concavity properties. Moreover, the order in which price and best response updates occur can be arbitrary.

There are, however, several limitations of the results presented here. First, the condition on utility functions that guarantees convergence excludes the achievable (Shannon) rate. In fact, the sum rate objective can have multiple local optima, so that convergence to a global optimum cannot be guaranteed. Still, convergence to a local optimum has always been observed in numerical experiments. Second, with MISO channels supermodularity can only be applied to a network with two users, although convergence is typically observed with more than two users. Finally, establishing conditions for the convergence of iterative pricing algorithms with MIMO channels and adaptive receivers remains an open problem.

Although the sum utility achieved with distributed pricing can be substantially larger than that corresponding to the Nash equilibrium (no information exchange), the model presented here has ignored the overhead associated with exchanging interference prices. The information exchange overhead might be substantially reduced by exchanging prices corresponding to a subset of strongest interferers (e.g., see [**?**]). In addition to the power and bandwidth needed to exchange those prices, resources are needed to estimate the cross-channel gains used to compute the best response updates. Furthermore, in practice the prices and cross-channel gains will contain estimation error, the magnitude of which depends on the resources allocated for channel estimation. An accurate assessment of the benefits of distributed pricing must ultimately take this overhead and estimation error into account.

Finally, our model has assumed that the users do not deviate from the specified algorithm, meaning that each user truthfully announces a set of interference prices. A non-cooperative user could improve her individual performance by falsely reporting a larger set of interference prices. (Cross-channel gains might also be similarly manipulated.) One approach to this problem is to implement policing and punishment strategies for detecting such behavior (e.g., see [5]). Alternatively, it is possible to design auction mechanisms that provide incentives for truthful reporting, albeit with an associated loss in efficiency (sum utility) (e.g., see [20], [21]). Understanding the tradeoffs among these approaches in the context of interference networks with multiple DoFs poses major challenges for future investigation.

Appendix A
Side-bar: Supermodular Games

Supermodular games represent a class of games in which the players exhibit *strategic complementarities*. Loosely, this refers to the fact that the players' actions can be ordered so that an increase in one player's action results in the best response of every other player also increasing (or more precisely, not decreasing). This notion of complementarity is formalized by requiring that the player's payoff functions have *increasing differences*. To be more concrete, consider a game in which each player chooses a real-valued strategy x_k. Each player seeks to maximize a payoff function $\Pi_k(x_k, x_{-k})$, where x_{-k} denotes the vector of strategies of all players except player k. Assuming the payoffs are twice differentiable, then player k's payoff will have an increasing difference if for all $j \neq k$,

$$\frac{\partial^2 \Pi_k}{\partial x_k \partial x_j} \geq 0, \tag{10}$$

for all x_k and x_j. The resulting game is said to be *supermodular* if each player's payoff has increasing differences and each player's set of allowable strategies is a compact subset of the real-line. This definition can be generalized to allow for non-differentiable payoffs and vector-valued strategies, see for example [22].

For our purposes, supermodular games have the following two useful properties:

1) If each agent's payoff is upper semi-continuous, the game must have a Nash equilibrium.
2) If the Nash equilibrium is unique, then best response updates globally converge to that Nash equilibrium from any initial choice of strategies.

Here, best response updates mean that players iteratively update their strategies to optimize their payoffs assuming that the other player's do not change their current strategy. For the preceding convergence result, this updating can be done in a general asynchronous manner, which includes synchronous updates as a special case. If the Nash equilibrium is not unique, then best response updates will still converge, provided that agents initialize to either their smallest or largest strategy choices; in these cases the set of Nash equilibria also has a *lattice* structure [22]. This should be contrasted with an arbitrary game in which a Nash equilibria may not exist, and even if a Nash equilibrium exists, best response updates need not converge.

REFERENCES

[1] "Notice of Proposed Rule Making and Order," Federal Communications Commission Report 03-322, Dec. 2003.

[2] C. W. Sung and W. S. Wong, "A Noncooperative Power Control Game for Multirate CDMA Data Networks," *IEEE Transactions on Wireless Communications*, vol. 2, no. 1, pp. 186–194, Jan. 2003.

[3] R. Menon, A. B. MacKenzie, J. E. Hicks, R. M. Buehrer, and J. H. Reed, "A Game-Theoretic Framework for Interference Avoidance," *to appear in IEEE Transactions on Communications*.

[4] T. Heikkinen, "Congestion Based Pricing in a Dynamic Wireless Network," in *Proc. IEEE Vehicular Technology Conference, 2001*, vol. 2, 2001, pp. 1073–1076 vol.2.

[5] R. Etkin, A. Parekh, and D. Tse, "Spectrum Sharing for Unlicensed Bands," *IEEE Journal on Selected Areas in Communications*, vol. 25, no. 3, pp. 517–528, April 2007.

[6] M. Chiang, "Balancing Transport and Physical Layers in Wireless Multihop Networks: Jointly Optimal Congestion Control and Power Control," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 1, pp. 104–116, Jan. 2005.

[7] M. Wiczanowski, S. Stanczak, and H. Boche, "Performance and Interference Control in Wireless Ad hoc and Mesh Networks - A Generalized Lagrangian Approach," *IEEE Transactions on Signal Processing*, vol. 56, no. 8, pp. 4039–4052, Aug. 2008.

[8] G. Scutari, D. Palomar, and S. Barbarossa, "Competitive Design of Multiuser MIMO Systems Based on Game Theory: A Unified View," *IEEE J. Sel. Areas Commun.*, vol. 26, no. 7, pp. 1089–1103, Sep. 2008.

[9] S. Ye and R. S. Blum, "Optimized Signaling for MIMO Interference Systems With Feedback," *IEEE Trans. Signal Process.*, vol. 51, no. 11, pp. 2839–2848, Nov. 2003.

[10] C. Zhou, M. Honig, and S. Jordan, “Utility-Based Power Control for a Two-Cell CDMA Data Network,” *IEEE Transactions on Wireless Communications*, vol. 4, no. 6, pp. 2764–2776, Nov. 2005.

[11] H. Boche and S. Stanczak, “Log-concavity of sir and characterization of feasible sir region for cdma channels,” in *Proc. Thirty-Seventh Asilomar Conference on Signals, Systems and Computers*, vol. 1, Nov. 2003, pp. 413–417 Vol.1.

[12] J. Huang, R. A. Berry, and M. L. Honig, “Distributed Interference Compensation for Wireless Networks,” *IEEE J. Sel. Areas Commun.*, vol. 24, no. 5, pp. 1074–1084, May 2006.

[13] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Englewood Cliffs, NJ: Prentice-Hall, 1989.

[14] J. Yuan and W. Yu, “Distributed Cross-layer Optimization of Wireless Sensor Networks: A Game Theoretic Approach,” in *Proc. IEEE Conf. Global Commun.*, Nov. 2006.

[15] C. Shi, R. A. Berry, and M. L. Honig, “Monotonic Convergence of Distributed Interference Pricing in Wireless Networks,” *submitted to IEEE International Symposium on Information Theory, 2009*.

[16] E. A. Jorswieck, E. G. Larsson, and D. Danev, “Complete Characterization of the Pareto Boundary for the MISO Interference Channel,” *IEEE Trans. Signal Process.*, vol. 56, no. 10, pp. 5292–5296, Oct. 2008.

[17] C. Shi, R. A. Berry, and M. L. Honig, “Distributed Interference Pricing with MISO Channels,” in *Proc. 46th Annual Allerton Conference 2008*, Sep. 2008.

[18] C. Hellings, D. A. Schmidt, and W. Utschick, “Optimized Beamforming for the Two-Stream MIMO Interference Channel at High SNR,” in *Proc. IEEE/ITG Workshop on Smart Antennas 2009*, Feb. 2009.

[19] C. Shi, D. A. Schmidt, R. A. Berry, M. L. Honig, and W. Utschick, “Distributed Interference Pricing for the MIMO Interference Channel,” *to appear in Proc. IEEE International Conference on Communications, 2009*.

[20] J. Huang, R. A. Berry, and M. L. Honig, “Auction-Based Spectrum Sharing,” *Mobile Networks and Applications*, vol. 11, no. 3, pp. 405–408, June 2006.

[21] J. Bae, E. Beigman, R. Berry, M. Honig, and R. Vohra, “Sequential Bandwidth and Power Auctions for Distributed Spectrum Sharing,” *IEEE Journal on Selected Areas in Communications*, vol. 26, no. 7, pp. 1193–1203, September 2008.

[22] D. M. Topkis, *Supermodularity and Complementarity*. Princeton, NJ: Princeton Univ. Press, 1998.