Auction-based Spectrum Sharing

Jianwei Huang* Randall A. Berry† Michael L. Honig‡

Abstract

We study auction mechanisms for sharing spectrum among a group of users, subject to a constraint on the interference temperature at a measurement point. The users access the channel using spread spectrum signaling and so interfere with each other. Each user receives a utility that is a function of the received signal-to-interference plus noise ratio. We propose two auction mechanisms for allocating the received power. The first is an auction in which users are charged for received SINR, which, when combined with logarithmic utilities, leads to a weighted max-min fair SINR allocation. The second is an auction in which users are charged for power, which maximizes the total utility when the bandwidth is large enough and the receivers are co-located. Both auction mechanisms are shown to be socially optimal for a limiting "large system" with co-located receivers, where bandwidth, power and the number of users are increased in fixed proportion. We also formulate an iterative and distributed bid updating algorithm, and specify conditions under which this algorithm converges globally to the Nash equilibrium of the auction.

Index Terms

CDMA, spectrum sharing, power control, game theory

This work was supported by the Northwestern-Motorola Center for Communications and by NSF CAREER award CCR-0238382. This paper was presented in part at *the 2nd Workshop on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt'04), Cambridge, UK, March 24-26, 2004*, and *the 42nd Annual Allerton Conference on Communication, Control and Computing, Monticello, IL, USA, September 29 - October 1, 2004*.

The authors are with the Department of Electrical and Computer Engineering, Northwestern University, Evanston, IL, USA 60208

* Corresponding author: e-mail: jianweih@ece.northwestern.edu, phone: (847) 491-5751, fax: (847) 491-4455

† e-mail: rberry@ece.northwestern.edu, phone: (847) 491-7074, fax: (847) 491-4455.

‡ e-mail: mh@ece.northwestern.edu; phone: (847) 491-7803, fax: (847) 467-3550.

I. INTRODUCTION

There has been growing interest in making more efficient use of spectrum by shifting from the conventional "command-and-control" spectrum usage model to more flexible "Exclusive Use" and "Commons" models (e.g. see [1]). In the Exclusive Use model, the licensee has exclusive rights to the spectrum, but could allow other users to purchase access rights to the spectrum when it is underutilized. In the Commons model, spectrum is unlicensed and an unlimited number of users can share spectrum with usage rights governed by technical standards. In either model, a basic question is how to share the available spectrum efficiently and fairly. A proposed requirement (e.g. see [1]) is that the *interference temperature* in the spectrum band be kept under some threshold, where interference temperature is defined to be the RF power measured at a receiving antenna per unit bandwidth.

In this paper, we study a spectrum allocation problem under such an interference temperature constraint. This model is motivated by the scenario in which users wish to purchase a local, relatively short-term data service. The spectrum to be used may be licensed to an independent entity (e.g., private company) or controlled by a government agency, either of which we refer to as a *manager*. Users may transmit to receivers at different locations, or to co-located receivers at a single access point. In both cases, the manager controls the amount of bandwidth and power assigned to each user in order to keep the interference temperature at a given measurement point below a certain threshold. We assume that all users adopt a spread spectrum signaling format, in which the transmitted power is evenly spread across the entire available band controlled by the manager. This allows efficient multiplexing of data streams from different sources corresponding to different applications, and reduces the combined power-bandwidth allocation problem to a received power allocation problem. Each user has a utility, which is a function of the received Signal-to-Interference plus Noise Ratio (SINR), reflecting his desired Quality of Service (QoS). The interference a user receives depends on the other users' transmission powers and the cross-channel gains, as well as the bandwidth.

In this setting, an interference temperature constraint is equivalent to a constraint on the received power at the measurement point. This allows us to view the received power as a divisible good; we study auction mechanisms for allocating this good. It is well known that a Vickrey-Clarke-Groves (VCG) auction can be used to achieve a socially optimal allocation, i.e., maximize the total utility [2]. However, as discussed in Sect. II-B, this may not be suitable here due to the required information from the users and the computational burden on the manager. Instead, we propose two auctions mechanisms that allocate the received power as a function of bids submitted by the users and the price announced by the manager. We model the resulting problem as a non-cooperative game [2], and characterize the Nash equilibria and related properties of the two auctions. We first analyze these auctions as a simultaneous move game, assuming all information (i.e., utilities and link channel gains) is available to the users (but not to the manager). We subsequently formulate an iterative and fully distributed algorithm, which only requires the users to obtain limited local information in order to converge globally to the Nash equilibrium (NE). This makes the auction mechanisms easily implementable and scalable with the population size.

Our approach is similar to a *share auction* (see [3]–[8] and the references therein), or *divisible auction*, where a perfectly divisible good is split among bidders whose payments depend solely on the bids. A common form of bids in a share auction is for each user to submit his demand curve (e.g., [3]–[5]), i.e., the amount of goods a user desires as a function of the price. The auctioneer can then compute a market clearing price based on the set of demand curves. However, in our problem, a user's demand curve for received power also depends on the demands of other users due to interference. On the other hand, if the demand curves are viewed in terms of SINR so that they are mutually independent, the market clearing price for SINR is not easy to find since the constraint is on received power. To overcome these difficulties, we adopt a signaling system similar to [6]–[8], where users submit one dimensional bids for the resource.

We assume a weighted proportional allocation rule in which a user's power allocation is proportional to his bid. This type of allocation rule has been studied in a wide range of applications (e.g., see [9], [10]) including network resource allocation (e.g., [6]–[8]). Given this allocation, the users participate in a game with the objective of maximizing their own benefit. It is well known that the NE typically does not maximize the total system utility [11]. This has been referred to as the *price of anarchy* (e.g., [6]). In order to achieve a more desirable NE, we allow the manager to announce a unit price (e.g., [12], [13]) either for received SINR (a SINR auction) or received power (a power auction). An SINR auction with logarithmic utilities leads to a weighted max-min fair SINR allocation. A power auction maximizes the total utility for a large enough bandwidth with co-located receivers. Both auctions maximize the total utility in a large enough system with co-located receivers if the total power and bandwidth are increased in fixed proportion to the number of users. Related work on uplink power control for CDMA has appeared in [13]–[16]. A key difference here is that there is a constraint on the total received power at all times[1]. Because of this, a user's interference depends on his own power allocation, which can make the problem non-convex.

We assume the user population is stationary, i.e., the users and their corresponding utilities stay fixed during the time period of interest. On a larger time-scale one can view time divided into periods, during which the number of users and each user's utility are fixed and the proposed auction algorithm is used. When a new period begins, users may join or leave the system. Remaining users may update their utilities to reflect changes in their QoS requirements. For example, a user with data that must be delivered by a deadline might increase his utility (as a function of SINR) as the deadline approaches. Here we do not consider mechanisms and associated dynamics over multiple periods.

The remainder of the paper is organized as follows. After introducing the auction mechanisms in Sect. II, we analyze the performance for a finite system and for a limiting "large system" in Sect. III and IV, respectively. In Sect. V we give an iterative and distributed bid updating algorithm, and show that it converges globally to the unique NE of the auction when one exists. Numerical results are given in Sect. VI and conclusions in Sect. VII. Several of the main proofs are given in the Appendix.

II. Auction Mechanisms

A. System Model

Spectrum with bandwidth B is to be shared among M spread spectrum users, where a user refers to a transmitter and an intended receiver pair. User $i's$ valuation of the spectrum is characterized by a utility $U_i(\gamma_i)$, where γ_i is the received SINR at user i's receiver. We primarily consider the case where each user's utility is given by $U_i(\gamma_i) = U(\theta_i, \gamma_i)$, where θ_i is a user-dependent parameter. As a particular example, we consider the *logarithmic utility* $U_i(\gamma_i) = \theta_i \ln(\gamma_i)$.[2]

[1]We assume that any transmission power constraint for each user is large enough so that it can be ignored.

[2]This approximates the weighted rate of user i in the high SINR regime.

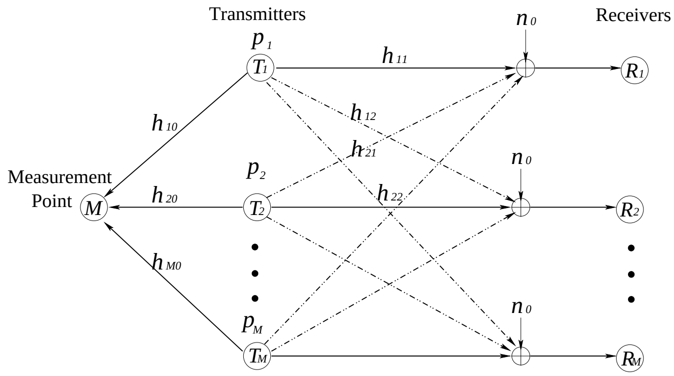

Fig. 1. System model for M transmitter-receiver pairs

Assumption 1: For each user i, $U_i(\gamma_i)$ is increasing, strictly concave, and twice continuously differentiable in γ_i.

Utilities that satisfy this assumption are commonly used to model "elastic" data applications [17]. For each i, the received SINR is given by

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

where p_i is user i's transmission power, h_{ij} is the channel gain from user i's transmitter to user j's receiver, and n_0 is the background noise power that is assumed to be the same for all users. To satisfy an interference temperature constraint, the total received power at a specified measurement point must satisfy

$$\sum_{i=1}^{M} p_i h_{i0} \leq P, \tag{2}$$

where h_{i0} is the channel gain from user i's transmitter to the measurement point. The system model is shown in Fig. 1. A power allocation is *Pareto optimal* if no user's utility can be increased without decreasing another user's utility.

Lemma 1: A power allocation scheme is Pareto optimal if and only if the total received power constraint is tight, i.e., $\sum_{i=1}^{M} p_i h_{i0} = P$.

This follows because if the power constraint is not tight, then each user can increase their power by a factor of $P/\sum_{i=1}^{M} p_i h_{i0}$, which increases the SINR for every user. Lemma 1 does not require Assumption 1; in particular, $U_i(\gamma_i)$ does not have to be concave in γ_i, although it must be strictly increasing. Note that Pareto optimality does not indicate how to split resources among users, only that the resource should be fully utilized. A stronger condition is *social optimality*, where the total utility $\sum_{i=1}^{M} U_i(\gamma_i)$ is maximized. Social optimality implies Pareto optimality, but the reverse is not true. Therefore, to achieve social optimality, the manager should always ensure that the received power constraint is tight.

A special case, on which we will focus, is when the receivers are co-located with the measurement point. This could model a situation where a service provider purchases the spectrum usage rights from the manager and provides service from a single access point. In this case, $h_{ij} = h_{i0}$ for all $i, j \in \{1, ..., M\}$, and we denote user i's received power as $p_i^r = p_i h_{i0}$. In a Pareto optimal allocation for this co-located receiver case, we have for each user i,

$$\gamma_i \equiv \gamma_i(p_i^r) = \frac{p_i^r}{n_0 + \frac{1}{B}(P - p_i^r)},$$

so that user i's utility $U_i(\gamma_i(p_i^r))$ under a Pareto optimal allocation does not depend on how the power is allocated among the interferers.

We assume that each user's utility is private information, i.e., only known to the user himself. The manager must then devise a mechanism for allocating power without having this knowledge *a priori*. Also the manager may not have *a priori* knowledge of the channel gains, h_{ij}'s. One such mechanism is the generalized VCG auction.

B. VCG Auction for Spectrum Sharing

A VCG auction results in a socially optimal outcome, and it is a (weakly) dominant strategy for users to bid truthfully (i.e., state their true utilities). In our context, a VCG auction can be described as follows: First, users are asked to submit their utilities $\{U_i(\gamma_i)\}$. The manager then computes the power allocation $\boldsymbol{p}^* = (p_1^*, \cdots, p_M^*)$ that maximizes the total utility, i.e., $U_{\max} = \sum_{j=1}^{M} U_j(\gamma_j(\boldsymbol{p}^*))$, given the received power constraint, and allocates power to the users accordingly. Furthermore, the manager computes the maximum total utility if user i is excluded from the auction, i.e., $U_{\max/i} = \max_{\{p_j\}/p_i} \sum_{j \neq i} U_j(\gamma_j)$ for each $i \in \mathcal{M}$. In total, the manager must solve $M+1$ optimization problems. The manager then charges user m the amount $U_{\max/i} - \sum_{j \neq i} U_j(\gamma_j(\boldsymbol{p}^*))$, which is the decrement in sum utility over all other users from including user i in the auction.

The VCG auction may not be suitable in this context for several reasons: (i) In order to completely specify the users' utilities, in particular, the SINR in (1), for each user i, the channel gains h_{ij} for all $i, j \in \{1, ..., M\}$ must be measured by the users and reported to the manager. This might be a heavy burden for the users in a large network. (ii) The manager must solve $M+1$ optimization problems, which are typically non-convex due to the interference. This becomes computationally expensive for large M, and may not be suitable for online allocations. For these reasons, we examine mechanisms that require less information exchange and less computation for the manager.

C. One-Dimensional Auctions with Pricing

We now describe two auctions (SINR- and power-based) in which users submit one-dimensional bids representing their willingness to pay, and the manager simply allocates the received power in proportion to the bids. The users then pay an amount proportional to their SINR (or power). The manager announces a nonnegative reserve bid β, and uses a corresponding reserve power that interferes with the other users. In contrast with the situation where the manager submits a reserve bid to extract more revenue from the other bidders [18], here the main purpose of the reserve bid is to guarantee a unique desirable outcome of the auction. We will show that the interference generated by the manager can be made arbitrarily small. Although the two auctions are relatively simple, we show that under some mild conditions they give power allocations that are arbitrarily close to the allocation from a VCG auction.

Regarding the information structure of the auction, we first assume that it is a complete information game, i.e., all users' utilities and all channel gains are known to all users. In Sect. V, we present a distributed algorithm that can achieve the NE of the auction with limited information, where each user i only needs to measure the background noise density n_0, the channel gain ratio $\hat{h}_{ii} = h_{ii}/h_{i0}$ and the SINR at his own receiver.

Simultaneous Auction Algorithm:

1) The manager announces a reserve bid $\beta \geq 0$, and a price $\pi^s > 0$ (in an SINR auction) or $\pi^p > 0$ (in a power auction).
2) After observing β, π^s (or π^p), user $i \in \{1, ..., M\}$ submits a bid $b_i \geq 0$.
3) The manager keeps reserve power p_0, and allocates to each user i a transmission power

p_i so that the received power at the measurement point is proportional to the bids, i.e.,

$$p_i h_{i0} = \frac{b_i}{\sum_{j=1}^{M} b_i + \beta} P, \text{ and } p_0 = \frac{\beta}{\sum_{j=1}^{M} b_i + \beta} P. \tag{3}$$

The resulting SINR for user i is

$$\gamma_i = \frac{p_i h_{ii}}{n_0 + \frac{1}{B}\left(\sum_{j\neq i} p_j h_{ji} + p_0 h_{0i}\right)}, \tag{4}$$

where h_{0i} is the channel gain from the manager (measurement point) to user i's receiver[3]. If $\sum_{i=1}^{M} b_i + \beta = 0$, then $p_i = 0$.

4) In an SINR (power) auction, user i pays $C_i = \pi^s \gamma_i$ $(C_i = \pi^p p_i h_{i0})$

A *bidding profile* is the vector containing the users' bids $\boldsymbol{b} = (b_1, ..., b_M)$. The *bidding profile of user i's opponents* is defined as $b_{-i} = (b_1, ..., b_{i-1}, b_{i+1}, ..., b_M)$, so that $\boldsymbol{b} = (b_i; b_{-i})$. In the preceding auctions, each user i submits a bid b_i to maximize his *surplus function*

$$S_i\left(b_i; b_{-i}\right) = U_i\left(\gamma_i\left(b_i; b_{-i}\right)\right) - C_i.$$

Here we omit the dependence on β and π.

An NE of the auction is associated with a bidding profile $\boldsymbol{b}^*$ such that $S_i(b_i^*; b_{-i}^*) \geq S_i(b_i'; b_{-i}^*)$ for any $b_i' \in [0, \infty)$ and any user i. Define user i's *best response* given b_{-i} as the set

$$\mathcal{B}_i\left(b_{-i}\right) = \left\{ \hat{b}_i \mid \hat{b}_i = \arg \max_{b_i \in [0,\infty)} S_i\left(b_i; b_{-i}\right) \right\},$$

i.e., the set of b_i's that maximize $S_i(b_i; b_{-i})$ given a fixed b_{-i}.[4] The NE bidding profile $\boldsymbol{b}^*$ is a fixed point, i.e., no user has the incentive to deviate unilaterally. The existence and uniqueness of an NE are shown in the following to depend on β and π^s (or π^p).

These auction mechanisms differ from some previously proposed auction-based network resource allocation schemes (e.g., [6], [7]) in that the bids here are not the same as the payments. Instead, the bids are signals of willingness to pay. The manager can therefore influence the NE by choosing β and π^s (or π^p). This alleviates the typical inefficiency of the NE, and allows us to reach Pareto optimal, and in some cases, socially optimal solutions.

[3]If $h_{0i} = 0$ for all $i \in \{1, ..., M\}$, then the manager does not interfere with the users and many of the results in the following section still hold. However, in the co-located case, we have $h_{0i} = 1$ for all i.

[4]In general the best response set may contain more than one element.

III. Finite System Analysis

A. SINR Auction

In this case, $C_i = \pi^s \gamma_i = \pi^s \frac{p_i h_{ii}}{n_0 + \frac{1}{B}\left(\sum_{j\neq i}^{M} p_j h_{ji} + p_0 h_{0i}\right)}$, so that each user's payment depends on both the transmission power and the interference.

Theorem 1: In an SINR auction:

(1) For $\beta > 0$, there exists a threshold price $\pi^s_{th} > 0$ such that a unique NE exists if $\pi^s > \pi^s_{th}$, and there is no NE if $\pi^s \leq \pi^s_{th}$.

(2) For $\beta = 0$, one of the following is true: (i) there is a unique NE with $b_i^* = 0$ for all i, (ii) there are an infinite number of Nash Equilibria, or (iii) there is no NE.

The proof is given in Appendix A; as shown there, when $\beta > 0$ and $\pi^s > \pi^s_{th}$, the best response for each user is unique, and the vector of best responses across users is given by

$$\mathcal{B}\left(\boldsymbol{b}\right) = \mathbf{K}\boldsymbol{b} + \boldsymbol{k}_0\beta, \tag{5}$$

where $\mathbf{K} = \left[k_{ij}\left(\pi^s\right)\right]_{i,j\in\{1,\ldots M\}}$ is a nonnegative matrix with $k_{ii}\left(\pi^s\right) = 0$ for all i and

$$k_{ij}\left(\pi^s\right) = \frac{g_i\left(\pi^s\right)\left(n_0 B + P\hat{h}_{ji}\right)}{PB\hat{h}_{ii} - g_i\left(\pi^s\right) n_o B} \geq 0, \forall j \neq i, \tag{6}$$

vector $\boldsymbol{k}_0 = \left(k_{10}, ..., k_{M0}\right)$ has nonnegative elements

$$k_{i0}\left(\pi^s\right) = \frac{g_i\left(\pi^s\right)\left(n_0 B + P h_{0i}\right)}{PB\hat{h}_{ii} - g_i\left(\pi^s\right) n_o B} \geq 0, \tag{7}$$

and $g_i\left(\pi^s\right)$ is a nonnegative and continuously nonincreasing function defined as

$$g_i\left(\pi^s\right) = \begin{cases} \infty, & 0 \leq \pi^s \leq U_i'\left(\infty\right), \\ U_i'^{-1}\left(\pi^s\right), & U_i'\left(\infty\right) < \pi^s < U_i'\left(0\right), \\ 0, & U_i'\left(0\right) \leq \pi^s. \end{cases} \tag{8}$$

The spectral radius of matrix $\mathbf{K}$, ρ_K, satisfies $0 \leq \rho_K < 1$. The unique NE is

$$\boldsymbol{b}^* = \left(\mathbf{I} - \mathbf{K}\right)^{-1} \boldsymbol{k}_0\beta = \sum_{n=0}^{\infty} \mathbf{K}^n \boldsymbol{k}_0\beta.$$

where $\mathbf{I}$ is the identity matrix.

Since we would like to avoid case (2) in Theorem 1, we assume $\beta > 0$ in the rest of the paper. Notice that the value of β does not affect the power allocation at the NE, since all equilibrium bids are proportional to β. Thus the manager only needs to announce an arbitrary $\beta > 0$. In general, π_{th}^{s} in Theorem 1 is difficult to find analytically. However, in the co-located receiver case with logarithmic utilities, we have a closed-form relation between π_{th}^{s} and the users' utility parameters. For $i \in \{1, ..., M\}$, define

$$k_i\left(\pi^s\right) = \frac{g_i\left(\pi^s\right)\left(P + Bn_0\right)}{PB - g_i\left(\pi^s\right)n_0 B}. \tag{9}$$

Proposition 1: In an SINR auction with co-located receivers and logarithmic utilities, $k_i\left(\pi_{th}^{s}\right) \geq 0$ for each user i and $\sum_{i=1}^{M} k_i\left(\pi_{th}^{s}\right) / \left(1 + k_i\left(\pi_{th}^{s}\right)\right) = 1$.

This follows from the proof of Theorem 1 by using the fact that with co-located receivers $k_{il}\left(\pi^s\right) = k_i\left(\pi^s\right)$ for all $l \in \{0, ..., M\}$, and explicitly solving for the NE. The bidding and power profiles at the NE are:

$$b_i^* = \frac{\frac{k_i(\pi^s)}{1+k_i(\pi^s)}}{1 - \sum_{j=1}^{M} \frac{k_j(\pi^s)}{1+k_j(\pi^s)}}\beta \ \text{ and } \ p_i^* = \frac{k_i\left(\pi^s\right)}{1 + k_i\left(\pi^s\right)}P \quad \text{for } i \in \{1, ..., M\}. \tag{10}$$

Given the existence of a unique NE, we next characterize the resulting resource allocation. We say an allocation $\{x_i\}_{i\in\{1,...,M\}}$ is *weighted max-min fair* with weights $\{w_i\}_{i\in\{1,...,M\}}$ if for each $i \in \{1, ..., M\}$, x_i can not be increased without decreasing some x_j, $j \in \{1, ..., M\}$, for which $x_j/w_j \leq x_i/w_i$.

Proposition 2: If a unique NE exists in an SINR auction with logarithmic utilities, the SINR allocation $\{\gamma_i^*\}_{i\in\{1,...,M\}}$ are weighted max-min fair with the weights $\{\theta_i\}_{i\in\{1,...,M\}}$ given a fixed reserve power p_0^*, and the payments $\{C_i^*\}_{i\in\{1,...,M\}}$ are proportional with the same weights.

Proof: User i's unique best response satisfies

$$\frac{\partial U_i\left(\gamma_i\left(\mathcal{B}_i\left(b_{-i}\right); b_{-i}\right)\right)}{\partial \gamma_i\left(\mathcal{B}_i\left(b_{-i}\right); b_{-i}\right)} = \frac{\theta_i}{\gamma_i\left(\mathcal{B}_i\left(b_{-i}\right); b_{-i}\right)} = \pi^s,$$

i.e., $\gamma_i^*/\theta_i = 1/\pi^s$ for all i. Clearly, no user's SINR can be increased without decreasing another user's SINR. User i's payment satisfies $C_i^*/\theta_i = \left(\pi^s \gamma_i^*\right)/\theta_i = 1$. ■

In [19], Kelly et al. consider an algorithm for rate allocation in a wire-line network with logarithmic utilities $w_i \log\left(x_i\right)$ for all users $i \in \{1, ..., M\}$. In that case, the socially optimal rate allocation $\{x_i\}_{i\in\{1,...,M\}}$ is *weighted proportional fair* with weights $\{w_i\}_{i\in\{1,...,M\}}$, i.e., for any other feasible rate allocation $\{x_i'\}_{i\in\{1,...,M\}}$, $\sum_{i=1}^{M} w_i\left(x_i' - x_i\right)/x_i \leq 0$. Their utility maximization problem is convex and separable since there is no externality (i.e., interference) among different users. Here, due to the interference among users, the problem is generally not separable (except in the co-located receiver case) and is typically not convex; thus the allocation achieved by the SINR auction with logarithmic utilities typically is not socially optimal or proportional fair.[5]

In a system with a unique NE, define the *system usage efficiency* by

$$\eta = \frac{\sum_{i=1}^{M} p_i^* h_{i0}}{P} = \frac{\sum_{i=1}^{M} b_i^*}{\sum_{i=1}^{M} b_i^* + \beta}.$$

For Pareto optimality $\eta = 1$, but the necessary condition for the uniqueness of a unique NE is $\eta < 1$ due to the required positive reserve bid β, i.e., Pareto optimality and a unique NE are conflicting objectives[6].

We define an ε-*system* as one with parameters $\left(P^\varepsilon, B^\varepsilon, M^\varepsilon, n_0^\varepsilon\right) = \left(P\left(1-\varepsilon\right), B, M, n_0 + \varepsilon P/B\right)$, where $\varepsilon \in (0,1)$. An ε-*Pareto optimal* allocation is defined as a Pareto optimal solution for the ε-system.

Proposition 3: In an SINR auction, there exists a price π^s for any $\varepsilon \in (0,1)$, such that the system has a unique NE and achieves an ε-Pareto optimal solution (i.e., $\eta = 1-\varepsilon$ in the original system).

Proof: From the proof of Theorem 1, it can be seen that as π^s increases from π_{th}^s to ∞, $\rho_K\left(\pi^s\right)$ decreases from 1 to 0, and is continuous and nonincreasing in the interval. Also, the bidding profile $\boldsymbol{b}^* = \left(\sum_{n=0}^{\infty} \mathbf{K}^n\right) \boldsymbol{k}_0 \beta$ changes from ∞ (for at least one user's bid) to 0 (for all users' bids), and is also continuous and nonincreasing in the interval. This implies the same for the summation $\sum_{i=1}^{M} b_i^*$, which means $\eta = \sum_{i=1}^{M} b_i^* / \left(\sum_{i=1}^{M} b_i^* + \beta\right)$ decreases from 1 to 0, and is continuous and nonincreasing in the interval. So there must exist a price $\pi^s \in \left(\pi_{th}^s, \infty\right)$ that achieves any $\eta = 1-\varepsilon \in (0,1)$. ∎

In practice, the manager can achieve a target η^* by adjusting π^s after observing the usage efficiency at the current NE: if it is too low, the price should be decreased. Note that if the price is decreased too much, there may not be an unique NE.

B. Power Auction

In this case $C_i = \pi^p p_i h_{i0}$. For the co-located receiver case with logarithmic utilities, Proposition 1 still holds, but with a different expression for $k_i\left(\pi^s\right)$ than that given in (9). The bidding

[5] Moreover, in this setting the socially optimal allocation with logarithmic utilities is not proportional fair.

[6] Here we are not including power used by the manager in our definition of Pareto optimality.

and power profiles at the NE are again given by (10), but it may be impossible to find a price $\pi^{p\varepsilon}$ that gives an arbitrary $\eta = 1 - \varepsilon$. This is because $U_i\left(\gamma_i\left(p_i^r\right)\right)$ is not always concave in the received power p_i^r, and so the p_i^r that maximizes user i's surplus may not be continuous with price π^p, i.e., it may jump from one local optimum to the other. As a result, $\eta = \sum_{i=1}^{M} p_i^r/P$ may be discontinuous at some values of π^p.

We say that a power allocation is *ε-socially optimal* if it maximizes the total utility of the ε-system. In the case of co-located receivers, the power auction can achieve an *ε-socially optimal* allocation for a more general class of utilities.

Assumption 2: For each $i \in \{1, ..., M\}$, $U_i\left(\gamma_i\right)$ satisfies Assumption 1 and

$$\frac{\left|U_i''\left(\gamma_i\right)\right|}{U_i'\left(\gamma_i\right)}\left(\gamma_i + B\right) > 2 \tag{11}$$

for any $\gamma_i \in [0, P/n_0]$.

Inequality (11) follows from setting $\partial^2 U_i\left(\gamma_i\left(p_i^r\right)\right)/\partial^2 p_i^r < 0$ for any $p_i^r \in [0, P/h_{i0}]$, i.e., the utility is strictly concave in the received power. For the case of logarithmic utilities, Assumption 2 is satisfied if $P/\left(Bn_0\right) < 0$ dB. For many utilities (e.g., $\theta_i \log\left(1+\gamma_i\right)$, $1 - e^{-\theta_i \gamma_i}$, and $\theta_i \gamma_i^{\alpha}$ $\left(\alpha \in (0,1)\right)$), Assumption 2 is satisfied when the bandwidth is large enough, so that the interference among users is relatively small.

Theorem 2: In a power auction with co-located receivers and Assumption 2, for any $\varepsilon \in (0,1)$ there exists a price $\pi^{p\varepsilon}$ such that the system has a unique NE, and the NE achieves ε-social optimality.

The proof is given in Appendix B. Theorem 2 implies that with large enough bandwidth, so that the externality effects among users are relatively small, the power auction with co-located receivers can achieve an allocation that is arbitrarily close to that produced by a VCG auction, and so is preferable to the SINR auction in terms of social optimality. When Assumption 2 is not satisfied, the power auction may not be able to achieve an η close to 1 (e.g., with logarithmic utilities); this can result in a lower total utility compared to the SINR auction, which can achieve any η.

C. Revenue Comparison between SINR and Power Auctions

From the manager's point of view, revenue maximization might be another important objective. Here we restrict our discussion to the two auctions previously described for co-located receivers.[7] Let R^p and R^s be the revenue derived from the power and SINR auctions, respectively. We first consider the case where users are identical (i.e., have the same utilities) and the utilities are concave in power.

Theorem 3: Given co-located receivers, identical utilities, and Assumption 2, suppose further that both auctions achieve the same system usage efficiency η. Then $R^p > R^s$, and $R^p/R^s \to 1$ as $M \to \infty$.

Proof: With identical utilities and same efficiency η, both auctions allocate the same received power p^{r*} to all users. Let $U\left(\gamma\left(p^r\right)\right) = U_i\left(\gamma_i\left(p_i^r\right)\right)$ for $1 \le i \le M$. From the first-order conditions for surplus maximization,

$$\pi^p = U'\left(\gamma\left(p^r\right)\right)\gamma'\left(p^r\right)|_{p^r=p^{r*}} \text{ and } \pi^s = U'\left(\gamma\left(p^r\right)\right)|_{p^r=p^{r*}} \tag{12}$$

so that

$$\frac{R^p}{R^s} = \frac{M\pi^p p^{r*}}{M\pi^s\gamma\left(p^{r*}\right)} = \frac{U'\left(\gamma\left(p^r\right)\right)\gamma'\left(p^r\right)|_{p^r=p^{r*}}p^{r*}}{U'\left(\gamma\left(p^r\right)\right)|_{p^r=p^{r*}}\gamma\left(p^{r*}\right)} = \frac{\frac{B(n_0B+P)}{(n_0B+P-p^{r*})^2}p^{r*}}{\frac{p^{r*}B}{n_0B+P-p^{r*}}} = \frac{n_0B+P}{n_0B+P-p^{r*}} > 1.$$

As $M \to \infty$, $p^{r*} \to 0$, and so $R^p/R^s \to 1$. ■

When Assumption 2 is not satisfied, the power auction may collect less revenue than the SINR auction, since the former might not be able to achieve η close to 1. However, for logarithmic utilities the relation between the revenues remains the same.

Proposition 4: Given co-located receivers with logarithmic utilities, assume there exists a $\bar{\theta}$ such that $\theta_i \le \bar{\theta}$ for $1 \le i \le M$. Then $R^p > R^s$ and $R^p/R^s \to 1$ as $M \to \infty$.

The proof is given in Appendix C. Notice that in Proposition 4 we do not require identical utilities or the same η in both auctions. Hence with logarithmic utilities the power auction always generates more revenue.

IV. Large System Analysis

In this section we consider the asymptotic behavior as P, B, M and β go to infinity, while keeping P/M, B/M and β/M fixed. We focus on co-located receivers and assume that each

[7] We note that other auction mechanisms may extract more revenue.

user i's utility parameter θ_i is independently chosen according to a continuous probability density $f(\theta)$ over $\left[\underline{\theta}, \bar{\theta}\right]$, where $0 \leq \underline{\theta} < \bar{\theta} < \infty$. The expected value of θ is denoted as $E[\theta]$.

Proposition 5: In an SINR auction with logarithmic utilities and co-located receivers, a unique NE exists in the large system limit if and only if

$$\pi^s > \pi^s_{th} = E[\theta]\left(n_0 + P/B\right)\frac{M}{P}. \tag{13}$$

In this case, the power and SINR allocations at the NE are weighted max-min fair with weights $\{\theta_i\}_{1 \leq i \leq M}$, and user i pays θ_i. If condition (13) is not satisfied, no NE exists.

The proof is given in Appendix D. The system usage efficiency at the NE is $\eta = \frac{E[\theta](n_0+P/B)}{\pi^s P/M}$. As $\eta \to 1$, the price π^s converges to π^s_{th}, which is proportional to the system load M/P. This coincides with the *congestion pricing scheme* proposed in [16], where the equilibrium price reflects the system congestion.

In the limiting system with co-located receivers, all users receive the same fixed noise plus interference level $(n_0 + P/B)$ at the NE, because each user gets a negligible proportion of the total power. This makes the SINR and power auctions equivalent if $\pi^s = (n_0 + P/B)\,\pi^p$. The socially optimal allocation maximizes the average utility per user. (Note that the total utility is infinite.)

Assumption 3: The utility $U(\theta, \gamma)$ is asymptotically sublinear with respect to γ, i.e.,

$$\lim_{\gamma \to \infty} \frac{1}{\gamma} U(\theta, \gamma) = 0, \ \forall \theta.$$

Theorem 4: In the limiting system with co-located receivers, if $U(\theta, \gamma)$ satisfies Assumptions 1 and 3, then both the SINR and power auctions can achieve ε-social optimality for any $\varepsilon \in (0, 1)$.

A sketch of the proof is given in Appendix E.[8] Assumption 3 is valid for common utilities, e.g., $\theta \ln(\gamma)$, $\theta \ln(1+\gamma)$, and $\theta\gamma^\alpha$ for any $\alpha \in (0,1)$, and any upper-bounded utility. Under this assumption, even if a finite number of users are allocated non-negligible proportions of the total power, their contributions to the average utility become negligible as the number of users increases. Because of this, the socially optimal allocation gives each user finite power, and so each user sees the same interference level $(n_0 + P/B)$. In that case, both auctions can achieve results that are arbitrarily close to that of a VCG auction.

[8] Theorem 4 can be generalized to the case of a non-collocated measurement point. Here we consider only the co-located case to simplify the proof.

V. ITERATIVE AND DISTRIBUTED BID UPDATING ALGORITHM

In Sect. II, we assumed that the users' utility functions and all the channel gains are public information, so that the auction can be analyzed as a simultaneous-move game with complete information. In practice, the users' utilities are likely to be private information, and it is difficult for user i to measure the channel gains associated with other users, i.e., h_{jk} for $j, k \neq i$. In that case, users cannot find the NE of the auction in one iteration. Next, we present an iterative and fully distributed algorithm that converges to the NE of the SINR auctions[9].

Suppose users update their bids according to the best response (5) simultaneously in iterations $t = 1, 2, \cdots$, i.e.,

$$\boldsymbol{b}^{(t)} = \mathbf{K}\boldsymbol{b}^{(t-1)} + \boldsymbol{k}_0\beta, \tag{14}$$

where $\boldsymbol{b}^{(0)}$ is an arbitrarily chosen feasible initial bidding profile.

Proposition 6: If there exists a unique NE in the SINR auction, then the update algorithm (14) globally converges to the NE from any positive $\boldsymbol{b}^{(0)}$.

Proof: For a unique NE we must have $\mathbf{K} \geq \mathbf{0}$ (component-wise), $\boldsymbol{k}_0 \geq \mathbf{0}$ and $\rho_K < 1$. Under this conditions iterating (14) gives

$$\lim_{t\to\infty} \boldsymbol{b}^{(t)} = \lim_{t\to\infty} \left[\mathbf{K}^t\right] \boldsymbol{b}^{(0)} + \lim_{t\to\infty} \left[\sum_{n=0}^{t-1} \mathbf{K}^n\right] (\boldsymbol{k}_0\beta) = \left(\mathbf{I} - \mathbf{K}\right)^{-1} \boldsymbol{k}_0\beta,$$

which is the unique NE. ■

Next, we show that (14) can be equivalently written in a distributed fashion, where each user only needs to measure the channel gain $\hat{h}_{ii} = h_{ii}/h_{i0}$, the background noise density n_0, and his received SINR $\gamma_i^{(t)}$ in each iteration t.

Proposition 7: In the SINR auction, (14) is equivalent to the following distributed updating algorithm for each user i in iteration $t = 1, 2, ...$

$$b_i^{(t)} = \begin{cases} \dfrac{g_i\left(\pi^s\right) P\hat{h}_{ii} - g_i\left(\pi^s\right)\gamma_i^{(t-1)} n_0}{\gamma_i^{(t-1)} P\hat{h}_{ii} - g_i\left(\pi^s\right)\gamma_i^{(t-1)} n_0} b_i^{(t-1)}, & \text{if } \gamma_i^{(t-1)} > 0, \\ 0, & \text{if } \gamma_i^{(t-1)} = 0, \end{cases} \tag{15}$$

with an arbitrary *positive* initial profile $\boldsymbol{b}^{(0)}$.

[9]Note that here we are still referring to the NE of the simultaneous move game as in Sect. II-C.

Proof: From the proof of Theorem 1, we know that by following the best response (14) in iteration t, each user i submits a bid $b_i^{(t)}$ in an attempt to achieve $\gamma_i\left(b_i^{(t)}; b_{-i}^{(t-1)}\right) = g_i\left(\pi^s\right)$, which maximizes his surplus during iteration t assuming the other bids are fixed at $b_{-i}^{(t-1)}$. Using (3) and (4), we have

$$b_i^{(t)} = \frac{g_i\left(\pi^s\right)\left(n_0\left(\sum_{j\neq i} b_j^{(t-1)} + \beta\right) + (P/B)\left(\sum_{j\neq i} b_j^{(t-1)}\hat{h}_{ji} + \beta h_{0i}\right)\right)}{P\hat{h}_{ii} - g_i\left(\pi^s\right) n_0}. \tag{16}$$

Again using (3) and (4) for the SINR at iteration $t-1$, we have

$$n_0\left(\sum_{j\neq i}^{M} b_j^{(t-1)} + \beta\right) + (P/B)\left(\sum_{j\neq i}^{M} b_j^{(t-1)}\hat{h}_{ji} + \beta h_{0i}\right) = b_i^{(t-1)}\left(P\hat{h}_{ii} - \gamma_i^{(t-1)} n_0\right)/\gamma_i^{(t-1)} \tag{17}$$

if $\gamma_i^{(t-1)} > 0$. By substituting this into (16) and noticing the fact that $\gamma_i^{(t-1)} = 0$ if and only if $b_i^{(t-1)} = 0$, we get the desired result. ∎

The update (15) requires only that user i measure $\hat{h}_{ii}$. There is no need to know the other users' bids. This makes the algorithm distributed and scalable.

The update (14) is similar to the Parallel Update Algorithm in [20] where users update their bids via a myopic strategy. Unlike Fig. 2 in [20], here the sequence of bids does not oscillate if each user i chooses an initial bid $b_i^{(0)}$ that is very small (close to zero). This is due to the nonnegativity of the matrix $\mathbf{K}$. Intuitively, this is because the users' best responses have "strategic complementarity" [21] – roughly, this means when one user submits a higher bid, the others want to do the same. In that case, gradient-based or random updates do not improve convergence.

The update (14) is mathematically similar to the power control algorithm proposed in [22] (see also [23], [24]) for a cellular network, where users adjust their powers (without any power constraints) to meet some preset target SINRs. In those papers, the matrix $\mathbf{K}$ depends only on the channel gains and the target SINRs, and so may not satisfy $\rho_K < 1$ (in which case there would not be a feasible allocation). There are several key differences between (14) and the algorithm in [22]: (1) We consider elastic data traffic without a preset target SINR; (2) We have a total received power constraint; (3) We use the algorithm to adjust bids instead of the power itself; and (4) We can adjust the price so that a unique NE always exists. The mathematical similarity arises from the fact that by designing appropriate auction mechanisms, we convert the constrained power allocation problem into an unconstrained non-cooperative game, in which each user updates his bid in an attempt to reach the desired equilibrium SINR level.

In practice, we would like to guarantee a unique NE, which requires $\pi^s > \pi^s_{th}$, and to achieve high efficiency η, which requires that π^s be close to π^s_{th}, without knowing the exact value of π^s_{th}. The manager must adaptively search for a suitable price. In our simulations, we use the following search method:

1) Initialization: Set $(\underline{\pi}, \overline{\pi}) = (0, \infty)$; choose an arbitrary initial price $\pi^{(0)} > 0$, and a maximum number of iterations T. Set $n = 0$.
2) Start the auction at price $\pi^{(n)}$, set $n = n + 1$.
   a) If the auction does not converge within T iterations, then stop. Let $\underline{\pi} = \pi^{(n-1)}$. If $\overline{\pi} = \infty$, set $\pi^{(n)} = 2\pi^{(n-1)}$; otherwise, set $\pi^{(n)} = (\underline{\pi} + \overline{\pi})/2$. Go to 2.
   b) If the auction converges within T iterations with $\eta < \eta^*$, then set $\overline{\pi} = \pi^{(n-1)}$ and $\pi^{(n)} = (\underline{\pi} + \overline{\pi})/2$. Go to 2.
   c) If the auction converges within T iterations with $\eta \geq \eta^*$, then stop.

Although we only discuss SINR auctions with logarithmic utilities, the bid updating algorithm also works for a power auction with co-located receivers and logarithmic utilities, as well as some other utilities such as $U_i(\gamma_i) = \theta_i \log(1 + \gamma_i)$. [10]

VI. Numerical Results

We first present some numerical results with logarithmic utilities and co-located receivers. In these simulations, $\{\theta_i\}_{i \in \{1,\ldots,M\}}$ are independently and uniformly distributed in $[1, 100]$. Each graph represents an average over 100 independent realizations.

Figures 2 and 3 show average utility per user for the two auctions along with the socially optimal allocation. In both auctions, we set the prices so that η is close to 1. From Theorem 2, the power auction achieves social optimality for $P/(Bn_0) < 0$ dB. Figure 2 shows that the difference in utilities achieved by the two auctions is negligible in this regime. For $P/(Bn_0) > 0$ dB, the utility is not concave with power, and the SINR auction achieves a higher utility higher than the power auction. In Fig. 3, we scale the system as in Sect. IV, and choose $P/(Bn_0) = 20$ dB so that the utility is not concave in power. When $M \leq 14$, the auctions do not achieve the socially optimal solution. For large M, the utilities for both auctions and the socially optimal solution converge to a constant. For this example, the asymptotic behavior is accurate for $M \geq 14$.

[10] Again, we note that in some cases a target η^* may not be achievable in the power auctions.

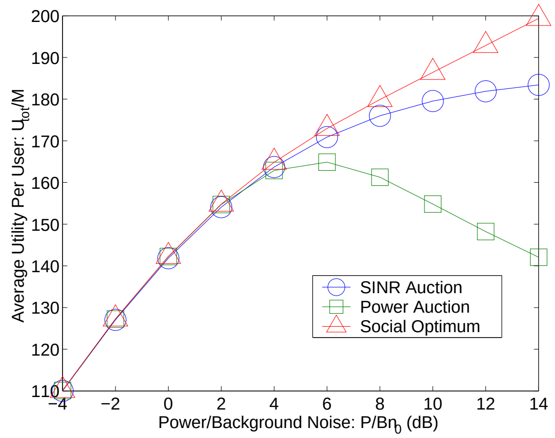

Fig. 2. Utility comparisons in a finite system of users with logarithmic utilities and co-located receivers: $(Bn_0, M) = (10^2, 4)$

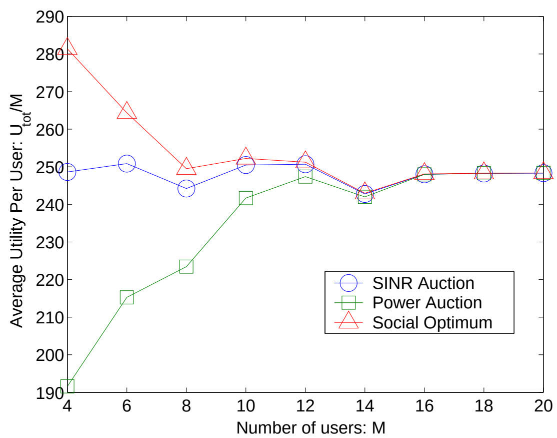

Fig. 3. Utility comparisons in the large system limit of users with logarithmic utilities and co-located receivers: $(P/n_0, B) = (10^4 M, 10^2 M)$

Figures 4 and 5 show the performance of the distributed bid updating algorithm. Figure 4 shows the users' bids starting from very small initial bids and monotonically converging to the unique NE bids. Figure 5 shows the performance of the updating algorithm as the system is scaled. The target system usage efficiency η^* is chosen to be 0.90, 0.95 and 0.98, respectively. We can see that the number of iterations needed for convergence increases with M and approaches a constant when M is large (i.e., $M > 20$). This shows that the algorithm scales well with the system size. The figure also shows that the number of iterations needed for convergence increases with η^*, implying that fast convergence and high system usage efficiency are generally conflicting objectives.

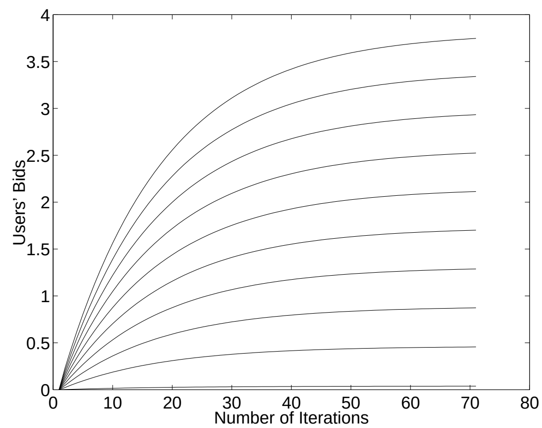

Fig. 4. Performance of the myopic bid updating algorithm with logarithmic utilities and co-located receivers: bids for each user vs. iterations for a finite system with $(P/n_0, B, M, \beta) = \left(10^2, 10^3, 10, 1\right)$ and $\eta^* = 0.95$

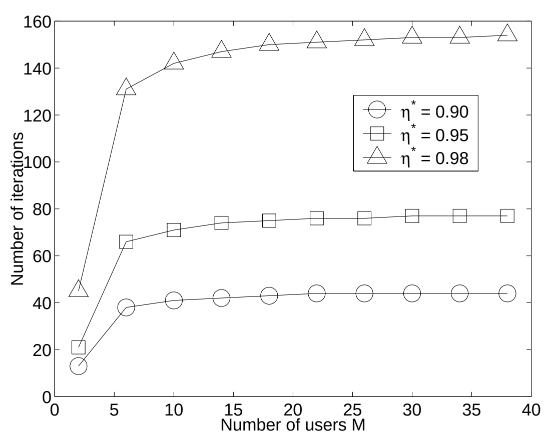

Fig. 5. Performance of the myopic bid updating algorithm with logarithmic utilities and co-located receivers: number of iterations required for a system with $(P/n_0, B) = \left(10^4 M, 10^2 M\right)$ and different target η^*

Next we show some numerical examples with non-collocated receivers. Figures 7 and 8 show the convergence of users' bids and transmit powers in an SINR auction using the distributed algorithm in Sec. V for the network shown in Fig. 6. The network has three users, with transmitters and receivers located at grid points. The link gains between nodes are inversely proportional to the square of the distance. All users have the same logarithmic utility with $\theta_i = 10$. Proposition 2 says that all users achieve the same SINR at the NE. The final bids and transmit powers depend on the distance between the users' transmitters and the measurement point. Since user 3's transmitter is furthest from the measurement point, user 3 can obtain a relatively high transmit power with a small bid. It is easy to see that if all users transmit with the same power, user 2 receives the most interference, and user 1 receives the least. Figure 8 shows that after compensating for the interference, user 2 transmits with the highest power, and user 1 transmits with the lowest power.

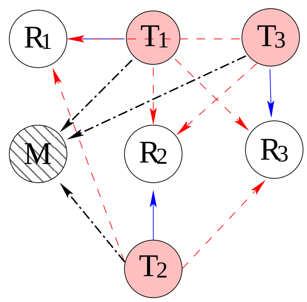

Fig. 6. A three-user network model

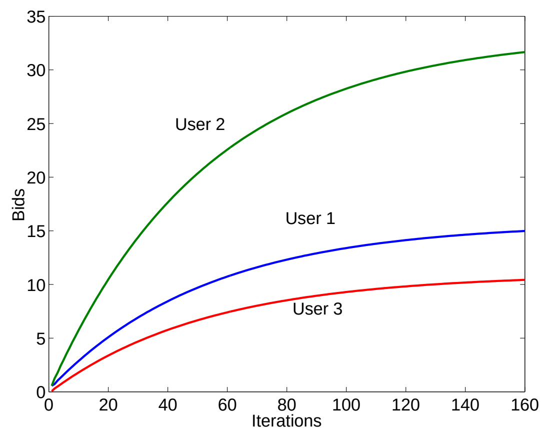

Fig. 7. Convergence of bids in the three-user network

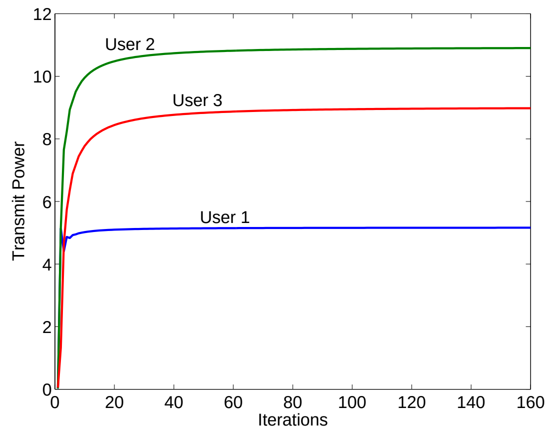

Fig. 8. Convergence of transmit power in the three-user network

VII. Conclusions

We have considered spectrum sharing among a group of spread spectrum users with a constraint on the total interference temperature at a particular measurement point. We proposed two auction mechanisms, SINR- and power-based, that allocate power using a simple proportional bidding rule. When combined with logarithmic utilities, the SINR auction leads to a weighted max-min fair SINR allocation. The following results were obtained for the special case in which the receivers are co-located with the measurement point. Namely, the power auction maximizes the total utility with large enough bandwidth. Also, subject to certain assumptions on the utility functions, the power auction generates more revenue than the SINR auction, although the difference in revenue collected by the two auctions vanishes as the number of users increases. Both auction mechanisms achieve social optimality (i.e., maximize utility per user) in the large system limit where bandwidth and power are increased in fixed proportion. We also presented an iterative, distributed bid updating algorithm, which for both auctions converges globally to the NE.

In this work we have assumed that the users and channels are static, and that the interference temperature is measured at a single location. Relaxing these assumptions leads to directions for future research. A related topic is how to assign bandwidth and power in the context of the Commons spectrum usage model, where there is no spectrum manager to preside over the resource allocation. In that situation, a primary goal is to avoid the “tragedy of commons”.

APPENDIX

A. Proof of Theorem 1

Case I ($\beta > 0$): We first specify the best response $\mathcal{B}_i\left(b_{-i}\right)$ for user $i \in \{1, ..., M\}$ with surplus

$$S_i\left(b_i; b_{-i}\right) = U_i\left(\gamma_i\left(b_i; b_{-i}\right)\right) - \pi^s \gamma_i\left(b_i; b_{-i}\right). \tag{18}$$

Define the normalized channel gain $\hat{h}_{ji} = h_{ji}/h_{j0}$ for all $j, i \geq 1$ so that

$$\gamma_i\left(b_i; b_{-i}\right) = \frac{b_i \hat{h}_{ii} P B}{n_0 B\left(\sum_{j=1}^{M} b_i + \beta\right) + P\left(\sum_{j \neq i} b_j \hat{h}_{ji} + \beta h_{0i}\right)}. \tag{19}$$

Notice that for any fixed b_{-i}, $\gamma_i\left(b_i; b_{-i}\right) \leq P\hat{h}_{ii}/n_0$ and equality is achieved when $b_i \to \infty$.

Differentiating (18) with respect to b_i yields

$$\frac{\partial S_i\left(b_i; b_{-i}\right)}{\partial b_i} = \left[\frac{\partial U_i\left(\gamma_i\left(b_i; b_{-i}\right)\right)}{\partial \gamma_i\left(b_i; b_{-i}\right)} - \pi^s\right] \frac{\partial \gamma_i\left(b_i; b_{-i}\right)}{\partial b_i}, \tag{20}$$

where

$$\frac{\partial \gamma_i\left(b_i; b_{-i}\right)}{\partial b_i} = \frac{\left(n_0 B\left(\sum_{j \neq i} b_i + \beta\right) + P\left(\sum_{j \neq i} b_j \hat{h}_{ji} + \beta h_{0i}\right)\right) \hat{h}_{ii} P B}{\left(n_0 B\left(\sum_{j=1}^{M} b_i + \beta\right) + P\left(\sum_{j \neq i} b_j \hat{h}_{ji} + \beta h_{0i}\right)\right)^2} > 0. \tag{21}$$

Since the term in brackets in (20) is strictly decreasing in b_i, $S_i\left(b_i; b_{-i}\right)$ is a strictly quasi-concave function of b_i, and there exists a unique best response for user i, $\mathcal{B}_i\left(b_{-i}\right)$, that satisfies

$$\begin{cases} \mathcal{B}_i\left(b_{-i}\right) = \infty, & \text{if } \pi^s \leq U_i'\left(\frac{P\hat{h}_{ii}}{n_0}\right) \\ \dfrac{\partial U_i\left(\gamma_i\left(\mathcal{B}_i\left(b_{-i}\right); b_{-i}\right)\right)}{\partial \gamma_i\left(\mathcal{B}_i\left(b_{-i}\right); b_{-i}\right)} = \pi^s, & \text{if } U_i'\left(\frac{P\hat{h}_{ii}}{n_0}\right) < \pi^s < U_i'(0) \\ \mathcal{B}_i\left(b_{-i}\right) = 0, & \text{if } U_i'(0) \leq \pi^s \end{cases} \tag{22}$$

If $\pi^s > \max_{1 \leq i \leq M} U_i'\left(\frac{P\hat{h}_{ii}}{n_0}\right)$, then $\mathcal{B}_i\left(b_{-i}\right) < \infty$, and can be shown to satisfy

$$\mathcal{B}_i\left(b_{-i}\right) = \sum_{j \neq i} k_{ij} b_j + k_{i0}\beta, \tag{23}$$

where k_{ij} is defined in (6), k_{i0} is defined in (7) and $g_i\left(\pi^s\right)$ is defined in (8). Therefore, if the auction has a unique NE $\boldsymbol{b}^*$, then it is the unique component-wise nonnegative solution to

$$\left(\mathbf{I} - \mathbf{K}\right)\boldsymbol{b} = \boldsymbol{k}_0\beta, \tag{24}$$

where $\mathbf{K} = \left[k_{ij}\right]_{i,j \in \{1,...M\}}$ with $k_{ii} = 0$ for all i, and $\boldsymbol{k}_0 = \left(k_{10}, ..., k_{M0}\right)$.[11] Define $\tilde{i} = \arg\max_{i \in \{1,...,M\}} U_i'\left(\frac{P\hat{h}_{ii}}{n_0}\right)$ and $\underline{\pi} = U_{\tilde{i}}'\left(\frac{P\hat{h}_{\tilde{i}\tilde{i}}}{n_0}\right)$ (i.e., $g_{\tilde{i}}\left(\underline{\pi}\right) = \frac{P\hat{h}_{\tilde{i}\tilde{i}}}{n_0}$). When $\pi^s > \underline{\pi}$, $\mathbf{K}$ is a

[11]We denote all vectors as row vectors. The need for transposition should be clear from the context.

nonnegative matrix (i.e., all entries are nonnegative) and $\boldsymbol{k}_0$ is also nonnegative component-wise. Let ρ_K be the spectral radius of matrix $\mathbf{K}$. If $\rho_K < 1$, then $\lim_{n\to\infty}\mathbf{K}^n = 0$, and $(\mathbf{I}-\mathbf{K})^{-1} = \sum_{n=0}^{\infty}\mathbf{K}^n$ exists and is nonnegative. In that case, there is a unique component-wise nonnegative solution to (24) given by

$$\boldsymbol{b}^* = \left(\sum_{n=0}^{\infty}\mathbf{K}^n\right)\boldsymbol{k}_0\beta, \tag{25}$$

which represents the unique NE of the auction. On the other hand, if $\rho_K \geq 1$, then $\sum_{n=0}^{\infty}\mathbf{K}^n = \infty$, and the auction has no NE.

To show the existence of π^s_{th}, as defined in the theorem, we will consider the following two subcases: (I.1) Only user $\tilde{\imath}$ has a positive best response at price $\underline{\pi}$, i.e., $g_l(\underline{\pi}) = 0$ for all $l \neq \tilde{\imath}$, and (I.2) There is at least one other user $l \neq \tilde{\imath}$ who has a positive best response at price $\underline{\pi}$.

Subcase I.1 ($g_l(\underline{\pi}) = 0$ for all $l \neq \tilde{\imath}$): Here we must have $\pi^s_{th} = \underline{\pi}$. This is because for any $\pi^s > \underline{\pi}$, $\mathcal{B}_l(b_{-l}) = 0$ for all $l \neq \tilde{\imath}$, and the unique NE $\boldsymbol{b}^* = (0,...,0,b^*_{\tilde{\imath}},0,...,0)$ where

$$b^*_{\tilde{\imath}} = k_{\tilde{\imath}0}\beta \geq 0. \tag{26}$$

For all $\pi^s \leq \underline{\pi}$, $\mathcal{B}_{\tilde{\imath}}(b_{-\tilde{\imath}}) = \infty$ and there exists no NE.

Subcase I.2 ($\exists l \neq \tilde{\imath}$ such that $g_l(\underline{\pi}) > 0$): To prove this subcase we first show the following two statements: (i) ρ_K is continuous and nonincreasing in π^s. (ii) There exists $\pi^s_H > \underline{\pi}$ such that $\rho_K(\pi^s_H) < 1$. Since $\rho_K(\underline{\pi}) \geq 1$, it then follows that there exists $\pi^s_{th} \in [\underline{\pi}, \pi^s_H)$ such that $\rho_K(\pi^s) \geq 1$ for any $\underline{\pi} \leq \pi^s \leq \pi^s_{th}$, and $\rho_K(\pi^s) < 1$ for any $\pi^s > \pi^s_{th}$. Additionally, we show that in this subcase, $\pi^s_{th} > \underline{\pi}$, i.e., there exists $\pi^s_L > \underline{\pi}$ such that $\rho_K(\pi^s_L) > 1$.

To show (i), let $\boldsymbol{x} = (x_1,...,x_M)$ be a nonnegative vector. From Corollary $8.3.3$ of [25] and the fact that a square matrix has the same eigenvalues as its transpose, we have

$$\rho_K(\pi^s) = \max_{\substack{\boldsymbol{x}\geq 0\\ \boldsymbol{x}\neq 0}} \min_{\substack{j\in\{1,...,M\}\\ x_j\neq 0}} \frac{1}{x_j}\sum_{i=1}^{M} k_{ij}(\pi^s)\,x_i, \tag{27}$$

where the dependence of ρ_K and k_{ij} on π^s are explicitly shown. Let $\boldsymbol{x}^*(\pi^s)$ be a vector that achieves $\rho_K(\pi^s)$ in (27). Note that $\boldsymbol{x}^*(\pi^s)$ must have more than one positive entry, otherwise $\rho_K(\pi^s) = 0$. Assume that $\underline{\pi} < \pi^s < \widetilde{\pi}^s$. From (6), $k_{ij}(\pi^s)$ is nonnegative, continuous and nonincreasing in $\pi^s > \underline{\pi}$. Hence,

$$\frac{1}{x_j}\sum_{i=1}^{M} k_{ij}(\pi^s)\,x_i \geq \frac{1}{x_j}\sum_{i=1}^{M} k_{ij}(\widetilde{\pi}^s)\,x_i \tag{28}$$

for any nonnegative $\boldsymbol{x}$ that has more than one positive entry and $x_j \neq 0$. This implies that

$$\max_{\substack{\boldsymbol{x}\geq 0\\ \boldsymbol{x}\neq 0}} \min_{\substack{j\in\{1,\ldots,M\}\\ x_j\neq 0}} \frac{1}{x_j}\sum_{i=1}^{M} k_{ij}\left(\pi^s\right) x_i \geq \max_{\substack{\boldsymbol{x}\geq 0\\ \boldsymbol{x}\neq 0}} \min_{\substack{j\in\{1,\ldots,M\}\\ x_j\neq 0}} \frac{1}{x_j}\sum_{i=1}^{M} k_{ij}\left(\widetilde{\pi}^s\right) x_i, \tag{29}$$

i.e., $\rho_K\left(\pi^s\right) \geq \rho_K\left(\widetilde{\pi}^s\right)$. Since each eigenvalue of a square matrix depends continuously upon its entries (see appendix D of [25]), $\rho_K\left(\pi^s\right)$ is continuous and nonincreasing in π^s for $\pi^s > \underline{\pi}$.

To show (ii), we have from Theorem 8.1.22 of [25],

$$\rho_K\left(\pi^s\right) \leq \max_{j\in\{1,\ldots,M\}} \sum_{i\neq j}^{M} k_{ij}\left(\pi^s\right). \tag{30}$$

Thus it is sufficient to show that

$$\max_{i,j\in\{1,\ldots,M\}} k_{ij}\left(\pi_H^s\right) < \frac{1}{M-1}. \tag{31}$$

Using (6), a sufficient condition for (31) is

$$\pi_H^s > \max_{i\in\{1,\ldots,M\}} U_i'\left(\frac{PB\min_{i\in\{1,\ldots,M\}}\hat{h}_{ii}}{MBn_0+(M-1)P\max_{i,j\in\{1,\ldots,M\}}\hat{h}_{ji}}\right) > \max_{i\in\{1,\ldots,M\}} U_i'\left(\frac{P\hat{h}_{ii}}{n_0}\right) = \underline{\pi}. \tag{32}$$

To show there exists $\pi_L^s > \underline{\pi}$ such that $\rho_K\left(\pi_L^s\right) > 1$, from (27) it is sufficient to show that there exists an $\boldsymbol{x} > 0$ and $\delta > 0$ such that $\pi_L^s = \underline{\pi} + \delta$ and

$$\sum_{i=1}^{M} k_{ij}\left(\pi_L^s\right)\frac{x_i}{x_j} > 1, \forall j \in \{1,..,M\}. \tag{33}$$

From (8) and the assumptions in Subcase I.2, both $1/g_{\tilde{\imath}}\left(\pi^s\right)$ and $1/g_l\left(\pi^s\right)$ are positive, continuous and strictly increasing functions for $\pi^s \in \left[\underline{\pi}, \underline{\pi}+\delta'\right)$ with $\delta' < \min\left(U_l'(0), U_{\tilde{\imath}}'(0)\right) - \underline{\pi}$. Then for any given $\delta_{\tilde{\imath}} > 0$ and $\delta_l > 0$, there exists a $\delta' > 0$ such that for any $\delta < \delta'$,

$$0 < \frac{1}{g_{\tilde{\imath}}\left(\underline{\pi}+\delta\right)} - \frac{1}{g_{\tilde{\imath}}\left(\underline{\pi}\right)} \leq \delta_{\tilde{\imath}}, \tag{34}$$

$$0 < \frac{1}{g_l\left(\underline{\pi}+\delta\right)} - \frac{1}{g_l\left(\underline{\pi}\right)} \leq \delta_l. \tag{35}$$

If we let $\delta_l = 1/g_l\left(\underline{\pi}\right) - n_0/\left(P\hat{h}_{ll}\right) > 0$, $\delta_{\tilde{\imath}} = n_0^2/\left(4\delta_l P^2\hat{h}_{\tilde{\imath}\tilde{\imath}}\hat{h}_{ll}\right) > 0$, $x_{\tilde{\imath}} = 1$ and $x_j = \left(n_o/P\hat{h}_{\tilde{\imath}\tilde{\imath}}\right)/\delta_{\tilde{\imath}}$ for all $j \neq \tilde{\imath}$, then

$$\frac{x_j}{x_{\tilde{\imath}}} = \frac{n_o/P\hat{h}_{\tilde{\imath}\tilde{\imath}}}{\delta_{\tilde{\imath}}} < \frac{\left(n_0+\frac{P}{B}\hat{h}_{j\tilde{\imath}}\right)/\left(P\hat{h}_{\tilde{\imath}\tilde{\imath}}\right)}{\frac{1}{g_{\tilde{\imath}}\left(\underline{\pi}+\delta\right)} - \frac{1}{g_{\tilde{\imath}}\left(\underline{\pi}\right)}} = k_{\tilde{\imath}j}\left(\underline{\pi}+\delta\right) = k_{\tilde{\imath}j}\left(\pi_L^s\right), \forall j \neq \tilde{\imath}, \tag{36}$$

where we have used the fact that $g_{\tilde{\imath}}(\underline{\pi}) = P\hat{h}_{\tilde{\imath}\tilde{\imath}}/n_0$ by definition. Thus $k_{\tilde{\imath}j}(\pi_L^s)\, x_{\tilde{\imath}}/x_j > 1$ for any $j \neq \tilde{\imath}$. Also

$$\frac{x_l}{x_{\tilde{\imath}}} = \frac{4\delta_l}{n_0/\left(P\hat{h}_{ll}\right)} > \frac{1/g_l(\underline{\pi}) - n_0/\left(P\hat{h}_{ll}\right) + \delta_l}{n_0/\left(P\hat{h}_{ll}\right)} > \frac{1/g_l(\underline{\pi}+\delta_\pi) - n_0/\left(P\hat{h}_{ll}\right)}{n_0/\left(P\hat{h}_{ll}\right) + \hat{h}_{\tilde{\imath}l}/\left(B\hat{h}_{ll}\right)} = \frac{1}{k_{l\tilde{\imath}}(\pi_L^s)}, \tag{37}$$

i.e., $k_{l\tilde{\imath}}(\pi_L^s)\, x_l/x_{\tilde{\imath}} > 1$. Combining (36) and (37) give (33), hence $\rho_K(\pi_L^s) = \rho_K(\underline{\pi}+\delta_\pi) < 1$.

Case II ($\beta = 0$): First, we observe that $\boldsymbol{b}^* = \mathbf{0}$ is an NE if and only if

$$U_i(0) \geq U_i\left(\frac{P\hat{h}_{ii}}{n_0}\right) - \pi^s \frac{P\hat{h}_{ii}}{n_0}, \forall i. \tag{38}$$

That is, if all other users bid zero, then user i's best response bid is also zero since a positive bid gives the change in surplus $\Delta S_i(b_i; b_{-i}) = U_i\left(\frac{P\hat{h}_{ii}}{n_0}\right) - \pi^s \frac{P\hat{h}_{ii}}{n_0} - U_i(0) \leq 0$. Furthermore, if there is a unique NE, then $\boldsymbol{b}^* = \mathbf{0}$. This is because if there exists a nonzero $\tilde{\boldsymbol{b}}^*$, which is a NE, then for any scalar $\upsilon > 0$, $\upsilon\tilde{\boldsymbol{b}}^*$ gives the same surplus values, hence is also a NE. Thus there are an infinite number of Nash Equilibria. Finally, there is no NE when π^s is too small (e.g., $\pi^s \leq U_i'\left(\frac{P\hat{h}_{ii}}{n_0}\right)$ for some user i).

B. *Proof of Theorem 2*

Given an $\varepsilon \in (0,1)$, it is straightforward to write out the Kuhn-Tucker (KT) conditions for the total utility maximization problem of the ε-system with co-located receivers:

$$\begin{aligned} \underset{\boldsymbol{p}^{r\varepsilon} \geq \mathbf{0}}{\text{maximize}} \quad & \sum_{i=1}^{M} U_i\left(\gamma_i\left(p_i^{r\varepsilon}\right)\right) \\ \text{subject to} \quad & \gamma_i\left(p_i^{r\varepsilon}\right) = \frac{p_i^{r\varepsilon}}{n_0 + \left(P - p_i^{r\varepsilon}\right)/B} \\ & \sum_{i=1}^{M} p_i^{r\varepsilon} \leq P(1-\varepsilon). \end{aligned} \tag{39}$$

Since problem (39) is a strictly convex maximization problem under Assumption 2, the KT conditions are necessary and sufficient for the unique ε-social optimal solution.

In the power auction, user i's surplus function $S_i(b_i; b_{-i}) = U_i\left(\gamma_i\left(p_i^r(b_i; b_{-i})\right)\right) - \pi^p p_i^r(b_i; b_{-i})$ is a strictly quasi-concave function in b_i. Hence there exists a unique value of b_i that maximizes $S_i(b_i; b_{-i})$ for fixed b_{-i}. By setting π^p equal to the Lagrange multiplier in the KT conditions for problem (39), the set of best responses for the users is the solution to the KT conditions. Thus the power profile at the NE achieves ε-social optimality for any $\varepsilon \in (0,1)$.

C. Proof of Proposition 4

With logarithmic utilities and co-located receivers, the first-order conditions for surplus maximization for user i gives

$$\pi^p = U_i'\left(\theta_i, \gamma_i\left(p_i^{r*}\right)\right)\gamma_i'\left(p_i^{r*}\right) = \frac{\theta_i\left(n_0 B + P\right)}{p_i^{r*}\left(n_0 B + P - p_i^{r*}\right)}. \tag{40}$$

Thus,

$$R^p = \sum_{i=1}^{M} \pi^p p_i^{r*} = \sum_{i=1}^{M} \frac{\theta_i\left(n_0 B + P\right)}{\left(n_0 B + P - p_i^{r*}\right)} > \sum_{i=1}^{M} \theta_i = R^s, \tag{41}$$

where the last equality is shown in the proof of Proposition 2. If $\theta_i \leq \bar{\theta}$ for each i, then as $M \to \infty$, $p_i^{r*} \to 0$ for each user i, and $R^p/R^s \to 1$.

D. Proof of Proposition 5

We obtain (13) by taking the limit of the conditions in Proposition $1,$ under the assumed scaling. Let Lim denote $\lim_{P,B,M\to\infty}$ with $P/B, P/M, \beta/M$ fixed. Thus,

$$Lim \sum_{i=1}^{M} \frac{k_i}{1+k_i} = Lim \sum_{i=1}^{M} \frac{\theta_i\left(P/B + n_0\right)}{P\left(\pi^s + \theta_i/B\right)} = \frac{1}{M} Lim \sum_{i=1}^{M} \frac{M\theta_i\left(P/B + n_0\right)}{P\pi^s} = \frac{P/B + n_0}{P/M\pi^s} E\left[\theta\right] \tag{42}$$

with probability 1. The first equality follows from the definition of k_i in (9), the second follows from the limit $B \to \infty$, and the third follows from the strong law of large numbers. Condition (13) then follows directly. The weighted max-min fair SINR allocation and payments stay fixed during the limiting process. Since every user sees the same noise plus interference at the NE, $n_0 + P/B$, we have $p_i^{r*} = \gamma_i^*(n_0 + P/B)$ for all i. This corresponds to a weighted max-min fair power allocation.

E. Proof of Theorem 4 (Sketch)

In the limiting system, the maximum average utility per user is the solution to:

$$\begin{aligned} \underset{p^r(\theta)\geq 0}{\text{maximize}} \quad & E_\theta\left[U\left(\theta, \frac{p^r\left(\theta\right)}{n_0 + \left(P - p^r\left(\theta\right)\right)/B}\right)\right] \\ \text{subject to} \quad & E_\theta\left[p^r\left(\theta\right)\right] = \frac{P}{M}\left(1-\varepsilon\right) \end{aligned} \tag{43}$$

The objective is the average utility per user, and the constraint corresponds to the total received power constraint in the ε-system. In both cases we have used the law of large numbers to express these in terms of expectations over θ.

The optimization is over all received power allocations, $p^r : \left[\underline{\theta}, \overline{\theta}\right] \to \mathcal{R}^+$. We first prove the following lemma:

Lemma 2: There exists a power allocation $p^r(\theta)$ that solves (43), which is finite everywhere, i.e.,

$$\lim_{P\to\infty} \frac{p^r(\theta)}{P} = 0, \forall \theta \in \left[\underline{\theta}, \overline{\theta}\right]. \tag{44}$$

This lemma implies that each user receives a negligible fraction of the total power as the system scales. The lemma can be proved by contradiction. If the lemma were not true, then at least one user would be allocated infinite power as the system scales. Because the utility is sublinear, this user would contribute a negligible amount to the average utility. Thus we could reallocate the user's power among the remaining users and strictly increase the average utility. This gives a contradiction, proving the lemma.

Lemma 2 ensures that at a solution to (43), each user receives the same interference plus noise $n_0 + P/B$. This makes (43) a strictly concave maximization problem. By using calculus of variations [26], we can solve for $p(\theta)$ in closed form, as well as for the corresponding positive Lagrange multiplier λ for the average power constraint. Letting $\pi^p = \lambda$ or $\pi^s = (n_0 + P/B)\lambda$ results in the same power allocation at the NE for the power and SINR auctions, respectively.

References

[1] "Spectrum policy task force report," Federal Communications Commission, US, Nov. 2002.

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

[3] K. J. Sunnevag, "Auction design for the allocation of emission permits," 2001, working paper.

[4] J. J. D. Wang and J. F. Zender, "Auctioning divisible goods," *Economic Theory*, no. 19, pp. 673–705, 2002.

[5] K. Back and J. F. Zender, "Auctions of divisible goods: On the rationale for the treasury experiment," *Review of Financial Studies*, vol. 6, pp. 733–764, 1993.

[6] R. Johari and J. N. Tsitsiklis, "Efficiency loss in a network resource allocation game," *Mathematics of Operations Research*, vol. 29, no. 3, pp. 407–435, Aug. 2004.

[7] R. Maheswaran and T. Basar, "Nash equilibrium and decentralized negotiation in auctioning divisible resources," *Group Decision and Negotiation*, vol. 12, no. 5, pp. 361–395, 2003.

[8] S. Yang and B. Hajek, "An efficient mechanism for allocation of a divisible good," 2004, working paper.

[9] R. Tijdeman, "The chairman assignment problem," *Discrete Mathematics*, vol. 32, pp. 323–330, 1980.

[10] C. A. Waldspurger and W. E. Weihl, "Stride scheduling: Deterministic proportional share resource management," Laboratory for CS, MIT, Technical Memorandum MIT/LCS/TM-528, July 1995.

[11] P. Dubey, "Inefficiency of Nash equilibria," *Mathematics of Operations Research*, vol. 11, pp. 1–8, 1986.

[12] D. G. D. Famolari, N. B. Mandayam and V. Shah, *A New Framework for Power Control in Wireless Data Networks: Games, Utility and Pricing*. Kluwer Academic Publishers, 1999, pp. 289–310.

[13] 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.

[14] 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.

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

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

[17] S. Shenker, "Fundamental design issues for the future internet," *IEEE Journal of Selected Areas in Communications*, vol. 13, pp. 1176–1188, 1995.

[18] P. Milgrom, *Putting Auction Theory to Work*. Cambridge University Press, 2004.

[19] 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.

[20] T. Alpcan and T. Basar, "Distributed algorithms for Nash equilibria of flow control games," *Annals of Dynamic Games*, vol. 7, 2003.

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

[22] G. J. Foschini and Z. Miljanic, "A simple distributed autonomous power control algorithm and its convergence," *IEEE Transactions on Vehicular Technology*, vol. 42, no. 4, Nov. 1993.

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

[24] T. Holliday, A. J. Goldsmith, and P. Glynn, "Distributed power control for time varying wireless networks: Optimality and convergence," in *Proc. of Allerton Conference*, 2003.

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

[26] M. I. Kamien and N. L. Schwartz, *Dynamic Optimization: The Calculus of Variations and Optimal Control in Economics and Management*. North-Holland, 1981.